

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 December 2007 (21.12.2007)

PCT

(10) International Publication Number
WO 2007/147084 A2

(51) International Patent Classification:
G06F 17/50 (2006.01) **H01L 21/82** (2006.01)

(21) International Application Number:
PCT/US2007/071262

(22) International Filing Date: 14 June 2007 (14.06.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/804,690 14 June 2006 (14.06.2006) US
60/804,826 15 June 2006 (15.06.2006) US
PCT/US2006/025294 28 June 2006 (28.06.2006) US
PCT/US2007/069214 18 May 2007 (18.05.2007) US
PCT/US2007/070814 8 June 2007 (08.06.2007) US
PCT/US2007/071050 12 June 2007 (12.06.2007) US

(71) Applicant (for all designated States except US): **LIGHT-SPEED LOGIC, INC.** [US/US]; Suite 2-102, 3255 Scott Blvd., Santa Clara, California 95054 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **BOSE, Subhasis** [US/US]; 11849 Sterling Panorama Terrace, Austin, TX 78738 (US).

(74) Agent: **SMITH, Walstein**; P.O. Box 1668, Georgetown, Texas 78627-1668 (US).

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

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

Declaration under Rule 4.17:

— of inventorship (Rule 4.17(iv))

Published:

— with declaration under Article 17(2)(a); without abstract; title not checked by the International Searching Authority

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: GENERALIZED CLOCK TREE SYNTHESIS

(57) Abstract:



WO 2007/147084 A2

1 **GENERALIZED CLOCK TREE SYNTHESIS**

2
3
4
5 **CROSS REFERENCE TO RELATED APPLICATIONS**

6
7 **[0001]** Priority benefit claims for this application are made in the accompanying
8 Application Data Sheet, Request, or Transmittal (as appropriate, if any). To the extent permitted
9 by the type of the instant application, this application incorporates by reference for all purposes
10 the following applications, all owned by the owner of the instant application:

11 U.S. Provisional Application Serial No. 60/804,690 (Docket No. LS.2006.13), filed June
12 14, 2006, first named inventor Subhasis Bose, and entitled **GENERALIZED**
13 **CLOCK TREE SYNTHESIS FOR STRUCTURED ARRAY FABRIC;**

14 U.S. Provisional Application Serial No. 60/804,826 (Docket No. LS.2006.14), filed June
15 15, 2006, first named inventor Geoffrey Mark Furnish, and entitled
16 **SIMULTANEOUS DYNAMICAL INTEGRATION APPLIED TO**
17 **DETAILED PLACEMENT;**

18 PCT Application Serial No. PCT/US2006/025294 (Docket No. LS.2006.01B), filed June
19 28, 2006, first named inventor Geoffrey Furnish, and entitled **METHODS AND**
20 **SYSTEMS FOR PLACEMENT;**

21 PCT Application Serial No. PCT/US2007/069214 (Docket No. LS-07-01PCT), filed
22 May 18, 2007, first named inventor Subhasis Bose, and entitled **METHODS**
23 **AND SYSTEMS FOR PLACEMENT;**

24 PCT Application Serial No. PCT/US2007/070814 (Docket No. LS-07-09PCT), filed
25 June 8, 2007, first named inventor Geoffrey Furnish, and entitled **MORPHING**
26 **FOR GLOBAL PLACEMENT USING INTEGER LINEAR**
27 **PROGRAMMING; and**

28 PCT Application Serial No. PCT/US2007/071027 (Docket No. LS-07-02PCT), filed
29 June 12, 2007, first named inventor Maurice LeBrun, and entitled **METHODS**
30 **AND SYSTEMS FOR PLACEMENT.**

BACKGROUND

[0002] Field: Advancements in integrated circuit design, including placement and routing of elements in a Computer Aided Design (CAD) context, are needed to provide improvements in performance, efficiency, and utility of use.

[0003] Related Art: Unless expressly identified as being publicly or well known, mention herein of techniques and concepts, including for context, definitions, or comparison purposes, should not be construed as an admission that such techniques and concepts are previously publicly known or otherwise part of the prior art. All references cited herein (if any), including patents, patent applications, and publications, are hereby incorporated by reference in their entireties, whether specifically incorporated or not, for all purposes.

SYNOPSIS

1
2
3 [0004] The invention may be implemented in numerous ways, including as a process,
4 an article of manufacture, an apparatus, a system, a composition of matter, and a computer
5 readable medium such as a computer readable storage medium or a computer network wherein
6 program instructions are sent over optical or electronic communication links. In this
7 specification, these implementations, or any other form that the invention may take, may be
8 referred to as techniques. The Detailed Description provides an exposition of one or more
9 embodiments of the invention that enable improvements in performance, efficiency, and utility
10 of use in the field identified above. The Detailed Description includes an Introduction to
11 facilitate the more rapid understanding of the remainder of the Detailed Description. The
12 Introduction includes Example Embodiments of one or more of systems, methods, articles of
13 manufacture, and computer readable media in accordance with the concepts described herein.
14 As is discussed in more detail in the Conclusions, the invention encompasses all possible
15 modifications and variations within the scope of the issued claims.

Brief Description of Drawings

[0005] Fig. 1 is a flow diagram illustrating selected details of an embodiment of placing, routing, analyzing, and generating fabrication data for any portion of an integrated circuit according to a Simultaneous Dynamical Integration (SDI)-based flow.

[0006] Fig. 2 is a flow diagram illustrating selected details of an embodiment of placing and routing any portion of an integrated circuit according to an SDI-based flow.

[0007] Fig. 3A is a flow diagram illustrating selected details of an embodiment of global placement according to SDI-based modeling and simulation.

[0008] Fig. 3B is a flow diagram illustrating selected details of an embodiment of initial placement operations for global placement.

[0009] Fig. 3C is a flow diagram illustrating selected details of an embodiment of density field based force component computation.

[0010] Fig. 3D is a flow diagram illustrating selected details of an embodiment of gate density accumulation.

[0011] Fig. 3E is a conceptual diagram illustrating an embodiment of two-point interpolation of node mass to grid points.

[0012] Fig. 3F is a conceptual diagram illustrating an embodiment of three-point interpolation of node mass to grid points.

[0013] Fig. 3G is a conceptual diagram illustrating an embodiment of applying boundary grid point masses to interior grid points.

[0014] Fig. 3H is a flow diagram illustrating selected details of an embodiment of digital density filtering.

[0015] Fig. 3I is a flow diagram illustrating selected details of an embodiment of interpolating gate fields to nodes.

1 **[0016]** Fig. 4 is a flow diagram illustrating selected details of an embodiment of SDI-
2 based modeling and simulation.

3

4 **[0017]** Fig. 5A is a flow diagram illustrating selected details of a first embodiment of
5 resource reconciliation, as a first example of legalization.

6

7 **[0018]** Fig. 5B is a flow diagram illustrating selected details of a second embodiment of
8 resource reconciliation, as a second example of legalization.

9

10 **[0019]** Fig. 5C is a flow diagram illustrating selected details of an embodiment of
11 partitioning.

12

13 **[0020]** Fig. 6 is a flow diagram illustrating selected details of an embodiment of
14 detailed placement (also referred to as detail placement elsewhere herein).

15

16 **[0021]** Fig. 7A is a flow diagram illustrating selected aspects of an embodiment of
17 delay path reduction and minimization, as an example of timing closure.

18

19 **[0022]** Fig. 7B illustrates a conceptual view of selected elements of an embodiment of
20 timing-driven forces.

21

22 **[0023]** Fig. 7C illustrates a spatial organization of the driver and the coupled loads of
23 Fig. 7B.

24

25 **[0024]** Fig. 7D illustrates an embodiment of Net Boundary Box (NBB) estimation of
26 routing to cover the driver and the loads of Fig. 7C.

27

28 **[0025]** Fig. 7E illustrates an embodiment of a rectilinear Steiner Route Tree (SRT)
29 estimation to cover the driver and loads of Fig. 7C.

30

31 **[0026]** Fig. 7F illustrates an embodiment of estimated RC parasitics associated with the
32 RST of Fig. 7E.

33

34 **[0027]** Figs. 8A and 8B collectively are a flow diagram illustrating selected details of
35 an embodiment of an integrated circuit Electronic Design Automation (EDA) flow using one or

more techniques including SDI-directed global placement, legalization, legalization-driven detailed placement, timing optimization, and routing.

[0028] Fig. 9 illustrates selected details of an embodiment of manufacturing integrated circuits, the circuits being designed in part based on SDI-directed design techniques.

[0029] Fig. 10 illustrates selected details of an embodiment of a computer system to execute EDA routines to perform SDI-directed place and route operations.

[0030] Fig. 11 illustrates an embodiment of an SDI-based detailed placement flow.

[0031] Figs. 12A and 12B illustrate concepts relating to an embodiment of netlist elaboration.

[0032] Fig. 13 illustrates an embodiment of detailed placement of a Q-block.

[0033] Fig. 14 illustrates an embodiment of an additional pass of detailed placement of a Q-block.

[0034] Fig. 15A illustrates a form of the form-level net of Fig. 12A. In this view the resource-level nodes are shown internal to the form.

[0035] Fig. 15B illustrates another form that uses different resources to implement the same function as the form of Fig. 15A. In at least one embodiment, the form of Fig. 15B is substituted for the form of Fig. 15A through a morphing process.

[0036] Fig. 15C illustrates a hierarchy of nodes, having hierarchical nodes, form-level nodes, and resource-level nodes.

[0037] Fig. 15D illustrates selected nets connected between selected nodes of Fig. 15C.

[0038] Fig. 15E illustrates the nodes and nets of Fig. 15D after augmentation with resource-level nodes.

[0039] Fig. 16A illustrates the supply and demand for resources R1 through R6 corresponding to target functions of an integrated circuit design having a first selection of forms

1 for the target functions. For at least some of the resources, the demand exceeds the available
2 supply.

3
4 **[0040]** Fig. 16B illustrates the supply and demand for resources R1 through R6 for the
5 same target functions, but using a second selection of forms for the target functions obtained by
6 morphing certain forms to use different resources. For each of the resources shown, the demand
7 is less than or equal to the supply.

8
9 **[0041]** Fig. 17A illustrates an example circuit with a plurality of critical paths.

10
11 **[0042]** Fig. 17B illustrates example computations relating to an embodiment of CPF
12 scoring.

13
14 **[0043]** Fig. 18 illustrates an embodiment of a cascade of buffers of increasing drive
15 strength.

16
17 **[0044]** Fig. 19 illustrates example computations relating to an embodiment of SDF
18 calculation.

19
20 **[0045]** Fig. 20A illustrates an overall procedural control flow in an illustrative relative
21 slack embodiment.

22
23 **[0046]** Fig. 20B illustrates the adjustment of timing driven weight in the relative slack
24 embodiment of Fig. 20A.

25
26 **[0047]** Fig. 21A illustrates a driver in the interior of a net bounding box region.

27
28 **[0048]** Fig. 21B illustrates a driver to one side of a net bounding box region.

29
30 **[0049]** Figs. 22A and 22B illustrate an example circuit excerpt before and after
31 processing according to an embodiment of timing driven buffering and resizing for an array
32 architecture.

33
34 **[0050]** Fig. 23 illustrates a flow diagram of an integrated circuit design flow including
35 an embodiment of processing in accordance with an embodiment of timing driven buffering and
36 resizing for an array architecture.

1

2 **[0051]** Fig. 24A illustrates a top-level view of an embodiment of timing driven
3 buffering and resizing for an array architecture.

4

5 **[0052]** Fig. 24B illustrates a detail view of selected details of an embodiment of timing
6 driven resizing for an array architecture.

7

8 **[0053]** Figs. 25A and 25B illustrate an example route tree as processed by an
9 embodiment of segmenting a portion of the route for timing driven buffering and resizing.

10

11 **[0054]** Fig. 26 illustrates example results of an embodiment of logic replication and
12 tunneling for an array architecture.

13

14 **[0055]** Fig. 27 illustrates a control flow in an illustrative embodiment, as used for
15 density modification.

16

17 **[0056]** Fig. 28 illustrates a control flow of an illustrative embodiment, as used to
18 determine the Steiner-cuts congestion term on the SDI grid.

19

20 **[0057]** Fig. 29 illustrates procedures of an illustrative embodiment, showing creation of
21 a congestion array.

22

23 **[0058]** Fig. 30 illustrates procedures of an illustrative embodiment, showing calculation
24 of a final congestion density enhancement array.

25

26 **[0059]** Fig. 31 illustrates an embodiment of a processing flow for node tunneling
27 through exclusion zones in an SDI-based integrated circuit design flow.

28

29 **[0060]** Fig. 32 illustrates an embodiment of SDI-related force calculations in a
30 tunneling congestion relief context.

31

32 **[0061]** Fig. 33 illustrates an embodiment of evaluation of tunneling transition criteria.

33

34 **[0062]** Fig. 34A illustrates an example clock tree suitable for input to a Clock Tree
35 Synthesis (CTS) tool for Structured Array Fabric (SAF)-based design flows.

36

- 1 **[0063]** Fig. 34B illustrates an example clock tree output from the CTS tool operating on
2 the input illustrated in Fig. 34A.
3
- 4 **[0064]** Fig. 34C illustrates an example clock tree network.
5
- 6 **[0065]** Fig. 35 illustrates an overview of an embodiment of a CTS flow.
7
- 8 **[0066]** Fig. 36A illustrates an example die floorplan of a design having embedded
9 Random Access Memory (RAM) or other Intellectual Property (IP) blocks.
10
- 11 **[0067]** Fig. 36B illustrates a portion of a clock net in a context of a portion of Fig. 36A.
12
- 13 **[0068]** Fig. 37A illustrates an example of timing driven pin swapping.
14
- 15 **[0069]** Fig. 37B illustrates an example of effects of clock tree partitioning.
16
- 17 **[0070]** Fig. 38 illustrates an analysis according to an embodiment of clock domain and
18 sub-domain partitioning.

DETAILED DESCRIPTION

[0071] A detailed description of one or more embodiments of the invention is provided below along with accompanying figures illustrating selected details of the invention. The invention is described in connection with the embodiments. It is well established that it is neither necessary, practical, or possible to exhaustively describe every embodiment of the invention. Thus the embodiments herein are understood to be merely exemplary, the invention is expressly not limited to or by any or all of the embodiments herein, and the invention encompasses numerous alternatives, modifications and equivalents. To avoid monotony in the exposition, a variety of word labels (including but not limited to: first, last, certain, various, further, other, particular, select, some, and notable) may be applied to separate sets of embodiments; as used herein such labels are expressly not meant to convey quality, or any form of preference or prejudice, but merely to conveniently distinguish among the separate sets. The order of some operations of disclosed processes is alterable within the scope of the invention. Wherever multiple embodiments serve to describe variations in process, method, and/or program instruction features, other embodiments are contemplated that in accordance with a predetermined or a dynamically determined criterion perform static and/or dynamic selection of one of a plurality of modes of operation corresponding respectively to a plurality of the multiple embodiments. Numerous specific details are set forth in the following description to provide a thorough understanding of the invention. These details are provided for the purpose of example and the invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the invention is not unnecessarily obscured.

INTRODUCTION

[0072] This introduction is included only to facilitate the more rapid understanding of the Detailed Description; the invention is not limited to the concepts presented in the introduction (including explicit examples, if any), as the paragraphs of any introduction are necessarily an abridged view of the entire subject and are not meant to be an exhaustive or restrictive description. For example, the introduction that follows provides overview information limited by space and organization to only certain embodiments. There are many other embodiments, including those to which claims will ultimately be drawn, discussed throughout the balance of the specification.

1

2 **[0073]** As described herein, “dynamic time-evolving SDI” refers to SDI techniques for
3 the modeling and simulation of elements for integrated circuit placement and routing. Dynamic
4 time-evolving SDI includes applying principles of Newtonian mechanics to an “analogy-system”
5 based on a netlist that is a specification of the integrated circuit as part of an EDA flow (such as
6 during physical design development of the integrated circuit). In some usage scenarios the
7 analogy-system (often referred to simply as “system”) includes a single point particle
8 corresponding to each device in the netlist. The system further includes a set of one or more
9 forces acting on each of the particles, in certain embodiments computed as a weighted sum.
10 Various numerical integration techniques are used to apply Newton’s second law of motion to
11 the system, forming a time-evolving representation of the system in state-space. In other words
12 a simulation determines paths of the particles in a plane (or three dimensions). Then resultant
13 locations of the point particles are mapped back into resultant placements of the corresponding
14 devices, thus providing SDI-directed placements.

15

16 **[0074]** Using dynamic time-evolving SDI, elements of the system are pushed
17 simultaneously forward in time through a smooth integration in which the model for the system
18 dynamics is an abstraction utilizing continuous variables and simultaneous exploration.
19 Departures from idealizations of continuous variables and simultaneity are artifacts of
20 techniques for solving the system of coupled simultaneous governing equations, such as that
21 occur with numerical integration on a digital computer. In such digital computer
22 implementations, the departures are limited to specifiable tolerances determined by the quality of
23 result goals and economic considerations (such as available solution time, supply of computing
24 power available, and other similar constraints).

25

26 **[0075]** The system forces include attractive and spreading components, used to model
27 effects of interconnect, resource usage (such as device area), and to drive various optimizations
28 (such as timing closure). Some of the system forces are directly expressed as functions of the
29 positions of other devices (such as attractive forces between connected devices), some of the
30 forces are indirect functions of the positions of other devices and are computed by way of
31 various fields (such as one or more density fields), and some of the forces that act on some of the
32 devices are independent of the positions of the other devices in the system. Computing selected
33 forces as fields in certain embodiments affords more computational efficiency.

34

35 **[0076]** SDI-directed placement is useful in various integrated circuit design flows and
36 related implementation architectures, including full custom, semi-custom, standard cell,

1 structured array, and gate array design flows and related implementation architectures. Several
2 variations in the context of structured array design flows enable efficient processing of numerous
3 constraints imposed by the paritally predetermined nature of the arrays. A library of composite
4 cells or “morphable-devices” is provided to a synthesis tool (such as Synopsys Design Compiler
5 or any other similar tool). The morphable-devices are used as target logic elements by the
6 synthesis tool to process a netlist (either behavioral or gate-level) provided by a user. A
7 synthesis result is provided as a gate-level netlist (such as a Verilog gate-level netlist) expressed
8 as interconnections of morphable-devices. The synthesis tool assumes the morphable-devices
9 represent the final implementation, subject to device sizing to resolve circuit timing issues.

10
11 **[0077]** The morphable-devices are, however, subject to additional modifications in the
12 structured array design flow context (see “Structured Arrays”, elsewhere herein), as each
13 morphable-device may be implemented in a plurality of manners using varying resources of the
14 structured array. During phases of resource reconciliation (where attempts are made to satisfy
15 required resources with locally available resources), one or more of the morphable-devices may
16 be transformed to a logically equivalent implementation. For example, an AND function may be
17 implemented by an AND gate, by a NAND gate and an Inverter, or by any other equivalent
18 formulation. Functionally equivalent alternatives are grouped according to implementation
19 function, and individual realizations within a given function are referred to as “forms”. Thus
20 any morphable-device may be implemented as any instance of any form having an equivalent
21 function. Subsequent operations account for variation between logically equivalent forms (such
22 as differences in area, timing behavior, routing resources used or provided, and any other
23 characteristic distinguishing one form from another). Operations relating to interchanging
24 implementations of morphable-devices to satisfy structured array resource limitations and
25 underlying topology, as well as meeting spatial organization constraints, are termed “morphing”.

26
27 **[0078]** The SDI-directed placement, in various contexts including structured array
28 design flows, includes several phases: global placement, legalization, and detailed placement.
29 Global placement in certain embodiments provides a first-cut location for each morphable-
30 device in a netlist. The first-cut location is subject to additional refinement by subsequent
31 processing (including legalization and detailed placement). Global placement is considered
32 complete when a configuration is attained that is determined to be sufficiently close to legality to
33 proceed to legalization, i.e. the configuration is likely to be reducible to a satisfactory
34 implementation. Legalization starts with the global placement configuration and produces a
35 final configuration in which demand for resources in every region is determined to be no greater
36 than corresponding supply in each region. Detailed placement starts with the legalized

1 placement configuration and assigns every element implementing a morphable-device to specific
2 resources in an implementation (such as a set of specific resource-slots in a structured array
3 architecture). Some simple functions may have degenerate forms requiring only a single
4 resource instance, but more complex forms are composite, requiring more than one physical
5 resource instance plus internal interconnect to correctly implement the function.

6
7 **[0079]** Various morphing and similar transformation operations may be used in any
8 combination of phases including global placement, legalization, and detailed placement,
9 according to various embodiments. Morphing techniques used in one phase may be distinct or
10 may be substantially similar to morphing techniques used in another phase, varying according to
11 implementation. In some embodiments, different processing phases proceed with morphing
12 operations operating according to respective morphing classes, i.e. a set of morphing classes for
13 global placement, a set of morphing classes for legalization, and set of morphing classes for
14 detailed placement. The morphing classes according to phases may be distinct or may be
15 substantially similar to one another, according to embodiment.

16
17 **[0080]** SDI-directed placement operations, when applied in a structured array design
18 flow context, may include specialized forces relating to various "morphing classes" representing
19 categories of structured array resources or related functionality. For example, resources for
20 combinational circuitry may be grouped in a combinational morphing class, while resources for
21 sequential circuitry may be grouped in a sequential morphing class. In some situations
22 morphable-devices are restricted to implementation by resources belonging to a limited set of
23 morphing-classes. Continuing with the example, combinational logic morphable-devices may
24 be restricted to implementation by resources of the combinational morphing class, while
25 sequential logic morphable-devices may be restricted to implementation by sequential morphing
26 class elements. One or more specialized forces relating to each of the morphing classes may be
27 used during global placement to effect spreading of morphable-devices according to
28 corresponding morphing classes. Continuing with the example, a combinational spreading force
29 may be selectively applied to combinational logic morphable-devices, while a sequential
30 spreading force may be selectively applied to sequential logic morphable-devices. In certain
31 embodiments, it is useful to subject all devices in the netlist (whether morphable or not) to a
32 single spreading force that acts to drive the circuit toward a density that is sustainable on the
33 implementation architecture, and augment the spreading force with the specialized resource-
34 class-specific spreading forces to further tune the placement.

Structured Arrays

[0081] In some usage scenarios structured arrays are implementation vehicles for the manufacture of integrated circuits, as described elsewhere herein. Structured arrays in certain embodiments include fundamental building blocks (known as “tiles”) instantiated one or more times across an integrated circuit substrate to form a Structured Array Fabric (SAF). In some embodiments structured arrays are homogeneous (i.e. all of the tiles are identical), while in some embodiments the arrays are heterogeneous (i.e. some of the tiles are distinct with respect to each other). Heterogeneity may occur as a result of tile type, arrangement, or other differences. Irregardless of tile number and arrangement, however, the SAF tiles are fixed (i.e. prefabricated) and independent of any specific design implemented thereupon.

[0082] SAF tiles, according to various embodiments, may include any combination of fully or partially formed active elements (such as transistors, logic gates, sequential elements, and so forth), as well as fully or partially formed passive elements (such as metallization serving as wires and vias providing interconnection between layers of metal). In some SAF embodiments “lower” layers of interconnect are included in SAF tiles (as the lower layers are formed relatively early in fabrication), while “upper” layers of interconnect are specific to a design (as the upper layers are formed relatively later in fabrication). Such SAF embodiments permit the lower prefabricated (and thus non-customizable) layers to be shared between different design implementations, while the higher/customizable layers provide for design-specific specialization or personalization.

[0083] SAF structures may be used to construct an entire chip, or may constitute only a portion of the floorplan of an encompassing circuit, allowing for design variation. The size of the SAF tiles is generally irrelevant to design flows, and a tile may be as small and simple as a single inverter or as large and complex as a Randomly Accessible read-write Memory (RAM) block or other large-scale Intellectual Property (IP) element.

[0084] EDA flows targeting designs based on structured array technology (such as the SDI-directed flow described elsewhere herein) account for the predetermined nature of the array, from gate-level netlist synthesis through subsequent implementation processing including layout of cells and interconnect. Such EDA flows enable realizing advantages of manufacture of integrated circuits including SAF tiles. The advantages include reduced manufacturing cost, as fewer mask layers (for example those corresponding to upper layers of interconnect) are

- 1 customized for each design, as well as reduced characterization cost (for example by re-use of
- 2 known structures such as the SAF tiles).

1 ILLUSTRATIVE COMBINATIONS

2
3 **[0085]** This introduction concludes with a collection of exemplary illustrative
4 combinations, including some explicitly enumerated as “ECs”, that tersely summarize
5 illustrative systems and methods in accordance with the concepts taught herein. Each of the
6 illustrative combinations or ECs highlights various combinations of features using an informal
7 pseudo-claim format. These compressed descriptions are not meant to be mutually exclusive,
8 exhaustive, or restrictive, and the invention is not limited to these highlighted combinations. As
9 is discussed in more detail in the Conclusion section, the invention encompasses all possible
10 modifications and variations within the scope of the issued claims.

11
12
13 General Exemplary Illustrative Combinations

14
15 **[0086]** A first embodiment for generating a clock tree, the embodiment including
16 determining locations of clock buffers and creating the clock tree based on the locations. The
17 foregoing embodiment wherein the clock buffers are resources of a structured array fabric
18 integrated circuit.

19
20 **[0087]** A second embodiment for generating a clock tree, the embodiment including
21 performing a floorplan driven clock partitioning and creating the clock tree based on the
22 floorplan driven clock partitioning. The foregoing embodiment wherein the clock tree is in
23 accordance with a structured array fabric integrated circuit.

24
25 **[0088]** A third embodiment for generating a clock tree, the embodiment including
26 updating timing of the clock tree, locating buffers of the clock tree according to an order, and
27 iterating the updating and the locating. The foregoing embodiment wherein the order is a top-
28 down order.

29
30 **[0089]** A fourth embodiment for generating a clock tree, the embodiment including
31 clustering leaf level clock pins of the clock tree, and recursively bipartitioning using results of
32 the clustering. The foregoing embodiment wherein the clustering is in accordance with a
33 structured array fabric integrated circuit.

1 Clock Tree Synthesis for SAF Exemplary Illustrative Combinations

2
3 **[0090]** EC1) A method comprising: synthesizing a clock tree for a structured array
4 fabric; wherein clock buffer resources are at fixed locations of an integrated circuit die.

5
6 **[0091]** EC2) The method of EC1, wherein clock buffers are restrained in at least one of
7 number and location.

8
9 **[0092]** EC3) The method of EC1, wherein the fixed locations correspond to sites of the
10 structured array fabric.

11
12 **[0093]** EC4) The method of EC1 further comprising automatic deskewing of multi-
13 level clock trees.

14
15 **[0094]** EC5) The method of EC4, wherein the multi-level clock trees comprise at least
16 one of clock gating structures, clock select hardware, and derived clock nets.

17
18 **[0095]** EC6) The method of any of EC1 to EC4, wherein the clock tree is implemented
19 on an integrated circuit having at least one of an embedded intellectual property block and a
20 memory array block.

21
22 **[0096]** EC7) The method of EC6, wherein the embedded intellectual property block
23 and the memory array block lack slots compatible with clock buffer placement.

24
25 **[0097]** EC8) The method of any of EC1 to EC7 further comprising synthesizing the
26 clock tree while accounting for user specified timing and/or constraints.

27
28 **[0098]** EC9) The method of EC8, wherein the constraints are compatible with
29 customized timing and placement tools.

30
31 **[0099]** EC10) The method of EC8, wherein the constraints are compatible with
32 industry-standard timing and placement tools.

33
34 **[0100]** EC11) The method of any of EC1 to EC7 further comprising using buffers to
35 create meandering buffer paths to deskew a clock tree.

- 1 **[0101]** EC12) The method of any of EC1 to EC7 further comprising using morphable
2 logic resources to create meandering buffer paths to deskew a clock tree.
3
- 4 **[0102]** EC13) The method of any of EC1 to EC7 further comprising automatic
5 processing of derived clock sub-domains.
6
- 7 **[0103]** EC14) The method of any of EC1 to EC7 further comprising iterative clock tree
8 optimization based on post-layout parasitic delays.
9
- 10 **[0104]** EC15) The method of any of EC1 to EC7 further comprising simultaneous clock
11 tree synthesis for multiple timing libraries corresponding to respective
12 process/temperature/voltage corners.

HIGH-LEVEL INTEGRATED CIRCUIT PHYSICAL DESIGN FLOW

[0105] Fig. 1 is a flow diagram illustrating selected details of an embodiment of placing, routing, analyzing, and generating fabrication data for any portion of an integrated circuit according to an SDI-based flow. A representation of all or any portion of the integrated circuit is provided (“Design Description” **120**), in certain embodiments including a gate-level netlist, placement constraints, timing requirements, and other associated design specific data. The gate-level netlist may be provided in any proprietary or standard format, or a hardware description language (such as Verilog).

[0106] A representation of fabrication flow is also provided (“Technology Description” **121**), in certain embodiments including information relating to fabrication material starting state and manufacturing flow. The fabrication material information may include data describing wafers and any associated predetermined processing on the wafers (for example fabrication of lower layers of devices). The predetermined processing may be associated with transistors, combinational logic gates, sequential logic devices, storage arrays, regular structures, power distribution, clock distribution, routing elements, and other similar portions of active and passive circuitry. The manufacturing flow information may include information relating to physical and electrical design rules and parameters for extraction of parasitic information for analyzing results during physical design flow processing.

[0107] Flow begins (“Start” **101**) and continues (“Pre-Process” **102**), where the design and technology descriptions are parsed and various design-specific data structures are created for subsequent use. The design description in certain embodiments includes a gate-level netlist describing interconnections of devices (morphable-devices, according to some embodiments), as well as constraints specific to implementation of the design (such as timing and placement requirements). The technology description includes information such as library definitions, fabrication technology attributes, and descriptions of manufacturing starting material (for example data describing SAF tile arrangement and composition of active and passive elements).

[0108] Physical locations of some or all of the devices are then determined (“SDI Place & Route” **103**), i.e. the design is placed, and wiring according to the netlist is determined (i.e. the design is routed). Place and route processing in certain embodiments includes multiple iterations of one or more internal processes (see “Place and Route Flow”, elsewhere herein). The placed and routed design is then analyzed (“Result Analysis” **104**), in certain embodiments with one or more analysis tools performing various functions such as parasitic extraction, timing

1 verification, physical and electrical rule checking, and Layout-Versus-Schematic (LVS) formal
2 verification.

3
4 **[0109]** Results of the analysis are examined by any combination of automatic (such as
5 software) and manual (such as human inspection) techniques (“OK?” **105**). If the results are
6 acceptable, then flow continues (“Yes” **105Y**) to produce information to manufacture the design
7 according to the results (“Generate Fabrication Data” **106**). The fabrication data varies by
8 embodiment and design flow context, and may include any combination of mask describing
9 data, FPGA switching-block programming data, and FPGA fuse/anti-fuse mapping and
10 programming data. Processing is then complete (“End” **199**).

11
12 **[0110]** If the results are not acceptable, then flow loops back (“No” **105N**) to repeat
13 some portion of the place and route operations. In some usage scenarios (not illustrated) one or
14 more modifications to any combination of the design and the technology may be made before
15 repeating some of the place and route operations. For example, synthesis may be repeated (with
16 any combination of changes to functionality as specified by behavioral or gate-level inputs and
17 synthesis commands), a different technology may be chosen (such as a technology having more
18 metal layers), or a different starting material may be selected (such as choosing a “larger”
19 structured array having more SAF tiles).

20
21 **[0111]** Processing functions (“Pre-Process” **102**, “SDI Place & Route” **103**, “Result
22 Analysis” **104**, “OK?” **105**, and “Generate Fabrication Data” **106**) are responsive to various
23 instructions and input data (“Commands and Parameters” **130**), according to various
24 embodiments. The effects of the commands and parameters on the processing are represented
25 conceptually in the figure (arrows **102C**, **103C**, **104C**, **105C**, and **106C**, respectively). In
26 various embodiments information is communicated between the processing functions (and other
27 processing elements not illustrated) in various forms and representations, as shown conceptually
28 (“Working Data” **131** and associated arrows **102D**, **103D**, **104D**, and **106D**, respectively). The
29 working data may reside in any combination of processor cache, system memory, and non-
30 volatile storage (such as disks), according to implementation and processing phase.

31
32 **[0112]** The illustrated placement, route, and analysis processing is applied, in various
33 embodiments, to integrated circuits implemented in various design flows or contexts, including
34 application specific, structured array (homogenous and heterogeneous varieties), mask-definable
35 gate array, mask-programmable gate array, Field-Programmable Gate Array (FPGA), and full
36 custom. The processing may be applied to an entire integrated circuit, or one or more portions

1 or sub-sections of an integrated circuit, according to various usage scenarios. For example, an
2 otherwise full custom integrated circuit may include one or more regions of standard cells, and
3 each of the standard cell regions may be processed according to all or portions of the illustration.
4 For another example, an Application Specific Integrated Circuit (ASIC) may include some
5 regions of standard cells and other regions of SAF tiles. Any combination of the standard cell
6 and SAF tile regions may be processed according to all or portions of the illustrated flow. These
7 and all similar variations are contemplated.

10 PLACE AND ROUTE FLOW

12 **[0113]** Fig. 2 is a flow diagram illustrating selected details of an embodiment of placing
13 and routing any portion of an integrated circuit, according to an SDI-based flow, such as
14 operations referred to elsewhere herein (“SDI Place & Route” **103**, of Fig. 1, for example).
15 Overall the flow includes determining approximate (i.e. subject to subsequent refinement)
16 locations for devices, reconciling resources, determining nearly final locations and
17 implementations for the devices, minimizing critical delay paths, and wiring the devices
18 according to a netlist. In certain embodiments each of the elements of the flow includes internal
19 functions to determine acceptability of results, iterate as necessary to improve the results, and to
20 direct feedback to earlier processing functions of the flow as needed.

22 **[0114]** Processing begins (“Start” **201**), in certain embodiments by receiving one or
23 more data structures and files describing a netlist having devices and associated connectivity,
24 along with manufacturing technology information. The structures and files may result from
25 parsing design and technology information (“Pre-Process” **102**, of Fig. 1, for example).
26 Approximate locations for the devices of the netlist are then determined (“SDI Global
27 Placement” **202**) according to the netlist, the technology, and commands/parameters (such as
28 those from “Commands and Parameters” **130**, of Fig. 1). If global placement results are
29 acceptable (i.e. suitable as a starting point for further processing), then flow proceeds (“OK”
30 **202Y**). If the global placement results are not acceptable, then flow loops back (“Not OK”
31 **202N**, “Repeat” **220**, and “Revise” **202R**) to repeat all or portions of the global placement.
32 Revised global placement processing (via “Revise” **202R**) in certain embodiments includes
33 modifying any combination of the netlist, global placement commands and parameters, and
34 manufacturing technology (such as specifying a larger die, or a denser device fabrication
35 process) based in part upon previous processing.

1 **[0115]** Subsequent to acceptable global placement, resources are reconciled according
2 to the global placement and manufacturing information (“Legalization” **203**), resulting in
3 elimination of areas of oversubscribed resources. In certain embodiments modifications are
4 made to the global placement results (effecting “movement” of placed elements) thus producing
5 a legalized placement. If legalization results are acceptable, then flow proceeds (“OK” **203Y**).
6 If the legalized placement is not acceptable (or not computed), then flow loops back for
7 additional processing (“Not OK” **203N**). In certain embodiments the additional processing is
8 based on previous processing, and may include repeating any portion of global placement
9 (“Revise” **202R** via “Repeat” **220**) and continuing onward, or repeating any portion of
10 legalization (“Revise” **203R** via “Repeat” **220**), according to various usage scenarios and
11 embodiments.

12
13 **[0116]** After acceptable legalization, then nearly final (or “exact”) locations and
14 implementations for the devices are determined (“(SDI) Detailed Placement” **204**). Relatively
15 small-scale adjustments are made to legalization results, via any combination of placed element
16 movement and placed element implementation, according to embodiment. In certain structured
17 array embodiments, the placed element implementation includes morphing of selected devices to
18 functionally equivalent alternatives. If detailed placement results are acceptable, then flow
19 proceeds (“OK” **204Y**). If the detailed placement is not acceptable (or not computed), then flow
20 loops back for additional processing (“Not OK” **204N**). In certain embodiments the additional
21 processing is based in part upon previous processing, and may include repeating any portion of
22 previous place and route functions and then continuing onward (such as via any of “Revise”
23 **204R**, “Revise” **203R**, and “Revise” **202R** by way of “Repeat” **220**).

24
25 **[0117]** Subsequent to detailed placement, delay paths are minimized (“Timing Closure”
26 **205**), in certain embodiments to meet user specified timing, in various ways according to
27 embodiment and/or user option or configuration. In certain embodiments the detailed placement
28 is analyzed and buffers (or buffer trees) are inserted in high fanout and timing-critical nets. In
29 some embodiments drivers are resized and optimized to meet maximum capacitance and/or
30 required time constraints with respect to timing critical receivers. In some embodiments clock
31 networks are synthesized, while in other embodiments the clock networks are predefined. In
32 either case the appropriate clock network elements are inserted into the netlist for clock
33 distribution and to meet clock skew constraints. Further according to embodiment and/or user
34 option or configuration, other timing closure driven optimizations are performed (see “Timing
35 Closure”, elsewhere herein). If the timing closure results are acceptable, then flow proceeds
36 (“OK” **205Y**). If the timing closure is not acceptable, then flow loops back for additional

processing (“Not OK” **205N**). The additional processing may include repeating any portion of previous place and route functions, based in part upon previous processing and then continuing onward (such as via any of “Revise” **205R**, “Revise” **204R**, “Revise” **203R**, and “Revise” **202R** by way of “Repeat” **220**). Note that in some embodiments flow loops back as a natural consequence of timing closure processing, rather than merely as a result of not-acceptable timing closure results. For example, certain timing closure techniques call for repetition of previous processing (such as one or more of “SDI Global Placement” **202**, “Legalization” **203**, and “(SDI) Detailed Placement” **204**), using various combinations of modified behaviors and parameters, along with optional changes to the netlist and constraints, according to various embodiments.

[0118] After timing closure is complete (or considered “close enough”), the resultant devices are wired together according to the resultant netlist (“Routing” **206**), and corresponding interconnect is generated. If the routing results are acceptable, then flow proceeds (“OK” **206Y**). Place and route processing is then complete (“End” **299**), and results are available for further use, such as any combination of analysis and mask generation (“Generate Fabrication Data” **106** of Fig. 1, for example). If the routing results are not acceptable, then flow loops back for additional processing (“Not OK” **206N**). In certain embodiments the additional processing is based in part upon previous processing, and may include repeating any portion of previous place and route functions and then continuing onward (such as via any of “Revise” **206R**, “Revise” **205R**, “Revise” **204R**, “Revise” **203R**, and “Revise” **202R** by way of “Repeat” **220**).

[0119] Various combinations of place and route processing functions (such as “SDI Global Placement” **202**, “Legalization” **203**, “(SDI) Detailed Placement” **204**, “Timing Closure” **205**, and “Routing” **206**) may include reading and writing shared information (such as references to “Working Data” **131**, of Fig. 1). Examples of working data include netlists, constraints, progress indicators, and other similar shared processing items. Various combinations of the aforementioned place and route processing functions also may include receiving one or more inputs specifying requested behaviors or processing (such as information from “Commands and Parameters” **130**, of Fig. 1). Examples of commands and parameters include scripts specifying iteration closure conditions, control parameters, goal descriptions, and other similar information to guide processing. The commands and parameters may be provided via any combination of scripts, command line inputs, and graphical user interfaces, according to various embodiments.

[0120] In some embodiments processing of one or more elements of Fig. 2 is optional, or performed only for selected iterations though the illustrated flow. For example, timing

closure operations may be operative in a first processing mode where legalization and detailed placement are skipped, and processing relating to timing closure is partially performed as part of global placement. Alternatively the first processing mode may be viewed as global placement operations being performed to a limited extent, then analyzed and further directed by timing closure operations (without legalization or direct placement), and then additional global placement operations being performed. Eventually a second mode of processing may be entered where legalization and detailed placement are performed, optionally followed by additional timing closure operating as in the first mode or operating in a manner specifically tailored to the second mode (see "Timing Closure", elsewhere herein).

SIMULTANEOUS DYNAMICAL INTEGRATION (SDI) DIRECTED GLOBAL PLACEMENT

[0121] Conceptually SDI may be understood as modeling each individual device of the netlist as a node, or point particle, having an associated mass, position (or location), and velocity. The nodes representing the devices of the netlist are coupled by and interact with each other via attractive and spreading forces. The forces may include attractive forces representing electrical connections between the devices (as specified by the netlist), and spreading forces modeling resource requirements versus availability (such as a density of logic gates needed versus a density of logic gates on hand). The nodes and effects of the coupling forces are simulated as evolving over time as governed by a system of coupled ordinary differential equations using continuous variables, according to classical Newtonian mechanics (i.e. force equals mass multiplied by acceleration, or $F=ma$). Thus locations of nodes (corresponding to device placements) evolve over time from initial positions to subsequent positions (corresponding eventually to the global placement result for the devices).

[0122] More specifically, the independent variables in the dynamical system simulation include configuration-space variables (position and velocity) of the nodes. In certain embodiments the position and velocity representations are multi-dimensional quantities (two or three dimensions, for example), according to usage scenario and embodiment. Force terms in the coupled equations of motion are related to any combination of the topology of the connections of the devices, timing analysis of evolving device locations (placement), obstructions, and region constraints (fixed and floating), according to embodiment. Force terms may also be related to any combination of partial node density, partial resource usage density, viscous damping, energetic pumping, interconnect congestion effect modeling, power or clock

1 distribution, and signal integrity representation, according to embodiment. Force terms may
2 include any function of the independent variables, provided commands and parameters, and
3 other similar mathematical devices useful in managing numerical behavior of continuous time
4 integration of the system of nodes and forces.

5
6 **[0123]** In certain embodiments the obstructions are represented as exclusion zones, and
7 arise as a result of architectural considerations, location-fixed (or predetermined) blocks (such as
8 large RAM arrays or IP elements), and other similar placement limiting conditions. In certain
9 embodiments the region constraints are represented as fixed, relative, or floating location
10 requirements on selected devices of the netlist. Corresponding position requirements (such as an
11 initial position with no subsequent change during system simulation time) are imposed for the
12 corresponding nodes in the dynamical simulation. Various combinations of region constraints
13 (relating to integrated circuit floorplan specifications, for example) may be developed by any
14 combination of automatic techniques (by software, for example) and manual techniques (by
15 users), according to usage scenarios and embodiments.

16
17 **[0124]** Conceptually the system of coupled simultaneous differential equations is
18 operational in continuous variables. While it is envisioned that certain embodiments will
19 perform at least some of the integration according to true analog integration techniques, in which
20 the state variables are actually continuous, in digital computer embodiments, the integration is
21 performed using digital integration techniques. Digital computers are limited to representing all
22 quanta with finite-precision variables and that continuous time integration may be implemented
23 on digital computers using "pseudo-continuous" numerical approximation techniques, a.k.a.
24 "numerical methods." Even when implemented using finite-precision approximations, the
25 "continuous variables" abstraction is a useful way to conceive and describe some of the
26 techniques described herein and to distinguish compared to other approaches using conceptually
27 discrete variables. Thus the term continuous as used throughout this disclosure should be
28 interpreted in accordance with the foregoing.

29
30 **[0125]** In digital computer embodiments, continuous state variables (including those
31 variables representing simulation time, mass, location, and velocity) are approximated as any
32 combination of single, double, or extended floating-point numbers. The continuous time
33 integration of the simultaneous coupled dynamical governing equations may be performed in
34 digital computer embodiments by any suitable digital integration technique, such as Runge-
35 Kutta, predictor-corrector, leap-frog, and any similar technique adaptable to continuous multi-
36 variable state space integration. In some embodiments the integration technique is chosen for

1 suitability based at least in part on adaptability to parallel processing (see “Computer System
2 Executing SDI-Directed EDA Routines”, elsewhere herein).

3
4 **[0126]** The forces acting in the system provide coupling between the nodes and act to
5 accelerate the nodes over time, resulting in movement of the nodes throughout the state-space
6 over time. A set of attractive forces (known as “net attractive forces”) is modeled to represent
7 connectivity between the devices of the netlist, or more specifically between pins (i.e. terminals
8 of circuit elements) of devices. In some embodiments the net attractive forces are modeled as
9 individual springs between a pin of one device and a pin of another device, with every
10 interconnection between any two pins being modeled as a corresponding spring. Force
11 associated with each spring is computed according to Hooke’s law (force is proportional to
12 distance between the pins). The net attractive force acting on each device is a vector sum of all
13 net attractive forces acting on all of the pins of the respective device.

