Introduction

Modern handheld applications provide consumers freedom from wall-power plugs or heavy battery packs. Yet for designers of light, handheld products, these same features pose some severe limitations on the choice of internal electronic components. Not only are application performance and battery life vital, but so are adaptability and re-usability for new designs, minimal board space, and low cost.

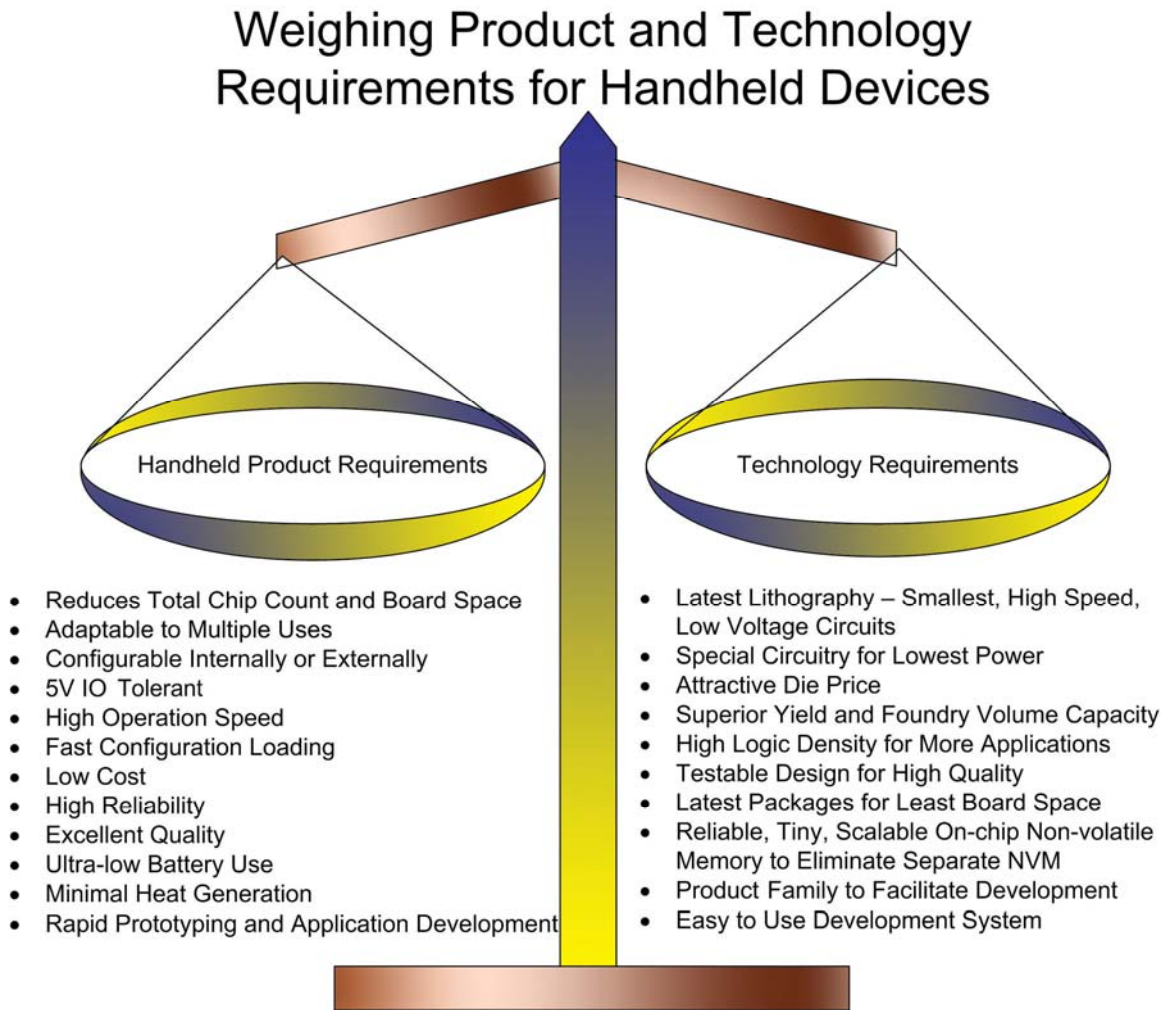

Figure 1Weighing Product and Technology Requirements for Handhelds

It is the excellent adaptability of FPGAs that makes them potentially attractive in handheld products. They can be used for "glue" logic, voltage translation between other components, IO expansion, interface and display control, and many other purposes. One

FPGA can even replace several existing board components, saving overall cost, weight, and complexity. Other factors, such as minimal board space and multiple voltage bank capability and 5V IO tolerance are important too for handheld applications. The left side of Figure 1 lists the attributes of an ideal FPGA for the handheld market.

Yet FPGAs available to date have not been widely used in handheld products. There are several reasons for this reluctance. Certainly low power consumption is a must in a handheld, so any FPGA for handheld applications must use the lowest battery power necessary. Performance cannot be sacrificed either – there are certain high-speed operations that the FPGA circuits must exercise. These are difficult tradeoffs. Those offering the highest priced FPGAs have generally focussed on obtaining the highest speed performance without concern for power consumption. Other FPGA offerings at the low price end may offer lower power by giving low logic cell and IO count. Neither choice is optimal for the handheld product designer, who really needs the full gamut of features.

