Jeff S. Pitzer, OSB No. 020846
Email: jpitzer@pitzerlaw.net
Peter M. Grabiel, OSB No. 171964
Email: pgrabiel@pitzerlaw.net
PITZER LAW
210 SW Morrison St., Suite 600
Portland, OR 97204
Telephone: 503-227-1477

Paul Richter (to be admitted *pro hac vice*)
DEVLIN LAW FIRM LLC
1526 Gilpin Avenue
Wilmington, Delaware 19806
Telephone: (302) 449–9010

David Sochia (to be admitted *pro hac vice*)
dsochia@McKoolSmith.com
Ashley N. Moore (to be admitted *pro hac vice*)
amoore@McKoolSmith.com
Richard A. Kamprath (to be admitted *pro hac vice*)
rkamprath@McKoolSmith.com
MCKOOL SMITH, P.C.
300 Crescent Court Suite 1500
Dallas, TX 75201
Telephone: (214) 978-4000

Attorneys for Plaintiff

IN THE UNITED STATES DISTRICT COURT

FOR THE DISTRICT OF OREGON

PORTLAND DIVISION

BELL SEMICONDUCTOR, LLC Plaintiff, v. AMPERE COMPUTING, LLC Defendant.	Case No. 3:22-cv-1903 **COMPLAINT FOR PATENT INFRINGEMENT** JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Bell Semiconductor, LLC ("Bell Semic" or "Plaintiff") brings this Complaint against Defendant Ampere Computing, LLC ("Ampere") for infringement of U.S. Patent No. 7,231,626 ("the '626 patent") and U.S. Patent No. 7,396,760 ("the '760 patent"). Plaintiff, on personal knowledge of its own acts, and on information and belief as to all others based on investigation, alleges as follows:

SUMMARY OF THE ACTION

1. This is a patent infringement suit relating to Ampere's unauthorized and unlicensed use of the '626 patent and '760 patent. The circuit design methodologies claimed in the '626 patent and '760 patent are used by Ampere in the production of one or more of its semiconductor chips, including its AC7-M128-30 Altra Max, which contains multiple ARM cores ("Ampere Accused Product").

2. Traditionally, the process flow for IC design is highly linear, with each phase of the design process depending on the previous steps. Accordingly, when revisions to portions of the physical design are made, as typically happens numerous times during the design process, all the subsequent steps typically need to be redone in their entirety for at least the layer, if not the entire device. This is because regardless of the size or extent of the revision to the physical design, the changes must be merged into a much larger integrated circuit design and then the remaining steps of the design process flow re-run.

3. Before the inventions claimed in the '626 patent, the typical turnaround time for implementing a change to the physical design for cutting edge devices was approximately one week regardless of the size of the change. This is extremely inefficient in most instances where the change relates to only a small fraction of the overall design. *See* Ex. A at 3:16–18 & Fig. 1.

4. The ’626 patent’s inventors solved this problem by defining a window that encloses a change specified by the revision to physical design. The window defines an area that is less than the area of the entire circuit design. Only the nets within that window are routed pursuant to the revision, leaving the remaining nets in the design unaffected. Then, the results of that incremental routing are inserted into a copy of the original IC design to produce a revised IC design that effects the physical design change without needing to redo the entire process flow.

5. Manufacturers use a process called Chemical Mechanical Planarization/Polishing (“CMP”) to smooth out the surface of the device to prepare the device for further processing, such as deposition of another layer. This allows subsequent layers to be built and connected more easily with fewer opportunities for short circuits or other errors that render the device defective. CMP functions best when there is a certain density and variance of the same material on the surface of the chip. This is because different materials will be “polished” away at different rates, leading to erosion or dishing on the surface.

6. To reduce this problem “dummy” material, also known as “dummy fill,” is typically inserted into low-density regions of the device to increase the overall uniformity of the structures on the surface of the layer and reduce the density variability across the surface of the device. However, dummy fill can increase capacitance if it is placed too close to signal wires, which slows the transmission speed of signals and degrades the overall performance of the device.

7. Just as unwanted capacitance can result from the interaction of elements within the layer of an integrated circuit, it can also result from interaction of elements across adjacent layers. While certain elements (such as signal lines and power lines) cannot be easily moved without affecting circuit performance, there is substantially more flexibility regarding placement, positioning, and spacing of non-signal carrying features such as dummy fill, even when certain

quantities of dummy fill are needed within layers and portions of layers to meet processing requirements.

8. Prior to development of the methodology described in the ’760 patent, the placement of dummy fill in the open areas of the interconnect layer was performed based primarily upon meeting density requirements. To the extent that timing and capacitance effects were considered in dummy fill dimensions, orientation, positioning, or otherwise in dummy fill placement, the conventional dummy fill tools at the time only considered intralayer effects—i.e., interactions between dummy fill features and other elements (such as signal nets) on that same layer. However, use of dummy fill that overlapped on successive layers could and often did create a substantial interlayer bulk capacitive effect that had a negative impact on circuit timing and performance, and which was not considered by the conventional dummy fill tools at the time even when they considered certain intralayer timing effects. *See* Ex. D at 1:43–2:6, 4:11–16.

9. Recognizing these drawbacks, as well as the importance of having a flat or planarized surface on the devices, the inventors of the ’760 patent set out to develop a design process that would also consider the interlayer bulk capacitance created by overlapping dummy fill and consider those intralayer effects in arranging dummy fill in the chip layout so as to minimize the unwanted bulk capacitance created by overlapping dummy fill features.

10. The inventors of the ’760 patent ultimately conceived of a method for addressing the interlayer capacitive effects of dummy fill by treating each successive set of layers as a pair and then rearranging the dummy fill in one or both layers so as to minimize their overlap. This was particularly advantageous in “intelligent dummy fill placement,” i.e., when timing impact is considered when placing dummy fill. *See* Ex. D at 2:10–19.

11. The inventions disclosed in the ’760 patent provide many advantages over the prior

art. In particular, rearranging the dummy fill features such that they do not align vertically in successive layers can reduce unwanted bulk capacitance introduced by dummy fill and thus minimize the interlayer capacitance. *See* Ex. D at 2:45–48, 2:47–59, 3:30–33, 5:19–39. This removed unwanted bulk capacitance that would otherwise slow down signals in the circuit and adversely affect timing in the IC, thus improving its speed and performance. *See* Ex. D at 2:3–6. These significant advantages are achieved through the use of the patented inventions and thus the ’760 patent presents significant commercial value for companies like Ampere.

12. Bell Semic brings this action to put a stop to Ampere’s unauthorized and unlicensed use of the inventions claimed in the ’626 and ’760 patent.

THE PARTIES

13. Plaintiff Bell Semic is a limited liability company organized under the laws of the State of Delaware with a place of business at One West Broad Street, Suite 901, Bethlehem, PA 18018.

14. Bell Semic stems from a long pedigree that began at Bell Labs. Bell Labs sprung out of the Bell System as a research and development laboratory, and eventually became known as one of America’s greatest technology incubators. Bell Labs employees invented the transistor in 1947 in Murray Hill, New Jersey. It was widely considered one of the most important technological breakthroughs of the time, earning the inventors the Nobel Prize in Physics. Bell Labs made the first commercial transistors at a plant in Allentown, Pennsylvania. For decades, Bell Labs licensed its transistor patents to companies throughout the world, creating a technological boom that led to the use of transistors in the semiconductor devices prevalent in most electronic devices today.

15. Bell Semic, a successor to Bell Labs’ pioneering efforts, owns over 1,900

worldwide patents and applications, approximately 1,500 of which are active United States patents. This patent portfolio of semiconductor–related inventions was developed over many years by some of the world's leading semiconductor companies, including Bell Labs, Lucent Technologies, Agere Systems, and LSI Logic and LSI Corporation ("LSI"). This portfolio reflects technology that underlies many important innovations in the development of semiconductors and integrated circuits for high–tech products, including smartphones, computers, wearables, digital signal processors, IoT devices, automobiles, broadband carrier access, switches, network processors, and wireless connectors.

16. The principals of Bell Semic all worked at Bell Labs' Allentown facility, and have continued the rich tradition of innovating, licensing, and helping the industry at large since those early days at Bell Labs. For example, Bell Semic's CTO was a LSI Fellow and Broadcom Fellow. He is known throughout the world as an innovator with more than 300 patents to his name, and he has a sterling reputation for helping semiconductor fabs improve their efficiency. Bell Semic's CEO took a brief hiatus from the semiconductor world to work with Nortel Networks in the telecom industry during its bankruptcy. His efforts saved the pensions of tens of thousands of Nortel retirees and employees. In addition, several Bell Semic executives previously served as engineers at many of these companies and were personally involved in creating the ideas claimed throughout Bell Semic's extensive patent portfolio.

17. On information and belief, Ampere has its principal place of business and headquarters at 4655 Great America Parkway, Suite 601, Santa Clara, CA 95054. On information and belief, Ampere develops, designs, and/or manufactures products in the United States, including in this District, according to the '626 and '760 patented processes/methodologies; and/or uses the '626 and '760 patented processes/methodologies in the United States, including

in this District, to make products; and/or distributes, markets, sells, or offers to sell in the United States and/or imports products into the United States, including in this District, that were manufactured or otherwise produced using the patented process. Additionally, Ampere introduces those products into the stream of commerce knowing that they will be sold and/or used in this District and elsewhere in the United States.

JURISDICTION AND VENUE

18. This is an action for patent infringement arising under the Patent Laws of the United States, Title 35 of the United States Code. Accordingly, this Court has subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a).

19. This Court has personal jurisdiction over Ampere under the laws of the State of Oregon, due at least to its substantial business in Oregon and in this District. Ampere has purposefully and voluntarily availed itself of the privileges of conducting business in the United States, in the State of Oregon, and in this District by continuously and systematically placing goods into the stream of commerce through an established distribution channel with the expectation that they will be purchased by consumers in this District. In the State of Oregon and in this District, Ampere , directly or through intermediaries: (i) performs at least a portion of the infringements alleged herein; (ii) develops, designs, and/or manufactures products according to the ’626 and/or ’760 patented process/methodology; (iii) distributes, markets, sells, or offers to sell products formed according to the ’626 and/or ’760 patented process/methodology; and/or (iv) imports products formed according to the ’626 and/or ’760 patented process/methodology.

20. On information and belief, venue is proper in this Court pursuant to 28 U.S.C. §§ 1391 and 1400 because Ampere has committed, and continues to commit, acts of infringement in this District and has a regular and established place of business in this District. For example,

Ampere maintains a regular and established place of business in the District at 1250 NW 9th Avenue, Suite 800, Portland, OR 97209.

21. Currently, Ampere is advertising many jobs in Oregon. *See Search Jobs*, Ampere Careers (https://amperecomputing.com/apply/) (last visited November 3, 2022). These positions include those that relate to the ’626 and ’760 patented technologies, such as positions for a Principal Verification Engineer CPU, a Principal Physical Design Engineer (Processor and Chip-level), and a Senior Principal CPU Design Engineer.

22. Venue is also convenient in this District. This is at least true because of this District’s close ties to this case—including the technology, relevant witnesses, and sources of proof noted above—and its ability to quickly and efficiently move this case to resolution. Further, Ampere has a regular and established place of business in Oregon and has purposely availed itself of the court system in this District on multiple occasions.

23. On information and belief, Bell Semic’s causes of action arise directly from Ampere’s circuit design work and other activities in this District. Moreover, on information and belief, Ampere has derived substantial revenues from its infringing acts occurring within the State of Oregon and within this District.

U.S. PATENT NO. 7,231,626

24. Bell Semiconductor owns by assignment the entire right, title, and interest in the ’626 patent, entitled “Method Of Implementing An Engineering Change Order In An Integrated Circuit Design By Windows.”

25. A true and correct copy of the ’626 patent is attached as Exhibit A.

26. The ’626 patent issued to inventors Jason K. Hoff, Viswanathan Lakshmanan, Michael Josephides, Daniel W. Prevedel, Richard D. Blinne, and Johathan P. Kuppinger.

27. The application that resulted in issuance of the’626 patent, United States Patent Application No. 11/015,123, was filed December 17, 2004. It issued on June 12, 2007 and expires on July 26, 2025.

28. The ’626 patent generally relates to “methods of implementing an engineering change order (ECO) in an integrated circuit design.” Ex. A at 1:1–13.

29. The background section of the ’626 patent identifies the shortcomings of the prior art. More specifically, the specification describes that the prior circuit design methodology was disadvantageous because “[i]n previous methods for implementing an engineering change order (ECO) request in an integrated circuit design, design tools are run for the entire integrated circuit design, even though the engineering change order typically is only a small fraction of the size of the integrated circuit design” Ex. A at 2:15–19.

30. The ’626 patent elaborates that because “cell placement, routing, design rule check validation, and timing closure run times typically scale with the size of the entire integrated circuit design,” Ex. A at 2:20–22, this produced a “typical turnaround time” of “about one week regardless of the size of the engineering change order. . . . because although the engineering change order may only have a size of a few cells, it must be merged with an integrated circuit design that typically has a much greater size.” *Id.* at 2:37–44. Certain of these steps “may be especially time consuming and resource intensive.” *Id.* at 3:16–17.

31. The inventions disclosed in the ’626 patent provide many advantages over the prior art. In particular, they provide a simple and efficient method for ensuring that revisions to the physical design of the IC do not unduly delay the completion of the design process. As the ’626 patent explains, “significant savings in the resources required to perform routing, design rule check verification, net delay calculation, and parasitic extraction may be realized by creating

windows in the integrated circuit design that include only the incremental changes to the overall integrated circuit design." Ex. A at 3:19–23.

32. As mentioned above, this is very beneficial because it substantially reduces the run time of the routing tools and related follow-on steps of the layout portion of the design process flow (such as calculation of net delay, design rule check, and parasitic extraction). Thus, it shortens the overall design timeline, and avoids cost overruns and delays, making it less costly to make changes later in the design process or more often. *See id.*

33. Given the aforementioned increased complexity of circuit designs and the corresponding delays from design changes, these efficiency gains have become more and more important in completing the design process without affecting time-to-market. These significant advantages are achieved through the use of the patented inventions and thus the '626 patent presents significant commercial value for chip designers.

34. In light of the drawbacks of the prior art, the '626 patent's inventors recognized the need for a circuit design methodology in which the time required to implement an ECO "depend[s] on the number of net changes in the [ECO] rather than on the total number of nets in the entire integrated circuit design." Ex. A at 2:51–53. The inventions claimed in the'626 patent address this need.

35. The '626 patent contains two independent claims and 8 total claims, covering a method and computer readable medium for implementing a change order in an integrated circuit design. Claim 1 reads:

> 1. A method comprising steps of:
>
> (a) receiving as input an integrated circuit design;

> (b) receiving as input an engineering change order to the integrated circuit design;
>
> (c) creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design;
>
> (d) performing an incremental routing of the integrated circuit design only for each net in the integrated circuit design that is enclosed by the window;
>
> (e) replacing an area in a copy of the integrated circuit design that is bounded by the coordinates of the window with results of the incremental routing to generate a revised integrated circuit design; and
>
> (f) generating as output the revised integrated circuit design.

36. This claim, as a whole, provides significant benefits and improvements to the function of the semiconductor device design process, *e.g.*, providing a novel and substantially more efficient process flow in which only the affected nets would be considered in the incremental routing. This results in substantial reduction in the expected time of the design portion of producing semiconductor devices.

37. The claims of the ’626 patent also recite inventive concepts that improve the functioning of the fabrication process, particularly as to post-ECO routing. The claims of the ’626 patent disclose a new and novel solution to specific problems related to improving semiconductor fabrication. As explained in detail above and in the ’626 patent specification, the claimed inventions improve upon the prior art processes by ignoring nets that are unaffected by an ECO in performing routing following the ECO. This has the advantage of substantially reducing the impact on design schedule of ECOs and other layout changes, thus increasing the efficiency of the design process and making it easier to improve the design and fix design errors without unduly delaying time-to-market. By making it easier to fix errors as they are found, and causing substantially less incremental delay upon finding and fixing errors, the claimed inventive

processes also increase the performance and reliability of the finished product. Because of the claimed inventive processes, individual less impactful design issues that still impact design performance (albeit not on a critical scale) can be caught and fixed without costing the same delay as more substantial errors.

U.S. PATENT NO. 7,396,760

38. Bell Semic is the owner by assignment of the ’760 patent. The ’760 patent is titled “Method and System for Reducing Inter-Layer Capacitance in Integrated Circuits.”

39. A true and correct copy of the ’760 patent is attached as Exhibit D.

40. The inventors of the ’760 patent are Kunal Taravade, Neal Callan, and Paul Filseth.

41. The ’760 patent issued on July 8, 2008 from an application filed on November 17, 2004.

42. The ’760 patent generally relates to “a method for reducing inter-layer capacitance” in integrated circuits “through dummy fill methodology.” Ex. D at 1:8–10.

43. The background section of the ’760 patent identifies the shortcomings of the prior art. More specifically, the specification describes that the prior dummy fill methodologies were disadvantageous because they typically focused on achieving uniformity of feature density and failed to sufficiently address adverse effects of the dummy fill on electric field and unwanted bulk capacitance. *See* Ex. D at 1:62–66. In addition, these dummy fill methodologies only considered intralayer effects of dummy fill, to the extent that they considered timing impact at all. *See* Ex. D at 1:66–2:3. Thus, placement of dummy fill, even if advantageous on each individual layer, could create problems when it overlapped with dummy fill features on successive layers, introducing an additional bulk capacitance component that could be substantial. *See id.* at 4:11–17, 4:25–28. These methodologies failed to consider interlayer effects

such as those caused by the overlap of dummy fill features in successive layers, which could have a substantial negative impact on timing. *See id.* at 2:3–6.

44. In light of the drawbacks of the prior art, the inventors of the ’760 patent recognized a need for “intelligent dummy fill placement to reduce interlayer capacitance caused by overlaps of dummy fill area on successive layers,” which would also “treat[] each consecutive pair of layers together when the intelligent dummy filling placement is performed.” Ex. D at 2:7–13. The inventions claimed in the ’760 patent address this need.

45. The ’760 patent contains two independent claims and 19 total claims. Claim 1 reads:

> 1. A method for placing dummy fill patterns in an integrated circuit fabrication process, comprising:
>
> (a) obtaining layout information of the integrated circuit, the integrated circuit including a plurality of layers;
>
> (b) obtaining a first dummy fill space for a first layer based on the layout information;
>
> (c) obtaining a second dummy fill space for a second layer, the second layer being placed successively to the first layer;
>
> (d) determining an overlap between the first dummy fill space and the second dummy fill space; and
>
> (e) minimizing the overlap by re-arranging a plurality of first dummy fill features and a plurality of second dummy fill features;
>
> (f) wherein the first dummy fill space includes non-signal carrying lines on the first layer and the second dummy fill space includes non-signal carrying lines on the second layer.

46. This claim, as a whole, provides significant benefits and improvements to the function of the semiconductor device, *e.g.*, minimizing interlayer bulk capacitance and thus improving the timing characteristics and performance of the IC while meeting interconnect

density requirements during processing. *See, e.g.*, Ex. D at 1:37–55, 5:19–39.

47. The claims of the ’760 patent also recite inventive concepts that improve the functioning of the fabrication process, particularly as to dummy filling. The claims of the ’760 patent disclose a new and novel solution to specific problems related to improving semiconductor fabrication. As explained in detail above and in the ’760 patent specification, the claimed inventions improve upon the prior art processes by considering successive layers rather than each layer on its own, and then determining the overlap between dummy fill features on successive layers before rearranging them to minimize their overlap and thus reduce interlayer bulk capacitance. This has advantages such as minimizing the parasitic capacitance of the interconnect layers, especially the bulk capacitance contributed by the interlayer effects of overlapping dummy fill features, while maintaining necessary interconnect density to meet fabrication requirements.

COUNT I – INFRINGEMENT OF U.S. PATENT NO. 7,231,626

48. Bell Semic re-alleges and incorporates by reference the allegations of the foregoing paragraphs as if fully set forth herein.

49. The ’626 patent is valid and enforceable under the United States Patent Laws.

50. Bell Semic owns, by assignment, all right, title, and interest in and to the ’626 patent, including the right to collect for past damages.

51. A copy of the ’626 patent is attached at Exhibit A.

52. On information and belief, Ampere has and continues to directly infringe pursuant to 35 U.S.C. § 271(a) one or more claims of the ’626 patent by using the patented methodology to design one or more semiconductor devices, including as one example the Ampere Accused Product, in the United States.

53. On information and belief, Ampere employs a variety of design tools, for example, Cadence, Synopsys, and/or Siemens tools, to perform incremental routing in implementing an ECO (the "Accused Processes") as recited in the ’626 patent claims. As one example, Ampere’s Accused Processes perform a method for only routing the nets affected by the ECO and merging that changed area into the overall circuit layout as required by claim 1 of the ’626 patent. Ampere does so by employing a design tool, such as at least one of a Cadence, Synopsys, and/or Siemens tool, to perform incremental routing as part of implementing an ECO for the Ampere Accused Product to generate a revised integrated circuit design.

54. Ampere’s Accused Processes also calculate and perform a parasitic extraction only for each net in the IC design enclosed by the window defining the ECO. (This parasitic extraction is also how the Accused Processes further calculate a net delay only for each net in the IC design enclosed by the window defining the ECO.) Ampere does so by employing a design tool, such as at least one of the Cadence, Synopsys, and/or Siemens tools, to perform the incremental routing during implementation of the ECO for the Ampere Accused Product’s circuit designs.

55. Ampere’s Accused Processes also perform a design rule check only for each net in the IC design enclosed by the ECO window. Ampere does so by employing a design tool, such as at least one of the Cadence, Synopsys, and/or Siemens tools, perform the incremental ECO and automatically perform a DRC for those nets to ensure that the ECO did not violate any design rules when it fixed other issues.

56. An exemplary infringement analysis showing infringement of one or more claims of the ’626 patent is set forth in Exhibit B. The declaration of Lloyd Linder, an expert in the field of semiconductor device design, is attached at Exhibit C and further describes Ampere’s infringement of the ’626 patent.

57. Ampere's Accused Processes infringe and continue to infringe one or more claims of the '626 patent during the pendency of the '626 patent.

58. On information and belief, Ampere has and continues to infringe pursuant to 35 U.S.C. § 271, *et. seq.*, directly, either literally or under the doctrine of equivalents, by using the Accused Processes in violation of one or more claims of the '626 patent. Ampere has and continues to infringe pursuant to 35 U.S.C. § 271, *et. seq.*, directly, either literally or under the doctrine of equivalents, by making, selling, or offering to sell in the United States, or importing into the United States products manufactured or otherwise produced using the Accused Processes in violation of one or more claims of the '626 patent.

59. Ampere's infringement of the '626 patent is exceptional and entitles Bell Semic to attorneys' fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

60. Bell Semic has been damaged by Ampere's infringement of the '626 patent and will continue to be damaged unless Ampere is enjoined by this Court. Bell Semic has suffered and continues to suffer irreparable injury for which there is no adequate remedy at law. The balance of hardships favors Bell Semic, and public interest is not disserved by an injunction.

61. Bell Semic is entitled to recover from Ampere all damages that Bell Semic has sustained as a result of Ampere's infringement of the '626 patent, including without limitation and/or not less than a reasonable royalty.

COUNT II – INFRINGEMENT OF U.S. PATENT NO. 7,396,760

62. Bell Semic re-alleges and incorporates by reference the allegations of the foregoing paragraphs as if fully set forth herein.

63. The '760 patent is valid and enforceable under the United States Patent Laws.

64. Bell Semic owns, by assignment, all right, title, and interest in and to the '760

patent, including the right to collect for past damages.

65. A copy of the ’760 patent is attached at Exhibit D.

66. On information and belief, Ampere has and continues to directly infringe pursuant to 35 U.S.C. § 271(a) one or more claims of the ’760 patent by using the patented methodology to design one or more semiconductor devices, including as one example the Accused Product, in the United States.

67. On information and belief, Ampere employs a variety of design tools, for example, Cadence, Synopsys, and/or Siemens tools, to rearrange dummy fill to minimize its overlap in successive layers (the “Accused Processes”) as recited in the ’760 patent claims. As one example, Ampere’s Accused Processes allow arrangement and rearrangement of dummy fill in a timing aware fashion, including with the ability to stagger the dummy fill in successive layers so as to minimize the interlayer bulk capacitance after determining their overlap as required by claim 1 of the ’760 patent. Ampere does so by employing a design tool, such as at least one of a Cadence, Synopsys, and/or Siemens tool, rearrange the dummy fill features in successive layers of its Accused Product.

68. Ampere’s Accused Processes also form the dummy fill features in a grid within one or more of the successive layers, provide square-shaped dummy fill features in one or more of the successive layers, determine the dummy fill space based on a local pattern density in one or more of the successive layers, and minimize total bulk capacitance and/or certain of its components. Ampere does so by employing a design tool, such as at least one of the Cadence, Synopsys, and/or Siemens tools, to implement dummy fill functionality in a timing-aware fashion and with consideration of interlayer capacitive effects in creation and design of its Accused Product.

69. An exemplary infringement analysis showing infringement of one or more claims of the ’760 patent is set forth in Exhibit E. The declaration of Dhaval Brahmbhatt, an expert in the field of semiconductor device design, is attached at Exhibit F and further describes Ampere’s infringement of the ’760 patent.

70. Ampere’s Accused Processes infringe and continue to infringe one or more claims of the ’760 patent during the pendency of the ’760 patent.

71. On information and belief, Ampere has and continues to infringe pursuant to 35 U.S.C. § 271, *et. seq.*, directly, either literally or under the doctrine of equivalents, by using the Accused Processes in violation of one or more claims of the ’760 patent. Ampere has and continues to infringe pursuant to 35 U.S.C. § 271, *et. seq.*, directly, either literally or under the doctrine of equivalents, by making, selling, or offering to sell in the United States, or importing into the United States products manufactured or otherwise produced using the Accused Processes in violation of one or more claims of the ’760 patent.

72. Ampere’s infringement of the ’760 patent is exceptional and entitles Bell Semic to attorneys’ fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

73. Bell Semic has been damaged by Ampere’s infringement of the ’760 patent and will continue to be damaged unless Ampere is enjoined by this Court. Bell Semic has suffered and continues to suffer irreparable injury for which there is no adequate remedy at law. The balance of hardships favors Bell Semic, and public interest is not disserved by an injunction.

74. Bell Semic is entitled to recover from Ampere all damages that Bell Semic has sustained as a result of Ampere’s infringement of the ’760 patent, including without limitation and/or not less than a reasonable royalty.

PRAYER FOR RELIEF

WHEREFORE, Bell Semic respectfully requests that this Court enter judgment in its favor as follows and award Bell Semic the following relief:

(a) a judgment declaring that Ampere has infringed one or more claims of the ’626 patent and ‘760 patent in this litigation pursuant to 35 U.S.C. § 271, *et seq.*;

(b) an award of damages adequate to compensate Bell Semic for infringement of the ’626 patent and ’760 patent by Ampere, in an amount to be proven at trial, including supplemental post-verdict damages until such time as Ampere ceases its infringing conduct;

(c) a permanent injunction, pursuant to 35 U.S.C. § 283, prohibiting Ampere and its officers, directors, employees, agents, consultants, contractors, suppliers, distributors, all affiliated entities, and all others acting in privity with Ampere from committing further acts of infringement;

(d) a judgment requiring Ampere to make an accounting of damages resulting from Ampere’s infringement of the ’626 patent and ’760 patent;

(e) the costs of this action, as well as attorneys’ fees as provided by 35 U.S.C. § 285;

(f) pre-judgment and post-judgment interest at the maximum amount permitted by law;

(g) all other relief, in law or equity, to which Bell Semic is entitled.

DEMAND FOR JURY TRIAL

Plaintiff hereby demands a jury trial for all issues so triable.

//

//

//

//

//

//

//

//

Dated: December 8, 2022

/s/ Peter M. Grabiel
Jeff S. Pitzer, OSB No. 020846
Email: jpitzer@pitzerlaw.net
Peter M. Grabiel, OSB No. 171964
Email: pgrabiel@pitzerlaw.net
PITZER LAW
210 SW Morrison St., Suite 600
Portland, OR 97204
Telephone: 503-227-1477

Paul Richter (*pro hac vice* forthcoming)
prichter@devlinlawfirm.com
DEVLIN LAW FIRM LLC
1526 Gilpin Avenue
Wilmington, Delaware 19806
Telephone: (302) 449–9010

David Sochia (*pro hac vice* forthcoming)
Texas State Bar No. 00797470
dsochia@McKoolSmith.com
Ashley N. Moore (*pro hac vice* forthcoming)
Texas State Bar No. 24074748
amoore@McKoolSmith.com
Richard A. Kamprath (*pro hac vice* forthcoming)
Texas State Bar No. 24078767
rkamprath@McKoolSmith.com
MCKOOL SMITH, P.C.
300 Crescent Court Suite 1500
Dallas, TX 75201
Telephone: (214) 978-4000

Attorneys for Plaintiff Bell Semiconductor, LLC

EXHIBIT A

US007231626B2

(12) United States Patent
Hoff et al.

(10) **Patent No.: US 7,231,626 B2**
(45) **Date of Patent: Jun. 12, 2007**

(54) **METHOD OF IMPLEMENTING AN ENGINEERING CHANGE ORDER IN AN INTEGRATED CIRCUIT DESIGN BY WINDOWS**

(75) Inventors: **Jason K. Hoff**, Houston, TX (US); **Viswanathan Lakshmanan**, Thornton, CO (US); **Michael Josephides**, Broomfield, CO (US); **Daniel W. Prevedel**, Fort Collins, CO (US); **Richard D. Blinne**, Ft. Collins, CO (US); **Johathan P. Kuppinger**, Windsor, CO (US)

(73) Assignee: **LSI Corporation**, Milpitas, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 221 days.

(21) Appl. No.: **11/015,123**

(22) Filed: **Dec. 17, 2004**

(65) **Prior Publication Data**

US 2006/0136855 A1 Jun. 22, 2006

(51) **Int. Cl.**
G06F 17/50 (2006.01)

(52) **U.S. Cl.** **716/13**; 716/14; 716/9; 716/10; 716/6

(58) **Field of Classification Search** 716/13, 716/14, 9, 10, 5, 6
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,416,717	A *	5/1995	Miyama et al.	703/14
5,761,080	A *	6/1998	DeCamp et al.	716/5
5,847,967	A *	12/1998	Asao	716/5
5,953,236	A *	9/1999	Hossain et al.	716/6
6,014,506	A *	1/2000	Hossain et al.	716/11
6,141,663	A	10/2000	Hunkins et al.	
6,243,715	B1	6/2001	Bogantz et al.	
6,317,865	B1 *	11/2001	Itoh	716/11
6,453,454	B1 *	9/2002	Lee et al.	716/11
6,536,028	B1 *	3/2003	Katsioulas et al.	716/17
6,651,239	B1 *	11/2003	Nikitin et al.	716/18
6,922,817	B2 *	7/2005	Bradfield et al.	716/1
7,051,307	B2 *	5/2006	Ditlow et al.	716/7
7,168,055	B2 *	1/2007	Hoff et al.	716/4
2004/0003366	A1 *	1/2004	Suzuki	716/11
2004/0199879	A1 *	10/2004	Bradfield et al.	716/1

(Continued)

OTHER PUBLICATIONS

Xiang et al., "ECO Algorithms for Removing Overlaps Between Power Rails and Signal Wires", IEEE/ACM International Conference on Computer Aided Design, Nov. 10, 2002, pp. 67-74.*

Primary Examiner—Phallaka Kik
(74) *Attorney, Agent, or Firm*—Eric J. Whitesell

(57) **ABSTRACT**

A method of implementing an engineering change order includes steps of: (a) receiving as input an integrated circuit design; (b) receiving as input an engineering change order to the integrated circuit design; (c) creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design; (d) performing a routing of the integrated circuit design that excludes routing of any net that is not enclosed by the window; (e) replacing an area in a copy of the integrated circuit design that is bounded by the coordinates of the window with results of the incremental routing to generate a revised integrated circuit design; and (f) generating as output the revised integrated circuit design.

8 Claims, 5 Drawing Sheets

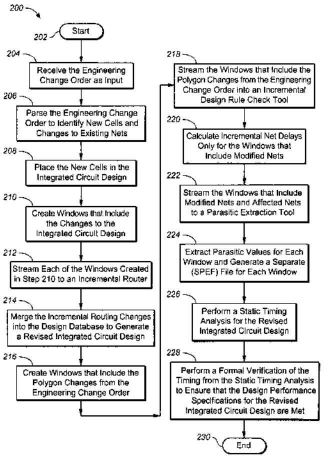

U.S. PATENT DOCUMENTS

2005/0015735	A1 *	1/2005	Tamura et al.	716/2
2005/0125756	A1 *	6/2005	Ditlow et al.	716/7
2005/0160391	A1 *	7/2005	Orita	716/13
2005/0166169	A1 *	7/2005	Kurzum et al.	716/10
2005/0235240	A1 *	10/2005	Tien	716/11
2006/0036982	A1 *	2/2006	Hoff et al.	716/5
2006/0288322	A1 *	12/2006	Zhang et al.	716/9

* cited by examiner

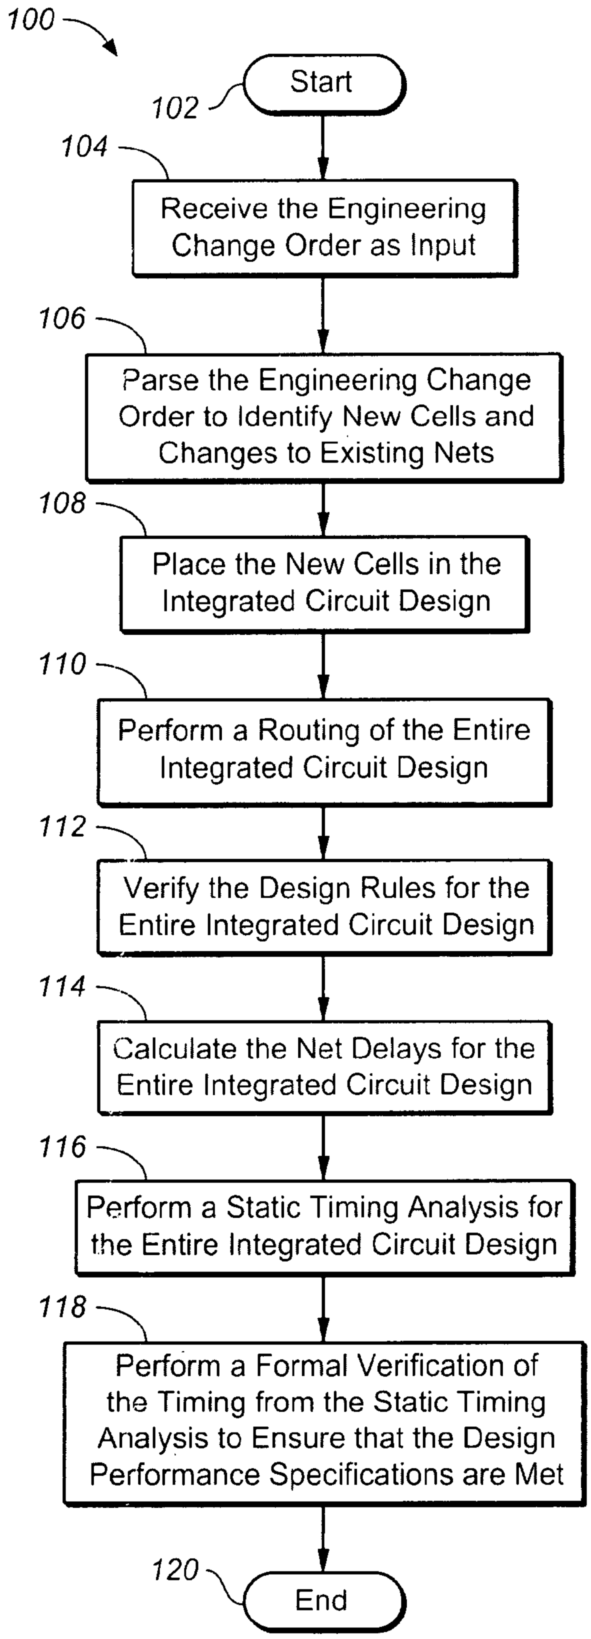
100
Start
102
104
Receive the Engineering Change Order as Input
106
Parse the Engineering Change Order to Identify New Cells and Changes to Existing Nets
108
Place the New Cells in the Integrated Circuit Design
110
Perform a Routing of the Entire Integrated Circuit Design
112
Verify the Design Rules for the Entire Integrated Circuit Design
114
Calculate the Net Delays for the Entire Integrated Circuit Design
116
Perform a Static Timing Analysis for the Entire Integrated Circuit Design
118
Perform a Formal Verification of the Timing from the Static Timing Analysis to Ensure that the Design Performance Specifications are Met
120
End

FIG. 1

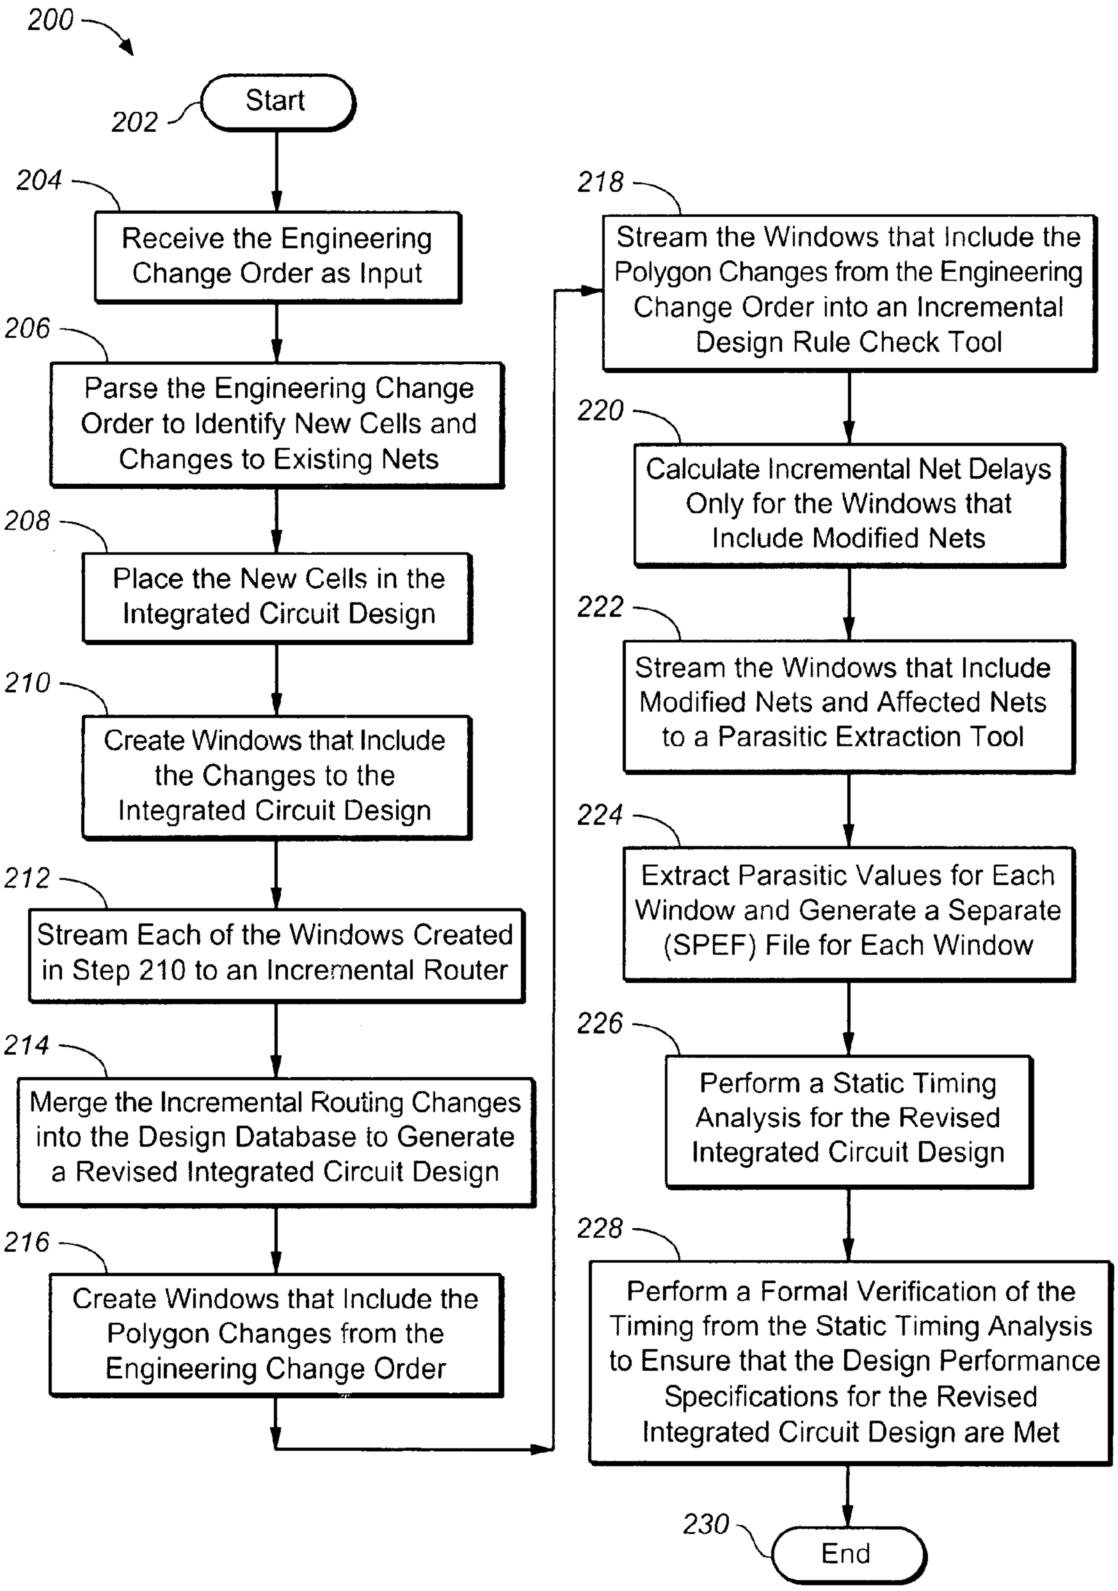
200
202
Start
204
Receive the Engineering Change Order as Input
206
Parse the Engineering Change Order to Identify New Cells and Changes to Existing Nets
208
Place the New Cells in the Integrated Circuit Design
210
Create Windows that Include the Changes to the Integrated Circuit Design
212
Stream Each of the Windows Created in Step 210 to an Incremental Router
214
Merge the Incremental Routing Changes into the Design Database to Generate a Revised Integrated Circuit Design
216
Create Windows that Include the Polygon Changes from the Engineering Change Order
218
Stream the Windows that Include the Polygon Changes from the Engineering Change Order into an Incremental Design Rule Check Tool
220
Calculate Incremental Net Delays Only for the Windows that Include Modified Nets
222
Stream the Windows that Include Modified Nets and Affected Nets to a Parasitic Extraction Tool
224
Extract Parasitic Values for Each Window and Generate a Separate (SPEF) File for Each Window
226
Perform a Static Timing Analysis for the Revised Integrated Circuit Design
228
Perform a Formal Verification of the Timing from the Static Timing Analysis to Ensure that the Design Performance Specifications for the Revised Integrated Circuit Design are Met
230
End

FIG. 2

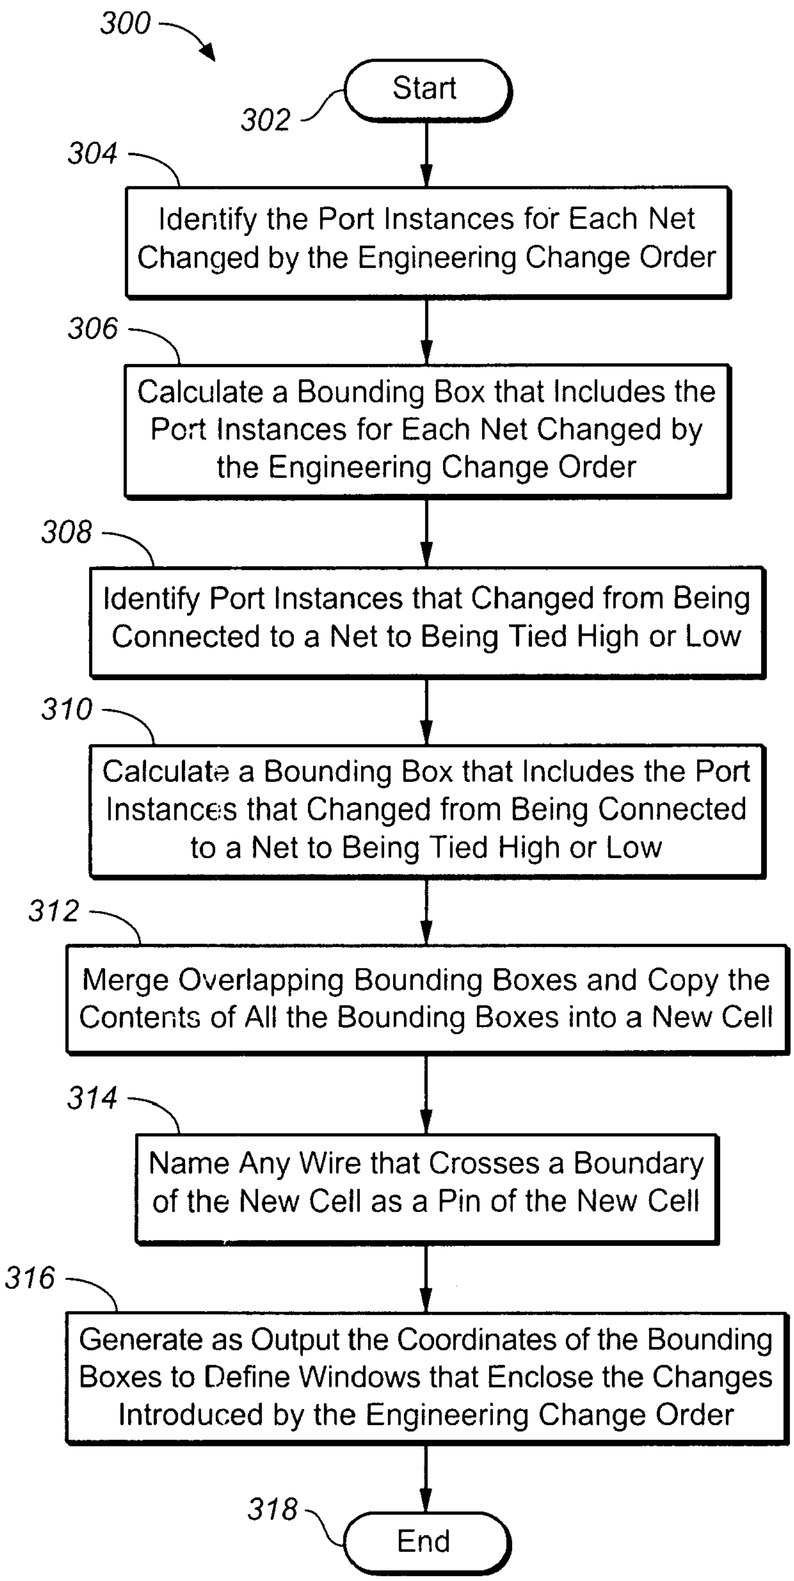
300
Start
302
304
Identify the Port Instances for Each Net Changed by the Engineering Change Order
306
Calculate a Bounding Box that Includes the Port Instances for Each Net Changed by the Engineering Change Order
308
Identify Port Instances that Changed from Being Connected to a Net to Being Tied High or Low
310
Calculate a Bounding Box that Includes the Port Instances that Changed from Being Connected to a Net to Being Tied High or Low
312
Merge Overlapping Bounding Boxes and Copy the Contents of All the Bounding Boxes into a New Cell
314
Name Any Wire that Crosses a Boundary of the New Cell as a Pin of the New Cell
316
Generate as Output the Coordinates of the Bounding Boxes to Define Windows that Enclose the Changes Introduced by the Engineering Change Order
318
End

FIG. 3

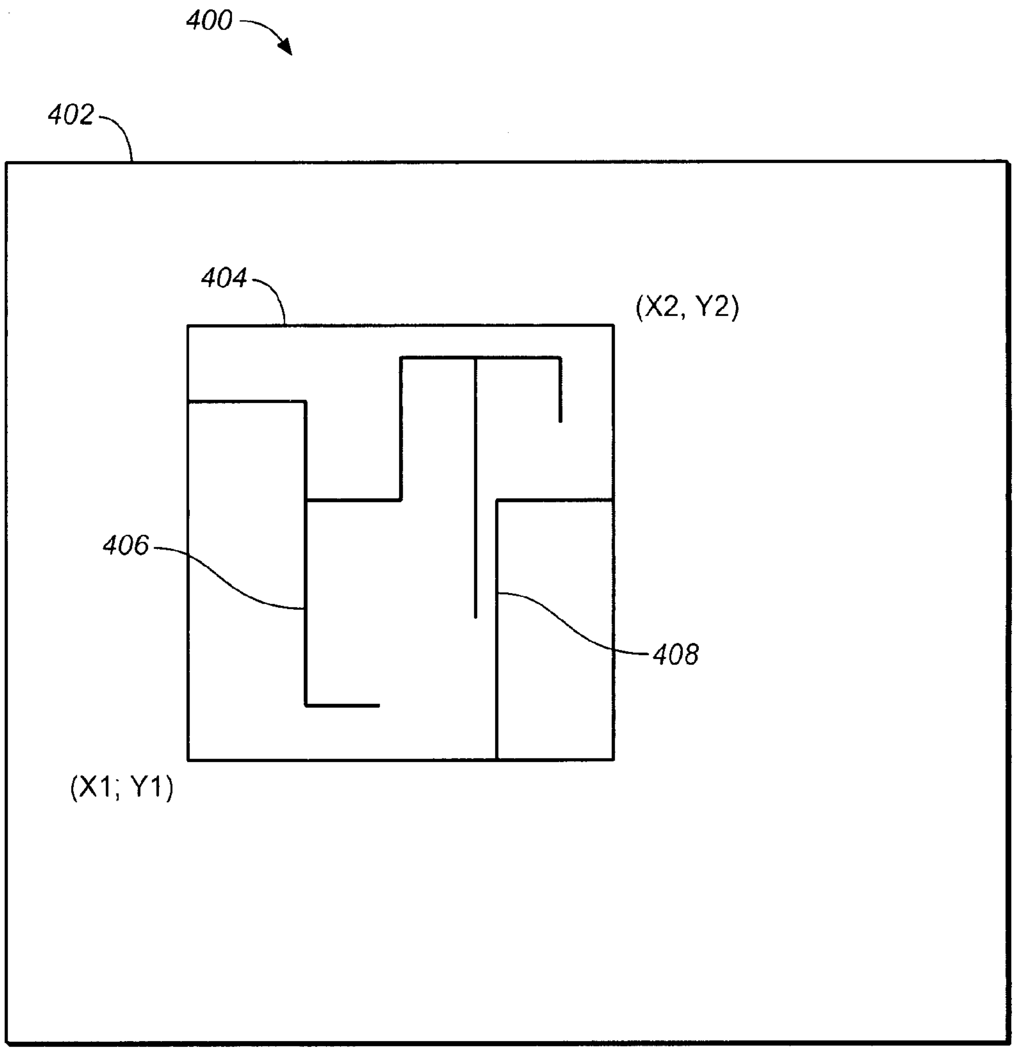

FIG. 4

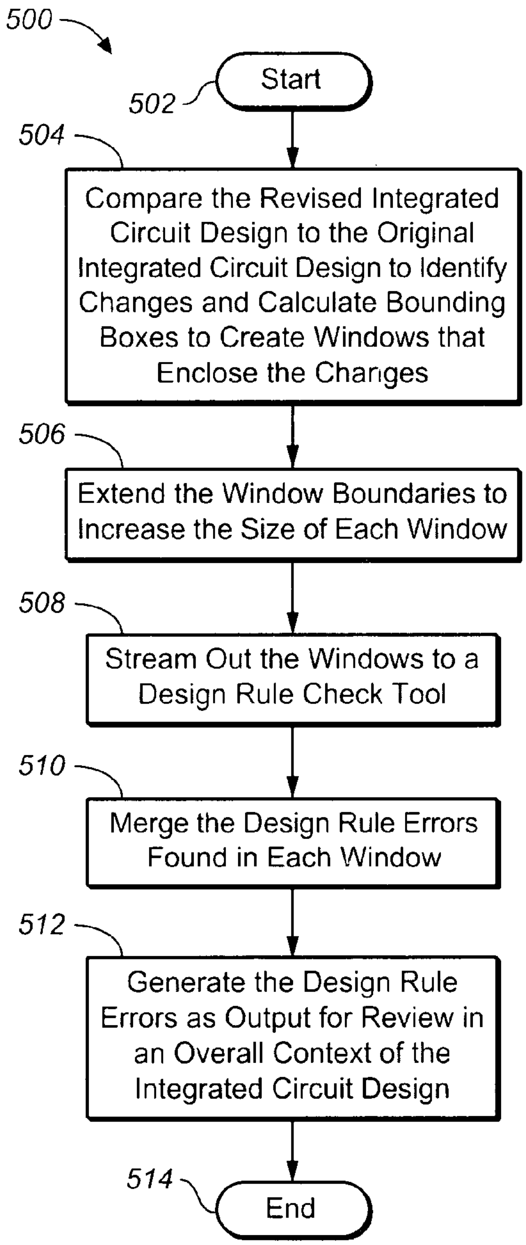
500
502
Start
504
Compare the Revised Integrated Circuit Design to the Original Integrated Circuit Design to Identify Changes and Calculate Bounding Boxes to Create Windows that Enclose the Changes
506
Extend the Window Boundaries to Increase the Size of Each Window
508
Stream Out the Windows to a Design Rule Check Tool
510
Merge the Design Rule Errors Found in Each Window
512
Generate the Design Rule Errors as Output for Review in an Overall Context of the Integrated Circuit Design
514
End

FIG. 5

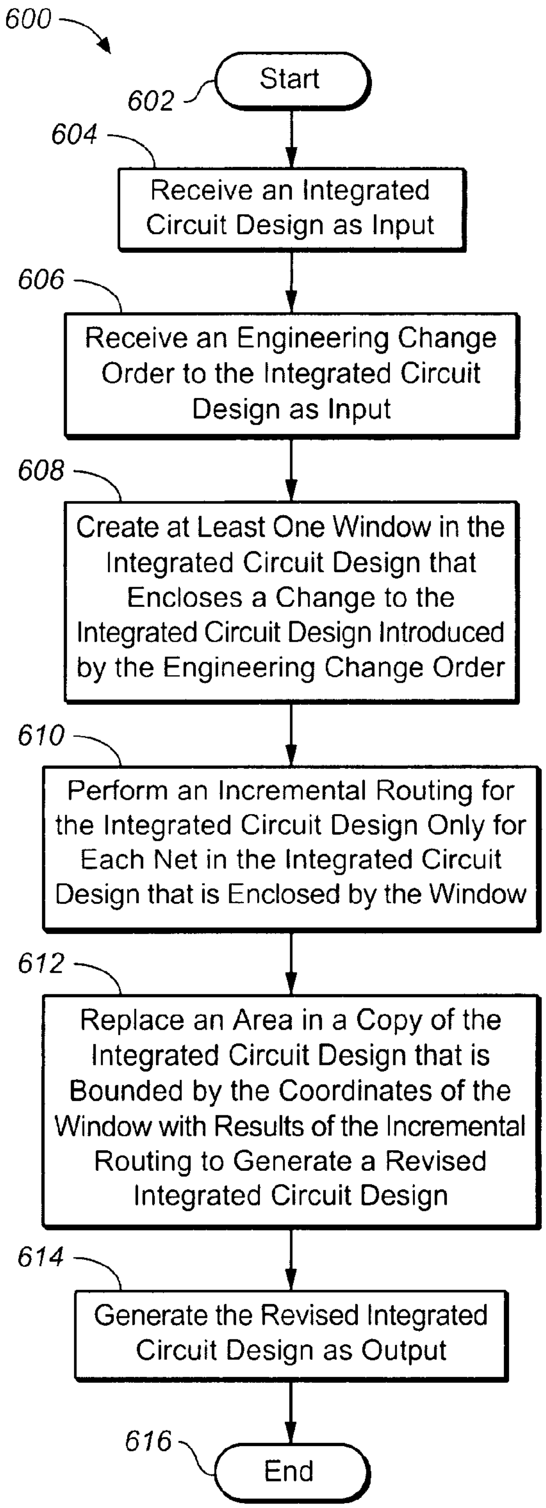
600
602
Start
604
Receive an Integrated Circuit Design as Input
606
Receive an Engineering Change Order to the Integrated Circuit Design as Input
608
Create at Least One Window in the Integrated Circuit Design that Encloses a Change to the Integrated Circuit Design Introduced by the Engineering Change Order
610
Perform an Incremental Routing for the Integrated Circuit Design Only for Each Net in the Integrated Circuit Design that is Enclosed by the Window
612
Replace an Area in a Copy of the Integrated Circuit Design that is Bounded by the Coordinates of the Window with Results of the Incremental Routing to Generate a Revised Integrated Circuit Design
614
Generate the Revised Integrated Circuit Design as Output
616
End

FIG. 6

METHOD OF IMPLEMENTING AN ENGINEERING CHANGE ORDER IN AN INTEGRATED CIRCUIT DESIGN BY WINDOWS

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to the design of integrated circuits. More specifically, but without limitation thereto, the present invention relates to methods of implementing an engineering change order (ECO) in an integrated circuit design.