14
15 **[0127]** In some embodiments the constant of proportionality used to calculate spring
16 force is identical for all springs. In some embodiments the constant of proportionality is
17 dependent on the fanout of a net (i.e. the number of pins connected together). In some
18 embodiments relatively high fanout nets are considered to be one or more drivers providing a
19 signal to one or more loads. Springs between the loads of the relatively high fanout nets are
20 eliminated (while springs from drivers to loads are retained). In some embodiments springs
21 between drivers and loads have a different constant of proportionality than other springs.
22 Modeling of net attractive forces is not restricted to ideal springs, and may instead be based on a
23 general linear or non-linear force model, according to various embodiments.

24
25 **[0128]** A set of spreading forces (known as “spatial spreading forces”) is modeled
26 based on one or more macroscopic density fields. In certain embodiments the density fields are
27 computed based on analysis of metrics associated with respective devices corresponding to the
28 nodes (and their locations) in the dynamical system. The metrics may include any combination
29 of standard cell area (in, for example, standard cell flow processing), fabric resource
30 consumption (in, for example, SAF flow processing), equivalent gate count, and other similar
31 functions of node properties. In some embodiments the spatial spreading forces (see “Field-
32 Based Force Components”, elsewhere herein) are with respect to a density field based on
33 resource utilization of corresponding nodes in a local region. In some embodiments resource
34 utilization may be evaluated using an area averaging or summation of nearby devices or an
35 equivalent-gate count rating (cost function) of spatially close devices.

1 **[0129]** In some embodiments a plurality of density fields are computed with respect to
2 a plurality of metrics. In some embodiments any combination of first, second, and third density
3 fields are computed with respect to first, second, and third categories of logic devices (such as
4 combinational logic devices, sequential logic devices, and total logic devices). In some
5 embodiments each of a plurality of partial density fields is computed according to a set of
6 respective non-interchangeable morphing classes (such as combinational and sequential
7 morphing classes) associated with an underlying SAF. In some embodiments (such as selected
8 standard cell based design flows) the density fields are computed based wholly or partially on
9 device area. In some embodiments (such as selected structured array based design flows) the
10 density fields are computed based wholly or partially on resource utilization as measured by
11 counts of the number of each type of resource needed to implement the function associated with
12 each device in the netlist.

13

14 **[0130]** Other attractive and spreading forces may also be included, according to usage
15 scenario and embodiment. Floorplan constraints, or various region constraints, may be
16 expressed as attractive or spreading forces, or as potential wells (with a tendency to retain nodes
17 in a region) or potential barriers (with a tendency to disperse nodes from a region), according to
18 usage scenario and embodiment. For example, boundaries of a die, or locations of input/output
19 (IO) rings may be expressed as fixed constraints that are mapped to attractive forces acting on
20 nodes having interconnect to the IO ring. For another example, a selected region of the die may
21 be excluded from use (such as for yield improvement or noise reduction) by fixed or relative (i.e.
22 floating) constraints that are mapped to spreading forces acting on nearby or all nodes (see
23 “Exclusion Zones”, elsewhere herein). In other embodiments or modes of operation, such
24 floorplan constraints may be implemented through coordinate clipping inside the integrator,
25 thereby preventing the motion of devices into disallowed regions.

26

27 **[0131]** User specified circuit timing constraints may warrant that certain pins in the
28 netlist be moved closer together to improve the performance of the design. A corresponding set
29 of attractive forces between drivers and select loads is fed into the system as attractive forces
30 with configurable binding strength.

Viscous Damping

[0132] Forces other than attractive and spreading forces between nodes or other elements may also be accounted for. As an example, a viscous damping force may be included as a way to (a) compensate for the effect of numerical errors (potentially incurred by the time integration and spatial differencing techniques used) contributing toward numerical heating, and (b) change the ratio between kinetic and potential energy of the node distribution. The damping serves to decelerate the ballistic motion of a node. One embodiment of such a force on a given node is a term proportional to the negative of the node velocity, with the proportionality constant being equal to μ , the global coefficient of viscosity. The value of μ may be supplied by direct manual input (by a user) or via automatic control, (under software control) according to embodiment, to provide partial control of the node distribution as a whole.

[0133] While μ is a global constant, it may have a local effect, and thus in some embodiments other parameters are selected for manipulation to provide control of the node distribution as a whole. For example, in some implementations a ratio of KE/TE, where KE is the kinetic energy of the node distribution and TE is the total energy of the system, is a convenient control parameter.

[0134] In some embodiments, the global viscosity coefficient is split into two terms, a gradually drifting term and a dynamically calculated term. The gradually drifting term enables the system to gradually adapt to time varying forces or parameter changes, while the dynamical term prevents runaway acceleration on a per-timestep basis.

[0135] Each timestep the total effective μ is adjusted in response to normalized kinetic energy (KE/TE) changes from a selected target value. In certain embodiments the adjustment to μ is given by:

If KE/TE > target then:

$$dm = cdm1 * ((KE/TE / target) - 1) + cdm2 * ((KE/TE / target) - 10)^2$$

$$\mu_{eff} = \mu * (1 + dm)$$

$$\mu *= (1 + \text{<small adjustment>})$$

If KE/TE < target then:

$$dm = cdm1 * ((target / KE/TE) - 1) + cdm2 * ((target / KE/TE) - 10)^2$$

$$\mu_{eff} = \mu / (1 + dm)$$

$$\mu /= (1 + \text{<small adjustment>})$$

1 where:

2 double mu_max = 1.e+8;

3 double cdm1 = 1.;

4 double cdm2 = 0.01; and

5 Note that "double" refers to double-precision variables used in some embodiments.

6

7 **[0136]** The <small adjustment> may vary with the relative difference between the
8 target and actual values of KE/TE, and tends to be small compared to 1. The term "mu_max"
9 limits μ to prevent numerical problems with a timestepper used for numerical integration. The
10 quadratic term contributes little until KE/TE differs from the target by a factor of 10, and
11 quenches runaway conditions.

12

13 **[0137]** By splitting the calculation of μ into a purely dynamical term and a slowly
14 varying term, the system remains generally stable while retaining an ability to react quickly to
15 energy spikes. Further, by using a constant μ during the course of the time integration,
16 performance may be enhanced, as operation counts are substantially reduced and adaptive
17 integrator timesteps (if relevant) may be allowed to increase.

18

19 **[0138]** In some embodiments a viscous damping proportionality constant is identical
20 for all nodes in the system, while in other embodiments one or more distinct proportionality
21 constants may be employed. For example, in certain embodiments the viscous damping
22 proportionality constant is modeled as a scalar field of position and the value of the constant at
23 the position of each circuit device is computed. Moreover, in certain embodiments the scalar
24 field is analytically specified, and selectively includes a dependence upon the independent time
25 variable. In other embodiments the scalar field is a derived quantity computed from other
26 numerical characteristics that may be evaluated for the time-evolving simulation. Additionally,
27 the viscous force is not limited to being proportional to the velocity of a node. In certain
28 embodiments the viscous force instead follows a functional form based on other selected state of
29 the system.

30

31 **[0139]** The aforementioned forces are merely representative examples. Forcing terms
32 may be associated with interactions between one or more nodes, and between one or more fixed
33 (or immovable) elements. Forcing terms may also be associated with fields that may be
34 dependent in some way upon one or more nodes, or with fields that are independent of nodes.
35 These and all similar types of forcing terms are contemplated in various embodiments.

36

1 **[0140]** Thus forces on the nodes of the system include direct interactions with
 2 topological neighbors (according to the netlist), collective interactions involving numerical
 3 constructs associated with temporal bulk properties of the node distribution, and with
 4 architectural features of the implementation. The result of the combination forces impinging on
 5 the system nodes is a complex dynamical interaction where individual nodes meander through
 6 the placement domain under the influence of the forces and wherein the forces vary continuously
 7 with the motion of all nodes in the netlist. The motion exhibits both chaotic and coherent
 8 behaviors. The motion of a given node may appear chaotic in the sense that the node trajectory
 9 may meander back and forth as a result of connections to other nodes. Yet the system may also
 10 exhibit coherent (or collective) motion in the sense that tightly connected nodes will tend to
 11 move in bulk and remain in proximity to topological neighbors even as the tightly connected
 12 nodes collectively move far from respective starting points.

13
 14 **[0141]** The integration of the governing equations of motion proceeds using standard
 15 techniques of numerical integration. (See for example, a reference describing numerical
 16 integration.) As an example, the next several paragraphs assume the use of a Runge-Kutta
 17 integrator.

18
 19 **[0142]** The computation of the forcing terms is referred to as "computing the
 20 derivatives". Differentiation is denoted with respect to time by ' (prime), so that $dx/dt = x'$,
 21 $d^2x/dt^2 = x''$, and so forth. The following variables are introduced to set up the governing
 22 equations for solution by numerical integration:

23 $v_{x,i} = (x_i)'$

24 $v_{y,i} = (y_i)'$

25 The subset of the system of equations relating to the i th node (for a two-dimensional layout
 26 application) is:

27 $(x_i)' = v_{x,i}$

28 $(y_i)' = v_{y,i}$

29 $(v_{x,i})' = F_{x,i}$

30 $(v_{y,i})' = F_{y,i}$

31
 32 **[0143]** Thus the system of simultaneous second order differential equations is
 33 transformed to a (larger) system of simultaneous first order differential equations, where the
 34 right hand side of each equation is the derivative of the respective left hand side. Conceptually
 35 computation of a derivative per se is not required (unless some element of the forcing terms is

1 itself expressed as a derivative of something else), but rather the right hand sides of the
2 equations are the derivatives.

3
4 **[0144]** There is time-varying complexity in the behavior (character of motion) of the
5 moveable nodes in the netlist when the forcing terms are time varying. In some embodiments a
6 time varying timestep is used to preserve numerical accuracy and to continue processing until
7 convergence criteria (error limits) are met during each timestep in the integration.

8
9 **[0145]** The integrator accepts as input a specification of a desired timestep, and then
10 processes the timestep in two ways: once directly, and once as two half-steps. If the results are
11 not close enough as determined by a specifiable error-norm, then the target timestep is reduced
12 until it is possible to perform the one-step plus the two-half-steps approaches with results within
13 an error norm. Besides new coordinate values for the independent variables, the integrator also
14 returns the length of the timestep just taken and the advised length for the next timestep. Thus
15 during periods of laminar motion when numerical convergence is readily achieved, the timestep
16 trends longer on successive calls to the integrator. But in periods of turbulent or chaotic motion,
17 where convergence requires more effort, the timesteps become as small as needed to ensure the
18 accuracy of the integration.

19
20 **[0146]** Fig. 3A is a flow diagram illustrating selected details of an embodiment of
21 global placement according to SDI modeling and simulation, such as operations referred to
22 elsewhere herein ("SDI Global Placement" **202**, of Fig. 2, for example). Overall the flow
23 includes various functions to enable and perform a series of dynamical simulations based on
24 Newtonian mechanics on a system representing the netlist and associated design constraints and
25 targets. The simulations use SDI techniques to orchestrate the interactions between particles
26 (representing netlist devices). The SDI techniques make use of fields that are calculated as
27 functions of the particle positions. The functions include determining a set of nodes
28 corresponding to the devices in the netlist, initialization of state variables (including mass,
29 location, and velocity associated with each node), adjusting forces, and evolving the resultant
30 system of simultaneous dynamical governing equations forward in time via integration. The
31 flow is repeated beginning at the adjustment processing until a suitable result is available, or it is
32 determined that a suitable result will not become available without further processing outside of
33 the illustrated flow.

34
35 **[0147]** Processing begins ("Start" **301**) with receipt of pre-processed information, in
36 certain embodiments as data structures representing the netlist and the associated devices and

connectivity (“Pre-Process” **102**, of Fig. 1, for example). Further data structures for representing a system of nodes and forces are created and initialized (“Determine Nodes and Forces” **302**), with each node in the system corresponding one-to-one with each device of the netlist, and with each node having a corresponding set of forces acting on it. State variables for the dynamical simulation are initialized (“Initialize State Variables” **303**), including determining starting values for mass, location, and velocity state variables for each node. The initial node locations correspond to initial placements of the corresponding netlist devices (see “Initial Placement”, elsewhere herein). Initial force values are also determined.

[0148] Large-scale goal-driven modifications to the forces in the system are then made (“Macro Adjust Forces” **304**). In some embodiments one or more attractive forces are over- or under-weighted for periods of time, and one or more spreading forces may also be reduced or increased in relative proportion to the attractive forces. For example, a “condensing” phase may inflate attractive forces and deflate spreading forces, and an “extending” phase may deflate attractive forces and inflate spreading forces. Operations associated with the macroscopic force adjustment track simulation time and change the forces according to condensing and extending phases. During the phases of system evolution, the coordinates of individual nodes continue to evolve separately based on the governing equations for each individual node. Consequently, the behavior of any individual node may vary from the bulk behavior of the collective system.

[0149] Other large-scale force adjustments may also be made, according to embodiment, including entirely removing one or more forces for a period of simulation time, and introducing a new force. The removal (or introduction) of a force may be at a predetermined point in simulation time, at a point in simulation time determined by computation of a test condition, any similar mechanism, and/or at the discretion of a human operator of the system, according to various embodiments. In certain embodiments the removal (or introduction) of a force is gradual, and the rate of change of the removal (or introduction) may vary over simulation time or be constant, according to implementation. In some embodiments the macroscopic force adjustments are in response to various force-control instructions and input data (such as represented conceptually by “Commands and Parameters” **130**, of Fig. 1).

[0150] Large-scale goal-driven modifications to the effects of masses in the system are then made (“Macro Adjust Masses” **305**). In certain embodiments the effects of masses are modified during phases where node densities are being adjusted to more evenly distribute resource consumption, or to more evenly match resources needed with resources available. For example, in usage scenarios including global placement of devices according to SAF tiles,

1 macroscopic mass adjustments may be made to “encourage” cells in over-subscribed regions to
2 “move” to less subscribed regions (see “Depletion Weighting”, located elsewhere herein). As in
3 the case of macroscopic force adjustments, macroscopic mass adjustments may be varied
4 according to simulation time phase, and may be gradually introduced (or removed) over the
5 course of system evolution throughout simulation time. In some embodiments the macroscopic
6 mass effect adjustments are in response to various mass-control instructions and input data (such
7 as represented conceptually by “Commands and Parameters” 130, of Fig. 1). Note that adjusting
8 the effects of mass, in certain embodiments, is with respect to densities and forces brought about
9 by the masses, while the momentum of each of the nodes having adjusted mass effects remains
10 unchanged.

11
12 [0151] A dynamical simulation of the nodes (as point particles) according to the mass,
13 location, velocity, force, and other state variables is performed (“SDI Simulation” 306) for some
14 amount of system simulation time. The time may be a predetermined interval, dependent on
15 specific completion criteria (as provided to the SDI simulation), and any similar interval
16 specification scheme, according to various embodiments. At the end of the simulation time the
17 system arrives at a new state. In certain embodiments the new state includes new locations for
18 one or more of the nodes, and the new locations of the nodes are interpreted as corresponding to
19 new locations for the devices being placed.

20
21 [0152] According to various embodiments, any combination of the system variables
22 (including simulation time and node mass, location, and velocity) and corresponding
23 interpretations of the system variables in the context of the netlist (including device location and
24 density) are examined to determine if portions of the flow should be repeated (“Repeat?” 307) or
25 if flow is complete (“OK Result?” 308). If repeating the flow would likely improve results, and
26 no other end condition has been met, then flow loops back (“Yes” 307Y) to macro adjustment of
27 selected forces and masses. In some embodiments configurable settings are adjusted prior to or
28 in conjunction with force and mass macro adjustments (such as settings associated with
29 “Commands and Parameters” 130, of Fig. 1). If the global placement is close enough (“No”
30 307N), then flow is complete (“OK” 202Y) and processing continues to legalization (see Fig. 2).
31 If there would likely be no benefit in iterating the global placement (“No” 307N), and the results
32 are not acceptable, then flow is also complete (“Not OK” 202N), but subsequent processing then
33 includes one or more revisions (see Fig. 2).

34
35 [0153] Tests to determine if the flow is to be repeated may be made for a predetermined
36 end condition, a predetermined rate of change, other similar criteria, and any combination

1 thereof according to assorted implementations. In some embodiments the flow is not repeated
2 even if improvement is likely possible (for example if an interval of simulation time has
3 expired).

4
5 **[0154]** Determinations (“Repeat?” **307** and “OK Result?” **308**) are according to any
6 combination of automatic (software program) and manual (human user) techniques, according to
7 various embodiments. For example, an automatic technique may include software determining
8 if the most recent iteration is a significant improvement over a previous iteration. If so, then
9 repeating the flow is beneficial. As another example, a manual technique may include a user
10 observing the time-evolving locations of devices and noticing that further improvements are
11 possible and that repeating the flow would be beneficial. Another manual technique may
12 include a user determining that the placement as changing over time is “stuck”, perhaps due to
13 some incorrectly specified constraints, and that additional iterations of the global placement flow
14 are not likely to be beneficial unless modifications are made to the constraints.

15
16 **[0155]** Any portion (or all) of global placement may be performed according to various
17 techniques, in addition to the aforementioned SDI directed technique. The additional techniques
18 include simulated annealing, objective minimization techniques such as conjugate-gradient,
19 chaotic processing, and other similar mechanisms to provide approximate or “close enough”
20 device coordinates, according to various embodiments.

21 22 23 Initial Placement

24
25 **[0156]** Fig. 3B is a flow diagram illustrating selected details of an embodiment of
26 initial placement operations for global placement, such as selected operations performed while
27 initializing state variables (as in “Initialize State Variables” **303** of Fig. 3A). Processing begins
28 (“Start” **310**) and then one of a plurality of starting location definition techniques is chosen
29 (“Select Technique” **310A**), based, in some embodiments, on instructions provided by a user
30 (such as information from “Commands and Parameters” **130**, of Fig. 1). A first technique
31 determines an initial placement based on a placement performed in the past (“Prior Solution”
32 **311**). A second technique formulates an initial placement based on randomization (“Random”
33 **312**). A third technique develops an initial placement according to any of a number of other
34 mechanisms (“Selected Algorithm” **313**), chosen by any combination of software and user input.
35 The chosen technique is then performed and processing is complete (“End” **314**).
36

Mass Determination

[0157] In some embodiments, determination of mass (as in “Determine Nodes and Forces” **302**, for example) is dependent on the design flow or implementation context (such as application specific, structured array, mask-definable gate array, mask-programmable gate array, FPGA, and full custom). For example, in a standard cell context, the mass of a node may be computed as a function (such as a linear function) of area occupied by the corresponding device in the netlist. For another example, in a structured array context, the mass of a node may be computed with respect to consumption of resources provided by the structured array, or with respect to local availability or scarcity of the resources, according to the corresponding device as implemented by the resources. For another example, in an FPGA context, the mass of a node may be computed according to consumption of Look Up Table (LUT) resources, or similar switching and/or routing resources.

[0158] In some embodiments the spatial spreading forces (see “Field-Based Force Components”, located elsewhere herein) are with respect to a density field based on resource utilization (such as an area averaging or summation of nearby devices or an equivalent-gate count cost function of spatially close devices) of corresponding nodes in a local region. In some embodiments first and second density fields are computed with respect to first and second categories of logic devices (such as combinational logic devices and sequential logic devices).

Field-Based Force Components

[0159] In some embodiments various elements of the spatial spreading forces are with respect to one or more resource usage based density fields, or other types of density fields. The density fields are managed independently, and may include any combination of all nodes, combinational nodes, and sequential nodes. Computation of density fields and resultant spreading forces conceptually includes calculating local densities according to a discrete grid, computing density fields, allocating field strengths according to the discrete grid to system nodes, and calculating resultant spatial spreading forces acting on the system nodes. In some embodiments the discrete grid is a uniform (or non-variable) grid, and in some embodiments the grid is a non-uniform (or variable) grid, the grid being implemented according to architectural considerations. Local density calculation includes summing resource usage computed in continuous spatial variables (i.e. node location and mass) according to the discrete grid and digitally filtering the resultant gridded scalar field. The local density calculation includes special

1 accounting for edges of the grid. The digital filter result is suitable for processing by a field
2 solver. Density field computation performed by the field solver includes determining density
3 fields (given density values on the grid) and digitally filtering the result. Allocating field
4 strengths includes interpolating field strengths to nodes (in continuous location space) while
5 accounting for edges of the grid. Repulsive (or spreading) forces are then computed according
6 to the allocated field strengths.

7
8 **[0160]** In some embodiments the grid is a unit grid, and the region enclosed by adjacent
9 grid lines is termed a “cell”. The grid may be two-dimensional (i.e. x and y) or the grid may be
10 three-dimensional (i.e. x, y, and z), according to implementation technology and other design-
11 flow related parameters. In some embodiments resource usage density is proportional to the
12 respective mass of each node, and the mass is in turn directly proportional to a “gate rating” that
13 is a measure of relative cost of implementing a logic function corresponding to the node. In
14 some embodiments the gate rating of the node is measured in “gate-equivalents” commonly
15 associated with design-flow device selection criteria.

16
17 **[0161]** Fig. 3C is a flow diagram illustrating selected details of an embodiment of
18 density field based force component computation, in a specific context of resource usage
19 densities expressed in certain embodiments as mass that is proportional to gate rating. The
20 operations of the flow are performed for each of a possible plurality of density fields, each field
21 having separate accounting. Flow begins (“Start” **330**), and proceeds to determine local resource
22 usage density by accumulating system node masses with respect to a scalar field organized as a
23 regular grid (in the illustrated embodiment) according to the SDI simulation spatial field
24 (“Accumulate Gate Densities” **331**). The grid is finite in size, completely covering space in the
25 system simulation corresponding to the area available for the devices of the netlist (either an
26 entire die or a portion thereof). The grid is extended, via one or more guard grid locations (or
27 grid cells) one or more units around each border of the area (the boundaries of the area) to more
28 accurately and efficiently model edge effects. The guard grid elements are then included in the
29 gate density calculation (“Fold Guard Cell Contributions” **332**). The single-unit guard-cell
30 buffer is used in some embodiments employing two and three-point allocation/interpolation
31 schemes, and a multi-unit guard-cell buffer is used in some embodiments having higher order
32 allocation schemes.

33
34 **[0162]** The resultant density values are then further optionally processed (“Digitally
35 Filter Density” **333**), according to embodiment, to smooth variations caused by grid element
36 representation inaccuracies. Density values for guard grid elements are then determined

1 (“Calculate Density Guard Cell Values” **334**) to enable straightforward and efficient field solver
2 implementations. Density field computations (“Solve Gate Fields” **335**) are then performed by
3 the field solver, determining the field value at each point as equal to minus the gradient at the
4 point (i.e. $\text{field} = -\text{Grad}(n)$). Any field solution technique applicable to calculating a derivative
5 with respect to a discrete grid may be used, such as a second order finite difference formula, or
6 any other suitable technique, according to embodiment. In some embodiments the second order
7 finite difference formula is given as the derivative at grid point “i”, and is equal to one-half the
8 quantity equal to the difference of the values at adjacent grid points along one of the orthogonal
9 dimensions (i.e. $\text{field}(i) = (\text{density}(i+1) - \text{density}(i-1)) / 2$). Derivatives are calculated for each
10 orthogonal dimension of the system node space (two or three dimensions, according to
11 embodiment). The result is a gridded vector field for each gridded density (such as all,
12 combinational, and sequential).

13
14 **[0163]** In some embodiments vector field values are stored in a data structure as a tuple.
15 Each member of the tuple corresponds to a value associated with an orthogonal dimension of the
16 vector field, and there is a tuple associated with each grid point. In some embodiments vector
17 field values are stored separately as scalar fields, according to each vector field orthogonal
18 component. Each respective scalar field represents all grid points. In some embodiments vector
19 field values are stored according to other arrangements that are mathematically equivalent to
20 tuples or scalar fields. In addition, vector fields may be stored in various combinations of tuple,
21 scalar field, and other forms, according to embodiment. The representation employed for the
22 vector fields may also change during processing to enable more efficient computations. Further,
23 during processing, any portion of vector field representations may be stored in any combination
24 of processor cache memory (or memories), processor main memory (or memories), and disk (or
25 other similar non-volatile long-term) storage, according to usage scenario and implementation.

26
27 **[0164]** The gridded vector fields are then processed according to a digital filter
28 (“Digitally Filter Fields” **336**). In some embodiments the filtering of the gridded vector fields is
29 according to operations identical, except for edge processing, to the smoothing performed on
30 density values (as in “Digitally Filter Density” **333**). The difference between the filter
31 operations is that for density filtering even parity is used when processing the boundaries, while
32 for field filtering even parity is used for field components parallel to the boundary and odd parity
33 is used for field components perpendicular to the boundary. The difference in parity accounts
34 for the differentiation operation performed between density and field domains, such that parity is
35 reversed from even (for density) to odd (for field) when differentiation is directed into a
36 boundary. For a (scalar) density, even parity means values associated with guard grid points are

1 added to interior grid points. For a (vector) field, even parity means the guard grid points are
2 equal to respective closest inner grid points for, and odd parity means that the guard grid points
3 are equal to the negative of respective closest inner grid points ("Calculate Field Guard Cell
4 Values" 337). Thus the average field directed into (or out of) a boundary vanishes at the
5 boundary. Assigning guard point field values enables subsequent efficient computation of field
6 values in the continuous location representation of nodes from the discrete field values
7 ("Interpolate Gate Fields to Nodes" 338). Corresponding forces may then be calculated
8 according to node field values and node masses. Processing is then complete ("End" 339).

9
10 [0165] Fig. 3D is a flow diagram illustrating selected details of an embodiment of gate
11 density accumulation, such as operations referred to elsewhere herein ("Accumulate Gate
12 Densities" 331, of Fig. 3C, for example). Conceptually mass associated with each node
13 (represented in continuous location space) is allocated to a local neighborhood portion of the
14 discrete grid points. Guard grid points are added around the boundary of the grid to efficiently
15 process edge conditions. In some embodiments a two-point linear spline, also known as a
16 Cloud-In-Cell (CIC) or area weighting technique, is used to allocate the mass of each node to
17 four neighboring grid points. In some embodiments a three-point spline technique is used to
18 allocate node mass to nine neighboring grid points.

19
20 [0166] More specifically, flow begins ("Start" 340) by initializing accumulation
21 variables (such as to zero), and then a check is made to determine if processing is complete for
22 all nodes in the simulated system ("Iterated Over All Nodes?" 341). If so, ("Yes" 341Y), then
23 gate field interpolation processing is complete ("End" 345). If not, then a first (and subsequently
24 a next) node is selected for processing, and flow continues ("No" 341N). Spline coefficients are
25 then determined for the node ("Determine Spline Weights" 342), based on distances from the
26 respective node to each field grid point (see the discussion of Fig. 3E, elsewhere herein).

27
28 [0167] After all of the spline weights for all of the grid points have been calculated, a
29 check is made to determine if all fields the respective node contributes to have been processed
30 ("Iterated Over all Fields" 343). If so ("Yes" 343Y), then processing loops back to check if all
31 nodes have been processed. If not, then a first (and subsequently a next) field is selected for
32 processing, and flow continues ("No" 343N). The effect of the node is then accumulated to the
33 respective field array at each of the grid points currently subject to interpolation ("Apply Node
34 Weight to Field Array" 344). Processing then loops back to determine if all fields have been
35 processed.

1 **[0168]** Fig. 3E is a conceptual diagram illustrating an embodiment of two-point
 2 interpolation of node mass to grid points, as performed during mass accumulation (such as
 3 “Determine Spline Weights” **342**, of Fig. 3D). Boundary **394** is shown to represent edges of the
 4 system simulation space (and corresponding edges of an integrated circuit region or die).
 5 Several points of the discrete grid are illustrated: interior point I_1 **381**, boundary points B_1 **371**,
 6 B_2 **372**, and B_3 **373**, and guard points G_1 **386**, G_2 **388**, and G_3 **389**. Mass from node N_1 **375** is
 7 shown accumulating to four grid points (G_1 , G_2 , G_3 , and B_2), according to distance along
 8 orthogonal dimensions of the system simulation location space (δx_1 **390** and δy_1 **392**).
 9 Conceptually grid points B_2 and G_1 together receive $(1-\delta x_1)$ of the mass of N_1 , while grid points
 10 G_2 and G_3 together receive δx_1 of the mass of N_1 . More specifically each dimension is processed
 11 in a geometric fashion, so the total mass contribution from N_1 to B_2 , for example, is $(1-\delta x_1) * (1-$
 12 $\delta y_1)$, and so forth. As illustrated in the figure, δx_1 is the projected distance along the x-axis from
 13 B_2 to N_1 , and similarly for δy_1 with respect to the y-axis, B_2 , and N_1 .
 14

15 **[0169]** The figure also illustrates mass allocation of node N_2 **376** to four neighboring
 16 grid points (B_1 , B_2 , B_3 , and I_1), none of which are guard points. The mass contribution from N_2
 17 to point B_2 is additive with the mass contribution from N_1 to B_2 . Also, there may be any number
 18 of other nodes (not illustrated) within the same grid cell as either of nodes N_2 and N_1 , and
 19 masses from the respective nodes are accumulated in the same manner as illustrated for N_2 and
 20 N_1 .
 21

22 **[0170]** Fig. 3F is a conceptual diagram illustrating an embodiment of three-point
 23 interpolation of node mass to grid points, as performed during mass accumulation (such as
 24 “Determine Spline Weights” **342**, of Fig. 3D). The figure is representative of operations similar
 25 to Fig. 3E, except the node being processed according to mass accumulation affects masses
 26 accumulating for nine nearest-neighbor grid points (B_0 **370**, B_1 **371**, B_2 **372**, B_3 **373**, B_4 **374**, I_4
 27 **384**, I_3 **383**, I_2 **382**, and I_1 **381**). The formula representing accumulation to a point (such as I_1) is
 28 implementation dependent.
 29

30 **[0171]** Fig. 3G is a conceptual diagram illustrating an embodiment of applying guard
 31 grid point masses to interior grid points, such as operations referred to elsewhere herein (“Fold
 32 Guard Cell Contributions” **332** of Fig. 3C, for example). The elements and representations are
 33 similar to Fig. 3E. In a first stage of processing, contributions of “right-hand column” guard
 34 elements (G_2 **388**, G_3 **389**, and G_4 **390**) are summed, or “folded” into corresponding guard and
 35 interior elements of the adjacent column (G_1 **386**, B_2 **372**, and B_3 **373**, respectively), as
 36 suggested conceptually by curved arrows **396**. In a second stage of processing, contributions of

1 “top row” guard elements (G_1 386 and G_0 385) are summed to (or folded into) corresponding
2 interior elements of the adjacent row (B_1 371 and B_2 372, respectively), as suggested
3 conceptually by curved arrows 395. The summation processing corresponds to even parity.
4 Similar processing is performed for the other two edges of the region.
5

6 [0172] Fig. 3H is a flow diagram illustrating selected details of an embodiment of
7 digital density filtering, such as operations referred to elsewhere herein (“Digitally Filter
8 Density” 333, of Fig. 3C, for example). Conceptually each density grid is filtered, alone or in
9 combination with other density grids, according to embodiment. Filtering each density grid may
10 include filtering all of the elements of the respective grid, although in certain embodiments
11 filtered elements may be selected. Applying the digital density filtering process includes
12 determining edge conditions for each grid element, “smoothing” temporary copies of elements
13 of the grid, and replacing the original grid elements with the smoothed elements.
14

15 [0173] More specifically, flow begins (“Start” 350) and a working copy of grid
16 elements is created. Then additional elements are added “outside” the spatial boundaries of the
17 temporary grid (“Populate Guard Cells” 351). The added guard elements enable more useful
18 smoothing results in some usage scenarios. Then a local averaging is performed on elements of
19 the temporary grid, including the guard elements (“Apply Spreading Function” 352). In some
20 implementations the spreading function reduces numerical artifacts associated with short-
21 wavelength density fluctuations. In some usage scenarios the numerical artifacts arise due to
22 inaccuracies in representation of a grid or grid elements.
23

24 [0174] Any combination of smoothing functions may be used, according to various
25 embodiments, including relatively conservative and relatively more aggressive techniques. In
26 some embodiments a binomial weighting function implementing a 1-2-1 spreading (with a
27 subsequent division by four to preserve total mass) over spatially neighboring grid element
28 values is used. In some embodiments the binomial weighting is performed in any number of
29 orthogonal dimensions, up to and including the maximum number of spatial dimensions
30 represented in the SDI simulation. After completing the spreading processing, the temporary
31 elements are used to replace the original array elements (“Copy to Original Array” 353) and
32 flow is complete (“End” 354).
33

34 [0175] In some embodiments all of the filtering operations for all of the elements of all
35 of density grids are completed before any of the associated temporary results replace the original
36 elements, as the original elements are required as inputs to respective filtering computations for

1 each grid. Alternatively, temporary copies of all of the original elements may be made, and the
2 copying may occur as filtering result are made available. Other similar arrangements of original
3 and temporary element management with respect to filtering computations are envisioned.

4
5 **[0176]** As mentioned elsewhere herein, processing according to the illustrated flow is
6 entirely optional, according to embodiment. In addition, in some embodiments multiple
7 iterations of the flow may be performed, in some usage scenarios using varying filter functions.
8 Consequently zero or more iterations of the illustrated flow are performed (the iterations are not
9 explicitly shown), according to application requirements and implementation.

10
11 **[0177]** Fig. 3I is a flow diagram illustrating selected details of an embodiment of
12 interpolating gate fields to nodes, such as operations referred to elsewhere herein (“Interpolate
13 Gate Fields to Nodes” **338**, of Fig. 3C, for example). Conceptually field components calculated
14 according to the (discrete) grid are mapped onto the continuous spatial coordinates of node
15 locations. In some embodiments the mapping is according to the node mass accumulation (such
16 as summations performed in “Accumulate Gate Densities” **331**). In other words, if an N-point
17 spline technique is used to accumulate densities, then an N-point spline technique is also used to
18 interpolate fields to nodes, and the value of N is the same for both techniques. Using matched
19 spline weights during accumulation and interpolation prevents “self-forces” that would
20 otherwise arise and spontaneously propel a node inconsistently with forces acting on the node.

21
22 **[0178]** More specifically, flow begins (“Start” **360**) by initializing node force values
23 (such as to zero), and then a check is made as to whether processing is complete for all nodes in
24 the simulated system (“Iterated Over All Nodes?” **361**). If so, (“Yes” **361Y**), then gate field
25 interpolation processing is complete (“End” **365**). If not, then a first (and subsequently a next)
26 node is selected for processing, and flow continues (“No” **361N**). Spline coefficients are then
27 determined for the node (“Determine Spline Weights” **362**), based in part on user input in some
28 embodiments (such as those from “Commands and Parameters” **130**, of Fig. 1). In some
29 embodiments the user input is chosen to drive balancing corresponding device distribution
30 throughout an integrated circuit die.

31
32 **[0179]** After all the spline weights for the respective node have been determined, a
33 check is made to determine if all fields affecting the respective node have been processed
34 (“Iterated Over all Fields” **363**). If so (“Yes” **363Y**), then processing loops back to check if all
35 nodes have been processed. If not, then a first (and subsequently a next) field is selected for
36 processing, and flow continues (“No” **363N**). The force contributed according to the respective

field is accumulated with forces associated with other fields ("Sum Field Contributions to Force on Node" 364). The accumulation is according to each orthogonal spatial dimension associated with force modeling (i.e. x and y for two-dimensional systems and x, y, and z for three-dimensional systems). Flow then loops back to determine if all fields have been processed.

Depletion Weighting

[0180] The effect a node has on local density and resultant forces may be "artificially" increased (or decreased) to expedite nodes moving to more satisfactory placements more quickly. Local density modification may be considered to be a result of manipulating a weighting associated with the mass of one or more nodes, and is referred to as depletion weighting. In other words, depletion weighting is a technique that may be used to drive the system to the point of legality in an SAF flow via dynamical means. By providing a dynamical solution to the problem, a higher quality result may be obtained in some usage scenarios. In certain embodiments depletion weighting operates by attaching a modifier to the density contributed by a node and the expansion field force acting upon it.

[0181] In some embodiments an expansion field without depletion weighting is used. In some embodiments an expansion field with depletion weighting is used. In some usage scenarios the depleting weighting improves any combination of actual node resource footprint, block footprint, and block capacity. In some usage scenarios the depleting weighting results in nodes being driven apart only as far as necessary to achieve legality.

[0182] In certain embodiments the depletion weight is calculated from a weighted sum of the differences between the available resources and the node resource footprint onto a quantization block, i.e. the amount of resource depletion caused by presence of the node in its current state. The depletion weight acts as a simple linear weight modification to both the density contributed by the node (in accumulation processing phases) and force acting on the node (in interpolation processing phases), and dependencies computed as:

$$dpwt = (1 + m)^{pdpwt}$$

where pdpwt is the power-law configuration parameter (that in certain embodiments defaults to 0, i.e. no modification), and the modifier "m" is as defined below. There is in addition a linear term and configuration parameter cdpwt (that in certain embodiments defaults to 1, i.e. no modification) that in some usage scenarios enables improved results compared to the power-law form alone.

1

2 **[0183]** The weights are computed differently if the quantization block is depleted in
 3 any one of the resources required for the node. For example, a node may be oversubscribed in
 4 only a single resource, but undersubscribed for others, leading to no net result unless resources
 5 are considered individually. Thus, if any resource appears depleted with respect to requirements
 6 for a node, then only the depleted resources are considered. In some usage scenarios the node is
 7 thus “coerced” out of a quantization block by depletion weighting related expansion forces.

8

9 **[0184]** The following equations are used when there is depletion for at least one
 10 resource. Nomenclature:

11 f_a node footprint for atom (a)

12 b_{f_a} block footprint for atom (a)

13 b_{c_a} block capacity for atom (a)

14 For overfull (i.e. depleted) quantization blocks, the modifier m is given by:

$$15 \quad m = \text{cdpwt} * \text{sum}_a \{ f_a * (b_{f_a} - b_{c_a}) / b_{c_a} \}$$

16 where only terms with $(b_{f_a} - b_{c_a}) > 0$ are considered, sum_a indicates a sum over all values
 17 of iteration variable “a”, and the term atom refers to a slot in an underlying SAF. The modifier
 18 ensures that (a) resources that are more limited are given higher weight, and (b) nodes
 19 possessing multiple depleted resources have higher weight.

20

21 **[0185]** For the case of no depletion, the modifier m is given by:

$$22 \quad m = \text{sum}_a \{ f_a * (b_{f_a} - b_{c_a}) / b_{c_a} / b_{c_a} \} / \text{sum}_a \{ f_a / b_{c_a} \}$$

23 where (compared to the depleted block case) additional terms serve to map the amount of
 24 depletion onto the range $[-1,0]$ (resulting in a weight in the range $[0,1]$). Thus $m = -1$ is the
 25 minimum when the block is completely empty and $m > 0$ when the block is full. In some
 26 embodiments depletion zones may be treated differently from one another.

27

28 **[0186]** In some embodiments a simpler normalization multiplier is used, i.e.
 29 $1/\text{sum}_a\{f_a\}$, having the effect of treating all depletion zones equally.

30

31 **[0187]** In some embodiments where depletion zones are treated differently from one
 32 another, depletion weighting tends to reduce density contributed by nodes that “fit” and to
 33 increase density for nodes that “don’t fit”. Also, nodes that fit tend to be affected by weaker
 34 expansion forces and nodes that don’t fit tend to be affected by stronger expansion forces. Thus
 35 the net effect of the depletion weighting is that nodes that easily fit contribute a smaller density
 36 and are affected by a lesser force from the expansion fields, but nodes that don’t fit contribute a

1 larger density and are affected by a stronger force. The variation in forces tends to contribute to
2 forward progress in several ways. The density differential between nodes that are fitting and
3 those that are not creates a situation where the system naturally (thermodynamically) evolves to
4 a lower energy state, where everything fits. Also, the force differential provides a direct
5 dynamical mechanism to cause non-fitting nodes to leave an overfull block (as a result of the
6 density surplus and the attendant local expansion field) before other nodes get a chance to leave
7 the block.

8
9 **[0188]** In some embodiments a depletion weight technique calculates the node
10 depletion weight at each of the nearest neighbor grid points used in the accumulation and
11 interpolation, so that nodes near a block boundary are subject to forces due to the inclusion of
12 the node in the neighboring block as well the block the node is included in. In certain usage
13 scenarios this prevents nodes from oscillating (or “sloshing”) between blocks when there is
14 likely no benefit to be gained from the oscillation.

15
16 **[0189]** The induced per-block expansion field tends to drive non-fitting nodes towards
17 the boundary where they may tend to cluster temporarily if the neighboring block does not have
18 the capacity to accept them. The cluster may be, however, a transient effect. Nodes that are
19 bunched near the edge of a block either slide along the edge until reaching an accepting block on
20 either side, or hover at the edge until conditions in the nearest neighboring block become
21 favorable for transit.

22 23 24 Exclusion Zones

25
26 **[0190]** In some embodiments various regions, or exclusion zones, may be defined that
27 are not allowed to include any morphable-devices, any placed elements, or any elements of
28 certain types, according to various usage scenarios. During later stages of global placement
29 iterations, exclusion zones may be processed to provide gradually growing regions of higher
30 density fields that result in repulsing forces that tend to expel nodes from the exclusion zones.
31 In certain embodiments the exclusion zones “grow” as simulation time moves forward, starting
32 out as point particles (like nodes), as miniature representations of the desired exclusion zone (the
33 miniature having an overall shape and aspect ratio equal or nearly equal to the desired exclusion
34 zone), or as two-dimensional lines, according to various usage scenarios. Subsequently the
35 starting representation evolves into an ever-growing object until the object matches the desired
36 exclusion zone in size and location. Similarly exclusion zones specified as strips across the

entire area being placed and routed begin as an exclusion line and grow over simulation time into an exclusion area equal in dimension and location to the required exclusion zone.

[0191] Exclusion zones (also referred to as "xzones") are a way to model architectural entities that explicitly prohibit inclusion of all non-qualifying node (or corresponding device) types, while preserving the SDI-based numerical model. In certain embodiments all adjacent xzones are collapsed into a single xzone, to simplify treatment.

[0192] In some embodiments simulation proceeds according to the laws of motion defined elsewhere herein, ignoring xzones, allowing the netlist a relatively large amount of time for detangling. Once the nodes are suitably spread, a transition is made to "exclusion mode" where the xzone constraints are obeyed.

[0193] A first technique to manage the transition is to explicitly move nodes out of the way, starting from the center of the exclusion zone and continuing outward. In some embodiments the outward processing is gradual to reduce disruption caused by spatial shifting of the nodes. The center of the xzone and moving xzone boundaries are defined to push nodes in a desired direction, i.e. in the direction of accessible final placement states. For exclusion zones that are in the form of a stripe along the entire chip area, nodes are moved to one or both sides as appropriate. For exclusion zones that are in the form of isolated rectangles, the nodes are moved in a ray from the center point to the affected node, to spread out the distribution in an isotropic manner.

[0194] A second technique is to apply an artificial density enhancement to the area inside the exclusion zone as it slowly expands. In this technique, twice the average density on the xzone boundary is imposed in the interior of the xzone during transition. This provides a dynamical shove against the nodes in advance of the approaching barrier.

[0195] After the xzone transition is complete, simulation continues as during the xzone transition, but with added constraints including:

- Nodes are snapped to xzone boundaries at the end of each timestep. A node may "tunnel" to the other side of an xzone if energetically favorable (see "Tunneling Congestion Relief" located elsewhere herein for additional information); and
- The density fields obey specified parity boundary conditions at the edge of each xzone, to ensure physically relevant behavior at the boundary. In some

1 implementations even parity is used, and in some implementations periodic parity is
2 used.

3 4 5 SIMULTANEOUS DYNAMICAL INTEGRATION (SDI) SIMULATION 6

7 **[0196]** SDI simulation (also known as Particle In Cell (PIC) simulation) provides
8 approximations to solutions of Newton's second law (i.e. force equals mass multiplied by
9 acceleration, or $F=ma$), as expressed by a system of coupled ordinary differential equations. For
10 each node, the sum of the forces (also known as forcing terms) acting on the respective node is
11 equal to the mass of the respective node multiplied by the second derivative with respect to time
12 of the state-space representation of the node. In some embodiments nodes are restricted to
13 planar (i.e. two-dimensional) movements, and there are four equations per node (x-position, y-
14 position, x-velocity component, and y-velocity component). In some embodiments nodes are
15 not so restricted (i.e. allowed three-dimensional movements), and there are six equations per
16 node (x, y, and z-positions, and corresponding velocity components).

