CHAPTER 8

Managing Differential Signal Placement

FAT WIRE CONCEPT APPLIED TO MCMs

8.0 General Thoughts

At the outset of the research it was understood that both the F-RISC/G chips and the multi-chip module that would serve as its package required full differential routing. As a natural progression in the research, the next step was to see if the insights gained from solving the chip level problem could be projected in the MCM domain. A re-examination of the problem constraints indicated that a considerable amount of what had been learned during chip level routing would map well into the MCM problem space. At the same time, certain conditions existed that prompted a different approach to bifurcation, which proved much more suitable to the MCM problem.

This chapter explores the differential signal adjacent placement problem at the MCM level. It begins by enumerating the constraint set imposed both by the project and the available core router. Then, taken in conjunction with the requirements, a revised implementation scheme is derived. The solution developed follows the architecture flow, yet differs in implementation from the chip level technique. It does, however, possess a flexibility and portability critical to the environment.

8.1 MCM Problem Specification

The original system specification called for a multi-chip module (MCM) to serve as both the package and the interconnect mechanism for the bit sliced F-RISC/G architecture. The immaturity of the technology being utilized forced the partitioning of the ALU components into manufacturable sizes. To successfully achieve the 1ns cycle time, the propagation time on the MCM substrate became a critical issue. Earlier research into thin-film MCMs utilizing low dielectric constant polymers indicated that copper wires patterned on layers of parylene, could meet the velocity and bandwidth requirements of the system[Dabr92]. A layout showing the originally envisioned chip placement on the MCM structure is provided in Fig. 8.1. Chip placement was performed by hand, seeking to insure the critical path distances were minimized.

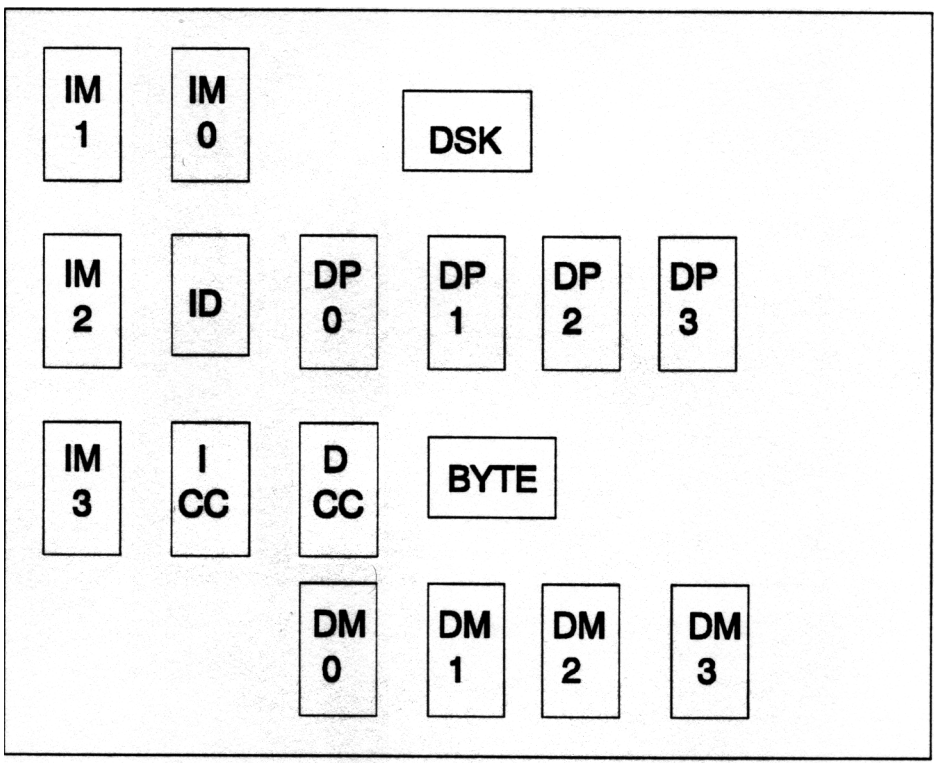

FIG. 8.1 ORIGINAL MCM CHIP PLACEMENT

Using three level current mode logic (CML), and retaining ECL voltage level compatibility, resulted in voltage swings on the differential lines of 250mV. In order to maintain as much noise immunity as possible, it was decided to utilize full differential signal routing on the MCM. With only a research version of a Wafer Scale Line Probe router available[Donl86], it was believed that considerable effort would have to be expended to provide the full differential capability desired. Certain differential routing functionality had been added to the system, but the technique used would most certainly result in net routing failures when the level of congestion approached 70 percent. Since the central area of the MCM was expected to exceed this, especially considering the block out regions for power and ground pillar vias, it was doubtful that it could meet our needs.

The high speed line probe router, when used in the single ended mode, produced very reasonable results. These results also conformed to the Manhattan specification with regards to metal layer orientation. If the basic functionality could be captured and used as the core router for the fat wire routing system, then the knowledge gained from solving the adjacent placement problem at the chip level could be quickly mapped into the MCM arena.

8.2 MCM UNKNOWNS

8.2.1 MOUNTING TECHNIQUE

Along with the specifications came a sizable list of unknowns. The actual MCM package definition was vacuous. It began with a C4 (flip chip) design. This was later replaced by the "Parker" package, that relied on tape automated bonding (TAB) to provide the connection from substrate vias to chip pads. The most important shortcoming of TAB was the minimum chip separation distance, which optimistically was claimed to be 3mm. This solution later gave way to a form of high density interconnect proposed by GE.

However, this too may be replaced by another technology at the time of actual implementation. The key point of this discussion is that the actual package specifications are in a state of flux, and this uncertainty must be provided for adequately.

8.2.2 Chip Geometry and Pad Out

The next major unknown is chip geometry and pad out. Although the primary architectural system signals have been relatively stable for some time, the corresponding location of the signal pads on the respective chips has been undetermined. Due to the high speed criteria, every picosecond of on chip signal time is budgeted. Consequently, it is not until the final logic block placement within each chip that the ideal pad out for every signal becomes known. This issue was compounded by the fact that full differential signals on the MCM doubled the number of pads, and it was not until recently that a staggered pad ring design was postulated and approved. The entire pad assignment and placement issue affected final chip size, since in the F-RISC/G version, the majority of chips are pad, not logic space limited. Once again, the lack of final details with regards to chip size and pad out predicated a routing mechanism with the flexibility to handle changing configurations.

8.2.3 Net List Specification

When the chip routing problem was undertaken, net lists evolved from schematic capture facilities. These automated systems were well defined and produced repeatable results that could be analyzed when attempting to decipher inversion information. The net lists utilized by the available line probe router were far less structured. Additionally, with pads on all four sides of each chip, and wafer I/O locations along the entire periphery of the substrate, the problem could be more aptly compared with a switch box than traditional standard cell or block routing problems[Ohts86]. This factor tended to unravel the feature vector recognition theory that performed so well for chip routing.

8.2.4 Wiring Size and Pitch

Wiring size and pitch can be varied as a benefit of the fat wire concept. But arriving at the optimum design tradeoff point requires much study. Chapter nine of the dissertation explores how the electro-magnetic characteristics vary as the wire spacing is changed. As wiring pitch increases, a penalty accrues either in overall chip area, or in wireability when assessing an MCM. With high core area congestion, too large a pitch will result in an inability to route the wafer. Yet if the wire size is too small, the current necessary to drive the substrate transmission lines may become excessive. But for the system to function, the farthest receiver on the longest net must see an identifiable signal. These factors are all in competition with one another, and it may take several actual routing iterations to see at what pitch the router fails to run to completion. But without well defined chips (size and pad out), a clear package mounting strategy (predicates minimum chip spacing), and a routing pitch that properly satisfies the competing needs of the drivers, receivers, and router, completing meaningful iterative cycles is difficult at best.

8.2.5 Revised System Chip Requirement

As cache studies progressed, it became apparent that with the original cache size, the system CPI figure would exceed 2.0. Since the viability of the entire project hinged on an optimal instruction rate, a CPI of 2.0 was determined to be too large. A major cache tradeoff analysis was conducted. Pushing the Smith curves to their extreme, a point of compromise was discovered[Smit80]. This new structure depended on a very wide block transfer capability, causing the MCM to undergo a significant modification. The revised layout, now consisting of twenty-five chips, is shown in Fig. 8.2.

