ESE 5700 Project 1, Part 1: The 6T SRAM Bitcell

Panos Dimtsoudis (panosdim) Dhruva Jindal (djindal)

10/13/26

AI usage disclaimer

CLAUDE AI was utilized in this report to generate tables, graph placeholders, photo subtext and took a major part in report organization. The MATLAB code for generating graphs from **our** csv data was done using Claude as well.

Team and contributions

Task	panosdim	djindal
1.1 Hand analysis (β, γ)	X	
1.1 Schematic β sweep		X
1.2 Layout, DRC, LVS	X	
1.3 Parasitic extraction		X
1.4 Stability characterization		X
1.5 Functional check	X	
Report writing	X	X

Both members have reviewed and understand the whole design and report.

Summary of results

All values post-layout at $V_{DD} = 0.8\,\mathrm{V}$ unless noted, to two significant figures.

Quantity	Value	Section
Fins PG / PD / PU ($L = 18\,\text{nm}$)	2 / 2 / 2	1.1
β, γ	1, 1	1.1
Cell area (μm^2, F^2)	0.53, 5.27×10^3	1.2
Density factor $D = A_{cell}/\sum W_i L_i$	79	1.2
C_{BL} per cell (fF)	0.27	1.3
C_{WL} per cell (fF)	0.78	1.3
SNM_{hold} (V)	0.29	1.4a
SNM_{read} (V)	0.13	1.4b
Read "0" node: hand / simulated (V)	0.16 / 0.14	1.1, 1.4b
$\text{WM}_{BL} = V_{trip}$ (V)	0.27	1.4c
I_{cell} (μA)	22	1.4d
Reads and writes at 0.5 V?	yes (SNM 0.081 V, WM 0.16 V)	1.4e

1 Sizing (1.1)

1.1 Cell schematic and device sizes

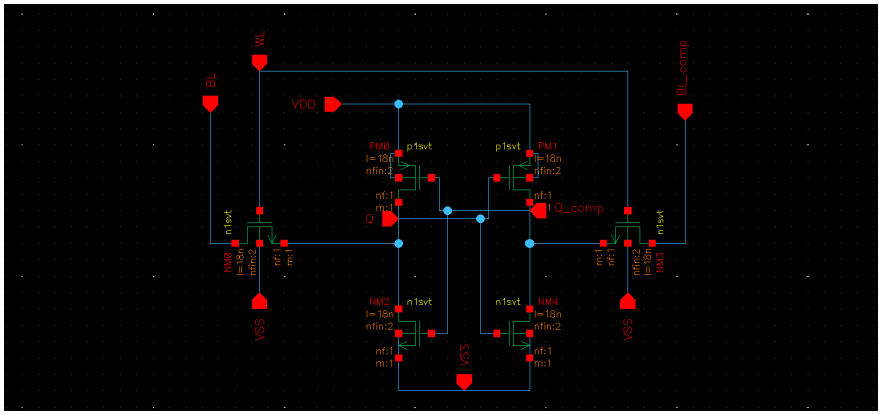


Figure 1: 6T cell schematic (`proj1_lib/sram6t`). All devices are standard-V_t FinFETs from `cds_ff_mpt` (`n1svt`, `p1svt`), $L = 18\,\text{nm}$, one finger. Final fin counts: PG = 2, PD = 2, PU = 2. All NMOS bodies are tied to VSS and both PMOS bodies to VDD.

Looking at the FinFET kit, we can see that each fin adds a fixed width to our devices, and L is fixed at 18 nm for the default poly pitch. Every device in the cell has the same L and one finger, so the sizing ratios reduce to fin-count ratios:

$$\beta = \frac{(W/L)_{PD}}{(W/L)_{PG}} = \frac{n_{fin,PD}}{n_{fin,PG}}, \qquad \gamma = \frac{(W/L)_{PU}}{(W/L)_{PG}} = \frac{n_{fin,PU}}{n_{fin,PG}}. \tag{1}$$

We first explored the ratios on a cell with a one-fin pull-up and a two-fin pass gate ($\gamma = 0.5$), sweeping the pull-down over 2, 3 and 4 fins ($\beta = 1, 1.5, 2$) which can be seen in the β sweep below. During layout, DRC showed that the kit does not allow one-fin devices. Specifically, we encountered rule OXIDE.W.1 many times, which requires an active height of at least 0.062 µm in steps of 0.048 µm. This means we need two fins or more. With two fins as the floor for every device, we chose the smallest cell the kit allows, PU = PG = PD = 2 fins, which gives $\beta = 1$ and $\gamma = 1$.