17
18 **[0197]** Fig. 4 is a flow diagram illustrating selected details of an embodiment of SDI
19 modeling and simulation, such as operations referred to elsewhere herein ("SDI Simulation"
20 **306**, of Fig. 3A, for example). Overall the illustrated processing serves to advance a dynamical
21 system simulation forward in time, updating state-space variables according to Newtonian
22 mechanics. Processing begins ("Start" **401**) and the system of coupled ordinary differential
23 equations is approximately solved by numerical integration for a short delta simulation time
24 interval ("Integrate Time Forward" **402**).

25
26 **[0198]** Changes to all of the state variables for all of the nodes are then simultaneously
27 processed ("Update State Variables" **403**), based on the numerical integration. In some
28 embodiments relatively small-scale changes are then made to one or more of the forces and
29 masses of the system ("Micro Adjust Forces" **404** and "Micro Adjust Masses" **405**), according to
30 a specified or a computed rate of change, in certain usage scenarios to provide more nearly
31 continuous changes to state-space variables than would otherwise be possible. The changes to
32 the force(s) are in addition to changes naturally arising due to the advancement of simulation
33 time. For example, in some embodiments large-scale force (and mass) changes (such as "Macro
34 Adjust Forces" **304** and "Macro Adjust Masses" **305**, of Fig. 3A) are partially effected by
35 incremental changes.

1 **[0199]** The new system state is examined (“Finished” **406**) to determine if the SDI
2 simulation is complete via a test of an end condition. An example termination condition is
3 completion of simulation of a specified time interval. If the SDI simulation is finished (“Yes”
4 **406Y**), then processing is complete (“End” **499**). If the end condition is not satisfied, then flow
5 loops back for further simulation forward in time (“No” **406N**). In some embodiments
6 configurable settings are adjusted prior to or in conjunction with continuing SDI simulation
7 (such as settings associated with “Commands and Parameters” **130**, of Fig. 1).

8
9 **[0200]** Numerical integration techniques compatible with the time-integration include
10 Runge-Kutta, predictor-corrector, leap-frog, and other similar integration techniques. Various
11 embodiments use any combination of integration techniques.

12
13 **[0201]** In some embodiments the time-integration is according to a fixed timestep,
14 while in other embodiments the integration is according to an adaptive timestep. The adaptive
15 timestep results in reduced integration costs during system simulation time periods of slowly
16 changing state variables and improved numerical accuracy during system simulation time
17 periods of rapidly changing state variables, or otherwise “stiff” governing equations. In some
18 embodiments the integrator (such as used in “Integrate Time Forward” **402**) receives an input
19 Delta-t (an amount to advance system simulation time). In some embodiments the integrator
20 provides an actual Delta-t (an amount system simulation time actually advanced during the
21 integration) and a suggested Delta-t for use in succeeding integration timesteps. In some of the
22 adaptive timestep embodiments one or more of the actual and suggested Delta-t values are used
23 to control the adaptive timestep.

24
25 **[0202]** While the discussion of SDI is specific to global placement, the technique is
26 applicable to other functions of the aforementioned place and route flow, including any
27 combination of global placement, legalization, detailed placement, and routing.

28 29 30 LEGALIZATION

31
32 **[0203]** Conceptually legalization determines if the global placement is likely to be
33 usable for a successful detailed place and route, and if not, legalization attempts to improve
34 placement before proceeding to detailed placement. The determination of suitability for detailed
35 placement includes assessing one or more metrics correlated with local solvability of placement
36 (and routing) problems not addressed by global placement. In some embodiments one of the

metrics includes sectioning all of the devices according to a grid (such as a regular grid) of analysis windows, and determining locally if within each analysis window resources exceed (or fall below) requirements. If all of the analysis windows are simultaneously solvable (i.e. available resources meet or exceed requirements), then detailed placement and routing is likely to succeed without additional refinements to the global placement. Improvements, or corrective actions, may take various forms including any combination of “moving” devices from one region to another, transforming devices from one implementation form to another, and partitioning-related strategies.

[0204] Fig. 5A is a flow diagram illustrating selected details of a first embodiment of resource reconciliation, as a first example of legalization (such as “Legalization” **203**, of Fig. 2). Overall the flow includes determining a size of an analysis window and allocating all devices in groups to their respective containing windows, and sub-dividing and transforming logic functions to reduce resource over-subscription. The flow also includes checks to determine if the devices allocated to each window may be implemented with the resources available in the window (i.e. no analysis window is over-subscribed), and if continued iterations are likely to provide improved results.

[0205] Processing begins (“Start” **501**) with global placement information (such as produced by “SDI Global Placement” **202**, of Fig. 2, for example). The global placement result may not be legal (i.e. in a standard cell flow devices may be overlapping, or in a structured array flow more resources may be used than are locally available), but is good enough to continue processing via refinement techniques implemented in legalization. An analysis window is determined (“Quantize” **502**), corresponding to a quantization block size, and conceptually replicated in a regular contiguous (but not overlapping) fashion such that all of the devices in the netlist are allocated to one (and only one) window (some windows may be devoid of devices). In some embodiments relating to a structured array design flow, the analysis window is a rectangular shape having a size that is an integer multiple of a corresponding SAF tile. In some embodiments the analysis window is aligned with respect to SAF tiles.

[0206] A first determination as to whether all of the analysis windows (also referred to as quantization blocks or simply “Q-Blocks”) are simultaneously legal, i.e. none are over-subscribed, is made (“All Q-Blocks OK?” **503**). If all of the Q-Blocks are legal, then legalization processing is complete (“OK” **203Y**) and processing continues to detailed placement (see Fig. 2). Otherwise (“No” **503N**) the devices are sub-divided (“Partition” **504**) via

1 partitioning strategies including any combination of fixed blocks, recursive bisection, and other
2 similar techniques, according to embodiment.

3
4 **[0207]** A second legalization check is performed (“All Q-Blocks OK?” **505**) that is
5 substantially similar to the first check. As in the first checking case, if all of the Q-Blocks are
6 legal, then processing is complete (“OK” **203Y**) and the legalized result is ready for detailed
7 placement. Otherwise (“No” **505N**) the devices are transformed (individually or in groups) to
8 logically equivalent formulations having reduced resource over-subscription (“Morph” **506**).
9 The transformation, or morphing, operations are directed to manipulate the netlist such that logic
10 functions requiring resources not available in a Q-Block are implemented as logic functions
11 using resources that are available. As an example, an OR function required in a Q-Block
12 exhausted of OR gates may instead be implemented as a NOR gate followed by an inverting
13 gate, if a NOR gate and an inverting gate are available in the Q-Block. Morphing may be used
14 in usage scenarios including structured array regions.

15
16 **[0208]** A third legalization check is performed (“All Q-Blocks OK?” **507**) that is also
17 substantially similar to the first check. As in the first checking case, if all of the Q-Blocks are
18 legal, then processing is complete (“OK” **203Y**) and the legalized result is ready for detailed
19 placement. Otherwise (“No” **507N**) a determination is made as to whether further legalization
20 iterations are likely to result in improvement (“Continue?” **508**). If continuing is potentially
21 beneficial (“Yes” **508Y**), then one or more adjustments are made to the analysis windows
22 (“Adjust Q-Blocks” **509**), and flow loops back to repeat processing starting with quantization.
23 In some embodiments the adjustments include increasing the Q-Block size in one or more
24 dimensions according to a granularity that is an integer multiple of a corresponding dimension of
25 an underlying SAF tile. For example, the Q-Block size may start out as “1 by 1” (i.e. equal in
26 size to the SAF tile), then be increased by one in the first dimension to “2 by 1” (i.e. twice the
27 SAF tile size in the first dimension), and then be increased by one in the second dimension to “2
28 by 2” (i.e. twice the SAF tile size in the second dimension). Alternatively, the Q-Block size may
29 be successively lowered, or may be increased in one dimension while being decreased in
30 another, according to various embodiments. More than one Q-Block size choice may result in
31 legal or otherwise useful results, according to various characteristics of the results (such as
32 minimum and maximum local resource utilization, and other similar metrics).

33
34 **[0209]** If it is determined that continuing legalization processing is not useful (i.e. not
35 likely to further a solution), then processing is also complete (“Not OK” **203N**) and subsequent
36 processing includes one or more revisions (see Fig. 2). In some embodiments checking if a Q-

Block size equals or exceeds a predetermined value (either before or after one or more adjustments) is part of the continuation determination, as legalization achieved with relatively smaller Q-Block sizes, in some usage scenarios, is more likely to result in successful detailed placement.

[0210] Fig. 5B is a flow diagram illustrating selected details of a second embodiment of resource reconciliation, as a second example of legalization (such as “Legalization” 203, of Fig. 2). Flow begins (“Start” 520) and proceeds to determine a window for quantizing (“Quantize at Specified Window Size” 521), binning elements into Q-blocks and optionally morphing selected elements to find a legal result. All Q-Blocks are then tested to determine if or to what extent resource conflicts exist (“All Q-Blocks Legal?” 522). If all Q-Blocks are simultaneously free of resource conflicts (“Yes” 522Y), then processing proceeds to mark the current state as a possible solution (“Nominate Current System State as Candidate Solution” 531). A test is then made to determine if the current Q-Block is a minimum size Q-Block (“Q-Block Window Size at Smallest Possible Dimensions?” 532). If so (“OK” 203Y), then processing is complete and the result is ready for detailed placement. If the current Q-Block is not the minimum size (“No” 532N), then processing proceeds with a smaller window (“Reduce Target Q-Block Window Size” 533). Flow then loops back (“Go to Start” 535) to attempt processing with the reduced window size.

[0211] If at least one Q-Block has a resource conflict (“No” 522N), then a determination is made as to the severity of the remaining conflicts (“Characterize Extent of Quantization Failure” 523). In some embodiments the determinations include “Easy”, “Hard”, and “Extreme” cases. Relatively simple conflicts (“Easy” 528) are processed by depletion weighting (“Activate / Tune Depletion Weighting” 524), and relatively more difficult cases (“Hard” 529) are processed by modifications to repulsive (or spreading) force sources (“Adjust Spreading Field Strengths” 525). Processing for the Easy and Hard cases then flows back to repeat all or portions of global placement (as revisions in the context of Fig. 2) according to depletion weighting activation/tuning or adjusted spreading strengths (“Back to Global Placement” 527 and then “Not OK” 203N). Substantially more difficult cases (“Extreme” 530) are processed by partitioning (“Go to Partitioning” 526).

[0212] The determination of conflict severity or difficulty may include examination of objective factors (such as a ratio of resources demanded compared to supplied in the Q-Blocks or other computable figures of merit), and may also include examination of subjective factors (such as how much processing time has already been expended during legalization, and other

similar progress indicators), according to various embodiments. In certain usage scenarios, upon entry to legalization, there may be a subjective perception that the system is far from legal due, for example, to over-concentration of nodes of one or more resource classes (such as Nand2, Nor2, Mux2, Inverter, and so forth) in certain regions. In some usage scenarios the strength of the spreading forces acting on the over-concentrated resource class is increased, and earlier processing (such as global placement processing with revisions via “Not OK” 203N of Fig. 2) is repeated. In other usage scenarios, if the resource imbalance is mild, then an attempt may be made to gently nudge the system with depletion weighting activated as revised global placement processing (such as via “Not OK” 203N of Fig. 2).

[0213] However, if extended time-evolution with increasingly powerful depletion weighting does not resolve the conflicts, then in certain embodiments the quantization failure may ultimately be deemed “Extreme” even though only a comparative paucity of Q-Blocks show only slightly over-subscribed resources. As the depletion weighting influencing factors become increasingly strong, the governing dynamical equations become stiff, and the overall assessment of legalization difficulty may be escalated to extreme, even though over-subscription is small. According to various embodiments assessment of legalization difficulty includes any combination of examining the system state, the netlist topology, the timing constraints and the architecture definition.

[0214] In some embodiments of the flow for standard cell implementation technologies, legalization may be pursued via modifications or adjustments to the spreading force strength. For example, the masses of nodes may be directly correlated to the areas of the standard cells, and the capacity of each Q-Block directly correlated to the respective Q-Block area. Thus spreading forces may be used to drive density so that area consumed by nodes within a Q-Block is no greater than the area of the Q-Block. When achieved, legalization is complete and flow proceeds to detail placement. In some embodiments legalization may be pursued via partitioning, optionally in combination with spreading force strength adjustments.

Partitioning

[0215] Fig. 5C is a flow diagram illustrating selected details of an embodiment of partitioning (such as processing performed as a result of “Go to Partitioning” 526, of Fig. 5B). Flow begins (“Start” 540) and then a technique for partitioning is chosen (“Select Partitioning Algorithm” 541) via any combination of manual (user directed) or automatic (software

determined) mechanisms, according to various embodiments. If a Q-Block technique is chosen (“Q-Block Edge Flow” **542**), then processing is performed for each Q-Block (“For Each Q-Block” **543**). If a Bi-Section technique is chosen (“Recursive Bi-Section” **548**), then processing is performed for each of a set of progressively smaller windows (“For Each Window” **549**), starting, in some embodiments, with a window size equal to the entire place and route region, and proceeding to progressively smaller and smaller windows.

[0216] Processing for each Q-Block according to the Q-Block edge flow technique includes determining nodes causing resource conflicts (“Identify Nodes Impinging on Over-Subscribed Resources” **544**), followed by choosing an exit edge (“Pick Edge to Flow Through” **545**) for the nodes that are impinging. Then the nodes are ranked, for example, by separation from the chosen exit (“Prioritize by Distance to Edge” **546**) and then moved across the exit edge (“Push Nodes Across Edge Until Legal or Balanced With Respect to Resource Class” **547**), thus entering a different Q-Block. After all Q-Blocks have been processed, a determination is made as to whether a legal result has been obtained (“Legal Result?” **559**). If a legal result has not been obtained, then one or more revisions are indicated and earlier processing is repeated ((No) “Not OK” **203N**). If a legal result has been obtained (“Yes” **559Y**), then the current configuration is nominated as a candidate solution, as in other legalization techniques (“Nominate Current State as Candidate Solution” **560**). Processing may then proceed to detailed placement (“OK” **203Y**), or may return for further legalization processing with a goal of achieving a legal result at a smaller Q-Block size (Not OK, **203N**), conceptually as a revision to legalization processing as described with respect to Fig. 2.

[0217] Processing for each window according to the recursive Bi-Section technique includes formulating two sections to break the window into (“Introduce Cut Line Across” **550**) and then determining resource requirements and availability in each of the sections (“Count Resource Supply / Demand in Each Region” **551**). Nodes are then moved between the sections (“Exchange Circuit Nodes Across Cut Lines Until Legal or Fail” **552**) until successful (“Legal” **557**) or no further improvements are possible (“Fail” **556**). If the result is legal, then the current state is marked as a possible result (“Nominate Current State as Candidate Solution” **553**) and then a determination is made as to whether a smaller Q-Block should be attempted (“Desired Q-Block Configuration?” **554**). If a target Q-Block size has not been reached, then flow returns back (“No” **558**) to continue bisecting windows. If the target Q-Block size has been reached, then processing is complete and flow may proceed to detailed placement (“OK” **203Y**).

1 **[0218]** In some embodiments the recursion operations are according to a tail recursion
2 formulation, and testing for the desired Q-Block configuration may include a tail recursion end
3 check (for example, if the next region is smaller than a predetermined end condition size) as an
4 iteration termination condition for recursive window processing. In some embodiments for use
5 in an SAF flow context the predetermined end size is equal to an SAF tile size.

6
7 **[0219]** If no further improvements are possible (via “Fail” **556**), then flow continues
8 (“Done” **555**) where a determination is made as to whether an acceptable candidate solution has
9 been found (“OK” **203Y**) and detailed placement may follow, or whether revisions and
10 repetition of earlier processing are indicated (“Not OK” **203N**).

11
12 **[0220]** Nodes may be selected for speculative migration across the cut line according to
13 any combination of various criteria, including proximity to an edge, a footprint onto over-
14 subscribed resources, and any other related reason, according to embodiment. In some
15 embodiments speculative expulsion of a node from one side of the cut line to the other side may
16 include morphing operations on any combination of nodes on the one side, the other side, and
17 both sides. The morphing operations are directed to discover suitable implementation forms for
18 all nodes such that nodes in each containing region may be fully implemented using only
19 resources in the respective containing region.

20 21 22 DETAILED PLACEMENT

23
24 **[0221]** Conceptually detailed placement serves to fine-tune placement as produced by
25 legalization, determining final placement of all the devices of the netlist. In certain
26 embodiments operations are relatively limited in scope, focusing on optimizations and
27 refinements generally limited to a region corresponding to a Q-Block.

28
29 **[0222]** Particular detail placement techniques are described in detail in the SAF
30 embodiments illustrated herein. Nevertheless, any of a variety of detail placement procedures
31 and techniques may instead be employed, as the specific mechanism for performing detail
32 placement (assignment of devices to specific, non-conflicting locations) is not a limiting aspect
33 of the SAF techniques described herein.

34
35 **[0223]** In some SAF embodiments illustrated herein legalization produces Q-Blocks
36 where supply is known to meet demand. Since the SAF already has the resources laid out in

1 some structured manner, there is thus certainty of the existence of a fitting assignment of
2 resource instances in the netlist to resource slots in the SAF. Consequently, there is no risk of
3 failure to find a detailed placement solution, and moreover the Q-Blocks can be detail placed
4 independently, including in certain embodiments, in parallel, concurrent operation.

5
6 **[0224]** Some embodiments use continuous variables during global placement to specify
7 placement position. Conceptually, the position coordinates determined by global placement in
8 these embodiments may be considered as “optimal” locations for each node, when interpreted as
9 being representative of the collective configuration of all circuit elements. Detail placement
10 attempts to find actual resource slots in the SAF for each resource instance in the netlist such
11 that all resource instances are simultaneously slotted as close as possible to the coordinate
12 calculated during SDI-directed global placement. Stated differently, a collective assignment of
13 all resource instances to resource slots is sought for each resource class in the SAF, such that the
14 overall variance from the coordinates assigned by global placement (and possibly modified
15 during legalization) is minimized or reduced. Some embodiments slot each node independently
16 in the closest available unoccupied slot (instead of prioritizing individual nodes).

17
18 **[0225]** Fig. 6 is a flow diagram illustrating selected details of an embodiment of
19 detailed placement useful in a variety of applications (such as processing performed in relation
20 to “Detailed Placement” **204** of Fig. 2). The illustrated flow may be used in design techniques
21 relating to SAFs. Overall the flow includes determining a prioritized order to satisfy resource
22 requirements and performing limited-scope optimizations, according to various embodiments.
23 The flow may iterate internally to provide successively more refined solutions, and terminates
24 when an acceptable result is found, or when it is determined that further iterations are not likely
25 to produce improved results.

26
27 **[0226]** Flow begins (“Start” **601**) upon receipt of placement information as produced by
28 legalization (such as “Legalization” **203** of Fig. 2, for example). As represented by “Assign
29 Resources” **602**, resources are prioritized by class. In an illustrative embodiment the
30 prioritization is in accordance with a function of demand for resources of a respective class and
31 supply of SAF resource slots, the slots being consumed by the resource instances of the
32 respective resource class. The prioritization is carried out such that as the percentage of
33 consumed slot supply increases, the priority of the respective resource class is increased, and as
34 the supply of resource slots increases (irrespective of demand), the priority of the respective
35 resource class is decreased. The function is used to evaluate the priority of each resource class,
36 and assignment of resource instances to resource slots is performed one resource class at a time,

in the determined priority order of resource classes. In some of embodiments the prioritization is done on a Q-Block basis. That is, the function is evaluated with respect to the demand, supply, and consumption local to each Q-Block.

[0227] Iterating through resource classes in priority order, within each resource class the resource instances impinging upon the respective resource class are identified, and an initial assignment of resource instances to resource slots is generated, with each resource instance drawing the closest still-unoccupied resource slot currently available. Closeness is measured in terms of distance from a slot center to the coordinate assigned by global placement (and possibly modified by legalization), for the node containing the resource instance.

[0228] Processing continues with a first form of limited-scope refinement (“Pairwise Interchange” **603**), where selected pairs of allocated resources are interchanged in an attempt to discover an improved solution. In certain embodiments, within the set of resource instances previously assigned slots, speculative interchanges are considered between every instance and every other slot (whether occupied or not). In other words, a resource instance may be swapped with the instance occupying another slot, or may simply be moved to an empty slot. Each speculative interchange is scored according to a function of the slot position and the preferred position of the occupying resource (as assigned by global placement and possibly modified by legalization). An example function is the sum of the squares of the distances between the slot centers and the preferred positions. Speculative interchanges are accepted with strictly greedy semantics, on the demonstration of a reduced sum of squared distances from instance to slot. The interchange process will eventually stall when the collective variance of resource instances from desired positions can no longer be strictly reduced.

[0229] In some embodiments pairwise interchanges may be evaluated according to a predicate:

$$D(p_i, s_j')^2 + D(p_i', s_j)^2 <? D(p_i, s_j)^2 + D(p_i', s_j')^2$$

where

p_i is the ideal position of node I ;

s_j is the actual location of slot j ; and

$D(p_i, s_j)$ is the distance between p_i and s_j .

The sum of $D(p_i, s_j')^2$ over all assignments ($i \rightarrow j$) is minimized, according to the predicate.

[0230] When the collective variance may no longer be reduced, any resource instances of other resource classes that are associated with composite forms (i.e. forms realizable from

resources of more than one slot, such as an And2 realized from a Nand2 slot and an Inverter slot) participating in the pairwise interchange are placed in an available slot (corresponding to an ancillary resource) that is closest to the resource instance of the respective composite form. The (ancillary) resource instance slot assignments are then marked as locked, and the ancillary instances are thereafter excluded from the set of assignable and revisable resource instances to be placed when a corresponding resource class is subsequently processed. When all resource classes in the SAF have been processed as described above, a complete and valid initial detail placement for one Q-Block has been rendered, and subsequent optimization processes are enabled.

[0231] In certain embodiments, the above processes (“Assign Resources” 602 and “Pairwise Interchange” 603) are used in combination with “Dynamic Morphing” 604. In some dynamic morphing embodiments note is made of resource instances that are placed farthest from a respective desired location and improved placement of the forms is attempted by morphing to a functionally equivalent available destination form having a more suitable placement configuration of resources instances. In certain dynamic morphing embodiments, such speculation over implementation form for netlist nodes is combined with iteration over slot assignment and pairwise interchange. In the latter dynamic morphing embodiments various visited states are scored according to collective variance from preferred locations (as described above) and the best state that can be found is taken as a result. In certain embodiments states visited are limited by a computational cost criteria.

[0232] Flow then continues to a third form of limited scope refinement (“Pin Swap” 605), where pin swapping directed to improve routability is performed. Here, speculation is performed over various functionally equivalent mappings of netlist nets to instance pins. As an example, the inputs of a NAND gate may be interchanged without changing the function implemented in the gate. This and other similar equivalent mappings for other gates and circuitry are selectively evaluated. By considering such netlist transformations, attempts are made to reduce the difficulty of achieving a fully routed circuit layout.

[0233] In some embodiments an optional first-cut attempt at improving timing paths is then performed (“Size Devices” 606). As an example, driver sizing is selectively performed by revising the netlist to employ forms composed of resources with higher drive strengths. Optimization is not limited to such up-sizing. Selective down-sizing of drivers on non-critical paths is also performed, to free up high drive strength resources (such as in an SAF) for use by paths that are more critical.

1
2 **[0234]** A determination is then made (“Repeat?” **607**) as to whether additional
3 iterations of all or part of the detailed placement flow is likely to improve results. If so (“Yes”
4 **607Y**), then processing loops back to resource assignment and continues forward again from
5 there. If further iterations are found to be unlikely to offer improvement (“No” **607N**), then a
6 determination is made as to whether the results are acceptable (“OK Result?” **608**). If so (“OK”
7 **204Y**), then processing is complete and ready for routing. If the results are not acceptable (“Not
8 OK” **204N**), then processing is also complete and subsequent processing includes one or more
9 revisions (see Fig. 2). The repeat and acceptable determinations are made by any combination
10 of automatic (such as software) and manual (such as human inspection) techniques, according to
11 various embodiments.

12
13 **[0235]** Fig. 6 is an illustrative example of detailed placement, as the order and/or
14 presence of operations **602** through **606** will vary according to embodiment. That is, many
15 combinations of “Assign Resources” **602**, “Pairwise Interchange” **603**, “Dynamic Morphing”
16 **604**, “Pin Swap” **605**, and “Size Devices” **606**, will have utility as embodiments of detailed
17 placement, including combinations reordering and/or omitting one or more of these operations.
18 As specific examples, some embodiments perform “Assign Resources” **602** and “Pairwise
19 Interchange” **603** but omit “Dynamic Morphing” **604** and “Pin Swap” **605**, while other
20 embodiments selectively perform “Dynamic Morphing” **604** and then subsequently perform
21 “Assign Resources” **602** and “Pairwise Interchange” **603**.

22
23 **[0236]** Another embodiment of detail placement re-employs SDI-directed placement
24 methodology (targeted at a resource-level netlist) optionally constrained to a reduced sub-circuit
25 present in a specific Q-Block. In the SDI-directed detail placement embodiment, the specific
26 forcing terms in the system of simultaneous governing equations are modified from that
27 described in global placement, and force models more appropriate to detail placement are
28 substituted. For example, in detail placement, once the Q-blocks are formed and legalized, there
29 is no further need to perform inter-Q-Block interchange of nodes. Consequently the bulk density
30 fields that were used in global placement to control unsustainable over-concentrations of specific
31 resource types are unnecessary by construction in the detail placement context. Thus the bulk
32 density fields are replaced by forcing terms that represent a spring drawing the resource-level
33 instances of each form toward the position assigned by global placement. Simultaneously,
34 overlap repulsions arising from pair-wise occupancy exclusions between resource instances of
35 each resource class act to drive the resource instances toward feasible slots while preserving the

1 topological disentanglement that was a key result of the global placement previously obtained by
2 SDI-directed techniques.

3
4 **[0237]** The illustrated SAF embodiments emphasize a conceptual separation between
5 global placement, legalization and detail placement, as facilitated by the described form-level
6 netlist abstraction and the technique of morphing and facilitating data structures and SAF
7 enabling properties. The approaches to detail placement used in the illustrative SAF
8 embodiments herein are not meant to be limiting and other detail placement approaches may be
9 substituted.

10
11 **[0238]** In some standard cell implementation technologies, there is no concept of
12 resource classes. In some usage scenarios “slots” correspond to tiled regions of a defined size.
13 Any standard cell may be positioned at any location on a so-called standard cell grid, with the
14 understanding that each standard cell consumes some number of contiguous abutting slots, and
15 that neighboring standard cell instances are non-overlapping.

16
17 **[0239]** In some implementations assessment of Q-Block legality by comparing demand
18 for standard cell slots to the capacity of the Q-Block (determined by counting the number of
19 contained standard cell slots), is an uncertain predictor of detail placement success. As an
20 example, consider a Q-Block that is 10 standard cell rows high by 100 standard cell columns
21 wide. The assigned standard cells in the Q-Block would be organized into no more than 10
22 rows, each row limited to 100 units (standard cell columns) in length. A detail placer may be
23 unable to construct row-sets of instances. Continuing the example, consider 11 standard cell
24 instances of a single cell type, the single cell requiring 51 standard cell columns. Then the Q-
25 Block would be infeasible, even though the slot supply comfortably exceeded demand.

26
27 **[0240]** As a result, standard cell embodiments may use a quantization (a Q-Block
28 sizing) that is enough larger than the largest frequently occurring standard cell (in certain usage
29 scenarios standard cells having sequential logic, a.k.a. “sequentials”) to improve the likelihood
30 that over-concentrations of unwieldy standard cells will succeed during the slot assignment
31 phase of detail placement. In some embodiments of a detail placer for standard cell design flows
32 the detail placer may include a mechanism for feeding back from detail placement to
33 legalization.

34
35 **[0241]** In one representative standard cell embodiment, the feedback includes operating
36 an iterative partitioner included in the detail placer. Solution of each Q-Block is attempted. If

1 any fail, then the capacity of the failing Q-Blocks is artificially depressed. The partitioner then
2 runs to attempt to redistribute the netlist nodes to distort the net topologies to the least possible
3 extent, while still achieving resource legality in each Q-Block, including the effect of the
4 artificially depressed capacity of certain Q-Blocks for the purpose of inducing the system to
5 move some cells to different neighboring Q-Blocks in the hopes of finding a soluble
6 configuration. Some embodiments targeting standard cell flows are based upon a conceptual
7 framework where the global-placement position coordinates assigned to each netlist node are
8 deemed ideal when considered as a collective determination, not as an individual determination.
9 Consequently, the standard cell embodiment partitioner preferably seeks to move across the
10 failing Q-Block edges whatever is already closest to the edge, and that can therefore be
11 displaced slightly with the least distortion in the overall netlist net topology.

12
13 **[0242]** In another representative standard cell embodiment, the cells in a Q-Block are
14 grouped into rows, determined through considering relative juxtaposition of the cells in the
15 coordinate that varies perpendicularly to the standard cell rows (such as the y coordinate). Thus
16 cells at higher y position coordinate will be promoted to the row above in preference to cells
17 with lower y position coordinate. Once the rows are formed and the contents optimized until
18 each row fits in the width of the containing Q-Block, layout within the rows proceeds in a
19 similar fashion. Specifically, cells are laid out horizontally within each row, and the global
20 placement assigned x position coordinates are used to determine relative packing order along the
21 standard cell row within each Q-Block.

22
23 **[0243]** In another representative standard cell embodiment, the detail placement is
24 solved via a re-employment of the SDI-directed techniques described previously for global
25 placement. The spreading fields of global placement are replaced with forcing terms modeling a
26 spring drawing each netlist cell instance toward the respective node position coordinate
27 determined by global placement. Moreover, pairwise overlap repulsion interactions between
28 neighboring nodes are included and tend to tile the nodes toward net disentanglement.

29
30 **[0244]** In variations of embodiments of detail placement for standard cells, further
31 optimizations may be performed through orientation speculation and pin swapping, e.g. to
32 reduce routing congestion. The optimizations are based upon the observation that each net that
33 crosses a given line contributes to demand for tracks crossing the line. If the demand for the
34 tracks crossing the line exceeds the supply of perpendicular-running tracks, then routing is more
35 difficult. However, the condition of over-demand for routing tracks may be highly localized. If
36 nets crossing the line from opposite directions to reach pins on either side can be swapped, then

the track demand is reduced by two. Techniques include pin swapping by exploitation of pin permutability semantics on an underlying standard cell (such as swapping inputs on a NAND gate) and by rotation and flipping a standard cell according to standard rules of the implementation architecture.

TIMING CLOSURE AND TIMING-DRIVEN PLACEMENT

[0245] Conceptually timing closure and timing-driven placement operate to reduce critical timing delays to facilitate higher-speed operation of an implementation of a netlist. A high fidelity timing kernel, in conjunction with precise modeling of interconnect parasitics, specifies timing-driven attractive forces, or modifies effects of one or more net attractive forces used during SDI-directed global placement. Timing-driven forces are derived from a snapshot of state variables of the time-evolving dynamical system simulation. As the dynamical system changes (due to influences of various forces, for example), electrical characteristics of a placement of the associated netlist also change, and effects of the new state variables (such as longer or shorter interconnects) are fed back into a timing kernel to reevaluate timing characteristics of a placement corresponding to the state variables. In some embodiments timing-driven forces are calculated and applied to nets selectively, in certain embodiments as a function of any combination of one or more slack coefficients, worst negative slack values, and total negative slack values. In some embodiments timing forces may also be derived using a path-based approach, where the paths include various critical and near-critical paths according to a placement of the netlist as indicated by the state variables.

[0246] Various quanta of SDI simulation time may advance between timing-driven force re-calculation, from as frequently as a single SDI iteration to as infrequently as an unbounded number of SDI iterations. For example, timing-driven forces may be adjusted on every iteration of the integration timestep or every N iterations, where N may be provided by a user, or determined by software, according to embodiment. In some embodiments, the frequency of timing update may be automatically computed by the timing kernel (in an “auto timing-directed-force update mode”) depending on the state of the dynamical system. For example, when the system is “hot” (i.e. has a relatively high ratio of kinetic energy to total energy), timing force updating is performed more frequently than when the system is “cold” (i.e. has a relatively low ratio of kinetic energy to total energy). In some embodiments the update frequency is determined in part by tracking system parameters including any combination of a cumulative node displacement since last update, a maximum displacement per net, and other

1 similar metrics to trigger an auto-update of timing forces. An incremental timing update is
2 performed on a timing graph when relatively small displacements of nodes are detected with
3 respect to the prior update. Iterative slack allocation and net delay budgets are computed on the
4 instantaneous placement every N iterations to adapt the timing budgets based on the time-
5 evolving placements.

6
7 **[0247]** Certain high fanout (or portions of high fanout) nets are identified as non-
8 critical with respect to timing and have little or no timing-driven forces associated with them.
9 False timing paths and non-critical multi-cycle timing paths are also identified as non-critical
10 and receive little or no timing-driven force enhancements. In some usage scenarios control nets
11 such as reset and one or more clocks may be recognized as timing non-critical.

12
13 **[0248]** Timing critical nets (or portions of nets) are identified and receive relatively
14 stronger timing-driven forces, in certain embodiments based on normalized timing slack
15 determined for the net. Thus a distinct timing-driven force component may be associated with
16 every pin on every net (or any sub-grouping thereof). In embodiments where the connectivity-
17 based net attractive force is equal for each pin on a net, the timing-driven force tends to enable
18 prioritizing resultant physical location according to greater timing criticality. At a macroscopic
19 level, timing-driven forces tend to keep timing critical and near timing critical devices in
20 relatively close physical proximity, thus reducing associated parasitics and improving timing
21 performance. The timing-driven forces also tend to guide placements toward solutions where
22 relatively higher drive strength devices are associated with relatively greater parasitic loads
23 (corresponding to longer wire lengths) and relatively lower drive strength devices are associated
24 with relatively lower parasitics (corresponding to shorter wire lengths).

25
26 **[0249]** In some embodiments parasitics (for example parasitics of relatively short
27 interconnects) are estimated using a simple bounding box model (i.e. net parasitics are estimated
28 as the product of a semi perimeter of a bounding box of the pins on the net multiplied by a
29 constant wire capacitance per unit length). In some embodiments transformations including
30 buffering, clock tree synthesis, driver resizing, timing-based restructuring, and incremental
31 timing post fixes are ignored during parasitic estimation, while in other embodiments the
32 transformations are accounted for by various estimation techniques.

33
34 **[0250]** In some embodiments parasitics (for example parasitics of relative long or
35 relatively high fanout interconnects) are estimated after inserting buffer trees and building
36 heuristically constructed near-Minimal Rectilinear Steiner Trees (MRST) of the high fanout nets

1 to accurately and efficiently estimate circuit timing. In some embodiments devices are modeled
2 as having an effective resistance that ignores input ramp time and non-linear timing response
3 effects of the device based on output capacitive load. In some embodiments a high fidelity
4 timing kernel propagates input ramp rise and fall times (treating them separately), and
5 simultaneously propagates circuit ramp time from various timing start points to various timing
6 end points. Timing exceptions (such as false and multi-cycle paths) are propagated through the
7 timing graph to account for effects of the exceptions.

8
9 **[0251]** In some embodiments, during placement, a lumped capacitive interconnect
10 delay model that ignores effects of distributed Resistance-Capacitance (RC) trees is used to
11 estimate selected parasitic effects. In some embodiments actual net routing information (or
12 approximations thereof) forms a basis for generation of one or more distributed RC trees for
13 estimating selected parasitic effects.

14
15 **[0252]** In some embodiments timing closure is implemented in a Timing Kernel (TK)
16 that dynamically updates a timing graph based on current placement state (that is in turn derived
17 from the locations of the nodes in the SDI simulation). Net and device delays are computed and
18 propagated to slack results on each pin, normalized slack coefficient(s) are determined, and then
19 updated timing-driven forces are generated for use by subsequent SDI simulation.

20
21 **[0253]** The timing graph is a graph data structure representing the netlist and includes
22 pre-computations and pre-propagations of user-defined constraints including any combination of
23 clock period, false path and multi-cycle path identifications, arrival times at primary inputs, and
24 required times at primary outputs. In certain embodiments the timing graph is organized as a
25 Directed Acyclic Graph (DAG) data structure. In certain embodiments the pre-computations
26 and pre-propagations are generated only when a new netlist is provided or modifications are
27 made to the current netlist. The timing graph includes timing node elements and timing edge
28 elements. A timing node element represents pins of a macro (such as a morphable-device), and
29 a timing edge element represents connectivity of timing node elements (such as a flattened or
30 non-hierarchical net of the netlist).

31
32 **[0254]** Timing delay through a timing node element (also known as a stage delay) is a
33 function of several parameters, including a cell delay (D_c) and a wire delay (D_w). The cell delay
34 is a function of input transition time and cell output loading. In some embodiments cell delay
35 values are determined via a cell delay table lookup. The cell delay table may be representative
36 of non-linear timing behavior and is specified in a timing library (such as a portion of

1 “Technology Description” **121** of Fig. 1). Cell output transition times are also a function of
2 input transition times and output loads, and are computed by the TK and propagated from inputs
3 to outputs.

4
5 **[0255]** A Steiner buffered tree constructor creates an interconnect tree based on
6 coordinates of pins of morphable-devices. RC parasitics are then computed from the
7 interconnect tree, and corresponding cell delays are computed according to pi-models of the RC
8 parasitics. Wire delays are computed using Elmore-Penfield-Rubenstein delay models according
9 to estimated net and pin parasitics.

10
11 **[0256]** Fig. 7A is a flow diagram illustrating selected aspects of an embodiment of
12 delay path reduction and minimization, as an example of timing closure (such as “Timing
13 Closure” **205**, of Fig. 2). As described with respect to Fig. 2, in some embodiments timing
14 closure is essentially operative within global placement, rather than, or in addition to, operative
15 external to global placement. In other words, in some embodiments timing closure operations
16 are performed intimately with operations of global placement (such as those illustrated in Fig.
17 3A). Flows having closely associated global placement and timing improvement are known as
18 having timing-driven global placement. For example, timing-driving forces may be adjusted
19 (such as in “Macro Adjust Forces” **304**) on every iteration (via “Repeat?” **307**), or the timing-
20 driven forces may be adjusted every N iterations, where N is computed or is provided by a user
21 (such as via “Commands and Parameters” **130**, of Fig. 1). The following discussion is according
22 to timing closure operation within global placement, however the technique is applicable in
23 other contexts.

24
25 **[0257]** Processing begins (“Start” **701**) with new morphable-device locations as derived
26 from SDI simulated time advancement and resultant node location evolution. Timing node
27 element locations and associated pin spatial positions are updated accordingly in a timing graph
28 (“Update Pin Coordinates” **702**). Approximate interconnect distributed resistance and
29 capacitance values are determined (“Estimate Parasitics” **703**) via any combination of an NBB
30 technique (such as for short interconnects) and a Steiner-route technique (such as for long
31 interconnects).

32
33 **[0258]** Driver trees are then added for long and high fanout nets, and nets exceeding a
34 specified maximum capacitance threshold (“Insert Buffers” **704**). In some embodiments the
35 driver trees are constructed according to recursive bipartition-based buffering, until a maximum
36 drive capacity has been met. If one or more new devices are added, thus changing the netlist,

1 then processing loops back to repeat parasitic estimation ("Changes", 704C). If no new devices
2 are added (for example since current buffering is sufficient or maximum drive capacity has been
3 met), then more nearly accurate parasitic approximations are determined, in certain
4 embodiments via Steiner-route techniques, and processing continues ("No Changes" 704N).

5
6 [0259] Delays are then disseminated through the timing graph, including computing
7 new timing edge element specific transition times ("Propagate" 705). Arrival times and required
8 times are also propagated through the timing graph in topological order. Arrival times are
9 propagated via a Depth-First Search (DFS) order while required times are propagated in reverse
10 DFS order. Spare delay time is then derived for each timing node element of the timing graph
11 ("Compute Slack" 706). The resultant slack times are then normalized and used to determine
12 revised timing weight coefficients and associated timing-driven forces for one or more pins
13 ("Normalize Slack" 707). In some embodiments timing-driven forces are reevaluated only for
14 pins participating in timing critical nets.

15
16 [0260] A determination is then made as to whether the timing closure is acceptable
17 ("OK Result?" 708). If so, then flow is complete ("OK" 205Y), and processing continues to
18 routing (see Fig. 2). If not, then flow is also complete ("No OK" 205N), but subsequent
19 processing then includes one or more revisions (see Fig. 2).

20
21 [0261] Fig. 7B illustrates a conceptual view of selected elements of an embodiment of
22 timing-driven forces, such as used during timing-driven global placement. Driver D 715 is
23 coupled to pins of three loads L_1 711, L_2 712, and L_3 713, and L_4 714. Each node is shown with
24 an associated timing slack in parentheses (-2, -1, 0, and -1, respectively). Corresponding timing-
25 driven forces are shown as F_1 721, F_2 722, F_3 723, and F_4 724 respectively. Since the timing
26 slack for L_1 711 is the most negative (-2), the corresponding timing-driven force F_1 721 is the
27 largest of the three illustrated. Similarly, since the timing slack for L_3 713 is the least negative
28 (0), the corresponding timing-driven force F_3 723 is the smallest of the three illustrated. During
29 SDI-directed placement, the action of timing forces F_1 721, F_2 722, F_3 723, and F_4 724 would be
30 such that the dynamical system nodes corresponding to D 715 and L_1 711 would experience a
31 stronger mutual attraction than that between D 715 and L_2 712, L_3 713, or L_4 714 other things
32 being equal. However, in a realistic circuit, many other factors would be simultaneously
33 considered, and moreover, more than one independent critical path could flow through any of
34 the participating nodes. Consequently, the actual motion of the nodes may not turn out to be the
35 same as might be indicated by such a consideration-in-isolation, as the full complexity of the
36 dynamical system may still overcome timing forces acting on any given node.

Steiner Route Tree Construction

[0262] In some embodiments Steiner-route tree construction is according to a heuristic-based modified Prim-Dijkstra algorithm, including elements of Prim's Minimum Spanning Tree (MST) algorithm and Dijkstra's Shortest Path Tree (SPT) algorithm, using a coefficient alpha that is between 0 and 1. As MST yields minimum wire length (or a spanning tree) and SPT yields a minimum radius tree, the coefficient alpha enables efficient trade-offs between MST and SPT.

Resistance/Capacitance (RC) Parasitic Estimation

[0263] In certain embodiments, interconnect delay, or wire delay, is determined by modeling a net as a distributed RC network, with load devices presenting a capacitive load on the net. Various approximation schemes may be used, according to embodiment, to estimate the eventual routing for the net before the routing is performed (during placement, for example). The estimated routing is used in turn to derive associated approximate RC network parameters, and the RC approximations are then used to estimate timing delays, as described elsewhere herein.

[0264] The RC network is divided into segments, and a wire segment delay is computed for each segment. In some embodiments the wire segment delay is computed according to an Elmore delay model (wire segment delay equals wire segment resistance multiplied by the sum of the wire segment capacitance and all of the associated input capacitances). In some embodiments the wire segment delay is computed according to a higher order moment delay calculation.

[0265] In some embodiments routing associated with large (or high fanout) nets is approximated by Steiner tree graph analysis. Delays from a driver to each respective load are then determined as the sum of resistance in series between the driver and the load multiplied by the sum of the capacitance between the driver and the load, where "between" refers to the tree graph segments coupling the driver to the load.

[0266] In some embodiments parasitics for short nets are estimated using net contributing factor heuristics. For example, wire capacitance from a driver to a load is equal to a

load contribution factor multiplied by a “NetMSRT” multiplied by a capacitance per unit length. NetMSRT is equal to a Net Semi-Perimeter (NSP) multiplied by an “NSP-FanOut-Scaling” factor. The NSP-FanOut-Scaling factor is equal to one-half the quantity equal to the square root of the number of net loads plus one. The load contribution factor describes a relative contribution of a load with respect to all of the loads on the net, and may be expressed as the distance to the load divided by the entire length of the net. Wire resistance is derived similarly to wire capacitance, except resistance per unit length is used instead of capacitance per unit length.

[0267] Fig. 7C illustrates a spatial organization (or topology) of driver D 715 and coupled loads L₁ 711, L₂ 712, and L₃ 713 and L₄ 714 of Fig. 7B.

[0268] Fig. 7D illustrates an embodiment of NBB estimation of routing to cover the driver and the loads of Fig. 7C. As shown, NBB 725 covers all of the loads and the driver, and is defined by the spatial locations of D 715, L₁ 711, and L₄ 714.

[0269] Fig. 7E illustrates an embodiment of a rectilinear SRT estimation to cover the driver and loads of Fig. 7C.

[0270] Fig. 7F illustrates an embodiment of estimated RC parasitics associated with the RST of Fig. 7E.

