

THE TECHNOLOGICAL REQUIREMENTS FOR APPLICATION SPECIFIC INTEGRATED CIRCUITS

R.W. Brodersen

Dept. of EECS
University of California, Berkeley California 94720

ABSTRACT

Application Specific Integrated Circuits are becoming an increasingly important class of devices for many semiconductor manufacturers. These circuits often contain highly optimized digital architectures as well as analog circuitry for interfacing to sensors, transducers and actuators. In order to reduce the design time of these circuits high level software tools are being developed. A discussion is given of the requirements of a process technology, which is compatible with this type of design methodology.

INTRODUCTION

The prices of the semiconductor industries' favorite product, memory, has fallen to the point where only a few of the largest manufacturers can afford to compete in that product area. The trend for the future is quite clear as well, in that the return that can be achieved from memory products will continue to erode. Over the last 10 years the cost/bit of RAM has dropped by about a factor of 100 and there is no reason to expect that trend to change. These cost reductions have been continuing in spite of the fact that each new stage in the technology development of RAM's has resulted in significantly increased production costs. The result of this quite unhealthy trend (if you are a RAM supplier) is that the margin on RAM's has decreased to the point that most semiconductor manufacturers are never able to recoup the enormous investment it takes to keep up in this super-competitive business.

The question that is therefore often discussed is what products can be made, on the ever increasingly expensive production lines, that have the characteristic that their margin goes up (instead of down) when significant more capability or performance is offered. One class of products, which appear to be an answer, are the Application Specific Integrated Circuits or ASIC's. These circuits are designed to efficiently implement only one application as opposed to the traditional multi-purpose semiconductor components, such as memory and microprocessors. In many cases ASIC's are able to replace significant portions of an existing system, so (for a while at least) the price of the circuit is not set entirely by competition, but by the value of the parts it replaces.

ASIC's, of course, are not without competitive pressures, but because they require extra system (application) knowledge, the time for competitive parts to appear, is considerably longer than for memory or microprocessors. Samples have already been disclosed for the next 3 generations of RAM components (1, 4 and 16 Megabits), but there is no similar prescience for the next 3 generations of ASIC devices. This has made the higher returns that are possible from ASIC's increasingly attractive for an ever growing number of companies. In fact it is fast becoming the conventional wisdom, (as measured by the popular electronics press and industry analysts), that ASIC devices will be the salvation of the semiconductor industry.

Throughout the history of the semiconductor industry, RAM's have been the driving force for technology development. For example, recently several of the largest manufacturers in Europe have formed a consortium to develop a next generation RAM process. One of the reasons given for that investment is the new technology is needed for their ASIC development, an area in which Europe has been traditionally very active.

However, since ASIC's have different technology requirements than memory, the need to be on the incredibly expensive RAM process development bandwagon, should be, if not abandoned, at least examined.

APPLICATION SPECIFIC I.C.'s

ASIC's are circuits that are dedicated to a single application and may either be meant for a single customer or may be a standard product. To make up for their lack of generality, they must provide higher performance, a more area efficient design or provide specialized I/O functions (often involving analog circuitry). Specialized designs have been used extensively in signal processing systems and it is expected that new signal processing applications will be a major focus of many ASIC designs.

In particular telecommunications applications have been a major application of ASIC designs. One example of the very profitable products in this area are the telephone line codecs (coder-decoders) which perform analog signal conditioning, A/D and D/A conversion. One of these circuits is required per telephone line to connect each end user to the new digital switches and approximately 10 million/year are now being made.

Another related area which has created considerable excitement, is the conversion of the telephone system to an all digital network, which is called the Integrated Services Digital Network or ISDN. This system will provide high bandwidth digital data, which will make possible interconnection of PC's, centralized mainframe computers, and home terminals over what will be the normal telephone network. This network will facilitate information distribution through various means such as voice, videotext, facsimile and direct database access. There is much debate currently on the standards to be imposed on this network, but there is no debate on the enormous variety and quantities of ASIC's which will ultimately be required in its implementation. Many other areas hold similar promise, such as factory automation with its need for sophisticated robots and control systems; office automation with voice mail and speech recognition; new communication techniques such as cellular telephones, satellite transmission, and optical fibers; automotive and consumer electronics and high speed local area networks.

