

A 130-nm 6-GHz 256×32 bit Leakage-Tolerant Register File

Ram K. Krishnamurthy, *Member, IEEE*, Atila Alvandpour, *Member, IEEE*,
Ganesh Balamurugan, *Student Member, IEEE*, Naresh R. Shanbhag, K. Soumyanath, *Member, IEEE*, and
Shekhar Y. Borkar, *Member, IEEE*

Abstract—This paper describes a 256-word \times 32-bit 4-read, 4-write ported register file for 6-GHz operation in 1.2-V 130-nm technology. The local bitline uses a pseudostatic technique for aggressive bitline active leakage reduction/tolerance to enable 16 bitcells/bitline, low- V_t usage, and 50% keeper downsizing. Gate-source underdrive of $-V_{cc}$ on read-select transistors is established without additional supply/bias voltages or gate-oxide overstress. 8% faster read performance and 36% higher dc noise robustness is achieved compared to dual- V_t bitline scheme optimized for high performance. Device-level measurements in the 130-nm technology show $703\times$ bitline active leakage reduction, enabling continued V_t scaling and robust bitline scalability beyond 130-nm generation. Sustained performance and robustness benefit of the pseudostatic technique against conventional dynamic bitline with keeper-up-sizing is also presented.

Index Terms—Bitline active leakage, dc noise robustness, dual threshold voltage, pseudostatic, register files.

I. INTRODUCTION

REGISTER files are performance-critical memory components in general-purpose microprocessors. They are required to meet two constraints: 1) single clock cycle read/write latency to support back-to-back read and read-after-write operations and 2) multiple read/write port capability, to enable the simultaneous access of several execution units in a super-scalar architecture. These requirements, coupled with the demand for a large number of word entries per port, force the usage of wide-OR style dynamic circuits for the local and global bitlines. With CMOS technology scaling, supply voltage V_{cc} and threshold voltage V_t are both scaled to maintain approximately the same V_{cc}/V_t ratio, to achieve high performance. However, aggressive V_t scaling results in an exponential increase in bitline active leakage currents [1] and poor bitline noise immunity scaling trend (Fig. 1). Therefore, alternate bitline circuit techniques that curtail this trend are required to achieve high noise immunity while sustaining high performance.




Fig. 1. Register file bitline currents and robustness scaling trend.

Previous works have involved either usage of negative word-line drivers [2] or dynamic threshold voltage adjustment via substrate/well bias control [3], [4]. However, these techniques require additional supply/bias voltages, dual gate oxides to overcome gate-oxide overstress induced on bitline transistors, and/or extra interconnect routing tracks for the control signals. These requirements either increase manufacturing cost or directly impact area due to the metal-limited nature of multiported register file layouts.

This paper describes a 256-word \times 32-bit 4-read, 4-write ported register file designed for 6-GHz operation in 1.2-V 130-nm dual- V_t bulk-CMOS technology [5]. A pseudostatic local bitline technique that establishes gate-source underdrive of $-V_{cc}$ on read-select transistors without additional supply/bias voltages or gate-oxide overstress is described. Aggressive bitline active leakage reduction enables 16 bitcells/bitline with all low- V_t transistors and good performance-robustness scaling trend for sub-130-nm register files. Simulation results on performance, robustness, and energy comparisons against conventional low- V_t and dual- V_t register files are presented. Device-level measurements that quantify the active leakage reduction benefits are also presented.

The remainder of this paper is organized as follows. Section II describes the organization of the 4-read, 4-write ported 256×32 bit register file, timing plan and clock generation circuits. In Section III, we review bitline active leakage robustness issues and scaling trends of conventional low- V_t and dual- V_t dynamic bitline schemes. Section IV discusses the proposed pseudostatic bitline technique and operation, bitline leakage impact, leakage benefit measurements in 130-nm technology, and area-performance tradeoffs. In Section V, performance-robustness comparisons and keeper sizing tradeoffs of the proposed technique against conventional low- V_t and dual- V_t dynamic bitlines are presented, followed by the conclusions in Section VI.

Manuscript received August, 2001; revised December, 2001.

R. K. Krishnamurthy, A. Alvandpour, K. Soumyanath, and S. Y. Borkar are with the Microprocessor Research Labs, Intel Corporation, Hillsboro, OR 97124 USA (e-mail: ram.krishnamurthy@intel.com; atila.alvandpour@intel.com; krishnamurthy.soumyanath@intel.com; shekhar.borkar@intel.com).

G. Balamurugan and N. R. Shanbhag are with the Department of Electrical and Computer Engineering, University of Illinois, Urbana-Champaign, IL 61801 USA (e-mail: balamuru@mail.icims.csl.uiuc.edu; shanbhag@uivlsi.csl.uiuc.edu).

Publisher Item Identifier S 0018-9200(02)03668-5.

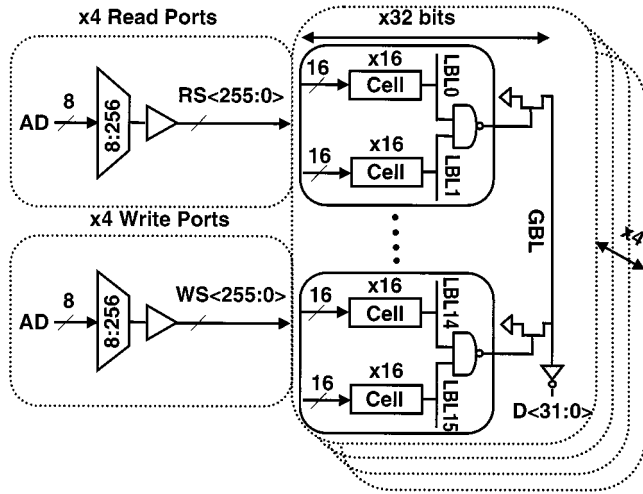


Fig. 2. 256×32 -bit register file organization.

