Foundations of Interconnect & Microstrip Design - 3rd Edition

by Terry Edwards & Michael Steer

The latest edition of this classic work by Terry Edwards of Engalco has now been fully revised and expanded with Professor Michael Steer from the University of Leeds and North Carolina State University as co-author and is due for publication by Wiley this month. We reproduce here, with kind permission of the authors, excerpts from the first two chapters which are new to this edition.

As with earlier editions, this book provides definitive treatment of microwave and millimeter-wave transmission lines. The coverage ranges from the principles of operation through to the design formulas, design guidelines and applications. Differential lines, critical to controlling noise in RFICs but also compatible with differential active circuit design, has expanded coverage. In digital circuits connections can sometimes be treated using simple lumped element connections but sometimes must be modelled as transmission lines. There is a continuous range of behaviours between these extremes and "interconnects" is the unifying terminology used to refer to these lines. As most readers are familiar with microwave transmission lines this book preview concentrates on interconnects in digital circuits and mixed signal systems.

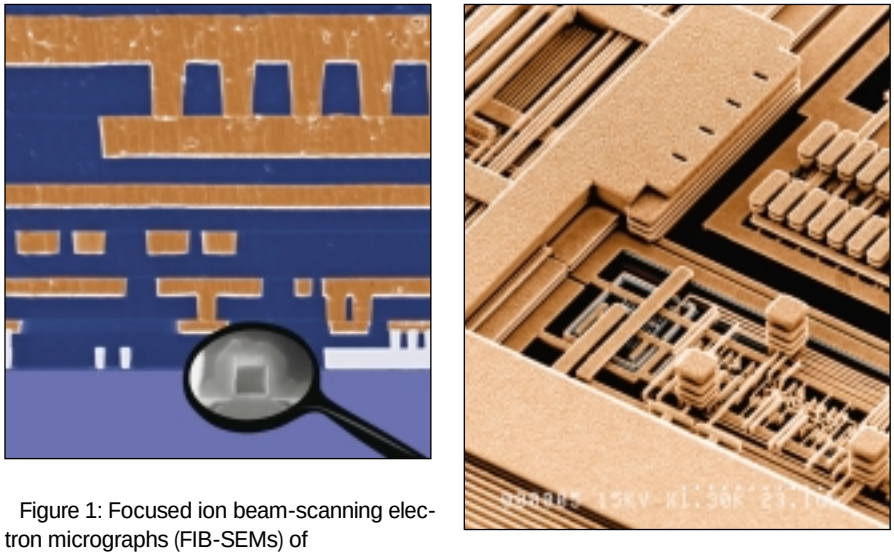

Figure 1: Focused ion beam-scanning electron micrographs (FIB-SEMs) of a six-level interconnect structure developed by IBM. The first (bottom) level is tungsten as are the vias. The other metal levels are copper: (a) cross-section; and (b) perspective after etching of oxide dielectric. © IBM, from IBM web page.

Electrical connections are part of a hierarchy that connects individual active devices at the lowest level to system-level connections at the highest [1-3]. One way of defining the interconnect hierarchy is as follows:

- Level 1 Interconnect: Chips
- Level 2 Interconnect: Multichip Modules (Package)
- Level 3 Interconnect: Printed Circuit Board PCB (Package) (Printed Wiring Board. (PWB) is the preferred usage in N. America while printed circuit board is used elsewhere.
- Level 4 Interconnect: Backplane (Package)
- Level 5 Interconnect: Rack, Connect Systems (Package)

The size of interconnects varies with interconnect level with both the smallest interconnects, both in cross-section and in length, being at the chip level. Also the way they are designed and modelled depends on the physical size of the connections, whether digital or analogue circuitry is being interconnected, and on the clock or operating frequency.

Interconnects and not the transistor switching speed now dominate the performance of many high-speed digital circuits and for RF and microwave circuits they have always been critical circuit components. The lower the level in the interconnect hierarchy generally the shorter the connection and the higher the performance in terms of lower delay to transmit a signal from one point to another.

A high performance system realizes a greater proportion of interconnections at the lower levels than does a lower cost system with less critical performance requirements. Not every system has all of the levels listed above and even some interconnect networks, especially at microwave frequencies, need not involve active devices.

Today's high performance digital systems have many interconnect issues common with microwave (generally 1 to 30GHz) and millimeter wave (above 30GHz) systems and it is appropriate to consider the interconnect issues as a continuum. For example the long haul interconnects on a chip (comparable in length to the edge dimensions of the chip) are conceived of and designed with the same criteria used to realize millimeter-wave circuits.

