

# Design Planning Strategies to Improve Physical Design Flows— Floorplanning and Power Planning

August 2003

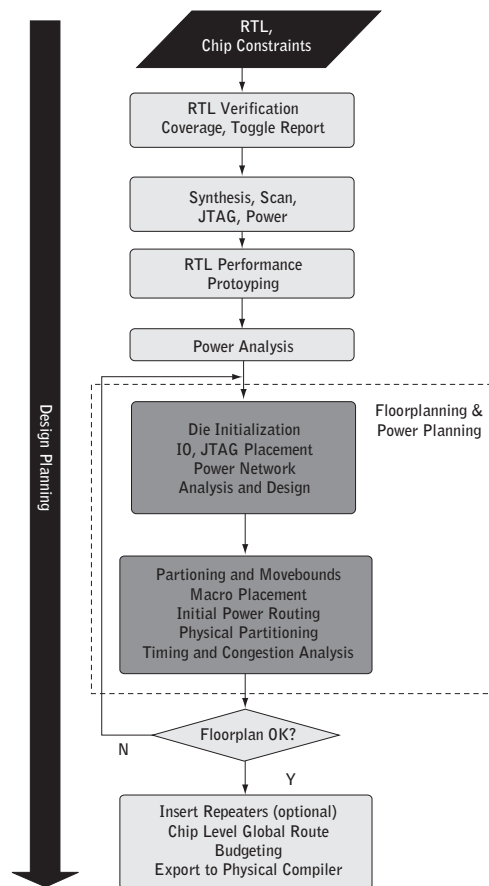
Sachin Idgunji, Staff Engineer, Synopsys Professional Services  
Steve Lloyd, Staff Engineer, Synopsys Professional Services  
Rick Mitchell, Staff Engineer, Synopsys Professional Services  
Ron Spillman, Staff Engineer, Synopsys Professional Services  
Jon Young, Design Center Manager, Synopsys Professional Services

The physical implementation of very deep sub-micron (VDSM) designs benefits from good, up-front planning processes, both in terms of schedule as well as quality of results. While the optimal planning strategy is design-dependent, there are a variety of design best practices that have proven successful in streamlining the physical design process.

While it is not the goal of this paper to be a complete reference for design planning, this paper describes general but practical strategies focused on improving the design planning process using two important practices, *floorplanning* and *power planning*.

## Why Plan?

Design planning is an integral part of an RTL to GDSII design process. In the design planning context, floorplanning is the process of placing cells and blocks in a way that makes later physical design steps more effective; the floorplan also provides a basis for better estimating interconnect timing for verification. Similarly, power planning helps ensure that all on-chip components have adequate power and ground connections. In many cases, especially for complex system-on-chip (SoC) designs, design planning should occur in parallel with RTL development, with the floorplans and power estimates being refined along with the RTL.



**Figure 1. Design planning occurs in parallel with the different stages of RTL maturity. Early floorplanning and power planning help guide RTL development, and vice versa.**

Before looking at guidelines for design planning, it is worth reviewing the reasons for taking time to create a good floorplan and power layout. Knowing that design planning can prevent major design problems is certainly one real motivator for improving floorplans and power structures (e.g., rings, meshes, etc.) Also, the insights gained from planning may inform high-level business decision-makers about whether a design will be financially viable.

Design planning has become crucial for large chips with hierarchical design flows because they are more likely to have long inter-block paths whose delays make timing closure difficult, leading to time-consuming and unpredictable tape-out schedules. In effect, hierarchical design moves design planning into the front-end of the design flow. Today, as mainstream chips migrate to 0.13 $\mu$ m technology and below, their size and complexity are making them more vulnerable to timing and power problems and more in need of careful front-end planning than ever.

