

NOSC

NAVAL OCEAN SYSTEMS CENTER San Diego, California 92152-5000

Technical Document 1114
July 1987

Ultra-Low Power Digital Filter

J. G. Nash
G. R. Nudd
Hughes Research Laboratories



Approved for public release,
distribution is unlimited

The views and conclusions contained in this report are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Naval Ocean Systems Center or the U.S. government

NAVAL OCEAN SYSTEMS CENTER
San Diego, California 92152-5000

E. G. SCHWEIZER, CAPT, USN
Commander

R. M. HILLYER
Technical Director

ADMINISTRATIVE INFORMATION

This report was prepared by Hughes Research Laboratories under contract N66001-82-C-0504 for Code 743 of Naval Ocean Systems Center, San Diego, CA 92152-5000.

Released by
J.M. Alsup, Head
Image Processing
and Display Branch

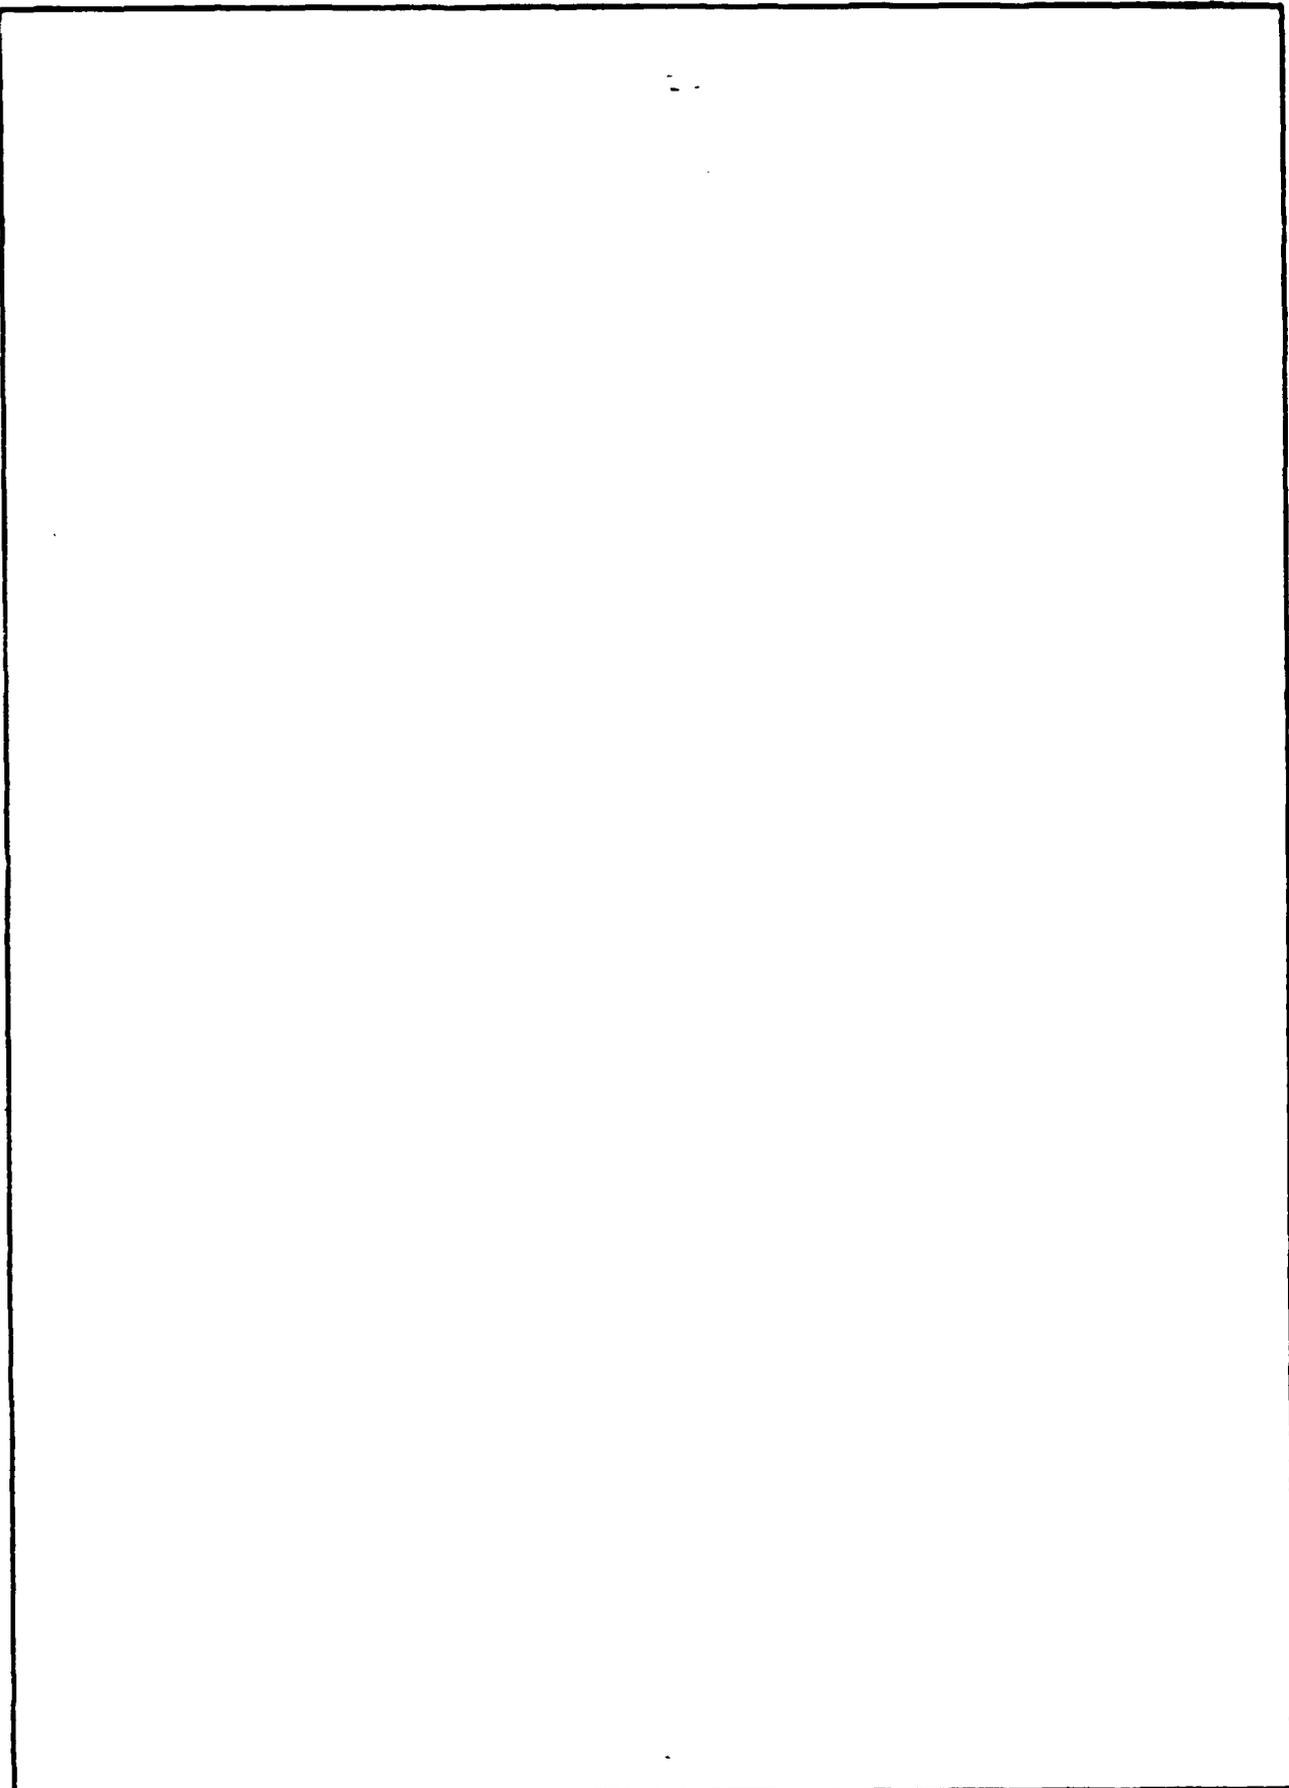
Under authority of
R.L. Petty, Head
Electromagnetic Systems
and Technology Division