Timing Weights Computation

[0271] In certain embodiments a timing weight is computed for all pins having a negative timing slack. All other pins are considered non-critical. Non-critical nets are marked as inactive nets and no timing forces are applied to them. Non-critical pins are assigned timing weights of zero (and thus affect no timing-driven forces). The timing weight of a pin may be modeled as a function of various timing parameters including pin slack, worst negative slack, total negative slack, interconnect length, and other similar parameters, according to implementation. In some embodiments the timing weight for a pin is equal to the square of the quantity equal to the slack of the pin divided by the worst negative slack of the entire netlist, and in various embodiments the timing weight is computed according to any number of linear and high-order calculations. The timing-driven forces are computed according to Hooke's law with

1 a coefficient equal to the respective timing weights (i.e. timing force equal to negative timing
2 weight multiplied by distance between driver node and load node).

3 4 5 Selected Timing Closure User Commands 6

7 **[0272]** Timing closure and timing-driven placement are automated to varying degrees
8 according to embodiment. In certain embodiments the automation is controlled or directed by a
9 plurality of control parameters provided in data files or scripts (such as via “Commands and
10 Parameters” 130, of Fig. 1). In some embodiments a relatively small number of control
11 parameters may be provided by a Graphical User Interface (GUI). Timing constraints are used
12 to perform timing closure and timing-driven placement, and the GUI may also provide for user
13 input of timing constraints files, such as Synopsys Design Constraint (SDC) compatible
14 information, via a “source SDC” command or menu item.
15

16 **[0273]** In some embodiments and usage scenarios design automation software
17 (including timing closure and timing-driven placement) may be operated in a batch mode. In the
18 batch mode any combination of selected switches may be specified in a file (such as a “schedule
19 file”, that may be included in “Commands and Parameters” 130, of Fig. 1). A first control
20 switch instructs SDI-driven (sometimes also referred to as force-driven) placement operations
21 (such as operations performed by a placement engine) to apply timing-driven forces at each
22 timestep. By default, the forces are turned off in some embodiments. Timing-driven forces are
23 recomputed at predefined intervals, or at a selected frequency with respect to timesteps, as
24 specified by another control switch.
25

26 **[0274]** A second control switch instructs SDI-driven placement to perform timing
27 analysis at predefined time intervals of the SDI simulation, and to report a specified number of
28 critical paths or selected critical paths. In certain usage scenarios the report includes some or all
29 of the most critical paths. If the first control switch is on, then the second control switch is
30 automatically turned on also. However, in some usage scenarios, users may keep the first
31 control switch off with the second control on to perform a timing analysis based on a current
32 system configuration. Selected critical paths may then be reported at predefined intervals during
33 SDI-driven placement. The interval may be user specified, and the reported paths may include a
34 selection of the most critical paths, with the report including worst-negative-slack information.
35

1 **[0275]** A third control switch controls how frequently a timing update is performed and
2 timing-driven force computation is performed in the SDI simulation (i.e. when the first control
3 switch is on). In some embodiments a default value for a parameter associated with the third
4 control switch is 50; i.e. every 50 timesteps timing-driven forces are determined anew. In
5 certain usage scenarios a larger value is specified for larger designs. For example if a design is
6 more than one million gates, then an iteration frequency of 100 may be specified. In some usage
7 scenarios the frequency may be adjusted dynamically (either manually by a user or automatically
8 by software). For example, at stages of placement where changes are relatively small (such as
9 later stages of placement), the interval may be increased.

11 **[0276]** In some embodiments GUI “radio buttons” may be provided to enable a user to
12 enable (or disable) any combination of the control switches. In some embodiments a command
13 window (either separate from or associated with the GUI) may be used to specify the third
14 control switch and the associated parameter.

17 SDI-DIRECTED ELECTRONIC DESIGN AUTOMATION (EDA) FLOW

19 **[0277]** Figs. 8A and 8B collectively are a flow diagram illustrating selected details of
20 an embodiment of an integrated circuit Electronic Design Automation (EDA) flow using one or
21 more techniques including SDI-directed global placement, legalization, legalization-driven
22 detailed placement, timing optimization, and routing. In the illustrations dashed-boxes represent
23 information provided in certain embodiments by users of the flow. In some embodiments
24 element **815** is provided by users of the flow while in other embodiments it is generated by
25 element **813**, and thus **815** is shown having a unique dashed-box patterning.

27 **[0278]** As a starting point, a design to be implemented is provided as a Hardware
28 Description Language (HDL) or Register Transfer Language (RTL) specification (“User Verilog
29 / VHDL RTL Design” **812**). Libraries are provided describing functional and timing
30 characteristics associated with all library cells that may be implemented on a base wafer, such as
31 a predetermined or prefabricated structured array wafer (“Cell Timing Models (.lib)” **811**). The
32 libraries may be accessed by various tools shown later in the flow.

34 **[0279]** The design is then converted to a specific implementation description according
35 to the library and the design specification (“Synthesis” **813**). Semiconductor vendor process
36 information such as the number and type of metal layers and via layers, process design rules, and

process parameters are provided (“Base Die Description” **814**). The die description also includes all die floorplan information associated with implementation as a structured array, i.e. descriptions of SAF tiles. The die description is processed (“Design Create Import Verilog / VHDL” **816**) in conjunction with a gate-level netlist produced by synthesis (“Gate-level Netlist (Verilog/VHDL)” **815**) resulting in a parsed netlist.

[0280] Selected improvements are performed, such as buffer deletion, dead logic removal, inverter pair elimination, and constant propagation (“Design Pre-optimization (buffer deletion, dead logic removal)” **817**). Then directives to guide the physical design are processed (“Load Floorplanning Constraints (IOs, RAMs, group, region constraints)” **818**). In certain usage scenarios the floorplan constraints are used to “lock” selected elements into desired regions of the die. For example IO pads may be assigned to the perimeter, and RAMs may be allocated to specific zones. Core logic may be guided to selected areas or grouped together as desired. In some embodiments the floorplan constraints are provided via one or more scripts (“Place Script; Floorplan Script” **822**).

[0281] Timing performance criteria are then processed (“Load Timing Constraints” **819**), in some embodiments according to timing libraries (“SDC Timing Libraries (.lib)” **823**). Information in the timing libraries may be according to an SDC format, and includes input arrival times, output required times, false path identification, and multi-cycle path notations. In certain embodiments subsequently locations are determined for all of the elements in the netlist (“Placement” **820**), guided by previously provided constraints. Timing performance improvements are then made to effect timing closure (“Buffering Clock Tree Synthesis Timing Driven Buffering/Resizing” **821**). Clock tree synthesis strives to meet desired clock skew constraints, and buffer resizing serves to meet user specified timing constraints.

[0282] Processing then flows (via **824**) to output post layout design data (“Export: DEF / Verilog” **831**). In certain usage scenarios a format compatible with Design Exchange Format (DEF) is used to facilitate interchange with various EDA tools. The output DEF (“DEF” **832**) specifies the structure of the design and all placement information. The output Verilog (“Verilog” **834**) specifies the post-layout gate-level netlist. The DEF output is provided along with information describing routing technology (“LEF” **833**) to compute interconnect details (“Router” **835**). The resultant geometry is output as DEF (“Routed DEF” **836**) that is processed (“3D Extractor” **837**) along with the routing technology information to determine connectivity and parasitic information (“SPEF” **839**). The parasitic information is according to a Standard Parasitic Exchange Format (SPEF).

1
2 **[0283]** A timing performance check is then made (“Timing Analysis” **840**) using the
3 parasitic information, the post-layout gate-level netlist, and device characterization information
4 (“StdCell Library” **838**). A correctness check is also made (“Formal Verification” **826**) by
5 comparing a pre-layout gate-level netlist (“Pre-layout Gate-level Netlist” **825**) with the intended-
6 to-correspond post-layout gate-level netlist. In some usage scenarios the pre-layout gate-level
7 netlist is identical to the netlist output from synthesis.

8
9 **[0284]** The illustrated EDA flow is an example only, as some of the illustrated
10 operations may be omitted or performed in slightly different orderings according to various
11 embodiments.

12 13 14 MANUFACTURE OF DEVICES DESIGNED VIA SDI-DIRECTED TECHNIQUES

15
16 **[0285]** Conceptually a structured array architecture is defined to satisfy a plurality of
17 user-specific designs. The architecture is optionally based on a pre-characterized standard cell
18 library. A plurality of user-specific designs are targeted for the defined architecture, and
19 physical layout is generated at least in part based on a SDI-directed place and route flow. An
20 inventory of wafers (or die) built according to the structured array architecture is used as a
21 starting point to manufacture instances of the user-specific designs. Thus a single structured
22 array architecture (and corresponding predetermined wafer inventory) serves to implement more
23 than one user-specific design via a SDI-directed placement and routing.

24
25 **[0286]** Fig. 9 illustrates an embodiment of selected details of manufacturing integrated
26 circuits, the circuits being designed in part based on SDI-directed design techniques. The
27 manufacturing flow begins (“Start” **901**) by receiving objectives for a design or a group of
28 designs (“Goals” **902**) along with optional information (“Standard Cell Library” **904**) regarding
29 relatively fixed-function elements previously manufactured and characterized according to a
30 selected integrated circuit production facility or “fab”. The received items are processed to
31 determine one or more SAF tiles to be arrayed to form a structured array integrated circuit
32 (“Define Structured Array” **903**). The standard cell library information may be used to develop
33 SAF tiles with lower cost than developing SAF tiles from “scratch”. Fabrication images are
34 produced from the structured array design (“Produce Lower Layer Masks” **905**).
35

1 **[0287]** The lower layer masks are combined with starting materials (“Wafers” **906**) to
2 produce an inventory of pre-fabricated structured array die (“Fabricate Lower Layers” **907**). A
3 first and a second device are designed according to a SDI-driven place and route flow, and the
4 resultant design databases are provided to the flow (“Device 1 SDI P&R Result” **908** and
5 “Device 2 SDI P&R Result” **909**). Each of the databases is then used to produce corresponding
6 sets of upper layer fabrication images (“Produce Device 1 Upper Layer Masks” **910** and
7 “Produce Device 2 Upper Layer Masks” **911**, respectively). The upper layer masks are used to
8 manufacture (“Fabricate Device 1 Upper Layers” **912** and “Fabricate Device 2 Upper Layers”
9 **913**, respectively) one or more integrated circuits according to each of the respective designs,
10 using portions of the previously developed inventory (“Fabricate Lower Layers” **907**). The
11 manufactured devices are then tested (“Test Device 1” **914** and “Test Device 2” **915**,
12 respectively) and the flow is complete (“End” **999**).

COMPUTER SYSTEM EXECUTING SDI-DIRECTED EDA ROUTINES

[0288] Fig. 10 illustrates an embodiment of selected details of a computer system to execute EDA routines to perform SDI-directed place and route operations. There are multiple sub-systems illustrated including computing and storage complexes (System **1001A** and System **1001B**) and workstations (local WS **1017B** and remote WS **1017C**). Similar elements have identifiers using the same numerical base, and a letter suffix is used to distinguish different instances. For brevity, unless there is a notable difference between the instances, only the first instance of similar elements is described.

[0289] A data processing machine (System **1001A**) includes a pair of computational elements (Processors **1014A** and **1015A**). Each processor includes a Central Processing Unit (CPUs **1010A** and **1011A**, respectively) as well as working memory (RAMs **1012A** and **1013A**, respectively). The machine is coupled to a storage array, such as disk **1018A**, that includes images of EDA software (SW **1019A**) and design database information (DD **1020A**). An interconnection resource (Local Area Network LAN **1016**) enables local communication between System **1001A**, System **1001B**, and workstation/PC (WS **1017B**) enables local users to access the facilities to direct and observe computations. Systems **1001A** and System **1001B** are also coupled to Wide Area Network WAN **1030**, such as a corporate intranet, the Internet, or both. Remote WS **1017C** communicates with any combination of System **1001A** and System **1001B** via WAN **1030**. In certain embodiments, WS **1017C** has a disk **1018C**, that includes images of EDA software (SW **1019C**) and design database information (DD **1020C**). In some embodiments at least part of the EDA software images may be compressed or encrypted while stored on disk.

[0290] SW **1019A** may include one or more machine-readable executable files corresponding to any combination of processing operations illustrated in Fig. 1, as well as any processing operations performed on behalf or under control of elements in Fig. 1. For example, global placement (such as SDI-directed global placement), legalization, detailed placement, timing closure, and routing operations may be encoded as portions of SW **1019A** for execution by System **1001A**. Similarly, design data (such as data corresponding to any combination of portions of "Commands and Parameters" **130** and "Working Data" **131**) may be stored in portions of DD **1020A**. In operation the CPUs (in conjunction with associated RAMs) execute portions of SW **1019A** to perform assorted EDA functions.

1 [0291] In some embodiments SW 1019A may include routines that are chosen (or
2 optimized) in part to facilitate parallel execution of EDA routines (such as SDI-directed global
3 placement, legalization, detailed placement, and routing) on CPUs 1010A and 1011A. In some
4 embodiments the parallel execution may be carried out on System 1001A simultaneously (or
5 overlapping) with System 1001B (via LAN 1016) such that CPUs 1010A, 1011A, 1010B, and
6 1011B are operating together to provide a SDI-directed EDA solution for a single user-specific
7 design. The parallel processing is not limited to two machines, nor to machines with multiple
8 internal processors. Rather, the parallel computation may be performed on a collection of
9 processors, however organized or subdivided amongst independent machines. For example, the
10 software may run on a massively parallel supercomputer, or on a network of multiprocessor
11 computers, or on a network of single processor computers.

12
13 [0292] In certain embodiments, each of System 1001A, WS 1017B, or WS 1017C may
14 have an associated removable media drive, represented respectively by drives 1040A, 1040B,
15 and 1040C. The removable media drives are used to load at least parts of the EDA software
16 images, such as those discussed above, from removable media, represented respectively by disks
17 1045A, 1045B, and 1045C. The removable media and the associated drives can take many
18 forms, including but not limited to optical, magnetic, and flash media, including such media as
19 floppy disks, CD-ROMs, DVD-ROMs, and flash disks.

20
21 [0293] In certain embodiments, WS 1017C transfers at least parts of EDA software
22 images SW 1019C from either or both of System 1001A and System 1001B via WAN 1030.
23 With or without a local EDA software image, according to various embodiments, WS 1017C
24 may interact with either or both of System 1001A and System 1001B for the purpose of locally
25 or remotely executing or controlling any of the global placement (such as SDI-directed global
26 placement), legalization, detailed placement, timing closure, and routing operations, as otherwise
27 taught throughout this disclosure. In various embodiments, WS 1017C selectively has control
28 interactions and/or data transfers (including data related to the design database information) with
29 respect to either or both of System 1001A and System 1001B. In various embodiments, the
30 transfers are selectively compressed or encrypted. At least parts of the EDA software images,
31 the control interactions, or the data transfers, are thus observable as propagated signals at points
32 that include signal observation point 1035C and point 1035A.

33
34 [0294] In various embodiments, the propagated signals selectively include interactions
35 related to enabling and/or licensing of WS 1017C (or a particular user of WS 1017C) to locally
36 and/or remotely execute and/or control any of the EDA operations taught herein. In certain

1 embodiments, an FTP service is made available to WS **1017C** for downloading of at least parts
2 of EDA software image **1019C** via WAN **1030**. In related embodiments, the downloaded
3 software is adapted to be a demonstration embodiment, with either limited functionality or that
4 functions only for a predetermined interval. In other related embodiments, a software key is
5 used by WS **1017C** (obtained via WAN **1030** or other means of distribution) to enable or restore
6 functionality of at least parts of the EDA software, whether the EDA software was loaded from
7 removable media **1045C** or propagated via WAN **1030**. In related embodiments, the
8 management and distribution of the software key is a component of the licensing process. The
9 licensing is not limited to workstations. In an analogous embodiment, at least part of System
10 **1001A** and System **1001B** are licensed using selective aspects of the above described
11 techniques.

12
13 **[0295]** In certain embodiments, executing EDA software, as otherwise taught herein,
14 selectively reports license related events via WAN **1030** to license management processes
15 running on at least one designated server. In related embodiments, the reported license related
16 events are evaluated in accordance with predetermined criteria and alerts, reports, control events,
17 and/or billings are selectively and/or automatically created and/or updated.

SDI-BASED DETAILED PLACEMENT EMBODIMENTS

[0296] Fig. 11 illustrates an embodiment of an SDI-based detailed placement flow useful in a variety of applications. The SDI-based detailed placement flow may replace and/or augment operations performed after global placement and before routing (such as any combination of processing relating to “Legalization” **203** and “(SDI) Detailed Placement” **204** of Fig. 2).

[0297] In **1101** a legal global placement is developed (such as via “SDI Global Placement” **202** of Fig. 2). In **1102** nodes are (optionally) prevented from moving between Q-blocks, thus increasing the likelihood that a fitting (i.e. legal) global placement is retained during continued system evolution. In some usage scenarios where circuit density is at or near a threshold of what can be supported in a structured ASIC architecture, the processing of **1102** is invoked. In some usage scenarios where the processing of **1102** is omitted, subsequent legalization processing is used.

[0298] In **1103** spreading force strengths are increased, and in some usage scenarios the spreading forces are substantially increased. According to various embodiments the spreading forces are increased by any combination of reducing digital filtering of fields sourcing the spreading forces, and increasing spatial resolution of a grid the spreading fields are calculated with respect to. In some usage scenarios the (substantial) increase in spreading forces does not result in (substantial) bulk motion since nodes (such as form-level nodes) are prevented from moving between Q-blocks. In some usage scenarios the (substantial) increase in spreading forces does add energy to the system, and various techniques for changing the ratio between kinetic and potential energy of the system may be employed (as described elsewhere herein).

[0299] In some usage scenarios processing in **1103** serves to overcome tight packing of form-level nodes that causes local density of the form-level nodes (on relatively short spatial length scales) to exceed slot density (i.e. supply) of the underlying SAF. In some usage scenarios the exceeding of supply increases effort required by a slot assigner to discriminate between alternate slot assignments. By spreading out the form-level nodes and reducing large density fluctuations on short spatial length scales, the form-level nodes within the Q-block are driven farther apart, and thus closer to coordinates of ultimate slot assignments. In some usage scenarios the reduction of density fluctuations serves to reduce dislocation during detail slot assignment, thus improving quality of the detail placement overall.

1 **[0300]** In **1104** morphing is optionally repeated, with new target locations for form-
2 centers. In some usage scenarios nodes demanding a resource may be unevenly distributed in a
3 region, and thus some of the resource-level nodes are moved a comparatively long distance to
4 reach a slot. The movement results in “cut inflation”, where nets are forced to be routed over
5 relatively longer distances and thus consume more routing resources than were anticipated by
6 the form-level placement. The cut inflation results in decreased routability. The cut inflation
7 may be overcome by the optional morphing, to improve the balance between spatial distribution
8 of resource slots and nodes. Nodes are then moved shorter distances during slot assignment,
9 reducing cut inflation and routability degradation.

10
11 **[0301]** In **1105** the netlist is elaborated with resource-level nodes and nets spanning
12 pins on the resource-level nodes (see the discussion relating to Fig. 12A and Fig. 12B). Forces
13 are included to tie resources to respective parent forms. In some embodiments information
14 relating to the resource-level nodes (and associated spanning nets) is retained in extended data
15 structures to facilitate SDI-based processing of the resource-level nodes.

16
17 **[0302]** In **1106** forces and interaction coefficients are initialized to relatively low values
18 for the new resource-level elements of the combined (i.e. elaborated) netlist. Integration is then
19 resumed in **1107**. The resumed integration is according to the forces and interaction coefficients
20 for the new elements in addition to the forces and the interaction coefficients “inherited” from
21 the global SDI-based processing. In some usage scenarios using the new and inherited forces
22 and coefficients together results in disentanglement of the resource-level nodes now present in
23 the netlist. Enabling the resource-level nodes to move independently of each other provides a
24 context for resources to move left (or right) or up (or down) with respect to sibling resources of
25 the same parent form. The movement of the resource-level forms enables more efficient slot
26 assignments otherwise indistinguishable when only the center of the parent form is examined.

27
28 **[0303]** In **1108** integration (i.e. time evolution of the system) is stopped according to
29 selected criteria. In some embodiments dampening effects are increased to drive the system
30 toward a new state reflecting separation of resource-level nodes and to prevent or reduce
31 thrashing. In some embodiments the dampening effects are auto-regulated.

32
33 **[0304]** The selected criteria may include any combination of a number of integrator
34 steps, an amount of “system time”, system kinetic energy (i.e. temperature) falling to a threshold
35 value, system kinetic energy falling by a threshold percentage with respect to an initial value,
36 and system kinetic energy falling by a threshold percentage in a single time step. The number,

1 the amount, the threshold value, and the threshold percentages may be predetermined or
2 programmatically varied according to various implementations and usage scenarios.

3
4 **[0305]** In **1109** all Q-blocks are processed. In some embodiments the processing for
5 each Q-block is according to functions described elsewhere herein with respect to Fig. 13. In
6 **1110** processing relating to **1109** is repeated until stopping criteria are met. In some
7 embodiments the criteria include full placement of all resource classes. In some embodiments
8 processing then continues according to functions described elsewhere herein with respect to Fig.
9 14.

10
11 **[0306]** Figs. 12A and 12B illustrate concepts relating to an embodiment of netlist
12 elaboration. Fig. 12A illustrates a portion of a system with three form-level nodes located on
13 computational grid **1210** and coupled by a plurality of form-level nets. Fig. 12B illustrates the
14 system of Fig. 12A with resource-level nodes (corresponding to resource-level forms) for each
15 of the form-level nodes “added” to the system. Also illustrated are connections between
16 resource-level nodes and corresponding parent nodes, as well as resource-level nets. The parent
17 connections and resource-level nets are representative of corresponding forces and interaction
18 coefficients that are added to the system as a result of elaboration and in preparation for SDI-
19 based detailed placement time evolution. The resource-level nodes and nets may be retained in
20 extended data structures for the SDI-based processing.

21
22 **[0307]** Fig. 13 illustrates an embodiment of detailed placement of a Q-block. In **1301**
23 priority of each resource class in a Q-block is assessed, based on a combination of factors
24 relating to resource supply and consumption. Less supply makes for higher priority, and more
25 consumption makes for higher priority. Note that prioritization results naturally vary from one
26 Q-block to another, as nodes (demand) and available slots (supply) vary from one Q-block to
27 another. Processing according to **1310**, **1320**, and **1330** is then performed for each resource
28 class in order according to the resource class prioritization.

29
30 **[0308]** In **1310** slot d^2 optimized slot assignment for nodes of the respective resource
31 class is performed via one or more techniques identical to or similar to processing associated
32 with elements illustrated or discussed with respect to Fig. 6 (such as “Pairwise Interchange”
33 **603**). In some embodiments the slot assignment is performed using an implementation
34 dependent technique.

1 **[0309]** In **1320** resource-level macros of the respective resource class are assigned to
2 computed (or destination) slots. The assignments are then “fixed” (i.e. prevented from moving
3 or being reassigned). According to various embodiments the fixing may be via any combination
4 of a variety of techniques. The techniques include:

- 5 • Instantaneous enactment, i.e. a node is moved directly to the destination slot and locked;
- 6 • Gradual enactment; i.e. a node is propelled toward the destination slot using a slow but
7 overwhelming force, stronger than all other forces acting on the node, so that the node
8 reaches the destination slot in an adiabatic motion over some reasonable number of
9 timesteps and is locked there; and
- 10 • Direct parametric motion; i.e. a line is drawn from the current position of the node to the
11 destination slot, and the node is moved directly along the line toward the destination slot
12 over a series of timesteps and is locked there.

13
14 **[0310]** In **1330** remaining unfixed elements are optionally enabled to relax according to
15 new coordinates corresponding to the destination slot assignments most recently made in **1320**.
16 In some embodiments (such as various embodiments using instantaneous enactment) processing
17 in **1330** is performed. In some embodiments (such as various embodiments using gradual
18 enactment or direct parametric motion) processing in **1330** is skipped.

19
20 **[0311]** Fig. 14 illustrates an embodiment of an additional pass of detailed placement of
21 a Q-block. Processing according to **1410**, **1420**, **1430**, and **1440** is performed for each resource
22 class in order according to the resource class prioritization determined in **1301** of Fig. 13. Each
23 resource class is unfixed in turn to enable additional relaxation. In some usage scenarios a
24 plurality of iterations of processing of all resource classes according to Fig. 14 is performed.
25 Unfixing each resource class enables higher priority resource classes (i.e. classes processed
26 ahead of other classes) to relax with respect to lower priority resource classes (i.e. classes
27 processed behind other classes).

1 ADDITIONAL MORPHING EMBODIMENTS

2
3 **[0312]** In at least some structured ASICs the supply of fundamental hardware resources
4 is predetermined and fixed. Careful apportionment of netlist nodes into function-realization-
5 entities (forms) can help to improve the quality of the physical solution of the EDA flow.
6 However, size and performance constraints cause the form selections of different nodes in the
7 netlist to be coupled, resulting in an extremely complex and thus potentially expensive
8 computational optimization problem. A procedural approach to generating a solution includes a
9 technique making use of Integer Linear Programming (ILP). Illustrative embodiments for circuit
10 placement are described.

11
12 **[0313]** A schema for representation of a circuit netlist when nodes of an initial (e.g.
13 synthesis- or schematic-derived) gate level netlist are interchangeable with functionally
14 equivalent alternatives implemented using different hardware resources is used. Herein, each
15 functionally equivalent realization is called a "form", and the initial gate level netlist is called the
16 form-level netlist. Exchanging a form instance in the form-level netlist with a functionally
17 equivalent alternate form is herein called "morphing". Fig. 12A illustrates a form-level net of
18 form-level nodes overlaid on a computational grid. Fig. 12B illustrates one type of view of an
19 elaboration of the form-level net of Fig. 12A to include resource-level nodes in a resource-level
20 net. Fig. 15A illustrates a form of the form-level net of Fig. 12A. In this view the resource-level
21 nodes are shown internal to the form. Fig. 15B illustrates another form that uses different
22 resources to implement the same function as the form of Fig. 15A. In at least one embodiment,
23 the form of Fig. 15B is substituted for the form of Fig. 15A through a morphing process.

24
25 **[0314]** Fig. 15C illustrates a hierarchy of nodes, having hierarchical nodes, form-level
26 nodes, and resource-level nodes. A top node (T) is the ancestor (parent, grand parent, great
27 grand parent, and so forth) of all nodes in the system. Nodes H1, H2, H3, H4, H5, H6 ... and
28 HN are hierarchical nodes. A hierarchical node is any node with other nodes as children. For
29 example, H4 is a child of H3, and H6 is a child of H4 (and a grandchild of H3). The hierarchical
30 arrangement of nodes is based on the structure of a circuit description, such as resulting from
31 synthesis (e.g. a Verilog or VHDL circuit description), in some usage scenarios and/or
32 embodiments. Alternatively all or any portion of the circuit description is manually coded by
33 designers. Nodes F1, F2, F3 ... and FN are form-level nodes representing instances of forms
34 that are compatible with implementation in a structured array, according to a library
35 specification. The top node, the hierarchical nodes and the form-level nodes are associated with
36 a form-level netlist.

1
 2 **[0315]** Nodes R1 through R8 are resource-level nodes illustrated grouped under
 3 corresponding form-level parent nodes. The exact number and type of the resource-level nodes
 4 is determined unambiguously by specification of each form type in the library, enabling
 5 computation (such as all or portions of placement and morphing) using the form-level netlist
 6 (instead of a resource-level netlist). An example of an “inflated” netlist is a netlist that has been
 7 augmented (“inflated”) with resource-level nodes via a process of “inflation”. An example of a
 8 “deflated” netlist is a netlist that has been stripped (“deflated”) of resource-level nodes (such as
 9 added via inflation) via a process of “deflation”. Individual, one or more pluralities, or an entire
 10 netlist of form-level nodes are inflated/deflated individually or in any combination, all at once or
 11 incrementally, according to various embodiments and/or usage scenarios.

12
 13 **[0316]** Fig. 15D illustrates selected nets connected between selected nodes of Fig. 15C.
 14 Two rows of entities are illustrated. The top row of entities (T, H1, H2 ... HN, F1, F2, F3, F4
 15 ... FN) illustrates the nodes. For clarity of description, the hierarchical arrangement of the nodes
 16 is omitted from Fig. 15D (but is illustrated in Fig. 15C). The bottom row of entities (N1, N2 ...
 17 NN) illustrates nets of the netlist. Note that while a specific interconnection of nodes by nets is
 18 illustrated in the figure, in general nodes are connected to any number of other nodes (zero, one,
 19 two, or more). In the figure, each line between a node and a net represents a pin, as discussed
 20 elsewhere herein. In general, nodes have pins to a plurality of nets, and a net that connects to a
 21 form-level node does so at a distinct port. For example, in Fig. 15A, the illustrated form type
 22 AND2_A has three form ports: A, B and Y.

23
 24 **[0317]** Fig. 15E illustrates the nodes and nets of Fig. 15D after augmentation (such as
 25 via inflation or elaboration) with resource-level nodes. The top row of entities (T, H1, H2 ...
 26 HN, F1, R1, R2, F2, R3, R4, R5, F3 ... FN, and RN) illustrates the nodes. The bottom row of
 27 entities (N1, N2, N3 ... NN) illustrates nets of the netlist. Compared to the pre-inflation
 28 scenario illustrated in Fig. 15D, the same nets connect to the same form ports of the same form
 29 instances, but as illustrated in Fig. 15E, there are also connections to some resource-level nodes,
 30 with extra pins being determined from the form specification. For example, one of the form-
 31 level nodes (such as F1 of Fig. 15E) is of a specific form type (such as AND2_A of Fig. 15A).
 32 A net connects to a port of the form-level node (such as a port A of F1 of Fig. 15E) and to a port
 33 of the specific form type resource-level node (such as port A of NAND2 15,105 of Fig. 15A).
 34 The resource-level node is associated with the form-level parent node via inflation of the netlist.

35

1 **[0318]** Various operations are performed when inflating a netlist. As a first operation,
2 resource-level nodes are added as children of respective form-level parent nodes. As a second
3 operation, nets connecting to form-level node ports are extended down to resource-level node
4 ports according to a form specification corresponding to the form-level node. As a first
5 operation, zero or more new nets are created that span pins of resource-level nodes, and span no
6 other elements.

7
8 **[0319]** As an example of operations that are performed when inflating a netlist,
9 consider Fig. 15A. Selected nets inside of form specification bubble **15,100A** are named. Name
10 n1 names a connection from form port A **15,101A** to NAND2 **15,105** resource port A. Name n2
11 names a connection from form port B **15,102A** to NAND2 resource port B. Name n4 names a
12 connection from INV **15,106** resource port Y to form port Y **15,103A**. Name n3 names a new
13 connection (a new net) from NAND2 resource port Y to INV resource port A. Thus a new
14 connection (as illustrated by n3) is added to a netlist whenever a form-level node of form type
15 AND2_A is inflated, while other connections (as illustrated by n1, n2, and n4) result from
16 adding pins to preexisting nets spanning form pins in a form-level netlist.

17
18 **[0320]** In various embodiments, synthesis is targeted to a library to produce a netlist of
19 form-level elements. The form-level netlist is then elaborated with additional implementation
20 detail (e.g. resource-level nodes and net connections) as determined by form specifications.
21 Placement (such as global or detail placement) and morphing are performed on a form-level
22 netlist (e.g. before elaboration or inflation), and alternatively on an inflated netlist. When
23 morphing operations are performed on a form-level netlist, there is no change to net connectivity
24 on form-level instances, thus providing enhanced computational efficiency in some usage
25 scenarios.

26
27 **[0321]** In a structured ASIC, the supply of hardware resources is predetermined and
28 fixed. The optimal selection of implementation form for each node in the form-level netlist is a
29 complex problem involving many coupled considerations. For example, certain hardware
30 resources in a structured ASIC might be faster than others, but if all form-level nodes were
31 morphed into forms that utilize the faster resource, then the total silicon area required to
32 implement a circuit could be greater than otherwise necessary, thus increasing cost of
33 manufacture. A denser placement may be obtained if the form-level instances in the netlist are
34 morphed amongst available forms so aggregate demand for each resource type across all form
35 instances in the netlist follows the same proportional relationship as the supply thereof in the
36 structured ASIC architecture being used to implement the circuit. However, since in such an

1 apportionment, many form instances will be implemented using forms that require slower
2 hardware resources, the circuit may perform slower overall. Careful apportionment of the forms
3 among the nodes of the netlist to optimize overall performance of the circuit is important. Each
4 change of a given form instance from one implementation form to another results in a change to
5 timing characteristics of all logic paths through the affected node, hence providing another
6 coupling pathway in the form determination process. Similarly, if resource exhaustion forces a
7 node to be implemented using a form such that the nearest available implementation resources
8 are far from the ideal location of the node, then routability degradation may occur.

9
10 **[0322]** There are many uses of morphing in structured ASIC EDA. The following list
11 of examples is provided for illustration only, and should not be taken as limiting.

12
13 **[0323]** As one illustrative example, consider the case of a netlist that is to be placed in a
14 structured ASIC logic array instance. Knowledge of whether the netlist can be packed to fit into
15 the available resource supply of the specified structured ASIC is desired. A simple tabulation of
16 the resources demanded by the forms in the initial gate level netlist can be performed and
17 compared to the supply of resources in the structured ASIC logic array instance. Fig. 16A
18 illustrates the supply and demand for resources R1 through R6 corresponding to target functions
19 of an integrated circuit design having a first selection of forms for the target functions. For at
20 least some of the resources, the demand exceeds the available supply. However, even if the
21 demand for any resource exceeds supply in the structured ASIC logic array instance, then a fit
22 may still be possible. It may be possible to morph some or all of the nodes in the form-level
23 netlist by exchanging selected form instances with functionally equivalent alternate forms, to
24 relieve the over demand for certain resources while increasing the demand for other
25 underutilized resources. Fig. 16B illustrates the supply and demand for resources R1 through R6
26 for the same target functions as for Fig. 16A, but using a second selection of forms for the target
27 functions obtained by morphing certain forms to use different resources. For each of the
28 resources shown, the demand is less than or equal to the supply. In this way, a morphing
29 operation can yield a determination of the feasibility of fitting a netlist into a structured ASIC
30 logic array instance.

31
32 **[0324]** As another illustrative example, consider the case of a netlist that is to be placed
33 into the smallest possible accepting logic array instance of the structured ASIC. In this situation
34 the size of the structured ASIC is not predetermined, but is to be an output of the netlist packing
35 optimization problem. Possible approaches include: A) A succession of structured ASIC logic
36 array instances of different sizes are individually evaluated using the fit-checking procedure

described in the preceding example. The smallest structured ASIC logic array instance that is large enough to hold the netlist is the result. B) Morph the form-level netlist until the stoichiometric ratios of the resources demanded by the forms matches as nearly as possible with the stoichiometric provisioning proportions in the structured ASIC. Then the ratio between the corresponding elements in the resource demand versus provisioning yields the required logic array size.

[0325] In yet another illustrative example, consider the case of the placement of a netlist within a specified structured ASIC logic array instance. In this case, in addition to determining if a netlist can fit, a complete final placement is sought, such that all resources consumed by forms of a form-level netlist are uniquely assigned to resource "slots" in the structured ASIC logic array instance. One approach is to divide available area into abutting blocks, and then attempt to find a morphing solution that fits a respective portion of the netlist over each block into the respective resource complement of the respective block. As with the netlist fit-checking operation described above, there may be an initial imbalance between the resources demanded by the forms and the structured ASIC logic array supply in a given region that can be relieved through morphing. Only a subset of the nodes in the netlist participate in the morphing operation, and only a portion of the resources of the structured ASIC logic array instance are available for utilization. The block morphing operation is performed on the subset of the netlist that is contained within each of the blocks. The blocks need not be of uniform shape or size. Of course, embodiments such as domain decomposition and netlist subsection morphing are not the only approaches to placement generation. As long as the whole netlist is morphed to fit within the resources of the whole structured ASIC logic array instance, there will be some way that the resources of the form instances in the netlist could be assigned to resource slots.

[0326] As an additional illustrative example, consider the case of placement of a netlist into a dynamically sized structured ASIC logic array instance, where the final size of the logic array is determined simultaneously with generation of a legal placement. Such a facility might work by "spreading" the netlist until nodal density fell to a point where block-based morphing (as described above) was successful for all domains containing circuit elements. The size of the final fitting configuration determines the size of the structured ASIC logic array to be used for the netlist. This example is distinct from the minimum logic array size determination example above, in that the former represents a theoretical maximum packing density determination, where all the netlist form-level nodes participate in the morph, whereas in this case there are many independent morphing problems where a reduced subset of the netlist nodes participate in

1 the morphing operation. The size of the logic array instance that can be obtained in this way will
2 in general be lower bounded by the former "theoretical maximum density" logic array size
3 described in the earlier example. In general, the fewer the number of form-level instances that
4 participate in a morphing operation, the less space-efficient the solution will be.

5
6 **[0327]** As an additional illustrative example, consider the case of a placement flow that
7 aims to generate a placement of a netlist using iterative refinement of morphing regions. In this
8 scenario, processing starts with a structured ASIC logic array instance size known to be big
9 enough to hold a morphed version of the netlist (at least as big as the minimum theoretical size
10 produced by the logic array size minimization example in the previous section). A morphing
11 window is defined, initially to be the size of the full structured ASIC logic array instance. The
12 netlist is globally placed within the window using any available global placement or partitioning
13 technique and morphing operations are attempted in subdomains (or subwindows) of the
14 (previous) morphing window. The subwindows may be constructed by bisection of the
15 enclosing window, or by any other suitable subdivision technique. When the global placement
16 has evolved to the point that each subwindow is morphing soluble, the netlist nodes are
17 constrained to stay within the subwindows, and the subwindows themselves are taken to define a
18 reduced spatial domain for further global placement or partitioning refinement. In this way, the
19 process proceeds by recursive subdivision of morphing windows, until some limiting cutoff
20 criteria is reached. For example, the process might terminate when the morphing windows reach
21 a size of 10 nanometers², or any other similar stopping criteria. Note in particular, that spatial
22 resolution of the recursively refined morphing window grid is not required to be spatially
23 uniform. Indeed, nonuniform spatial resolution refinement grids may be of special utility in
24 situations with complex floorplans.

25 26 27 Morphing Techniques

28
29 **[0328]** Now consider a detailed description of some specific techniques for
30 implementing morphing according to various embodiments.

31 32 33 Morphing Techniques: Interchange Morpher

34
35 **[0329]** An illustrative interchange morphing (problem) solver uses three specification
36 components:

1

2 **[0330]** 1) A library. The library is a statement of available forms, the function each
3 form implements, and quantity of each resource that is utilized by each form.

4

5 **[0331]** 2) Netlist nodes, each node of some particular initial form type. The netlist
6 nodes may be a subset of the netlist.

7

8 **[0332]** 3) Capacity of resources provided by the structured ASIC. The capacity may be
9 a subset of total resources available for placement. In some usage scenarios the capacity is
10 specified as an array of integers, indexed by an identifier of the resources in the structured ASIC
11 logic array architecture.

12

13 **[0333]** Interchange morphing proceeds in stages, as follows:

14

15 **[0334]** 1) Assess initial demand for resources by accumulating demand for each
16 resource type by the form of each participating node. In pseudo-code:

17

18 for_each node r do:

19 footprint(r) = 0

20 for_each node n do:

21 f = n.form

22 for_each resource r do:

23 footprint(r) = footprint(r) + library.resource_demand(f, r)

24

25 **[0335]** If footprint(r) <= capacity(r) for each r, then the nodes fit on entry and no
26 additional morphing is required in order to achieve a fit. In some usage scenarios additional
27 morphing may be desirable, since there are many factors of interest besides just placement
28 feasibility.

29

30 **[0336]** 2) Take forms without alternates. Depending on the specific construction
31 details of the structured ASIC library, there may be forms with no alternates, i.e., functions with
32 only one way to be implemented in the structured ASIC architecture (that is specified in the
33 library). Forms without alternates will not be morphing since there are no interchange
34 possibilities, so the forms without alternates are taken as is. One way to do this is to remove the
35 forms from the morphing participation set, and remove resources consumed by the removed

forms from the resource capacity vector. Alternatively other bookkeeping strategies may be used.

[0337] 3) Register balancing. In some structured ASIC architecture configurations, the forms implementing sequential (register) functions are restricted, having much reduced morphability (fewer alternate implementation forms) compared to combinational forms. For example, there may be only one or two sequential resources (flip flops) in the structured ASIC architecture, from which the sequential forms can be built. Often there is only a single sequential form per sequential resource type, for the sequential functions. In contrast, it is not uncommon for combinational functions to have a dozen alternate implementation forms, with corresponding resource demand touching each non-inverter resource type. Because of the reduced implementation flexibility, it may be desirable to resolve sequential balancing next.

[0338] This can be done, for example, by the following procedure. Score sequential nodes according to respective footprints onto oversubscribed resources. Sort the nodes by the scores, so the higher scoring nodes are considered first for morphing into alternate forms. For each sequential node with a footprint onto an oversubscribed resource, score each respective alternate form according to an objective function, and select the best scoring form. If the selected form is different from the current form, then a morph is performed. After each morph, check to see if the sequential resources have been brought into alignment with the resource supply. If so, then exit the register balancing processing, and otherwise continue to the next node.

[0339] Aspects of certain objective functions will now be detailed. Other objective functions may also be used, thus these embodiments are merely illustrative and not limiting. For scoring sequential forms, founding some usage scenarios it may be useful to accumulate 1 (one) for each combinational resource utilized, plus 10 times the number of any oversubscribed resources used by the form. Lower scores are thus preferable. For combinational forms, in some usage scenarios it may be useful to accumulate for each resource 'a' utilized by the form, the quantity:

$$\text{double sa} = (100. * \text{cfpa} * \text{tfpa}) / \text{capacity_a} * \\ (\text{tfpa} > \text{capacity[a]} ? (100. * \text{tfpa} / \text{capacity_a}) : 1.);$$

where cfpa is the form footprint onto resource a, tfpa is the total footprint onto resource a if the form were to be chosen, capacity[a] is available supply for resource 'a' in the current morphing context, and capacity_a is the same as capacity[a], unless capacity[a] equals zero, in which case capacity_a is .01 (to avoid division by zero). The formula has the property of heavily costing

1 demand for oversubscribed resources, and of accentuating the cost of using forms with a
2 footprint onto resources that are provided in smaller proportions by the structured ASIC
3 architecture. In some embodiments alternate mathematical formulas provide similar behavior.
4

5 **[0340]** 4) Morph combinational nodes. Similar to register balancing, remaining as yet
6 unmorphed non-sequential (e.g. combinational) nodes that have a footprint onto an over
7 subscribed resource are identified. The alternate forms are scored according to the objective
8 function, and the best (lowest cost) morph selected.
9

10 **[0341]** In some usage scenarios the combinational node morphing results in a collection
11 of nodes that have been morphed to fit within a resource supply of a specified problem context.
12 In some usage scenarios the combinational node morphing results are insufficient, and the
13 following additional procedures may be optionally invoked.
14

15 **[0342]** 5) A morph away from an oversubscribed resource may be blocked because
16 alternate forms all have a footprint onto some resource that will become oversubscribed if the
17 morph is taken. Thus ways to "make room" for forms that will be coming out of oversubscribed
18 resources and impinging upon different resources than a starting configuration are searched for.
19 One technique is to "extract" inverter forms. Since the inverter function can be implemented
20 with essentially any (inverting) combinational resource, there is really no danger of an inverter
21 being unable to be reinserted, if there is room. The technique comprises extracting inverters,
22 scoring forms with a footprint onto oversubscribed resources using the objective function, and
23 then taking the best scoring alternate form. Finally, the inverters (the forms implementing the
24 inverter function) are added back in, morphing as necessary to attempt to achieve a fit.
25

26 **[0343]** In some usage scenarios 5) is run after the procedures 1 through 4, although this
27 is not required.
28

29 **[0344]** 6) Building on 5), morphing may be inhibited whenever a destination resource is
30 fully occupied. Thus in addition to extracting the inverters, any forms that impinge on almost-
31 full resources are also extracted. The extracting opens up additional space so that when iterating
32 through the forms impinging on over-subscribed resources, there is more room in resources that
33 previously appeared full. Then the full set of removed nodes are reinserted, morphing as
34 needed.
35

36 **[0345]** In some usage scenarios 6) is run after 5), but this is not required.

Morphing Techniques: Integer Linear Programming Based Morphing

[0346] Some morphing embodiments use integer linear programming. A linear program is constructed comprising a system of equations and constraints specified over a set of state variables representing the number of forms of each form-type. The formulation includes:

1) Function instance conservation constraint equations

2) Resource capacity constraints

3) An objective function

[0347] The independent system variables are taken to be the number of each form to be utilized. The system variables are constrained to be non-negative integers. The count of instances of a given form type cannot be either fractional (a given netlist node is implemented exactly and entirely using one specific form in any valid morph state) or negative.

