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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
3 July 2014 (03.07.2014)

WIPO | PCT

(10) International Publication Number
WO 2014/106098 A1

(51) International Patent Classification:
G06F 17/50 (2006.01)

(21) International Application Number:
PCT/US2013/078073

(22) International Filing Date:
27 December 2013 (27.12.2013)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/747,973 31 December 2012 (31.12.2012) US
14/139,568 23 December 2013 (23.12.2013) US

(71) Applicant: **SYNOPSYS, INC.** [US/US]; 700 E. Middlefield Road, Mountain View, California 94043 (US).

(72) Inventors: **TSAI, Jeng-Liang**; 1360 Poe Lane, San Jose, California 95130 (US). **HO, Pei-Hsin**; 2572 NW Jonathon Place, Portland, Oregon 97229 (US).

(74) Agent: **SAHASRABUDDHE, Laxman**; 2820 Fifth Street, Davis, California 95618 (US).

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

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: A PRIORI CORNER AND MODE REDUCTION

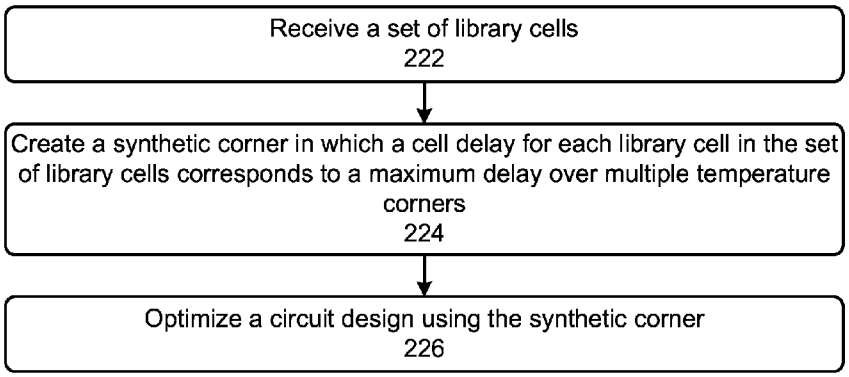


FIG. 2B

(57) Abstract: Systems and techniques are described for performing a priori corner and mode reduction. Some embodiments create a synthetic corner in which (1) a cell delay for each library cell in a set of library cells corresponds to a maximum delay over multiple temperature corners, and/or (2) a cell delay for each library cell in a set of library cells corresponds to a maximum delay over multiple parasitic corners. Some embodiments can identifying, for a given corner, a portion of the circuit design that is common across multiple modes, and then replace the multiple modes with a single mode for optimizing and verifying timing constraints of the portion of the circuit design that is common across the multiple modes. The circuit design can then be optimized over the reduced set of modes and/or corners.

WO 2014/106098 A1

A PRIORI CORNER AND MODE REDUCTION

Inventors: Jeng-Liang Tsai and Pei-Hsin Ho

BACKGROUND

Technical Field

[0001] This disclosure generally relates to reducing the number of corner and mode combinations over which a circuit design needs to be optimized. More specifically, this disclosure relates to a priori corner and mode reduction.

Related Art

[0002] Rapid advances in computing technology have made it possible to perform trillions of computational operations each second on data sets that are sometimes as large as trillions of bytes. These advances can be largely attributed to the exponential increase in the size and complexity of integrated circuits.

[0003] To ensure that a manufactured chip will perform as desired, the circuit design is usually checked for compliance with a large number of design requirements. In fact, a circuit design is typically not manufactured until it satisfies the design requirements which can include, but are not limited to, timing requirements, noise requirements, leakage requirements, etc.

[0004] In a typical circuit design flow, a compliance checking tool is used to identify design requirement violations in a circuit design. Incremental adjustments are then made to the circuit design to fix the design requirement violations. Once the circuit design is violation free, the circuit design can be readied for manufacturing.

[0005] The advance to smaller semiconductor geometries has increased the variability in process and operating conditions of integrated circuits. A process corner is a particular set of process conditions under which the circuit design may be manufactured. An operating corner is a particular set of operating conditions (e.g., temperature, voltage) under which the circuit design may operate.

[0006] Modern circuit designs typically have multiple operating modes (e.g., stand-by, mission, test, etc.). Specifically, some parts of the circuit design may be completely switched off in certain operating modes. Further, different operating modes can have drastically different voltages, clock speeds, etc.

[0007] It is desirable to ensure that a circuit design meets the functional and performance goals under all combinations of process corners, operating conditions, and/or operating modes. The term "Multi-Corner/Multi-Mode (MCMM) scenario," or "scenario" for short, refers to a combination of a particular process corner, a particular operating condition, and/or a particular operating mode.

[0008] The amount of time required to fix design requirement violations is increasing rapidly because the size and complexity of circuit designs is increasing rapidly and because the number of scenarios over which circuit designs need to be checked is increasing rapidly.

