

---

## **CHAPTER 11**

### **Managing Differential Signal Placement**

#### **DISCUSSION AND CONCLUSIONS**

---

##### **11.0 GENERAL THOUGHTS**

This chapter brings our journey to a close. Research can be viewed as analogous to working a gold mine. One begins by examining small surface nuggets to see if they are worth pursuing. Then when a potential vein of gold is identified, efforts are expended to exploit it and bring forth all of the gold one can. While following the vein, branches occur. Selecting the true course is often difficult. Many times promising avenues must be bypassed in order to conclude the original work. So too is it with research. Many solid and important findings were uncovered. But many possible avenues of other research were identified but could not be pursued. This chapter serves three purposes. First, it serves to provide closure to the investigation of adjacent placement management of differential signal pairs. Where chapter ten provided the details of the prototype implementations, this one postulates the optimal system based on the knowledge gained from the research. Next, it highlights what I believe to be the significant discoveries that have emerged during this investigation. Finally, it lists those other veins of gold, or alternative areas of potential research, that had to be passed up along the way.

## 11.1 HOW THINGS OUGHT TO BE<sup>2</sup>

The research began with a clear goal: provide a sound theoretical solution to the adjacent placement problem of differential signal pairs, that was implementable from an engineering perspective. It was motivated by both personal desire and real world necessity. A thorough literature search was conducted. Although it did not provide specific help, it laid the foundation for further study. Original problem constraints on time (inability to construct a new router from scratch), availability of a commercial tool (no known tool handles the problem of placing differential signals in adjacent tracks), and an actual system requiring differential routing (F-RISC/G), interacted to produce an excellent solution.

The theoretical algorithms associated with each of the three problem spaces are sound. Upper bounds on their time complexity are excellent. The standard cell and block routing solutions are equivalent to the core router plus a linear time bifurcation factor. The MCM algorithm is equal to the core router plus a quadratic factor associated with the polygon chaining. This factor occurs as part of the CIF bifurcation. The solutions presented for all regimes are optimal. For chip and block level routing, the differential pair is contained within the blockout of the fat wire. At most, that encompasses two tracks from source to terminus. Any required inversion is handled in that space. The same is true for the MCM router. Alternative solutions using a two track block out scheme will ultimately generate a three track by two track exclusion area every time an inversion cross over is introduced. If a three track grid technique is utilized, one out of every three tracks is completely wasted.

The solutions also effectively deal with the software testing issue. CAD software for the VLSI arena ranks just behind life critical software in importance of verification and

---

<sup>2</sup> Not taken from the title of Rush Limbaugh's Book.



testing. The huge capital investments to manufacture chip sets cannot be put in jeopardy by incompletely tested software. The general solution to bifurcation suffers from an inability to exhaustively test it. At best, a path cover test plan could be developed[Beiz86]. Partitioning the problem, using feature vectors and net categories, permits each sub-solution to be thoroughly tested. This factor is crucial, regardless of the solution developed.

Implementation specifics for each problem space were tailored to cope with core router deviations. Had a completely controllable core router been available, the theory could have been implemented in a straight forward manner. Having uncovered the underlying theory, and gaining experience from overcoming various implementation obstacles, I believe an optimal differential router construct has emerged for each region. In the next three sections, the “built from the ground up” solution will be postulated using all of the knowledge gained.

#### **11.1.1 DIFFERENTIAL STANDARD CELL ROUTER**

The system architecture from chapter three is the starting point. To facilitate bifurcation and make the solution more general, an intermediate file format standard would be formulated. It could be modeled on the VLSI Tool's *CP* file but would carry additional information. To deal with the two classes of uncuttable configurations, the over constrained via, and the single segment net connecting two standard cell rows, an additional guidance mechanism would be built into the router. Since the net recognition routines run in linear time, little penalty would accrue by performing the recognition function at the time of routing. If a routed net is not recognized, it indicates one of the two offending classes. Over constrained vias are easily identified by using the graph theory concept of node degree. Any via with a port segment degree greater than one is over constrained. The single segment net can be identified trivially. In either case, a guided local

rip up and re-route should occur. The goal would be to re-connect the net without an over constrained via. This can be done through the introduction of a jog. The penalty for the jog is three additional vias. However, the cost in vias can be eliminated by automated instantiation of generic standard cells. An explanation of the concept is provided in the future research section.

With guided local re-route, the fat wire output file will be one hundred percent bifurcatable without any human intervention. This is not to say that the router should not permit human interaction. On the contrary, one of the important learning points of the research is that all future router design must allow interaction at all stages. Important nets, either on the critical path or clock nets, may need to be routed by hand. Final cell placement may need to be altered due to congestion[Rice82]. Only the designer's full understanding of the system allows meaningful changes to be made at that point.