For those reasons, SiliconBlue Technologies has designed a family of special FPGA products to meet handheld market needs left unaddressed to date. Our products offer advanced technology and speed, yet are carefully engineered to use the lowest battery power. Their advanced transistor dimensions enable fast, high logic gate count which in turn makes many diverse applications possible on the same chip. In addition to low power, circuit speed, and high logic gate count, these FPGAs also contain their own non-volatile permanent programmable configuration memory on the same chip. Having on-chip, non-volatile configuration memory monolithically integrated, our FPGAs configure rapidly at power-up, and save the board space and expense of external flash.

In creating its single-chip solution, SiliconBlue employed a range of unique technology enhancements. These enhancements overcome technical issues that arise from the latest CMOS logic processes, and which hamper the use of these technologies in battery operated applications. The solutions are adaptable to future CMOS technologies as they emerge, providing a foundation for even more advanced SiliconBlue products. This White Paper describes these technical issues and describes solutions that SiliconBlue has implemented.

Key features of SiliconBlue FPGAs:

- Latest technology
  - Special power-saving 65LP CMOS process
  - High density FPGA fabric for multiple applications
  - Small BGA package (~12 mm^2) frees board space
- Advanced Circuitry
  - Power-saving; cuts current to unused circuits
  - Minimal transistor gate and off-state leakage
  - Low core voltage for higher reliability and power savings
  - Multi-voltage 5V tolerant IOs with LVDS and MDDR support
  - High reliability logic and memory with on-chip error-correction
- Special on-chip non-volatile configuration memory (NVCM)
  - Fast, secure power-on configuration
  - No separate configuration memory chip needed
  - Small cell area: 0.2um^2 on 65nm
  - Scalable to future CMOS processes
  - No process or mask adders; built on standard CMOS
  - Proven yield and reliability

Technology Details

Figure 2 Required Technology Elements

In this section, we discuss various technological improvements in FPGA design employed by SiliconBlue Technologies in our high logic capacity, battery-saving products. The speed and density improvements provided by 65nm and smaller CMOS processes also come with issues that can limit their use in handheld products. These issues have been analyzed and overcome by SiliconBlue engineering, and we discuss them below. We also review the basic operating methods of the special NVCM memory that SiliconBlue has employed to store configuration on chip.

IC standby current

Principally because its area and cost can be made predictably smaller year by year, and its reliability and performance actually improve as this shrinking is achieved, MOS

transistor technology has long dominated the semiconductor logic market. The reason that MOS technology dominates, rather than bipolar technology, is that an existing, working MOS circuit can be made smaller, faster, and cheaper with predictable effort. MOS technology is basically a process of printing or lithographing regions of material carefully registered to underlying layers. Year after year, the electronics industry finds ways to lithographically print more and more transistors into a silicon chip of the same area – "scaling" the transistors to smaller size.

Yet as a consequence of scaling to smaller dimensions, the latest semiconductor processes have encountered a problem not seen in earlier process generations –significant current is wasted. Two new mechanisms are to blame: off-state source-to-drain transistor leakage, and leakage from the gate-to-channel directly through the thin dielectric. In the past, the current leaking unproductively in CMOS transistor circuits arose mainly from DC power lost in reverse-biased PN junctions, and from AC capacitive losses during transistor switching. The PN junction leakage current lessened as the transistors and associated PN junction areas became smaller with each generation of CMOS process. Reductions in AC switching power losses were found by using special low-K dielectrics, and lowering operating voltage. But it has been the scaled-down dielectric thickness causing gate leakage, and the scaled-down transistor length resulting in off-state transistor current leakage, that required new technical solutions.

Let us take the issue of gate dielectric leakage first. The cause is that charge can move or "tunnel" directly through the "insulating" dielectric that separates the gate terminal from the underlying source and drain terminals. Figure 3 illustrates the issue. In an NMOS transistor which is turned on, the positive gate voltage pulls electrons from the substrate silicon to the oxide-substrate surface. But as shown in the diagram, in earlier process generations the gate dielectric was simply too thick, and its energy barrier too high, for the charges to move to the upper gate material from the silicon channel.

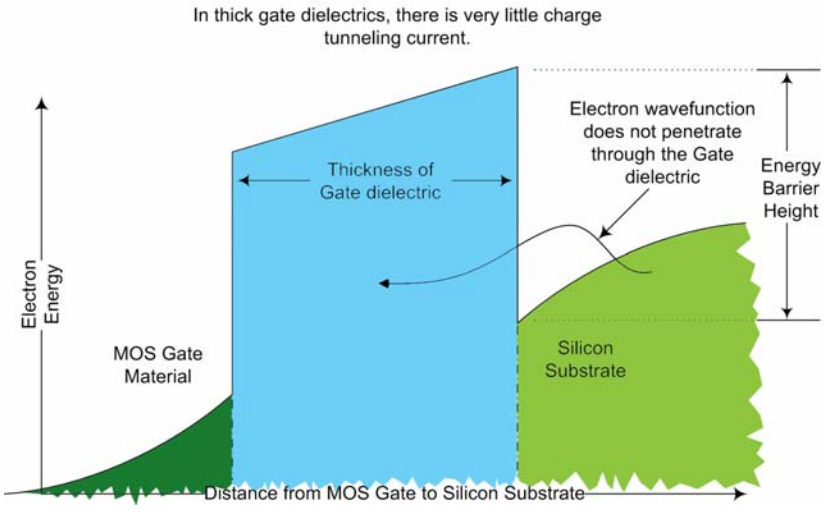

Figure 3 Negligible tunneling current if gate dielectric is thick

But in the latest CMOS technologies, the "oxide" material is so thin, and its energy barrier lower, and the electric fields across it so high, that by the tunneling process, charges can disappear from one side and reappear on the other. Figure 4 shows this situation. The electrons in the substrate have a wave-like nature, shown in the diagram by the illustrated wavefunction. A tail of the wavefunction extends into the gate dielectric, and that tail comes out the other side into the MOS gate material. The penetration of the wavefunction tail means that some fraction of the electrons in the substrate will jump into the gate material. Once in the gate material, these electrons are pulled out by the positive gate bias supply and result in a current. As this current serves no purpose, it will just waste battery power in a handheld product.