SUMMARY

[0009] Systems and techniques are described for performing a priori corner and mode reduction. After the number of modes and/or corners has been reduced, the circuit design can then be optimized over only the reduced set of modes and/or corners, thereby significantly improving runtime performance.

[0010] Specifically, some embodiments receive a set of library cells. Next, the embodiments can create a synthetic corner in which a cell delay for each library cell in the set of library cells corresponds to a maximum delay over multiple temperature corners. The embodiments can then optimize a circuit design using the synthetic corner.

[0011] In some embodiments, a synthetic corner can be created in which a cell delay for each library cell in the set of library cells corresponds to a maximum delay over multiple parasitic corners. The embodiments can then optimize a circuit design using the synthetic corner.

[0012] Some embodiments can identify, for a given corner, a portion of the circuit design that is common across multiple modes. Then, the multiple modes can be replaced with a single mode, and the portion of the circuit design can be optimized only for the single mode. Note that timing information for the portion of the design is computed and checked against the constraints in only the single mode (as opposed to being computed and checked in each of the multiple modes).

BRIEF DESCRIPTION OF THE FIGURES

[0013] FIG. 1 illustrates how a net delay can be synthesized for a two-pin net in accordance with some embodiments described herein.

[0014] FIG. 2A illustrates a process for performing a priori mode reduction in accordance with some embodiments described herein.

[0015] FIG. 2B illustrates a process for performing a priori corner reduction in accordance with some embodiments described herein.

[0016] FIG. 2C illustrates a process for performing a priori corner reduction in accordance with some embodiments described herein.

[0017] FIG. 3 illustrates a computer system in accordance with some embodiments described herein.

DETAILED DESCRIPTION

[0018] The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein. In this disclosure, when the term "and/or" is used with a list of entities, it refers to all possible combinations of the list of entities. For example, the phrase "X, Y, and/or Z" covers the following cases: (1) only X; (2) only Y; (3) only Z; (4) X and Y; (5) X and Z; (6) Y and Z; and (7) X, Y, and Z. Additionally, in this disclosure, the term "based on" means "based solely or partially on."

Overview of an electronic design automation (EDA) flow

[0019] An EDA flow can be used to create a circuit design. Once the circuit design is finalized, it can undergo fabrication, packaging, and assembly to produce integrated circuit chips. An EDA flow can include multiple steps, and each step can involve using one or more EDA software tools. Some EDA steps and software tools are described below. These examples of EDA steps and software tools are illustrative purposes only and are not intended to limit the embodiments to the forms disclosed.

[0020] Some EDA software tools enable circuit designers to describe the functionality that they want to implement. These tools also enable circuit designers to perform what-if planning to refine functionality, check costs, etc. During logic design and functional verification, the HDL (hardware description language), e.g., SystemVerilog, code for modules in the system can be written and the design can be checked for functional accuracy, e.g., the design can be checked to ensure that it produces the correct outputs.

[0021] During synthesis and design for test, the HDL code can be translated to a netlist using one or more EDA software tools. Further, the netlist can be optimized for the target technology, and tests can be designed and implemented to check the finished chips. During

netlist verification, the netlist can be checked for compliance with timing constraints and for correspondence with the HDL code.

[0022] During design planning, an overall floorplan for the chip can be constructed and analyzed for timing and top-level routing. During physical implementation, circuit elements can be positioned in the layout (placement) and can be electrically coupled (routing).

[0023] During analysis and extraction, the circuit's functionality can be verified at a transistor level and parasitics can be extracted. During physical verification, the design can be checked to ensure correctness for manufacturing, electrical issues, lithographic issues, and circuitry.

[0024] During resolution enhancement, geometric manipulations can be performed on the layout to improve manufacturability of the design. During mask data preparation, the design can be "taped-out" to produce masks which are used during fabrication.

A priori corner and mode reduction

[0025] In a place and route design flow, the runtime of pre-route optimization is extremely important. This is because designers usually perform pre-route optimization multiple times to identify timing issues in the design and make design modifications accordingly, and then move on to next stage only after satisfactory timing has been achieved in the pre-route stage.

[0026] Designs from 65nm and below all have to be optimized over multiple corner and mode combinations and the number of active scenarios (active corner/mode combinations) typically dictates the optimization runtime. If the number of active scenarios can be reduced without loss of quality of results (QoR), it will greatly improve the optimization runtime and the total turn-around-time (TAT) for customers.

[0027] From 65nm onwards, designs are usually timed at both high and low temperatures at a given operating voltage to capture temperature inversion effect. Normally, logic gates switch slower at higher temperatures. Temperature inversion is a phenomenon where some logic gates (especially high voltage threshold (HVT) cells under low supply voltage) can switch slower at a low temperature than at a high temperature. Due to temperature inversion effect, the worst negative slack (WNS) may be worse at a low temperature than at a high temperature. Therefore, an optimization engine cannot guarantee to improve WNS without looking at both high and low temperature corners. One approach to ensure that an optimization engine captures this effect is to consider a pair of corners (for the same process and voltage parameters). However, this can significantly slow down pre-route optimization (because the optimization has to be performed over twice as many scenarios).

