

Designing for Next Generation Wireless Applications With SiGe

By
Richard Sfeir
Director of Marketing
CommQuest Technologies / An IBM Company

Over the past decade, the semiconductor technologies used in communications systems have been undergoing something of a forced divergence – between the high-integration capabilities of silicon-based processes and the high-performance possibilities of exotic processes like gallium arsenide (GaAs). Largely because of the inherent disparities between these processes, it even appeared that communications-oriented semiconductors might be finally approaching their practical limits in terms of both size reduction and performance improvement. For many applications, low-cost, high-volume silicon processes have been successfully used throughout the 1 to 2 GHz frequency domain, however for new RF applications that require much higher speed circuit operation, such as 30 GHz, standard silicon processes fall far short. On the other hand, compound III-V semiconductors such as GaAs, have been successfully implemented in these ranges, however at significant additional expense due to their exotic process requirements.

As a cost-driven market arena that intrinsically requires both performance and a high-level of integration, the next generation of mobile wireless devices was literally dependent upon finding a cost-effective way to re-converge these capabilities into a unified semiconductor process. Many industry experts believe that the answer has now arrived in the form of Silicon Germanium (SiGe) process technologies.

Standing at the Great Divide

Over many years of continuous process refinements, the cost advantages of silicon have been both proven and perfected. Initial substrate materials are much less expensive than III-V materials and can economically be manufactured in wafer sizes as large as 200mm to maximize production efficiency. Silicon's basic process stability and compatibility with enhancements such as CMOS, have led to the continued refinement of robust, proven fabrication methods that seamlessly combine high levels of chip-level integration with cascading cost curves. Overall silicon has proven to be a very high-yielding material, as compared to substrates made from less stable materials, such as GaAs, and therefore established itself over two decades ago as the mainstream semiconductor process standard. The very fact of silicon's popularity over the years has served to focus most industry efforts on continued process improvement, thus making it even more efficient and further widening the cost gap between silicon and other more exotic processes.

Conversely, GaAs has presented a number of cost-challenging process problems from the outset. First, it makes use of a basic elemental combination that is extremely rare on the one hand and quite toxic on the other. Secondly, the mechanically-fragile nature of GaAs limits the practical wafer-size on the front-end and often requires special packaging methods on the back-end. As with most III-V compounds, GaAs is hard to grow in large diameter, high-purity crystalline forms, with its inherent higher number of defects per mm^3 consequently driving up production costs. From a circuit design standpoint, the lower heat transfer rates of arsenic-based substrates limits both GaAs' packing densities and power handling capacity.

However, despite its higher cost, GaAs has steadily gained widening usage in both communications and military applications because of its significantly higher performance capabilities. Because GaAs inherently provides greater low-field electron mobility and has a lower saturation field than silicon, heterojunction bipolar transistors (HBTs) created in GaAs can run at much higher speeds than silicon circuits. In addition, the fact that GaAs substrates have an energy gap four orders of magnitude greater than silicon, allows for dramatically reduced parasitic capacitance, further boosting its speed potential.

Because of their significantly higher frequency capabilities, GaAs processes first gained popular usage in many military applications, where performance demands clearly overrode cost considerations. Subsequently, as GaAs processes began to edge down the cost curve, a number of targeted applications began to make their way into commercial wireless designs. For instance, in cellular handsets GaAs has been found to be particularly well-suited for specific radio front end ICs, such as Power Amplifiers (PAs) and Low-Noise Amplifiers (LNAs), while other circuits, such as mixers, modulators/demodulators, and VCOs are generally still produced via more cost-effective Si Bipolar processes.

Even though these inter-device functional divisions might not have always been optimal, the basic incompatibilities between GaAs and Si processes pretty much dictated the architectural partitioning. Because GaAs is inherently too costly for use throughout the handset and silicon comes up short on some specific performance requirements, the only available path has been to opportunistically-optimize wherever possible, leading to the evolution of process-driven rather than functionality-driven architectural decisions.


For previous generation handsets, this deployment of various different ICs using separate process technologies has served reasonably well to balance the inherent tradeoffs between performance and cost requirements. However, the coming third generation of mobile wireless devices will have to push the limits of both high-performance and IC integration well beyond the boundaries of today's processes. Market requirements for dual-band and even tri-band capabilities will become common place to meet user requirements for regional and global roaming. At the same time the relentless user demand of longer battery life will push power limitations even further downward.

To meet these demands in next generation handsets, the only viable path forward will be to use higher levels of IC integration, pulling together complete functional blocks of the handset architecture, such as the baseband and radio sections. The refinement of SiGe processes to their current production-capable levels is certain to play an important role in the near-term achievement of these integration goals, without sacrificing any of the required performance objectives.

