ASICRCFWP
Rev. 1, 11/2004

ASIC Versus Reconfigurable Compute Fabric (RCF) Solutions

By Roman Robles

Design managers often accept application-specific integrated circuit (ASIC)-based solutions as their least expensive option in developing telephony infrastructure solutions, including modems. For early-generation wireless base stations, the demands of code division multiple access (CDMA) air interfaces and the limitations of programmable digital signal processors (DSPs) made ASICs a comfortable solution. However, there may be a price to pay for this comfort. This paper explores the total cost of ASIC ownership and challenges the conclusion that ASICs are always the least expensive choice. It enumerates the benefits of Freescale's Reconfigurable Compute Fabric (RCF) technology, which can greatly reduce engineering development time and project costs while yielding a longer lasting revenue stream.

CONTENTS

1 The ASIC-Based Solution

This section considers the ASIC-based baseband solution for the 3G infrastructure market by taking a look at the phases in the design flow, estimating the total cost of these development phases, and then considering these costs in the context of the cost of the ASIC itself and incremental costs of any re-spin of the device.

1.1 ASIC Design Flow

When the cost of an ASIC-based solution is considered, the base of reference is usually the direct silicon cost in dollars per mm^2. In consumer market applications in which the total development cost can be amortized across very large volumes, this approximation is generally useful. In contrast, infrastructure markets, which are characterized by low-to-modest volumes and demand greater longevity of product (that is, less churn) amid a landscape of ever-changing standards, can present a radically different business situation. First, let us explore an arguably representative case for the development of a large ASIC (~4 to 6 million gates) in an advanced process (0.13µ to 90 nm). **Figure 1** shows a simplified depiction of the steps required to deliver an ASIC-based modem subsystem. This is the design flow referenced in this paper. The following sections describe the tasks in this flow and consider the general resources and time frames required to complete these tasks.

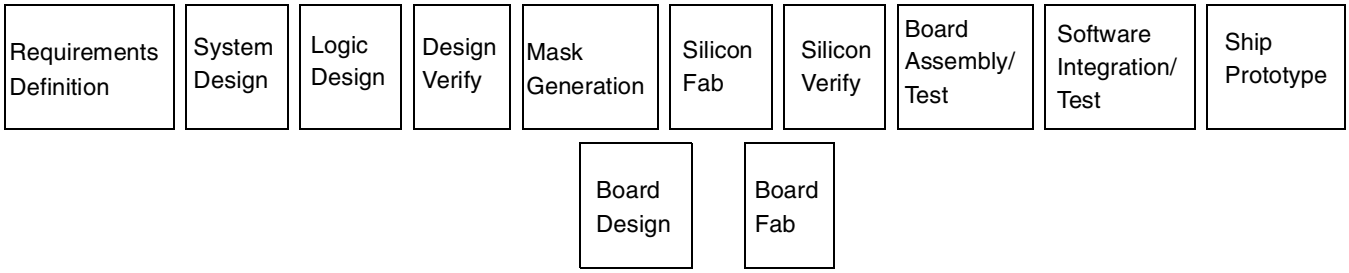

Figure 1. Typical ASIC Design Flow

1.1.1 Requirements Definition

A product life cycle begins with a marketing and/or engineering proposal for the product concept. Ideally, a product proposal gets the benefit of marketing, engineering, and customer inputs during the process of reviewing, revising, and refining the target behavior and performance of the product. Assuming that these teams have previous experience in this market and/or with products of related applications, this task can be completed in a couple of months, though such an aggressive schedule is not typical.

1.1.2 Systems Design: Developing the Modem Architecture

Once the concept is documented and any necessary marketing analyses have warranted that further effort is justified, systems engineers and systems architects are assigned the task of describing the solution in detail. They evaluate the trade-offs between different approaches, analyze competing algorithmic options, and survey the multiple choices in hardware/software partitioning. It may be necessary to develop new algorithms or search for innovative implementations of existing ones. This stage of development may require significant investment in research, and the level of resources required to complete this effort varies greatly.

