

18.1 Implementation of the CELL Broadband Engine™ in a 65nm SOI Technology Featuring Dual-Supply SRAM Arrays Supporting 6GHz at 1.3V

J. Pille¹, C. Adams², T. Christensen², S. Cottier³, S. Ehrenreich¹, F. Kono⁴, D. Nelson², O. Takahashi³, S. Tokito⁵, O. Torreiter¹, O. Wagner¹, D. Wendel¹

¹IBM, Boeblingen, Germany, ²IBM, Rochester, MN, ³IBM, Austin, TX
⁴Toshiba, Austin, TX, ⁵Sony, Austin, TX

The two processing elements of the CELL Broadband Engine™ [1,2] drive different memory requirements: the cache system of the Power Processor Element (PPE) includes separate 32kB L1 data and instruction caches, and a 512kB unified L2. The 256kB local store (LS) is the main memory element of the Synergistic Processor Element (SPE). Several supporting arrays are required. While PPE and SPE arrays work in the core clock domain, the L2 and support functions run at half the core frequency. The high frequency of operation and adherence to strict cycle boundaries at the macro level result in pipelined arrays. The macro cycle boundaries are at the address and data inputs, the wordline driver and the array data out. This scheme results in a full clock cycle between wordline select and data out [2].

To improve manufacturability in 65nm technology, all designs (core and nest) use the same 0.700μm² SRAM cell. A ripple-domino sense scheme with a local bitline of 16 cells reduces the impact of device variations in the cell (V_t scatter), improves stability and allows tuning for speed to keep up with logic performance. Array core supporting circuits (wordline/precharge driver with integrated level shift, cross-coupled NAND, latches) underlie strict design rules. To balance between evaluation and restore phases and to maximize the functional window, detailed duty-cycle analysis is performed on all array macros. Clocked signals are directly derived from the well-tuned global clock through a local clock buffer (LCB) to minimize variation and maximize tracking. The chip features a duty-cycle corrector (DCC) to optimally center the global clock mid-cycle edge for maximum performance and yield. These rigorous enforcements allow optimal technology tuning since all arrays behave the same.

Figure 18.1.1 shows the read/write circuit, also called local evaluation circuit (loceval), with a minimum device solution. It connects to two groups of 16 cells and can be independently controlled through the two precharge signals (prch0_b, prch1_b), only one of which can be selected at a time. The precharge signals restore the local bitlines (blt0, blc0, blt1, blc1) to V_{DD} and trigger the read/write access, timed synchronous to the wordline. The two global write bitlines (wt_b, wc) are used to switch between read and write.

When writing a '0' to the bottom group of cells, the true bitline (blt0) is pulled down by the write devices N01, N02 (wc = 1, prch0_b = 1) while blc0 is held high (N00, wt_b = 1). The two PMOS devices (P02, P12) are turned off. When writing a '1', blc0 is driven low by wt_b. The blt0 bitline is held high through the PMOS hold device P02 to mask any false switching on the local bitline blt0 to the global read bitline (rbt).

During read, the write devices are inactivated by setting wc = 0 and wt_b = 1. When reading a '0' from the bottom group of cells, prch0_b goes high, the selected cell pulls down the blt0 line activating the global read device N2 via a two-input NAND stage (P03, P13, N03, N13). When reading a '1', blt0 is held high by the selected cell fighting the leakage of the other 15 cells connected to the bitline. The bitline settles a V_t below V_{CS} , so the NAND switching point needs to be below that level.

Traditional SRAM designs use the core voltage (V_{DD}) for SRAM arrays. Due to loss in stability in recent technologies (low voltages, V_t scatter) an additional array-cell specific voltage (V_{CS}) is introduced to increase the cell stability and performance. Different dual-supply schemes have been proposed. Connecting the whole array to the elevated power supply improves stability

and performance but also increases DC and AC power consumption. Connecting the cell only to V_{CS} leaves the bitlines at the lower V_{DD} , which further reduces the stress to the cell and reduces the power consumption on the bitlines. Stability improves while the difference between the two voltages increases. The drawback of this scheme is that the cell at the higher voltage needs to be overwritten with the lower voltage, making it difficult to write at large offsets between V_{DD} and V_{CS} . This design uses a scheme where V_{CS} also controls the drive of the write devices, which allows the write-ability to track with the stability improvement of the cell at higher voltages. Figure 18.1.2 shows the two voltage domains: the cell, the wordline and precharge signals are connected to V_{CS} , while the loceval stage itself and the global write lines (wt_b, wc) are connected to V_{DD} . The higher voltage in the cell increases stability, cell read performance is improved by overdriving the wordline. Keeping the local bitlines at V_{DD} further reduces stress to the cell and lowers power. Figure 18.1.3 shows the resulting shmoo plot.

A typical read path is shown in Fig. 18.1.4. In the read '0' case, the 6T SRAM cell discharges the local bitline and the signal propagates through the loceval onto the global bitline. Depending on the specific array, a column select and/or redundancy stage follows. The cross-coupled NAND stage converts the dynamic input into a static output, followed by another muxing and/or redundancy stage (array specific). The data is captured in the output latch. In this scheme the cell performance directly impacts the access time.