The Fundamentals of SiGe

The key advantage of Silicon Germanium is that it is fundamentally a higher speed silicon process, thereby offering maximum leverage from existing silicon fabrication processes. By doping the silicon (Si) substrate with germanium (Ge), SiGe creates supercharged HBTs that can operate at 65 GHz as compared to 15-25 GHz for best-of-breed silicon-only processes. Because SiGe uses only one additional processing step beyond Si CMOS, SiGe is able to match or even exceed the performance of III-V materials, such as GaAs at a fraction of their cost. Some industry comparisons have predicted as much as a 70% cost savings, with SiGe epitaxy costing in the range of \$0.60 per mm² and GaAs epitaxy running about \$2.00 per mm² (as compared to about \$0.01 mm² for Si CMOS). *(These are fairly old costing numbers from about 1995 – do we have anything newer or should we just omit the reference?)*

The basic advantage of SiGe lies in the ability to selectively re-engineer the band gap of ordinary silicon, thereby increasing the transistors' inherent switching speeds. Conventional silicon has a fixed band gap of 1.12eV (electron-volt), however by grading the Ge content across the transistor base region, the band gap performance can be significantly and selectively enhanced. Figure 1 demonstrates the effects of this graded Ge doping by comparing a) the re-engineered band gap profile that results in a ramped potential surface and b) the linear grading of Ge atoms across the base region to produce the re-engineered band gap.




The SiGe HBT also has higher current gain and improved output conductance, which makes it particularly adaptable for the high speed analog circuits needed in microwave and RF applications. SiGe transistors also have demonstrated excellent noise characteristics over a very wide range of temperatures, making them ideal for applications such as PAs and LNAs. Additionally, since SiGe HBTs typically are not required to run at full 65 GHz frequencies, there is ample performance headroom, which can be traded off for operation at lower power levels. Such energy-saving potential can be vital in meeting user demands for longer handset battery life.

From a production standpoint, the potential semiconductor benefits of SiGe compounds were clearly apparent to researchers as much as a decade ago, however the evolution of that potential into a viable, qualified production process required the resolution of some critical issues. Because germanium has a lattice-like crystalline state structure that is quite similar to silicon, it presents a natural opportunity for doping onto a silicon substrate while still maintaining a highly stable structure. Although germanium is slightly harder and more brittle than silicon, the most significant challenge to growing reliable structures is the fact that the germanium atom is 4% larger. This size difference means that the SiGe compound layers are naturally strained and must therefore be carefully grown within critical thickness levels. On one hand, the strain between layers is an important aspect of varying SiGe's semiconductor properties while on the other hand, if uncontrolled it has the potential for mechanical instability.

The technique most suited to effectively growing SiGe structures was ultimately proven to be ultra-high vacuum chemical vapor deposition (UHV-CVD), developed at IBM Research in the early 1980s. Today, IBM's refinement of the SiGe fabrication process allows new UHV-CVD techniques to epitaxially grow and integrate robust SiGe HBTs in conjunction with standard silicon Bi-CMOS circuitry, according to sub 0.5 micron design rules. The ability to incorporate both deep and shallow trench isolation allows for maximum design flexibility, such as merging RF and logic onto a single die, while also greatly improving overall device density and performance.

Figure 2 shows a SiGe HBT cross-section with both deep and shallow trench isolation and two levels of planar metal interconnection



These advanced SiGe process technologies have been pioneered and production qualified on 200mm diameter wafers at the IBM Advanced Semiconductor Technology Center in Hopewell Junction, NY. Focusing upon SiGe's applicability to the wireless design arena, several RF-oriented circuit designs have already been created and qualified, including voltage-controlled oscillators (VCOs), power amplifiers (PAs), low-noise amplifiers (LNAs), mixers, digital delay lines, etc. In parallel with these component level developments, communications-specific developers such as CommQuest Technologies have been working with both IBM and handset manufacturers to optimally leverage SiGe advances into system-level handset architectures.

Real-world SiGe Applications in Mobile Handset Architectures

A good example of SiGe's capacity for enhancing mobile handset architectures is CommQuest's recent introduction of new SiGe-based Tri-band LNA and PA/VCO components. Earlier in 1998, the introduction of a Tri-Band GSM chip had already significantly reduced overall parts counts. However the continued need to use inflexible GaAs processes for creating individual PAs and LNAs, dictated that a dual or tri-band handset design still had to use completely redundant radio transceiver pipelines for each band, including all of the associated support circuitry. By leveraging SiGe's combination of performance and silicon integration capabilities, these new chip designs now allow the merger of circuitry for all three bands onto single LNA and PA chips.

