Advanced Routing in Changing Technology Landscape

Hardy Kwok-Shing Leung
Magma Design Automation, Inc.
2 Results Way
Cupertino, California 95014
ksleung@magma-da.com

ABSTRACT

As process technology continue to advance, the operating environment for routing tools has changed significantly. While the general concept of routing and techniques employed remain the same, the complexities and challenges that modern-day routers face are not well understood or addressed by the research community. In this paper, we will examine a handful of interesting nanometer effects that have significant impact on the behavior of routers, and discuss several opportunities in which routers can play a more important role in improving the manufacturability of nanometer designs.

Categories and Subject Descriptors

B.7.2 [**Integrated Circuits**]: Design Aids – *Placement and routing*; J.6 [**Computer Applications**]: Computer-Aided Engineering – *computer-aided design (CAD)*

General Terms

Algorithms, designs

Keywords

Physical design, routing, design rules, manufacturability

1. INTRODUCTION

Traditionally, routing was considered a solved problem in which the basic algorithmic issues were well-studied, and elegant solutions proposed. Routing can be further broken down into the main categories of global and detailed routing, with an optional step of track routing (or pin assignment) bridging the gap between the two. A decidedly incomplete sample of the literature includes [1, 2, 3, 4, 5, 6]. In the past, understandably, academia focused on solving high-level routing problems such as performance-driven topology generation, congestion reduction, routing-based congestion analysis for placement, and routing for crosstalk avoidance. Researchers shied away from getting into details of design rules, leaving the "grunt work" to the industry.

ISPD'03, April 6–9, 2003, Monterey, California, USA.

In this paper, we will examine several important details of routing that are important in nanometer design technology. In most academic work, often these details were ignored, "assumed" away for the sake of simplicity, or dismissed outright as implementation details. Yet, they have tremendous implication on the implementation of nanometer routing solutions, hence deserving a closer look.

2. ROUTING CHALLENGES

It is a misconception that process scaling only alters the parameters (such as pitches, spacing requirements, and coupling factors), without changing the fundamentals of routing as an optimization problem. The belief was that routers were merely required to be extended to handle new design rules. However, as we move into nanometer territory, some of the requirements such as spacing rules, reliability rules, and process antenna rules (most of which were old requirements invented generations ago) impose severe constraint on the routing algorithms, and even render certain well-established routing techniques and assumptions obsolete[1].

In this section, we will focus on the following three topics: (1) complex spacing rules, (2) transitional pitches, and (3) process antenna effects. Though important and representative, these topics are by no mean comprehensive. Other topics pertaining to nanometer routing such as timing- and noise-driven routing, routing in uncertainty, density rules, and 90-nanometer rules are arguably as important but will be skipped due to time and page limit.

2.1 Complex Spacing Rules

Routing of a single net in the presence of obstacles is a well-defined optimization problem: given a set of pin shapes, the objective is to construct a topology using wires and vias that minimizes the total wire-length, subject to constraints such as per-layer width and spacing design rules. Previously, the problem was further simplified by restricting the solution space to a rigid routing grid graph. Classical graph algorithms such as Dijkstra's shortest path algorithm can be employed to solve the problem.

Unfortunately, manufacturing processes are becoming increasingly complex, and non-trivial nanometer effects are more difficult to control and model. Much of the complexity trickles down to the domain of design implementation. Existing design rules were tightened, and new design rules created.

[1]For example, rectangular bloating based on pre-determined width and spacing requirement.

Table 1: Spacing Requirement as a Function of Both Width and Length (1u is one micron)

Spacing Required	Condition
0.19u	default width (0.21u)
0.23u	width $\geq$ 0.32u and length $\geq$ 0.95u
0.65u	width $\geq$ 10.05u, with a halo of 1u

One such example can be found in the evolution of spacing rules. Started as a constant per-layer requirement, spacing rules were extended to be width-dependent (given two polygons, the minimum spacing required is a function of the larger of the widths of the two polygons). Then, length dependency was added (the minimum spacing required is also a function of the *parallel run-length* of the two polygons). Finally, halo rules were added to refine the meaning of "width" (any thin wire within a certain vicinity of the fatwire is considered a fatwire for the purpose of spacing calculation).

Interestingly, while these special complication in spacing rules first appeared many technology generations ago, they had virtually no impact on place-and-route tools since only the extremely wide geometries (such as power rings and meshes) require extra spacing. For example, in 0.18-micron technology, wider spacing may only be needed for geometries wider than 10 microns; it was practically impossible for signal router to generate geometries that violate such rule, and all such special spacing requirement are static and can be determined prior to routing. As such, only minimal support for "fatwire" was needed.

