



- (51) **International Patent Classification:**
G06F 19/00 (2006.01) *G06F 17/50* (2006.01)
- (21) **International Application Number:**
PCT/US2009/052354
- (22) **International Filing Date:**
31 July 2009 (31.07.2009)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
12/210,887 15 September 2008 (15.09.2008) US
- (71) **Applicant** (for all designated States except US): **SYNOPTICS, INC.** [US/US]; 700 E. Middlefield Road, Mountain View, CA 94043 (US).
- (72) **Inventors; and**
- (75) **Inventors/Applicants** (for US only): **BIST, Manoj** [IN/US]; 551 Columbia Avenue, #c, Sunnyvale, CA 94085 (US). **MEHROTRA, Sandeep** [IN/US]; 5482 Ora Street, San Jose, CA 95129 (US).
- (74) **Agent:** **SUZUE, Kenta**; Haynes Beffel & Wolfeld LLP, P.O. Box 366, Half Moon Bay, CA 94019 (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, 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, 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, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, 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, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) **Title:** METHOD AND APPARATUS FOR MERGING EDA COVERAGE LOGS OF COVERAGE DATA

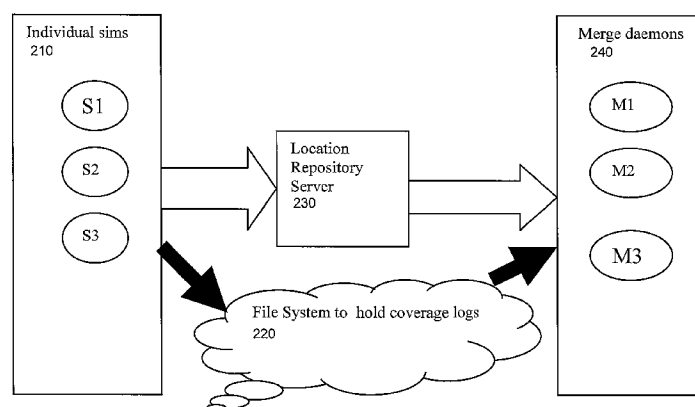


FIG. 2

(57) **Abstract:** An electronic design automation technology merges coverage logs. The coverage logs are generated by verification of a hardware description language circuit design. The coverage logs are merged as the coverage logs are generated, without waiting for all pending coverage logs. Another electronic design automation technology also merges coverage logs. The merged coverage logs include a first coverage log of a first simulation of a hardware description language circuit design and a second coverage log of a second simulation of the hardware description language circuit design. The first simulation is based on a first hardware verification language coverage model of the hardware description language circuit design. The second simulation is based on a second hardware verification language coverage model of the hardware description language circuit design. The second hardware verification language coverage model is newer and different than the first hardware verification language coverage model.

METHOD AND APPARATUS FOR MERGING EDA COVERAGE LOGS OF COVERAGE DATA

BACKGROUND

5 [0001] Electronic design automation EDA is applied in the semiconductor industry for virtually all device design projects. After an idea for the product is developed, EDA tools are utilized to define a specific implementation. The implementation defined using EDA tools is used to create mask data used for production of masks for lithographic use in the production of the finished chips, in a process referred to as tape-out. The masks are then created and used with
10 fabrication equipment to manufacture integrated circuit wafers. The wafers are diced, packaged and assembled to provide integrated circuit chips for distribution.

[0002] An exemplary procedure for design using EDA tools begins with an overall system design using architecture defining tools that describe the functionality of the product to be implemented using the integrated circuit. Next, logic design tools are applied to create a high
15 level description based on description languages such as Verilog or VHDL, and functional verification tools are applied in an iterative process to assure that the high-level description accomplishes the design goals. Next, synthesis and design-for-test tools are used to translate the high-level description to a netlist, optimize the netlist for target technology, and design and implement tests that permit checking of the finished chip against the netlist.

20 [0003] A typical design flow might next include a design planning stage, in which an overall floor plan for the chip is constructed and analyzed to ensure that timing parameters for the netlist can be achieved at a high level. Next, the netlist may be rigorously checked for compliance with timing constraints and with the functional definitions defined at the high level using VHDL or Verilog. After an iterative process to settle on a netlist and map the netlist to a cell library for the
25 final design, a physical implementation tool is used for placement and routing. A tool performing placement positions circuit elements on the layout, and a tool performing routing defines interconnects for the circuit elements.

[0004] The components defined after placement and routing are usually then analyzed at the transistor level using an extraction tool, and verified to ensure that the circuit function is
30 achieved and timing constraints are met. The placement and routing process can be revisited as needed in an iterative fashion. Next, the design is subjected to physical verification procedures, such as design rule checking DRC, layout rule checking LRC and layout versus schematic LVS

checking, that analyze manufacturability, electrical performance, lithographic parameters and circuit correctness.

[0005] After closure on an acceptable design by iteration through design and verify procedures, like those described above, the resulting design can be subjected to resolution enhancement techniques that provide geometric manipulations of the layout to improve manufacturability. Finally, the mask data is prepared and taped out for use in producing finished products.

SUMMARY

[0006] One aspect of the technology is an electronic design automation method that merges coverage logs. The coverage logs are generated by verification of a hardware description language circuit design. The coverage logs are merged as the coverage logs are generated, without waiting for all coverage logs of pending verification. Examples of pending verification are pending dynamic simulation (e.g., pure random simulation, directed random simulation, and pure directed simulation) and pending formal verification.

[0007] In some embodiments, at least one of the coverage logs is generated by simulation of the hardware description language circuit design. In some embodiments, merging coverage logs results in a merged coverage log including formal verification coverage data.

[0008] Various embodiments have results that are responsive to the merging of coverage logs. One such result is altering a condition of a pending simulation of the hardware description language circuit design. Examples of changing a condition include altering an input parameter, such as altering an input condition (e.g., changing the input configuration file and/or parameters to simulate the chip in a different mode/configuration, changing the control parameters of the testbench that would in turn add more constraints and/or relax existing constraints thereby causing different input stimulus to be generated) or changing a random seed condition.

[0009] Another such result is freeing at least part of mass storage.

[0010] Another such result – responsive to not just the merging of coverage logs but also to a predetermined condition being met – is generating a coverage report.

[0011] Another such result is determining that, completing a pending verification is expected to insufficiently improve verification coverage of the hardware description language circuit design. One way to determine that the pending verification is expected to insufficiently improve verification coverage, is determining that an attribute of the hardware description language circuit design to be simulated by a pending simulation has been simulated already. In one

embodiment, when it is determined that the pending verification is expected to insufficiently improve verification coverage, then the pending verification is stopped.

[0012] Another such result updating a coverage metric of verification coverage of the hardware description language circuit design. The coverage metric is inclusive of coverage data from merging the coverage logs. Example coverage metrics consider assertion coverage, functional coverage, and code coverage (e.g., line coverage, condition coverage, branch coverage, path coverage, toggle coverage, assign toggle coverage).

[0013] In some embodiments, the technology changes from a status of unreadiness, to readiness to merge coverage logs. In response to this status change, coverage log data are requested that was generated prior to the status change.

[0014] Some embodiments create multiple running instances. These multiple running instances perform merging of coverage logs of simulations of the hardware description language circuit design, or verification generally.

[0015] Another aspect of the technology is an electronic design automation method that merges coverage logs. The merged coverage logs include a first coverage log of a first simulation of a hardware description language circuit design and a second coverage log of a second simulation of the hardware description language circuit design. The first simulation is based on a first hardware verification language coverage model of the hardware description language circuit design. The second simulation is based on a second hardware verification language coverage model of the hardware description language circuit design. The second hardware verification language coverage model is newer and different than the first hardware verification language coverage model.

[0016] The hardware description language of the hardware description language circuit design includes any of Verilog, SystemVerilog, and VHDL.

[0017] The hardware verification language of the first and second hardware verification language models includes any of SystemVerilog, Native Testbench, E, and Vera.

[0018] Various embodiments have various conditions that control whether coverage data is deleted and/or retained, responsive to merging the coverage logs.

[0019] One such condition includes, responsive to the first coverage log and the second coverage log including different maximum numbers of automatically created bins, retaining coverage data of the cover point of the second coverage log and deleting coverage data of the cover point from the first coverage log.

[0020] Another such condition includes, responsive to the first coverage log including a bin having a bin name absent from the second coverage log, deleting coverage data of the bin after said merging.

[0021] Another such condition includes, responsive to the first coverage log including a bin having a bin name present in the second coverage log, retaining coverage data of the bin after said merging.

[0022] Another such condition includes, responsive to the first coverage log having a cover point with a first expression width and the second coverage log having the cover point with a second expression width different from the first expression width, retaining coverage data of the cover point with the second expression width and deleting coverage data of the cover point with first expression width.

[0023] Another such condition includes, responsive to the first coverage log having a bin with a first bin definition and the second coverage log has the bin with a second bin definition different from the first bin definition, then after said merging, retaining coverage data of the bin with the second bin definition and deleting coverage data of the first bin definition.

[0024] Another such condition includes, responsive to the first coverage log including a cross coverage point name of a cross coverage point naming at least two coverage point identifiers, and the second coverage log includes the cross coverage point name of the cross coverage point naming said at least two coverage point identifiers, then deleting coverage data of the cross coverage point of the first coverage log responsive to deletion of coverage data of at least one of the identified coverage points from the first coverage log.

[0025] Another such condition includes, responsive to the first coverage log including a first user-defined cross bin of a first cross coverage point, and the second coverage log including a second user-defined cross bin of a second cross coverage point, the first cross coverage point and the second cross coverage point having a same name, then despite the first user-defined cross bin and the second user-defined cross bin having different definitions, retaining coverage data of the first user-defined cross bin and coverage data of the second user-defined cross bin.

[0026] Another such condition includes, responsive to the first coverage log including first autocross bins of a first cross coverage point, and the second coverage log including a second user-defined cross bin subsuming the autocross bins of the first cross coverage point, retaining coverage data of the first autocross bins in the second user-defined cross bin of the second coverage log.