[0028] Some embodiments described in this disclosure can capture the temperature inversion effect without increasing the optimization run time. Specifically, a synthetic corner capability is described in this disclosure to capture temperature inversion effect without the need to activate two corners. In some embodiments, the optimization tool analyzes and predicts which types of gates are expected to switch slower at a low temperature in the final optimized design. The optimization tool can then activate only the high temperature corner, but instruct the timer to calculate cell delays for those gate types that are expected to exhibit temperature inversion using low temperature delay tables. This allows the optimization tool to achieve a significant runtime gain while still being able to model the temperature inversion effect, thus achieving a QoR that is similar to the QoR that would have been achieved if both the high temperature corner and the low temperature corner had been activated.

[0029] The runtime of an MCMM optimization engine is proportional to the number of active scenarios that determines the amount of timing update computation. Carefully managing the number of active scenarios will direct impact the runtime of an MCMM optimization engine.

[0030] Some conventional techniques use dominant scenarios to reduce the number of MCMM scenarios. In these techniques, a subset of scenarios is identified based on timing analysis. The scenarios in this subset are called dominant scenarios, and these scenarios are more timing critical than other scenarios. MCMM optimization is then performed over only this subset of scenarios. The MCMM optimization engine is able to capture most timing violations by activating only the dominant scenarios. In some of these approaches, a timing update is performed for all scenarios, endpoint violations are analyzed, and a binate covering problem is constructed to identify the dominant scenarios, which are then used for optimization. The advantage of this approach is that the approach does not depend on specific timer features (except for an API to get endpoint violation information efficiently). However, this approach also has a few drawbacks. First, each dominant scenario analysis requires a full timing update for all scenarios. Second, the dominant scenario set may change as optimization progresses (in fact, in some implementations, the dominant scenario set has to be updated twice, and therefore, these approaches can require three full timing updates for all scenarios). Third, the dominant scenario set may still contain too many scenarios (in one implementation of this approach, some scenarios can be periodically deactivated to save more runtime, but this can potentially create convergence problem). Fourth, some implementations of this approach can only handle one setup scenario at a time. Therefore, an optimization solution needs to be timed in all activated scenarios (potentially with different modes) to capture design changes that improve timing in one scenario but degrade timing in another scenario. This can lead to high rejection rate and runtime.

[0031] To summarize, these techniques adopt a "dynamic" timing analysis approach to identify dominant scenarios, which requires a full timing update and the analysis has to be performed multiple times. Coupled with the infrastructure limitation in some of the implementations that can look at only one setup scenario, many solutions generated by these approaches can get rejected and the work wasted, leading to long runtime.

[0032] In contrast to such "dynamic" techniques, some embodiments described herein use a "static" analysis approach to reduce number of active corners/modes. If the timing tool already separates corner and mode, it is natural to decompose the scenario reduction problem into two orthogonal problems, corner reduction and mode reduction. If modes are carefully created with no redundancy, all modes (e.g. functional and test modes) will likely need to be considered simultaneously and there is less opportunity for mode reduction. It does not mean mode reduction is not possible. For example, if a large chunk of the logic has the same boundary timing across multiple modes for a given corner, the delay calculation for the common logic may be shared.

[0033] Most data paths are cell delay dominant, and cell delay is affected by process (P) variations, voltage (V) variations, temperature (T) variations, as well as the input slew and the output load. It is easy to identify which corners have the slowest process (P) and lowest voltage (V). If multiple corners have the same (P, V) values, a cell will be slower in the corner with the largest wire capacitance due to the larger output load. Therefore, it is easy to narrow down the corners that are critical for setup optimization to a pair of corners that are identical in P and V but differ in T, for the temperature inversion effect mentioned before.

[0034] Note that the "static" corner analysis described herein is very fast because it does not rely on a full timing update. Also, the result does not change as the optimization progresses. The ultimate goal of corner reduction is one-corner-per-mode. This can be achieved if temperature inversion effect can be captured without activating the low temperature corner. In static timing analysis (STA), it is a common technique to apply a "MAX" operation (this operation computes the maximum value from a given set of values) to solve the complexity and runtime problem. For example, STA usually takes the max transition from all cell arcs at an output pin, and propagates the maximum value forward. This results in some pessimism but is a commonly accepted approach due to its runtime advantage. Similarly, some embodiments described herein "create" a synthetic corner, in which the cell delay of any cell instance is the max cell delay across high and low temperature corners. Some embodiments then perform setup optimization using the synthetic corner, where critical paths from both corners will be seen by these embodiments.