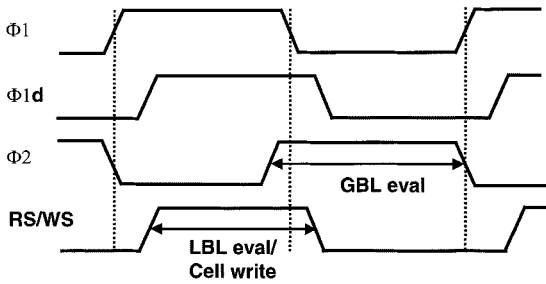


Fig. 3. 2Φ domino timing plan.




Fig. 4. Clock generator circuits.

II. 256×32 BIT REGISTER FILE ORGANIZATION

Fig. 2 shows the organization of the 4-read, 4-write ported $256\text{-word} \times 32\text{-bit/word}$ register file. The complete read operation is performed in two clock cycles. An 8-bit read/write address per port is decoded in the first cycle to deliver the read/write select signals into the register file array. The decode cycle is noncritical and therefore the decoder is implemented in conventional static CMOS using three stages of 2- and 3-input NAND and NOR gates. In the next cycle, which is performance critical, the actual bitline read operation is conducted. Fig. 3 shows the 2Φ domino timing plan for this cycle. Fig. 4 shows the clock generation and driver circuits used. At the beginning of this cycle, 256 D1 footed-domino buffers per port are triggered by $\Phi1$, buffered from the incoming 50% duty-cycle core clock CCLK and drive the decoded read/write select signals across the 32-bit array width, as shown in Fig. 5. LBL evaluation or bitcell write operation occurs during the $\Phi1$ portion of the cycle. Fig. 6 shows the register file bitcell, with symmetric loading of two read ports on each side of the storage cell for optimal cell write stability [6]. Matched pass




Fig. 5. Read/write select D1 domino drivers.

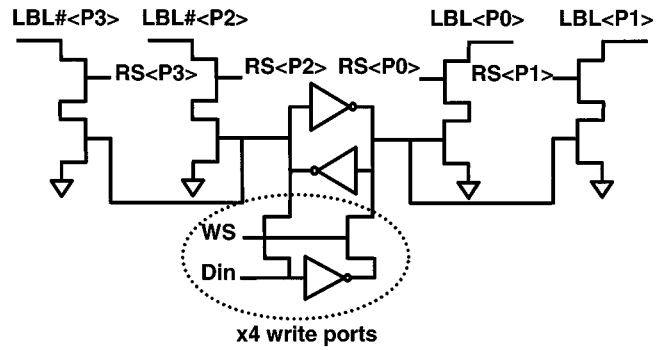


Fig. 6. Symmetric register file bitcell.

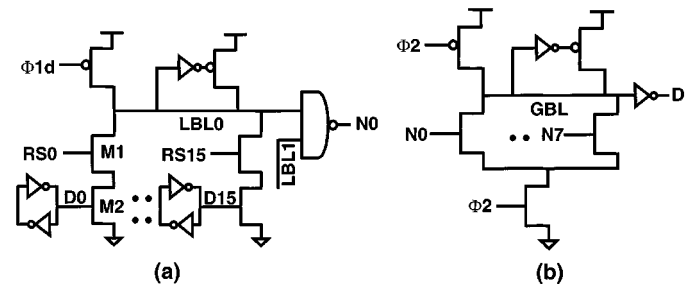


Fig. 7. (a) Local bitline scheme. (b) Global bitline scheme.

transistors are used on each side of the storage cell to enable single-ended full-swing write operation. The input data, D_{in} , is locally inverted to get its complement for the write operation. Fig. 7(a)–(b) shows the full-swing local bitline (LBL) and global bitline (GBL) scheme. Each LBL (1 per read port) supports 16 bitcells and a two-way merge via a static NAND gate. Data from the storage cell is read by two access transistors per word (M1 and M2) on each LBL and thus forms a dynamic 16-way NAND–NOR. The LBL is triggered by a delayed version




Fig. 8. 130-nm 256 × 32-bit register file layout.




Fig. 9. LBL setup for worst-case read-select input noise.

of Φ_1 , shown in Fig. 3 as Φ_{1d} and is fully time borrowable enabling the read/write select signals to arrive into the LBL evaluate phase. The GBL is a dynamic 8-way OR which merges the LBL 2-input static NAND outputs to deliver a 32-bit word per read port. GBL operation is conducted during Φ_2 of the cycle and to avoid precharge races and crowbar currents at the phase boundary, the GBL domino is footed by the clock transistor. GBL phase is also fully time borrowable enabling the LBL NAND-merge outputs to arrive into Φ_2 . The Φ_2 clock is locally generated and is an inverted version of the Φ_1 clock, offering an automatic stretching capability (Fig. 4). Thus, when Φ_1 core clock period is reduced for slow-frequency test and debug, the Φ_2 clock stretches out in proportion. Φ_2 's rising edge is positioned slightly ahead of Φ_1 's falling edge to achieve good tolerance to Φ_1 clock skew, jitter, and duty-cycle variations.