[0027] Another such condition includes, responsive to the first coverage log including a first cross coverage point, and the second coverage log including a second cross coverage point, then despite the first cross coverage point and the second cross coverage point having different auto cross coverage space, retaining i) auto cross coverage data of the first cross coverage point that is
5 subsumed by auto cross coverage space of the second cross coverage point, and ii) coverage data of the second cross coverage point.

[0028] Another such condition includes, responsive to the second coverage log including a cross coverage point absent from the first coverage log, retaining the cross coverage point.

[0029] Another such condition includes, responsive to the first coverage log including a
10 cross coverage point absent from the second coverage log, deleting the cross coverage point.

[0030] Another such condition includes, responsive to the first coverage log including property coverage absent from the second coverage log, retaining the property coverage.

[0031] Another such condition includes, responsive to the second coverage log including property coverage absent from the first coverage log, retaining the property coverage.

[0032] Other aspects of the technology are directed to a computer readable medium with
15 computer readable instructions for circuit design, including computer instructions performing the technology described herein. Other aspects of the technology are directed to a computer with a memory storing computer instructions performing the technology described herein.

20 BRIEF DESCRIPTION OF THE DRAWINGS

[0033] Figure 1 is a simplified process flow of making integrated circuits, from the idea stage, through electronic design automation (EDA) stages, and ending with the integrated circuits.

[0034] Figure 2 is a simplified block diagram of a system performing simulation of a
25 hardware verification language circuit design.

[0035] Figure 3 is a simplified state diagram showing the operation of the simulation system of Figure 2.

[0036] Figure 4 is a simplified directory structure of the simulation system of Figures 2 and
3.

[0037] Figures 5A and 5B are simplified flow charts contrasting an improved simulation
30 flow that merges, rather than discards, prior coverage logs.

[0038] Figure 6 is a simplified block diagram of a computer system of various embodiments.

DETAILED DESCRIPTION

[0039] Figure 1 is a simplified process flow of making integrated circuits, from the idea stage, through electronic design automation (EDA) stages, and ending with the integrated circuits.

5 [0040] Fig. 1 shows a simplified representation of an illustrative digital integrated circuit design and test flow. As with all flowcharts herein, it will be appreciated that many of the steps in Fig. 1 can be combined, performed in parallel or performed in a different sequence without affecting the functions achieved. In some cases a re-arrangement of steps will achieve the same results only if certain other changes are made as well, and in other cases a re-arrangement of
10 steps will achieve the same results only if certain conditions are satisfied. Such re-arrangement possibilities will be apparent to the reader.

[0041] At a high level, the process of Fig. 1 starts with the product idea (step 100) and is realized in an EDA (Electronic Design Automation) software design process (step 110). When the design is finalized, the fabrication process (step 150) and packaging and assembly processes
15 (step 160) occur resulting, ultimately, in finished integrated circuit chips (result 170). Some or all of the finished chips are tested in step 180 on a tester machine using predefined test vectors and expected responses.

[0042] The EDA software design process (step 110) is actually composed of a number of steps 112-130, shown in linear fashion for simplicity. In an actual integrated circuit design
20 process, the particular design might have to go back through steps until certain tests are passed. Similarly, in any actual design process, these steps may occur in different orders and combinations. This description is therefore provided by way of context and general explanation rather than as a specific, or recommended, design flow for a particular integrated circuit.

[0043] A brief description of the components steps of the EDA software design process (step
25 110) will now be provided.

[0044] System design (step 112): The designers describe the functionality that they want to implement, they can perform what-if planning to refine functionality, check costs, etc. Hardware-software architecture partitioning can occur at this stage. Example EDA software products from Synopsys, Inc. that can be used at this step include Model Architect, Saber,
30 System Studio, and DesignWare® products.

[0045] Logic design and functional verification (step 114): At this stage, the VHDL or Verilog code for modules in the system is written and the design is checked for functional accuracy. More specifically, the design is checked to ensure that it produces the correct outputs

in response to particular input stimuli. Example EDA software products from Synopsys, Inc. that can be used at this step include VCS, VERA, DesignWare®, Magellan, Formality, ESP and LEDA products. While some designs might at this stage already include certain design-for-test features such as scan chains and associated scan compression or decompression circuitry, these are not included in the terms "logic design" and "circuit design" as they are used herein.

[0046] Synthesis and design for test (DFT) (step 116): Here, the VHDL/Verilog is translated to a netlist. The netlist can be optimized for the target technology. Additionally, the implementation of a test architecture occurs in this step, to permit checking of the finished chips. Example EDA software products from Synopsys, Inc. that can be used at this step include Design Compiler®, Physical Compiler, Test Compiler, Power Compiler, FPGA Compiler, TetraMAX, and DesignWare® products. A current product for implementing a test architecture, with a few user-specified configuration settings as described above, is DFT MAX. DFT MAX is described in Synopsys, DFT MAX Adaptive Scan Compression Synthesis, Datasheet (2007), incorporated herein by reference.

[0047] Netlist verification (step 118): At this step, the netlist is checked for compliance with timing constraints and for correspondence with the VHDL/Verilog source code. Example EDA software products from Synopsys, Inc. that can be used at this step include Formality, PrimeTime, and VCS products.

[0048] Design planning (step 120): Here, an overall floor plan for the chip is constructed and analyzed for timing and top-level routing. Example EDA software products from Synopsys, Inc. that can be used at this step include Astro and IC Compiler products.

[0049] Physical implementation (step 122): The placement (positioning of circuit elements) and routing (connection of the same) occurs at this step. Example EDA software products from Synopsys, Inc. that can be used at this step include the Astro and IC Compiler products.

[0050] Analysis and extraction (step 124): At this step, the circuit function is verified at a transistor level, this in turn permits what-if refinement. Example EDA software products from Synopsys, Inc. that can be used at this step include AstroRail, PrimeRail, Primetime, and Star RC/XT products.

[0051] Physical verification (step 126): At this step various checking functions are performed to ensure correctness for: manufacturing, electrical issues, lithographic issues, and circuitry. Example EDA software products from Synopsys, Inc. that can be used at this step include the Hercules product.

[0052] Tape-out (step 127): This step provides the “tape-out” data for production of masks for lithographic use to produce finished chips. Example EDA software products from Synopsys, Inc. that can be used at this step include the CATS(R) family of products.

[0053] Resolution enhancement (step 128): This step involves geometric manipulations of the layout to improve manufacturability of the design. Example EDA software products from Synopsys, Inc. that can be used at this step include Proteus, ProteusAF, and PSMGen products.

[0054] Mask preparation (step 130): This step includes both mask data preparation and the writing of the masks themselves.

[0055] Aspects of the present technology relate particularly to “logic design and functional verification” described above.

[0056] Design and functional verification is the process of proving that the hardware description language circuit design satisfies the design specification of the circuit. Although both static and dynamic methods of verification are available, static methods tend to be limited to relatively small circuit blocks due to the limitations of performing formal proofs of the hardware description language circuit design. Dynamic methods rely on simulation of the hardware description language circuit design, and cannot achieve complete verification. To accomplish a reasonably thorough verification of the hardware description language circuit design with a modern complexity of millions of gates, simulation farms run numerous simulations with different conditions, such as different input conditions or different random seeds. Such simulation continues until the aggregate coverage metric(s) is/are satisfactory.

[0057] Due to the trend of decreasing cost for increased computation resources, each time the simulation model is adjusted, one of ordinary skill in the art would tend to re-run simulations to reach a satisfactory coverage metric(s) result.

[0058] The resulting coverage data from completed simulations of the hardware description language circuit design are stored in mass storage as coverage logs. The various Electronic Design Automation vendors generally have their own proprietary coverage logs that store only Electronic Design Automation coverage data and do not store general purpose data as with databases by Oracle, IBM, MySQL, etc. Coverage data is a set of interesting conditions, or attributes, that are expected to happen during simulation.

[0059] In some embodiments, the coverage logs are merged. The coverage model is a set of interesting conditions that we want to observe during simulation. Each simulation creates a coverage log that represents which portion of the coverage model was covered by that

simulation. The merging operation aggregates the coverage logs to compute which parts of the coverage model are collectively covered by all the simulations.

[0060] The simulations are performed according to the hardware verification language coverage model, which includes functional coverage and assertion coverage. Other
5 embodiments include various combinations of one or more of functional coverage, assertion coverage, and code coverage.

[0061] Coverage data are discussed in “IEEE Standard for SystemVerilog — Unified Hardware Design, Specification, and Verification Language”, IEEE Std 1800, hereby
10 incorporated by reference. The following discussion of coverage data is largely taken from this IEEE standard.

[0062] A coverage group, or covergroup, can contain one or more coverage points, or coverpoints. A coverage point can be an integral variable or an integral expression. Each
15 coverage point includes a set of bins associated with its sampled values or its value transitions. The bins can be explicitly defined by the user or automatically created by the hardware verification language environment. Bins for a coverage point can be automatically created by
the hardware verification language environment or explicitly defined.

[0063] A coverage point bin, or bin, associates a name and a count with a set of values or a
20 sequence of value transitions. If the bin designates a set of values, the count is incremented every time the coverage point matches one of the values in the set. If the bin designates a sequence of value transitions, the count is incremented every time the coverage point matches
the entire sequence of value transitions.

[0064] Thus, a bin is a unit identifying an attribute of the hardware description language
25 circuit undergoing simulation, such that after simulation the coverage metric of the model is calculated based on the bins. The coverage metric is an assessment of progress of verifying a hardware description language circuit design. The coverage model is the aggregate of all bins,
including cross points of bins and cover points of bins. In some embodiments, the same model is used for all simulations, and the resulting coverage data are merged. In other embodiments,
different models are used for at least some simulations, and the resulting coverage data are
selectively merged.

30 [0065] A cover group can specify cross coverage between two or more coverage points or variables. Cross coverage of a set of N coverage points is defined as the coverage of all combinations of all bins associated with the N coverage points, that is, the Cartesian product of

the N sets of coverage point bins. An autocross coverage is an automatically generated cross coverage of a coverage point with another coverage point.

[0066] Coverage of a coverage point item can be computed differently depending on whether the bins of the coverage point are explicitly defined by the user or automatically created by the tool.

