



William D. Mack (S'75-M'79) was born in San Francisco, CA, on December 12, 1955. He received the B.S. and M.S. degrees in electrical engineering from the University of California, Berkeley, in 1977 and 1979, respectively.

He was a Teaching Assistant at U.C. Berkeley in 1978 and a Teaching Associate during 1979. Summer positions included the Space Sciences Laboratory at U.C. Berkeley, Ampex Corporation, Redwood City, CA, and the Tektronix IC Design Group, Beaverton, OR. In July 1979, he joined the Signetics Analog Circuit Research Group, Sunnyvale, CA. He is presently engaged in the design and characterization of high speed and high accuracy bipolar integrated circuits.

Mr. Mack is a member of HKN.



Mark Horowitz (S'78) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978. He is currently with the Department of Electrical Engineering at Stanford University, Palo Alto, CA, where his Ph.D. research is on computer aids for integrated circuit design.

Since 1979, he has been employed part time by Philips Research Labs, Sunnyvale, CA, working on device modeling and digital circuit design.

Mr. Horowitz is a member of Eta Kappa Nu and Sigma Xi.

Correspondence

Minimum Propagation Delays in VLSI

CARVER MEAD AND MARTIN REM

Abstract—Conditions are outlined under which propagation delays in VLSI circuits can be achieved that are logarithmic in the wire lengths. These conditions are imposed by area requirements and the velocity of light.

I. INTRODUCTION

With feature sizes decreasing and chip area increasing it becomes more and more time consuming to transport signals over long distances across the chip [5]. Designers are already introducing more levels of metal connections using wider and thicker paths for longer distances. Another recent development is the introduction of an additional level of connections between the chip and the printed circuit board; multilayer ceramic chip carriers. The trend is undoubtedly towards even more connecting levels.

In this paper we demonstrate that it is possible to achieve propagation delays that are logarithmic in the lengths of the wires, provided the connection pattern is designed to meet rather strong constraints. These constraints are, in effect, satisfied only by connection patterns that exhibit a hierarchical structure. We also show that even at the ultimate physical limits of the technology the propagation for reasonably sized

Manuscript received July 15, 1981. This work was supported by the Office of Naval Research under Contract N00014-76-C-0367 and by the Defense Advanced Research Agency under ARPA order number 3771, and monitored by the Office of Naval Research under Contract N00014-79-C-0597.

C. Mead is with the Department of Computer Science, California Institute of Technology, Pasadena, CA 91125.

M. Rem is with the Eindhoven University of Technology, Eindhoven, The Netherlands, and the Department of Computer Science, California Institute of Technology, Pasadena, CA 91125.

VLSI chips is dominated by these considerations rather than by the velocity of light.

II. PROPAGATION DELAY

We compute the time it takes a minimum-sized transistor to drive a wire of length l with width s . We assume the wire to have a distance s to its neighboring wires and a thickness and spacing to the layer beneath it of σs where σ is a constant smaller than one. This model is representative of current technology at all levels from the smallest transistors through printed circuit boards. Let s_0 be the minimal width of a wire on the chip so that a minimal transistor has area s_0^2 .

The following equation is a good approximation of the total time T required to drive the wire:

$$T \approx (R_t + R_w) C_w \quad (1)$$

where R_t is the resistance of the minimal transistor, R_w is the resistance of the wire, and C_w is its capacitance. The resistance of a wire is proportional to its length and inversely proportional to its cross section

$$R_w = \rho \frac{l}{\sigma s^2} \quad (2)$$

The capacitance of a wire is inversely proportional to the distance to its underlying layer and it is proportional to the area of the side facing that layer:

$$C_w \approx \epsilon \frac{sl}{\sigma s} = \frac{\epsilon l}{\sigma} \quad (3)$$

We note that the product of R_w and C_w is already quadratic in l . Thus the time it takes to drive a wire is at least quadratic in the wire length. However, things are not as bad as they look: R_t , the resistance of a minimal transistor, is the dominant term in (1). We can decrease that term by fitting a larger driver to the wire. But that driver must then in its turn be charged by the transistor driving it. The optimal arrangement is the well-known exponential sequence of drivers. The first

one is the minimal transistor; the next one is bigger by a factor α . It drives another driver that is again bigger by a factor α , etc., until we finally reach a driver that is large enough to drive the whole wire in a sufficiently short time.

