Floorplanning Principles

Richard Rodgers
Agilent, APD
richard_rodgers@agilent.com

Kevin Knapp
Synopsys Professional Services
knapp@synopsys.com

Chris Smith
Synopsys Professional Services
cfsmith@synopsys.com

ABSTRACT

With the increasing complexity of today's deep submicron designs the need for identification of best practices and default methodologies becomes even greater; not so as to restrict the creativity of designers but to provide a context in which creativity can be exercised while ensuring that critical design issues are addressed.

In addition, the scope of responsibility placed upon individual chip designers has increased significantly. The distinction between front and back-end designers has not only been blurred but has been all but eliminated. Chip designers today must understand the requirements of the physical implementation of their design and take them into account from the very beginning of the design process.

These two elements of the swift evolution of the design process over the last few years have manifest themselves in two key ways. First, the need for floorplanning expertise has been thrust upon the well-balanced chip designer, lest the independence of the two place the project in peril. The second manifestation is the automation of hierarchical design via a divide-and-conquer approach.

This paper provides guidelines and best practices regarding floorplanning principles for 90nm design, and how they are used within a larger RTL2GDSII design system developed by Synopsys Professional Services and employed at Agilent's ASIC Products Division.

Table of Contents

Table of Figures

1 Introduction

The complexity of today's designs and the increasing breadth of issues that must be accounted for as geometries shrink, have thrust the well-rounded chip designer into the world of physical design.

The physical design concerns for the classic ASIC designer come into play primarily at the beginning of the process via floorplanning and towards the end of the overall flow via place and route. This paper focuses on best practices and default methodologies within floorplanning to enable successful 90nm design implementation. It addresses many of the primary floorplanning issues associated with hierarchical SOC design. Most of the practices addressed here are general and should apply regardless of the actual point tool used for floorplan implementation.

There are several key elements driving the need for best practices and default methodologies. Among the most critical is the need to account for the impact of physical design early in the implementation process. Concerns include signal integrity issues, both within a given block and across hierarchical boundaries, voltage drop analysis across the power grid and its impact on timing, available routing and placement resources, antenna rules, and diode insertion. All of these concerns have an impact on the timing, power consumption, and area of the design.

In addition, there are various, often orthogonal, ways in which these concerns might be addressed; so some measure of uniformity is desired across multiple design teams. Standardization facilitates ease of use and familiarity since engineers may serve on multiple projects. It provides uniform resolution for manufacturing concerns and can also provide "correct by construction" resolution for some of these issues.

The final element driving the need for best practices is the use of hierarchical design techniques which employ a divide and conquer approach. Hardware and software limits are continually being raised such that many designs - previously only achievable by the use of hierarchal techniques - can now be completed in a flat manner through placement and routing. However for extremely large designs (>10M gates), or those which use many separate SOC functional blocks, it is still necessary to adopt a hierarchical approach in order to reduce turn around time and to localize the effect of ECO's to individual sub designs.

1.1 Hierarchical Design Flows

There are essentially two approaches to hierarchical design and the techniques outlined here are applicable to both. Some SOC designs may make use of a mixture of the two approaches where necessary.

The traditional, bottom-up, approach is to use the top level floorplanner to manipulate a collection of sub designs where each has been, or will be, implemented in isolation and their results fed back for iteration of the top level floorplan. The top level floorplan sees only a simplified abstract of the sub blocks and hence the floorplan is created without full visibility into global level optimizations that may be possible between sub designs. Interactions between the various sub designs and the top level floorplan change as a result of the sub-designs going through their various iterations. This can result in a very long development cycle for achieving the final floorplan.

A new approach, employed over the past several years, is to utilize what is typically termed a Virtual Flat (VF) flow. VF is a top-down approach whereby the implementation of the SOC design, and its sub designs, is guided by physical partitions that are determined by the floorplanner while having access to the complete design netlist. A structural netlist for the entire SOC design is provided to the floorplanner, along with top level timing requirements. The floorplanner then automates the partitioning of the design into smaller sub designs, often referred to as child blocks, and generates a floorplan, as well as timing and physical constraints for each. These child blocks can then be independently optimized and implemented all the way through place and route, guided and constrained by the physical and timing constraints generated from the top level floorplan. The child blocks are then abstracted into timing and physical models which are used at the top level to complete the top level physical synthesis and routing. Various types of hard macro intellectual property (e.g. memories, microprocessor cores) are typically instantiated within the top level and within child blocks.

Figure 1 below shows a graphic comparison of the two hierarchical approaches. The key difference to note is the iteration between block and top floorplans for the bottom up approach, versus the "single pass" approach using the top-down VF flow. In practice there are iterations in the VF floorplan but they are all contained, and controlled, inside of the VF floorplanning tool and only ever feed-forward to the place and route steps of the design flow.

This divide and conquer approach adds additional physical design concerns of its own to the overall methodology. Resolution of each of the physical design concerns mentioned earlier may take different forms depending on whether the block being implemented is hard macro intellectual property (IP), a soft macro child block, or the top level. This paper will endeavor to account for these varying requirements.