[0067] For user-defined bins, the coverage of a coverage point is computed as the ratio of: i) the cardinality of the covered bins, i.e. the subset of all (defined) bins that are covered, and ii) the cardinality of the set of bins defined.

[0068] For automatically generated bins, the coverage of a coverage point is computed as the ratio of: i) the cardinality of the covered bin, i.e. the subset of all (auto-defined) bins that are covered, and ii) the lesser of a) the minimum number of bits needed to represent the coverage point and b) auto_bin_max.

[0069] Cumulative coverage considers the union of all significant bins; thus, it includes the contribution of all bins (including overlapping bins) of all instances.

[0070] To determine whether a particular bin of a coverage group is covered, the cumulative coverage computation considers the value of the at_least option of all instances being accumulated. Consequently, a bin is not considered covered unless its hit count equals or exceeds the maximum of all the at_least values of all instances. Use of the maximum represents the more conservative choice.

[0071] Figures 2-4 address an embodiment of EDA verification, in particular coverage driven verification.

[0072] This discussed embodiment does not require that all coverage logs be created before they can be merged. This approach scales particularly well in a random regression test environment where random tests are run 24 hours a day, 7 days a week, with numerous input conditions or random seeds. Because coverage logs are merged prior to the creation of the last coverage log, in this approach the amount of disk space required is no longer directly proportional to the number of tests. Similarly, one does not need to wait until the last simulation has finished running, to sample a merged coverage log. A pipeline is maintained of the unprocessed (unmerged) coverage logs. Early aggregation of coverage logs also allows assists the identification and analysis of coverage holes well before all the tests have finished running. "Bad" stimulus is uncovered sooner rather than later.

[0073] Figures 2-4 are discussed below.

[0074] Figure 2 is a simplified block diagram of a system performing simulation of a hardware verification language circuit design. The system includes individual simulation clients 210, a file system 220 to store coverage logs, a location repository server 230, and merge daemons 240. The simulation clients acting as multiple producers, and “urg” merge daemons act as multiple consumers and provide the synchronization primitives to maintain consistency of the queue.

[0075] Individual simulation clients 210 are the individual random simulations that generate functional coverage logs. Each test client dumps out a coverage log representing parts of the coverage model it targeted. Each client ensures the following:

[0076] 1) It does not overstep on a different simulation client’s database directory.

[0077] 2) It saves the log on a file system accessible from the machine where the “Urg merge daemons” are running for unified report generators.

[0078] 3) These clients talk to the location repository server using an interface, for example in Perl.

[0079] The clients inform the location repository server 230 of the directory where the coverage logs are located on the file system 220. An example implementation is an object oriented perl interface e.g. VtgSimClientIntf, with the following exemplary methods :

[0080] 1) new: Constructor to create the object. Takes two mandatory arguments and one optional argument.

[0081] i) server: name of the machine where the location repository server is running

[0082] ii) port: port number on which the location repository server is running

[0083] iii) newmodel: This argument is optional. The default value is 0. When set to 1, it switches the master model used for flexible merging at the merge server.

[0084] 2) dispatchLocation : Method to inform the location repository server regarding the location of the coverage log. Returns a small xml string on success and undef on failure.

[0085] Example 1:

[0086] use VtgSimClientIntf;

[0087] my \$simcli = VtgSimClientIntf->new(server=>'vgamd84', port=>7070);

[0088] my \$retVal = \$simcli->dispatchDbLocation("/u/foo/bar.vdb");

[0089] If successful the value of \$retVal would be:

[0090] <dbsummary curldb="/u/foo/bar.vdb" />

[0091] Example 2:

[0092] use VtgSimClientIntf;

[0093] my \$simcli = VtgSimClientIntf->new(server=>'vgamd84', port=>7070,
newmodel=>1);

[0094] my \$retVal = \$simcli ->dispatchDbLocation("/u/foo/bar.vdb");

[0095] This example switches the master coverage model used by mergeserver to the one
5 stored under /u/foo/bar.vdb . If successful the value of \$retVal would be:

[0096] <dbsummary curldb="/u/foo/bar.vdb" />

[0097] The file system 220 to hold coverage logs collectively represents enormous amounts
of mass storage, due to the demands of storing the coverage logs from the large number of
individual simulation clients 210.

10 [0098] Location repository server 230 is a lightweight server which keeps track of the
location of individual coverage log directories generated by individual simulations. It keeps track
of the top directory location. Some embodiments do not do data movement.

[0099] Merge daemons 240 are computation intensive processes that query the location
repository server 230 to fetch the list of coverage logs to process. The actual coverage logs are
15 then accessed from the file system 220 and processed. An aggregation tool such as urg(unified
report generator) merges coverage logs generated by individual simulations to compute which
parts of the coverage model are collectively covered by all tests in the verification suite. The
merge daemons provide hooks to sample the accumulated coverage logs before completion of all
coverage logs.

20 [00100] Figure 3 is a simplified state diagram showing the operation of the simulation system
of Figure 2.

[00101] The early state is start merge server 310. The merge server starts and is capable of
receiving coverage logs, and then fires the 'serv_post_start_call_back'. This callback is used for
submitting any coverage log that has been generated while the merge server was offline. If
25 multiple merge servers run on the same machine, option "daemon_port_start" specifies the port
of the merge server behind which the merge daemons run.

[00102] The intermediate state is active state 310. As simulations run, when a particular
simulation job finishes, it uses the simulation client interface to submit the coverage log location
to the merge server. At this state, checkpoint callbacks occur. After merging a predetermined
30 number of coverage logs (configurable by checkpoint_granularity in the config file) or at regular
intervals (configurable by time_granularity), it fires the 'checkpoint callback' for each merge
daemon. This callback deletes coverage logs which do not add any value or keep track of how
the coverage metrics are changing.

[00103] The later state is stop the merge server 330. At this point it fires daemon_post_call_back for each merge daemon. This performs clean up activities, sends email etc.

[00104] Details regarding merge server hooks/callbacks follow below, listed from A)-E):

5 [00105] A) Coverage log submission hook: This hook allows coverage logs to submitted to the merge server for aggregation.

[00106] B) Location server post start callback: ('serv_post_start_call_back' in the diagram) This callback gets fired once the merge server starts and is capable of receiving coverage logs. This callback should be used for submitting any coverage logs (using the coverage log submission hook) that have been generated while the merge server was offline.

10 [00107] C) Checkpoint callback: (checkpoint_callback in the diagram) This callback gets fired at regular intervals, configured via the following parameters:

[00108] 1) time_granularity: The user can set this parameter to a value (in seconds) to indicate when this callback should be fired. E.g if the value is 1800 sec, the callback will be fired after every 1800 sec.

[00109] 2) checkpoint_granularity: The user can set this parameter to indicate how often this callback should be fired. For e.g if the value is set to 900, then this callback will be fired after merging every 900 coverage logs.

[00110] In addition to the predetermined conditions above, other predetermined conditions that fire this callback and generate a coverage report are a coverage percentage or other coverage metric amount reached by the merged coverage log.

[00111] If both the parameters are set, then the callback will be fired whichever event happens first. The checkpoint_callback has the following info:

[00112] 1) Base score for each metric in the last callback

25 [00113] 2) For each coverage log merged, the checkpoint callback has the following info: 2a) name of the coverage log 2b) aggregated coverage score for each metric

[00114] This callback can be used to implement custom actions. Some of them are mentioned below:

[00115] 1) Incremental grading: If the incremental coverage score for a coverage log is zero, then it can be deleted. Incremental coverage added by a test is built into the hooks exposed by the merge server. So, if a coverage log does not add any value it can be deleted promptly, thereby bounding the required amount of mass storage by about the total number of coverable objects being targeted.

[00116] 2) Monitor how coverage score is changing with time: This callback can be used to plot how the coverage score is changing with time. If the coverage score does not change or changes very little, the client can attempt to change the stimulus or indicate to a real person (via email/page) that an action needs to be taken.

5 [00117] 3) Backing up of good tests: The client can determine if a specific coverage log indicates a change in coverage score or covers a difficult to hit/cover coverable object and back up that test.

[00118] Generally, checkpoint callback command is capable of running asynchronously with the merge daemon. The merge daemon will not wait for the script to finish. The checkpoint
10 callback interface is a command capable of processing an xml file. If the command is 'cpcb', it would be invoked as 'cpcb path-to-a-valid-xml-file' from the merge daemon after it creates a checkpoint.

[00119] An example schema of the checkpoint callback interface file follows:

[00120] <!ELEMENT checkpointed_databases (basescore) (dir*)>

15 [00121] <!ELEMENT basescore (score*) >

[00122] <!ATTLIST basescore EMPTY>

[00123] <!ELEMENT dir (test*) >

[00124] <!ATTLIST dir path CDATA #REQUIRED>

[00125] <!ATTLIST dir error (yes|no) #REQUIRED>

20 [00126] <!ATTLIST dir newmodel (yes|no) #OPTIONAL>

[00127] <!ELEMENT test (score*) >

[00128] <!ATTLIST test name CDATA #REQUIRED>

[00129] <!ELEMENT score >

[00130] <!ATTLIST score metric (assert|testbench|line|cond|fsm|branch) # REQUIRED >

25 [00131] <!ATTLIST score numerator CDATA #REQUIRED>

[00132] <!ATTLIST score denominator CDATA #REQUIRED>

[00133] The "basescore" element contains the coverage data score for the test onto which all the coverage log directories were incrementally merged.

[00134] The "dir" element represents the top level coverage log directory that was processed.
30 The attribute "path" indicates the location of the directory. The attribute "error" indicates if an error occurred while processing that directory. The error details are redirected to file "err.log".

[00135] "score" represents the coverage score for each test on a metric by metric basis.