There exists a simple rule to determine the time required to have a driver charge another driver [2]. Let τ be the time it takes a minimal transistor to charge the gate of another minimal transistor. The time required to have a driver with capacitance C_1 drive another driver with capacitance C_2 is

$$\tau \frac{C_2}{C_1}. \quad (4)$$

Let C_t be the capacitance of a minimal transistor. We have it drive a driver with capacitance αC_t . This second one drives a driver with capacitance $\alpha^2 C_t$, etc., until the last driver has a gate capacitance of about C_w/α . The number of drivers (including the initial transistor) required is

$$\log_{\alpha} \frac{C_w}{C_t}. \quad (5)$$

The capacitance C_t of a minimal transistor is equal to $(\epsilon s_0^2)/d$ in which d is the thickness of the gate insulator. The number of drivers is then $\log_{\alpha} ld/s_0^2 \sigma$ and we get for the time T_d spent in driving a zero resistance wire through the sequence of drivers

$$T_d = \alpha \tau \log_{\alpha} \frac{ld}{s_0^2 \sigma}. \quad (6)$$

We may replace (1) by

$$T \approx T_d + R_w C_w. \quad (7)$$

From (2), (3), (6), and (7) we conclude

$$T \approx \alpha \tau \log_{\alpha} \frac{ld}{s_0^2 \sigma} + \rho \epsilon \frac{l^2}{s^2 \sigma^2}. \quad (8)$$

We now have a formula for the propagation delay with both a logarithmic and quadratic term. One can see why a longer wire requires a larger s : that decreases the quadratic term. Actually, we wish to restrict the lengths of wires to values of l that are sufficiently small to assure that the quadratic term does not dominate. We therefore restrict ourselves to values of l for which the quadratic term is smaller than the logarithmic one. If a signal must go a distance l we chose a path of width s and thickness σs such that

$$\frac{l^2}{s^2} \lesssim \frac{\sigma^2}{\rho \epsilon} \alpha \tau \log_{\alpha} \frac{ld}{\sigma s_0^2}. \quad (9)$$

With this choice we assure that the total delay will never be larger than twice that required if there were no wire delay

$$T_d < T < 2T_d. \quad (10)$$

We have assumed that the values of s could be chosen from a continuous range. Although this is a good conceptualization of the increasing number of different connection layers, in practice we will have to choose s from a discrete set. The connecting wires will be placed at different levels. The widths of the paths at the next level will be some factor β times the widths at the preceding level. Given a distance l the signal has to travel (9) gives us the ideal s and we choose the next level at which the widths of the wires are larger than s . This leads to an interesting observation, the "magnifying glass phenomenon." Not only will the widths of the wires at any given level be the same but their lengths will also be about equal. The patterns at different levels are similar; at the next level the features are just magnified by a factor β .

Velocity of Light

Asymptotically, no signal can travel faster than the velocity of light. We must ask under what conditions the above considerations will set a limit which is more stringent, i.e., when the velocity of light is not attainable. In (6) and (10) we can substitute $\tau = s_0/v$ where v is the limiting velocity of electrons in the channel (a few 10^6 cm/s in silicon)

$$T < \frac{2\alpha s_0}{v} \log_{\alpha} \frac{ld}{s_0^2 \sigma}. \quad (11)$$

The minimum additional time ΔT required to propagate a signal an additional distance Δl is

$$\frac{\Delta T}{\Delta l} \approx \frac{dT}{dl} \geq \frac{\alpha s_0}{v l \ln \alpha}. \quad (12)$$

The domain of validity of the above results is $\Delta l/\Delta T < c$ in which c is the velocity of light in SiO_2

$$l < \frac{c \alpha s_0}{v \ln \alpha}. \quad (13)$$

For typical technology today, $s_0 = 4 \mu$, $\alpha/\ln \alpha$ is about 6, and l should be less than about 10 cm. Off chip the velocity of light can be higher, however, and the approximation is good to about a foot. Hence the velocity of light cannot be reached using the best MOS technology in the most optimal way within a typical small card bay but will be important at larger dimensions. Even for the ultimate technology ($s_0 = 0.25 \mu$), the results given above will dominate over velocity-of-light considerations for chips up to about a centimeter across.