Up to the early 1990's on-chip digital signals had components much below 1GHz; relatively short run lengths (constrained by the dimensions of the chip); and widths of several microns; and height of a micron or more.

The first two factors resulted in the electrical lengths of on-chip interconnects being much less than a wavelength of the

signals present (i.e. of the highest frequency components present).

Then an interconnect could be adequately modelled as a shunt lumped capacitance. A series resistance completed the model (a so-called RC model) but this was small because of the relatively large cross-sections of the interconnect.

This situation has changed because of three main developments, primarily for digital circuits but affecting analogue circuitry because of the rise of mixed signal systems.

The main developments are

(a) Faster clocks, of a gigahertz and above.
(b) Longer interconnects as the lateral dimensions of large chips are around 2cm.
(c) Fine lithography enabling interconnects to have cross-sectional dimensions of less than a micron.

The consequence of these developments are that RC modelling is not always adequate. However on-chip digital interconnects are highly irregular and densely packed and often do not even approximate interconnections of uniform transmission line segments.

The typical arrangement of on-chip interconnects is shown in figure 1.

The first metal layer, here, is tungsten and provides local interconnect with short interconnect lengths of small cross-section. Higher level interconnects use lower resistivity metal, here copper, and have larger cross sections thus further reducing line resistance. Greater attention is made to the provision of signal return paths at the higher metal levels. The higher level metal lines also have lower coupling to the silicon which, having finite conductivity (but not very high) is a significant source of loss.

Copper diffuses into silicon oxide (SiO_2) and into silicon (Si) with disastrous effects so that barriers are required. Currently titanium nitride (TiN) and tantalum (Ta) are used [4].

The provision of signal return paths for digital and mixed signal chips has only been a recent consideration and then only for a few nets. The exceptions are clock distribution nets and other long nets (such as certain data buses) where signal integrity is paramount. Another important characteristic was that little attention was given to providing signal return paths. In contrast, with RF and microwave chips it has always been considered necessary to provide current return paths so that the electromagnetic field produced by a signal on an interconnect is uniform contributing to well defined electrical characteristics, hence the term controlled impedance interconnect. Such an interconnect could be conveniently modelled as a transmission line. With RF and microwave chips there are relatively few active devices and so the provision of interconnects large enough (to effectively minimize interconnect resistance) and of defined ground planes for good current return paths can be accommodated at reasonable cost.

The characterization of and modelling required for these interconnects is a major topic. The main issues to be resolved are

(a) When is it necessary to model an interconnect as an RLC circuit?
(b) When is it necessary to use a transmission line model?
(c) What is the significance of the current return path?
(d) How do the characteristics of interconnects vary with frequency or clock frequency?

Types of on-chip interconnects

On-chip interconnects have several distinguishing characteristics that can be used in determining the type of modelling required.

It needs to be emphasized that an interconnect is not just a wire connection from one point in a circuit to another. There is always a current return path and this path (and its proximity to the signal path) is equally important in determining the electrical characteristics of the interconnect.

Electromagnetic fields between the two conductors making up a single interconnection between two points and anything that interferes with this field will have an effect on signals being transmitted. Two obvious physical conjectures establish that there must be a return path: without it charges would build up at some point in the circuit to very high levels; also the electric field supported by the charges (and induced by the voltage source originally) on the signal conductor of an interconnect must terminate on matching charges and these are located on the signal return path. A signal interconnect is usually a conductor which is constrained in its transverse dimensions but this is not the case with the signal return path. If a ground plane carries the return signal then the return path can be as wide as the ground plane at DC, but as the frequency of the signal increases the width of the current return path narrows as the charges on the signal return path have limited time to redistribute. The return path at high frequencies tends to be as close to the interconnect bearing the signal as possible. The proximity, and metal width available, can significantly affect the resistance, capacitance and inductance of the interconnect.

On-chip interconnects can be broadly classified by the type of signals conveyed, and geometrical and circuit characteristics as described below.

Interconnect networks (or nets) on digital chips can be categorised by the type of connections made:

■ Local connections with maximum lengths of 1-3mm.

The majority of connections on a chip are of this type. The utilization of space is a premium and so minimum line widths are used as dictated by the technology and these are in the range of 0.1 to 0.5µm, known as deep sub-micron. These lines have high resistance and are driven by transistors of minimum dimensions, and hence of low current drive capability, and have high output (or drive) impedance Z_{DRV}. It is not economical to provide metal for individual signal return paths.