[0348] Once the constraint equations and the objective function are specified, the ILP solver returns with the number of each form to be utilized, which optimizes the objective function and satisfies the constraints. Of course, it is possible that no solution exists, if for example, the number of form instances assigned to a region is so great that the forms cannot be packed in, or if there is inadequate morphability in any of the functions. If there is no solution, then the ILP solver returns a status indicating that no solution could be found.

[0349] The function instance conservation constraint equations state that the result will have the same number of instances of each function type as were in the original configuration of the subset of the netlist participating in the morph. Stated another way, the intent of morphing is to select alternate forms implementing the same circuit function, so the action of the morpher on a set of nodes should preserve the number of instances implementing each function. Within a function, the distribution of nodes implemented in different forms can change, but the total number of nodes in all the forms implementing the function is the same in the output as in the input. Morphing per se does not change the Boolean algebraic structure of the form-level netlist. (Other optimization technologies unrelated to morphing do that, and use of morphing does not preclude use of the other technologies.)

[0350] For example, suppose that the number of form instances implementing the NAND2 function is 5, apportioned on input as 3 form instances using form NAND2_1 and 2

using form NAND2_2, and that the number of form instances implementing a MUX4 function is 7, apportioned as 3 form instances using MUX4_1, 2 using MUX4_2 and 2 using MUX4_3.

Further assume that the state variables x_0 , x_1 , x_2 , x_3 , x_4 represent the number of form instances of the forms NAND2_1, NAND2_2, MUX4_1, MUX4_2 and MUX4_3 respectively.

Then the following two constraint equations would be among the set of function instance conservation equations:

$$1*x_0 + 1*x_1 + 0*x_2 + 0*x_3 + 0*x_4 + 0*x_5 + \dots = 5$$

$$0*x_0 + 0*x_1 + 1*x_2 + 1*x_3 + 1*x_4 + 0*x_5 + \dots = 7$$

[0351] The resource capacity constraints are inequalities that state that the resources utilized by a given form allocation may not exceed resources that are available. There is one respective constraint inequality for each resource in the structured ASIC architecture. In the respective inequality constraint for each resource, the coefficient of each state variable is the number of that resource consumed by the corresponding form. The right hand side is the capacity of that resource in the current region context.

[0352] For example, consider a morphing problem for a structured ASIC architecture containing NAND2, NOR2 and INV resources (among others). There are INV_INV, INV_ND2 and INV_NR2 implementing an inverter function each using one of the INV, NAND2 and NOR2 resources respectively. There is a form XNOR2_1 implementing an XNOR2 function using three NAND2 resources and one NOR2 resource. There is a form XNOR2_2 implementing an XNOR2 function using two NAND2 and two NOR2 resources. In the current region there are 400 INV, 100 NAND2, and 150 NOR2 resources. Then the resource capacity constraints would include terms like these:

$$1*x_0 + 1*x_1 + 1*x_2 + 0*x_3 + 0*x_4 + \dots \leq 400$$

$$0*x_0 + 0*x_1 + 0*x_2 + 3*x_3 + 2*x_4 + \dots \leq 100$$

$$0*x_0 + 0*x_1 + 0*x_2 + 1*x_3 + 2*x_4 + \dots \leq 150$$

where x_0 represents the number of INV_INV forms, x_1 the number of INV_ND2 forms, x_2 the number of INV_NR2 forms, x_3 the number of XNOR2_1 forms and x_4 the number of XNOR2_2 forms.

[0353] Some structured ASIC architectures have resources that can be reconfigured to switch between different primitive Boolean functions. For example, in some structured ASIC architectures, a mask reconfiguration might allow an abstract resource to be switched between implementing either a NAND2 function or a NOR2 function. Morphing support for such architectures can be accommodated in variations of the integer linear programming formulation

by including combination constraint inequalities to constrain the sum of forms implemented using the reconfigurable resources to be no larger than the total possible. For example, posit a structured ASIC architecture such that within a given region there are 100 NAND2 resources, 100 NOR2 resources, and 100 NAND2/NOR2 combinational resources. Label the NAND2 resource 0, the NOR2 resource 1, and the NAND2/NOR2 combinational resource 2. Further, represent the footprint of form i onto resource j as R_{ij} and the supply of resource i as S_i . Then constraint inequalities would include terms like:

$$R_{00}x_0 + R_{10}x_1 + R_{20}x_2 + \dots \leq S_0 + S_2$$

$$R_{01}x_0 + R_{11}x_1 + R_{21}x_2 + \dots \leq S_1 + S_2$$

$$(R_{00} + R_{01})x_0 + (R_{10} + R_{11})x_1 + (R_{20} + R_{21})x_2 + \dots \leq S_0 + S_1 + S_2$$

[0354] The above formulation enables exploration of solutions where the combinational resources are allocated flexibly between either resource behavior, but simultaneously excludes solutions that oversubscribe the simple plus combinational resource supply.

Morphing Techniques: Objective Function

[0355] In some usage scenarios an ILP solver package allows a user to specify an objective function of the system variables to optimize, as there may be many solution vectors that satisfy the various constraint equations. Without the ILP solver, the best choice of the many available solutions may not be apparent. An objective function is a function specified as a linear combination of the system state variables. The ILP solver then returns the best solution found, as measured by the objective function. That is, of the range of solutions satisfying the constraint equations, the chosen solution will be the one that maximizes the objective function.

$$F = \sum_i O_i x_i$$

where i ranges over the number of variables in the system, x_i is the i^{th} system variable, and O_i is the coefficient to be applied to the i^{th} system variable. More specifically, $0 \leq i < N_{\text{forms}}$, where N_{forms} is the number of forms in the library and x_i is the number of the corresponding form in the solution.

[0356] One particularly useful objective function to use is a so-called "form efficiency". The form efficiency measures efficiency of implementation of each form in terms of respective Boolean computational work that the respective form performs divided by a similar measure of Boolean computational work that could be performed using resources consumed implementing

the respective form. In some usage scenarios the efficiency of a form varies between 0 and 1, although the normalization is immaterial to the optimization problem.

[0357] Other embodiments use optimization objectives other than form efficiency.

Morphing Techniques: Software Implementation

[0358] An illustrative usage scenario of form morphing follows.

[0359] The structured ASIC logic array is divided into regions, and a global placer apportions circuit nodes to the various regions. A morphing controller function then cycles through the regions, identifies respective resource capacities and respective netlist nodes contained within each region, and calls the morpher, passing in the resource capacities, nodes (with the current form assignments), possibly a choice of objective function, and possibly also an indication of the priority of the nodes, and possibly also a function for evaluating the suitability of any given form for any given node.

[0360] The morpher evaluates the number of nodes implementing each function present in the set of participating nodes as respective function instance counts according to a library. The function instance counts, along with the resource capacities, are used to formulate the system of equations and inequality constraints, as described above. The coefficients of the objective function are supplied, and the ILP solver is invoked.

[0361] If a solution is found, then the resulting quota of forms (i.e., a particular distribution of form types determined by the ILP solver) is apportioned to the participating nodes in some manner. One illustrative example technique is to pass through the nodes, and test to see if the full quota of the respective current form has been reached yet. If not, take the form, and move to the next node. If so, morph this node to the next not-yet-exceeded form type within its function group.

[0362] An additional illustrative, but not limiting, example technique for apportioning forms is as follows. Order input nodes according to a priority indicator supplied by a caller. Assign each node to a "preferred" form type (for example, whatever form type the node was assigned by the tool (e.g. a timing-driven synthesis tool) that produced the original form-level

1 structural netlist), if available. If unavailable, then assign to one of the other forms in the
2 function group (e.g. a lower or higher drive strength logically equivalent form).

3
4 **[0363]** An additional illustrative, but not limiting, example technique for apportioning
5 forms is as follows. When a preferred form quota for a node is exhausted, then instead of
6 assigning the node, push the node back onto a queue for subsequent consideration. After all
7 nodes have been visited once, and either assigned or queued, the queue of blocked nodes is
8 reprocessed. Each node of the queue is assigned any of the available alternate forms in a
9 corresponding function group.

10
11 **[0364]** An additional illustrative, but not limiting, example technique for apportioning
12 forms is as follows. Use the supplied evaluator function to evaluate the form-ranking on a per
13 node basis, thus enabling factors outside the scope of the ILP formulation to affect determination
14 of the apportionment of the quota of forms developed by the ILP based morpher. In other
15 words, the morpher is responsible for determining a fitting set of form quotas, but other systems
16 or techniques are responsible for apportioning available forms based on more detailed per-node
17 considerations. For example, timing critical path nodes may receive special treatment.

18
19 **[0365]** As a specific illustrative, but not limiting, example technique, the externally
20 supplied evaluator function returns a measure of the timing criticality of each node, enabling the
21 order of visitation for form assignment to be in order of timing criticality of the nodes. As nodes
22 are visited in timing criticality order, each is assigned a respective preferred form, if available.
23 The preferred form is determined, for instance, by the timing driven synthesis tool that produced
24 the original structural netlist. If the preferred form for a given node is no longer available (e.g.,
25 because the quota for that form determined by the ILP solver has been fully depleted), then the
26 node is either assigned an available alternate form at that time, or is queued for subsequent form
27 assignment after the full node list has been visited once for preferred-form assignment
28 disposition.

29
30 **[0366]** In some usage scenarios and/or embodiments, A problem that is sometimes
31 encountered with the aforementioned approach to timing driven morphing and form assignment,
32 is that the objective function used in the ILP morpher is unable to adequately discriminate
33 between alternate solutions with respect to timing performance of the resulting circuit. The
34 inability is because the objective function is strictly a linear combination of system variables that
35 are quantities of each form-type in a solution set, whereas the timing criticality is ordinarily a
36 complex and potentially nonlinear function of various different variables. Thus, in some

1 circumstances there is no way to directly model the timing behavior of the circuit in the ILP
2 objective function. As a consequence, the optimized solution vector returns a reduced quota for
3 some particular form type that is in relatively high demand by timing critical nodes. Thus
4 achievable timing performance of the circuit is curtailed due to inadequate provisioning of
5 desirable forms, even though adequate provisioning is possible given resource capacities of a
6 morphing context.

7
8 **[0367]** A first technique to address incidents of inadequate provisioning employs
9 critical path preservation. Nodes with pins with negative timing slack are prioritized according
10 to criticality (e.g. magnitude of negative timing slack), and some selectable percentage of the
11 nodes are selected and initially granted respective preferred form assignments. The selected
12 nodes are then subtracted out of the node set of the morphing problem, and the resources
13 consumed by the granted preferred form assignments are subtracted from the resource capacities
14 of the morphing problem. The ILP solver then tries to determine a fitting solution for the
15 remaining nodes from the remaining resource supply. If the solver fails, then the percentage of
16 critical path nodes marked for preservation is reduced, or alternatively the global placement is
17 directed to further spread the netlist to accommodate resource needs of the nodes on the critical
18 paths.

19
20 **[0368]** A second technique to address inadequate provisioning incidents employs
21 headroom analysis with incremental tracking. Nodes are granted respective preferred forms,
22 exceeding the form quota returned by the ILP solver if necessary. The over-quota granting is
23 accounted for by tracking the resource utilization and comparing the tracked resource utilization
24 to resource capacities in the morphing region. For example, after the ILP solver returns a result,
25 the resource requirements of the resulting form quota solution vector are tabulated. Then, as
26 nodes are visited for form assignment, if the preferred form of a node is unavailable due to form
27 quota exhaustion, the node is still granted the preferred form as long as headroom remains
28 between the resource requirements tabulation of the ILP form quota result vector, and the
29 resource capacity vector of the morphing region. Thus a form of “over-allocation form
30 assignment” results. A mechanism to determine if the headroom remains is by incrementally
31 tracking the resource requirements vector (the resource footprint of the nodes in this morphing
32 problem) and comparing the resource requirements vector to the resource capacities. As long as
33 there is headroom, the over-allocation of form assignments is accommodated. The incremental
34 tracking adds in the footprint of a chosen form, and subtracts out the footprint of a nominated
35 alternate, thus determining a net change to resource utilization as a result of the over-allocation

- 1 form assignment. As long as the capacity of each resource in the morphing region is not
- 2 exceeded by the cumulative form assignments, then the placement is feasible.

1 PATH-BASED TIMING FORCE EMBODIMENTS

3 Timing Driven Force Computation

5 **[0369]** Timing driven SDI-based placement uses timing forces to systematically
6 influence and optimize timing performance of a placement of elements such as in a design for an
7 integrated circuit. In some embodiments timing characteristics of a circuit are modeled in a
8 timing graph from a time-evolving placement and timing forces are applied by a placement
9 engine as a feedback mechanism. A timing graph may be a Directed Acyclic Graph (DAG) that
10 has nodes that represent pins of a circuit netlist (e.g. provided by a user of the engine and an
11 associated design flow) and edges that represents timing arcs within a library cell and
12 interconnections of the netlist. The timing forces are applied in conjunction with net
13 connectivity forces and spreading forces to improve placement quality as measured by circuit
14 timing performance and routability.

16 **[0370]** One approach for modeling timing force for use in a timing driven SDI-based
17 placement flow is known as a Path-Based Timing Force (PBTF) model. PBTF heuristics apply
18 proportionate timing forces on each node (or element) of various critical paths, so that when
19 spreading forces are applied according to each critical path, the elements are pushed away or
20 held together based on respective contribution to overall circuit performance.

22 **[0371]** In various embodiments of a PBTF system, any combination of factors may be
23 used in determining timing force on an element. The factors include:

24 Critical Paths influence Factor (CPF);
25 Drive Resistance Factor (DRF); and
26 Stage Delay Factor (SDF).

29 Critical Paths influence Factor (CPF)

31 **[0372]** CPF models contributions of a node to all or any portion of critical paths of a
32 circuit. In various embodiments of a PBTF model usage scenario a timing driven placement
33 seeks to improve any combination of the Worst Negative Slack (WNS) and the Total Negative
34 Slack (TNS) of the circuit. Contributions of a node to the critical paths of the circuit are
35 accounted for to improve the TNS of the circuit.

[0373] Fig. 17A illustrates an example circuit with a plurality of critical paths. The critical paths include:

Path 1, $P_1 = \{N_0, N_2, N_3\}$;

Path 2, $P_2 = \{N_0, N_2, N_4\}$;

Path 3, $P_3 = \{N_1, N_2, N_3\}$; and

Path 4, $P_4 = \{N_1, N_2, N_4\}$.

[0374] Node N_2 is common to all the paths, while all the other nodes are present in two of the four paths. Thus in some embodiments a CPF computation for node N_2 will be higher than CPF computations for the other nodes. In some usage scenarios all critical paths of the circuit are explicitly enumerated. In some usage scenarios not all critical paths of the circuit are explicitly enumerated, since there are an exponential number of timing paths, and CPF modeling builds a heuristic based CPF model for each node of a timing graph.

[0375] A CPF score is computed by topologically traversing nodes of the timing graph in forward Depth-First-Search (DFS) order and reverse DFS order. Two scores are computed for each node: transitive FanIn CPF (FICPF) and transitive FanOut CPF (FOCPF). The respective CPF score of each node is the product of FICPF and FOCPF.

[0376] FICPF is computed during the forward DFS traversal as a sum of FICPFs of all immediate predecessor nodes of a node if the respective predecessor node is a critical node:

$$\text{node_FICPF} = \text{Sum (critical fanin_FICPF)}.$$

[0377] Similarly, during reverse DFS traversal, an FOCPF of each timing graph node is computed as a sum of FOCPF of all immediate successor nodes if the respective successor node is a critical node:

$$\text{node_FOCPF} = \text{Sum (critical fanout_FOCPF)}.$$

[0378] Then each node CPF score is computed by multiplying the respective FICPF and the respective FOCPF:

$$\text{node CPF score} = \text{node_FICPF} * \text{node_FOCPF}.$$

[0379] CPF is then normalized by dividing the CPF score by the maximum CPF of the timing graph:

$$\text{normalized_node_CPF} = (\text{node CPF score}) / \text{Max (node CPF score)}. \quad (\text{Eq. 1})$$

[0380] Fig. 17B illustrates example computations relating to an embodiment of CPF scoring. Tuples in the figure represent (FICPF, FOCPF) pairs, and underlined numbers represent slack on each node.

Drive Resistance Factor (DRF)

[0381] DRF models contributions of each node on a critical path based on drive resistances of node drivers. In some usage scenarios drive resistance of a node driver is a significant delay contributor to overall path timing. In one modeling equation that considers first-order effects, stage delay of a gate is computed as follows.

$$\text{gate delay} = T_i + R_d * C_l; \quad (\text{Eq. 2})$$

where

T_i : intrinsic delay of the gate;

R_d : drive resistance of the gate; and

C_l = interconnect capacitance + pin capacitances (i.e. total capacitive load on the output of a gate).

[0382] In some embodiments pin capacitances are fixed (or unchanged) during timing driven placement, and thus the timing driven force model is directed to influence interconnect capacitance. According to Eq. 2, improving the product of drive resistance and total output load tends to improve stage delay of a critical path node. The product may be improved by arranging for drivers with relatively higher driver resistance (R_d) to drive relatively lower capacitive loads, resulting in drivers having relatively low driver resistance (such as some drivers on critical paths) driving higher capacitive loads (such as relatively long wires). In some usage scenarios an incremental delay cost associated with driving a "stretched" wire with a strong driver is less than with a weak driver.

[0383] Fig. 18 illustrates an embodiment of a cascade of buffers of increasing drive strength (i.e. decreasing drive resistance). Five levels of buffer are illustrated with relative drive strengths of x1, x2, x4, x8, and x16 (i.e. each stage provides a factor of two more drive than the preceding stage). Nodes driven by the buffers are illustrated respectively as N_1 , N_2 , N_3 , N_4 , and N_5 .

[0384] Overall delay of the path illustrated in Fig. 18 is minimized if all the logic levels have equal delay. Ignoring intrinsic gate delays, the delay for each element of the path is balanced by equalizing respective products of $R_d * C_l$.

Since

$$R_d(x1) > R_d(x2) > R_d(x4) > R_d(x8) > R_d(x16)$$

the PBTf system attempts to maintain the following relative capacitive loading ordering:

$$C_l(x1) < C_l(x2) < C_l(x4) < C_l(x8) < C_l(x16).$$

Since C_l is directly proportional to wire length, and higher timing force tends to result in shorter wire lengths, timing forces are made proportionate to drive resistance.

[0385] Relative DRF is normalized by dividing a respective DRF weight of each node by the DRF weight of the node having the least drive resistance:

$$\text{node DRF} = (\text{node_DRF_weight}) / \text{Min}(\text{node_DRF_weights of all nodes}) \quad (\text{Eq. 3})$$

where

node_DRF_weight = Drive resistance of the driver gate for the node under consideration.

Stage Delay Factor (SDF)

[0386] Stage Delay Factor (SDF) models stage delay contributions of each driver on a critical net (or net on a critical path) and accounts for the maximum path length of each load pin on the critical net. The SDF combines stage delay and maximum path length factors to assign an SDF force component to each load pin. An SDF force is proportional to the maximum path length associated with the load pin.

[0387] The SDF is computed as follows:

$$\text{SDF Factor} = \text{dcoeff} * \exp(\text{lpwpd} / \text{min_cycle} - 1) \quad (\text{Eq. 4})$$

where

lpwpd = load pin: worst path delay;

min_cycle = clock period delay of the clock controlling the net; and

dcoeff = driver stage delay coefficient.

[0388] The dcoeff is computed as follows:

$$\text{dcoeff} = (\text{dgsd} / \text{dpwpd}) * \text{path_levels}$$

where

1 dgsd = stage delay of the driver gate;
 2 dpwpd = driver pin: worst path delay; and
 3 path_levels = number of logic levels in the path.

4

5 **[0389]** Load pin: worst path delay is computed as follows:

6 $lpwpd = AT(load_pin) + clock_cycle - RT(load_pin)$

7 Driver pin: worst path delay is computed as follows:

8 $dpwpd = AT(driver_pin) + clock_cycle - RT(driver_pin)$

9 where

10 AT: Arrival time; and

11 RT: Required time.

12

13 **[0390]** Fig. 19 illustrates example computations relating to an embodiment of SDF
 14 calculation. In the figure:

15 $lpwpd(L_1) = 12$;

16 $lpwpd(L_2) = 11$;

17 $lpwpd(L_3) = 7$;

18 $dpwpd = 12$;

19 $clock_cycle = 10$;

20 $dgsd = 1$;

21 $SDF(L1) = dcoeff * \exp(12/10 - 1)$;

22 $SDF(L2) = dcoeff * \exp(11/10 - 1)$; and

23 $SDF(L3) = 0$.

24

25 **[0391]** A stage delay of a driver gate is the sum of the driver gate delay and the
 26 interconnect wire delay that is driven by the driver. The driver gate stage delay discriminates the
 27 load based on criticality by factoring in the worst path delay of the load pin.

28

29 **[0392]** If a load pin is part of a slower critical path, then a higher force coefficient is
 30 associated with the load pin than a load pin that is part of a relatively faster critical path. The
 31 exponential term provides discrimination between two critical paths of unequal lengths. For
 32 example, if a first critical path is missing by a target by 2ns while a second critical path is
 33 missing the target by 1ns, then a higher multiplying factor is associated with the first path
 34 (compared to the second path) due to the exponential term. Thus critical paths with worse
 35 violations are weighted more.

36

Bounding Box Based Pin Force

[0393] In some embodiments timing forces are not applied in association with non-critical loads that fanout from a critical driver, thus enabling some relaxation of some (non-critical) loads so that more critical load pins of a net may be pulled closer to the driver. In some embodiments timing forces are applied for non-critical pins, if the pins form any portion of a bounding box of a critical net. A bounding box is defined as a rectangle around all the pins of a net. If a non-critical pin is on the edge of the bounding box, then an attractive force is applied to the load pin, thus in some cases reducing total interconnect capacitance (or at least preventing an increase in capacitance).

Path based timing Force

[0394] A first variant of a path-based-timing-force is:

$$PBTF_1 = CPF * RDF + SDF$$

where

CPF: Normalized_node_CPF (as in Eq. 1);

RDF: Normalized_node_DRF (as in Eq. 3); and

SDF: Normalized_node_sdf (as in Eq. 4).

[0395] A second variant of a path-based-timing-force is:

$$PBTF_2 = CPF * RDF + RSF$$

where

CPF: Normalized_node_CPF (as in Eq. 1);

RDF: Normalized_node_DRF (as in Eq. 3);

RSF: Normalized_node_RSf; and

Normalized_node_RSf = node_slack / Minimum slack of timing graph.

1 RELATIVE-SLACK-BASED TIMING FORCE EMBODIMENTS

2
3 [0396] The SDI technique of optimizing chip placement relies on a variety of forces
4 affecting nodes in a dynamical fashion, integrated forward in time. These forces are chosen to
5 simultaneously improve metrics that constitute a desirable placement, including routability and
6 timing performance, while achieving a physically realizable (legal) configuration. An approach
7 to timing-driven placement is described in what are referred to herein as "relative slack"
8 embodiments. Relative slack embodiments provide improved results (in both absolute
9 performance as well as numerical behavior) in some usage scenarios.
10

11 [0397] In a first illustrative, but not limiting, class of relative slack embodiments, forces
12 affecting pins on a critical path (as well as pins on shared nets) are increased or decreased in an
13 incremental fashion, rather than being directly calculated by a timing kernel. In the first class of
14 embodiments, pin-to-pin forces (so-called timing-based or timing-driven forces) affecting nets
15 (e.g. timing-critical nets) are governed by a force law equation having a linear increase with
16 distance (Hooke's law) and a driver-to-load connectivity model. Other classes of relative slack
17 embodiments may employ any arbitrary functional variation with distance, as well as alternate
18 connectivity models. A set of weights governing the timing-based force attraction are
19 periodically updated, and adjusted to result in successively better relative node configurations
20 with regard to overall circuit performance.
21

22 [0398] Relative slack embodiments assume existence of a timing kernel that is called
23 during an SDI run to provide relative slack data used in updating the timing driven forces.
24 Specific details of the timing kernel implementation are irrelevant since only data from a timing
25 graph and propagated pin slacks analysis are needed. The frequency of update can be controlled
26 in a variety of ways: e.g. at regular timestep intervals, in response to a triggering event
27 (dynamical or otherwise), or in response to external (user, script, or graphical) input. Each
28 update provides a "snapshot" of the critical path analysis for every net and pin in the system at
29 that moment of time.
30

31 [0399] The relative slack as calculated for each pin, as well as the position of connected
32 pins (to handle boundary box effects as noted below), results in an adjustment in the "timing
33 weight" associated with each pin. The timing weight is then used as a multiplier in the force law
34 equation governing pin-to-pin attraction. Pins that need to be moved closer together to satisfy
35 timing constraints tend to have weights increased (modulo possible normalization, noted below),
36 in some usage scenarios in a manner varying with the amount of slack available. That is, the

less slack (or more negative slack), the greater the positive adjustment to the attraction. Pins that have excess slack tend to have weights decreased. The reduction in weight on pins that have become "over-tightened" creates additional room for relaxation towards an optimal timing state.

[0400] At least some relative slack embodiments seek to improve timing of nets that do not meet target slack through "bounding box" (or bbox) contraction. Because increases to total net length result in increased capacitance, the associated timing can be negatively impacted by long distance nets -- even if the associated load pin is not on the critical path. The long distance net effect may be especially pronounced on large designs. The bounding box contraction considers a range of distances from the net bounding box, to help ensure that the bounding box is continuously contracted (otherwise pins on the bounding box may merely trade places).

[0401] The incremental approach to change in timing forces provides a quiet and consistent approach to timing closure during the course of an SDI run. In some cases where the timing constraints have been unrealistically set, it may be necessary to introduce a maximum to the total timing forces exerted by the system (for example, adding an upper limit to the ratio of timing net energy to total net energy, through a normalization term). A wide variety of other tunable controls are possible, including but not limited to:

- baseline relative tightening factor (typically small compared to unity);
- target min pin slack (typically zero);
- positive pin slack where relaxation may occur;
- minimum change in pin slack to consider it in an "improving state";
- distance between driver and load pins when no further tightening occurs;
- distance from net bounding box where tightening starts to occur;
- min bounding box size when no further "bbox" tightening occurs; and
- relative strength of bounding box vs. critical path tightening terms.

ILLUSTRATIVE DETAILED RELATIVE SLACK PROCEDURE

[0402] An illustrative, but not limiting, relative slack procedural flow is as follows.

[0403] First, in at least some embodiments, a pre-processing phase is performed (in other embodiments this might occur as a post-processing phase), where timing weight adjustment criteria or timing weights themselves are adjusted to control properties of distribution

of the timing weights as a whole. The pre-processing permits balancing resulting timing-driven forces with other effects in the system, such as connectivity forces (affecting routability) and expansion fields (affecting routability as well as utilization).

[0404] Second, update a timing graph using a Timing Kernel (TK). Using the updated timing graph, for every pin on every timing critical net, the slack associated with the respective pin is calculated (See **20,200** of Fig. 20A).

[0405] Third, iterate over all timing critical nets **20,300**, and all load pins on the nets **20,400**. Fourth, for each load pin on a respective timing critical net, calculate a respective pin timing weight adjustment (**20,500** of Fig. 20A and the entirety of Fig. 20B):

[0406] 1. Calculate worst slack on the respective net and find bounding box pins. The pins are taken from some region around the bounding box of the net (the size of which is determined by performance tuning, scaling by system size).

[0407] 2. Determine if the respective driver pin needs to be factored into the bbox calculation. That is, when the driver pin determines the bounding box position, increasing the attraction to nearby pins that are farther from the bbox may be counterproductive. The attraction to pins on the far side of the bbox is likely more influential in decreasing the overall capacitance. Fig. 21A illustrates a driver D in the interior of a net bounding box region determined by loads L_1 , L_2 , and L_4 . Fig. 21B illustrates a driver D to one side of a net bounding box region determined by the driver and loads L_1 , L_2 , and L_4 .

[0408] 2a. To focus on connections of loads to the driver, the effect of a driver on a bbox is indirectly applied to the loads themselves, through a multiplication factor on any tightening term.

[0409] 3. For each pin, modify a respective timing weight as needed (see Fig. 20B).

[0410] 4. For pins that meet target slack (Yes-path from **21,210** to **21,250**):

[0411] 4a. If the slack for the associated is net is negative (No-path from **21,250** to **21,270**), then to continue to make positive progress

1 bounding box effects are considered. By taking into account a range of
2 distances from the bbox, rather than a hard boundary, sloshing
3 (oscillations) as pins move onto or off of the bbox is reduced.
4

5 [0412] If (see decision **21,270**) a net is near or on the bounding
6 box of a critical net, then determine how much to tighten up the
7 connection.
8

9 [0413] If (see decision **21,280**) a load pin is within a
10 specified (small) distance from the driver, do nothing
11 (End **21,285**), as further tightening of the connection is
12 counterproductive (e.g. result in increased oscillatory
13 motion between the load and driver).
14

15 [0414] Otherwise, strategies for tightening (increase
16 weight **21,290**) include:

- 17 - if the bbox size is sufficiently small, then do nothing;
- 18 - if a pin is on bbox, then tighten at full strength;
- 19 - if a pin is farther than a specified distance from the
20 bbox, then do nothing; and
- 21 - otherwise (in between), then tighten from 0-1 times
22 full strength, depending linearly on distance.
23

24 [0415] 4b. If the pin was not tightened (Yes-path from **21,250** to
25 **21,260**), then the pin may be considered as a candidate for relaxation
26 (**21,260**). By allowing connections to either strengthen or weaken, the
27 ability of the system to evolve and relax to an optimal configuration is
28 improved.
29

30 [0416] 4b1. The amount of relaxation allowed for the pin
31 connection is dependent on the worst slack for the net. If the
32 pin has positive slack, but the worst case slack on the net is
33 negative, then the amount of relaxation allowed is reduced.
34 Recall that the pin was not tightened, so little is added to the
35 total capacitance on the net.
36

1 [0417] 4b2. Further, the relaxation is subject to a reasonable
2 upper bound. Otherwise the weights may drop from substantial
3 to nonexistent in a single pass.

4
5 [0418] 4b3. In both of these cases, by moderating the
6 relaxation allowed during one update cycle, we help prevent
7 sudden movement away from what was potentially a fairly
8 optimal solution. This is manifested as increased sloshing in
9 the overall timing performance.

10

11 [0419] 5. For pins having negative slack (No-path from 21,210 to 21,220):

12

13 [0420] 5a. If (see decision 21,220) slack of a constrained pin is
14 improving according to a specified criterion, then let the pin continue to
15 evolve without change (Yes-path to End 21,225).

16

17 [0421] 5b. If (see decision 21,230) the driver and load are within a
18 critical distance, then no tightening is performed (Yes-path to End
19 21,235). Otherwise tighten the connection (increase weight 21,240), in
20 a manner varying with the ratio of the slack on the pin and the worst
21 negative slack, thus pins most affecting the critical path are likely
22 affected the most.

TIMING DRIVEN BUFFERING EMBODIMENTS

Timing Driven Buffering Overview

[0422] Timing driven buffering and resizing for integrated circuit designs, e.g. structured array architectures, provides increased performance, reduced cost, or both. Nets having high capacitance and/or fanout and timing critical nets are preferentially processed to reduce maximum delay and/or transition time, enabling allocation of limited structured array resources to more important nets. Timing driven buffering is performed to generate trees of buffers. Timing driven sizing is performed to upsize selected elements. During the buffering Steiner tree routes are segmented and various buffering options are evaluated for each segment according to buffer cost, required time, and lumped capacitance. The options are sorted and partitioned according to the sort. Computational efficiency is improved by eliminating all but a topmost portion of each partition. Options are further evaluated according to performance including timing and routing costs. Displacement coefficients of macros are computed during the sizing to evaluate desirability of reallocating resources implementing less critical macros to more critical macros. A plurality of low-level implementations of each macro are evaluated and compared. Logic replication and tunneling may be performed according to timing improvements and routing costs. Hold time fixes may be implemented by delaying clocks and/or replacing a fast FlipFlop (FF) with a slower element.

[0423] In some embodiments of design flows relating to array architecture based integrated circuits (e.g. structured arrays or other similar Application Specific Integrated Circuit (ASIC) implementations), timing driven buffering is used to “reconstruct” or “re-synthesize” nets having high capacitive loads or having high fanouts. In some usage scenarios modifying the nets reduces a maximum capacitive load driven by any buffer or driver, or group of elements. In some usage scenarios the modifying reduces a maximum fanout associated with any net or group of nets. In some embodiments a high capacitive load may be driven by a dedicated buffer, or a dedicated tree of buffers. In various embodiments any combination of maximum transition time, maximum rise/fall time, and maximum delay are minimized when performing timing driven buffering.

[0424] In some embodiments the timing driving buffering is according to fixed resources available in various structured array architectures. In some embodiments the timing driven buffering is iterative (e.g. to achieve timing closure). In some embodiments the timing driven buffering accounts for any combination of local and global congestion. In some

embodiments the timing driven buffering includes morphing non-buffer resources and allocating the morphed resources as buffers.

[0425] In some embodiments of array architecture design flows, timing driven gate resizing is used to improve performance of various combinations of highly capacitive and high fanout nets. Logic gates are upsized (i.e. replaced with a gate having an equivalent logic function but greater drive strength) as necessary to reduce maximum delay and/or transition times. In some embodiments the upsizing is via so-called “form replacement” or replacing a form-level macro with an alternate form-level macro (such as substituting a gate with a higher drive strength for a gate with a lower drive strength).

[0426] In some embodiments timing driven gate resizing is constrained according to fixed resources available in various structured array architectures. In some embodiments a plurality of resources are simultaneously “swapped” (i.e. deallocated from a first use and reallocated to a second use) to improve critical path timing. In some embodiments the timing driven gate resizing includes morphing non-buffer resources and allocating the morphed resources as “upsized” gates or buffers.

[0427] In various embodiments of timing driven buffering and resizing for structured array architectures, timing driven hold time fixes are implemented by any combination of morphing, delaying clock signals, and buffering. In some embodiments any combination of logic replication and tunneling are used to improve circuit performance of designs implemented according to a structure array fabric.

[0428] Figs. 22A and 22B illustrate, respectively, an example circuit excerpt before and after processing according to an embodiment of timing driven buffering and resizing for an array architecture. Fig. 22A illustrates critical load C_2 driven by buffer b_2 that is driven by buffer b_1 that is in turn coupled to Driver. Thus there are two buffers between the driver and the critical load. Non-critical loads NC_1 and NC_2 are also driven by buffer b_2 . Loads on a critical path from Driver to C_2 include c_0 driven by Driver and C_1 driven by buffer b_1 . Fig. 22B illustrates a result of timing driven buffering and resizing, as applied to the topology of Fig. 22A, where critical load C_2 is driven from new/modified buffer b_1' that is directly coupled to Driver. Thus there is only one buffer between the driver and the critical load, providing enhanced arrival time for the critical load compared to the topology of Fig. 22A.

1 Structured ASIC Timing Closure

2
3 **[0429]** Fig. 23 illustrates a flow diagram of an integrated circuit design flow including
4 an embodiment of processing in accordance with an embodiment of timing driven buffering and
5 resizing for an array architecture, e.g. a structured ASIC.

6
7
8 Timing Driven Buffering

9
10 **[0430]** Fig. 24A illustrates a top-level view of an embodiment of timing driven
11 buffering and resizing for an array architecture. In some usage scenarios timing driven buffering
12 and resizing serves to reduce delays of critical path elements and decrease transition times
13 associated with drivers (or nets or both). Routing-aware buffering is used to reduce maximum
14 congestion in otherwise heavily congested regions.

15
16 **[0431]** In some embodiments an initial buffering phase is performed ignoring timing-
17 driven constraints, while in other embodiments the initial buffering accounts for timing-driven
18 constraints. According to various implementations timing-driven buffering and resizing includes
19 any combination of net prioritization, global Steiner tree routing, evaluating multiple route trees,
20 computing buffering options, pruning, and determining and selecting a solution.

21
22 **[0432]** In some embodiments a buffering subsystem processes nets individually,
23 prioritizing the nets according to timing criticality, enabling preferential treatment for more
24 critical nets. The preferential treatment is according to any combination of buffering resources,
25 wiring resources, and routing congestion (measured according to a metric). In structured array
26 usage scenarios, buffer resources are finite and several nets may be simultaneously competing
27 for the same resources. Ordering nets and processing the most critical nets (or the nets having
28 the highest negative slack) first provides the more critical nets with access to the buffer
29 resources first. In addition, as more nets are processed, the most critical of the remaining nets
30 have access to wire routing regions most beneficial to routing the remaining nets through. Less
31 critical nets are relegated to more meandering routes to meet region congestion constraints.

32
33 **[0433]** In some embodiments the buffering subsystem initially constructs global Steiner
34 tree routes for all nets to estimate heavily congested regions. Routing and/or congestion
35 hotspots that should be avoided while buffering (at least for non-critical nets) are identified.

1 **[0434]** In some embodiments the buffering subsystem initially builds multiple route
2 trees for each driver that couple the respective driver to all loads of the driver. The route trees
3 are heuristic based, and the heuristics include prioritizing critical loads differently than non-
4 critical loads and operating with an awareness of the previously identified hotspots. The route
5 tree building includes any combination of shortest path weight and net spanning factor
6 techniques, enabling results having different topologies.

7
8 **[0435]** In one embodiment of one of the route tree heuristics, loads are first grouped
9 into multiple partitions based on load (or pin) criticality. More critical loads are prioritized for
10 Steiner tree route construction first. Then less critical loads are processed, enabling the more
11 critical loads to have a more direct route from driver to load. In addition, the more critical loads
12 are presented with higher shortest path weight, thus reducing branching of the route tree from
13 the more critical loads to the less critical loads.

14
15 **[0436]** In some implementations a Steiner tree based route is decomposed into several
16 segments, such as according to a global cell granularity used when constructing the Steiner tree
17 based route. A dynamic programming technique is used to compute a buffer solution for each of
18 the route trees. The dynamic technique includes maintaining several solutions for each segment
19 to be considered for use to implement a sub-tree of the respective route tree. The respective
20 route tree is processed bottom-up, i.e. all of the load terminals of the tree are visited before the
21 driver. Buffering options at a segment are computed by combining solutions of all predecessor
22 sub-trees with a current solution.

23
24 **[0437]** Fig. 25A illustrates a portion of a route tree having several branches
25 decomposed into segments according to processing by an embodiment of timing driven
26 buffering. Child options are a function of downstream options. For example:

27 Options at S_0 = Product(Options at S_1 , Options at S_2).

28
29 **[0438]** Fig. 25B illustrates several segment options for segment S_0 of Fig. 25A. The
30 options include no buffering (Opt1), a buffer before the branch to segment S_2 (Opt2), a buffer on
31 segment S_1 (after the branch as Opt3), a buffer on segment S_2 (after the branch as Opt4), and two
32 buffers after the branch, one on each of segments S_1 and S_2 (Opt5).

33
34 **[0439]** If a segment currently being processed is a branch point, then the current
35 segment has multiple sub-trees below it, and each of the sub-trees contains an array of options.
36 The options are merged by performing a cross product of option sets. After computing the cross

1 product, each feasible solution for the sub-tree is combined with a buffering solution for the
2 current segment.

3
4 **[0440]** Multiple segment options are computed for each segment. The number of
5 options produced is proportional to the number of buffer types (or buffer electrical
6 characteristics) available according to technology associated with an integrated circuit design
7 (such as a standard cell library). In some implementations various options are computed for
8 each segment, including a non-buffered option, a high-drive strength buffer option, and a low-
9 drive strength buffer option.

10
11 **[0441]** For each option, several parameters are determined, including Buffer Cost (BC),
12 Required Time (RT), and lumped Capacitive Load (CL). The parameters are subsequently used
13 to determine option cost and feasibility. BC measures cost according to the buffering solution
14 for the entire sub-tree “underneath” the segment being evaluated. RT measures expected
15 required time for a signal at the input of the segment. CL measures cumulative capacitive load
16 of the segment and all associated child segments.

17
18 **[0442]** Pruning techniques are used to limit computation, maintaining selected options
19 for each route segment. The selected options chosen are those most likely to result in a “good”
20 solution according to the root of the route tree. A first pruning technique includes deleting any
21 infeasible solutions, such as a buffering option that has accumulated capacitance exceeding the
22 maximum drive capability according to available buffers. A second pruning technique removes
23 redundant options. An option having higher BC and smaller RT, or an option having higher BC,
24 smaller RT and higher CL, compared to an existing option, is considered redundant. A third
25 pruning technique includes trimming the number of options according to an upper bound. In
26 some embodiments the upper bound is variable, while in other embodiments the upper bound is
27 predetermined (at a value such as 10, 20, 50, or 100). In some implementations the options are
28 sorted in order of RT (highest RT first). In some embodiments a contiguous portion of the top of
29 the sorted options is retained, the portion being equal in number to the upper bound (i.e. the
30 “best” options are kept). In some embodiments the sorted options are partitioned into four
31 quarters, and a number of options are preserved from each quarter. In some embodiments the
32 number is chosen to be one-fourth of the upper bound. In some usage scenarios the preserving
33 according to partitions enables discovery of solutions that appear locally inferior, but when
34 combined with parent segments appear superior.

1 **[0443]** In some embodiments determining and selecting a buffering solution includes
2 evaluating options according to performance (such as arrival time) and (estimated) routing
3 congestion. A disproportionately higher weighting is applied to timing cost when evaluating a
4 critical net. A buffering solution having lower hotspot (i.e. congestion) cost is preferentially
5 chosen for non-critical nets.

6
7
8 Timing Driven Sizing
9

10 **[0444]** Fig. 24B illustrates a detail view of selected details of an embodiment of timing
11 driven resizing for an array architecture. Timing-driven form sizing (or resizing) selects
12 alternate forms to improve any combination of drive capability and stage delay, for example by
13 replacing a lower drive strength gate with a relatively higher drive strength gate. In some usage
14 scenarios macro or form sizing is preferred over buffering when cost of upsizing a driver is less
15 than buffering a net. In some structured ASIC usage scenarios buffer sites are predetermined
16 according to block tiles, and thus the fixed locations of buffer sites may result in relatively high
17 intrinsic buffer cost or associated congestion cost. In some situations there may be no available
18 sites (or slots) near a macro targeted for resizing.

19
20 **[0445]** In some embodiments a form-sizing subsystem attempts to discover nearby sites
21 by (re)implementing the macro using a different set of primitives. According to various
22 embodiments the primitives correspond to standard cells, structured array tile elements, or other
23 similar low-level resources. In some implementations the form-sizing subsystem is enabled to
24 “displace” (or “move”) selected forms (such as forms on non-critical paths) that are initially near
25 the macro that is to be resized. In structured array integrated circuit designs, strictly speaking
26 the forms are not moved, but instead fixed-location sites are deallocated in one area and
27 reallocated in another area.

28
29 **[0446]** A Displacement Coefficient (DC) of a macro is computed as follows:
30 DC of macro = Sum (DC of each morphable form within the macro); and
31 DC of a morphable form = Product(primitive densities of all the primitives within the
32 morphable form).

33
34 **[0447]** The DC is a conceptual measurement of “placeability” or ease of placement of
35 an element when the element is currently unplaced. A macro is more placeable if it may be
36 implemented with more morphable alternatives. A morphable alternative is more placeable if

1 the primitives of the morphable alternative are placeable (or relatively more placeable), such as
2 when there are available (or unused) sites for the primitives.

3
4 **[0448]** The primitive densities relating to the DCs of morphable forms are computed as
5 follows. A site density grid is constructed that is a two-dimensional matrix of grid resource
6 usage. For each element of the density grid, a number of available resources and used resources
7 are computed for each resource type. Relatively sharp density gradients are smoothed by
8 accumulating density from eight neighboring grid elements to a respective grid element. Thus
9 the computed density at each grid element is an average density at the element in conjunction
10 with eight nearest neighboring elements. The site density grid values are then used to determine
11 the DCs of the morphable forms.

12
13 **[0449]** The DC of a morphable form is computed by looking up the density of each of
14 the primitives of the morphable form, within the site density grid and according to respective
15 primitive types. The morphable form DC computation continues by multiplying the look up
16 results (i.e. primitive densities) together. If a particular resource or resource type is depleted (or
17 nearly depleted) within the grid, then the morphable form DC is zero (or nearly zero). Thus the
18 resource depletion results in the placeability of the morphable form being low.

19
20 **[0450]** Resizing a macro includes selecting a form from a plurality of implementation
21 choices. Each of the choices is speculatively selected and evaluated with respect to the macro
22 being resized. A timing score is computed that is equal to arrival time at an output of the macro
23 assuming the macro is implemented with the speculatively selected form. If the timing score is
24 poorer than previously saved possible implementation choices, then the current choice is
25 rejected. If the timing score is better, and the drive strength of the speculatively selected form is
26 sufficient to drive the capacitive load at the output, then the speculatively selected form is saved
27 as a possible implementation choice.