III. AREA

The arrangements outlined in the preceding section allowing us to treat propagation delays as being logarithmic will only work if we can allot enough area at the lowest level for the drivers and at the higher levels for the wires.

A minimal transistor has area s_0^2 . The next driver in the sequence requires an area αs_0^2 , the third one $\alpha^2 s_0^2$, etc. The total area A of the drivers thus becomes

$$A \approx s_0^2 (1 + \alpha + \alpha^2 + \dots) (\log_{\alpha} l \text{ terms}) \quad (14)$$

$$A \approx \frac{s_0^2 (l - 1)}{\alpha - 1}, \quad (15)$$

or approximately

$$A \approx \frac{s_0^2 l}{\alpha - 1}. \quad (16)$$

Notice that we can trade area for time. By increasing α the area of the drivers decreases, cf. (16), but the propagation delay increases, cf. (8).

A transistor that has to drive a wire of length l requires area $s_0^2 l / (\alpha - 1)$ at the lowest level. This area is proportional to the length of the wire. That is fortunate; if we double both the length and the width of a chip we also double the lengths of the longest (cross-chip) wires and the areas of their drivers. But the total area of the chip will quadruple and we will thus be able to double the number of wires as well.

The longer wires come on higher levels on which the wires are wider thereby consuming more area. Each level, however, has the same total area. As a result, we can accommodate the wires at the higher levels only if we do not have too many of them. Assume again that at the next level the wires are β times thicker, longer, and wider. Call the lowest level number 0 and let N_i be the number of wires at level i ($i \geq 0$); then we must have

$$N_i < N_0 \beta^{-2i}. \quad (17)$$

The number of wires as a function of their lengths must decrease exponentially fast. This is a strong restriction. It suggests that efficient chips must have a hierarchical structure [2], [4]. If a design does not meet this exponential rule the best we can achieve is a propagation delay linear in the wire length by inserting repeaters at equidistant positions along the wires. The consequences of linear wire delays are discussed in [1].

One may also see complexity computations that assume that wires have no delay. Thompson, e.g., writes in [6]

The propagation time can be made independent of the length of the wire, by fitting larger drivers to longer wires. Larger drivers of course occupy more area, but need not take more than 10 percent of the area of the wire they drive. By fudging λ upwards by 5 percent, the area of the driver is thus absorbed into the area of its wire.

We have seen that the area of the driver is indeed proportional to the wire length but Thompson neglects the fact that charging the gate of the larger driver will also take time. Our choice of the sequences of exponentially growing drivers allowed us to do this in a time that is logarithmic in the wire length, a technique that can work only if we have very few long wires. Thompson's model also neglects the fact that the drivers have to be at the lowest level, in polysilicon and diffusion, independent of the level of the wire.

REFERENCES

- [1] B. Chazelle and L. Monier, "Optimality in VLSI," in *VLSI 81: Very Large Scale Integration*, J. P. Gray, Ed. London: Academic, 1981, pp. 269-278.
- [2] C. Mead and L. Conway, *Introduction to VLSI Systems*. Reading, MA: Addison-Wesley, 1980.
- [3] C. Mead and M. Rem, "Cost and performance of VLSI computing structures," *IEEE J. Solid-State Circuits*, vol. SC-14, no. 2, pp. 455-462, Apr. 1979.
- [4] M. Rem, "Mathematical aspects of VLSI design," in *Proc. Caltech Conf. VLSI*, C. L. Seitz, Ed., Dep. Comput. Sci., Calif. Inst. Technol., Pasadena, CA, Jan. 1980, pp. 55-64.
- [5] C. L. Seitz, "Self-timed VLSI systems," in *Proc. Caltech Conf. VLSI*, C. L. Seitz, Ed., Dep. Comput. Sci., Calif. Inst. Technol., Pasadena, CA, Jan. 1980, pp. 345-355.
- [6] C. D. Thompson, "Area-time complexity for VLSI," in *Proc. 11th Annu. ACM Symp. Theory of Computing*, ACM Special Interest Group on Automata and Computing Theory with IEEE Computer Society Technical Committee, Atlanta, GA, May 1979, pp. 81-88.