This register file organization enables the complete LBL and GBL operation to be performed in four gate stages. An equivalent static CMOS LBL and GBL implementation will require seven stages of 2- and 3-input NAND-NOR gates. Single-ended read/write selects and bitline signaling are used throughout to reduce wiring congestion and enable a dense layout occupying $356 \mu\text{m} \times 89 \mu\text{m}$ (Fig. 8), with maximum LBL and GBL lengths of $115 \mu\text{m}$ and $1092 \mu\text{m}$, respectively.

III. BITLINE LEAKAGE ROBUSTNESS

A. Low- V_t Leakage Impact

LBL and GBL dynamic ORs are susceptible to noise due to high active leakage during evaluation when the precharged domino node should stay high. LBL is particularly more sensitive than GBL due to smaller domino node stored charge ($0.1 \times$ compared to GBL) and a wider dynamic OR structure (16-way for LBL versus 8-way for GBL). Fig. 9 shows the LBL setup for worst-case noise scenario with all low- V_t transistors, i.e., optimized for best read performance. After the LBL has been precharged to V_{CC} and the clock transitions high, the bitline




Fig. 10. LBL dc droop in 130-nm process at 1.2 V, 110 °C.

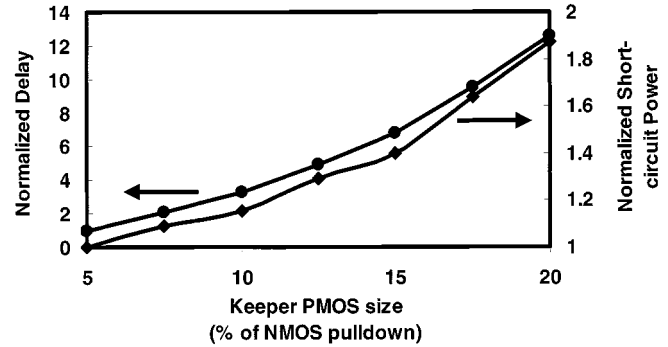
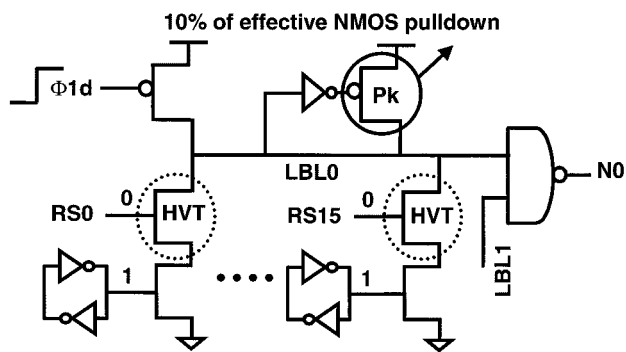
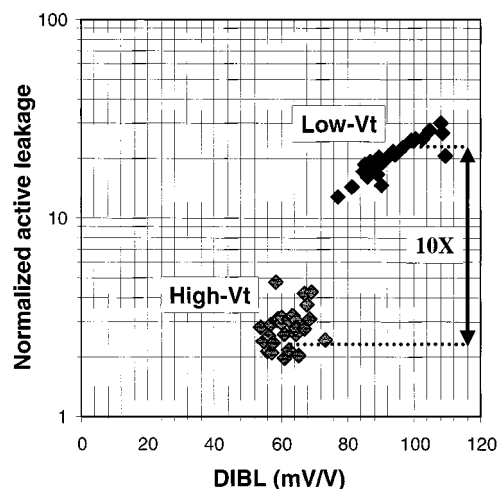


Fig. 11. Keeper upsizing versus delay and short-circuit power in 130-nm process at 1.2 V, 110 °C.

leakage is maximized when all bitcell data inputs D0-D15 are held high (V_{CC}) and all read-select inputs RS0-RS15 are at GND, i.e., the bitline is deselected. The pMOS keeper transistor (P_k) size is 10% of the effective nMOS pulldown strength. Fig. 10 shows the LBL domino node behavior when the clock stays high for a time window significantly longer than the operating frequency, a situation encountered during stop-clock mode or slow-frequency testing. During this condition, the precharged bitline is expected to retain state at V_{CC} , held actively by the pMOS keeper. However, even without any read-select input noise, a 50-mV dc droop (4% of V_{CC}) is observed on the bitline at the worst-case process leakage corner in the 1.2-V 130-nm technology [7]. With dc noise on the read-select inputs, the LBL unity gain dc noise margin (monitored at the NAND gate output) is 86 mV (7.2% of V_{CC}) and does not meet the minimum noise margin floor requirement of 150 mV (12.5% of V_{CC}). This minimum noise margin floor stems from realistic 130-nm supply/ground bounce and coupling noise specifications. Low- V_t on the GBL pulldown transistors achieves sufficient noise immunity for the reasons mentioned above, and we will consequently focus on the LBL robustness in the remainder of this paper.

B. PMOS Keeper Upsizing Tradeoff



A straightforward solution to improving bitline robustness is to strengthen the pMOS keeper transistor for a given effective nMOS pulldown strength. This compensates for the bitline leakage by increasing the dynamic node replenished charge. However, this also proportionately increases contention on the bitline during pulldown which degrades read delay and increases bitline short-circuit power [8]. Fig. 11 shows the delay and short-circuit power increase (at the 1.2-V 110 °C

Fig. 12. Dual- V_t LBL scheme.Fig. 13. 130-nm low- V_t versus high V_t active leakage measurements.

operating corner) as the LBL keeper is upsized from 5% to 20% of the effective nMOS pulldown strength in the 130-nm technology: nearly 3% delay and 0.3% short-circuit power penalty is observed for every 1% increase in keeper pMOS size. A more effective solution is to replace the pulldown transistors with less leaky high- V_t transistors in a dual- V_t technology.