[00136] An example checkpoint callback file follows:

[00137] <checkpointed_databases>
[00138] <basescore >
[00139] <score metric="assert" numerator="0" denominator="0" />
[00140] <score metric="testbench" numerator="0" denominator="0" />
5 [00141] </basescore>
[00142] <dir path="/remote/vtghome6/manojb/demo/pr/junk/db_0" , error="no">
[00143] <test name="/remote/vtghome6/manojb/demo/pr/junk/db_0/results" >
[00144] <score metric="assert" numerator="1" denominator="1" />
[00145] <score metric="testbench" numerator="0" denominator="0" />
10 [00146] </test>
[00147] <test name="/remote/vtghome6/manojb/demo/pr/junk/db_0/t3" >
[00148] <score metric="assert" numerator="1" denominator="1" />
[00149] <score metric="testbench" numerator="0" denominator="0" />
[00150] </test>
15 [00151] <test name="/remote/vtghome6/manojb/demo/pr/junk/db_0/test" >
[00152] <score metric="assert" numerator="1" denominator="1" />
[00153] <score metric="testbench" numerator="47" denominator="50997" />
[00154] </test>
[00155] </dir>
20 [00156] <dir path="/remote/vtghome6/manojb/demo/pr/junk/db_1" , error="no"
newmodel="yes">
[00157] <test name="/remote/vtghome6/manojb/demo/pr/junk/db_1/results" >
[00158] <score metric="assert" numerator="1" denominator="1" />
[00159] <score metric="testbench" numerator="47" denominator="50997" />
25 [00160] </test>
[00161] <test name="/remote/vtghome6/manojb/demo/pr/junk/db_1/t3" >
[00162] <score metric="assert" numerator="1" denominator="1" />
[00163] <score metric="testbench" numerator="47" denominator="50997" />
[00164] </test>
30 [00165] <test name="/remote/vtghome6/manojb/demo/pr/junk/db_1/test" >
[00166] <score metric="assert" numerator="1" denominator="1" />
[00167] <score metric="testbench" numerator="47" denominator="50997" />
[00168] </test>

[00169] </dir>

[00170] </checkpointed_databases>

[00171] D) Daemon pre start call back: (daemon_pre_call_back in the diagram). This callback is fired once before the merge daemon start processing coverage logs. This callback can
5 be used to collect statistics for each merge daemon.

[00172] E) Daemon post start call back: (daemon_post_call_back in the diagram) This callback gets fired when the merge server is stopped. One such callback is fired for each merge daemon. This can be used to take actions like the following:

[00173] 1. Generate custom reports

10 [00174] 2. Send emails etc.

[00175] Figure 4 is a simplified directory structure of the simulation system of Figures 2 and 3.

[00176] By default each merge daemon creates a sub directory inside the directory specified by directory_prefix. For example, for the default configuration file specified by

15 "/aggregated_data", the directory structure of Figure 4 is created.

"UrgDaemon_1/UserDir/MergeDb" is the aggregated coverage log created by the first merge daemon. "UrgDaemon_1/UserDir/err.log" is the error log file for the coverage log aggregation done by the first daemon. Errors (corrupt coverage logs, bugs/crashes in the tool) are redirected to this file.

20 [00177] The following is an example syntax of launching a merge server on location repository server 230.

[00178] mergeserv [--config=<config_file>] --start|--stop

[00179] config_file is the [optional] configuration file read by the merge server. An example format of this configuration file is provided below.

25 [00180] --start: Starts the merge server. If the merge server is already running, this option does not do anything.

[00181] --stop: Stops the merge server. This initiates the stop sequence. The merge server exits after all the submitted databases have been merged.

[00182] An example merge server configuration file follows, in xml format.

30 [00183] <mserve_config>

[00184] <location_serv port="7070" datadir="/u/regress_usr/locationdb" />

[00185] <serv_post_start_call_back command="submit_again" />

[00186] <merge_daemon num="2" directory_prefix="/aggregated_data">

[00187] <daemon_pre_call_back command="pre.pl"/>

[00188] <merge_daemon_options checkpoint_granularity="300"
time_granularity="900" />

[00189] <check_point_call_back command="foo.pl" />

5 [00190] <daemon_post_call_back command="post.pl"/>

[00191] </merge_daemon>

[00192] </mserv_config>

[00193] In this example default configuration:

[00194] 1) Location repository server runs on port 7070.

10 [00195] 2) A single instance of merge daemon is fired.

[00196] 3) All merge daemon clients create the aggregated data under directory
"aggregated_data".

[00197] 4) Merge server maintains the list of unprocessed coverage logs in a log that is stored
under the directory specified by the attribute datadir. In many embodiments, only the user
15 running the merge server has write permissions to this directory. The default value of datadir is
".".

[00198] Running multiple instances of the merge daemon is a performance optimization to
handle the case in which coverage logs are generated quickly. To run multiple merge servers on
the same machine, the config file used for firing the individual merge servers should not conflict
20 with each other. Specifically the following entries in the config file should not conflict.

[00199] i) location_serv port

[00200] ii) datadir

[00201] iii) directory_prefix

[00202] iv) daemon_port_start

25 [00203] Also, any callback script used with the merge servers should not modify the same
system resources. For example, if the checkpoint callback saves useful test information in some
file/directory, then the file/directory used by the individual callback scripts should be different
for each merge server.

[00204] Here is an example of sample config files for running two merge servers on the same
30 machine.

[00205] <mserv_config>

[00206] <location_serv port="7070" datadir="/u/regress_usr/locationdb" />

[00207] <serv_post_start_call_back command="submit_again" />

[00208] <merge_daemon num="2" directory_prefix="/aggregated_data">
 [00209] <daemon_pre_call_back command="pre.pl"/>
 [00210] <merge_daemon_options checkpoint_granularity="300"
 time_granularity="900" />
 5 [00211] <check_point_call_back command="foo.pl" />
 [00212] <daemon_post_call_back command="post.pl"/>
 [00213] </merge_daemon>
 [00214] </mserv_config>
 [00215] <mserv_config>
 10 [00216] <location_serv port="7071" datadir="/u/regress_usr/locationdb.second" />
 [00217] <serv_post_start_call_back command="submit_again" />
 [00218] <merge_daemon num="2" directory_prefix="/aggregated_data.second">
 [00219] <daemon_pre_call_back command="pre.pl"/>
 [00220] <merge_daemon_options checkpoint_granularity="300"
 15 time_granularity="900" daemon_port_start="14449" />
 [00221] <check_point_call_back command="foo.pl" />
 [00222] <daemon_post_call_back command="post.pl"/>
 [00223] </merge_daemon>
 [00224] </mserv_config>
 20 [00225] Whereas the preceding text is generally directed to merging coverage logs, the
 following text is generally directed to techniques of merging particular coverage logs.
 [00226] Figures 5A and 5B are simplified flow charts contrasting an improved simulation
 flow that merges, rather than discards, prior coverage logs. Due to the trend of decreasing cost
 for increased computation resources, each time the simulation model is adjusted, one of ordinary
 25 skill in the art would tend to re-run simulations to reach a satisfactory coverage metric(s) result,
 as shown in Figure 5A, rather than preserving any coverage data from coverage logs of old
 simulation models.
 [00227] The process flow of Figure 5A is as follows. In 501, a functional coverage model is
 created. In 503, random simulations are run, starting from some input condition or some first
 30 random seed. In 505, coverage logs generated by simulations are merged. In 507, the merged
 coverage log is analyzed. In 509, the functional coverage model is refined. The process flow
 repeats, after 611, in which the previous coverage logs are deleted.

[00228] The process flow of Figure 5B modifies the process flow of Figure 5A. In 504, rather than starting the random simulations from some input condition or some first random seed as in 503, the random simulations are run with new input conditions or new random seeds. In 512, rather than deleting the coverage logs as in 511, the previous coverage logs are imported.

5 [00229] The technology of Figure 5B can be applied to either the case where coverage logs are merged when all simulations are completed, or prior to completion of the last simulation as discussed above in connection with Figures 2-4.

[00230] The following examples use SystemVerilog style syntax to represent coverage logs, but are applicable to any hardware verification language with comparable syntax. Log T1 represents an older coverage log. Log T2 represents a newer coverage log. The Merged Log represents the result of merging the respective coverage logs. Each example is followed by explanation.

[00231] Example 1:

Log T1	Log T2	Merged Log
cp1: coverpoint firstsig; option.auto_bin_max = 64;	cp1: coverpoint firstsig; option.auto_bin_max = 32;	cp1: coverpoint firstsig; option.auto_bin_max = 32;

15 [00232] According to Example 1, differing auto_bin_max values are not merge equivalent.

[00233] Example 2:

Log T1	Log T2	Merged Log
cp2: coverpoint secondsig { bin first = [0:63]; bin mid = [71:82]; }	cp2: coverpoint secondsig { bin first = [0:63]; bin second = [65:128]; }	cp2: coverpoint secondsig { bin first = [0:63]; bin second = [65:128]; }

[00234] According to Example 2, "bin mid = [71:82]" of "coverpoint secondsig" of Log T1 is deleted because it is missing in the definition of "coverpoint secondsig" in Log T2.

20 [00235] Example 3:

Log T1	Log T2	Merged Log
cp3: coverpoint thirdsig;	cp3: coverpoint thirdsig;	cp3: coverpoint thirdsig;

[00236] According to Example 3, cp3 of log T1 is merge equivalent to cp3 of log T2 and they are merged.

[00237] Example 4:

Log T1	Log T2	Merged Log
Bit[7:0]signal; cp4: coverpoint signal;	Bit[15:0]signal; cp4: coverpoint signal;	cp4: coverpoint signal;

25

[00238] According to Example 4, “cp4: coverpoint signal” of Log T1 is deleted because the width of the coverpoint expression (‘signal’) has changed from 8 bits to 16 bits.

[00239] Example 5:

Log T1	Log T2	Merged Log
cc1: cross cp1, cp2;	cc1: cross cp1, cp2;	cc1: cross cp1, cp2;

- 5 **[00240]** According to Example 5, “cc1: cross cp1, cp2;” of Log T1 is deleted due to coverpoint “cp1” of Log T1 being deleted (refer to example 1 for cp1)

[00241] Example 6:

Log T1	Log T2	Merged Log
cc2: cross cp2, cp3 { bins mybin = binsof(cp2) intersect [0:255]; }	cc2: cross cp2, cp3 { bins mybin = binsof(cp2) intersect [0:255]; bins yourbin = binsof(cp2) intersect [256:511]; }	cc2: cross cp2, cp3 { bins mybin = binsof(cp2) intersect [0:255]; bins yourbin = binsof(cp2) intersect [256:511]; }

- 10 **[00242]** According to Example 6, the coverage data for user bins mybin is merged for T1 and T2. The auto crosses of cc2 in T1 which are also in the auto cross coverage space of cc2 in T2 are merged.