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b RESTRICTIVE MARKINGS	
2a SECURITY CLASSIFICATION AUTHORITY		3 DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.	
2b DECLASSIFICATION/DOWNGRADING SCHEDULE			
4 PERFORMING ORGANIZATION REPORT NUMBER(S)		5 MONITORING ORGANIZATION REPORT NUMBER(S) NOSC TD 1114	
6a NAME OF PERFORMING ORGANIZATION Hughes Research Laboratories	6b OFFICE SYMBOL <i>(if applicable)</i>	7a NAME OF MONITORING ORGANIZATION Naval Ocean Systems Center	
6c ADDRESS (City, State and ZIP Code) 3011 Canyon Road Malibu, CA 90265		7b ADDRESS (City, State and ZIP Code) San Diego, CA 92152-5000	
8a NAME OF FUNDING/SPONSORING ORGANIZATION Department of Defense	8b OFFICE SYMBOL <i>(if applicable)</i> DoD	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N66001-82-C-0504	
8c ADDRESS (City, State and ZIP Code) 9800 Savage Road Ft. Meade, MD 20755		10 SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO RDDA	PROJECT NO EE93
		TASK NO	AGENCY ACCESSION NO DN488 839
11 TITLE (Include Security Classification) Ultra-Low Power Digital Filter			
12 PERSONAL AUTHOR(S) J.G. Nash, G.R. Nudd			
13a TYPE OF REPORT Final	13b TIME COVERED FROM Oct 83 TO May 83	14 DATE OF REPORT (Year, Month, Day) July 1987	15 PAGE COUNT
16 SUPPLEMENTARY NOTATION			
17 COSATI CODES		18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
			finite impulse response (FIR) adder accumulator chip-to-chip communication multiplier
19 ABSTRACT (Continue on reverse if necessary and identify by block number) This report describes a study of techniques for building an ultra-low power digital filter.			
20 DISTRIBUTION AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21 ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a NAME OF RESPONSIBLE INDIVIDUAL W. McKnight		22b TELEPHONE (include Area Code) 619-225-7439	22c OFFICE SYMBOL Code 743

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)



UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

TABLE OF CONTENTS

I. Report Synopsis	
A. Introduction	1
B. Summary	1
C. Conclusion	8
II. Detailed Description of FIR Filter	
A. Introduction	10
B. Filter Organization	12
C. Coefficient and Data Storage Schemes	14
D. Multiplier/Accumulator	17
E. Calculation of Total Power Consumption	26
F. Technology Issues	27
G. Implementation With CCDs	27

I. REPORT SYNOPSIS

I-A. Introduction

In this report we describe a study of techniques for building an ultra-low power digital filter. The basic goal was to design a 1024 equivalent tap filter, which is programmable, linear phase, operates at a sample rate of 8KHz, and consumes a maximum of 1.4mA at 3.6V (5mW). Input data word length was specified as between 8 and 12 bits. It was also desirable that the circuit be a single chip implementation. Any clock, control, or refresh circuitry was to be included on chip and be counted as part of the power budget.

In Section I-B. of this report we present a brief summary of the technical work performed and in Section I-C conclusions are drawn from the data presented. Section II contains the technical supporting details.

I-B. Summary

The filter organization chosen for investigation was a direct 1024 tap FIR (finite impulse response) filter implemented as a single tap operating at 1024 times the 8KHz sampling rate. The choice of a FIR filter provided the required linear phase for the unspecified transform coefficients. In addition the single tap approach provides a more area and power efficient realization of a 1024 tap filter than numerous taps operating at lower speeds.

In order to determine power consumption, the FIR filter was decomposed into its component cells shown in Table I. The parameters b and s correspond to the word length and the feature size. The three main parts of the filter were the memory (data and coefficient storage), the arithmetic section and the

CELL	POWER (mW)
Shift Register	8.7s 10 ⁻³
Full Adder	2.5s 10 ⁻²
Select/Clear	2.1s 10 ⁻²
Select Line Driver	7.7bs 10 ⁻⁴
RAM	3.4s 10 ⁻³

Table I. Power Consumption Associated with Important Cells in FIR Filter.

control section. The power associated with each of these parts is categorized in Table II, based on the power calculations of Table I. The control section was not included because it did not contribute substantially to the power budget.

The technology used as a basis for the calculations was CMOS/SOS, which has the unique characteristic that power consumption tends to be very low due to the lack of substrate parasitics and low DC power consumption. For this technology we assumed that all power consumption, P , was dynamic switching power or

$$P = C V^2 f$$

where C is the circuit capacitance, V is the supply voltage and f is the switching frequency. In all the calculations the supply voltage was taken as 3.6V and the frequency f was assumed to be 8MHz. The capacitance was calculated by adding the gate capacitances for the particular circuit being analyzed. The calculations were done assuming a nominal feature size of 5 microns, but a scaling factor s ($s=1.0$ for 5 micron features) was included to estimate power consumption for feature sizes less than 5 microns, e.g., $s=0.5$ for 2.5 micron features.

Note that in estimating power consumption by the formula above, parasitics were neglected. We feel that this is an adequate approximation for the circuits considered at feature sizes of 5 microns. For one circuit (full adder), we laid out the circuit (approximately 30 devices) and calculated the expected parasitics associated with layer to layer overlaps. The added parasitics per device were less than 10% of the gate capacitance. However, at reduced feature sizes intralayer capacitances would become important. Further

FUNCTIONAL BLOCK	NUMBER OF CELLS	TYPE CELL					POWER (mW)	
		SR	FA	S	SD	RAM	RAM	SR
Data Storage	1024b	X				X	3.4bs	8.7bs
Coefficient Storage	512b	X				X	1.8bs	4.3bs
Multiplier	$b(\frac{b}{2} + 1)$		X				$2.5b(\frac{b}{2} + 1)s 10^{-2}$	
	$b(b + \frac{1}{2})$	X					$8.7b(b + \frac{1}{2})s 10^{-3}$	
	$\frac{b}{2}(b + 3)$			X			$b(b + 3)s 10^{-2}$	
	$\frac{5b}{2}$				X		$1.9b^2s 10^{-3}$	
Adder/Acc	b		X				$2.5bs 10^{-2}$	
	48	X					0.42s	
	b	X					$8.7bs 10^{-3}$	

Table II. Breakdown of FIR Filter Power Consumption by Cells. Here SR, FA, S, SD, and RAM correspond to Shift Register, Full Adder, Select/Clear, Select Line Driver, and RAM cells, respectively.

study is required to estimate such effects.