Figure 18.1.5 shows a hardware shmoo plot with V_{DD} on the x-axis and V_{CS} on the y-axis at a fixed frequency. The dotted line shows the case $V_{DD} = V_{CS}$. According to this plot, for $V_{DD} = V_{CS}$, the minimum V_{DD} is 0.875V. If V_{CS} is set to 1.0V, V_{DD} can be reduced to 0.8V. The advantage is 75mV for V_{minf} (minimum voltage to pass a given frequency), reducing the overall chip power by about 19%. Since V_{CS} is only connected to the cell and wordline/precharge driver, the load is predominantly DC leakage power, thus the V_{CS} power distribution can be less dense. The total V_{CS} power is only 9% of the total chip power. There are three distinct regions on the V_{DD}/V_{CS} plane. For $V_{DD} > 0.775V$, performance is V_{DD} -limited and no V_{CS} setting can speed it up. For $0.8V < V_{DD} < 0.9V$ and $0.9V < V_{CS} < 1.0V$ a mixed V_{DD}/V_{CS} path limits speed and any higher V_{DD} or V_{CS} makes it pass. For $V_{CS} \geq 1.0V$, the cell speeds up so much that it no longer limits the overall performance and the array is limited by V_{DD} only. For $V_{DD} \gg V_{CS}$, the stress to the cell is too high (bitline voltage higher than cell voltage) and DC stability fails are observed. At $V_{CS} = 0.7V$, the cell is below the stability limit. Due to the enhanced write driver, no writeability limit is observed at the measured offsets between the two voltages. The cell can be written in time even at higher voltage differences.

Figure 18.1.6 shows the V_{minf} versus on-chip ring oscillator speed for a population of chips. At $V_{DD} = V_{CS}$, the cell is in the critical path thus its variation causes V_{minf} scatter. At $V_{CS} = V_{DD} + 200mV$, the cell is no longer limiting. V_{minf} is driven by V_{DD} -only circuits having larger devices with less variation, resulting in reduced scatter. Furthermore, the capability to reduce V_{CS} during test helps find slow cells, which can be fixed with redundancy, further improving performance and reducing power consumption. The maximum measured lab frequency with chip workload is 6GHz at $V_{DD} = V_{CS} = 1.3V$. The die micrograph and magnifications of the major arrays are shown in Fig. 18.1.7.

References:

- [1] D. Pham, S. Asano, M. Bolliger, et al., "The Design and Implementation of A First-Generation CELL Processor," *ISSCC Dig. Tech. Papers*, pp. 184-185, Feb., 2005.
- [2] D. Pham, T. Aipperspach, D. Boerstler, et al., "Overview of the Architecture, Circuit Design and Physical Implementation of a First-Generation CELL Processor," *IEEE J. Solid-State Circuits*, vol. 41, pp. 1692-1706, Aug., 2006
- [3] J. Davis, J. Plass, P. Bunce, et al., "A 5.6GHz 64kB Dual-Read Data Cache for the POWER6™ Processor," *ISSCC Dig. Tech. Papers*, pp. 622-623, Feb., 2006.
- [4] K. Zhang, K. Hose, V. De, et al., "The Scaling of Data Sensing for High Speed Cache Designs in Sub-0.18mm Technologies," *Symp. on VLSI*, pp. 226-227, Jun., 2000.

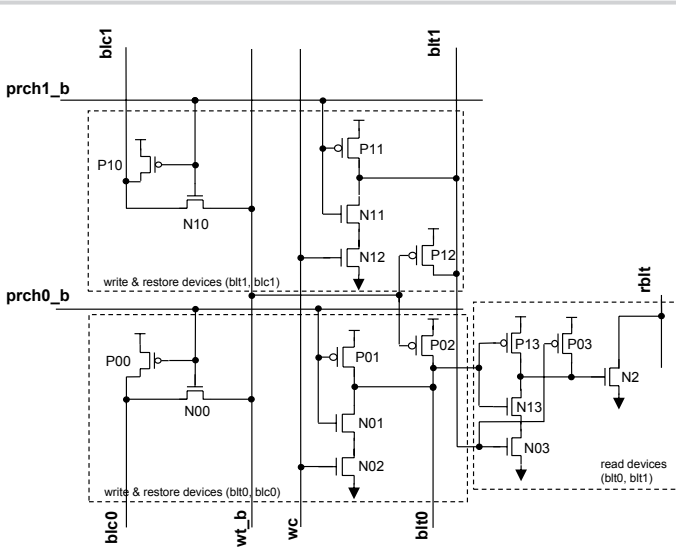


Figure 18.1.1: Local read and write circuit (loceval).

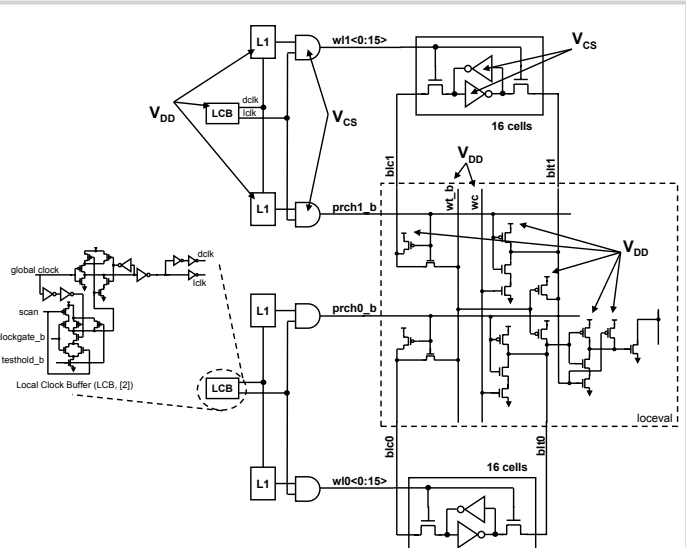


Figure 18.1.2: Voltage domains in 32-cell block including controls.