[00243] Example 7:

Log T1	Log T2	Merged Log
cc3: cross cp2, cp3 { bins mybin = binsof(cp2) intersect [0:255]; }	cc3: cross cp2, cp3 { bins mybin = binsof(cp2) intersect [0:8191]; }	cc3: cross cp2, cp3 { bins mybin = binsof(cp2) intersect [0:255]; bins mybin = binsof(cp2) intersect [0:8191]; }

- 15 **[00244]** According to Example 7, autocrosses seen in log T1 will still be imported and merged with the ones seen in log T2, despite the fact that autocross coverage space of cc3 is different in T1 and T2 due to differing definitions of the user bin ‘mybin’. Some of the autocross bins of cc3 in T1 are subsumed in the definition of user bin mybin in log T2. These autocross bins from log T1 are merged into user bin mybin of log T2. Remaining autocross bins in log T1 will be merged with corresponding autocross bins in log T2.

- 20 **[00245]** Example 8:

Log T1	Log T2	Merged Log
--------	--------	------------

	cc4: cross cp1, cp2, cp3	cc4: cross cp1, cp2, cp3
--	-----------------------------	-----------------------------

[00246] According to Example 8, cross point cc4 absent from log T1 but present in T2 is retained after merging.

[00247] Example 9:

Log T1	Log T2	Merged Log
first_prop:cover property(.....); second_prop:cover property(.....);	first_prop:cover property(.....); third_prop:cover property(.....);	first_prop:cover property(.....); second_prop:cover property(.....); third_prop:cover property(.....);

5

[00248] According to Example 7, cover property 'third_prop' absent from log T1 but present in T2 is retained after merging; and cover property 'second_prop' absent from log T2 but present in T1 is retained after merging.

[00249] Figure 6 is a simplified block diagram of a computer system of various embodiments.

10 **[00250]** Fig. 6 is a simplified block diagram of a computer system 610 that can be used to implement software incorporating aspects of the present invention. While the flow charts and other algorithms set forth herein describe series of steps, it will be appreciated that each step of the flow chart or algorithm can be implemented by causing a computer system such as 610 to operate in the specified manner.

15 **[00251]** Computer system 610 typically includes a processor subsystem 614 which communicates with a number of peripheral devices via bus subsystem 612. Processor subsystem 614 may contain one or a number of processors. The processor subsystem 614 provides a path for the computer system 610 to receive and send information described herein, including within the processor subsystem 614, such as with a multi-core, multiprocessor, and/or virtual machine
20 implementation. The peripheral devices may include a storage subsystem 624, comprising a memory subsystem 626 and a file storage subsystem 628, user interface input devices 622, user interface output devices 620, and a network interface subsystem 616. The input and output devices allow user interaction with computer system 610. Network interface subsystem 616 provides an interface to outside networks, including an interface to communication network 618,
25 and is coupled via communication network 618 to corresponding interface devices in other computer systems. Communication network 618 may comprise many interconnected computer systems and communication links. These communication links may be wireline links, optical links, wireless links, or any other mechanisms for communication of information. While in one

embodiment, communication network 618 is the Internet, in other embodiments, communication network 618 may be any suitable computer network. The communication network 618 provides a path for the computer system 610 to receive and send information described herein.

[00252] The physical hardware component of network interfaces are sometimes referred to as network interface cards (NICs), although they need not be in the form of cards: for instance they could be in the form of integrated circuits (ICs) and connectors fitted directly onto a motherboard, or in the form of macrocells fabricated on a single integrated circuit chip with other components of the computer system.

[00253] User interface input devices 622 may include a keyboard, pointing devices such as a mouse, trackball, touchpad, or graphics tablet, a scanner, a touch screen incorporated into the display, audio input devices such as voice recognition systems, microphones, and other types of input devices. In general, use of the term "input device" is intended to include all possible types of devices and ways to input information into computer system 610 or onto computer network 618.

[00254] User interface output devices 620 may include a display subsystem, a printer, a fax machine, or non visual displays such as audio output devices. The display subsystem may include a cathode ray tube (CRT), a flat panel device such as a liquid crystal display (LCD), a projection device, or some other mechanism for creating a visible image. The display subsystem may also provide non visual display such as via audio output devices. In general, use of the term "output device" is intended to include all possible types of devices and ways to output information from computer system 610 to the user or to another machine or computer system.

[00255] Storage subsystem 624 stores the basic programming and data constructs that provide the functionality of certain embodiments of the present invention. For example, the various modules implementing the functionality of certain embodiments of the invention may be stored in storage subsystem 624. These software modules are generally executed by processor subsystem 614.

[00256] Memory subsystem 626 typically includes a number of memories including a main random access memory (RAM) 630 for storage of instructions and data during program execution and a read only memory (ROM) 632 in which fixed instructions are stored. File storage subsystem 628 provides persistent storage for program and data files, and may include a hard disk drive, a floppy disk drive along with associated removable media (illustratively shown as computer readable medium 640 storing circuit design 680), a CD/DVD ROM drive, an optical drive, or removable media cartridges. The databases and modules implementing the functionality

of certain embodiments of the invention may have been provided on a computer readable medium such as one or more CD/DVD-ROMs, and may be stored by file storage subsystem 628. The host memory 626 contains, among other things, computer instructions which, when executed by the processor subsystem 614, cause the computer system to operate or perform functions as described herein. As used herein, processes and software that are said to run in or on "the host" or "the computer", execute on the processor subsystem 614 in response to computer instructions and data in the host memory subsystem 626 including any other local or remote storage for such instructions and data.

[00257] Bus subsystem 612 provides a mechanism for letting the various components and subsystems of computer system 610 communicate with each other as intended. Although bus subsystem 612 is shown schematically as a single bus, alternative embodiments of the bus subsystem may use multiple busses.

[00258] Computer system 610 itself can be of varying types including a personal computer, a portable computer, a workstation, a computer terminal, a network computer, a television, a mainframe, a parallel processing system, a network of more than one computer, or any other data processing system or user device. Due to the ever changing nature of computers and networks, the description of computer system 610 depicted in Fig. 6 is intended only as a specific example for purposes of illustrating the preferred embodiments of the present invention. Many other configurations of computer system 610 are possible having more or less components than the computer system depicted in Fig. 6.

[00259] The following is a mathematical appendix discussing exemplary conditions of merging coverage logs.

[00260] **1. Merge equivalence**

[00261] Coverable objects across logs are merged if they are merge equivalent. The notion of merge equivalence is different for coverpoints and crosspoints. We use the notation $A1 \equiv_{\text{merge-eq}} A2$ to indicate that A1 and A2 are merge equivalent.

[00262] **1.1 Merge equivalence of coverpoints**

[00263] **1.1.1 Merge equivalence of autobinned coverpoints**

[00264] **1.1.1.1. Default merge equivalence**

[00265] $P1 \equiv_{\text{merge-eq}} P2$ iff

[00266] $\text{name}(P1) = \text{name}(P2)$

[00267] $\text{auto_bin_max}(P1) = \text{auto_bin_max}(P2)$

[00268] $\text{width}(P1) = \text{width}(P2)$

[00269] where P1 and P2 are auto-binned coverpoints.

[00270] **1.1.1.2 Flexible merge equivalence**

[00271] This is same as default merge equivalence for autobinned coverpoints.

5 [00272] **1.1.2 Merge Equivalence of userbinned coverpoints**

[00273] **1.1.2.1 Default merge equivalence**

[00274] $P1 \equiv_{\text{merge-eq}} P2$ iff

[00275] $\text{name}(P1) \equiv \text{name}(P2)$

[00276] $\text{width}(P1) = \text{width}(P2)$

10 [00277] $B_u(P1) = B_u(P2)$

[00278] where P1 and P2 are user-binned coverpoints and $B_u(P)$ indicates the set of user defined bins.

[00279] **1.1.2.2 Flexible merge equivalence**

[00280] $P1 \equiv_{\text{merge-eq}} P2$ iff

15 [00281] $\text{name}(P1) \equiv \text{name}(P2)$

[00282] $\text{width}(P1) = \text{width}(P2)$

[00283] Note that the requirement to have same set of user-defined bins is dropped.

[00284] **1.2 Merge Equivalence of crosspoints**

[00285] Consider a crosspoint,

20 [00286] $C = P1 \times P2 \dots \times Pn$

[00287] We define the following:

[00288] B_u = Set of user defined bins in cross C

[00289] For bin $b \in B_u$, we define the following:

[00290] S_b = Auto cross bins comprising the user defined bin b

25 [00291] B_a = Set of auto cross bins = $\prod B_j - S_{B_u}$

[00292] B_{ign} = Set of illegal/ignore/unreachable auto cross bins

[00293] \prod = Set cross product operator

[00294] B_j = Set of bins in jth coverpoint, P_j .

[00295] The notation $B_u(C)$ is used to indicate the set of user defined bins for cross C.

30 [00296] Similar notation is extended for other sets like B_a etc.

[00297] Consider two cross points:

[00298] $C1 = A1 \times A2 \dots \times A_n$, where $A1, A2, \dots, A_n$ are component coverpoints.

[00299] $C2 = B1 \times B2 \dots \times B_m$, where $B1, B2, \dots, B_m$ are component coverpoints

[00300] **1.2.1 Default merge equivalence**

[00301] $C1 \equiv_{\text{merge-eq}} C2$ if the following holds.

5 [00302] - Both crosses have the same number of coverpoints.

[00303] $n == m$

[00304] - Component coverpoints are pair wise merge equivalent.

[00305] $A1 \equiv_{\text{merge-eq}} B1, A2 \equiv_{\text{merge-eq}} B2 \dots, A_n \equiv_{\text{merge-eq}} B_m$

[00306] - $B_u(C1) = B_u(C2)$

10 [00307] - If D is the set of user defined bins in $C1$ and $C2$, then the following holds

[00308] $\forall d \in D \ S_d(C(T1)) = S_d(C(T2))$

[00309] **1.2.2 Flexible merge equivalence**

[00310] $C1 \equiv_{\text{merge-eq}} C2$ if the following holds.