2. Description of Related Art

In previous methods for implementing an engineering change order (ECO) request in an integrated circuit design, design tools are run for the entire integrated circuit design, even though the engineering change order typically is only a small fraction of the size of the integrated circuit design. For example, cell placement, routing, design rule check validation, and timing closure run times typically scale with the size of the entire integrated circuit design.

SUMMARY OF THE INVENTION

In one embodiment of the present invention, a method of implementing an engineering change order in an integrated circuit design by windows includes steps of:

(a) receiving as input an integrated circuit design;
(b) receiving as input an engineering change order to the integrated circuit design;
(c) creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design;
(d) performing a routing only for each net in the integrated circuit design that is enclosed by the window;
(e) replacing an area in a copy of the integrated circuit design that is bounded by coordinates of the window with results of the incremental routing to generate a revised integrated circuit design; and
(f) generating as output the revised integrated circuit design.

In another embodiment of the present invention, a computer program product for implementing an engineering change order in an integrated circuit design by windows includes:

a medium for embodying a computer program for input to a computer; and
a computer program embodied in the medium for causing the computer to perform steps of:
(a) receiving as input an integrated circuit design;
(b) receiving as input an engineering change order to the integrated circuit design;
(c) creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design;
(d) performing a routing of the integrated circuit design only for each net in the integrated circuit design that is enclosed by the window;
(e) replacing an area in a copy of the integrated circuit design that is bounded by coordinates of the window with results of the incremental routing to generate a revised integrated circuit design; and
(f) generating as output the revised integrated circuit design.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the accompanying figures, in which like references indicate similar elements throughout the several views of the drawings, and in which:

FIG. **1** illustrates a flow chart for a method of implementing an engineering change order in an integrated circuit design according to the prior art;

FIGS. **2**A and **2**B illustrate a flow chart for a method of implementing an engineering change order in an integrated circuit design by windows;

FIG. **3** illustrates a flow chart for creating an engineering change order window for FIGS. **2**A and **2**B;

FIG. **4** illustrates a diagram of a window in an integrated circuit design;

FIG. **5** illustrates a flow chart for performing an incremental design rule check for FIGS. **2**A and **2**B; and

FIG. **6** illustrates a flow chart of a computer program for implementing an engineering change order in an integrated circuit design by windows.

Elements in the figures are illustrated for simplicity and clarity and have not necessarily been drawn to scale. For example, the dimensions of some elements in the figures may be exaggerated relative to other elements to point out distinctive features in the illustrated embodiments of the present invention.

DESCRIPTION OF THE ILLUSTRATED EMBODIMENTS

In previous methods for implementing a functional or timing engineering change order (ECO) to an integrated circuit design, the typical turnaround time is typically about one week regardless of the size of the engineering change order. This is because although the engineering change order may only have a size of a few cells, it must be merged with an integrated circuit design that typically has a much greater size. For example, if an engineering change order for five cells may be required for an integrated circuit design that includes five million cells. As a result, design tool run times generally scale with the size of the entire integrated circuit design for routing, design rule check verification, net delay calculation, and parasitic extraction. Preferably, the time required to implement an engineering change order should depend on the number of net changes in the engineering change order rather than on the total number of nets in the entire integrated circuit design.

FIG. **1** illustrates a flow chart **100** for a method of implementing an engineering change order according to the prior art.

Step **102** is the entry point for the flow chart **100**.

In step **104**, the engineering change order is received as input.

In step **106**, the engineering change order is parsed to identify new cells and changes to existing nets.

In step **108**, the new cells are placed in the integrated circuit design by a software place and route design tool.

In step **110**, a routing of the entire integrated circuit design is performed by the place and route design tool.

In step **112**, the design rules for the technology used to manufacture the integrated circuit are verified for the entire integrated circuit design, for example, by design rule check software such as Mentor Calibre™.

In step **114**, the net delays are calculated for the entire integrated circuit design.

In step **116**, a parasitic extraction is performed for the entire integrated circuit design to determine the values of net coupling capacitance and parasitic resistance.

In step **118**, a static timing analysis is performed for the entire integrated circuit design to determine the effect of net delay including net parasitic capacitance and resistance on the integrated circuit design.

In step **120**, a formal verification of the timing is performed from the static timing analysis to ensure that the design timing specifications are met.

Step **122** is the exit point of the flow chart **100**.

In the method of FIG. **1**, steps **110**, **112**, **114** and **116** may be especially time consuming and resource intensive, depending on the complexity of the integrated circuit design. A significant savings in the resources required to perform routing, design rule check verification, net delay calculation, and parasitic extraction may be realized by creating windows in the integrated circuit design that include only the incremental changes to the overall integrated circuit design as follows.

In one embodiment of the present invention, a method of implementing an engineering change order in an integrated circuit design includes steps of:

(a) receiving as input an integrated circuit design;

(b) receiving as input an engineering change order to the integrated circuit design;

(c) creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design;

(d) performing a routing only for each net in the integrated circuit design that is enclosed by the window;

(e) replacing an area in a copy of the integrated circuit design that is bounded by coordinates of the window with results of the incremental routing to generate a revised integrated circuit design; and

(f) generating as output the revised integrated circuit design.

FIGS. **2**A and **2**B illustrate a flow chart **200** for a method of implementing an engineering change order in an integrated circuit design by windows.

Step **202** is the entry point for the flow chart **200**.

In step **204**, the engineering change order is received as input as in FIG. **1**.

In step **206**, the engineering change order is parsed to identify new nets and changes to existing nets that constitute the changes to the integrated circuit design in the same manner as in FIG. **1**.

In step **208**, the new cells are placed in the integrated circuit design, for example, by a software place and route design tool in the same manner as FIG. **1**.

In step **210**, windows are created that include the new cells and the net changes that constitute the changes to the integrated circuit design that are to be routed. The term "window" as used herein is defined as a rectilinear boundary that encloses an area of the integrated circuit design that is less than the entire area of the integrated circuit design. For example, a window may include a subset of nets that have been changed by the engineering change order. Alternatively, a window may include polygons that have been introduced or changed by the engineering change order. The window boundaries are calculated from the coordinates of the new polygons and the changed nets in the integrated circuit design database so that each of the changes to the integrated circuit design is enclosed by a window.

In step **212**, each of the windows created in step **210** is streamed to an incremental router. The incremental routing may be performed by the same routing tool used in FIG. **1**, however, only the nets that are modified by the engineering change order are routed, in contrast to routing the entire integrated circuit design as in FIG. **1**. Windows that do not overlap may be routed in parallel, while windows that do overlap are routed serially so that any duplicated routing may be removed. If any nets in a window are found open, that is, not all of the net connections are included in the window, then the net is "frozen", which means that the net may not be changed by the router. In addition, a partition manager is preferably included in the incremental router that allows the user to expand the size of the windows and merge overlapping windows.

In step **214**, the incremental routing changes are merged into the design database, for example, by replacing the contents enclosed by the coordinates of each window in a copy of the original integrated circuit design by the contents of the window to generate a revised integrated circuit design.

In step **216**, windows are created that include the polygon changes from the engineering change order.

In step **218**, the windows that include the polygon changes from the engineering change order are streamed into an incremental design rule check tool. The incremental design rule check tool checks only the polygons that were changed, advantageously avoiding unnecessary re-checking of the all the polygons in the integrated circuit database.

In step **220**, incremental net delays are calculated only for the windows that include modified nets, advantageously avoiding unnecessary re-calculation of all the net delays in the integrated circuit design.

In step **222**, the windows that include modified nets and affected nets are streamed to a parasitic extraction tool. An affected net is a net that has a coupling capacitance with a modified net that exceeds a predefined coupling capacitance threshold.

In step **224**, parasitic values are extracted for each window, and a separate standard parasitic extraction format (SPEF) file is generated by the parasitic extraction tool for each window.

In step **226**, a static timing analysis is performed for the revised integrated circuit design in the same manner as in FIG. **1**.

In step **228**, a formal verification of the timing is performed from the static timing analysis to ensure that the design performance specifications for the revised integrated circuit design are met in the same manner as in FIG. **1**.

Step **230** is the exit point of the flow chart **200**.

FIG. **3** illustrates a flow chart **300** for creating an engineering change order window for FIG. **2**.

Step **302** is the entry point of the flow chart **300**.

In step **304**, the port instances for each net changed by the engineering change order are identified. A net change may be, for example, a net that has moved or has different connections.

In step **306**, a bounding box that includes the port instances for each net changed by the engineering change order is calculated from the net coordinates in the design database of the original integrated circuit design.

In step **308**, port instances that changed from being connected to a net to being tied high or low are identified.

In step **310**, a bounding box that includes the port instances that changed from being connected to a net to being tied high or low is calculated from the coordinates in the design database.

In step **312**, overlapping bounding boxes are merged, and the contents of all the bounding boxes are copied into a new cell.

In step **314**, any wire that crosses a boundary of the new cell is named as a pin of the new cell.

In step **316**, the coordinates of the bounding boxes are generated as output to define windows that enclose the changes introduced by the engineering change order.

Step **318** is the exit point of the flow chart **300**.

FIG. **4** illustrates a diagram of a window in an integrated circuit design. Shown in FIG. **4** are an integrated circuit design **402**, a window **404**, a changed net **406**, and an affected net **408**.

In FIG. **4**, The window **404** is bounded by the coordinates (X**1**, Y**1**):(X**2**, Y**2**) that enclose the changed net **406** and the affected net **408**. The area enclosed by the window **404** is less than the entire area of the integrated circuit design **402**, thereby reducing the number of calculations required to implement the engineering change order.

In this example, the changed net **406** has been moved, resulting in a coupling capacitance with the affected net **408** that exceeds a threshold coupling capacitance. The affected net **408** is therefore included in the window for calculating the incremental net delay and for performing the parasitic extraction.

FIG. **5** illustrates a flow chart **500** for performing an incremental design rule check for FIGS. **2**A and **2**B

Step **502** is the entry point of the flow chart **500**.

In step **504**, the revised integrated circuit design is compared to the original integrated circuit design to identify changes, and bounding boxes are calculated to create windows that enclose only the physical changes and not the entire nets.

In step **506**, the window boundaries are extended to increase the size of each window, forming a halo margin or region around each window. The halo region allows the design rule check software to examine objects that are nearby each window. The size of the halo margin may be determined, for example, by the design rule check tool or by the user to ensure that there is sufficient room for the spacing rules to work correctly.

In step **508**, the windows are streamed out, for example, in GDSII format, to a design rule check tool, which may be the same as that used in FIG. **1**.

In step **510**, the design rule errors found in each window are merged together, removing any errors in the halo region.

In step **512**, the design rule errors are generated as output for review in an overall context of the integrated circuit design.

Step **514** is the exit point of the flow chart **500**.

The flow chart described above may also be implemented by instructions for being performed on a computer. The instructions may be embodied in a disk, a CD-ROM, and other computer readable media according to well known computer programming techniques.

In another aspect of the present invention, a computer program product for analyzing noise for an integrated circuit design includes:

a medium for embodying a computer program for input to a computer; and

a computer program embodied in the medium for causing the computer to perform steps of:

(a) receiving as input an integrated circuit design;

(b) receiving as input an engineering change order to the integrated circuit design;

(c) creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design;

(d) performing a routing only for each net in the integrated circuit design that is enclosed by the window;

(e) replacing an area in a copy of the integrated circuit design that is bounded by coordinates of the window with results of the incremental routing to generate a revised integrated circuit design; and

(f) generating as output the revised integrated circuit design.

FIG. **6** illustrates a flow chart **600** of a computer program for implementing an engineering change order in an integrated circuit design by windows.

Step **602** is the entry point of the flow chart **600**.

In step **604**, an integrated circuit design is received as input.

In step **606**, an engineering change order to the integrated circuit design is received as input.

In step **608**, at least one window is created in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order. The window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design.

In step **610**, an incremental routing is performed for the integrated circuit design only for each net in the integrated circuit design that is enclosed by the window, advantageously avoiding repeating calculations for nets that are not changed or affected by the engineering change order.

In step **612**, an area in a copy of the integrated circuit design that is bounded by the coordinates of the window is replaced with results of the incremental routing to generate a revised integrated circuit design.

In step **614**, the revised integrated circuit design is generated as output.

Step **616** is the exit point of the flow chart **600**.

Although the methods illustrated by the flowchart descriptions above are described and shown with reference to specific steps performed in a specific order, these steps may be combined, sub-divided, or reordered without departing from the scope of the claims. Unless specifically indicated herein, the order and grouping of steps are not limitations of the claims.

The specific embodiments and applications thereof described above are for illustrative purposes only and do not preclude modifications and variations that may be made thereto by those skilled in the art within the scope of the following claims.

What is claimed is:

1. A method comprising steps of:

(a) receiving as input an integrated circuit design;

(b) receiving as input an engineering change order to the integrated circuit design;

(c) creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design;

(d) performing an incremental routing of the integrated circuit design only for each net in the integrated circuit design that is enclosed by the window;

(e) replacing an area in a copy of the integrated circuit design that is bounded by the coordinates of the window with results of the incremental routing to generate a revised integrated circuit design; and

(f) generating as output the revised integrated circuit design.

2. The method of claim **1** further comprising a step of calculating a net delay only for each net in the integrated circuit design that is enclosed by the window.

3. The method of claim **1** further comprising a step of performing a design rule check only for each net in the integrated circuit design that is enclosed by the window.

4. The method of claim **1** further comprising a step of performing a parasitic extraction only for each net in the integrated circuit design that is enclosed by the window.

5. A computer readable storage medium tangibly embodying instructions for a computer that when executed by the computer implement a method for implementing an engineering change order in an integrated circuit design by windows, the method comprising steps of:

(a) receiving as input an integrated circuit design;

(b) receiving as input an engineering change order to the integrated circuit design;

(c) creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design;

(d) performing an incremental routing only for each net in the integrated circuit design that is enclosed by the window;

(e) replacing an area in a copy of the integrated circuit design that is bounded by coordinates of the window with results of the incremental routing to generate a revised integrated circuit design; and

(f) generating as output the revised integrated circuit design.

6. The computer readable storage medium of claim **5** wherein the method further comprises a step of calculating a net delay only for each net in the integrated circuit design that is enclosed by the window.

7. The computer readable storage meduim of claim **5** wherein the method further comprises a step of performing a design rule check only for each net in the integrated circuit design that is enclosed by the window.

8. The computer readable storage medium of claim **5** wherein the method further comprises a step of performing a parasitic extraction only for each net in the integrated circuit design that is enclosed by the window.

* * * * *

EXHIBIT B

BELL SEMICONDUCTOR LLC'S ANALYSIS OF INFRINGEMENT

U.S. Patent No. 7,231,626

Claims 1-4

Bell Semiconductor, LLC ("Bell Semiconductor") provides evidence of infringement of exemplary claims 1-4 of U.S. Patent No. 7,231,626 ("the '626 patent") by the AC7-M128-30 Altra Max produced by Ampere Computing, LLC ("Ampere"). In support thereof, Bell Semiconductor provides the following claim charts.

"Accused Products" as used herein refers to at least devices produced or sold by Ampere that are or include semiconductor integrated circuit devices made using a design tool, that are made, produced, and/or processed by a design tool, such as a Cadence Design Systems, Inc.("Cadence"), Synopsys, Inc. ("Synopsys"), and/or Siemens Digital Industries Software (formerly Mentor Graphics) ("Siemens") tool,[1] by implementing an engineering change order through a window that is less than the entire area of the integrated circuit design. On information and belief, these design tools all function similarly with respect to the functionality described herein. For simplicity, the Cadence tool will be the primary tool cited herein to illustrate infringement of the claimed methods. These claim charts demonstrate infringement by comparing each element of the asserted claims to corresponding components, aspects, and/or features of the Accused Products. These claim charts are not intended to constitute an expert report on infringement. These claim charts include information provided by way of example, and not by way of limitation.

The analysis set forth below is based only upon information from publicly available resources regarding the Accused Products, as Ampere and relevant third parties have not yet provided any non-public information. An analysis of non-public technical documentation may assist in further identifying all infringing features and functionality. Accordingly, Bell Semiconductor reserves the right to supplement this infringement analysis once such information is made available to Bell Semiconductor. Furthermore, Bell Semiconductor reserves the right to revise this infringement analysis, as appropriate, upon issuance of a court order construing any terms recited in the asserted claims or as other circumstances so merit.

Bell Semiconductor contends that each element of each claim asserted herein is literally met, and would also be met under the doctrine of equivalents, as there are no substantial differences between the Accused Products and the elements of the patent claims in function, way, and result. Ampere directly infringes the asserted claims of the '626 patent by performing each of the limitations. If Ampere attempts to argue that there is no literal infringement and/or if Ampere attempts to draw any distinction between the claimed functionality and the Accused Products, then Bell Semiconductor reserves the right to rebut the alleged distinction as a matter of literal infringement and/or as to whether any such distinction is substantial under the doctrine of equivalents.

Unless otherwise noted, the cited evidence applies across each of Ampere's products that were made, produced, or processed from a circuit design using windows, including but not limited to the AC7-M128-30 Altra Max. Bell Semiconductor reserves the right to amend this infringement analysis based on other products made, produced, or processed in the same or similar manner to that identified herein.

---

[1] Ampere is a customer of at least Cadence, as demonstrated here: https://www.youtube.com/watch?v=CUB2Xn6Erd0.

Ampere is the producer and/or seller of the referenced above, as demonstrated by the following package images for the AC7-M128-30 Altra Max.

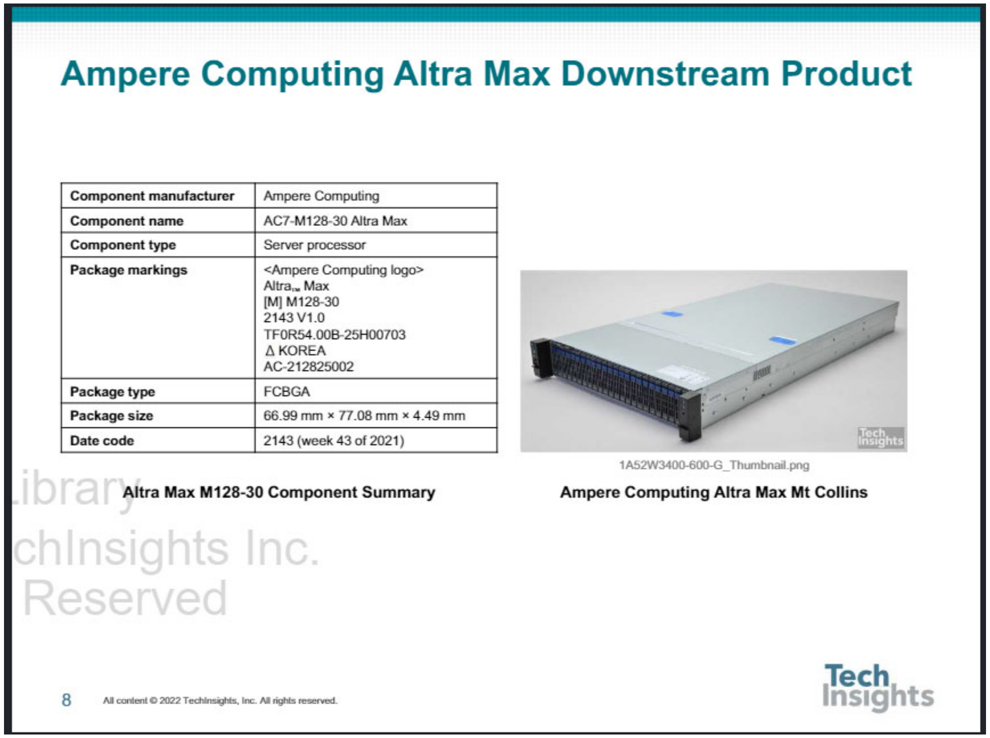

Ampere Computing Altra Max Downstream Product

Component manufacturer	Ampere Computing
Component name	AC7-M128-30 Altra Max
Component type	Server processor
Package markings	<Ampere Computing logo> Altra™ Max [M] M128-30 2143 V1.0 TF0R54.00B-25H00703 Δ KOREA AC-212825002
Package type	FCBGA
Package size	66.99 mm × 77.08 mm × 4.49 mm
Date code	2143 (week 43 of 2021)

Altra Max M128-30 Component Summary

1A52W3400-600-G_Thumbnail.png

Ampere Computing Altra Max Mt Collins

Claim 1	Accused Product
1. A method comprising steps of:	To the extent the preamble is limiting, the Accused Products are produced by performing the method steps outlined in the remaining claim elements. 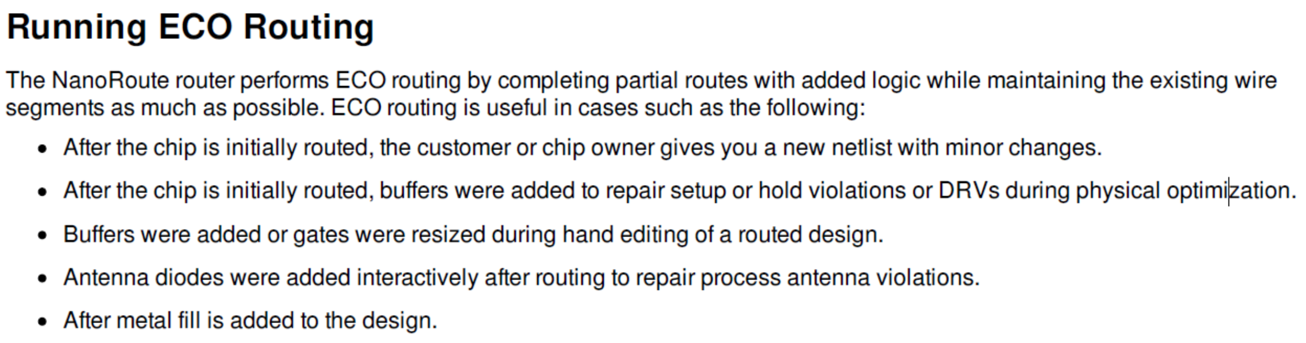 **Running ECO Routing** The NanoRoute router performs ECO routing by completing partial routes with added logic while maintaining the existing wire segments as much as possible. ECO routing is useful in cases such as the following: • After the chip is initially routed, the customer or chip owner gives you a new netlist with minor changes. • After the chip is initially routed, buffers were added to repair setup or hold violations or DRVs during physical optimization. • Buffers were added or gates were resized during hand editing of a routed design. • Antenna diodes were added interactively after routing to repair process antenna violations. • After metal fill is added to the design.

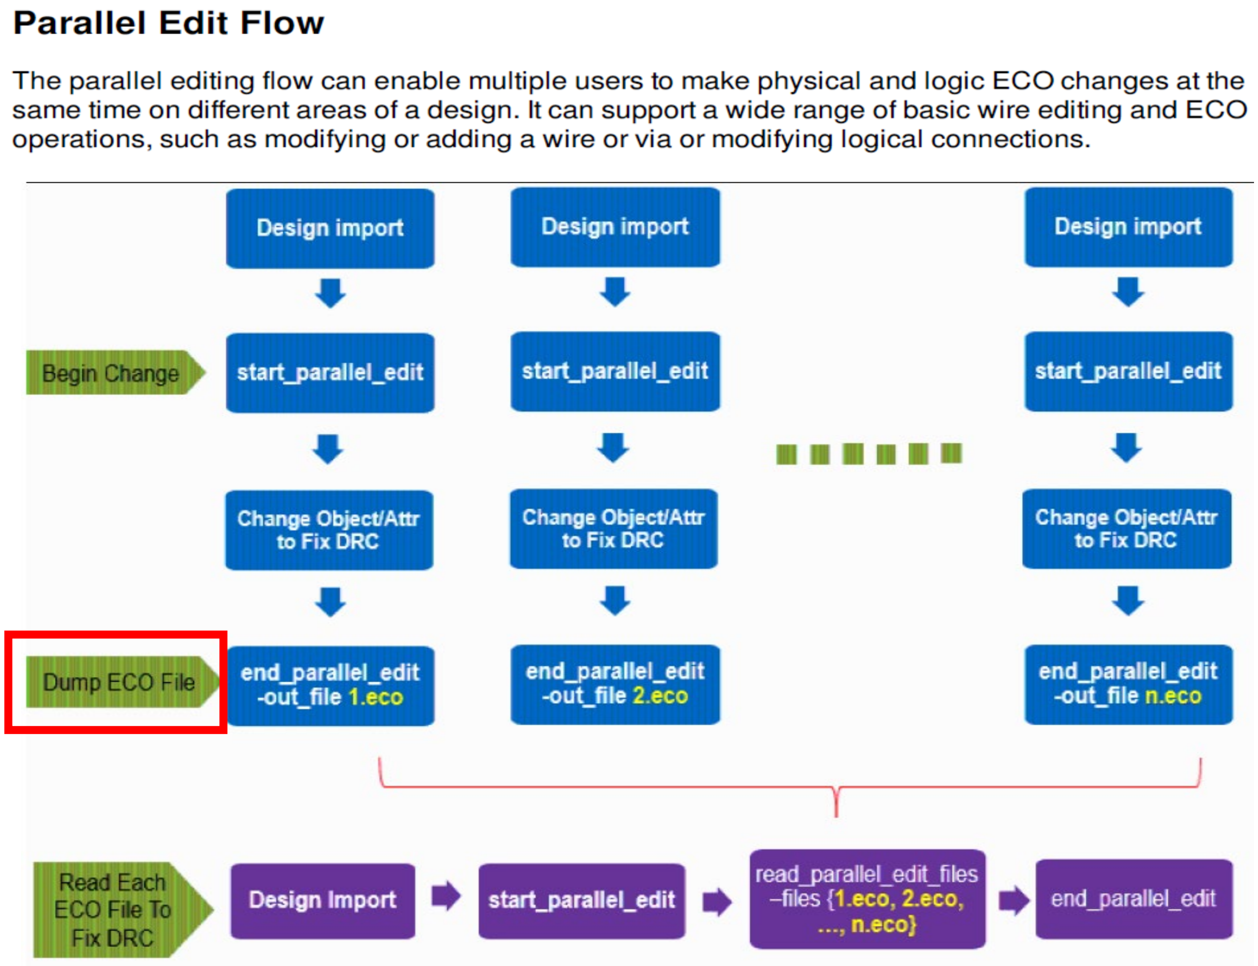

Parallel Edit Flow

The parallel editing flow can enable multiple users to make physical and logic ECO changes at the same time on different areas of a design. It can support a wide range of basic wire editing and ECO operations, such as modifying or adding a wire or via or modifying logical connections.

See Innovus User Guide product version 20.10, March 2020, pages 682 and 1583.

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from that circuit design created using one or more of the above-identified and described design tools that use windows. This is explained by semiconductor expert Lloyd Linder (“Linder”) in Exhibit C cited herein.

<table>
<tr>
<td>(a) receiving as input an integrated circuit design;</td>
<td>
The Accused Products are made, produced, or processed from a circuit design that is created by receiving as input an integrated circuit design.

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from that circuit design that is created using one or more of the above-identified and described design tools such that it receives a design through its import and/or loading procedures. <i>See</i> Ex. C at ¶¶ 44-47.

The steps in the parallel edit flow are detailed below:
1. Import or restore the design in which you want to make parallel edits along with other users.
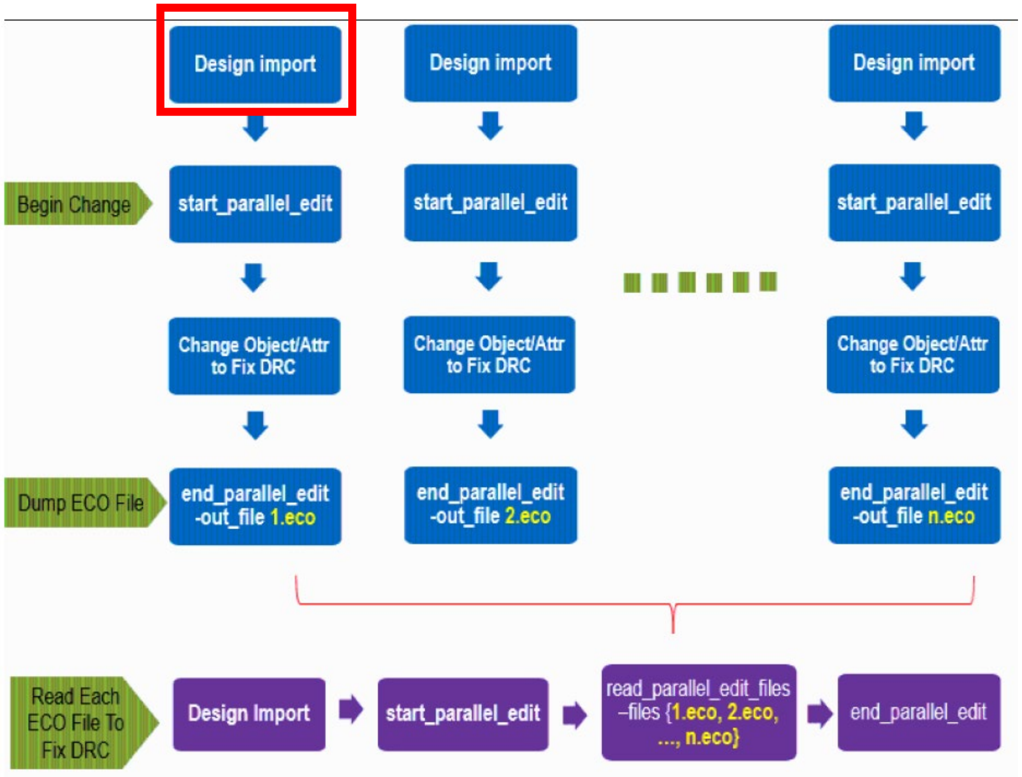

</td>
</tr>
</table>

	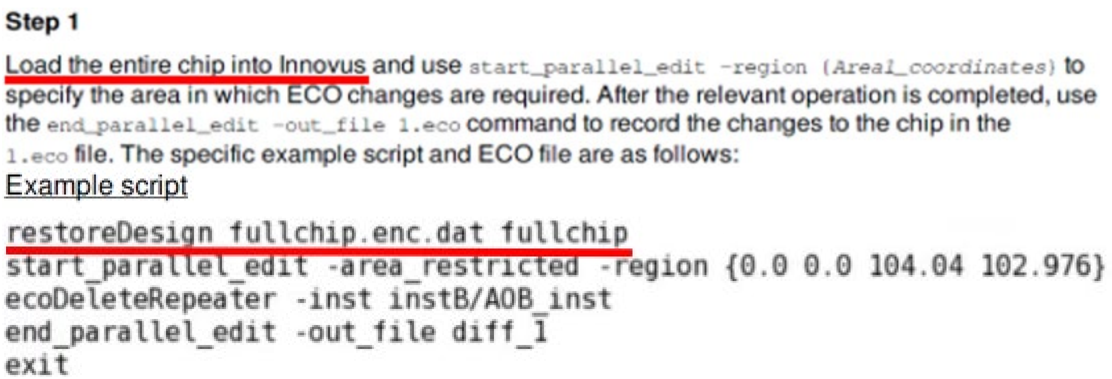
	Step 1
	Load the entire chip into Innovus and use `start_parallel_edit -region {Area1_coordinates}` to specify the area in which ECO changes are required. After the relevant operation is completed, use the `end_parallel_edit -out_file 1.eco` command to record the changes to the chip in the `1.eco` file. The specific example script and ECO file are as follows:
	Example script
	`restoreDesign fullchip.enc.dat fullchip`
	`start_parallel_edit -area_restricted -region {0.0 0.0 104.04 102.976}`
	`ecoDeleteRepeater -inst instB/AOB_inst`
	`end_parallel_edit -out_file diff_1`
	`exit`
	See Innovus User Guide product version 20.10, March 2020, pages 1583 and 1590.

(b) receiving as input an engineering change order to the integrated circuit design;	The Accused Products are made, produced, or processed from a circuit design that is created by receiving as input an engineering change order to the integrated circuit design. 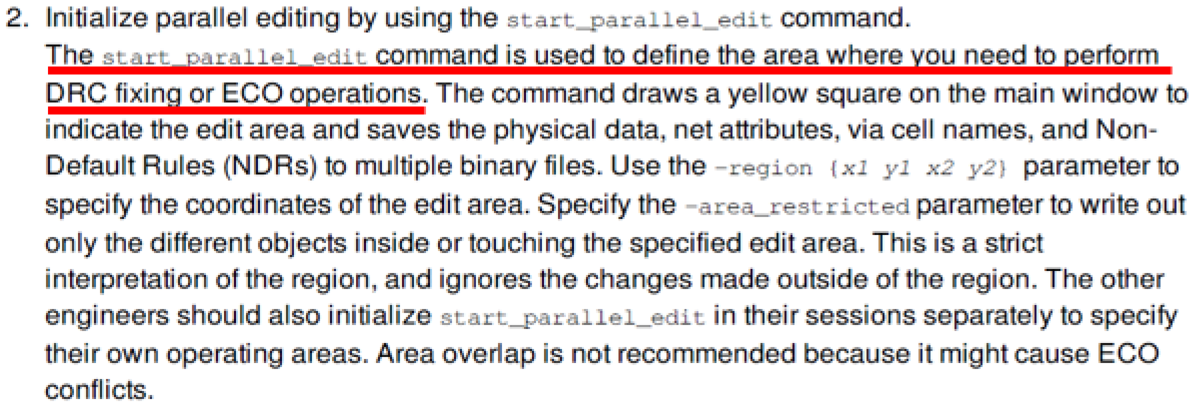 2. Initialize parallel editing by using the `start_parallel_edit` command. The `start_parallel_edit` command is used to define the area where you need to perform DRC fixing or ECO operations. The command draws a yellow square on the main window to indicate the edit area and saves the physical data, net attributes, via cell names, and Non-Default Rules (NDRs) to multiple binary files. Use the `-region {x1 y1 x2 y2}` parameter to specify the coordinates of the edit area. Specify the `-area_restricted` parameter to write out only the different objects inside or touching the specified edit area. This is a strict interpretation of the region, and ignores the changes made outside of the region. The other engineers should also initialize `start_parallel_edit` in their sessions separately to specify their own operating areas. Area overlap is not recommended because it might cause ECO conflicts.

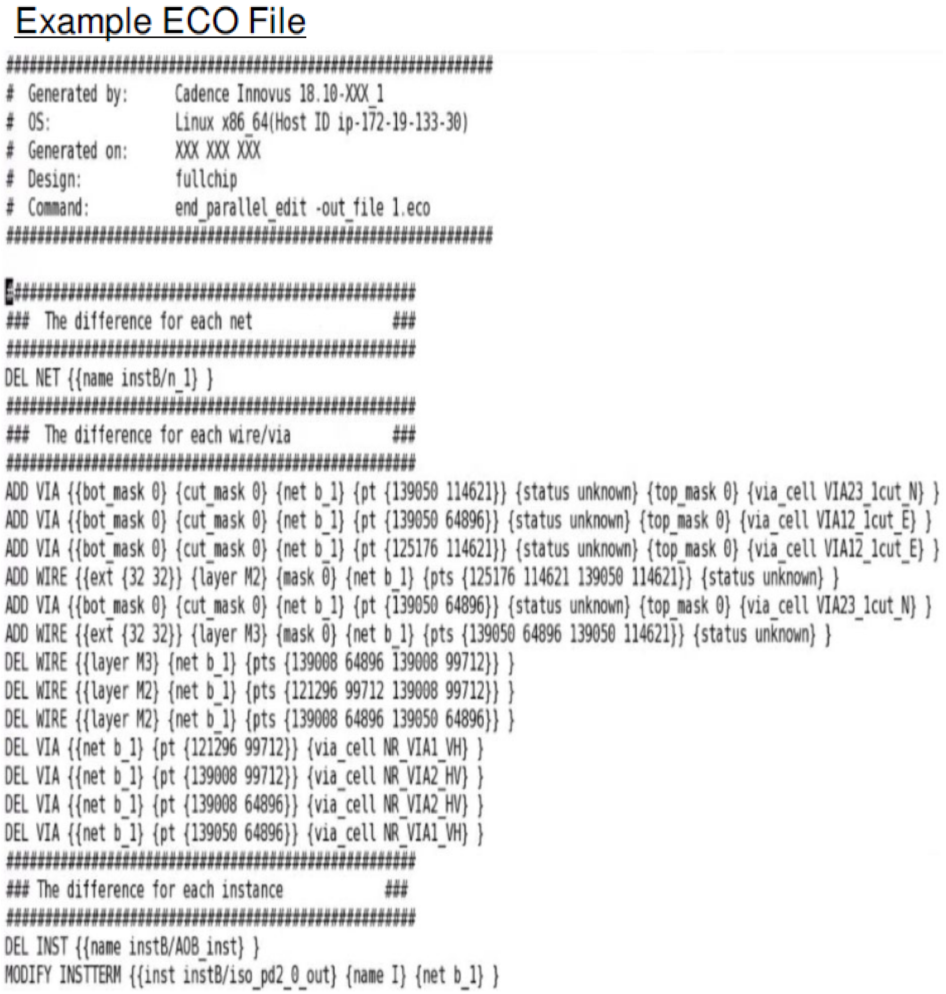

Example ECO File

```
############################################################
#  Generated by:      Cadence Innovus 18.10-XXX_1
#  OS:                Linux x86_64(Host ID ip-172-19-133-30)
#  Generated on:      XXX XXX XXX
#  Design:            fullchip
#  Command:           end_parallel_edit -out_file 1.eco
############################################################

##################################################
###  The difference for each net               ###
##################################################
DEL NET {{name instB/n_1} }
##################################################
###  The difference for each wire/via          ###
##################################################
ADD VIA {{bot_mask 0} {cut_mask 0} {net b_1} {pt {139050 114621}} {status unknown} {top_mask 0} {via_cell VIA23_1cut_N} }
ADD VIA {{bot_mask 0} {cut_mask 0} {net b_1} {pt {139050 64896}} {status unknown} {top_mask 0} {via_cell VIA12_1cut_E} }
ADD VIA {{bot_mask 0} {cut_mask 0} {net b_1} {pt {125176 114621}} {status unknown} {top_mask 0} {via_cell VIA12_1cut_E} }
ADD WIRE {{ext {32 32}} {layer M2} {mask 0} {net b_1} {pts {125176 114621 139050 114621}} {status unknown} }
ADD VIA {{bot_mask 0} {cut_mask 0} {net b_1} {pt {139050 64896}} {status unknown} {top_mask 0} {via_cell VIA23_1cut_N} }
ADD WIRE {{ext {32 32}} {layer M3} {mask 0} {net b_1} {pts {139050 64896 139050 114621}} {status unknown} }
DEL WIRE {{layer M3} {net b_1} {pts {139008 64896 139008 99712}} }
DEL WIRE {{layer M2} {net b_1} {pts {121296 99712 139008 99712}} }
DEL WIRE {{layer M2} {net b_1} {pts {139008 64896 139050 64896}} }
DEL VIA {{net b_1} {pt {121296 99712}} {via_cell NR_VIA1_VH} }
DEL VIA {{net b_1} {pt {139008 99712}} {via_cell NR_VIA2_HV} }
DEL VIA {{net b_1} {pt {139008 64896}} {via_cell NR_VIA2_HV} }
DEL VIA {{net b_1} {pt {139050 64896}} {via_cell NR_VIA1_VH} }
##################################################
### The difference for each instance           ###
##################################################
DEL INST {{name instB/AOB_inst} }
MODIFY INSTTERM {{inst instB/iso_pd2_0_out} {name I} {net b_1} }
```

See Innovus User Guide product version 20.10, March 2020, pages 1584 and 1591.

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created by receiving as input an engineering change order ("ECO") to the integrated circuit design. *See* Ex. C at ¶¶ 44-47.

(c) creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design;	The Accused Products are made, produced, or processed from a circuit design that is created by creating at least one window in the integrated circuit design that encloses a change to the integrated circuit design introduced by the engineering change order wherein the window is bounded by coordinates that define an area that is less than an entire area of the integrated circuit design. 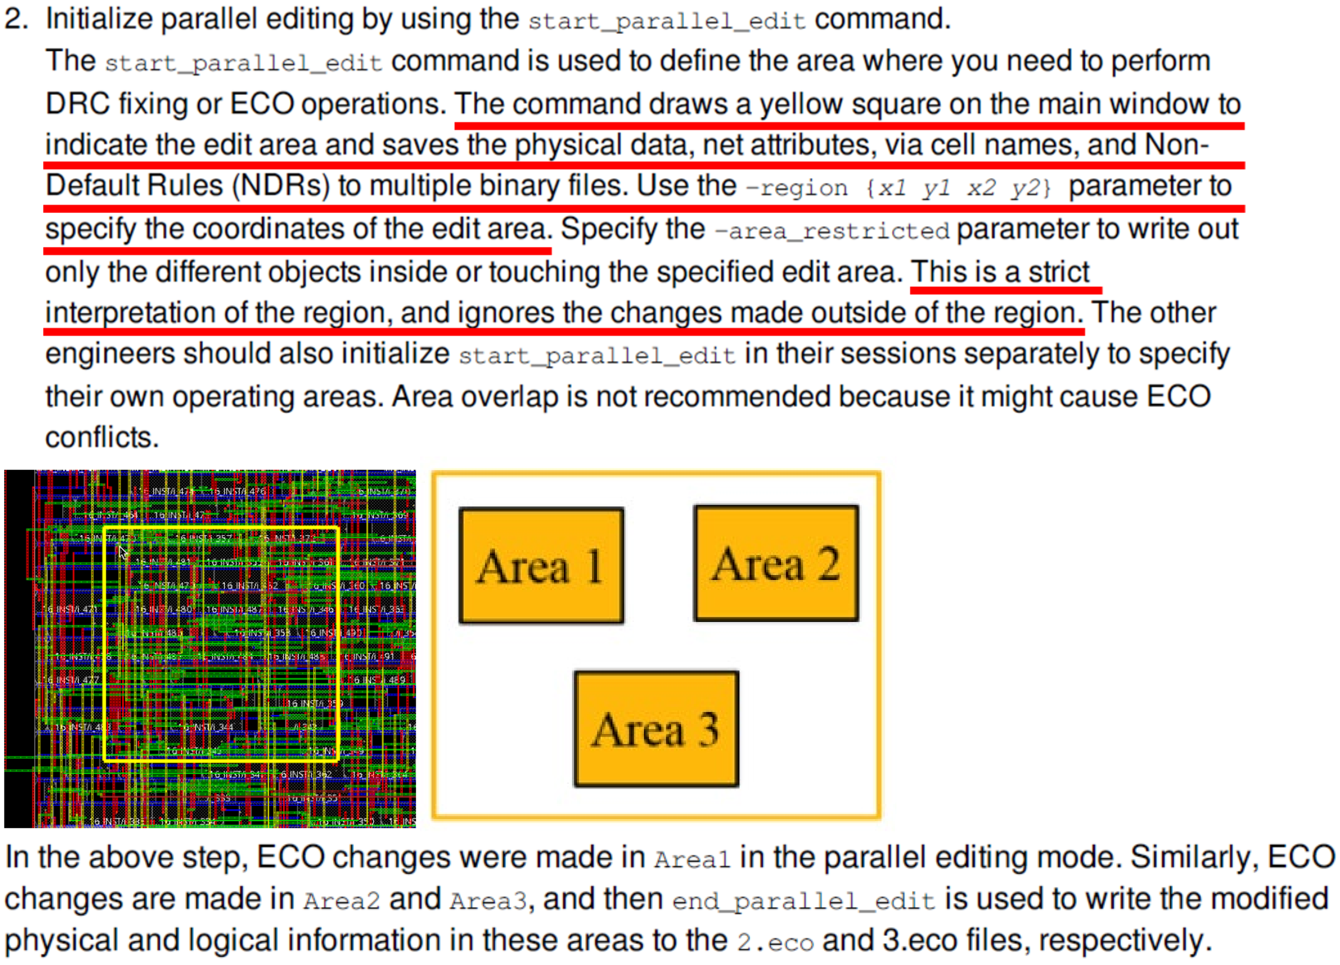 2. Initialize parallel editing by using the `start_parallel_edit` command. The `start_parallel_edit` command is used to define the area where you need to perform DRC fixing or ECO operations. The command draws a yellow square on the main window to indicate the edit area and saves the physical data, net attributes, via cell names, and Non-Default Rules (NDRs) to multiple binary files. Use the `-region {x1 y1 x2 y2}` parameter to specify the coordinates of the edit area. Specify the `-area_restricted` parameter to write out only the different objects inside or touching the specified edit area. This is a strict interpretation of the region, and ignores the changes made outside of the region. The other engineers should also initialize `start_parallel_edit` in their sessions separately to specify their own operating areas. Area overlap is not recommended because it might cause ECO conflicts. In the above step, ECO changes were made in `Area1` in the parallel editing mode. Similarly, ECO changes are made in `Area2` and `Area3`, and then `end_parallel_edit` is used to write the modified physical and logical information in these areas to the `2.eco` and 3.eco files, respectively. *See Innovus User Guide product version 20.10, March 2020, pages 1584 and 1590-91.*

	For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from that circuit design that is created by a using a parallel edit command that draws a square to enclose the edit area for the ECO. The edit area is less than an entire area of the integrated circuit design. *See* Ex. C at ¶¶ 44-47.
(d) performing an incremental routing of the integrated circuit design only for each net in the integrated circuit design that is enclosed by the window;	The Accused Products are made, produced, or processed from a circuit design that is created by performing an incremental routing of the integrated circuit design only for each net in the integrated circuit design that is enclosed by the window.