[0035] Note that there may not be any runtime gain if the timer has to calculate the cell delays from both corners (i.e., high and low temperature corners) to find out the max delay of the two. Therefore, some embodiments pre-characterize the technology parameters to determine the nominal input transition and output load of each gate type, and determine whether the max gate delay will occur at high or low temperature. The timer can then use this information to realize a synthetic corner, where gate delays are (approximately) the worst from two corners.

[0036] With the synthetic corner support from timer, some embodiments are able to perform setup optimization in a one-synthetic corner-per-mode setting. Some embodiments then look at all timing numbers from those corner/mode combinations simultaneously during optimization to avoid the solution rejection and convergence issues seen in some conventional approaches.

[0037] While performing pre-route optimization, if chemical-mechanical polishing (CMP) variation is a concern, then multiple corners can be created that differ only in the resistance and capacitance values that are generated for nets. Usually, the same cell library will be paired with a maximum capacitance model and a maximum resistance model to form two corners, which are hereinafter referred to as the "maximum parasitic capacitance corner" and the "maximum parasitic resistance corner," respectively. Since data paths are usually cell-delay dominant, most paths will have larger path delay in the maximum parasitic capacitance corner due to larger load seen by cells. However, paths containing long wires may have worse path delays in the maximum parasitic resistance corner due to larger RC delay.

[0038] To achieve the best pre-route optimization runtime, optimization in some embodiments aims to activate only one corner per mode. However, without activating the maximum parasitic resistance corner, some embodiments would not be able to identify and buffer the long wires that could create timing or transition issues in the maximum parasitic resistance corner. Therefore, some embodiments feature a synthetic net delay capability to capture the timing impact from both the maximum parasitic capacitance and resistance corners without the need to activate two corners. Specifically, some embodiments determine which corners correspond to the same process, voltage, and temperature values, but have different parasitic corners. The embodiments then activate only one of the corners, but instruct the net estimation engine to look at net delays and capacitances in all corners. The net estimation engine then generates synthesized net delays and capacitances for the active corner from data collected across corners. For example, the net estimation engine may present to the timer a net capacitance from the maximum parasitic corner and a net delay from the larger of the two in both the maximum parasitic capacitance and resistance corners.

[0039] To capture the timing effect from multiple parasitic corners without introducing too much timing pessimism, the net estimation engine may synthesize a net delay or a net capacitance that is not equal to the values in either the maximum parasitic resistance or capacitance corners. However, while synthesizing the net delay, care must be taken so that the assumptions that the timer currently has on net delay are not violated. For example, if the timer assumes that a net delay would not change when its driver is sized, then this assumption must not be violated when synthesizing the net delay.

[0040] An example of how the net delay can be synthesized is now described. In different parasitic corners, the resistance and capacitance values usually vary in opposite directions (i.e., when the resistance increases from a first corner to a second corner, then the capacitance usually decreases, and vice versa). For example, when wire thickness or width is reduced due to dishing or erosion from CMP, resistance will increase while capacitance will decrease due to reduced wire surface area.

[0041] In the virtual routing stage, one can roughly assume that if a wire has resistance *Rw* and capacitance *Cw* in the maximum parasitic capacitance corner, then the same wire will have resistance $kr \times Rw$ and capacitance $kc \times Cw$ in the maximum parasitic resistance corner, where $kr > 1$ and $kc < 1$ are coefficients that can be selected depending on the circuit design and/or technology. In some circuit designs, *kr* was roughly equal to 4 while *kc* was roughly equal to 0.5.

[0042] FIG. 1 illustrates how a net delay can be synthesized for a two-pin net in accordance with some embodiments described herein. An output of gate 102 is electrically connected to an input of gate 104 (gates 102 and 104 can generally be any logic gates). Without loss of generality, let *Rw* and *Cw* be the resistance and capacitance values for the wire in the maximum parasitic capacitance corner, and let *Cp* be the pin load at the end of the wire, e.g., let *Cp* be the input capacitance of gate 104.

[0043] Let $D(C_{tot})$ be the driver delay when it drives a load of C_{tot}. The stage (driver and net) delay in the maximum parasitic capacitance corner can be approximated as follows :

$$SD_{maxC} = D(Cw + Cp) + Rw \cdot (0.5 \cdot Cw + Cp) = D(C_{tot1}) + ND_1, \qquad (1)$$

where the subscript "*maxC*" refers to the maximum parasitic capacitance corner, capacitance C_{tot1} is equal to $Cw + Cp$, and ND_1 is the net delay (*Cw* is multiplied by 0.5 in the above equation because the total wire capacitance *Cw* can be approximated as two lumped capacitances of $0.5 \times Cw$ at the two ends of the wire).

[0044] Likewise, the stage delay in the maximum parasitic resistance corner can be approximated as follows:

$$SD_{maxR} = D(kc \cdot Cw + Cp) + kr \cdot Rw \cdot (0.5 \cdot kc \cdot Cw + Cp) = D(C_{tot2}) + ND_2, \quad (2)$$

where the subscript "*maxR*" refers to the maximum parasitic resistance corner, capacitance C_{tot2} is equal to $kc \times Cw + Cp$, and ND_2 is the net delay.