A good floorplan helps ensure timing closure in many ways: by arranging blocks such that critical paths are short, by preventing routing congestion that can lead to longer paths, by integrating hard IP in an efficient way, by eliminating the need for over-the-top routing for noise-sensitive blocks, etc. The challenge is to create a floorplan with good area efficiency, to conserve silicon real estate and leave sufficient routing area—and do it as early in the design process as possible. Critical inter-block connections and on-chip buses require special allowances to ensure they will be able to meet timing constraints. Signal integrity is also important, especially for buses, where crosstalk from simultaneously switching drivers is more likely to cause signal glitches. Analog IP raises additional signal integrity issues due to noise sensitivity. Failure to deal adequately with any of these issues during floorplanning can lead to numerous design iterations and/or chip respins.

Besides the constraints on the floorplan, there are power issues that are also more challenging in today's chips due to factors such as additional metal layers and thinner wires. Power planning can help avoid iterations too, as well as reliability problems caused primarily by IR drop and electromigration.

Steady-state IR drop is caused by the resistance of the metal wires comprising the power distribution network. By reducing the voltage difference between local power and ground, steady-state IR drop reduces both the speed and noise immunity of the local cells and macros. Further, dynamic IR drop occurs when the simultaneous switching of on-chip components such as clocks, clocked elements, bus drivers and memory decoder drivers causes a dip or spike in the power/ground grid. The current pulled by simultaneously rising edges leads to a dip in the power grid, while a similar phenomenon on falling edges leads to a spike in the voltage of the ground grid. Both of these phenomena are referred to as IR drop and can reduce logic gate noise margins. The resulting logic faults or timing errors are extremely difficult to anticipate with traditional signal integrity and timing analysis.

Electromigration occurs when large current densities cause a flow of metal atoms from the negative- to the positive-biased end of a length of interconnect. This flow can result in catastrophic failures by either creating voids (opens in the metal line) or extrusions (shorts with neighboring metal lines). Electromigration has become a bigger problem as interconnect dimensions shrink, causing current densities to rise.

Floorplanning helps avoid IR drop and electromigration problems through strategies such as placing the most power-hungry blocks near the periphery of the die and preventing concentrations of such blocks in any one area. Power planning prevents problems by making sure that each part of the power mesh has sufficient current-carrying capacity.

Design planning is an obvious part of a developer's COT ("customer-owned tooling") design flow, but since it encompasses both front-and back-end design, it impacts ASIC manufacturing handoffs as well. ASIC vendors have traditionally handled floorplanning and power planning as part of their back-end design, but the complexity of today's designs have forced physical planning earlier in the flow, in parallel with RTL refinement.

## The Design Planning Flow

Several types of design planning flows can be applied, depending on the tools and methodologies that suit a specific design.

Choosing between a flat or hierarchical approach has a major influence on design planning. Synopsys Professional Services works with flat or hierarchical design flows as well as a "virtual" flat flow. The traditional flat ASIC flow avoids the effort to set up hierarchy but carries a higher risk of having to iterate all the way back to the architecture design to correct timing problems that appear after routing. Establishing a hierarchical flow helps prevent timing surprises. Determining factors between flat and hierarchical design include the design size and performance requirements.

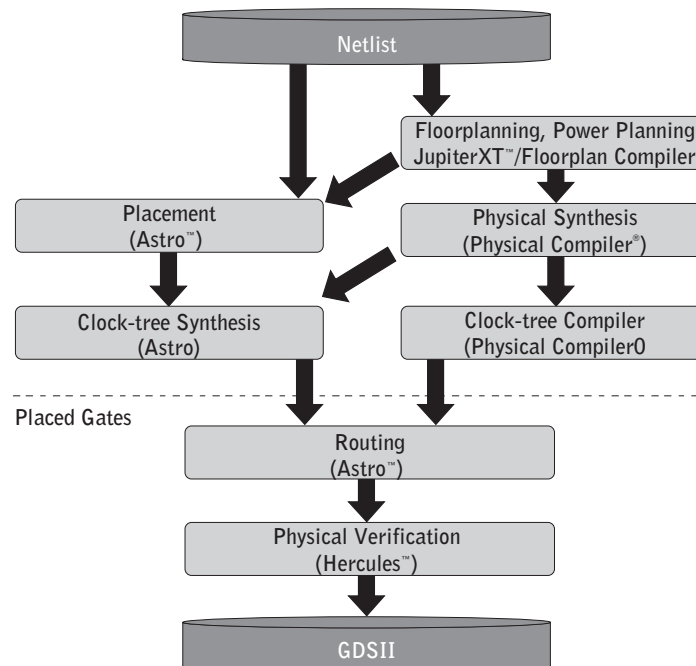


Figure 2. Netlist to GDSII processes supporting flat, hierarchical, and virtual flat flows