28
29 **[0451]** In some embodiments placing a macro after determining an implementation
30 according to one or more morphable forms proceeds as follows. New coordinates of the (now
31 form-level) macro are computed based on all of the connections of the form-level macro. The
32 coordinates of drivers of nets connected to all of the input pins of the form-level macro as well
33 as associated net fanouts are used to compute the new coordinates.

34
35 **[0452]** In some embodiments a form placing sub-system performs an attempted
36 placement of each of the possible implementation choices determined during the resizing of the

1 macro. The underlying morphable forms are already prioritized based on the respective timing
2 scores (based on an idealized placement), and the attempted placements are, in various
3 embodiments, performed in priority order (i.e. morphable forms resulting in better arrival times
4 are tried first). Unplaceable morphable forms are not considered further. After a respective
5 placement is found for each placeable one of the morphable forms, the placed morphable form is
6 scored based on timing in the respective placement. After attempting to place and after scoring
7 all of the morphable forms, the one of the morphable forms with the best score, if any, is
8 selected. In some embodiments, if there is no placeable one of the morphable forms, the
9 window size is increased and attempted placement of the morphable forms is repeated.
10 Attempted placement of one of the morphable forms succeeds if individual attempted
11 placements of each of the respective primitives of the morphable form all succeed. Attempted
12 placement of a particular one of the respective primitives of a particular one of the morphable
13 form proceeds as follows.

14
15 **[0453]** A site locator (or slot locator) searches all possible sites around a given
16 coordinate within a certain window size and returns a list of all sites within the window
17 assignable to the particular primitive. In some embodiments, the list is sorted in Manhattan
18 distance from the given coordinate. The list is then examined. In some embodiments, a first
19 acceptable site is selected. In other embodiments, all sites in the list are processed and scored,
20 such as by scoring on DC, and the best-scoring acceptable site is selected. In some
21 embodiments, an acceptable site is one that has a respective DC above a threshold. In various
22 embodiments, unoccupied sites have a respective DC above the threshold, such as a DC of 1.
23 The respective DC of an occupied site is obtained by looking up the DC of the parent form-level
24 macro of the occupied site. If an occupied site is selected, then the parent macro is tentatively
25 scheduled to move from the site (i.e. a primitive in the site will be placed elsewhere) and the
26 particular primitive is tentatively assigned to the site. The tentative move of the parent macro
27 and the tentative assignment of the particular primitive to the site are actualized if the particular
28 morphable form is selected as the one of the morphable forms with the best score. Any parent
29 macros that are scheduled to move are queued to be visited later based on criticality of the parent
30 macros.

31
32
33 Timing Driven Logic Replication and Tunneling
34

35 **[0454]** In some embodiments a driver node is logically replicated for nets having high
36 capacitive loading or high fanout. The replication is performed selectively according to

1 evaluations of timing improvements and routing costs. In some embodiments tunneling is
2 performed to move the driver closer to a cluster of loads. In some embodiments the tunneling is
3 performed after evaluating the timing improvements and routing costs.

4
5 **[0455]** Fig. 26 illustrates example results of an embodiment of logic replication and
6 tunneling for an array architecture. The example illustrates a single FF driving three clusters of
7 load (C_1 , C_2 and C_3). After replication and tunneling (shown in the lower portion of the figure),
8 the FF is replicated as FF_1 , FF_2 , and FF_3 . Each of the replicated FFs is then placed near the
9 respective cluster driven by the FF.

10
11
12 Timing Driven Hold Time Fixes

13
14 **[0456]** In some embodiments timing driven hold time fixes proceed as follows. Excess
15 setup time (or slack setup time) is determined for each launch FF that is a root of a hold time
16 violation. If there is excess setup time, then in some embodiments the clock signal feeding the
17 launch FF is delayed. In some implementations the delay is via addition of a dummy load. In
18 other embodiments a hold time violation is addressed by morphing the launch FF to a slower FF.
19 In some implementations the morphing is via swapping the (original) launch FF with an unused
20 (or available) slower FF.

1 DENSITY ENHANCEMENT EMBODIMENTS

2
3 [0457] Node density in various SDI embodiments is influenced by a variety of effects,
4 including netlist connectivity, circuit performance, and expansion fields. The former two exert
5 an attractive force between nodes that depends upon netlist topology considerations or critical
6 path analysis. For brevity these are referred to as "connectivity forces". Without the presence of
7 expansion fields, the connectivity forces tend to draw together nodes into a highly clumped
8 configuration that may exceed local slot resource supply. Spreading of nodes by the expansion
9 fields then serves a twofold purpose: (a) provide solutions to slot assignment over some suitably
10 chosen sub-discretization of a die, and (b) enhance routability, since localized clumping of nodes
11 implies greater local demand for routing resources.

12
13 [0458] In a chip floorplan that is free of obstructions, very strong expansion fields
14 result in a node distribution that is almost perfectly uniform. However this situation may not be
15 desirable, since some amount of clumping may be beneficial. Once the node distribution
16 reaches the point of routability and is compatible with local resource supply, further increases to
17 the expansion field strength may only worsen the routing problem by forcing nodes further apart
18 than is optimal, seen by examining cutscores or circuit performance as a function of expansion
19 field strength.

20
21 [0459] Further, the demand for routing resources may exceed supply only in very
22 localized regions, while the bulk of the node distribution presents a tractable routing problem.
23 The localized regions may occur due to netlist (topological) or floorplan effects. Increasing the
24 expansion field strength to compensate for the "clumpiness" of the node distribution in selected
25 regions affects the distribution as a whole, and in some usage scenarios may be suboptimal. In
26 cases where the floorplan contains obstructions, the supply of routing resources can be a
27 complex function of position on the die, and here a global approach can fail to have the desired
28 effect entirely.

1 **[0460]** The illustrative but not limiting density-driven approaches presented here for
2 addressing the problem of routing congestion in SDI can be categorized as follows:

3 1. Region based

4 a. By factor

5 b. By function

6 2. Steiner cuts based

7 a. Relative

8 b. Absolute (i.e. supply vs demand)

9
10 **[0461]** In the illustrative density enhancement embodiments, the density enhancement
11 is inserted between the "fold" (such as "Fold Guard Cell Contributions" **332** of Fig. 3C) and
12 "filtering" (such as "Digitally Filter Fields **336** of Fig. 3C) phases of node density computation.

13
14 **[0462]** The flow **27,100** for density modification is illustrated in Figure 27. Note
15 effects introduced by procedures **27,100b**, **27,100c**, and **27,100d** are completely independent of
16 each other and can therefore be applied in any combination.

17
18 **[0463]** In procedure **27,100a**, the normalization factor is typically taken as the average
19 density, not counting that in excluded regions.

20
21 **[0464]** In procedure **27,100b**, for each defined region that possesses a density scaling
22 factor, the density is multiplied by the associated factor at each density field grid point contained
23 within the region. Note this technique is essentially the same as increasing the effective mass, as
24 pertaining to the node density, for each node contained therein.

25
26 **[0465]** Given a statistically uniform node distribution to start with, the scale factor
27 density enhancement tends to drive nodes out of the specified region, ultimately resulting in a
28 node density on the order of (average density) / (scale factor) there, edge effects
29 notwithstanding. Any number of such regions and scale factors can be defined. Regions may
30 overlap if so desired.

31
32 **[0466]** In procedure **27,100c**, for each defined region that possesses a density
33 enhancement function, the associated spatial variation multiplied by the normalization factor is
34 added to the existing density. The spatial variation is evaluated at each density field grid point
35 contained within the region. In some embodiments an arbitrary functional variation is supported

1 by expressing the function in Tcl (Tool Command Language) and using an embedded Tcl
2 interpreter to return the result of the given expression at each grid point.

3
4 **[0467]** The functional variation enhancement may be well suited for the case where the
5 node density becomes very rarefied, e.g. in small channels between obstructions. In rarefied
6 density situations, the scale factor approach becomes less effective for pushing nodes out of the
7 region, since there are fewer nodes to "push against". The functional variation serves essentially
8 as a background density, only depending on the existing node density through the normalization
9 factor (which is global). As in procedure **27,100b**, there is no limit to the number of regions and
10 functions that can be defined, and regions may overlap if desired.

11
12 **[0468]** In procedure **27,100d**, a Steiner-cuts congestion density enhancement term is
13 added. At this point in the flow, for this density enhancement embodiment, a congestion
14 enhancement value at each grid point is available (described in detail below). Adding the
15 congestion enhancement term (suitably normalized or scaled, for example times a suitable
16 normalization factor, e.g. the average density) for each grid point gives a final result.

17
18 **[0469]** The flow **28,200** used to determine the Steiner-cuts congestion term on the SDI
19 grid in the density enhancement embodiment is given in Figure 28.

20
21 **[0470]** In procedure **28,200a**, a so-called "congestion array" is generated that is a
22 measure of routing congestion output, taken from a Steiner cuts measurement. Since the
23 calculation of routing congestion may be computationally expensive, the congestion array need
24 only be calculated initially and at specified intervals as a run proceeds. An intermediate grid is
25 used to assert the independence of the congestion array from the particular form of the routing
26 congestion diagnostic, as well as from the SDI grid resolution. The procedures used to create
27 the congestion array are illustrated in Figure 29.

28
29 **[0471]** In procedure **28,200b**, the congestion array is run-length averaged according to
30 a specified relaxation factor. This helps prevent sudden "shock" to the system (which can cause
31 unphysical fluctuations) every time the congestion array is recalculated, by phasing the change
32 in gradually. The relaxation parameter is chosen to vary from zero (static; congestion array
33 never changes) to unity (congestion array changes instantaneously).

1 **[0472]** In procedure **28,200c**, a final congestion density enhancement array is
2 calculated. The calculation may be performed once each timestep, in response to configuration
3 changes, or both. Further details are illustrated in Figure 30.

4
5 **[0473]** In procedure **29,300a**, the Steiner-cuts array is fetched from the generator. In
6 some embodiments a timing kernel (TK) performs procedure **29,300a**. The calculation may
7 include an idealized buffer tree, at implementor or user discretion.

8
9 **[0474]** In procedure **29,300b**, the Steiner-cuts array is subject to a filtering operation to
10 increase smoothness, which helps improve accuracy of a subsequent interpolation procedure. In
11 some embodiments a number of binomial digital filter passes are used.

12
13 **[0475]** In procedure **29,00c**, the value at each grid point in the intermediate grid
14 discretization is calculated using a linear spline approach.

15
16 **[0476]** In procedure **30,400a**, the congestion array is smoothed using filtering similar
17 procedure **29,300b**, in part to improve the accuracy of the interpolation. But filtering is also
18 considered the "final smoothing" phase of the field and is subject to the most user and/or
19 programmatic control, to improve the quality of the final result. The smoothing is most effective
20 when the scale lengths associated with the variation of the density enhancement are "semi-
21 global", e.g. small compared to the die size, but large compared to motion of a node in a single
22 timestep.

23
24 **[0477]** In procedure **30,400b**, the congestion array is normalized as needed. First it is
25 clipped at a pre-determined value of maximum congestion, to constrain resulting density
26 gradients within reasonable limits. In relative spreading mode, a normalization of unity is
27 imposed, thus inducing a density-driven outflow from congested areas without regard to actual
28 routing supply.

29
30 **[0478]** In absolute spreading mode, the routing demand versus supply is compared to
31 the maximum allowable relative demand (e.g. 80% capacity). Only at grid points where
32 congestion exceeds the allowed limit does the enhancement field take on substantial values
33 (while enforcing a reasonably smooth variation). In the case of a density-gradient model for
34 calculating the expansion fields, the congestion density field that results is flat everywhere that
35 routing supply satisfies demand, rising smoothly into elevated "mounds" at locations where the
36 demand exceeds supply.

- 1
- 2 **[0479]** The congestion array is then modified according to desired strength of the
- 3 density enhancement effect. Both multiplicative and power-law transformations may be applied.
- 4 The strength of the enhancement may be increased over time to allow for the gradual movement
- 5 of nodes out of congested areas.
- 6
- 7 **[0480]** In procedure **30,400c**, the value of the congestion array at each SDI grid point is
- 8 calculated using a linear spline approach.

1 TUNNELING CONGESTION RELIEF EMBODIMENTS

2
3 **[0481]** In some SDI-based integrated circuit design flow embodiments “tunneling” is
4 used to relieve congestion at boundaries. Tunneling governs transitions of nodes through one or
5 more obstructed regions not available for node placement, i.e. xzones, of a chip (or portion
6 thereof). In some embodiments the transition is according to a mathematical criterion. In some
7 embodiments nodes are selected as tunneling candidates according to node attraction into one of
8 the obstructed regions. In some embodiments the criterion is affected by node density. In some
9 embodiments the criterion is affected by node interconnections (or connectivity). In some
10 embodiments the criterion is affected by circuit performance (i.e. timing).

11
12 **[0482]** Tunneling enables further placement progress, according to selected metrics
13 such as routability and circuit performance, while taking into account xzones. Tunneling has
14 several aspects including candidate node selection, nodal move speculation, and node tunneling
15 criteria (i.e. keep move/tunnel or reject).

16
17 **[0483]** In some embodiments tunneling is performed at the end of an SDI timestep.
18 Any intervening sub-steps taken by the time integrator (e.g. part steps taken by a Runge-Kutta
19 (RK) integrator) are not considered. During the course of a timestep (and any associated sub-
20 steps) the nodes are allowed to drift into xzones to allow the time integrator to proceed at full
21 speed, since in some usage scenarios a smooth trajectory in a numerical simulation enables more
22 accurate integration, and thus may enable a longer timestep (given a predetermined accuracy
23 target). At the end of one full timestep, only nodes that have been coerced into xzones are
24 considered for tunneling speculation.

25
26 **[0484]** Fig. 31 illustrates an embodiment of a processing flow for node tunneling
27 through exclusion zones in an SDI-based integrated circuit design flow. In some
28 implementations any combination of the illustrated elements are performed by software routines
29 known collectively as a “node mover”. In **31,100a** nodes are selected as candidates for
30 tunneling based on respective positions. Nodes that have moved into an xzone are included in a
31 set of all transiting nodes. Each respective node will have arrived at the respective position (or
32 point) due to (discounting inertial effects) the vector sum of all forces acting on the respective
33 node. For example, some of the forces may be due to netlist connectivity (i.e. the respective
34 node is drawn towards topologically close nodes) and some of the forces may be due to a local
35 overabundance of nodes (density buildup). In some usage scenarios, selecting nodes in xzones

1 for tunneling consideration is an efficient selection criteria that discriminates nodes likely to
2 benefit from a tunneling transition to another side of an xzone or multiple abutting xzones.

3
4 **[0485]** In **31,100b**, having determined candidate nodes, per-node initialization is
5 performed. In some usage scenarios, total tunneling candidate nodes are a small fraction of total
6 nodes, and for efficiency a secondary set of data structures is used to process the candidate
7 nodes. A transiting node class contains a node id (that maps to an original node entry) and any
8 ancillary data required for accurate tunneling speculation. Henceforth, the class of all node
9 candidates for tunneling is referred to as "transiting nodes".

10
11 **[0486]** In **31,100c**, all transiting nodes are snapped to the nearest xzone boundary. The
12 snapped position is identical to the resulting node position were no tunneling to occur, to provide
13 a baseline for field computation and comparison to the post-transit result.

14
15 **[0487]** In **31,100d**, the forces on transiting nodes at the current positions (pre-
16 speculation) are evaluated. See the discussion relating to Fig. 32 located elsewhere herein for
17 further information.

18
19 **[0488]** In **31,100e**, the position of the transiting node is restored to the positions before
20 processing relating to **31,100c**. The node mover then finds the intercept on the xzone boundary
21 that results from application of the force vector components on the node. In some embodiments
22 node inertia is also taken into account when determining the xzone boundary intercept. The
23 node is speculatively moved to just past the intercept position, outside the original xzone. In the
24 event that multiple abutting xzones exist and the node lands in yet another xzone, the mover is
25 invoked again using the original trajectory to direct the move. The speculative movement
26 procedure is continued as many times as necessary for the node to arrive in a region external to
27 any xzone.

28
29 **[0489]** In **31,100f**, the forces on transiting nodes at the new positions (post-speculation)
30 are evaluated. See the discussion relating to Fig. 32 located elsewhere herein for further
31 information.

32
33 **[0490]** In **31,100g**, the transition criteria are evaluated and examined. If the transition
34 is accepted, then the node associated with the transiting node acquires the new coordinates.
35 Otherwise the coordinates as determined in **31,100c** are retained. See the discussion relating to
36 Fig. 33 located elsewhere herein for further information.

1

2 [0491] Fig. 32 illustrates an embodiment of SDI-related force calculations in a
3 tunneling congestion relief context. In 32,200a, forces on the node are cleared and preparations
4 are made for the field calculation.

5

6 [0492] In 32,200b, forces on each node due to all non-field interactions are summed,
7 including all connectivity and timing based pin to pin forces, as well as any other nodal
8 interaction forces present.

9

10 [0493] In 32,200c, gate field components are computed. The first time through (pre-
11 speculation phase), a full field calculation is performed. The pre-speculation phase is with the
12 nodes snapped to the nearest xzone boundary, so the result represents a result assuming no nodes
13 transit. The second time through (post-speculation phase), the field calculation from the first
14 phase is used, but applied to the speculative nodal coordinates. That is, it is assumed that the
15 fields are not significantly changed on a global scale as a result of tunneling. In some usage
16 scenarios, since only a small number of transitions are considered relative to the total number of
17 nodes, the assumption serves as a reasonable approximation, and may be beneficial for
18 computational efficiency since field computations for each individual speculation are avoided.

19

20 [0494] Fig. 33 illustrates an embodiment of evaluation of tunneling transition criteria.
21 In 33,300a, the speculative node coordinates are examined to see if there are violations of any
22 node region constraints and if nodes fall into a legal logic area. If there is any violation, then the
23 transition is rejected.

24

25 [0495] In 33,300b, a statistical window on how many transitions are considered is
26 applied. In some embodiments, the window is small (such as 1%, 2%, 5%, or 10%) compared to
27 unity, but not so small that an inordinate number of passes through speculator routines are
28 required to process all qualifying nodes. The windowing helps prevent sloshing, where many
29 nodes tunnel from a high to a low density region at once, altering density so much that nodes
30 tunnel back later. In other words, the statistical window helps to ensure that approximations
31 made with respect to 32,200c (of Fig. 32) are valid.

32

33 [0496] In 33,300c, a variety of biasing factors are applied. In some embodiments, the
34 factors are multiplied together, producing a resultant bias. If the resultant bias is less than unity,
35 then transiting is more likely (e.g. encouraged). If the resultant bias is greater than unity, then
36 transiting is less likely (e.g. discouraged). According to various embodiments, one or more of

1 the factors is less than unity, or one or more of the factors is greater than unity. The factors
2 include any combination of the following.

- 3 • A default biasing factor.
- 4 • A bias against multiple transitions in a row, to ensure longer relaxation time.
- 5 • A distance based biasing, to make it more difficult to travel long distances. The distance
6 based biasing may involve either a hard limit or a functional dependence on distance
7 traveled (e.g. linear or quadratic).
- 8 • A distance based biasing specific to timing critical nodes. Nodes on a critical path may
9 have an unpredictable effect on timing due to tunneling, so the critical path nodes may
10 be selectively more further constrained than other nodes.

11

12 **[0497]** In **33,300d**, the magnitude of the forces on the node at the old and the new
13 positions are computed. If the new force magnitude after biasing is less than the old force
14 magnitude, then the transition is considered to be energetically favorable and therefore accepted.
15 Otherwise the transition is rejected.

CLOCK TREE SYNTHESIS (CTS) EMBODIMENTS

[0498] CTS is a process for creating a clock network in an Integrated Circuit (IC) physical design flow. CTS has general applicability to design flows having limited placement options for clock buffers, such as SAF-based design flows. Note that although CTS is described herein within a general context of an SDI-based flow, there are applications to other types of design flows using conventional EDA tools. In some usage scenarios a structured ASIC design has one or more clock signals that fan out to many (perhaps thousands) of register clock pins. A register clock pin may be a clock pin of a flip-flop, a latch, or clock pins of embedded memory and other IP blocks.

[0499] Clock nets produced by logic synthesis or derived from schematics act as placeholders for CTS-produced clock nets. Each of the logic synthesized clock nets drives a high drive strength buffer (an ideal clock). Each of the CTS-produced clock nets includes one or more levels of buffers, interconnect wires, and other gating logic such as clock_and, clock_or, clock_mux, and other similar clock manipulation elements. In some embodiments CTS is run post placement so that precise coordinates of clock pins driven by each clock net are known (such as portions of processing performed in conjunction with “Buffering Clock Tree Synthesis Timing Driven Buffering/Resizing” **821** of Fig. 8A).

[0500] In some implementations a CTS tool builds a clock network that strives to optimize characteristics of the clock network including skew and latency. Clock skew is the difference of signal arrival times at clock pins of two registers. The CTS tool optimizes a maximum clock skew of the circuit, i.e. the largest clock skew between any pair of registers that have timing paths (setup/hold) between them is minimized.

[0501] Clock latency is delay from a root of a clock tree to a clock input pin of a register. The CTS tool optimizes the maximum latency, i.e. the largest delay is minimized. In addition to skew and latency metrics, there are other considerations such as power and routing congestion addressed by the CTS tool. The CTS tool attempts to optimize (i.e. minimize) the buffers and wire resources used for clock distribution since the resources directly impact circuit routing congestion and dynamic power usage.

[0502] In some embodiments CTS is performed in a post detail placement phase to enable building optimized clock networks, based on locations of clock leaf pins. Gating logic enables power reduction by selectively turning on and off sub-trees of a clock tree. Clock

selector logic (such as using a clock_mux) multiplexes multiple user clocks and test clocks. A clock tree may have several levels of clock selector logic gates and several levels of clock gating logic gates. In some usage scenarios clock gating circuitry is pre-designed by the user at a hardware description level and is then synthesized into gates by a synthesis tool. The CTS tool balances clock networks while taking into consideration delays of various gating logic, thus treating the gating logic transparently and automatically.

[0503] Fig. 34A illustrates an example clock tree suitable for input to a CTS tool for SAF-based design flows. Primary clock domains are illustrated as pclk0 and pclk1. Gated clock sub-domains are illustrated as gclk0 and gclk1. A clock selector based clock sub-domain is illustrated as mclk. Clocks pins of registers are illustrated as ckp0, ckp1, ... ckpn; kkg0, ... ckgn; cks0, cks1, ... cksn; and cksg0, ... cksgn. Register clock pins ckp0, ... ckpn and cksg0, ... cksgn are associated with gated clocks. Register pins cks0, cks1, ... cksn are associated with selected clocks. Register clock pins cksg0, ... cksgn are associated with two levels of clock processing (select and gate functions).

[0504] Fig. 34B illustrates an example clock tree output from the CTS tool operating on the input illustrated in Fig. 34A. In the illustrated output various Clock Networks produced by the CTS tool (according to the input illustrated by Fig. 34A) are shown driving the register clock pins.

[0505] Fig. 34C illustrates an example clock tree network. Leaf buffers are illustrated as b_1 , b_2 , b_3 , and b_4 . Each of the buffers are shown driving (or fanning out to) a respective sea of clock pins as illustrated conceptually by the triangular element at each respective driver output. Terminals of the clock network are illustrated as t_1 , t_2 , and t_3 . Selected terminal buffers are illustrated as tb_1 and tb_2 . A clock root is illustrated as CT. The illustrated clock tree network is representative of some implementations of the Clock Networks of Fig. 34B. For example, consider the Clock Network of Fig. 34B driving register clock pins ckp0, ckp1, ... ckpn. CT of Fig. 34C corresponds to the element driving pclk0. Leaf buffer b_1 drives ckp0, leaf buffer b_2 drives ckp1, and so forth.

[0506] Fig. 35 illustrates an overview of an embodiment of a CTS flow. According to various embodiments the CTS flow includes any combination of floorplan driven clock partitioning, topological clock sorting, top-down recursive bi-partitioning, clock domain (and sub-domain) processing, bottom-up skew minimization, and top-down buffer placement.

1 **[0507]** Floorplan driven clock partitioning (such as illustrated in Fig. 35) may be used
2 when a die floorplan has extensive arrays of RAM and/or IP structures that lack suitable sites or
3 slots for clock tree buffer elements. When the CTS tool builds a clock tree, buffer sites at
4 intermediate points of each clock network are used to drive two sub-trees “underneath” the
5 respective intermediate point. Having large rows(columns) of RAMs/IP blocks implies that
6 there are extensive die regions that are either completely devoid of clock buffer sites or have the
7 sites at sub-optimal locations. Therefore, CTS preprocesses the clock network and embeds
8 Pseudo-clock Sub-Domains (PSDs) that are first balanced within each row(column).
9 Subsequently, the clock sub-domains are deskewed across logic rows(columns). The first level
10 PSDs can be deskewed by buffer resources within a row(column), thus alleviating the need to
11 find sites over RAM and/or IP regions.

12
13 **[0508]** Fig. 36A illustrates an example die floorplan of a design having embedded
14 RAM or other IP blocks. Regions **36,300a** represent an I/O ring. Regions **36,300b1**, **36,300b2**,
15 and **36,300b3** represent rows of embedded RAMs. Regions **36,300c1**, **36,300c2**, and **36,300c3**
16 represent rows of logic blocks. CTS clock preprocessing proceeds as follows. Within each
17 PSD, all clock leaf pins in each contiguous logic region (such as each of regions **36,300c1**,
18 **36,300c2**, and **36,300c3**) are merged so the leaf pins fan out from a single Root Clock
19 row(column) Buffer (RCB). The RCB is optimally placed at the centroid of the bounding-box
20 encompassed by all the leaf clock pins within the respective logic region.

21
22 **[0509]** All RAM clock pins are then combined with logic clock pins by drawing a
23 partitioning line through the middle of each RAM region. For example, if there are RAM clock
24 pins in region **36,300b2**, then each one is merged with clock pins of one of adjacent regions
25 **36,300c1** or **36,300c2** depending on proximity of the respective RAM clock pin to the adjacent
26 regions (i.e. the closest one of the regions is chosen).

27
28 **[0510]** Then each of the region PSDs are deskewed individually. In some usage
29 scenarios the deskewing is by combining even and odd row(column) RCBs separately. In other
30 words, every other row(column) is processed together. In situations where RAM (or IP)
31 rows(columns) are alternated with logic block rows(columns), and the rows(columns) are of
32 approximately the same size, processing even/odd rows(columns) separately simplifies
33 equidistant placement of RCB pairs, since the center of each RCB pair will be in a logic block
34 row(column). For example, RCBs associated with region **36,300c1** are processed with RCBs
35 associated with region **36,300c3**, and equidistant placement may be satisfied by region **36,300c2**,
36 a logic region. Note that the RCBs associated with a logic region may include RAM clock pins

1 from adjacent RAM regions, such as region **36,300c1** RCBs include merged RAM clock pins
2 from the upper half of region **36,300b2** and the lower half of region **36,300b1**.

3
4 **[0511]** Subsequently, the even and odd RCBs are deskewed at the clock root. The
5 aforementioned merging, partitioning, and RCB placement processing is performed for each
6 primary clock. The leaf clock pins driven by gated-clocks and clock selectors cells are treated
7 transparently during the processing. If a gated-clock or clock-selector logic drives leaf clock
8 pins in multiple logic regions, then the gating logic is replicated in each of the respective regions
9 the gated clock fans out to, thus enabling transparent merging of un-gated and gated-clock leaf
10 pins.

11
12 **[0512]** Fig. 36B illustrates a portion of a clock net in a context of a portion of Fig. 36A.
13 Clock net "clk" feeds both un-gated and gated clock pins that span out to logic regions **36,300c1**
14 and **36,300c2**. The gated clock is replicated in region **36,300c2** so that the RCB in each region is
15 enabled to independently drive both the un-gated and the gated branches of the clock trees. The
16 replication technique reduces multi-level clock balancing across RAM regions and introduction
17 of skew uncertainties.

18
19 **[0513]** Topological clock sorting, or domain ordering (such as illustrated in Fig. 35) is
20 performed so that the CTS tool visits the clock domains in an order that facilitates deskewing of
21 lower level sub-domains prior to higher level sub-domains. In some embodiments various clock
22 sorting functions are performed by software routines implemented in a topological sorter.

23
24 **[0514]** In some usage scenarios a primary clock has several gated-clock and select-
25 clock logic based sub-domains. As shown in Fig. 34A, main clock (clk) fans out to several leaf
26 level clock pins after several levels of gating (gclk0, mclk, and gclk1).

27
28 **[0515]** The sub-domains gclk0, mclk, and gclk1 carry the same primary clock (clk), but
29 are gated (controlled) by user logic to enable selectively turning off for one or more clock
30 cycles. Clock distribution information of Fig. 34A is processed by the topological sorter to
31 produce sub-domain ordering: gclk1 -> mclk -> gclk0 -> pclk0 -> pclk1 -> clk. The ordering
32 ensures that when the un-gated leaf level pins of clk nets are being deskewed with the gated-
33 clock pin (gclk0), the gated clock pin has already been completely processed (expanded) and any
34 associated clock network latency is determined.

35

1 **[0516]** Clock domains (and sub-domain) processing (such as illustrated in Fig. 35)
2 includes processing the domains according to the topological domain ordering. A Clock
3 Domain Processor (CDP) of the CTS tool first collects all clock pins to be deskewed. A user
4 may mark pins to be excluded for deskewing and the CDP obeys the marking. The CDP forms
5 leaf level clusters, and, in some embodiments, subsequent recursive bipartitioning operates on
6 buffers of the leaf level clusters (rather than directly on clock pins). For all the leaf clock pins
7 that are pins of a leaf level register (such as flipflops, latches, and RAM blocks), leaf level
8 clustering forms bottom-up clusters that may be driven by a leaf level clock buffer.

10 **[0517]** Clustering of leaf level clock pins (such as illustrated in Fig. 35) is performed
11 via a partitioning of all the leaf level clock pins, and forms small well-formed clock pin clusters
12 that may be driven by leaf level clock buffers, thus reducing complexity of leaf level clock
13 buffer implementation. The partitioning uses recursive bipartitioning with an objective function
14 that minimizes the diameter of the polygon formed by all pins in a partition. In some
15 embodiments, the leaf level clustering operates similarly to a subsequent recursive
16 bipartitioning.

18 **[0518]** As the diameter of the polygon computation has polynomial complexity, in
19 some implementations a fast heuristic technique with linear complexity is used. The linear
20 complexity technique computes an NSP of a bounding box of all leaf level pins in a partition.
21 Clusters are also formed to increase relative “closeness” to other clusters having common setup
22 and hold paths. Cluster closeness of two clusters is the number of clock buffers common to the
23 clusters. In other words, tightly coupled leaf clock pins are grouped to share relatively many
24 clock buffers, thus enabling more efficient skew reduction.

26 **[0519]** Fig. 37A illustrates an example of timing driven pin swapping. As illustrated, it
27 is preferable to partition clusters as $P1 = \{La, Ca\}$, $P2 = \{Lb, Cb\}$ instead of $P1 = \{La, Lb\}$ and
28 $P2 = \{Ca, Cb\}$. The former promotes sharing of clock buffers between launch and capture flip-
29 flops thereby reducing the skew between launch and capture flip-flops since unshared clock
30 buffers may be subject to separate process, voltage, and temperature variations and thus may
31 introduce skew. According to various embodiments, timing driven pin swapping is performed
32 either before or after recursive bipartitioning.

34 **[0520]** In some embodiments, timing driven pin swapping is performed as pairwise
35 swapping of edge pins based on timing relationships of the pins within respective clusters. The
36 swapping is directed to achieve maximal common launch and capture paths for a pair of clock

pins that have either a setup path or a hold path in common. For example, a pin of a first cluster that is close to a second cluster (so as to not increase the interconnect cost of the second cluster by more than a specified amount) is considered for swapping with a pin of the second cluster that is close to the first cluster, if the swap improves a timing cost by a specified threshold. In some embodiments, pins are considered for pairwise swapping if swapping the pins improves setup and/or hold time constraints, such as by reducing skew. In various embodiments, performing timing driven pin swapping subsequent to recursive bipartitioning provides information as to skew between clusters, based on results of the recursive bipartitioning.

[0521] During recursive bipartitioning, each partition is scored based on timing relationships between each pin and every other pin of the partition. Cluster cost is a weighted sum of interconnect wiring cost and cluster-closeness cost. The interconnect wiring cost is determined from the NSP of the bounding box of all the pins constituting the cluster. For example, partition cost may be given by:

$$\text{Part_cost} = 0.5 * \text{cic} * \text{cic} + 0.5 * \text{ctc} * \text{ctc}$$

where

cic: is cluster interconnect cost, given by

$$\text{cic} = (1 - \text{part_interconnect_cost} / \text{best_interconnect_cost}); \text{ and}$$

ctc is cluster timing cost given by

$$\text{ctc} = (1 - \text{part_timing_cost} / \text{best_timing_cost}).$$

In some embodiments, the partitions are analyzed in a single pass to find the partition with the best Part_cost. The best_interconnect_cost is a best one of the part_interconnect_costs among the partitions analyzed so far. The best_timing_cost is a best one of the part_timing_costs among the partitions analyzed so far. In other embodiments, a first pass over the partitions determines the best_interconnect_cost and the best_timing_cost among all of the partitions, and then a second pass analyzes the partitions to find the one of the partitions with the best Part_cost.

[0522] Fig. 37B illustrates an example of effects of (top-down) clock tree partitioning.

A random distribution of clock pins is illustrated in the upper portion of the figure. Results of clock tree partitioning and cluster formation are illustrated in the lower portion of the figure.

The CDP performs top-down partitioning using leaf-level buffer input pins and any clock sub-domain clock pins. Clock sub-domain clock input pins include input pins of gated clock cells, clock selector cells, and derived clock pins of flip-flops. The clock sub-domains are processed top-down instead of being clustered with leaf level clock pins, thus enabling insertion delay of the clock sub-domain to be utilized to balance the sub-domains. As illustrated, results of a first recursive partitioning pass are shown as **37,100**. Results of a pair of (independent) second

recursive partitioning passes are shown as **37,200a** and **37,200b**. Results of a third recursive partitioning pass are shown as **37,300b**. Note that although the recursive partitioning results are illustrated as straight cut-lines splitting contiguous regions, various embodiments and starting conditions may result in cut-lines of any shape or form, such as zig-zags, curves, and so forth. Further note that the split regions may be non-contiguous; i.e. form one or more “islands” driven by a single leaf-level buffer.

[0523] Fig. 38 illustrates an analysis according to an embodiment of clock domain and sub-domain partitioning. A clock “Root” is illustrated with relationships to leaf buffers lb1, lb2, lb3, lb4, lb5, lb6, and lb7. A tree of clock terminals is illustrated by t_1 , t_2 , t_3 , t_4 , t_5 , t_6 , and t_7 . In some embodiments edges are added to represent timing relationships (such as setup and hold times) between leaf level buffers. One type of timing relationship between first and second buffers is when the first buffer drives a first storage element, the second buffer drives a second storage element, and the two storage elements are coupled via a path having a setup (or hold) timing constraint. An example setup(hold) timing relationship between a flip-flop driven by lb1 and a flip-flop driven by lb4 is represented conceptually as dashed-line **38,100**. As illustrated, skew is minimized between the two flip-flops by driving lb1 and lb4 via the same clock terminal (t_4).

[0524] The CDP creates distinct clock domains for the following types of clock nets:

- Primary clock nets;
- Clock nets driven by gated clock cells;
- Clock nets driven by clock selector cells;
- Pseudo clock domains (if floorplan driven clock partitioning has been performed); and
- Derived clock nets.

[0525] Timing relationships between the leaf level buffers are used to create optimum timing driven partitions. A scoring function for a partition is a function of interconnect cost and timing cost. To determine setup/hold timing relationships between leaf level buffers, an abstract graph is used as illustrated in the figure, having edges between two leaf level buffers if a setup/hold path exists between elements driven by the two leaf level buffers. The weight of the edge is the number of setup/hold paths between the two leaf level buffers.

[0526] As a result of top-down partitioning, the clock tree has two types of nodes, terminals and paths. A terminal represents a graph node that is processed by downstream modules for skew minimization. Each of the terminals has a pair of paths that represent the

1 respective buffer path from the respective parent terminal to the respective pair of child
2 terminals.

3
4 **[0527]** Clock domain edges are analyzed so that clock clusters recursively propagate
5 the clock edge (e.g. a rising edge or a falling edge) used by the clock pin clusters at leaf level.
6 Thus only used ones of rise time and/or fall time are propagated for all intervening levels of
7 logic cells (including buffers and non-buffers). In some embodiments, the edges optimized by a
8 particular terminal are the union of the edges optimized by child terminals of the particular
9 terminal. For example, a first leaf-level cluster and a second leaf-level cluster both have pins
10 that use only rising edge clocks. If the two clusters are children of a first terminal, then the first
11 terminal as a whole only optimizes rise times (and advantageously does not optimize fall times).
12 If a third leaf-level cluster has at least some pins that use falling-edge clocks, and the third leaf-
13 level cluster is joined with the first terminal at a higher-level terminal (such as by the higher-
14 level terminal being a parent of the first terminal and of a second terminal, and the second
15 terminal being a parent of the third leaf-level cluster), then the higher-level terminal optimizes
16 both rise and fall times.

17
18 **[0528]** During skew minimization (such as illustrated in Fig. 35) each internal terminal
19 of a clock network is analyzed in a bottom-up visitation order and an ideal delay for each
20 respective buffer pair is determined that will minimize the skew of the terminal. Skew
21 minimization uses a successive approximation approach to determine types of buffer(s) and
22 interconnect lengths associated with each of the buffers.

23
24 **[0529]** During a first pass skew optimization (or minimization), a default input
25 transition time is used to compute delays of all clock buffers. For each terminal, respective
26 locations of buffer pairs to be placed are determined that would minimize skew for an entire sub-
27 tree. If the skew cannot be minimized by placing the buffer pair between two child terminals,
28 then an amount of meandering interconnect/buffers to minimize the skew is determined.

29
30 **[0530]** An iterative skew improver portion of the CTS tool performs multi-pass skew
31 computation and idealized delay allocation for each stage of a buffer tree. The skew improver
32 performs a multi-pass optimization because skew minimization is done bottom-up but input
33 transition is propagated top-down. Therefore during the first pass, a skew minimizer uses a
34 default input transition for each buffer stage of a clock network and performs skew minimization
35 at each level. Subsequently, a clock network timing update is performed that updates transition

1 times at each level, top-down, using an estimated output load on each of the buffers of the
2 network.

3

4 **[0531]** A second pass of skew minimization is performed that uses the propagated input
5 transition time at each of the clock buffers. Subsequent passes are performed (such as 1, 2, 3, 4,
6 or 5 iterations) seeking convergence of the skew minimizer.

7

8 **[0532]** Clock network timing is updated in accordance with buffer placement, delays of
9 buffer gates, and interconnect delays. Since cell delays are functions of input transition time and
10 phase, the clock network timing (arrival time and transition time) is relatively accurate to ensure
11 that buffer cell delays computed at each of the terminals matches closely with post-routed-
12 extracted delays.

13

14 **[0533]** Transition time at an output of a gate is a function of input transition time at the
15 gate and effective output load (C_{eff}) driven by the gate. The proper phase of transition times is
16 propagated down a clock network to accurately estimate transition times and cell delays at a next
17 level of the clock network. In some usage scenarios (such as an SAF-based design flow) buffers
18 may not be placed at ideal locations (i.e. there is no logic block in a proper buffer position).
19 Thus clock buffer placement is performed iteratively. Whenever a buffer is placed at a
20 somewhat non-ideal location, the effect of that buffer placement is propagated throughout the
21 clock sub-tree.

22

23 **[0534]** A buffer placer module of the CTS tool inserts a pair of buffers at each terminal
24 of a clock network. Unlike standard cell design flows where a buffer may be placed anywhere
25 in a row of standard cell logic, structured ASICs are constrained in where buffer resources may
26 be placed.

27

28 **[0535]** Buffer placement is performed recursively down the clock tree. At each
29 terminal, the buffer placer evaluates a finite number of buffer pair sites for suitability as a buffer
30 pair of the respective terminal. The buffer pairs are located by using a search window around an
31 ideal buffer pair location. In some embodiments, such as in a structure array fabric, a choice of a
32 buffer site is also a choice of a buffer (with fixed characteristics, such as drive strength), since
33 buffers are in fixed locations of the structured array fabric.

34

35 **[0536]** The buffer placer uses a speculative scoring function to score each pair of
36 buffers. Each buffer pair is scored on the basis of the objective function, such as:

1 $\text{buf_pair_cost} = .9 * \text{buf_delay_cost} + .1 * \text{buf_dist_cost};$

2 where

3 $\text{buf_delay_cost} = \text{dd0} * \text{dd0} + \text{dd1} * \text{dd1} + \text{dd2} * \text{dd2};$

4 where

5 $\text{dd0} = (1 - \text{est_delay} / \text{ideal_delay})$ for the respective parent terminal;

6 $\text{dd1} = (1 - \text{est_delay} / \text{ideal_delay})$ for the respective left terminal; and

7 $\text{dd2} = (1 - \text{est_delay} / \text{ideal_delay})$ for the respective right terminal.

8 where the ideal_delay is a delay for the respective buffer pair assuming placement at an ideal
9 location. According to various embodiments, weighting of buf_delay_cost and buf_dist_cost in
10 computation of bur_pair_cost varies.

11

12 **[0537]** Similarly,

13 $\text{buf_dist_cost} = \text{dbb} * \text{dbb} + \text{msd1} * \text{msd1} + \text{msd2} * \text{msd2};$ and

14 dbb = normalized Manhattan distance between the pair of buffer.

15 Ideally the pair of buffers should be as close as possible to reduce any delay uncertainty between
16 a parent buffer and respective buffer pairs. Using a dbb term penalizes any pair of buffers that
17 are far apart.

18 $\text{msd1}(2)$ = normalized distance between left/right buffer and merging segment.

19 At a lowest level where a parent terminal of the clock tree drives child leaf level clock buffers
20 (as illustrated by parent terminal t_1 in Fig. 34C), a respective merging segment of the parent
21 terminal is a line that has a first, fixed distance from a first child leaf level clock buffer (driving
22 a first sub-tree), and a second, fixed distance from a second child leaf level clock buffer (driving
23 a second sub-tree). At higher levels where a parent terminal of the clock tree drives child
24 terminals of the clock tree (as illustrated by parent terminal t_3 in Fig. 34C), a respective merging
25 segment of the parent terminal is a line that has a first, fixed distance from a respective merging
26 segment of the first child terminal, and a second, fixed distance from a respective merging
27 segment of the second child terminal. The merging segment thus specifies a set of ideal sites for
28 a respective buffer pair of a parent terminal to drive child terminals or child leaf level clock
29 buffers, the ideal sites chosen to balance skew between the children. The distance of the buffer
30 location and the merging segment are measured. The idealized buffer locations for the
31 downstream sub-tree are computed with the parent buffer being ideally placed on the merging
32 segment. If actual placement of the buffer deviates too much from the idealized line segment
33 then the estimates for the downstream terminal are no longer valid. In some embodiments, the
34 dbb term and/or the msd1(2) terms are normalized. According to various embodiments, the
35 normalization is one or more of: with respect to a fixed distance, such as 1K microns; with

respect to a technology-dependent and/or design-dependent fixed distance; and with respect to a maximum of the pre-normalization values among placements considered.

[0538] When two sub-trees have considerable differences in accumulated insertion delays then delay buffers are inserted to match insertion delay at a parent terminal. Differences in insertion delays may occur in some usage scenarios where one branch of the clock sub-tree is a (relatively large) gated-clock sub-domain and remaining branches are relatively smaller gated or un-gated clock-sub-domains.

[0539] Delay buffers are scored using an objective scoring function:

$$\text{delay_buf_cost} = 0.70 * \text{dcost} * \text{dcost} + 0.2 * \text{ncost} * \text{ncost} + 0.1 * \text{pcost} * \text{pcost};$$

where

$$\text{dcost} = (1 - (\text{accum_delay} + \text{incr_delay}) / \text{ideal_delay});$$

$$\text{ncost} = (1 - \text{actual_length} / \text{ideal_length}); \text{ and}$$

$$\text{pcost} = (1 - \text{path_remaining_length} / \text{path_ideal_remaining_length}).$$

[0540] Besides the delay cost (which has the highest weighting), delay_buf_cost uses two other metrics to evaluate a candidate delay buffer. Ncost factors in any deviation from ideal length of an interconnect for a respective path, and pcost factors in deviation of path length from a respective ideal path length.