Thus, quantum-induced current leakage through the gate dielectric can dramatically increase the chip current. The mechanism of quantum tunneling has been understood since Fowler and Nordheim first explained it in 1928 [i], but it has proven to be a very complex and delicate problem for both modern CMOS foundries and for electronic circuit designers to solve. Ordinarily, to eliminate the tunneling current, one could (1) reduce the electric field inside the gate dielectric by lowering gate voltage, or (2) make the dielectric physically thicker, or (3) increase the energy barrier height. But we also want to minimize AC current use which comes from ½ C V^2 – so we want to lower core Voltage V and interconnect capacitance C to do that.

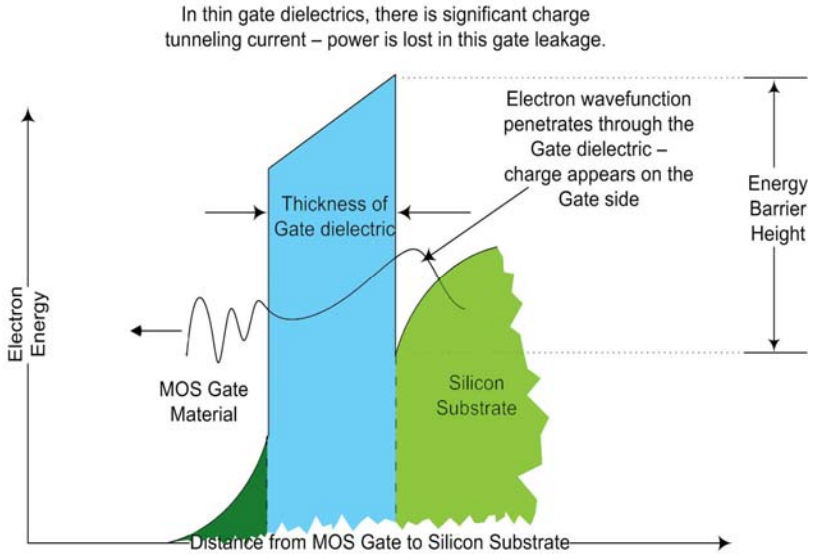

Figure 4 High tunneling current if gate dielectric is thin

To optimize circuit speed with battery conservation, three factors must be balanced. These are the thickness of the dielectric, its energy barrier height, and its dielectric constant K. Raising K permits the gate dielectric to be made thicker. However, the chemical elements which raise K also tend to lower the energy barrier height. The

challenge is to find the proper combination of these factors. The CMOS process that SiliconBlue has chosen contains elements in the gate dielectric that increase the dielectric constant yet retain a high energy barrier height. As a consequence, gate leakage current is greatly reduced, and the transistor speed is held high.

Lowering Transistor Off-State Leakage

Above, we explained SiliconBlue Technologies choice of CMOS process to minimize gate dielectric leakage current. Careful control of off-state leakage is also critical in reducing product leakage current. Recall that transistors which are nominally "off" can still leak current between their source and drain if these are too close together. The distance between the source and drain is known as the gate length. SiliconBlue has carefully studied the contribution of transistor off-state leakage in comparison to both gate leakage current and PN junction leakage. As Figure 5 below shows, off-state leakage is strongly influenced by the gate length. There is an optimal gate length that minimizes product off-state leakage. SiliconBlue has carefully studied these factors and utilized gate lengths that minimize product power consumption.

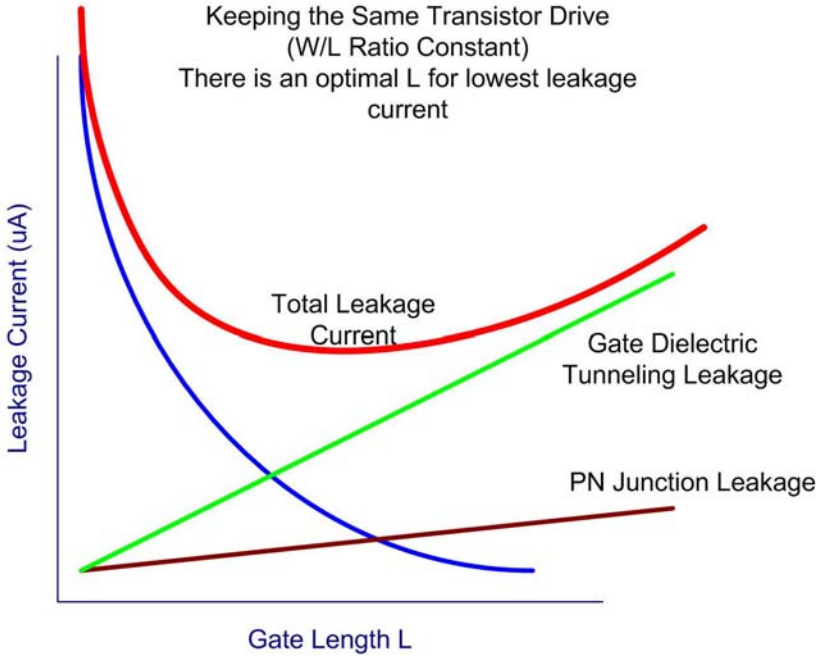

Figure 5 Minimizing Off-State Leakage Current

These combined improvements in CMOS process to lower dielectric leakage, and careful choice of transistor lengths to minimize off-state leakage, result in the very low power consumption of SiliconBlue products. For example, the SiliconBlue iCE65L04 product achieves a typical standby current of only 25µA. At the 65nm lithography node, circuits are condensed and SiliconBlue is able to construct very advanced FPGAs. Just as

importantly, SiliconBlue has been able to integrate a special non-volatile memory to be used for on-chip configuration (NVCM). We discuss this NVCM memory in the following section.