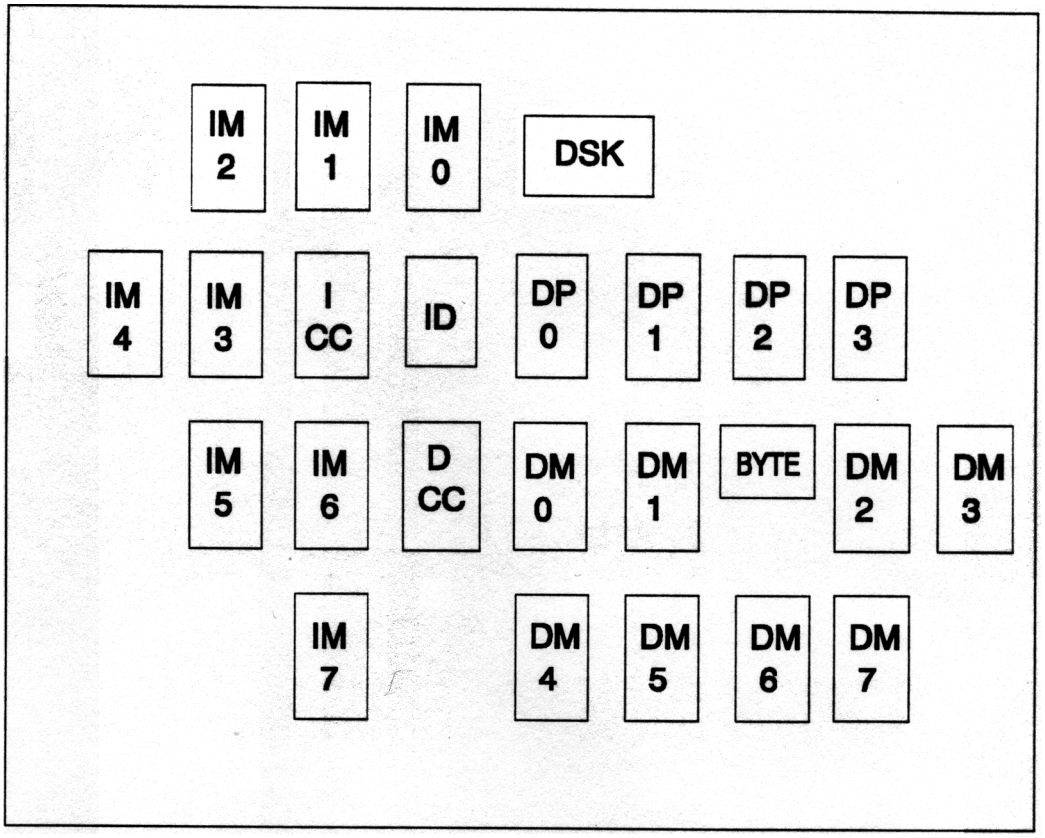

FIG. 8.2 MODIFIED F-RISC/G SYSTEM

8.3 GUIDED IMPLEMENTATION

The system requirements outlined in a previous section along with the many uncertainties noted above actually conspired to yield a very precise and portable implementation. As with the original chip level research, the stages of the fat wire router architecture were again assessed. Stage I had varying levels of difficulty. Since the system net lists were input by hand, composing the fat wire input list would be time consuming but relatively simple. Extracting inversion information would be considerably more complex. Much of the process can be automated once the final pad library is completed, and the chip pad outs finalized. Then, as the chip pad definition files are automatically generated, the relative location of the 0/1 pads for a given signal can be identified. Using

the basic net list to identify logical inversions, the pad location information and signal ordering can then be extracted.

Stage II, or the routing stage can be handled by the line probe router operating in its single ended mode. Provided the density level does not get too high, and the wiring pitch is not specified too large, a solid fat wire routing solution can be produced. This will serve as the input to the final stage.

When comparing methodologies for accomplishing bifurcation (Stage III) on the MCM, the obvious place to begin was the feature vector recognition approach. However, since an intermediate file containing information similar to the *CP* files is not available, another method had to be discovered. The first output available from the Donlan router is a *CIF* version of the routing solution. After carefully studying the options and the many uncertainties surrounding the MCM definition, it was found that the best solution was to bifurcate the *CIF* output.

8.4 CIF CUTTING MECHANICS

There are two main areas of discussion surrounding *CIF* file cutting. One involves the mechanics, and provides the computational complexity information. This data is necessary for assessing both the feasibility and the desirability of cutting *CIF* files. The second area of discussion surrounds the advantages of this method of implementation. In order to provide some of the factual information for the latter discussion topic, the mechanics of implementation will be presented first.

8.4.1 MECHANICS OF CIF BIFURCATION

The general flow through the MCM bifurcation process follows the diagram in Fig. 8.3. It begins with the *CIF* output, generated as the fat wire routing solution. This solution adheres to the fat wire pitch selected at the beginning of the process. As with any system

designed to output *CIF* in the most efficient manner, the router dumps all horizontal segments, followed by all vertical segments, and finally the vias. In this form, all net information is lost. A considerable amount of reverse engineering was conducted to modify the source code to retain net information when outputting *CIF* polygons. This effort epitomized the old "code" maintenance problem[Pres82]. After the modifications, *CIF* output could be requested in net cluster sequence. A net name, along with all of its associated polygons were clustered in sections in the output file.

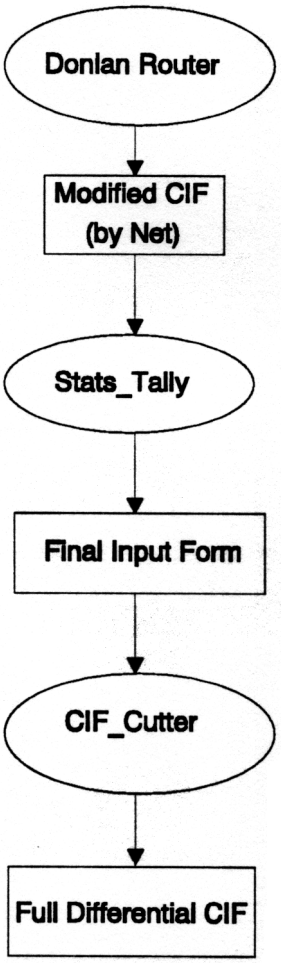

FIG. 8.3 CIF CUTTING FLOW

The next phase takes the modified *CIF* file and computes individual net and full route statistics. These are important for determining the longest net, congestion, and general solution quality when compared to the ideal routing length. In the process of tallying the net data, the file undergoes one final transform, readying it for the *CIF* cutting program.

Actual bifurcation of the *CIF* file is next. With the file structured in net order, the individual polygons comprising a net could be re-linked to form the full net topology. The re-linking is necessary to "understand" the bends and turns of the net at the segment junctions. In this way, when the major segments are cut, the appropriate vias and filler segments can be added. An upper bound on time complexity for re-linking the components is given by equation 8-1.

$$O(Linking) = N^2 \tag{8-1}$$

Since N represents the number of polygons in a given net structure, the time necessary to re-link all nets is given by a summation of the individual net linking times. The exact nature and characteristics of the MCM routing problem tend to produce nets that have very few segments. Also, the fact that the core router uses a line probe algorithm, acts to minimize unnecessary bends in nets, with the effect of reducing final *CIF* polygons. Typical nets possess anywhere from three to eight segments.

Once the linkage of the segments is established, the splitting portion of the operation commences. First, all major horizontal and vertical segments are regenerated as new *CIF* boxes. The new width corresponds to the wire geometry desired, and the new center coordinates are offset to produce the selected wiring pitch. A example demonstrating this operation is given in Fig. 8.4. One thing must be noted. The original endpoint calculations for segment length brought the original boxes together at a grid intersection point. When the fat polygons are bifurcated and then spaced about the old

center line, the wire ends no longer mate at the original junction. The highlight box on the bifurcated junction on the right, shows how the wire spacing from the old center line aggravate the via placement problem.

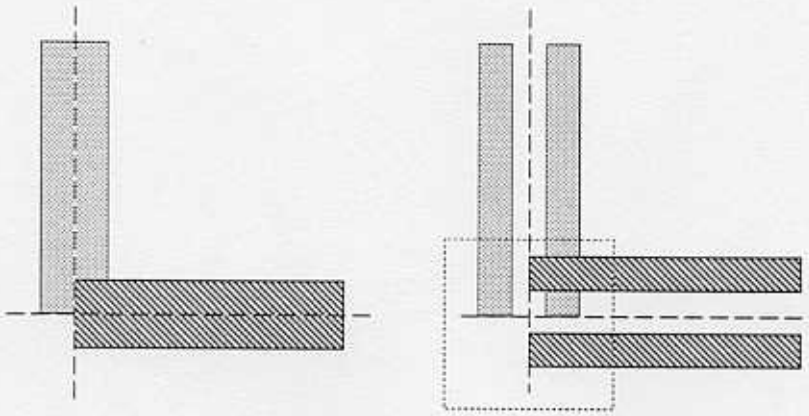

FIG. 8.4 CIF JUNCTION BIFURCATION

Instead of being able to directly calculate the modified via location, adjustments tailored to the wire geometry and pitch must be introduced. In the example above, the original fat wire intersection was laid out on a 40μm grid. After bifurcation, the differential pair-to-pair pitch will still be 40μm, but the wire to wire spacing will be 10μm with 10μm wide wires. The major grid is established by the core router, but all of the other dimensions can be varied directly as the *CIF* cutting process takes place. In this way, the *CIF* bifurcation approach permits even greater flexibility than the chip implementation.

Since MCM bifurcation was being handled in a different fashion than for the chips, the inversion problem was abstracted away along a slightly different path. The decision was reached to strictly apply the zero fat via cut (Table 6.1) at all segment intersections of a net. Originally it was believed that the fat wire bend vias should be bifurcated to coincide with the direction of bend. This would tend to save on wire crossings between the wires of a given pair. An example of the proposed technique is provided in Fig. 8.5. It is apparent, that by accommodating the bend direction, two crossings could be eliminated.