The results of the power calculations are shown in Tables II and III. Two cases were considered: a filter with data and coefficient storage entirely in shift registers or entirely in static RAMs. Calculations for power associated with the RAM were based on a Hughes Newport Beach 16K static RAM.⁽¹⁾

The possibility also exists for alternate memory storage schemes, that offer the potential for considerably reduced power consumption. For example, the shift register for data storage can be built as two separate circuits each containing half the capacity and operating at half the speed. A multiplexer could be used to selectively obtain outputs from the two circuits so that the net data rate from the two circuits remained 8MHz. However, since each of the circuits is operating at half the original 8MHz speed, the shift register power consumption would be reduced by a factor of 2. Similarly, four separate shift registers, each operating at 2MHz, would consume a minimum of one-fourth the original power. Of course power and area overhead requirements would increase when this approach is used. It would be expected that there would be some point of optimal reduction in size of the shifter register storage. The same approach could be taken with the static RAM memory. However, we would expect that the additional penalties associated with increased power and area overhead would be much more severe for the RAM than the shift register.

The schedule for SOS technology development at Hughes Newport Beach is shown in Table IV. Here, we are using the VHSIC program as an estimate for when various feature sizes will be come available. Of important note is that

(1) A. Gupta, M.F. Li, K.K. Yu, S.C. Su, P. Pandya, and H.C. Yeh, "Radiation-Hard 16K CMOS/SOS Clocked Static RAM," Proc. International Conference on Electron Devices, Washington, DC, Dec. 1981, pp.616-619.

	b =	8	12	12	12
	s =	s	s	0.25 (VHSIC I)	0.14 (VHSIC II)
Power (mW)	RAM Approach	45s	70s	18	10
	Shift Register Approach	110s	160s	41	23

Table III. Total Power Consumption of FIR Filter for Various Values of b and s.

TIME SCHEDULE	
PROCESS	AVAILABILITY
5 μ	Now
2.5 (SOS II)	Will possibly leapfrog direct to VHSIC I
VHSIC I: 1.25 μ	May 1984
VHSIC II: 0.7 μ	May 1986

Table IV. Schedule for Progress in CMOS/SOS Technology Based on Hughes Participation in VHSIC Program.

the standard voltage level used in VHSIC I technology is 5V rather than the 3.6V used in all our calculations. It is not yet clear what voltage level will be associated with VHSIC II technology.

Since most of the power consumed in the filter is in memory storage and since the data is accessed serially, a CCD technology is also a possibility for building the filter. CCD's are well known for their low power and high packing density. In addition new advances have been made in building logic circuits using CCD's.⁽²⁾

For the memory sections alone we estimate that a CCD approach might provide a factor of 2 decrease in power consumption. Further study is necessary to determine overall savings in power using CCDs.

I-C. Conclusions

From the results shown in Tables II and III it can be seen that neither of the two FIR digital filter approaches, single shift register or single static RAM data and coefficient memory storage, meet the desired goal of 5mW power consumption. However, since our results indicate that most of the power for the filter is consumed in memory storage, rather than in the multiplier, it appears possible that alternate memory organizations offer potential for considerably less power consumption than shown in Tables II and III. With further study of possible memory organizations and their associated overhead, it would be possible to suggest an optimal design for the 1024 tap FIR filter in terms of power consumption.

The SOS technology progress being made as part of the VHSIC effort at

(2) J. Greg Nash, "Digital CCD Logic Circuits For Signal Processing," Proc. Government Microcircuit Applications Conf., Orlando, Fla., Nov. 2-4, 1982.

Hughes projects that 1.25 micron channel lengths will be available in May 1984 and that 0.5 to 0.7 micron channel lengths will be available in May 1986. Thus, the reduced power levels associated with these smaller feature sizes will further enhance the possibilities for meeting the 5mW power specification as shown in Table III. We note, however, that the 3.6V power supply is not necessarily compatible with either the VHSIC I or VHSIC II technologies. Further fabrication process analyses are necessary to evaluate the extent of this problem.

II. Detailed Description of FIR Filter

II-A. Introduction

In this section we will present a detailed technical description of the FIR filter on which we have based our power consumption numbers. We consider two schemes for organizing the data and coefficient storage. In addition we describe other alternate memory storage schemes which offer potential for considerably decreasing overall power consumption.

Our choice of an FIR approach to the the filter requirements was based primarily on the consideration that linear phase was required, regardless of the choice of coefficients. In other words, it would be difficult build a more power efficient IIR type filter and still meet the linear phase requirements, because the coefficients or applications of the filter are unspecified at this point.

Our approach to calculating power consumption was made considerably easier because CMOS/SOS technology has the unique property that circuits are built on an insulating substrate. As a result, ~~parasitic~~ parasitic capacitances associated with junctions and substrates are very small. We have laid out a full adder using a minimum gate configuration and calculated the parasitics based on layer to layer overlaps and junction capacitances. These calculations have shown that it is a reasonable assumption to neglect parasitic capacitances for circuits with 5um features and local interconnects (e.g., pipelined circuits). However, as feature sizes shrink and aspect ratios of interconnects increase, this assumption will become progressively less valid.

We have assumed that all power consumption is due to dynamic switching of

the gates on and off, or

$$P = C V^2 f,$$

where C is the capacitance being switched, V is the supply voltage (3.6V) and f is the switching frequency (8MHz). Standby power is generally in the microwatt range because there are no DC conducting paths. Power can then be calculated by multiplying the number of transistors in a logic stage by the capacitance per transistor gate to get the total capacitance that is switched.

All calculations were done at a nominal feature size of 5 μ m, so that the gate capacitance of a transistor is approximately 0.01pf. In order to account for smaller feature sizes, we have included a parameter s ($s=1.0$ for 5 μ m feature sizes), indicating how capacitances will scale as feature sizes are reduced. For the gate of a transistor, the area is reduced by s^2 , but the oxide thickness decreases as s , so that the actual gate capacitance is $0.01s$ pf. By feature size here we are referring to channel length. The physical gate length and channel length can become considerably different as dimensions shrink.

By using the gate oxide capacitance of transistors in a circuit as a measure of the total circuit capacitance being charged and discharged, we are overestimating the average power consumption. For example if the input to a circuit was the same every clock cycle, there would be no switching energy consumed.

Because we expect to overestimate the power based on the considerations mentioned above, we have neglected the power consumption associated with the control circuits and clock drivers. We expect the amount of control circuitry to be very small compared to the memory and arithmetic circuitry, and hence not

be an important factor in calculating power. There is little control necessary because of the regularity of data flow and the minimal amount of decision making required.

We have also neglected any power consumption associated with a multichip implementation of the FIR filter. Until design rules are specified it would not be possible determine the number of chips required. However, using VHSIC I technology as a guide, we would expect that the entire filter could fit on one chip. If additional chips were required, we estimate that approximately 1-2mW extra power per chip would be consumed driving the pads at the required 8MHz rate.

