


ElectronicsWeekly



First transistor to work at -90°C to be mass produced

■ Sub 1 Kelvin cryo-mosfets dissipate only 2pW at 500MHz so transistors could run from a 40mV supply at 1GHz at mK temperatures

STEVE BUSH

SemiQon of Finland has claimed “the world’s first transistor capable of functioning efficiently in cryogenic conditions”.

This is not just for 4K ‘liquid helium’ temperatures, this is for the <1K environments reached only using helium-3 cooling.

“The transistors can be mass produced using existing CMOS fabs, with no new infrastructure required,” said the company. In quantum computers the “transistor allows for control and readout electronics to be

placed directly inside a cryostat, alongside the processors, but without causing the disruption that dissipation of heat brings to these systems”.

They are FD-SoI (fully-depleted silicon-on-insulator) mosfets, designed particularly for sub-1K operation by SemiQon, and made on 150mm SOI wafers.

Charge carriers are extended from doped contact regions into a 20nm thick intrinsic channel, which can be tuned by a front and back gate, it said.

The company estimated that a qubit

controller made from 100nm wide, 30nm long transistors could operate at 500MHz with each mosfet dissipating 2pW. It also said that transistors could be made to run from a 40mV supply at 1GHz.

Given the amount of cooling available with 3He refrigeration, between 107 and 1,010 of the mosfets could run at mK temperatures, the company it said.

SemiQon expects to deliver its first cryo-CMOS transistors to customers in 2025. Use is foreseen in quantum computing, HPC (supercomputing) and in space vehicles.



Claim your free subscription

NEWS

- 7** Swedish all-electric hydrofoil ferry sets sail for Lake Tahoe

BLOGS

- 17** An Engineer in Wonderland ends the year on a high: creating a PCB using a CNC machine

INDUSTRY PROFILE

- 18** Margaret Schmitt shares her love for new technology, startups and mentoring

MEMORY DESIGN

- 26** Design guidance for using ReRAM to save energy in battery operated devices

Performance gets a reboot with partitioning using ReRAM

■ Embedded NVM such as ReRAM, coupled with smart memory partitioning, can improve the energy efficiency of battery-operated and energy-harvesting ICs, observes **Ilan Sever**

Low power consumption is a key design consideration, especially for battery-operated IoT devices with Bluetooth Low Energy or energy-harvesting technology, and for medical devices such as wearables and implantables. These devices must ensure that data gathered by tiny sensors is regularly and reliably delivered, from often inaccessible locations. For many applications reliability can have life or death consequences. Devices must be designed for long battery life – supporting applications that last up to 10 to 15 years on one battery – in an extremely space-constrained area.

Embedded non-volatile memory (NVM) for these devices requires low power operation, high endurance and reliability over the product's lifetime to support continuous monitoring, logging and communicating of small amounts of data.

System performance

Choosing low power embedded NVM such as resistive RAM (ReRAM) can help improve system energy efficiency and enable longer use times between recharges or battery replacements. It can also lead to better thermal



By reducing power consumption, the memory subsystem can allocate more power to other critical components

management and overall 'greener' technology. Importantly, by reducing power consumption, the memory subsystem can allocate more power to other critical components to improve overall system performance.

While the factors impacting NVM power consumption can vary based on application ReRAM is shown to consume significantly less read, write and standby power than embedded flash. This is due to the inherent low-power attributes of ReRAM as a technology, as well as smart power saving techniques we can implement. The low voltage levels used for memory transactions, coupled with ReRAM's fast memory access time, greatly reduce the overall power consumed. In addition, with programming, standby, sleep and very deep power-down

ReRAM modes, as well as rapid wake-up from deep power-down mode, designers can enable near-zero leakage power of internal and external NVM.

Designers can take this power advantage even further by implementing smart, power-aware system memory partitioning strategies. This includes dividing data intelligently across volatile and NVM resources to reduce the size of system SRAM.