[0045] The net estimation engine can control the stage delay by providing the appropriate C_{synth} and ND_{synth} values to the timer (e.g., the piece of software that computes timing in the circuit design) so that the timer computes the stage delay as $SD_{synth} = D(C_{synth}) + ND_{synth}$.

[0046] For example, in some embodiments, the net estimation engine can synthesize C_{synth} and ND_{synth} as follows: $C_{synth} = \max(C_{tot1}, C_{tot2})$ and $ND_{synth} = \max(ND_1, ND_2)$. Note that, when $kr \times kc > 1$, this strategy is equivalent to $C_{synth} = C_{tot1}$ and $ND_{synth} = ND_2$. However, this will lead to a stage delay that is greater than the stage delay in either of the two corners, i.e., $SD_{synth} > \max(SD_{maxC}, SD_{maxR})$. In other embodiments that are less pessimistic, the net estimation engine provides C_{synth} and ND_{synth} values so that $SD_{synth} = \max(SD_{maxC}, SD_{maxR})$.

Process for performing a priori corner and mode reduction

[0047] FIG. 2A illustrates a process for performing a priori mode reduction in accordance with some embodiments described herein. The process can begin by identifying, for a given corner, a portion of the circuit design that is common across multiple modes (operation 202). Next, the process can replace the multiple modes with a single mode for optimizing the portion of the circuit design that is common across the multiple modes, wherein the single mode captures all timing requirements across the multiple modes (operation 204). Recall that different portions of the circuit design may be "switched off" or deactivated in different modes. However, it is possible that a portion of the circuit design is active in multiple modes. If so, the multiple modes can be replaced by a single mode as far as the common portion of the circuit design is concerned, thereby enabling the common portion of the circuit design to be optimized (which involves computing the timing information and checking if all timing constraints are satisfied) across all modes at once. Note that the single mode captures all timing constraints across all of the multiple modes. Therefore, the embodiments only have to verify timing for the optimized circuit design once in the single mode instead of verifying timing constraints in all of the modes separately. This saving in the timing computation can significantly speed up the runtime for the optimization process.

[0048] Of course the "mode merging" described above can only be performed for a given corner. For example, if the circuit design needs to be optimized over modes $M1$ and $M2$ for corner $C1$ (i.e., over scenarios ($M1$, $C1$) and ($M2$, $C1$)), and if logic L is common to modes $M1$ and $M2$, then logic L can be optimized and the timing constraints can be checked only once for corner $C1$, instead of requiring the timing constraints to be computed and checked separately for

scenarios (*M*1, *C*1) and (*M*2, *C*1). However, a similar type of mode reduction cannot be performed between modes for two different corners, e.g., between scenarios (*M*1, *C*1) and (*M*2, *C*2).

[0049] FIG. 2B illustrates a process for performing a priori corner reduction in accordance with some embodiments described herein. The process can begin by receiving a set of library cells (operation 222). Next, the process can create a synthetic corner in which a cell delay for each library cell in the set of library cells corresponds to a maximum delay over multiple temperature corners (operation 224). The process can then optimize a circuit design using the synthetic corner (operation 226).

[0050] FIG. 2C illustrates a process for performing a priori corner reduction in accordance with some embodiments described herein. The process can begin by receiving a set of library cells (operation 232). Next, the process can create a synthetic corner in which each library cell in the set of library cells corresponds to a maximum delay over multiple parasitic corners (operation 234). The process can then optimize a circuit design using the synthetic corner (operation 236).

[0051] In some embodiments, the above-described corner and mode reductions can be combined. For example, in some embodiments the process can perform one or more of the following operations: operations 202, 224, and 234. Specifically, in some embodiments, the process can create a synthetic corner in which each library cell in a set of library cells is associated with a maximum delay over multiple combinations of temperatures and parasitic corners. Next, for a given synthetic corner, the process can merge modes that share common logic. Once the modes and/or corners have been reduced, the process can perform optimization using the merged modes and/or synthetic corners.

Computer system

[0052] FIG. 3 illustrates a computer system in accordance with some embodiments described herein. A computer or a computer system can generally be any system that can perform computations. Specifically, a computer system can be a microprocessor, an application specific integrated circuit, a distributed computing system, a cloud computing system, or any other computing system now known or later developed. Computer system 302 comprises processor 304, memory 306, and storage 308. Computer system 302 can be coupled with display 314, keyboard 310, and pointing device 312. Storage 308 can generally be any device that can store data. Specifically, a storage device can be a magnetic, an optical, or a magneto-optical storage device, or it can be based on flash memory and/or battery-backed up memory. Storage 308 can store application 316, operating system 318, and data 320.