II-B. Filter Organization

There are numerous possible direct, cascade, parallel and serial approaches to implementing a 1024 tap FIR digital filter. Since it appears that the power consumption associated with data and coefficient storage will dominate the power budget for any FIR implementation, we feel that the organization with the greatest possibilities for different memory storage schemes would be the best on which to base our analysis. The organization chosen as a basis for power calculations is a high speed, single tap arrangement shown in Figure 1, with single memories for data and coefficient storage. This single tap filter operates at the rate of approximately 8MHz or 1024 times the 8KHz sampling frequency. After each new data sample is obtained, the filter will perform 1024 multiplications and additions to produce an updated correlation coefficient. A multitap approach might require more chips, with a correspondingly larger proportion of the power consumption going to chip-to-chip communication.

The filter consists of three sections, a memory, an arithmetic section and

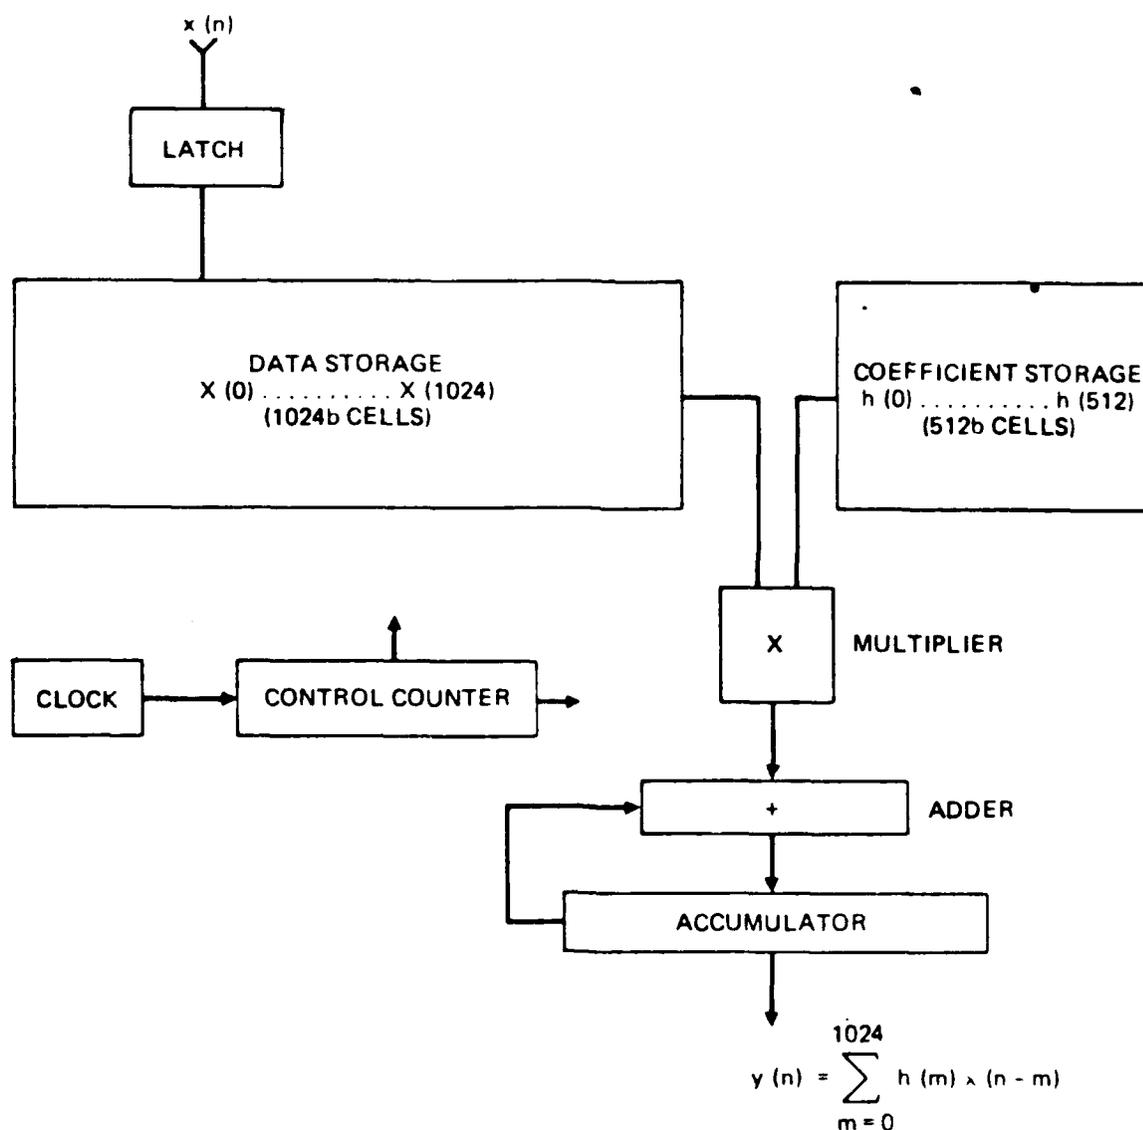


Figure 1. Block diagram of FIR 1024 tap filter implemented as a high speed single tap. Note that the coefficient memory is half the data memory because the linear phase filter has a symmetrical impulse response.

a control unit. The data $x(n)$ and coefficients $h(n)$ are stored in separate memory modules which are RAM, shift registers, or some combination of both. The arithmetic section consists of a multiplier and an adder accumulator. The multiplier is used to multiply $h(m)$ and $x(n-m)$ and the adder/accumulator performs the summation

$$y(n) = \sum_{m=1}^{1024} h(m) x(n-m)$$

to obtain each filter output $y(n)$. The control unit consists primarily of counter type circuits to organize the flow of data/coefficients to and from the memories and to determine when to zero the accumulator.

II-C. Coefficient and Data Storage Schemes

We will consider two types of memory organizations in this report: a pure shift register implementation and a pure RAM implementation. The shift register approach, shown in Figure 2, has the advantages of simplicity of design and ease of control, i.e., no address generation circuitry is necessary. However, power consumption can be higher because all shift register cells are clocked together even though only two words (data and coefficient) are received each clock cycle. Use of a RAM would require more design effort (e.g. row and column decoders, sense circuits, pre-charge circuits, address latches, and possibly separate timing circuits), but would use less power because the entire array is not activated each clock cycle.

CMOS/SOS Shift Register Cell

The most efficient design for a shift register cell in terms of power and area is the dynamic logic arrangement shown in Figure 2. This cell consists of two inverter stages connected by two sets of pass transistors, each driven by one of the two clock phases. The total cell capacitance is simply 8 times the