FPGAs require configuration

Modern FPGA devices are configured into a logic system using on-chip SRAM memory. The advantage of the SRAM is it provides very fast operation. Unfortunately, when the chip is powered off, the SRAM configuration is also lost. The configuration may be imported from outside the FPGA, or it may be stored in a separate chip residing inside the same package, or it may be integrated on-chip. In the first two cases, it takes two chips to make the FPGA operate. The resulting two-chip approach adds die and assembly cost, and raises packaging and inventory costs. Due to inherent factors associated with assembly and testing, the dual die approach also increases manufacturing risk and lowers yields, leading to a higher overall cost. It also makes the application less secure from cloning or bit-stream snooping, as communication between the two chips can be picked off and studied.

The third approach is more ideal - have the FPGA itself capable to turn off its power, yet retain its memory configuration. There are several ways to do this. One method is to utilize an embedded Flash or EEPROM CMOS process to make the FPGA, so that on-chip memory for FPGA configuration can be programmed and erased. But the CMOS processes needed for either EEPROM or Flash memory fabrication require special wafer processing and consequent engineering, which are costly. These charge storage memory types also pose certain reliability issues. Altogether, the approach of using on-chip Flash or EEPROM memory consequently results in a more expensive process, with associated delays in product development and resulting time-to-market.

SiliconBlue On-Chip Configuration

SiliconBlue has also incorporated an alternate way to provide power-on configuration which does not require modification of the CMOS logic process, nor adds assembly or interconnection costs. This method is to construct a non-volatile memory from the existing CMOS process. SiliconBlue Technologies has been fortunate to have license to the Kilopass XPM ™ non-volatile memory, which is of this type[ii]. Kilopass has pioneered and patented the use of this technology since 2001, and has established and demonstrated that it is durable and reliable. SiliconBlue Technologies has implemented this memory efficiently using an internal serial interface, specially suited to FPGA integration, and is referred to as our "Non-Volatile Configuration Memory" or "NVCM"™ for short.

SiliconBlue NVCM Principles

Before we delve more deeply into the method of NVCM operation, we must first review the general principles of how dielectrics change from insulating to conductive.

It is the role of semiconductor engineering to measure, mathematically model, and finally to assure that the transistor gate oxides endure their intended use voltage and lifetime. Before the 1980s, semiconductor engineers commonly thought that a transistor gate dielectric would break down instantly, if the voltage across it exceeded a certain electric field limitation[iii]. That is, below that "breakdown field" the oxide would never break, and at or above it the oxide would break instantly. Since the 1980s, however, research has shown that the breakdown does not take place instantly. And in fact, the oxide will break over some time with any given electric field across the oxide – it just takes longer and longer as the voltage is lowered. Several mathematical functions have been proposed to relate the survival time before breakdown to voltage and temperature. One of the common formulas in use today is a simple power-law for the effect of voltage, so that the lifetime depends inversely on applied voltage to a certain power. For temperature, a Boltzmann factor would express the result that the lifetime is also shortened as temperature rises.

Moreover, several different behaviors can be observed in the final dielectric conductance after breakdown occurs. The behaviors fall into several general regions depending on the energy discharged in the programming pulse, as illustrated in Figure 6 below. In the 1980s, when aluminum was used in researching gate dielectric behavior, one could actually see where the discharge occurred – the aluminum gate conductor would be vaporized. One could see small aluminum holes through a microscope – each hole was created when the electrical discharge would penetrate the silicon dioxide gate dielectric.