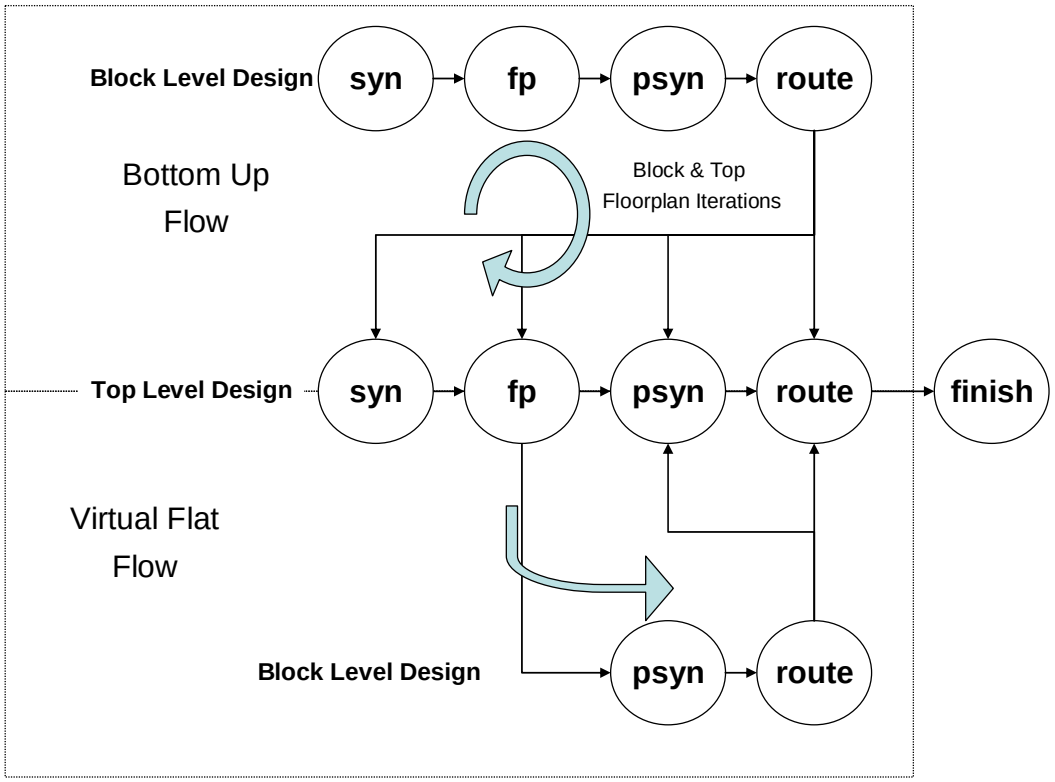

Figure 1 : Top Down Vs Bottom Up Design Flow

2 Power Planning

Creation of the power network within a design is one of the most important, and often time consuming, parts of the implementation. It is important that the power network be robust enough to support the power requirements of the design while ensuring that it does not consume too many routing resources such that the design becomes unroutable.

Power planning is not an isolated sub-flow; it is integrated with the overall design flow and must be taken into account early in the design process for a number of reasons:

- The total number of pads, which consists of signal pads as well as power and ground pads for the core and IO, may determine the overall design's physical size (i.e. the design may be pad limited). In this case, the number of power pads directly affects the chip's size and floorplan.
- The power structures within the core area consume physical area; this affects floorplan area and block placement. Inserting power structures as an afterthought to floorplanning can lead to many floorplan changes late in the design cycle.
- The power grid topology effects top level routability. For example, a full-chip fine grid structure using two layers of metal will severely limit available routing resources on those layers. This needs to be taken into account when planning the routing for high fanout signals, such as clock nets, which may prefer to use the higher level metal layers
- The power structure topology also affects placement and routing within the child blocks, and thus should be in place immediately after top level synthesis and prior to final child block partitioning in the VF flow.
- The power structure effects functionality and reliability. Excessive voltage drop at a given cell instance limits its performance and may lead to unpredictable circuit behavior. Furthermore, problems associated with electromigration directly lead to device failure.

As a result, it is necessary that the power structure for the full chip be developed early, and analyzed often throughout the design flow.

2.1 Power and Ground Insertion and Routing

The exact architecture of the power network is extremely design dependent and no one single approach will work on all designs. However some general principles can be outlined as a starting point for power network design. This section gives a simple overview of the terminology used in power network design. The power network is made up of four basic structural elements:

Power Pads: I/O pads that supply power to the chip.

Power Rings: Form complete rectangular or rectilinear rings around the periphery of the die, around the standard cell core area, around individual hard macros, or inside of hierarchical blocks. The rings are typically created in higher-level routing layers to leave the lower routing layers free for signal routing

Power Straps/Trunks: Horizontal and vertical metal wires placed in an array across the entire die or sections of the die. The horizontal wires are often referred to as straps; while the vertical wires are referred to as trunks. Again, these are typically created in the higher level routing layers. Power straps and trunks typically start uniformly spaced across the die, but modifications to allow for hard macro power rings, wiring keepout areas, and other restricted areas for power mean that the end result is not always strictly uniform.

Power Rails: These are the low level, typically metal 1, routing that is used to connect the standard cell power rails together, and to tie them into the power trunks. These rails will only be created within standard cell placement areas that are not already blocked by hard macro placements or wiring keepouts.

In the following discussions reference will be made to terms that can be identified on the very simplistic overview of a SOC design power structure architecture as shown in Figure 2 below.

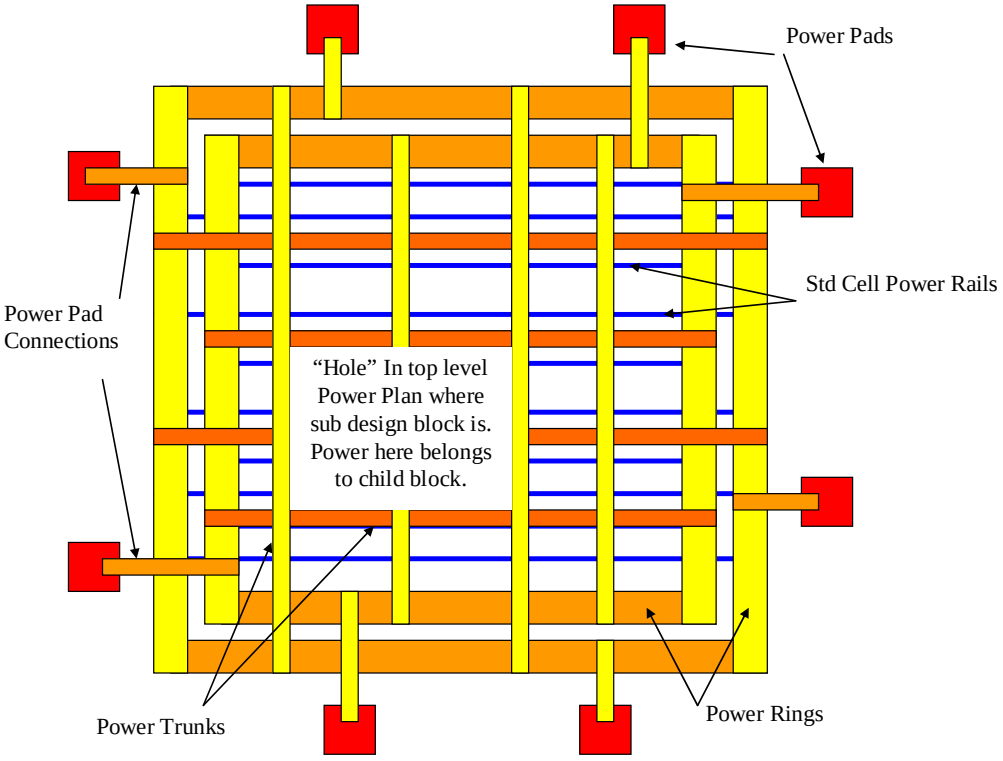

Figure 2 : Simplified Power Distribution Architecture

2.2 Top Level Power Rings

Assuming that perimeter IO is being used, at the top level power and ground rings should be placed around the perimeter of the standard cell core of the design. While floorplanning tools will accomplish the insertion of these rings, the user typically must specify their width and spacing.