[00311] - Both crosses have the same number of coverpoints.

15 [00312] $n == m$

[00313] - Component coverpoints are pair wise merge equivalent.

[00314] $A1 \equiv_{\text{merge-eq}} B1, A2 \equiv_{\text{merge-eq}} B2 \dots, A_n \equiv_{\text{merge-eq}} B_m$

[00315] The above rules have the following implication:

[00316] - If a coverpoint cannot be merged, any crosspoint involving that coverpoint will

20 simply be added to the merged log. In short $\neq_{\text{merge-eq}}$ of coverpoints has a ripple effect to all the crosspoints.

[00317] **1.2 Merge Equivalence of assertions**

[00318] **1.2.1 Default merge equivalence**

[00319] For assertions $A1$ and $A2$, $A1 \equiv_{\text{merge-eq}} A2$ if the following holds.

25 [00320] $\text{name}(A1) = \text{name}(A2)$

[00321] **1.2.2 Flexible merge equivalence**

[00322] These rules are same as default merge equivalence.

[00323] **2. Rules for merging**

[00324] The rules for default merging are straightforward since default merge equivalence
30 ensures that the coverage space is the same across logs being merged.

[00325] The following sections describe the rules for flexible merging.

[00326] **2.1 Rules for merging coverpoints and crosspoints**

[00327] Let $M(T1)$ be a coverpoint/crosspoint in $T1$ and $M(T2)$ be a coverpoint/crosspoint in $T2$.

5 [00328] 1. If $M(T1) \equiv_{\text{merge-eq}} M(T2)$, after merging according to the rules defined above, bins defined only in the older functional coverage model will be dropped.

[00329] 2. If $M(T1) \not\equiv_{\text{merge-eq}} M(T2)$ then the merged log contains $M(T2)$, where $T2$ is the log with a more recent functional coverage model.

[00330] **2.2 Rules for merging assertions**

10 [00331] If $A(T1)$ is an assertion in log $T1$ and $A(T2)$ is an assertion in log $T2$, then they are merged according to the following rules.

[00332] 1. If $A(T1) \equiv_{\text{merge-eq}} A(T2)$ then the merged log contains $A(T1)$ (or $A(T2)$).

[00333] 2. If $A(T1) \not\equiv_{\text{merge-eq}} A(T2)$ then the merged log contains $A(T2)$, where $T2$ is the more recent log.

15 [00334] While the present invention is disclosed by reference to the preferred embodiments and examples detailed above, it is to be understood that these examples are intended in an illustrative rather than in a limiting sense. It is contemplated that modifications and combinations will readily occur to those skilled in the art, which modifications and combinations will be within the spirit of the invention and the scope of the following claims. What is claimed
20 is:

CLAIMS

- 1 1. An electronic design automation method, comprising:
2 merging coverage logs generated by verification of a hardware description language
3 circuit design, as the coverage logs are generated, without waiting for all coverage logs of
4 pending verification.
- 1 2. The electronic design automation method of claim 1, wherein at least one coverage log of
2 said coverage logs is generated by simulation of the hardware description language circuit
3 design.
- 1 3. The electronic design automation method of claim 1, further comprising:
2 responsive to said merging, altering a condition of a pending simulation of the hardware
3 description language circuit design.
- 1 4. The electronic design automation method of claim 1, further comprising:
2 responsive to said merging, altering an input parameter of a pending simulation of the
3 hardware description language circuit design.
- 1 5. The electronic design automation method of claim 1, further comprising:
2 responsive to said merging, freeing at least part of mass storage.
- 1 6. The electronic design automation method of claim 1, further comprising:
2 responsive to said merging and to a predetermined condition being met, generating a
3 coverage report.
- 1 7. The electronic design automation method of claim 1, further comprising:
2 responsive to said merging, determining that completing a pending verification is
3 expected to insufficiently improve verification coverage of the hardware description language
4 circuit design.
- 1 8. The electronic design automation method of claim 1, further comprising:
2 responsive to said merging, determining that completing a pending verification is
3 expected to insufficiently improve verification coverage of the hardware description language

4 circuit design, by determining that an attribute of the hardware description language circuit
5 design to be simulated by a pending simulation has been simulated already.

1 9. The electronic design automation method of claim 1, further comprising:
2 responsive to said merging, determining that completing a pending verification is
3 expected to insufficiently improve verification coverage of the hardware description language
4 circuit design, and causing the pending verification to stop.

1 10. The electronic design automation method of claim 1, wherein said merging results in a
2 merged coverage log including formal verification coverage data.

1 11. The electronic design automation method of claim 1, further comprising:
2 responsive to readiness to perform said merging, requesting coverage log data that was
3 generated prior to said readiness to perform said merging.

1 12. The electronic design automation method of claim 1, further comprising:
2 responsive to said merging, performing:
3 updating a coverage metric of verification coverage of the hardware description
4 language circuit design, said coverage metric inclusive of coverage data from said
5 merging.

1 13. The electronic design automation method of claim 1, further comprising:
2 creating multiple running instances, the multiple running instances performing merging
3 of coverage logs of simulations of the hardware description language circuit design.

1 14. A computer readable medium with computer readable instructions for circuit design,
2 comprising:
3 computer instructions merging coverage logs generated by verification of a hardware
4 description language circuit design, as the coverage logs are generated, without waiting for all
5 pending coverage logs.

1 15. The computer readable medium of claim 14, further comprising:

2 computer instructions altering a condition of a pending simulation of the hardware
3 description language circuit design, responsive to said merging.

1 16. The computer readable medium of claim 14, further comprising:
2 computer instructions freeing at least part of mass storage, responsive to said merging.

1 17. The computer readable medium of claim 14, further comprising:
2 computer instructions generating a coverage report, responsive to said merging and to a
3 predetermined condition being met.

1 18. An apparatus, comprising:
2 a computer with a memory storing computer instructions merging coverage logs
3 generated by verification of a hardware description language circuit design, as the coverage logs
4 are generated, without waiting for all pending coverage logs.

1 19. The apparatus of claim 18, wherein the memory further stores:
2 computer instructions altering a condition of a pending simulation of the hardware
3 description language circuit design, responsive to said merging.

1 20. The apparatus of claim 18, wherein the memory further stores:
2 computer instructions freeing at least part of mass storage, responsive to said merging.

1 21. The apparatus of claim 18, wherein the memory further stores:
2 computer instructions generating a coverage report, responsive to said merging and to a
3 predetermined condition being met.