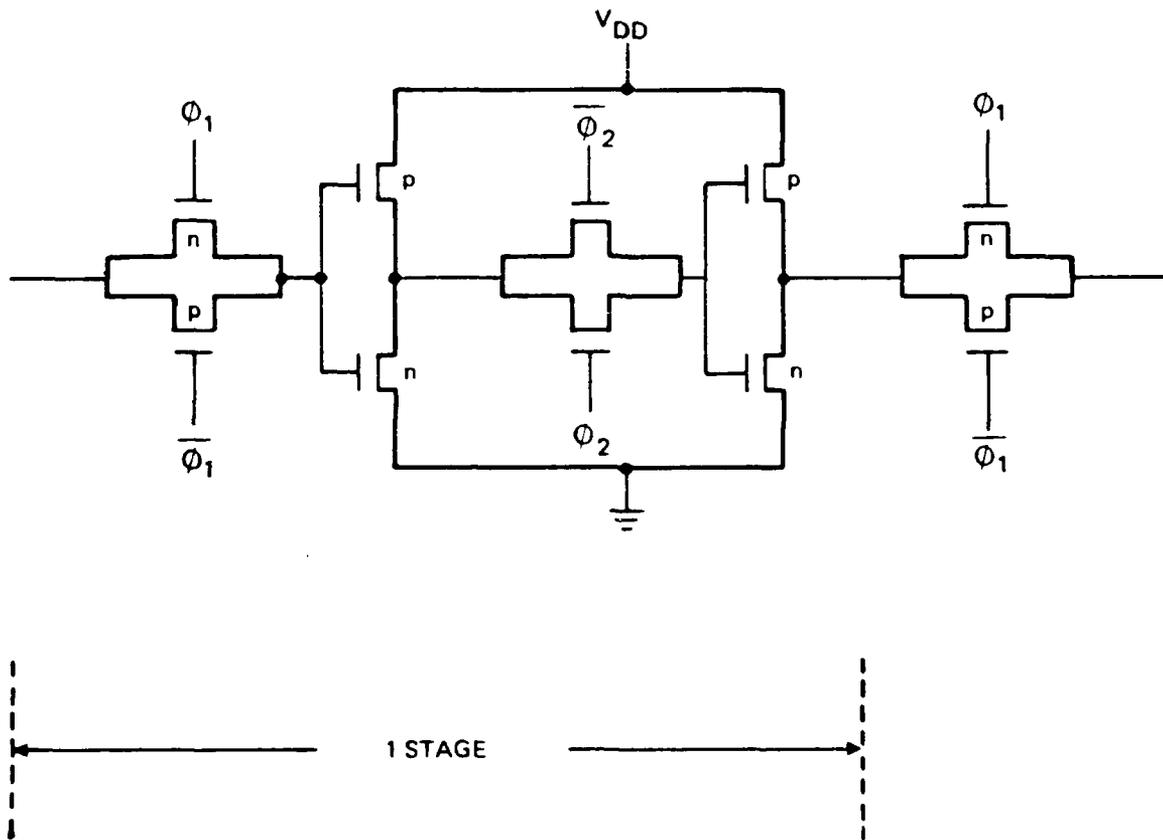


Figure 2. Shift register cell consisting of two inverters connected by two sets of pass transistors.

capacitance of one square of gate oxide or

$$\text{Total Cell Capacitance} = 8 \times 0.015 \text{ pf}$$

and the cell power consumption (CV^2f) is

$$\text{Power/Cell} = 8.75 \times 10^{-3} \text{ mW}$$

We assume that the capacitive clock loading due to the pass transistors in each of the cells is much greater than the parasitic capacitance of the clock interconnect lines within each cell. In this way it is not necessary to consider the power consumed in charging parasitic capacitances of global clock lines. (This is in contrast to non-SOS technologies, where the dominant clock line capacitance comes from the parasitics associated with its global distribution.)

Static RAM

Hughes has recently developed a static 16K CMOS/SOS clocked static RAM memory for high speed, low power applications.⁽¹⁾ Since this device contains the approximate storage requirements desired, we will use it as a basis for the static RAM approach to the design of a low power filter.

The 16K Hughes RAM presently operates from a 5V power supply with typical access times of 110nsec. Static power dissipation is 35uW and operating power is 20mW at its maximum clock speed of 3MHz. Feature sizes are 4.0um (drawn) corresponding to a channel length of 2.5um and die size is 5.5 x 6.5 mm².

The RAM is organized as 4096 words x 4 bits per word (note that this is not the organization we would use), with the array split into two 128 x 32

blocks on each side separated by a decoder as shown in Figure 3. A clocked approach is used so that the RAM does not consume bias power in either the enabled or disabled states. All the power used is CV^2f dynamic power associated with precharging the bit lines and charging the row lines and decoders.

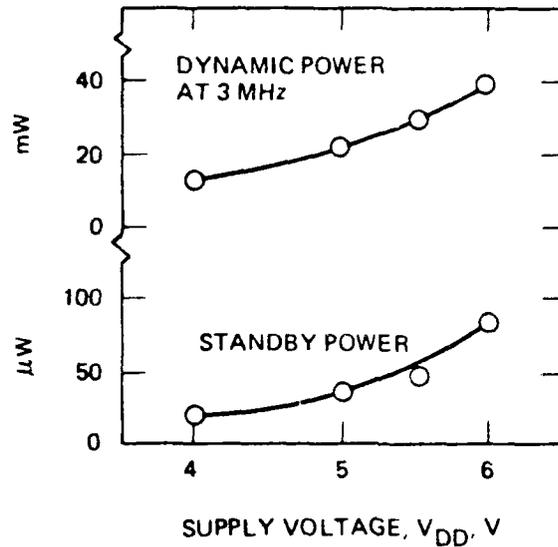
In order to estimate the power used by the low power filter RAM we need to multiply the power dissipation of the Hughes version by a factor of 8/3 to account for the higher clock speed needed, by a factor of 1.5b/16 to account for the different storage requirements, by a factor of $(3.6/5.0)^2$ to account for the different voltage levels, and by a factor of 2.0s to account for scaling. With these modifications we have approximately

$$\text{RAM Power Dissipation} = 5.2sb \text{ mW}$$

For the FIR filter, we would build separate RAMs for data and coefficient storage. We expect that these smaller sizes could then run at the required 8 MHz.

II-D. Multiplier/Accumulator

For the proposed low power filter application there are a number of possible approaches to performing multiplication. The most efficient approach in terms of area is a serial/parallel (shift-and-add) organization; however, the disadvantage of this approach is that it requires a separate set of high-speed clocks and it will be limited in speed. For $b=12$ a speed of approximately 50MHz would be required. Depending upon the design rules used, this could be pushing the state of art. Slower speed operation is possible using two serial/parallel multipliers operating in parallel, but at the expense of increased control overhead. A parallel array multiplier built with



(a)

PC - PRECHARGE
 SL - SENSE LATCH
 CD - COLUMN DECODER

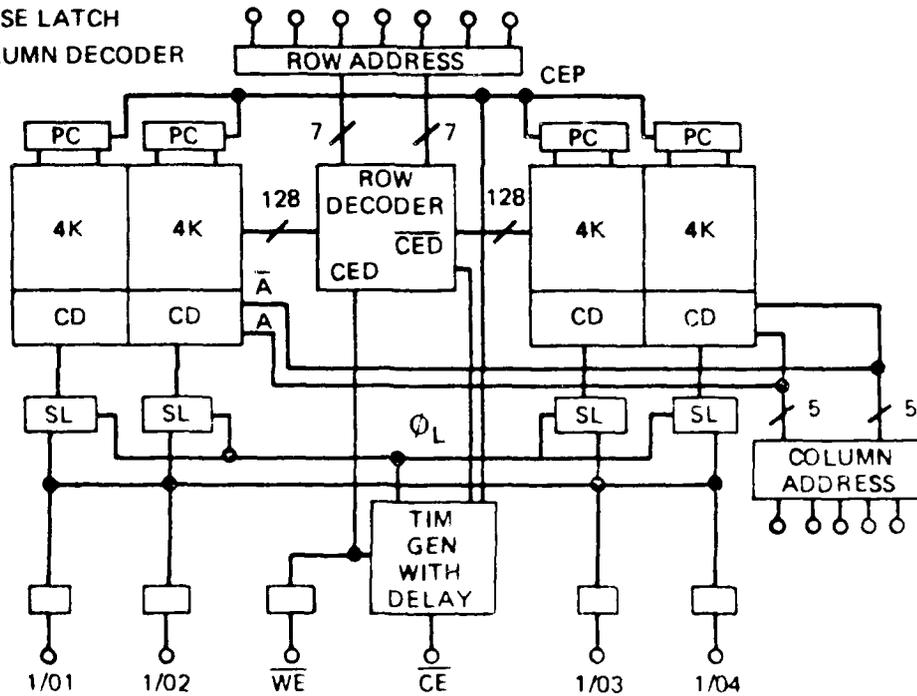


Figure 3. Block diagram and power consumption of Hughes Newport Beach static RAM. The circuit contains approximately 100,000 devices.