A good rule of thumb is as follows. Assume that each side of the ring must carry a quarter of the current associated with the power. The overall power budget for the chip is divided by four and then converted to current using the primary voltage of the core. The current density for the metal layer(s) used for the power and ground rings can then be used to determine the required width. If possible, this width should be limited in order to avoid the need for wide metal slotting. Perimeter power rings are typically created in a lower layer metal to ease the connection to pads and standard cell power and ground, and to minimize the required via structures. Thus, metal 3 and metal 4 are often used. See the example below.

Sample power ring width calculation
Values will vary for your design. These numbers are only for illustration:
Power budget for chip: 800mW
Routing layer for power rings: metal3 and metal4
Max current density for metal: 26mA/micron
Primary voltage of core: 1V

Width of perimeter P/G ring: (800mW/4) = 200mW/side
(200mW/side) / 1V = 200mA/side
(200mA/side) x (1micron/26mA) = 7.69micron/side
Hence we would use 8micron/side for margin.

2.2.1 Hard Macro IP Perimeter Power Rings

It is also best to create power and ground rings around any hard macro IP present in the design. This enables orientation independence for the IP and eliminates the need for the power structure of the chip to specifically align with the power structure of the hard macro. Using the power budget of the hard macro, the same rule of thumb can be used as for the top in order to determine the width of the rings. Once the hard macro has been ringed in power and ground, the floorplanning tool can be used to accomplish the connection of the power and ground pins of the hard macro to the power rings.

2.3 Standard Cell Power Rails

Once the perimeter and hard macro power rings have been established power and ground must be routed to the standard cell rows. This connection is accomplished in part by abutment of the cells themselves. The floorplanner is used to add additional rails aligned with the power rails inside of the standard cells to ensure continuity of the standard cell power rails across the design. For some technologies and/or floorplans it is sometimes necessary to temporarily insert filler cells in order to achieve a complete grid. Following insertion of the rails the filler cells are removed.

The floorplanner accomplishes the rail connections with spacing based upon the standard cell height, but the user must specify the width of the rails. Most standard cell core layouts use mirrored placement rows resulting in abutment of alternating power and ground rails. Thus the

width of the standard cell power and ground rails is typically twice the width of the straps contained within the cells themselves and should be done on the lowest horizontal metal layer. Power rails that are wider can be used but this will reduce placement area within the design since it will often result in the standard cell rows being spaced further apart.

2.4 Power and Ground Trunks

In addition to the standard cell power and ground routing, additional horizontal and vertical straps and trunks are required in the power grid in order to adequately distribute power across the chip and minimize voltage drop across the power network.

As discussed above, power rings are usually determined by the available space in the I/O area and/or based upon typical rule of thumb approaches intended to maximize their current carrying ability. Similarly, standard cell power rails are usually determined by the standard cell technology being used. We can consider that these two parts of the power structure have very little flexibility and are somewhat pre-determined. Thus, it is in the area of additional power straps and trunks that the designer has the most control and flexibility; this is also the most important means to address detailed IR drop across the power network
.
Properly sized and spaced horizontal and vertical trunks will decrease the IR drop, but will also consume routing resources. Thus, a balance must be established between the need to retain routing resource for signal routes and the need to minimize IR drop across the power network of the design. Power and ground trunks are usually the dominant controlling factor on IR drop.

2.5 IR Drop

IR drop occurs in the design due to the resistive nature of the power routing from the power pads to the power pins of all the cell instances in the design. The IR drop for a given instance depends upon the current that the power network must deliver in order for the cells in that area of the design to operate at their targeted frequency. Taking these two factors into account, the IR drop across the power network will vary across the design. Typically this IR drop analysis is viewed at the top level by means of a contour or heat map type of display.

The IR drop is calculated as a change in the nominal voltage occurring at the cell instance power pins. The figure below shows a typical IR drop contour map display and a schematic of a single cell instance.

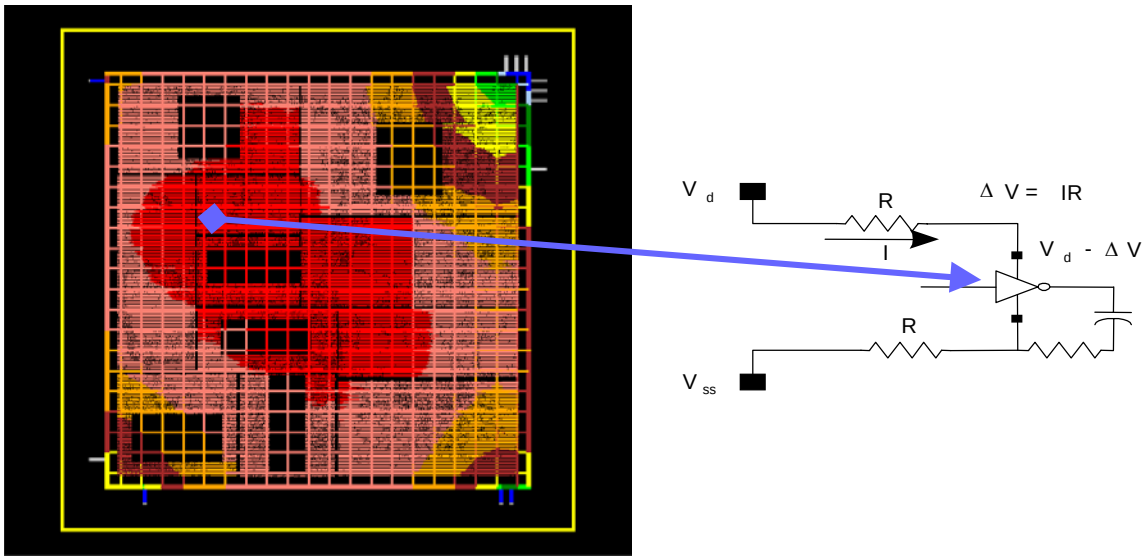

Figure 3 : IR Drop Display

2.5.1 IR Drop Effects on Timing

Since the delay through a cell is a function of the supply voltage, anything less than nominal voltage at a given cell instance will reduce its performance and potentially cause timing violations. It is likely that a design that passes timing may fail timing if IR drop effects are taken into account during the timing analysis. The diagram below shows how the effects of IR drop, modeled as a per-instance supply voltage, can affect the endpoint slack of a design. The endpoints that have only a small positive timing margin are at the highest risk of failing when IR drop is taken into account.

