

Simulation and Experimental Characterization of a Unified Memory Device with Two Floating-Gates

Neil Di Spigna, Daniel Schinke, Srikant Jayanti, Veena Misra, and Paul Franzon

North Carolina State University, Department of Electrical and Computer Engineering,
Raleigh, NC USA

{nhdispig,djschink,sjayant2,vmisra,paulf}@ncsu.edu

Abstract. The operation of a novel unified memory device using two floating-gates is described through experimental characterization of a fabricated proof-of-concept device and confirmed through simulation. The dynamic, nonvolatile, and concurrent modes of the device are described in detail. Simulations show that the device compares favorably to conventional memory devices. Applications enabled by this unified memory device are discussed, highlighting the dramatic impact this device could have on next generation memory architectures.

Keywords: memory, nonvolatile, dynamic, volatile, unified, floating-gate, FLASH, DRAM, high-k dielectric, simulation.

1 Introduction

This chapter is in part based off previously published work on the demonstration of a novel double floating-gate unified memory device [1]. In this paper, that work is extended through device simulations and additional details on the fabrication, operation, and design of circuits based on such a device. Such a unified memory device could store both volatile (dynamic) and nonvolatile states simultaneously. This could have a dramatic impact on traditional memory hierarchies [2-4]. For example, the data stored in the nonvolatile mode of the device when the computer is powered down could quickly be written to the dynamic state when the power is turned on, allowing for instant-on computing. This data transfer could also operate in reverse as dynamic data could be written to nonvolatile states to allow for full or partial hibernation of the memory fabric. Alternatively, writing dynamic data quickly to nonvolatile data could enable fast in-situ checkpointing. Finally, there are a number of novel logic applications for such a device that could impact numerous areas of computation [4].

Two floating-gates (FGs) have been used previously for enhancing memory operation [5-8]. However, these designs typically have been used in an effort to increase the memory window and data retention compared to single FG devices. For example, the size of the nanocrystals in the two FG layers can be engineered to exploit the Coulomb Blockade effect [8]. In this research, however, two FGs are used to enable a device which can store both dynamic and nonvolatile states concurrently. The two modes of operation

are distinguished by the charge condition of the two FGs. The device is in the dynamic mode when charge is simply redistributed between the two FGs. It is not until charge is drawn up from the substrate that the device enters its nonvolatile mode. Therefore, the operation of the device requires the existence of a window between when charge is merely redistributed between the FGs, and when charge is drawn up from the channel; allowing for the coexistence and selective control of both states. This requires engineering the vertical stack to create and fine tune such a window. The fabrication of the proof-of-concept device is discussed in the next section and the design tradeoffs based on those decisions are considered throughout.

2 Device Fabrication and Modeling

The double FG MOSCAPs shown in Fig. 1 were fabricated to experimentally demonstrate and confirm the device operation. The process flow is outlined in Fig. 1a. The wafers were cleaned before a SiO_2 gate oxide was grown through thermal oxidation. For the bottom FG, palladium was deposited through e-beam evaporation and patterned using liftoff. HfO_2 was used as the inter-FG dielectric and deposited through Atomic Layer Deposition (ALD). The top FG was then fabricated once again using palladium. ALD was used to deposit the control dielectric of HfAlO , which has been previously shown to have low leakage that is required of ultra-scaled FLASH memory technologies [9]. A palladium control gate was deposited and patterned, followed by a backside etch and aluminum deposition. A Transmission Electron Micrograph (TEM) of the cross section of the fabricated device is shown in Fig. 1b.

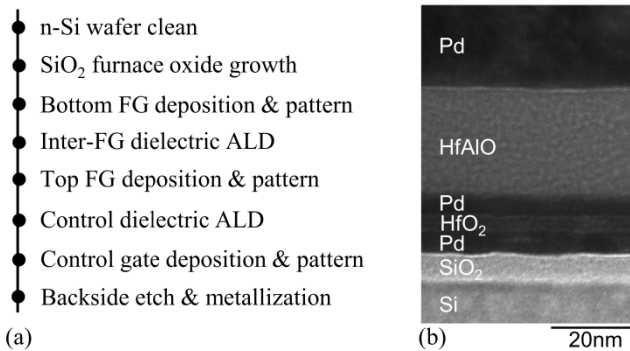


Fig. 1. Device Fabrication. (a) Recipe and (b) TEM cross section.