combinatorial logic would be straightforward, but area inefficient (only a small part of the array is working at any given time) and relatively slow (only 125nsec per multiplication is available)

Since memory storage requirements will dominate the area usage in any case we think a parallel multiplier organization, with pipelining to increase speed, is the most appropriate approach. In a pipelined multiplier the logic is split into a number of stages so that only a few gate delays are involved each clock cycle. The penalty paid is the increased latency through the multiplier (equal to the number of stages times the time period associated with one clock cycle), which isn't an important criteria for the FIR filter.

A block diagram of a pipelined, carry-save, radix-4 parallel multiplier is shown in Figure 4. As can be seen, it consists of an array of full adder cells which take three binary operands and produce a sum and a carry bit. In the carry-save approach the sum and carry bits are transmitted to the next logic stage. Carry propagation is delayed until the last stage.

The multiplicand and multiplier registers, shown at the top and side of Figure 4, accept one operand each clock cycle. The multiplier word is shifted two bit positions each clock cycle and the lowest three digits are decoded using Booth's algorithm in order to reduce the number of partial product additions by one-half (equivalent to radix-4 multiplication). The output of the recoder is a select/clear signal which indicates whether a shifted, complemented or zero multiplicand should be added to the partial product. Shift registers, shown along with the full adders, are used to shift the multiplicand down through the logic stages of the multiplier, one full adder stage per clock cycle. The carry propagation is done in the ripple adder (last stage).

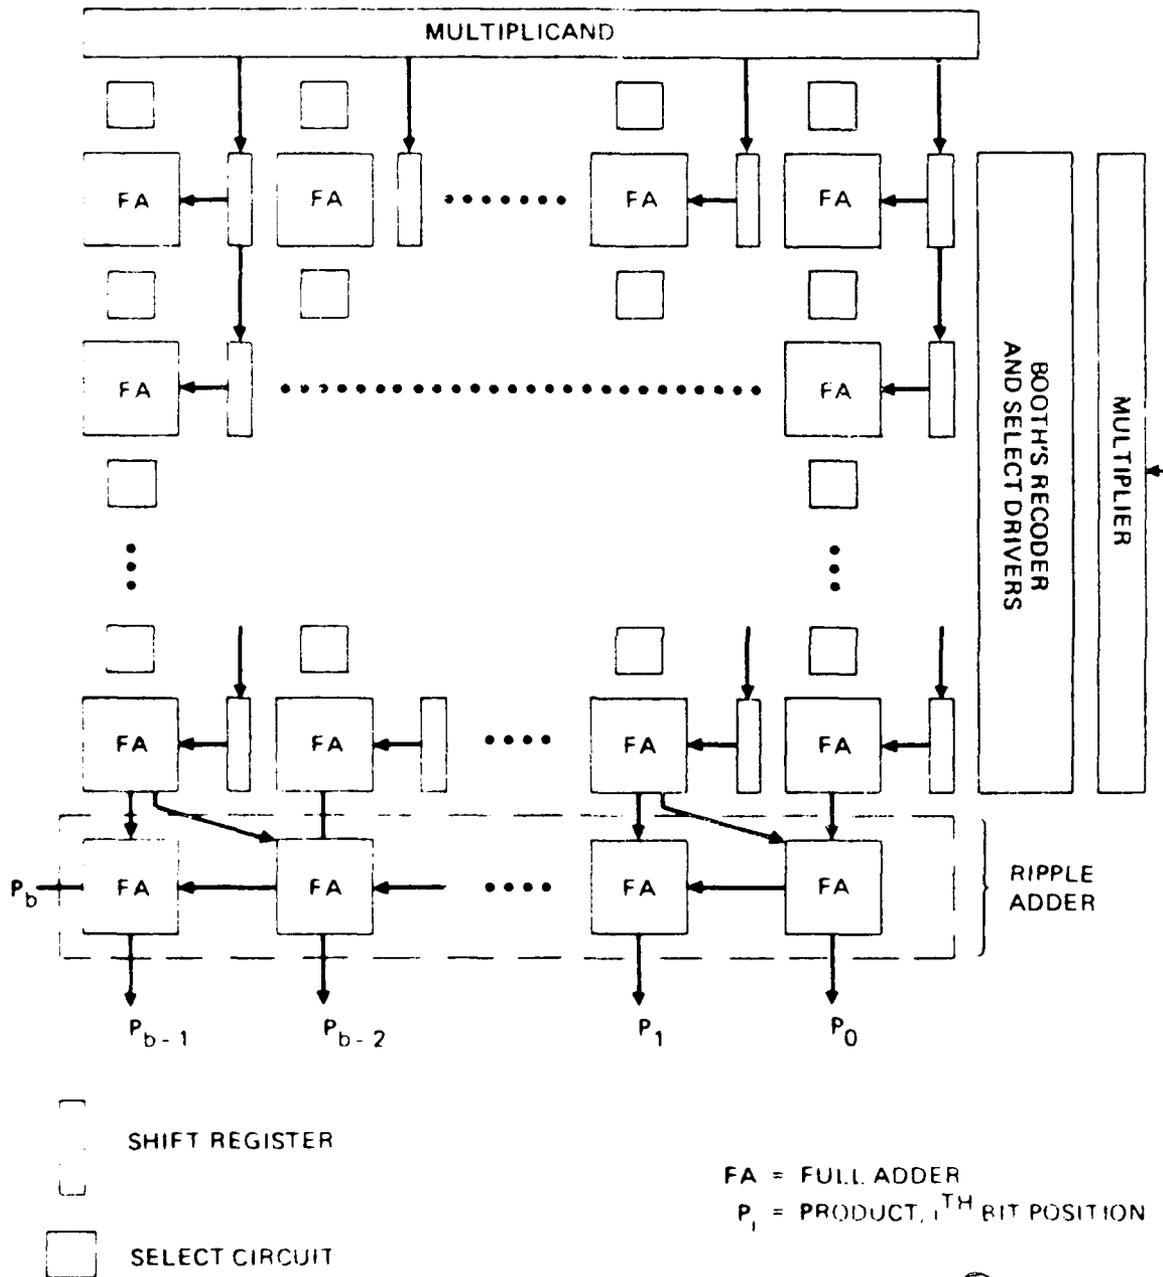


Figure 4. Block diagram of a pipelined, radix-4, carry/save parallel multiplier.