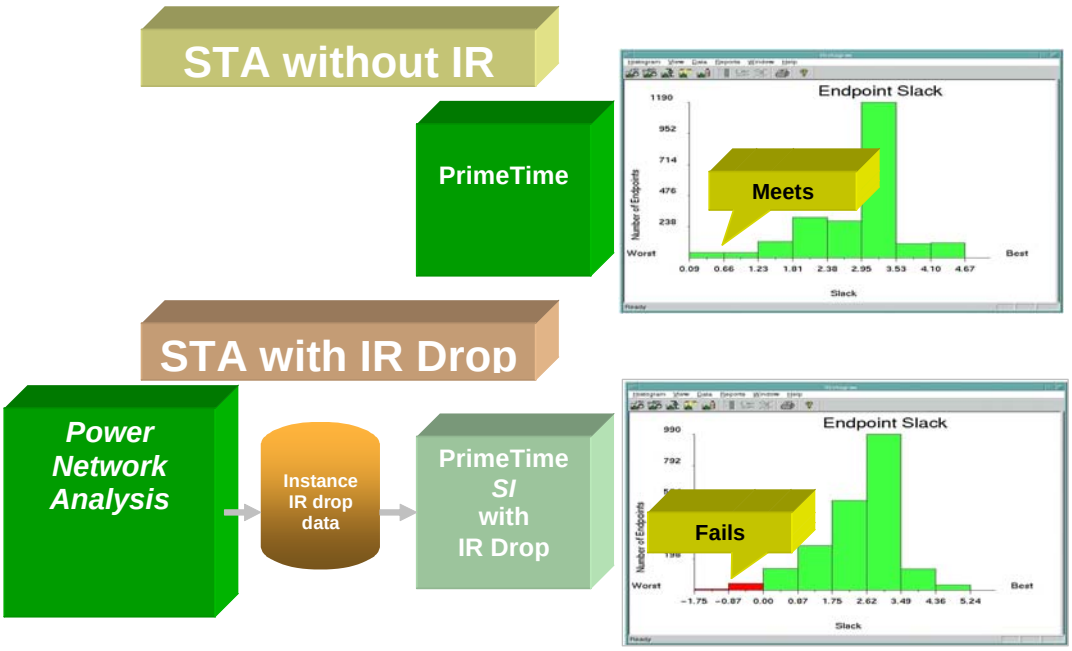

Figure 4 : Static Timing Analysis With IR Drop Effects

2.6 Power Network Analysis

Historically, IR drop information has been calculated following the final place and route of a design. In situations where IR drop issues exist this may mean having to iterate through large, and time consuming, parts of the design flow all the way back to floorplanning. As a result, chip designers have tended to over-design the power network at the expense of routing.

It is now possible, however, to extract such information during the floorplanning stage of the design flow, and thus provide IR drop feedback at an early stage where it is easily addressed. Analyzing IR drop during floorplanning also helps to alleviate the problems associated with over-designing the power network. It thus frees up more routing resources and reduces congestion concerns.

IR Drop analysis during floorplanning is referred to as Power Network Analysis or PNA; this is to distinguish it from the final, post-route IR drop analysis which is typically referred to as Rail Analysis. PNA uses faster, and potentially less accurate, resistance extraction engines to extract the power network for the IR drop analysis; however, when correctly used it has been shown to be within 15-20% of the final "sign-off" quality Rail Analysis. The diagram below shows some example results from a comparison of a floorplanning level PNA versus a full analysis using Astro Rail after detailed route.

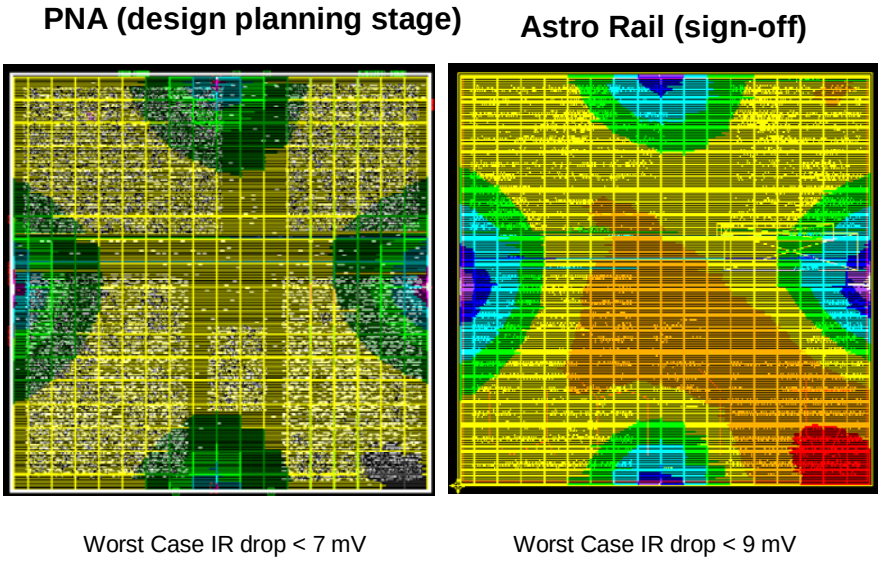

Figure 5 : PNA Versus Rail Analysis

For PNA to be successful it is imperative that the power structure be “as complete as possible” during the floorplanning stage. Typically this will mean insertion of all rings, rails, and trunks along with their required via structures. As a result, the power structure for the design is effectively completed during floorplanning and should require no additional work during the place and route stages.

For this to be a valid approach DRC checking must be performed on the power grid and all issues resolved during floorplanning. This task minimally ensures that there are no opens or shorts in the power grid, all power and ground pads are connected, and that all power and ground pins of standard cells and hard macros are connected. To proceed with DRC violations in the power grid will most likely prove catastrophic later in the flow requiring the user to return to floorplanning to correct the problem. Proceeding with DRC problems will also likely result in a larger mismatch between the PNA and the sign-off Rail Analysis.

Since the power consumed by a cell instance is dependent upon the load it has to drive, it is important that the design be prototype or global routed within the floorplanner in order to achieve the best correlation between PNA and sign-off Rail Analysis. This route should be performed after the power grid has been fully established in order to give the most accurate representation. This, in turn, will enable the most accurate assessment of the IR drop via PNA to assure tight correlation with Astro Rail once the final place and route is accomplished.

2.7 Electromigration

Electromigration problems may occur as a result of several reasons, but all are associated with the current flowing through the power rails. Excessive current density over a long period of time and the high power requirements of high frequency designs can lead to electromigration unless care is taken in the design of the power network.