Being a proof-of-concept structure, the layer thicknesses of the vertical stack were not aggressively scaled, which as will be shown later in this paper have led to relatively high operating voltages. Further device scaling combined with engineering the materials has been shown through simulation to lower operating voltages, optimize device performance and achieve high storage density [4]. Towards these ends, a 65-nm gate length MOSFET, with the properties listed in Table 1, was modeled in Sentaurus TCAD.

Table 1. Device Model Properties

Layer	Material	Thickness
Control Gate	Molybdenum	10 nm
Control Dielectric	HfO ₂	18 nm
Top FG	Platinum	3 nm
Inter-FG Dielectric	HfSiO	3.2 nm
Bottom FG	Magnesium	3 nm
Gate Dielectric	SiO ₂	4 nm
Substrate	Bulk Si	-

This model is used to confirm the characterized device operation and is the link between the proof-of-concept structure and the circuit simulations discussed in Section 6. In addition to more aggressive thickness scaling, some important material distinctions can be made between the fabricated device and the simulated device. In contrast to the fabricated device, different materials were chosen for the two FGs, creating an asymmetry across the inter-FG dielectric, as can be seen in the energy band diagrams for the simulated and fabricated devices shown in Fig. 2. A high work function metal, Pt, was used for the top FG, whereas a low work function metal, Mg, was used for the bottom FG. This results in fast dynamic programming as electrons tunnel from the bottom FG to the top FG relatively easily. Once trapped, the deep potential well of the top FG sustains sufficiently long retention times but comes at the expense of dynamic erasing, as will be shown later in the circuit simulations. The characterization and operation of this device is discussed in the following sections, starting with the dynamic mode.

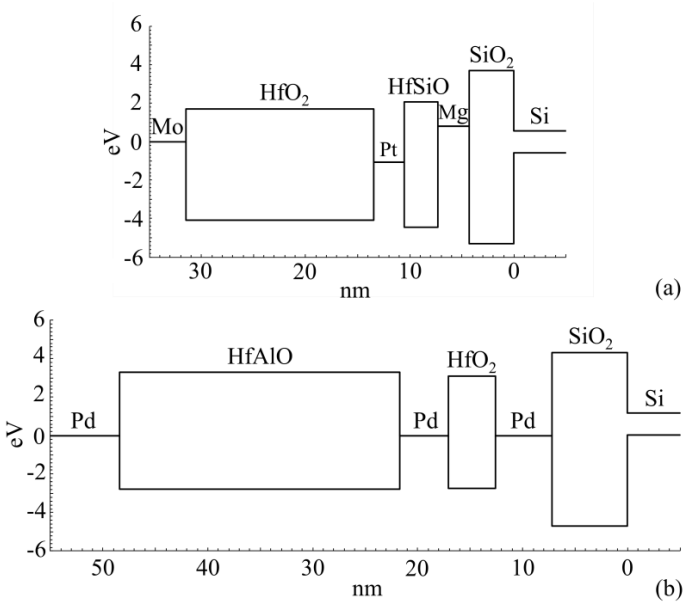


Fig. 2. The energy band diagrams of the (a) simulated and (b) fabricated devices

3 Dynamic Mode Operation

The mode of the device is determined by the applied voltage envelope. For a relatively small bias and short duration, the device's dynamic mode is programmed/erased. The dynamic operation of the device is illustrated in Fig. 3. The device is swept from a negative voltage, to a positive voltage, and then back to the negative voltage. Initially, as the device has a small negative voltage applied to the control gate, electrons will move to the bottom FG leaving behind a positive charge on the top FG, as pictured on the right-side of Fig. 3. The negative charge on the bottom FG, closer to the substrate, will cause a slight shift of the flat-band voltage to the right, as can be seen in the measured CV characteristic of the fabricated device. As the sweep continues, the voltage applied to the gate becomes positive, and the opposite charge condition results. Electrons now move to the top FG resulting in a positive charge closer to the substrate, as pictured on the left-side of Fig. 3. This positive charge closer to the substrate will cause a slight shift of the flat-band voltage to the left, once again demonstrated by the measured CV characteristic. Thus, for dynamic operation, the hysteresis is counter-clockwise, which would be the opposite direction anticipated for traditional single FG devices. Notice that for dynamic mode operation, the voltage applied to the control gate is insufficient to draw up charge from the substrate, but rather is only strong enough to simply redistribute charge on the FGs across the relatively thin inter-FG dielectric. Thus, there is no net increase in the charge on the FGs. This condition is what distinguishes the mode of operation of the device.