The rest of this section shows, derived by both hand analysis and simulation, that this cell reads and writes correctly. Before making our approach of a 2/2/2 cell concrete, however, we had laid out a 2/4/6 cell. We will compare our results throughout the report with the $\beta = 1.5$, $\gamma = 0.5$ alternative (PU = 2, PG = 4, PD = 6 fins). This approach is about 5.9× larger and we concluded that the 2/2/2 cell has the better read margin and the smaller write margin of the two.

		Fins per finger			
Device	Cell	β sweep	Final cell	Fingers	L (nm)
PG1, PG2 (pass gate)	`n1svt`	2	2	1	18
PD1, PD2 (pull-down)	`n1svt`	2 / 3 / 4	2	1	18
PU1, PU2 (pull-up)	`p1svt`	1	2	1	18
Sweep: $\beta = 1,\ 1.5,\ 2$ and $\gamma = 1/2$.			Final: $\beta = 2/2 = 1$ and $\gamma = 2/2 = 1$.		

Table 1: Device sizes. The β sweep uses a one-fin pull-up, which DRC does not allow in layout; the final cell puts every device at the two-fin minimum. The design point is discussed at the end of this section.

1.2 Device parameters used in the hand analysis

After reading the PDK documentation, we found no threshold voltage or transconductance, and the BSIM-CMG FinFET model has no single parameter for either. We therefore simulated one single-fin device of each type (test bench `tb_dev`, Fig. 2). Both devices are held in saturation with $|V_{DS}| = V_{DD} = 0.8\,\mathrm{V}$ while the gate drive v is swept from 0 to 0.8 V (NMOS $V_{GS} = v$; PMOS $V_{SG} = v$), at the `tt` corner.

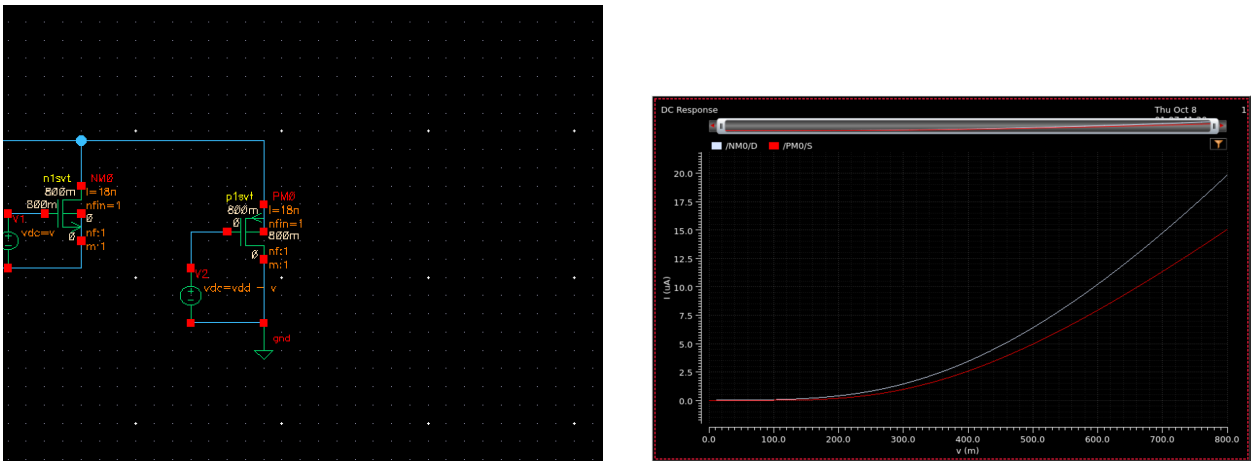


Figure 2: Left: device test bench. Right: drain current of a single-fin `n1svt` and `p1svt` versus gate drive, both in saturation.