Electromigration issues may manifest themselves as performance degradation over time as metal migration increases the resistance of the power nets; or, as a sudden catastrophic open or short circuit. Electromigration is calculated and displayed in a similar way to IR Drop except that it is a measure of the current density throughout the power network. Areas of the power network that exceed the electromigration current density can be considered as potential hot-spots for failure. In the same way that PNA can be performed during floorplanning, it is also possible, and recommended, that electromigration analysis be performed during floorplanning.

Correlation of electromigration analysis performed during floorplanning with the final signoff analysis is typically not as tight as seen in IR drop analysis. Final sign-off electromigration analysis considers a larger set of design rules and thus has a longer run time. Attempting to assess all these same design rules and incurring the run time hit when the design is not fully implemented, but only in the floorplanning stage, is not prudent. Thus, accuracy is traded for run time given the level of abstraction of the route information in the floorplanner. The diagram below shows a comparison for electromigration analysis during floorplanning and during final signoff with Astro Rail.

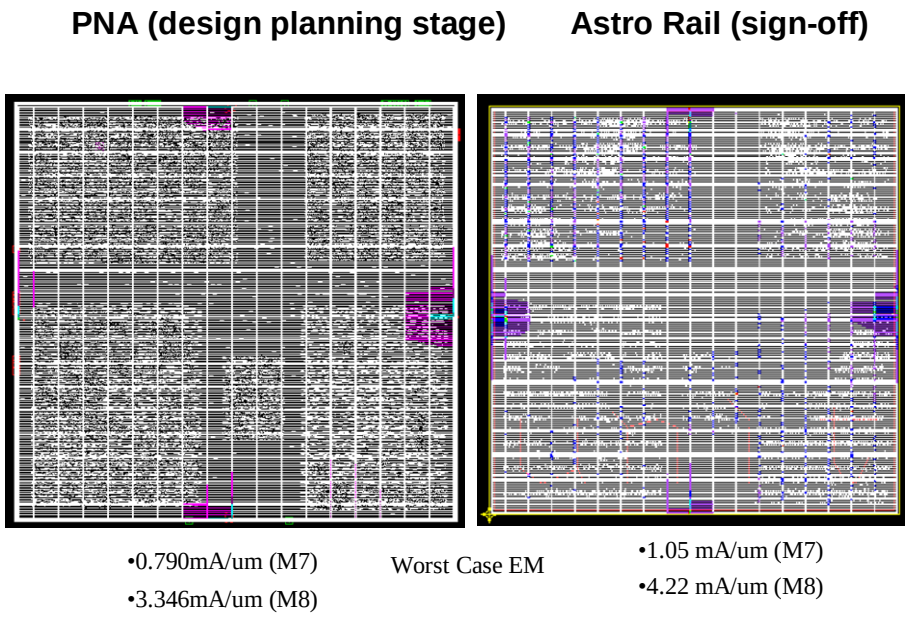

Figure 6 : Electromigration Comparison

2.8 How Many Trunks Are Enough?

Earlier we discussed the fact that the insertion of additional power straps and trunks is the major tool available to the user in creating a robust power network. In the process of developing the power network the user must decide the spacing, width, and layer of the additional straps and trunks. There has been much discussion over whether it is better to use many thin power routes or fewer wide routes. In general, many thin routes produce better routing porosity for signal routes and present less concern for congested areas. The main problem with the manual process of power trunk and strap insertion is that it ends up being somewhat trial and error. Even with the ability to analyze the power network using PNA the user is still faced with having to create the straps, run the analysis, and then iterate as necessary back through additional strap insertion or strap modification.

Jupiter XT has automated Power Network Synthesis or PNS which relieves many of the difficulties of the manual process. It has several degrees of freedom enabling users to customize it for their particular needs. One or all of the following basic constraints may be given: maximum IR drop, maximum number of straps/trunks, minimum metal width, metal layers to use. If the user already has a power grid established that is acceptable in general, but requires enhancement in a particular area, PNS can be run on a specific region of the design.

PNS will perform its analysis and synthesis of the power grid. It will then perform Power Network Analysis (PNA) under the hood and display the results. If the results are acceptable, the user may commit them to the design. If not, the constraints can be refined and the synthesis re-run.

As an alternative to full PNS, the user may establish a basic initial power grid and then explicitly perform power network analysis (PNA) to determine the areas requiring refinement based upon IR drop. PNS may also be used simply to size an existing power network structure to meet the IR Drop constraints.

3 Signal Integrity

Signal integrity issues can result in two major failure modes:

Timing Failures – crosstalk between nets can increase or decrease the delays depending upon whether the two coupled nets are switching in the same direction or different directions

Functional Failures – noise coupling between nets and/or cells can induce glitches that result in unintended logic transitions.

3.1 Hierarchical SI Avoidance

The hierarchical divide and conquer approach results in some special concerns for signal integrity issues. While implementing the child soft macros there is no detailed visibility into the top level routing regarding wires or cells close to the child soft macro that might represent victims or aggressors. Similarly, when performing the top level route with abstracted models representing child blocks, the top level does not have visibility into the children to know of potential victims or aggressors.

The diagram below shows the types of cross coupling that can occur. The first is related to the capacitive coupling between nets N4 and N5 located inside and outside of the block respectively. The 2nd is related to noise coupling between the instances U3 and U4.

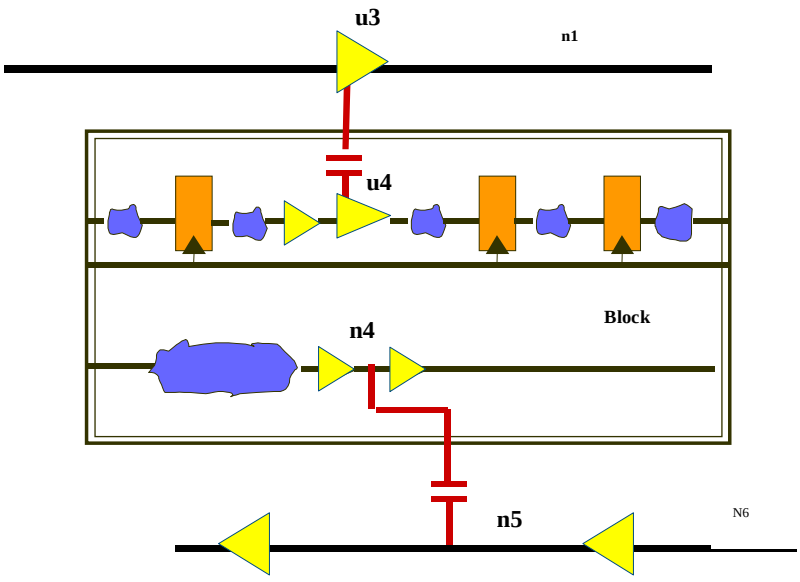

Figure 7 : Hierarchical Cross Talk Issues

Since cross coupling victims and aggressors generally result from coupling between long parallel routes, resolution of this issue is equivalent to the elimination of long parallel routes both near the edge and at the top metal layer of a child block. In the case of parallel routes near the edge

of a child block, this is accomplished by establishing a wiring keepout halo around each child block. In total, this halo should be partially within the child block and partially at the top level, and should only be on routing layers parallel to the particular side of the block.