[0053] Application 316 can include instructions that when executed by computer 302 cause computer 302 to perform one or more processes that are implicitly or explicitly described in this disclosure. Data 320 can include any data that is inputted into or outputted by application 316.

[0054] The above description is presented to enable any person skilled in the art to make and use the embodiments. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein are applicable to other embodiments and applications without departing from the spirit and scope of the present disclosure. Thus, the present invention is not limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

[0055] The data structures and code described in this disclosure can be partially or fully stored on a non-transitory computer-readable storage medium and/or a hardware module and/or hardware apparatus. A non-transitory computer-readable storage medium includes all computer-readable storage mediums with the sole exception of a propagating electromagnetic wave or signal. Specifically, a non-transitory computer-readable storage medium includes, but is not limited to, volatile memory, non-volatile memory, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs), DVDs (digital versatile discs or digital video discs), or other media, now known or later developed, that are capable of storing code and/or data. Hardware modules or apparatuses described in this disclosure include, but are not limited to, application-specific integrated circuits (ASICs), field-programmable gate arrays (FPGAs), dedicated or shared processors, and/or other hardware modules or apparatuses now known or later developed.

[0056] The methods and processes described in this disclosure can be partially or fully embodied as code and/or data stored in a non-transitory computer-readable storage medium or device, so that when a computer system reads and executes the code and/or data, the computer system performs the associated methods and processes. The methods and processes can also be partially or fully embodied in hardware modules or apparatuses. Note that the methods and processes can be embodied using a combination of code, data, and hardware modules or apparatuses.

[0057] The foregoing descriptions of embodiments of the present invention have been presented only for purposes of illustration and description. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.

What Is Claimed Is:

1. A method for performing a priori corner and mode reduction, the method comprising:

receiving a set of library cells;

creating a synthetic corner in which a cell delay for each library cell in the set of library cells corresponds to a maximum delay over multiple temperature corners; and

optimizing, by computer, a circuit design using the synthetic corner.

2. The method of claim 1, wherein the method further comprises:

identifying, for a corner, a portion of the circuit design that is common across multiple modes; and

replacing the multiple modes with a single mode for optimizing and verifying timing constraints of the portion of the circuit design that is common across the multiple modes.

3. The method of claim 1, wherein the cell delay in the synthetic corner is equal to a cell delay in a high temperature corner when temperature inversion does not occur, and wherein the cell delay in the synthetic corner is equal to a cell delay in a low temperature corner when temperature inversion occurs.

4. A method for performing a priori corner and mode reduction, the method comprising:

receiving a set of library cells;

creating a synthetic corner in a cell delay for each library cell in the set of library cells corresponds to a maximum delay over multiple parasitic corners; and

optimizing, by computer, a circuit design using the synthetic corner.

5. The method of claim 4, wherein the method further comprises:

identifying, for a corner, a portion of the circuit design that is common across multiple modes; and

replacing the multiple modes with a single mode for optimizing and verifying timing constraints of the portion of the circuit design that is common across the multiple modes.

6. The method of claim 4, wherein the cell delay in the synthetic corner is equal to a maximum of the following two cell delays: a cell delay in a maximum parasitic capacitance corner and a cell delay in a maximum parasitic resistance corner.

7. A method for performing a priori corner and mode reduction, the method comprising:

receiving a circuit design;

identifying, for a corner, a portion of the circuit design that is common across multiple modes;

replacing the multiple modes with a single mode; and

optimizing, by computer, the portion of the circuit design for the single mode.

8. The method of claim 7, wherein the corner is a synthetic corner in which a cell delay for each library cell in a set of library cells corresponds to a maximum delay over multiple temperature corners.

9. The method of claim 7, wherein the corner is a synthetic corner in which each library cell in a set of library cells corresponds to a maximum delay over multiple parasitic corners.

