

**Kilopass XPM Technology:
High Density CMOS Logic Non-Volatile Memory
Enables Embedded Programmable Firmware**

Author: Craig Rawlings, Kilopass Technology Inc.



Kilopass Technology Inc.
3333 Octavius Drive, Suite 101
Santa Clara, CA 95054

Abstract

Chip and system designers have traditionally been faced with large tradeoffs for choosing a firmware storage solution from traditionally available non-volatile technology options. These tradeoffs are divided between implementing on-chip Masked-ROMs, or storing firmware in off-chip NVM (non-volatile memory) in conjunction with on-chip SRAM blocks. Simply stated, the newly available Kilopass XPM technology is the only standard CMOS-NVM solution that meets the requirements of firmware storage without the drawbacks of mask-based (ROM) technologies and without the drawbacks of a non-standard CMOS process.

Author Biography

Craig Rawlings has held senior marketing and sales management responsibilities at companies including Hewlett Packard, Actel, and now Kilopass. He has over 14 years of experience in the EDA and semiconductor industries. Craig has also held roles that include both international and domestic marketing and sales responsibilities; Kilopass is Craig's third start-up experience.

Chip Designer's Dilemma

Chip and system designers have traditionally been faced with large tradeoffs for choosing a firmware storage solution from traditionally available non-volatile technology options. These tradeoffs are divided between implementing on-chip Masked-ROMs, or storing firmware in off-chip NVM (non-volatile memory) in conjunction with on-chip SRAM blocks, transferring the firmware for execution into the on-chip SRAM blocks. These two choices came at the substantial cost of flexibility in the Masked case, or at the cost of extra silicon for the SRAM block plus additional board area for the external chips. A third choice available for some applications is the use of integrated FLASH blocks on-chip, resulting in a cost premium to the whole chip of between 30%-50%.

A field or factory programmable firmware solution for embedded on-chip CPU, MCU or DSP based CMOS SoCs has not previously been available. In other words, there has been no CMOS NVM solution for the integrated storage of an SoC's firmware. While there are a number of new technologies that enable non-volatile memory functionality in CMOS, only Kilopass provides a path to achieve the high bit density, area efficiency, and programmable requirements within an NVM block needed for embedded on-chip firmware storage applications on all current process nodes.

Embedded memory has always been a critical requirement for SoC, ASIC, and ASSP designers. For integrated designs requiring code storage, as well as serialization and permanent data storage (e.g. encryption keys), SoC architects have had two primary decision paths: 1) masked-ROM in CMOS, or 2) embedding EPROM, EEPROM, or FLASH in a non-standard CMOS process technology, or 3) an off chip NVM solution. This limited decision path has for some time posed architects with undesirable trade-offs. The NVM (non-volatile memory) requirements for small bit density functions may be satisfied with a number of viable NVM technologies; however, currently there is only one NVM technology that satisfies the embedded on-chip requirements of firmware.

This is where the tradeoffs come into play: On one hand, an architect will encounter high cost and larger die area in a non-standard CMOS process, or, on the other hand, choosing a masked-ROM in a standard CMOS process results in a hardwired design that is subject to NREs and will result in severe schedule penalties for any changes to firmware. In order to give SoC architects the ultimate in flexibility, a standard CMOS single transistor NVM cell is needed that provides simple programmability in either wafer or packaged form.

Kilopass Technology Primer

A distinctive feature of the Kilopass technology is that it uses standard CMOS technology compatible with any wafer source or process node. The programming mechanism is embedded in a standard CMOS transistor. The programming mechanism is programmed closed—it is open in its un-programmed state.

Flash and other charge types of floating-gate NVM technologies are limited in their ability to store charge due to the fact that as the logic oxides get thinner, direct tunneling occurs and the charge tunnels off. Thus, Flash charge storage technologies have scaling limitations between 80 and 85 Angstroms for the wafer's tunnel oxide thickness. That is the thickness of the tunnel oxide Flash must have to reliably hold the charge once it is stored. Logic technology gate oxides have long since become thinner, now in the range of 30 Angstroms. Moreover, the lower voltage, required to achieve tunneling, limits the spacing between storage cell elements—the transistors. In other words, with Flash, there must be some defined amount of isolation between transistors which limits the implementation of FLASH to trailing process nodes.