It was the severe tightening of these rules that changed the game. Table 1 describes typical spacing requirements for the first few metal layers of a 0.13-micron process technology[2]. Given a default width of only 0.21 microns, a 0.32-micron "fatwire" can no longer be considered extremely fat.

The consequences of these changes:

- Fatwire will be created and complex spacing requirements will be triggered during signal routing. Hence they must be modeled correctly by place-and-route tools, rather than checked and corrected as an after thought.
- Notch filling can no longer be done as a post-processing step, since filling a notch may create fatwire, thereby triggering a larger spacing requirement. Notches must be either avoided, or modeled and filled on the fly.
- Spacing violation check must be done based on polygon analysis (even if rectangles are used as the underlying representation), and must be done actively as part of routing.

We shall further illustrate the problem of the tightened spacing rules with two examples. Figure 1 shows an example in which filling a notch caused unexpected fatwire spacing violation, even though prior to notch filling, no spacing violation between geometries of different nets was detected. This example illustrates that notches can not be filled casually as a post-processing step, but instead must be modeled during routing. Moreover, such modeling must be based on accurate haloed polygon analysis. In this example, one of the spacing violations (DRC#2) appeared not only far away from the location of the notch, but also at the tip of a "thin" halo rectangle.

[2]The parameters in this table, as well as those in the rest of the paper, are shown for illustrative purpose only. However, they are representative of actual 0.13-micron processes technologies.

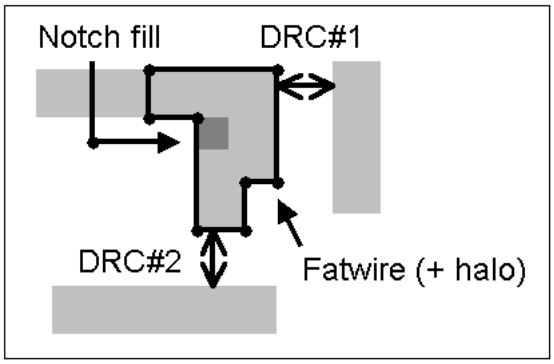

Figure 1: An example to illustrate that notch filling can cause unexpected spacing violation.

Figure 2 shows an example where a net $N1$ with default width requirement and *quadruple* via requirement[3] was routed in two different ways. Notches may be created between via overhangs of $N1$ if the vias are too close to each other, and a subsequent filling of the notch will trigger width-length-dependent spacing rule, causing a spacing violation against geometries of $N2$. This spacing violation may come as a total surprise to a detailed router that does not fully understand same-net spacing violations (such as ones based solely on Dijkstra's algorithm). Even if the router can detect such spacing violation after the fact – which already requires on-the-fly notch filling followed by polygon-based design rule checking – it will repeat the same mistake were it to reroute the same net.

2.2 Transitional Pitches

In modern process technologies more metal layers are available for routing. Signal routing layers are often divided into local layers (for example, the first four layers), intermediate layers (the next two), and global layers (the last two). Layers within the same group have similar routing pitches, design rules, and parasitics. For example, the intermediate layers may have a pitch that is 2X that of the local layers. An interesting problem arises at the highest layer of a given group (called a transitional layer in the following): the overhang of the "up via" can be significantly larger than that of the "down via" (see Figure 3). Normally, the routing pitch of a given layer must be defined to be no less than the "line-to-via" distance ($\frac{width}{2} + spacing + \frac{via}{2}$)[4]. In the case of a transitional layer, however, the line-to-up-via distance can be much larger than the line-to-down-via distance.

For competitive reason, achieving better routing density is

[3]This is not an uncommon electro-migration requirement for nets driven by high drive-strength cells.

[4]Routing with a pitch less than "line-to-via" distance may cause substantial congestion, since each via will block more than one track. Note also that routing pitch is not a concept specific only to grid-based routing; gridless routing also utilizes a routing pitch for alignment purpose to avoid needless fragmentation of routing space.

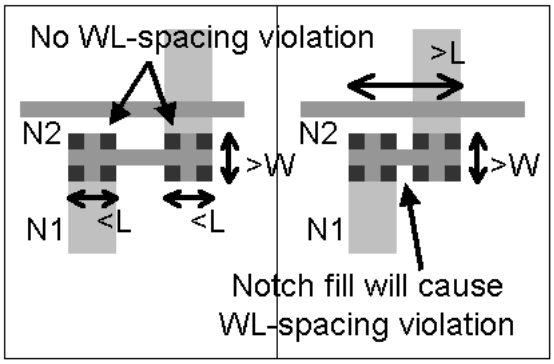

Figure 2: Another example to illustrate that notch filling can cause unexpected spacing violation. In this example, there is a wider spacing requirement if the width exceeds W and the parallel run-length exceeds L. Individual via overhang will not trigger the WL-spacing rule because both dimensions are shorter than L. However, if two via overhangs are too close, the resulting polygon after notch filling may be wide enough so that the dimensions exceed W and L respectively, causing a fatwire spacing violation.