C. Dual Threshold Voltage Solution

The dual- V_t LBL scheme optimized for high-performance uses high- V_t on the read-select transistors (M1) and low- V_t on the bitcell data transistors (M2), as shown in Fig. 12. Low- V_t is used on all other transistors in the register file for best performance. There is no layout area penalty since the high- and low- V_t transistors occupy the same layout footprint. Fig. 13 shows low- V_t and high- V_t normalized active leakage measured across several dies in 130-nm technology at 1.2-V 80 °C for transistor dimensions comparable to the bitline nMOS pulldown transistors. The high- V_t transistor active leakage is nearly 10 \times lower than for low- V_t for the same dimensions. From the LBL inputs set up for the worst-case noise scenario shown in Fig. 9, the high- V_t M1 transistors now limit the bitline leakage. This implies that the worst-case domino node active leakage is reduced by 10 \times , offering considerable improvement in LBL robustness. However, this benefit is achieved at the cost of degraded performance due to reduced drive currents of high- V_t transistors. This performance-robustness impact will be evaluated in Section V.

Fig. 14. Low- V_t and high- V_t robustness scaling trend.Fig. 15. Pseudostatic low- V_t bitline technique.

Although the dual- V_t scheme offers a good one-time fix for bitline leakage in 130-nm generation, future process generations require both the high- and low- V_t 's to scale with supply voltage. The motivation to scale both V_t 's stems from the fact that *both* high- and low- V_t circuit performance improvement should track with process scaling. If one of the V_t 's, e.g., high- V_t , is un-scaled or scaled less aggressively than low- V_t , corresponding high- V_t circuits manifest themselves as speed critical, limiting overall performance scaling. This V_t scaling trend in turn results in nearly 3–5 \times increase in both the high- V_t and low- V_t transistor leakage [9]. Fig. 14 shows the sub-180-nm scaling trend of dc robustness, defined as the dc unity gain noise margin normalized to supply voltage, averaged across several representative dynamic circuits. The desired trend is for robustness to remain constant with technology scaling, i.e., dc noise margin scales linearly with supply voltage, similar to conventional static CMOS circuits. High- V_t usage clearly offers the desired robustness benefit in 180-nm \rightarrow 130-nm scaling, actually improving the robustness by 3% over 180-nm generation. Sub-130-nm robustness continues to degrade by 15%–30% for both high- V_t and low- V_t , and we are back on the same treadmill.

IV. PSEUDOSTATIC BITLINE TECHNIQUE

A. Circuit Description and Operation

Fig. 15 shows the proposed pseudostatic leakage-tolerant LBL technique. There are three main modifications to a conventional dynamic bitline topology.

- 1) Read-select input and bitcell data locations on the bitline stack are swapped. Read-select signals feed the lower (M2) transistors of the LBL.

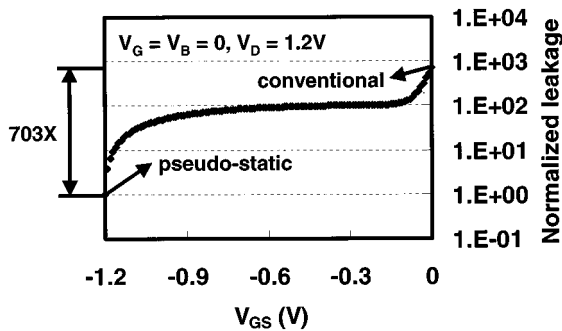


Fig. 16. 130-nm low- V_t nMOS V_{GS} underdrive active leakage measurements.

- 2) Static-precharge transistors (P_x) driven actively by the read-select signals are introduced. These transistors anchor the bitline stack nodes (V_s) at V_{CC} when read-selects are at GND.
- 3) Static 2-input NOR gates are introduced, whose inputs are the stack node and bitcell data#. The NOR-gate outputs drive the upper (M1) transistors of the LBL.

Overall, five additional transistors per read port are added to the bitcell.

When the read-select inputs are at GND, the NOR-gate outputs force the leakage-limiting M1 transistor inputs to GND. Any of the NOR-gate outputs conditionally transitions high only when the appropriate read-select input transitions high. Thus, for a precharged bitline, when any one of the 16 bitline stacks is deselected, i.e., read-select is at GND, a $-V_{CC}$ gate-source underdrive on the leakage-limiting M1 transistor of that stack is established. This effectively cuts off the bitline subthreshold active leakage current path, since both drain and source of the M1 transistor are at V_{CC} . In addition, the body effect on the M1 transistor is maximized due to full V_{CC} of source-body bias, which further elevates their V_t .

B. Bitline Leakage Impact

To evaluate the benefit of the pseudostatic technique and the contribution of each of the above factors to total bitline leakage reduction, we first measure this impact on 130-nm transistors. Fig. 16 shows device-level measurements across several dies of normalized active leakage reduction achieved on low- V_t nMOS transistors with gate-source underdrive in 130-nm technology at 1.2-V 80°C for transistor dimensions comparable to the bitline nMOS pulldown transistors. Similar to the leaking nMOS pulldowns of a deselected precharged bitline, the gate voltage V_G and body voltage V_B are set up at GND and the drain voltage V_D is set up at V_{CC} . Leakage is measured between the drain terminal and GND. In a conventional bitline, the source voltage would be at GND, whereas in the pseudostatic bitline, the source voltage would be V_{CC} . As source voltage increases from GND to V_{CC} , the gate-source underdrive V_{GS} changes from 0 (conventional dynamic bitline situation) to $-V_{CC}$ (pseudostatic bitline situation). At the same time, the drain-source voltage also decreases from V_{CC} to 0, effectively shutting off the drain-source subthreshold leakage path. The body effect is also maximized since $V_{BS} = -V_{CC}$ at this condition. Note that the drain junction diode leakage still exists and therefore drain leakage never converges to zero. The compounded leakage reduction effects re-




Fig. 17. Pseudostatic versus dual- V_t dynamic LBL robustness scaling trend.