The "virtual" flat approach is an advanced floorplanning methodology for hierarchical design. It is based on making critical floorplanning decisions using the full context of the chip. After importing the netlist into the floorplanning tool, hierarchy is temporarily ignored while the floorplan is generated. After timing budgets and area allocations have been made, hierarchy is restored to the netlist to continue with the hierarchical design process.

## Floorplanning Considerations

When creating a floorplan, it's important to consider some basic characteristics of the process technology. For example, a typical standard cell library defines cell rows as horizontal, and the routing direction of each layer follows the alternating pattern: metal1 horizontal, metal2 vertical, etc. Since metal1 is used for routing within cells and/or for providing power buses for standard cell rows, the ability to use this level for general routing is limited. If a chip is to be manufactured with a four-layer process, for example, there is more routing resource in the vertical direction than in the horizontal. A five-layer process offers equal access in both directions.

At any level, creating “non-preferred” routing (i.e., not using the preferred routing direction for that level) is not recommended. When creating metal rings around cores and blocks, it is essential to allow room for routing access to pins. When pushing metal down into blocks, designers can avoid congestion at the block corners by looking at blockage effects inside the block.

When placing blocks, four-way intersections in top-level channels should be avoided; “T” intersections create much less congestion. This consideration can be critical to leave adequate space for routing channels, especially if there is not much opportunity for over-the-cell routing. Using flylines can help determine optimized placement and orientation, but when the flylines are numerous enough to “paint the area” between blocks, designers must rely on their best judgment for block placement, and later evaluate the results for possible modification.



**Figure 3. Flylines display the flow of nets and the effect of block placement and orientation**

Once blocks are placed, block-level pins may be placed. It is necessary to determine the correct layer for the pins and spread the pins out to reduce congestion. Placing pins in corners where routing access is limited should be avoided, instead multiple pin layers for less congestion should be used.

Placing cells within the perimeter of hard macros is not recommended. To keep from blocking access to signal pins, it is a good idea to avoid placing cells under power straps unless the straps are on high metal layers (i.e., higher than metal2). Density constraints or placement blockage arrays may be used to reduce congestion since these strategies will help spread cells over a larger area, thereby reducing the routing requirements in that area.

When placing top-level buffers, the corresponding regions for power and ground connectivity should be created. Supplying power and ground to any areas that might be useful in placing buffers or repeaters in post-placement timing convergence optimization can be beneficial. Any blockage that increases congestion should be avoided.

In any physical-design work, it is essential to understand the requirements of the target process technology. The lower utilization would result in a larger chip, but the chip is less likely to have problems in routing. For example, most processes now require the insertion of holes in large metal areas in a step known as “slotting” or “cheesing.” Slotting relieves stress-related problems in the metal due to thermal effects but may change the metal’s current-carrying characteristics. It is imperative to consult the design rule document for this and many other physical variables.

## Practical Planning Steps

Along with the multitude of design considerations to keep in mind when creating a floorplan, the designer has to start with a philosophy that meets business goals. For example, determining which parameters to optimize – e.g., speed, schedule, power consumption, die size, etc. – impacts the margins to set for area utilization and other parameters.

For instance, if the choice is to optimize the project schedule, setting utilization lower than might be chosen if optimizing for die size is an obvious strategy. The lower utilization would result in a larger chip but it is less likely to have problems routing. Similarly, increasing the power-planning margin so that, for example, considerably more metal is used in the power mesh to avoid failure under DC power loads also increases die area but improves schedule by avoiding downstream power problems. (Filling open spaces with power-mesh metal also reduces the risk of power issues.) However, care must be taken to ensure the extra metal does not push signal wires so close together that capacitance, power consumption and signal integrity problems result.