V_t is the model's own `vth` from the DC operating point at $|V_{GS}| = |V_{DS}| = 0.8\,\mathrm{V}$ (both devices in region 2, saturation). The transconductance parameter per fin is fitted to the square law at the

same bias, $k' = 2I_D/(|V_{GS}| - |V_t|)^2$:

$$k'_n = \frac{2(19.88\,\mu\text{A})}{(0.8 - 0.257)^2} = \frac{39.76\,\mu\text{A}}{(0.543\,\text{V})^2} \approx 135\,\mu\text{A/V}^2, \tag{2}$$

$$k'_p = \frac{2(15.07\,\mu\text{A})}{(0.8 - 0.203)^2} = \frac{30.15\,\mu\text{A}}{(0.597\,\text{V})^2} \approx 84.6\,\mu\text{A/V}^2, \tag{3}$$

$$\frac{k'_p}{k'_n} = \frac{I_p}{I_n}\left(\frac{V_{DD} - V_{tn}}{V_{DD} - |V_{tp}|}\right)^2 = 0.758 \times 0.827 \approx 0.63. \tag{4}$$

Parameter (1 fin, $L = 18\,\text{nm}$, `tt`)	NMOS	PMOS
$\lvert V_t\rvert$ (V)	0.26	0.20
$\lvert I_D\rvert$ at $\lvert V_{GS}\rvert = \lvert V_{DS}\rvert = 0.8\,\text{V}$ (µA)	20	15
Overdrive $V_{DD} - \lvert V_t\rvert$ (V)	0.54	0.60
k' per fin (µA/V^2)	130	85
k'_p/k'_n	0.63	

Table 2: Extracted device parameters (two significant figures; the calculations use the unrounded values).

The raw current ratio $I_p/I_n = 0.76$ is higher than k'_p/k'_n because the PMOS has the smaller $|V_t|$ and therefore more overdrive at the same gate drive. The ratio of 0.63 plays the role of μ_p/μ_n in the lecture analysis. It is higher than the ≈ 0.5 typical of planar devices, which makes the pull-up relatively stronger and writing harder. At an 18 nm gate length the square law is only a first-order model (velocity saturation flattens the measured curves), and k' is fitted at a single bias. The estimates below are therefore approximate; Section 1.4b compares them with simulation.

1.3 Read stability

First assume $Q = 0$ is stored and both bitlines are at V_{DD} with WL $= V_{DD}$. PG1 has $V_{GS} = V_{DS} = V_{DD} - V_Q$, so it is always saturated. PD1 has $V_{GS} = V_{DD}$ and a small $V_{DS} = V_Q$, so it is in the linear region. Then neglecting body effect on PG1, the currents balance as

$$\frac{k_{PG}}{2}(V_{DD} - V_Q - V_{tn})^2 = k_{PD}\left[(V_{DD} - V_{tn})V_Q - \frac{V_Q^2}{2}\right]. \tag{5}$$

With $V_{ov} = V_{DD} - V_{tn} = 0.543\,\text{V}$ and $\beta = k_{PD}/k_{PG}$, this becomes $V_Q^2 - 2V_{ov}V_Q + V_{ov}^2/(1+\beta) = 0$, whose physical root is

$$V_Q = V_{ov}\left(1 - \sqrt{\frac{\beta}{1+\beta}}\right). \tag{6}$$

The node must stay below the switching threshold V_M of the other inverter (PU2/PD2). From the square law, with $r = \sqrt{k_{PU}/k_{PD}} = \sqrt{(k'_p/k'_n)\,n_{fin,PU}/n_{fin,PD}}$,

$$V_M = \frac{V_{tn} + r\,(V_{DD} - |V_{tp}|)}{1 + r}. \tag{7}$$

β	$n_{fin,PD}/n_{fin,PU}$	$\sqrt{\beta/(1+\beta)}$	V_Q (V)	r	V_M (V)
1	2	0.707	0.16	0.56	0.38
1.5	3	0.775	0.12	0.46	0.36
2	4	0.816	0.10	0.40	0.35
1 (final)	1	0.707	0.16	0.79	0.41

Table 3: Hand estimate of the read-disturb voltage on the "0" node and the switching point of the opposite inverter ($V_{DD} = 0.8\,\mathrm{V}$, $V_{tn} = 0.257\,\mathrm{V}$). The first three rows are the sweep cell (PU = 1 fin); the last row is the final cell (PU = PD = 2 fins).

For example, at $\beta = 1$: $V_Q = 0.543 \times (1 - 0.707) = 0.159\,\mathrm{V}$. In the final cell the pull-up has as many fins as the pull-down, so $r = \sqrt{0.63} = 0.79$ and $V_M = (0.257 + 0.79 \times 0.597)/1.79 = 0.41\,\mathrm{V}$, about 0.03 V higher than in the sweep cell at the same β. The stronger pull-up therefore raises the switching point and recovers part of the read margin that $\beta = 1$ gives up. All three points satisfy both $V_Q < V_M$ (by at least 0.22 V) and the stricter $V_Q < V_{tn} = 0.26\,\mathrm{V}$. So, by hand analysis, even $\beta = 1$ does not flip the cell during a read. A larger β lowers V_Q further, which widens the read SNM, at the cost of area and bitline capacitance. Neglecting body effect makes the estimate slightly pessimistic: as V_Q rises, PG1's threshold increases and its current drops, so the true V_Q should be a little lower.