1.1.3 Logic Design

During the logic design phase, the block diagrams and algorithms are converted into source code in a hardware definition language. This phase consumes much time and many resources. There may be some iteration between this step and the previous two development steps. Portions of this phase and the verification phase can overlap.

1.1.4 Design Verification

This critical phase of development is frequently underestimated. This phase has been estimated to require more than twice the effort of the design phase, but added head count can mitigate some of the schedule impact. The severe cost of re-spin dictates thoroughness of verification. Some of the design verification task can overlap with logic design, but a significant portion of this task remains after the design is complete.

1.1.5 Mask Generation

Although mask generation is a well-defined and relatively brief task, it is a milestone that very clearly delineates the end of the design effort and the beginning of the silicon fabrication. Today's leading edge process technologies (for example, 0.13µ or 90 nm geometries) call for significantly expensive mask sets. These masks alone can cost as much as $1M. Mask generation should be completed in 1–2 weeks, depending on the vendor and the situation.

1.1.6 Silicon Foundry Fabrication (FAB)

Silicon fabrication, in which the wafers are processed and the die are packaged, is another task that is unambiguously defined but not brief. It can be difficult to influence the delivery schedule of a silicon foundry significantly. This phase of development can occupy ~8–14 weeks, depending on the technology and the foundry.

1.1.7 Silicon Verification/Functional Check out

Considerable effort goes into confirming that the packaged silicon devices behave as the simulation predicted. Functional behavior must be verified, and performance must be measured. Thoroughly exhaustive testing can consume much time and may not be practical. However, it is desirable that a large portion of this work be complete before system integration and system test proceed. Verification can continue in parallel with other efforts. A separate board may be designed and built specifically for silicon verification. In all cases, this step depends on the availability of some form of platform. Additional expenses are associated with this approach, and, though relatively minor, should be added into the calculations.

1.1.8 Board Design and Fabrication

Ideally, board design and fabrication occur in parallel with silicon fabrication to minimize its impact on the project schedule. Board design cannot start until a reasonably stable functional specification is available, and it cannot progress very far until all device pinouts are frozen. In addition to the actual cost of the board fabrication, project expenses should include the engineering, technical, and administrative head count and the required CAD tools. Although this total board effort can consume a few months, it can run concurrently with the silicon fabrication activities.

1.1.9 Board Assembly and Test

Obviously, board assembly cannot be completed until all silicon is received—and it usually must wait for a moderate degree of silicon verification. After the first few boards are assembled, a progressive series of tests is necessary to confirm the board design and to verify that manufacturing and assembly steps have not introduced problems. In some cases, this board is used for silicon verification as well. If a separate board has been designed for silicon verification, then the same conditions apply. In either case, integration and test depend on assembly and test of the target board (as opposed to any interim boards designed for silicon verification).

1.1.10 Software Integration and Test

There is an assumption that a modem will include some form of processor to provide the necessary control plane functions. While this processor may be embedded on the ASIC, doing so has a tremendously adverse impact on the project schedule and greatly increases the complexity, size, power, and so on of the ASIC. This paper does not address this significantly more complex option.

Efficiency dictates that all control software be written and tested through simulation prior to receipt of the prototype subsystem hardware. During this stage of development, the ASIC functionality is verified as part of a larger system. Behavior and performance of the entire subsystems are measured and compared to the original target. The effort and time required in this phase of the project are large and frequently underestimated. In addition, once the complete subsystem is verified to be functional and (hopefully) to meet the target specifications, it must be demonstrated that the subsystem performance complies with all associated standards. This process of verification with compliance can be extremely expensive, especially if problems must be corrected or if the user does not have ready access to necessary test facilities. Any deficiencies in design that appear at this stage impose the risk that changes must be implemented and that the process must return to the logic design phase for a second iteration.