Fig. 3. Dynamic Mode Operation

The flat-band voltage shift of the fabricated device relative to the applied voltage envelope is shown in Fig. 4. As greater negative voltage envelopes are applied to the device, there is a greater positive shift in the flat-band voltage; whereas increasing positive voltage envelopes causes a greater negative shift in the flat-band voltage. The symmetry in the characteristics is indicative of the use of the same metal for the two FGs. As shown in Fig. 2b, this results in a symmetric energy barrier between the two FGs such that the program and erase characteristics are also symmetric.




Fig. 4. Dynamic Program/Erase Characteristics

The simulations of the dynamic mode operation of the device are shown in Fig. 5. The initial uncharged device characteristics are shown (1). A 5 V pulse is applied to the control gate for 50 ns causing the threshold voltage to shift about -330 mV (2). After about 300 ms, the threshold voltage decays about 220 mV back towards the initial state of the device (3). For these simulations, it is assumed that a 100 mV difference is needed to distinguish between the two distinct states, thus a refresh is required. A 5 V refresh pulse is applied to the control gate. Since the device has not fully decayed back to the initial state, this refresh pulse only needs to be applied for 40 ns, rather than the initial 50 ns applied to redistribute charge in the fresh device. As can be seen in Fig. 5a, this refresh returns the device back to the charged state threshold voltage (4). This volatile cycle continues, requiring a refresh period of about 300 ms to retain the charged state, and thus demonstrating the dynamic mode of the device.




Fig. 5. Dynamic Mode Simulations. The (a) drain current and (b) capacitance vs. control gate voltage.

The volatile nature illustrated by the device simulations was confirmed in the fabricated device, as shown in the dynamic retention characteristics of Fig. 6. A +10 V pulse was applied to the fabricated device which caused a flat-band voltage shift to the left, as shown in Fig. 6a. This is directly analogous to the simulations shown in Fig. 5b. Once the bias to the control gate was removed, the CV characteristics would decay back to the original curve. This is illustrated in Fig. 6b in which the capacitance at 0.5 V is measured over time. As the charge difference between the two FGs decays, the flat-band voltage shifts to the right and the capacitance at 0.5 V decays from ~8.5 pF to ~6.4 pF. At this point, after ~22 s, the +10 V is reapplied to the control gate, once again refreshing the charge difference between the two FGs and causing the flat-band voltage to shift back to the left. This is shown for five cycles in Fig. 6b, successfully demonstrating and confirming the predicted volatile nature of the dynamic mode operation.

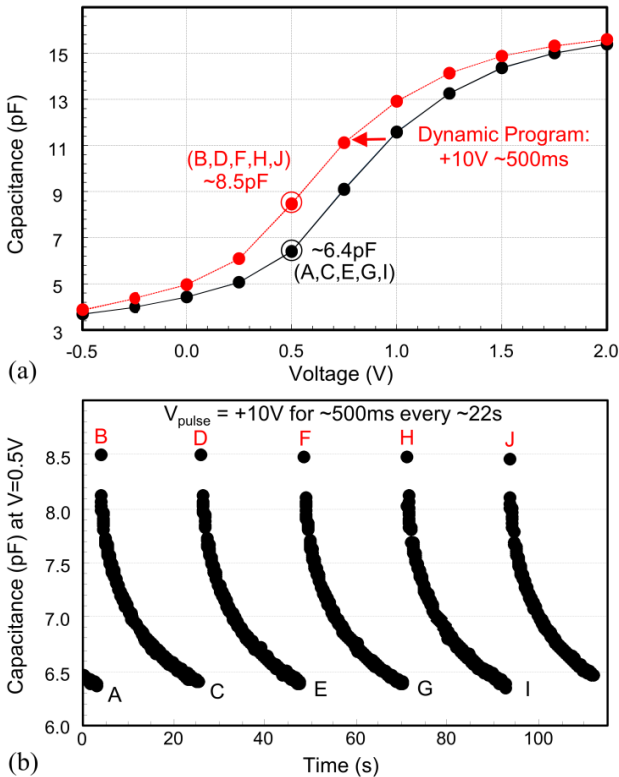


Fig. 6. Dynamic Mode Retention Characteristics. (a) A +10 V; ~500 ms pulse is applied to the control gate resulting in a negative flat-band voltage shift. (b) The capacitance was measured at 0.5 V over time for 5 cycles showing the charge difference between the two FGs decay back to the original CV curve.




Fig. 7. Dynamic Mode Endurance Characteristics. The device characteristics are shown for 200,000 cycles with a ± 7.5 V for ~200 ms pulse.