■ Medium length connections with minimum lengths of 1-3mm

The maximum usable length is determined by the characteristics of the technology and of the interconnect itself (e.g. width and hence resistance). Typically the maximum usable length is around 2-5mm. The longer usable length is achieved principally by making the line wider but also by paying attention to the current return path. Higher current drivers with lower Z_{DRV} are required.

■ Long connections up to a chip-edge in length.

There are three families of connections that are in this category, each with different characteristics. However what they have in common is that the drivers of these connections must have high current capability, and low Z_{DRV} (at the end of the line) so that receivers switch correctly when the signal makes its first transition down the line. The families are:

Data buses: These are often wide (meaning that there are many parallel interconnections, perhaps 128 for example) and can be about half the chip-edge in length. Their purpose is to convey data from one part of a chip to another such as from the central processor to an on-chip data cache. Such lines switch at the same time and coupling of the lines is a considerable concern, also the delay on the lines is affected by the switching pattern. Another characteristic of these lines is that they have a small load at the end as they are often driving one or a few receivers.

Control lines: these transmit signals that must be widely distributed. A good example is a reset signal. There are relatively few of these types of lines with speed issues being less important.

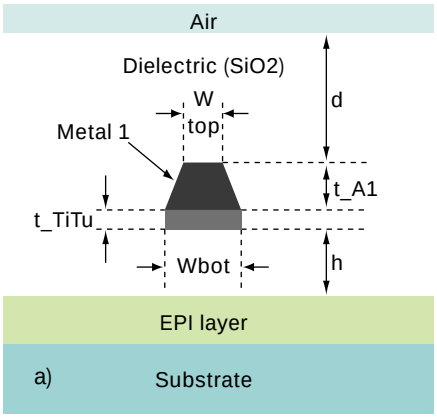

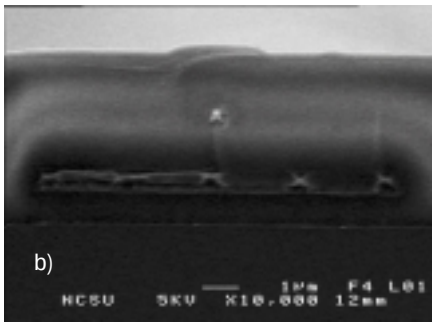

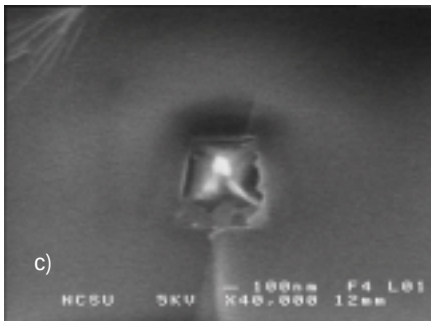

Figure 2: Cross-section of a deep-submicron interconnect: (a) line-diagram of an interconnect with the EPI layer acting as the ground (b) SEM cross-section of a two metal layer line; and (c) close-in SEM cross-section of a top metal line.

Clock signal: this must be very widely distributed with integrity of the signal being paramount. As well as ensuring first incidence switching the clock skew, that is the difference between the arrival times of the clock edges in different parts of the chip, must be precisely controlled. The clock distribution net can consume a very large proportion of the chip power both because of the gates that are driven at the end of the clock distribution net and because of efforts required to minimize skew.

Many long nets have intermediate buffers along them to boost the signal to compensate for attenuation of the signal.

Power and ground distribution buses. It is perhaps unusual to think of the power and ground buses as interconnects, as the voltage on them is, hopefully, constant. However these buses carry large switching currents that change at the clocking frequency. The pulsing current transmits a signal just as real as does a voltage variation. With fast chips it is necessary to model and design these as high speed interconnections. Ideally these nets are of very low impedance so that the current fluctuations have little effect on voltage levels. Low impedance can be achieved using metal planes for supply and ground and by large capacitors between the supply and ground nets.

Interconnects for digital signals can be contrasted to interconnects for other types of signals. The density of analogue and RF circuits does not approach that of digital circuits and in any event controlling the characteristics of transistors (impedance levels, pole frequencies etc.) is critical to the performance of these frequency sensitive circuits.

Microwave and millimetre-wave on-chip circuits (usually called MMICs for Microwave Monolithic Integrated Circuits) usually have interconnections that are a substantial portion of a wavelength, but also have ground planes or similar structures forming a highly regular signal return path. Thus transmission line modelling is necessary and is also convenient.