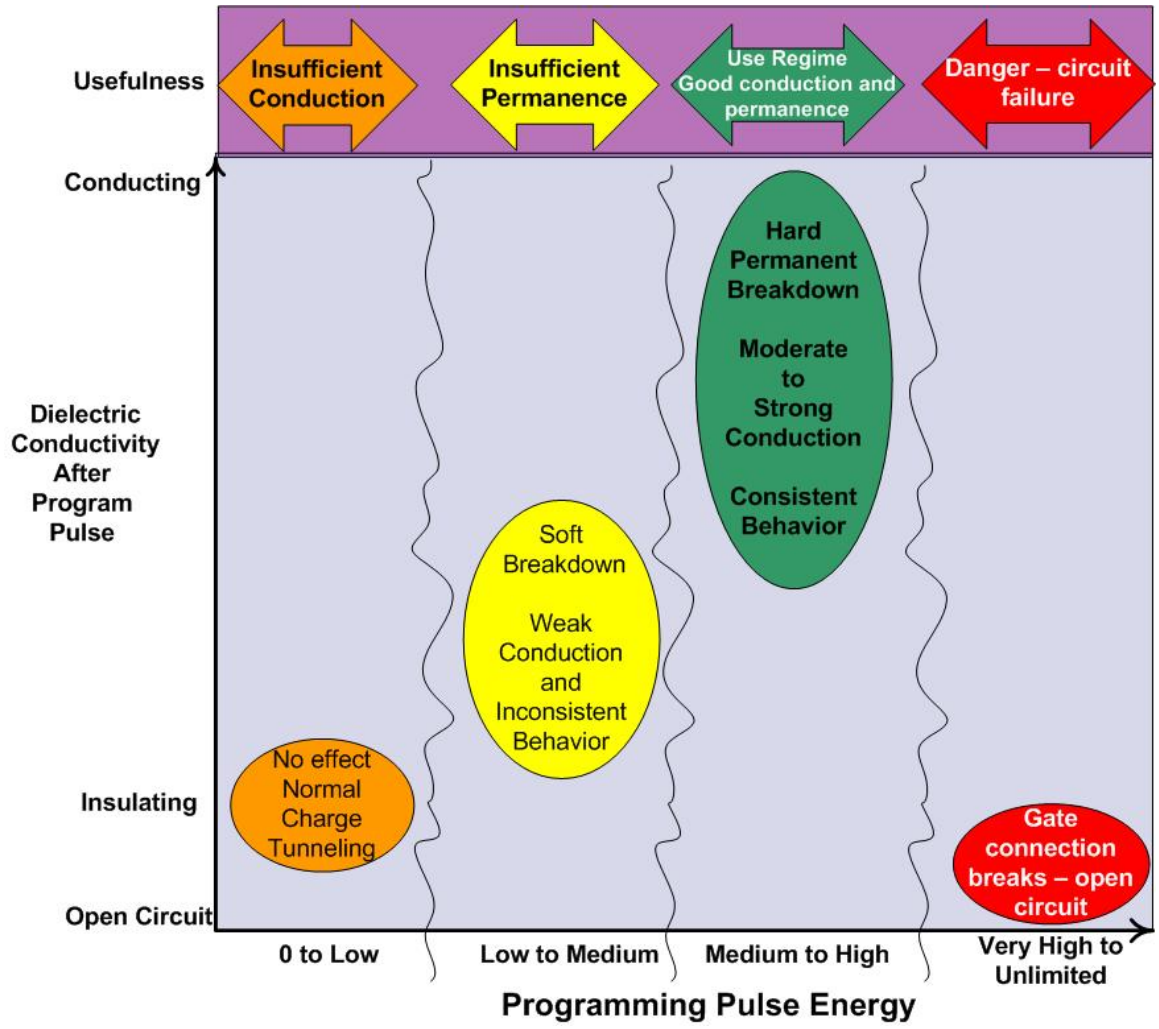

Figure 6 Optimizing Programming Conditions

In modern technologies, aluminum is no longer used for a gate conductor. Instead, highly doped silicon coated with a high temperature refractory-metal silicide conductor is used today. The silicon and its refractory-metal silicide coating have much higher melting and vaporization temperatures than aluminum. Consequently, when an electrical discharge occurs, these materials may either remain intact, or may melt, or may vaporize, depending on the level of energy per unit volume they absorb.

In Figure 6 above, we show there is a range of optimal programming energy that can be used to program a logical "1." Too low a programming energy has no effect on dielectric conductivity – the cell will remain in its "0" state. Raising the programming energy over some threshold can damage the dielectric into a very weakly conductive or "Soft Breakdown" state. That state is not desirable for a memory as it can fluctuate in resistivity and its read current is weak. A third region exists in which the memory cell is

highly conductive and stable – that is the desired state for a memory application. Finally, at too high a programming energy, the gate or its wiring may melt or vaporize, leading to an open circuit. This latter case must be avoided – the memory sensing circuit will detect no current flow and in that case fail to register a "1" value properly. Consequently, the energy dose delivered during cell programming must be accurately determined and controlled.

In order to direct the programming energy to just the right location, an array of two-transistor memory cells is used. A small part of an array is shown in Figure 7. One of the transistors, labeled WP, is used to store the digital information. If the data is "0" this transistor is left intact; if a "1" it's gate is programmed to be conductive to the underlying transistor channel. A second transistor, labeled WR, serves to isolate the WP transistor from the column bitline, shown vertically in Figure 7. During programming and reading, the biases on all rows and columns are carefully controlled to prevent damage.