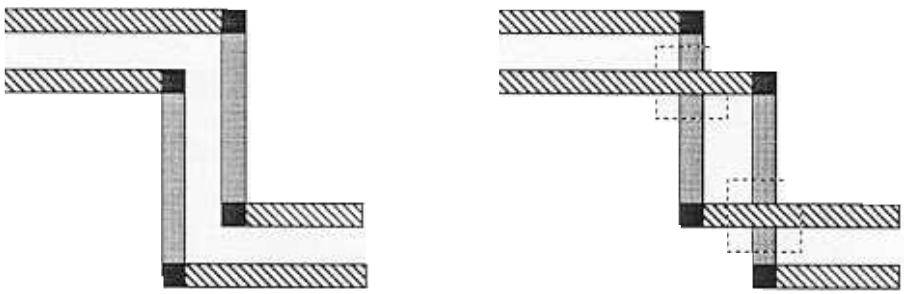

FIG. 8.5 BEND TAILORED BIFURCATION

Further analysis indicated that this impact was not as severe as anticipated. For a net that happened to bend in other ways, the zero cut approach matched what bend tailoring would have done. Additionally, if the junction was a multiple segment one, then the bend tailored concept has no meaning. At that point, either a polarity based decision or the zero cut technique would have to be employed. To accommodate the inversion abstraction, the zero cut method guaranteed that at any point along a net, the polarity was known. This has the effect of transforming the net into a transparent connection conduit, just as the multiple layers of metalization did in the chip routing domain.

An interesting discovery was made while solving the via placement problem. After partitioning the fat polygons into the two respective thinner ones, it became apparent that the segment endpoints were no longer synchronized. Reviewing Fig. 8.4, the end of the top horizontal wire and the right vertical wire actually overlap and are each too long. On the other extreme, the left vertical wire and the bottom horizontal wire are now too short to meet at the projected center line coordinates where the via would be located. An imaginary bounding box was constructed to surround the fat wire via site. To insure all wire ends observed this initial no-route zone, the split segments were all shortened by the dimension of the fat wire via. This would take into account the necessary reduction in

length at each end. Applying this technique to a sample junction produces the effect shown in Fig. 8.6.

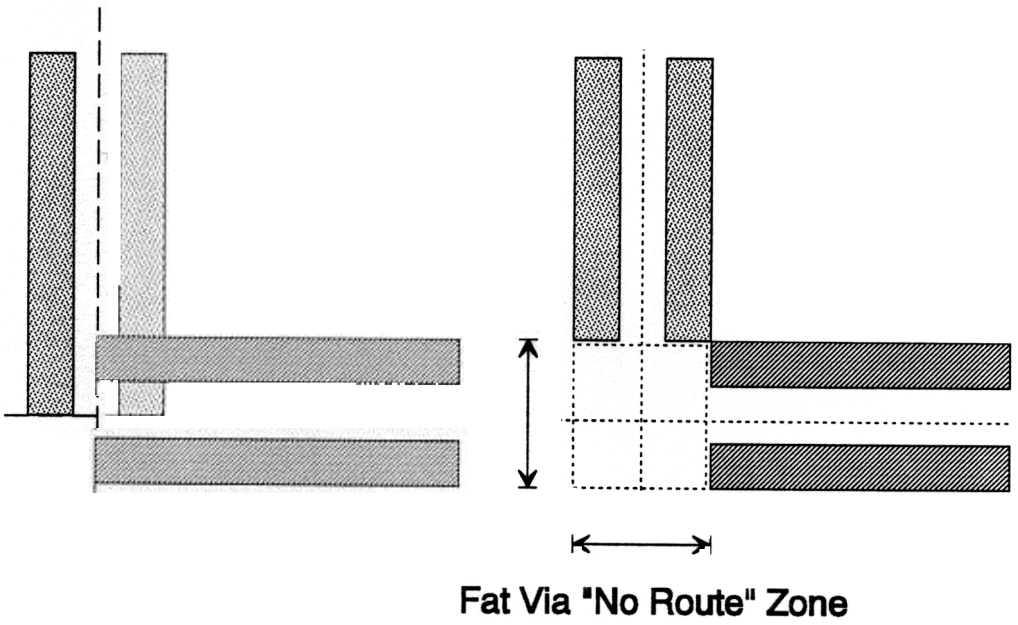

FIG. 8.6 CIF JUNCTION WITH SHORTENED WIRES

The problem now involves placing the vias on the extended center lines of each split wire, and then determining how the segments which do not meet can be joined. An initial idea proposed a staggered offset for the center coordinates of the two wire boxes. This overlooked the fact that the net ends should terminate at the same point, not appear staggered. It also neglected to take into account the fact that escape lines are constructed from chip pad locations out to routing channels. When chain nets are involved, this meant that you could actually have two different horizontal or vertical segments joining one another. This formed a straight line meeting with no requirement for offset. So although the idea possessed merit with respect to the standard two segment junction, there were numerous special cases that would have to be handled. The search continued for a robust solution.

The segment joining possibilities could easily be enumerated. A horizontal segment could possess from one to three other segment connections at each end. Likewise, the vertical segments also could have anywhere from one to three segments attached at their ends. This is graphically portrayed in Fig. 8.7. Preliminary analysis indicated that each individual situation could be handled by pre-constructing the appropriate via linking segments, and instantiating them when needed. However, it seemed that certain combinations of segment linkages might require the introduction of segment extension components that would conflict with one another.

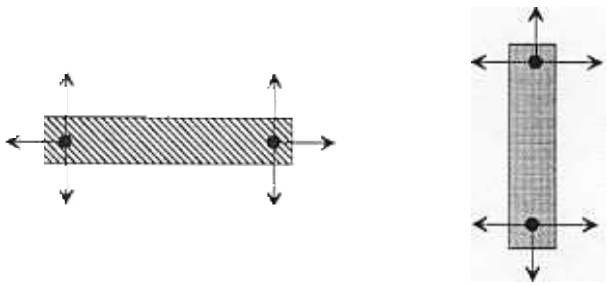

FIG. 8.7 END LINK POSSIBILITIES

The problem is easy to visualize, and a sample of how two segment junctions generate differing extension segments is shown in Fig. 8.8. The base segment is the vertical segment. If it were to have a connection to another vertical segment, the necessary connection segments are highlighted. If the same vertical wire had to be connected to a horizontal segment going to the right, then the connection segments to complete that junction are shown on the left. Since any combination of one or more of any type could exist, the possibility of conflicts, if each instantiation were handled in isolation, seemed very possible.

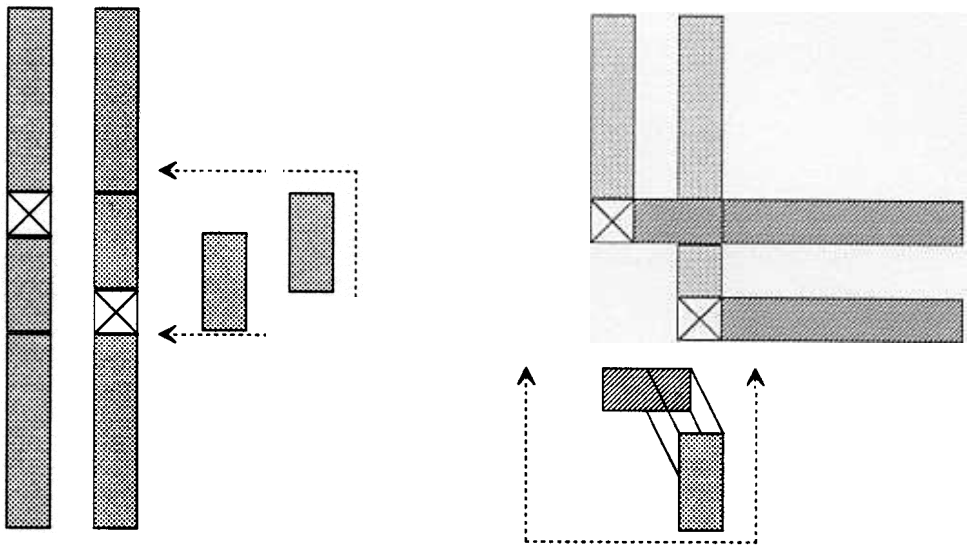

FIG. 8.8 CONNECTING SEGMENT ADDITIONS

Drawing upon experience in computer graphics, computational geometry and the experience of working with the zero cut phenomenon in chip bifurcation it was determined that conflicts would not occur. Additionally, the connecting segments for each joining operation could be pre-computed and instantiated based on what segment connections existed. The final solution segments are the logical OR function of the single connection junctions. This reduced the combinatorics of pre-computation of connecting segments for complex junctions.

If a composite junction between a vertical segment, another vertical segment, and a horizontal segment were to be formed(Fig. 8.8), the pre-computed segments for each of the one-to-one junctions could be OR'd and added to the data base. A separate database entry for that particular type of junction was not needed. An example of this technique is shown in Fig. 8.9, where the stored solution segments from the two separate junctions are merged through a logical OR, and the resulting three components are appended to the *CIF* file to complete the junction cut.