1.1.11 Ship Prototype Product

Finally, when the prototype product is shipped, the revenue stream can begin. Of course, if errors in the architecture or design are encountered, this stage is delayed for a re-spin and the process repeats itself, hopefully at an accelerated pace.

1.2 Total Cost of ASIC-Based Design

Table 1 shows an estimate for the nonrecurring engineering expenses of an ASIC-based design. This estimate is based on moderate assumptions about the effort required for the phases, and it reveals the high cost of getting the product ready to ship to customers: 18 months of time and more than $10M, assuming no re-spins.

Table 1. Estimate of Nonrecurring Engineering Expenses for ASIC-Based Solutions

Task	Schedule (Weeks)	Headcount (Engineering Months)	Cost in Dollars ($180K/Engineering Year)	Comments
Requirements Definition	8	8	120K	
Systems Design	8	16	240K	
Logic Design	16	150	2250K	Logic design and verification = 25 weeks
Design Verification	32	300	4500K	
Mask Generation	2	—		~$750K–$1M for masks
Fabrication	14	8	120K	
Silicon Verification	4	10	150K	
Board Design and Fabrication		12	144K	Parallel with silicon fabrication
Board Assembly and Test	4	12	144K	
Software Integration and Test	10	40	600K	
NRE Total	**75 Weeks**		**~$8.3M**	**Calendar time and headcount only**

Table 1. Estimate of Nonrecurring Engineering Expenses for ASIC-Based Solutions (Continued)

Task	Schedule (Weeks)	Headcount (Engineering Months)	Cost in Dollars ($180K/Engineering Year)	Comments
CAD tools			~$3M	
Masks			~$1M	
Boards and Comp			~$0.25M	Small quantity only
TOTAL			**~$12.5M**	

1.3 Amortization of ASIC Development Costs

The cost of engineering development detailed in **Table 1** must be added to the cost of the ASIC to estimate the real ownership cost of the ASIC solution. Moreover, we must assess the incremental cost of any re-spin of the device, including both the cost of producing the second pass silicon and the additional engineering development time. Assessing the cost of ASIC development to be in the range of $12.5 M and assuming that everything will work as needed on the first try (both optimistic assumptions), we distribute this $12.5M across the production quantity of this device. Making another highly optimistic assumption, we expect to ship 10 million voice channel equivalents based on this solution. We have now added over $1.00 per channel to the cost of the ASIC solution—beyond the system material costs. Note carefully that this adder is multiplied by the number of unique ASICs required in the system (for example, if the solution requires two unique ASICS, almost $2.00 per channel is added, and so on.)

1.4 Cost of Being Late to Market

A McKinsey and Company study concluded that the cost of delayed entry into a market can be estimated as shown in **Figure 2**.