Finally, the dynamic mode endurance of the fabricated device is shown in Fig. 7. A consistent 300 mV window between the programmed and erased states is maintained as the device was cycled over 10^5 times, though a cycling drift is present. The inter-FG dielectric is critical in ensuring the stable operation over the extensive number of cycles required of DRAM. As this dielectric is further scaled, this will permit the reduction of voltages and fields, and thus the use of lower energy tunneling mechanisms that will reduce the stress on this dielectric. Choosing an appropriate dielectric for this inter-FG is actively being investigated.

4 Nonvolatile Mode Operation

The nonvolatile mode of the device is entered when a voltage pulse is applied to the control gate that is sufficient enough to draw up a net charge from the substrate to the FGs, as illustrated in Fig. 8. Once again the device is swept from a negative voltage, to a positive voltage, and then back to the negative voltage. However, unlike in the dynamic mode operation, the bias is large enough to draw up charge from the substrate. As the voltage applied to the gate starts out negative, electrons are repelled towards the substrate leaving behind a positive charge on the FGs. A net positive charge on the FGs causes a negative shift in the flat-band voltage, as can be seen in the measured CV characteristic. As the sweep continues, the voltage applied to the gate becomes positive, and electrons are now drawn up from the substrate resulting in a negative charge on the FGs. The voltage is then swept in reverse, and the negative charge on the FGs results in a positive flat-band voltage shift, once again witnessed in the measured CV characteristic. This results in a clockwise hysteresis, which is expected for traditional single FG nonvolatile devices.




Fig. 8. Nonvolatile Mode Operation

The dynamic mode operation depicted in Fig. 3 and the nonvolatile mode operation depicted in Fig. 8 are combined to demonstrate the program/erase characteristics shown in Fig. 9. This clearly illustrates how the mode of the device is determined by the applied voltage envelope. The negative flat-band voltage window (CCW




Fig. 9. Program/Erase Characteristics. The dynamic mode of the curve is enlarged in the inset.

hysteresis) is evidenced for low voltages; while the more traditional positive flat-band voltage window (CW hysteresis) occurs at higher voltages. Control devices fabricated without two FGs but with the same dielectrics did not show the dynamic mode.

The nonvolatile program/erase characteristics of the fabricated device are shown in Fig. 10. The asymmetry between the program and erase voltages is expected since, as shown in the band diagram of Fig. 2b for the fabricated device, the electrons can more easily tunnel onto the FGs than they can tunnel back to the substrate. In Fig. 10, both the initial and after-60 second flat-band voltage shifts are plotted, and in every case, the after-60 second shift is more pronounced than the initial shift. This is due to the fact that as the bias is being applied to the control gate, more of the charge is being drawn up to the top FG, relative to the bottom FG. When the external bias is removed, the charge redistributes between the two FGs resulting in an increase in the charge on the bottom FG. Since the bottom FG is closer to the substrate, this charge redistribution leads to a greater flat-band voltage shift over time as the charge settles. This is confirmed by the simulations of the nonvolatile mode of the modeled device shown in Fig. 11.




Fig. 10. Nonvolatile Program/Erase Characteristics

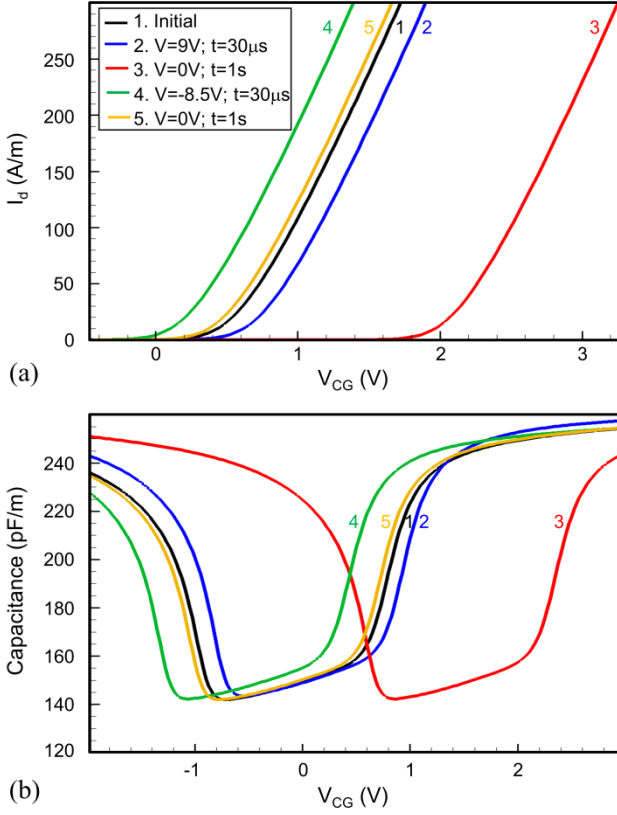


Fig. 11. Nonvolatile Mode Simulations. The (a) drain current and (b) capacitance vs. control gate voltage.