As these considerations imply, floorplanning and power planning are part of an integrated process. The following suggestions will help achieve good results in typical wire-bond packages:

### Chip-level Floorplanning

- The RTL should be examined for logical models to break out into hierarchical physical elements. If there are multiple instances of any logical hierarchical element, these elements can be grouped to form one physical element. It is easier to floorplan with same-size blocks, so small blocks should be grouped and large blocks divided when appropriate. Working with “medium-sized” blocks is typically best; six to twelve roughly equivalent-sized blocks is a reasonable target.
- Typically, floorplans should be started with I/Os at the periphery (depending on the package design.)
- It is best to place parts of the design that have special layout requirements – e.g., memories, analog circuitry, PLLs, logic that works with a double-speed clock, blocks that require a different voltage, any exceptionally large blocks, etc. – first to ensure that their needs are accommodated. Design blocks with special needs must be understood at the beginning; for example, flash memory has a high-voltage programming input that must be within a certain distance of an I/O pin, so it is best to place this element first.
- If there are two or more large blocks or other features that make a reasonable floorplan impossible, it may be necessary to increase the die size or re-arrange I/Os. Finding this problem early in the flow makes it easier to make a business decision about whether the chip will be financially viable with a larger, more expensive die. If any of the large blocks are soft (synthesizable) IP or otherwise available as RTL, it might be possible to avoid going to a larger die by repartitioning that block into smaller pieces.
- The floorplan can be completed by arranging the rest of the blocks in the remaining space based on their I/Os and power consumption. Other considerations being equal, it is best to avoid placing blocks that consume a lot of power near the chip’s center.

## Block-level Floorplanning

- Initial synthesis should be run to determine the total area of the cells in a block.
- Determining the area of a block beyond the area of its cells is a factor of utilization. Utilization varies depending on the library, technology, and characteristics of the design implemented, but for the typical library, the sweet spot is usually about 70 percent utilization. An unusually high percentage of registers or hard IP will increase this number; large numbers of multiplexers or other small, pin-dense cells will decrease it.

## Power Budgeting

- Calculating the power dissipation at the block level throughout the design process is important to determine if the design is meeting the specified power budget, and to estimate the size of the power grid needed in the floorplan. Early in the design process, manual calculations or spreadsheets are often used to estimate power; as the RTL matures, design tools can be employed to refine power estimates (+/- 30% is a reasonable target.) As the RTL migrates to gates and transistors, the power estimates can be further improved.
- For final power sign-off of the floorplan, actual netlist, net switching activity, and annotated parasitics should be used.

## Power Planning and Analysis

- The I/O power pins should be analyzed for the effects of simultaneously switching I/Os, focusing on the outputs because they draw the most current. Most I/O library suppliers will have performed extensive characterization of the I/Os, and will recommend a ratio of power/ground pad to simultaneously switching output pad.
- When creating a power mesh for the floorplan, the chip's layout should be considered. Even if power consumption were spread evenly over the entire chip when viewed from the block level, IR drop would still be worse in the chip's center due to the length of the wires. IR drop in the center of the chip causes the logic there to run slightly slower, and this effect will become more important in overall timing as threshold voltages decrease. Some wires in the mesh consequently carry more current than others, so the current on every wire, junction and via should be calculated. Vias that are not big enough act like fuses, ready to blow when the current is too high, so via arrays should be analyzed for IR, current density and electromigration.
- Electromigration problems in power meshes can be avoided by meeting the maximum current density limits for the process. Foundries publish these limits, and advanced design tools can estimate current density throughout the power mesh and highlight areas of concern.
- To achieve a given steady-state, worst-case power drop and maintain current densities within acceptable limits, a certain total width of metal for power bussing is necessary. Getting that total metal width involves tradeoffs between the number of metal straps versus the spacing between straps. Related issues include:
  - Straps wider than the process slotting size will be slotted at some point in the design process, thus reducing their conductivity from the value previously calculated for them.
  - Many failures due to electromigration are caused by metal atoms migrating along metal grain boundaries that are parallel to the direction of current flow. As metal wire width decreases, fewer grain boundaries are parallel to the direction of current. When wire widths are in the "bamboo region," where almost no grain boundaries are parallel to the direction of current flow, the perpendicular grain boundaries give the wire an appearance similar to a stalk of bamboo. These wires can handle up to an order of magnitude greater current density before electromigration failures compared to wider wires.
  - Power and ground wires both create return paths for inductive coupling effects; the more power and ground wires used, the more signal wires are protected from inductive noise effects.