This wiring keepout should be wide enough to ensure no cross-coupling between nets located immediately outside of the keepout. For reasons discussed later in this paper, it is recommended that a wiring keepout as wide as the height of one placement site be established around the outside of the perimeter of the child block and another wiring keepout of the same width be established around the interior of the perimeter of the child block (See Example 3.0.) A placement site will be approximately 8 to 11 wiring tracks high. These keepouts will thus result in an overall keepout width of 16 to 22 wiring tracks. This should be more than sufficient to ensure no cross coupling. A smaller overall width may be acceptable and the user may make these keepouts narrower as a particular block may require, while still ensuring no cross-coupling.

Eliminating parallel runs at the top metal layer of a child block can be accomplished using one of two methods. Over the block routing requires that an entire metal layer or layers be reserved for top level routing resources. This eliminates long, parallel routes between the child block and the top level routing above since routes in adjacent routing layers will primarily be orthogonal. Through-the-block routing requires proper shielding of top level signals from routes in the child block. One effective method to accomplish the desired shielding is to reserve all routing resources between a pair of power straps for top level routing only. Routing resources not specifically reserved between pairs of power straps for top level routes are available for routing in the child block. This method also relies upon orthogonal routing between adjacent metal layers.

Another part of the solution is to place buffers inside of the block as close as possible to the ports of the block. Thus, for top level signals the amount of wire that exists inside of the block is minimized. When the top level is routed one of the algorithms used to help prevent SI issues is to limit the maximum allowable route length. However without knowledge of the route length inside of the block it is possible that the algorithm may be "fooled" by wire length inside of the block. By placing buffers as close as possible to the ports of the block this potential error is reduced.

Example 3.0 – Wiring keepouts for Signal Integrity
For a technology with placement sites that are 0.28um wide and 2.24um tall, and whose preferred direction for metal one is vertical, wiring keepouts 2.24um wide for odd-numbered metal layers should be placed along the entire length of the left and right sides of the child block perimeter. Wiring keepouts 2.24um wide for even-numbered metal layers should be placed along the entire length of the top and bottom sides of the child block. These keepouts should be placed one each along the exterior and interior of the perimeter of the block resulting in an overall halo width of 4.48um.

As a result, per the diagram below, nets N1 and N2 would be spaced at least 4.48um apart which is more than sufficient to reduce cross coupling to zero in a typical 90nM technology. Also U1 would be placed very close to the ports on the block thus ensuring that route N3 is correctly handled by the route length limits used for SI avoidance.

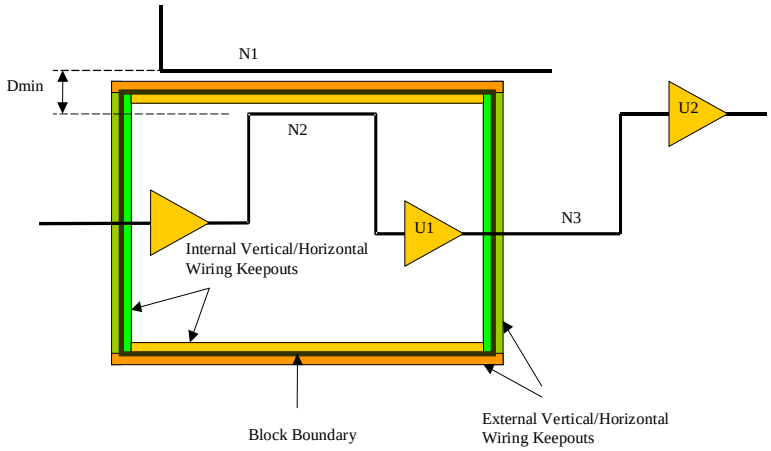

Figure 8 : Hierarchical SI Avoidance

Note that since the keepouts are only for wire layers that run parallel to the particular side of the block, they still allow for connection to the top level pins of the block in the other metal layers.

Astro has the ability to perform routing in the non-preferred direction for a given metal layer. And thus, even in the presence of the wiring keepouts described in the example above, Astro could create a long metal 2 wire along the left or right side of the block, for example. Such a route could pose a cross-coupling problem. Thus, another component is required of the overall solution. In Astro's route costing function, additional cost must be added for wrong-way routing so as to minimize the likelihood that Astro would establish a long route in the non-preferred direction for any given metal layer. This is accomplished via the following command:

axSetIntParam droute wrongWayExtraCost 50

A value of 50 is somewhat arbitrary and there is no hard and fast rule for the appropriate limit. However, through use, this value has been found to be effective for this purpose while still allowing Astro to perform short, non-preferred routing when needed in "tough" situations.

3.2 Wiring Keepouts Versus Ground Rings

An alternative approach to minimize or remove cross-hierarchy cross-coupling concerns is to create a ground ring around the outside of the macro block. This has the effect of coupling the internal and external potential long routes to ground instead of to each other. Depending upon the technology used this may result in a smaller overall block area and required keepout area. It should also serve to act as a placement keepout area for at least one placement site around the outside of the block, as referenced later in this paper.

3.3 Hard Macro IP

When implementing IP for use in other designs and projects, the designer has no control over the user of the IP block. It is unknown whether the user will include the wiring keepouts around the exterior of the block as described above. As a result, the design in which the IP is instantiated may present a cross-coupling problem to the IP block. Thus, it is recommended that the IP design include the full width of the wiring keepout halo within the perimeter of IP block.

Conversely when using an IP block in a design, the user may or may not know what wiring keepouts, if any, exist inside the perimeter of the IP block. If the user is able to obtain such information from the IP vendor he may respond accordingly with any required wiring keepout halos around the exterior of the IP block perimeter. If such information cannot be obtained, it is recommended that the full width of the wiring keepout halo, as described above, be place around the exterior of the IP block perimeter. While somewhat conservative, these measures will eliminate cross-coupling concerns between hard macro IP and the block within which they are instantiated.