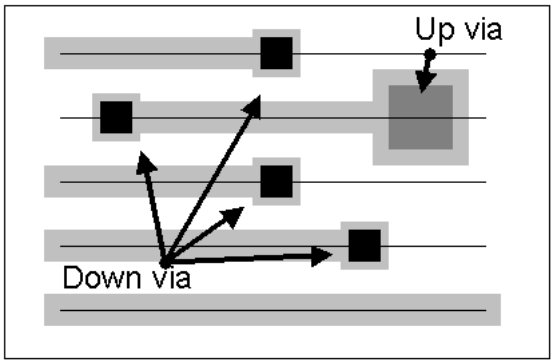

Figure 3: An example to illustrate a significant difference in the size of the up-via overhang and down-via overhang. With a pitch defined based on line-to-down-via distance, a down-via can be placed on a track without blocking the adjacent track. However, an up-via will block both adjacent tracks if it is centered on the routing grid.

almost always the overriding consideration. Pitches should be defined based on line-to-down-via distance unless the design is uncongested. With the use of the smaller pitch, the main concern is that each up-via will block multiple tracks and may cause local congestion problem if not modeled correctly. Fortunately, for transitional layers, down-vias usually significantly outnumber up-vias, making these surprise resource consumption an exception. Even so, all routers must comprehend and model transitional pitches correctly. Global router must account for the extra resource consumed by up-vias, track router needs to order tracks such that up-vias share tracks and are never sandwiched between other wires, and detailed router needs to shift up-vias to avoid blocking both adjacent tracks. These are important considerations that need to be taken seriously. Failure to do so results in poor routing quality even in uncongested designs.

2.3 Process Antenna Rules

Process antenna rules model the design requirement that the total charge accumulated on metal connected to a polysilicon gate during any stage of metalization cannot exceed a certain threshold, beyond which the excessive charge accumulation may permanently damage the gate. Antenna rules are generally of the form $\frac{(\text{wire area})_D}{(\text{gate area})_D} \leq (\texttt{ratio})_D$. The presence of diffusion area D, and hence a discharge path, is factored into the equation as a modifier to the wire, gate, or diffusion areas. For example, one can express the ratio as a monotonic, piecewise-linear function of the diffusion area, signifying higher tolerance when diffusion is present.

In 0.18-micron technology and above, process antenna effect is considered an "annoying but solved" problem. The ratio is typically large (of the order of 1000), and antenna violation is relatively rare. It is unlikely to see designs with large number of antenna violations, and the solution to the problem is relatively straight-forward. Both metal "jumpers" insertion and diode insertion (to force discharge path) worked well in practice.

Let `gate-strength(`g`, `L`)` be the maximum length of a wire of minimum width on layer L that can be directly connected to the gate g without causing an antenna violation. The larger the values of `gate-strength`, the easier it is to fix antenna violation. In 0.18-micron technology and above, `gate-strength` of 1000 microns and above is not uncommon, and fixing by post-processing suffices. In 0.13-micron and below, however, the average and worst-case `gate-strength`'s are substantially reduced, partially due to the use of cells with small gate areas (for example, extensive use of low-power cells) and a tightening of the antenna ratio. When the worst-case `gate-strength` is merely a handful of cellrows, antenna fixing becomes very challenging.

Besides from reduced `gate-strength`, a couple of other effects also makes the task of antenna fixing non-trivial:

- Transitional pitches – Adding jumpers may be infeasible because up-vias from transitional layers may cause spacing violations with adjacent tracks.
- Weakened diffusion effect – In the past, diffusion was modeled as an infinite discharge path, hence a panacea to any antenna violations. A common hierarchical antenna methodology was to diode-protect all macro pins and to waive all antenna checks through macro internals. In 0.13-micron and below, diodes have limited antenna-fixing capacity. Often process antenna violations were only flagged if the entire chip was flattened (or checked hierarchically).
- Extensive use of wide power meshes for IR-drop – if a routing topology runs along a wide power for a long distance, there may not be opportunity for jumpers to be created.

3. OPPORTUNITIES

Despite all the challenges and complexities, the changing nanometer landscape also brings about new opportunities in routing. In this section we shall examine a few of them.