Experimental characterization of an on-chip interconnect

Biswas et al. [5-6] experimentally characterized interconnects fabricated at SEMATECH using a deep submicron process. The cross-section of the interconnect is shown in figure 2. The epi-layer is highly doped and forms a ground plane. On this a layer of dielectric is grown and then a thin layer of Titanium Tungsten (TiW) is deposited before a thick layer of aluminium (Al). Dielectic is then deposited and then the surface is planarized (the CMP-process), so that the tops of the dielectric and metal interconnect form a plane. The process is repeated until finally a passivation layer is deposited.

The metalization consists of multiple metal types including, in addition to the basic aluminium interconnect metalization, buffer layers of, for example, titanium, and tungsten vias. Higher performance interconnects use copper as the main metalization but this must be coated with a buffer layer to prevent reaction of the copper with the dielectric. Further complicating our ability to model this structure are changes in the density of the dielectric, and thus the dielectric constant, around the metal. Biswas et al. used microwave frequency domain measurement techniques. Time domain techniques enable alternative experimental characterization using signals that are much closer to the types of signals encountered on a chip but without the dynamic range of frequency domain measurements [7-11]. The conductors were aluminium and a Chemical-Mechanical Planarization (CMP) process was used to ensure planarization - this is an important step when the aspect ratio (height-to-width) is appreciable. A particularity of the measurements is that it is not possible to extract all four frequency dependent parameters, either the four R, L, G, C parameters or the four components (i.e. the real and imaginary parts) of Z_0 and γ, from measurements. Generally it is assumed that the capacitance per unit length of the line does not vary much with frequency and this has been extensively investigated [12]. The characteristic impedance and propagation constant of an interconnect of nominal width 0.25µm is shown in figure 3 and its R and L parameters shown in figure 4.

The parameters are arrived at using two lengths of line of the same cross-section. The through line (TL) technique was used which involves just two calibration structures and reduces the test area required otherwise substantially [13-14]. This is recommended as a standard technique for the frequency domain characterization of on-chip interconnects.

In figures 3 and 4 two sets of results are presented - obtained by considering three lines so that two independent extractions are possible. (This was done for purposes of cross-checking results. Any discrepancies that were found were always traced back to contact problems as reliable probe contact up to 20GHz is a potential problem.)

The irregular results shown at low frequencies can be ignored as the lines of necessity must have short electrical lengths here and hence resolution is poor. However the complete electrical properties of the interconnects are contained in Z_0 and γ and these are smooth. The per-unit-capacitance, C, was obtained using separate low frequency measurements and was 827fF/cm. G was taken as negligible. One notable feature is the high resistance of the line due to the small cross-section.

RC modelling on-chip interconnects

The reason to model interconnects is to verify signal integrity and provide insight to allow redesign if necessary. The ultimate