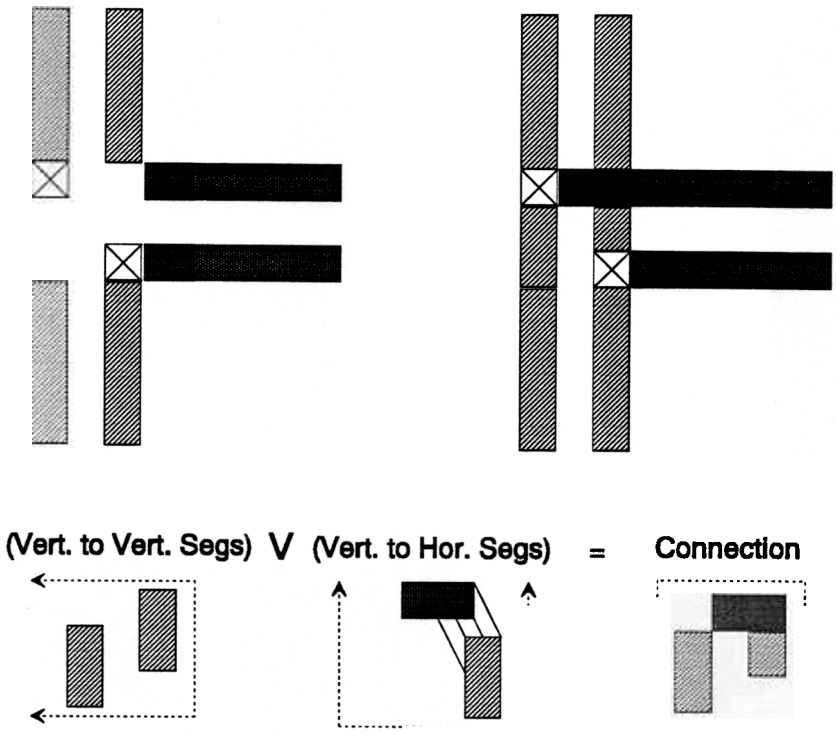

FIG. 8.9 OR'ING OF CONNECTION SEGMENTS

The remaining issues of *CIF* splitting deal with re-introduction of inversions, and the adaptability to system package definition. It was discovered that these two independent problems could be solved as one in a very elegant manner. For any net that connected a single chip to a wafer I/O location, inversions were not a factor. By maintaining the zero cut via splitting philosophy, the net appears transparent, and the wire ordering at the chip will appear at the wafer I/O locations. The inversion is solved by rewriting the wafer pinout-to-signal table, swapping the perceived signal ordering at the wafer edge. For chained nets involving multiple chips, a different methodology was applied.

With differential signal ordering standardized at the pads, the net lists can be scanned for inversions. If an inversion is required at a chip pad site, a creative and flexible solution evolved. All wires that connect to chip pads, terminate some delta from the actual

pads. Two pre-defined *CIF* macros are constructed. One extends the wires directly to the pads without inversion. The other contains the wire crossover necessary to effect the signal inversion. This crossover is handled in the spreading layer, so that normal routing tracks are not impeded. This has a very significant added benefit. Since the final package is still in flux, and the mounting technology not settled, a variable spacing capability is introduced. The connecting macros can be scaled to adapt to actual package mounting technology. This permits near final routing studies to be conducted, independent of the mounting technology decision. Samples of such chip connecting macros are contained in Fig. 8.10.

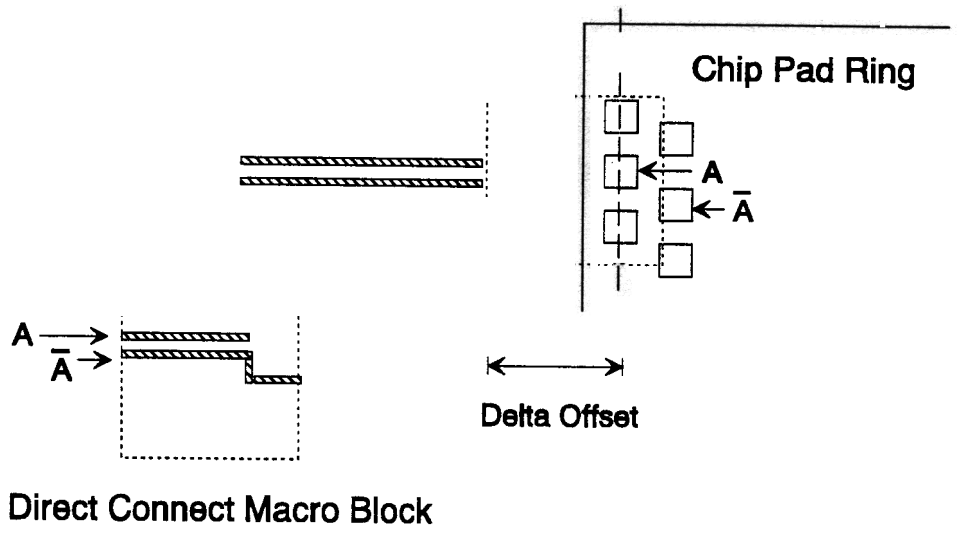

FIG. 8.10 CIF MACRO ATTACH EXAMPLE

Not specifically characterized as a *CIF* splitting function, but directly related to constructing the completely routed MCM system is the issue of critical net routing. In an ideal environment, the core router should permit human guidance and interaction. This is especially important with critical paths and clock nets. Since our core router lacked this capability, an alternative approach became imperative. Preliminary routing analysis indicated that the density in the central area of the wafer would come seriously close to

exceeding the abilities of the line probe router. With the package definition in a state of transition, and the need for extra power and ground planes becoming evident, an additional wiring layer was postulated. This layer would be used specifically for the system's clock deskew scheme. Two differential pairs must be routed from the deskew chip to each of the other clocked chips of the system. These lines must be as nearly balanced in length as possible. Noise minimization along these routes is of paramount importance. By assigning a specific layer for their use, the clock deskew nets could be routed by hand at the end of the bifurcation process. Using the state of the art *CIF* layout editor within VLSI Tools, the deskew nets can easily be laid out to meet performance requirements.

Insuring the proper blockout of power and ground pillar vias was the last hurdle. The core router must have a mechanism for specifying no route regions so that the power and ground pillar vias can be patterned without concern for signal routing. The Donlan router, with its origin as a wafer scale router, was designed to bypass specified regions on a wafer that were determined to possess defects. By extracting power and ground pad locations from the designers files, block out regions could be generated and input to the router. As currently implemented, the router makes no distinction between layers when avoidance regions are described. If a zone is established, all routing layers in the zone are blocked from use. In a production router, it is important to be able to specify block out zones by layer. Since power and ground vias will only cut through to their respective planes, signal planes not involved would have free use of those layers. For high density interconnect, this is very important.

As part of this investigation, the idea of a macro-via came to light[1]. With proper

[1] This idea was conceived by Professor Rymaszewski.

planning, critical signals could be further shielded while transiting the pillar vias, by incorporating a power and ground combination in the same macro-via structure. A diagram of how such a structure might look is provided in Fig. 8.11.

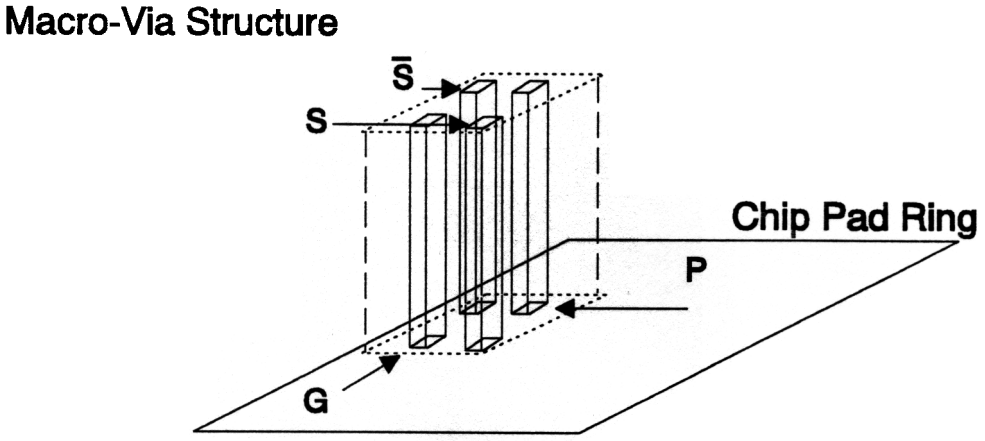

FIG. 8.11 MACRO-VIA DIAGRAM

8.5 SOLUTION CHARACTERISTICS

CIF file cutting as the bifurcation methodology provided many of the characteristics of the feature vector approach. The inversion problem could be abstracted out of the problem space, and only re-introduced at the final step. The time complexity associated with *CIF* splitting is directly tied to the net linking algorithm. Even applying brute force, the computational complexity is no worse than N^2. For the average number of segments of an MCM net, this cost is still very small. Following this approach again skirts the computational complexity and verification issues associated with a generalized bifurcation algorithm.

The unique macro cell connection operation effectively deals with the inversion problem, while at the same time providing enormous flexibility in package mounting technology. This feature allows MCM development to proceed in parallel with chip

design. Final packaging decisions can be postponed until full system simulation is performed, thereby waiting until correct power dissipation and signal propagation properties have been determined.