[0541] If the skew minimizer determines that the path requires some amount of meandering interconnect to add extra delay at the buffer, then a dummy-load insertion technique is used to implement the incremental meandering wire resource. A dummy load inserter portion of the CTS tool searches for optimal dummy load sites (typically a low drive strength inverter) on an SAF-based chip and connects the buffer to the dummy load.

[0542] The CTS tool balances for max and min corners simultaneously as optimum skew for a max corner is not the optimum skew for min corner. In some usage scenarios skew at the max corner typically affects the setup timing paths whereas clock skew for the min corner affects the hold time paths. During deskewing monitored by the CTS tool, timing for both max and min corners (also known as mixed mode) is considered, and the CTS tool uses scoring functions (as described elsewhere herein) that uses a weighted sum of max and min scoring functions.

1 **[0543]** Post-routed-extracted parasitics are used to perform clock tree optimization.
2 The clock optimization is used to achieve timing closure in designs having correlation issues
3 with predicted clock skew and post-extracted clock skew. In some usage scenarios, the CTS tool
4 achieves a high degree of correlation with post-extracted skew using several techniques as
5 described elsewhere herein. The CTS tool performs several clock tree optimizations, such as
6 replacement of a clock gating cell, replacement of terminal buffers, dummy load insertion, and
7 swapping a CTS buffer for some other morphable element that may be implemented as a buffer.

CONCLUSION

[0544] Certain choices have been made in the description merely for convenience in preparing the text and drawings and unless there is an indication to the contrary the choices should not be construed per se as conveying additional information regarding structure or operation of the embodiments described. Examples of the choices include: the particular organization or assignment of the designations used for the figure numbering and the particular organization or assignment of the element identifiers (i.e., the callouts or numerical designators) used to identify and reference the features and elements of the embodiments.

[0545] Although the foregoing embodiments have been described in some detail for purposes of clarity of description and understanding, the invention is not limited to the details provided. There are many embodiments of the invention. The disclosed embodiments are exemplary and not restrictive.

[0546] It will be understood that many variations in construction, arrangement, and use are possible, which are consistent with the description and are within the scope of the claims of the issued patent. For example, interconnect and function-unit bit-widths, clock speeds, and the type of technology used are variable according to various embodiments in each component block. The names given to interconnect and logic are merely exemplary, and should not be construed as limiting the concepts described. The order and arrangement of flowchart and flow diagram process, action, and function elements are variable according to various embodiments. Also, unless specifically stated to the contrary, value ranges specified, maximum and minimum values used, or other particular specifications (such as integration techniques and design flow technologies), are merely those of the described embodiments, are expected to track improvements and changes in implementation technology, and should not be construed as limitations.

[0547] Functionally equivalent techniques known in the art are employable instead of those described to implement various components, sub-systems, functions, operations, routines, and sub-routines. It is also understood that many functional aspects of embodiments are realizable selectively in either hardware (i.e., generally dedicated circuitry) or software (i.e., via some manner of programmed controller or processor), as a function of embodiment dependent design constraints and technology trends of faster processing (facilitating migration of functions previously in hardware into software) and higher integration density (facilitating migration of functions previously in software into hardware). Specific variations in various embodiments

1 include, but are not limited to: differences in partitioning; different form factors and
2 configurations; use of different operating systems and other system software; use of different
3 interface standards, network protocols, or communication links; and other variations to be
4 expected when implementing the concepts described herein in accordance with the unique
5 engineering and business constraints of a particular application.
6

7 **[0548]** The embodiments have been described with detail and environmental context
8 well beyond that required for a minimal implementation of many aspects of the embodiments
9 described. Those of ordinary skill in the art will recognize that some embodiments omit
10 disclosed components or features without altering the basic cooperation among the remaining
11 elements. It is thus understood that much of the details disclosed are not required to implement
12 various aspects of the embodiments described. To the extent that the remaining elements are
13 distinguishable from the prior art, components and features that are omitted are not limiting on
14 the concepts described herein.
15

16 **[0549]** All such variations in design comprise insubstantial changes over the teachings
17 conveyed by the described embodiments. It is also understood that the embodiments described
18 herein have broad applicability to other computing and networking applications, and are not
19 limited to the particular application or industry of the described embodiments. The invention is
20 thus to be construed as including all possible modifications and variations encompassed within
21 the scope of the claims of the issued patent.

1 WHAT IS CLAIMED IS:

1 1. A method comprising:

2 determining locations of clock buffer resources of a structured array fabric
3 integrated circuit; and
4 creating a clock tree using, at least in part, the locations.

1 2. The method of claim 1,

2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.

1 3. The method of claim 2, wherein the placing is via simultaneous dynamical integration.

1 4. The method of claim 2, wherein the creating is subsequent to the placing.

1 5. The method of claim 4, wherein the placing is a detailed placing.

1 6. The method of claim 1, wherein the creating optimizes a maximum clock skew of the clock
2 tree.

1 7. The method of claim 1, wherein the creating optimizes a maximum latency of the clock tree.

1 8. The method of claim 1, wherein the clock tree comprises one or more levels of clock selector
2 logic gates.

1 9. The method of claim 8,

2 further comprising synthesizing clock selector circuitry into synthesized gates;
3 and
4 wherein the creating optimizes the synthesized gates to produce the one or more
5 levels of clock selector logic gates.

1 10. The method of claim 9, wherein the clock selector circuitry is designed, at least in part, by a
2 user.

1 11. The method of claim 1, wherein the clock tree comprises one or more levels of clock gating
2 logic gates.

1 12. The method of claim 11,
2 further comprising synthesizing clock gating circuitry into synthesized gates;
3 and
4 wherein the creating optimizes the synthesized gates to produce the one or more
5 levels of clock gating logic gates.

1 13. The method of claim 12, wherein the clock gating circuitry is designed, at least in part, by a
2 user.

1 14. The method of claim 1, wherein the locations are at fixed positions in the structured array
2 fabric integrated circuit.

1 15. The method of claim 1,
2 further comprising performing top-down buffer placement; and
3 wherein the performing top-down buffer placement comprises the determining.

1 16. The method of claim 1, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing,
6 bottom-up skew minimization,
7 updating clock tree timing, and
8 top-down buffer placement.

1 17. A method comprising:
2 performing floorplan driven clock partitioning; and
3 creating a clock tree in a structured array fabric integrated circuit using, at least
4 in part, at least some results of the floorplan driven clock partitioning.

- 1 18. The method of claim 17,
2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 19. The method of claim 18, wherein the placing is via simultaneous dynamical integration.
- 1 20. The method of claim 18, wherein the creating is subsequent to the placing.
- 1 21. The method of claim 20, wherein the placing is a detailed placing.
- 1 22. The method of claim 17, wherein the floorplan driven clock partitioning comprises one or
2 more of
3 determining regions of the structured array fabric integrated circuit, zero or
4 more of the regions corresponding to Intellectual Property (IP) blocks
5 and one or more of the regions corresponding to logic blocks,
6 at each respective one of the regions corresponding to the logic blocks, placing a
7 respective root clock buffer at a centroid of a bounding box
8 encompassed by all leaf clock pins within the respective region,
9 combining each clock pin of each respective one of the regions corresponding to
10 the IP blocks with the leaf clock pins of a nearest one of the regions
11 corresponding to the logic blocks, and
12 defining one or more pseudo-clock sub-domains, each of the pseudo-clock sub-
13 domains comprising one or more of the regions corresponding to the
14 logic blocks.
- 1 23. The method of claim 22, wherein the floorplan driven clock partitioning further comprises
2 deskewing each of the pseudo-clock sub-domains individually.
- 1 24. The method of claim 22,
2 wherein the floorplan driven clock partitioning further comprises deskewing a
3 first one of the pseudo-clock sub-domains and a second one of the
4 pseudo-clock sub-domains together; and
5 wherein a third one of the pseudo-clock sub-domains is in between the first
6 pseudo-clock sub-domain and the second pseudo-clock sub-domain.

- 1 25. The method of claim 22, wherein the floorplan driven clock partitioning further comprises
2 deskewing a first one of the pseudo-clock sub-domains and a second one of the pseudo-
3 clock sub-domains together, the first pseudo-clock sub-domain and the second pseudo-
4 clock sub-domain separated by a third one of the pseudo-clock sub-domains.
- 1 26. The method of claim 25, further comprising, prior to the deskewing together, deskewing the
2 first pseudo-clock sub-domain and the second pseudo-clock sub-domain individually.
- 1 27. The method of claim 25, wherein the first pseudo-clock sub-domain and the second pseudo-
2 clock sub-domain are even ones of the pseudo-clock sub-domains, and the third pseudo-
3 clock sub-domain is an odd one of the pseudo-clock sub-domains.
- 1 28. The method of claim 17, wherein the floorplan driven clock partitioning comprises
2 determining regions of the structured array fabric integrated circuit, one or more of the
3 regions corresponding to logic blocks.
- 1 29. The method of claim 28, further comprising replicating at least one gated-clock logic circuit
2 in any one of the regions having a clock pin in fan out of the at least one gated-clock
3 logic circuit.
- 1 30. The method of claim 28, further comprising replicating at least one clock-selector logic
2 circuit in any one of the regions having a clock pin in fan out of the at least one clock-
3 selector logic circuit.
- 1 31. The method of claim 17, further comprising one or more of
2 topological clock sorting,
3 top-down recursive bipartitioning,
4 clock domain processing,
5 bottom-up skew minimization,
6 updating clock tree timing, and
7 top-down buffer placement.
- 1 32. A method comprising:
2 updating timing of a clock tree;
3 locating buffers of the clock tree according to a top-down order; and
4 iterating the updating and the locating over one or more iterations.

- 1 33. The method of claim 32, wherein the locating is in accordance with a structured array fabric
2 integrated circuit.
- 1 34. The method of claim 32,
2 further comprising placing elements of all or portions of an integrated circuit
3 design; and
4 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 35. The method of claim 34, wherein the placing is via simultaneous dynamical integration.
- 1 36. The method of claim 34, wherein the updating, the locating, and the iterating are subsequent
2 to the placing.
- 1 37. The method of claim 36, wherein the placing is a detailed placing.
- 1 38. The method of claim 32, further comprising synthesizing the clock tree in a structured array
2 fabric integrated circuit using, at least in part, at least some results of the locating.
- 1 39. The method of claim 32, wherein the updating comprises propagating transition times.
- 1 40. The method of claim 39, wherein the updating operates in a top-down order.
- 1 41. The method of claim 32, wherein the updating uses a respective estimated output load on
2 each of the buffers.
- 1 42. The method of claim 41, further comprising adjusting the respective estimated output loads
2 based, at least in part, on the locating.
- 1 43. The method of claim 42, wherein the respective estimated output load on each of the buffers
2 is according to a respective estimated wire length driven by the respective buffer.
- 1 44. The method of claim 32, wherein the updating uses post-route-extracted parasitics, at least
2 in part, to determine input transition times of at least some of the buffers.
- 1 45. The method of claim 32, wherein the updating uses post-route-extracted parasitics, at least
2 in part, to determine delays of at least some of the buffers.

- 1 46. The method of claim 32, wherein a first one of the iterations uses respective default input
2 transition times of each of the buffers.
- 1 47. The method of claim 46, wherein at least some of the respective default input transition
2 times are specified by a user.
- 1 48. The method of claim 46, wherein subsequent ones of the iterations use respective
2 propagated input transition times of each of the buffers.
- 1 49. The method of claim 32, further comprising performing bottom-up skew minimization.
- 1 50. The method of claim 49, wherein the iterating further iterates the bottom-up skew
2 minimization.
- 1 51. The method of claim 49, wherein the bottom-up skew minimization determines merging
2 segments that are respective ideal sets of locations of each of a plurality of pairs of the
3 buffers.
- 1 52. The method of claim 51, wherein the locating is based, at least in part, on the merging
2 segments.
- 1 53. The method of claim 51, wherein each of the buffers of a particular one of the pairs has a
2 common parent terminal.
- 1 54. The method of claim 51, wherein each of the pairs corresponds to one of a plurality of
2 terminals of the clock tree.
- 1 55. The method of claim 32, wherein the locating comprises placing the buffers.
- 1 56. The method of claim 32, wherein the locating comprises determining locations of pairs of
2 the buffers, each of the pairs corresponding to a terminal of the clock tree.
- 1 57. The method of claim 56, wherein respective sites of each of the pairs are scored according to
2 an objective function of delay cost and distance cost.

- 1 58. The method of claim 57, wherein the objective function is a weighted sum of a max corner
2 function of delay cost and distance cost, and a min corner function of delay cost and
3 distance cost.
- 1 59. The method of claim 57, wherein the objective function is a weighted sum of the delay cost
2 and the distance cost.
- 1 60. The method of claim 57, wherein the delay cost of a particular one of the pairs is based, at
2 least in part, on estimated versus ideal delay with respect to a respective parent one of
3 the terminals of the particular pair.
- 1 61. The method of claim 57, wherein the delay cost of a particular one of the pairs is based, at
2 least in part, on estimated versus ideal delays of each of the buffers of the particular pair.
- 1 62. The method of claim 57, wherein the distance cost of a particular one of the pairs is based,
2 at least in part, on a normalized distance between the buffers of the particular pair.
- 1 63. The method of claim 62, wherein the normalized distance is a normalized Manhattan
2 distance.
- 1 64. The method of claim 57, wherein the distance cost of a particular one of the pairs is based,
2 at least in part, on a normalized distance of each of the buffers of the particular pair from
3 a merging segment.
- 1 65. The method of claim 64, wherein the merging segment specifies a set of ideal sites for the
2 particular pair.
- 1 66. The method of claim 64,
2 further comprising performing bottom-up skew minimization; and
3 wherein the bottom-up skew minimization determines the merging segment.

- 1 67. The method of claim 32, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing, and
6 bottom-up skew minimization.
- 1 68. The method of claim 32, further comprising one or more of
2 floorplan driven clock partitioning,
3 clock domain processing, and
4 bottom-up skew minimization.
- 1 69. A method comprising:
2 clustering at least some of a plurality of leaf level clock pins of a clock tree; and
3 performing recursive bipartitioning using, at least in part, results of the
4 clustering.
- 1 70. The method of claim 69, wherein the clustering is compatible with a structured array fabric
2 integrated circuit.
- 1 71. The method of claim 69,
2 further comprising placing elements of all or portions of an integrated circuit
3 design; and
4 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 72. The method of claim 71, wherein the placing is via simultaneous dynamical integration.
- 1 73. The method of claim 71, wherein the clustering and the performing are subsequent to the
2 placing.
- 1 74. The method of claim 73, wherein the placing is a detailed placing.
- 1 75. The method of claim 69, further comprising synthesizing the clock tree in a structured array
2 fabric integrated circuit using, at least in part, at least some results of the performing.

- 1 76. The method of claim 75, wherein the clustering comprises a recursive partitioning of the at
2 least some of the leaf level clock pins.
- 1 77. The method of claim 69,
2 wherein the recursive bipartitioning is a first recursive bipartitioning; and
3 wherein the clustering partitions the at least some of the leaf level clock pins
4 using a second recursive bipartitioning.
- 1 78. The method of claim 77, wherein the second recursive bipartitioning uses an objective
2 function that minimizes a diameter of a polygon formed by ones of the at least some of
3 the leaf level clock pins in a respective partition.
- 1 79. The method of claim 69, wherein the clustering forms bottom-up clusters, each of the
2 bottom-up clusters comprising a portion of the at least some of the leaf level clock pins.
- 1 80. The method of claim 79, further comprising assigning a respective leaf level clock buffer to
2 each of the bottom-up clusters.
- 1 81. The method of claim 80, wherein the recursive bipartitioning operates to partition the
2 respective leaf level clock buffers.
- 1 82. The method of claim 80, further comprising placing the respective leaf level clock buffers.
- 1 83. The method of claim 82, wherein the placing comprises
2 updating timing of the clock tree,
3 locating the respective leaf level clock buffers according to a top-down order,
4 and
5 iterating the updating and the locating over one or more iterations.
- 1 84. The method of claim 79, further comprising performing timing driven pin swapping.
- 1 85. The method of claim 84, wherein the performing timing driven pin swapping is subsequent
2 to the performing recursive bipartitioning.

1 86. The method of claim 84, wherein the performing timing driven pin swapping comprises
2 swapping a first one of the leaf level clock pins of a first one of the bottom-up clusters
3 with a second one of the leaf level clock pins of a second one of the bottom-up clusters.

1 87. The method of claim 86, wherein the timing driven pin swapping improves a hold time of
2 either the first leaf level clock pin or the leaf level second clock pin.

1 88. The method of claim 69, wherein the recursive bipartitioning operates according to a
2 function of partition interconnect cost and partition timing cost.

1 89. The method of claim 88, wherein the function is a weighted sum of a max corner version of
2 the function, and a min corner version of the function.

1 90. The method of claim 69,
2 further comprising clock domain processing; and
3 wherein the clock domain processing comprises the clustering.

1 91. The method of claim 69, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 bottom-up skew minimization,
5 updating clock tree timing, and
6 top-down buffer placement.

1 92. The method of claim 69, further comprising one or more of
2 floorplan driven clock partitioning,
3 bottom-up skew minimization, and
4 top-down buffer placement.

- 1 93. A method comprising:
2 performing recursive bipartitioning of buffer clock pins of at least some of a
3 plurality of leaf level clock buffers of a clock tree according to a
4 function of a respective partition interconnect cost and a respective
5 partition timing cost of each of a plurality of partitions; and
6 optimizing the clock tree in accordance with a structured array fabric integrated
7 circuit using, at least in part, at least some results of the recursive
8 bipartitioning.
- 1 94. The method of claim 93,
2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 95. The method of claim 94, wherein the placing is via simultaneous dynamical integration.
- 1 96. The method of claim 94, wherein the optimizing is subsequent to the placing.
- 1 97. The method of claim 96, wherein the placing is a detailed placing.
- 1 98. The method of claim 93, wherein the function is a weighted sum of a max corner version of
2 the function, and a min corner version of the function.
- 1 99. The method of claim 93, wherein the partition interconnect cost of a particular one of the
2 partitions is a function of a net semi-perimeter of a bounding box of all of the buffer
3 clock pins in the particular partition.
- 1 100. The method of claim 93, wherein the recursive bipartitioning comprises scoring ones of the
2 partitions to choose a best one of the partitions.

- 1 101. The method of claim 93,
2 wherein the clock tree comprises one or more sub-domains, each of the sub-
3 domains having a respective sub-domain clock buffer; and
4 wherein the performing further partitions buffer clock pins of the sub-domain
5 clock buffers along with the buffer clocks pins of the at least some leaf
6 level clock buffers.
- 1 102. The method of claim 101, further comprising performing recursive bipartitioning of ones
2 of the buffer clock pins of the sub-domain clock buffers of each of the sub-domains.
- 1 103. The method of claim 101, wherein one of the sub-domains is a gated-clock sub-domain.
- 1 104. The method of claim 101, wherein one of the sub-domains is a clock-selector sub-domain.
- 1 105. The method of claim 93, further comprising clustering at least some of a plurality of clock
2 pins of the clock tree into one or more clusters, each of the clusters associated with a
3 respective one of the at least some leaf level clock buffers.
- 1 106. The method of claim 105, wherein the clustering is prior to the performing.
- 1 107. The method of claim 105, further comprising performing timing driven pin swapping.
- 1 108. The method of claim 107, wherein the performing timing driven pin swapping is
2 subsequent to the performing recursive bipartitioning.
- 1 109. The method of claim 107, wherein the performing timing driven pin swapping comprises
2 swapping a first one of the clock pins of a first one of the clusters with a second one of
3 the clock pins of a second one of the clusters.
- 1 110. The method of claim 109, wherein the performing timing driven pin swapping improves a
2 hold time of either the first clock pin or the second clock pin.

- 1 111. The method of claim 93, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 clock domain processing,
5 bottom-up skew minimization,
6 updating clock tree timing, and
7 top-down buffer placement.
- 1 112. A method comprising:
2 performing recursive bipartitioning of buffer clock pins of buffers of a clock
3 tree, the buffers comprising one or more leaf level clock buffers and one
4 or more sub-domain clock buffers; and
5 optimizing the clock tree in a structured array fabric integrated circuit using, at
6 least in part, at least some results of the recursive bipartitioning.
- 1 113. The method of claim 112, wherein at least one of the sub-domain clock buffers drives a
2 gated-clock sub-domain.
- 1 114. The method of claim 112, wherein at least one of the sub-domain clock buffers drives a
2 clock-selector sub-domain.
- 1 115. The method of claim 112,
2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 116. The method of claim 115, wherein the placing is via simultaneous dynamical integration.
- 1 117. The method of claim 115, wherein the optimizing is subsequent to the placing.
- 1 118. The method of claim 117, wherein the placing is a detailed placing.
- 1 119. The method of claim 112, further comprising inserting dummy loads.

- 1 120. The method of claim 112,
2 wherein the clock tree comprises one or more sub-domains, each of the sub-
3 domains associated with one of the sub-domain clock buffers; and
4 further comprising performing recursive bipartitioning of buffer clock pins of
5 respective leaf level clock buffers of each of the sub-domains.
- 1 121. The method of claim 112, further comprising performing timing driven pin swapping.
- 1 122. The method of claim 121, wherein the performing timing driven pin swapping is
2 subsequent to the performing recursive bipartitioning.
- 1 123. The method of claim 112, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing,
6 bottom-up skew minimization,
7 updating clock tree timing, and
8 top-down buffer placement.
- 1 124. A method comprising:
2 performing bottom-up skew minimization; and
3 optimizing a clock tree in a structured array fabric integrated circuit using, at
4 least in part, at least some results of the bottom-up skew minimization.
- 1 125. The method of claim 124,
2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 126. The method of claim 125, wherein the placing is via simultaneous dynamical integration.
- 1 127. The method of claim 125, wherein the optimizing is subsequent to the placing.
- 1 128. The method of claim 127, wherein the placing is a detailed placing.

- 1 129. The method of claim 124, wherein the bottom-up skew minimization comprises visiting
2 terminals of the clock tree in a bottom-up order.
- 1 130. The method of claim 129, wherein the bottom-up skew minimization further comprises
2 computing, at each of the terminals, a respective ideal delay of a respective buffer pair.
- 1 131. The method of claim 130, wherein the bottom-up skew minimization further determines,
2 based at least in part on the respective ideal delay of a particular one of the respective
3 buffer pairs, at a particular one of the terminals, a respective merging segment.
- 1 132. The method of claim 131, wherein the respective merging segment is used, at least in part,
2 to place the particular buffer pair.
- 1 133. The method of claim 130, wherein the respective ideal delay of a particular one of the
2 terminals is based, at least in part, on skew between a first child of the particular
3 terminal and a second child of the particular terminal.
- 1 134. The method of claim 133, wherein the first child is a leaf level clock buffer of the clock
2 tree.
- 1 135. The method of claim 133, wherein the first child is one of the terminals of the clock tree.
- 1 136. The method of claim 124,
2 further comprising clock domain processing; and
3 wherein the clock domain processing determines a respective set of clock edges
4 used by a respective sub-tree under each of a plurality of terminals of
5 the clock tree.
- 1 137. The method of claim 136, wherein each of the respective sets comprises one or more of a
2 rising clock edge and a falling clock edge.

1 138. The method of claim 136,
2 wherein a particular one of the terminals is a parent of a first child one of the
3 terminals and a second child one of the terminals; and
4 wherein the respective set of clock edges used by the particular terminal is a
5 union of the respective sets of clock edges of the first child terminal and
6 the second child terminal.

1 139. The method of claim 136,
2 wherein a particular one of the terminals is a parent of a leaf level clock buffer;
3 and
4 wherein the respective set of clock edges used by the particular terminal
5 comprises a set of all clock edges of all devices in fan out of the leaf
6 level clock buffer.

1 140. The method of claim 136, wherein a particular one of the terminals is a parent of a gated-
2 clock sub-domain.

1 141. The method of claim 136, wherein a particular one of the terminals is a parent of a clock-
2 selector sub-domain.

1 142. The method of claim 124, further comprising propagating input transition times top-down.

1 143. The method of claim 142, wherein the propagating is based, at least in part, on post-routed-
2 extracted parasitics.

1 144. The method of claim 142, further comprising iterating the bottom-up skew minimization
2 and the propagating.

1 145. The method of claim 124, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing,
6 updating clock tree timing, and
7 top-down buffer placement.

- 1 146. A method comprising:
2 performing a top-down placement of terminals of a clock tree, each of the
3 terminals comprising a respective pair of buffers, each of the terminals
4 associated with a respective merging segment;
5 optimizing a clock tree in a structured array fabric integrated circuit using, at
6 least in part, at least some results of the performing; and
7 wherein the performing is according to the merging segments.
- 1 147. The method of claim 146, wherein the top-down placement is compatible with the
2 structured array fabric integrated circuit.
- 1 148. The method of claim 146,
2 further comprising placing elements of all or portions of an integrated circuit
3 design; and
4 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 149. The method of claim 148, wherein the placing is via simultaneous dynamical integration.
- 1 150. The method of claim 148, wherein the performing is subsequent to the placing.
- 1 151. The method of claim 150, wherein the placing is a detailed placing.
- 1 152. The method of claim 146, further comprising synthesizing the clock tree using, at least in
2 part, at least some results of the performing.
- 1 153. The method of claim 146, further comprising performing a bottom-up skew minimization
2 to determine the merging segments.
- 1 154. The method of claim 153, wherein, at each of the terminals, the performing the bottom-up
2 skew minimization determines a respective first ideal delay to a respective first child of
3 the terminal, and a respective second ideal delay to a respective second child of the
4 respective terminal.
- 1 155. The method of claim 154, wherein the respective first ideal delay and the respective second
2 ideal delay are used, at least in part, to determine the respective merging segment.

- 1 156. The method of claim 154, wherein the respective first ideal delay and the respective second
2 ideal delay are in accordance with minimizing skew between the respective first child
3 and the respective second child.
- 1 157. The method of claim 154, wherein the respective first child is a leaf level clock buffer of
2 the clock tree.
- 1 158. The method of claim 154, wherein the respective first child is one of the terminals of the
2 clock tree.
- 1 159. The method of claim 154, wherein a first buffer of the respective pair is used to drive the
2 respective first child, and a second buffer of the respective pair is used to drive the
3 respective second child.
- 1 160. The method of claim 154, further comprising inserting, at a particular one of the terminals,
2 one or more delay buffers to match, at least in part, the respective first ideal delay and
3 the respective second ideal delay.
- 1 161. The method of claim 160, wherein each of the delay buffers is scored using an objective
2 scoring function of a delay cost and a length cost.
- 1 162. The method of claim 154, further comprising interconnecting, between a particular one of
2 the terminals and the respective first child of the particular terminal, meandering wire to
3 match, at least in part, the respective first ideal delay and the respective second ideal
4 delay.
- 1 163. The method of claim 162, further comprising inserting dummy loads on the meandering
2 wire.
- 1 164. The method of claim 146, wherein the performing comprises, at each of the terminals of
2 the clock tree, evaluating one or more buffer pair sites of the respective pair, each of the
3 buffer pair sites comprising a first buffer site and a second buffer site.
- 1 165. The method of claim 164, wherein the performing further comprises using a scoring
2 function to score each of the buffer pair sites.

- 1 166. The method of claim 165, wherein the scoring function is a weighted sum of a max corner
2 function of delay cost and distance cost, and a min corner function of delay cost and
3 distance cost.
- 1 167. The method of claim 165, wherein the scoring function is a weighted sum of a delay cost
2 and a distance cost.
- 1 168. The method of claim 167, wherein the delay cost of a particular one of the buffer pair sites
2 is based, at least in part, on estimated versus ideal delay with respect to a respective
3 parent terminal of the particular buffer pair site.
- 1 169. The method of claim 167, wherein the delay cost of a particular one of the buffer pair sites
2 is based, at least in part, on estimated versus ideal delays of each of the buffer sites of
3 the particular buffer pair site.
- 1 170. The method of claim 167, wherein the distance cost of a particular one of the buffer pair
2 sites is based, at least in part, on a normalized distance between the first and the second
3 buffer sites of the particular buffer pair site.
- 1 171. The method of claim 167,
2 wherein a particular one of the terminals is associated with a particular one of
3 the respective pairs, and is associated with a particular one of the
4 respective merging segments; and
5 wherein the distance cost of a particular one of the buffer pair sites of the
6 particular pair is based, at least in part, on a normalized distance of each
7 of the buffer sites of the particular buffer pair site from the particular
8 merging segment.
- 1 172. The method of claim 146, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing, and
6 updating clock tree timing.

- 1 173. The method of claim 146, further comprising one or more of
2 floorplan driven clock partitioning, and
3 clock domain processing.
- 1 174. A computer readable medium having a set of instructions stored therein which when
2 executed by a processing device causes the processing device to perform procedures
3 comprising:
4 determining locations of clock buffer resources of a structured array fabric
5 integrated circuit; and
6 creating a clock tree using, at least in part, the locations.
- 1 175. The computer readable medium of claim 174,
2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 176. The computer readable medium of claim 175, wherein the placing is via simultaneous
2 dynamical integration.
- 1 177. The computer readable medium of claim 175, wherein the creating is subsequent to the
2 placing.
- 1 178. The computer readable medium of claim 177, wherein the placing is a detailed placing.
- 1 179. The computer readable medium of claim 174, wherein the creating optimizes a maximum
2 clock skew of the clock tree.
- 1 180. The computer readable medium of claim 174, wherein the creating optimizes a maximum
2 latency of the clock tree.
- 1 181. The computer readable medium of claim 174, wherein the clock tree comprises one or
2 more levels of clock selector logic gates.

1 182. The computer readable medium of claim 181,
2 further comprising synthesizing clock selector circuitry into synthesized gates;
3 and
4 wherein the creating optimizes the synthesized gates to produce the one or more
5 levels of clock selector logic gates.

1 183. The computer readable medium of claim 182, wherein the clock selector circuitry is
2 designed, at least in part, by a user.

1 184. The computer readable medium of claim 174, wherein the clock tree comprises one or
2 more levels of clock gating logic gates.

1 185. The computer readable medium of claim 184,
2 further comprising synthesizing clock gating circuitry into synthesized gates;
3 and
4 wherein the creating optimizes the synthesized gates to produce the one or more
5 levels of clock gating logic gates.

1 186. The computer readable medium of claim 185, wherein the clock gating circuitry is
2 designed, at least in part, by a user.

1 187. The computer readable medium of claim 174, wherein the locations are at fixed positions
2 in the structured array fabric integrated circuit.

1 188. The computer readable medium of claim 174,
2 further comprising performing top-down buffer placement; and
3 wherein the performing top-down buffer placement comprises the determining.

1 189. The computer readable medium of claim 174, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing,
6 bottom-up skew minimization,
7 updating clock tree timing, and
8 top-down buffer placement.

- 1 190. A computer readable medium having a set of instructions stored therein which when
2 executed by a processing device causes the processing device to perform procedures
3 comprising:
4 performing floorplan driven clock partitioning; and
5 creating a clock tree in a structured array fabric integrated circuit using, at least
6 in part, at least some results of the floorplan driven clock partitioning.
- 1 191. The computer readable medium of claim 190,
2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 192. The computer readable medium of claim 191, wherein the placing is via simultaneous
2 dynamical integration.
- 1 193. The computer readable medium of claim 191, wherein the creating is subsequent to the
2 placing.
- 1 194. The computer readable medium of claim 193, wherein the placing is a detailed placing.
- 1 195. The computer readable medium of claim 190, wherein the floorplan driven clock
2 partitioning comprises one or more of
3 determining regions of the structured array fabric integrated circuit, zero or
4 more of the regions corresponding to Intellectual Property (IP) blocks
5 and one or more of the regions corresponding to logic blocks,
6 at each respective one of the regions corresponding to the logic blocks, placing a
7 respective root clock buffer at a centroid of a bounding box
8 encompassed by all leaf clock pins within the respective region,
9 combining each clock pin of each respective one of the regions corresponding to
10 the IP blocks with the leaf clock pins of a nearest one of the regions
11 corresponding to the logic blocks, and
12 defining one or more pseudo-clock sub-domains, each of the pseudo-clock sub-
13 domains comprising one or more of the regions corresponding to the
14 logic blocks.

- 1 196. The computer readable medium of claim 195, wherein the floorplan driven clock
2 partitioning further comprises deskewing each of the pseudo-clock sub-domains
3 individually.
- 1 197. The computer readable medium of claim 195,
2 wherein the floorplan driven clock partitioning further comprises deskewing a
3 first one of the pseudo-clock sub-domains and a second one of the
4 pseudo-clock sub-domains together; and
5 wherein a third one of the pseudo-clock sub-domains is in between the first
6 pseudo-clock sub-domain and the second pseudo-clock sub-domain.
- 1 198. The computer readable medium of claim 195, wherein the floorplan driven clock
2 partitioning further comprises deskewing a first one of the pseudo-clock sub-domains
3 and a second one of the pseudo-clock sub-domains together, the first pseudo-clock sub-
4 domain and the second pseudo-clock sub-domain separated by a third one of the pseudo-
5 clock sub-domains.
- 1 199. The computer readable medium of claim 198, further comprising, prior to the deskewing
2 together, deskewing the first pseudo-clock sub-domain and the second pseudo-clock
3 sub-domain individually.
- 1 200. The computer readable medium of claim 198, wherein the first pseudo-clock sub-domain
2 and the second pseudo-clock sub-domain are even ones of the pseudo-clock sub-
3 domains, and the third pseudo-clock sub-domain is an odd one of the pseudo-clock sub-
4 domains.
- 1 201. The computer readable medium of claim 190, wherein the floorplan driven clock
2 partitioning comprises determining regions of the structured array fabric integrated
3 circuit, one or more of the regions corresponding to logic blocks.
- 1 202. The computer readable medium of claim 201, further comprising replicating at least one
2 gated-clock logic circuit in any one of the regions having a clock pin in fan out of the at
3 least one gated-clock logic circuit.

- 1 203. The computer readable medium of claim 201, further comprising replicating at least one
2 clock-selector logic circuit in any one of the regions having a clock pin in fan out of the
3 at least one clock-selector logic circuit.
- 1 204. The computer readable medium of claim 190, further comprising one or more of
2 topological clock sorting,
3 top-down recursive bipartitioning,
4 clock domain processing,
5 bottom-up skew minimization,
6 updating clock tree timing, and
7 top-down buffer placement.
- 1 205. A computer readable medium having a set of instructions stored therein which when
2 executed by a processing device causes the processing device to perform procedures
3 comprising:
4 updating timing of a clock tree;
5 locating buffers of the clock tree according to a top-down order; and
6 iterating the updating and the locating over one or more iterations.
- 1 206. The computer readable medium of claim 205, wherein the locating is in accordance with a
2 structured array fabric integrated circuit.
- 1 207. The computer readable medium of claim 205,
2 further comprising placing elements of all or portions of an integrated circuit
3 design; and
4 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 208. The computer readable medium of claim 207, wherein the placing is via simultaneous
2 dynamical integration.
- 1 209. The computer readable medium of claim 207, wherein the updating, the locating, and the
2 iterating are subsequent to the placing.
- 1 210. The computer readable medium of claim 209, wherein the placing is a detailed placing.

- 1 211. The computer readable medium of claim 205, further comprising synthesizing the clock
2 tree in a structured array fabric integrated circuit using, at least in part, at least some
3 results of the locating.
- 1 212. The computer readable medium of claim 205, wherein the updating comprises propagating
2 transition times.
- 1 213. The computer readable medium of claim 212, wherein the updating operates in a top-down
2 order.
- 1 214. The computer readable medium of claim 205, wherein the updating uses a respective
2 estimated output load on each of the buffers.
- 1 215. The computer readable medium of claim 214, further comprising adjusting the respective
2 estimated output loads based, at least in part, on the locating.
- 1 216. The computer readable medium of claim 215, wherein the respective estimated output load
2 on each of the buffers is according to a respective estimated wire length driven by the
3 respective buffer.
- 1 217. The computer readable medium of claim 205, wherein the updating uses post-route-
2 extracted parasitics, at least in part, to determine input transition times of at least some
3 of the buffers.
- 1 218. The computer readable medium of claim 205, wherein the updating uses post-route-
2 extracted parasitics, at least in part, to determine delays of at least some of the buffers.
- 1 219. The computer readable medium of claim 205, wherein a first one of the iterations uses
2 respective default input transition times of each of the buffers.
- 1 220. The computer readable medium of claim 219, wherein at least some of the respective
2 default input transition times are specified by a user.
- 1 221. The computer readable medium of claim 219, wherein subsequent ones of the iterations
2 use respective propagated input transition times of each of the buffers.

1 222. The computer readable medium of claim 205, further comprising performing bottom-up
2 skew minimization.

1 223. The computer readable medium of claim 222, wherein the iterating further iterates the
2 bottom-up skew minimization.

1 224. The computer readable medium of claim 222, wherein the bottom-up skew minimization
2 determines merging segments that are respective ideal sets of locations of each of a
3 plurality of pairs of the buffers.

1 225. The computer readable medium of claim 224, wherein the locating is based, at least in part,
2 on the merging segments.

1 226. The computer readable medium of claim 224, wherein each of the buffers of a particular
2 one of the pairs has a common parent terminal.

1 227. The computer readable medium of claim 224, wherein each of the pairs corresponds to one
2 of a plurality of terminals of the clock tree.

1 228. The computer readable medium of claim 205, wherein the locating comprises placing the
2 buffers.

1 229. The computer readable medium of claim 205, wherein the locating comprises determining
2 locations of pairs of the buffers, each of the pairs corresponding to a terminal of the
3 clock tree.

1 230. The computer readable medium of claim 229, wherein respective sites of each of the pairs
2 are scored according to an objective function of delay cost and distance cost.

1 231. The computer readable medium of claim 230, wherein the objective function is a weighted
2 sum of a max corner function of delay cost and distance cost, and a min corner function
3 of delay cost and distance cost.

1 232. The computer readable medium of claim 230, wherein the objective function is a weighted
2 sum of the delay cost and the distance cost.

1 233. The computer readable medium of claim 230, wherein the delay cost of a particular one of
2 the pairs is based, at least in part, on estimated versus ideal delay with respect to a
3 respective parent one of the terminals of the particular pair.

1 234. The computer readable medium of claim 230, wherein the delay cost of a particular one of
2 the pairs is based, at least in part, on estimated versus ideal delays of each of the buffers
3 of the particular pair.

1 235. The computer readable medium of claim 230, wherein the distance cost of a particular one
2 of the pairs is based, at least in part, on a normalized distance between the buffers of the
3 particular pair.

1 236. The computer readable medium of claim 235, wherein the normalized distance is a
2 normalized Manhattan distance.

1 237. The computer readable medium of claim 230, wherein the distance cost of a particular one
2 of the pairs is based, at least in part, on a normalized distance of each of the buffers of
3 the particular pair from a merging segment.

1 238. The computer readable medium of claim 237, wherein the merging segment specifies a set
2 of ideal sites for the particular pair.

1 239. The computer readable medium of claim 237,
2 further comprising performing bottom-up skew minimization; and
3 wherein the bottom-up skew minimization determines the merging segment.

1 240. The computer readable medium of claim 205, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing, and
6 bottom-up skew minimization.

- 1 241. The computer readable medium of claim 205, further comprising one or more of
2 floorplan driven clock partitioning,
3 clock domain processing, and
4 bottom-up skew minimization.
- 1 242. A computer readable medium having a set of instructions stored therein which when
2 executed by a processing device causes the processing device to perform procedures
3 comprising:
4 clustering at least some of a plurality of leaf level clock pins of a clock tree; and
5 performing recursive bipartitioning using, at least in part, results of the
6 clustering.
- 1 243. The computer readable medium of claim 242, wherein the clustering is compatible with a
2 structured array fabric integrated circuit.
- 1 244. The computer readable medium of claim 242,
2 further comprising placing elements of all or portions of an integrated circuit
3 design; and
4 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 245. The computer readable medium of claim 244, wherein the placing is via simultaneous
2 dynamical integration.
- 1 246. The computer readable medium of claim 244, wherein the clustering and the performing
2 are subsequent to the placing.
- 1 247. The computer readable medium of claim 246, wherein the placing is a detailed placing.
- 1 248. The computer readable medium of claim 242, further comprising synthesizing the clock
2 tree in a structured array fabric integrated circuit using, at least in part, at least some
3 results of the performing.
- 1 249. The computer readable medium of claim 248, wherein the clustering comprises a recursive
2 partitioning of the at least some of the leaf level clock pins of.

- 1 250. The computer readable medium of claim 242,
2 wherein the recursive bipartitioning is a first recursive bipartitioning; and
3 wherein the clustering partitions the at least some of the leaf level clock pins
4 using a second recursive bipartitioning.
- 1 251. The computer readable medium of claim 250, wherein the second recursive bipartitioning
2 uses an objective function that minimizes a diameter of a polygon formed by ones of the
3 at least some of the leaf level clock pins in a respective partition.
- 1 252. The computer readable medium of claim 242, wherein the clustering forms bottom-up
2 clusters, each of the bottom-up clusters comprising a portion of the at least some of the
3 leaf level clock pins.
- 1 253. The computer readable medium of claim 252, further comprising assigning a respective
2 leaf level clock buffer to each of the bottom-up clusters.
- 1 254. The computer readable medium of claim 253, wherein the recursive bipartitioning operates
2 to partition the respective leaf level clock buffers.
- 1 255. The computer readable medium of claim 253, further comprising placing the respective
2 leaf level clock buffers.
- 1 256. The computer readable medium of claim 255, wherein the placing comprises
2 updating timing of the clock tree,
3 locating the respective leaf level clock buffers according to a top-down order,
4 and
5 iterating the updating and the locating over one or more iterations.
- 1 257. The computer readable medium of claim 252, further comprising performing timing driven
2 pin swapping.
- 1 258. The computer readable medium of claim 257, wherein the performing timing driven pin
2 swapping is subsequent to the performing recursive bipartitioning.

1 259. The computer readable medium of claim 257, wherein the performing timing driven pin
2 swapping comprises swapping a first one of the leaf level clock pins of a first one of the
3 bottom-up clusters with a second one of the leaf level clock pins of a second one of the
4 bottom-up clusters.

1 260. The computer readable medium of claim 259, wherein the timing driven pin swapping
2 improves a hold time of either the first leaf level clock pin or the leaf level second clock
3 pin.

1 261. The computer readable medium of claim 242, wherein the recursive bipartitioning operates
2 according to a function of partition interconnect cost and partition timing cost.

1 262. The computer readable medium of claim 261, wherein the function is a weighted sum of a
2 max corner version of the function, and a min corner version of the function.

1 263. The computer readable medium of claim 242,
2 further comprising clock domain processing; and
3 wherein the clock domain processing comprises the clustering.

1 264. The computer readable medium of claim 242, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 bottom-up skew minimization,
5 updating clock tree timing, and
6 top-down buffer placement.

1 265. The computer readable medium of claim 242, further comprising one or more of
2 floorplan driven clock partitioning,
3 bottom-up skew minimization, and
4 top-down buffer placement.

1 266. A computer readable medium having a set of instructions stored therein which when
2 executed by a processing device causes the processing device to perform procedures
3 comprising:
4 performing recursive bipartitioning of buffer clock pins of at least some of a
5 plurality of leaf level clock buffers of a clock tree according to a
6 function of a respective partition interconnect cost and a respective
7 partition timing cost of each of a plurality of partitions; and
8 optimizing the clock tree in accordance with a structured array fabric integrated
9 circuit using, at least in part, at least some results of the recursive
10 bipartitioning.

1 267. The computer readable medium of claim 266,
2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.

1 268. The computer readable medium of claim 267, wherein the placing is via simultaneous
2 dynamical integration.

1 269. The computer readable medium of claim 267, wherein the optimizing is subsequent to the
2 placing.

1 270. The computer readable medium of claim 269, wherein the placing is a detailed placing.

1 271. The computer readable medium of claim 266, wherein the function is a weighted sum of a
2 max corner version of the function, and a min corner version of the function.

1 272. The computer readable medium of claim 266, wherein the partition interconnect cost of a
2 particular one of the partitions is a function of a net semi-perimeter of a bounding box of
3 all of the buffer clock pins in the particular partition.

1 273. The computer readable medium of claim 266, wherein the recursive bipartitioning
2 comprises scoring ones of the partitions to choose a best one of the partitions.

1 274. The computer readable medium of claim 266,
2 wherein the clock tree comprises one or more sub-domains, each of the sub-
3 domains having a respective sub-domain clock buffer; and
4 wherein the performing further partitions buffer clock pins of the sub-domain
5 clock buffers along with the buffer clocks pins of the at least some leaf
6 level clock buffers.

1 275. The computer readable medium of claim 274, further comprising performing recursive
2 bipartitioning of ones of the buffer clock pins of the sub-domain clock buffers of each of
3 the sub-domains.

1 276. The computer readable medium of claim 274, wherein one of the sub-domains is a gated-
2 clock sub-domain.