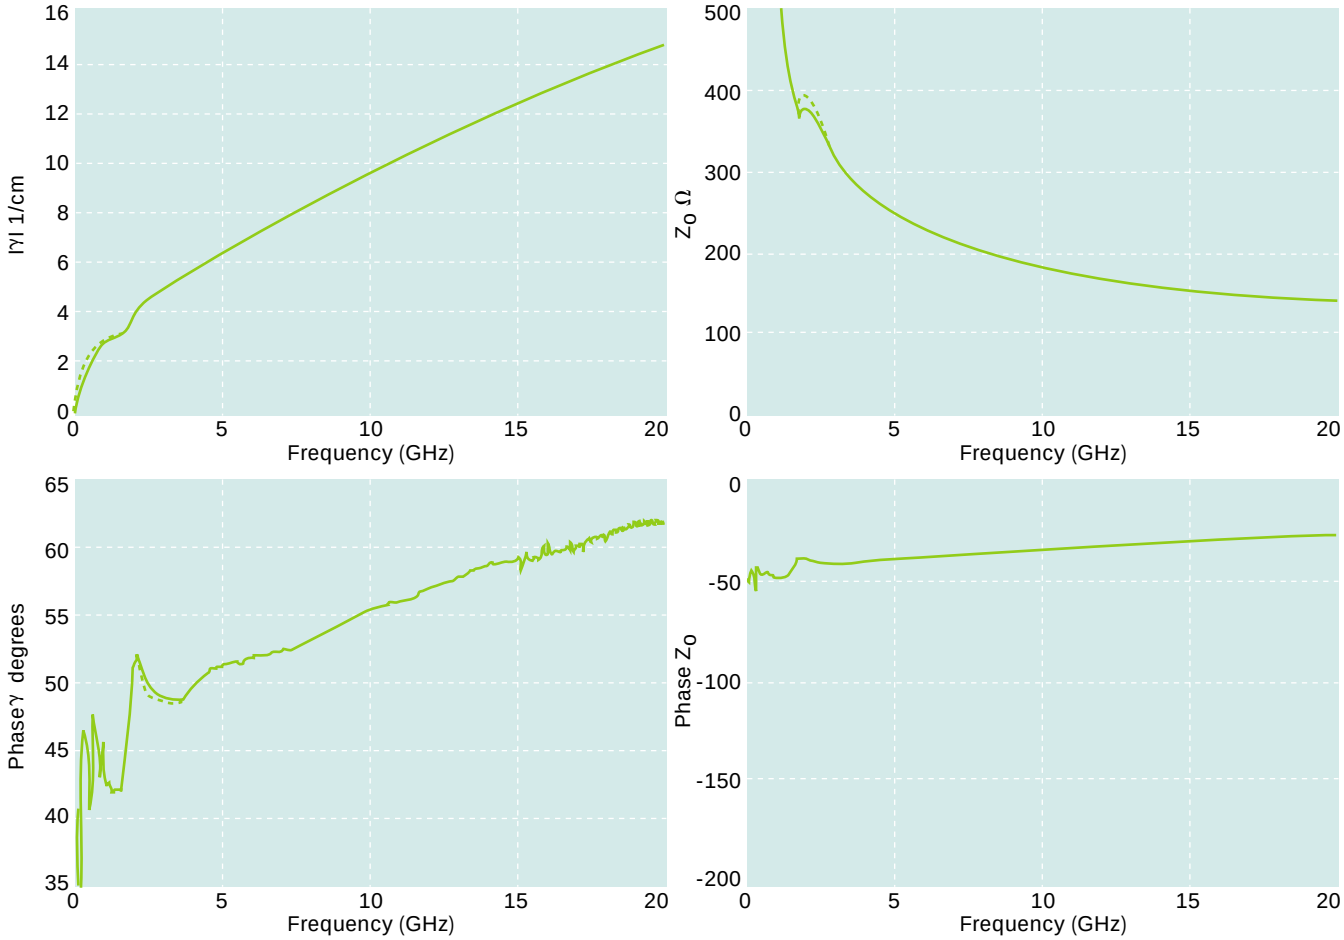

Figure 3: Experimentally derived characteristic propagation constant and characteristic impedance of a nominal 0.35µm wide interconnect: (a) magnitude and (b) phase of propagation constant; and (c) and (d) (a) magnitude and (b) phase of characteristic impedance.}

goal of course is to take electrical specifications and to develop the physical layout requirements but this is not as well developed as would be liked. The main signal integrity issues are

(a) Will the chip operate at the rated speed? Is the timing correct, that is, do the signals at the receiver end of a net arrive at the required switching levels within the timing budget?

(b) Is signal integrity adequate? That is, is crosstalk within acceptable levels?

(c) Will it perform reliably? Reliability includes electromigration and fuse effects as well as signal integrity issues.

Modelling on-chip interconnect effects is daunting as there are millions of nets on a typical submicron IC and the ultimate clock speed of a circuit is determined by the slowest of these interconnects and the nets that have the highest crosstalk. An idea of the complexity can be seen by examining the interconnects involved in a single CMOS inverter shown in figure 5.

It is clearly not possible to incorporate the interaction of every net on every other other one. So the signal integrity verification problem is driven by the need to minimize the modelling complexity as much as possible. This is done by first identifying those nets, the critical nets, that are most likely to cause a problem and to model these using simple modelling. These nets are not assured of actually being problems but if there is a problem we can have a very high confidence that the problem net will be one of the critical nets. Critical nets are identified using relatively loose criteria but still sophisticated enough that only a small fraction of the total number of nets are so identified. The critical nets are then examined in considerable detail to determine if there is a signal integrity violation.

Thus in a second pass the at risk nets are identified. Generally the at risk nets must be examined by a human to evaluate whether or not there is a signal integrity violation.

The simplest electrical model that will capture the effect is used, for example, in increasing complexity: a simple delay model; an RC model consisting of resistors and capacitors only; an RLC model consisting of inductors as well; and finally a distributed transmission line model. These types of models are discussed below.