Though the applications themselves vary widely, the circuits themselves have some common elements. Quite often a particular important aspect of these designs is the incorporation of specialized interface circuits. In some cases the interface will be to low bandwidth analog signals, such as required for modems and speech circuits or in other cases will involve complex digital signaling protocols, as required in local area networks and the ISDN. Typical analog circuit functions which must be performed are A/D and D/A conversion, amplification and filtering.

In order to obtain adequate performance and silicon area efficiency in the digital portions of these circuits, the digital architectures are highly optimized. Speed is gained through pipelining and parallelism and area efficiency gained by only providing the minimum required word widths and storage and by the use of dedicated controllers.

The difficulty of predicting the exact system configurations and system level tradeoffs, which stems from a lack of standards in many areas of interest to ASIC designers, results in very intense pressure for fast turnaround of the design. Often the first chips developed for a given application become the defacto standard because customers will use the circuit

which is available first. Once a customer has gone through the trouble of designing an ASIC into a system and has worked out testing specifications, reliability qualification, packaging, and the myriad of other details required to use a complex systems chip, he may be unlikely to repeat the process for the subsequent suppliers of the same function.

The complex functionality of many ASIC circuits, the long product definition phase and the variety of circuits (analog and digital) that are employed in an ASIC design has resulted in the cost of the design and layout being a significant portion of the final cost of the circuit as well as taking a significant fraction of the total chip turnaround time. This in the past has been a major reason that design of "custom" integrated circuits, has been a very expensive undertaking. Fortunately, there has recently been the development of several design approaches which are being used to overcome this difficulty.

Gate Arrays

Gate array technologies are commonly used to design ASIC's, in which a fixed array of transistors are interconnected by several levels of metal conductors. The design and layout of these arrays is very fast and in fact can be completely automated, because of the very limited degrees of freedom in the layout. This also makes possible automated verification and accurate performance predictions. Since only a few process steps are required for customization, the time required for fabrication of the circuit can be extremely short. The problems with this approach stem from its limited flexibility, which result in a substantial loss in area efficiency and performance. In addition, analog circuits are very difficult to cast into this rigid mold and some common required elements such as memory are implemented inefficiently.

Standard cells

A more efficient approach is the use of standard cells in which arrays of cells of common height are placed in rows and the adjacent rows are used for interconnection (wiring channels). The cells which are available in systems using this approach are often just augmented versions of standard TTL circuits. For this reason the subsequent chips appear to be MOS versions of TTL printed circuit boards. Again the very structured form of the design makes possible a high degree of computer automation, so the input need only be a logical or the TTL equivalent description of what is desired. Since the placement of the cells varies with each chip, the entire fabrication time is part of the chip design turnaround time. Though more efficient than gate arrays, the small size of the cells and the constraints on the cell interconnect still severely limit what could be achieved if the full capability of a given technology is exploited.

Macrocells

The approach that seems to have the advantages of hand crafted design in terms of area efficiency and flexibility, but has the capability of a high degree of computer automation is based on macrocells. These are cells that are substantially more complex than the TTL equivalent cells of the standard cell systems, but are generated automatically from computer programs. These cell generators are being developed that range from being very sophisticated programs that build up the macrocells from the transistor level to simple programs that abut previously designed subcells into arrays. Since the macrocells are of different sizes and aspect ratios, the placement of these cells and their wiring is much more complex than gate arrays and standard cells. The optimization algorithms to do this placement and routing are under active development. Particularly difficult problems are the routing of special wires, such as power, clocks and the sensitive analog signals. Systems, however, do exist that can take a high level language input, generate the macrocells and assemble circuits composed of tens of thousands of transistors in a few minutes [1].