1 277. The computer readable medium of claim 274, wherein one of the sub-domains is a clock-
2 selector sub-domain.

1 278. The computer readable medium of claim 266, further comprising clustering at least some
2 of a plurality of clock pins of the clock tree into one or more clusters, each of the
3 clusters associated with a respective one of the at least some leaf level clock buffers.

1 279. The computer readable medium of claim 278, wherein the clustering is prior to the
2 performing.

1 280. The computer readable medium of claim 278, further comprising performing timing driven
2 pin swapping.

1 281. The computer readable medium of claim 280, wherein the performing timing driven pin
2 swapping is subsequent to the performing recursive bipartitioning.

1 282. The computer readable medium of claim 280, wherein the performing timing driven pin
2 swapping comprises swapping a first one of the clock pins of a first one of the clusters
3 with a second one of the clock pins of a second one of the clusters.

1 283. The computer readable medium of claim 282, wherein the performing timing driven pin
2 swapping improves a hold time of either the first clock pin or the second clock pin.

- 1 284. The computer readable medium of claim 266, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 clock domain processing,
5 bottom-up skew minimization,
6 updating clock tree timing, and
7 top-down buffer placement.
- 1 285. A computer readable medium having a set of instructions stored therein which when
2 executed by a processing device causes the processing device to perform procedures
3 comprising:
4 performing recursive bipartitioning of buffer clock pins of buffers of a clock
5 tree, the buffers comprising one or more leaf level clock buffers and one
6 or more sub-domain clock buffers; and
7 optimizing the clock tree in a structured array fabric integrated circuit using, at
8 least in part, at least some results of the recursive bipartitioning.
- 1 286. The computer readable medium of claim 285, wherein at least one of the sub-domain clock
2 buffers drives a gated-clock sub-domain.
- 1 287. The computer readable medium of claim 285, wherein at least one of the sub-domain clock
2 buffers drives a clock-selector sub-domain.
- 1 288. The computer readable medium of claim 285,
2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 289. The computer readable medium of claim 288, wherein the placing is via simultaneous
2 dynamical integration.
- 1 290. The computer readable medium of claim 288, wherein the optimizing is subsequent to the
2 placing.
- 1 291. The computer readable medium of claim 290, wherein the placing is a detailed placing.

1 292. The computer readable medium of claim 285, further comprising inserting dummy loads.

1 293. The computer readable medium of claim 285,
2 wherein the clock tree comprises one or more sub-domains, each of the sub-
3 domains associated with one of the sub-domain clock buffers; and
4 further comprising performing recursive bipartitioning of buffer clock pins of
5 respective leaf level clock buffers of each of the sub-domains.

1 294. The computer readable medium of claim 285, further comprising performing timing driven
2 pin swapping.

1 295. The computer readable medium of claim 294, wherein the performing timing driven pin
2 swapping is subsequent to the performing recursive bipartitioning.

1 296. The computer readable medium of claim 285, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing,
6 bottom-up skew minimization,
7 updating clock tree timing, and
8 top-down buffer placement.

1 297. A computer readable medium having a set of instructions stored therein which when
2 executed by a processing device causes the processing device to perform procedures
3 comprising:
4 performing bottom-up skew minimization; and
5 optimizing a clock tree in a structured array fabric integrated circuit using, at
6 least in part, at least some results of the bottom-up skew minimization.

1 298. The computer readable medium of claim 297,
2 further comprising placing elements of all or portions of an integrated circuit
3 design in accordance with the structured array fabric integrated circuit;
4 and
5 wherein the clock tree interconnects clock pins of at least some of the elements.

- 1 299. The computer readable medium of claim 298, wherein the placing is via simultaneous
2 dynamical integration.
- 1 300. The computer readable medium of claim 298, wherein the optimizing is subsequent to the
2 placing.
- 1 301. The computer readable medium of claim 300, wherein the placing is a detailed placing.
- 1 302. The computer readable medium of claim 297, wherein the bottom-up skew minimization
2 comprises visiting terminals of the clock tree in a bottom-up order.
- 1 303. The computer readable medium of claim 302, wherein the bottom-up skew minimization
2 further comprises computing, at each of the terminals, a respective ideal delay of a
3 respective buffer pair.
- 1 304. The computer readable medium of claim 303, wherein the bottom-up skew minimization
2 further determines, based at least in part on the respective ideal delay of a particular one
3 of the respective buffer pairs, at a particular one of the terminals, a respective merging
4 segment.
- 1 305. The computer readable medium of claim 304, wherein the respective merging segment is
2 used, at least in part, to place the particular buffer pair.
- 1 306. The computer readable medium of claim 303, wherein the respective ideal delay of a
2 particular one of the terminals is based, at least in part, on skew between a first child of
3 the particular terminal and a second child of the particular terminal.
- 1 307. The computer readable medium of claim 306, wherein the first child is a leaf level clock
2 buffer of the clock tree.
- 1 308. The computer readable medium of claim 306, wherein the first child is one of the terminals
2 of the clock tree.

1 309. The computer readable medium of claim 297,
2 further comprising clock domain processing; and
3 wherein the clock domain processing determines a respective set of clock edges
4 used by a respective sub-tree under each of a plurality of terminals of
5 the clock tree.

1 310. The computer readable medium of claim 309, wherein each of the respective sets
2 comprises one or more of a rising clock edge and a falling clock edge.

1 311. The computer readable medium of claim 309,
2 wherein a particular one of the terminals is a parent of a first child one of the
3 terminals and a second child one of the terminals; and
4 wherein the respective set of clock edges used by the particular terminal is a
5 union of the respective sets of clock edges of the first child terminal and
6 the second child terminal.

1 312. The computer readable medium of claim 309,
2 wherein a particular one of the terminals is a parent of a leaf level clock buffer;
3 and
4 wherein the respective set of clock edges used by the particular terminal
5 comprises a set of all clock edges of all devices in fan out of the leaf
6 level clock buffer.

1 313. The computer readable medium of claim 309, wherein a particular one of the terminals is a
2 parent of a gated-clock sub-domain.

1 314. The computer readable medium of claim 309, wherein a particular one of the terminals is a
2 parent of a clock-selector sub-domain.

1 315. The computer readable medium of claim 297, further comprising propagating input
2 transition times top-down.

1 316. The computer readable medium of claim 315, wherein the propagating is based, at least in
2 part, on post-routed-extracted parasitics.

1 317. The computer readable medium of claim 315, further comprising iterating the bottom-up
2 skew minimization and the propagating.

1 318. The computer readable medium of claim 297, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing,
6 updating clock tree timing, and
7 top-down buffer placement.

1 319. A computer readable medium having a set of instructions stored therein which when
2 executed by a processing device causes the processing device to perform procedures
3 comprising:
4 performing a top-down placement of terminals of a clock tree, each of the
5 terminals comprising a respective pair of buffers, each of the terminals
6 associated with a respective merging segment;
7 optimizing a clock tree in a structured array fabric integrated circuit using, at
8 least in part, at least some results of the performing; and
9 wherein the performing is according to the merging segments.

1 320. The computer readable medium of claim 319, wherein the top-down placement is
2 compatible with the structured array fabric integrated circuit.

1 321. The computer readable medium of claim 319,
2 further comprising placing elements of all or portions of an integrated circuit
3 design; and
4 wherein the clock tree interconnects clock pins of at least some of the elements.

1 322. The computer readable medium of claim 321, wherein the placing is via simultaneous
2 dynamical integration.

1 323. The computer readable medium of claim 321, wherein the performing is subsequent to the
2 placing.

1 324. The computer readable medium of claim 323, wherein the placing is a detailed placing.

- 1 325. The computer readable medium of claim 319, further comprising synthesizing the clock
2 tree using, at least in part, at least some results of the performing.
- 1 326. The computer readable medium of claim 319, further comprising performing bottom-up
2 skew minimization to determine the merging segments.
- 1 327. The computer readable medium of claim 326, wherein, at each of the terminals, the
2 performing the bottom-up skew minimization determines a respective first ideal delay to
3 a respective first child of the terminal, and a respective second ideal delay to a respective
4 second child of the respective terminal.
- 1 328. The computer readable medium of claim 327, wherein the respective first ideal delay and
2 the respective second ideal delay are used, at least in part, to determine the respective
3 merging segment.
- 1 329. The computer readable medium of claim 327, wherein the respective first ideal delay and
2 the respective second ideal delay are in accordance with minimizing skew between the
3 respective first child and the respective second child.
- 1 330. The computer readable medium of claim 327, wherein the respective first child is a leaf
2 level clock buffer of the clock tree.
- 1 331. The computer readable medium of claim 327, wherein the respective first child is one of
2 the terminals of the clock tree.
- 1 332. The computer readable medium of claim 327, wherein a first buffer of the respective pair is
2 used to drive the respective first child, and a second buffer of the respective pair is used
3 to drive the respective second child.
- 1 333. The computer readable medium of claim 327, further comprising inserting, at a particular
2 one of the terminals, one or more delay buffers to match, at least in part, the respective
3 first ideal delay and the respective second ideal delay.
- 1 334. The computer readable medium of claim 333, wherein each of the delay buffers is scored
2 using an objective scoring function of a delay cost and a length cost.

- 1 335. The computer readable medium of claim 327, further comprising interconnecting, between
2 a particular one of the terminals and the respective first child of the particular terminal,
3 meandering wire to match, at least in part, the respective first ideal delay and the
4 respective second ideal delay.
- 1 336. The computer readable medium of claim 335, further comprising inserting dummy loads
2 on the meandering wire.
- 1 337. The computer readable medium of claim 319, wherein the performing the top-down
2 placement comprises, at each of the terminals of the clock tree, evaluating one or more
3 buffer pair sites of the respective pair, each of the buffer pair sites comprising a first
4 buffer site and a second buffer site.
- 1 338. The computer readable medium of claim 337, wherein the performing further comprises
2 using a scoring function to score each of the buffer pair sites.
- 1 339. The computer readable medium of claim 338, wherein the scoring function is a weighted
2 sum of a max corner function of delay cost and distance cost, and a min corner function
3 of delay cost and distance cost.
- 1 340. The computer readable medium of claim 338, wherein the scoring function is a weighted
2 sum of a delay cost and a distance cost.
- 1 341. The computer readable medium of claim 340, wherein the delay cost of a particular one of
2 the buffer pair sites is based, at least in part, on estimated versus ideal delay with respect
3 to a respective parent terminal of the particular buffer pair site.
- 1 342. The computer readable medium of claim 340, wherein the delay cost of a particular one of
2 the buffer pair sites is based, at least in part, on estimated versus ideal delays of each of
3 the buffer sites of the particular buffer pair site.
- 1 343. The computer readable medium of claim 340, wherein the distance cost of a particular one
2 of the buffer pair sites is based, at least in part, on a normalized distance between the
3 first and the second buffer sites of the particular buffer pair site.

1 344. The computer readable medium of claim 340,
2 wherein a particular one of the terminals is associated with a particular one of
3 the respective pairs, and is associated with a particular one of the
4 respective merging segments; and
5 wherein the distance cost of a particular one of the buffer pair sites of the
6 particular pair is based, at least in part, on a normalized distance of each
7 of the buffer sites of the particular buffer pair site from the particular
8 merging segment.

1 345. The computer readable medium of claim 319, further comprising one or more of
2 floorplan driven clock partitioning,
3 topological clock sorting,
4 top-down recursive bipartitioning,
5 clock domain processing, and
6 updating clock tree timing.

1 346. The computer readable medium of claim 319, further comprising one or more of
2 floorplan driven clock partitioning, and
3 clock domain processing.

1 347. A system comprising:
2 means for determining locations of clock buffer resources of a structured array
3 fabric integrated circuit; and
4 means for creating a clock tree based, at least in part, on the locations.

1 348. The system of claim 347,
2 further comprising means for placing elements of all or portions of an integrated
3 circuit design in accordance with the structured array fabric integrated
4 circuit; and
5 wherein the clock tree interconnects clock pins of at least some of the elements.

1 349. The system of claim 348, wherein the means for placing is according to simultaneous
2 dynamical integration.

1 350. The system of claim 348, wherein the means for creating is operable subsequent to the
2 means for placing.

1 351. The system of claim 350, wherein the means for placing is a means for detailed placing.

1 352. The system of claim 347, wherein the means for creating is enabled to optimize a
2 maximum clock skew of the clock tree.

1 353. The system of claim 347, wherein the means for creating is enabled to optimize a
2 maximum latency of the clock tree.

1 354. The system of claim 347, wherein the clock tree comprises one or more levels of clock
2 selector logic gates.

1 355. The system of claim 354,
2 further comprising means for synthesizing clock selector circuitry into
3 synthesized gates; and
4 wherein the means for creating is enabled to optimize the synthesized gates to
5 produce the one or more levels of clock selector logic gates.

1 356. The system of claim 355, wherein the clock selector circuitry is designed, at least in part,
2 by a user.

1 357. The system of claim 347, wherein the clock tree comprises one or more levels of clock
2 gating logic gates.

1 358. The system of claim 357,
2 further comprising means for synthesizing clock gating circuitry into
3 synthesized gates; and
4 wherein the means for creating is enabled to optimize the synthesized gates to
5 produce the one or more levels of clock gating logic gates.

1 359. The system of claim 358, wherein the clock gating circuitry is designed, at least in part, by
2 a user.

1 360. The system of claim 347, wherein the locations are at fixed positions in the structured array
2 fabric integrated circuit.

1 361. The system of claim 347,
2 further comprising means for top-down buffer placement; and
3 wherein the means for top-down buffer placement comprises the means for
4 determining.

1 362. The system of claim 347, further comprising one or more of
2 means for floorplan driven clock partitioning,
3 means for topological clock sorting,
4 means for top-down recursive bipartitioning,
5 means for clock domain processing,
6 means for bottom-up skew minimization,
7 means for updating clock tree timing, and
8 means for top-down buffer placement.

1 363. A system comprising:
2 means for floorplan driven clock partitioning; and
3 means for creating a clock tree in a structured array fabric integrated circuit
4 based, at least in part, on at least some results of the means for floorplan
5 driven clock partitioning.

1 364. The system of claim 363,
2 further comprising means for placing elements of all or portions of an integrated
3 circuit design in accordance with the structured array fabric integrated
4 circuit; and
5 wherein the clock tree interconnects clock pins of at least some of the elements.

1 365. The system of claim 364, wherein the means for placing is according to simultaneous
2 dynamical integration.

1 366. The system of claim 364, wherein the means for creating is operable subsequent to the
2 means for placing.

1 367. The system of claim 366, wherein the means for placing is a means for detailed placing.

1 368. The system of claim 363, wherein the means for floorplan driven clock partitioning
2 comprises one or more of
3 means for determining regions of the structured array fabric integrated circuit,
4 zero or more of the regions corresponding to Intellectual Property (IP)
5 blocks and one or more of the regions corresponding to logic blocks,
6 means for placing, at each respective one of the regions corresponding to the
7 logic blocks, a respective root clock buffer at a centroid of a bounding
8 box encompassed by all leaf clock pins within the respective region,
9 means for combining each clock pin of each respective one of the regions
10 corresponding to the IP blocks with the leaf clock pins of a nearest one
11 of the regions corresponding to the logic blocks, and
12 means for defining one or more pseudo-clock sub-domains, each of the pseudo-
13 clock sub-domains comprising one or more of the regions
14 corresponding to the logic blocks.

1 369. The system of claim 368, wherein the means for floorplan driven clock partitioning further
2 comprises means for deskewing individually each of the pseudo-clock sub-domains.

1 370. The system of claim 368,
2 wherein the means for floorplan driven clock partitioning further comprises
3 means for deskewing together a first one of the pseudo-clock sub-
4 domains and a second one of the pseudo-clock sub-domains; and
5 wherein a third one of the pseudo-clock sub-domains is in between the first
6 pseudo-clock sub-domain and the second pseudo-clock sub-domain.

1 371. The system of claim 368, wherein the means for floorplan driven clock partitioning further
2 comprises means for deskewing together a first one of the pseudo-clock sub-domains
3 and a second one of the pseudo-clock sub-domains, the first pseudo-clock sub-domain
4 and the second pseudo-clock sub-domain separated by a third one of the pseudo-clock
5 sub-domains.

1 372. The system of claim 371, further comprising, means for deskewing individually the first
2 pseudo-clock sub-domain and the second pseudo-clock sub-domain, the means for
3 deskewing individually operable prior to the means for deskewing together.

- 1 373. The system of claim 371, wherein the first pseudo-clock sub-domain and the second
2 pseudo-clock sub-domain are even ones of the pseudo-clock sub-domains, and the third
3 pseudo-clock sub-domain is an odd one of the pseudo-clock sub-domains.
- 1 374. The system of claim 363, wherein the means for floorplan driven clock partitioning
2 comprises means for determining regions of the structured array fabric integrated circuit,
3 one or more of the regions corresponding to logic blocks.
- 1 375. The system of claim 374, further comprising means for replicating at least one gated-clock
2 logic circuit in any one of the regions having a clock pin in fan out of the at least one
3 gated-clock logic circuit.
- 1 376. The system of claim 374, further comprising means for replicating at least one clock-
2 selector logic circuit in any one of the regions having a clock pin in fan out of the at
3 least one clock-selector logic circuit.
- 1 377. The system of claim 363, further comprising one or more of
2 means for topological clock sorting,
3 means for top-down recursive bipartitioning,
4 means for clock domain processing,
5 means for bottom-up skew minimization,
6 means for updating clock tree timing, and
7 means for top-down buffer placement.
- 1 378. A system comprising:
2 means for updating timing of a clock tree;
3 means for locating buffers of the clock tree, the means for locating operable
4 according to a top-down order; and
5 means for iterating the means for updating and the means for locating over one
6 or more iterations.
- 1 379. The system of claim 378, wherein the means for locating is in accordance with a structured
2 array fabric integrated circuit.

- 1 380. The system of claim 378,
2 further comprising means for placing elements of all or portions of an integrated
3 circuit design; and
4 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 381. The system of claim 380, wherein the means for placing is according to simultaneous
2 dynamical integration.
- 1 382. The system of claim 380, wherein the means for updating, the means for locating, and the
2 means for iterating are operable subsequent to the means for placing.
- 1 383. The system of claim 382, wherein the means for placing is a means for detailed placing.
- 1 384. The system of claim 378, further comprising means for synthesizing the clock tree in a
2 structured array fabric integrated circuit based, at least in part, on at least some results of
3 the means for locating.
- 1 385. The system of claim 378, wherein the means for updating comprises means for propagating
2 transition times.
- 1 386. The system of claim 385, wherein the means for updating is operable according to a top-
2 down order.
- 1 387. The system of claim 378, wherein the means for updating is operable according to a
2 respective estimated output load on each of the buffers.
- 1 388. The system of claim 387, further comprising means for adjusting the respective estimated
2 output loads based, at least in part, on the means for locating.
- 1 389. The system of claim 388, wherein the respective estimated output load on each of the
2 buffers is according to a respective estimated wire length driven by the respective buffer.
- 1 390. The system of claim 378, wherein the means for updating is operable according to post-
2 route-extracted parasitics, and is enabled to determine input transition times of at least
3 some of the buffers.

- 1 391. The system of claim 378, wherein the means for updating is operable according to post-
2 route-extracted parasitics and is enabled to determine delays of at least some of the
3 buffers.
- 1 392. The system of claim 378, wherein a first one of the iterations uses respective default input
2 transition times of each of the buffers.
- 1 393. The system of claim 392, wherein at least some of the respective default input transition
2 times are specified by a user.
- 1 394. The system of claim 392, wherein subsequent ones of the iterations use respective
2 propagated input transition times of each of the buffers.
- 1 395. The system of claim 378, further comprising means for bottom-up skew minimization.
- 1 396. The system of claim 395, wherein the means for iterating is further enabled to iterate the
2 means for bottom-up skew minimization.
- 1 397. The system of claim 395, wherein the means for bottom-up skew minimization is enabled
2 to determine merging segments that are respective ideal sets of locations of each of a
3 plurality of pairs of the buffers.
- 1 398. The system of claim 397, wherein the means for locating is operable according to the
2 merging segments.
- 1 399. The system of claim 397, wherein each of the buffers of a particular one of the pairs has a
2 common parent terminal.
- 1 400. The system of claim 397, wherein each of the pairs corresponds to one of a plurality of
2 terminals of the clock tree.
- 1 401. The system of claim 378, wherein the means for locating comprises means for placing the
2 buffers.

- 1 402. The system of claim 378, wherein the means for locating comprises means for determining
2 locations of pairs of the buffers, each of the pairs corresponding to a terminal of the
3 clock tree.
- 1 403. The system of claim 402, further comprising means for scoring respective sites of each of
2 the pairs, the means for scoring operable according to an objective function of delay cost
3 and distance cost.
- 1 404. The system of claim 403, wherein the objective function is a weighted sum of a max corner
2 function of delay cost and distance cost, and a min corner function of delay cost and
3 distance cost.
- 1 405. The system of claim 403, wherein the objective function is a weighted sum of the delay
2 cost and the distance cost.
- 1 406. The system of claim 403, wherein the delay cost of a particular one of the pairs is based, at
2 least in part, on estimated versus ideal delay with respect to a respective parent one of
3 the terminals of the particular pair.
- 1 407. The system of claim 403, wherein the delay cost of a particular one of the pairs is based, at
2 least in part, on estimated versus ideal delays of each of the buffers of the particular pair.
- 1 408. The system of claim 403, wherein the distance cost of a particular one of the pairs is based,
2 at least in part, on a normalized distance between the buffers of the particular pair.
- 1 409. The system of claim 408, wherein the normalized distance is a normalized Manhattan
2 distance.
- 1 410. The system of claim 403, wherein the distance cost of a particular one of the pairs is based,
2 at least in part, on a normalized distance of each of the buffers of the particular pair from
3 a merging segment.
- 1 411. The system of claim 410, wherein the merging segment specifies a set of ideal sites for the
2 particular pair.

1 412. The system of claim 410,
2 further comprising means for bottom-up skew minimization; and
3 wherein the means for bottom-up skew minimization is enabled to determine the
4 merging segment.

1 413. The system of claim 378, further comprising one or more of
2 means for floorplan driven clock partitioning,
3 means for topological clock sorting,
4 means for top-down recursive bipartitioning,
5 means for clock domain processing, and
6 means for bottom-up skew minimization.

1 414. The system of claim 378, further comprising one or more of
2 means for floorplan driven clock partitioning,
3 means for clock domain processing, and
4 means for bottom-up skew minimization.

1 415. A system comprising:
2 means for clustering at least some of a plurality of leaf level clock pins of a
3 clock tree; and
4 means for recursive bipartitioning based, at least in part, on results of the means
5 for clustering.

1 416. The system of claim 415, wherein the means for clustering is compatible with a structured
2 array fabric integrated circuit.

1 417. The system of claim 415,
2 further comprising means for placing elements of all or portions of an integrated
3 circuit design; and
4 wherein the clock tree interconnects clock pins of at least some of the elements.

1 418. The system of claim 417, wherein the means for placing is according to simultaneous
2 dynamical integration.

1 419. The system of claim 417, wherein the means for clustering and the means for recursive
2 bipartitioning are operable subsequent to the means for placing.

- 1 420. The system of claim 419, wherein the means for placing is a means for detailed placing.
- 1 421. The system of claim 415, further comprising means for synthesizing the clock tree in a
2 structured array fabric integrated circuit based, at least in part, on at least some results of
3 the means for recursive bipartitioning.
- 1 422. The system of claim 421, wherein the means for clustering comprises a means for recursive
2 partitioning of the at least some of the leaf level clock pins.
- 1 423. The system of claim 415,
2 wherein the means for recursive bipartitioning is a first means for recursive
3 bipartitioning; and
4 wherein the means for clustering is enabled to partition the at least some of the
5 leaf level clock pins based, at least in part, on at least some results of a
6 second means for recursive bipartitioning.
- 1 424. The system of claim 423, wherein the second means for recursive bipartitioning is operable
2 according to an objective function that minimizes a diameter of a polygon formed by
3 ones of the at least some of the leaf level clock pins in a respective partition.
- 1 425. The system of claim 415, wherein the means for clustering is enabled to form bottom-up
2 clusters, each of the bottom-up clusters comprising a portion of the at least some of the
3 leaf level clock pins.
- 1 426. The system of claim 425, further comprising means for assigning a respective leaf level
2 clock buffer to each of the bottom-up clusters.
- 1 427. The system of claim 426, wherein the means for recursive bipartitioning is enabled to
2 partition clock pins of the respective leaf level clock buffers.
- 1 428. The system of claim 426, further comprising means for placing the respective leaf level
2 clock buffers.

1 429. The system of claim 428, wherein the means for placing comprises
2 means for updating timing of the clock tree,
3 means for locating the respective leaf level clock buffers according to a top-
4 down order, and
5 means for iterating the means for updating and the means for locating over one
6 or more iterations.

1 430. The system of claim 425, further comprising means for timing driven pin swapping.

1 431. The system of claim 430, wherein the means for timing driven pin swapping is operable
2 subsequent to the means for recursive bipartitioning.

1 432. The system of claim 430, wherein the means for timing driven pin swapping comprises
2 means for swapping a first one of the leaf level clock pins of a first one of the bottom-up
3 clusters with a second one of the leaf level clock pins of a second one of the bottom-up
4 clusters.

1 433. The system of claim 432, wherein the means for timing driven pin swapping is enabled to
2 improve a hold time of either the first leaf level clock pin or the leaf level second clock
3 pin.

1 434. The system of claim 415, wherein the means for recursive bipartitioning is operable
2 according to a function of partition interconnect cost and partition timing cost.

1 435. The system of claim 434, wherein the function is a weighted sum of a max corner version
2 of the function, and a min corner version of the function.

1 436. The system of claim 415,
2 further comprising means for clock domain processing; and
3 wherein the means for clock domain processing comprises the means for
4 clustering.

1 437. The system of claim 415, further comprising one or more of
2 means for floorplan driven clock partitioning,
3 means for topological clock sorting,
4 means for bottom-up skew minimization,
5 means for updating clock tree timing, and
6 means for top-down buffer placement.

1 438. The system of claim 415, further comprising one or more of
2 means for floorplan driven clock partitioning,
3 means for bottom-up skew minimization, and
4 means for top-down buffer placement.

1 439. A system comprising:
2 means for recursive bipartitioning of buffer clock pins of at least some of a
3 plurality of leaf level clock buffers of a clock tree, the means for
4 recursive bipartitioning operable according to a function of a respective
5 partition interconnect cost and a respective partition timing cost of each
6 of a plurality of partitions; and
7 means for optimizing the clock tree in accordance with a structured array fabric
8 integrated circuit based, at least in part, on at least some results of the
9 means for recursive bipartitioning.

1 440. The system of claim 439,
2 further comprising means for placing elements of all or portions of an integrated
3 circuit design in accordance with the structured array fabric integrated
4 circuit; and
5 wherein the clock tree interconnects clock pins of at least some of the elements.

1 441. The system of claim 440, wherein the means for placing is operable according to
2 simultaneous dynamical integration.

1 442. The system of claim 440, wherein the means for optimizing is operable subsequent to the
2 means for placing.

1 443. The system of claim 442, wherein the means for placing is a means for detailed placing.

- 1 444. The system of claim 439, wherein the function is a weighted sum of a max corner version
2 of the function, and a min corner version of the function.
- 1 445. The system of claim 439, wherein the partition interconnect cost of a particular one of the
2 partitions is a function of a net semi-perimeter of a bounding box of all of the buffer
3 clock pins in the particular partition.
- 1 446. The system of claim 439, wherein the means for recursive bipartitioning comprises means
2 for scoring ones of the partitions, the means for recursive bipartitioning enabled to
3 choose a best one of the partitions.
- 1 447. The system of claim 439,
2 wherein the clock tree comprises one or more sub-domains, each of the sub-
3 domains having a respective sub-domain clock buffer; and
4 wherein the means for recursive bipartitioning is enabled to partition buffer
5 clock pins of the sub-domain clock buffers along with the buffer clocks
6 pins of the at least some leaf level clock buffers.
- 1 448. The system of claim 447, wherein the means for recursive bipartitioning is further enabled
2 to independently partition ones of the buffer clock pins of the sub-domain clock buffers
3 of each of the sub-domains.
- 1 449. The system of claim 447, wherein one of the sub-domains is a gated-clock sub-domain.
- 1 450. The system of claim 447, wherein one of the sub-domains is a clock-selector sub-domain.
- 1 451. The system of claim 439, further comprising means for clustering at least some of a
2 plurality of clock pins of the clock tree into one or more clusters, each of the clusters
3 associated with a respective one of the at least some leaf level clock buffers.
- 1 452. The system of claim 451, wherein the means for clustering is operable prior to the means
2 for recursive bipartitioning.
- 1 453. The system of claim 451, further comprising means for timing driven pin swapping.

1 454. The system of claim 453, wherein the means for timing driven pin swapping is operable
2 subsequent to the means for recursive bipartitioning.

1 455. The system of claim 453, wherein the means for timing driven pin swapping comprises
2 means for swapping a first one of the clock pins of a first one of the clusters with a
3 second one of the clock pins of a second one of the clusters.

1 456. The system of claim 455, wherein the means for timing driven pin swapping is enabled to
2 improve a hold time of either the first clock pin or the second clock pin.

1 457. The system of claim 439, further comprising one or more of
2 means for floorplan driven clock partitioning,
3 means for topological clock sorting,
4 means for clock domain processing,
5 means for bottom-up skew minimization,
6 means for updating clock tree timing, and
7 means for top-down buffer placement.

1 458. A system comprising:
2 means for recursive bipartitioning of buffer clock pins of buffers of a clock tree,
3 the buffers comprising one or more leaf level clock buffers and one or
4 more sub-domain clock buffers; and
5 means for optimizing the clock tree in a structured array fabric integrated circuit
6 based, at least in part, on at least some results of the means for recursive
7 bipartitioning.

1 459. The system of claim 458, wherein at least one of the sub-domain clock buffers drives a
2 gated-clock sub-domain.

1 460. The system of claim 458, wherein at least one of the sub-domain clock buffers drives a
2 clock-selector sub-domain.

- 1 461. The system of claim 458,
2 further comprising means for placing elements of all or portions of an integrated
3 circuit design in accordance with the structured array fabric integrated
4 circuit; and
5 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 462. The system of claim 461, wherein the means for placing is according to simultaneous
2 dynamical integration.
- 1 463. The system of claim 461, wherein the means for optimizing is operable subsequent to the
2 means for placing.
- 1 464. The system of claim 463, wherein the means for placing is a means for detailed placing.
- 1 465. The system of claim 458, further comprising means for inserting dummy loads.
- 1 466. The system of claim 458,
2 wherein the clock tree comprises one or more sub-domains, each of the sub-
3 domains associated with one of the sub-domain clock buffers; and
4 further comprising means for recursive bipartitioning of buffer clock pins of
5 respective leaf level clock buffers of each of the sub-domains.
- 1 467. The system of claim 458, further comprising means for timing driven pin swapping.
- 1 468. The system of claim 467, wherein the means for timing driven pin swapping is operable
2 subsequent to the means for recursive bipartitioning.
- 1 469. The system of claim 458, further comprising one or more of
2 means for floorplan driven clock partitioning,
3 means for topological clock sorting,
4 means for top-down recursive bipartitioning,
5 means for clock domain processing,
6 means for bottom-up skew minimization,
7 means for updating clock tree timing, and
8 means for top-down buffer placement.

- 1 470. A system comprising:
2 means for bottom-up skew minimization; and
3 means for optimizing a clock tree in a structured array fabric integrated circuit
4 based, at least in part, on at least some results of the means for bottom-
5 up skew minimization.
- 1 471. The system of claim 470,
2 further comprising means for placing elements of all or portions of an integrated
3 circuit design in accordance with the structured array fabric integrated
4 circuit; and
5 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 472. The system of claim 471, wherein the means for placing is according to simultaneous
2 dynamical integration.
- 1 473. The system of claim 471, wherein the means for optimizing is operable subsequent to the
2 means for placing.
- 1 474. The system of claim 473, wherein the means for placing is a means for detailed placing.
- 1 475. The system of claim 470, wherein the means for bottom-up skew minimization comprises
2 means for visiting terminals of the clock tree in a bottom-up order.
- 1 476. The system of claim 475, wherein the means for bottom-up skew minimization further
2 comprises means for computing, at each of the terminals, a respective ideal delay of a
3 respective buffer pair.
- 1 477. The system of claim 476, wherein the means for bottom-up skew minimization is enabled
2 to determine, based at least in part on the respective ideal delay of a particular one of the
3 respective buffer pairs, at a particular one of the terminals, a respective merging
4 segment.
- 1 478. The system of claim 477, further comprising means for placing the particular buffer pair,
2 the means for placing operable according to the respective merging segment.

1 479. The system of claim 476, wherein the respective ideal delay of a particular one of the
2 terminals is based, at least in part, on skew between a first child of the particular
3 terminal and a second child of the particular terminal.

1 480. The system of claim 479, wherein the first child is a leaf level clock buffer of the clock
2 tree.

1 481. The system of claim 479, wherein the first child is one of the terminals of the clock tree.

1 482. The system of claim 470,
2 further comprising means for clock domain processing; and
3 wherein the means for clock domain processing is enabled to determine a
4 respective set of clock edges used by a respective sub-tree under each of
5 a plurality of terminals of the clock tree.

1 483. The system of claim 482, wherein each of the respective sets comprises one or more of a
2 rising clock edge and a falling clock edge.

1 484. The system of claim 482,
2 wherein a particular one of the terminals is a parent of a first child one of the
3 terminals and a second child one of the terminals; and
4 wherein the respective set of clock edges used by the particular terminal is a
5 union of the respective sets of clock edges of the first child terminal and
6 the second child terminal.

1 485. The system of claim 482,
2 wherein a particular one of the terminals is a parent of a leaf level clock buffer;
3 and
4 wherein the respective set of clock edges used by the particular terminal
5 comprises a set of all clock edges of all devices in fan out of the leaf
6 level clock buffer.

1 486. The system of claim 482, wherein a particular one of the terminals is a parent of a gated-
2 clock sub-domain.

- 1 487. The system of claim 482, wherein a particular one of the terminals is a parent of a clock-
2 selector sub-domain.
- 1 488. The system of claim 470, further comprising means for propagating input transition times
2 top-down.
- 1 489. The system of claim 488, wherein the means for propagating is operable according to post-
2 routed-extracted parasitics.
- 1 490. The system of claim 488, further comprising means for iterating the means for bottom-up
2 skew minimization and the means for propagating.
- 1 491. The system of claim 470, further comprising one or more of
2 means for floorplan driven clock partitioning,
3 means for topological clock sorting,
4 means for top-down recursive bipartitioning,
5 means for clock domain processing,
6 means for updating clock tree timing, and
7 means for top-down buffer placement.
- 1 492. A system comprising:
2 means for top-down placement of terminals of a clock tree, each of the terminals
3 comprising a respective pair of buffers, each of the terminals associated
4 with a respective merging segment;
5 means for optimizing a clock tree in a structured array fabric integrated circuit
6 based, at least in part, on at least some results of the means for top-down
7 placement; and
8 wherein the means for top-down placement is operable according to the merging
9 segments.
- 1 493. The system of claim 492, wherein the means for top-down placement is compatible with
2 the structured array fabric integrated circuit.

- 1 494. The system of claim 492,
2 further comprising means for placing elements of all or portions of an integrated
3 circuit design; and
4 wherein the clock tree interconnects clock pins of at least some of the elements.
- 1 495. The system of claim 494, wherein the means for placing is according to simultaneous
2 dynamical integration.
- 1 496. The system of claim 494, wherein the means for top-down placement is operable
2 subsequent to the means for placing.
- 1 497. The system of claim 496, wherein the means for placing is a means for detailed placing.
- 1 498. The system of claim 492, further comprising means for synthesizing the clock tree based,
2 at least in part, on at least some results of the means for top-down placement.
- 1 499. The system of claim 492, further comprising means for bottom-up skew minimization, the
2 means for bottom-up skew minimization enabled to determine the merging segments.
- 1 500. The system of claim 499, wherein, at each of the terminals, the means for bottom-up skew
2 minimization is enabled to determine a respective first ideal delay to a respective first
3 child of the terminal, and a respective second ideal delay to a respective second child of
4 the respective terminal.
- 1 501. The system of claim 500, wherein the means for bottom-up skew minimization is further
2 enabled to determine the respective merging segment based, at least in part, on the
3 respective first ideal delay and the respective second ideal.
- 1 502. The system of claim 500, wherein the respective first ideal delay and the respective second
2 ideal delay are in accordance with minimizing skew between the respective first child
3 and the respective second child.
- 1 503. The system of claim 500, wherein the respective first child is a leaf level clock buffer of
2 the clock tree.

- 1 504. The system of claim 500, wherein the respective first child is one of the terminals of the
2 clock tree.
- 1 505. The system of claim 500, wherein a first buffer of the respective pair is used to drive the
2 respective first child, and a second buffer of the respective pair is used to drive the
3 respective second child.
- 1 506. The system of claim 500, further comprising means for inserting, at a particular one of the
2 terminals, one or more delay buffers to match, at least in part, the respective first ideal
3 delay and the respective second ideal delay.
- 1 507. The system of claim 506, further comprising means for scoring each of the delay buffers,
2 the means for scoring based, at least in part, on an objective scoring function of a delay
3 cost and a length cost.
- 1 508. The system of claim 500, further comprising means for interconnecting, between a
2 particular one of the terminals and the respective first child of the particular terminal,
3 meandering wire to match, at least in part, the respective first ideal delay and the
4 respective second ideal delay.
- 1 509. The system of claim 508, further comprising means for inserting dummy loads on the
2 meandering wire.
- 1 510. The system of claim 492, wherein the means for top-down placement comprises means for
2 evaluating, at each of the terminals of the clock tree, one or more buffer pair sites of the
3 respective pair, each of the buffer pair sites comprising a first buffer site and a second
4 buffer site.
- 1 511. The system of claim 510, wherein the means for top-down placement further comprises
2 means for scoring each of the buffer pair sites.
- 1 512. The system of claim 511, wherein the means for scoring is operable according to a
2 weighted sum of a max corner function of delay cost and distance cost, and a min corner
3 function of delay cost and distance cost.

- 1 513. The system of claim 511, wherein the means for scoring is operable according to a
2 weighted sum of a delay cost and a distance cost.
- 1 514. The system of claim 513, wherein the delay cost of a particular one of the buffer pair sites
2 is based, at least in part, on estimated versus ideal delay with respect to a respective
3 parent terminal of the particular buffer pair site.
- 1 515. The system of claim 513, wherein the delay cost of a particular one of the buffer pair sites
2 is based, at least in part, on estimated versus ideal delays of each of the buffer sites of
3 the particular buffer pair site.
- 1 516. The system of claim 513, wherein the distance cost of a particular one of the buffer pair
2 sites is based, at least in part, on a normalized distance between the first and the second
3 buffer sites of the particular buffer pair site.
- 1 517. The system of claim 513,
2 wherein a particular one of the terminals is associated with a particular one of
3 the respective pairs, and is associated with a particular one of the
4 respective merging segments; and
5 wherein the distance cost of a particular one of the buffer pair sites of the
6 particular pair is based, at least in part, on a normalized distance of each
7 of the buffer sites of the particular buffer pair site from the particular
8 merging segment.
- 1 518. The system of claim 492, further comprising one or more of
2 means for floorplan driven clock partitioning,
3 means for topological clock sorting,
4 means for top-down recursive bipartitioning,
5 means for clock domain processing, and
6 means for updating clock tree timing.
- 1 519. The system of claim 492, further comprising one or more of
2 means for floorplan driven clock partitioning, and
3 means for clock domain processing.

1/56




Fig. 1

2/56




Fig. 2

3/56




Fig. 3A

4/56




Fig. 3B

5/56




Fig. 3C

6/56




Fig. 3D

7/56




Fig. 3E




Fig. 3F




Fig. 3G

8/56




Fig. 3H

9/56




Fig. 3I

10/56




Fig. 4

11/56




Fig. 5A

12/56




Fig. 5B

13/56




Fig. 5C

14/56




Fig. 6

15/56




Fig. 7A

16/56




Fig. 7B




Fig. 7C




Fig. 7D




Fig. 7E




Fig. 7F

17/56




Fig. 8A

18/56




Fig. 8B

19/56




Fig. 9

20/56




Fig. 10

21/56




Fig. 11

22/56




Fig. 12A




Fig. 12B

23/56




Fig. 13

24/56


1400


Fig. 14

25/56




Fig. 15A




Fig. 15B

26/56




Fig. 15C

27/56




Fig. 15D




Fig. 15E

28/56




Fig. 16A




Fig. 16B

29/56




Fig. 17A




Fig. 17B




Fig. 18

31/56




Fig. 19

32/56




Fig. 20A

33/56




Fig. 20B

34/56




Fig. 21A




Fig. 21B

35/56




Fig. 22A




Fig. 22B

36/56




Fig. 23

37/56




Fig. 24A




Fig. 24B

38/56




Fig. 25A

39/56




Fig. 25B

40/56




Fig. 26

41/56




Fig. 27

42/56




Fig. 28

43/56




Fig. 29

44/56




Fig. 30

45/56




Fig. 31

46/56




Fig. 32

47/56




Fig. 33

48/56




Fig. 34A

49/56




Fig. 34B

50/56




Fig. 34C

51/56




Fig. 35

52/56




Fig. 36A

53/56






Fig. 36B

54/56



Partitions, $P1 = \{La, Lb\}$
 $P2 = \{Ca, Cb\}$



Partitions, $P1 = \{La, Ca\}$
 $P2 = \{Lb, Cb\}$

Fig. 37A

55/56




Fig. 37B

56/56





Fig. 38

PATENT COOPERATION TREATY

PCT

DECLARATION OF NON-ESTABLISHMENT OF INTERNATIONAL SEARCH REPORT (PCT Article 17(2)(a), Rules 13ter.1(c) and (d) and 39)

Applicant's or agent's file reference LS-07-13PCT	IMPORTANT DECLARATION	Date of mailing (<i>day/month/year</i>) 05 NOVEMBER 2007 (05.11.2007)
International application No. PCT/US2007/071262	International filing date (<i>day/month/year</i>) 14 JUNE 2007 (14.06.2007)	(Earliest) Priority date (<i>day/month/year</i>) 14 JUNE 2006 (14.06.2006)
International Patent Classification (IPC) or both national classification and IPC G06F 17/50(2006.01)i, H01L 21/82(2006.01)i		
Applicant LIGHTSPEED LOGIC, INC. et al		

This International Searching Authority hereby declares, according to Article 17(2)(a), that **no international search report will be established** on the international application for the reasons indicated below.

1. ☐ The subject matter of the international application relates to:
 - a. ☐ scientific theories.
 - b. ☐ mathematical theories.
 - c. ☐ plant varieties.
 - d. ☐ animal varieties.
 - e. ☐ essentially biological processes for the production of plants and animals, other than microbiological processes and the products of such processes.
 - f. ☐ schemes, rules or methods of doing business.
 - g. ☐ schemes, rules or methods of performing purely mental acts.
 - h. ☐ schemes, rules or methods of playing games.
 - i. ☐ methods for treatment of the human body by surgery or therapy.
 - j. ☐ methods for treatment of the animal body by surgery or therapy.
 - k. ☐ diagnostic methods practised on the human or animal body.
 - l. ☐ mere presentation of information.
 - m. ☐ computer programs for which this International Searching Authority is not equipped to search prior art.
2. ☒ The failure of the following parts of the international application to comply with prescribed requirements prevents a meaningful search from being carried out:

☐ the description
 ☒ the claims
 ☐ the drawings
3. ☐ A meaningful search could not be carried out without the sequence listing; the applicant did not, within the prescribed time limit:

☐ furnish a sequence listing on paper complying with the standard provided for in Annex C of the Administrative Instructions, and such listing was not available to the International Searching Authority in a form and manner acceptable to it.
☐ furnish a sequence listing in electronic form complying with the standard provided for in Annex C of the Administrative Instructions, and such listing was not available to the International Searching Authority in a form and manner acceptable to it.
☐ pay the required late furnishing fee for the furnishing of a sequence listing in response to an invitation under Rule 13ter.1(a) or (b)
4. ☐ A meaningful search could not be carried out without the tables related to the sequence listings; the applicant did not, within the prescribed time limit, furnish such tables in electronic form complying with the technical requirements provided for in Annex C-bis of the Administrative Instructions, and such tables were not available to the International Searching Authority in a form and manner acceptable to it.
5. Further comments:

Name and mailing address of ISA/KR  Korean Intellectual Property Office 920 Dunsan-dong, Seo-gu, Daejeon 302-701, Republic of Korea Facsimile No. 82-42-472-7140	Authorized officer SUN Dong Guk Telephone No. 82-42-481-8248 
---	---