Delay modelling

Nearly all of the drivers on a chip are sized near their minimum functional dimensions and so have limited current capability. The consequence of this is that a driver, driving a net, does not appear like a linear driver but instead as sourcing or sinking essentially a fixed current until the desired voltage level is obtained. Thus the voltage on a short interconnect only rises according to the rate at which charge is sourced or sunk from the line and the time required to charge to the desired voltage levels is proportional to the capacitance of the interconnect. Thus the delay on a line is directly related to the time constant of the line which is equal to the time for a single pole circuit, an RC circuit for example, to

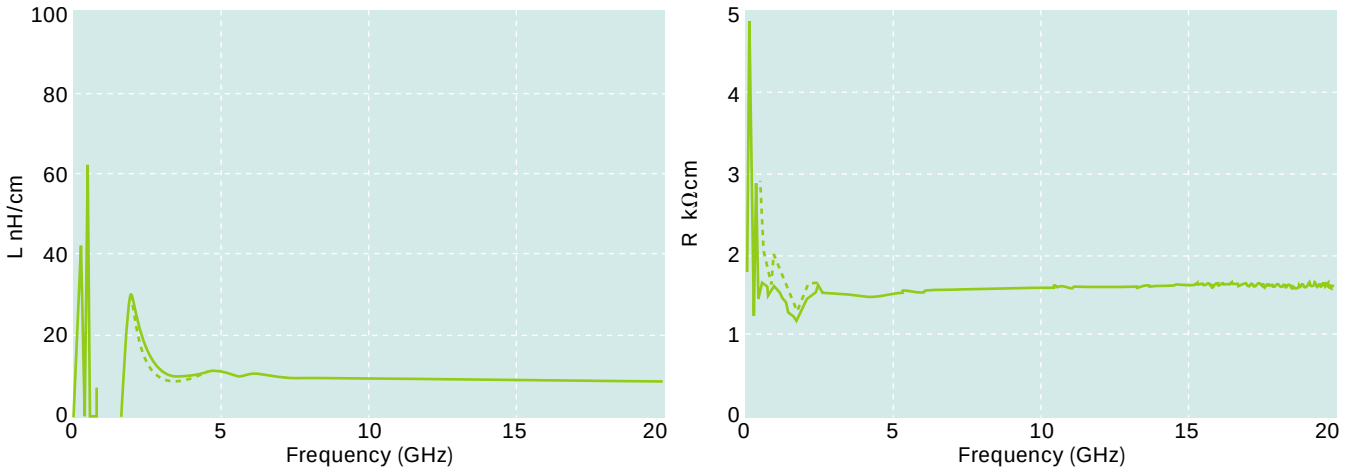

Figure 4: Experimentally derived inductance (a) and resistance (b) unit length parameters. The capacitance was 827 fF/cm.

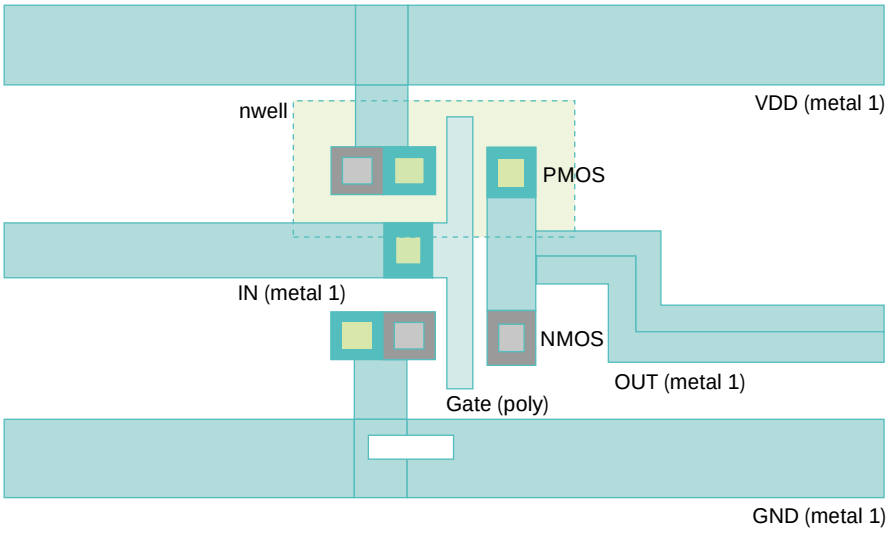

Figure 5: The layout of a CMOS inverter

change by (1-1/e)≈65% of the total change that will occur. Typically this is about the voltage change required to alter the logic state of a digital circuit. The total effective time constant, τ_T, of a series cascaded circuit with K poles with time constant τ_k for the k th pole is,

$$\tau_T = \sum_{k=1}^{K} \tau_k \qquad (1)$$

where the k th time constant

$$\tau_k = R_k C_k \qquad (2)$$

R_k is the resistance of the k th segment of interconnect, and C_k is the capacitance of the segment. So that an estimate for the delay t_d through such a circuit is

$$t_d = \tau_T = \sum_{k=1}^{K} \tau_k \qquad (3)$$

Thus a good delay model of the cascaded circuit is

$$t_d = \sum_{k=1}^{K} R_k C_k \qquad (4)$$

Equation 1 gives us only an estimate of the effective overall time constant of a circuit and there are restrictions on its validity. The primary restriction is that it is useful when the individual time constants τ_k are comparable. However it provides the best simple estimate of timing.