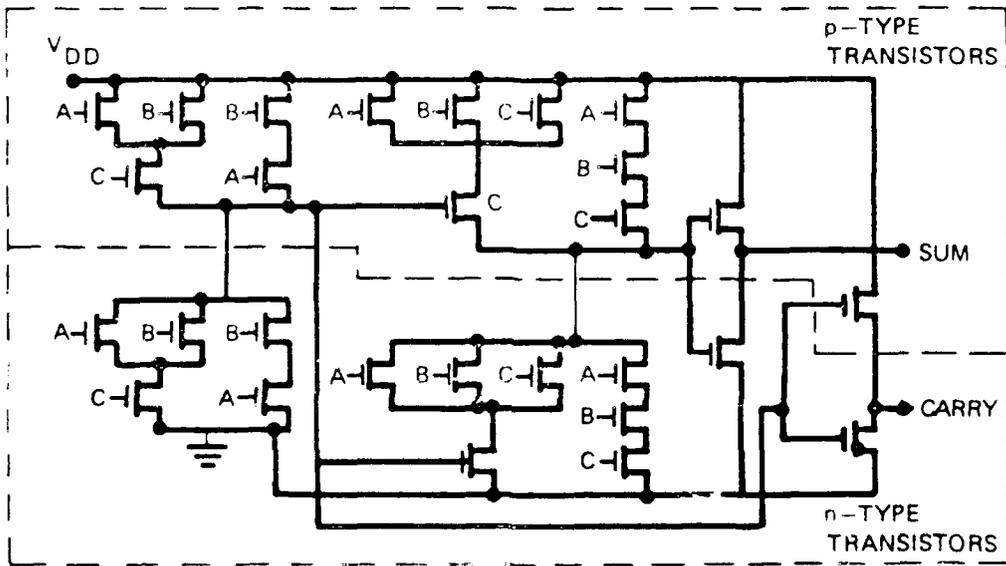
To estimate the power consumption it is only necessary to add the power consumption of the appropriate cells. There are only three basic cells in the multiplier array: a full adder, a shift register, and a select/clear circuit. There is a small amount of random logic in Booth's recoder which we neglect. Finally, the drivers that charge the select/clear circuits must also be included in the power budget. In Table III we break down the number of components in the multiplier associated with each of the above parts. These are parameterized according to the bit length of the operands being multiplied. As can be seen the parts count goes approximately as b^2 , and thus the power will be proportional to b^2 as well. This is in contrast to memory storage power which is proportional to b .

To estimate power consumption associated with the full adder cell we refer to Figure 5, which shows a minimum device, CMOS circuit. Here, A, B, and C are the three inputs to the cell. Power can be calculated based on the number of transistors and the gate capacitance per transistor (0.01s pf for 5um feature sizes) or

$$\text{Power/Full Adder Cell} = 2.5s \cdot 10^{-2} \text{ mW}$$

In this calculation we are assuming equal gate areas for both n and p-channel transistors.

The select/clear circuit, shown in Figure 6, is used to select the inputs to the full adders. The circuit inputs come from the select/clear control lines of the Booth recoder and from the multiplicand, X. The term 2X refers to the shifted version of the multiplicand and CLR indicates that no partial product addition is to take place. As before the total circuit capacitance is equal to the number of transistors times the capacitance per gate (0.01s pf) or



(b)

Figure 5. Full adder cell with a minimum number of devices in order to minimize power consumption.

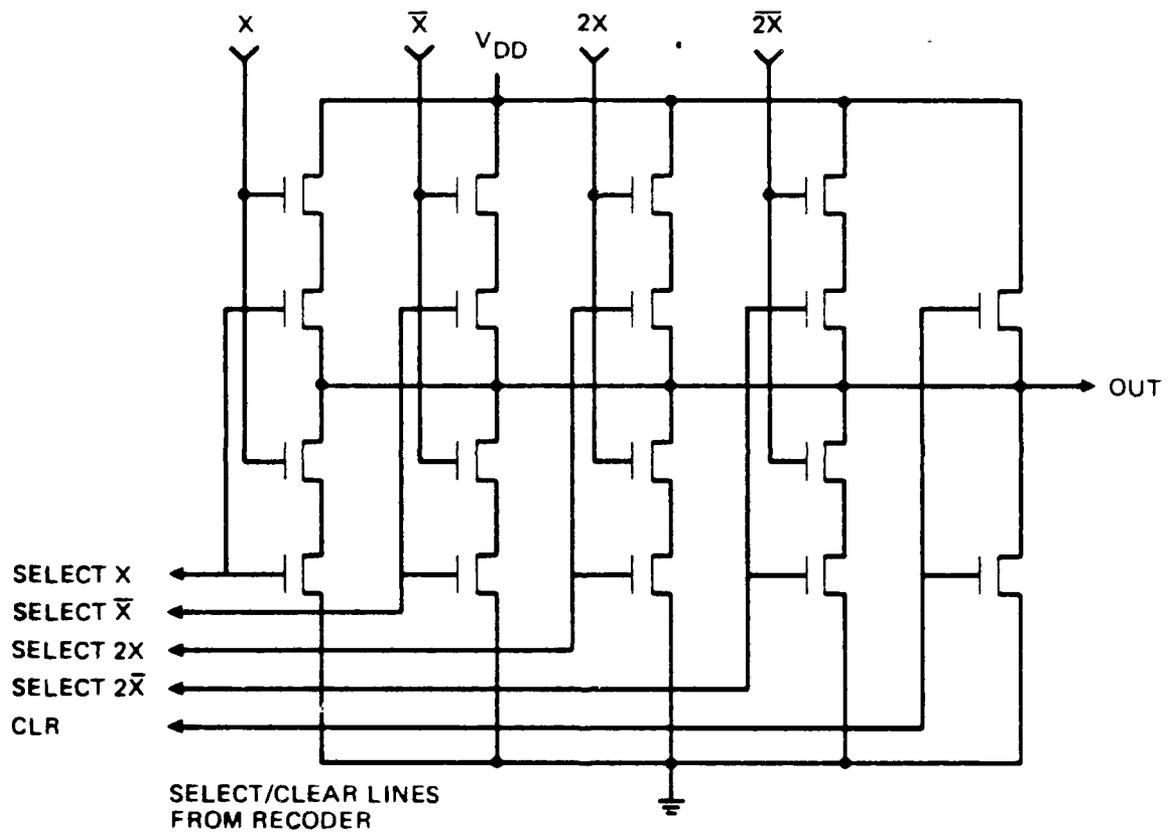


Figure 6. Select/Clear circuit to control addends to partial product.

$$\text{Power/Select-clear Cell} = 2.1s \cdot 10^{-2} \text{ mW}$$

Since each of the select/clear control lines is connected to many transistor gates in the multiplier array, the power consumed by the control line drivers is likely to be significant. Note that we have already calculated the power associated with the capacitance they drive (i.e., the select/clear control lines). Therefore we must only add the power dissipated in the driver gates themselves. We will assume that the gate capacitance of each of the control line drivers (5, total) is approximately $1/e$ of the control line they drive. (Multistage amplifiers typically increase drive capability by a factor of e each stage for minimum propagation delay through the circuit.) The control line capacitance is $0.02bs \text{ pf}$, and therefore,

$$\text{Power/Driver} = 7.7bs \cdot 10^{-4} \text{ mW}$$

The adder/accumulator circuit shown in Figure 7 is the last stage in the arithmetic section. The ripple adder is broken into three pipelined sections of $b/3$ bits in order to reduce by a factor of three the carry propagation delay required. As can be seen the most significant bits from the multiplier must be delayed in shift register stages in order that they arrive at the full adders in synchronism with the carry bits from the less significant bit positions. In addition the outputs of the full adders in the least significant bit positions must be delayed in shift register stages in order to arrive at the accumulator inputs in synchronism with the outputs of the full adders in the most significant bit positions.