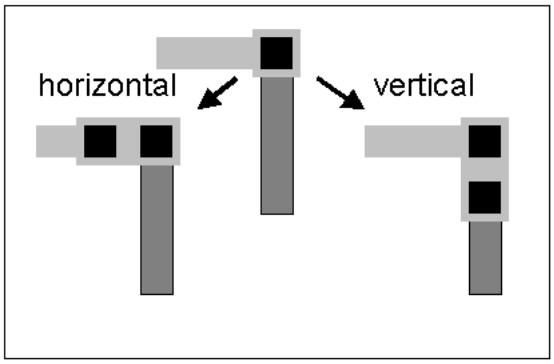

Figure 4: A redundant via can replace a single via in many different ways. Different redundant via orientation may have different routability impact on the two layers.

3.1 Redundant Vias

To improve yield and reliability, foundaries encourage the use of redundant vias (vias with multiple cuts) so that the random failure of single cuts can be tolerated. This is further emphasized in 0.13-micron technology and below. Our study shows that 70–80% of single vias of even the most congested designs can be made redundant without causing additional design rule violations. Roughly speaking, redundant vias can be added if *one* of the two layers has some room adjacent to the via, even if the other layer is fully occupied[5] (see Figure 4).

While redundant vias can be added as a post-processing step outside place-and-route tools, there is significant advantage in making redundant via addition an integral part of the routing flow. We estimate that proper tool support for redundant vias may help improve the success rate to 90% for average designs, while considering timing and design rules:

- Proper planning during global and track routing can significantly improve the success rate by avoiding local congestion or at-capacity routing on adjacent layers.
- Detailed router can detour wires to allow the addition of redundant vias in otherwise infeasible regions.
- Redundant vias can be added while considering other important design rules such as antenna rules for vias.
- Any negative impact on timing can be immediately flagged and repaired.

3.2 Wire Spreading, Widening, and Filling

To further improve yield, foundaries recommend that wires be spread and widened wherever room is available. Like the case of redundant via addition, proper planning during global and track routing can produce significantly better results – in terms of quality, timing, and ultimately yield – than do approaches based on post-place-and-route processing.

By the same token, metal filling needs to be part of the layout-implementation flow. Post-processing approaches are undesirable because they are either too conservative (failing difficult metal density requirements) or too aggressive (causing unexpected timing failure even for design that were supposedly closed in timing). A more viable approach is to consider timing and design-rule impact during metal filling.

[5]Unless if redundant vias require significantly bigger via overhangs, in which case the success rate will be much lower.

3.3 OPC- and PSM-Aware Routing

Very often, complex and tight design rules (such as end-of-rules, protrusion rules, and halo rules) were devised as "work-arounds" or guardbands to discourage or prohibit undesirable features, so that layout enhancement techniques such as Optical Proximity Correction (OPC) will work correctly. However, there will be a point (at 90- or 65-nanometer technology nodes) when the indirect approach will run out of steam. Imposing design rules may not always be appropriate either because certain practices help improve yield but are too expensive to be enforced as hard design rules. Hence, routing tools need to have more awareness and understanding of the intent of OPC and Phase Shift Mask (PSM) effects to produce layouts that are most friendly for layout enhancement techniques and yield improvement, possibly through a set of yield-enhancing "soft" design rules. In fact, many foundaries have emphasized the importance of supporting routers that are OPC- and PSM-aware for 65-nanometer technology. Router will need to comprehend design rules for PSM when phase shifting becomes widely applied to not only cell internals but also general signal routing[6].

4. CONCLUSIONS

In this paper we have discussed several nanometer effects that are important from the perspective of routing implementation for nanometer layouts. There are challenges, but the opportunities are even more exciting. It is hoped that this work sparks interest in improving and extending state-of-the-art routing to cope with the brave new physical world of nanometer technology.

5. REFERENCES

[1] S. Akers. A modification of lee's path connection algorithm. *IEEE Trans. on Computers*, EC(16):97–98, February 1967.

[2] S. Batterywala, N. Shenoy, W. Nicholls, and H. Zhou. Track assignment: A desirable intermediate step between global and detailed routing. In *Proc. Int'l Conf. on Computer-Aided Design*, pages 59–66, 2002.

[3] C.-C. Chang and J. Cong. Pseudo pin assignment with crosstakl noise control. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 20:598–611, March 2001.

[4] J. Cong, M. Xie, and Y. Zhang. An enhanced multilevel routing system. In *Proc. Int'l Conf. on Computer-Aided Design*, pages 51–58, November 2002.

[5] D. Hightower. A solution to line routing problems on the continuous plane. In *Proc. 6th Design Automation Workshop*, pages 1–24, 1969.

[6] J. Soukup. Fast maze router. In *Proc. 15th Design Automation Conf.*, pages 100–102, 1978.

[6]Due to the presence of predominant preferred routing directions on metal layers, it might be possible to guarantee phase-conflict-free layout by enforcing extra width and spacing requirements in the preferred directions and hence ensuring that no wrong-way features need to be phase-shifted.