RC modelling

The first step in developing an RC model is to develop the capacitance model of a subset of nets. As the density of nets is very large it is not sufficient to consider a net in isolation. Neither is it feasible to consider the interaction of a net with every other net on the chip. A successful approach is to focus on the signal integrity of one net due to both the delay effects and to crosstalk from other nets. Fortunately the electromagnetic properties help us here. A voltage and current carrying interconnect generates electric and magnetic, or electromagnetic, fields which fall off with distance as $1/R^N$ where R is distance and N≥1. (N does not fall off much faster as the electric and magnetic fields are essentially confined in a plane.) This is the justification for considering a region of influence such as that shown in figure 6. Here the region of influence (the box) is centred on an individual net, the victim net. It is assumed that every conductor within the region of influence will be coupled to the victim net. The capacitances describing this coupling are generally developed using a number of template models. For example if the victim net and another net cross each other (the nets having X- and Y-orientations respectively) a parallel plate capacitor model is used for the capacitance.

The templates of a large number of simple structures that can commonly occur are generally developed in advance for a particular process. The templates could consist of precise data, each template developed for a structure of a particular geometry, and a very large number of templates are needed. Alternatively fewer templates of the polynomial type with geometry variables such as separation distance could be used and the coefficients of the polynomials developed by fitting to electromagnetic simulation of several variants of a particular template structure. With the multiple metal composition of interconnects and fabrication-dependent dielectric density variations, these templates should be calibrated with actual measurements. Pattern recognition is used to identify candidate structures (that match a template). If the two nets run par-

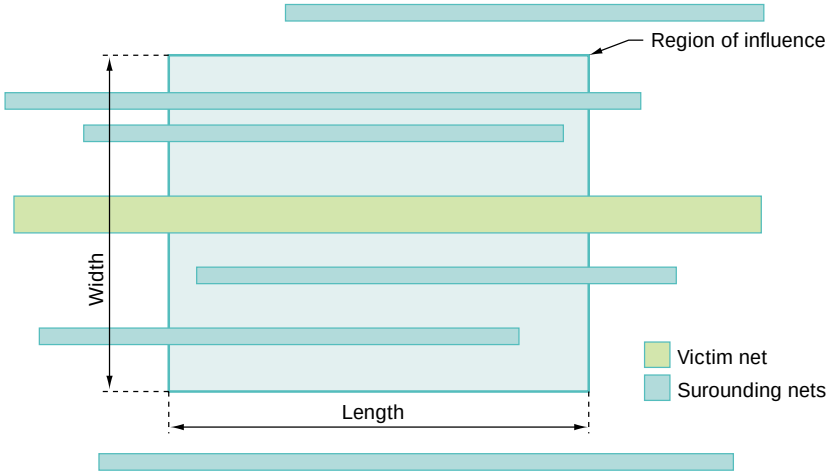

Figure 6: Neighbouring nets and the region of influence used in capacitance extraction

allel to each other then a standard capacitance calculation is performed — again essentially another parallel plate capacitance calculation but using fringing capacitance corrections. These capacitance templates are typically developed using three dimensional electromagnetic simulation codes to build up a standard library.

Modelling inductance

Modelling the inductance of long high performance digital interconnects is now required. Nets of this type are relatively few but they include the on-chip clock distribution net. For these nets, LRC modelling (an excellent approximation of a transmission line as far as digital inteconnects are concerned) must be used. The aggressive nets are either constructed as signals lines over a ground plane (a microstrip structure) or with shield lines between neighbouring signal lines to form a coplanar waveguide-like structure. In both cases the aim is to achieve a well controlled signal return line. It is possible to achieve very high performance from nets without a reference conductor following in tandem with each signal net, but this requires extremely good design discipline and model extraction tools. For structures with reference conductors, extracting inductances is much easier than it would be otherwise. Complicated structures without the reference conductors require full-wave three-dimensional electromagnetic field solutions [15].