The CIF format is well defined and almost universally accepted in industry. By attacking the bifurcation problem in this arena, the solution proposed in the dissertation is immediately implementable by others. And although by-hand work to route critical nets is frowned upon, high speed systems under development today use it for their important signals[Bari92][Bade92]. The interaction allowed is only limited by the *CIF* editing routines at ones disposal.

8.6 SUMMARY

Early analysis indicated that mapping the feature vector idea from the chip domain into the MCM routing domain would provide the optimal solution. However, further study revealed that the MCM net problem more closely resembled the switch box routing problem than the channel routing problem for which feature vectors were intended. Many uncertainties with regards to the MCM specification, and actual chip geometry and pad outs dictated a flexible. After searching the design space and evaluating the tradeoffs, it was decided that the most suitable bifurcation technique for MCMs was a *CIF* splitting approach. Although appearing somewhat restrictive in nature, the benefits of the methodology mounted. The added ability to vary wire geometry and spacing beyond what could be done in the chip regime will undoubtedly be very important. Flexibility in dealing with package mounting technology through the use of *CIF* connection cells is imaginative. Finally, the use of macro-vias for shielding critical signals in the pillar vias may grow in importance as signal levels continue to decline and noise margins are further reduced.

An interesting research note should be mentioned here. Even though conceptual solutions seem to map well from domain to domain, the implementation techniques may not. In moving from the chip regime to the MCM regime, had the feature vector methodology been forced upon the problem, the more advantageous approach of *CIF* splitting may never have been recognized.

CHAPTER 9

Managing Differential Signal Placement

ELECTRO-MAGNETIC ANALYSIS OF DIFFERENTIAL PAIR SPACING

9.0 General Thoughts

One of the key benefits provided by the fat wire routing system is the ability to vary two independent wire spacing parameters. Traditional routing systems accommodate wire spacing through the use of a technology specific parameter established in the router guidance mechanism. This parameter is usually dependent on the minimum lithography capability of the fabrication process. Once the device geometry and minimum spacing criteria are determined, it is used to guide wire placement. The most obvious way of employing it involves the grid based routing mechanism. Horizontal and vertical routing tracks are established on a large grid system. The point to point spacing on the grid is directly related to what the ultimate wire spacing will be. Since all wires must fall on a horizontal or vertical grid, the grid intersection points serve as intersections on which to begin or end interconnect segments. With this traditional grid system, the point spacing is consistent throughout. The only way of varying spacing would be to lay wires on various

multiples of tracks. Clearly, this would be extremely wasteful. With the premium given to chip area, such a solution would be discarded rapidly.

The ability to vary two parameters opens up the possibility of conducting a thorough electro-magnetic analysis of the pair spacing parameters and their effects on capacitance, and chip area. This chapter begins by examining the parameter spacing relationships that exist in both the chip and MCM routing domain. The discussion continues by looking at the relative effects on chip and wafer area as differential wiring pitch is modified. Then a methodology is proposed that allows both chip and MCM designers to utilize their process specific parameters to compute optimal design tradeoff regions. Extracts from the F-RISC/G project MCM analysis are provided as numerical examples.

9.1 WIRE SPACING PARAMETERS

Preliminary results from chip routing indicated that the technique could generate wire pairs whose wire to wire spacing differed from the pair to pair spacing. Instead of being tied to a completely rigid grid, where all wires are evenly spaced, the fat wire approach to handling differential routing overcomes this physical roadblock, and allows for varying two spacing parameters independently. The spacing between the wires in a differential pair is adjusted when the fat wire geometry is set. This geometry is governed by equation 9-1, where FWW (Fat Wire Width) is the edge to edge dimension of the fat wire, NWW (Normal Wire Width) is the fabrication wire geometry, IWG (Inter Wire Gap) is the space between the inner edge of the two wires of a pair, and BS (Buffer Space) is the distance from the outer edge of a normal wire to the boundary established by the fat wire.

$$FWW = IWG + (2 * NWW) + (2 * BS) \qquad (9\text{-}1)$$

The spacing between outer edges of wires composing adjacent pairs, D_D (Differential to Differential spacing) is presented in equation 9-2.

$$D\_D = FP - \{(2 * NWW) + (2 * BS) + IWG\} \tag{9-2}$$

These spacings and respective dimensions are depicted in Fig. 9.1.

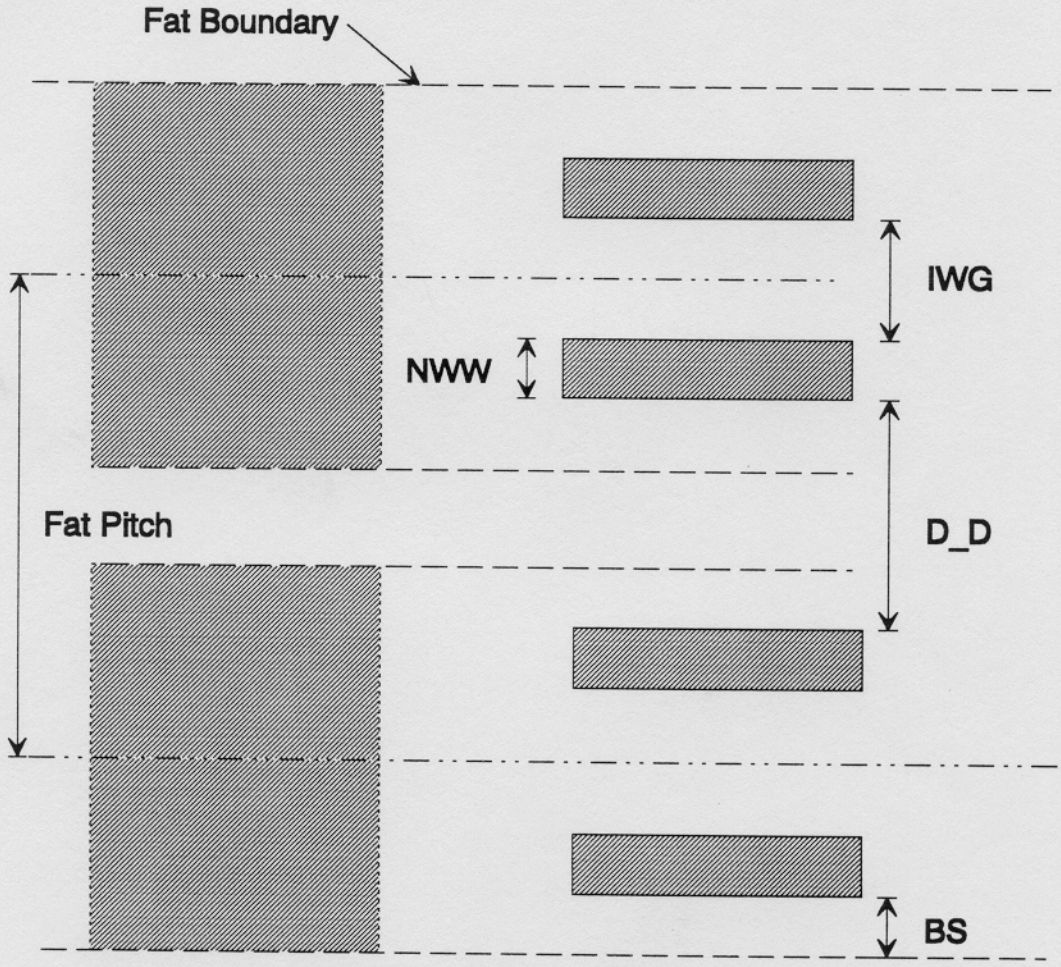

FIG. 9.1 WIRE SPACING DIAGRAM

By appropriately varying the design parameters, the differential-to-differential spacing does not have to match the inter-wire gap. If a grid based routing approach to routing differential wires is used, the wire-to-wire and pair-to-pair spacing would have to be the same.

9.2 E-M CHARACTERISTICS OF WIRE SPACING

If a traditional approach to differential routing is taken, the wire spacings would all be set using available lithography limits. The fat wire approach opens the door to a more thorough EM analysis. Ideally, this analysis should be conducted during the early stages of the chip and MCM design process, so that the results obtained can be integrated into the solution.

The first factor which comes to mind concerning differential pairs is the added wire capacitance. In CMOS circuits, efforts are made to preclude two signals from paralleling one another for long distances to reduce the capacitive coupling effects. Since the wires of a differential pair are deliberately placed in adjacent tracks throughout their run, the capacitive effects are noticeable. In the F-RISC/G chip domain, the effects are even more pronounced, since the GaAs substrate is semi-insulating and does not appear electrically as a normal ground plane. By increasing the spacing between the wires, the capacitance between them can be reduced.

The next factor to consider is crosstalk. It can occur between the two wires of a pair and also between the pairs. The worst case situation is a quiet differential pair located between two other pairs, which each transition in opposite ways. In this case, the quiet pair would see the cumulative effects of both nearby wires undergoing like changes.

Finally, signal loss must be considered. This is especially important on the MCM. Wiring length gets longer as the individual wire spacing increases. Thus, for a given MCM line driving ability, the wire spacing can affect the ultimate distance over which a given signal can be successfully transmitted. The relationship between both chip size and MCM as a function of wire spacing will be developed in a following section.