TECHNOLOGY REQUIREMENTS

From the above discussion, four characteristics of ASIC's that differentiate them from memories or conventional general purpose microprocessors can be defined:

- 1) Fastest possible chip design and fabrication
- 2) Highly automated tools for design and layout
- 3) Combined analog and digital circuitry
- 4) High performance through specialized parallel digital architectures

The requirements on the fabrication technology which support these characteristics will now be discussed.

Fast Turnaround

One of the most important characteristics for a technology for ASIC's is the need for fast fabrication turnaround. This is in direct contrast to the processes which have been developed for the RAM's, in which each new "advance" has usually resulted in a lengthened processing time, since they require more masking steps, use lithography techniques (wafer steppers) that result in severe process line bottlenecks and have additional steps that are targeted solely at increasing RAM density. The importance of fast turnaround for those designing ASIC's is made obvious by the success of gate array techniques which in spite of their poor performance and low density compared to standard cell or macrocell techniques, make up for it by the fast turnaround they can achieve. The fabrication time has only recently become such an issue since the design tools did not previously exist which would allow designs to be made significantly faster than the fabrication time.

Long fabrication times have an effect on the type of circuits that are designed, because instead of optimizing a chip for each application, an attempt is made to define a more generic chip which can span several applications. This approach significantly complicates the design task, as well as results in circuits which are less efficient. This is particularly true for analog interface circuitry which is very difficult to make programmable. Also, very efficient special purpose digital architectures are difficult to generalize, so considerable compromise is required in the digital portion of the circuitry as well.

It may be thought that even though turnaround is sacrificed with the high density RAM processes that something is gained - namely a high density RAM, which should be useful in ASIC designs. However, these cells are often inefficient to use because of the high overhead in their peripheral circuitry and the difficulty of modifying these cells for specialized operation, such as simultaneous read-write or multi-porting.

Another feature of fast turn-around is that the concept of a silicon breadboard becomes feasible. It is possible to have test chips to test portions of the circuits (eg. the individual macrocells) which substantially eases the testing and debugging tasks. Also this avoids discrete breadboarding which often doesn't accurately simulate the critical parts of the integrated chip architecture and is very time consuming.

Automatic Layout Generation

As ASIC's become a major product area of the semiconductor industry, highly automated design tools will become commonplace. There are several technological developments that can improve the efficiency of these computer generated designs as well as make the optimization tasks within these programs more manageable. It is expected that the macrocell design approach will be the preferred approach in the future and thus most comments are directed towards this type of design.

Multi-level Wiring One common characteristic of most automatic layout programs is the separation of the task into two distinct regions. One area contains the active circuitry (standard cells or macrocells) while the remaining region contains the interconnect wiring (wiring channels). In typical standard cell designs approximately half of the chip area is taken up in signal wiring. The clocks and power are not a problem in standard cell designs since they are generally passed directly between the cells. A minimum of two levels of interconnect wiring is needed to perform the random signal wiring, but if three levels were available then the routing areas could be reduced by one half, resulting in a chip size decrease of 25 percent [2]. It is also very desirable for these interconnect levels to have very low sheet resistivity in order to insure that delay problems are not introduced in the wiring channels.

The macrocell design approach has the additional problem of routing power and clocks in addition to signals. Also, if analog macrocells are used then there is the additional problem of protecting sensitive analog nodes. The problem of routing signal wires, power and clocks in the same channel is very difficult if only two interconnect levels are available and one of them is poly, since it is generally very undesirable to have crossovers in the power lines, and only slightly less desirable in the clock lines. Clearly more levels of metallization, simplifies this problem. For example, four interconnect levels (three metal and one poly) levels could be used as follows: the top level would bus the power around and have relatively coarse features; the lower two levels and the poly could be used for 3 level routing except in areas in which there are clocks and then 2 level routing could be used, with the clocks routed over the top of the signal wires. The availability of multi-levels of metal would appear to be a much less expensive approach to obtaining increased chip functionality than the strategy of scaling down the feature size.