The initial uncharged device characteristics are shown (1). A 9 V pulse is applied to the control gate for 30 μs (2). This pulse is large enough to pull up charge from the channel, resulting in a net increase of charge on the FGs. Initially, most of the charge is drawn up to the top FG, limiting the impact on the channel. Thus, only a relatively minor positive threshold voltage shift occurs immediately after the pulse, as shown in Fig. 11a (2). However, after the voltage is removed from the control gate, the charge on the FGs redistributes resulting in a much more pronounced positive threshold voltage shift of about 1.52 V after about 1 s (3). This is the same phenomenon that occurred in the fabricated devices, though not to the same extent. The device does not reach its stable state until after some time passes. The relationship between the initial applied pulse and the charge redistribution settling time is currently being investigated and will have to be accounted for at the circuit level. Finally, as shown in the simulation, a -8.5 V pulse applied for 30 μs (4) followed by a charge settling period of 1 s (5) returns the device approximately back to its uncharged state.

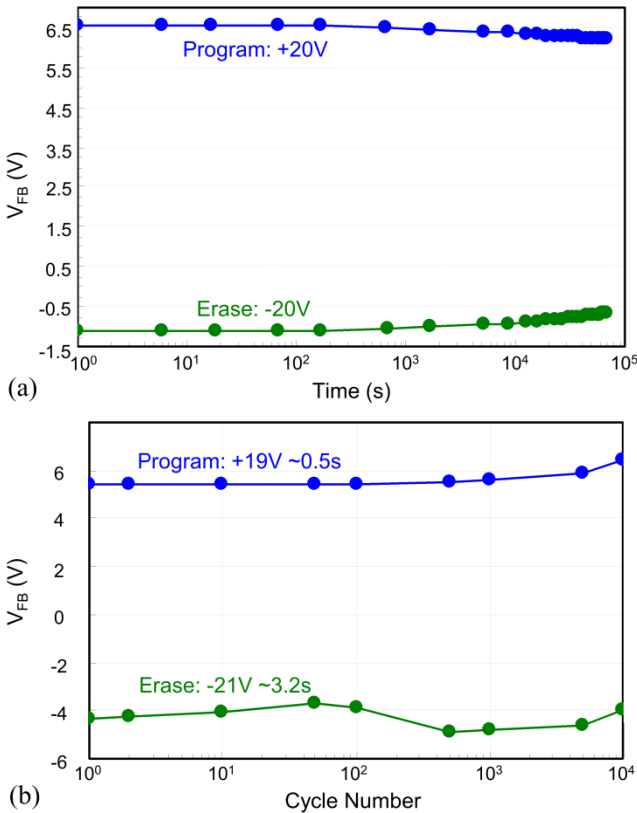


Fig. 12. The nonvolatile mode (a) retention and (b) endurance characteristics

To verify the nonvolatile nature of the fabricated device, the retention of the nonvolatile mode is plotted in Fig. 12a. A window of at least 4.5 V is maintained by extrapolating the data out to 10 years. Finally, the nonvolatile endurance of the device was demonstrated through over 10,000 cycles, as shown in Fig. 12b.

5 Concurrent Mode Operation

The device is not limited to operation in either the dynamic mode or the nonvolatile mode, but rather it can operate in both the dynamic and nonvolatile modes at the same time. The experimental verification of concurrent mode operation is shown in Fig. 13. The device is first programmed into the nonvolatile state using a +17 V sweep, as shown in Fig. 13a, which results in a positive flat-band voltage shift as negative charge is drawn up from the substrate into the FGs. To demonstrate concurrent mode operation, a dynamic state is then embedded on top of the programmed nonvolatile state by the application of a dynamic pulse of +10 V. As a response to this pulse, some of the negative charge on the bottom FG is drawn up to the top FG, leaving