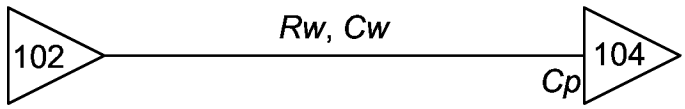


FIG. 1

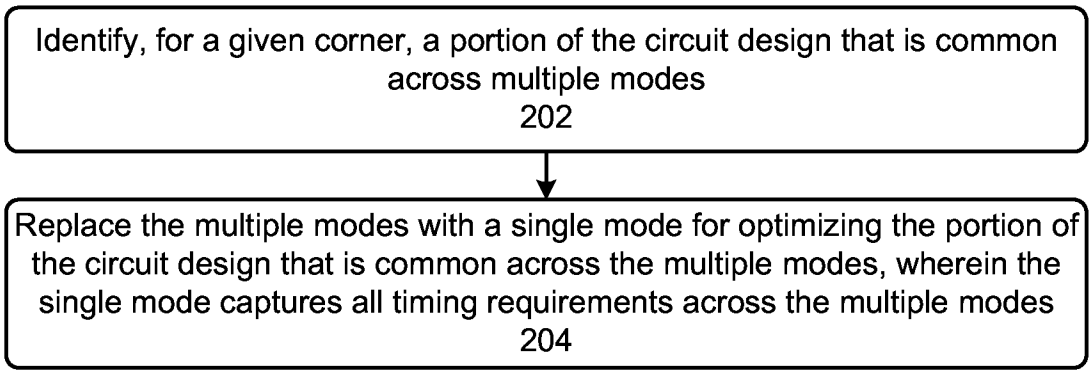


FIG. 2A

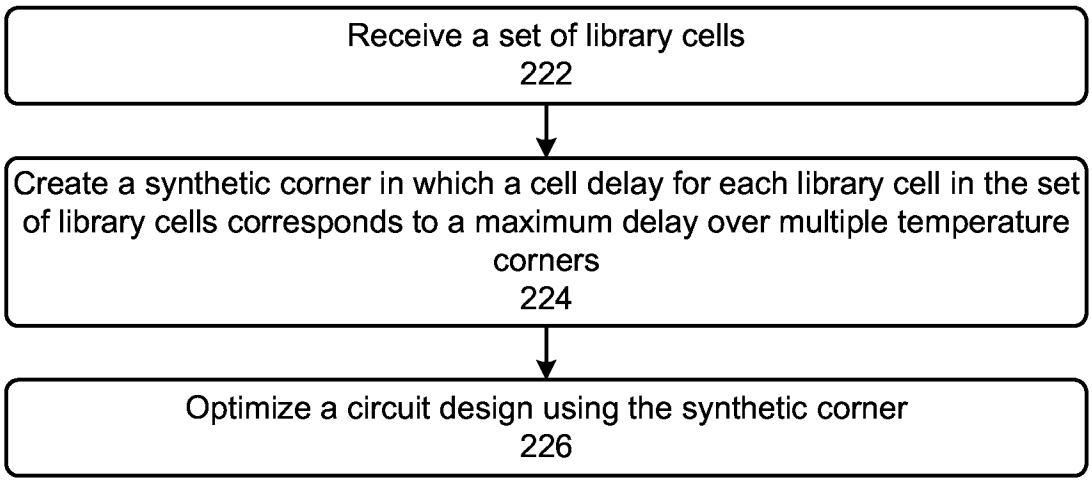


FIG. 2B

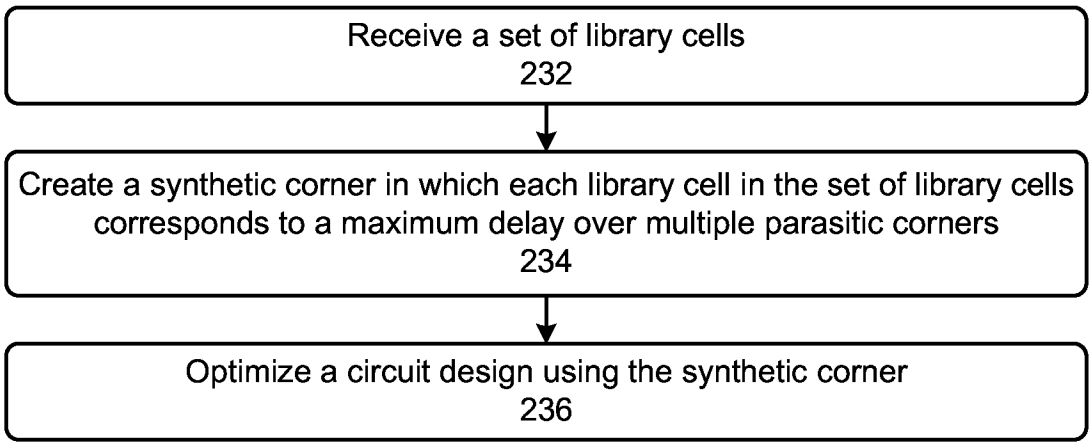

Receive a set of library cells
232
Create a synthetic corner in which each library cell in the set of library cells corresponds to a maximum delay over multiple parasitic corners
234
Optimize a circuit design using the synthetic corner
236