Medical IoT devices

An example of a use case is a medical sensor application. This could be a wearable sensor designed to monitor a specific health parameter. The device has a low-power 32-bit microcontroller, 256Kbit NVM code storage, 256Kbits data logging SRAM storage consuming 1.8uW in retention mode, and always-on logic. The power source is a CR2016 battery with a capacity of 90mAh at 3V.

Such a medical device wakes up every second to measure sensor data, handling ~20,000 MCU execution commands and logging four words of intermediate data before going back to sleep. That process


takes about 10msec of wake-up time. The device also wakes up every 10 minutes to calibrate and run algorithms to process the logged data. Each time it does this it handles ~two million MCU execution commands, and logs 1,000 words (32-bit) of processed data, staying awake for around two seconds before it goes back to sleep.

The write cycles of an NVM are not unlimited, so it is important to consider the endurance needs of this application. ReRAM can be programmed at least 100K times, we are looking at ~50K hours or ~six years of logging. To log 1,000 words every 10 minutes, we are looking at ~130K hours or ~15 years of logging before wearing out.

Architecture for MCU power

In a typical system it is common to store code on external flash and then load code onto the local code SRAM from which the MCU fetches the code. Each time the system wakes up to log and process data there is MCU power consumption related to executing the write cycles as well as time and energy needed to load the code from the external flash into the code SRAM, and for the MCU to fetch the code. Power is required to maintain the code SRAM or to keep on the always-on logic for these operations.

Figure 1: Example medical device logging system using data and code ReRAM



One alternative methodology would be to use an eXecute in Place (XiP) architecture. In this way an on-chip NVM such as ReRAM can be used to store code instead of the code SRAM and the MCU can fetch the code directly from the ReRAM. This increases the speed and decreases the power since there is no need to access the external flash. It is also possible to turn off the code ReRAM, further reducing power.

Such an architecture can result in a 30% power savings over the traditional architecture.

A typical system uses external flash and an MCU with code SRAM. It logs the temporary data into the data SRAM and when the logged data is taken into NVM it is stored back to

the external flash. This is a fairly power-hungry process. (The lifetime of a device with this topology would be around one and a half years, which is not sufficient for an implantable medical device, for example.)

Replacing the off-chip memory with on-chip ReRAM in a XiP architecture means there is no need to transfer to the external flash (Figure 1).

In addition to XiP, log data can be stored in

on-chip ReRAM, eliminating the external flash (both for code and logging data). Replacing the on-chip code SRAM as well as part of the data SRAM with ReRAM (Figure 1) can achieve a 60% power reduction and a device that can last up to four years.

Finally, instead of logging the data into SRAM and storing processed data onto ReRAM logging processed data directly into ReRAM eliminates most of the on-chip SRAM (Figure 2). This can extend the lifetime for this application to more than five years, largely due to mitigation of the cell leakage problems associated with

SRAM technology during the inactive states of the system. With NVM such as ReRAM there is close to zero power consumption used for retaining the data during inactive states.


New technologies

In the case of a device that does not have a battery, such as an energy-harvesting device, a traditional architecture can be prohibitive. Using a combination of logging data in SRAM and uploading it to flash can consume more power than is available.

Small, low-power devices that require over-the-air (OTA) firmware updates can also benefit. Hearing aids, wireless earphones, pacemakers and other medical and wearable appliances use small rechargeable batteries and can contain firmware code in Mbyte capacities. Performing a chip erase and then programming the new code required for the OTA update requires more time and energy than is available using a standard off-the-shelf low-power flash device.

It is also worth noting that ReRAM technology is now available in 22nm processes where embedded flash is unable to scale. □

Figure 2: Example of a medical device logging system with on-chip ReRAM for code and data



About the author

Ilan Sever is vice-president of R&D for Weebit Nano

Web: [linkedin.com/in/ilansever](https://www.linkedin.com/in/ilansever)