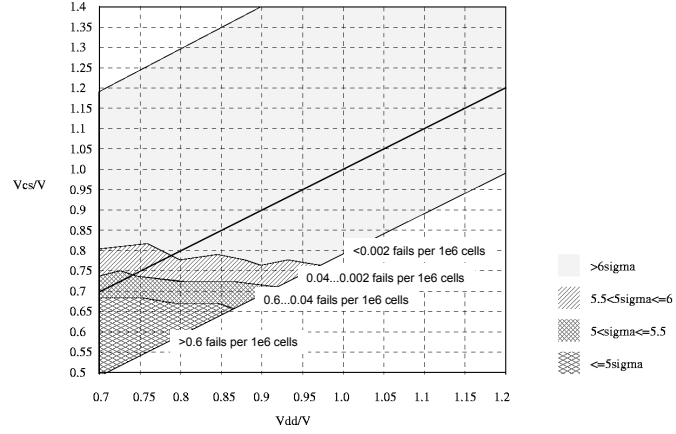


Figure 18.1.3: Cell write-ability analysis, statistical simulation assuming $\pm 6\sigma$ device variation (V_t , length, width).

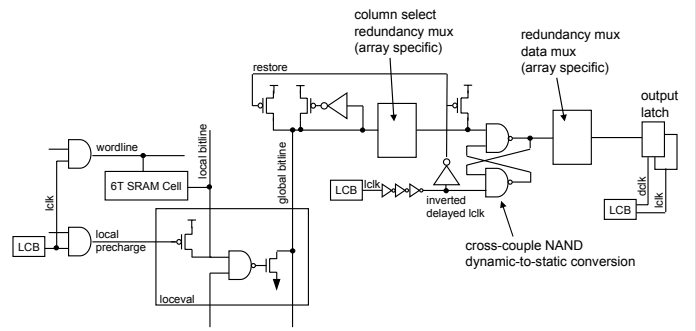


Figure 18.1.4: Typical ripple-domino read path (wordline to data out latch).

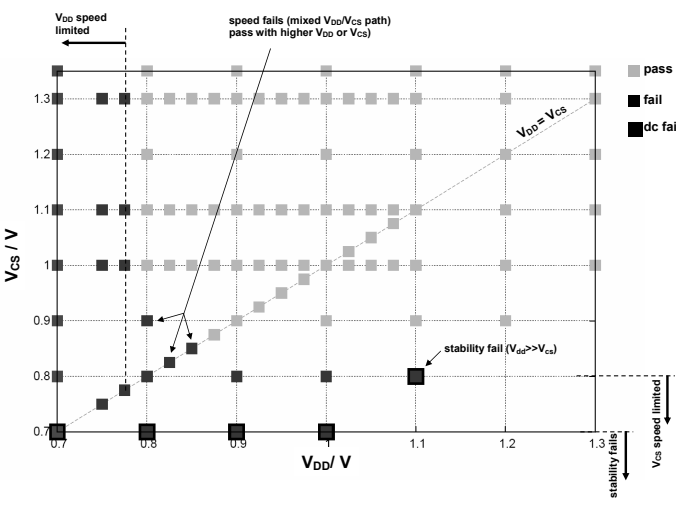


Figure 18.1.5: At-speed V_{DD}/V_{CS} shmoo plot (100mV steps).




Figure 18.1.6: V_{minf} scatter improves with $V_{CS} > V_{DD}$.

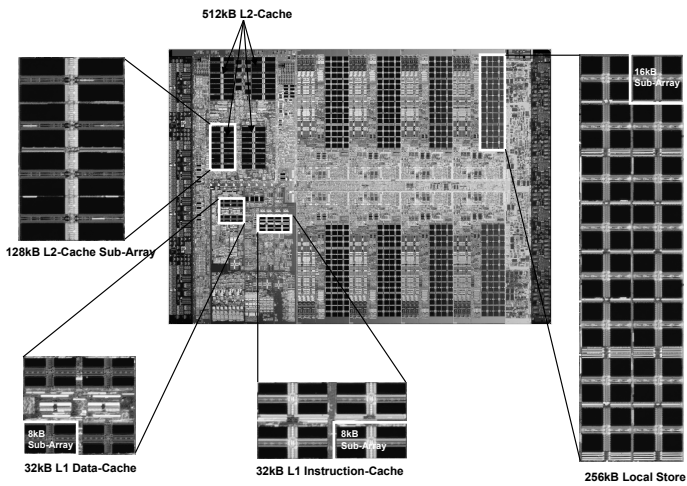


Figure 18.1.7: 65nm CELL Broadband Engine™ die micrograph and magnifications of the major arrays.