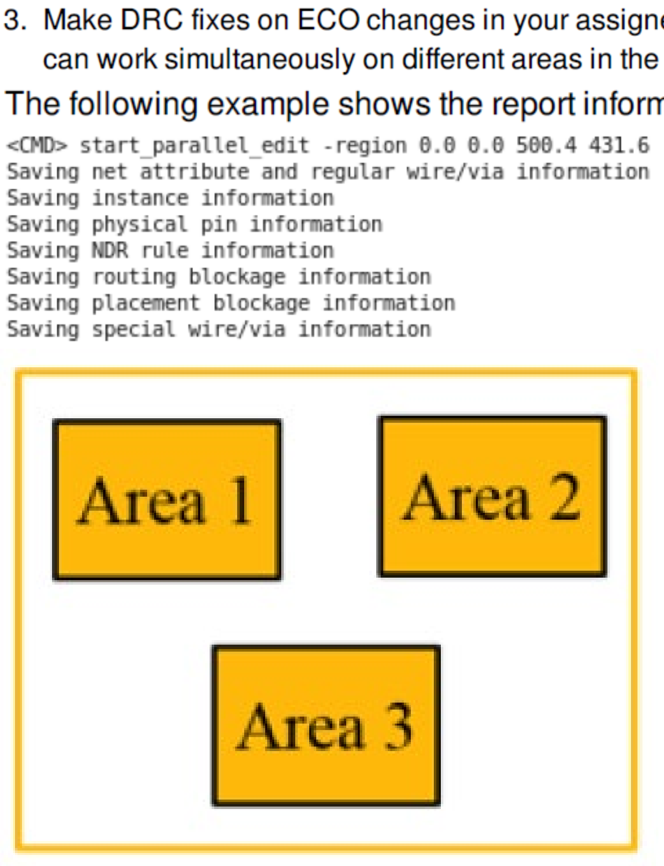

3. Make DRC fixes on ECO changes in your assigned area in the design. Multiple engineers can work simultaneously on different areas in the design in the parallel edit mode.

The following example shows the report information generated by `start_parallel_edit`.

```
<CMD> start_parallel_edit -region 0.0 0.0 500.4 431.6
Saving net attribute and regular wire/via information
Saving instance information
Saving physical pin information
Saving NDR rule information
Saving routing blockage information
Saving placement blockage information
Saving special wire/via information
```

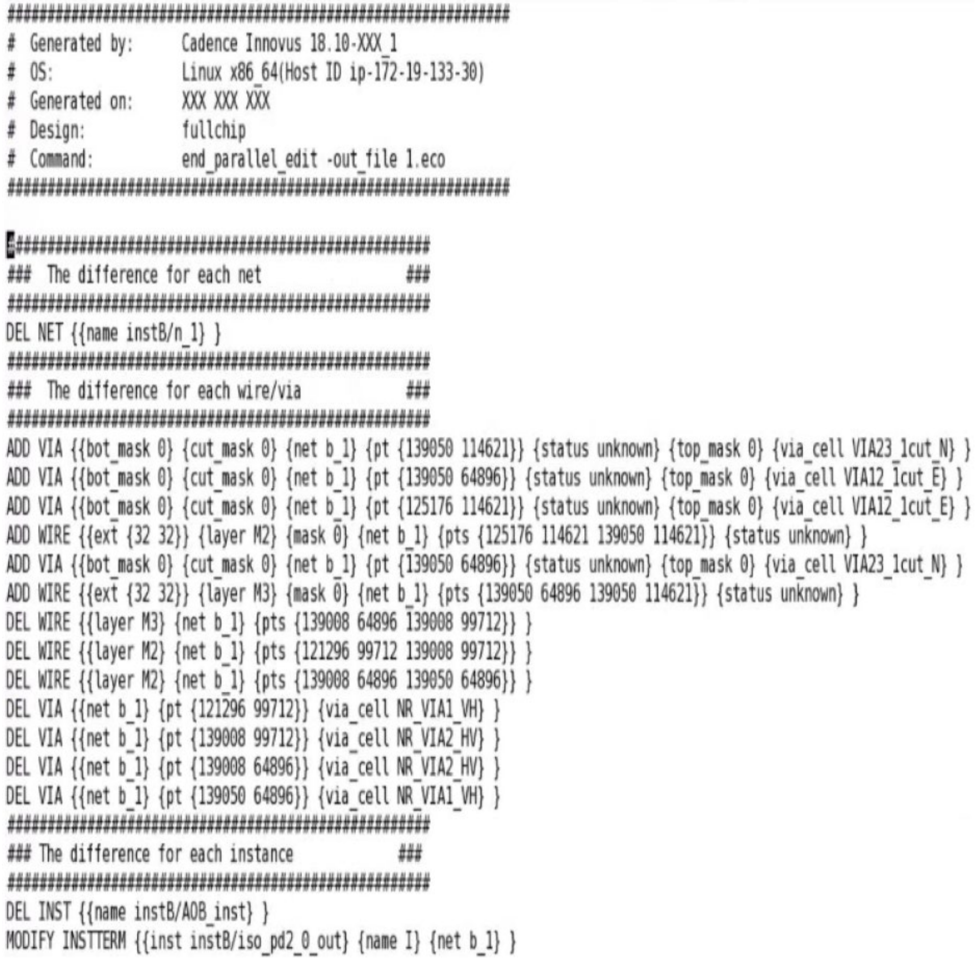

```
##############################################################
#  Generated by:       Cadence Innovus 18.10-XXX_1
#  OS:                 Linux x86_64(Host ID ip-172-19-133-30)
#  Generated on:       XXX XXX XXX
#  Design:             fullchip
#  Command:            end_parallel_edit -out_file 1.eco
##############################################################

####################################################
###  The difference for each net                 ###
####################################################
DEL NET {{name instB/n_1} }
####################################################
###  The difference for each wire/via            ###
####################################################
ADD VIA {{bot_mask 0} {cut_mask 0} {net b_1} {pt {139050 114621}} {status unknown} {top_mask 0} {via_cell VIA23_1cut_N} }
ADD VIA {{bot_mask 0} {cut_mask 0} {net b_1} {pt {139050 64896}} {status unknown} {top_mask 0} {via_cell VIA12_1cut_E} }
ADD VIA {{bot_mask 0} {cut_mask 0} {net b_1} {pt {125176 114621}} {status unknown} {top_mask 0} {via_cell VIA12_1cut_E} }
ADD WIRE {{ext {32 32}} {layer M2} {mask 0} {net b_1} {pts {125176 114621 139050 114621}} {status unknown} }
ADD VIA {{bot_mask 0} {cut_mask 0} {net b_1} {pt {139050 64896}} {status unknown} {top_mask 0} {via_cell VIA23_1cut_N} }
ADD WIRE {{ext {32 32}} {layer M3} {mask 0} {net b_1} {pts {139050 64896 139050 114621}} {status unknown} }
DEL WIRE {{layer M3} {net b_1} {pts {139008 64896 139008 99712}} }
DEL WIRE {{layer M2} {net b_1} {pts {121296 99712 139008 99712}} }
DEL WIRE {{layer M2} {net b_1} {pts {139008 64896 139050 64896}} }
DEL VIA {{net b_1} {pt {121296 99712}} {via_cell NR_VIA1_VH} }
DEL VIA {{net b_1} {pt {139008 99712}} {via_cell NR_VIA2_HV} }
DEL VIA {{net b_1} {pt {139008 64896}} {via_cell NR_VIA2_HV} }
DEL VIA {{net b_1} {pt {139050 64896}} {via_cell NR_VIA1_VH} }
####################################################
### The difference for each instance            ###
####################################################
DEL INST {{name instB/AOB_inst} }
MODIFY INSTTERM {{inst instB/iso_pd2_0_out} {name I} {net b_1} }
```

See Innovus User Guide product version 20.10, March 2020, pages 1585 and 1590-91.

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from that circuit design that is created by performing an incremental routing of the integrated circuit design only for each net in the integrated circuit design that is enclosed by the square of the edit area. *See* Ex. C at ¶¶ 44-47.

(e) replacing an area in a copy of the integrated circuit design that is bounded by the coordinates of the window with results of the incremental routing to generate a revised integrated circuit design; and	The Accused Products are made, produced, or processed from a circuit design that is created by replacing an area in a copy of the integrated circuit design that is bounded by the coordinates of the window with results of the incremental routing to generate a revised integrated circuit design.  **Step 2** Load the entire chip first and then read the eco files generated in the previous step by using the `read_parallel_edit_files` command. All the changes in different areas will be reflected in the chip. *See Innovus User Guide product version 20.10*, March 2020, pages 1585 and 1591. For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from that circuit design that is created by replacing an area in a copy of the integrated circuit design that is bounded by the square of the edit area with results of the incremental routing to generate a revised integrated circuit design. *See* Ex. C at ¶¶ 44-47.

(f) generating as output the revised integrated circuit design.	The Accused Products are made, produced, or processed from a circuit design that is created by generating as output the revised integrated circuit design.
	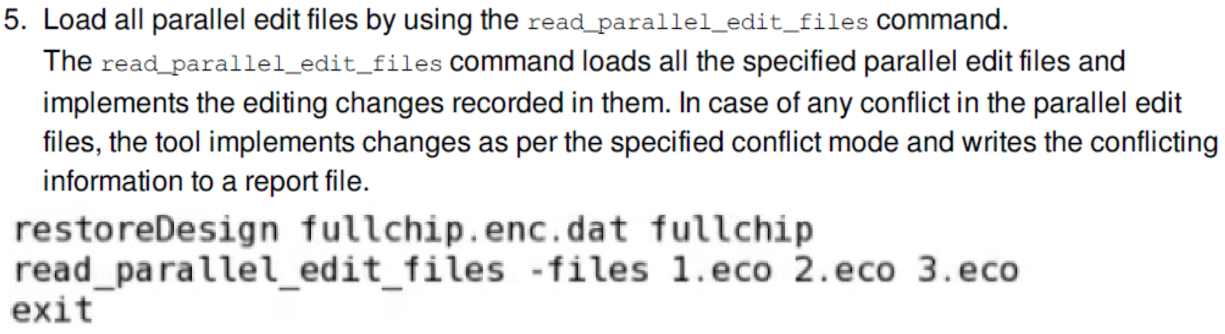
	5. Load all parallel edit files by using the `read_parallel_edit_files` command.
	The `read_parallel_edit_files` command loads all the specified parallel edit files and implements the editing changes recorded in them. In case of any conflict in the parallel edit files, the tool implements changes as per the specified conflict mode and writes the conflicting information to a report file.
	`restoreDesign fullchip.enc.dat fullchip`
	`read_parallel_edit_files -files 1.eco 2.eco 3.eco`
	`exit`
	See Innovus User Guide product version 20.10, March 2020, pages 1585 and 1592.
	For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from that circuit design that is created by implementing the revised integrated circuit design as an output. *See* Ex. C at ¶¶ 44-47.

2. The method of claim 1 further comprising a step of calculating a net delay only for each net in the integrated circuit design that is enclosed by the window.	The Accused Products are made, produced, or processed from a circuit design that is created by the method of claim 1 further comprising a step of calculating a net delay only for each net in the integrated circuit design that is enclosed by the window. 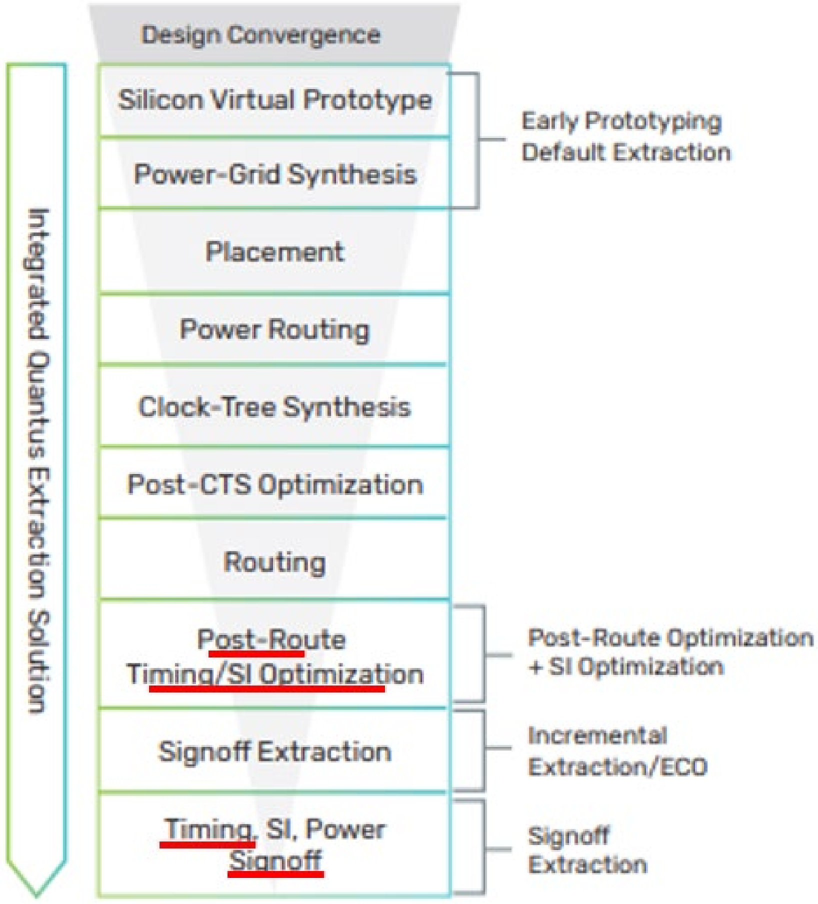 *Figure 1: Enabling in-design in the Innovus environment* https://www.cadence.com/content/dam/cadence-www/global/en_US/documents/tools/digital-design-signoff/quantus-extraction-ds.pdf, page 2.

Post-Route Timing Optimization and Timing Signoff after an ECO both require calculation of net delays for the nets in the IC design enclosed by the incremental window.

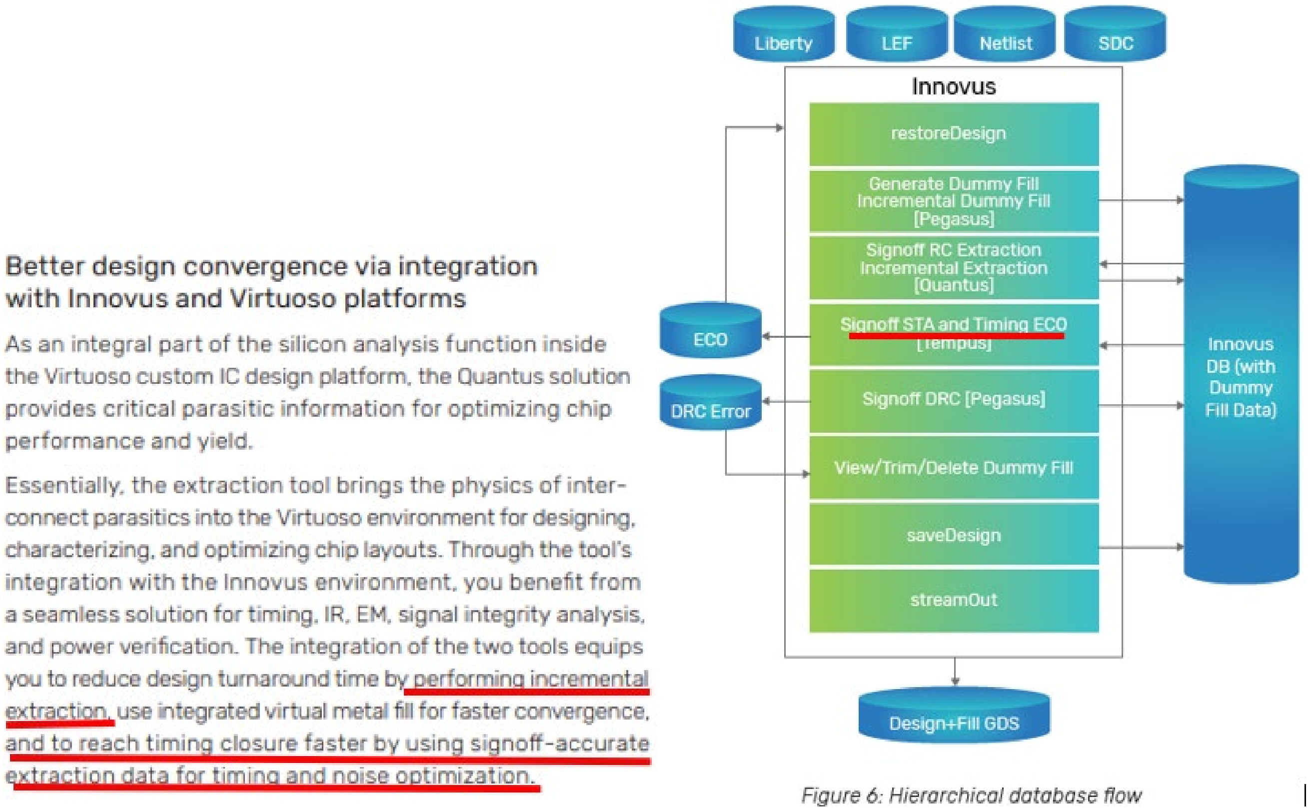

Better design convergence via integration with Innovus and Virtuoso platforms

As an integral part of the silicon analysis function inside the Virtuoso custom IC design platform, the Quantus solution provides critical parasitic information for optimizing chip performance and yield.

Essentially, the extraction tool brings the physics of interconnect parasitics into the Virtuoso environment for designing, characterizing, and optimizing chip layouts. Through the tool's integration with the Innovus environment, you benefit from a seamless solution for timing, IR, EM, signal integrity analysis, and power verification. The integration of the two tools equips you to reduce design turnaround time by performing incremental extraction, use integrated virtual metal fill for faster convergence, and to reach timing closure faster by using signoff-accurate extraction data for timing and noise optimization.

Figure 6: Hierarchical database flow

https://www.cadence.com/content/dam/cadence-www/global/en_US/documents/tools/digital-design-signoff/pegasus-tb.pdf, pages 2, 3, 4

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was further made, produced, or processed from that circuit design for which a net delay was calculated only for the nets enclosed by the window, as indicated by Cadence's signoff of Timing ECO step. *See* Ex. C at ¶¶ 44-47.

3. The method of claim 1 further comprising a step of performing a design rule check only for each net in the integrated circuit design that is enclosed by the window.	The Accused Products are made, produced, or processed from a circuit design that is created by the method of claim 1 further comprising a step of performing a design rule check only for each net in the integrated circuit design that is enclosed by the window. 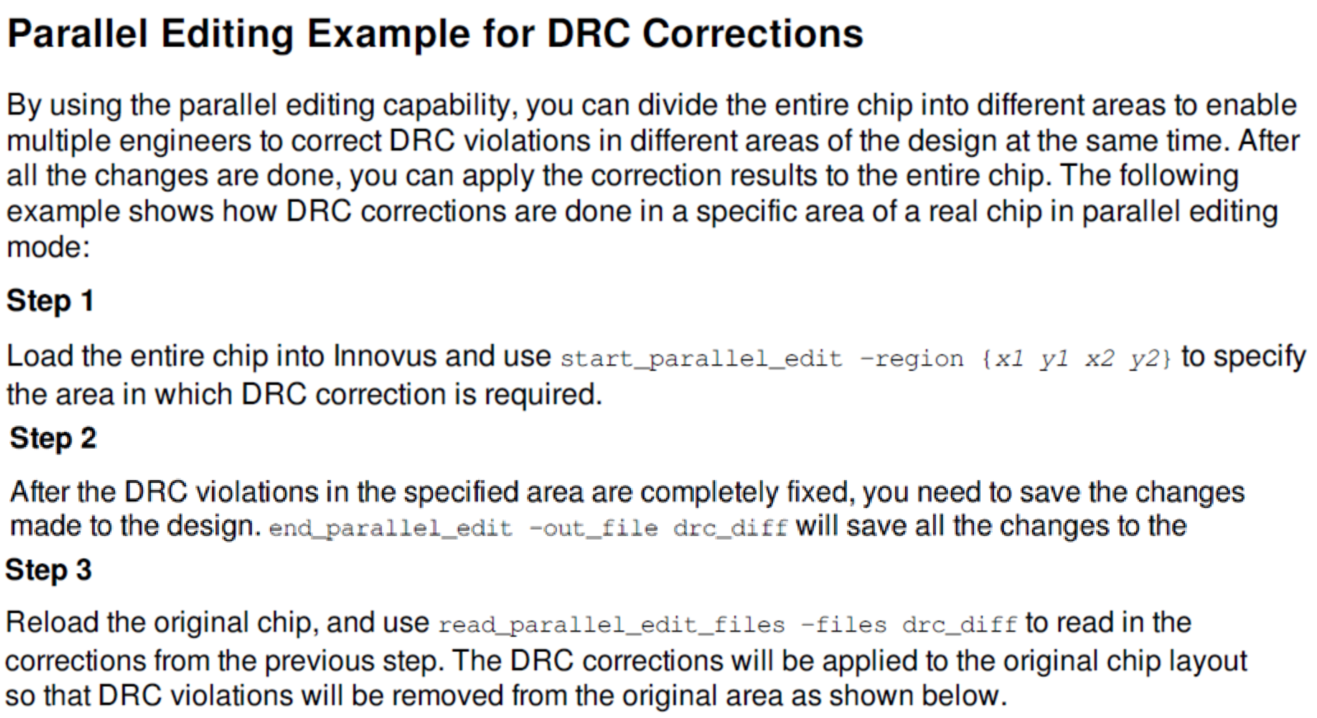 **Parallel Editing Example for DRC Corrections** By using the parallel editing capability, you can divide the entire chip into different areas to enable multiple engineers to correct DRC violations in different areas of the design at the same time. After all the changes are done, you can apply the correction results to the entire chip. The following example shows how DRC corrections are done in a specific area of a real chip in parallel editing mode: **Step 1** Load the entire chip into Innovus and use `start_parallel_edit -region {x1 y1 x2 y2}` to specify the area in which DRC correction is required. **Step 2** After the DRC violations in the specified area are completely fixed, you need to save the changes made to the design. `end_parallel_edit -out_file drc_diff` will save all the changes to the **Step 3** Reload the original chip, and use `read_parallel_edit_files -files drc_diff` to read in the corrections from the previous step. The DRC corrections will be applied to the original chip layout so that DRC violations will be removed from the original area as shown below.

The specified area is highlighted in the GUI at this point. The `verify_drc` command can report all DRC violations in the highlighted area as shown below:

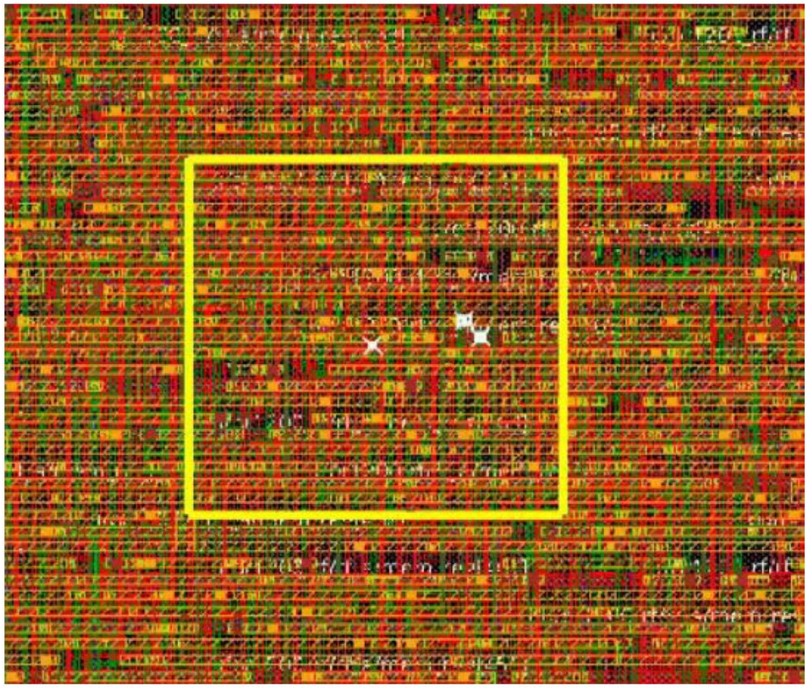

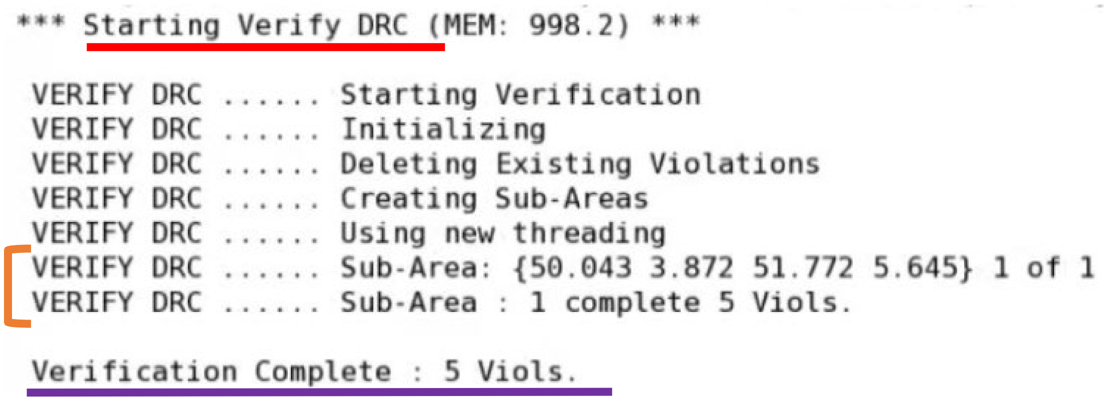

```
*** Starting Verify DRC (MEM: 998.2) ***

 VERIFY DRC ...... Starting Verification
 VERIFY DRC ...... Initializing
 VERIFY DRC ...... Deleting Existing Violations
 VERIFY DRC ...... Creating Sub-Areas
 VERIFY DRC ...... Using new threading
 VERIFY DRC ...... Sub-Area: {50.043 3.872 51.772 5.645} 1 of 1
 VERIFY DRC ...... Sub-Area : 1 complete 5 Viols.

 Verification Complete : 5 Viols.
```

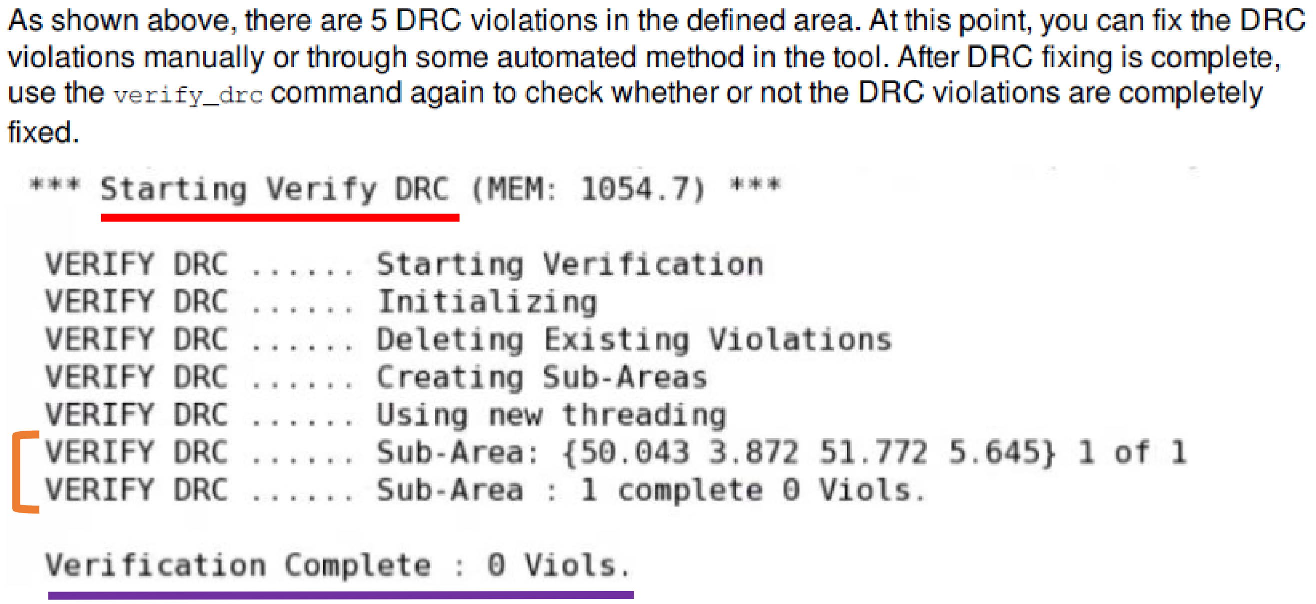

As shown above, there are 5 DRC violations in the defined area. At this point, you can fix the DRC violations manually or through some automated method in the tool. After DRC fixing is complete, use the `verify_drc` command again to check whether or not the DRC violations are completely fixed.

```
 *** Starting Verify DRC (MEM: 1054.7) ***

  VERIFY DRC ...... Starting Verification
  VERIFY DRC ...... Initializing
  VERIFY DRC ...... Deleting Existing Violations
  VERIFY DRC ...... Creating Sub-Areas
  VERIFY DRC ...... Using new threading
  VERIFY DRC ...... Sub-Area: {50.043 3.872 51.772 5.645} 1 of 1
  VERIFY DRC ...... Sub-Area : 1 complete 0 Viols.

  Verification Complete : 0 Viols.
```

See Innovus User Guide product version 20.10, March 2020, pages 1585-88.

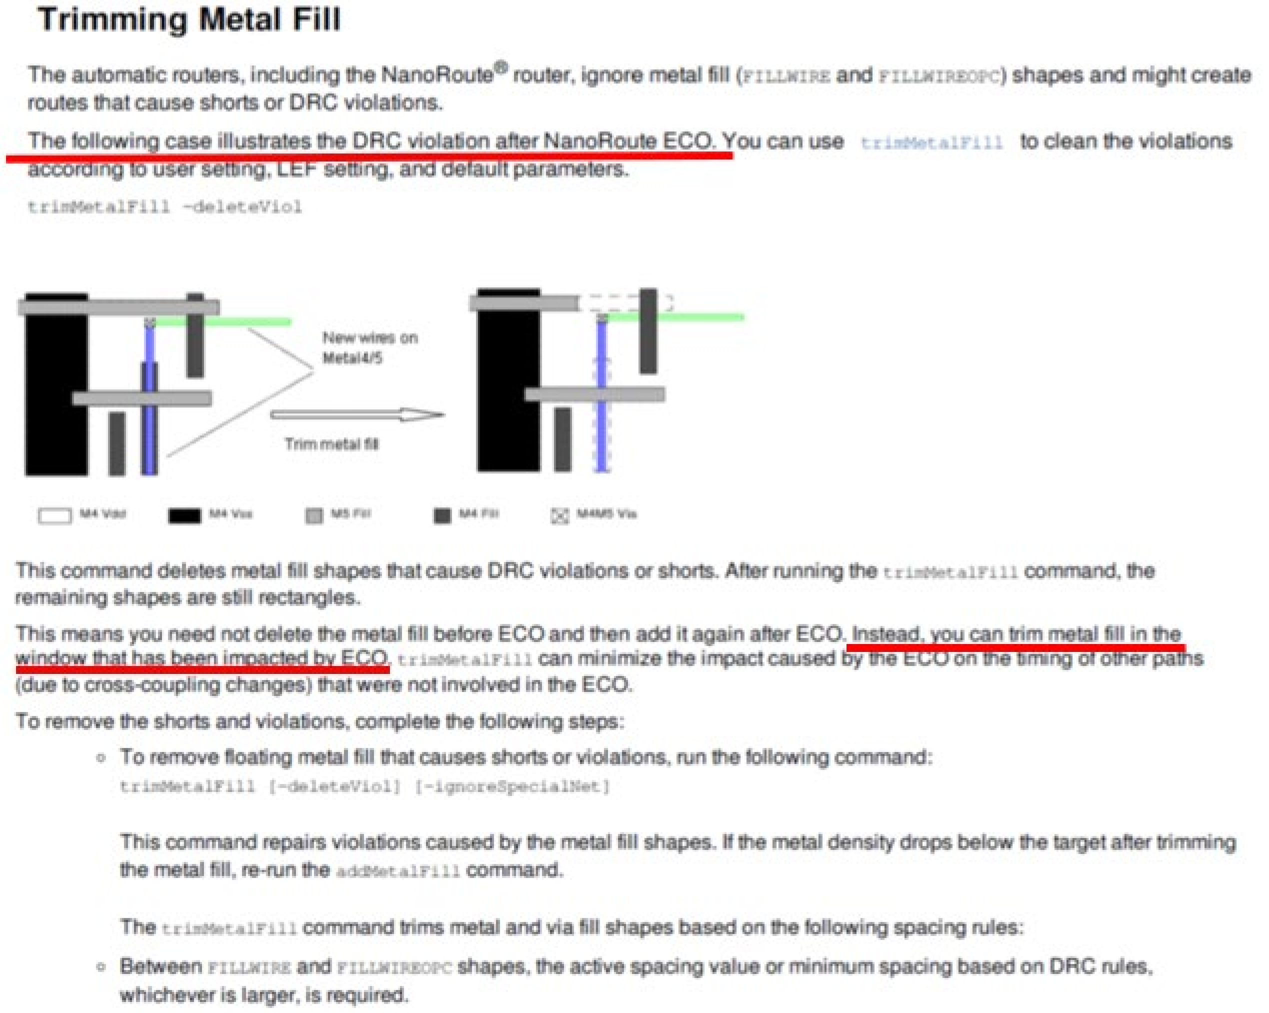

Trimming Metal Fill

The automatic routers, including the NanoRoute® router, ignore metal fill (`FILLWIRE` and `FILLWIREOPC`) shapes and might create routes that cause shorts or DRC violations.

The following case illustrates the DRC violation after NanoRoute ECO. You can use `trimMetalFill` to clean the violations according to user setting, LEF setting, and default parameters.

```
trimMetalFill -deleteViol
```

This command deletes metal fill shapes that cause DRC violations or shorts. After running the `trimMetalFill` command, the remaining shapes are still rectangles.

This means you need not delete the metal fill before ECO and then add it again after ECO. Instead, you can trim metal fill in the window that has been impacted by ECO. `trimMetalFill` can minimize the impact caused by the ECO on the timing of other paths (due to cross-coupling changes) that were not involved in the ECO.

To remove the shorts and violations, complete the following steps:

- To remove floating metal fill that causes shorts or violations, run the following command:

  ```
  trimMetalFill [-deleteViol] [-ignoreSpecialNet]
  ```

  This command repairs violations caused by the metal fill shapes. If the metal density drops below the target after trimming the metal fill, re-run the `addMetalFill` command.

  The `trimMetalFill` command trims metal and via fill shapes based on the following spacing rules:

- Between `FILLWIRE` and `FILLWIREOPC` shapes, the active spacing value or minimum spacing based on DRC rules, whichever is larger, is required.

See Innovus User Guide product version 20.10, March 2020, pages 722-23.

	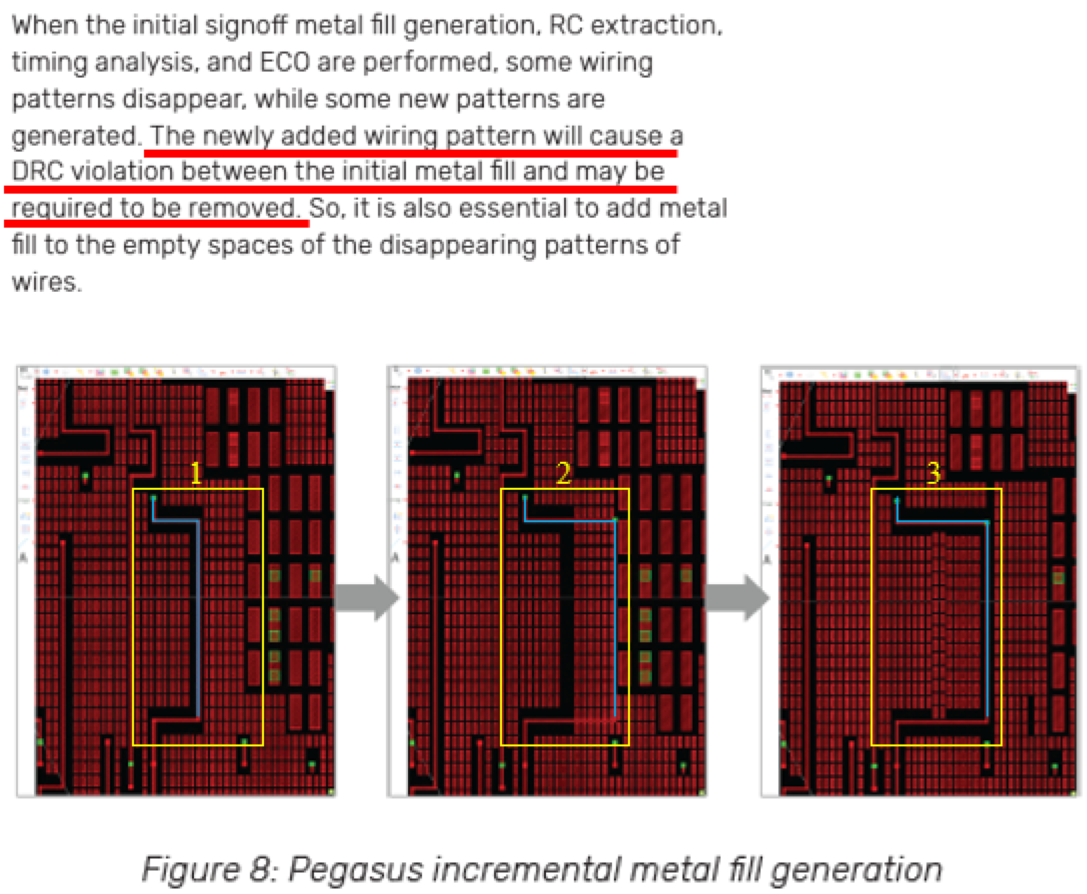 When the initial signoff metal fill generation, RC extraction, timing analysis, and ECO are performed, some wiring patterns disappear, while some new patterns are generated. The newly added wiring pattern will cause a DRC violation between the initial metal fill and may be required to be removed. So, it is also essential to add metal fill to the empty spaces of the disappearing patterns of wires. *Figure 8: Pegasus incremental metal fill generation* *https://www.cadence.com/content/dam/cadence-www/global/en_US/documents/tools/digital-design-signoff/pegasus-tb.pdf, pages 4 and 5*

	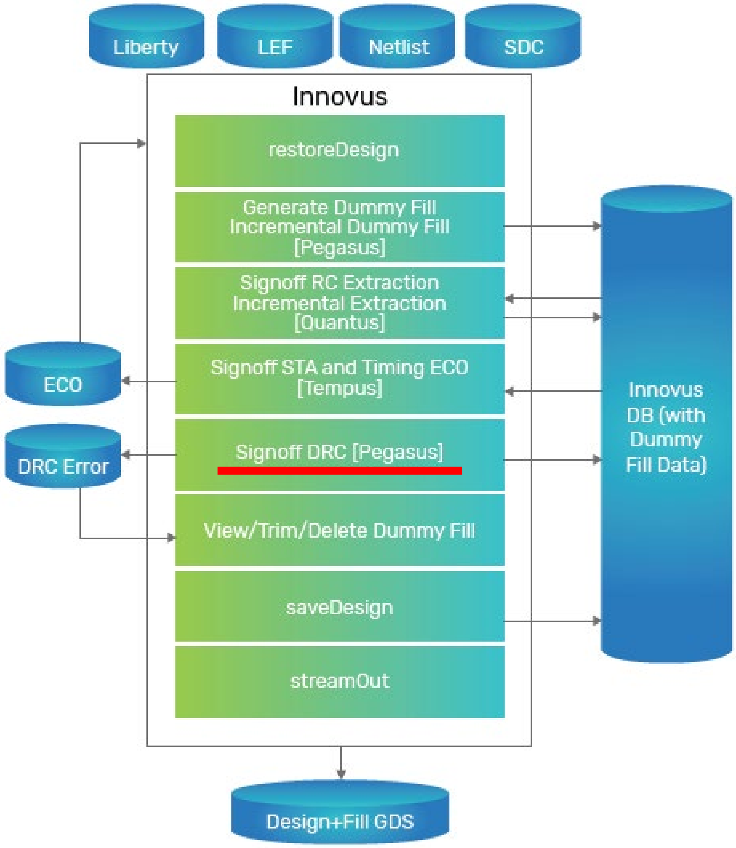 *Figure 6: Hierarchical database flow* *https://www.cadence.com/content/dam/cadence-www/global/en_US/documents/tools/digital-design-signoff/pegasus-tb.pdf, page 4*

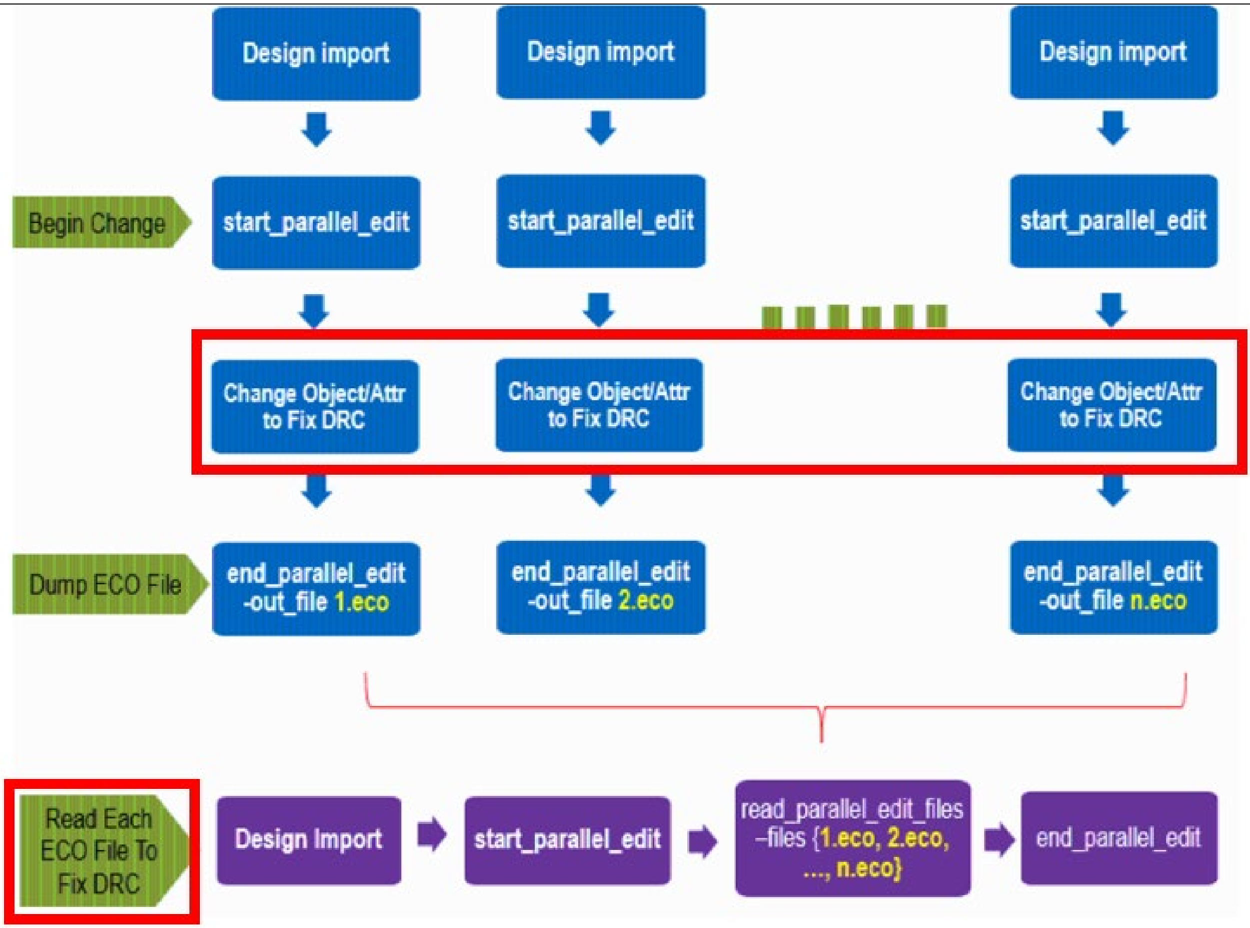

See Innovus User Guide product version 20.10, March 2020, pages 1583 and 1590.

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was further made, produced, or processed from that circuit design that was further created by performing a design rule check separately for each net in the integrated circuit design that is enclosed by the window. *See* Ex. C at ¶¶ 44-47.

