

developing a 1ns RISC processor. They combined to create both an environment conducive to meaningful research, and a sense of urgency that pushes one to the limits.

This dissertation attempts to present the findings of the research in a chronological manner. Feature vectors, finite state machine theory, and actual implementation steps were not isolated and independent as the ordering of the chapters might indicate. Instead, the organization was selected to promote understanding, and present dense theoretical ideas in digestible portions.

Chapters one and two set the stage by providing the motivation, goals and historical background of the research. Chapter three presents the novel router architecture that is responsible for effectively managing the adjacent placement of differential signal pairs. Key to its success is the ability to bifurcate fat wire nets in linear time. Chapters four and five develop the feature vectors and finites state machine theory necessary for effective bifurcation presented in chapter six. Chapter seven examines how extendible the theory is when applied to technologies that permit multiple layers of metalization. Chapter eight extends the approach to the MCM domain. Chapter nine looks at proper differential pair spacing, and chapter ten reviews system performance. Chapter eleven concludes the work, highlighting the findings and pointing out potential areas of future research.

Each chapter begins with a General Thoughts section. This acts as a pseudo abstract for the chapter, focusing the reader's attention to the key ideas to be presented. Chapters conclude with a Summary section, designed to recap the findings of the chapter.

ABSTRACT

The interaction of reduced device voltage levels and ever increasing clock rates has had the effect of lowering noise margins. The use of current mode logic (CML) in conjunction with full differential signal routing can delay these effects and thereby preserve the noise margins of future, high speed digital systems. Providing full differential routing, placing complimentary wires of a pair in adjacent tracks throughout their run from source to terminus, is an interesting problem. Underlying routing algorithms are NP-complete as a function of the number of nets. Doubling the number of nets to accomplish differential routing significantly degrades the quality of solution. This thesis proposes an alternative solution to the problem which encompasses both the chip and multi-chip module (MCM) design regimes.

Employing the concept of logical entity nets, where each differential pair is viewed as a single logical signal, the fundamental routing algorithm complexity remains equivalent to the single ended problem. After routing the logical nets using oversized design rules, the resulting nets are bifurcated. This yields a full differential solution. A net-to-finite state machine (FSM) mapping process is put forth. Using FSMs as recognizers, all bifurcatable net topologies can be identified. To facilitate implementation, a feature vector approach to net recognition was developed. It is demonstrated to be equivalent to the FSM process through the use of regular expressions. The result is a linear time algorithm for net bifurcation. The concept is proven to be optimal with respect to time complexity and equal to or better than any known algorithm with regards to layout space.

The solution is robust and extendible to multi-layer metalization. It has also been adapted for use in the MCM arena. An ARPA sponsored research project to produce a 1ns RISC processor has provided an ideal test bed. Experimental results confirm the theoretical predictions.

CHAPTER 1

Managing Differential Signal Placement

INTRODUCTION

1.0 GENERAL THOUGHTS

As a Chinese philosopher once said, "Even the longest journey begins with the first step." This saying holds special symbolism for this endeavor. As the introduction to the research undertaking, this is the first step of our journey. Not only is this dissertation an attempt to capture the critical information accrued during a long research journey, but its focus, in a manner of speaking, deals with the art of making journeys shorter by laying out the optimal path prior to embarking.

A good introduction serves several key purposes. It should specify the goal of the undertaking, and present the motivation for the research. It should describe how the solution will advance the body of scientific knowledge. Finally, it should provide a road map to the presentation and discussion so that the information presented unfolds in a smooth and understandable fashion. In an effort to follow this scheme we will begin by laying out the goal of the research.

1.1 RESEARCH GOAL

Designers of printed circuit boards (PCB), large scale integrated circuits (LSI), and very large scale integrated circuits (VLSI), have been grappling with the wire routing problem for the last three decades. With the exception of handling a few special bus layouts, the concern has always been on single ended routing between destinations. Comparisons among solutions has focused on the optimality metrics of via counts, total interconnect length and overall chip area. The goal of my research is to explore the concept of differential signal routing, and facilitate the layout of complimentary signals in adjacent paths between their source and destinations. In doing so it is incumbent upon me to determine the best solution to the problem, and understand how that solution behaves in two major design regimes of interest: (1) chip level routing, and (2) Multi-chip Module routing. As an outcome of this research, I will present solutions to this problem which encompass standard cell routing, block routing and Multi-Chip Module (MCM) routing.

It is important to provide a clear definition of terms, especially in an area of new research.