## Core Hardening Example

The hardening of an ARM946E core using TSMC 0.13  $\mu\text{m}$  technology serves as an example of floorplanning and power planning practices. This core typically includes 20 RAMs, which occupy about half of the die area, and about 50,000 standard cells.




Figure 4. An ARM946E-S™ Core

Core hardening involves steps similar to those of any planning flow, except that the plan needs to be more flexible to allow for variations in the way the core will be used in a specific design. In the ARM core hardening described here, for example, the Synopsys Professional Services team configured the layout to allow the core to be rotated through 90 degrees.

The team implemented a power grid over the entire core on the top three metal layers using 1.2-, 1.0- and 1.6- $\mu\text{m}$  straps placed every 24, 22 and 32  $\mu\text{m}$ , respectively. One of the floorplanning challenges in this project was to arrange the horizontal straps such that they did not compromise the standard cell pre-routing. Perl scripts were used to automate this process based on a set of user parameters.

Another floorplanning challenge involved the provisions for rotating the core. All the power and signal pins were duplicated on a second layer and suitable via arrays generated to connect the two layers together inside the core. The customer could therefore use the core in any orientation and still be able to connect to it using preferred-direction routing for both signal and power supply pins.

The two sets of pins were created by starting with a single set of pins on one layer. A Perl script routine was then used to collect the information about this first set of pins and duplicate the pins on another layer. Another Perl script was also used to create the via arrays to connect the two sets of pins, the size of the arrays for the supply connections being chosen such that current density was low enough to prevent electromigration.

All the signal pins were made wide enough for access at the chip level without the need for off-grid routing, no matter how the core was aligned to the grid. The team also generated the RAMs such that the supply rings inside them were on layers chosen to suit the orientation of the RAMs in the floorplan, thus making it easier to connect the RAMs into the power grid.

Clock nets were routed double-width, double-spaced, and with double-vias to minimize signal integrity and electromigration issues. The clock trees were built using integrated clock-gating cells inserted by Power Compiler.™

All signal inputs and outputs to the macro included protection from charge-collecting antenna problems by the addition of a reverse-biased diode cell near the I/O pin. In addition, hardened cores are characterized within Astro™ so that the amount of charge-collecting metal and protection diodes on each input and output are known when the macro is instantiated hierarchically into a design.

## An ASIC Example

A brief ASIC design example from a Synopsys Professional Services customer engagement illustrates the benefits of careful floorplanning and power planning in advance of a handoff to an ASIC vendor. The design was a 5-Mgate chip targeted at an ASIC vendor's 0.18 μm process and included a large macro cell. The macro was placed and routed at a lower level of the design hierarchy. In Figure 5, the white area outside of the big box is the routing at the top level. The image thus shows the top level of a two-level overall physical hierarchy, although the macro also contains multiple levels.




Figure 5. ASIC example

Prior to the placed-gates handoff, design planning, place and route, clock tree synthesis, and timing and congestion closure were performed on the chip. A major challenge in this physical design was the congestion around the corner of the large macro. Many signals going to logic (on the left in the figure) had to fit through a small space. The floorplanning effort began by pushing cells away from this area to create more routing room. Spreading the cells apart led to timing problems reported by static timing analysis, so buffers needed to be added to speed up critical paths. The cells were re-arranged to make room for the buffers. Through this process, timing was achieved and congestion levels were brought within reasonable limits.