Scalable Design Rules Scaling, of course, can give enormous improvements in density and circuit performance, but it comes at a very high cost. In addition to the equipment and process development costs, there is another cost related to increased software effort in the automatic layout programs. These programs attempt to cope with varying design rules in a number of ways, none of which have been particularly successful if the design rules vary in a way not expected by the program developer. It would be extremely helpful if the technologist in developing the scaled process, would develop the new process which had design rules which improved in a way that was predictable. Ideally the process development would go hand in hand with the development of the design tools, so that the new process modifications would not obsolete the design software.

Keeping the process as simple as possible is one way to reduce the complexity of the design rules, which implies that the number of "clever and novel" process modifications be kept to a minimum. This has the added advantage of potentially allowing more sources to provide the processing.


Predictable Model Parameters Another problem with aggressive scaling is the increase in the variation in the device model parameters. It considerably complicates the automatic layout software to insure that adequate performance will be achieved in all the potential variations of circuits that are allowed. In addition, better physical models are required for the 2 and 3-d effects, so that predictions can be made of the behavior of the circuit in future processes. These models also need to be computationally efficient since larger circuits are being simulated.

Analog Circuits

The combination of analog and digital circuitry is one particularly challenging aspect of ASIC circuit design. The requirements on the technology to implement high quality analog circuits can range from very small modifications to requiring complex bipolar/CMOS processes, depending on the circuit approach used to implement the analog circuits. Designers can do much to simplify the requirements on the technology, so a technologist should be a little skeptical when a long shopping list of specialized process modifications are requested in support of analog designs.

An approach to analog interface circuits (A/D's, D/A's and filters), that are quite compatible with standard digital processes, are based on the concept of oversampling [3]. This approach will become important as the

technologies are scaled, since they are able to trade off speed and digital complexity with analog accuracy. In fact, a recent audio band 12 bit A/D converter with anti-aliasing filter was recently implemented using this technique without requiring any precision analog components [4]. A block diagram of the analog portion of one of these converters is shown in Figure 1.



(1) Block diagram of an A/D converter which doesn't require precision analog components

The comparator in this circuit performs a 1 bit A/D conversion at a rate 256 times the final desired sample rate (2MHz for a 8kHz final rate). The feedback circuit and integrator around this comparator moves the quantization noise from the very coarse 1 bit conversion to frequencies above the audio range. A digital filter which filters out this noise, follows this analog circuit and also reduces the sample rate to the desired final value. In the configuration of Figure 1 the A/D accuracy is almost independent of the noise and offset of the comparator, the capacitor ratios, and the gain of the integrator over a very wide range of parameters. The one requirement of this circuit is that their be at least one linear capacitor. Since the capacitor ratio is not critical and there are only a few of these capacitors a low capacitance/area capacitor is adequate, such as the poly metal capacitor in a standard single metal process. The requirement for 12 bits of accuracy is that the voltage coefficient of this capacitor be less than 20 parts/million.

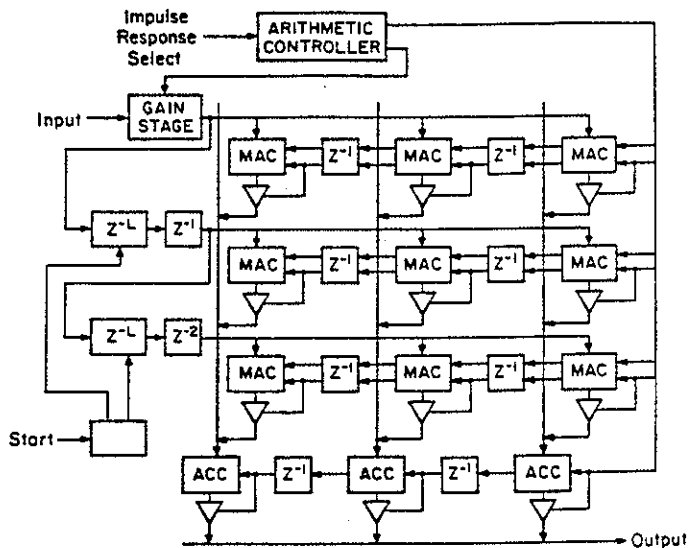
A variety of circuit techniques, such as voltage references, comparators, switched capacitor filters and charge redistribution A/D's have been developed which can perform wide variety of analog functions in a standard digital MOS processes. However, most of these techniques require processes with at least 10 volt supplies and a high quality capacitor, such as exists between two levels of polysilicon with a thermally grown oxide. Unfortunately, the new 1 to 2 micron processes are mostly limited to 5 volts and a drop to lower supply voltages is inevitable in the future. It is quite difficult to design complex analog circuits (such as operational amplifiers) at such reduced voltages and quite often the large voltages are required for driving displays, actuators, etc.. The two metal, single poly process, which is quite widespread, does not have a high quality capacitor, since the deposited insulators between the interconnection levels contain hysteresis, which result in significant voltage coefficients. Therefore, if the very successful analog techniques of the past several years are to be still used, then some additional technology development is required.