4. The method of claim 1 further comprising a step of performing a parasitic extraction only for each net in the integrated circuit design that is enclosed by the window.	The Accused Products are made, produced, or processed from a circuit design that is created by the method of claim 1 further comprising a step of performing a parasitic extraction only for each net in the integrated circuit design that is enclosed by the window. 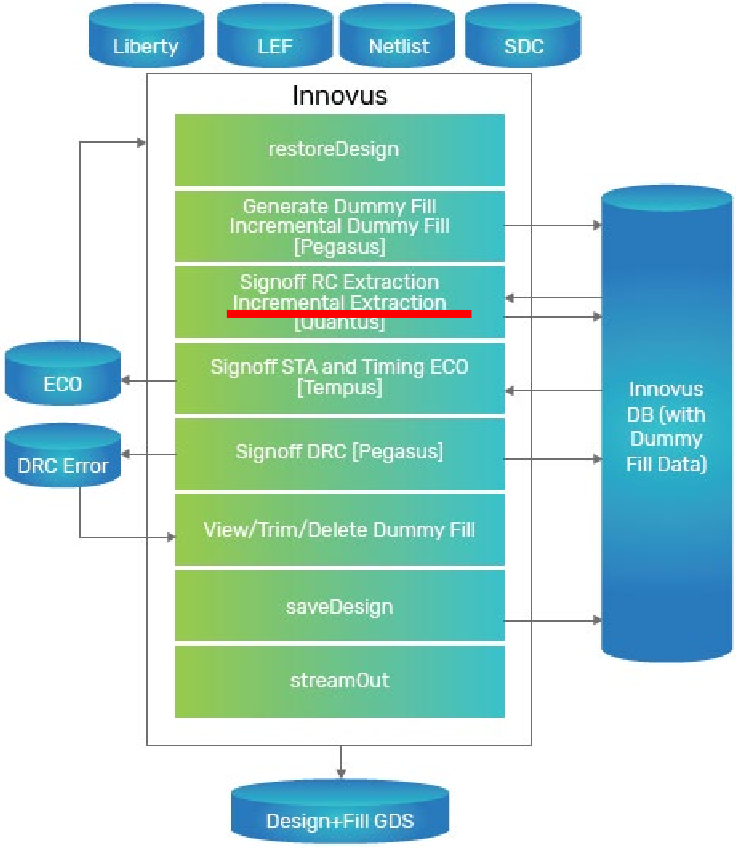 *Figure 6: Hierarchical database flow* *https://www.cadence.com/content/dam/cadence-www/global/en_US/documents/tools/digital-design-signoff/pegasus-tb.pdf, page 4*

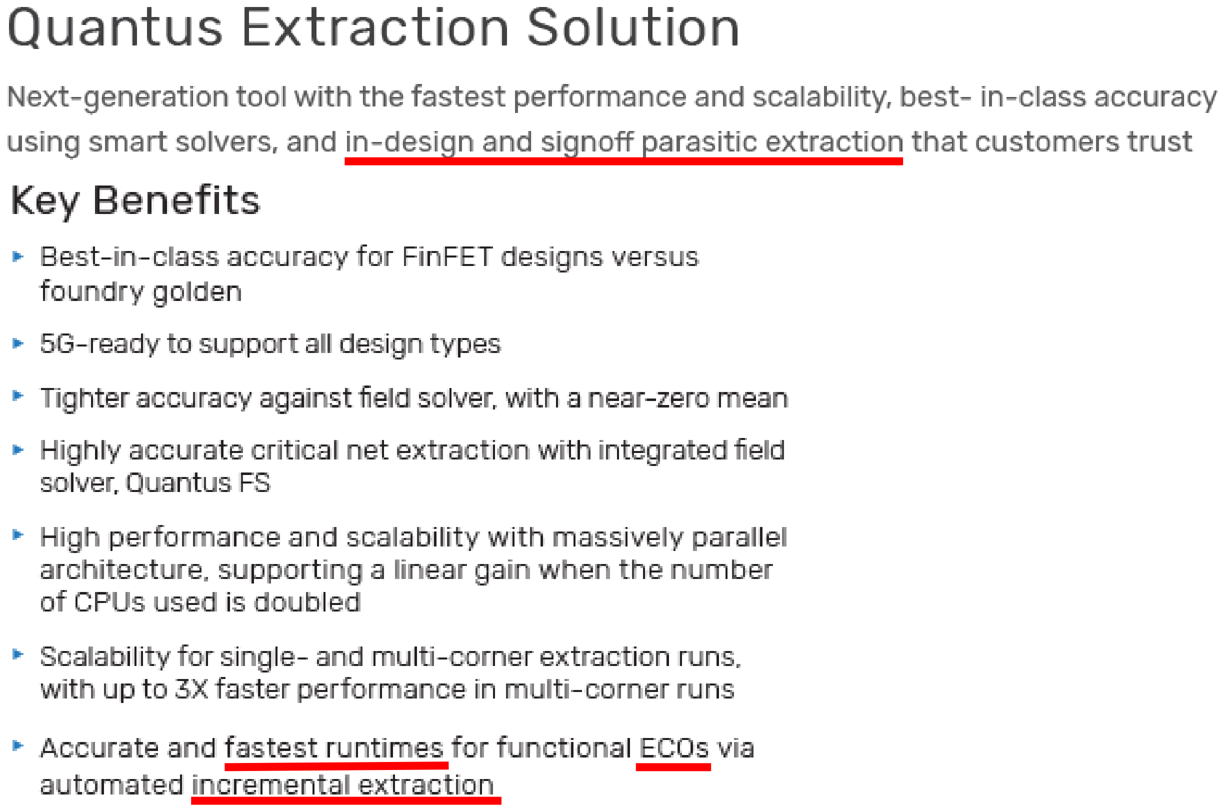

Quantus Extraction Solution

Next-generation tool with the fastest performance and scalability, best- in-class accuracy using smart solvers, and in-design and signoff parasitic extraction that customers trust

Key Benefits

- Best-in-class accuracy for FinFET designs versus foundry golden
- 5G-ready to support all design types
- Tighter accuracy against field solver, with a near-zero mean
- Highly accurate critical net extraction with integrated field solver, Quantus FS
- High performance and scalability with massively parallel architecture, supporting a linear gain when the number of CPUs used is doubled
- Scalability for single- and multi-corner extraction runs, with up to 3X faster performance in multi-corner runs
- Accurate and fastest runtimes for functional ECOs via automated incremental extraction

	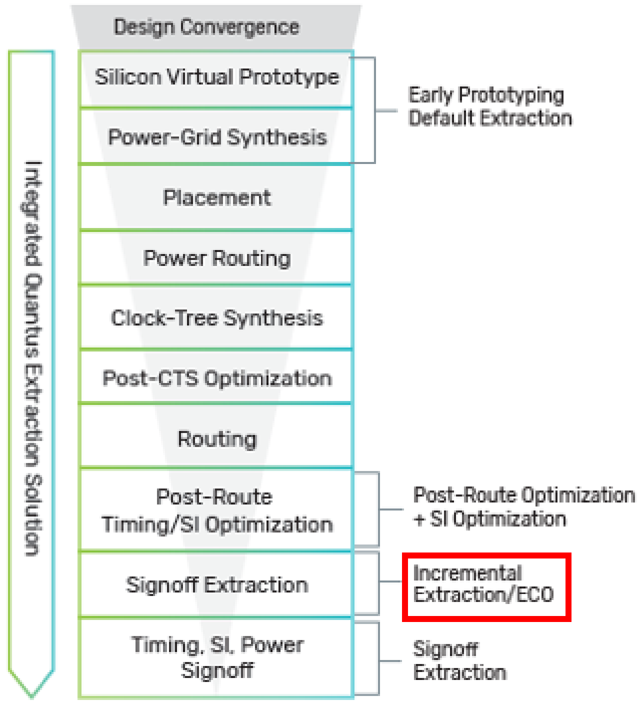 *Figure 1: Enabling in-design in the Innovus environment* 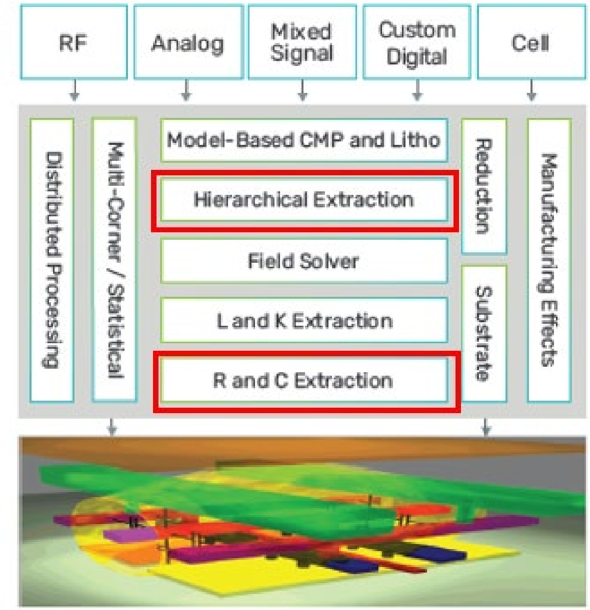 *Figure 2: Key functionalities of Quantus Extraction Solution*

	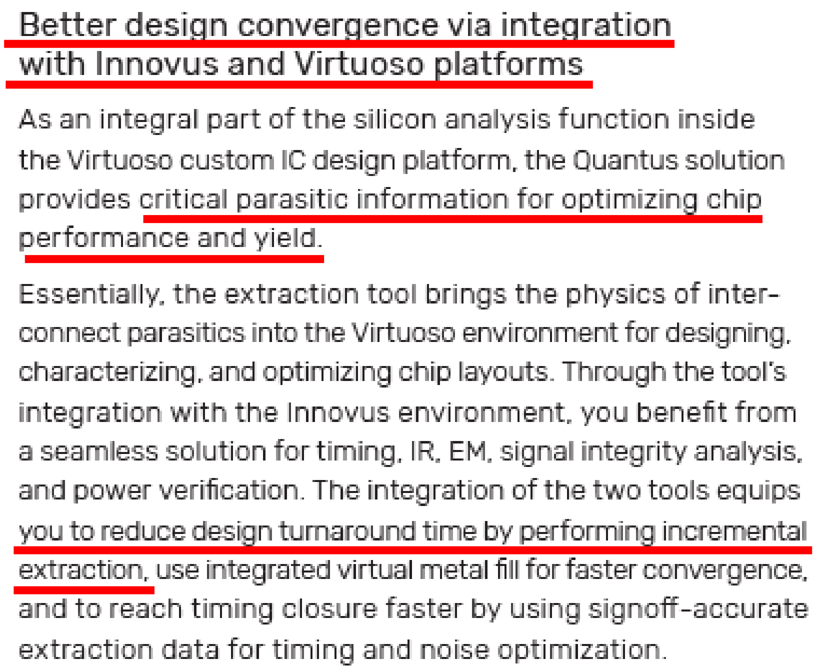 **Better design convergence via integration with Innovus and Virtuoso platforms** As an integral part of the silicon analysis function inside the Virtuoso custom IC design platform, the Quantus solution provides critical parasitic information for optimizing chip performance and yield. Essentially, the extraction tool brings the physics of interconnect parasitics into the Virtuoso environment for designing, characterizing, and optimizing chip layouts. Through the tool's integration with the Innovus environment, you benefit from a seamless solution for timing, IR, EM, signal integrity analysis, and power verification. The integration of the two tools equips you to reduce design turnaround time by performing incremental extraction, use integrated virtual metal fill for faster convergence, and to reach timing closure faster by using signoff-accurate extraction data for timing and noise optimization. https://www.cadence.com/content/dam/cadence-www/global/en_US/documents/tools/digital-design-signoff/quantus-extraction-ds.pdf, pages 1-3 For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was further made, produced, or processed from that circuit design that was created by performing a parasitic extraction for each net in the integrated circuit design that is enclosed by the window, as indicated by Cadence's incremental parasitic extraction functionality. *See* Ex. C at ¶¶ 44-47.

Caveat: The notes and/or cited excerpts utilized herein are set forth for illustrative purposes only and are not meant to be limiting in any manner. For example, the notes and/or cited excerpts, may or may not be supplemented or substituted with different excerpt(s) of the relevant reference(s), as appropriate. Further, to the extent any error(s) and/or omission(s) exist herein, all rights are reserved to correct the same.

EXHIBIT C

IN THE UNITED STATES DISTRICT COURT

FOR THE DISTRICT OF OREGON

PORTLAND DIVISION

BELL SEMICONDUCTOR, LLC Plaintiff, v. AMPERE COMPUTING, LLC Defendant.	Case No. ______________ **COMPLAINT FOR PATENT INFRINGEMENT** JURY TRIAL DEMANDED

DECLARATION OF LLOYD F. LINDER

1. I make this declaration on behalf of Bell Semiconductor, LLC ("Bell Semic"). I understand that Bell Semic will offer my declaration as evidence in support of its contemporaneously-filed complaint for patent infringement in the above-captioned case.

2. My qualifications to testify concerning the relevant technology are set forth in my curriculum vitae, which is attached hereto as **Exhibit 1**.

3. I received my Bachelor of Science degree in Electrical Engineering (BSEE) from UCLA in 1985. I received my Master of Science degree in Electrical Engineering (MSEE), also from UCLA, in 1987. Thereafter, I continued studying Electrical Engineering at USC, where I received an Engineer's Degree in 1989 and researched and completed my thesis towards a doctoral degree in 2002. Following the completion of my BSEE, I began work at Hughes Aircraft, where I worked for 12 years.

4. When Hughes Aircraft was acquired by Raytheon in 1997, my title was "Senior Scientist." At Hughes Aircraft, I was the technical lead for RF/analog/mixed signal IC development and was a subject matter expert (SME) in integrated circuits, serving as a company-wide resource for review of integrated circuit designs and technical support of new business.

5. My next position was as an Engineering Fellow at Raytheon from 1997–2002, where I again was the technical lead for RF/analog/mixed signal IC development and was a subject matter expert (SME) in integrated circuits, serving as a company-wide resource for review of integrated circuit designs and technical support of new business.

6. In February 2002, I began a role as Director of Technology at TelASIC Communications, a company that I also founded. In this role, I served as the technical lead for the development of state-of-the-art ADC (analog-to-digital converter) and DAC (digital-to-analog converter) commercial products for the cellular base station market.

7. In 2006, I began working under "Lloyd Linder Consulting" as an Independent Integrated Circuit Design, Systems, Intellectual Property, and Wireless Consultant, a role that continues to this day. In this role, I have served as a consultant to over 100 companies in the commercial and military contractor semiconductor component market space and have served as an expert witness in semiconductor cases.

8. From 2008–2009, overlapping with my consulting, I took a position at Menara Networks for approximately 10 months, where I was involved with the development of an electronic dispersion compensation (EDC) IC and the development of quad transceiver for next-generation 10 Gb/s ASICs with integrated FEC / EFEC in CMOS.

9. I have received various honors over the course of my education and career. In 1985, I was named UCLA's most outstanding senior electrical engineering student, graduating Phi Beta Kappa and Summa Cum Laude. I am an IEEE Senior Member and served as a Judge for the San Fernando Valley Section Entrepreneurial Business Plan Competition in 2008. I am a named inventor on over 100 issued United States Patents (with several currently pending) and over 300 international patents, and have published over a dozen journal and conference papers focusing on

semiconductor design and layout. I am a two-time Hughes Aircraft Division Patent Award Winner, and was named by Hughes as a Masters Fellow, Engineers Fellow, and Doctoral Fellow.

10. I have reviewed U.S. Patent No. 7,231,626 to Hoff et al. ("Hoff '626"), which is asserted in the Complaint, and its file history. In addition, I have reviewed the claim chart(s) accompanying the Complaint supported by this Declaration.

11. My college education over 15 years and 35 years of knowledge and experience in integrated circuit design, layout, and fabrication provides the necessary experience to support my stated conclusions set forth below.

Background on Integrated Circuit Manufacture, and Specifically the Layout Process Flow Segment of the Manufacturing Process and Subsequent Changes to the Physical Design

12. Semiconductor manufacture begins with the creation of a set of specialized electronic files that dictate the three-dimensional structure and features of the semiconductor device. These files, which are normally referred to as Graphic Design System (GDSII) files, are specifically formatted for and serve as necessary inputs for the devices that build the semiconductor device layer-by-layer according to the instructions contained in the GDSII files. Any changes to the structures in the GDSII files will result in changes to the structures in the fully fabricated device.[1] The manufacturing process ends with the wafer containing the individual semiconductor devices being fully fabricated and sawed into individual semiconductor dies.

---

[1] The physical design validation of an integrated circuit design ensures that all spatial constraints are satisfied for the traces and devices formed in various layers of an integrated circuit die. The structures formed in the several layers of an integrated circuit die are represented in a GDSII format file that contains the chip topological information for creating the masks used in manufacturing the integrated circuit dies. This is also called the "layout," and which patents in this area typically call a "design". The GDSII format is an industry standard used by commercially available physical verification tools to represent physical design data. All structures affecting the performance of the circuit die must and will be present in the layout.

13. I have created the image below, which provides a simplified schematic showing, at a high level, a commonly-used integrated circuit design flow process that is representative of many (if not most) process flows in current use for creation of circuit layouts:

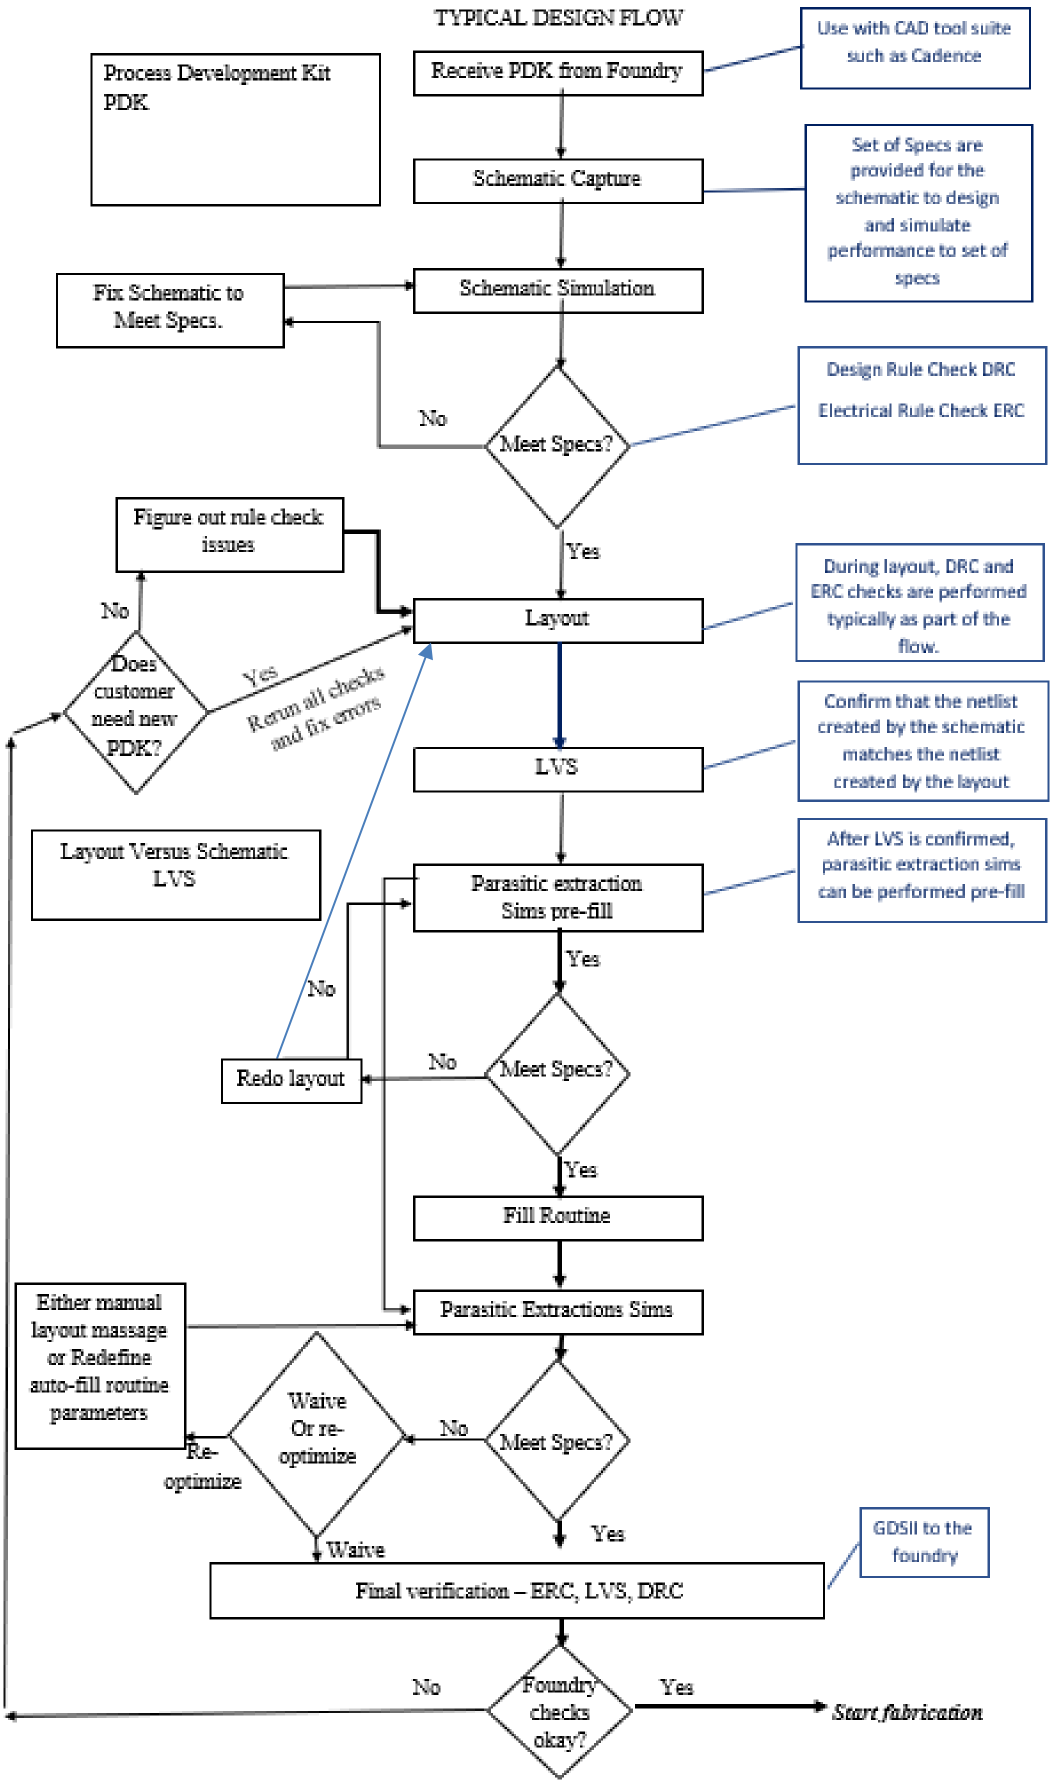
TYPICAL DESIGN FLOW
Process Development Kit PDK
Receive PDK from Foundry
Use with CAD tool suite such as Cadence
Schematic Capture
Set of Specs are provided for the schematic to design and simulate performance to set of specs
Fix Schematic to Meet Specs.
Schematic Simulation
No
Meet Specs?
Design Rule Check DRC
Electrical Rule Check ERC
Yes
Figure out rule check issues
No
Layout
During layout, DRC and ERC checks are performed typically as part of the flow.
Does customer need new PDK?
Yes
Rerun all checks and fix errors
LVS
Confirm that the netlist created by the schematic matches the netlist created by the layout
Layout Versus Schematic LVS
Parasitic extraction Sims pre-fill
After LVS is confirmed, parasitic extraction sims can be performed pre-fill
Yes
No
Redo layout
No
Meet Specs?
Yes
Fill Routine
Parasitic Extractions Sims
Either manual layout massage or Redefine auto-fill routine parameters
Waive Or re-optimize
No
Meet Specs?
Re-optimize
Waive
Yes
GDSII to the foundry
Final verification – ERC, LVS, DRC
No
Foundry checks okay?
Yes
Start fabrication

14. The integrated circuit design flow process includes a design engineer, using design tools, to create a design for an integrated circuit to be processed.

15. Design tools from vendors such as Cadence, Synopsys, or MentorGraphics (now Siemens) will then be used to design, simulate, and lay out integrated circuits. The typical design tool suite includes[2] schematic capture, simulation, layout, verification (layout versus schematic (LVS) and design rule check (DRC)), and fill generation routines. These fill routines can be automated or manual, and can be provided by the design tool company in whole or in part.

16. To be sure, the precise capabilities of each design tool available to a particular design engineer may differ within a company (based on what options in the design suite are available to a particular user or on a particular device), and between different design tool suites. However, based on my experience, at a high level, the design tools used by design engineers in the semiconductor industry, all operate in substantially similar fashion for schematic capture, simulation, layout, verification, design rule check, and fill-generation. In particular, based on my experience as a consultant, the design tools commonly used in the industry to place dummy fill operate in substantially similar fashion in providing incremental and timing-aware fill generation for integrated circuit layouts. This also applies to the design tools commonly used to identify texted metal shorts, which likewise operate in substantially similar fashion.

17. In the design process, the schematic is created first. The layout design tool is used to place and route all of the active (i.e., transistors) and passive components (i.e., resistors, capacitors, and inductors), and the interconnections between devices (represented as wires) in the schematic. It represents the circuit function that is to be physically implemented in the silicon. The schematic is created and simulated, using the CAD tools, to confirm that the circuit functions to a desired specification.

---

[2] Sometimes electrical rule check (ERC) is also included in design tool suite capabilities.

18. Once that performance specification is confirmed from the schematic simulation, the layout of the circuit is performed to physically place each of the individual elements necessary to implement the circuit functions set forth in the schematic in the GDSII file. During layout, layout rules for active and passive devices must be followed, but conformance is not checked until a DRC is run (typically at least as part of the final verification, though it can be run at any point or points in the layout process). Instead of only checking conformance at the end, it is possible to use a subset of the DRC deck to check for texted metal shorts in the layout of the schematic at an early or incremental point in the process flow. This allows the top-level routing to be completed in parallel with the block-level schematic and layout. Doing this will help accelerate the design timeline and avoid any delays occasioned by only finding such rule violations at the end of the design process flow for that particular schematic.

19. Once the layout is completed, it is compared to the schematic of the circuit using layout-versus-schematic (LVS) tool to confirm that the two are identical. From the schematic, a netlist (a list of devices and the associated nodes) is generated. From the netlist, the schematic could be re-generated by hand by drawing the devices and connecting the device nodes. From the layout of devices and associated nodes, a corresponding netlist is generated, from which a similar schematic could be generated by hand by drawing the devices and connecting the device nodes from the layout netlist. Then the schematic netlist is compared to the layout netlist using the LVS tool. The LVS tool compares the schematic netlist to the layout netlist to see if they match—i.e., whether they contain the same devices connected in the same fashion. If they do not match, the discrepancies between the two must be found and corrected, and LVS re-run. Any violations of layout rules must be corrected and DRC re-run for the layout.

20. After passing LVS, the process of performing parasitic extraction simulations before the fill has been placed (pre-fill) can be performed on an extracted netlist created from the

layout. If parasitic simulations are performed prior to the fill placement, the designer can get an idea of the impact on circuit performance from the basic layout parasitics pre-fill. From the layout, a netlist is extracted that includes any of parasitic resistance (R), parasitic inductance (L), parasitic capacitance (C), or any combination of the three. Additionally, the parasitic extraction can include what is termed "coupled" capacitance (parasitic capacitance between metal lines) as well as the parasitic capacitance to the substrate. The extracted netlist, with the selected added parasitics, can be used to run simulations on the baseline layout to determine if there is any performance degradation due to the baseline layout routing.

21. The simulated performance of the layout, which includes the parasitics, needs to be as close as possible to the specification that was already satisfied by the schematic. That is why parasitic extraction is performed, and why it is iterated pre-fill and post-fill. So if there is performance degradation due to the baseline layout, the layout is redone until its performance is at acceptable parameters. Ideally, the extracted simulation results match the schematic simulation results, which means that the layout parasitics had no impact on the circuit performance.

22. Once the layout passes pre-fill, the design tool is used to insert dummy fill at appropriate locations in the layout that do not contain devices or other features. As is well-known in the industry, the purpose of adding dummy fill is to achieve a higher and more uniform density of interconnect across the surface of each layer of the chip, to improve the outcomes of the chemical-mechanical polishing/planarization (CMP) step during fabrication. If individual pieces of fill are below a certain minimum size, they may not planarize properly during CMP, which will result in the dielectric material deposited on top of those too-small features not planarizing properly,[3] which will produce in dishing in the dielectric and result in a non-planarized surface.

---

[3] The effect on the dielectric from underlying interconnect is known as the deposition bias. A "positive bias" or "positive deposition" bias is when the width of the protrusion in the dielectric is

Thus, in practice, the fill pieces added cannot be below a certain minimum feature size. Adding dummy fill at or exceeding the minimum feature size and to achieve a higher and more uniform density of interconnect lowers the likelihood of defects caused by the CMP process step and thus improves the yield of modern integrated circuits.

23. Once all components of the integrated circuit design have been placed and routed, a physical design validation is typically performed at the very end of the design cycle. This ensures that all spatial constraints are satisfied for the traces and devices in each layer of an IC, that the die complies to all process rules, and that any additional required steps specific to manufacturability for a selected technology have been performed (e.g., metal utilization).

24. Even after a physical design validation, the physical design may change for any one of a number of reasons, including but not limited to timing delays, performance, or functionality. In such instances, the various steps in the process flow will have to be redone to accommodate the changes in the physical design. This includes placement of dummy fill as well.

25. Once the fill routine is completed, the fill checks are done, and final verification is performed again (LVS, DRC). The fill checks are performed based on percentage requirement on a specified area in the layout.

26. Once the layout database has been verified, it is sent for fabrication in the form of a GDSII database, which is the industry standard format for delivery of the chip database. As previously mentioned, fill is required to be included as part of the GDSII database.

---

greater than that of the underlying active interconnect feature. Conversely, a "negative bias" or "negative deposition bias" is when the width of the protrusion in the dielectric is less than that of the underlying active interconnect feature. In either case, large density variations of the active interconnect features will typically result in interconnect that is insufficiently planarized during CMP, and thus, overpolishing of the dielectric that produces significant dishing. This is particularly detrimental in fabrication of multi-layer chips and packages.

27. The design resource is provided with a process design kit (PDK), which includes all of the information necessary to capture a schematic, run a simulation, do a layout, and perform all of the checks on the layout to make sure that the final GDSII is in an acceptable form to be ready for fabrication. It is the design resource / customer's responsibility to make sure that the designed chip meets all of the expected requirements for fabrication and bears the risk of failing to follow any steps in the design flow. For example, if the circuit does not work, that is the customer's responsibility. If the layout does not match the schematic, that is the customer's responsibility. The GDSII does have to meet all of the DRCs in order to be fabricated.

28. In order to develop an integrated chip product, tools are needed to develop the schematic, the layout, verification of the layout, and the final GDSII database for fabrication. Many companies use different tools (from different vendors) to accomplish this process either typically due to cost or preference of internal proprietary tools. Regardless of the process and specific tools that are used, the GDSII database goes through an internal DRC after it is received and before fabrication of the integrated chip:

a. The design resource receives a PDK that contains all of the information is included to create a GDS database to release for fabrication. This includes circuit symbols for the creation of the schematic, models for the circuit symbols to run simulation, and associated layout devices that have been created with all of the process layers needed.

b. Additionally, there are what are known as "rule decks" in the PDK that allow for LVS and DRC. A rule deck is typically a file that specifies all of the available rules (for example, minimum feature sizes such as line width and minimum fill dimensions), the layers to process on each rule, and the parameters of each rule. The LVS deck compares the schematic to the layout, and the DRC deck covers all of the design rules for placing and routing devices. For LVS, a netlist of the layout is created. This netlist is compared to a netlist created for the schematic. The LVS tool compares the two to determine if they match or not.

c. Additionally, there is a parasitic extraction deck that extracts all of the parasitics of the layout that is used to run simulations to close timing or to confirm that the layout still meets all of the chip requirements.

d. There can also be an electrical rule check (ERC) deck as well, depending on the fabrication involved.

29. If the DRC rules at pre-fabrication do not match those at the design resource, it is possible that there will be DRC errors. This could be due to a number of reasons, including the DRC in the provided process design kit (PDK) is not up to date, and so the PDK will be updated with the updated DRC and the design resource will have to redo everything and fix the DRC errors, providing a new GDSII database before fabrication can begin. These DRC checks at pre-fabrication will include checks for the fill on all layers to confirm that the fill requirement is met, on a granular level, for all tiles at the chip boundary level.

Even Minor Design Changes Can Lead to Substantial Delays

30. Traditionally, each step in the design process flow was performed on a layer-by-layer basis, including but not limited to routing, parasitic extraction, calculation of delays, DRC, and placement of dummy fill. Any change would result in needing to redo the process steps for the design as a whole, as these changes usually came at the end of the process flow after all the layout and rule-checking steps were performed. Thus, it made sense to do any and all changes after all the layout steps have been performed.

31. However, it is rarely (if ever) that any modern IC proceeds through the design process flow without changes to the physical design along the way, especially as designs have

become more complex. This can occur, for example, when interconnect density is not met, when a connection error is discovered, when timing characteristics are not met, or for one or more of any number of other reasons. These changes can occur at various points throughout the design process.

32. The problem with such revisions to the physical design during or after the layout process flow is not just that the particular step in the process flow will need to be redone, but also that all subsequent steps (which depend on the prior step) will also need to be redone, even if the change was extremely minor and only affected a small segment of the device and/or relatively few nets. For example, any revisions to routing after dummy fill has been placed will require at least revisions to the dummy fill locations.

33. And, as mentioned above, each such step and all subsequent steps would traditionally need to be done on a layer-by-layer basis. Thus, the amount of time it would take to implement changes to the physical design of the chip would not scale with the size of the change itself (which could be relatively minor), but rather with the complexity and size of the IC as a whole. This started to become especially burdensome by the early 2000s as cutting-edge devices grew to encompass numerous layers and the number of nets and complexity grew exponentially.

34. For example, the typical turnaround time for a design change implementing a functional or timing ECO was, with traditional design tools and process flow, typically on the order of one week regardless of the size of the design change. (Hoff ’626 at 2:37–40.) That is because even a design change of just a few cells would still need to be merged into the overall design of substantially greater size, with the routing, DRC verification, net delay, and parasitic extraction, among others, scaling with the overall size of the IC. (*Id.* at 2:40–49.)

35. To make matters worse, each time an ECO is received, the process flow will need to be repeated. This sort of iterative process could and did significantly impact the design schedule and typically resulted in substantial cost overruns.

36. Prior to Hoff ’626, every revision to the physical design or an implementation of another ECO during the design process typically required turnaround time on the order of approximately one week regardless of the extent of the revision. (Hoff ’626 at 2:36–44.) That is because even small revisions needed to be merged with the overall design and all the parameters (including cell placement, routing, DRC and physical validation, parasitic extraction, timing closure, and dummy fill) rechecked for every net in the design. (*Id.* at 2:44–52.) This prior art process flow is depicted in Figure 1 of Hoff ’626:

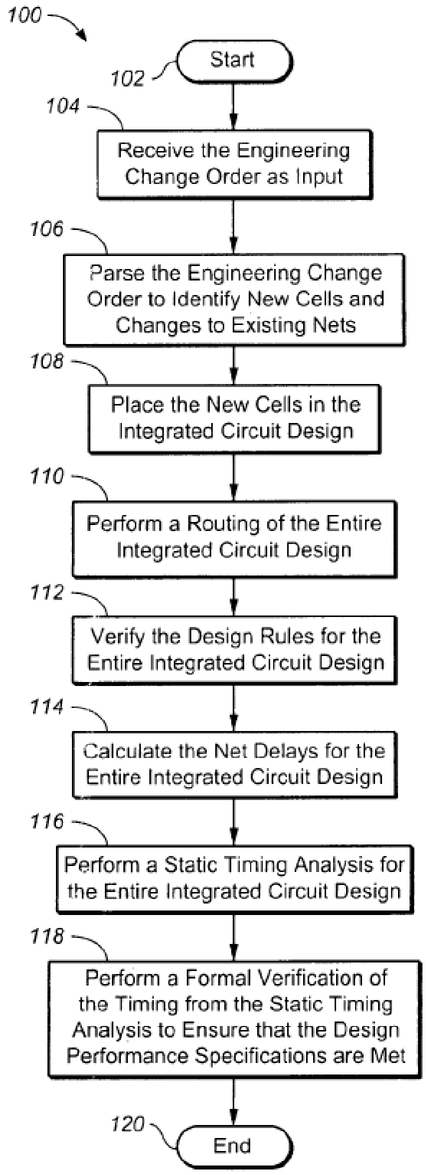

FIG. 1

Explanation of Hoff ’626

37. In place of having to deal with redoing all of the subsequent process steps for the entire layer (or design as a whole) at a time whenever a change to the physical design is implemented, Hoff ’626 teaches a technique that permits changes to the physical design that proceed incrementally, rather than for the layer or design as a whole. (1:25–2:3.) Specifically, Hoff ’626 teaches that by enclosing the physical design change in a window that is a subset of the IC design, that change can be implemented incrementally by performing routing only for the nets

enclosed by the window encompassing the physical design change. (*Id.* & Fig. 2.) Thereafter, that revised routing in that window replaces the original area in a copy of the IC design to generate a revised IC design. (*Id.* & Fig. 2)

38. Optionally, other steps in the design process flow can also be performed incrementally following the incremental routing, i.e., only for the nets encompassed by the window needed to implement the physical design change. Those can include calculation of net delay for each such net, DRC for each such net, and/or parasitic extraction for each such net. (*See, e.g.*, Fig. 2, 4:32–52.)

39. The significant efficiency gains from the ability to implement ECOs without having to rerun the routing tool for the entire design are repeatedly described within Hoff ’626. (*E.g.*, 1:15–22, 2:37–53, 3:19–24, 4:26–32, 6:29–33.) This helps IC manufacturers “avoid[] repeating calculations for nets that are not changed or effected by the [ECO].” (*Id.* at 6:31–33.) These substantial time savings allow designers to meet aggressive design schedules. Based on my experience in semiconductor layout and design I agree that this new and improved incremental process flow results in substantial efficiency gains in the design process for modern semiconductor devices. These gains are so substantial, and these incremental methods of routing and subsequent layout steps are so widely used today that it is hard to quantify just how important the inventions claimed by Hoff ’626 are to achieving current time-to-market for modern chip designs. It is also difficult to quantify how much the typical IC production schedule would be delayed for modern chip designs without the use of the methods claimed in Hoff ’626.

40. Based on my experience in semiconductor layout and design, it was not well-understood, routine, or conventional at the time of Hoff ’626 to perform routing and subsequent layout process steps in an incremental fashion following receipt of an ECO. In particular, it was not well-understood, routine, or conventional to create a window in the IC design, defining an area

less than the entire IC design, that enclosed a revision introduced by the ECO; nor was it well-understood, routine, or conventional to perform an incremental routing of the IC design only for each net enclosed by that window. Accordingly, it was also not well-understood, routine, or conventional to create a revised integrated circuit design by replacing the area bounded by the window in the original IC design with the results of the incremental routing. These features, central to Hoff '626, is required by every claim in the patent and recited explicitly in both independent claims. This is not only true in considering those elements by themselves, but also in an ordered combination with the other recited claim elements in creating a novel process flow that did not require a full-layer re-routing and subsequent process steps simply to account for the relatively few nets affected by an ECO.

41. This also applies to the dependent claims, which similarly recite performing subsequent process flow steps that would be affected by performing the ECO on an incremental basis (i.e., only for the nets enclosed by the window). Thus, the features recited by these dependent claims were also not well-understood, routine, or convention because the conventional design tools available at the time did not perform such process flow steps on an incremental basis.

42. Hoff '626 explains that "[i]n previous methods for implementing an [ECO] request . . . design tools are run for the entire integrated circuit design." (1:15–17.) As a result, "the typical turnaround time is typically about one week regardless of the size of the [ECO]." (2:38–41.) As I have explained, the typical process flows and all the routing and layout tools of which I was aware lacked the capability to accommodate ECOs without a full-layer re-run of the routing tool and subsequent calculations such as net delay and parasitic extraction. Nor am I aware of any "unconventional" tools or process flows dating to the time of Hoff '626 that had such capabilities. During prosecution of the application that matured into Hoff '626, the applicant specifically amended the claims to clarify that the routing of the incremental circuit within the window was an

incremental routing, rather than for the circuit design as a whole. The applicant explained that incremental routing was not disclosed in the prior art cited by the examiner, and the examiner agreed with applicant's arguments in allowing the claims. I agree with the examiner and the applicant.

Claim Charts

43. I have reviewed the Complaint supported by this Declaration, along with the Claim Charts showing infringement of Hoff '626 For at least the reasons set forth below, I agree that the Claim Charts establish use of at least one of the methods recited by the claims of Hoff '626.

44. I have used design tools from different vendors in my career. As a consultant, I use the tools to review schematics and layouts, which has included industry-standard tools for detection of texted metal shorts. Based on the requirements for the latest process technology nodes, and the yield requirements for these technologies, the latest fill tools that are used by designers use timing-aware fill routines with minimum fill dimensions to meet timing as well as yield requirements simultaneously. These tools typically operate on an incremental basis; although they can and do perform routing and insert dummy fill on a layer-by-layer basis, they typically operate in incremental fashion thereafter so that the impact of layout changes and ECOs to the overall design schedule is minimal and does not require rerunning the routing and dummy fill tool for the entire layer for each and every change. As I have done as a consultant, I can review either layouts post-fill or reverse engineering ("RE") of semiconductor die to confirm that these tools have been used to construct the layout or the die.

45. Given the aggressive schedules for bringing modern semiconductor devices to market, and the availability of incremental dummy fill in common design tools like Cadence's Innovus product, it is unlikely (if not implausible) that most chip designers would not have access to design tools that practice the inventions claimed in Hoff '626. I am aware that at least Cadence

provides this functionality. It is even less likely that such designers would not use the incremental dummy fill features that allow ECO without a time-consuming (and design-freezing) repeat of the dummy fill insertion process for the entire layer whenever the layout changes or another ECO is implemented. Especially because these typically happen late in the design process, and often happen more than once, any entity who declined to use these features would be at a substantial competitive disadvantage in bringing its products to market in a timely fashion. As such, based on my experience in semiconductor layout and design, and my review of designs and supervision of designers that used such tools, I believe that it is highly unlikely that such functionality was not used in creating most modern semiconductor devices.

46. Even when the full history of the GDSII database for a particular integrated circuit is not available, my experience in semiconductor design and layout gives me sufficient basis to opine whether one or more of the methods claimed in Hoff ’626 have likely been used in creating integrated circuits.

47. Given the multiple dependencies in the semiconductor processing design flow and the reality of ECO after layout has been completed, I believe that it is highly unlikely that anyone using the Cadence Innovus tool (or another design tool with similar functionality for incremental routing and dummy metal fill) to create a modern IC would not have used at least one of the inventions recited in Hoff ’626 to minimize the delay from having to re-route every net following an ECO, and then re-do the dummy fill. The delays from having to manually re-run the routing tool after each ECO or layout change is so impactful that failure to use these now-commonly available tools would result in a severe competitive disadvantage and substantial delays in bringing products to market.

48. By contrast, based on my experience in semiconductor layout and design, I would only assume that relatively simple IC designs would have been made in recent years without employing at least one of the methods claimed in Hoff '626.

49. In addition, based on my experience, it can be assumed with a high degree of confidence that modern components in the same family or product line made by the same producer and used by the same customer in the same product line share similar features and were designed and laid out in similar fashion. For example, I would expect with a high degree of confidence that two RYZEN processors made by AMD (such as the Ryzen 7 1700 and 5 5500U processors) or two Qualcomm Snapdragon transceivers (such as the Snapdragon 865 and 665 transceivers) would have similar performance specifications, be used for similar (if not identical) purposes by similar clients, be built on the same knowledge base and design history of prior generation devices, and thus share substantially similar design philosophies. Given that these are cutting-edge, modern devices, I expect that their production similarly involved incremental layout and routing methodologies following at least one of the methods claimed in Hoff '626.

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Dated: November 9, 2022

Lloyd Linder

EXHIBIT 1 to
EXHIBIT C

cell 818.632.9660
lflinder@yahoo.com
09/22/2021

Lloyd F. Linder

Skills summary

Extensive experience in high performance / high dynamic range analog mixed signal, custom digital, and RF integrated circuit design, layout, and test, from concept to production for commercial, military, and space IC products. Have knowledge of IC process development, characterization, and modeling. Have significant amount of experience in obtaining new military, space, and commercial IC business, as well as obtaining funding for start-up activities. Technical oversight for large IC design teams (20-40 people). Have ability to contribute creatively to the solution of difficult technical problems. Have 99 U.S. patents issued with twenty U.S. patents pending. Have 200-300 issued international patents. Have experience at the discrete analog / RF / mixed signal board level design and layout, system analysis, link budget analysis.

Specialties:

IP analysis / technical due diligence for M&A
IP portfolio management and creative contribution to new IP generation
Prior art searches and technical support for patent office action amendments
Perform simulations, review IP, and provide expertise in support of patent litigation
System Design/Architecture/Analysis/Block Specifications
New business development/capture
Winning proposals for small (SBIR Phase I and II) and large businesses
Product road maps
VC funding pursuits
Technical due diligence for VCs, angel investors, and M&A
IP creation and protection
Client deposition
Technical lead for IC design groups and product development
AMS/RFIC Design Chip/Circuit/System Architect
High speed, high performance ADC, Sample/Hold, ADC drvier amplifier, and DAC architectures
RF/AMS/SOC BIST/DFT architectures and methodologies
PLL and DDS
Digitally programmable RF transceivers/SDR/GPS/cellular/wireless transceiver architectures
RF TxRx, optical TxRx, modulator driver, LDD, TIA
Flash Ladar, active/passive imaging ROIC
regulators, high voltage/high current switches, ATE electronics
SiGe BiCMOS, CMOS, SOI, bipolar, Complementary bipolar, GaAs, InP
Digital beamforming
Cadence tools - schematic capture, SPECTRE, layout review

Solutions to production problems
Secret clearance

Objective

Looking for consulting opportunities to utilize my experience and provide technical leadership in all aspects of analog / mixed signal / RF integrated circuit and discrete circuit design: architectural conception, design, simulation, layout, test, and measurement / simulation correlation.

Experience

April 2006 – Present

Lloyd Linder Consulting

Consultant

Black Forest Engineering / Luminar Colorado Springs, CO.

- Architecture development for next generation automotive lidar ROIC

Strategic IP Initiatives Inc. Morgan Hill, CA.

- Review of patent portfolio for PIC applications in SFP optical modules

GreyB Service Pte Ltd Shaw Centre, Singapore.

- Provide expertise for many different invalidity cases for many different clients

Upwork Santa Clara, CA.

- Technical writer for application notes and VNA user’s guide

Syrnatec Inc. Newington, CT.

- Consult on ROICs for SBIR proposals

Kenney&Sams, P.C. Boston, MA.

- Expert witness in software defined radio development case

Microchip Technology Inc. Burnaby, Canada

- Participated in ADC architecture study for 28 and 56 GBps SERDES products.
- Did webinar for world-wide IC design staff on ADC design issues.

ElevATE Semiconductor Carlsbad, CA.

- Write white paper on ATE products.

GoodIP GmbH Munich, Germany

- Review and analysis of GaN and LED patent portfolio for potential acquisition.
- Spoke as the GaN / LED IP expert at webinar to discuss IP auction.

Nevada Nanosystems Inc. Reno, NV.

- Review and analysis of ASIC architecture and requirements in support of MEMS control circuitry. Evaluate IC design houses in the down-select process.

Axzon Inc. Austin, TX.

- Contribute to the analysis and review of CMOS transceiver architecture for UHF RFID reader.

Microcosm Inc. Torrance, CA.

- Architect the discrete RF transceiver and analog base band signal processing solution for digital beamformer. Develop packaging concepts for the beamfomer and interface to the antenna elements.

Otava Inc. Moorestown, NJ.

- Contribute to the definition of overall transceiver and circuit architectures for 5G beam formimg solution for 28-40 GHz 5G applications.

Analog Circuit Works Boston, Mass.

- Review of specification for fiber optic SOC application. Contribute to system level analysis and architecture.

Second Sight Medical Sylmar, CA.

- Debug of current eyewear for blind people. Contribute to reduction of noise and coherent spurs in current product.

Linear Microsystems Irvine, CA.

- Architect for high volume SOC for VR headset application. System analysis for SOC proposal for fiber optic communications applications.

Teqnovations Colorado Springs, CO.

- System level architect for receiver for phased array radar applications.

Analog Devices Inc. Colorado Springs, CO.

- Work with design team on a DAC buffer IC to improve performance in the lab and remove oscillation. Review existing architecture for redesign and developed new architecture.

Facebook Woodland Hills, CA.

- Schematic and board layout reviews of electronics for flight system

GHB Intellect San Diego, CA.

- Review of Patent claims for client for multiple issued patents in the wireless communications area.

BINJ Labs Scituate, MA.

- Definition of top level Software Defined Radio architecture for wide band operation. Work with board development effort and SDR development company.

Sentinel Monitoring Systems Inc. Tucson, AZ.

- Schematic review of high speed data converter and timing board schematic.

FlexPowerControl Woodland Hills, CA.

- Consultation on system level requirements document in-home energy control product development. Develop IP for the company.

SpectraResearch Dayton, OH.

- Consultation on integration considerations for discrete X-band and Ka / Ku-band transponder architecture for SWAP-C improvements.

McKool Smith Dallas, TX.

- Perform simulations of IP for RFIC receiver, and provide expert opinion for patent litigation in a report for case :2:15-cv-00011-JRGRSP in the United States District Court for the Eastern District of Texas.

Quantum Semiconductor LLC San Jose, CA.

- Consultation on architecture and simulations for proprietary ADC architecture.

Maven Research San Francisco, CA.

- Consultation on use of design house by third party for products.

Faraday Technology Corporations Santa Clara, CA.

- Review of third party transceiver RFIC schematic design, layout, and testability.

Irunway Dallas, TX.

- Review of IP for legal firm in patent litigation. Simulation of IP for transceiver RFIC. Simulation results included in expert report.

Alphacore Inc. Phoenix, AZ.

- Architecture review and enhancements of high speed CMOS ADC and visible monolithic imaging chips. Develop DROIC architectures for SBIR proposal pursuits.

Brady Worldwide Inc. Milwaukee, WI.

- IP review of start-up for potential investment / acquisition. Develop sensitivity analysis for present products and future CMOS technology scaling.

InPhi Corporation Westlake Village, CA.

- Review and contribute to architecture refinements for single channel and multi-channel EAM drivers for 2 level and 4 level PAM. Perform architecture study for low power EAM driver.

DRS RSTA Inc. Cypress, CA.

- Help with EDA methodology for AMS design for ROICs.

Ridgetop Group Tuscon, AZ

- Architected wide band RF front end for multiple current SBIR proposals.

Space Micro San Diego, CA

- Review of 0.18 µm CMOS quadrature DAC design and layout.

Teradyne Agoura Hills, CA

- Analog DFT / BIST / testability architect for 40 nm CMOS SOC for next generation Teradyne tester. Architect transistor level circuit solutions for power supply IC for DUT testing. Investigate high voltage CMOS and complementary bipolar process technologies for internal design development with unique current clamping architecture. Define baseline circuit topology and perform simulations.

Lockheed Martin Moorestown, New Jersey
Littleton, Colorado
Deer Creek, Colorado

- Involved in high voltage driver IC development for GaN PA.
- Involved in architecture development and review of wide band band receiver integrated circuits from DC to Ka / Ku in SiGe.
- Involved in overall ADC architecture development and definition, and transistor level circuit architectures, for next generation RFIC transceivers and data converters. Architect of next generation transceiver architectures for SOC / heterogeneous applications. Involved in study for government customer of data converter architectures for next generation digital beam forming applications. Architecture review and development for high voltage analog driver array. Involved in architecture for multiple receiver RFIC developments for S, C, and Ka bands.

HRL Malibu, CA

- Perform market survey of component technologies for multi-mode, multi-mode commercial software defined radio applications. Summarize capability for data converters, front end modules, antennas, transceivers, and base band processors. Compare the available technology to an architecture based on custom chip development for the solution.

Micrel San Jose, CA

- Involved in the testing, debug, and redesign of a 65 nm CMOS 2.4 G / 5 G WiFi transceiver product. Provide technical guidance for redesign of RF front end and the 2.4 G / 5 G LO clock distribution.

Key2Mobile Westlake Village, CA

- Developed RF transceiver concepts for multi-band, multi-mode remote radio head system. This is included transceiver architectures based on digital beam-forming and direct RF sampling and direct RF synthesis data converters.

Hittite Microwave Colorado Springs, CO

- Consultant on the architecture for the high dynamic range, high speed IBM SiGe 8HP BiCMOS DAC and complementary bipolar ADC driver amplifier products for the cellular base station market.

Semtech Redondo Beach, CA

- Technical consultant on IBM SiGe 8HP BiCMOS interleaver IC and IBM 32 nm SOI monolithic coherent detection transceiver SOC for 100G coherent optical detection systems. Involved at the architectural level for the SOI interleaver, 8 bit, 64 GSPS ADC, and 8 bit 64 GSPS DAC circuits, and overall system calibration.