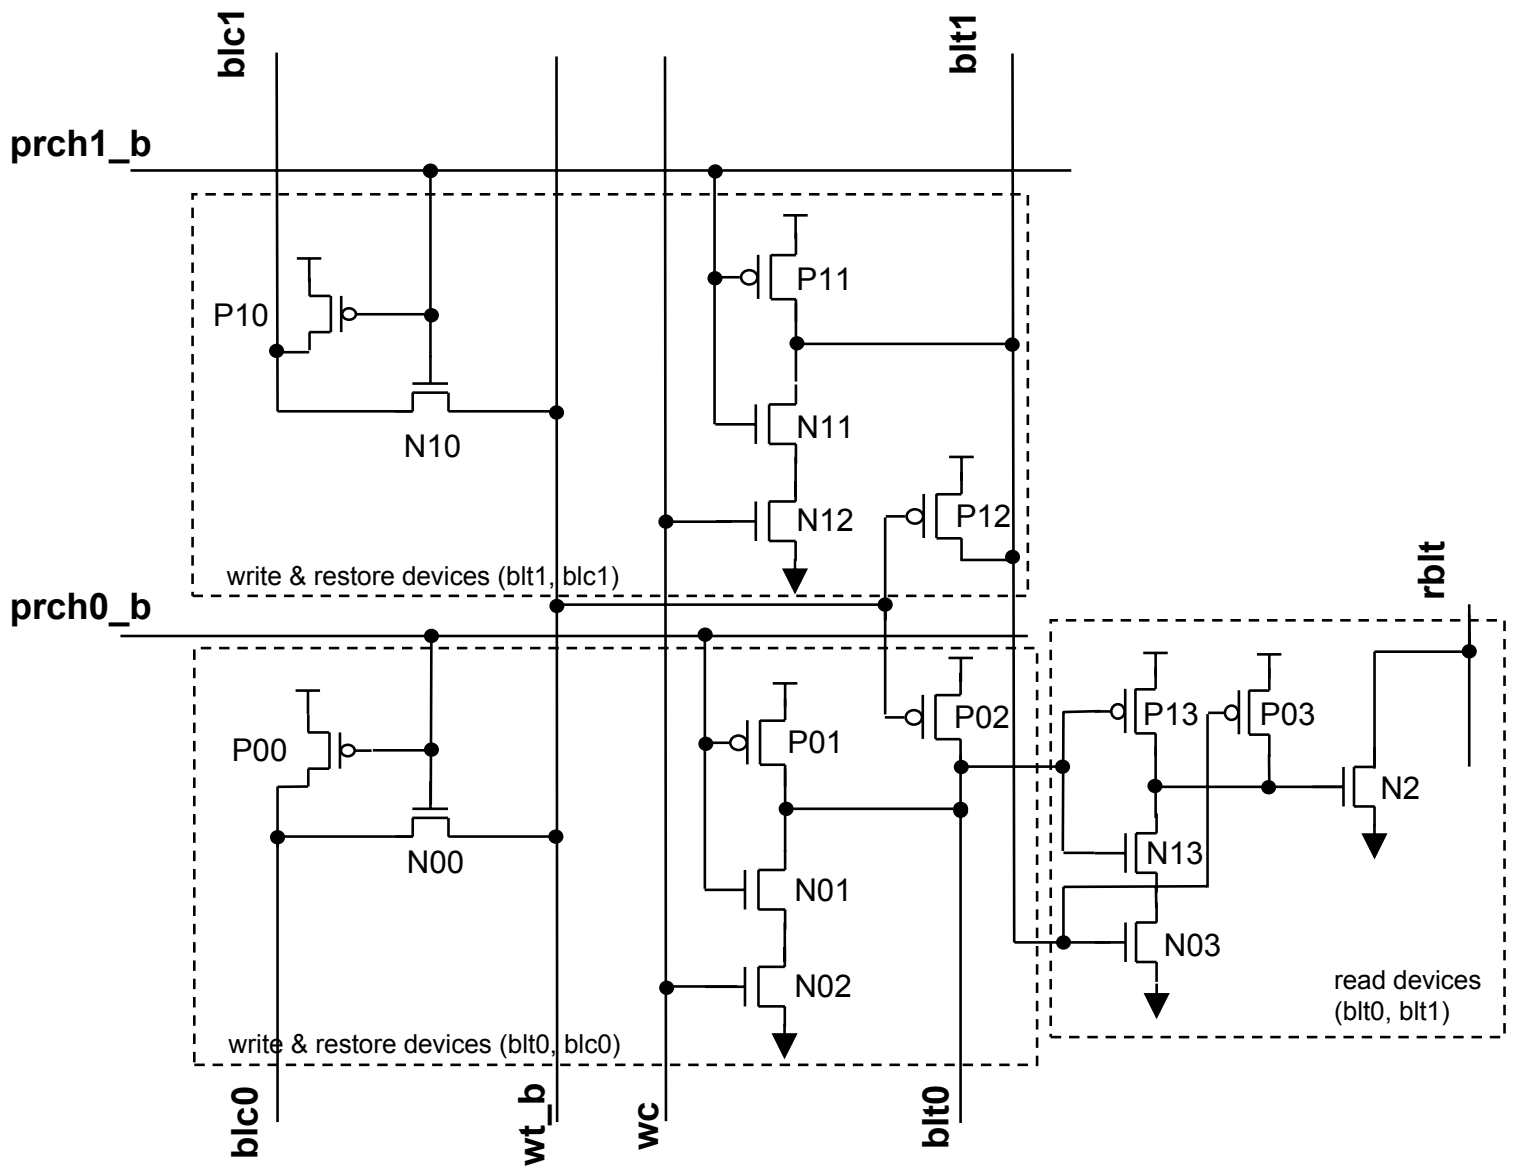


Figure 18.1.1: Local read and write circuit (loceval).