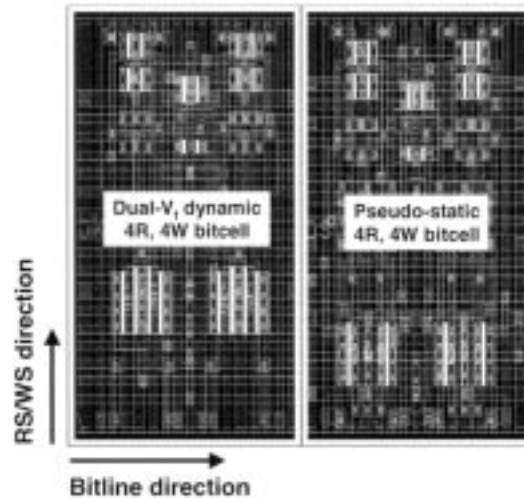


Fig. 18. 130-nm dual- V_t dynamic and low- V_t pseudostatic bitcell layouts.

sult in a total active leakage reduction of 703 \times , even with using low- V_t transistors (70 \times better than the benefit of replacing with dual- V_t). This implies that when the bitline is deselected, the bitline leakage corresponds to that of about four process generations before, whereas when the bitline is selected for evaluate operation, the pulldown transistor performance is that of the current process generation. This demonstrates the good scaling behavior beyond 130-nm technology, while allowing the use of performance-setting low- V_t transistors. Fig. 17 captures this sub-130-nm scaling trend: while the conventional low- V_t and dual- V_t dynamic bitlines display 15%–30% robustness degradation per generation, the pseudostatic bitline scheme displays less than 5% robustness degradation per generation. Consequently, the robustness benefit improves substantially with technology scaling, reaching 106% at the 70-nm node.

C. Area Impact

No explicit supply/bias voltages or additional interconnect tracks are distributed into the register file, since the $V_{GS} = -V_{CC}$ effect is achieved only via circuit means. Further, the worst-case gate-oxide stress between the gate and any terminal of the read-select transistors does not exceed $|V_{CC}|$ —a valuable cost benefit since the process does not have to support multiple gate-oxide thicknesses. Overall, the transistor active area is increased by 10% due to the static precharge and NOR gate devices, which are typically much smaller than the bitline nMOS pulldown devices. However, this overhead is easily embedded within the existing low- V_t and dual- V_t register file layout, which is interconnect

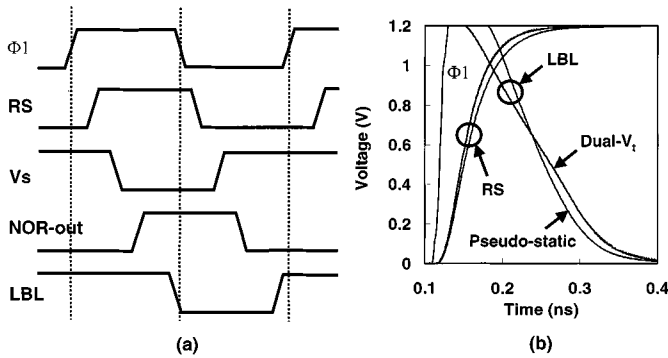


Fig. 19. (a) Pseudostatic LBL evaluate operation. (b) LBL pulldown edge-rate comparisons versus dual- V_t LBL.

limited. Fig. 18 shows the layout of the conventional dynamic and pseudostatic 4-read, 4-write ported register file bitcells in 130-nm technology. Fully shielded, aggressive pitch read/write select and bitline wiring is employed. The pseudostatic bitcell fits within the same layout template as the conventional dynamic bitcell, demonstrating that the pseudostatic bitline technique can be implemented as a drop-in replacement to existing register files.

D. Performance Impact

There are two main factors that degrade pseudostatic bitline evaluate performance compared to conventional dynamic bitlines. First, the static precharge transistor increases read-select fan-in capacitive load by 6%. This slows down the read-select drivers and the associated delay penalty directly impacts clock period. Although the read-select buffers can be upsized to absorb this additional load, this was not done to keep the comparisons against conventional dynamic bitline technique fair. It is important to note that the LBL nMOS pulldown stack height is still two, as opposed to other proposed leakage-tolerant techniques which increase the stack height [10]. Such techniques require read-select signals to fanin into two or more of the nMOS stacked transistors and result in $> 50\%$ increase in read-select fanin load. Second, the 2-input NOR pullup directly adds to the critical path delay. Fig. 19(a) shows the pseudostatic LBL timing for evaluate operation. Initially, RS transitions high causing V_s to discharge to GND. This in turn, triggers the 2-input NOR gate to pullup. The NOR output turns on M1, triggering the LBL discharge—a total of three “inversions” for the pulldown operation. In a conventional LBL evaluate operation the RS transitioning high directly triggers the LBL discharge—a single inversion operation, i.e., two inversions fewer. Both these delay-degrading penalties are more than compensated by two delay-improving circuit changes enabled by the aggressive bitline leakage reduction: 1) low- V_t usage throughout, including bitline pulldown transistors, and 2) 50% downsized keeper pMOS transistor, i.e., downsized to 5% of effective nMOS pulldown strength, thereby mitigating the bitline contention during evaluate. Both these factors result in a fast bitline pulldown operation as captured in Fig. 19(b), which compares the low- V_t pseudostatic and dual- V_t dynamic LBL discharge waveforms. The pseudostatic LBL has a 36% pulldown edge-rate improvement for the same pulldown transistor sizes as the dynamic LBL. This improved edge-rate