9.3 ANALYSIS METHODOLOGY

The analysis must consider two regimes: Chip Level and MCM Level. The basic analysis techniques are general enough in nature that they can be described prior to investigating the peculiarities associated with each individual regime.

For a given technology, dielectric stacking and wire dimension, a thorough analysis can be conducted to independently determine the capacitive effects of first varying the wire to wire spacing within a pair. Then, using some of the preferred results, the effects of pair-to-pair spacing can be analyzed. In the context of this research only the coupling effects have been studied. A full analysis must include a study of the magnetic effects

Having investigated and plotted curves showing wire spacing versus capacitive effects, chip area and MCM routing track curves can be super-imposed. The region where these curves intersect represents the best design tradeoff point.

9.3.1 CHIP DOMAIN ANALYSIS

For each chip technology there will be a relatively inflexible cross section. This cross section is typically dictated by the fabrication technology. The substrate material, interlayer materials, and encapsulants all have different electrical properties. It is this cross section that must be simulated in order to understand the capacitive relationship between wires of a pair, and between differential pairs.

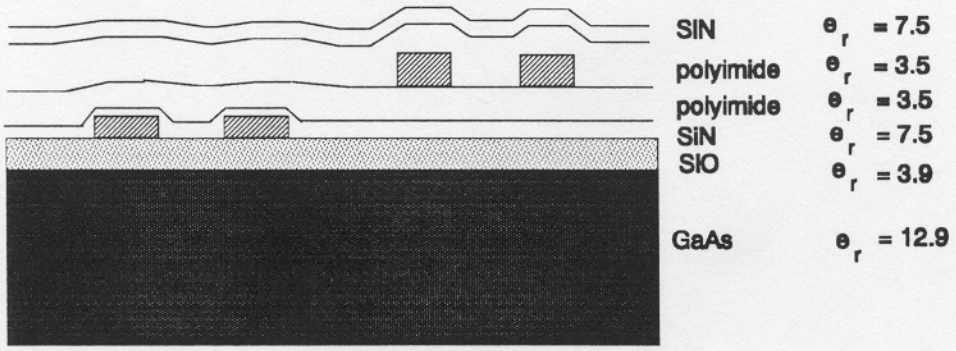

FIG. 9.2 SAMPLE ROCKWELL CHIP CROSS SECTION

A sample cross section taken from the Rockwell GaAs process, is shown in Fig. 9.2. In this case, the minimum design rule spacing of 3μm has been utilized. Since this is the minimum design rule size, the only option is to space the wires further apart. Beginning with this minimum spacing, the capacitive coupling can be computed using simulation programs. A diagram showing the coupling interactions between two differential pair is provided at Fig. 9.3.

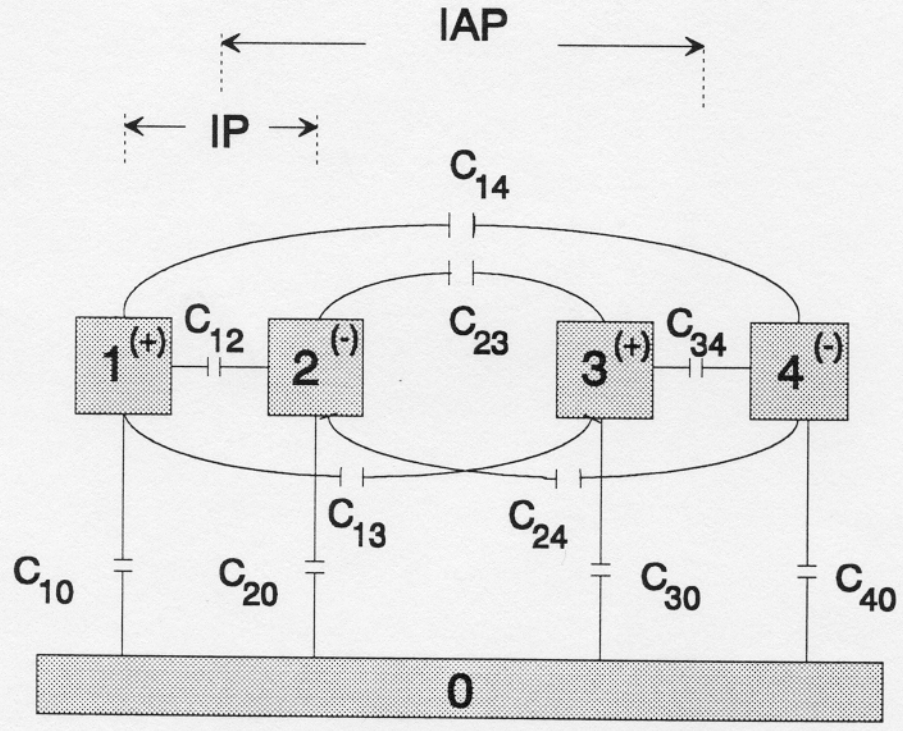

FIG. 9.3 PAIR COUPLING MODEL

Since the cross couplings appear in parallel, the total resultant can be computed as a sum of the component capacitances. The results of a basic cross section analysis for a single pair is given in Table 9.1. The capacitive values are per unit length of wiring. Using this factor as input to a spreadsheet, average net lengths can be provided to compute capacitance loading effects due to the pair interaction.

TABLE 9.1

ROCKWELL CHIP GEOMETRY - WIRE CAPACITANCE

Wire Spacing	Wire Thickness	Capacitance
3 microns	0.5 microns	0.116fF/micron
3 microns	1.0 microns	0.059fF/micron

Developing chip growth rates as functions of wire spacing are very problem specific. One way of estimating the effects is to approximate the number of horizontal and vertical routing tracks. If this can be done without full cell library development, and a detailed placement and routing solution, then a rough area function can be formulated. To compute a floor on chip area, the minimum wire spacing dictated by lithography and fabrication technology is used. This value is first multiplied by the number of horizontal wiring tracks. It is then added to the number of cell rows multiplied by the standard cell row height. Next, the number of vertical tracks not contained within the bounds of the standard cell region must be multiplied by the minimum spacing. This result should be added to the average standard cell row length. Taking the results of these two calculations as the height and width of the chip, the area can be approximated. From this point, as the wire spacing in increased, a revised estimate of the chip size can be calculated.

However, the growth function is not necessarily a smooth one. For the wiring to correctly mate to the standard cells, the port locations along the cell boundary must match the fat wire routing grid. Cell designers, in an effort to minimize chip size, strive to compact the standard cells as much as possible. Consequently, the cells grow in quantum steps. If all of the logic can be contained within a given multiple of the routing grid, then stretching the port spacing may not be possible without a fat grid jump in size. This concept is shown pictorially in Fig. 9.4. The cell on the left can just be completed at a

multiple of the routing grid. If the differential pair spacing must be increased, there is not enough room along the edge of the cell to reposition the ports. To stretch the cell to accommodate the new port locations, it must grow in a multiple of the fat wire grid. This is shown by the cell on the right, where the adjusted size is a full grid step larger.

If room exists along the cell boundary to accommodate spreading, wire spacing may increase many microns before the next step limit is reached. Once reached, then the next increment in wire spacing will boost the cell size by another grid spacing. This all means that the horizontal growth curve will be a stair step like function. Plotting the vertical and horizontal growth functions, the product function representing chip area can be estimated, Fig. 9.5.

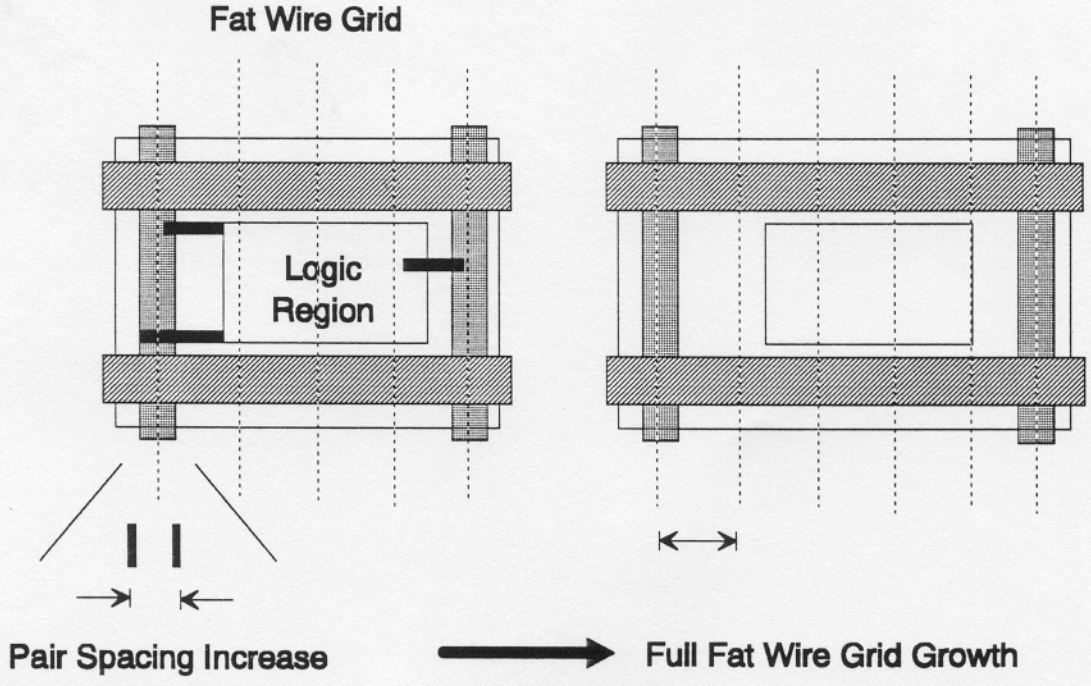

FIG. 9.4 CELL GROWTH BY STEPS

The capacitance curves should be constructed varying the spacing of a single wire pair. Next, using selected steps from the single pair analysis, pair to pair curves should be generated. This will produce a family of curves that show coupling length, wire spacing

and capacitance. Naturally, as normal plate capacitance falls off as function of 1/x, the curves should all be similar, with their bend locations slightly offset from one another.