FIG. 2C

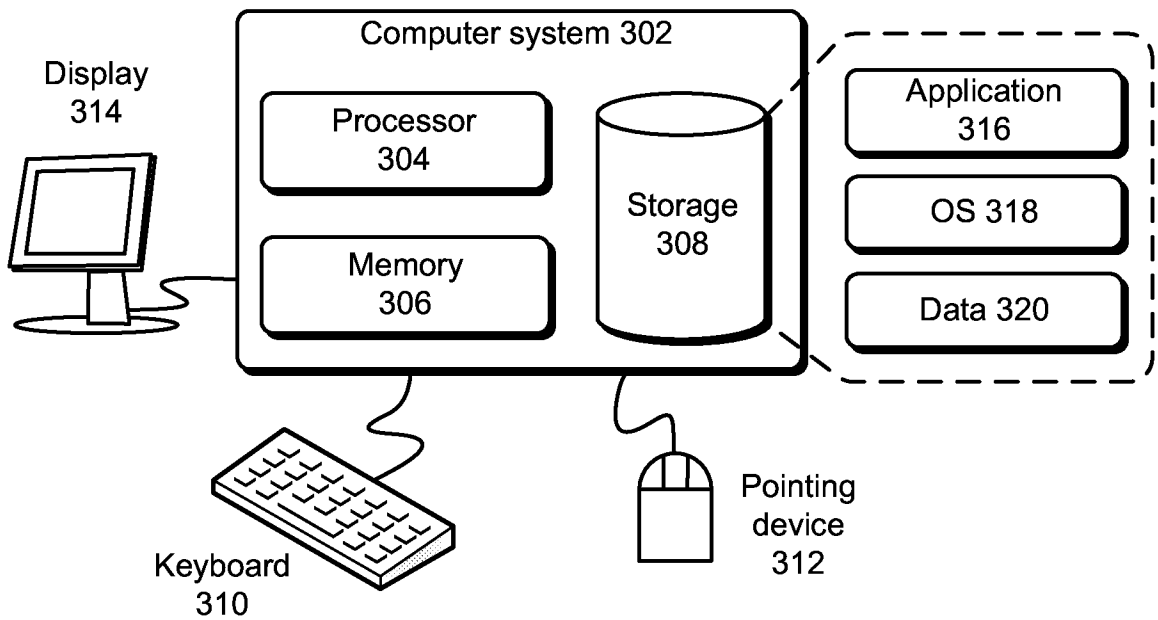

Computer system 302
Display
314
Processor
304
Memory
306
Storage
308
Application
316
OS 318
Data 320
Keyboard
310
Pointing
device
312

FIG. 3

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/078073

A. CLASSIFICATION OF SUBJECT MATTER

G06F 17/50(2006.01)i

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F 17/50

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & keywords: circuit design, delay, temperature, common, replace, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2011-0209113 A1 (KHALID RAHMAT et al.) 25 August 2011 See paragraphs [0045], [0047], [0085], [0093], [0132]; and figures 1-8.	1-9
A	US 2011-0231811 A1 (QIAN-YING TANG et al.) 22 September 2011 See paragraphs [0073]-[0081]; and figures 3-5.	1-9
A	US 2011-0078639 A1 (XI-WEI LIN et al.) 31 March 2011 See paragraphs [0044]-[0052], [0074]-[0080]; and figures 5, 10-12.	1-9
A	US 2010-0115477 A1 (CHRISTOPH ALBRECHT et al.) 06 May 2010 See paragraphs [0061]-[0070]; and figure 2.	1-9
A	US 2010-0169846 A1 (PUNEET GUPTA et al.) 01 July 2010 See paragraphs [0110]-[0113]; and figures 12-14.	1-9

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents,such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
26 March 2014 (26.03.2014)	**26 March 2014 (26.03.2014)**

Name and mailing address of the ISA/KR
International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea
Facsimile No. +82-42-472-7140

Authorized officer
KIM, Seong Woo
Telephone No. +82-42-481-3348

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.
PCT/US2013/078073

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0209113 A1	25/08/2011	US 2008-168406 A1	10/07/2008
		US 2014-053124 A1	20/02/2014
		US 7366997 B1	29/04/2008
		US 7941779 B2	10/05/2011
		US 8572535 B2	29/10/2013
US 2011-0231811 A1	22/09/2011	CN 102870119 A	09/01/2013
		EP 2548148 A2	23/01/2013
		US 2013-174115 A1	04/07/2013
		US 8359558 B2	22/01/2013
		WO 2011-116056 A2	22/09/2011
		WO 2011-116056 A3	24/11/2011
US 2011-0078639 A1	31/03/2011	CN 101681878 A	24/03/2010
		EP 2203935 A1	07/07/2010
		EP 2203935 A4	25/04/2012
		JP 05416700 B2	12/02/2014
		JP 2010-541280 A	24/12/2010
		TW 200921849 A	16/05/2009
		US 2009-0113368 A1	30/04/2009
		US 2013-332893 A1	12/12/2013
		US 7895548 B2	22/02/2011
		US 8504969 B2	06/08/2013
		WO 2009-055113 A1	30/04/2009
US 2010-0115477 A1	06/05/2010	US 2008-0276210 A1	06/11/2008
		US 7624364 B2	24/11/2009
		US 8589845 B2	19/11/2013
US 2010-0169846 A1	01/07/2010	US 2010-0169847 A1	01/07/2010
		US 7441211 B1	21/10/2008
		US 8127266 B1	28/02/2012
		US 8185865 B2	22/05/2012
		US 8490043 B2	16/07/2013