Kilopass technology by contrast does not require this spacing. It obeys standard logic technology design rules with regard to transistor spacing. Additionally, it scales with any standard CMOS transistor for each new process node. In fact, a primary benefit is the fact that the XPM cell's density and performance characteristics actually improve with process scaling.

Note: Anti-fuse technology has existed for some time, but the anti-fuse is composed of a non-standard type of dielectric, typically placed between two metal layers. An anti-fuse closes a connection rather than opening a connection when programmed. Kilopass utilizes a CMOS anti-fuse technology that is a closed connection after programming, but unlike prior anti-fuse technologies, the XPM cell's programmable link is, in fact, the CMOS transistor's gate-oxide such that it scales seamlessly with the standard CMOS transistor as the industry scales to new process nodes. We consider the Kilopass technology an enabling technology innovation providing for high density CMOS-based non-volatile memory applications in SOC design.

The Kilopass Breakthrough

The Kilopass innovation is a very short channel MOS transistor with its gate connected to its drain. It has a finite turn-on voltage, in the low hundreds of millivolts, and a finite impedance on the order of tens of kilohms. The memory block can be programmed with an externally provided programming voltage supply or an on-chip charge pump. Importantly to SoC designers, the XPM cell requires low programming current—under 60 μ A for each programming event. Due to its inherently rapid programming time, a 1Mb memory completes programming in a matter of a few seconds.

SoC Firmware Project Risks

While the SoC's project risks depend on the specifics of each project, there are some common risks that may be reviewed and assessed as they relate to firmware. In the case of implementing firmware utilizing a masked-ROM, anything committed to ROM will involve higher risks for subsequent mask changes. Of course, the NRE for ROM changes may be minimized by containing all ROM changes to a single mask. Even though the foundry's NRE charges will be significantly reduced, there are still significant hidden costs involved with engineering the changes to the ROM's mask, performing verification,

testing, and incurring substantial schedule delays prior to production. From an operations standpoint, the prospect of having fixed, or even worse, dead inventory introduces high risk. Also, as process nodes advance there continues to be an increasing gap between simulated results and the SoC device's physical characteristics. A programmable embedded firmware storage capability as a replacement of a masked-ROM provides both a prototyping and production solution for proving firmware without risking schedule delays and incurring NREs (see Table 1). Related to well accepted experience, embedded on-chip firmware programmability may reduce the SoC project's time-to-market by as much as six months.

	Kilopass XPM	Masked-ROM	Flash/EEPROM
Programmable	√	X	√
No NRE	√	X	√
No Mfg Lead Time	√	X	√
Low Power	√	√	√
Low Cost	√	√	X
Hi-Performance	√	√	X
IP Security	√	X	X
Long Life	√	√	X
Hi Reliability	√	√	X
Scales w/ Process	√	√	X
Standard CMOS Rules	√	√	X

Table 1. Comparison of Firmware Storage Technologies

Even if the firmware is correct the first time, the ASIC, ASSP, or SoC's ROM mask has a 10-12 week manufacturing lead time. Generally, the firmware is directly in the system design's critical path. Note that it is very difficult to complete architecting of the firmware until the silicon device's hardware is golden. For platform chip solutions with PUs and/or DSPs, any new products based on new firmware designs will have a requisite manufacturing lead time of up to three months. While some microcontrollers and microprocessors—these are not the processors or controllers typically used on a SoC design—now include Flash or EEPROM firmware storage capability, these are based on more expensive non-standard CMOS processes that lag leading process node geometries. Table 1 above provides a comparison of the benefits of each of the currently available firmware storage technologies.

As another option, chip level system architects may choose to move to a non-standard CMOS process in order to embed a Flash or EEPROM memory function on-chip. This decision carries with it a net wafer cost premium on the order of an additional 30% die cost (on like nodes). This premium is due to the non-standard CMOS process adder, additional yield degradation, and lagging process node. When factoring in next generation CMOS availability, the SoC FLASH premium on individual die costs may be as high as 50% over standard CMOS. System performance could also be affected due to the older geometry solutions for the FLASH process. There is also a security risk for firmware IP (Intellectual Property) stored in Flash or EEPROM devices since features that operate on voltage differences may be mapped and deciphered with low cost

equipment and without much effort. Additionally, in order to meet the higher performance required by today's CPUs and DSPs, the firmware code must be read from the Flash or EEPROM into high speed SRAM. This is true whether the Flash/EEPROM is located on-chip or off-chip.