TABLE I
130-NM SUMMARY OF TRANSISTOR CHARACTERISTICS [7]

Parameter	Value
T_{OX} [nm]	1.5
I_{OFF} high- V_t [nA/ μm]	10
I_{DSAT-N} high- V_t [mA/ μm]	1.02
I_{DSAT-P} high- V_t [mA/ μm]	0.5
I_{OFF} low- V_t [nA/ μm]	100
I_{DSAT-N} low- V_t [mA/ μm]	1.17
I_{DSAT-P} low- V_t [mA/ μm]	0.6

TABLE II
130-NM SIMULATION RESULTS AT 1.2 V, 110 °C

LBL Scheme	Read Delay	DC robustness (DC noise margin / V_{cc})	Energy/transition
Low- V_t	158ps	0.072	3.54pJ
Dual- V_t	178ps	0.157	3.36pJ
This work	165ps	0.214	3.61pJ

further speeds up the 2-input NAND gate merge immediately following the LBL. These factors result in the low- V_t pseudostatic LBL technique achieving faster read performance than conventional dual- V_t dynamic LBL technique.

V. PERFORMANCE-ROBUSTNESS COMPARISONS AND DISCUSSION

A. Performance Comparisons

The conventional low- V_t and dual- V_t dynamic and the low- V_t pseudostatic register files are all optimally sized and simulated in the 130-nm technology at 1.2-V 110°C. In all cases, a Lagrange-multiplier based quadratic optimizer is used for optimal transistor sizing to make the comparisons fair and to achieve a robust optima centered around the “knee” of their respective energy-delay curves [11]. Although the low- V_t dynamic LBL does not meet minimum noise margins (for reasons mentioned in Section III), comparisons against this register file are included for completeness. All simulations include full layout extracted and back-annotated device and interconnect parasitics. Performance and energy simulations are conducted at the nominal process leakage corner. All robustness comparisons are performed at worst-case process leakage corner, since full functionality is required at this corner.

Table I shows the 130-nm high- V_t and low- V_t transistor characteristics. Saturation drive currents for high- V_t transistors at I_{OFF} value of 10 nA/ μm are 12%–17% lower than the saturation drive currents for low- V_t transistors at I_{OFF} value of 100 nA/ μm . Also, the nMOS low- V_t /high- V_t I_{DSAT} 's are approximately $2\times$ higher than the corresponding pMOS low- V_t /high- V_t I_{DSAT} 's. A typical low- V_t /high- V_t fanout of 4 inverter delay in this technology at the 1.2-V 110°C operating corner is 25 ps/30 ps. Table II shows the read delay (this effectively

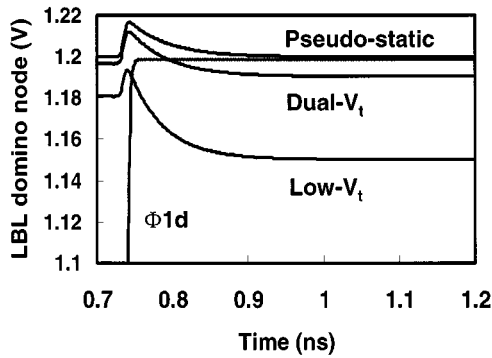


Fig. 20. LBL dc droop comparisons in 130-nm process at 1.2 V, 110 °C.

sets the clock period), dc noise robustness, and energy/transition comparisons between low- V_t dynamic, dual- V_t dynamic and low- V_t pseudostatic register files. The complete read-path, which includes read-select buffers, LBL and GBL delays, operates at 165 ps with the proposed technique, i.e., 8% faster than dual- V_t and only 4% slower than the all low- V_t best-speed solution. As discussed in Section IV, the speed improvement of the pseudostatic technique over dual- V_t is attributed to the low- V_t nMOS pulldown benefit, low bitline keeper contention, and improved NAND merge delay. Of course, in comparison against a low- V_t dynamic LBL, the pseudostatic technique's reduced contention, downsized keeper, and better NAND merge delay is inadequate to offset the inherent delay penalty, resulting in the 4% slowdown. The energy/transition, which includes active leakage, is 7% higher than dual- V_t and 2% higher than low- V_t for the pseudostatic register file. This is due to the increased switched capacitance penalty of the additional transistors and their associated active leakage.

B. Robustness Comparisons

To compare dc robustness of the LBL techniques, worst-case read-select dc noise conditions are set up for each. For the low- V_t and dual- V_t dynamic LBL, bitcell data inputs (D0-D15 in Fig. 9) are set to V_{cc} since that turns on the lower nMOS pull-down (M2) transistors and grounds all the stack nodes. For the pseudostatic LBL, bitcell data# inputs (D0#-D15# in Fig. 15) are set to GND, since that turns on the upper pMOS transistors of the 2-input NOR gates. This condition enables any noise on read-select signals to propagate onto the NOR outputs. The dc noise offset above GND is applied on all the read-select inputs (RS0-RS15), representing input-referenced ground bounce and coupling noise propagated from the previous read-select driver stages. DC robustness is evaluated as unity-gain noise margin at the output of the LBL-merge 2-input NAND gate normalized to V_{cc} .