### **11.1.2 DIFFERENTIAL BLOCK ROUTING**

The greatest obstacle to implementation was encountered in the block routing regime. The primary cause was the deviation of the core router from the Manhattan stipulations. As with the standard cell solution, if the router can be appropriately monitored and directed, the solution will directly follow from the theory of net recognition and bifurcation. However, in this space, major power distribution nets must also be routed.

In the near term, where full Manhattan compliance might place too much of a burden on the router, with the side effect of a less than optimal fat wire routing solution, a hybrid approach is most reasonable. Guided re-route similar to that used for standard cell areas can be attempted. Where possible, uncuttable configurations are modified. Next, the normal recognizers examine the nets and everything in the original taxonomy is cut. Finally, the remaining nets are bifurcated using the polarity push forward technique. This



implies that the segments must first be chained to assess turn direction. Using this information, bend following bifurcation operations, along with zero cut vias push any inversion action out to the pad. There, the correct pad orientation is identified and a revised instantiation requested if the placed pad is not of the desired polarity.

With these modifications, reasonable solutions can be expected. There is a slight increase in time complexity due to net chaining. Even though the operation is quadratic at worst, the number of segments is typically quite small. Thus, the usual time factor is only marginally worse than the linear factor for recognition and direct bifurcation. As multi-layered metalization becomes the norm, a much more elegant technique is envisioned. By handling power distribution nets in two restricted layers, the traditional two routing layers will be available for all signal nets. At that time, with a compliant router, the complete standard cell approach should be adopted. This would again produce a one hundred percent fully automated solution. This domain is clearly the one where optimal theory cannot overcome implementation obstacles at this juncture.

### 11.1.3 DIFFERENTIAL MCM ROUTING

The fat wire routing concept extends perfectly into the MCM regime. However, the net recognition approach is not the best for this problem space. Initial attempts to directly apply the theory from standard cells to MCMs encountered enormous difficulties. Port connectors appeared on all sides instead of just two. Reviewing the problem definition and comparing it with the literature revealed that the MCM arena most closely resembled the switch box routing problem. The constraints on horizontal connector polarity in the block routing space no longer existed. A re-thinking of the complete problem was in order. The final result pointed to conducting the bifurcation phase in the *CIF* domain.

The MCM solution should directly apply the fat wire concept to actual routing. Comparative statistics from the single chip, five chip and twenty five chip system routes show that as congestion rises, alternative approaches such as the three track block out cannot complete a route. Half of the inversion problem can be readily satisfied at the wafer I/O pads. For nets that only connect chip to chip, inversions should be handled with callable *CIF* macros.

The advantages of this technique are numerous. Package independence can be maintained until very late in the design cycle. Chip pad design and placement variations can readily be accommodated. Should the mounting technique go from TAB to C4, a normal wafer design would be rendered useless, unless chip designers reworked their logic. With the *CIF* approach, the entire substrate solution can be reflected to handle flip chip mounting.

## 11.2 KEY LEARNING POINTS

Using the gold miner's analogy, many small nuggets were discovered over the span of this research. What are believed to be the most important ones are itemized in bullet form. Since each has been completely described in the dissertation, only a summary note is provided where necessary.

- There is a fundamental Taxonomy tree from which all readily bifurcatable nets can be generated.
- FSM theory provided both recognizers and assurance of bifurcatability.
- Empirically derived feature vectors were proven correct through the link to regular expressions.
- Linear time net recognition and bifurcation algorithms were developed.



- An elegant state variable XOR formula for computing fat wire via cuts, so as to correctly re-apply inversions in linear time was constructed.
- CIF splitting routines were built that correctly compensated for corner cuts through polygon OR'ing.
- The extendibility of the concept to  $N$ -layers of metalization has been explored and defined.
- There must be domain specific implementation of the fat wire concept.
- A thread of continuity exists for iterative MCM design. Using wiring pitch as the integrating factor, routability can be assessed, EM analysis conducted, chip driver capability calculated, and package thermal dissipation characteristics matched.
- Variable wiring pitch is available as a side effect of the fat wire concept.
- Techniques for retaining the engineer in the almost totally automated design loop, and capitalizing on his expertise must be reviewed.
- MCM package independent solutions are made available by bifurcating *CIF* files.

To the casual listener, many of these ideas may be viewed as either obvious or straightforward at this point in the dissertation. But in reality, they only evolved over many months of struggling with the problem. An old quote that says, "Everything is clear, once it has been explained," is very apropos to the problem space in which I worked.<sup>3</sup>

### 11.3 FUTURE RESEARCH

Many interesting questions arose along the journey. Some of these were on the primary research path and have been answered. Others were more tangential and had to be bypassed. The next several sections outline avenues of potential research.