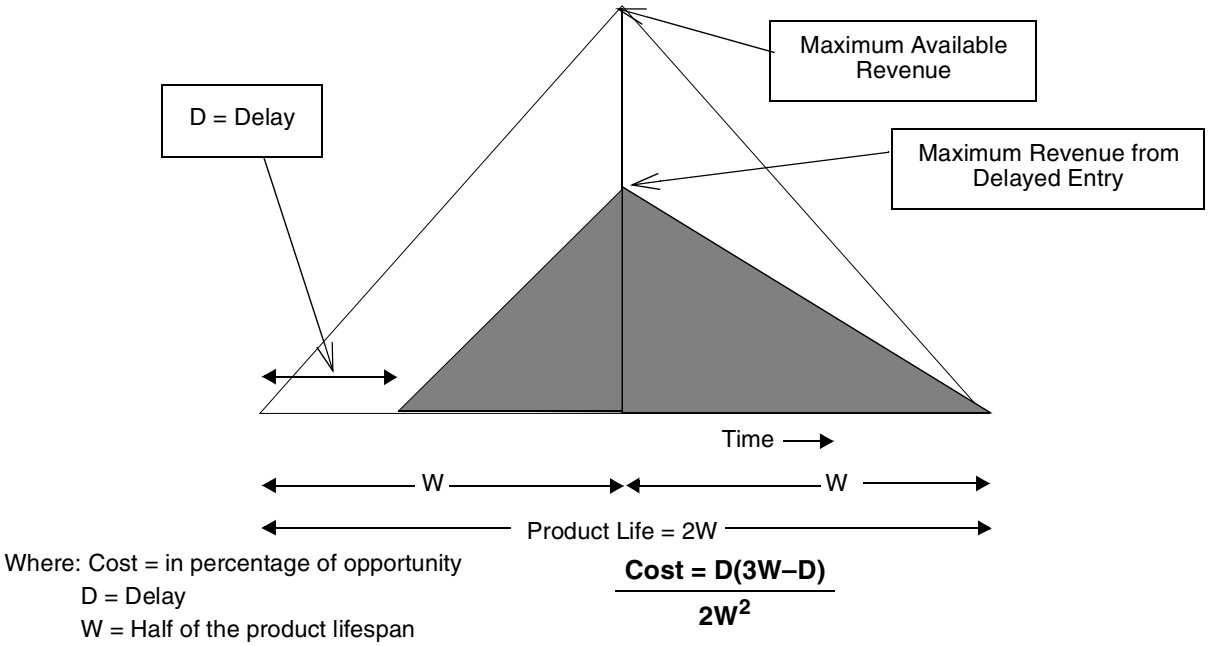

Figure 2. Cost of Delayed Entry into a Market

Assuming that a Node-B product lifetime is 5 years (60 months) and a delay of 6 months in getting to market, the cost is a loss of 28 percent of the original revenue estimates. Apparently, the stated advantages of ASIC-based solutions are much less obvious than many of us have been led to believe. Are there any viable alternatives?

2 Freescale Solution: Reconfigurable Compute Fabric

Efficiently combining an array of compute elements and a configurable interconnect fabric, the Freescale RCF technology can deliver performance at the levels required by next-generation air interfaces while avoiding the expenses and schedule delays associated with developing large ASICs.

2.1 Accelerated Development Methodology

Systems designed around Freescale RCF devices permit the user to take advantage of the benefits afforded by an off-the-shelf solution. Since these solutions are based on standard products instead of custom ASICs, they avoid the expensive and time-consuming steps of logic design and verification, silicon manufacture, and device test. The prototype boards can be populated as soon as the boards can be designed and manufactured. Board delivery no longer occurs at a leisurely pace behind the silicon fabrication. Board delivery now becomes the critical path. Unlike silicon fabrication, board delivery is a task that offers many opportunities for acceleration—under the control of a project manager.

The RCF methodology greatly reduces development time because it offers numerous opportunities for parallel activity. Whereas ASIC-based product development necessarily follows a sequential progression, RCF product development permits considerable overlap between system design, software development, board design/manufacture/assembly/test, and system integration/test. The ASIC methodology of **Figure 1** is modified as shown in **Figure 3** to yield the RCF methodology shown in **Figure 4**.