- Developed concepts for 10-12 bit, 4-8 GSPS ADC architectures for digital array radar and digital beam-forming applications, as well as high performance sample and hold architecture for military applications.

FBI Westwood, CA

- Technical consultant on matters of national security. Awarded medal for service to the country.

Nu-Trek Inc. San Diego, CA

- Developed an RF BIST architecture for characterization of a RF transceiver. Helped the company win Air Force Phase I and Phase II awards.
- Involved with test evaluation of pipeline ADC for cryogenic applications, L1 / L2 band GPS receiver, and Universal Reliability SOC development for lifetime testing of X-Band and L-Band transceivers.
- Responsible for the development of the company’s product roadmaps.
- Contributed to SBIR and STTR proposals on nonlinear coupled oscillators for active array applications, active sonar signal processing, high dynamic range ADC, GPS, RFIC transceivers, RIICs, and ROICs. Involved in testing of ADCs for ROICs, and architecture development for RF BIST.

FLIR Electro-Optical Components Ventura, CA

- Technical lead for the development on ROIC architectures for NASA, Navy, Air Force, MDA, and Army SBIR Phase I and Phase II programs. Development of novel active and passive unit cell architectures.

SYS Technologies / Kratos Defense San Diego, CA

- Review of SiGe BiCMOS class E power amplifier design to improve reliability of operation. Reviewed circuit, simulation results, test results (dynamic and DC), and performed thermal analysis based on self-heating.
- Review of schematic and layout for small form factor PCB that contains GPS transceiver RFIC and companion digital ASIC. Review of schematic and layout for PLL, IF / baseband test chip evaluation boards. Suggestions for characterization / test debug**.**
- Review Transmitter and Receiver schematics and block layouts of GPS transceiver IC in IBM 7HP 0.18 µm SiGe BiCMOS for redesign effort. Review of improved receiver IC design.
- Review of test results for the evaluation board and suggestions for characterization and debug. Review of board design and layout revisions for improved electrical and thermal drift performance for successful demo. Help with board yield and manufacturing issues.

Aerius Photonics LLC Ventura, CA

- Technical lead for STTR Phase I ROIC circuit design partner.
- Developed ROIC architectures for laser vibrometer SBIR Phase I proposal.

LinearChip Inc. Aliso Viejo, CA

- Developed complete single chip CMOS 802.11 a/b/g/n/ac transceiver architecture, with on-chip T/R switch, for proposal to commercial test equipment house.
- Developed CMOS MSK transmitter for patient temperature monitor and medical equipment tracking ASIC. Developed overall quadrature transmitter architecture, circuit topologies for the DAC, active filter, quadrature mixer, power amplifier, and temperature stabilized reference oscillator. Performed noise / distortion budgets, duty-cycled power calculations for extended battery life, determined PLL phase noise requirements, and defined the circuit block specifications.

Technical lead for the circuit simulations in XFAB XH018 0.18µm CMOS RF process.

- Participating in pursuit of new IC design business. Development of PHEMT TIA gain block concept, satellite receiver / de-multiplexer architecture, CMOS and SiGe RF receiver architectures, and CMOS analog AFE for hand-held controller for telescope, including 16 bit audio DAC architecture. Involved with multiple RFIC proposals for military and commercial applications. Developed multi-channel AM/FM receiver for location positioning application.
- Developed unique TIA / AGC / output amp circuit concept for cable TV over fiber market based on JAZZ 0.18µm SBCH18XL SiGe BiCMOS process.
- Review of battery charger circuit architectures for XFAB XC06 0.6 µm CMOS IC. Architecture suggestions for capacitor charging loop and low battery indicator circuit.

Wistron Corporation Taipei, Taiwan

- Contributed to discrete quadrature receiver architecture for a new module business proposal to DirectTV.
- Involved with review of existing ODU L-Band module specifications, and application to new digital L-Band module. Contribute to the definition of ASIC requirements and discrete component requirements for new ODU architecture.

Arete Associates Sherman Oaks, CA

- Phase I and internal IRAD programs for advanced TIA architecture development for high speed laser pulse return processing.
- Review of schematics, simulation results, and layout floor plan, and layout for a 64 channel TIA / OTA wavelength converter in IBM 5HPE 0.5µm SiGe BiCMOS. Function is to receive 1 ns laser pulse and provide linear gain for retransmission. Made suggestions for additional simulations for design robustness and concerns, as well as improved performance based on layout changes. Also made layout suggestions for improved isolation and performance.
- Developing unique 3D FPA and digitally programmable ROIC pixel unit cell architectures for inclusion in STTR proposal.
- Review of 2 TIA input / 128 channel output VCSEL driver, 2 GHz pulsed sampling IC in IBM 7WL 0.18 µm SiGe BiCMOS.

Raytheon Corp. Dallas, Tx. / El Segundo, CA

- Involved in the development of 40 GSPS ADC architecture in 32 nm IBM SOI.
- Involved in development of on-chip mechanical stress measurement circuit for thermal imaging sensor IC in IBM 9SF 90 nm CMOS process. Submitted patent disclosure on stress sensor circuit. Reviewed the pixel circuit design and the overall ROIC circuitry and timing.
- Participated in winning MDREX proposal effort to AFRL to develop high performance 10 bit, 2.2 GSPS ADC, 14 bit, 3 GSPS DDS, and 1:4 DMUX ICs in IBM 8HP 0.12µm SiGe BiCMOS, and integrate these functions with 8HP receiver and transmitter circuits developed by AFRL, as well as 8HP bias and control circuits developed by Raytheon. Involved with Phase I kick-off meeting.
- Participate in SBIR proposals development with partners Nu-Trek and Crossfield. Developed design methodology for CMOS GPS SOC integration. Developed sub-threshold CMOS RF circuits for low power GPS receiver concept. Generated phase adjust circuit for RF clock distribution for antenna array electronics, and submitted patent disclosure. Contributed write-up on high speed, high performance DAC topologies for DDS proposals. 1 Phase I SBIR proposal awarded.

- Schematic and layout review of IBM 7HP 0.18µm SiGe BiCMOS digital control and bias regulation circuits for MMIC common leg circuit. Reviewed regulator loading problems, solutions, and simulation results. Made simulation recommendations for verification of design robustness. Reviewed metal mask fixes.
- Review of test data, process parameters, and circuit design for production HRL G1.5 2µm InP band-pass Σ-Δ IC for yield enhancement. Developed vendor RFQ for multi-phase IC redesign effort.
- Involved in metal mask effort for 14 bit, 3 GSPS DDS / DAC IC in IBM 7HP 0.18µm SiGe BiCMOS through the Trusted Foundry program. Contributed to design, layout changes in the high speed digital DLL section to improve performance for higher clock applications. Documented circuit design limitations for future redesign for clock rate enhancement.
- Analyzing test data, PCM data for the metal mask DDS chip. Correlate sensitivity of test results to process parameters, and simulation results, to improve yield on future fabrications. Investigate design improvements for possible new mask set. Provide suggestions for performance improvements at the board level. Involved in improvements of the DLL section for next all-layer mask release.
- Member of team analyzing radiation induced latch-up effects in commercially available data converter for space application. Involved in focused laser beam testing to create single events in the converter. Analyzing radiation data, reliability data, and design guides for end-of-life DC power estimates for ICs designed in NS 0.8µm ABIC-IV BiCMOS, MAXIM GST-1 bipolar, IBM 5SF CMOS, and Honeywell HX3000 SOI.
- Review of link budget and block performance for IF and base-band ICs implementing high dynamic range receiver, including IF amp, mixer, VGA / attenuator, and ADC driver amp in IBM 5HPE 0.5µm SiGe BiCMOS. Made architectural suggestions for improved performance. Reviewed test results, and made suggestions for debug.
- Design review of IBM 5HPE 0.5µm SiGe BiCMOS IF receiver ASIC including mixer, gain control amp, and ADC buffer amplifier. Review of measurements for test circuits.
- Jitter analysis for ADC module clock path using divider circuits.
- Design reviews of IBM 5DM 0.5µm SiGe BiCMOS L1 / L2 GPS dual channel receiver IC: LNA1, LNA2, RF amp, RF mixer, IF mixer / AGC, gain calibration loop, baseband amp, PLL, and ADC. Defined layout floor plan, channel layout isolation techniques, and package requirements to meet isolation requirements. Reviewed final GPS receiver layout and developed plan for simulation of layout, package, and external component parasitics. Contributed to correlation of test results. Review of socketed test board layout.
- Review of quad channel IQ detector hybrid design and layout. Made suggestions for debugging of existing hybrid oscillation and improved performance. Review of TI BiCOM-2 0.7µm CBiCMOS IQ IC performance and correlation to hybrid measurement.
- Review of simulations and test results for IBM 8HP 0.12µm SiGe BiCMOS millimeter wave transceiver chipset module performance from outside vendor.
- Review of M/N PLL hybrid design. Review of TI BiCOM-2 0.7 µm CBiCMOS M/N PLL IC performance and correlation to hybrid measurement.
- Performed BOL and EOL noise analysis to estimate the jitter of on-chip clock receiver circuits for the SPT7760 ADC IC based on the MAXIM GST-1 process parameters, and off-chip COTS components for satellite application.

Menara Networks Irvine, CA

- Reviewed schematics and simulation results for high frequency, multi-GHz active RF low pass filter in JAZZ SBC18HXL 0.18µm SiGe BiCMOS. Involved in top level layout review of first generation EDC IC.
- Contributing to the development of Nyquist 10 GSPS sample / hold circuit, for hold mode feed-through reduction, low distortion, and low power. Contributed to the DC AGC circuit to generate coefficients for Gilbert multiplier as well as developed a new summing circuit for the sampled signals in the analog sampled FIR filter. Involved in solving layout-induced performance problems.
- Involved in the clock distribution architecture and feedback loop to reduce DC offset effects on zero crossing. Review of regulator circuit power-up issues. Created new ADC-based residue architecture to generate the appropriate transfer function for a phase detector circuit for consideration on a higher performance, lower power version of the existing Fiber Optic Receiver IC in JAZZ SBC18HXL 0.18µm SiGe BiCMOS.
- Preliminary review of third party design of 10 Gb/sec CDR PLL schematic blocks for SOC application in STM 0.13µm SiGe BiCMOS. Reviewed metal mask circuit design changes.

MOSIS / ISI Marina del Rey, CA

- Review of MOSIS website. Provide report consisting of ways to improve technical and business information on the website, as well as customer support. Used client questionnaire to provide additional insight into website, technical information, customer service and support.

Tangea Semiconductor Manhattan Beach, CA

- Participated in business plan development, facility requirements, engineering manpower requirements, technical concepts, and potential investor meetings.

Technoconcepts Inc. / Terocelo Van Nuys, CA

- Involved in the design, development, and review of RF receiver and transmitter ICs for multi-band, multi-mode applications using JAZZ SBC18HXL 0.18 µm SiGe BiCMOS. Versions of the Receiver IC include RF VGA, mixer, RZ (Return-to-Zero) and NRZ (non-Return-to-Zero) 6 GSPS 1 bit Σ-Δ ADC, with 2 feedback DAC and 3 feedback DAC versions, PLL, 1:16 DMUX, and base band decimation and filtering. Activities include circuit-level architectures to improve existing performance, schematic and simulation reviews, top level layout floor-planning, layout review, and evaluation board review.
- Involved in testing and correlation of packaged part performance, to simulation, for noise density, harmonic distortion, and intermodulation distortion products. Involved in architecture improvements and enhancements, based on test results, for a metal mask effort. Improvements include extension of VGA architecture to extend the input 1 dB compression point, IIP3, and bandwidth. Added provision to sample very low input frequencies. Improved the dynamic range of the transconductance amplifiers in the modulator. Added dither DAC at the comparator input to improve idle tone performance. Extended the dynamic range of the present RZ Σ-Δ design. Improved the DAC settling and noise response. Added an LNA, and a single-ended to differential converter to interface to the existing VGA.
- Development and layout of test chips for the TSMC 90nm CRN90LP CMOS RF process. Designs include single-ended, wideband and tuned versions of high IIP3 LNA architectures using novel distortion cancellation techniques for operation in the 1- 6 GHz frequency bands. Also developing versions of a 1.8V bidirectional, tri-stated CMOS / LVDS I/O driver with resistive and active back terminations,

and a single input, multi-band, multiple output, switched / tuned RF front end. Involved with CASCADE probing, and correlation to simulation results.

- Development of a wide-band, software defined, multi-band, multi-mode, receiver chip for the TSMC 65nm CRN65GP mixed signal / RF process. Architected the circuits for the wideband LNA with novel AGC and attenuator, the RF DMUX / on-chip RF filter bank, the RF MUX / single-ended-to-differential converter, the quadrature mixer with AGC, RSSI log amp, and the programmable base-band filters. All the RF blocks utilize IM3 cancellation. Target services include AM, FM, DTR, DTV, WiFi, and WiMAX for automotive application for Japanese customer. Involved in the link budget analysis, block specifications, and the requirements for the associated base band ADC and PLLs for the high band LO, low band LO, and ADC LO ICs. Also proposed this architecture for NSF proposal in order to obtain additional development funding.
- Involved in the design, layout, and review of customer demonstration board, based on existing transceiver chipset, for WiMAX applications, using Picochip base-band processor.
- Developing transceiver architecture concepts for multi-service / simultaneous /switched reception off of single multi-band antenna, diversity, and MIMO with self-calibration. Developed low power architecture for multi-service receiver summing into single ADC. Developed transmit / receive architecture for multi-band, multi-mode military radio. Architectures are baseline for new business opportunities.
- Developed IBM 0.12µm 8HP SiGe BiCMOS low power LNA topologies for 10-20 GHz for SBIR proposal. Developed RF interference cancellation architecture for SBIR proposal.

Red Dot Wireless Milpitas, CA

- Performed transceiver architecture study for TD-SCDMA / EVDO / WiFi / MIMO WiMAX application. Involved in presentations to investors.

Ubidyne Ulm, Germany

- Schematic, simulation, and layout review of IHP SG25H1 SiGe BiCMOS S1.0 Receiver RFIC. Review of characterization test results.
- Schematic, simulation, and layout review of JAZZ SBCH18XL SiGe BiCMOS S2.0 receiver, transmitter, and PA driver IC designs. Schematic and layout review of Toshiba 90nm CMOS high speed custom digital IC design.
- Review of test data and circuit design for S2.0 receiver to determine yield issues.

Dynamic Research Corporation San Diego, CA

- Review of existing top level IC issues and proposed packaging approach for GPS transceiver in IBM 7HP 0.18 µm SiGe BiCMOS.

Q3Web Wideband Wireless Inc. Harbor City, CA

- Constructed white paper on the development of IP libraries in the IBM CMOS and SiGe BiCMOS technologies for use by military contractors. Being submitted to government for funding consideration.

August 2009 – May 2011 Aerius Photonics Ventura, CA.

Senior Systems Engineer

- Directly responsible for winning new SBIR business for ROIC development. Captured $2.6M of ROIC development money. Have written multiple winning proposals for Army, Navy, MDA, NASA, Air Force, and NSF SBIR Phase I efforts. Have won Navy and ARMY Phase II efforts. Additionally, supported other proposal wins for other technology developments.

- Responsible for the development of new and novel circuit concepts for CMOS and SiGe BiCMOS passive and active imaging ROICs. Have found and engaged ROIC design partnerships. Involved in new business development.
- Have defined ROIC specifications and system requirements in conjunction with BALL Aerospace, Sensor Creations, Raytheon, BAE Systems, Arete and Associates, Tetravue, IDEO, Microvision, Velodyne, and other military and commercial contract partners.
- Technical lead for the conceptual phases of the following ROIC developments: JAZZ SBC35 ROIC for laser vibrometry, ON Semiconductor C5 0.5µm ROIC for 3-D FLASH LADAR for beach zone / surf zone applications, ON Semiconductor C5 0.5µm CMOS ROIC for dual well application (patent pending), JAZZ SBC18 CMOS monolithic imaging ROIC with integrated SiGe APD, and ON Semiconductor 0.18µm CMOS 1920 X 1080 SWIR ROIC.
- Contributed to the development of new laser range finder receiver architecture with unique time programmable gain.
- Developed digitally programmable pixel architecture and specifications for linear direct / coherent detect ROIC for STTR Phase I effort.
- Technical lead in development of single pixel discrete board level customer demo for STTR Phase I LADAR application, and board developments for: 10 Gb/second hexagonal InGaAs detector array, custom ROSA and CDR, 4 X 4 array with fan-out electronics in support of sub-ns laser returns, and new laser range finder receiver.

April 2008- January 2009 Menara Networks Irvine, CA

Director of ASIC Development

- Involved in simulations of existing Electronic Dispersion Compensation (EDC) IC to correlate to measured performance to define metal mask effort for JAZZ SBC18HXL 0.18µm SiGe BiCMOS IC. Helped define testing for chip characterization.
- Patent application on interleaved FIR with unique sample / hold for extended high dynamic range over previous implementations.
- Contributed to architectural improvements for low power version of EDC IC in IBM 8HP 0.13 µm SiGe BiCMOS.
- Performed simulation trade-off study to compare circuit performance for transmitter application in TSMC CMN65LP 65 nm CMOS, IBM 10LPE 65 nm CMOS, and IBM 8HP 0.12µm SiGe BiCMOS.
- Simulation and design of transmitter pre-driver and output driver for quad 10 Gb/s transceiver with integrated EFEC in IBM 10LPE 65 nm CMOS.
- Simulated low power EML driver concept in IBM 8HP 0.12µm SiGe BiCMOS.
- Involved in the investigation and development of new architectures in support of 100 Gbit/s optical links.
- Review of board schematics and layouts for EDC IC evaluation board, OTN XFP module, and OTN 300 pin module.

February 2002- April 2006 TelASIC Communications El Segundo, CA

Director of Technology / Founder

- Technical lead for next generation single supply (+5V) data converter ICs in IBM 0.5 µm 5AM SiGe BiCMOS: RF sampling 10 bit, 1 GSPS ADC, on-chip sample / hold, and up-converter IC with integrated 14 bit, 1 GSPS DAC, IF amplifier / mixer. Developed self-heating error correction for quantizer preamp.

- Technical lead for new business pursuits in military and commercial integrated circuits. These included RF handset transceivers, spot BTS, fiber optics, analog transmit cancellation receiver for Gigabit Ethernet, automotive radar, data converters, sample / hold, FPA imaging sensor readout, and arbitrary waveform generators. Technical lead in new business pursuits in board-level DPD-based transceivers based on COTS components. Involved in proposals with Raytheon for DARPA programs, including RHBD (Rad Hard by Design) and Team Phase II.
- Technical advisor for the conception, development, and commercial production of the TC2412 14 bit, 737 MSPS DAC IC and TC1412 14 bit, 250 MSPS ADC IC in IBM 5AM (data sheets available at www.telasic.com/website/products). Architectural contributor to base station transceiver chipset concept. Data converter IC products are being sold to PMC for the NTT DoCoMo base station market. Involved in development of IBIS models of DAC and ADC IC for customer board level simulations.
- Technical contributions to various IC designs including IBM 5AM ADC driver amps, Re-sampler, IBM 7HP test chips, including sample / hold, amplifier, mixer, ring oscillator, and 14 bit, 3 GSPS DDS. IBM 8T chips including re-sampler and 3 bit / 40 GSPS sample / hold / Quantizer / DAC.
- First version of the ADC IC developed, TC1411, received analogZone 2003 Product of the Year Award.
- Involved in business plan development, process to obtain first round of funding.
- Performed technical due diligence on potential new investments for venture fund.

1999– February 2002 Raytheon Advanced Products El Segundo, CA

Engineering Fellow

- Technical lead for new business pursuits in military / commercial ICs, working with organizations across the company. Programs won include: SIMBAW, TEAM, ADRT, ULTRACOMM, ACN, 3D Flash Ladar image sensor, and APLA.
- Involved in FPA multi-sampling ROIC developments in AMI 0.5μm CMOS: 10 X 10, 64 X 48 arrays. Unit cell development in IBM 0.13 μm CMOS for 256 X 256 array. Developed fan-out concept to maintain small cell pitch while allowing for multi-sampling analog architecture. Evolved multi-sampling concepts for programmable sample time stamp and number of samples.
- Technical advisor for integrated circuit development across Raytheon. This includes design review / debug of existing IC developments.
- Involved in the design and development of data converters, IBM 5HP IF sampling Band-pass Σ-Δ, compact DDS, and RF transceivers for the military and commercial markets. Circuit concepts include fast frequency hopping PLLs, active biquad filters, and a sine weighted DAC.
- Developed Tondelayo 802.11a half duplex transceiver chipset in IBM 5HP for start-up (Systemonic) acquired by Philips. The transceiver demonstrated 802.11a frequency bands in the 5- 6 GHz frequency range using a PCMCIA format.
- Technical contributor to conversion of 0.6μm to 0.5μm NS CMOS for 14 bit, 10 MSPS radiation hard, algorithmic ADC. Involved in design and review process.
- Cooperative effort with LUCENT for next generation design improvements and completed plastic / ceramic package performance evaluation trade-off of CSP1152A CMOS 14 bit, 65 MSPS ADC for Sirius radio application.
- Technical evaluator for BOEING 0.35μm CMOS 1.0625 Gbit/sec Multi-channel Fibre Channel Transceiver for Raytheon AESA application. Contributed to

design reviews, layout reviews, design changes / iterations, suggested simulations to run, and evaluated test results.

- Developed concepts for the integration and packaging of RF MEMS devices with integrated circuits. Circuit architectures included RF front ends with tunable capacitors, tunable RF band-pass and notch filters, as well as active MUX circuits for frequency hopping between filters.

1993–1999 Hughes Communications Products El Segundo, CA

Senior Staff Engineer / Senior Scientist

- Technical lead of the development ICs for the Digital Receiver Program. This included NS ABIC-IV 0.8µm BiCMOS ICs with the following functions: LNA, mixer, LO driver, DAGC, fractional frequency hopping low phase noise PLL, IF amplifier, video amplifier, serial interface, log detection / blanking, and control logic, and a LUCENT CBIC-V2 summing amplifier IC. Involved in the packaging and test of plastic packaged parts. Digital Receiver board contained 2 chip UHF receiver, single chip GPS receiver IC, and master PLL IC.
- Technical lead for research and development of RF and analog ICs in CMOS, Hughes NB SOS / SOI, silicon (NS ABIC-IV, V, MAXIM SHPi), and IBM SiGe bipolar / BiCMOS process technologies. Functions included LNAs, RF LNA, mixer, VCO, ring oscillator, video amplifier, IF amplifier, sample / hold, high speed 10 bit ECL-to-CMOS translator / latch, 12 bit DAC, and LP Σ-Δ.
- Development of AM / FM LNAs in IC Delco 1.2µm CMOS for automotive radio.
- Involved in the chipset development for NS for the 1.0625 Gbit/sec ANSI X3T11 8B/10B standard using ABIC-IV process. Involved in laser diode driver, TZA, SIPO, PISO, Transceiver ICs.
- Involved with process development / modeling for CMOS, SOS, and silicon / SiGe BiCMOS process technologies.
- Team member of the NS ABIC-IV ADC chip development based on requirements for the ICO satellite system
- Involved in the debug, characterization, and production of hybrids and modules for airborne radar and AMRAAM missile applications.

1991–1993 Radar Systems Group, Hughes Aircraft El Segundo, CA

Staff Engineer

- Lead the IC development of high performance data converter components including sample / hold, summing amplifier, timing generator, band-gap, ADC reference amplifier, video amplifier, DAC, flash quantizers, and buffer amplifier in LUCENT CBIC-U / U2 and MAXIM SHPi and CPi processes.
- Involved in the development of mixer ICs for high dynamic range radar using HRL InP, AlGaAs, and Litton and Hughes D-MESFET technologies.

Fall Semester 1990 Electrophysics Department, USC Los Angeles, CA

Lecturer

- Taught EE448, Senior Electronics. Generated syllabus, created homework problems, tests, and solutions. Handed out final grades.

1986–1991 Radar Systems Group, Hughes Aircraft El Segundo, CA

Member of Technical Staff

- Involved in the design of components for high performance data converters.

- Contributed to the design of high dynamic range buffer amplifiers, sample / holds, integrators, and gain stages in Hughes 2 µm CBiCMOS, Fairchild 1.25 µm Fast-Z Fineline, and NS Aspect-II, Aspect-III, and ABIC-IV.
- Designed 2 bit adaptive threshold ADC IC in ORBIT 2µm CMOS for EPLRS radio.
- Contributed to architecture for Hughes NB 2 µm CMOS gate array ASIC.
- Involved with process development / modeling for bipolar, complementary bipolar, CMOS, BiCMOS, CBiCMOS.
- Generated analog tile array for the Hughes Carlsbad CBiCMOS process.

July 1985–1986 Radar Systems Group, Hughes Aircraft El Segundo, CA

Member of Technical Staff

- Involved in design of high speed digital ICs for the VHSIC (Very High Speed Integrated Circuit) program using Fairchild Fast-Z Fineline 1.25µm bipolar.
- Involved in design of high speed 64 X 16, 124 X 24, and 1K X 24 SRAM ICs.
- Contributed to the merged junction bipolar SPICE model.
- Designed current mode logic for UHSBL (Ultra High Speed Bipolar Logic) cell library in NS Aspect-III. Cells used in high speed MUX, DMUX, and DDS ICs.
- Developed digital IC architectures for high speed radar signal processing.

Education

1980-1985 UCLA Westwood, CA

- BS Electrical Engineering Summa Cum Laude, Phi Beta Kappa
- UCLA's Most Outstanding Senior Electrical Engineering Student

1985-1987 UCLA Westwood, CA

- MS Electrical Engineering

1987-1989 USC Los Angeles, CA

- Engineer's Degree Electrical Engineering

1990-Present USC Los Angeles, CA

- PHD Candidate under the advisement of Prof. John Choma Jr.
- Completed Thesis: "Nonlinear Error Correction for the Bipolar Canonic Cells."
- Designed, fabricated, and DC wafer probed amplifier circuits in MAXIM SHPi bipolar process to verify theory developed.

Publications

- M. Chambers and L. Linder, "A Precision Monolithic Sample- And-Hold for Video Analog-to-Digital Converters," ISSCC Feb. 1991.
- B. Felder, et al., "A Low Noise 13 Bit 10 MSPS ADC Hybrid with High Dynamic Range," GOMAC 1994.
- W. Cheng, et al., "A 3 Bit, 40GSPS ADC- DAC in 0.12µm SiGe," ISSCC Feb. 2004.
- O. Panfilov, et al., "Direct Conversion Transceivers as a Promising Solution for Building Future Ad-hoc Networks," International Conference on Next Generation Tele-traffic and Wired / Wireless Advanced Networking September 2007.
- O. Panfilov, et al., "Test Results of the Direct Conversion Transceiver Demo Board", November 2007 SDR Forum Technical Conference.

- O. Panfilov, et al., “Overcoming Challenges of Direct Conversion Software Radio,” IEEE International Design and Test Workshop December 2007.
- A. Varghese and L.F. Linder, “Software Defined Radios for Wireless Handsets,” April 2008 Wireless Design & Development Magazine.
- S. Elahmadi, et al., “A Monolithic One-Sample / Bit Partial-Response Maximum Likelihood SiGe Receiver for Electronic Dispersion Compensation of 10.7 GB / s Fiber Channels,” OFC / NFOEC March 2009.
- S. Elahmadi, et al., “A 50 dB Dynamic Range, 11.3 GSPS, Programmable Finite Impulse Response (FIR) Equalizer in 0.18µm SiGe BiCMOS Technology for High Speed Electronic Dispersion Compensation (EDC) Applications,” RFIC Symposium, June 2009.
- J. Edwards, et al., “A 12.5 Gbps Analog timing Recovery System for PRML Optical Receivers,” RFIC Symposium, June 2009.
- D. Baranauskas, et al., “A 6th Order 1.6 to 3.2 GHz Tunable Low-Pass Linear Phase gm-C Filter for Fiber Optic Adaptive EDC Receivers,” RFIC Symposium, June 2009.
- S. Elahmadi, et al., “An Analog PRML Receiver for up to 400km of Uncompensated OC-192 Fiber-Optic Channels,” ESSCIRC September 2009.
- S. Elahmadi, et al., "An 11.1 Gbps Analog PRML Receiver for Electronic Dispersion Compensation of Fiber Optic Communications," IEEE JSSC, vol.45, no. 7, July 2010.
- Montierth, D., Strans, T., Leatham, J., Linder, L., and Baker, R. J., "Performance and Characteristics of Silicon Avalanche Photo detectors in the C5 Process," 2012 IEEE Midwest Symposium on Circuits and Systems, Boise, Idaho.
- Rauch M. and Linder L., "Collaborating with Nu-Trek," 2012 HiRev Industry Day, December 2012, Los Angeles, California.

Awards / Achievements

- IEEE Senior Member
- Two-time Hughes Aircraft Division Patent Award winner
- 55 issued US patents, over 300 international patents, several US patents pending
- Numerous IC design team awards at Hughes and Raytheon
- Hughes Masters Fellow, Engineers Fellow, Doctoral Fellow
- 2008 IEEE San Fernando Valley Section Entrepreneurial Business Plan Competition Judge
- Menara Networks EDC IC Team Award – world’s first error free operation over 400 km of uncompensated fiber

IC Process Experience

- CMOS, silicon / SiGe bipolar, silicon / SiGe BiCMOS, complementary bipolar, CBiCMOS, SOI, SOS, GaAs D-mode / E-mode MESFET, AlGaAs / InP HBT

Skills

- CADENCE Analog Artist Schematic Composer, Spectre Simulator

Security Clearance

- Secret clearance is presently supported by Nu-Trek.
- US Citizen

EXHIBIT D

US007396760B2

(12) **United States Patent**
Taravade et al.

(10) **Patent No.: US 7,396,760 B2**
(45) **Date of Patent: Jul. 8, 2008**

---

(54) **METHOD AND SYSTEM FOR REDUCING INTER-LAYER CAPACITANCE IN INTEGRATED CIRCUITS**

(75) Inventors: **Kunal N. Taravade**, Portland, OR (US); **Neal Callan**, Lake Oswego, OR (US); **Paul G. Filseth**, Los Gatos, CA (US)

(73) Assignee: **LSI Corporation**, Milpitas, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 388 days.

(21) Appl. No.: **10/991,107**

(22) Filed: **Nov. 17, 2004**

(65) **Prior Publication Data**

US 2006/0105564 A1 May 18, 2006

(51) **Int. Cl.**
H01L 21/4763 (2006.01)

(52) **U.S. Cl.** **438/626**; 438/622; 438/12; 438/17

(58) **Field of Classification Search** 438/10, 438/12, 14, 17, 618, 622, 626, 631, 645
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,763,955	A *	6/1998	Findley et al.	257/775
6,751,785	B1	6/2004	Oh	716/10
6,998,653	B2 *	2/2006	Higuchi	257/758
7,015,582	B2 *	3/2006	Landis	428/623
7,089,522	B2 *	8/2006	Tan et al.	716/11
7,152,215	B2 *	12/2006	Smith et al.	716/4
7,188,321	B2 *	3/2007	Wong et al.	716/2
2002/0190382	A1 *	12/2002	Kouno et al.	257/758
2005/0186751	A1 *	8/2005	Beach et al.	438/384

OTHER PUBLICATIONS

Hierarchical Dummy Fill for Process Uniformity; Yu Chen, Andrew B. Kahng, Gabriel Robins and Alexander Zelikovsky; Computer Science Department, UCLA, Los Angeles, CA 90095-1596; UCSD CSE and ECE Departments, La Jolla, CA 92093-0114; Department of Computer Science, University of Virginia, Charlottesville, VA 22903-2442; Department of Computer Science, Georgia State University, Atlanta, GA 30303.

Using Smart Dummy Fill and Selective Reverse Etchback for Pattern Density Equalization; Brian Lee, Duane S. Boning, Dale L. Hetherington and David J. Stein; Massachusetts Institute of Technology, Cambridge, MA, Sandia National Laboratories, Albuquerque, NM; Mar. 2000.

* cited by examiner

Primary Examiner—Kevin M Picardat
(74) *Attorney, Agent, or Firm*—Suiter Swantz PC LLO

(57) **ABSTRACT**

The present invention is directed to a method and system of intelligent dummy filling placement to reduce inter-layer capacitance caused by overlaps of dummy filling area on successive layers. The method and system treats each consecutive pair of layers together so as to minimize dummy filling overlaps between each layer. In particular, dummy fill features on each layer may be placed in a checkerboard pattern to avoid overlaps. As such, the present invention may eliminate large overlap area of the dummy patterns on consecutive layers by utilizing intelligent dummy filling placement.

19 Claims, 5 Drawing Sheets

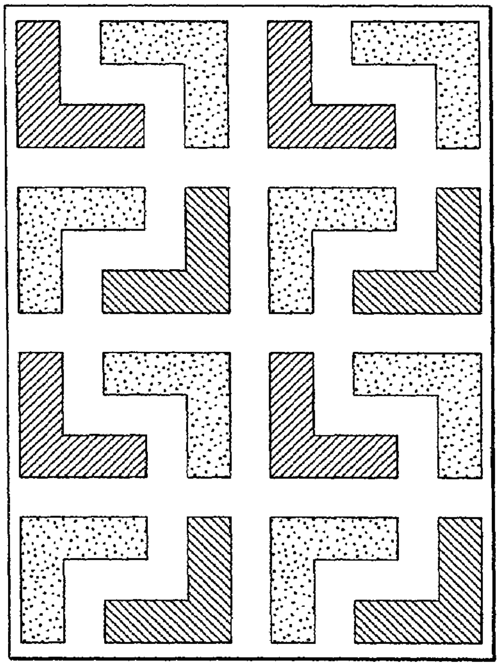

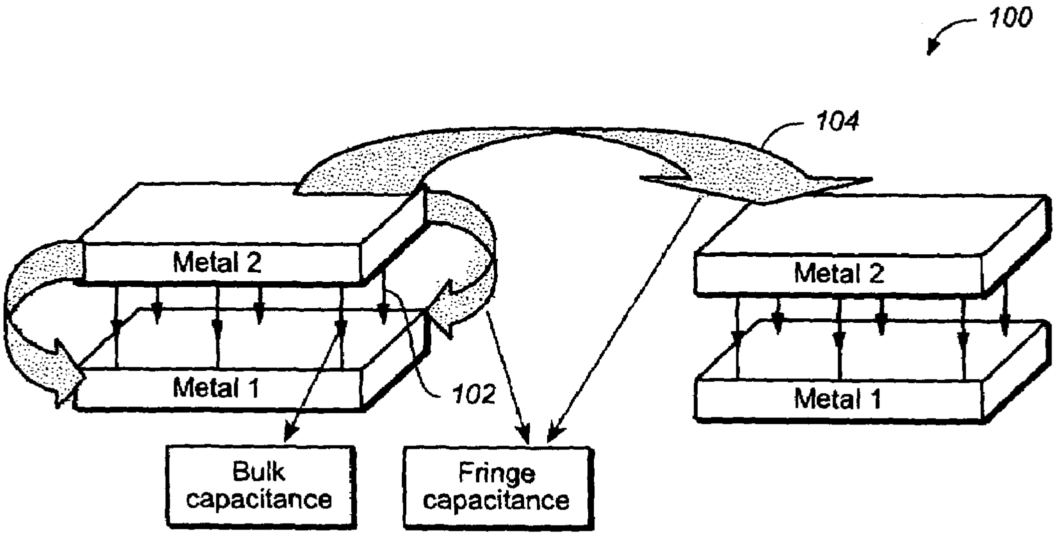
100
104
Metal 2
Metal 1
102
Metal 2
Metal 1
Bulk capacitance
Fringe capacitance

FIG._1

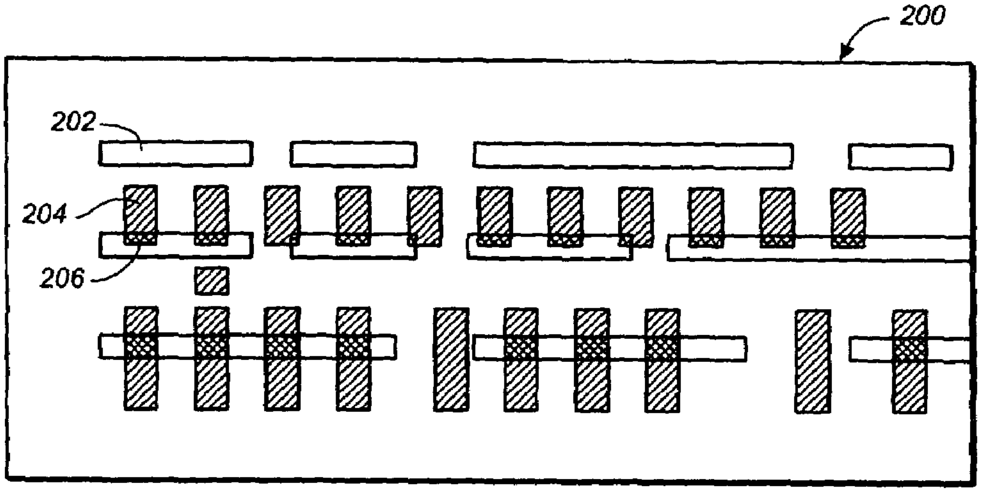
200
202
204
206

FIG._2

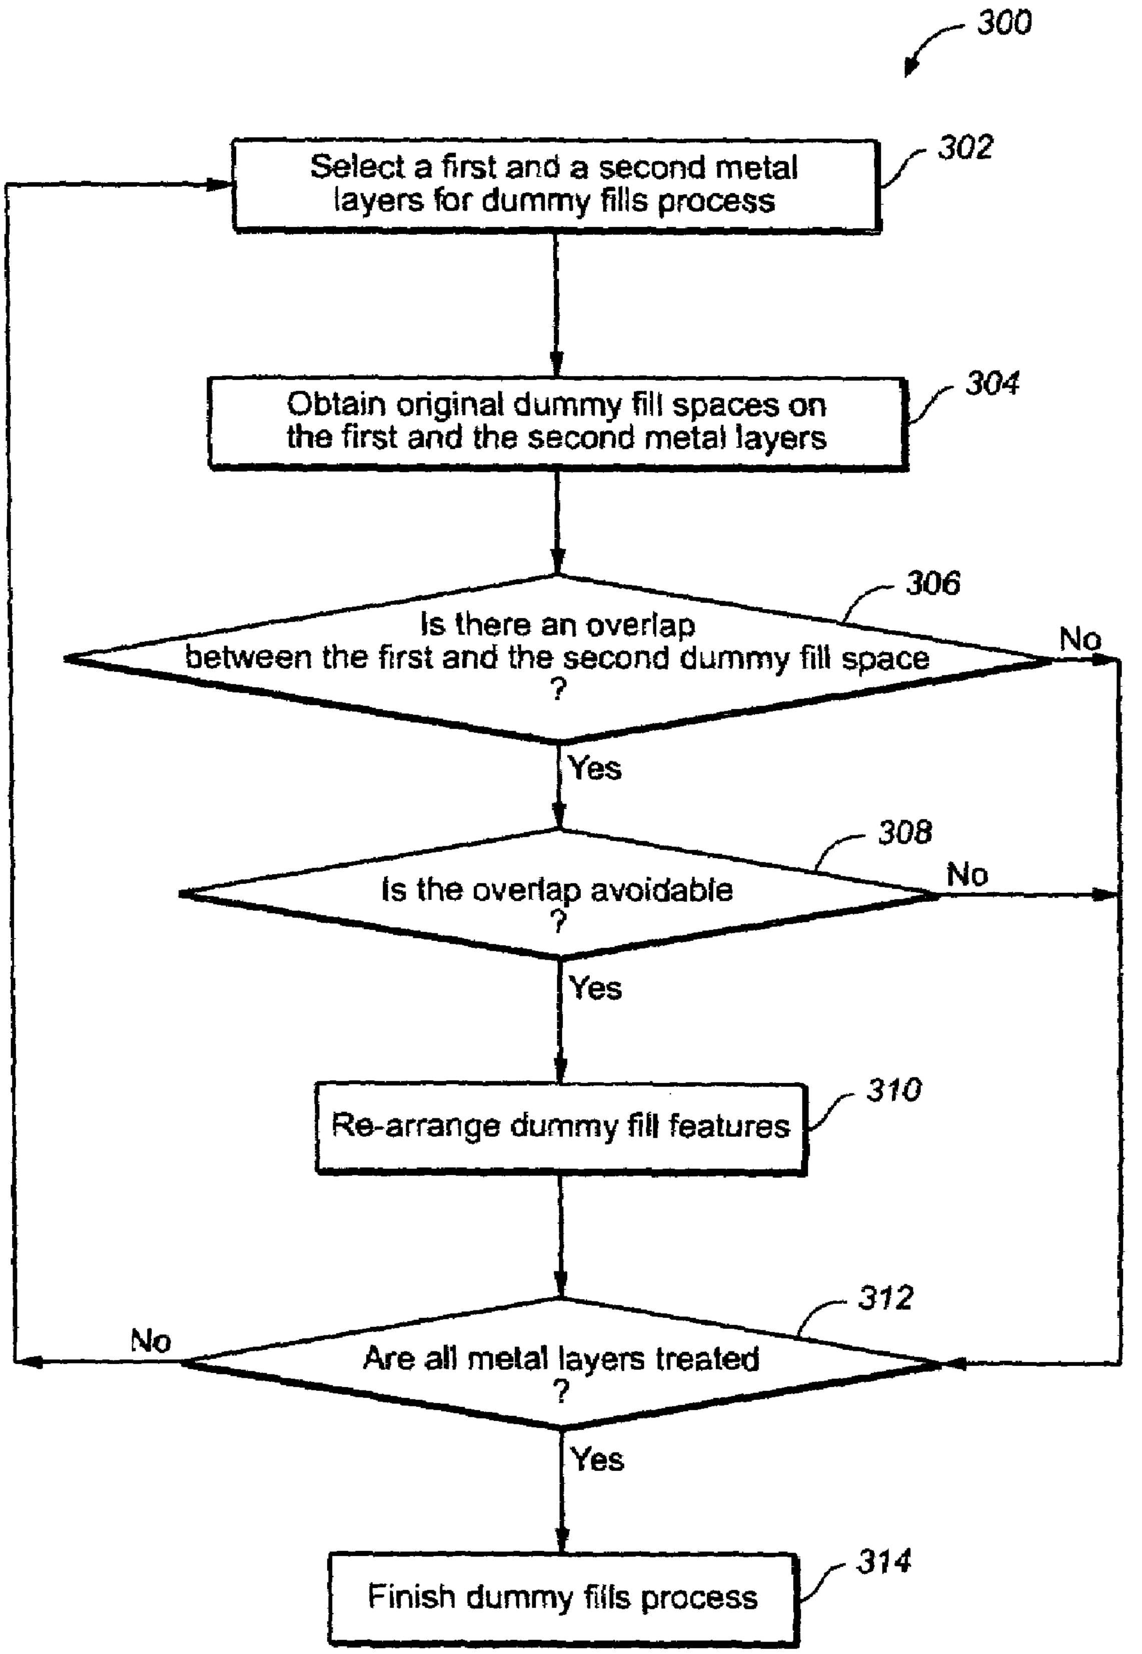
300
Select a first and a second metal layers for dummy fills process
302
Obtain original dummy fill spaces on the first and the second metal layers
304
Is there an overlap between the first and the second dummy fill space ?
306
No
Yes
Is the overlap avoidable ?
308
No
Yes
Re-arrange dummy fill features
310
Are all metal layers treated ?
312
No
Yes
Finish dummy fills process
314

FIG._3

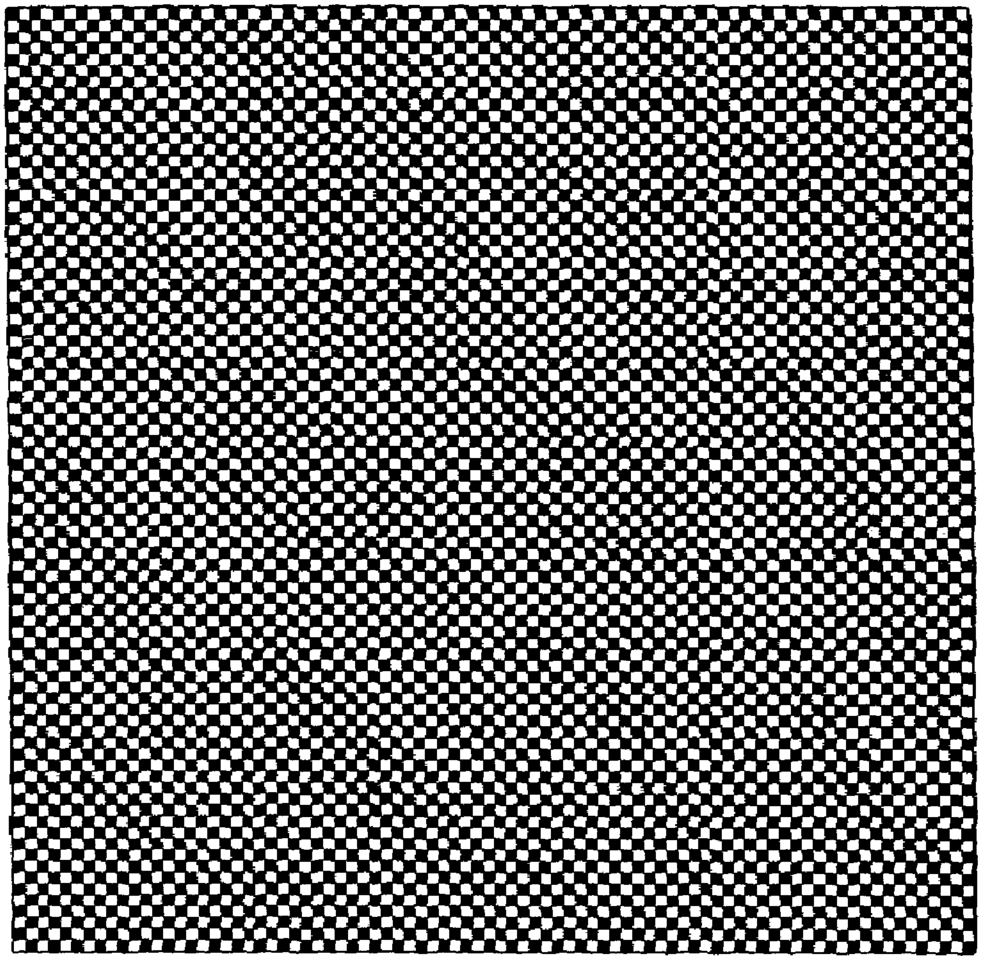

FIG._4

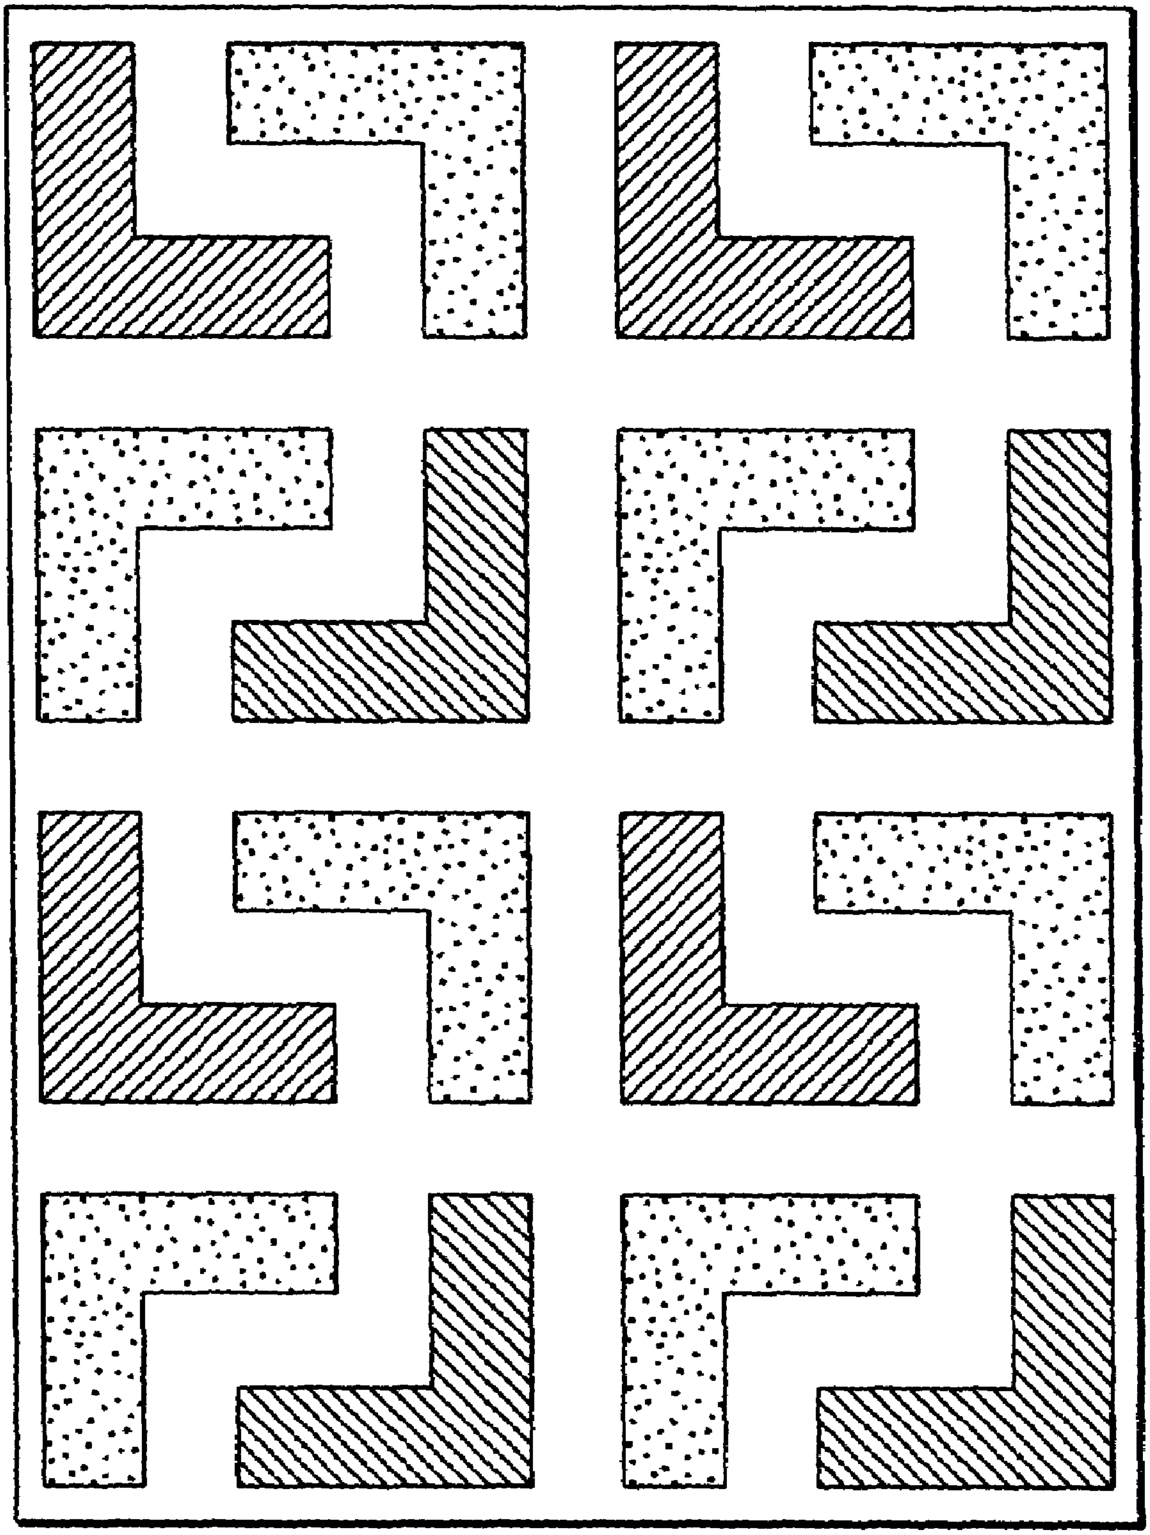

FIG._5

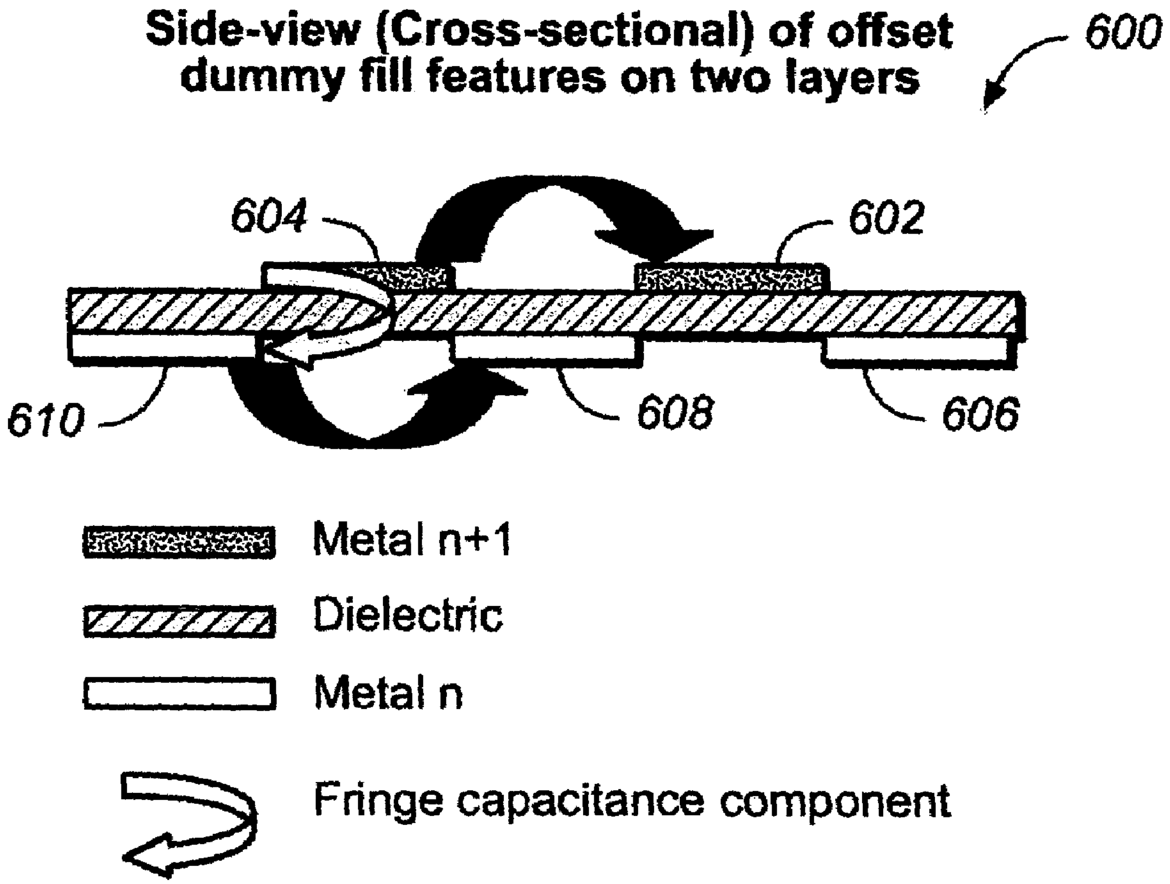
Side-view (Cross-sectional) of offset dummy fill features on two layers
600
604
602
610
608
606
Metal n+1
Dielectric
Metal n
Fringe capacitance component

FIG._6

METHOD AND SYSTEM FOR REDUCING INTER-LAYER CAPACITANCE IN INTEGRATED CIRCUITS

FIELD OF THE INVENTION

The present invention generally relates to the field of integrated circuit fabrication, and particularly to a method for reducing inter-layer capacitance through dummy fill methodology.