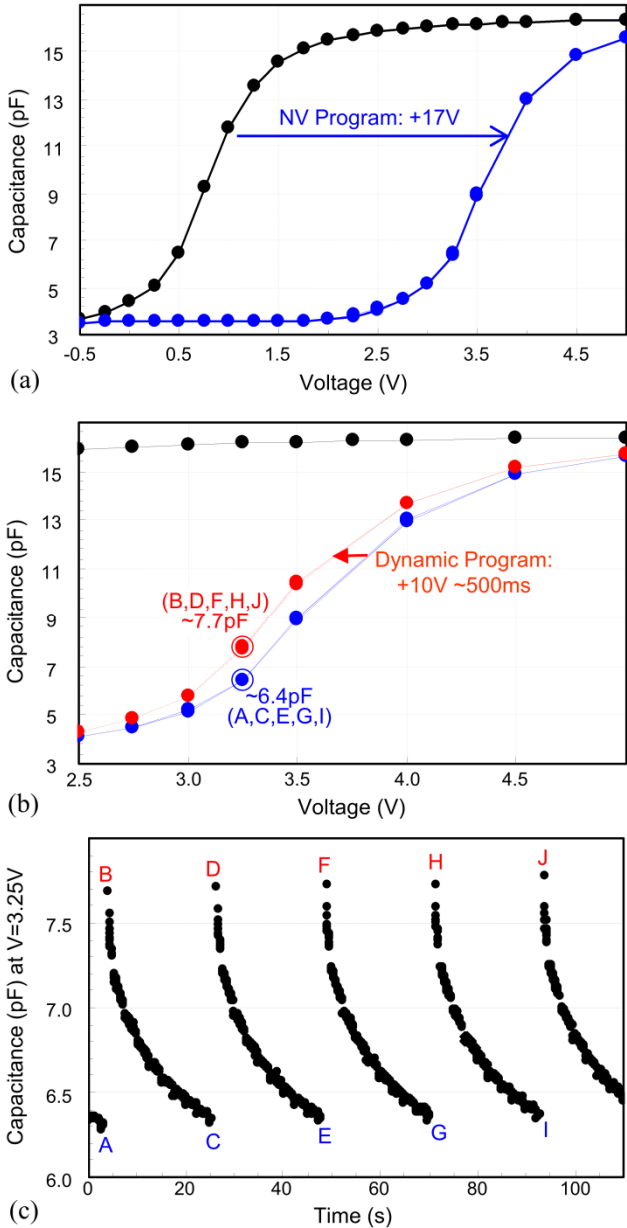


Fig. 13. Concurrent Mode Operation. (a) The device is charged into the programmed nonvolatile mode with a +17 V sweep. (b) A +10 V pulse embeds a dynamic state on top of the programmed nonvolatile state. (c) The capacitance is measured at 3.25 V over 5 cycles demonstrating the retention of the embedded dynamic state. As the dynamic state decays, the device returns to the programmed nonvolatile state, requiring a refresh.

behind a less negatively charged bottom FG, resulting in a slightly negative flat-band voltage shift relative to the original charged nonvolatile state, as shown in Fig. 13b. Once this dynamic bias is removed, the charge difference between the two FGs decays, and the flat-band voltage shifts back to the original nonvolatile programmed state. This cycling of the dynamic state embedded on top of the programmed nonvolatile state is repeated five times, as shown in Fig. 13c. Combining this data with that shown in Fig. 6, which represents the dynamic state embedded on top of the nonvolatile erased state, successfully demonstrates that the dynamic state can be embedded on both the programmed and erased nonvolatile states. Thus, concurrent mode operation of the fabricated device is experimentally verified.

The concurrent mode simulations are shown in Fig. 14. The device is first programmed into its charged nonvolatile state (1). A dynamic pulse of 5 V for 50 ns results in about a -330 mV threshold voltage shift (2). Upon cessation of the bias, the charge redistributes and the threshold voltage starts decaying back to the charged nonvolatile state (3). The CV curve of Fig. 14b is directly analogous to the experimental characterization shown in Fig. 13b. A relatively small dynamic pulse shifts the flat-band voltage in the negative direction, at which point it begins to decay back to the charged nonvolatile state flat-band voltage, as shown in Fig. 13c. Upon a refresh, the flat-band voltage once again shifts to the left (4). Thus, it is shown through simulation that a dynamic mode can be embedded on the charged nonvolatile state.




Fig. 14. Concurrent Mode Simulations. The (a) drain current and (b) capacitance vs. control gate voltage.

The dynamic, nonvolatile, and concurrent mode operation have been experimentally demonstrated. The characterization of the fabricated devices has been confirmed through device simulation, successfully verifying the operation of this novel unified memory device. A memory array using this device is discussed in the next section.