Kilopass' first family of embedded XPM products perform approximately 20% faster than their Flash/EEPROM counterparts on like geometries. While the XPM's current performance does not match the performance of high speed SRAM for CPU/DSP applications, the underlying technology supports high speed. Kilopass includes in its roadmap support for single instance high performance firmware applications in the near future. From a design standpoint, because the Kilopass XPM technology scales, this new technology inherently mitigates risks associated with forward looking competitive pressures that drive products to leading process nodes.

SoC System Firmware Trends

As is evident from the general industry, embedded systems are positioned to play a dominant role in electronics markets for a long time to come. Although the PC started the digital revolution three decades ago, PCs have now evolved into a mature market. The core of new activity in 2005 is in embedded systems of all types, but the main focus of this activity will be in wireless and consumer electronics (see Table 2 below).

Due to the recent trending convergence of communications, computing, and entertainment applications, consumer applications have become a strong driver of semiconductor trends. Specifically, digital media will create and drive fast growing markets for SoC architects and designers.

These consumer SoC applications will continue to utilize both embedded MPU and DSP semiconductor circuits in 2005. It is interesting to note that although the SIA (Semiconductor Industry Association) has forecast only around 1% growth rate for the entire industry, the stand-alone DSP market is forecasted to grow at close to 20% (so says Will Strauss, president of Forward Concepts, a DSP market and technical analysis firm). This indicates a very strong trend toward digital signal processing.

DSP Market by Segment
(FY2005E in Millions)

Wireless	68%	\$8,840
Consumer	11%	\$1,430
Multipurpose	8%	\$1,040
Wireline	6%	\$780
Computer	4%	\$520
Automotive	3%	\$390

Table 2. DSP Market Forecast (Source: Forward Concepts Co.)

Digital Revolution

With the digitalization of media over the past few years, we've seen several trends emerge in chip design. First, the processing of digital media drives requirements for DSPs necessitated by algorithmic calculations. Second, MCUs and CPUs have increasingly been incorporated with DSPs by successfully delivering higher performance, more peripherals, and lower power in conjunction with the DSP.

Additionally, in 2005 the industry will see the emergence of the DSC (Digital Signal Controller). DSC applications will continue to blossom beyond motor control into virtually anything that moves and is controlled electronically. Leading-Edge motor control designs have taken the lead in combining the advantages of DSP and MCU.

Firmware Storage Requirements

Development tools will advance with a focus toward making signal-processing development look as much like a CPU/MCU as possible (e.g. automatic code generators). While performance requirements for firmware will depend heavily on the specific application, there will be a continued drive toward higher performance. A large portion of SoC designs with embedded CPU features will require firmware to operate at clock rates between 75 and 250 MHz. High speed applications will continue to push clock rate requirements well over 250MHz. DSP functionality has surpassed these requirements for years; however, the more current trend is to migrate DSP functionality to dedicated function blocks, providing enhanced performance with lower power consumption. For an increasingly large segment of these designs, low power will be a very important consideration as many of these products will be battery powered.