1.4 Writability

To write a 0 into a cell storing $Q = 1$, BL is driven to 0 with WL = V_{DD}. At the moment of the write $\bar{Q} = 0$, so PU1 is fully on ($V_{SG} = V_{DD}$) and in saturation. PG1 has $V_{GS} = V_{DD}$ and a small $V_{DS} = V_Q$, so it is in the linear region. The currents balance as

$$k_{PG}\left[(V_{DD} - V_{tn})V_Q - \frac{V_Q^2}{2}\right] = \frac{k_{PU}}{2}\left(V_{DD} - |V_{tp}|\right)^2, \tag{8}$$

which, with $k_{PU}/k_{PG} = (k'_p/k'_n)\,\gamma$, gives

$$V_Q = V_{ovn} - \sqrt{V_{ovn}^2 - \frac{k'_p}{k'_n}\,\gamma\,V_{ovp}^2}. \tag{9}$$

The write succeeds if V_Q falls below a target V_{tgt} (the opposite inverter's switching point, or V_{tn} for margin). Solving for γ:

$$\gamma_{max} = \frac{2V_{ovn}V_{tgt} - V_{tgt}^2}{(k'_p/k'_n)\,V_{ovp}^2}, \qquad (k'_p/k'_n)\,V_{ovp}^2 = 0.63 \times (0.597)^2 = 0.224\,\mathrm{V}^2. \tag{10}$$

Target	V_{tgt} (V)	γ_{max}
V_{tn} (strict)	0.26	0.95
V_M ($\beta = 1$, sweep cell)	0.38	1.2
V_M (final cell, PU = PD)	0.41	1.23

Table 4: Largest pull-up ratio for which the pass gate can pull the "1" node below the target.

The final cell has PG = PU = 2 fins, so $\gamma = 1$. The "1" node is then pulled to $V_Q = 0.543 - \sqrt{0.295 - 0.224 \times 1} = 0.28\,\mathrm{V}$. That is 0.13 V below the opposite inverter's switching point ($V_M = 0.41\,\mathrm{V}$ for PU = PD = 2 fins, $\gamma_{max} = 1.23$), so the write succeeds, but it fails the strict V_{tn} target ($\gamma_{max} = 0.95 < 1$). At $V_Q = 0.28\,\mathrm{V}$, PU1 has $V_{SD} = 0.52\,\mathrm{V}$, below $V_{SG} - |V_{tp}| = 0.60\,\mathrm{V}$, so it has left saturation. Its true current is then below the saturation value assumed above, and the true V_Q is lower: the estimate is conservative. For comparison, $\gamma = 0.5$ (PG = 4 fins with the two-fin minimum pull-up) pulls the node to 0.12 V, but it doubles the pass gates and, to keep $\beta \geq 1$, the pull-downs as well. Because FinFET P and N devices are close in strength ($k'_p/k'_n = 0.63$), a pass gate the size of the pull-up gives the smallest write margin of the sizings considered. We decided to accept that margin in exchange for a much smaller minimum area. Post-layout simulations confirm and showcase our decisionin 1.4c. We also note that β does not enter the write condition, so the β sweep does not affect writability.

1.5 Schematic-level β sweep

β	$\mathrm{SNM_{read}}$ (V)	"0" node (V)
1.0	0.10	0.13
1.5	0.12	0.098
2.0	0.13	0.078

Table 5: Read SNM at the schematic level, $V_{DD} = 0.8\,\mathrm{V}$, half-size cell (PG = 2 fins, PU = 1 fin, PD = 2, 3, 4 fins). The "0" node column is the worse of the two stable points on the read butterfly.