6 Circuit Simulations

The memory array shown in Fig. 15 was designed in Cadence Virtuoso 2010 using a BSIM4.0 MOSFET model with a 45-nm gate length and device parameters similar to the simulated device described in Table 1. However, instead of bulk silicon, the substrate was SOI with a thickness of 13 nm; an SiO₂ back gate dielectric of 1.2 nm was used, and the control and back gates were composed of aluminum. The operation of this memory array is described in Table 2.




Fig. 15. Simulated memory array architecture

Table 2. Memory Array Operation

Operation	Bias Across Gate Stack	Duration
Dynamic Program	5 V	50 ns
Dynamic Erase	-5 V	10 μ s
Dynamic Refresh	5 V	40 ns
Nonvolatile Program	9 V	30 μ s
Nonvolatile Erase	-9 V	14 μ s
Low V_t Read	1.2 V	2.2 ns
High V_t Read	2.7 V	2.2 ns

To dynamically program a target device, 3 V is placed on the appropriate WL and -2 V is placed on the appropriate SL. This results in a 5 V bias across its gate stack, which when applied for 50 ns results in the target device being dynamically programmed, as previously described in the device simulation. However, to prevent inadvertent programming of non-target devices on that WL, the non-target SLs need to be biased to 2 V such that there is only a 1 V bias across their gate stack. This represents the dynamic retain condition.

As previously discussed, the device was engineered to have a low work function metal for the bottom FG and a high work function metal for the top FG. This resulted in an asymmetric barrier that allowed for fast dynamic programming and increased dynamic retention as charge tunneled easily from the bottom FG into the deeper energy well of the top FG, as shown in Fig. 2a. This resulted in a dynamic retention of 300 ms. However, this came at the expense of the dynamic erase; which as shown in Table 2 takes 10 μ s. This is much longer than conventional DRAM. If, on the other hand, the materials are chosen to be symmetric, as was the case for the fabricated device in which palladium was used for both the top and bottom FGs, the dynamic erase time would reduce to 200 ns. Of course with a symmetric barrier, there is no longer the deeper potential well for the charge in the dynamically programmed state and so the retention time would also be reduced. However, for this device the retention time would only reduce from 300 ms to 100 ms, which could prove a wise tradeoff for reducing the dynamic erase time from 10 μ s to 200 ns. Certainly further work function engineering can be performed to tailor the device performance towards target applications.

Another advantage of the device is that it operates more like an SRAM than a DRAM, and thus the read operation takes only 2.2 ns, which is much faster than DRAM. The read is also nondestructive, unlike DRAM. The memory array should also have a higher density than DRAM due to the difficulty of scaling the DRAM capacitor and maintaining sufficient charge sharing with the bitline. This device is scalable, in bulk form, to at least the 16-nm node. Through stacking, it has the potential to reach densities equivalent to the 8-nm node.

Overall, the device offers several advantages compared to conventional DRAM [3]. However, such a comparison is ill-conceived. The device may not be wholly superior to DRAM, nor to a similarly scaled single FG nonvolatile device, since it requires an extra FG and the addition of an ultra-thin inter-FG dielectric layer; but the device offers a tremendous advantage that neither of the other devices do singularly; it can store both dynamic (DRAM) and nonvolatile (FLASH) states concurrently. Such a unified memory device has enormous potential to impact next generation memory architectures.

7 Applications

There are a number of applications for such a unified memory device. For example, the device could be used to enable instant-on computing. The computer could quickly be powered down by simply moving all of the dynamic states into their nonvolatile

states. If the entire memory array is written to its nonvolatile state in parallel, this would take only about 30 ms. When the user wants to power the computer back on, the memory controller simply needs to write back all of the nonvolatile data into the dynamic state. Once again, when performed in parallel, this would take only about 14 ms. In theory, the user could power up and power down the computer in only a fraction of a second. Beyond user convenience, this could allow for the operating system to power down during moments of inactivity. For example, if the user walked away from their computer to get a drink or take a phone call, the operating system could power down and conserve battery life. When the user returned, the power up penalty would only be a fraction of a second.

This device could also enable partial hibernation. For parts of the memory that are not currently being used, those arrays could be written to the nonvolatile state in the background as the user continues to operate their computer. This could enable a flexible memory fabric that could be selectively powered down which could have a significant impact on energy-proportional computing. An example application for this would be Google servers. Recently, a study on their server power usage showed that at utilization workloads that were common (20-30%), the servers operated at less than half their peak energy efficiency performance [10]. Given the nature of their utilization, current solutions to transfer to inactive modes are impractical because of both a time latency and energy penalty. The device described in this chapter could make such transitions practical by significantly reducing the wake-up penalties. Alternatively, partial hibernation enabled by this device could be used to further enhance active energy-saving schemes.