'Foundations of Interconnect & Microstrip Design - 3rd Edition' (ISBN 0471607010) can be purchased at a pre-publication offer price of £44.00 through the MEE online bookshop at www.mwee.com

References

1] Bakoglu, H. B., Circuits, Interconnections and Packaging for VLSI, Addison-Wesley Publishing Co., New York, 1990
2] Doane, D. A. and Franzon, P. D., Multichip Module Technologies and Alternatives: The Basics, Van Nostrand Reinhold, New York, 1993.
3] Harper, C. A., Electronic Packaging & Interconnection Handbook, Second Ed., McGraw-Hill Inc, New York, 1997.
4] Gutmann, R. J., Advanced silicon IC interconnect technology and design: present trends and RF wireless implications. IEEE trans, Microwave Theory and Techn., Volume 47, June 1998, pp 667-674.
5] Biswas, B., Glasser, A., Steer, M., Franzon, P., Griffis, D., and Russel, P., Experimental electrical characterization of on-chip interconnects, Proc. IEEE 6th Topical Meeting on Electrical Performance of Electronic Packaging, Oct 1997, pp 57-59.
6] Biswas, B., Modeling and Simulation of High Speed Interconnects, M.S. Thesis, North Carolina State Univ.1998.
7] Deutsch, A., Kopscay, G., Ranieri, V. A., Cataldo, J. K., Galligan, E. A., Graham, W. S., McGouey, R. P., Nunes, S. L., Paraszczak, Ritsko, J. J. Serino, R. J., Shih, D. Y., & Wilczynski, J. S., High speed signal propagation on lossy transmission lines, IBM J. Res Devel., Vol 34, 1990, pp. 601-615
8] Deutsch, A., Kopscay, G., Surovic, C. W., Rubin, B. J., Terman, L. M., Dunne, Jr., R. P. Gallo, T. A. and Dennard R. H., Modeling and characterization of long on-chip interconnections for high-performance micorprocessors, IBM J. Res. Devel, Vol. 39, 1996, pp. 547-567.
9] Jong, J-M., Janko, B. and Tripathi, V., Equivalent circuit modeling of interconnects from time-domain measurements, IEEE Trans. Components, Hybids and Manufacturing Techn, Vol 16, Feb 1993, pp.119-126.
10] Restle, P. J., Measurement and modeling of on-chip transmission line effects in a 400MHz microprocessor, IEEE J. of Solid-State Circuits, Vol. 33, April 1998, pp 662-665.
11] Deutsch, A., Electrical characteristics of interconnections for high-performance systems, Proc. IEEE, Vol. 86, Feb 1998, pp 315-355.
12] Engen, G. F. and Hoer, C. A., Thru-reflect-line: an improved technique for calibrating the dual six-port automatic network analyser, IEEE Trans on Micorwave Theory and Tech, Vol 27, Dec 1979 pp. 987-993.
13] Steer, M.B., Goldberg, S., Rinne, G., Franzon, P. D., Turlik, I., and Kasten, J. S., Introducing the through-line deembedding procedure, 1992 IEEE MTT-S Int. Microwave Symp. Dig., June 1992, pp.1455-1458.
14] Goldberg, S. B., Steer, M.B., Franzon, P.D. and Kasten, J.S., Experimental electrical characterization of interconnects and discontinuites in high speed digital systems, IEEE trans. on Components Hybrids and Manufacturing Techn. Vol 14, Dec 1991 pp761-765.
15] Lipa, S., Steer, M. B., Cangellaris, A. C. and Franzon, P. D., Experimental characterization fo transmission lines in thin-film multichip modules, IEEE trans. Components, Packaging and Manufacturing Techn., Vol. 19, Mar. 1996, pp 122-126.

Matched interconnect design using ground-surrounded microstrip line. Access Full Text. Matched interconnect design using ground-surrounded microstrip line. Author(s): M. Khalaj-Amirhosseini and A. Cheldavi.Â This structure is a microstrip line surrounded by two strips grounded at some points along their lengths. This structure will be used to design matched Gâ€“Sâ€“G (groundâ€“signalâ€“ground) type interconnects (maximum delivered power to the main load in a wide frequency band). Inspec keywords: integrated circuit interconnections; microstrip lines.Â T. Edwards . (1992) Foundations for microstrip circuit design. 8). D. Homentcovschi , R. Oprea . Analytically determined quasi-static parameters of shielded or open multiconductor microstrip lines. IEEE Trans. Microw.