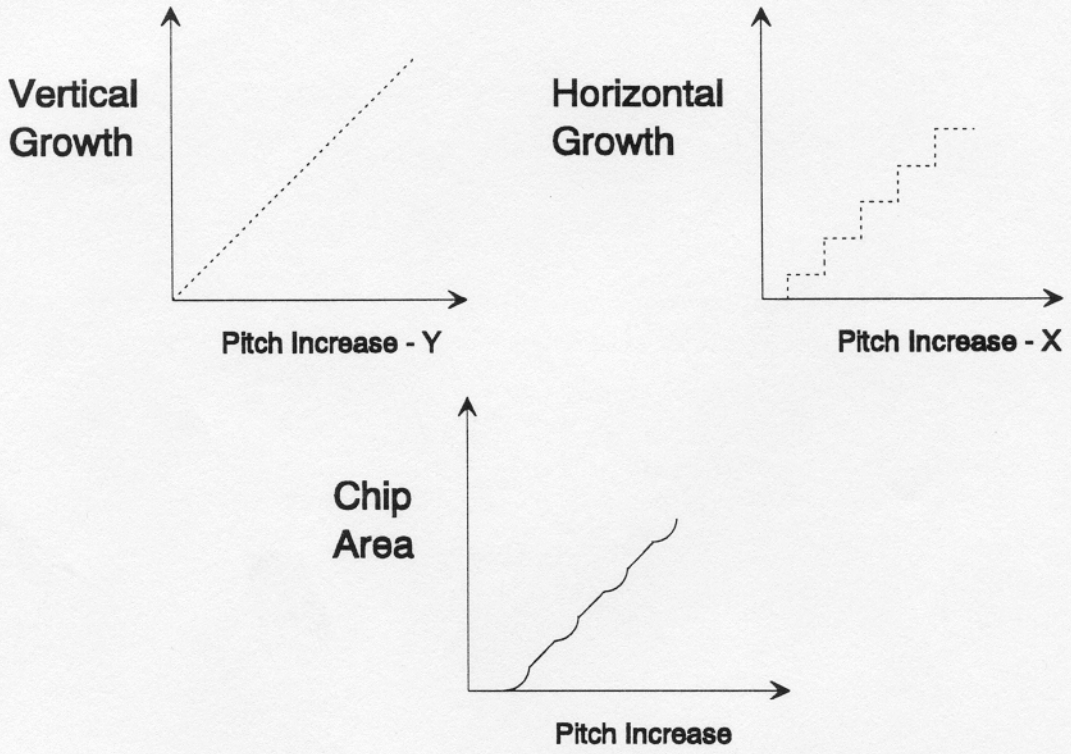

FIG. 9.5 WIRE SPACING VS CHIP GROWTH

For established design facilities, historical data for cell size, average row length, and overall number of horizontal and vertical routing tracks can be utilized. This data can be manipulated to produce pseudo-chip growth curves as a function of wire spacing. By overlaying the chip growth curve on the family of capacitance curves, an optimum design tradeoff region can be identified. Then, based on design requirements and specifications, the actual wire spacing is determined.

The calculated wire spacing is then utilized throughout the design cycle. Standard cell design criteria are established so as to attempt to fall within grid multiples. The system router is modified to understand the selected fat wire pitch. The translators whose job it is

to fold differential ports into single ones can be constructed. And finally, the bifurcation software can be adjusted to produce the final patterned wires. For chips that will ultimately be used in MCMs, pad spacing on the chips, which is typically dealt with in a vacuum, must be reviewed from an MCM routing perspective.

9.3.2 MCM Domain Analysis

The methodology for assessing and determining the proper wire spacing for differential pairs is similar to that used for chip evaluations. However, there are some subtle differences. Instead of evaluating the effect of chip growth as a function of increases in wire spacing, the MCM size is more constant. Typically it will be fixed, based on the chip mounting technology. In the MCM space, increases in wire spacing result in fewer overall wiring tracks. Fig. 9.6 shows this phenomenon. If a wafer of 10cm by 10cm is going to be used as a substrate for a system, then the wiring grid will be overlaid on it.

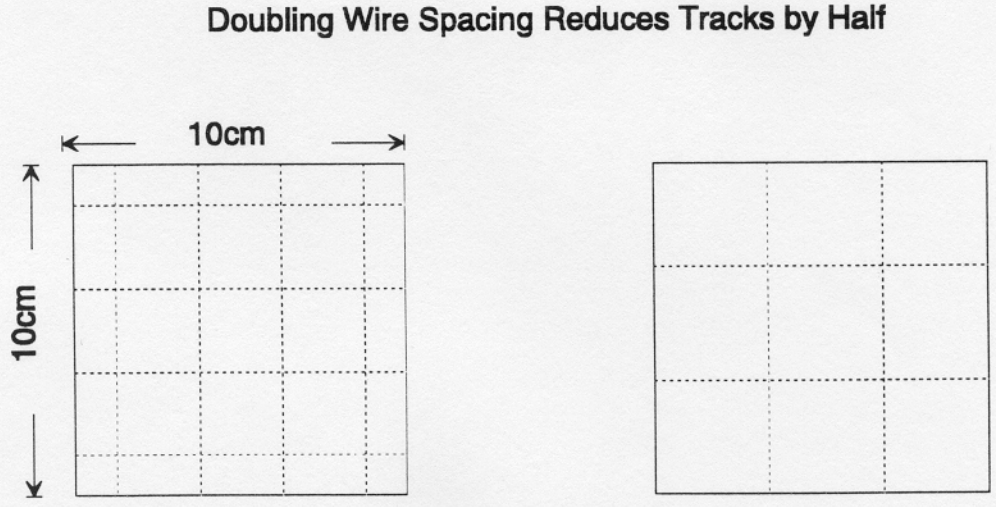

Fig. 9.6 Wire Spacing Vs Wire Tracks

By taking the wafer size and dividing it by the grid spacing, the total number of horizontal and vertical wiring tracks can be calculated. Chip connectivity parameters along with basic wiring congestion influence the routability of the system. A systolic array layout with

regular connection patterns, and relatively few long connections may be routable in fewer tracks than a tightly coupled bit sliced processor that actually occupies less total wafer area. But as a general rule, as available routing tracks are reduced, the routability percentage will decrease.

As with chip design, coupling between wires must be balanced against available routing tracks. In addition to coupling, the signal loss on MCM wires must be considered. Even with perfectly terminated wires behaving as transmission lines, differential drivers attempting to communicate to receivers 10cm away, must consume 8mA of current. This assumes a copper wire with a ten by four micron cross section. Consequently, in addition to wire spacing, calculating the tradeoff wire geometry is also of paramount importance.

The forces in opposition in the MCM regime are: (1) capacitive and magnetic coupling which would tend to push wire spacing further apart, (2) line drive capability (either current or voltage), which calls for larger geometry wires and the commensurate increase in wire spacing, (3) available routing tracks, which are a necessity for completing the interconnect system, and tend to push for smaller wire size and spacing, and (4) package heat removal capability. These factors must all be considered and satisfied simultaneously. This calls for a flexible and interactive wafer design system, since most of the analysis is carried out through an iterative cycle. The four analysis functions and sample curves are provided in Fig. 9.7.