4 Antenna Issues

4.1 Child Soft Macros

Similar issues as noted above for signal integrity exist for antenna prevention within a hierarchical design flow. When implementing child blocks and performing antenna checking and fixing, there may be no visibility into what additional routing metal exists at the top level on the nets connecting the top level pins of the child soft macro. This additional routing metal may pose an antenna problem. Similarly when performing top level route using abstracted models for the children, there is limited visibility into the child block to know what additional routing metal exists within the child for nets connecting the top level pins.

If the entire top level and child block designs are completed within Astro then it is possible to annotate the child ports with information relating to how much routing exists inside of the block. When top level antenna fixing is performed it is possible to access this information such that the router is able to avoid introducing, or is able to fix, antenna issues.

However when the top level is not completed in Astro, or if the end application is not known it is important that antenna prevention diodes be placed on the nets that connect the top level pins of child soft macros. A methodology has been developed in Jupiter/Astro whereby these diodes are placed within the child soft macros on all of the input, output or bidirectional pins of the macro as specified by the user. The diodes are then set as fixed placed in close proximity to the port they serve. While not a hard and fast number, it is recommended that these diodes be placed within 20um of the port they serve. The diodes are later connected by Astro. Similar strategies have been successfully implemented using Physical Compiler.

4.2 Hard Macro IP

Similar comments apply here as with signal integrity issues above for hard macro IP. When implementing the IP for use in other design projects using non-Astro tools, the length and area of the signal route outside of the IP is not known, and may thus pose an antenna problem. Therefore, it is necessary to place antenna protection diodes within the macro near the pin they serve.

When using hard macro IP received from other sources, you may or may not know if antenna diodes were used within the IP. If such information can be obtained, respond accordingly. If such information is not available or if it is determined that no antenna protection diodes were used within the IP, it is important to place antenna protection diodes outside the ports of the hard macro IP since the signal route inside of the macro may incrementally pose a problem.

4.3 General Hierarchical Flow Concerns

There are several miscellaneous issues pertaining to antenna prevention diodes within a hierarchical flow that are worth mentioning.

The concerns above are unique to nets that traverse block boundaries between the top level, child soft macros, or hard macro IP. Antenna issues also need to be addressed within each of these

blocks. For example, Jupiter, Astro and Hercules have several different ways of specifying antenna rules. The rules specify the allowable signal routing area as a multiple of the gate area of the minimum geometry transistor of the process. These limits can be defined for a given routing layer or can be cumulative for all layers of the net. It is important that the same limits and rule types be used across Jupiter, Astro, and Hercules to avoid any surprises and associated re-spin through the flow.

Furthermore, many libraries may not have CLF files for the diode cells. Thus, when generating a Verilog netlist out of Astro to be used by Hercules LVS it may be important not to include the diode cells.

Similar concerns also apply for Verilog generation and RC extraction for use in PrimeTime. While posing minimal RC loading, it is best to include the diodes in the RC extraction performed by STARRCXT. Thus the Verilog netlist used by PrimeTime must also have the diodes included for successful parasitic back annotation.

5 Additional Concerns

5.1 Child Soft Macros and Hard Macro IP

In addition to the power grid pre-routes, signal integrity, and antenna concerns mentioned above there is an additional issue for child soft macro blocks. A placement halo should be established interior to the block around its perimeter. This halo serves two purposes. The main purpose is to prevent DRC violations around the perimeter between the structures of the soft macro and objects at the top level. In addition, a placement halo frees up routing resources around the perimeter for access to the ports of the macro. Similar comments apply when implementing hard macro IP for use in other designs. It is recommended that the width of this placement keepout be equal to the height of a placement site.

5.2 Top Level

Similar comments as in section 5.1 above apply at the top level for a placement keepout halo around the exterior of the perimeter of child soft macros and hard macro IP.

When all of the above concerns have been addressed in the top level floorplan and its associated child soft macro and hard macro IP blocks, it is recommended that a final virtual flat placement, Power Network Analysis and power DRC run be performed. This will ensure that the power network is clean and that IR drop targets are met while honoring the various keepouts.
In addition, a final proto route is recommended in order to make a final congestion assessment. At this stage of the design flow, the user should identify and eliminate any congestion hot spots. Furthermore, the utilization of all child soft macros and the top level should be no higher than 65% as a rule of thumb. These congestion and utilization guidelines are set because there are several topics yet to be addressed prior to the final implementation, including clock tree synthesis, scan chain stitching and lock-up latches, physical optimization, and hold time fixing.

6 Conclusion

Swift changes in the evolution of the ASIC design flow require today's designer to be familiar with physical design issues. The impact of these concerns must be accounted for throughout the design flow to ensure the overall success of the project. This paper has addressed some prominent floorplanning concerns pertaining to 90nm design, and has provided some best-practice methodologies for handling these issues in the context of a hierarchical design flow.

7 Acknowledgements

Synopsys DFCOE developers of the TIGER flow.

8 Appendix

8.1 Checklist for Top Level Floorplanning Concerns

Following is a checklist that may be used during floorplan development for top level blocks to track the various issues identified above and ensure that they are properly addressed. Note that there are other floorplanning concerns that need to be addressed and this list should not be considered as exhaustive. Rather, it is intended to cover the issues addressed in this paper.

Floorplan Development Checklist for Top Level Blocks

List hard macro IP instantiated at the top level:

_____
_____
_____

List the names of all child soft macros:

_____
_____
_____
_____

Power and Ground Routing:

1. Perimeter power and ground rings for core:
    - Power budget for chip: _____ (A)
    - Routing layers for power rings: _____
    - Maximum current density for these layers: _____ (B)
    - Primary voltage of core: _____(C)

    Resulting width of perimeter power and ground ring = (A/4) / C / B: _____

Note: If the core utilizes regions of different voltages each region should be ringed and this calculation performed for each region.

2. Power and ground rings for hard macro IP:
    - Any hard macro IP instantiated at the top level? _____
    - If not move on to next section.

    Each hard macro should be ringed. Perform the same calculation as above for each:
    - Power budget for hard macro: _____
    - Routing layers for power rings: _____
    - Maximum current density for these layers: _____
    - Primary voltage of hard macro: _____

Resulting width of perimeter power and ground ring = (A/4) / C / B:

_____

3. Standard Cell Power and Ground:

Width of power and ground rails within standard cells: _____

Routed standard cell power and ground lines should be 2x this width for mirrored placement rows resulting in abutment of alternating power and ground rails.