BACKGROUND OF THE INVENTION

In any integrated circuit, there is an inevitable capacitance that is introduced from electromagnetic interaction between electrical conductors, such as interconnect layers (metals). There are two components of such capacitance, a bulk (area) component and a fringe (peripheral) component. The bulk component is proportional to the overlap area of interconnect layers and the fringe component depends on the separation and the perimeter of adjacent interconnect layers. Referring now to FIG. **1**, the bulk capacitance **102** and the fringe capacitance **104** between Metal **1** and Metal **2** of an exemplary integrated circuit **100** are shown. The bulk capacitance generated due to the overlap of signal carrying lines on Metal **1** and Metal **2** may not be easily avoided since the placement of signal carrying lines is dictated by circuit functionality. However, the bulk capacitance introduced due to the overlap of non-signal carry lines may be reduced by changing the placement of non-signal carry lines.

An example of non-signal carry lines includes “dummy” fills which are utilized to even the topography and pattern density across the chip, prevent etch, or the like. “Dummy” fills refer to additional features to an integrated chip layout. In a typical integrated chip layout, there are unused areas on a layer after the signal, power and clock segments have been routed. These unused areas can be large enough such that additional features (metals) should be added to satisfy minimum metal coverage requirements for manufacturing. The “dummy” fills may be added to the unused areas such that subsequent layers on the integrated circuit are substantially planar.

For example, a dummy fills methodology is utilized in chemical mechanical polishing or planarization (CMP) process. Often, the planer profile resulting from the CMP process is dependent on the pattern density of the underlying layer. The density may vary and thus result in CMP planer profile variation. Such CMP planer profile variation may be reduced by employing the dummy fills methodology. In particular, dummy fills (dummy features) are inserted into a wafer prior to the CMP process so as to make the pattern density more uniform in IC chips. Uniform feature density improves wafer-processing uniformity for certain operations such as CMP. Dummy fills are typically placed according to conventional dummy fills methodologies that locate dummy fills where space is available. However, the conventional dummy fills methodologies allow a large planer profile variation. Some sophisticated dummy fills methodologies are utilized to reduce the large planer profile variation by selectively inserting dummy fills to achieve an effective density to within a predetermined range.

While most dummy fills methodologies have focused on uniform feature density, the problems created by the inserted dummy fills such as adverse effects on the electric field, unwanted bulk capacitance, and the like have not been addressed. Further, the existing dummy fill methodologies treat each layer independently which results in a large overlap over dummy fill areas on successive layers. Referring now to FIG. **2**, the overlaps **206** between Metal **1** dummy fill area **202** and Metal **2** dummy fill area **204** are shown. If the overlaps **206** are large, the unwanted bulk capacitance may be increased, thereby slowing down signals in the circuit and adversely affecting timing.

Therefore, it would be desirable to provide a method and system of intelligent dummy fill placement to reduce inter-layer capacitance caused by overlaps of dummy fill area on successive layers. It would be also desirable to provide a method and system for treating each consecutive pair of layers together when the intelligent dummy filling placement is performed.

SUMMARY OF THE INVENTION

Accordingly, the present invention provides a method and system for reducing inter-layer capacitance utilizing an intelligent dummy filling placement in integrated circuits.

In a first aspect of the present invention, a system for locating dummy fill features in an integrated circuit fabrication process is provided. The system may comprise an input for obtaining circuit layout information which provides initial signal lines on layers of the integrated circuit. The system may treat each successive pair of layers (a first layer and a second layer) together. The system may comprise a means for defining dummy fill features including small squares within the dummy fill space. The dummy fill spaces are suitable to have dummy fill features inserted. The dummy fill spaces may include areas where dummy patterns are intended to be placed on the first layer and the second layer. Then, the system may assign alternating dummy fill features to each layer in order to avoid overlaps between dummy fill features on each layer.

In a second aspect of the present invention, a method of placing dummy fill patterns to minimize inter-layer capacitance in an integrated circuit fabrication process is provided. The integrated circuit may include many interconnect layers (metals). The method may treat each consecutive pair of layers (a first layer and a second layer) together. Layout information of the integrated circuit may be obtained to determine an initial dummy fill space for a first layer and a second layer. Whether there are overlaps between the initial dummy fill space on the first layer and the initial dummy fill space on the second dummy fill space may be determined. If the overlaps are found and avoidable by re-arranging dummy fill patterns, a first dummy fill pattern and a second dummy fill pattern may be re-arranged to minimize the overlaps.

Additionally, the first dummy fill pattern may be placed to form a checkerboard pattern. If the first layer is already arranged in the form of a checkerboard pattern, the first dummy fill pattern may not be re-arranged. Then, the second dummy fill pattern may be placed to form a checkerboard pattern so as to be offset from the first dummy fill pattern. In this manner, each of the dummy fill features on the first layer may not be placed directly above dummy fill features on the second layer. Consequently, the unwanted bulk capacitance introduce by the dummy fill may be reduced and thus the inter-layer capacitance is minimized.

It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention as claimed. The accompanying drawings, which are incorporated in and constitute a part of the specification, illustrate an embodiment of the invention and together with the general description, serve to explain the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

The numerous advantages of the present invention may be better understood by those skilled in the art by reference to the accompanying figures in which:

FIG. **1** is an illustration of fringe and bulk capacitance components in an exemplary integrated circuit having Metal **1** and Metal **2** layers;

FIG. **2** is an illustration of layout image showing overlaps of dummy fill areas of Metal **1** and Metal **2** layers in FIG. **1**;

FIG. **3** is a flow diagram illustrating a method implemented in accordance with an exemplary embodiment of the present invention wherein two consecutive layers are treated;

FIG. **4** is a top view of a layer showing a checkerboard pattern formed by the method described in FIG. **3**;

FIG. **5** is a top view of two layers showing an alternative pattern with which the present invention can be embodied; and

FIG. **6** is a cross-sectional view of two layers showing offset dummy fill features inserted by the method described in FIG. **3**.

DETAILED DESCRIPTION OF THE INVENTION

Reference will now be made in detail to the presently preferred embodiments of the invention, examples of which are illustrated in the accompanying drawings.

Referring generally now to FIGS. **3** through **5**, exemplary embodiments of the present invention are shown.

The present invention is directed to a method and system of intelligent dummy filling placement to reduce inter-layer capacitance caused by overlaps of dummy fills on successive layers. Generally, dummy fill refers to the addition of features to a layout for the purpose of raising the density of specific regions on the layout of the integrated circuit. The method and system treats each consecutive pair of layers together so as to minimize the overlaps of dummy fills between each layer. In particular, dummy fill features on each layer may be placed in a checkerboard pattern to avoid overlaps. As such, the present invention may eliminate large overlap areas of the dummy fills on consecutive layers by utilizing intelligent dummy fill placement. In the following description, numerous specific descriptions are set forth in order to provide a thorough understanding of the present invention. It should be appreciated by those skilled in the art that the present invention may be practiced without some or all of these specific details. In some instances, well known process operations have not been described in detail in order to prevent obscurity of the present invention.

Referring now to FIG. **3**, a flow diagram **300** illustrating a method implemented in accordance with an exemplary embodiment of the present invention wherein a dummy fill process is performed on each layer of an integrated circuit is shown. Generally, an integrated circuit fabrication process involves a series of layering processes in which metallization, dielectrics, and other materials are applied to the surface of a semiconductor wafer to form a layered interconnected structure (an interconnect layer). The integrated circuits generally include inter-layered circuits comprising a plurality of metal lines across multiple layers that are interconnected by metal-filled vias. The method begins in step **302** in which a first layer and a second layer are selected for dummy fill process. The first layer and the second layer are a consecutive pair of layers of the IC.

Generally, dummy fills are utilized to improve planer profile uniformity by helping to level the feature density across the layout during an integrated circuit fabrication process. For example, dummy fills are utilized in chemical mechanical polishing or planarization (CMP) process. Often, the planer profile resulting from the CMP process is dependent on the pattern density of the underlying layer. The dependency may vary and thus offset the CMP planer profile variation. Such variation may be reduced by employing the dummy fills methodology. In particular, dummy fills (dummy features) are inserted into a wafer prior to the CMP process so as to make the pattern density more uniform in IC chips. Uniform feature density improves wafer-processing uniformity for certain operations such as CMP. Placement of the dummy fills is typically made according to conventional dummy fill methodologies that locate the uniform-densities dummy where space is available. However, the inserted dummy fills may create problems such as adverse effects on the electric field, unwanted bulk capacitance, and the like.

In Step **304**, the original (initial) dummy fill spaces of the first layer and the second layer may be obtained based on layout information. The layout information may be provided by a user, an IC fabrication process system, a CAD tool, or the like. The original dummy fill space may include areas where dummy fill patterns are intended to be placed on layers. Then, in Step **306**, whether there is any overlap between the original dummy fill space of the first layer and the original dummy fill space of the second layer may be determined. The overlaps of dummy fill areas between the first layer and the second layer are undesirable since the unwanted bulk capacitance may be introduced by the overlaps. Thus, in step **308**, whether the overlap can be avoided by re-arrangement of dummy features may be checked. Then, dummy fill patterns on the first layer and the second layer may be re-arranged to minimize the overlaps in Step **310**. In a particular embodiment of the present invention, a grid (composed of small squares) may be defined within the dummy fill spaces. The method may assign alternating squares (dummy fill features) in the grid to each layer. In this manner, dummy fill features on the first layer are not placed directly above the ones on the second layer but offset from each other. It is to be noted that the dummy fill features may be placed to form various predefined patterns designed to prevent overlaps on adjacent layers. Referring now to FIG. **4**, an exemplary top view of a layer showing a checkerboard pattern formed by the present invention is shown. As shown in FIG. **4**, dummy fill features placed in a checkerboard pattern may avoid overlap, thereby reducing the bulk capacitance component of the total capacitance. Preferably, the dummy fill features are placed to form a checkerboard pattern. Referring now to FIG. **5**, an exemplary top view of two layers showing a different pattern with which the present invention can be embodied is shown.

Referring back to FIG. **3**, if there is no overlap found, the method may proceed to check whether all interconnect layers in the IC have been treated in Step **312**. If all interconnect layers have been treated, the method may finish the dummy fill pattern placement in step **314**. If all interconnect layers have not been treated, the method may proceed to step **302** by selecting the next pair of consecutive layers.

Additionally, the method may check whether the first layer is already arranged in the form of a checkerboard. If the first layer includes dummy fill pattern in the form of a checkerboard, the dummy fill pattern on first layer may not be re-arranged. The dummy fill pattern on the second layer may be re-arranged to form a checked board pattern by offsetting against the dummy fill pattern on the first layer.

One of skill in the art will appreciate that there are various ways to check the form of the dummy fill pattern. In a particular embodiment, numbers may be assigned to dummy features in order to check whether the dummy fill pattern is already in the form of a checkerboard pattern. For example, a dummy feature may have a row number, a column number, and a layer number. The dummy fill pattern may be checked by implementation of a simple Boolean check as follows: Pattern checking number=row number+column number+layer number. Each dummy feature may have a pattern checking number. The numbering scheme for the simple Boolean check may be assigned such that the pattern checking number is always odd for given row number, column number and layer number. As such, the dummy fill features on the first layer and the second layer are placed on alternating row and column combinations. Additionally, the simple Boolean check may be utilized to determine whether to re-arrange dummy features on the layer.

In FIG. **6**, a cross-sectional view **600** of two layers showing offset dummy fill features inserted by the present invention is shown. The first dummy features **602**, **604** is arranged to offset the second dummy features **606**-**610**. The checkerboard style layout of the dummy fill pattern prevents situations in which dummy patterns on successive layers overlap, thereby increasing parasitic capacitance of the circuit by adding bulk (area) capacitance of the chip in proportion to the overlap area of the dummy patterns on consecutive layers. By offsetting the dummy patterns in a checkerboard fashion, the large bulk capacitance component may be eliminated. As a result, the total capacitance for an integrated circuit may be reduced.

Generally, the total capacitance for an integrated circuit composed of interconnect layers (metals) may be given by:

$$C_{TOTAL}=C_{BULK}+C_{FRINGE}$$

where C_{BULK}=Bulk intra-layer capacitance (bulk capacitance of metal lines on the same layer)+Bulk inter-layer Capacitance (bulk capacitance of metal lines on adjacent layers) and

C_{FRINGE}=Fringe intra-layer capacitance (fringe capacitance of metal lines on the same layer)+Fringe inter-layer Capacitance (fringe capacitance of metal lines on adjacent layers).

In a particular embodiment of the present invention, the above-described method and system may be implemented through various commercially available polygon manipulation languages. An example of the commercially available polygon manipulation languages may include, but are not limited to, Mentor Graphics® Calibre®, Synopsys® Hercules® or the like.

It should be noted that the method and system of the present invention may be utilized for wafer processing operations such as CMP. However, the method and the system of the present invention may be utilized for any suitable integrated circuit fabrication process.

In the exemplary embodiments, the methods disclosed may be implemented as sets of instructions or software readable by a device. Further, it is understood that the specific order or hierarchy of steps in the methods disclosed are examples of exemplary approaches. Based upon design preferences, it is understood that the specific order or hierarchy of steps in the method can be rearranged while remaining within the scope and spirit of the present invention. The accompanying method claims present elements of the various steps in a sample order, and are not necessarily meant to be limited to the specific order or hierarchy presented.

It is believed that the method and system of the present invention and many of its attendant advantages will be understood by the forgoing description. It is also believed that it will be apparent that various changes may be made in the form, construction and arrangement of the components thereof without departing from the scope and spirit of the invention or without sacrificing all of its material advantages. The form herein before described being merely an explanatory embodiment thereof. It is the intention of the following claims to encompass and include such changes.

What is claimed is:

1. A method for placing dummy fill patterns in an integrated circuit fabrication process, comprising:

obtaining layout information of the integrated circuit, the integrated circuit including a plurality of layers;

obtaining a first dummy fill space for a first layer based on the layout information;

obtaining a second dummy fill space for a second layer, the second layer being placed successively to the first layer;

determining an overlap between the first dummy fill space and the second dummy fill space; and

minimizing the overlap by re-arranging a plurality of first dummy fill features and a plurality of second dummy fill features,

wherein the first dummy fill space includes non-signal carrying lines on the first layer and the second dummy fill space includes non-signal carrying lines on the second layer.

2. The method as described in claim **1**, wherein the plurality of first dummy fill features forms a grid within the first dummy fill space.

3. The method as described in claim **1**, wherein the plurality of second dummy fill features forms a grid within the second dummy fill space.

4. The method as described in claim **1**, wherein the first dummy fill space is determined based on a local pattern density for the first layer.

5. The method as described in claim **1**, wherein the second dummy fill space is determined based on a local pattern density for the second layer.

6. The method as described in claim **2**, wherein the grid includes a plurality of squares.

7. The method as described in claim **1**, the minimizing the overlap step further comprising:

determining whether the plurality of first dummy fill features form a predefined pattern; and

re-arranging the plurality of first dummy fill features to form the predefined pattern if the plurality of first dummy fill features are not arranged in the predefined pattern.

8. The method as described in claim **7**, further comprising:

re-arranging the plurality of second dummy fill features based on the plurality of first dummy features if the plurality of first dummy fill features are already arranged in the predefined pattern.

9. The method as described in claim **8**, wherein the plurality of second dummy fill features are re-arranged so as to be offset from the plurality of first dummy fill features.

10. The method as described in claim **7**, wherein the predefined pattern is a checkerboard pattern.

11. The method as described in claim **1**, wherein a total bulk capacitance is minimized.

12. The method as described in claim **11**, wherein the total bulk capacitance includes a bulk inter-layer capacitance.

13. The method as described in claim **11**, wherein the bulk inter-layer capacitance is a bulk capacitance created by overlaps between the first layer and the second layer.

14. A method of filling dummy patterns for pattern density equalization in an integrated circuit fabrication process, comprising:

obtaining a local density pattern of a first layer, the local density pattern obtained based on an initial layout design of the integrated circuit;

determining a second layer, the second layer being placed successively to the first layer;

obtaining a local density pattern of the second layer, the local density pattern obtained based on the initial layout design of the integrated circuit;

designing a plurality of dummy fill features on the first layer and the second layer, the plurality of dummy fill features being suitable for increasing pattern density in low density spaces on the first layer and the second layer;

determining whether there is an overlap between the plurality of dummy fill features on the first layer and the plurality of dummy fill features on the second layer; and

minimizing the overlap by re-arranging the plurality of dummy fill features on the first layer and the second layer,

wherein a total inter-layer capacitance of the integrated circuit is minimized.

15. The method as described in claim **14**, the minimizing the overlap step further comprising:

determining whether the plurality of first dummy fill feature form a checkerboard pattern; and

placing the plurality of first dummy fill features to form the checkerboard pattern base through a mathematical check if the plurality of first dummy fill features are not a form of the checkerboard pattern,

wherein the mathematical check is applied to numeric values of each of the plurality of first dummy fill features and the numeric values of each of the plurality of first dummy fill features are determined based on the location on the checkerboard pattern.

16. The method as described in claim **14**, further comprising:

placing the plurality of second dummy fill features based on an arrangement of the plurality of first dummy features if the plurality of first dummy fill features form a checkerboard pattern.

17. The method as described in claim **16**, wherein the plurality of second dummy fill features are placed so as to form an alternate checkerboard pattern against the checkerboard pattern of the plurality of first dummy fill features.

18. The method as described in claim **16**, wherein the plurality of second dummy fill features are placed so as to be offset from the plurality of first dummy fill features.

19. The method as described in claim **14**, further comprising:

placing the plurality of second dummy fill features to form the checkerboard pattern base through a mathematical check,

wherein the mathematical check is applied to numeric values of each of the plurality of second dummy fill features and the numeric values of each of the plurality of second dummy fill features are determined based on the location on the checkerboard pattern.

* * * * *

EXHIBIT E

BELL SEMICONDUCTOR LLC'S ANALYSIS OF INFRINGEMENT

U.S. Patent No. 7,396,760

Claims 1–6 & 11–13

Bell Semiconductor ("Bell Semic") provides evidence of infringement of exemplary claims 1–6 & 11–13 of U.S. Patent No. 7,396,760 ("the '760 patent by the AC7-M128-30 Altra Max produced by Ampere Computing, LLC ("Ampere"). In support thereof, Bell Semic provides the following claim charts.

"Accused Products" as used herein refers to at least devices produced or sold by Ampere that are or include semiconductor integrated circuit devices made using a design tool, that are made, produced, and/or processed by a design tool, such as a Cadence Design Systems, Inc. ("Cadence"), Synopsys, Inc. ("Synopsys"), and/or Siemens Digital Industries Software (formerly Mentor Graphics) ("Siemens") tool, by rearranging dummy fill features to minimize their overlap when viewed across adjacent layers.[1] On information and belief, these design tools all function similarly with respect to the functionality described herein. For simplicity, the Cadence tool will be the primary tool cited herein to illustrate infringement of the claimed methods. These claim charts demonstrate infringement by comparing each element of the asserted claims to corresponding components, aspects, and/or features of the Accused Products. These claim charts are not intended to constitute an expert report on infringement. These claim charts include information provided by way of example, and not by way of limitation.

The analysis set forth below is based only upon information from publicly available resources regarding the Accused Products, as Ampere and relevant third parties have not yet provided any non-public information. An analysis of non-public technical documentation may assist in further identifying all infringing features and functionality. Accordingly, Bell Semic reserves the right to supplement this infringement analysis once such information is made available to Bell Semic. Furthermore, Bell Semic reserves the right to revise this infringement analysis, as appropriate, upon issuance of a court order construing any terms recited in the asserted claims or as other circumstances so merit.

Bell Semic contends that each element of each claim asserted herein is literally met, and would also be met under the doctrine of equivalents, as there are no substantial differences between the Accused Products and the elements of the patent claims in function, way, and result. Ampere directly infringes the asserted claims of the '760 patent by performing each of the limitations. If Ampere attempts to argue that there is no literal infringement and/or if Ampere attempts to draw any distinction between the claimed functionality and the Accused Products, then Bell Semic reserves the right to rebut the alleged distinction as a matter of literal infringement and/or as to whether any such distinction is substantial under the doctrine of equivalents.

Unless otherwise noted, the cited evidence applies across each of Ampere's products that were made, produced, or processed from a circuit design using windows, including but not limited to the AC7-M128-30 Altra Max, which includes multiple ARM cores. Bell Semic reserves the right to amend this infringement analysis based on other products made, produced, or processed in the same or similar manner to that identified herein.

---

[1] Ampere is a customer of at least Cadence, as demonstrated here: https://www.youtube.com/watch?v=CUB2Xn6Erd0.

Ampere is the producer and/or seller of the referenced above, as demonstrated by the following package images for the AC7-M128-30 Altra Max.

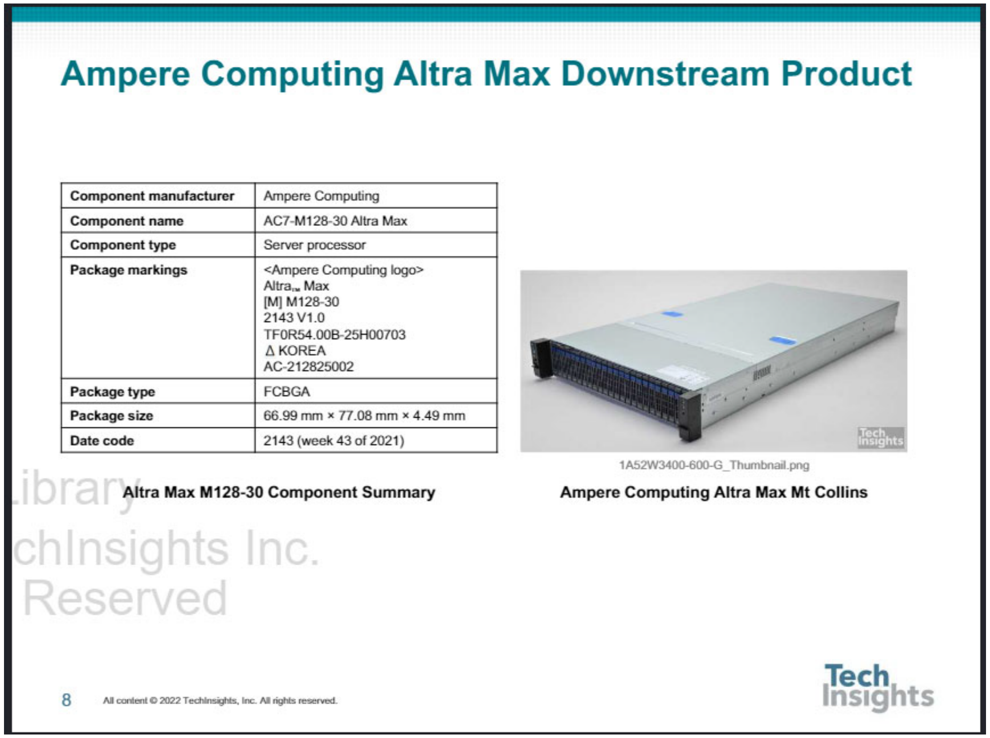

Ampere Computing Altra Max Downstream Product

Component manufacturer	Ampere Computing
Component name	AC7-M128-30 Altra Max
Component type	Server processor
Package markings	<Ampere Computing logo> Altra™ Max [M] M128-30 2143 V1.0 TF0R54.00B-25H00703 Δ KOREA AC-212825002
Package type	FCBGA
Package size	66.99 mm × 77.08 mm × 4.49 mm
Date code	2143 (week 43 of 2021)

Altra Max M128-30 Component Summary

1A52W3400-600-G_Thumbnail.png

Ampere Computing Altra Max Mt Collins

Claim 1	Accused Products
1. A method for placing dummy fill patterns in an integrated circuit fabrication process, comprising:	To the extent the preamble is limiting, the Accused Products are produced by performing a method for placing dummy fill patterns in an integrated circuit fabrication process: 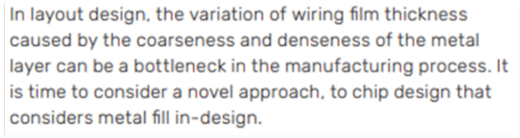 In layout design, the variation of wiring film thickness caused by the coarseness and denseness of the metal layer can be a bottleneck in the manufacturing process. It is time to consider a novel approach, to chip design that considers metal fill in-design. *See https://www.cadence.com/content/dam/cadence-www/global/en_US/documents/tools/digital-design-signoff/pegasus-tb.pdf, page 1.* 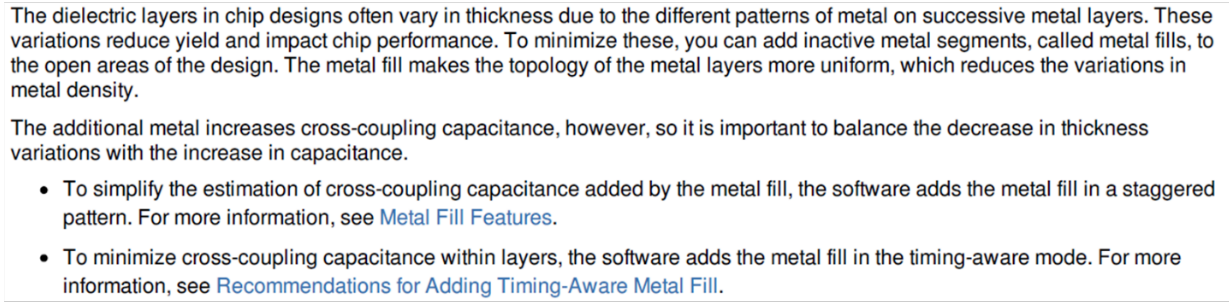 The dielectric layers in chip designs often vary in thickness due to the different patterns of metal on successive metal layers. These variations reduce yield and impact chip performance. To minimize these, you can add inactive metal segments, called metal fills, to the open areas of the design. The metal fill makes the topology of the metal layers more uniform, which reduces the variations in metal density. The additional metal increases cross-coupling capacitance, however, so it is important to balance the decrease in thickness variations with the increase in capacitance. • To simplify the estimation of cross-coupling capacitance added by the metal fill, the software adds the metal fill in a staggered pattern. For more information, see Metal Fill Features. • To minimize cross-coupling capacitance within layers, the software adds the metal fill in the timing-aware mode. For more information, see Recommendations for Adding Timing-Aware Metal Fill. *See Innovus User Guide product version 20.10, March 2020, page 705.* For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that places dummy fill patterns in an integrated circuit fabrication process. This is further explained by semiconductor expert Dhaval Brahmbhatt ("Brahmbhatt") in Exhibit F cited herein, *See* Ex. F at ¶¶ 72-76.
obtaining layout information of the integrated circuit, the integrated circuit including a plurality of layers;	The Accused Products are made, produced, or processed from a circuit design that is created by obtaining layout information of the integrated circuit, the integrated circuit including a plurality of layers.

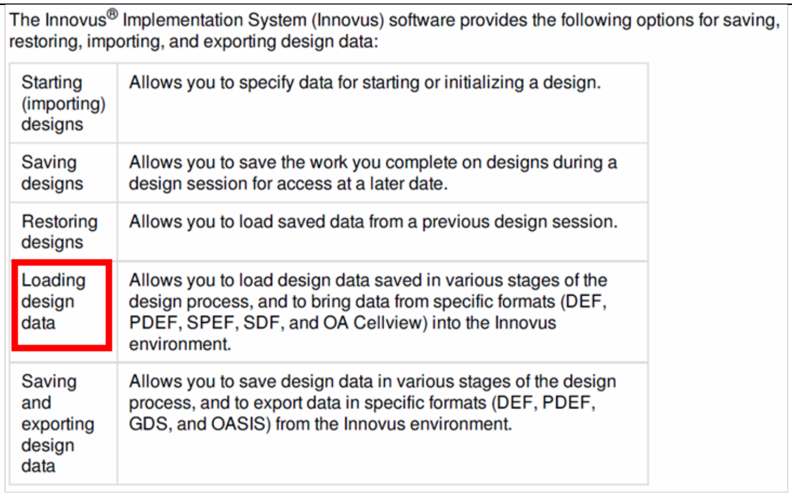

The Innovus® Implementation System (Innovus) software provides the following options for saving, restoring, importing, and exporting design data:

Starting (importing) designs	Allows you to specify data for starting or initializing a design.
Saving designs	Allows you to save the work you complete on designs during a design session for access at a later date.
Restoring designs	Allows you to load saved data from a previous design session.
Loading design data	Allows you to load design data saved in various stages of the design process, and to bring data from specific formats (DEF, PDEF, SPEF, SDF, and OA Cellview) into the Innovus environment.
Saving and exporting design data	Allows you to save design data in various stages of the design process, and to export data in specific formats (DEF, PDEF, GDS, and OASIS) from the Innovus environment.

See Innovus User Guide product version 20.10, March 2020, page 219.

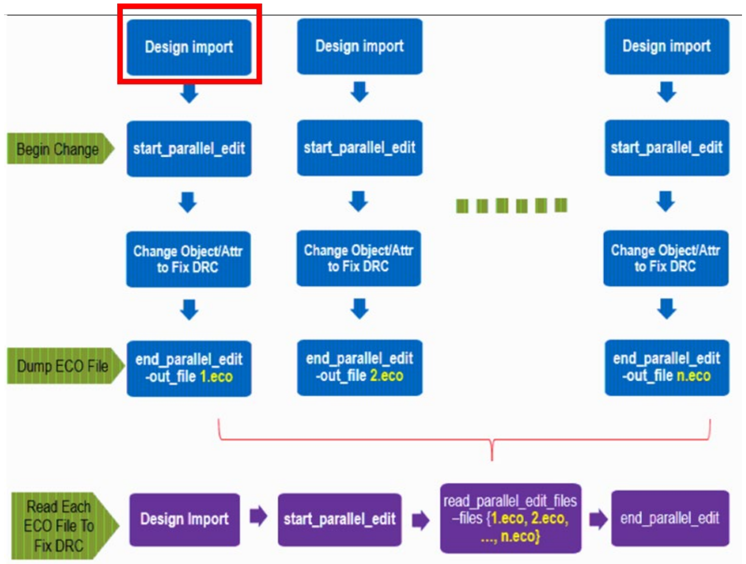

See Innovus User Guide product version 20.10, March 2020, page 1583.

	For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created by loading design information of the integrated circuit. The integrated circuit includes multiple layers. *See* Ex. F at ¶¶ 70-72.

obtaining a first dummy fill space for a first layer based on the layout information;	The Accused Products are made, produced, or processed from a circuit design that is created by obtaining a first dummy fill space for a first layer based on the layout information.
	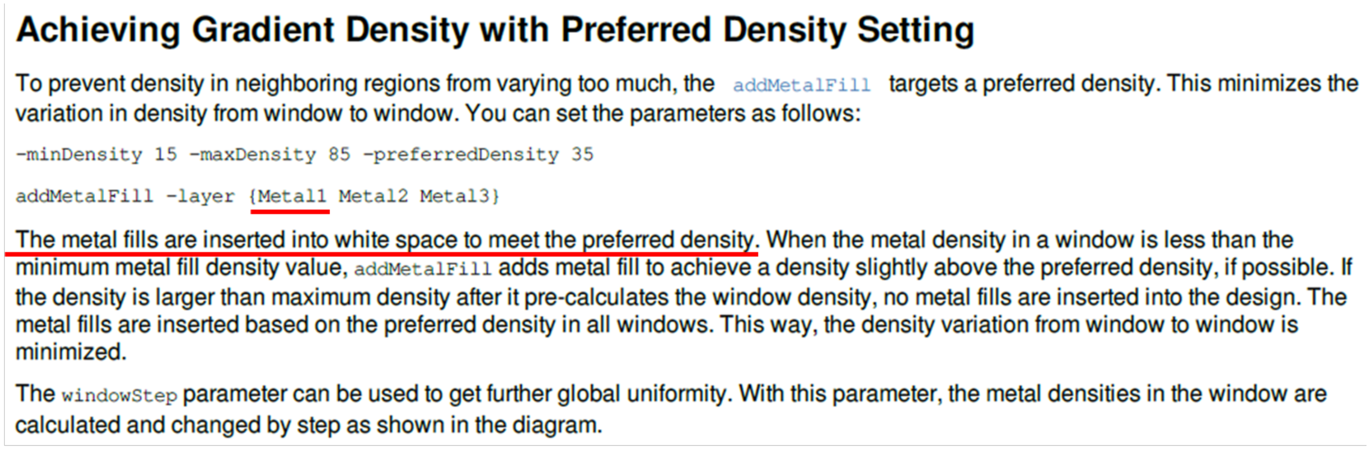
	Achieving Gradient Density with Preferred Density Setting
	To prevent density in neighboring regions from varying too much, the `addMetalFill` targets a preferred density. This minimizes the variation in density from window to window. You can set the parameters as follows:
	`-minDensity 15 -maxDensity 85 -preferredDensity 35`
	`addMetalFill -layer {Metal1 Metal2 Metal3}`
	The metal fills are inserted into white space to meet the preferred density. When the metal density in a window is less than the minimum metal fill density value, `addMetalFill` adds metal fill to achieve a density slightly above the preferred density, if possible. If the density is larger than maximum density after it pre-calculates the window density, no metal fills are inserted into the design. The metal fills are inserted based on the preferred density in all windows. This way, the density variation from window to window is minimized.
	The `windowStep` parameter can be used to get further global uniformity. With this parameter, the metal densities in the window are calculated and changed by step as shown in the diagram.
	See Innovus User Guide product version 20.10, March 2020, page 719 .
	For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created by obtaining a metal fill space for a first layer based on the loaded design information. *See* Ex. F at ¶ 75.

obtaining a second dummy fill space for a second layer, the second layer being placed successively to the first layer;	The Accused Products are made, produced, or processed from a circuit design that is created by obtaining a second dummy fill space for a second layer placed successively to the first layer.

	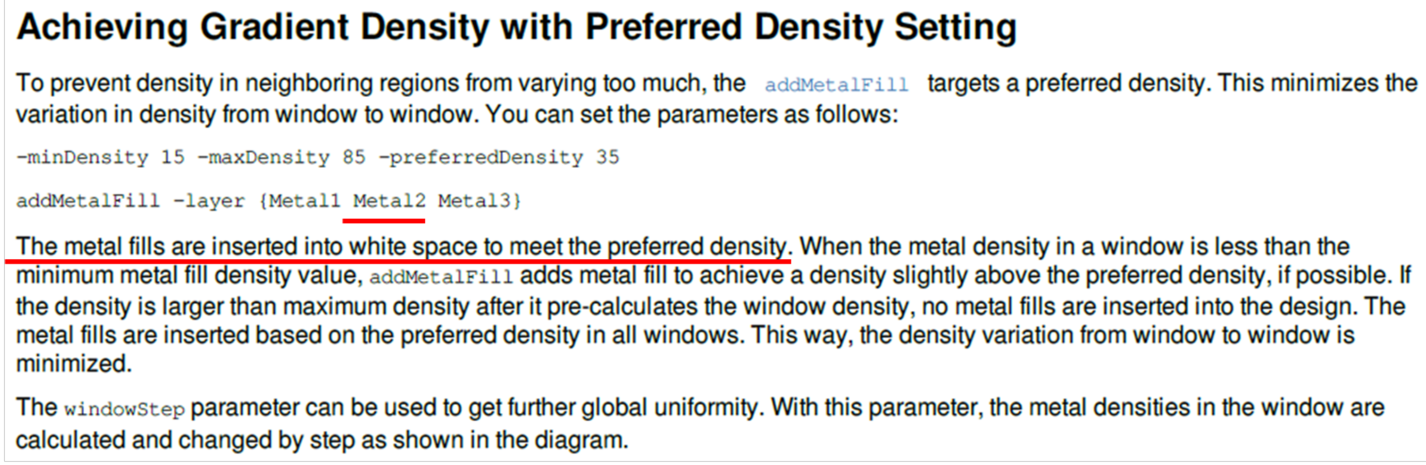
	Achieving Gradient Density with Preferred Density Setting
	To prevent density in neighboring regions from varying too much, the `addMetalFill` targets a preferred density. This minimizes the variation in density from window to window. You can set the parameters as follows:
	`-minDensity 15 -maxDensity 85 -preferredDensity 35`
	`addMetalFill -layer {Metal1 Metal2 Metal3}`
	The metal fills are inserted into white space to meet the preferred density. When the metal density in a window is less than the minimum metal fill density value, `addMetalFill` adds metal fill to achieve a density slightly above the preferred density, if possible. If the density is larger than maximum density after it pre-calculates the window density, no metal fills are inserted into the design. The metal fills are inserted based on the preferred density in all windows. This way, the density variation from window to window is minimized.
	The `windowStep` parameter can be used to get further global uniformity. With this parameter, the metal densities in the window are calculated and changed by step as shown in the diagram.
	See Innovus User Guide product version 20.10, March 2020, page 719 .
	For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created by obtaining a metal fill space for a second layer, where the second layer is placed successively to the first layer. *See* Ex. F at ¶ 75.

determining an overlap between the first dummy fill space and the second dummy fill space; and	The Accused Products are made, produced, or processed from a circuit design that is created by determining an overlap between the first dummy fill space and the second dummy fill space. 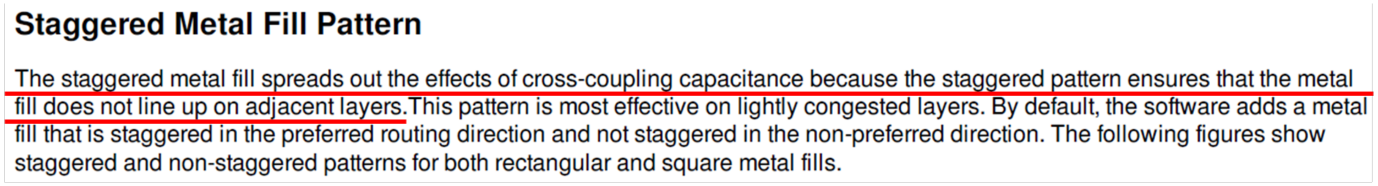 **Staggered Metal Fill Pattern** The staggered metal fill spreads out the effects of cross-coupling capacitance because the staggered pattern ensures that the metal fill does not line up on adjacent layers.This pattern is most effective on lightly congested layers. By default, the software adds a metal fill that is staggered in the preferred routing direction and not staggered in the non-preferred direction. The following figures show staggered and non-staggered patterns for both rectangular and square metal fills. *See Innovus User Guide product version 20.10, March 2020, page 706 .* For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that created by determining overlap between the first and second metal fill spaces. The only way to stagger metal fill is to first determine where there is overlap in metal fill and then to rearrange it to be staggered. *See* Ex. F at ¶ 75.

minimizing the overlap by re-arranging a plurality of first dummy fill features and a plurality of second dummy fill features,	The Accused Products are made, produced, or processed from a circuit design that is created by minimizing the overlap by re-arranging a plurality of first dummy fill features and a plurality of second dummy fill features. 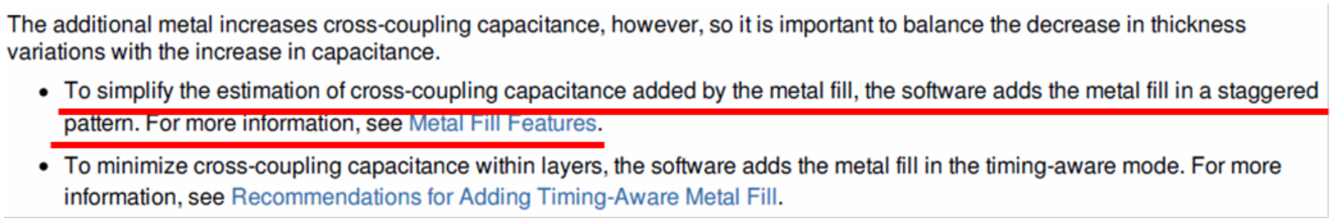 The additional metal increases cross-coupling capacitance, however, so it is important to balance the decrease in thickness variations with the increase in capacitance. • To simplify the estimation of cross-coupling capacitance added by the metal fill, the software adds the metal fill in a staggered pattern. For more information, see Metal Fill Features. • To minimize cross-coupling capacitance within layers, the software adds the metal fill in the timing-aware mode. For more information, see Recommendations for Adding Timing-Aware Metal Fill. *See Innovus User Guide product version 20.10, March 2020, page 705* .

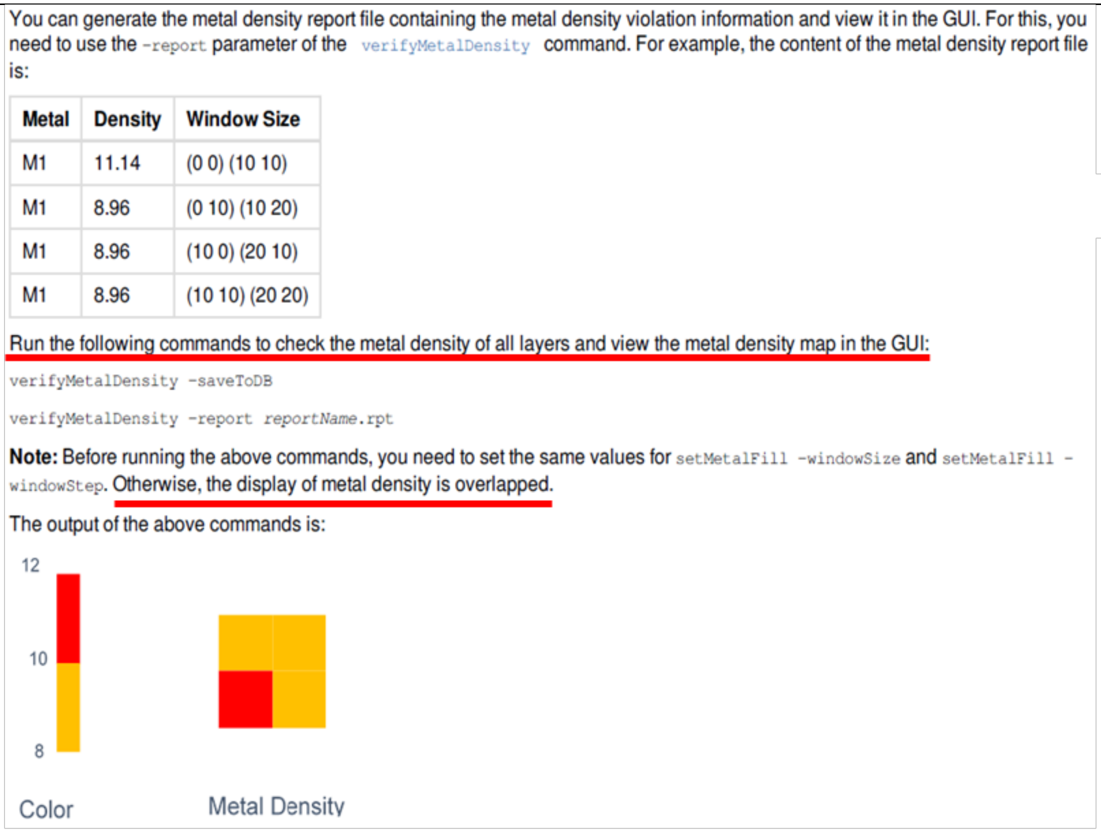

You can generate the metal density report file containing the metal density violation information and view it in the GUI. For this, you need to use the `-report` parameter of the `verifyMetalDensity` command. For example, the content of the metal density report file is:

Metal	Density	Window Size
M1	11.14	(0 0) (10 10)
M1	8.96	(0 10) (10 20)
M1	8.96	(10 0) (20 10)
M1	8.96	(10 10) (20 20)

Run the following commands to check the metal density of all layers and view the metal density map in the GUI:

```
verifyMetalDensity -saveToDB

verifyMetalDensity -report reportName.rpt
```

Note: Before running the above commands, you need to set the same values for `setMetalFill -windowSize` and `setMetalFill -windowStep`. Otherwise, the display of metal density is overlapped.

The output of the above commands is:

See Innovus User Guide product version 20.10, March 2020, page 728 .

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created by minimizing the overlap by re-arranging a plurality of first and second metal fill features to be staggered. Given the near-certainty that ECOs are implemented during the design process, and the layout is altered (and thus, dummy metal fill is repositioned), it is necessary to minimize the resulting overlap between dummy fill features on successive layers. *See*Ex. F at ¶¶ 71, 74–76.

wherein the first dummy fill space includes non-signal carrying lines on the first layer and the second dummy fill space includes non-signal carrying lines on the second layer.

The Accused Products are made, produced, or processed from a circuit design that is created such that the first dummy fill space includes non-signal carrying lines on the first layer and the second dummy fill space includes non-signal carrying lines on the second layer.