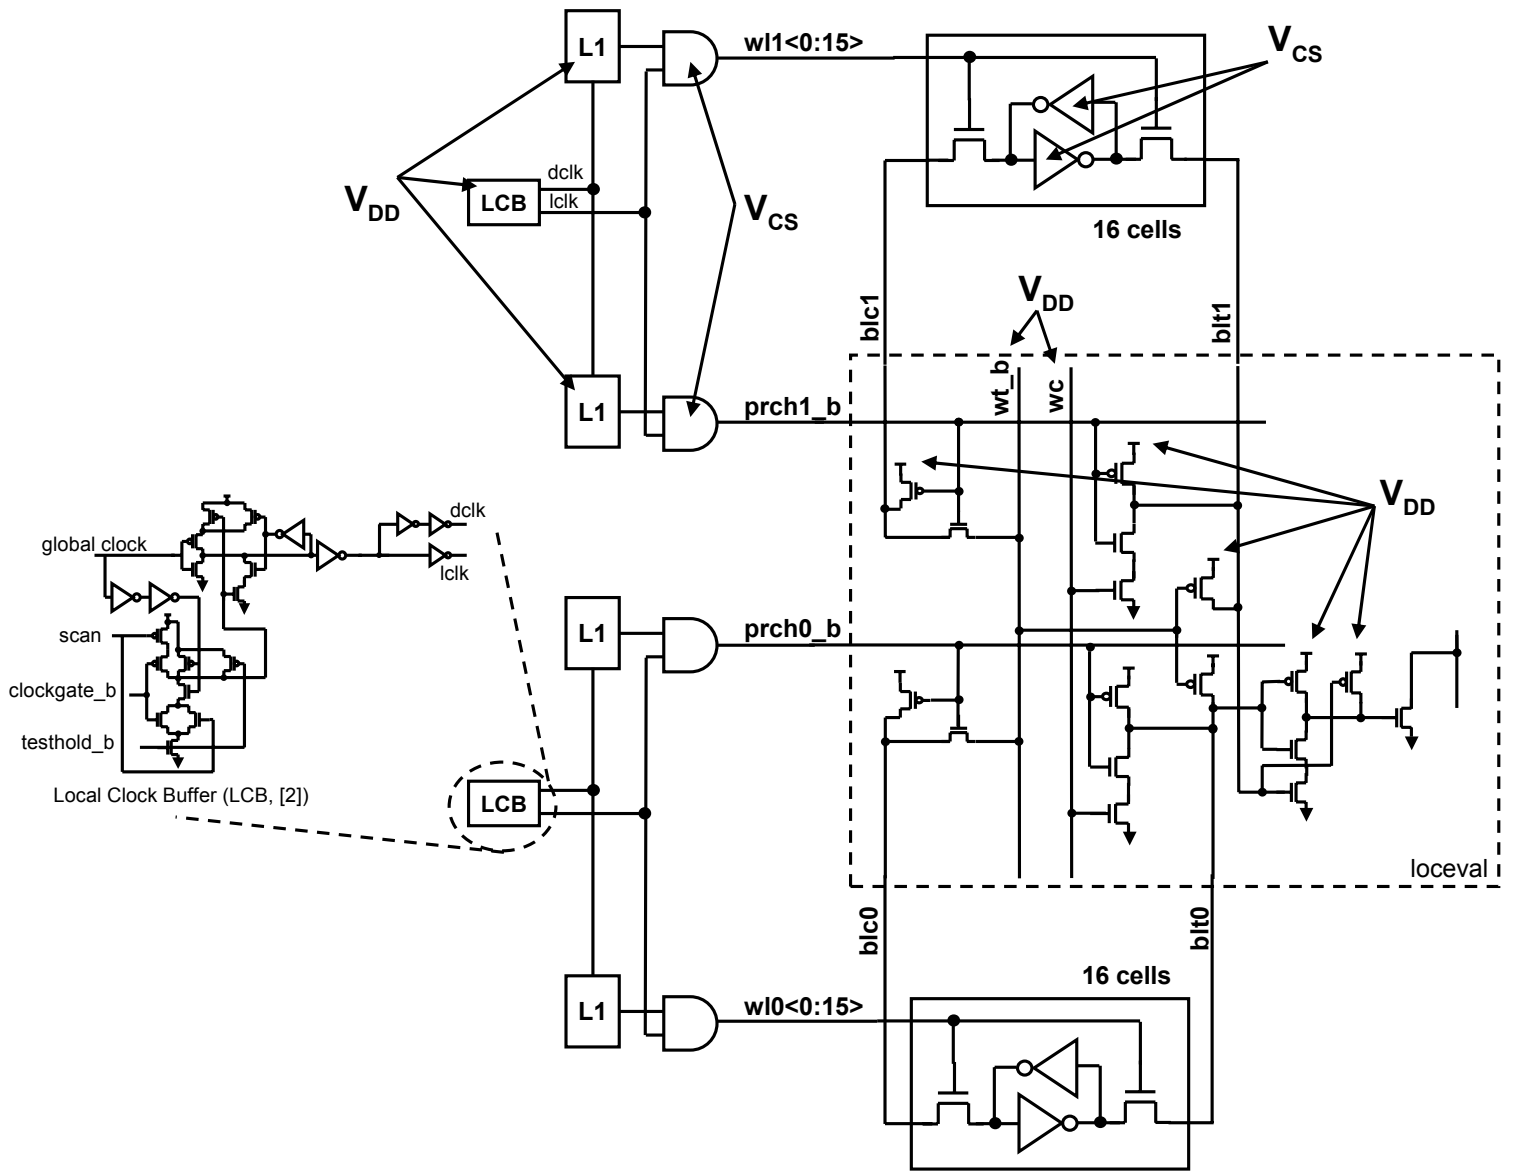


Figure 18.1.2: Voltage domains in 32-cell block including controls.

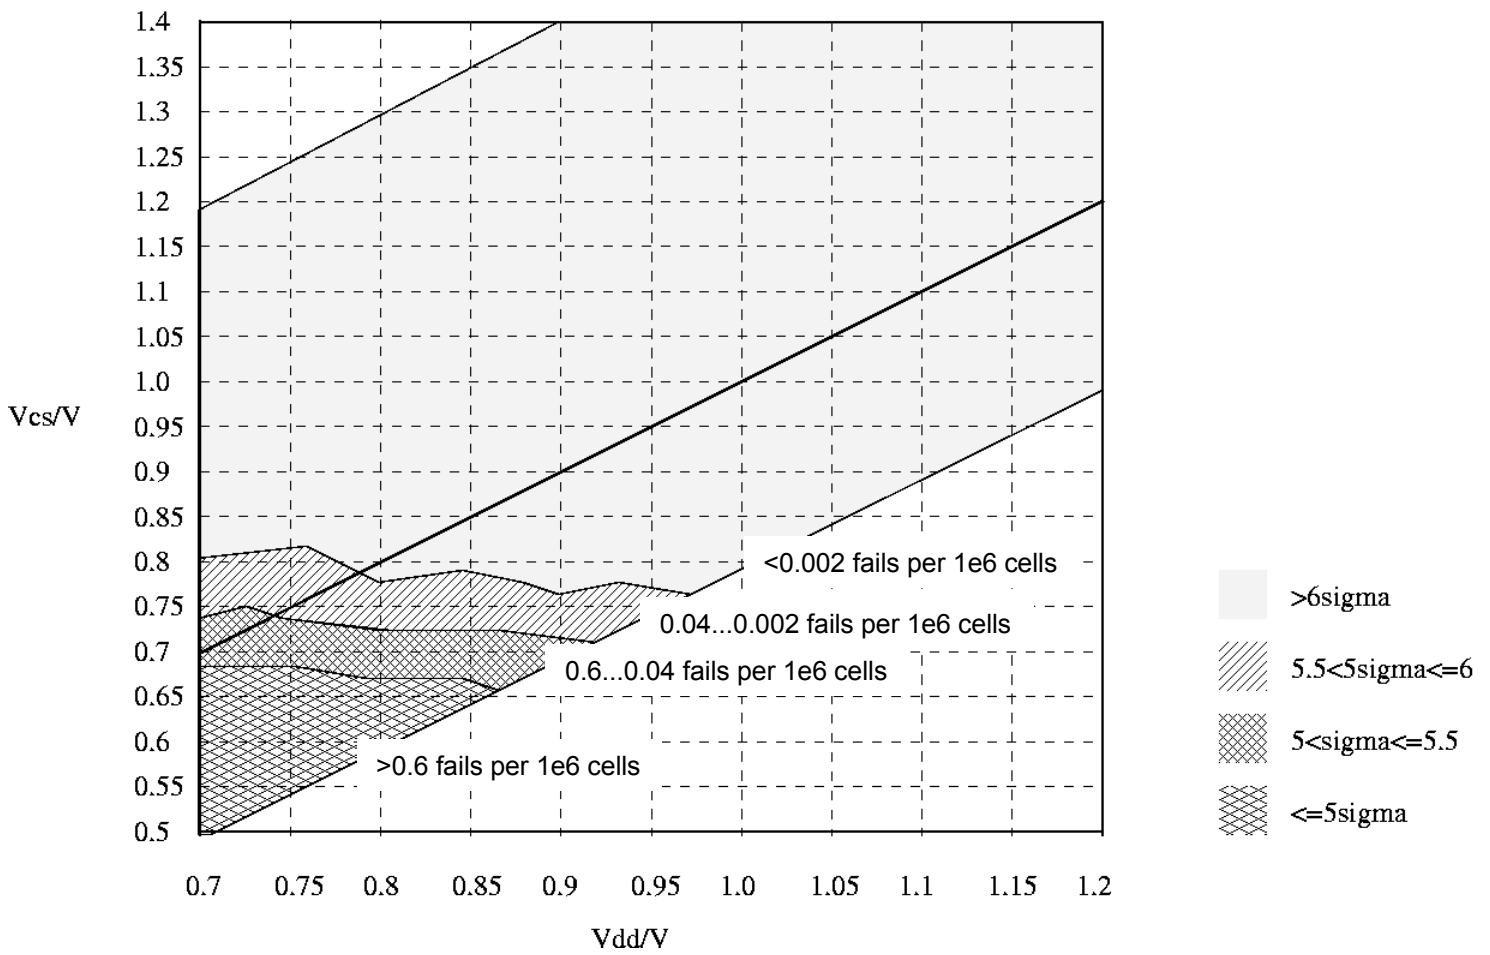


Figure 18.1.3: Cell write-ability analysis, statistical simulation assuming $\pm 6\sigma$ device variation (V_t , length, width).