4. Additional Horizontal Straps and Vertical Trunks:

Maximum IR drop target: _____

Maximum IR drop achieved: _____

Line widths and layers used:

_____

Special concerns and configurations:

_____

_____

5. Power and Ground Verification:

Proto-route performed? _____

Final PNA performed?

Maximum IR drop target: _____

Maximum IR drop achieved: _____

Power DRC performed and all issues resolved?

Notes (for lessons learned or special considerations):

_____

_____

_____

_____

_____

Signal Integrity Concerns:

1. Wiring Keepouts Around Child Soft Macros:

Height of standard cell placement site: _____

Preferred direction for metal 1: _____

Wiring keepouts that are equal to the placement site height in width should be placed around the exterior of the perimeter of all child soft macros. The keepouts along the left and right sides of the macros should only be for vertical metal layers. The keepouts along the top and bottom should only be for horizontal metal layers.

Wiring keepouts established around all child soft macros? _____

Notes (for lessons learned or special considerations):

_____

_____

_____
_____

2. Wiring Keepouts Around Hard Macro IP:

For each hard macro list known cross-coupling abatement measures within the interior perimeter of each hard macro:

_____
_____
_____
_____

A total wiring halo width of twice a standard cell placement site height is required around all hard macro IP. Create wiring keepouts for each hard macro to ensure this requirement, adjusting for the known cross-coupling abatement measures within the interior of the hard macro. The keepouts along the left and right sides of the macros should only be for vertical metal layers. The keepouts along the top and bottom should only be for horizontal metal layers.

List hard macro name and width of exterior wiring keepout halo:

_____
_____
_____

Antenna Issues:

1. Antenna Prevention Diodes Around Hard Macro IP:

For each hard macro list known antenna prevention measures within the perimeter of each hard macro:

_____
_____
_____
_____

For all digital signals, if no antenna diodes exist for each pin within the perimeter of the hard macro place one on the exterior near the pin it serves. Mark it as fixed placed.

General Concerns:

1. Placement Keepouts Around Child Soft Macros:

Height of standard cell placement site: _____
Preferred direction for metal 1: _____
Placement keepouts that are equal to the placement site height in width should be placed around the exterior of the perimeter of all child soft macros.
Placement keepouts established around all child soft macros? _____
Notes (for lessons learned or special considerations):

_____

_____
_____
_____

2. Placement Keepouts Around Hard Macro IP:

Placement keepouts that are equal to the placement site height in width should be placed around the exterior of the perimeter of all hard macro IP blocks.
Placement keepouts established around all hard macro IP blocks? _____
Notes (for lessons learned or special considerations):

_____
_____
_____
_____

3. Once all issues within the child soft macros have been addressed via their checklists the following should be performed:

Final virtual flat placement performed to honor all keepouts? _____

Note utilization of top level and of each child soft macro. All should be at or below 65% maximum. Give reasons for exceptions:

_____
_____
_____

Final proto-route performed? _____

All congestion issues addressed? _____

Note exceptions here. Give reason and plan for resolution:

_____
_____
_____

Final Power Network Analysis performed? _____

Maximum IR drop achieved: _____

If any changes made to power grid to address issues above a final power DRC run must be performed.

8.2 Checklist for Child Soft Macro Floorplanning Concerns

Following is a checklist that may be used during floorplan development for child soft macro blocks to track the various issues addressed above and ensure that they are properly addressed. This checklist assumes that these child soft macros are being implemented in the context of a Virtual Flat hierarchical design flow. This checklist should be applied for each child soft macro block.

Note that there are other floorplanning concerns that need to be addressed and this list should not be considered as exhaustive. Rather, it is intended to cover the issues addressed in this paper.

Floorplan Development Checklist for Child Soft Macro Blocks

Name of child soft macro block: _____

List hard macro IP instantiated in the child soft macro, if any:

_____

_____

Power and Ground Routing:

1. Power and ground rings for hard macro IP:

    Any hard macro IP instantiated at the top level? _____

    If not move on to next section.

    Each hard macro should be ringed. Perform the following calculation for each:

        Power budget for hard macro: _____

        Routing layers for power rings: _____

        Maximum current density for these layers: _____

        Primary voltage of hard macro: _____

        Resulting width of perimeter power and ground ring = (A/4) / C / B:

_____

    List hard macro and it associated ring width:

_____

    _____

    _____

Signal Integrity Concerns:

1. Wiring Keepouts Within Child Soft Macros:

    Height of standard cell placement site: _____

    Preferred direction for metal 1: _____

Wiring keepouts that are equal to the placement site height in width should be placed around the interior of the perimeter of all child soft macros. The keepouts along the left and right sides of the macros should only be for vertical metal layers. The keepouts along the top and bottom should only be for horizontal metal layers.

Wiring keepouts established around the interior of the perimeter of this child soft macro?

_____

Notes (for lessons learned or special considerations):

_____
_____
_____
_____

2. Wiring Keepouts Around Hard Macro IP:

For each hard macro list known cross-coupling abatement measures within the interior perimeter of each hard macro:

_____
_____
_____
_____

A total wiring halo width of twice a standard cell placement site height is required around all hard macro IP. Create wiring keepouts around the exterior of each hard macro to ensure this requirement, adjusting for the known cross-coupling abatement measures within the interior of the hard macro. The keepouts along the left and right sides of the macros should only be for vertical metal layers. The keepouts along the top and bottom should only be for horizontal metal layers.

List hard macro name and width of exterior wiring keepout halo:

_____
_____
_____

<u>Antenna Issues:</u>

1. Antenna Prevention Diodes Around Hard Macro IP:

For each hard macro list known antenna prevention measures within the perimeter of each hard macro:

_____
_____
_____
_____

For all digital signals, if no antenna diodes exist for each pin within the perimeter of the hard macro place one on the exterior near the pin it serves. Mark it as fixed placed.

General Concerns:

1. Placement Keepouts Within Child Soft Macros:

   Height of standard cell placement site: _____

   Preferred direction for metal 1: _____

   Placement keepouts that are equal to the placement site height in width should be placed around the interior of the perimeter of all child soft macros.

   Placement keepouts established within this child soft macro? Specify width: _____

   Notes (for lessons learned or special considerations):

   _____
   _____
   _____
   _____

2. Placement Keepouts Around Hard Macro IP:

   Placement keepouts that are equal to the placement site height in width should be placed around the exterior of the perimeter of all hard macro IP blocks.

   Placement keepouts established around all hard macro IP blocks? _____

   Notes (for lessons learned or special considerations):

   _____
   _____
   _____
   _____