From the handset manufacturers' standpoint, this usage of SiGe has dramatically reduced chip counts by over 60% (from between 250-300 components down to 100 or less). In addition by leveraging SiGe's inherent low-power operation, the new Tri-Band PA can achieve better than 60% power-added efficiency, as compared to approximately 50% for best case GaAs-based PAs. These overall energy savings can directly translate into end-user talk-time increases of as much as 20% or more.

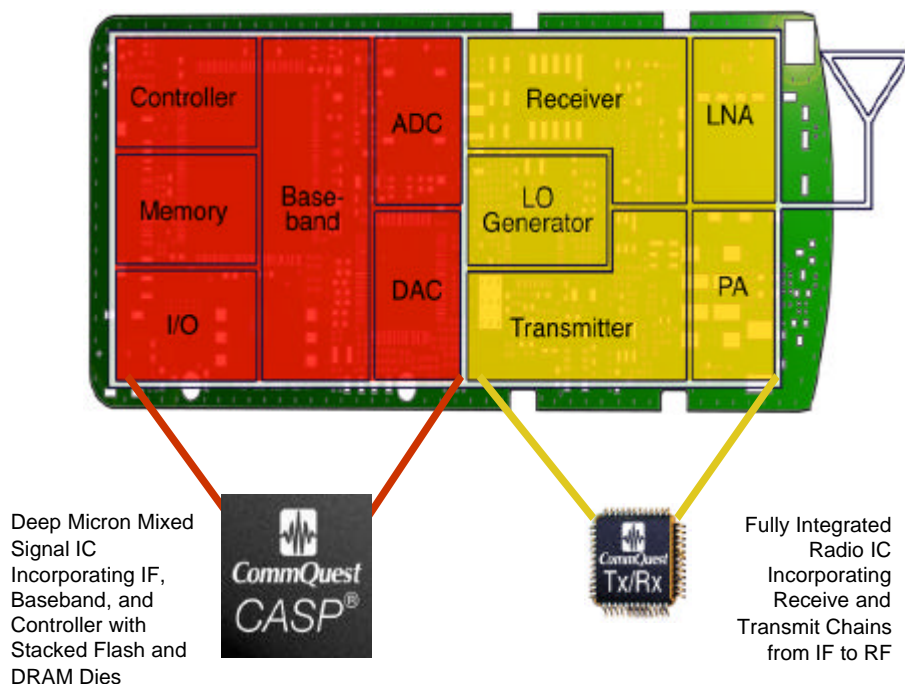
It is clear that the next steps in IC-level integration of wireless handsets will continue to make use of the performance and cross-integration capabilities of SiGe, however it will also require a fundamental shift in overall architectural partitioning and design. Obviously, a successful evolution to the oft-stated design goal of handsets built completely around one or two chips cannot be achieved solely through the benefits of SiGe.

One of the critical requirements will be simplifying the entire baseband processing section so that SiGe's integration leverage can be optimally brought into play. Traditionally, handset baseband designs have been built around general-purpose DSPs, in conjunction with a separate microcontroller and mixed-signal devices for handling IF functions. With next-generation digital wireless systems requiring even more and faster internal processing by handsets, the limitations of segmenting designs around general-purpose DSPs have now been surpassed. Instead a new breed of communications-specific IC architectures are now emerging, such as CommQuest's CASP

(Communications Application Specific Processor). By using modular designs that combine communications-specific functional blocks, these new IC-level architectures are able to better balance the system requirements for performance, power-conservation, size, and cost.

Optimally, the architectural decisions regarding the chip-level integration of these functional building blocks should be completely process independent. However, in the past, the incompatibilities between silicon-based and GaAs-based processes could often force sub-optimal IC partitioning decisions in order to achieve needed performance through separate GaAs components. Now, with SiGe's potential for seamlessly integrating both high-speed RF and digital logic into single-die implementations, designers can make optimal architecturally-driven choices rather than live with sub-optimal, process-driven dictates.

By opening the door for more rational IC integration decisions along with the ability to cost-effectively implement them using proven silicon-based processes in existing 200mm wafer fabs, SiGe may truly usher in the era of 1-2 chip phone designs. In the near term, the most likely partitioning will be between 1) a fully integrated radio IC, making liberal usage of SiGe to implement all receive and transmit chains, and 2) a mixed-signal IC handling all of the intermediate frequency (IF), baseband, controller and memory functions. (See Figure 3).



As SiGe moves further into mainstream usage for commercial wireless handset integration, the breadth of its process capabilities will also be expanded and refined, building on top of the rich base of already proven silicon processes. SiGe HBT designs

will be scaled both vertically and horizontally, allowing for greater overall design flexibility and circuit densities. At the same time, the combination of future SiGe and CMOS advances will continue to reduce power requirements for heterogeneous RF and logic devices, ultimately leading to the ability to combine all required handset circuitry into cost-effective, power-efficient, single-chip implementations.