Fig. 20 shows the LBL domino node dc droop comparisons with no input noise on the read-selects and at the worst-case leakage corner. The dc droop of the proposed scheme is $5\times$ ($25\times$) lower than dual- V_t (low- V_t), exhibiting an almost ideal dynamic node behavior. This droop benefit is directly due to the bitline leakage reduction during this circuit condition. When dc noise is applied on the read-select inputs at the worst-case leakage corner, in the dynamic LBL this noise directly appears as V_{GS} on the read-select transistors and triggers a bitline mi-

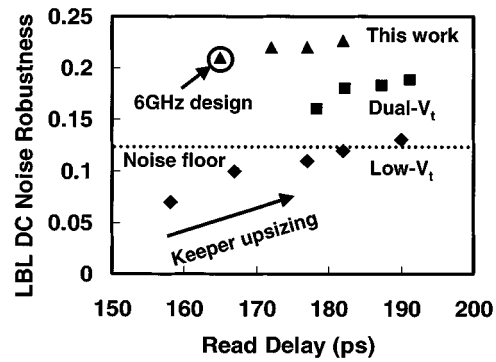


Fig. 21. LBL keeper upsizing comparisons in 130-nm process at 1.2 V, 110 °C.

sevaluation. In the pseudostatic LBL, noise on the read-select inputs goes through two stages of attenuation—from input to stack node and from stack node to NOR output—before the input noise manifests itself as V_{GS} on the M1 transistors. This, coupled with the inherently lower dc droop to begin with, results in the pseudostatic LBL demonstrating 36% (197%) higher dc robustness than dual- V_t (low- V_t) scheme. As technology scales, both the dc droop and dc robustness benefits of the pseudostatic technique amplifies due to the exponential scaling trend of transistor leakage.

C. Keeper Sizing Versus Robustness Tradeoffs

As mentioned in Section III, the bitline keeper pMOS transistor can be upsized to trade off delay for better robustness on all LBL schemes. Fig. 21 shows the LBL dc robustness versus the critical read delay tradeoff for LBL keeper upsizing beyond their respective baseline sizes on all the three register files. Low- V_t LBL shows a poor tradeoff of 4% robustness improvement for every 1% delay penalty and requires the keeper pMOS upsized to 20% of the effective nMOS pulldown strength to meet minimum noise margin requirements. The pseudostatic register file sustains both robustness and speed advantages over low- V_t and dual- V_t techniques even with keeper upsizing.

VI. CONCLUSION

A pseudostatic bitline scheme is described that establishes $-V_{cc}$ gate-source underdrive on the leakage-limiting read-select transistors without distributing additional supply/bias voltages or gate-oxide overstress. Using this technique, we have presented a 6-GHz 4-read, 4-write ported 256-word \times 32-bit register file in 1.2-V 130-nm technology. 8% faster read performance with a simultaneous 36% increase in dc noise robustness is achieved when compared to a dual- V_t bitline scheme optimized for high performance. The aggressive bitline active leakage reduction enabled 16 bitcells/bitline, low- V_t usage, and 50% keeper downsizing, and demonstrates good performance-robustness scaling trend for sub-130-nm register files.

ACKNOWLEDGMENT

The authors wish to thank M. Patyra, V. De, S. Hsu, and C. Webb for discussions, B. Bloechel, S. Narendra, and M. Anders for measurements, C. Lobo, L. Irvan, and S. Apperson for layout help, and R. Hofsheier and J. Rattner

for their encouragement and support. The authors also wish to thank the reviewers for several valuable suggestions for improving the quality of this paper.

REFERENCES

- [1] K. Agawa, H. Hara, T. Takayanagi, and T. Kuroda, "A bitline leakage compensation scheme for low-voltage SRAMs," *IEEE J. Solid-State Circuits*, vol. 36, pp. 726–734, May 2001.
- [2] H. Tanaka *et al.*, "A precise on-chip voltage generator for a gigascale DRAM with a negative wordline scheme," *IEEE J. Solid-State Circuits*, vol. 34, pp. 1084–1090, Aug. 1999.
- [3] H. Kawaguchi, Y. Itaka, and T. Sakurai, "Dynamic leakage cut-off scheme for low-voltage SRAMs," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 1998, pp. 140–141.
- [4] T. Kuroda *et al.*, "A 0.9 V 150 MHz 10 mW 4 mm² 2-D discrete cosine transform core processor with variable threshold voltage scheme," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1770–1779, November 1996.
- [5] R. Krishnamurthy, A. Alvandpour, G. Balamurugan, N. Shanbhag, K. Soumyanath, and S. Borkar, "A 0.13 μ m 6-GHz 256 × 32 bit leakage-tolerant register file," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 2001, pp. 25–26.
- [6] M. Golden and H. Partovi, "A 500-MHz write-bypassed 88-entry 90-bit register file," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 1999, pp. 105–108.
- [7] S. Tyagi *et al.*, "A 130-nm generation logic technology featuring 70-nm transistors, dual-V_t transistors and 6 layers of Cu interconnects," in *IEDM Tech. Dig.*, 2000, pp. 567–570.
- [8] A. Alvandpour, R. Krishnamurthy, K. Soumyanath, and S. Borkar, "A conditional keeper technique for sub-130-nm wide dynamic gates," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 2001, pp. 29–30.
- [9] M. Anders, R. Krishnamurthy, K. Soumyanath, and R. Spotten, "Robustness of sub-70-nm dynamic circuits: Analytical techniques and scaling trends," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 2001, pp. 23–24.
- [10] G. Balamurugan and N. Shanbhag, "A twin-transistor noise-tolerant dynamic circuit technique," *IEEE J. Solid-State Circuits*, vol. 36, pp. 273–280, Feb. 2001.
- [11] G. Anderson, S. Duvall, R. Gatlin, and J. Eiles, *OPTIM2 v. 1.5.1 User's Guide*, 2001.