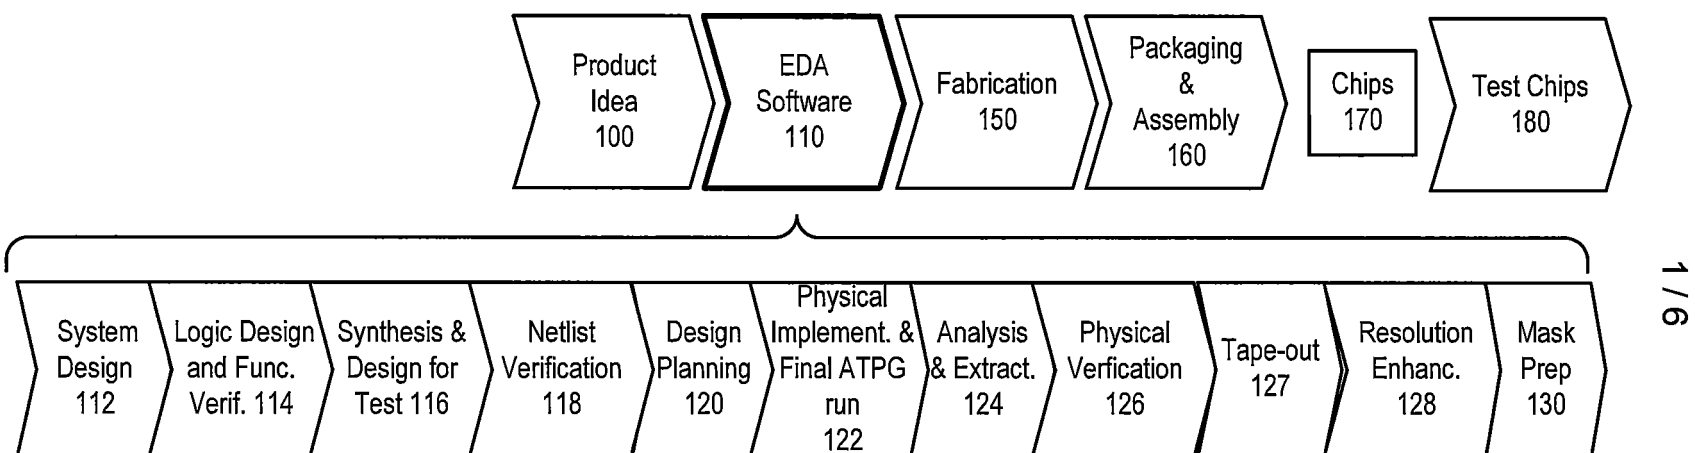


FIG. 1

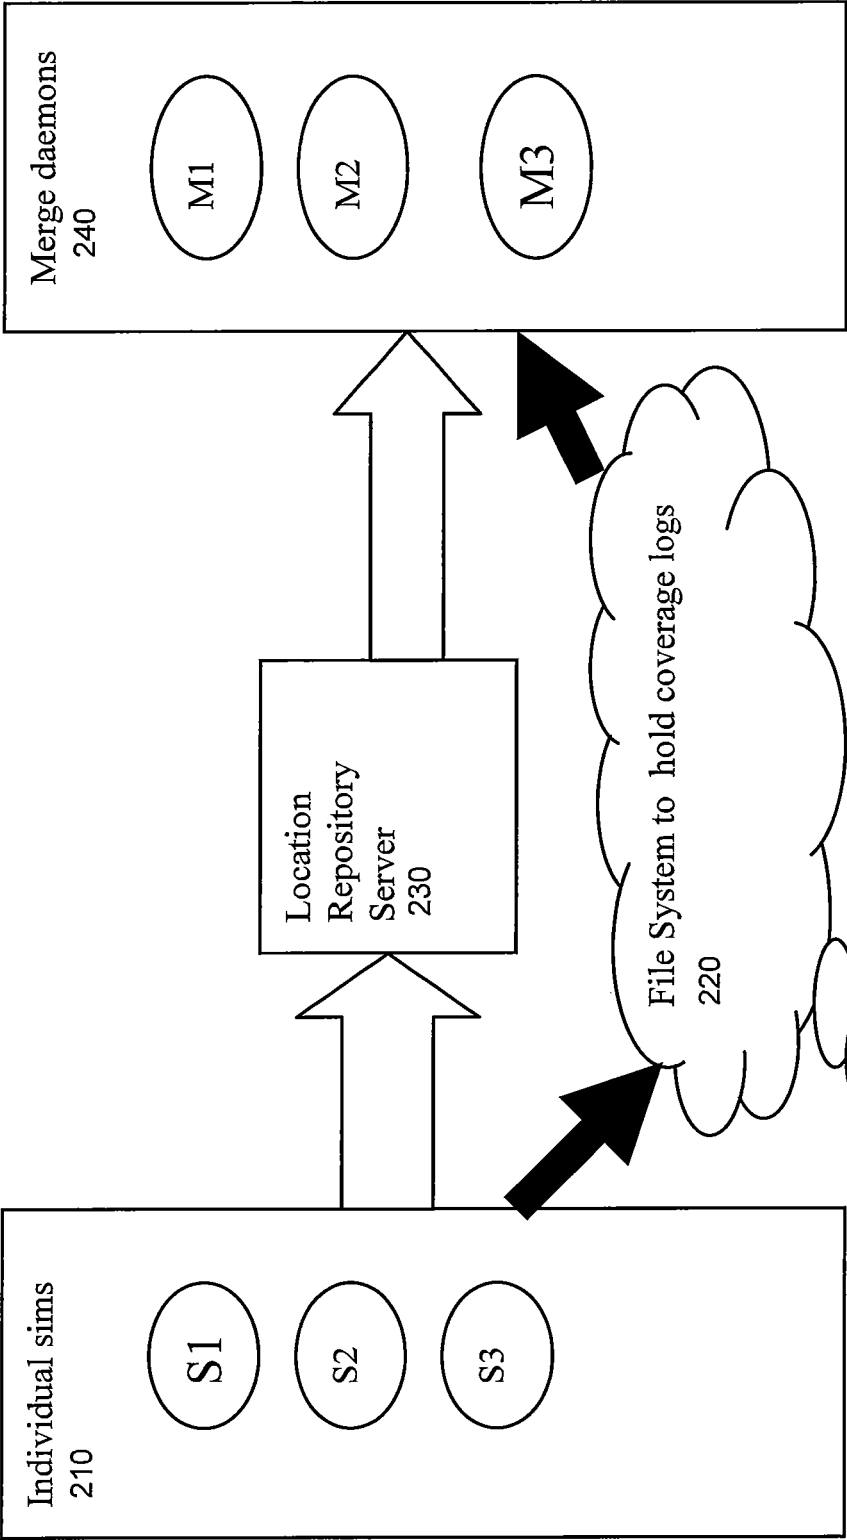


FIG. 2

3 / 6

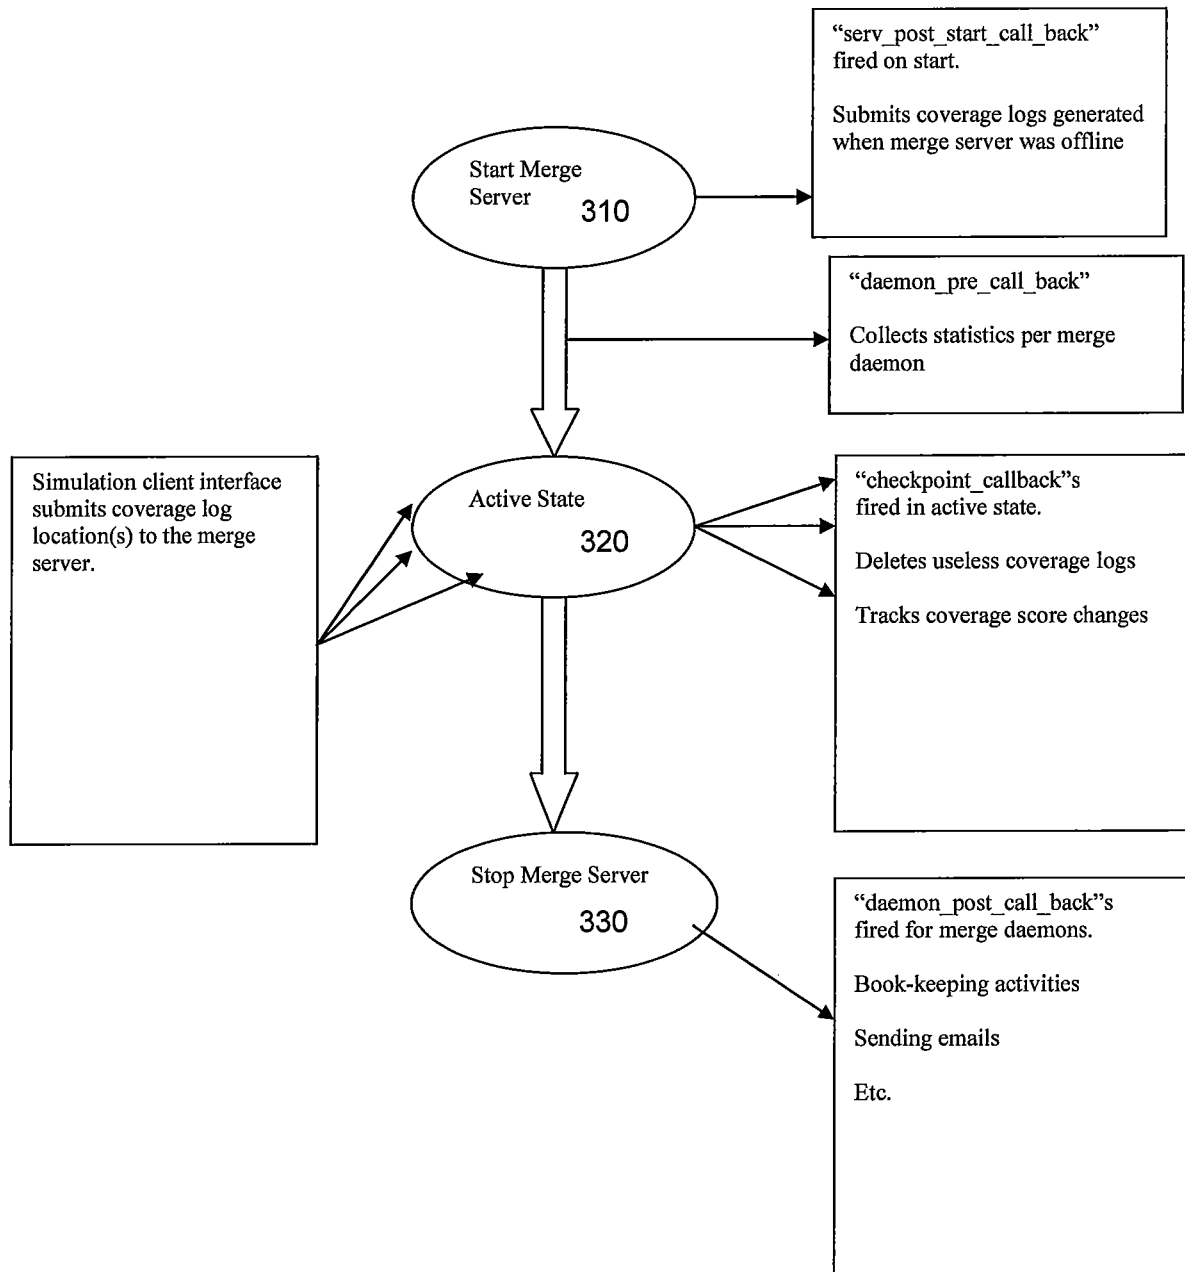


FIG. 3

4 / 6

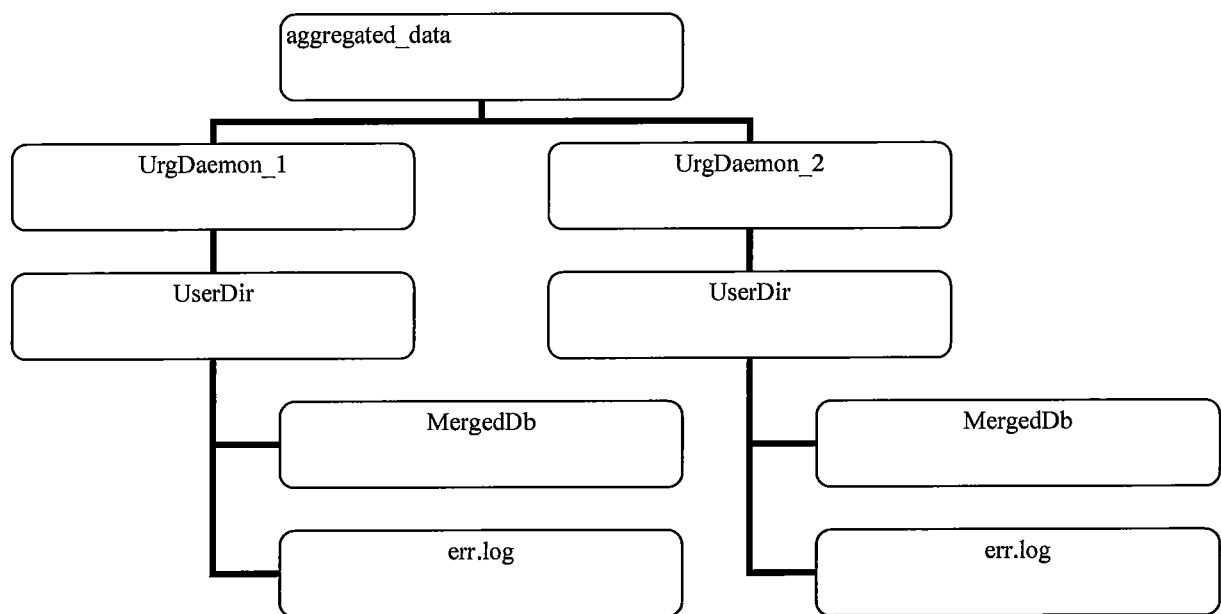