Differential Signal Routing

Placing both signals of a differential pair (a signal and its complement) in adjacent routing tracks throughout their course from source to destination, swapping their ordering if necessary to achieve logic inversions, as specified by the net list. Ideally, the routes employed should be as close to optimal as to what is achieved by current single ended router technology.

To provide a common frame of reference, additional terminology and acronyms used throughout the dissertation can be found at Appendix A.

1.2 IMPETUS FOR THE RESEARCH

My research gains its impetus from two sources. First, the trends in technology indicate that differential routing will grow in importance for high speed digital applications in the future. Second, an actual project at RPI, the F-RISC/G, requires differential routing in order to achieve its operational specifications. This combination of a search for new knowledge and an urgency derived from an actual project, proved an ideal setting for stimulating such research.

1.2.1. TRENDS IN TECHNOLOGY

Advances in computers and the underlying microelectronics over the past three decades are staggering. Technology trend lines are clear. They point toward higher clock rates, reduced power consumption, and increased functionality. Intimately linked to these trends are device geometry, device speed, current levels, power consumption, and the system package. Over the past thirty years, minimum feature sizes of devices have decreased by eleven percent per year. Chip sizes have increased nineteen percent per year, and packing efficiency has improved by a factor of two every ten years. Together these factors have produced a one hundred fold increase in devices per chip per decade[Bako90].

Through the interactions of these factors, logic level voltage swings had the potential of being scaled with other device dimensions as a factor of $1/S$, where S is the scale factor of the principal device dimensions. Until recently, logic voltage levels have not been reduced with feature sizes. Reasons for this included retaining TTL compatibility, avoiding the need for multiple power supplies, and in many cases, reducing the device gate delay by a factor of S . On the down side, retaining this voltage level resulted in increased electric fields, since device geometries were shrinking. This in turn affected reliability issues, such as oxide breakdown, and hot electron injection into gate oxide.

In the past few years, the pressure to reduce logic voltage levels has grown. One of the key factors is the advent of the laptop computer. Forced to operate on batteries to be useful, system power consumption had to be minimized. The May, 1992, issue of Spectrum has articles dealing with lower powered CMOS circuits, and 3.3V TTL operation[Spec92]. For high speed devices such as HBTs, the key concern is overall power dissipation. To keep chip temperatures within operational limits, without employing significant artificial cooling mechanisms, dynamic power dissipation must be reduced.

Noise margins are measured as the difference between the highest voltage clearly recognized as a low logic level, and the lowest voltage identified as a high logic level. By reducing the overall swing, the noise margin interval is also reduced. The impact of this is to reduce the noise immunity of the overall system. Differential logic, such as current mode logic (CML), or source coupled FET logic (SCFL), employing current steering techniques, can help maintain noise immunity at low voltage swings. CML generates very little switching noise. By maintaining a nearly constant current flow through one of the two or more paths through the current switches, and redirecting this current to achieve logic changes, current surges associated with conventional logic are avoided. Use of differential logic, such as cascode voltage switch logic (CVSL), can help improve noise immunity even in rail-to-rail switching CMOS circuits. Employed in concert with differential signal routing they can act to preserve, or at least delay, the decay in system noise immunity. IBM's ES/9000 system already employs a combination of ECO and multifunction DCS cascode circuits, but does not insist on adjacent track routing[Bari92]. Digital Equipment's new macropipelined VAX uses differential cascode logic on critical circuits, but relies on handcrafting to handle that portion of the layout[Bade92].

Given the trends in technology, and the reduction in noise margins, it seems certain that differential signal routing will grow in importance in the coming years. From this perspective, it is certainly a worthwhile area of research.

1.2.2 F-RISC/G

In the real world, resources are always limited. The current body of scientific knowledge can be viewed as a sphere. The surface of this sphere represents the limit of current knowledge, and the space beyond, the problems that must be solved to advance the total body of knowledge. With the ever increasing number of open questions, researchers must discriminate between them when allocating limited resources in order to find their solutions. Although trend lines indicate that differential routing may be needed in the future, energizing research in such a domain is difficult unless an immediate need for the capability exists.

Ongoing ARPA sponsored research at Rensselaer Polytechnic Institute, aimed at producing a Spartan RISC architecture with a target cycle time of 1ns or less, has provided that real world motivational vehicle for solving the differential routing problem. To achieve the specified cycle time, GaAlAs/GaAs heterojunction bipolar transistor (HBT) technology was chosen. Operating in a current mode logic configuration, unloaded gate delays of approximately 20ps are possible[Greu91]. To retain ECL supply voltage compatibility of -5.2 volts with a three level CML configuration, resistor current sourcing was selected, Fig. 1.1. This fact in conjunction with the small voltage swings (approximately 250mV) significantly reduced the noise margin of the system. To retain a minimum acceptable level of noise immunity, the effects of both current mode logic and adjacent signal routing were determined to be necessary.