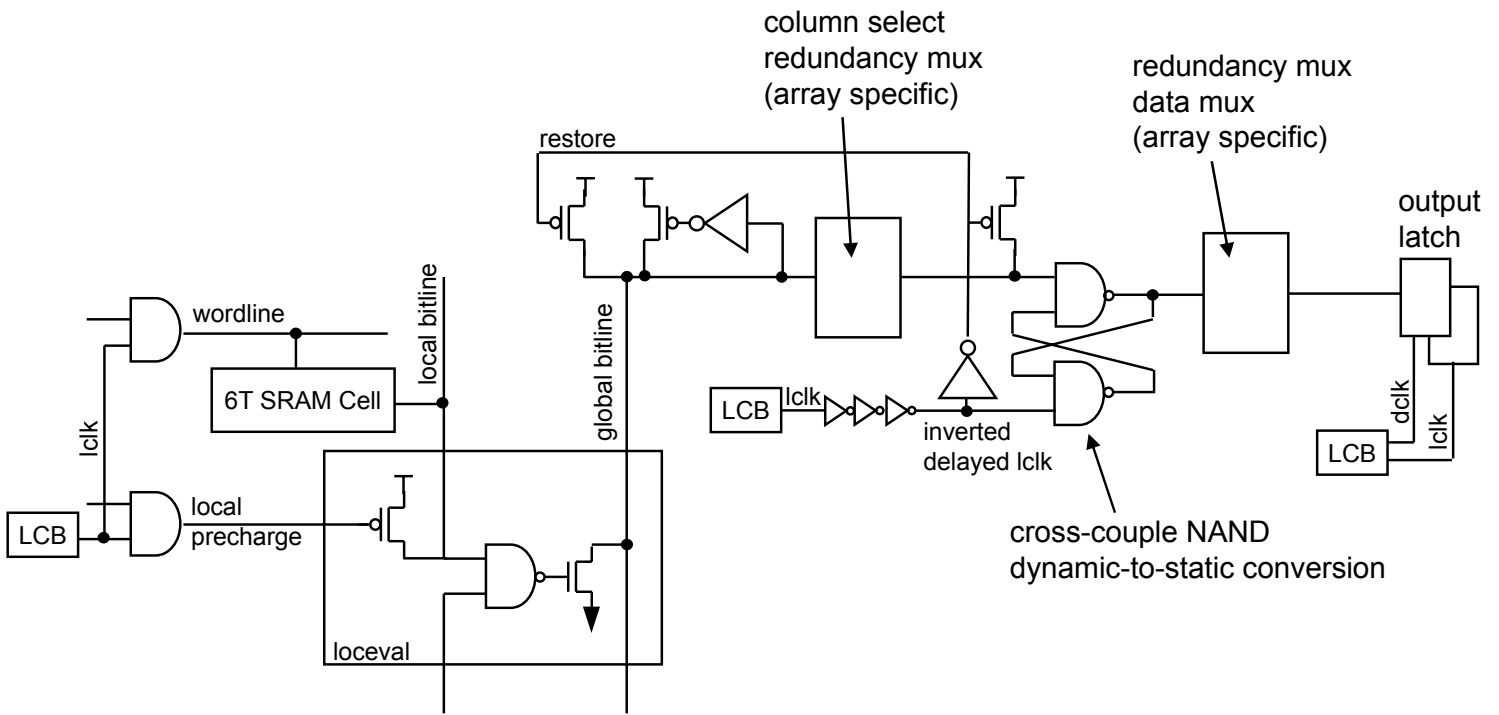


Figure 18.1.4: Typical ripple-domino read path (wordline to data out latch).


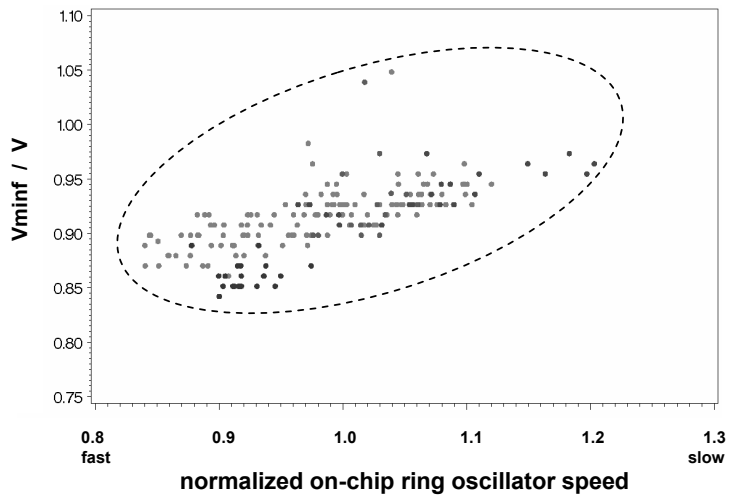


Figure 18.1.5: At-speed V_{DD}/V_{CS} shmoo plot (100mV steps).

$$V_{CS} = V_{DD}$$



$$V_{CS} = V_{DD} + 200mV$$




Figure 18.1.6: V_{minf} scatter improves with $V_{CS} > V_{DD}$.