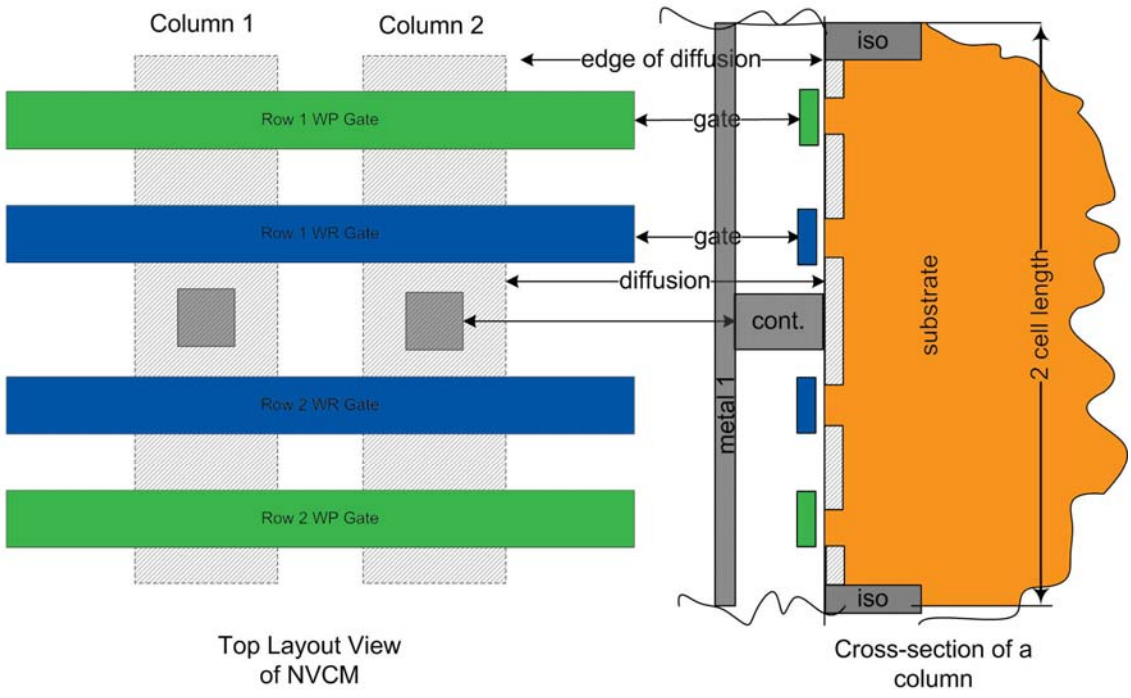

Figure 7 NVCM Memory Cell Array Structure

There were two major considerations to developing the SiliconBlue NVCM memory successfully. For one, the memory cell must program rapidly. This implies that the gate dielectric breakdown time is achieved in a short time and which requires a particular voltage. Another requirement is that all other cells must remain unaffected. That means the voltage and time exposure of other array cells must be low. To accomplish these

goals, the memory cell nodes are carefully biased, so that only the cell being programmed receives the required voltage and time needed for programming.

In order to assure these conditions, gate dielectric breakdown research is utilized. This research has established the very strong relationship between gate voltage, temperature, and the time that a particular dielectric will last without measureable damage. The time through which the dielectric remains insulating and intact, until it becomes conducting, is stated as the "TDDB" or time dependent dielectric breakdown time. Figure 8 shows the breakdown behavior of a range of dielectric thickness and temperature. We see from the graph how rapidly the TDDB lifetime falls as voltage is raised. Focusing on the leftmost curve, we observe that dropping from a lifetime of 10 years (3E8 seconds) to 1 microsecond (1E-6 seconds) only takes doubling the applied gate voltage from 2.3 to 4.6 volts.