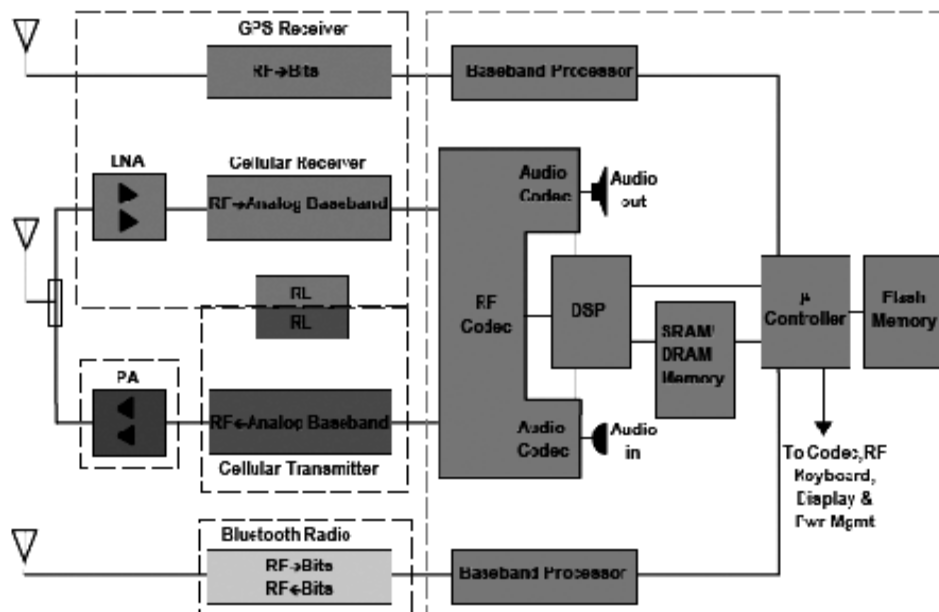


Figure 1. Architecture Design Example for a 3G Cell Phone

The baseline for firmware performance will be compared to today's high speed on-chip SRAM specifications and solutions. For low power applications, baseline power consumption characteristics will be compared to masked-ROM solutions. For firmware storage, a technology that offers the area and power efficiencies of a masked-ROM and the performance of an SRAM provides the ideal solution.

With regards to bit density and the memory configuration itself, most 8-bit MCUs have now moved up to 16 bit data paths, while 16 bit CPUs are now 32 bits. In either case, a 16 bit wide memory configuration for an embedded memory block would be able to serve both 16 bit and 32 bit (using two instances) CPU/DSP applications. Even with two 16 bit memory blocks little area efficiency would be lost as area is dominated by the actual memory array and output data circuitry.

Firmware storage requirements tend to be application dependent and the organization of the overall firmware's system design vary. Some applications may choose to locate the system's BIOS or boot firmware on-chip with more application specific firmware stored off-chip. Most SoC firmware requirements typically fall into the range of between 16KB and 128KB of storage capacity depending on the application and the chip design.

Enabling Technology for Embedded Programmable Firmware

Outside of Kilopass' new XPM technology, there is no other technology today that enables a CMOS non-volatile memory block to be embedded in a SoC design for the purpose of storing firmware. The chip architect would otherwise be forced to store the chip's firmware in a masked-ROM or to move the chip design to a non-standard CMOS process.

There are sizable trade-offs for either alternative:

Masked-ROM Alternative

From a design specification and requirements standpoint, masked-ROMs provide very good area and power efficiency, while maintaining leading-edge performance. The drawback is the same that has led to the emerging prevalence of FPGAs over ASICs for many applications, as well as other programmable NVM technologies. The time-to-market advantages of being mask free are indisputable. Additionally, one device may contain any number of firmware images, paving the way for platform chip solutions based on firmware-centric architectures. Also, the life of the product may be extended through feature enhancements via the system's firmware without the lead time or NRE expenses associated with mask changes. Due to these mask related drawbacks, there is a very visible trend toward technology companies reducing their engineering dependency on fixed mask solutions.

Preeminence of Standard CMOS

The urgent need for on-chip configurable firmware is readily apparent in the fact that a number of system architects have pursued non-CMOS processes for the sole purpose of a Flash or EEPROM repository for firmware. It is important to note that there are hidden costs with going to a non-CMOS process. These include increased technology and manufacturing risks, more limited wafer sources, increased complexity in ramping volume silicon production, and significantly higher wafer costs. These all add up to increased engineering and economic risk. These risks have only increased as contract foundries have increasingly standardized on a standard CMOS process. Most device architects and chip designers refer to a portable definition of CMOS as being “T-Like”, short for TSMC-like for the purposes of CMOS portability to other contract foundries. They have coined this definition because they understand very well the benefits of being compatible with standard CMOS processes.

For all those chip designers who have experimented with a non-CMOS process in order to obtain configurability for their firmware, the prospect of having configurable firmware in a standard CMOS process must appear like an oasis in a vast desert. The fact is that firmware is complex as it is the marriage between hardware and software. There has long been a need to be able to test firmware embedded on-chip, prior to committing it to a long lead time mask process. Kilopass has succeeded in making firmware mask free within a standard CMOS process, enabling configurable firmware without the downside risks of going to a non-standard CMOS process.