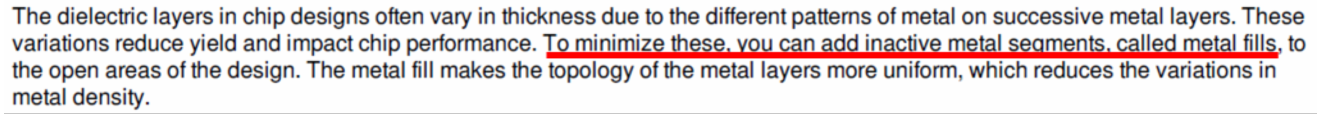

The dielectric layers in chip designs often vary in thickness due to the different patterns of metal on successive metal layers. These variations reduce yield and impact chip performance. To minimize these, you can add inactive metal segments, called metal fills, to the open areas of the design. The metal fill makes the topology of the metal layers more uniform, which reduces the variations in metal density.

See Innovus User Guide product version 20.10, March 2020, page 705 .

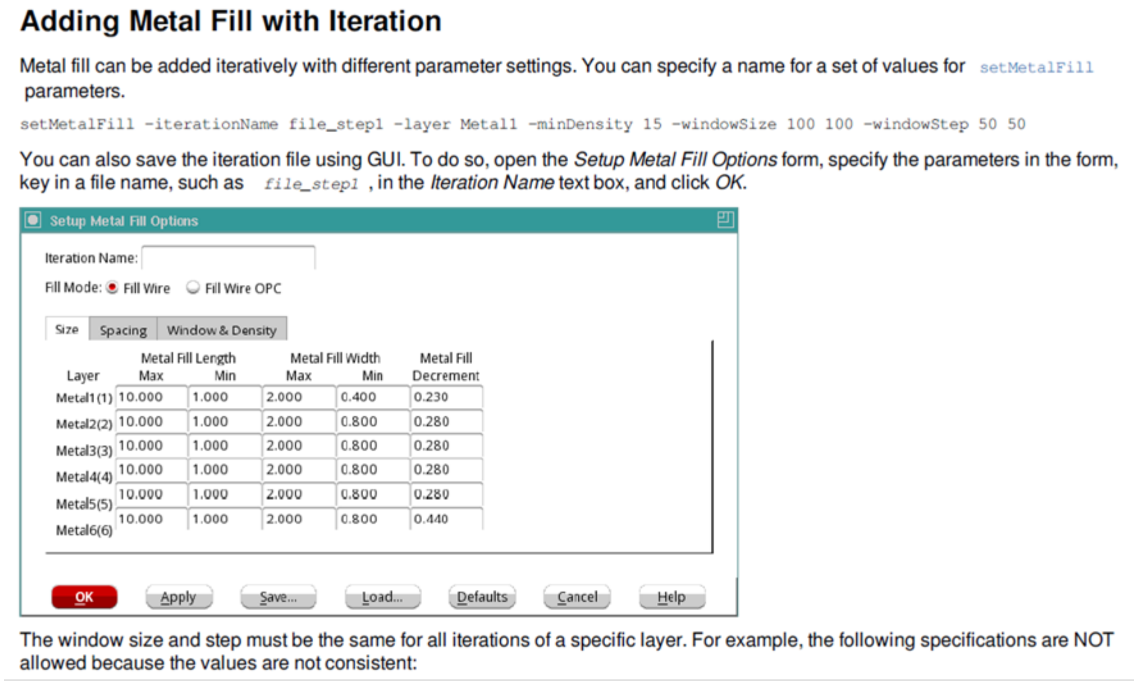

Adding Metal Fill with Iteration

Metal fill can be added iteratively with different parameter settings. You can specify a name for a set of values for `setMetalFill` parameters.

```
setMetalFill -iterationName file_step1 -layer Metal1 -minDensity 15 -windowSize 100 100 -windowStep 50 50
```

You can also save the iteration file using GUI. To do so, open the *Setup Metal Fill Options* form, specify the parameters in the form, key in a file name, such as *`file_step1`* , in the *Iteration Name* text box, and click *OK.*

Iteration Name:

Fill Mode: ● Fill Wire ○ Fill Wire OPC

Layer	Metal Fill Length Max	Min	Metal Fill Width Max	Min	Metal Fill Decrement
Metal1(1)	10.000	1.000	2.000	0.400	0.230
Metal2(2)	10.000	1.000	2.000	0.800	0.280
Metal3(3)	10.000	1.000	2.000	0.800	0.280
Metal4(4)	10.000	1.000	2.000	0.800	0.280
Metal5(5)	10.000	1.000	2.000	0.800	0.280
Metal6(6)	10.000	1.000	2.000	0.800	0.440

The window size and step must be the same for all iterations of a specific layer. For example, the following specifications are NOT allowed because the values are not consistent:

See Innovus User Guide product version 20.10, March 2020, page 726 .

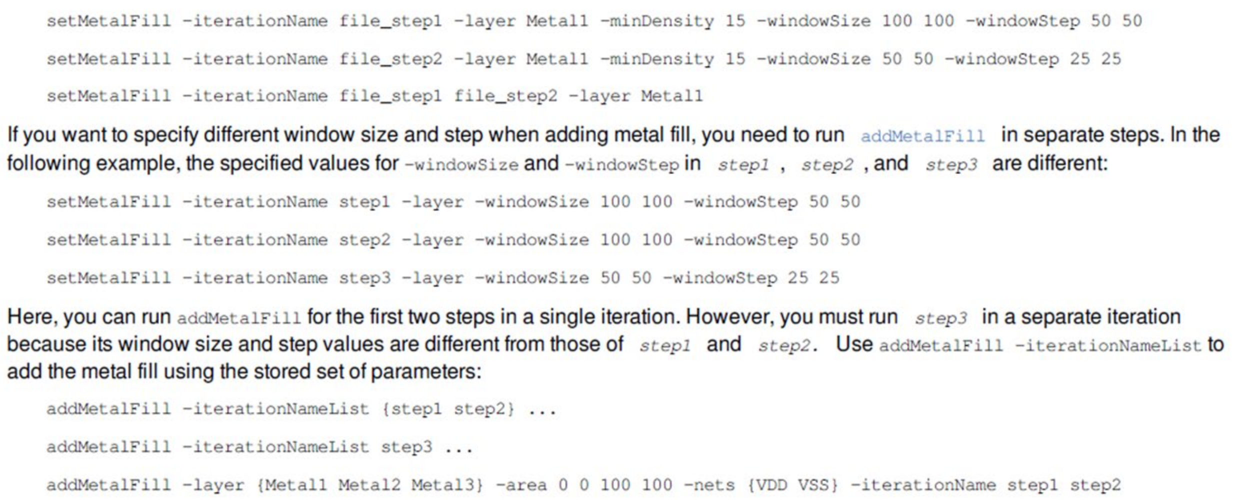

```
setMetalFill -iterationName file_step1 -layer Metal1 -minDensity 15 -windowSize 100 100 -windowStep 50 50

setMetalFill -iterationName file_step2 -layer Metal1 -minDensity 15 -windowSize 50 50 -windowStep 25 25

setMetalFill -iterationName file_step1 file_step2 -layer Metal1
```

If you want to specify different window size and step when adding metal fill, you need to run `addMetalFill` in separate steps. In the following example, the specified values for `-windowSize` and `-windowStep` in *step1* , *step2* , and *step3* are different:

```
setMetalFill -iterationName step1 -layer -windowSize 100 100 -windowStep 50 50

setMetalFill -iterationName step2 -layer -windowSize 100 100 -windowStep 50 50

setMetalFill -iterationName step3 -layer -windowSize 50 50 -windowStep 25 25
```

Here, you can run `addMetalFill` for the first two steps in a single iteration. However, you must run *step3* in a separate iteration because its window size and step values are different from those of *step1* and *step2*. Use `addMetalFill -iterationNameList` to add the metal fill using the stored set of parameters:

```
addMetalFill -iterationNameList {step1 step2} ...

addMetalFill -iterationNameList step3 ...

addMetalFill -layer {Metal1 Metal2 Metal3} -area 0 0 100 100 -nets {VDD VSS} -iterationName step1 step2
```

See Innovus User Guide product version 20.10, March 2020, page 727 .

You can also do the same through the GUI by using the *Route - Metal Fill - Add* command.

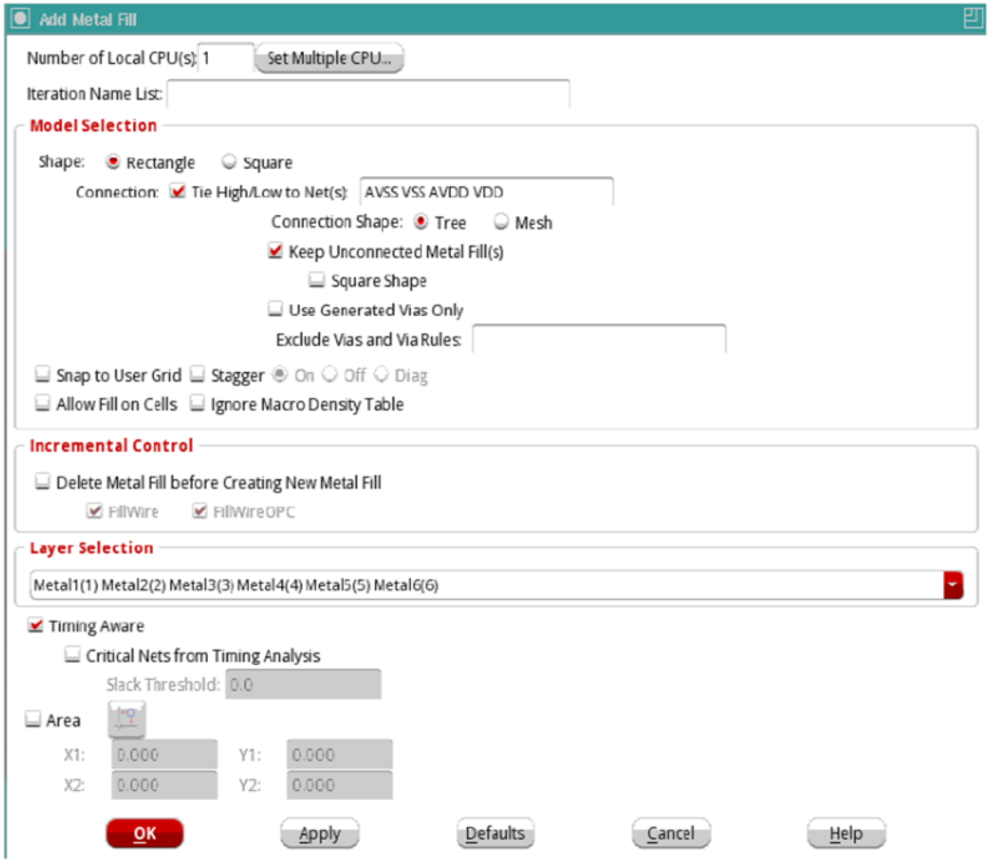

Add Metal Fill

Number of Local CPU(s): 1

Iteration Name List:

Model Selection

Shape: ● Rectangle ○ Square

Connection: ☑ Tie High/Low to Net(s): AVSS VSS AVDD VDD

Connection Shape: ● Tree ○ Mesh

☑ Keep Unconnected Metal Fill(s)

☐ Square Shape

☐ Use Generated Vias Only

Exclude Vias and Via Rules:

☐ Snap to User Grid ☐ Stagger ● On ○ Off ○ Diag

☐ Allow Fill on Cells ☐ Ignore Macro Density Table

Incremental Control

☐ Delete Metal Fill before Creating New Metal Fill

☑ FillWire ☑ FillWireOPC

Layer Selection

Metal1(1) Metal2(2) Metal3(3) Metal4(4) Metal5(5) Metal6(6)

☑ Timing Aware

☐ Critical Nets from Timing Analysis

Slack Threshold: 0.0

☐ Area

X1: 0.000 Y1: 0.000

X2: 0.000 Y2: 0.000

Key in the existing file list in *Iteration Name List* text box in the *Add Metal Fill* form and then click *OK*.

The engine processes the iterations in the order listed and stops when the preferred density is reached in any iteration.

See Innovus User Guide product version 20.10, March 2020, page 727 .

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created such that the first dummy fill space includes non-signal carrying lines on the first layer and the second dummy fill space includes non-signal carrying lines on the second layer. *See* Ex. F at ¶¶ 70-75.

Claim 2	Accused Products
2. The method as described in claim 1, wherein the plurality of first dummy fill features forms a grid within the first dummy fill space.	The Accused Products are further made, produced, or processed from a circuit design that is created such that the plurality of first dummy fill features forms a grid within the first dummy fill space.

You can also do the same through the GUI by using the *Route - Metal Fill - Add* command.

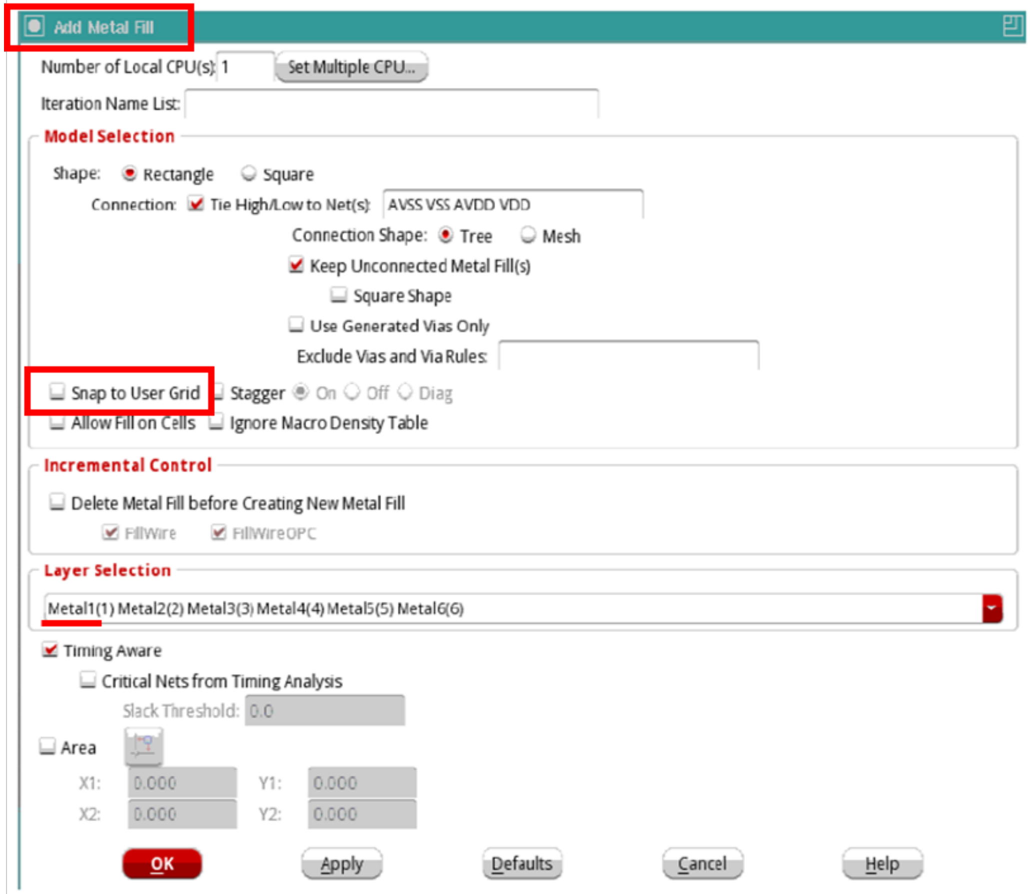

Number of Local CPU(s): 1
Iteration Name List:
Shape: ● Rectangle ○ Square
Connection: ☑ Tie High/Low to Net(s): AVSS VSS AVDD VDD
Connection Shape: ● Tree ○ Mesh
☑ Keep Unconnected Metal Fill(s)
☐ Square Shape
☐ Use Generated Vias Only
Exclude Vias and Via Rules:
☐ Snap to User Grid ☐ Stagger ● On ○ Off ○ Diag
☐ Allow Fill on Cells ☐ Ignore Macro Density Table
☐ Delete Metal Fill before Creating New Metal Fill
☑ FillWire ☑ FillWireOPC
Metal1(1) Metal2(2) Metal3(3) Metal4(4) Metal5(5) Metal6(6)
☑ Timing Aware
☐ Critical Nets from Timing Analysis
Slack Threshold: 0.0
☐ Area
X1: 0.000 Y1: 0.000
X2: 0.000 Y2: 0.000

Key in the existing file list in *Iteration Name List* text box in the *Add Metal Fill* form and then click *OK*.

The engine processes the iterations in the order listed and stops when the preferred density is reached in any iteration.

See Innovus User Guide product version 20.10, March 2020, page 727 .

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created such that the plurality of first dummy fill features forms a grid within the first dummy fill space. *See* Ex. F at ¶ 75.

Claim 3	Accused Products
3. The method as described in claim 1, wherein the plurality of second dummy fill features forms a grid within the second dummy fill space.	The Accused Products are further made, produced, or processed from a circuit design that is created such that the plurality of second dummy fill features forms a grid within the second dummy fill space. You can also do the same through the GUI by using the *Route - Metal Fill - Add* command. 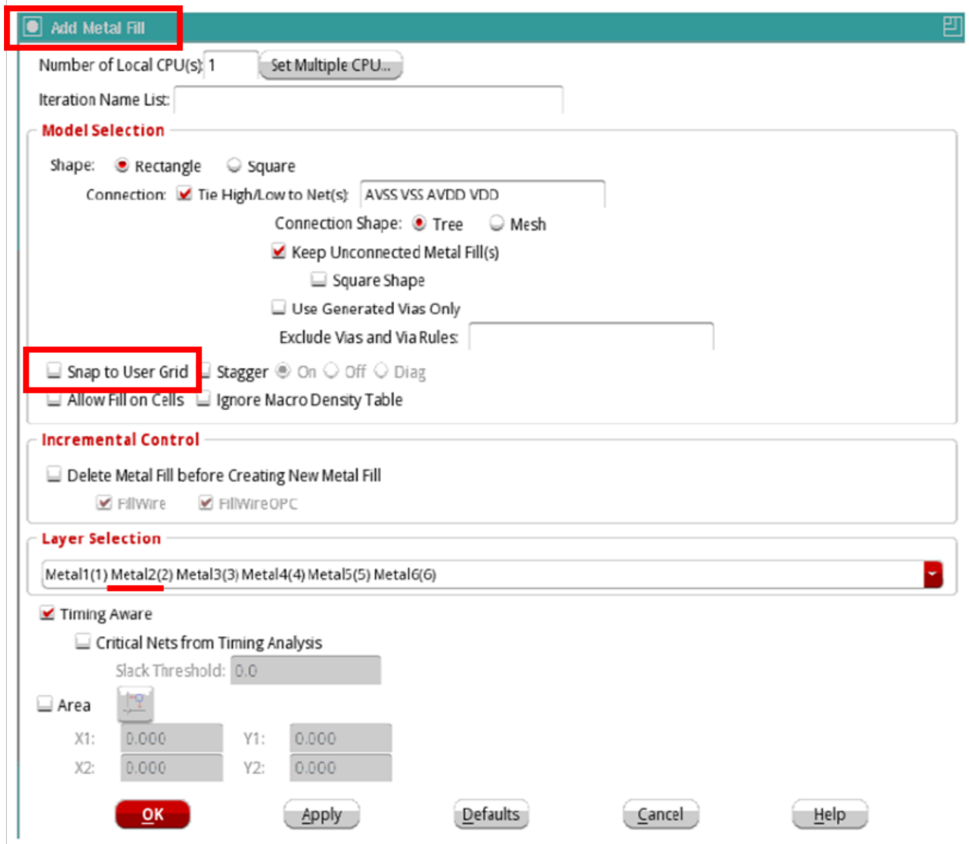 Number of Local CPU(s): 1 Iteration Name List: Model Selection Shape: ● Rectangle ○ Square Connection: ☑ Tie High/Low to Net(s): AVSS VSS AVDD VDD Connection Shape: ● Tree ○ Mesh ☑ Keep Unconnected Metal Fill(s) ☐ Square Shape ☐ Use Generated Vias Only Exclude Vias and Via Rules: ☐ Snap to User Grid ☐ Stagger ● On ○ Off ○ Diag ☐ Allow Fill on Cells ☐ Ignore Macro Density Table Incremental Control ☐ Delete Metal Fill before Creating New Metal Fill ☑ FillWire ☑ FillWireOPC Layer Selection Metal1(1) Metal2(2) Metal3(3) Metal4(4) Metal5(5) Metal6(6) ☑ Timing Aware ☐ Critical Nets from Timing Analysis Slack Threshold: 0.0 ☐ Area X1: 0.000 Y1: 0.000 X2: 0.000 Y2: 0.000 Key in the existing file list in *Iteration Name List* text box in the *Add Metal Fill* form and then click *OK*. The engine processes the iterations in the order listed and stops when the preferred density is reached in any iteration. *See Innovus User Guide product version 20.10, March 2020, page 727 .*

	For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created such that the plurality of second dummy fill features forms a grid within the second dummy fill space. *See* Ex. F at ¶ 75.
Claim 4	Accused Products
4. The method as described in claim 1, wherein the first dummy fill space is determined based on a local pattern density for the first layer.	The Accused Products are made, produced, or processed from a circuit design that is created such that the first dummy fill space is determined based on a local pattern density for the first layer. 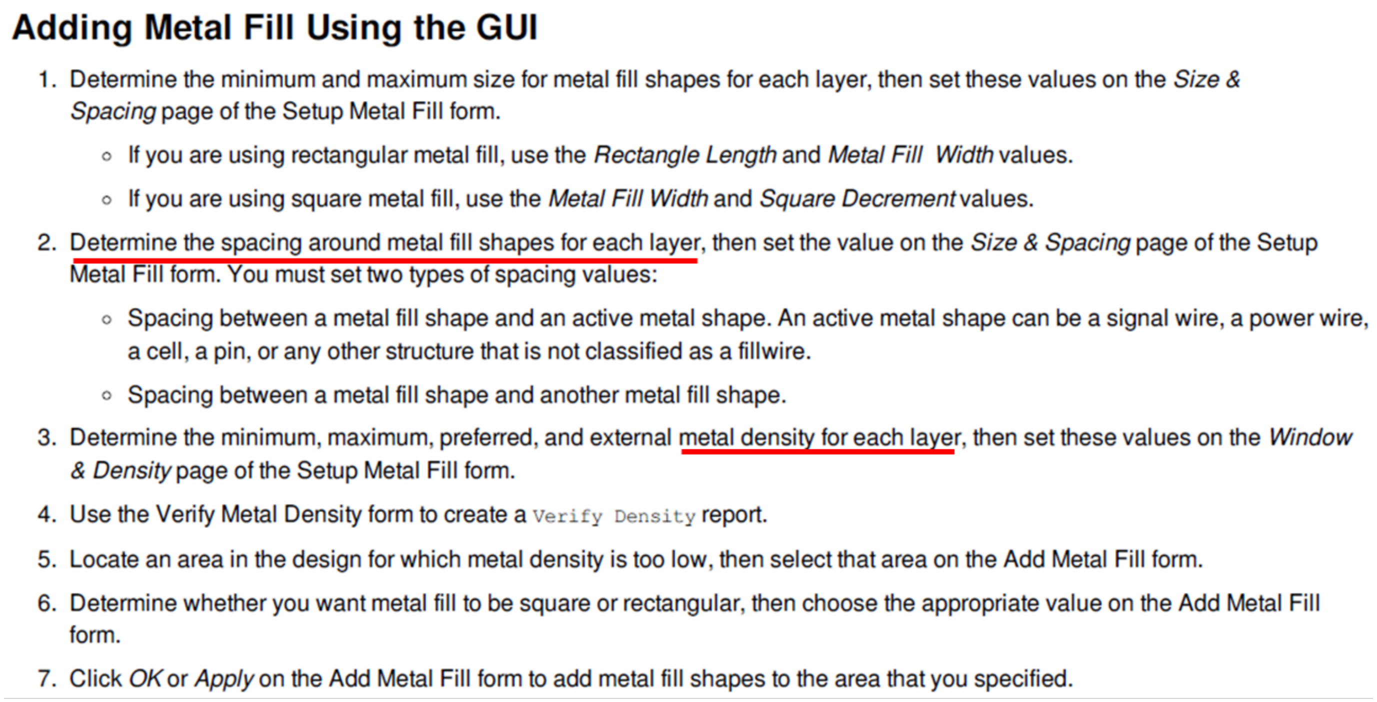 **Adding Metal Fill Using the GUI** 1. Determine the minimum and maximum size for metal fill shapes for each layer, then set these values on the *Size & Spacing* page of the Setup Metal Fill form. ◦ If you are using rectangular metal fill, use the *Rectangle Length* and *Metal Fill Width* values. ◦ If you are using square metal fill, use the *Metal Fill Width* and *Square Decrement* values. 2. Determine the spacing around metal fill shapes for each layer, then set the value on the *Size & Spacing* page of the Setup Metal Fill form. You must set two types of spacing values: ◦ Spacing between a metal fill shape and an active metal shape. An active metal shape can be a signal wire, a power wire, a cell, a pin, or any other structure that is not classified as a fillwire. ◦ Spacing between a metal fill shape and another metal fill shape. 3. Determine the minimum, maximum, preferred, and external metal density for each layer, then set these values on the *Window & Density* page of the Setup Metal Fill form. 4. Use the Verify Metal Density form to create a `Verify Density` report. 5. Locate an area in the design for which metal density is too low, then select that area on the Add Metal Fill form. 6. Determine whether you want metal fill to be square or rectangular, then choose the appropriate value on the Add Metal Fill form. 7. Click *OK* or *Apply* on the Add Metal Fill form to add metal fill shapes to the area that you specified. *See Innovus User Guide product version 20.10, March 2020, page 726 .* For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created such that the first dummy fill space is determined based on a local pattern density for the first layer. *See* Ex. F at ¶ 75.

Claim 5	Accused Products
5. The method as described in claim 1, wherein the second dummy fill space is determined based on a local pattern density for the second layer.	The Accused Products are made, produced, or processed from a circuit design that is created such that the second dummy fill space is determined based on a local pattern density for the second layer. 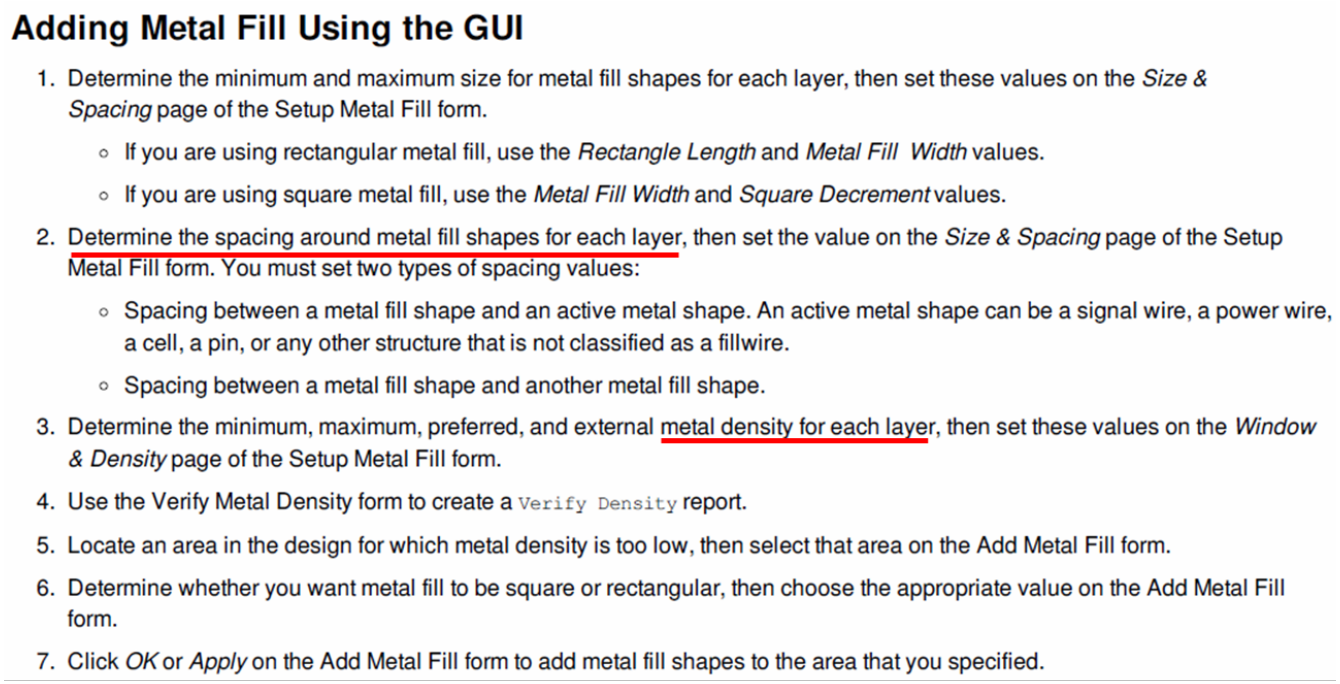 **Adding Metal Fill Using the GUI** 1. Determine the minimum and maximum size for metal fill shapes for each layer, then set these values on the *Size & Spacing* page of the Setup Metal Fill form. ◦ If you are using rectangular metal fill, use the *Rectangle Length* and *Metal Fill Width* values. ◦ If you are using square metal fill, use the *Metal Fill Width* and *Square Decrement* values. 2. Determine the spacing around metal fill shapes for each layer, then set the value on the *Size & Spacing* page of the Setup Metal Fill form. You must set two types of spacing values: ◦ Spacing between a metal fill shape and an active metal shape. An active metal shape can be a signal wire, a power wire, a cell, a pin, or any other structure that is not classified as a fillwire. ◦ Spacing between a metal fill shape and another metal fill shape. 3. Determine the minimum, maximum, preferred, and external metal density for each layer, then set these values on the *Window & Density* page of the Setup Metal Fill form. 4. Use the Verify Metal Density form to create a `Verify Density` report. 5. Locate an area in the design for which metal density is too low, then select that area on the Add Metal Fill form. 6. Determine whether you want metal fill to be square or rectangular, then choose the appropriate value on the Add Metal Fill form. 7. Click *OK* or *Apply* on the Add Metal Fill form to add metal fill shapes to the area that you specified. *See Innovus User Guide product version 20.10, March 2020, page 726 .* For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created such that the second dummy fill space is determined based on a local pattern density for the second layer. *See* Ex. F at ¶ 75.

Claim 6	Accused Products
6. The method as described in claim 2, wherein the grid includes a plurality of squares.	The Accused Products are made, produced, or processed from a circuit design that is created such that the grid includes a plurality of squares. 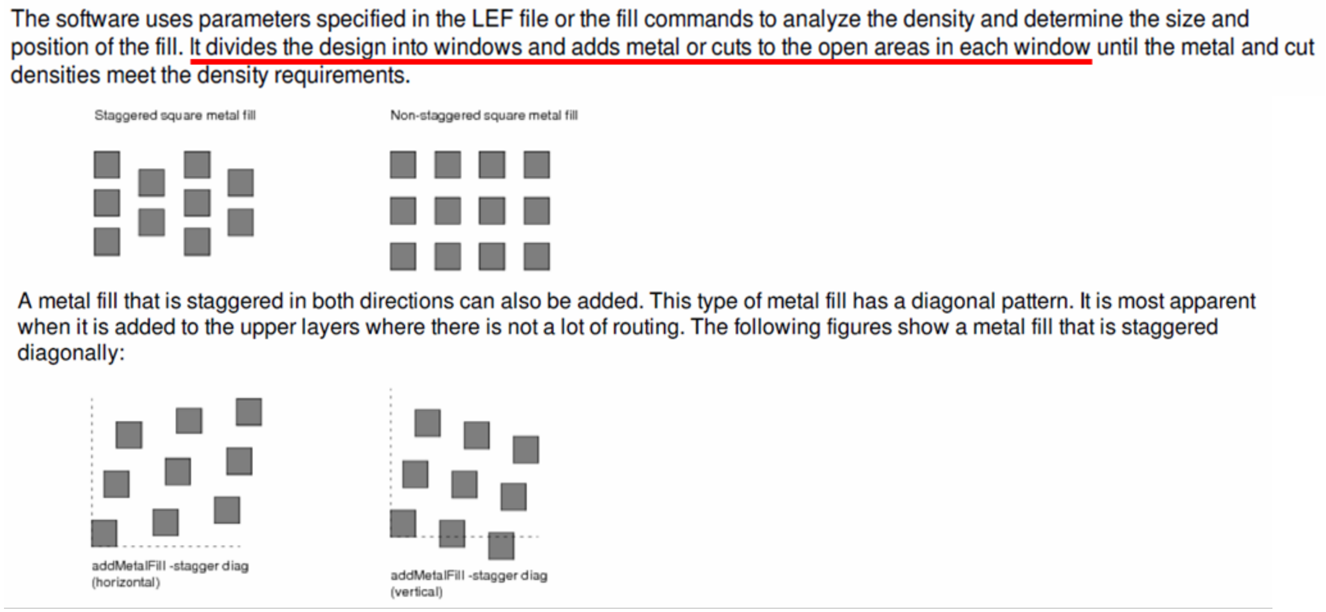 The software uses parameters specified in the LEF file or the fill commands to analyze the density and determine the size and position of the fill. It divides the design into windows and adds metal or cuts to the open areas in each window until the metal and cut densities meet the density requirements. A metal fill that is staggered in both directions can also be added. This type of metal fill has a diagonal pattern. It is most apparent when it is added to the upper layers where there is not a lot of routing. The following figures show a metal fill that is staggered diagonally: *See Innovus User Guide product version 20.10, March 2020, pages 705-6*.

You can also do the same through the GUI by using the *Route - Metal Fill - Add* command.

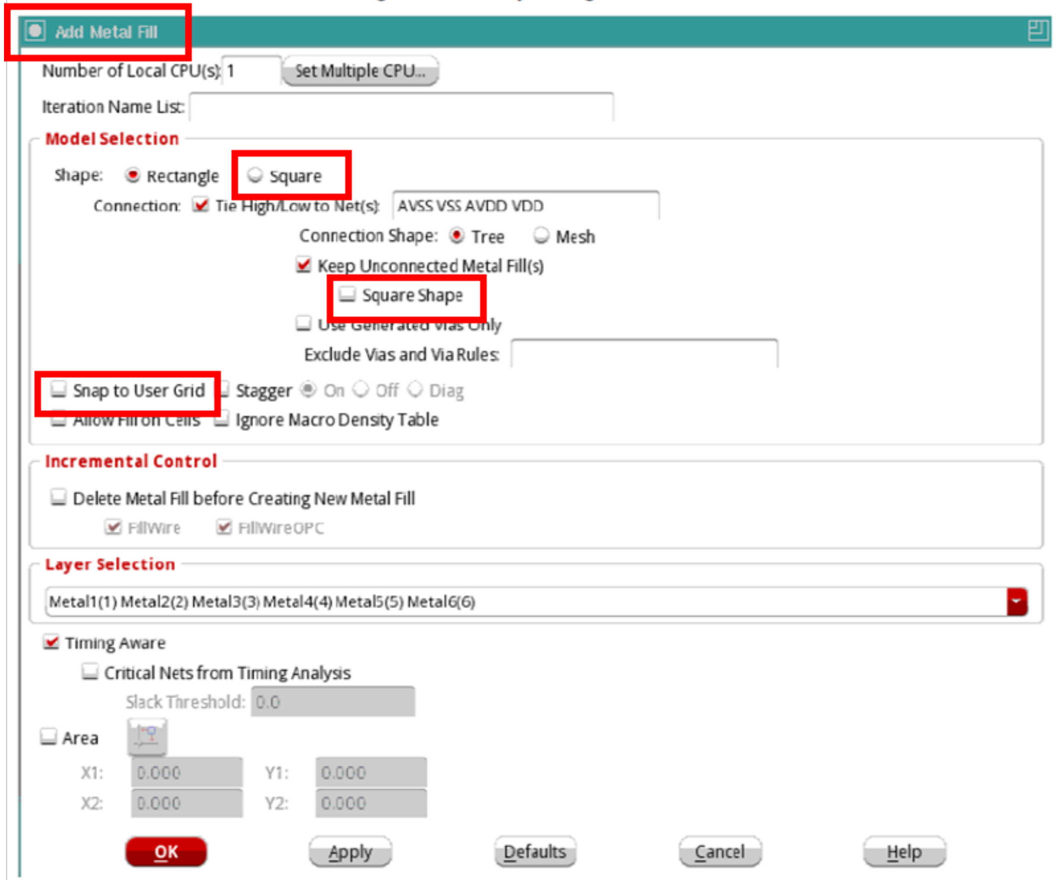

Add Metal Fill

Number of Local CPU(s): 1

Iteration Name List:

Model Selection

Shape: ● Rectangle ○ Square

Connection: ☑ Tie High/Low to Net(s): AVSS VSS AVDD VDD

Connection Shape: ● Tree ○ Mesh

☑ Keep Unconnected Metal Fill(s)

☐ Square Shape

☐ Use Generated Vias Only

Exclude Vias and Via Rules:

☐ Snap to User Grid ☐ Stagger ● On ○ Off ○ Diag

☐ Allow Fill on Cells ☐ Ignore Macro Density Table

Incremental Control

☐ Delete Metal Fill before Creating New Metal Fill

☑ FillWire ☑ FillWireOPC

Layer Selection

Metal1(1) Metal2(2) Metal3(3) Metal4(4) Metal5(5) Metal6(6)

☑ Timing Aware

☐ Critical Nets from Timing Analysis

Slack Threshold: 0.0

☐ Area

X1: 0.000 Y1: 0.000

X2: 0.000 Y2: 0.000

Key in the existing file list in *Iteration Name List* text box in the *Add Metal Fill* form and then click *OK*.

The engine processes the iterations in the order listed and stops when the preferred density is reached in any iteration.

See Innovus User Guide product version 20.10, March 2020, page 727 .

For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that is created such that the grid includes a plurality of squares. *See* Ex. F at ¶ 75.

Claim 11	Accused Products
11. The method as described in claim 1, wherein a total bulk capacitance is minimized.	The Accused Products are further made, produced, or processed from a circuit design that is created such that a total bulk capacitance is minimized. Bulk capacitance is the area capacitance between the two adjacent metal layers. 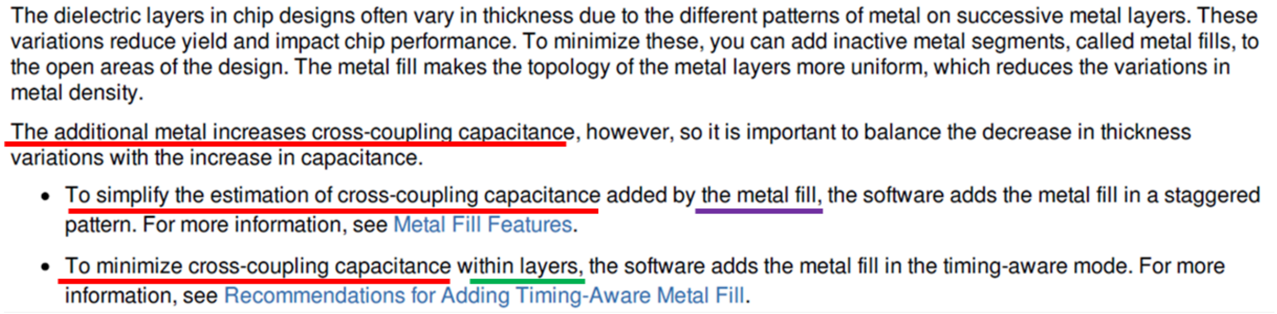 The dielectric layers in chip designs often vary in thickness due to the different patterns of metal on successive metal layers. These variations reduce yield and impact chip performance. To minimize these, you can add inactive metal segments, called metal fills, to the open areas of the design. The metal fill makes the topology of the metal layers more uniform, which reduces the variations in metal density. The additional metal increases cross-coupling capacitance, however, so it is important to balance the decrease in thickness variations with the increase in capacitance. • To simplify the estimation of cross-coupling capacitance added by the metal fill, the software adds the metal fill in a staggered pattern. For more information, see Metal Fill Features. • To minimize cross-coupling capacitance within layers, the software adds the metal fill in the timing-aware mode. For more information, see Recommendations for Adding Timing-Aware Metal Fill. *See Innovus User Guide product version 20.10, March 2020, page 705.* 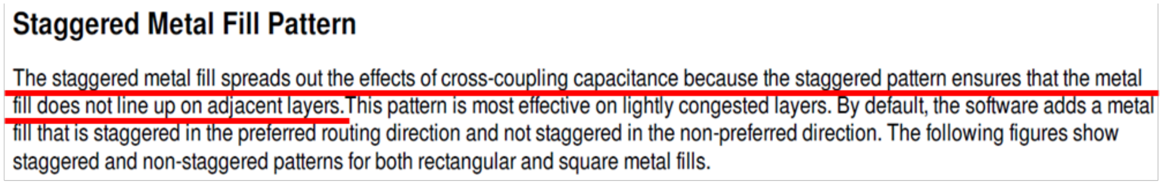 **Staggered Metal Fill Pattern** The staggered metal fill spreads out the effects of cross-coupling capacitance because the staggered pattern ensures that the metal fill does not line up on adjacent layers.This pattern is most effective on lightly congested layers. By default, the software adds a metal fill that is staggered in the preferred routing direction and not staggered in the non-preferred direction. The following figures show staggered and non-staggered patterns for both rectangular and square metal fills. *See Innovus User Guide product version 20.10, March 2020, page 706 .* Definition of Bulk capacitance from Column 5 of Taravade 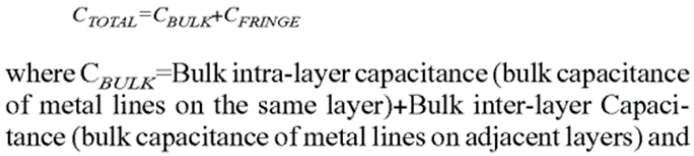 $C_{TOTAL}=C_{BULK}+C_{FRINGE}$ where C_{BULK}=Bulk intra-layer capacitance (bulk capacitance of metal lines on the same layer)+Bulk inter-layer Capacitance (bulk capacitance of metal lines on adjacent layers) and For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design that minimized total bulk capacitance. *See* Ex. F at ¶ 71.

Claim 12	Accused Products
12. The method as described in claim 11, wherein the total bulk capacitance includes a bulk inter-layer capacitance.	The Accused Products are further made, produced, or processed from a circuit design that is created such that the total bulk capacitance includes a bulk inter-layer capacitance. Coupling capacitance between signals and dummies of multiple layers 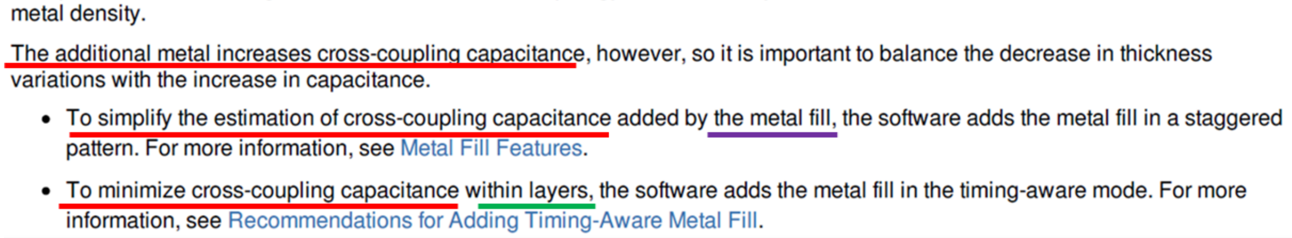 The dielectric layers in chip designs often vary in thickness due to the different patterns of metal on successive metal layers. These variations reduce yield and impact chip performance. To minimize these, you can add inactive metal segments, called metal fills, to the open areas of the design. The metal fill makes the topology of the metal layers more uniform, which reduces the variations in metal density. The additional metal increases cross-coupling capacitance, however, so it is important to balance the decrease in thickness variations with the increase in capacitance. • To simplify the estimation of cross-coupling capacitance added by the metal fill, the software adds the metal fill in a staggered pattern. For more information, see Metal Fill Features. • To minimize cross-coupling capacitance within layers, the software adds the metal fill in the timing-aware mode. For more information, see Recommendations for Adding Timing-Aware Metal Fill. *See Innovus User Guide product version 20.10, March 2020, page 705.* Bulk inter-layer capacitance is the bulk capacitance of metal lines on adjacent layers (5:34) For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design in which the minimized total bulk capacitance included a bulk inter-layer capacitance. *See* Ex. F at ¶ 71.

Claim 13	Accused Products
13. The method as described in claim 11, wherein the bulk inter-layer capacitance is a bulk capacitance created by overlaps between the first layer and the second layer.	The Accused Products are further made, produced, or processed from a circuit design that is created such that the bulk inter-layer capacitance is a bulk capacitance created by overlaps between the first layer and the second layer. 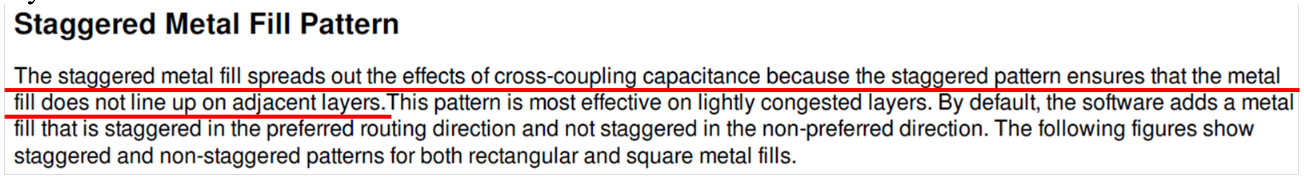 **Staggered Metal Fill Pattern** The staggered metal fill spreads out the effects of cross-coupling capacitance because the staggered pattern ensures that the metal fill does not line up on adjacent layers. This pattern is most effective on lightly congested layers. By default, the software adds a metal fill that is staggered in the preferred routing direction and not staggered in the non-preferred direction. The following figures show staggered and non-staggered patterns for both rectangular and square metal fills. *See Innovus User Guide product version 20.10, March 2020, page 706.* For example, Ampere creates a circuit design for the AC7-M128-30 Altra Max., which was made, produced, or processed from a circuit design in which the bulk inter-layer capacitance is created by overlaps between the first layer and the second layer. *See* Ex. F at ¶¶ 71, 75.

Caveat: The notes and/or cited excerpts utilized herein are set forth for illustrative purposes only and are not meant to be limiting in any manner. For example, the notes and/or cited excerpts, may or may not be supplemented or substituted with different excerpt(s) of the relevant reference(s), as appropriate. Further, to the extent any error(s) and/or omission(s) exist herein, all rights are reserved to correct the same.

EXHIBIT F

IN THE UNITED STATES DISTRICT COURT

FOR THE DISTRICT OF OREGON

PORTLAND DIVISION

BELL SEMICONDUCTOR, LLC Plaintiff, v. AMPERE COMPUTING, LLC Defendant.	Case No. _______________ **COMPLAINT FOR PATENT INFRINGEMENT** JURY TRIAL DEMANDED

DECLARATION OF DHAVAL BRAHMBHATT

1. I make this declaration on behalf of Bell Semiconductor, LLC ("Bell Semic"). I understand that Bell Semic will offer my declaration as evidence in support of its contemporaneously-filed complaint for patent infringement in the above-captioned case.

2. My qualifications to testify concerning the relevant technology are set forth in my curriculum vitae, attached hereto as **Exhibit 1**.

3. I hold a Master of Science (M.Sc.) in Physics, with a specialization in Solid State Electronics, from Gujarat University in India, which I received in 1977. I also hold a Master of Science in Electrical Engineering (M.S.E.E.) from University of Cincinnati in the United States, which I received in 1978. My continuing education included a certificate in Executive Program for Small Companies in summer of 1993 and a certificate in Marketing Management in summer of 1994, both from Stanford University. I received additional certifications in International Marketing at the University of London in 1995 and certification as a Trained Nanotechnologist in 2007.

4. I have over 30 years of experience with integrated circuit design, semiconductor processing, semiconductor manufacturing, and product quality and reliability. Since 2002, I have

served as the Founder, President, and CEO of PHYchip Corporation, a company focused on memory and physical layer (PHY) chips as well as modules and sub-systems. I also occasionally serve as a technical expert in a variety of patent litigation lawsuits involving IC memory, CMOS Analog IC, I/O interface, SIMM/DIMM memory modules, and high-speed physical layer chips.

5. From 1978 to 1980, I served as a Senior Design Engineer at Fairchild Semiconductor Corporation where I was responsible for the memory design, debug, and production of the 32K bit EPROM memory and placing this non-volatile memory on a micro-processor in collaboration with the microprocessor design team. After my time at Fairchild, I moved to Synertek Inc. where I served as a Design Project Manager from 1980 to 1982. At Synertek, I was responsible for the design and development of an industry first 256-bit single power supply, 5V ONLY NMOS EEPROM with on-chip high-voltage generation. I received multiple technology pioneering patents for this invention.

6. In 1982, I decided to join National Semiconductor as a Design Manager. There, I was in charge of high-density single power supply 64 k-bit EEPROM memory. I left National Semiconductor in 1983 to start my own company, ICT, Inc., a semiconductor startup company in the area of high-speed programmable logic and programmable memory integrated circuits (“IC”). As Vice President of ICT, I personally designed leading nonvolatile memory and logic IC chips for the company, supervised engineering, and managed all design and product development in the company. As a Founder and Vice-President, I managed collaboration between ICT and its Japanese collaboration partner Asahi-Kasei Corporation, its Korean technology partner Hyundai Electronics, and its U.S. collaboration partners American Microsystems and Advanced Micro Devices. ICT, Inc. eventually went public and was thereafter acquired.

7. In 1989, I left ICT to join National Semiconductor again as a Product Line Director. I was in charge of the business unit in the memory IC product line where I supervised close to 100

employees and oversaw product development and P&L, amongst other responsibilities. As a Senior Product Line Director, I managed collaborations between National Semiconductor and several partners, including the Japanese company Toshiba, and visited Japan frequently to both help develop new technologies with our collaborators and address suspected defects in products manufactured by National Semiconductor and incorporated into the products of Japanese customers and other customers worldwide.

8. In early 1996, I decided to leave National Semiconductor to join Smart Modular Corporation, a recognized leader in SIMM/DIMM memory modules, as a Vice President of Technology & Business Development. As Vice President, I developed and managed product development in IC memory based sub-systems such as PCMCIA, CompactFlash, and other memory cards/modules. This company also went public and was later acquired by a major worldwide manufacturer of electronic products named Solectron Inc.