Specialized Digital Architectures

Scaling is one way to achieve high performance, but as discussed above, it is not without its disadvantages. Another way to achieve high performance is to use digital architectures which are optimized for a given task. Video processing is an area in which it is commonly believed that the 1-2 micron technologies are necessary, because the sample rate for these systems is typically 10MHz. However, recently a set of ASIC's circuits have been developed that implement a variety of image processing algorithms in real time, but only require 4 micron technology [5].

One of these circuits implements a two dimensional filter (3x3 convolution) of an image. At the 10MHz sample rate, 9 multiply and 9 add operations are required in addition to four memory reads and writes. The computation rate of this chip is therefore on the order of 200 million operations/second. The extensive use of pipelining and parallelism is indicated in the block diagram shown in Figure 2. Each of the 9 MAC blocks performs a multiply accumulate operation in parallel and there are a

number of stages of pipelining in the memory as well as in the MAC units.



(2) Block diagram of an image convolver circuit

CONCLUSIONS

Application Specific Integrated Circuits have many characteristics which are different from RAM's and general purpose microprocessors. The most important of these differences have been discussed along with the requirements they place on the fabrication technology. These differences are sufficiently fundamental that the use of RAM designs to drive the process development is probably no longer appropriate.

A great premium is placed on the shortest possible time for the chip design and fabrication cycle for ASIC's. Automated design tools are reducing the time required for design, so that fast turnaround through the processing is becoming a key issue. These automated tools impose their own requirements on the technology with multi-level metallization, scalable design rule descriptions and stable and predictable device models, as being of key importance.

Though scaling of the feature size yields significant benefits for ASIC's, there are tradeoffs which must be analyzed and it is likely that ASIC's do not require the state-of-the-art feature definition that is required to be competitive in RAMs. For example, specialized pipelined chip architectures can achieve high performance without using the latest technology. As an example an image processing circuit was described which achieves 200 million operations/second in 4 micron technology.

The ability to implement analog circuit functions is often a critical part of an ASIC design. To continue the use of existing techniques, a high quality capacitor is required as well as higher supply voltages than usually predicted for future processes. There are new approaches, however, which can exploit future scaled processes and they will likely become increasingly important. The existence of these design approaches, indicate that it is reasonable for the technologist to question the designer's "shopping list" of process features.

REFERENCES

- (1) J. M. Rabaey, S.P. Pope and R.W. Brodersen "An Integrated Automated Layout System for DSP Circuits", *IEEE Trans. on CAD*, Vol CAD-4, No. 3, July 1985, pp. 285-296.
- (2) Private Communication with A. V. Sangiovanni-Vincentelli
- (3) J.C. Candy, W.H. Ninke and B. A. Wooley, "A Per-Channel A/D Converter Having 15 Segment Companding", *IEEE Trans. on Communications*, Vol. COM-24, No. 1, Jan. 1976, pp. 33-42.
- (4) M.W. Hauser, P.J. Hurst and R.W. Brodersen, "MOS A/D-filter Combination That Doesn't Require Precision Analog Components", *1985 IEEE ISSCC Digest of Tech. Papers*, pp. 80-81, Feb. 1985
- (5) Private Communication with P.A. Ruetz, PhD Thesis (to be published)