Another example application in which these devices could be beneficial is in-situ checkpointing. The device could be running continuously in dynamic mode, and then upon desire for a check-point, the entire memory array could be quickly written to the nonvolatile state in only about 30 ms. This would be much more efficient than writing through narrow channels to disk. Thus, more check-points could be efficiently taken, improving the resiliency of the computer. Upon detection of an error, the correct state could be recovered much faster than traditional memory hierarchies would permit. Instant-on computers, energy-proportional computing, and in-situ checkpointing are just a few examples of the potential that could be realized with a memory array composed of this new unified memory device.

8 Conclusion

New unified memory devices using two FGs were modeled, simulated, fabricated and characterized. The operation of the devices in dynamic, nonvolatile, and concurrent modes were demonstrated in proof-of-concept MOSCAPs and confirmed through device simulations. The programming, retention, and endurance characteristics were demonstrated for the different modes. A memory array based on these devices was designed and simulated. It was shown that these devices compare favorably to both conventional DRAM and FLASH devices. However, the true potential of these devices is not in their use as either a DRAM or FLASH replacement, but rather as a

new unified memory device that can store both dynamic and nonvolatile states concurrently. Applications for such a device were discussed that highlight the significant impact this device could have on next generation memory architectures.

Acknowledgements. This article is in part based on works supported by the National Science Foundation under award nos. 0811582 and 1065458. We thank Dr. Eric Rotenberg, Steve Lipa, W. Shepherd Pitts, Shivam Priyadarshi, Vinodh Kotipalli and Narayanan Ramanan for their valuable contributions to aspects of this effort. Thanks to Dr. Dale Batchelor of AIF at NCSU for TEM analysis and to Jonathan Pierce for FIB preparation of the TEM cross sections.

References

1. Di Spigna, N., Schinke, D., Jayanti, S., Misra, V., Franzon, P.: A Novel Double Floating-Gate Device. In: IEEE/IFIP 20th International Conference on VLSI SoC, pp. 53–58 (2012)
2. Park, K.-H., Park, C.M., Kong, S.H., Lee, J.-H.: Novel Double-Gate 1T-DRAM Cell Using Nonvolatile Memory Functionality for High-Performance and Highly Scalable Embedded DRAMs. *IEEE Transactions on Electron Devices* 57(3), 614–619 (2010)
3. Han, J.-W., Ryu, S.-W., Kim, D.-H., Choi, Y.-K.: Polysilicon Channel TFT With Separated Double-Gate for Unified RAM (URAM)-Unified Function for Nonvolatile SONOS Flash and High-Speed Capacitorless 1T-DRAM. *IEEE Transactions on Electron Devices* 57(3), 601–607 (2010)
4. Schinke, D., Di Spigna, N., Shiveshwarkar, M., Franzon, P.: Computing with Novel Floating-Gate Devices. *Computer* 44(2), 29–36 (2011)
5. Afshari, K.: Nonvolatile Memory with Multi-Stack Nanocrystals as Floating Gates. In: 2007 NNIN REU Research Accomplishments, pp. 38–39 (2007)
6. Lee, C., Gorur-Seetharam, A., Kan, E.C.: Operational and Reliability Comparison of Discrete-Storage Nonvolatile Memories: Advantages of Single- and Double-Layer Metal Nanocrystals. In: *IEDM 2003 Technical Digest*, pp. 557–560 (2003)
7. Singh, P.K., Bisht, G., Hofmann, R., Singh, K., Krishna, N., Mahapatra, S.: Metal Nanocrystal Memory with Pt Single- and Dual-Layer NC With Low-Leakage Al₂O₃ Blocking Dielectric. *IEEE Electron Device Letters* 29(12), 1389–1391 (2008)
8. Ohba, R., Sugiyama, N., Uchida, K., Koga, J., Toriumi, A.: Nonvolatile Si Quantum Memory With Self-Aligned Doubly-Stacked Dots. *IEEE Transactions on Electron Devices* 49(8), 1392–1398 (2002)
9. Jayanti, S., Yang, X., Suri, R., Misra, V.: Ultimate Scalability of TaN Metal Floating Gate with Incorporation of High-K Blocking Dielectrics for Flash Memory Applications. In: *IEDM 2010 Technical Digest*, pp. 106–109 (2010)
10. Barroso, L.A., Holzle, U.: The Case for Energy-Proportional Computing. *Computer* 40(12), 33–37 (2007)