Astro-Xtalk was used to detect and fix crosstalk problems. Astro-Rail was used to check for IR drop, power consumption and electromigration. In such cases, the power consumption values are obtained by either running test vectors or by setting a certain level of switching activity. Setting the switching level at 30-40 percent is pessimistic, so if power levels are acceptable with that activity, the design will likely not be plagued with power issues. If any power levels are close to exceeding the budgeted limits, a more detailed analysis with test vectors should be run.

When the design team had resolved all power issues, minimized interconnect delays and clock skew, and achieved timing closure, the team did a placed-gates handoff to the ASIC vendor. The advanced methodologies employed cut 3-4 months off the design cycle, compared to the schedule offered by the ASIC vendor. Even though the customer had a clear idea of the expected die size, because the design was a new version of an existing chip, the design team also reduced the expected die size by 10 percent.

## Power Planning for Multiple Voltages

The total power consumed by CMOS circuits is generally composed of two components. The first is active power or switching power, which is power consumed when capacitive loads in a circuit charge and discharge during operation. Active power density has increased over the years due to the increase in performance with every technology node and the growing number of active circuits in a given die area.

The second power component is leakage power caused by current flowing through a device in its "off" state. Leakage power is growing exponentially because of the decreasing threshold voltage in advanced technology nodes, which causes an increase in sub-threshold conduction. The other components of leakage power are gate tunneling and leakage via reverse-biased diffusion junctions (source and drain).

Several approaches are available for managing power in these advanced technology nodes. One good method is to use multiple voltage supplies for sections of logic that are grouped according to their functionality. This method introduces the concept of voltage islands.

Using this multi-voltage approach requires setting up separate voltage domains on the chip that are fed by different power supplies. Signals traveling across domains may need level shifters—cells that ensure a proper rail shift—though level shifters may be unnecessary when a signal goes to a domain that has a lower voltage. Evaluating these factors accurately requires an understanding of the cell library at the transistor level because the gate source voltage must be evaluated within 100 mV.

Powering-down sections of logic requires the use of power isolation cells. These cells can be simple NAND gates; disabling one input turns off the output. The outputs of a powered-down section should not be allowed to float. Driving all inputs prevents glitches on the section's outputs; otherwise it is important to use a gating signal (through a NAND gate, for instance) to disable the outputs. For power planning with multiple-voltage domains, it is necessary to evaluate whether interfaces need level shifters, power isolation cells and/or gating cells.

The library must also have the right characterizations to enable multi-voltage designs. The necessary timing information is available from Scalable Polynomial Delay Models (SPDMs), preferably with voltage and temperature taken into account in timing characterizations. Non-Linear Delay Models (NLDMs) are also suitable. Multiple-NLDM libraries are characterized at different voltages. Synopsys Liberty models support these modeling methods.

Additionally, it is crucial to bear in mind that reducing the voltage decreases the logic's timing performance. As a result, it may be necessary to add parallel logic paths to meet performance goals. Identifying logic that can run at a lower frequency can help identify potential candidates for voltage scaling. Many times doubling a datapath logic section and reducing the frequency by half will result in a power savings (if it is practical to reduce the voltage), since power has a squaring function to voltage and only a linear function for frequency and switching. Having various logic sections run at different speeds may cause variations in signal arrival times, increasing skews, so it is important to make sure that the entire design will continue to meet timing specifications. Advanced tools help address this voltage effect on timing.

**SYNOPSYS®**

**700 East Middlefield Road, Mountain View, CA 94043 T 650 584 5000 [www.synopsys.com](http://www.synopsys.com)  
For product related training, call 1-800-793-3448 or visit the Web at [www.synopsys.com/services](http://www.synopsys.com/services)**

Synopsys, the Synopsys logo and Physical Compiler are registered trademarks and Power Compiler, Astro, Astro-Rail, Hercules and Jupiter-XT are trademarks of Synopsys, Inc. All other products or service names mentioned herein are trademarks of their respective holders and should be treated as such. All rights reserved. Printed in the U.S.A.

©2003 Synopsys, Inc. 8/03.TM WO