400

FIG. 4

5 / 6

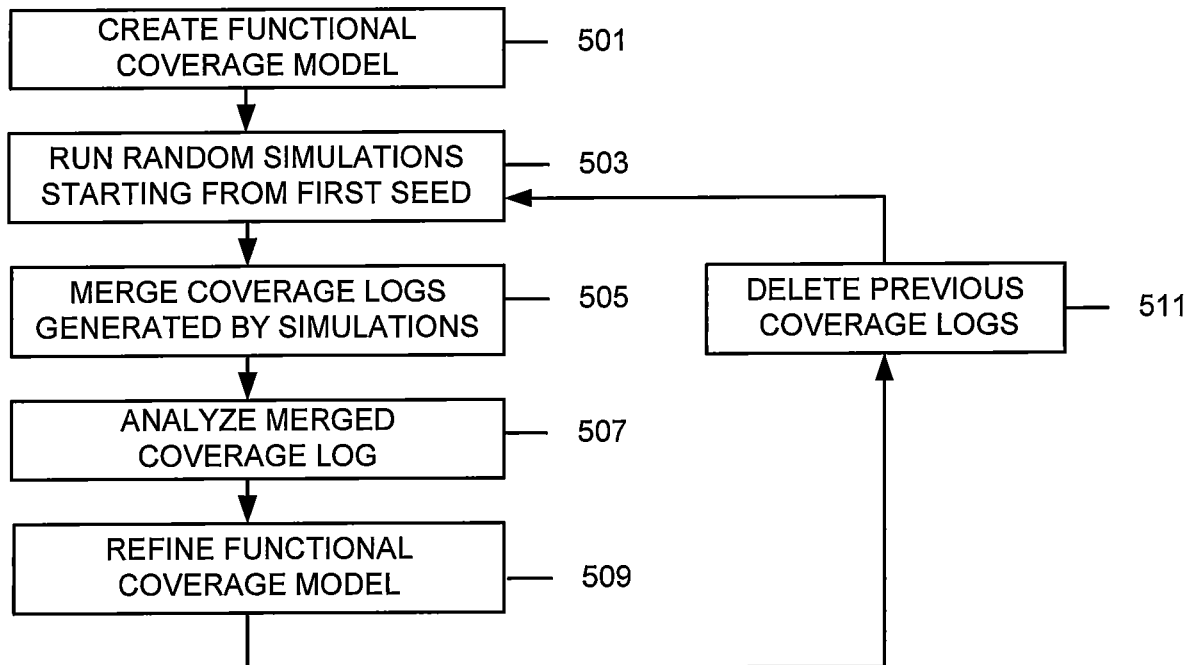


FIG. 5A

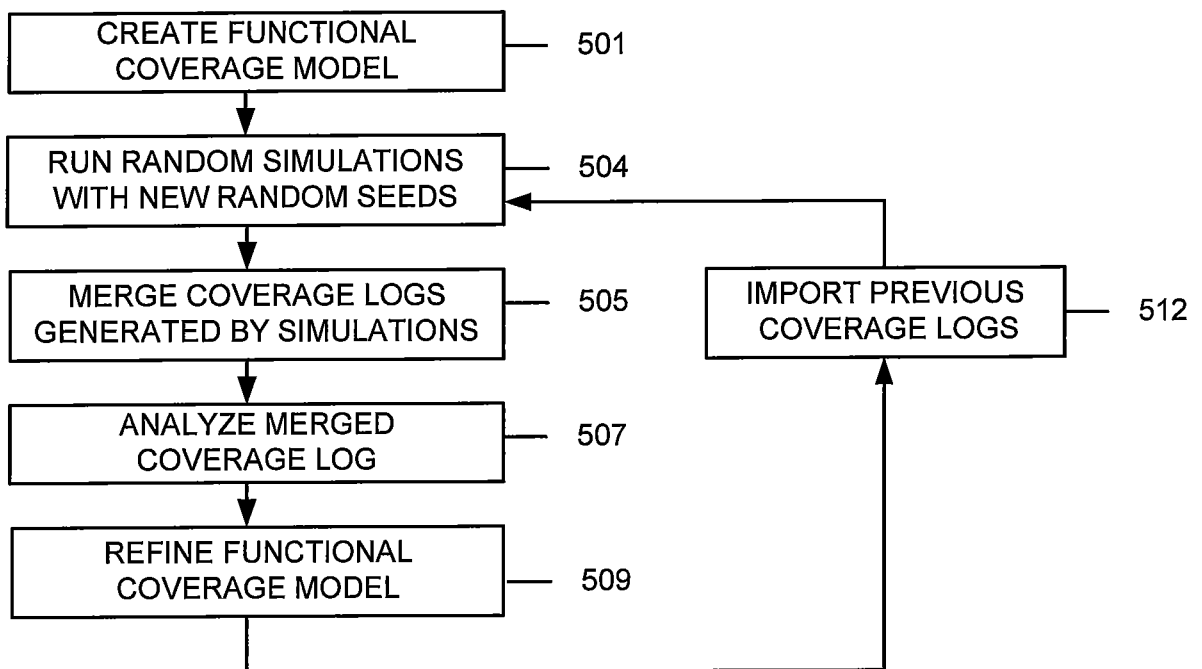


FIG. 5B

6 / 6

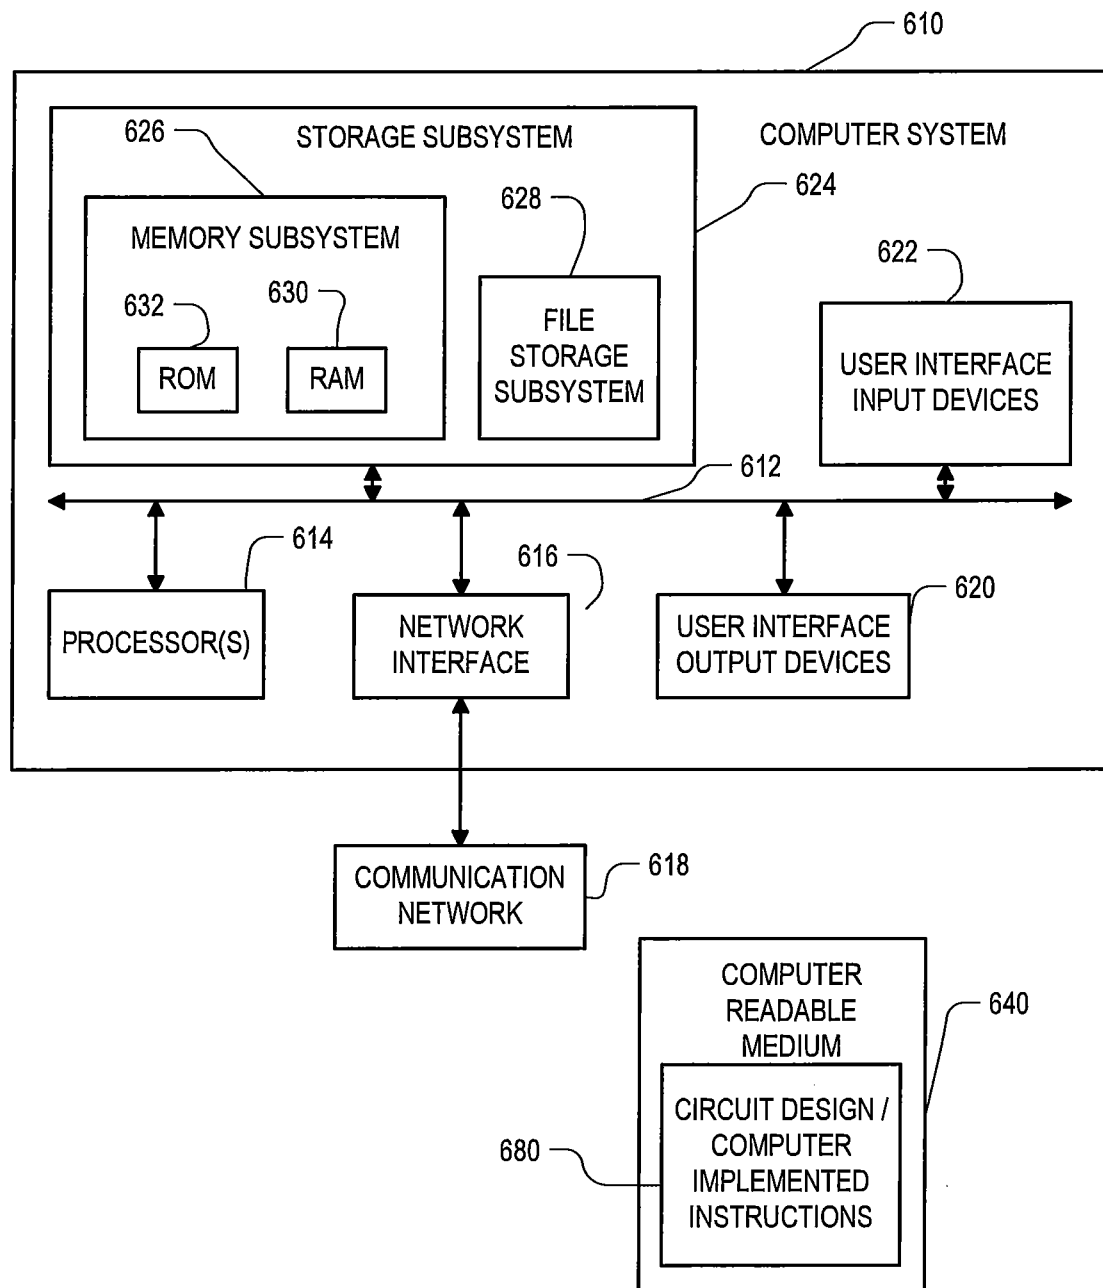


FIG. 6