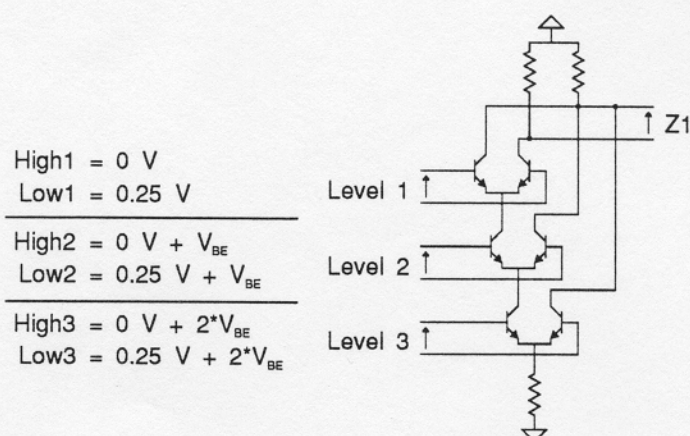


FIG. 1.1 BASIC THREE LEVEL CURRENT TREE

1.3 BENEFITS OF DIFFERENTIAL LOGIC

By routing a signal and its complement in adjacent tracks throughout their run from source to terminus, the ability to reject common mode noise is enhanced. If a signal pair routed in this fashion encounters noise during transit, then both signals of the pair will be disturbed by a similar amount. The destination sense amplifiers will detect the same level difference that existed at the source, discounting transmission attenuation.

Besides noise immunity, differential signal routing provides two other important benefits. It gives the architect an added degree of design flexibility. Inverters are available to him for free. There is no cost in logic gate count, time delay or power consumption. Wherever an inversion is required, swapping the signal pair ordering along the route accomplishes it. This in turn results in a simplification of the standard cell gate library, since NAND, NOR and OR gates are actually just copies of AND gates, with instructions to the router to provide either input and/or output inversions.

Differential routing also eliminates the need to provide reference voltages to each of the standard cells. In the case of three level CML it obviates the need to distribute three reference voltages throughout the chip. Although this can be accomplished by routing additional "rails" through each standard cell row, considerable chip area can be consumed.

1.4 COST OF DIFFERENTIAL ROUTING

It has long been recognized that differential routing will result in increases in chip dimensions, interconnect length and the capacitance associated with interconnect growth, over single ended logic designs of an equivalent complexity[Bako90].

The decision to employ differential signal routing immediately doubles the number of signal nets in a design, over a single ended version. This doubling of the nets produces a delta-x and delta-y increase in chip dimension. A full analysis of the increases and how they can be estimated is carried out in chapter nine. Using the increases in the x and y dimensions, an estimate for the revised chip area can be calculated. From these results, the percent increase in chip size can be computed.

As the chip size grows, the average signal length for an arbitrary interconnection also increases. This growth in interconnect length carries with it an increase in capacitive loading. The ripple effect is a slowing of signal transmissions for a given drive capability.

1.5 DECISION

For the F-RISC/G project, the benefits of differential routing had to be weighed against the long recognized costs that included: (1) increases in chip dimensions, (2) interconnect length, and (3) the capacitance associated with interconnect growth, over single ended logic designs of an equivalent complexity[Bako90]. Having analyzed these actual alternatives, it was determined that the most crucial factor facing the design was noise immunity. Given this prioritization of issues the decision was made that differential routing was a necessity.

1.6 RESEARCH CONSTRAINTS

Several constraints existed when the research was begun. They served to fashion and guide the research direction. First, no commercial tools possessed the capability of handling the differential routing task as it is defined in this document. Second, modern day routers, with sophisticated heuristics, take many man-years to construct and verify. Resources for such an endeavor were not available. As a result, any attempt to construct a new router from scratch was out of the question. Finally, source code was not available for the commercial design tools, including the single ended router, to which the F-RISC/G design group had access. These constraints acted to set the stage for the research that ensued.

Initially, one might believe that constraints of this nature would limit the solution space. Instead, they forced a complete rethinking of the entire problem. This, in turn, led to a solution that might have otherwise eluded discovery. In hindsight, the constraints actually helped in guiding the research to its solid conclusions.