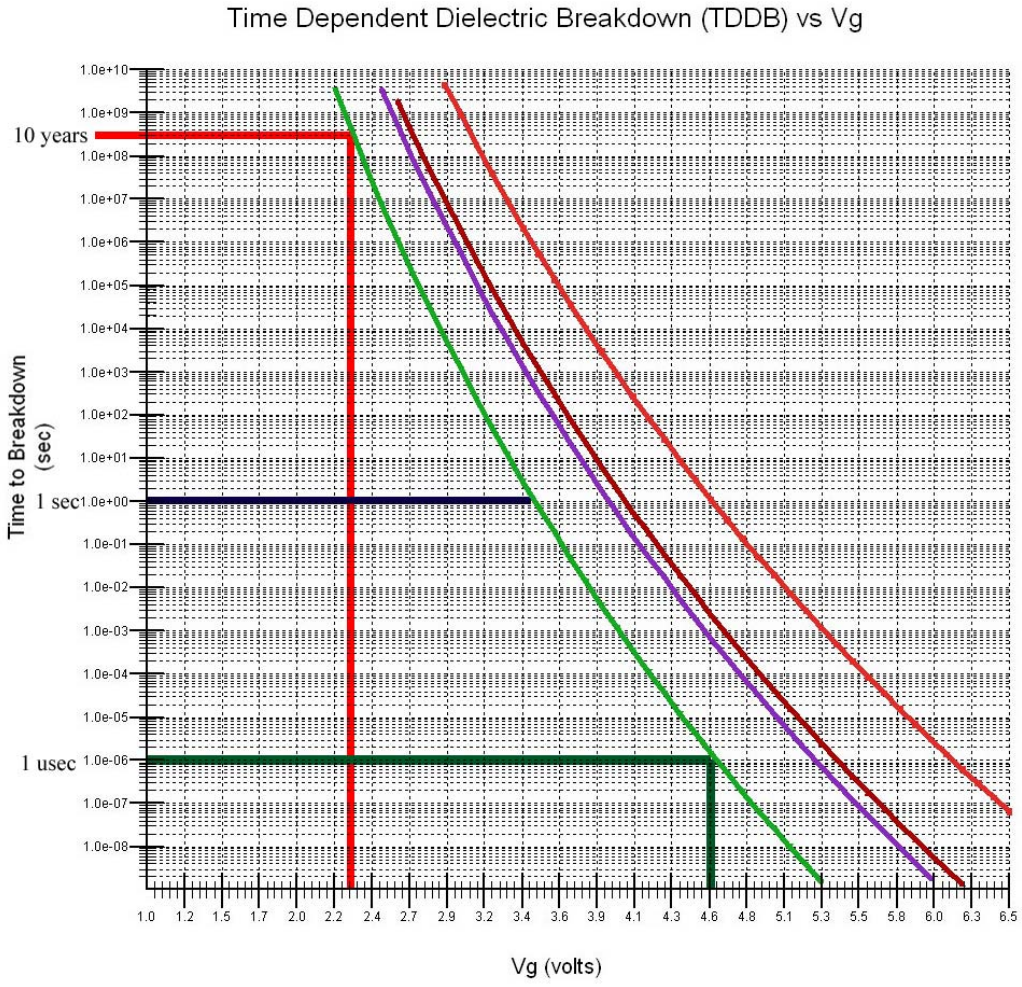

Figure 8 Gate Dielectric Time to Breakdown Predictions

As we have stated, SiliconBlue Technologies NVCM memory cell uses controlled electrical change of transistor gate dielectric from insulator to conductor as the basis of the memory. A highly controlled electrical discharge through the gate dielectric material causes it to change from insulating to conducting. Mechanisms for this change have been studied by earlier research groups. The change can be invisible to optical detection. For example, published research shows that an electrical discharge can modify the atomic bonding of an insulator, so that it can behave more conductively, without significantly moving the insulator atoms[iv]. Chemical bonding calculations by M. Kimura show that rearrangement of Si-O bonds into Si-Si and O-O bonds can dramatically reduce the silicon dioxide energy gap by about 4eV. As the atoms of the insulator need not move appreciably in changing conductivity from insulating to conducting states, the programmed cells can appear physically identical to unprogrammed cells by optical and SEM means. Having undetectable data provides strong configuration data security.

A typical NVCM memory cell cross section and current-voltage characteristics are illustrated in Figure 9. The graph of Figure 9 shows the cell current-voltage behavior in its "0" and "1" states. In the case of a "0" bit, the dielectric under the WP gate is intact and insulating. In the "1" state, a small conductive region exists between the WP gate and the underlying silicon. The conductive region allows current to flow from the WP gate, under the WR transistor, to the bitline. Circuits connected to the bitline sense the current and decipher the level into a logical "0" or "1." Essentially zero current flows in the "0" state, while substantial current is present in the "1" state. Voltages used for reading the current are kept low, so that cells remain reliable for well past 10 years of constant reading. Likewise, programming voltages are applied through a patented process, carefully engineered to create the conductive channel only in the single cell being programmed.

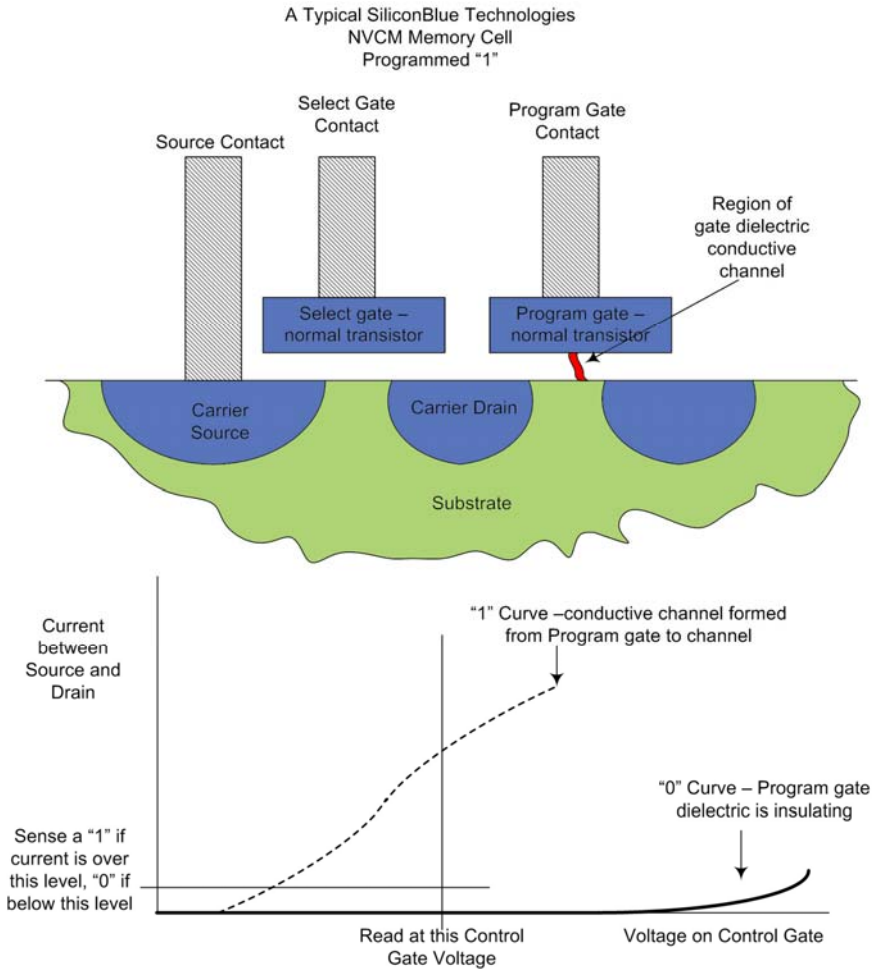

Figure 9 NVCM Transistor Cross-Section and Read Currents

Reliability

Having a well-performing and cost-effective product is really of no practical value unless the product also performs as expected throughout its lifetime. Both the initial performance (Quality) and the operation at the end of product life (Reliability) are critical factors in the success of the higher complexity hand-held products built upon SiliconBlue devices. For this reason, SiliconBlue has paid close attention to the inherent reliability of its products from the process and design phases to final product testing.

Studies show that the NVCM memory cell and its associated circuitry are more reliable than alternative NVM memory. Tests show that programmed and unprogrammed cells will not switch electrical state during the product lifetime. As NVCM memory does not store charge, it is immune to mobile ion contamination, UV light, and similar ionic and radiation exposure which can affect floating gate and SONOS type memory. The

conductive dielectric region is constrained to very small volumes, as the gate dielectric is below 40 angstroms thick, and in practice the conductive region cannot be detected by optical or SEM means, configuration data is highly secure.