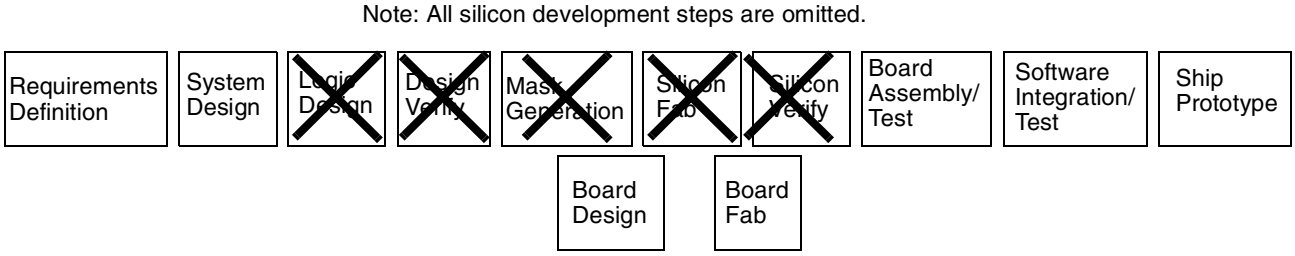

Figure 3. RCF Modifications to ASIC Methodology

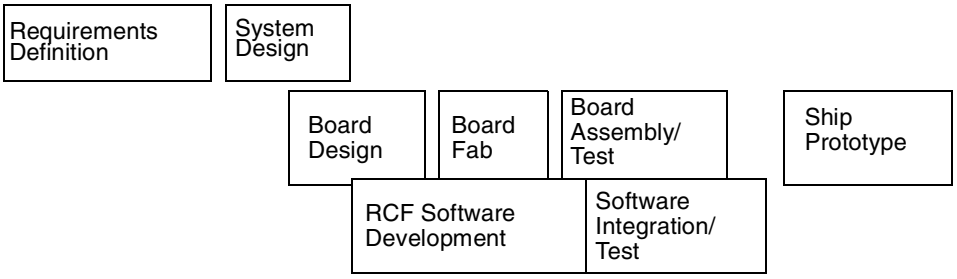

Figure 4. Resulting RCF Methodology

2.2 Impact of RCF Methodology on Development Cost

As **Figure 4** shows, software development can be completed in parallel with the fabrication of the prototype boards. As with ASIC-based projects, it remains true that system integration and test can consume significant amounts of time. However, RCF makes it highly probable that any anomalous behavior observed at this stage can be remedied with a software patch—even after prototype systems ship. Eliminating the silicon manufacturing step and developing software in parallel with the prototype boards, RCF can reduce project development costs from the ~$12.5M estimated for ASICs to ~$3.9M and also reduce the schedule from ~18 months to ~9 months, as shown in **Table 2**.

Table 2. Estimate of Nonrecurring Engineering Expenses for RCF Solutions

Task	Schedule (Weeks)	Headcount (Engineering Months)	Cost in Dollars ($180K/Engineering Year)	Comments
Requirements Definition	8	8	120K	
Systems Design	8	16	240K	
Software Design	16	100	1500K	Software design concurrent with board development
Board Design and Fabrication	8	12	144K	
Board Assembly and Test	4	12	144K	
Software Integration and Test	10	50	750K	
NRE Total	**38 Weeks**		**~$2.9M**	**Calendar time and headcount only**
CAD tools			~$0.75M	
Boards and Comp			~0.25M	Small quantity only
TOTAL	**38 Weeks**		**~$3.9M**	

2.3 Benefits of RCF Methodology

The many benefits of the RCF methodology are as follows:

- *Reduction in Development Time*. Board delivery is the critical path item that determines when integration and test can complete. With RCF devices, Freescale makes a library of integrated RCF software modules available. This library of baseband processing functions can form the basis of a user's implementation, reducing the development time for basic functionality so developers can focus on integrating their own intellectual property (IP) into the basic suite where appropriate.
- *Confidence through tested silicon and guaranteed performance*. Since the silicon is tested at the factory with AC and DC performance parameters guaranteed by Freescale, the project team can start system-level testing immediately upon receipt of the prototype hardware. The team can focus on developing the solution instead of testing and verifying the silicon components.
- *Reconfigurability supports multiple standards*. In today's cellular market, multiple standards (W-CDMA, cdma2000, TD-SCDMA, EDGE, GPRS, etc.) compete for acceptance and position in a dynamic business environment. The differences in these standards and the inflexibility of ASICs have forced original equipment manufacturers (OEMs) to design unique modems specifically for each air interface. Each of the projects has the associated expenses of developing new ASICs.