1.7 LITERATURE SEARCH

As with any journey, it is best to check one's map prior to embarking. Likewise, with solid research, the scientific literature must be reviewed. Seldom, if ever in today's complex world does a researcher make some form of discovery, without the benefit of knowledge gained by others who have previously studied the same or similar problems. Most steps forward are incremental in nature. Typically, they are rooted in the findings of others, which have been shown to be true over time.

With this underlying philosophy, I conducted a thorough and comprehensive search of the professional literature pertaining to the problem. I began by reviewing some of the summary works on the topic of routing. From there, I traced the key evolutionary

steps, from Lee's maze running algorithm and Hightower's line probe approach to the more sophisticated heuristics used by today's intelligent routers. On-line library search facilities were of immense help. Hundreds of journals, spanning many years could be reviewed for subject matter and relevancy to the problem. Also, in an area that has not been well studied, as is the case of differential routing, this is nearly the only way to insure that someone has not already made a contribution in the area or provided a solution to the problem.




FIG. 1.2 INFORMATION CUBE HIGHLIGHTING KNOWLEDGE VOID

While scanning the literature for terms associated with differential routing, over six thousand articles dealing with routing and/or LSI and/or VLSI circuits were scanned. Of these, none mentioned anything about differential routing or differential routing terms. Fig. 1.2 shows this graphically. With the three principal axis labeled as data structure, algorithm, and application, there is an obvious void with respect to differential routing. An INSPEC/COMPENDEX review conducted on April 23, 1993, using the terms, Printed Circuit Board(PCB), LSI, VLSI and Routing returned over two thousand entries. With the additional qualifier of Differential, there were no returns.

This result has both an up and down side to it. On the down side, it meant that no one had performed any research on which I could build. However, on the positive side, it meant that any meaningful discoveries that I might happen upon would be both original and hopefully of use to others who might have to travel this route in the future.

1.8 ROAD MAP TO RESEARCH

Although it is possible to review the table of contents to see what is covered, I believe it is important to layout the framework of the dissertation. This not only forces me to structure the undertaking, but hopefully present it in a way that assists in comprehension of both the complexities of the problem and the key insights of the solution.

After this introduction, I intend to review the traditional techniques used to deal with the single ended routing problem. Standard cell routing, block routing, and Multi-chip module (MCM) routing will all be reviewed. Similarities and distinctions between the regimes will be highlighted. This development, although it does not specifically address differential routing, provides the starting point for the remainder of the research. It should also be noted that it is for these same design spaces that I will propose a differential routing solution.

A novel architecture designed to handle the differential routing problem is then presented. The architecture makes a conceptual leap, and in so doing guarantees a solution whose optimality is equivalent to that of the core router. Crucial to its success is a bifurcation process that operates in linear time.

Chapter 4 traces the evolution of preliminary feature vectors and the development of a net taxonomy. Chapter 5 covers finite state machine mapping which provides the theoretical underpinnings that prove the correctness of the feature vector approach.

Chapter 6 outlines the linear time bifurcation algorithm. Having developed this algorithm, the router architecture is proven optimal in both space and time. It is this algorithm that makes the entire approach viable.

Since multi-layer metalization is rapidly becoming a reality, Chapter 7 looks at extending the solutions presented in Chapter 3, into n -layers. With a generalization to an arbitrary number of layers, the proposed solution is shown to be extendible as technology advances.

Chapter 8 investigates how the "fat wire" concept can best be extended into the MCM arena. The implementation in that problem space takes a very different approach from the feature vector methodology utilized in the chip domain. The benefits of the alternative implementation are significant.

Chapter 9 explores the electro-magnetic effects of differential pair spacing. One of the advantages of the architecture is the ability to vary both inter-pair spacing and intra-pair spacing independently of one another. This permits the designers to select an optimal tradeoff point between electro-magnetic effects and chip area increase. This chapter describes the methodology used in developing a family of curves that can be utilized to select appropriate inter-pair and intra-pair spacing parameters. Both chip and MCM analyses are conducted.

Chapter 10 provides the results of applying the architecture to the three regimes mentioned. Results from routes of sub-components and full chips taken from the FRISC/G processor chip set are provided. Multiple MCM test vehicles are presented. They extend from a single chip module to the full twenty-five chip F-RISC/G system, modified to reduce the cycles per instruction (CPI) figure of the processor.

A summary of the research findings is set forth in Chapter 11. The chapter concludes by postulating areas of future study aimed at complete CAD support for differential systems. These opportunities provide excellent points for follow on research.