The accumulator is basically a set of storage registers which supply an input to the ripple adder each half clock cycle and receive a result from the ripple adder the other half of the clock cycle. The component parts of the

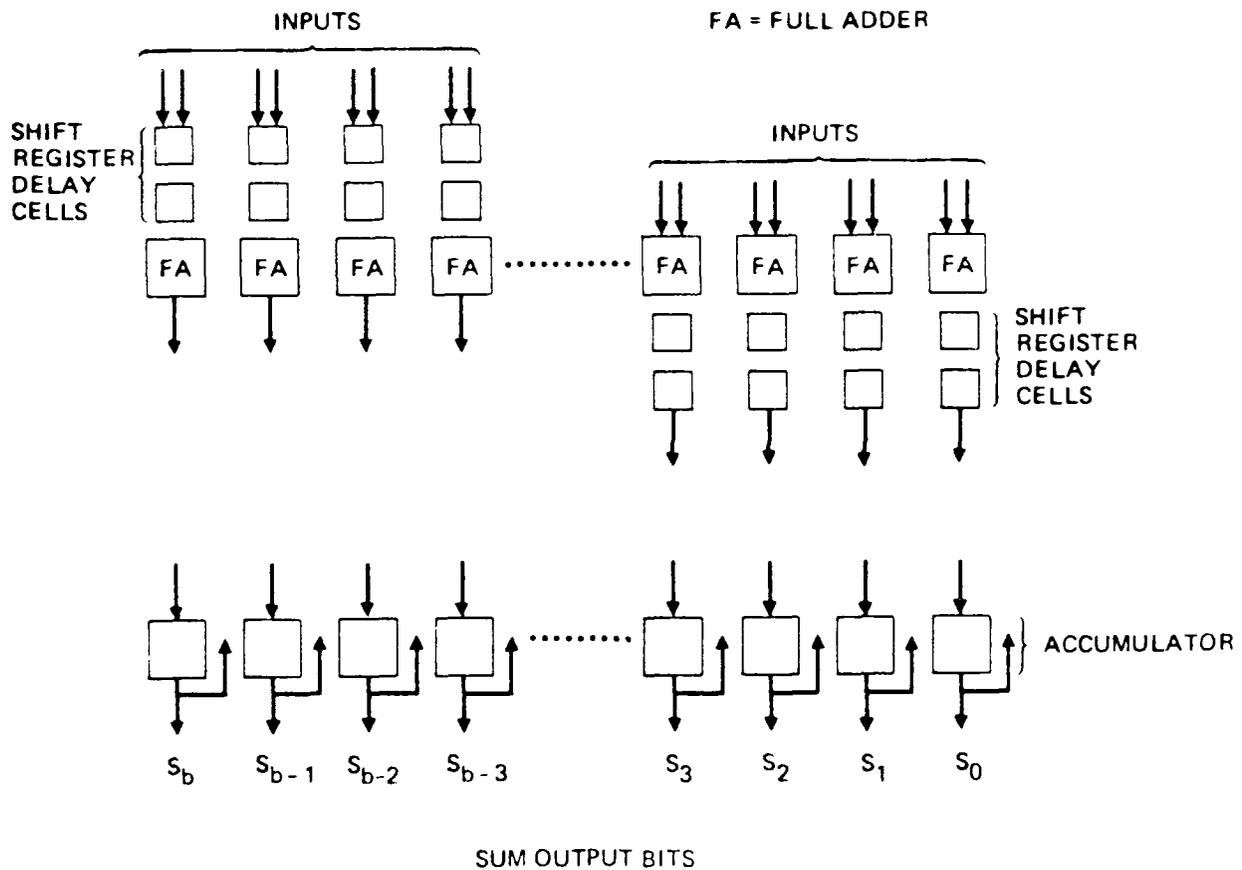


Figure 7. Adder/Accumulator, pipelined for higher speed. Here, it is shown with three sections of 4 bits each corresponding to $b=12$.

adder/accumulator are broken down below in terms of circuits already analyzed for power consumption

<u>Function</u>	<u>Breakdown</u>
Ripple Adder	b Full Adders
" "	48 Shift Register Cells
Accumulator	b Storage Register Cells

We will assume in future analyses (e.g. Table III) that the storage register cells in the accumulator are equivalent in complexity to shift register cells.

II-E. Calculation of Total Power Consumption

The total filter power has been obtained by adding all the power obtained for all the cells used and summing the result. The results are tabulated in Table III as a function of b and s and the total power can be approximated by the expression

$$P_{TOTAL} \approx (13b + 0.054b^2)s \text{ mW} \quad (\text{Shift Register Memory})$$

and

$$P_{TOTAL} \approx (5.2b + 0.054b^2)s \text{ mW} \quad (\text{RAM Memory})$$

We can see from these expressions that one term is proportional to b and one term proportional to b^2 , corresponding to power consumed in the memory and in the multiplier, respectively. A quick calculation shows that multiplier power does not begin to dominate the power consumption until b reaches approximately 16. Since for this application b is less than or equal to 12, the memory power will dominate the power budget.

Because the memory power consumption is so important, the techniques for reducing memory power described in Section II-C will be of great value. In any case the minimum filter power would be that consumed by the multiplier.

The expression above for total power consumption is also proportional to s , reflecting the reduced capacitances at smaller feature sizes.

II-F. Technology Issues

The time table for introduction of the capabilities for various feature sizes is shown in Table IV, based on the CMOS/SOS VHSIC I work presently under development at Hughes Newport Beach. The SOS II technology listed is that used to build their 16K static RAM.⁽¹⁾

We feel that the VHSIC technology offered will be directly applicable to the filter we are considering with the exception of the difference in the supply voltage of 3.6V. The VHSIC I program has already standardized on a voltage of 5.0V and the VHSIC II program has not yet standardized on a voltage level. The incompatibility of voltage levels could be a serious issue and requires further study. Even if it were possible to run circuits at 3.6V with a 5V SOS technology, there could be considerably reduced drive power if the turn-on voltages of the p and n-channel transistors were of the order of one volt. This could reduce the possible operating speed so that the 8MHz rate used in our calculations would be too high.

II-G. Filter Implementation Using CCDs

Charge coupled devices are well known for their low power and high packing density. Thus, there are possibilities for application of a digital CCD technology to the FIR filter. The use of serial memory organizations and pipelined multiplier logic is particularly suited to a CCD approach.⁽³⁾ Although further study is required to definitively map out possible advantages

(3) G. G. Nash, "An 8-Bit Parallel CCD Digital Multiplier," Proc. Custom Integrated Circuits Conf., Rochester, N.Y., May 1982, pp.155-160.

to use of this technology, we have done some preliminary estimates on possible savings in the memory section alone. It appears that there might be a power savings by a factor of 2 and an area saving by a factor of 3 to 4. At speeds of 8MHz we do not think that leakage will be of major importance and memory sizes of 16k have been built in many versions before.

The main drawback to the use of CCDs is the 3.6V supply voltage. CCDs are generally run at higher voltage levels in order to provide adequate transfer margins. However, it is possible that with appropriate bootstrap circuits and proper scaling of circuit parameters the 3.6V supply could be used.

**END
DATE
FILMED**

8-12-87