Ram K. Krishnamurthy (S'92–M'98) received the B.E. degree in electrical engineering from Regional Engineering College, Trichy, India, in 1993 and the Ph.D. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 1998. His Ph.D. research was focused on low-power DSP circuit design.

Since 1998, he has been with Intel Corporation's Microprocessor Research Labs, Hillsboro, OR, where he is currently leading the high-performance and low-voltage circuits research group. He is an adjunct faculty of Department of Electrical and Computer Engineering, Oregon State University, Corvallis, where he teaches VLSI system design. He serves on Intel's SRC review committee and the ASIC, CICC, and ISCAS program committees. He holds ten patents and has published 20 papers.



Atila Alvandpour (M'99) received the M.Sc. and Ph.D. degrees from Linköping University, Linköping, Sweden, in 1995 and 1999, respectively. His Ph.D. research was on low-power high-performance circuit techniques and power estimation methods.

Since 1999, he has been with the high-performance circuits research group at Intel Corporation's Microprocessor Research Labs, Hillsboro, OR. He has published over 20 papers in VLSI and holds three patents.



Ganesh Balamurugan (S'99) received the B.Tech. degree in electronics and communication engineering from the Indian Institute of Technology, Madras, India, in 1996 and the M.S. degree in electrical and computer engineering from the University of Texas at Austin in 1998. He is currently working toward the Ph.D. degree in the Electrical and Computer Engineering Department of the University of Illinois at Urbana-Champaign.

During the summer of 1999, he was with Sun Microsystems working on power distribution modeling and on-chip inductance effects. In the summer of 2000, he worked at Intel Corporation on high-performance leakage-tolerant circuit design techniques. His current research interests include high-speed I/O circuits and noise-tolerant circuit design.



Naresh R. Shanbhag received the B.Tech. degree from the Indian Institute of Technology, New Delhi, India, in 1988, the M.S. degree from Wright State University, Dayton, OH, in 1990, and the Ph.D. degree from the University of Minnesota, Minneapolis, in 1993, all in electrical engineering.

From July 1993 to August 1995, he was with AT&T Bell Laboratories, Murray Hill, NJ, where he was responsible for the development of VLSI algorithms, architectures, and implementation of broadband data communications transceivers. In particular, he was the lead chip architect for AT&T's 51.84 Mb/s transceiver chips over twisted-pair wiring for asynchronous transfer mode (ATM)-LAN and broadband access chip sets. Since August 1995, he has been with the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory, University of Illinois, Urbana-Champaign, where he is currently an Associate Professor and the Director of the Illinois Center for Integrated Microsystems. At the University of Illinois, he founded the VLSI Information Processing Systems (ViPS) Group, whose charter is to develop technologies for the design of low-power, high-performance and reliable integrated circuit implementations of broadband communications and digital signal processing systems. He has published more than 90 journal articles, book chapters, and conference publications in this area and holds three U.S. patents. He is also a co-author of the research monograph *Pipelined Adaptive Digital Filters* (Norwell, MA: Kluwer, 1994).

Dr. Shanbhag received the 2001 IEEE TRANSACTIONS ON VLSI SYSTEMS Best Paper Award, the 1999 IEEE Leon K. Kirchmayer Best Paper Award, the 1999 Xerox Faculty Award, the National Science Foundation CAREER Award in 1996, and the 1994 Darlington Best Paper Award from the IEEE Circuits and Systems Society. Since July 1997, he has been a Distinguished Lecturer for the IEEE Circuits and Systems Society. From 1997 to 1999, he served as an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS: Part II. He is currently an Associate Editor for IEEE TRANSACTIONS ON VLSI SYSTEMS, and has served on the technical program committees of various international conferences.



K. Soumyanath (M'93) received the B.E. degree in electronics and communication engineering from the Regional Engineering College, Tiruchirappalli, India, in 1979, the M.S. degree in electronics from the Indian Institute of Science, Bangalore, in 1985, and the Ph.D. degree in computer science from the University of Nebraska at Lincoln in 1993.

He was a member of the faculty at Tufts University, Medford, MA, until 1995, where he served as the Director of the ARPA supported program in Mixed Signal IC design for the Department of Defense. Since 1996, he has been with Intel Corporation, Hillsboro, OR, where he is a Principal Engineer and currently leads the CMOS communications circuits program for Intel's Internet Systems and Circuits Research Labs. His previous responsibilities included leading several high-speed digital research projects for the Circuits Research Laboratory. In 1998, he served as the Chair of the Design Sciences task force for the Semiconductor Research Corporation and currently serves on the Technical Program committee for ICCD. He has published over 15 papers in VLSI and holds six patents.



Shekhar Y. Borkar (M'97) was born in Mumbai, India. He received the M.S. degree in electrical engineering from the University of Notre Dame, Notre Dame, IN, in 1981 and the B.S. and M.S. degrees in physics from the University of Bombay, India, in 1979.

He joined Intel Corporation, Hillsboro, OR, in 1981. He is an Intel Fellow and Director of Circuit Research at Intel's Intel Architecture Group. He worked on the design of the 8051 family of microcontrollers, iWarp multi-computer and high-speed signaling technology for Intel supercomputers. He is an adjunct member of the faculty of the Oregon Graduate Institute. He has published ten articles and holds eleven patents.