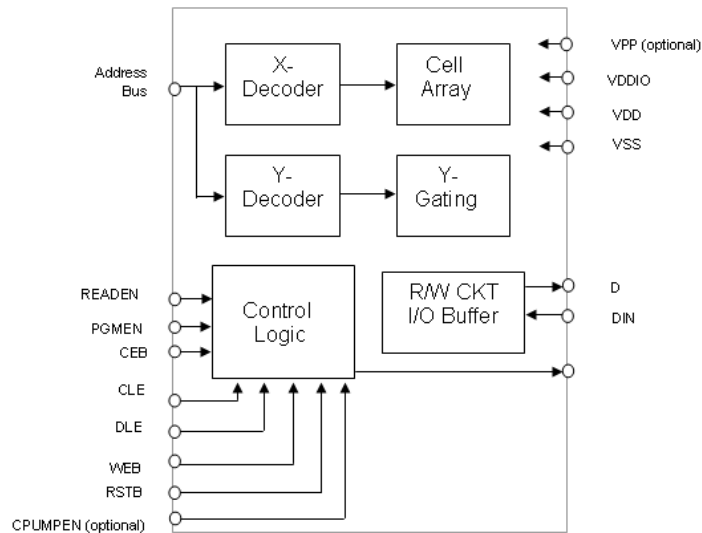


Figure 2. Typical XPM Memory Block

The opportunity to improve XPM technology for current and prospective Kilopass customers increases dramatically as process nodes advance, bit densities scale to 100 megabits, and performance catches up to today’s system clock speeds. Additionally, each customer’s investment in firmware is made secure through Kilopass’ Totally Secure™ technology. It is economically impractical for all but a very large government to reverse engineer an embedded Kilopass NVM block. Through the

invention of the XPM memory cell, Kilopass has enabled a new generation of configurable firmware in CMOS system design.

Kilopass XPM Technology & NVM IP

The Kilopass XPM product family is comprised of a collection of field and/or factory programmable memory blocks. These NVM blocks include a technology license for Kilopass' patented XPM technology as well as memory IP (Intellectual Property) in the form of a GDSII layout. These blocks come in a wide range of memory sizes, ranging from a very small 8-bit register up to blocks multiple megabits in size. The XPM memory arrays offer high-performance with a read access time of approximately 40ns typical access times for a 1 Mb memory block implemented in 0.18-micron technology with higher system performance achievable using interleaved instances. The memory array operates with standard CMOS logic operating voltages: 1.8V for memory blocks implemented in 0.18-micron processes or 1.2V for memory blocks implemented in 90-nm processes. In 90nm CMOS, the 1Mbit XPM memory has low power properties for both standby and active power consumption. For standby the block will consume a low 2 μ A, with active current a respectable 2mA. See Fig. 2 above for a typical XPM memory block configuration.




Figure 3. XPM Bit Register Configuration

XPM cells are also offered in a bit register configuration with 16 bits in a block; they may also be cascaded to form a multiple of 16 bits of storage. See Fig. 3 above for a diagram of three 16 bit XPM registers cascaded together. These small bit size memory functions may be used for other NVM features in conjunction with larger density memory blocks targeting firmware. Several typical small bit density functions include memory repair, serial ID, encryption key storage, and digital trimming for analog circuits.

Summary

The invention of the XPM cell in standard CMOS provides a new disruptive, enabling technology, freeing SoC designers from traditional mask-based non-volatile storage circuits. This makes possible the concept of firmware-based design utilizing a single SoC hardware design as a platform solution. The primary benefits are reduced cost and time-to-market as well as significantly reduced project risk while maintaining system performance and low power requirements. Additionally, as business continues to globalize with design efforts moving outside North America, design security becomes everyone's business. Whereas traditionally Defense based applications were primarily concerned about design security, it is now a requirement for most business enterprises. Firmware stored in a Kilopass memory array is virtually impregnable to reverse engineering and IP theft.



www.kilopass.com

Tel: (408) 980-8808

Email: klptinfo@kilopass.com