Switched-Capacitor Frequency Control Loop

T. R. VISWANATHAN, S. MURTUZA, V. H. SYED, J. BERRY,
AND M. STASZEL

Abstract—A novel approach to digital transduction of physical quantities using a switched-capacitor frequency control loop is presented. The operation of the loop as well as its properties are explained. A few applications are outlined with supporting experimental results.

Manuscript received May 4, 1981; revised November 17, 1981.

T. R. Viswanathan is with the Department of Electrical Engineering, Carnegie-Mellon University, Pittsburgh, PA 15213.

S. Murtuza, V. H. Syed, J. Berry, and M. Staszl are with the Department of Electrical Engineering, University of Michigan-Dearborn, Dearborn, MI 48128.

I. INTRODUCTION

Digital transduction of physical quantities is extremely important to fully exploit the microprocessor evolution for real-time process control applications or robotics. The modern trend [1] is to integrate the transducer and the associated signal processing circuits on a single chip of silicon along with the microprocessor.

In many transducers, the physical quantity to be measured causes a change in the value of a resistor or a capacitor giving an analog electrical signal which is sampled and digitized using an analog-to-digital converter. The switched-capacitor frequency control loop provides a unified approach to obtain a digital output proportional to either resistance R or capacitance C by generating a periodic signal whose period is equal to the product RC . The loop locks and tracks the value of the circuit element by adjusting the period which is easily digitized by means of a precise clock. The dynamics of the tracking process is fast enough for a wide range of industrial applications.

II. THE FREQUENCY CONTROL LOOP

In Fig. 1(a) the resistor R and the switched-capacitor (SWC) C operate in parallel between the reference source V_R and the node P of the operational amplifier (op amp). The op amp in conjunction with C_2 forms an integrator. The output voltage (v) of the op amp is fed to a loop filter which in turn drives a voltage-controlled oscillator (VCO). A nonoverlapping two-phase clock is derived from the output of the VCO.

Switches S_1 , S'_1 , S_2 , and S'_2 are open when the clock is low and they close when the clock goes high. S_1 and S'_1 are operated by clock phase ϕ_1 and S_2 and S'_2 are operated by phase ϕ_2 . This is a well-known SWC arrangement [2] which is insensitive to the parasitic capacitances.

Assuming that R and C are fixed, a current $I = (V_R/R)$ flows into the inverting input of the op amp which is at virtual ground. The corresponding integrator output $v(t)$ will be a negative going ramp voltage. This voltage, filtered through the loop filter, controls the frequency of the VCO. In VCO's based on relaxation oscillators, frequency f is related to the average value \bar{v} of the input voltage $v(t)$ by the relationship $f = f^* [1 - (\bar{v}/V^*)]$ where f^* and V^* can be preprogrammed. As \bar{v} decreases, f increases and the corresponding period ($T \triangleq 1/f$) decreases. Since the switch pairs (S_1, S'_1) and (S_2, S'_2) operate once for each period of the VCO, a quantum of charge CV_R is removed from the input of the integrator in each period T . This causes the output of the integrator to jump up by a step of magnitude equal to $V_R(C/C_2)$. Thus in the steady state, $v(t)$ will be a sawtooth with constant period. If the continuous charge supply into the integrator exceeds the discrete charge removal, \bar{v} decreases. This in turn reduces T , thereby speeding up the charge removal process. Thus there is a negative feedback in the loop. If the free-running frequency f^* of VCO is chosen properly, the loop locks and adjusts f such that the charge fed into the integrator during a period T through the resistor equals the charge CV_R removed by the SWC in the same period giving the charge balance condition $(V_R/R)T = CV_R$ or $T = CR$. Thus T is precisely equal to the time constant CR because any small charge unbalance would vary \bar{v} which will, in turn, readjust the value of T to restore the balance condition. Any soft nonlinearity or temperature dependence in the input-output relationship of the VCO does not affect the operation of the loop. In essence, the loop matches with precision a continuous process with a discrete process. The use of the SWC in the inverting mode enables the use of a single reference voltage which does not appear in the charge balance condition. In its simplest form the loop works without any loop filter, i.e., the output of the integrator feeds the VCO.