To perform MCM signal drive analysis, MAGICAD was utilized. First a cross section extraction was performed using the thin film substrate model [Dabr92]. With the electrical parameters extracted, a network calculation was then set up. For the given wire geometry's, dielectrics, wire length and drive current, a correctly terminated set of differential transmission lines were simulated to see whether or not the signal at the receivers could be correctly interpreted and proper noise margins maintained. A sketch of

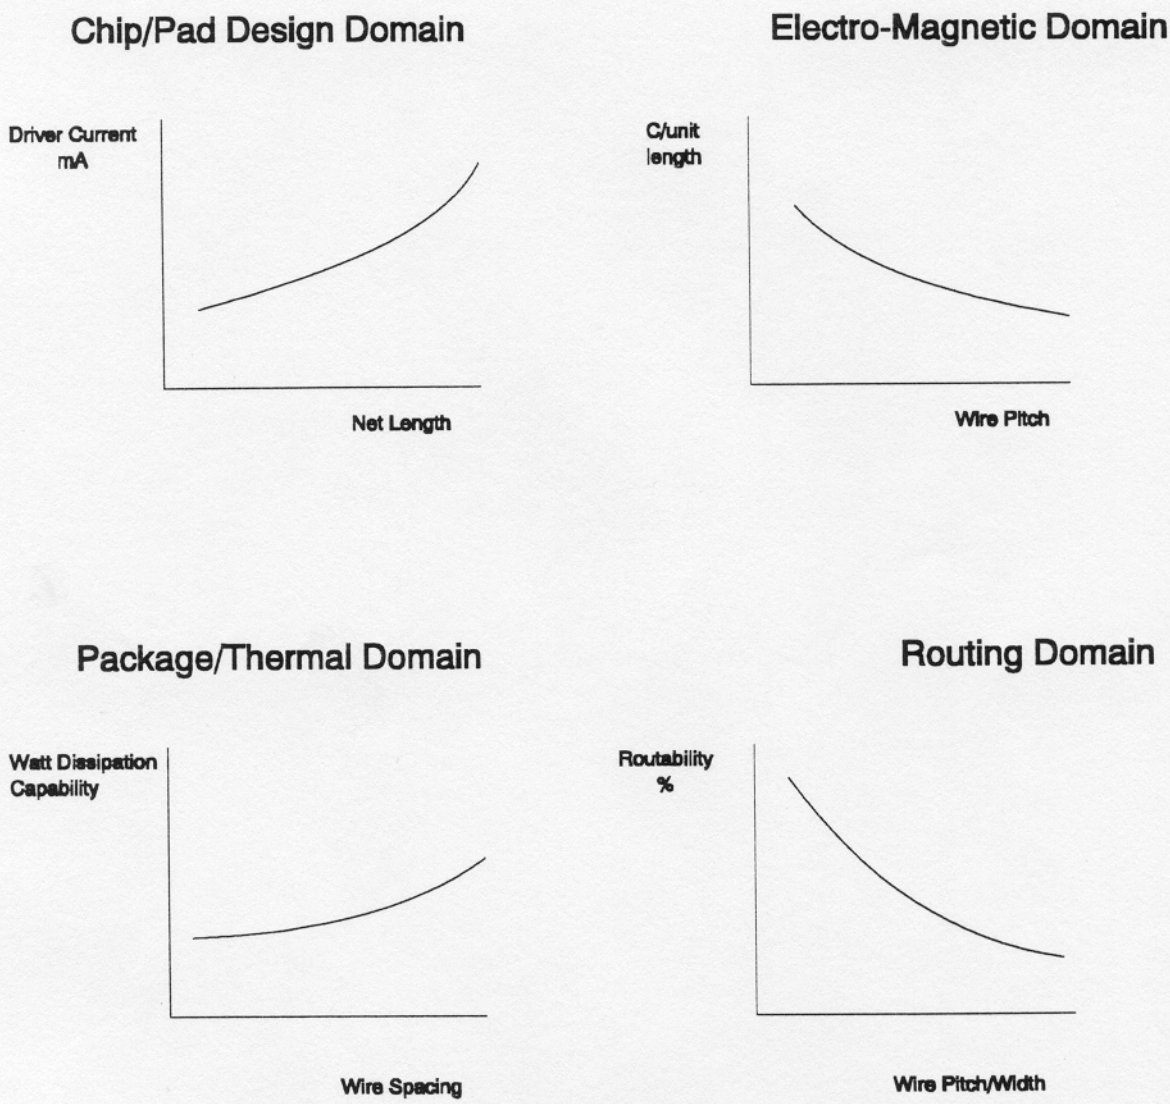

FIG. 9.7 MCM ANALYSIS DOMAINS

the pair cross section is shown in Fig. 9.8. The extracted electrical parameters are shown in Table 9.2. The network diagram for the simulations is provided in Fig. 9.9, with the input and output signals displayed in Fig. 9.10.

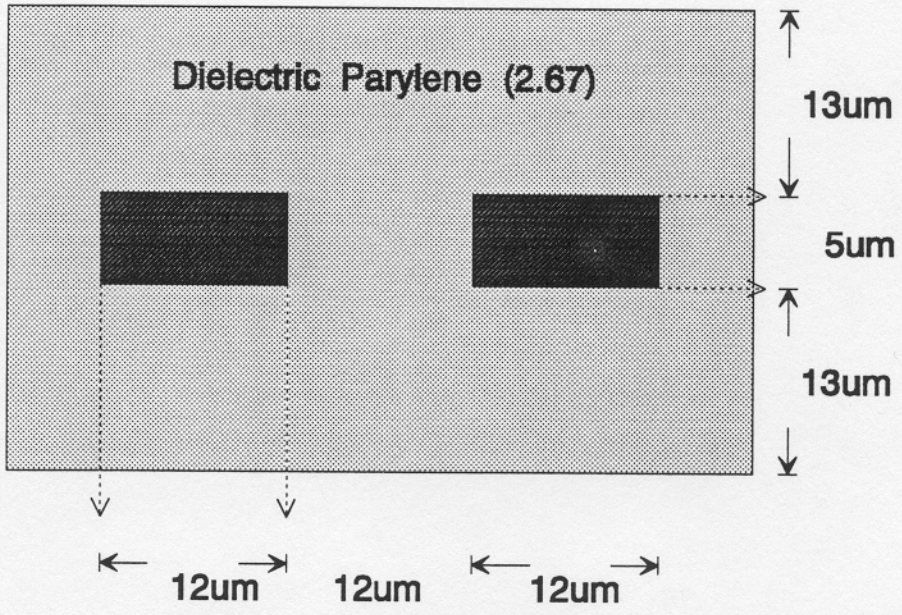

FIG. 9.8 MCM EM CROSS SECTION

TABLE 9.2

CROSS SECTION ELECTRICAL PARAMETERS

12μm x 5μm	# 1	# 2	# 1 - # 1	# 1 - # 2	# 2 - # 1	# 2 - # 2
C (F/m)			1.12E-10	1.52E-11	1.52E-11	1.12E-10
L (H/m)			2.69E-07	3.64E-08	3.64E-08	2.69E-07
$\mathbf{Z_0}$	49.00	49.00				
$\mathbf{R_{dc}}$ (Ω)			557	0	557	0
V (inch/sec)	7.22E+09	7.22E+09				

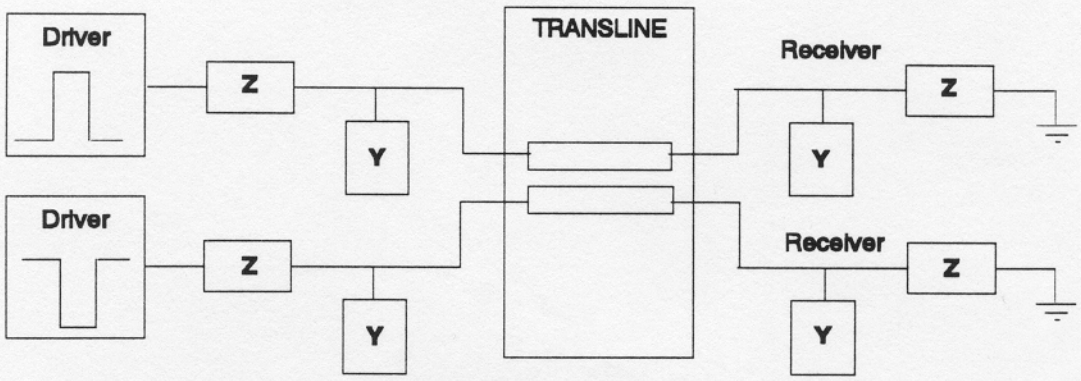

FIG. 9.9 NETWORK SIMULATION DIAGRAM

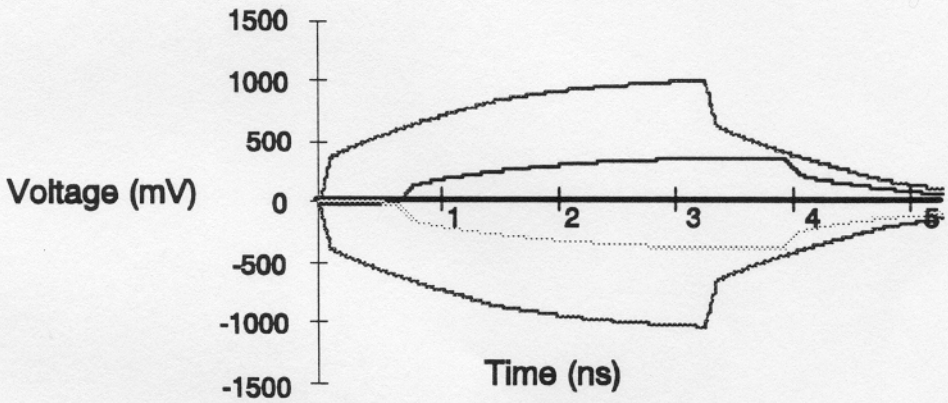

FIG. 9.10 INPUT AND OUTPUT WAVE FORMS

At the same time that the driver study was being conducted, the routability tests were being carried out. Using an anticipated grid spacing of twenty microns and wires ten