---

<sup>3</sup> Source of Quote Unknown.

### 11.3.1 AUTOMATED INSTANTIATION OF GENERIC STANDARD CELLS

The only shortcoming of the net recognition and bifurcation approach as applied to standard cells, was the inability to deal with over constrained vias, and the degenerate single segment net. Guided local re-route can overcome this but there is an associated cost in vias. Ideally, the solution should not introduce metric penalties of this nature. Examination of the GaAs standard cell library for the F-RISC/G test bed, uncovered an intriguing possibility[Nah91].

When a given standard cell is instantiated in the layout, a fixed polarity for the differential port pair exists. Opposite configurations coming together at a via produce the over-constrained situation. One possibility is to generate all variations for each cell in the library. This is unfeasible due to enormous storage requirements. The combinatorics involved in generating all variations of a cell are given by equation 11-1. For a cell with two inputs and a single output, there would be 64 variations.

$$Cell\_X\_Variations = 2^{\left(\sum_{Cell\_X\_Fat\_Ports}\right)} \quad (11-1)$$

Instead, a different view is taken. Create a generic version of each standard cell by logic function and power level. Cells should be crafted so that the bulk of the logic is contained in a given area, and a “connection zone” provided in another, Fig. 11.1.




FIG. 11.1 GENERIC CELL CONFIGURATION



The fat wire port locations are identified and provided to the layout and core routing tools. During the bifurcation process, when connections are made to the cell ports, the differential wires are extended into the interior of the cell so as to establish the connection in a manner that agrees with the required polarity. Fig. 11.2 shows the fat wire net polarities approaching the over-constrained via. Given a limited cell library, and our research implementation constraints, the circuit cannot be bifurcated.




FIG. 11.2 POLARITY APPROACHES OVER CONSTRAINED VIA

Fig. 11.3 shows the generic instantiation approach with connections completed to Cell X. As the first connection from a fat wire via, the degree of freedom available permits either polarity. Once this one has been chosen, there is now only one way to extend the wires from the via to the ports on Cell Y. With the connection zone inside the cell, the

wires can be extended into the cell, and appropriate connections made as if another via existed. This is shown in Fig. 11.4




FIG. 11.3 CONNECTION MADE TO CELL X

This technique is analogous to the use of objects in a software environment. It greatly simplifies standard cell libraries. Only a single generic version of each cell is necessary. It solves the over-constrained via problem without the introduction of any additional vias. A penalty in standard cell area is incurred by requiring a connection zone in each cell. However, a review of the test bed library indicates that creative crafting and design of the cells can minimize this impact since many cells tend to follow this scheme in a general sense.






FIG. 11.4 CONNECTIONS COMPLETED TO CELL Y

### 11.3.2 CHIP PAD PLACEMENT THROUGH ITERATIVE MCM DESIGN

As MCM design and routing occurred, the necessity to integrate chip pad placement decision making with MCM system routing requirements became obvious. Chip designers working in isolation from one another and the MCM layout, tend to place pads based on their proximity to the interior logic that generates the signal. To meet chip timing constraints, they may have no other alternative. Yet to meet the demands of ever increasing system clock rates, the MCM routing structure must be considered. Using high speed line probe routers, such as the one available to the F-RISC/G team, iterative design cycles can be set up. With preliminary chip pad locations and a system net list, the MCM can be routed. An examination of the critical nets along with the average net length can

serve to focus concentration. Chip pad placement for these signals should be examined. If a revised signal location on a given chip facilitates MCM routing, or is on a critical timing path that would be shortened, the chip designer should be provided that information. Using proven engineering tradeoff analysis techniques, a decision can be reached as to whether or not to re-locate the pad. The earlier in the design cycle this information can be generated, the more useful and easier to apply it will be.

Power and ground pad locations are also important. By necessity a minimum number will be required, but their location can significantly affect the routability of the MCM. Study has shown, that whenever possible, these pad locations should be symmetrical on the chip with respect to the axis of interest. East/West pads should lie on identical y coordinates, while North/South pads should lie on matching x coordinates. By arranging the pads in this fashion, only a single track experiences a major blockage. Random placement tends to block multiple tracks, seriously impeding the routability of the core area. Tied directly to these issues is the concept of the macro-via introduced in chapter eight. Further study of this concept may yield very interesting results in noise immunity.

### **11.3.3 MCM DESIGN METHODOLOGY**

The final area where future research has a long way to go is in the area of MCM design methodology. Most of the design issues have been recognized. They include: (1) packaging techniques, (2) thermal dissipation, (3) electro-magnetic analysis of substrate wiring, (4) chip driver capability, and (5) routability[Comp93]. Although loose linkages have been known to exist between each issue space, there has not been a unifying thread tying them together. Exploring variable differential wire spacing revealed a connection. Substrate wiring pitch can be viewed as the integrating factor linking the various problem



spaces. Using it in conjunction with a high speed router, an MCM design station can be constructed.