- *Common hardware brings larger volumes*. While the challenge of building a universal base station yet remains, RCF provides the flexibility to deliver a single, commercially viable modem design that can address the needs of all of the air interfaces. RCF therefore permits the OEM to take advantage of the economies of scale afforded by the combined volumes of multiple markets—a feat that is virtually impossible with alternative solutions. The benefits of these increased volumes should be taken into account when the cost of RCF-based solutions is considered. Not only does RCF permit a lower development cost at an accelerated schedule, but much of the development expense can be distributed over a greater potential product volume.
- *Competitive system cost, power, and size*. The benefits of a reconfigurable approach would be lost if the target system's material cost, power consumption, physical size, or channel density were not competitive. Freescale has put considerable effort into all of these areas, and our RCF devices are comfortably competitive with alternate approaches in all of these key areas of comparison.
- *Risk Reduction*. The standards for next-generation wireless air interfaces are constantly changing and will continue to do so for many years to come. The market demands must adapt to unforeseen applications and usage patterns. The flexible architecture of RCF reduces the risk that a design will not be able to adapt to the new standards/market demands
- *Faster time-to-market*. RCF devices are standard, off-the-shelf components. Once the system architecture and partition are selected, hardware development can began almost immediately. OEMs can start manufacturing their systems in parallel with the software development effort. Prototype ideas can be tested in the lab or even in the field and deployed in weeks. Cycle times for development of a new ASIC are measured in months. ASICs must go through long periods of design validation, fabrication, silicon verification and testing before system manufacturing can continue. RCF-based solutions can be developed jointly with systems engineering/software delivery.
- *Longer time in market*. As markets evolve and new applications develop, existing ASIC-based systems require replacement. Their inflexibility renders them unable to serve these new market segments. The reconfigurability of RCF permits the design of a system to continue as a viable solution because it can be modified in the field to implement the new services.
- *Potential for new revenue streams*. The changes in standards and unforecast market demands for new features can become a welcome opportunity for OEMs to generate new sources of revenue. Service providers can offer new services without the need to deploy new hardware. OEMs can enjoy new revenue streams from existing customers by offering software upgrades to address the new services.

3 Conclusion

The real cost of ASIC ownership extends far beyond the actual cost of the few mm^2 of silicon that is sometimes used as a metric. A candid and complete comprehensive accounting of all the costs associated with ASIC development reveals the high costs that often remain unconsidered. These costs must be multiplied by the number of different ASICs in a system.

The development cycle of ASIC-based solutions follows a serial progression through concept, design, design verification, fabrication, device test and verification, board/system assembly, system test, and (finally) shipment to customers (where the revenue stream begins). In contrast, the RCF development cycle eliminates many of these phases and allows others to be completed in parallel—greatly reducing development time and project costs. ASIC solutions are inflexible, precluding their applicability to multiple air interfaces unless considerable additional design effort and superfluous silicon is included in the development project—all at significant incremental expense and with associated schedule impact.

RCF technology lowers the risks incurred from market churn and lowers development costs. RCF technology yields earlier generation of revenue, longer time generating revenue, the potential for higher volumes through addressing multiple air interfaces with common hardware. Moreover, RCF technology opens the door to new sources of revenue for both the service provider and the OEM.

RCF technology yields a competitive edge. It provides a solution that competes with the price and performance of hardware solutions while offering the benefits of programmability.

NOTES:

NOTES:

How to Reach Us:

Home Page:
www.freescale.com

E-mail:
support@freescale.com

USA/Europe or Locations not listed:
Freescale Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
+1-800-521-6274 or +1-480-768-2130
support@freescale.com

Europe, Middle East, and Africa:
Freescale Halbleiter Deutschland GMBH
Technical Information Center
Schatzbogen 7
81829 München, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
support@freescale.com

Japan:
Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku,
Tokyo 153-0064, Japan
0120 191014 or +81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:
Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate
Tai Po, N.T. Hong Kong
+800 2666 8080

For Literature Requests Only:
Freescale Semiconductor Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
1-800-441-2447 or 303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com

ASICRCFWP
Rev. 1
11/2004

freescale™
semiconductor