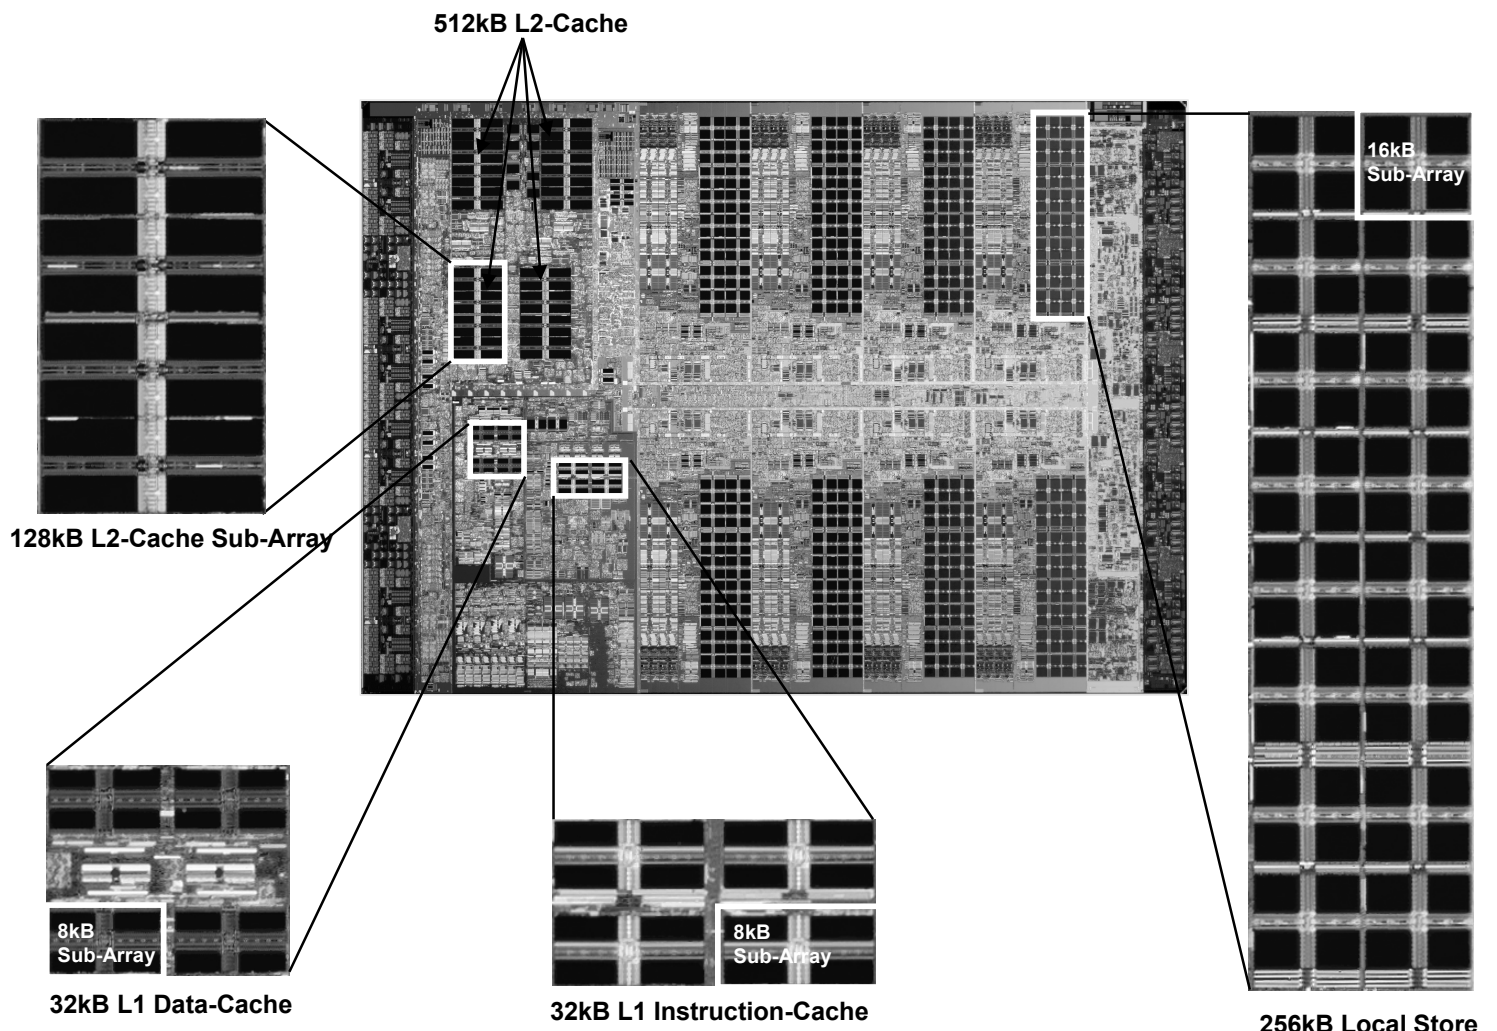


Figure 18.1.7: 65nm CELL Broadband Engine™ die micrograph and magnifications of the major arrays.