Historically popular NVM memory methods include floating gate memories (Eprom, EEprom, NOR and NAND Flash), and oxide-nitride-oxide (SONOS) types. All these types of memory utilize electrical charge stored on or near a transistor gate in order to affect the transistor conductivity. By this means, different transistor conductivity levels can be used to represent "1"s and "0"s.

Consequently, mechanisms that allow trapped charges to escape, or to become negated by opposite charges, reduce the reliability of the memory over time or at elevated temperature. Several mechanisms are known to cause charge loss – among them are direct thermal excitation, net charge loss by mobile ion contamination, UV light exposure, and quantum tunneling losses. The quantum tunneling concern alone constrains the insulators around the charged node to stay over certain thicknesses to prevent charge loss, and this requirement hampers scaling of the memory to smaller dimensions. Accordingly, the use temperature range and guaranteed data integrity time have been lowered in the most recent generations of non-volatile memory products.[v]. So unfortunately, as temperature increases, or if mobile ion contaminants are present, or there are leakage paths in the insulators surrounding the trapped charge, stored charge necessary for digital data can become lost. And as the physical structures become smaller in each technology generation, the insulators surrounding the data charge storage capacitor become thinner and can become more prone to leakage.

In contrast to charge storage based non-volatile memory (NVM), the SiliconBlue NVCM depends instead on dielectric conductivity modification. The NVCM dielectric conductivity mechanism is free from mobile ion contamination concerns. Extensive reliability tests show that unmodified gate dielectric, utilized to create the logical "0" state, has outstanding reliability under NVCM read conditions. Studies also show the NVCM logical "1" state is also reliable. Research shows that programmed memory cells become slightly more conductive during a high temperature bake tests, and withstand thousands of hours at 150C – they do not revert back to the insulating "0" state.

Further ensuring the reliability of its configuration memory, SiliconBlue Technologies prevents against unforeseen random errors using well-established error correction methods. These methods ensure that single bit errors of either the "0" or "1" state are corrected. This additional engineering for reliability assures that throughout the intended life of our products, the NVCM memory used to configure the FPGA will remain accurate, and the FPGA will function properly.

Summary

SiliconBlue is the first company to combine several advanced technologies especially for the portable, hand-held market. The first advancement is the adaptation of a very high density, high speed 65nm low power CMOS process. That low power process was

specifically developed to help eliminate excessive leakage currents inherent in the general high-speed 65nm process. The second innovation is our adoption of a novel non-volatile configuration memory (NVCM) built directly on that same low leakage CMOS processes. On-chip configuration memory gives our products an "instant-on" capability and further saves board space and component count. The third element of our design is an advanced static RAM (SRAM) FPGA technology. At the 65nm process node, our FPGA devices possess many times the logic gates of competing FPGA products. In turn, the high gate count enables customers to incorporate a wide range of applications on a single chip. These diverse applications on a single SiliconBlue chip can replace many single-applications chips, making mobile products smaller and longer-lasting. Separate descriptions of various SiliconBlue applications can be found on our website.

In the above discussion, we have outlined in detail the numerous technology and design advancements built into SiliconBlue Technologies ultra-low power FPGA products. These included the special Non-Volatile Configuration Memory (NVCM) which permits SiliconBlue not only to incorporate "instant-on" configuration, but also to rapidly advance to the latest CMOS logic processes as they become available. Moreover, the special circuit techniques and design methods that make our products the most suitable for battery – operated applications are also extendable to future FPGA products. Finally, our use of enhanced features, like error-correction circuitry (ECC) assures customers that our components will remain functional and reliable throughout product lifetimes. In total, these design and process advancements provide SiliconBlue FPGA products maximum value to customers.

---

[i] R. H. Fowler and L. W. Nordheim, "Electron emission in intense electric fields," **Royal Soc. Proc**., 1928, **119**, pp 173-181.

[ii] For an explanation of the Kilopass XPM memory, see for example, "A novel embedded OTP NVM using standard foundry CMOS logic technology," J. Peng, G. Rosendale, M. Fleisler, D. Fong, J. Wang, C. Ng, Z.S. Liu, H. Luan, 2006 IEEE Non-Volatile Memory Workshop, Monterey, CA. pp 24-26.

[iii] Breakdown field is defined as the voltage across the dielectric, divided by the thickness of the dielectric. Textbooks such as the 1981 edition of S. M. Sze "Physics of Semiconductor Devices" still contained a stated dielectric strength (breakdown field) of silicon dioxide as 10 million volts per cm. So a dielectric that is 50 angstrom thick would last if the voltage across it were below 5V, and would break if the voltage exceeded 5V. The concept of a breakdown field has been rejected, as we now understand that time and temperature play critical roles in breakdown as well.

[iv] M. Kimura, "Oxide breakdown mechanism and quantum physical chemistry for time-dependent dielectric breakdown," 1997 Reliability Physics Symposium, April 1997, pp 190-200.

[v] One may refer to the datasheets of major suppliers of semiconductor memory for the upper temperature limits. Spansion, for example, which employs the SONOS type of memory, or of Intel for NOR Flash, or of SanDisk which employs NAND flash. Spansion's S29JL064H for example is a 64Mb floating gate Flash memory on 0.13um, and has a maximum temperature range from -40C to +85C. Earlier technologies

based on 0.23um or larger lithography show a maximum range from -55C to +125C. Likewise the Spansion S29AL016J is a SONOS "Mirror bit" memory on 0.11um lithography, with the same -40C to +85C range. SanDisk's Extreme memory product, based on NAND memory, is rated from -25C to +85C.