9. In late 1997, I ventured to start yet another business, Modern Media Memory, Inc., where I served as the CEO until 1998. There, I designed and consulted on PCMCIA flash memory cards using NAND/NOR flash memory IC and CompactFlash cards using NAND flash IC.

10. Following Modern Media Memory, and around 1999, I joined MARS Technologies as Chief Operating Officer. This company designed and developed advanced network communications IC components focused on physical layer chips. MARS had a close technology collaboration relationship with Panasonic. MARS was acquired and the combined company eventually became a part of Broadcom.

11. Around 2000, I founded Modern Telecom that focused on advanced compound semiconductor (InP, GaAs) based technologies for telecommunications systems.

12. As an Adjunct Professor, I have taught graduate and undergraduate courses in Nanotechnology at Santa Clara University Graduate School of Engineering and at The Ohlone

College. I also taught full day courses on Nanotechnology at the Society of Photo-Optical Instrumentation Engineers (SPIE).

13. From 2006 to 2017, I was invited to serve on SBIR/STTR panels by NIH and NSF where I worked with other industry and academia experts to help the U.S. Government agencies decide on technology development funding awards for small companies in excess of tens of millions of dollars annually.

14. Over the years, I have received 11 U.S. patents in design and development of semiconductor devices. In ten of these eleven patents I was named as a sole inventor and as the lead inventor of each of the aforementioned patents. Four of these patents went on to have international counterparts.

15. I was named a Fellow under the National Scholarship Scheme by the Government of India and as a Fellow by Rotary International.

16. Currently I serve as the Co-Chairman of IEEE Region 6, Central Area. I am also the founder of the IEEE San Francisco Bay Area Nanotechnology Council and the former Chairman of the IEEE San Francisco Bay Area Vehicle Technology Society. I have received numerous awards by IEEE in 2007, 2008, 2012, and 2020. IEEE (Institute of Electrical and Electronics Engineers) is over 100 years old, and the biggest and most recognized worldwide organization of its kind.

17. Around 2007, I was appointed by then-Congressman Mike Honda and State Controller of California Steve Westly on the Blue Ribbon Task Force on Nanotechnology. Around that same time, I also made a presentation to the science sub-committee of the United States Congress.

18. I have reviewed U.S. Patent No. 7,396,760 to Taravade et al. ("Taravade '760"), which is asserted in the Complaint, and its file history. In addition, I have reviewed the claim charts accompanying the Complaint supported by this Declaration.

19. I have also reviewed various declarations of Lloyd Linder in support of other complaints filed by Bell Semic on patents relating to various aspects of dummy metal fill.[1] I agree with the substance of those declarations, and have reused their accurate descriptions of the background technology in this Declaration to help contextualize the innovations captured by Taravade '760. The portions incorporated from the Linder Declaration are identified by italicized text.

20. My college education over 7 years and 30-plus years of knowledge and experience in integrated circuit design, layout, and fabrication provides the necessary experience to support my stated conclusions set forth below.

Background on Integrated Circuit Manufacture, Including the Layout Process Flow Segment of the Manufacturing Process

21. Semiconductor manufacture begins with the creation of a set of specialized electronic files that dictate the three-dimensional structure and features of the semiconductor device. These files, which are normally referred to as Graphic Design System (GDSII) files, are specifically formatted for and serve as necessary inputs for the devices that build the semiconductor device layer-by-layer according to the instructions contained in the GDSII files. Any changes to the structures in the GDSII files will result in changes to the structures in the fully

---

[1] These Declarations relate to the Shrowty '259, Cwynar '807, Lakshmanan '803, and Hoff '626 patents.

fabricated device.[2] *The manufacturing process ends with the wafer containing the individual semiconductor devices being fully fabricated and sawed into individual semiconductor dies.*

22. The image below (although not in italics, is borrowed from Linder Declaration) *provides a simplified schematic showing, at a high level, a commonly-used integrated circuit design flow process that is representative of many (if not most) process flows in current use for creation of circuit layouts:*

---

[2] *The physical design validation of an integrated circuit design ensures that all spatial constraints are satisfied for the traces and devices formed in various layers of an integrated circuit die. The structures formed in the several layers of an integrated circuit die are represented in a GDSII format file that contains the chip topological information for creating the masks used in manufacturing the integrated circuit dies. This is also called the "layout," and which patents in this area typically call a "design". The GDSII format is an industry standard used by commercially available physical verification tools to represent physical design data. All structures affecting the performance of the circuit die must and will be present in the layout.*

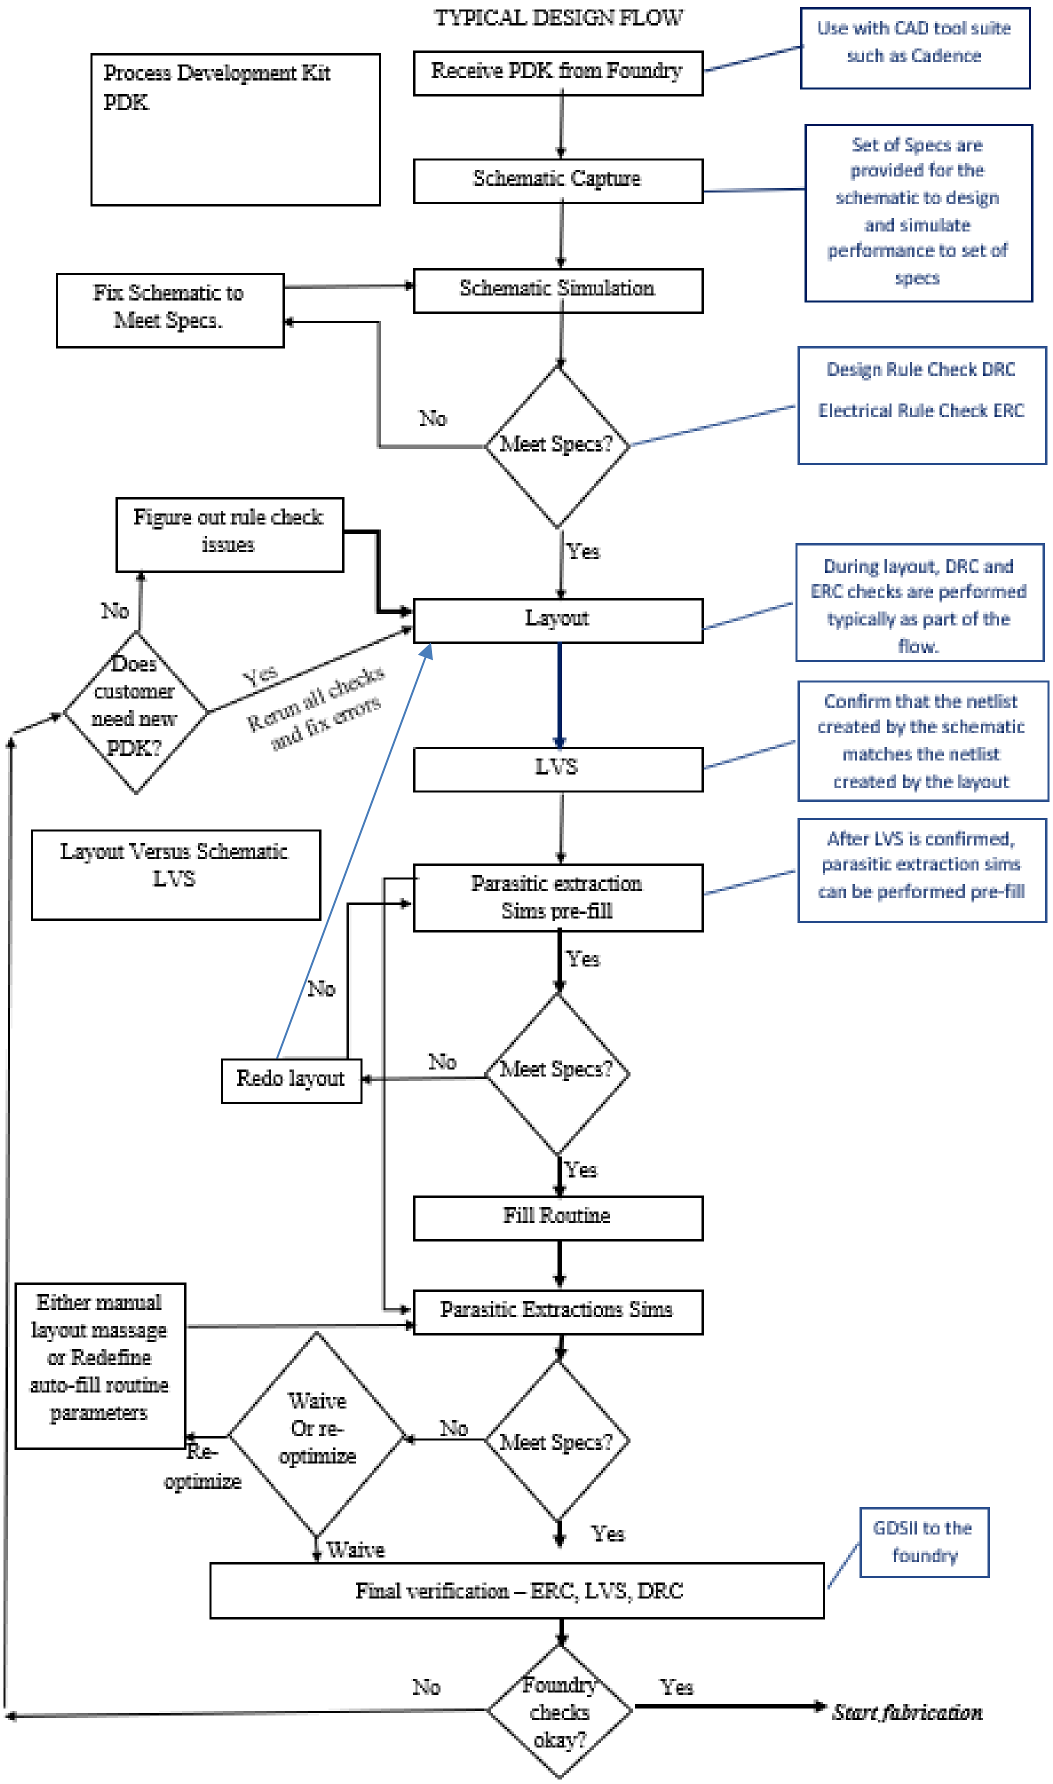

TYPICAL DESIGN FLOW

23. The integrated circuit design flow process includes a design engineer, using design tools, to create a design for an integrated circuit to be processed.

24. Design tools from vendors such as Cadence, Synopsys, or MentorGraphics (now Siemens) will then be used to design, simulate, and lay out integrated circuits. The typical design tool suite includes[3] *schematic capture, simulation, layout, verification (layout versus schematic (LVS) and design rule check (DRC)), and fill generation routines. These fill routines can be automated or manual, and can be provided by the design tool company in whole or in part.*

25. *To be sure, the precise capabilities of each design tool available to a particular design engineer may differ within a company (based on what options in the design suite are available to a particular user or on a particular device), and between different design tool suites. However, based on my experience, at a high level, the design tools used by design engineers in the semiconductor industry, all operate in substantially similar fashion for schematic capture, simulation, layout, verification, design rule check, and fill-generation.* In particular, based on my experience, I agree with Lloyd Linder that, the design tools commonly used in the industry to place dummy fill operate in substantially similar fashion in providing incremental and timing-aware fill generation for integrated circuit layouts, including the tools used for calculating the additional interlayer and intralayer capacitance in the placement or adjustment of dummy fill.

26. In the design process, the schematic is created first. The layout design tool is used to place and route all of the active (i.e., transistors) and passive components (i.e., resistors, capacitors, and inductors), and the interconnections between devices (represented as wires) in the schematic. It represents the circuit function that is to be physically implemented in the silicon. The schematic is created and simulated, using the CAD tools, to confirm that the circuit functions to a desired specification.

---

[3] *Sometimes electrical rule check (ERC) is also included in design tool suite capabilities.*

27. Once that performance specification is confirmed from the schematic simulation, the layout of the circuit is performed to physically place each of the individual elements necessary to implement the circuit functions set forth in the schematic in the GDSII file. During layout, layout rules for active and passive devices must be followed, but conformance is not checked until a DRC is run (typically at least as part of the final verification, though it can be run at any point or points in the layout process).

28. Once the layout is completed, it is compared to the schematic of the circuit using layout-versus-schematic (LVS) tool to confirm that the two are identical. From the schematic, a netlist (a list of devices and the associated nodes) is generated. From the netlist, the schematic could be re-generated manually *by drawing the devices and connecting the device nodes. From the layout of devices and associated nodes, a corresponding netlist is generated, from which a similar schematic could be generated by hand by drawing the devices and connecting the device nodes from the layout netlist. Then the schematic netlist is compared to the layout netlist using the LVS tool. The LVS tool compares the schematic netlist to the layout netlist to see if they match—i.e., whether they contain the same devices connected in the same fashion. If they do not match, the discrepancies between the two must be found and corrected, and LVS re-run. Any violations of layout rules must be corrected and DRC re-run for the layout.*

29. After passing LVS, the process of performing parasitic extraction simulations before the fill has been placed (pre-fill) can be performed on an extracted netlist created from the layout. If parasitic simulations are performed prior to the fill placement, the designer can get an idea of the impact on circuit performance from the basic layout parasitics pre-fill. From the layout, a netlist is extracted that includes any of parasitic resistance (R), *parasitic inductance (L), parasitic capacitance (C), or any combination of the three. Additionally, the parasitic extraction can include what is termed "coupled" capacitance (parasitic capacitance between metal lines) as* *well as the parasitic capacitance to the substrate. For maximum accuracy, this should not only include intralayer effects (i.e., interactions between metal elements on the same layer, such as between dummy fill and signal lines) but also interlayer effects (i.e., interaction between parallel or overlapping metal features on adjacent layers), The extracted netlist, with the selected added parasitics, can be used to run simulations on the baseline layout to determine if there is any performance degradation due to the baseline layout routing.*

30. The simulated performance of the layout, which includes the parasitics, needs to be as close as possible to the specification that was already satisfied by the schematic. That is why parasitic extraction is performed, and why it is iterated pre-fill and post-fill. So if there is performance degradation due to the baseline layout, the layout is redone until its performance is at acceptable parameters. Ideally, the extracted simulation results closely *match the schematic simulation results, which means that the layout parasitics had no* significant *impact on the circuit performance.*

31. Once the layout passes pre-fill, the design tool is used to insert dummy fill at appropriate locations in the layout that ideally *do not contain devices or other features. As is well-known in the industry, the purpose of adding dummy fill is to achieve a higher and more uniform density of interconnect across the surface of each layer of the chip, to improve the outcomes of the chemical-mechanical polishing/planarization (CMP) step during fabrication. If individual pieces of fill are below a certain minimum size, they may* give rise to planarization issues *during CMP, which will result in the dielectric material deposited on top of those too-small features not planarizing properly,*[4] *which will produce ~~in~~ dishing in the dielectric and result in a non-*

---

[4] The effect on the dielectric from underlying interconnect is known as the deposition bias. A "positive bias" or "positive deposition" bias is when the width of the protrusion in the dielectric is greater than that of the underlying active interconnect feature. Conversely, a "negative bias" or "negative deposition bias" is when the width of the protrusion in the dielectric is less than that

planarized surface. Thus, in practice, the fill pieces added cannot be below a certain minimum feature size. Adding dummy fill at or exceeding the minimum feature size and to achieve a higher and more uniform density of interconnect lowers the likelihood of defects caused by the CMP process step and thus improves the yield of modern integrated circuits.

32. Once all components of the integrated circuit design have been placed and routed, a physical design validation is typically performed at the very end of the design cycle. This ensures that all spatial constraints are satisfied for the traces and devices in each layer of an IC, that the die complies to all process rules, and that any additional required steps specific to manufacturability for a selected technology have been performed (e.g., metal utilization).

33. Even after a physical design validation, the physical design may change for any one of a number of reasons, including but not limited to timing delays, performance, or functionality. In such instances, the various steps in the process flow will have to be redone to accommodate the changes in the physical design. This includes placement of dummy fill as well.

34. *As the pre-fill step confirms that parasitics of the baseline layout, pre-fill, do not degrade the performance of the integrated circuit, it is desirable that the fill likewise does not degrade performance. However, depending on its placement, dummy fill can also degrade the performance of the integrated circuit, which is undesirable. To minimize this, the design suites include timing-aware fill tools that minimize, if not prevent, any degradation to circuit performance caused by dummy fill insertion. These tools also incorporate details on fill density, size, and position necessary to meet the requirements of the fabrication process and allow the user to specify the minimum and maximum dimensions of the dummy fill.*

---

of the underlying active interconnect feature. In either case, large density variations of the active interconnect features will typically result in interconnect that is insufficiently planarized during CMP, and thus, overpolishing of the dielectric that produces significant dishing. This is particularly detrimental in fabrication of multi-layer chips and packages.

35. Based on my experience, I agree with Lloyd Linder that the use of such timing-aware fill tools has become standard practice in designing modern integrated circuits. *In fact, modern integrated circuit designs are required to have fill included as part of the database submitted for fabrication. Due to the complicated nature of these designs, such as SoCs and highly integrated circuits with many layers, the fill process cannot be manual at least for the practical reason of there being far too many locations and options for fill position and dimension to designate by hand for fill insertion. Moreover, the chip has many critical nets (i.e., important timing-sensitive signal lines), so there is a need for the fill-placement to be aware of any impact on the timing and resulting performance impact of the circuit. Timing-aware fill tools are used to attempt to simultaneously meet interconnect density (including feature size) and timing closure requirements, but they are not guaranteed to do so 100% of the time. When this occurs, a decision must be made to compromise performance at the expense of yield, or vice-versa.*

36. Once the fill routine is completed, the fill checks are done, and final verification is performed again (LVS, DRC). The fill checks are performed based on percentage requirement on a specified area in the layout.

37. Once the layout database has been verified, it is sent for fabrication in the form of a GDSII database, which is the industry standard format for delivery of the chip database. As previously mentioned, fill is required to be included as part of the GDSII database.

38. The design resource is provided with a process design kit (PDK), which includes all of the information necessary to capture a schematic, run a simulation, do a layout, and perform all of the checks on the layout to make sure that the final GDSII is in an acceptable form to be ready for fabrication. It is the design resource / customer's responsibility to make sure that the designed chip meets all of the expected requirements for fabrication and the design resource / customer *bears the risk of failing to follow any steps in the design flow. For example, if the circuit* *does not work, that is the customer's responsibility. If the layout does not match the schematic, that is the customer's responsibility. The GDSII does have to meet all of the DRCs in order to be fabricated.*

39. In order to develop an integrated chip product, tools are needed to develop the schematic, the layout, verification of the layout, and the final GDSII database for fabrication. Many companies use different tools (from different vendors) to accomplish this process either typically due to cost or preference of internal proprietary tools. Regardless of the process and specific tools that are used, the GDSII database goes through an internal DRC after it is received and before fabrication of the integrated chip:

a. The design resource receives a PDK that contains all of the information is included to create a GDS database to release for fabrication. This includes circuit symbols for the creation of the schematic, models for the circuit symbols to run simulation, and associated layout devices that have been created with all of the process layers needed.

b. Additionally, there are what are known as "rule decks" in the PDK that allow for LVS and DRC. A rule deck is typically a file that specifies all of the available rules (for example, minimum feature sizes such as line width, line spacing, *and minimum fill dimensions), the layers to process on each rule, and the parameters of each rule. The LVS deck compares the schematic to the layout, and the DRC deck covers all of the design rules for placing and routing devices. For LVS, a netlist of the layout is created. This netlist is compared to a netlist created for the schematic. The LVS tool compares the two to determine if they match or not.*

c. *Additionally, there is a parasitic extraction deck that extracts all of the parasitics of the layout that is used to run simulations to close timing or to confirm that the layout still meets all of the chip requirements.*

d. *There can also be an electrical rule check (ERC) deck as well, depending on the fabrication involved.*

40. If the DRC rules at pre-fabrication do not match those at the design resource, it is possible that there will be DRC errors. This could be due to a number of reasons, including the DRC in the provided process design kit (PDK) is not up to date, and so the PDK will be updated with the updated DRC and the design resource will have to redo the necessary portions or even *everything and fix the DRC errors, providing a new GDSII database before fabrication can begin. These DRC checks at pre-fabrication will include checks for the fill on all layers to confirm that the fill requirement is met, on a granular level, for all tiles at the chip boundary level.*

Dummy Fill is Required in Design and Layout of Multi-Layer Semiconductor Chips

41. I agree with Lloyd Linder that, to the best of my knowledge, adding dummy fill is a requirement for every integrated circuit using the latest technology nodes. Certain older nodes still in fabrication (>350nm) may not require fill, but I believe that even some of these older technology nodes have incorporated fill requirement to enhance yield.

42. As mentioned above, it is required that the GDS database include fill within the database submitted for fabrication. In particular, most fabrication processes used in modern semiconductor chip designs require both a minimum density and a minimum feature size for the interconnects (i.e., pieces of metal or semiconductor) placed on each layer of a multi-layer chip design. This is the case both for each of the layers as a whole and for individual subunits of each layer, and is fundamental to the creation of consistent fabrication of multi-layer devices with minimal defects.

43. Fabrication processes typically partition each layer of the chip design into rectangular regions called tiles, each of which must also meet a minimum density requirement. For any given region of the chip, the interconnect density is the area of all of the interconnect in that region divided by the total area of that region.

44. Sufficient interconnect density and substantial uniformity of interconnect are required for the chemical mechanical polishing (CMP) portion of the chip fabrication process. CMP is crucial to achieve planarity, which allows for multi-layer chip designs and high yield of functional devices. Insufficient interconnect density and/or insufficient uniformity of interconnect between various regions will increase the likelihood of defects during the chip manufacturing process, which will resultantly degrade the yield.

45. Once the functional features of the chip design (such as power lines, signal nets, vias, and the like) have been laid out as needed in the first instance, there will usually be substantial portions of the chip design that have insufficient interconnect density to permit CMP without incurring substantial likelihood of defects.

46. To increase the interconnect density of the layer as a whole, and of regions within each layer, numerous individual pieces of interconnect are inserted into available space in low-density regions of the chip until the minimum interconnect density specified for the particular fabrication process is achieved for each tile. Because these pieces of interconnect are not intended to carry signal or power, but instead are added to provide structural stability to the chip during processing*, they are generally known as "dummy fill."*

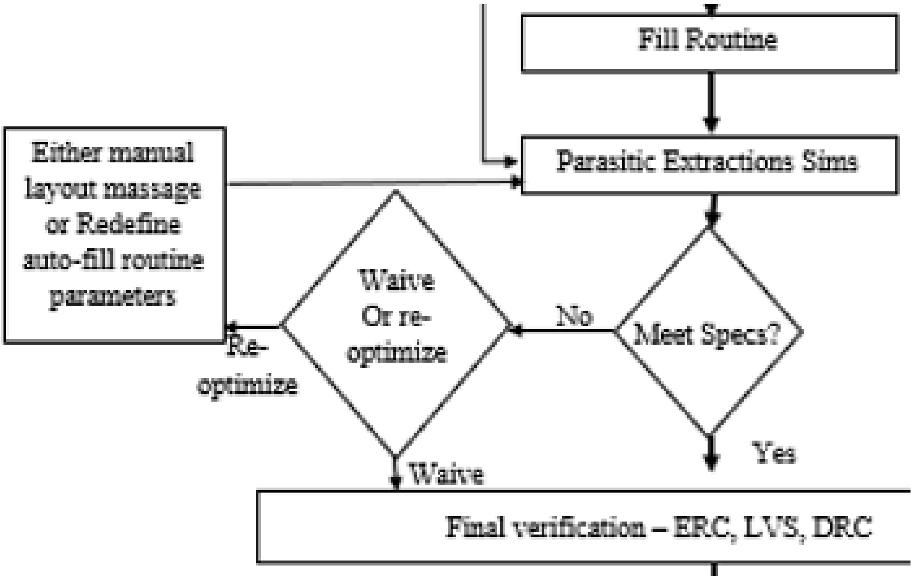

47. Placement of dummy fill is typically performed by a dummy fill software tool, and is one of the last steps in the chip design flow, with its extent and placement typically occurring after routing and timing closure. The time it takes the dummy fill tool to complete its task depends on the complexity of the circuit layout, and correspondingly, the size of the design database. If dummy fill must be run (or re-run) for the entire layer, even small changes in layout can result in significant delays while the dummy fill tool runs each time the layout changes.

48. In operation, the dummy fill software tool typically partitions each layer of the design into rectangles called tiles, which it examines in each layer of the design. If the interconnect density in each tile does not meet (or exceed) the specified minimum interconnect density for the fabrication process, the dummy fill tool inserts dummy fill into free regions of that tile where no interconnect is present.

49. The dummy fill software tool typically allows the user to specify the shape (rectangular or square) and dimensions (maximum and minimum) for the dummy fill to be inserted into open areas of the layout. In addition or alternatively, fill dimensions, shape, and position can be (and typically are) supplied separately from the fabricator in a format such as a LEF file, which *the dummy fill software tool then incorporates and uses to place dummy fill in open areas of the layout*

50. For large integrated circuits, commonly called system-on-a-chip (SoC) with either large analog content and small digital content ("big A, little D") or large digital content and small analog content ("big D, little A"), it is not practical to manually add dummy fill, so automated fill routines are almost always used. Because there are so many critical signals in a large SoC, the process cannot be done manually due to the time and trained human resources *it would require. Thus, the design timelines and practical realities require that the automated fill routines are used instead.*

51. However, placing the dummy fill that is too large in size, too extensive, and/or too close to signal nets increases capacitance between the signal wires and the dummy fill in the physical device if fabricated without taking additional measures. That increase in capacitance in the fabricated physical device would in turn slow the transmission speed of signals and degrades the overall performance of the integrated circuit. This effect between the signal wires and the dummy fill (or dummy fill and other dummy fill) *is undesirable and is* caused by what is generally *known as "parasitic capacitance."*

52. The added parasitic capacitance will degrade parameters, such as operating frequency and rise/fall time, for a critical clock or signal, and this must be avoided in order for the circuitry to work properly. The manufacturers often would be required to sell units that are slow but fully functional otherwise at a significantly lower average selling price (ASP).

53. The parasitic capacitance within a layer is inversely proportional to the distance between the dummy fill and the signal wire. Thus, parasitic capacitance from dummy fill will be minimized if the dummy fill is placed far from signal nets.

54. *In other words, the higher the required interconnect density, the closer it must be placed to signal nets, with increasingly higher parasitic capacitance and negative impact on timing and circuit performance. Conversely, the more sensitive the timing requirements for the circuit, the less the parasitic capacitance can be tolerated near crucial signal nets and the lower the interconnect density can be for tiles that include such signal nets. This tradeoff is further complicated when multiple metal layers are involved, which can be ten or even more.*

55. However, parasitic capacitance also arises from interlayer effects, as shown in Figure 1 of Taravade '760, which depicts how overlapping metal elements (both signal-carrying and non-signal-carrying) on different layers can still produce unwanted and undesired capacitance:

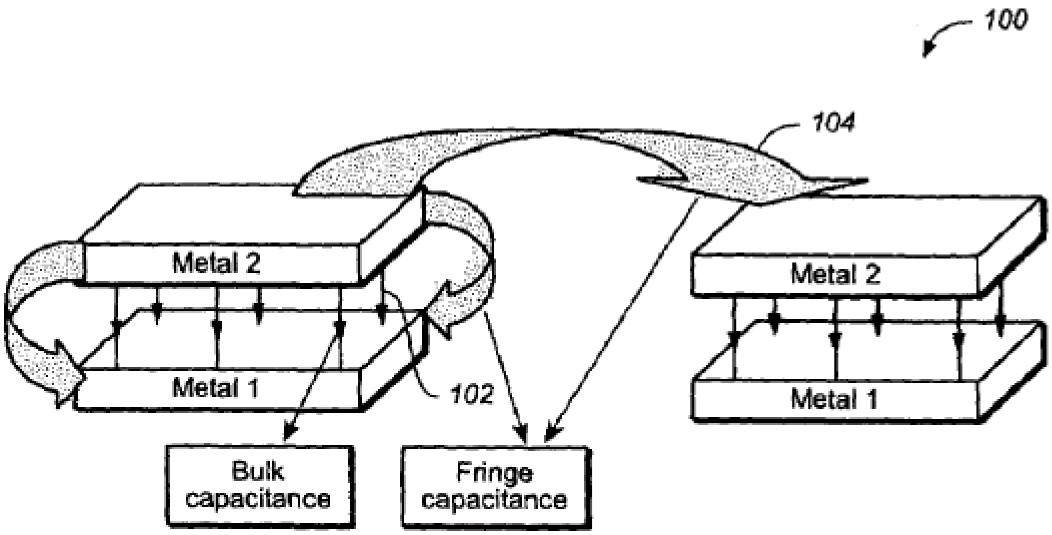

56. Some portion of the bulk capacitance, such as that due to the overlap of signal lines, may not be easily addressed to a meaningful extent because the placement is dictated by circuit functionality and circuit layout realities. Accordingly, their overlap on adjacent layers (and thus, their contribution to unwanted bulk capacitance) may be difficult to reduce.

57. However, the interlayer bulk capacitive effects contributed by other features, such as that resulting from overlapping dummy fill features on adjacent layers, is more readily addressed. Especially compared to signal lines, the specific positioning of particular dummy fill

features (which do not carry signal) is not dictated nearly as strongly by circuit functionality demands. Rather, as discussed in greater detail above, the considerations for dummy fill placement primarily involve reaching sufficient density of interconnect for each tile on a layer and the layer as a whole, uniformity of interconnect density, and minimizing timing impact on crucial signal nets. Thus, it is possible to consider the capacitive interactions between dummy fill features on adjacent layers and to mitigate their negative effects by minimizing their overlap (thereby reducing the interlayer bulk capacitance) by repositioning the dummy fill features in one layer relative to the dummy fill features in an adjacent layer. Given the relatively small size of dummy fill features relative to signal nets (and especially crucial signal nets), and the typical spacing provided between individual features, this can readily be accomplished with no more than a miniscule impact on intralayer effects and without reducing interconnect density or uniformity on the tile or the layer.

58. It may be that the timing requirements cannot be met without a revision to the fill placement, density, positioning, and/or sizing, and re-extraction of the layout parasitics to determine if the timing requirements are met. If they are not, then a decision would have to be made to either (i) continue the iteration process, or (ii) apply for a waiver and bear the risk of lower yield or (ii) accept decreased performance that could significantly impact the ASP as was explained earlier.

59. Balancing these tradeoffs started to become particularly problematic by the early 2000s, as new processing technologies with smaller and smaller features demanded increasingly higher minimum interconnect density values at the same time that chip designs became much more aggressive in the circuit timing requirements. In such cases, it was often almost impossible to insert sufficient dummy fill into a tile such that the higher minimum density requirements could be met without also reducing the large “stay-away” distance, and thereby raising the timing impact of the dummy fill to levels that affected the performance of the chip. One potential solution was

for the chip designer to waive the minimum interconnect density specified by a particular fabrication process. However, because invoking this waiver would not comply with the fabrication process requirements, the yield of the produced devices would not be guaranteed in such cases, which rendered this alternative not viable in practice.

60. Even when dummy fill placement on an individual layer of the device was not problematic by itself, its interactions with overlapping dummy fill features on adjacent layers could and still did result in substantial undesired capacitance from interlayer effects. That is because even "timing-aware" or "smart" dummy fill tools conventional prior to the time of Taravade '760 focused primarily on solving the problems of feature density and uniformity *within a layer* or portions of a layer. *See* Taravade '760 at 1:62–67, 4:11–16. While they may have considered the timing impact of dummy fill, that impact was typically limited to *intra*layer effects, such as on adjacent signal nets.

61. Accordingly, these tools and methodologies for inserting dummy fill generally treated each layer independently. Because they did not typically consider *inter*layer capacitance even when applying timing-aware methodologies and techniques, they tended to produce substantial overlaps in dummy fill features between adjacent layers.

62. This unwanted bulk capacitance would tend to slow down signals in the IC and adversely affect its timing. *See* Taravade '760 at 2:1–6. Adjustment of layers individually and manually to reduce overlap in dummy fill features to mitigate interlayer capacitance was an involved and time-consuming iterative process that could produce substantial delays in meeting design schedules. Especially as features became smaller and performance demands increased, it became both increasingly important and increasingly difficult to remove additional sources of unwanted capacitance from the ICs.

Taravade ’760

63. Even when dummy fill could be placed in such a fashion that it would simultaneously satisfy interconnect density requirements for each tile and minimize any impact on critical nets within a layer, prior to Taravade ’760, the contribution of adjacent layers’ overlapping dummy fill to interlayer capacitance could and did still have a substantial negative impact on timing . *See* Taravade ’760 at 4:14–16.

64. Taravade ’760 teaches a technique and a system for reducing the bulk capacitance caused by overlapping dummy fill in adjacent layers by repositioning the dummy fill features so as to minimize their overlap by not only considering each layer on its own, but also with respect to each of its adjacent/successive layers, by treating “each consecutive pair of layers together.” Taravade ’760 at Abst. Once the circuit layout is provided, suitable spaces for dummy fill insertion are identified. *See id.* at 2:28–31 & 2:41–43. Overlaps or potential overlaps are determined and then avoided. *See id.* at 2:31–34 & 2:40–48.

65. Dummy fill can be arranged initially to minimize overlaps and/or rearranged to minimize overlap in features once avoidable overlaps are discovered. *See* Taravade ’760 at 2:28–34 & 2:43–48. Either way, in considering each individual pair of layers as a unit, the final placement of dummy fill features on the top layer will not be placed directly above dummy fill features on the lower layer; they will be offset in order to reduce the unwanted bulk capacitance and thus minimize the inter-layer capacitance. *See id.* at 2:49–59, 4:47–49. For square-shaped dummy fill features, this will typically result in a checkerboard-like pattern. *See id.* at 2:49–55.

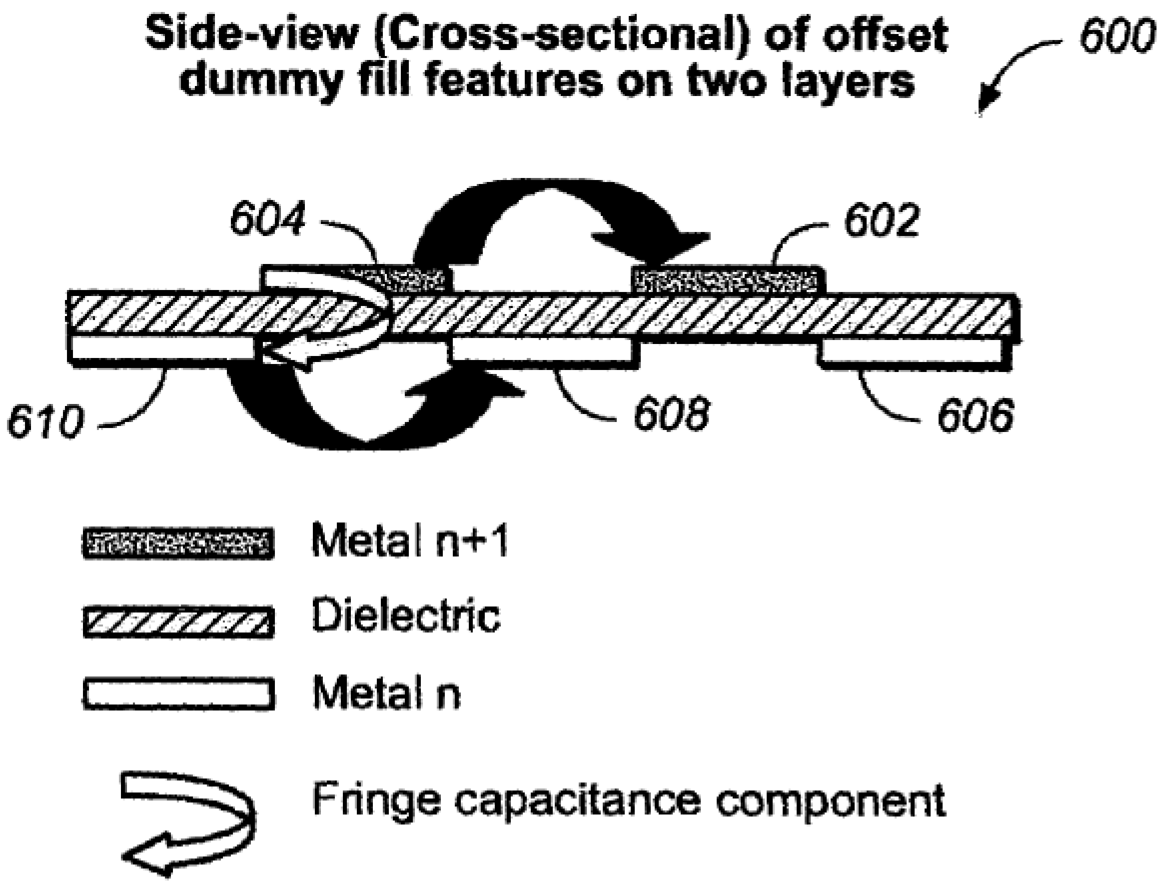

66. The significant bulk capacitance reduction (and thus, increased ability to meet demanding performance requirements and operating speed) are repeatedly described within *See, e.g.*, Taravade '760 at 1:24–30, 2:3–6, 2:57–59, 3:30–33, 4:43–45, 4:47-49, 5:18–39. This helps IC manufacturers eliminate the large bulk capacitance component and reduce the total capacitance of an IC. *See id.* at 5:23–27.

67. Based on my experience in semiconductor layout and design, I agree that this new and improved technique of offsetting dummy fill features in adjacent layers results in substantial bulk capacitance reduction in an integrated circuit, and is crucial to meeting the aggressive performance demands of modern ICs. These gains are so substantial, and in my experience, the offsetting of dummy fill features in adjacent layers to prevent their overlapping is so widely used today that it is hard to quantify just how important the inventions claimed by Taravade '760 are to

achieving the market-demanded performance and the resulting financial gains from the marketing/sales of modern chip designs.

68. Based on my experience in semiconductor layout and design, it was not well-understood, routine, or conventional at the time of Taravade ’760 to identify overlap in dummy fill features in adjacent layers in multi-layer IC designs. Likewise, it was not well-understood, routine, or conventional to rearrange one or both sets of dummy fill features to minimize their overlap. These aspects of the technique, recited in claim 1 of Taravade ’760, are central to the invention and required by every claim of the patent. This is true not only considering each of these elements by themselves, but also in combination with each other and as an ordered combination with the other recited claim elements. As Taravade ’760 explains, “the problems created by the inserted dummy fills such as adverse effects on the electric field, unwanted bulk capacitance, and the like have not been addressed.” (1:63–65.)

Claim Charts

69. I have reviewed the Complaint supported by this Declaration, along with the Claim Charts showing infringement of Taravade ’760. For at least the reasons set forth below, I agree that the Claim Charts establish use of at least one of the methods recited by the claims of Taravade ’760.

70. I have used design tools from different vendors in my career. As a consultant, I use the tools to review schematics and layouts and design and simulate circuits. Based on the requirements for the latest process technology nodes,and the yield requirements for these technologies, the latest fill tools that are used by designers and/or foundries use timing-aware fill routines with minimum fill dimensions to meet timing as well as yield requirements simultaneously. These include rearranging dummy fill features to minimize overlap in adjacent layers and eliminating another source of unwanted capacitance from the IC.

71. In particular, these tools allow rearrangement of dummy fill features to minimize overlap that would otherwise occur as a result of the inevitable and frequent ECOs and/or other layout changes during the design process. In my experience, layout changes in at least one layer are a near certainty in all recent process nodes given the complexity of the chips, aggressive timing and performance requirements, and small feature sizes. As a result of these layout changes, existing dummy fill will need to be adjusted or repositioned not just to account for the new intralayer effects, but also to minimize any interlayer effects as a result of layout changes (and corresponding changes to dummy fill spacing, positioning, and dimensions) on adjacent layers as well.

72. Based on my work history in industry and as I have done as a consultant, I can review reverse engineering ("RE") of semiconductor die to confirm that at least one of these tools (or similar tools) have been used to construct the layout or the die.

73. Even when the full history of the GDSII database for a particular integrated circuit is not available, my experience in semiconductor design and layout gives me sufficient basis to opine whether one or more of the methods claimed in Taravade '760 have likely been used in creating integrated circuits.

74. Given the aggressive schedules for bringing modern semiconductor devices to market, and the availability of incremental dummy fill in common design tools like Cadence's Innovus product, it is unlikely (if not implausible) that most chip designers would not have access to design tools that practice the inventions claimed in Taravade '760. I am aware that at least Cadence provides this functionality.

75. Among other things, it is my understanding from the Cadence Innovus User Guide that when Cadence applies a staggered metal fill, it is by default only staggered in the preferred routing direction; it is not staggered (and thus overlaps) in the non-preferred direction. Thus, in order to minimize the overlap between the dummy fill features, it is necessary to assess the extent

of the overlap in the non-preferred dimension (where the fill is not staggered) and then further rearrange the default staggered metal fill applied by Cadence to create dummy metal fill that is fully offset in adjacent layers on both horizontal axes, rather than just one as is necessary to fully minimize interlayer bulk capacitance.

76. As such, based on my experience in semiconductor layout and design, and my review of designs, I believe that it is highly likely that such functionality was used in creating most modern semiconductor devices given the importance of reducing capacitance, including interlayer capacitances, in achieving timing closure and modern performance requirements.

77. By contrast, based on my experience in semiconductor layout and design, I would only assume that relatively simple IC designs would have been made in recent years without employing at least one of the methods claimed in Taravade ’760. Simply put, there is no reason to accept substantial interlayer parasitic capacitance if it is relatively straightforward and easy to rearrange the dummy fill patterns between layers to minimize their overlap, and thus, their capacitance.

78. In addition, based on my experience, it can be assumed with a high degree of confidence that modern components in the same family or product line made by the same producer and used by the same customer in the same product line share similar features and were designed and laid out in similar fashion. This includes offsetting dummy fill features between adjacent layers.

79. *The Cadence paper “New Metal Fill Considerations for Nanometer Technologies” demonstrates several things. First, the use of the word “new” is justified in that it is a new approach, as documented here.* Secondly, it reinforces the importance of formulating “a comprehensive methodology surrounding metal fill . . . in order to minimize impact on design timing as well as to cut down on design iterations.” The paper explains that “sometimes the dummy

metal fill geometries that were added to the original design must be deleted to make room for the ECO process to succeed.” Overall, this indicates that, at least following ECO, the Cadence tool suite is used for offsetting dummy metal fill following ECO to minimize overlap of features (and thus, interlayer capacitance), as claimed in Taravade ’760.

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Dated: December 1, 2022

DHAVAL BRAHMBHATT

EXHIBIT 1
to
EXHIBIT F

Dhaval J. Brahmbhatt ---- Confidential Resume ----

President & CEO, PHY*chip* Corporation

At a Glance:

- Expert witness in patent litigation nearly 20 years (ITC, PTAB, USDC, State Courts)
- Collaborated with JEDEC, CompactFlash Association, SEMI, etc.
- Experienced IC Product Development & Design Engineer, Modules/Sub-systems
- Awarded 11 US patents in IC device/devl., of these four became international patents
- Two graduate degrees (Physics, Electrical Engineering) & a number of certifications
- Entrepreneur - Started companies, went through IPO, acquisition, etc.
- Programmable memory IC and Modules/sub-systems (Flash, EEPROM, EPROM)
- Communications IC, DRAM/SRAM memory IC and SIMM/DIMM cards experience
- Device debug and product development experience
- Founder & Chairman Emeritus, IEEE S F Bay Area Nanotechnology Council *
- United States Citizen

IEEE – Institute of Electrical & Electronics Engineers – prestigious worldwide organization over 100 years old.

Professional Summary

Mr. Brahmbhatt has worked as Expert Witness with leading IP law firms, multi-national companies and patent holding firms for almost 20 years. Brings ITC, PTAB, US District Court experience. Over the past 20 years he has been involved in a variety of IC design, semiconductor manufacturing equipment and sub-systems IP matters. Additionally, has the rare experience of modifying claims at the USPTO of an existing US patent. He regularly visits events at the newly opened US PTO office in San Jose and has been trained in the proprietary prior art search tools offered by the US PTO.

Mr. Brahmbhatt is a creative engineer awarded 11 US patents (four became Int'l patents) and brings years of hands-on product design/development experience in IC memories (standard & custom), memory cards/modules (Flash Memory Modules, SIMM/DIMM), interface buses, logic circuits, driver circuits, micro-controllers, programmable logic, remote keyless entry for cars, tags for inventory control, etc. Mr. Brahmbhatt has participated in standardization committees for Flash Memory Cards. As Vice-President at the recognized memory module maker Smart Modular Technologies, he developed memory modules while at National Semiconductor, Mr. Brahmbhatt was in charge of managing Flash Memory IC alliance with Toshiba Corporation and National Semiconductor. As Sr. Product Line Director, he doubled the revenues and made his product line profitable by systematically addressing design, production and yield issues. He brings years of experience managing collaborations with international partners such as Asahi-Kasei (Japan), Hyundai (Korea), Toshiba (Japan), Intel, AMD, etc.

Mr. Brahmbhatt is a serial entrepreneur listed on the prestigious Silicon Valley Genealogy Tree for the 1983 IC Programmable Memory and Programmable Logic high-tech startup, ICT Inc., this company made IPO in 1989. He was Vice President responsible for memory sub-systems/modules/Flash memory cards employing a micro-controller at Smart Modular Technologies, this SIMM/DIMM company also made an IPO and later got acquired. Following Smart Modular, Mr. Brahmbhatt was COO at MARS technologies (worked on high-speed transceivers) that was acquired by Globespan-Virata, which itself got acquired by Broadcom. Mr. Brahmbhatt then developed ultra-high-speed transceivers (40 Gb/s) using compound semiconductors at Modern Telecom, Inc. Mr. Brahmbhatt's engineering experience spans from being an individual contributor to Engineering VP to "C" level executive in tech companies. He has worked in all aspects of IC starting from masks to fab to packaging and yields.

Mr. Brahmbhatt champions emerging transportation technologies at the San Francisco Bay Area unit of IEEE Vehicle Technologies Society of which he was the Chairman for 2013-2018. In addition to being the Co-Chairman of Region 6, Central Area of IEEE, he is also the Founder and Chairman Emeritus of the IEEE SF Bay Area Nanotechnology Council and was a member of the Blue-Ribbon Task Force on Nanotechnology for the State of California. He is a Senior Member of the IEEE and has been recognized numerous times by the IEEE with citations and awards. He is the past President of prestigious Silicon Valley Engineering Council which represents various engineering societies in the Silicon Valley. Mr. Brahmbhatt has participated as a Program Chair at the Flash Memory Summit held every year at Santa Clara Convention Center for the past several years. He was also a keynote speaker for the International Conference of the IEEE Vehicle Technology Society held at Santa Clara Convention Center in 2013.

Mr. Brahmbhatt has been an Adjunct Faculty at Santa Clara University Graduate School of Engineering and Ohlone College and has taught full day courses at Society of Photo-Optical Instrumentation Engineers (SPIE). Mr. Brahmbhatt holds an M. Sc. Physics from India and M. S. E. E. from Ohio (USA), he has received numerous scholarships & fellowships during his distinguished education career.

------------------------ Current Employment & Other Information ------------------------

Current Employment

From:	2002	PHY*chip* Corporation
To:	Present	Milpitas, CA
	Position:	Founder, President & CEO Design consultation, commercialization of research, and Intellectual Property support for IC & modules. Worked as a technical expert/expert witness in a variety of patent litigation involving IC fab equipment, semiconductor process/device & design, JEDEC standards, Non-Volatile & Volatile Memory, data communications, USB Flash, CMOS Analog IC, I/O interface & drivers, SIMM/DIMM memory modules, data security/integrity in Flash memory, etc. Appointed by US District Judge as his technical consultant on a DRAM patent litigation.

Mr. Brahmbhatt is sole inventor on 10 and the lead inventor on all 11 patents listed:

Patent #	Issued	Title
5,646,886	1997	Flash Memory Having Segmented Array for Improved Operation
5,583,808	1996	EPROM Array Segmented for High Performance and Method for Controlling Same
5,457,652	1995	Low Voltage EEPROM
5,341,342	1994	Flash Memory Cell Structure
5,016,217	1991	Logic Cell Array Using CMOS EPROM Cells Having Reduced Chip Surface Area
4,910,471	1990	CMOS Ring Oscillator Having Frequency Independent of Supply Voltage (this is the only patent involving a second inventor)
4,885,719	1989	Improved Logic Cell Array Using CMOS EEPROM Cells
4,831,589	1989	EEPROM Programming Switch Operable at Low VCC
4,823,317	1989	EEPROM Programming Switch
4,460,979	1984	Memory Cell
4,442,481	1984	Low Power Decoder Circuit

EDUCATION & CERTIFICATIONS:

Year	College/University	Degree
1978	University of Cincinnati, Ohio (USA)	M. S. Electrical Engineering
1977	Gujarat University, India	M. Sc. Physics, specialization in Solid State Physics & Electronics
1993	Small Company Management, Stanford Univ	Certificate
1994	Marketing Management, Stanford University	Certificate
1995	Marketing Excellence, Univ of London, Canada	Certificate
2007	Certified Trained Nanotechnologist California Institute of Nanotechnology	Certificate
2010	Certified Green Building Professional	Certificate
2010	Judicial Council State of California, Registered Language Interpreter – 5 languages.	Certificate & Registered with State
2011	Energy Efficiency Professional (CA Workforce Devl. & San Jose City College)	Certificate

Professional Associations and Achievements:

- Co-Chairman, IEEE Region 6, Central Area
- Current Vice-Chair & former Chair of IEEE SF Bay Area Vehicle Tech. Society
- Received Special IEEE Appreciation Award 2007 & IEEE Outstanding Leadership and Professional Service Award 2008, 2012 and 2020.
- Appointed on State of California's "Blue Ribbon Task Force on Nanotechnology"
- Former President, Silicon Valley Engineering Council
- Past Chairman Economic Development Commission, City of Milpitas
- Member of Consultant Network Silicon Valley (IEEE-CNSV)
- Rotary International Fellowship Award
- Government of India, National Fellowship Award
- Eta-Kappa-Nu (HKN) Member

* * *