1.9 SUMMARY AND GLIMPSE OF THINGS TO COME

The professional literature provides extensive coverage of traditional routing topics, but is virtually devoid of information with regards to differential routing. The absence of information in the literature, along with the lack of a commercially available tool to handle the task, sparked an aggressive undertaking to solve this problem. These efforts have produced significant results. Among the key findings to be developed are:

- A novel router architecture for solving the adjacent placement routing problem of differential signal pairs.
- A linear time recognition algorithm, which categorizes nets into one of the basis categories by means of a feature vector.
- A theoretical basis set of readily bifurcatable net categories, based on the development of a net to finite state machine (FSM) mapping operation.
- A linear time algorithm for bifurcating logical nets.
- Recognition that this architecture brings the full single signal routing body of knowledge to bear on the problem, thereby making an enormous stride up the differential routing learning curve.
- An extension of the theory into n-layers.
- A methodology for assessing inter-pair and intra-pair spacing.
- Extension of the fat wire concept into the MCM regime.

- Actual chip and MCM design results that objectively demonstrate interconnect length savings, while reducing total chip area, over non-adjacent routing of the differential wire pairs.

These key findings will be developed as the research journey unfolds in the remaining chapters.

CHAPTER 2

Managing Differential Signal Placement

HISTORICAL REVIEW

2.0 GENERAL THOUGHTS

The design of modern VLSI chips consists of several stages. First, the system is partitioned into major functional sections. These major sections must be placed or distributed on the chip to facilitate interconnection. Finally, a detailed routing of the circuits is accomplished. None of these is dealt with in isolation. Rather, an iterative cycle takes place, until the final solution is achieved.

The focus of this dissertation is on the detailed routing phase, in particular, the efficient management of differential pair placement. To provide a thorough foundation for the ensuing research it is important to review the background and discoveries leading to present day techniques used to deal with the problem. As was pointed out in chapter one, literature pertaining to the specific problem at hand is non-existent. But a wealth of material is available with respect to the routing problem in general. This chapter takes a look at the historical evolution of the routing problem.

Mankind has been solving the general routing problem both knowingly and unknowingly since earliest times. Cavemen attempted to layout the minimum route when going out to search for food. A lack of suitable footwear, along with exposure to the elements and other threatening species motivated their efforts. North American fur trappers unwittingly applied reasonable heuristics to find near optimal solutions to the "traveling salesman" problem of walking or touring their trap lines.

In the realm of printed circuit board (PCB), LSI and VLSI technology, routing solutions provided by early maze running approaches gave way to line probe algorithms. These were improved upon by artificial intelligence techniques, and rule based systems. For each of the major techniques, the computational complexity and the quality of solution provided is examined. Using this information as a measuring stick, the results obtained in this dissertation for handling differential pairs can be objectively assessed.

2.1 FACTORS TO CONSIDER

When assessing routing solutions there are certain characteristics or factors to consider. The four that have emerged as key are: (1) percentage of automated routing, (2) minimum chip area, (3) minimum number of vias, and (4) minimum total wire length[Joob86]. A short discussion of the importance of each, along with their interaction is provided in the following paragraphs.

2.2.1 PERCENT ROUTED

An ideal system would route one hundred percent of the nets, satisfying in so far as possible the other three optimality factors. However, with systems consisting of thousands and tens of thousands of nets, the ability to attain one hundred percent completion in an automated manner is a tenuous goal. When a system fails to completely route all nets, the final few must be dealt with by hand. In some cases the engineers can quickly identify the

net or nets that are blocking or precluding completion, and re-route them so that the final few nets can be successfully connected. On the other hand, when automated systems begin to fail, the density of the area in question has typically reached a point where human interaction requires many hours, if not days, to route the remaining nets. Certainly systems that guarantee completion are important, but ones that facilitate human interaction at all stages of the process may be even more important.

This is especially true where critical nets, such as clock signals or special power and ground busses, must be optimized. To effectively deal with such nets, they should be routed first, and in some cases done by hand[Bade92]. Systems should allow such interaction. Chapter eight of the dissertation discusses this particular problem when studying the MCM routing problem.

2.1.2 MINIMUM CHIP AREA

The higher the quality of the routing solution, the smaller will be the overall chip area. Reduction in chip size has been a design criteria for many decades. One of the principle reasons has been the defect probability associated with interconnect length.

As system clock rates continue to increase, propagation time through the circuits becomes a limiting factor. By decreasing chip size, and shrinking devices through improvements in lithography, the travel time can be kept within design constraints. If chip sizes are larger than optimal, the critical paths will experience a corresponding increase. The resultant effect can be a slow down in system performance. Hence, the push for optimal routing continues.