Current tools for thermal mapping, routing, and EM analysis could all be linked through this factor. The feedback scheme presented in Fig. 9.13 can serve as the template.

#### 11.3.4 BUS COLLAPSING

By employing the fat wire concept to differential pairs, the basic routing problem size is reduced by a factor of two. If this approach is taken to the next logical step, it can be applied to major bus trunks. If busses are envisioned as the next step up the hierarchical ladder in the fat wire regime, it may be possible to achieve a greater than two factor of problem space reduction.

Applying this to equation 6-3, the resulting complexity would be an  $m$ th root instead of the square root, where  $m$  reflects the bus compression ratio. Certainly, the fat wire busses would have to stop short of their destination in order to allow for matchup fan out, but the potential for research is extremely enticing.

### 11.4 CONCLUSION

This undertaking has been fascinating and extremely rewarding. The fat wire routing concept, which originated as a work around solution to an engineering shortcoming, has come into its own. It has been demonstrated to be optimal in both time complexity and space utilized for the routing result. Beginning with net recognition through the use of finite state machine theory, a linear time bifurcation algorithm was developed.

Prototype systems were constructed to validate the theory. The F-RISC/G system provided an ideal test bed for comparative experiments. The working system has routed the RPI Testchip, the Instruction Decoder and the DataPath chip. Of these, the Testchip

has been fabricated and tested. The concepts have been extended into the MCM domain and produced an excellent solution that permits great flexibility. A full twenty-five chip system solution, out to the second level of cache, has been routed as a proof of concept. Comparative statistics between the fat wire MCM approach and the three track block out technique had to be given up when the alternative method failed at the five chip test.

Theoretical analysis has been conducted to see how well the theory holds as multi-layer metalization becomes available. The study in chapter seven shows that it is extendible by a single layer to a limited degree, but flourishes if layers are added in pairs. This is a clear sign of a robust solution.

Finally, there has been an underlying knowledge mapping function operating in the background throughout. Whenever the fat wire architecture is mentioned, there is always the generic block labeled -- "**core router.**" This is understood to be the best available single ended router. Since the quality of the fat wire solution is guaranteed to have an optimality equivalent to the core router, employing the soundest and most creative techniques to the underlying single ended routing problem will produce an optimal result. This effect can be viewed graphically in Fig. 11.5, where all of the knowledge gained in the single ended routing domain is immediately projected into the differential routing space.



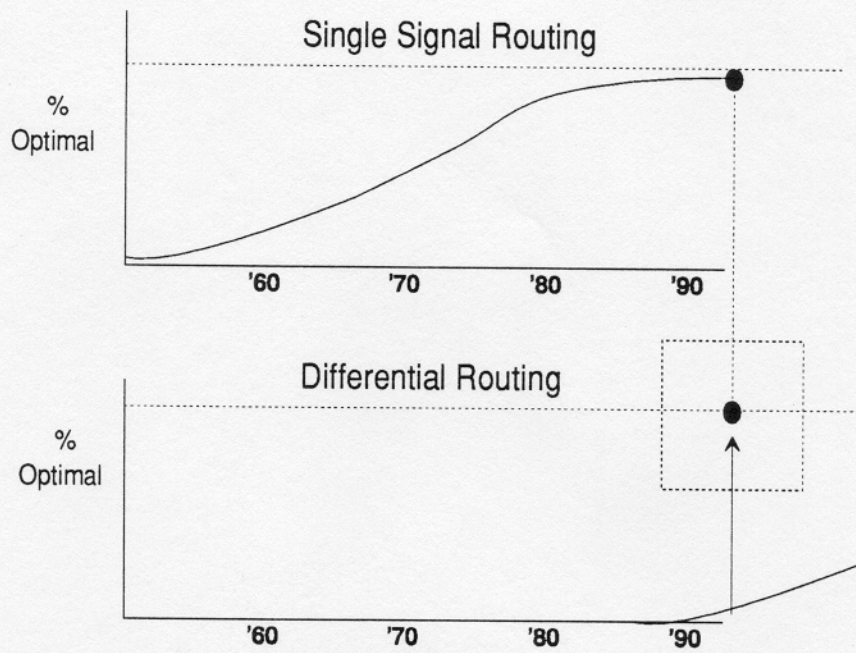


FIG. 11.5 ROUTER KNOWLEDGE MAPPING

It is my hope that the solution provided will enhance the basic body of scientific knowledge, and that it will stand up well against the tests of time and the dawning of the differential routing era.