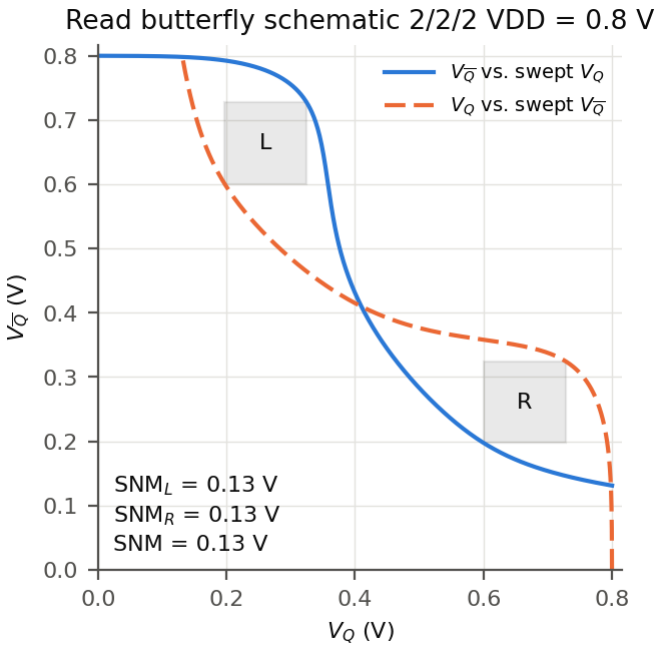


Figure 3: Schematic-level read butterfly of the final cell (PU = PG = PD = 2 fins, WL = BL = $\overline{\text{BL}}$ = V_{DD} = 0.8 V). The schematic cell is symmetric, so SNM$_L$ = SNM$_R$ = 0.127 V. The sweep butterflies are in `fig/butterfly_read_beta1p0.png`, `fig/butterfly_read_beta1p5.png` and `fig/butterfly_read_beta2p0.png`.

The sweep uses test bench `tb_snm`. It contains two copies of the cell, each with one storage node forced by an ideal source v_{in}, swept from 0 to V_{DD}, which breaks the feedback loop. The other storage node is recorded, giving the two halves of the butterfly in a single DC sweep. SNM is the side of the largest square inscribed in each lobe, extracted with the 45° rotation method (Appendix A). It agrees with an independent brute-force search to within 0.5 mV.

			"0" node (V)		V_M (V)	
β	PD fins (sweep)	SNM$_{read}$ (V)	hand	sim	hand	sim
1	2	0.10	0.16	0.13	0.38	0.39
1.5	3	0.12	0.12	0.098	0.36	0.37
2	4	0.13	0.10	0.078	0.35	0.36
1 (final, PU = 2)	2	0.13	0.16	0.13	0.41	0.41

Table 6: Hand analysis against schematic simulation. "sim" values are the stable and metastable intersections of the simulated read butterfly. The first three rows are the sweep cell (PU = 1 fin); the last row is the final 2/2/2 cell.

The hand estimate of the "0"-node voltage is approximately 25% higher than simulation at every β. The hand estimate of V_M is within 0.02 V of the simulated switching point, so the first-order model captures the trend well, though not the exact values.

Design point. We choose the minimum-area cell, **PU = PG = PD = 2** fins ($\beta = 1$, $\gamma = 1$). In the sweep, read SNM is 0.10 V at $\beta = 1$, against 0.12 V at $\beta = 1.5$ and 0.13 V at $\beta = 2$, all with a one-fin pull-up. With the two-fin minimum, any $\beta > 1$ needs a pull-down of at least 3 fins, and keeping $\gamma = 0.5$ as well needs PG = 4 and PD = 6 fins. That cell, which we laid out first, is 3.08 µm^2 against 0.53 µm^2 for 2/2/2. The 2/2/2 cell is the smallest cell the kit allows: every device is at the two-fin minimum and every NMOS active has the same height, so the four NMOS devices abut on one straight strip with shared diffusions. The two-fin pull-up also recovers the read margin that $\beta = 1$ gives up, so the cost of the minimum area falls on the write margin and the read current instead. (next paragraph)

Final cell at schematic level. The final cell differs from the $\beta = 1$ sweep point only in its pull-ups (2 fins instead of 1). A schematic-level simulation of the final cell gives $\mathrm{SNM}_{read} = 0.127$ V and a "0" node of 0.131 V, against 0.10 V and 0.13 V for the $\beta = 1$ sweep point. The "0" node is unchanged, as expected, since it is set by PG and PD alone. The stronger pull-up raises the opposite inverter's switching point from 0.39 V to 0.41 V, exactly the hand estimate, and that widens both lobes by more than a quarter. The final cell's read SNM is therefore slightly higher than that of the $\beta = 1.5$ sweep point (0.12 V) and of the schematic 2/4/6 cell (0.118 V). The cell layout (1.2) and all post-layout results (1.3–1.5) use these final fin counts.

2 Layout (1.2)

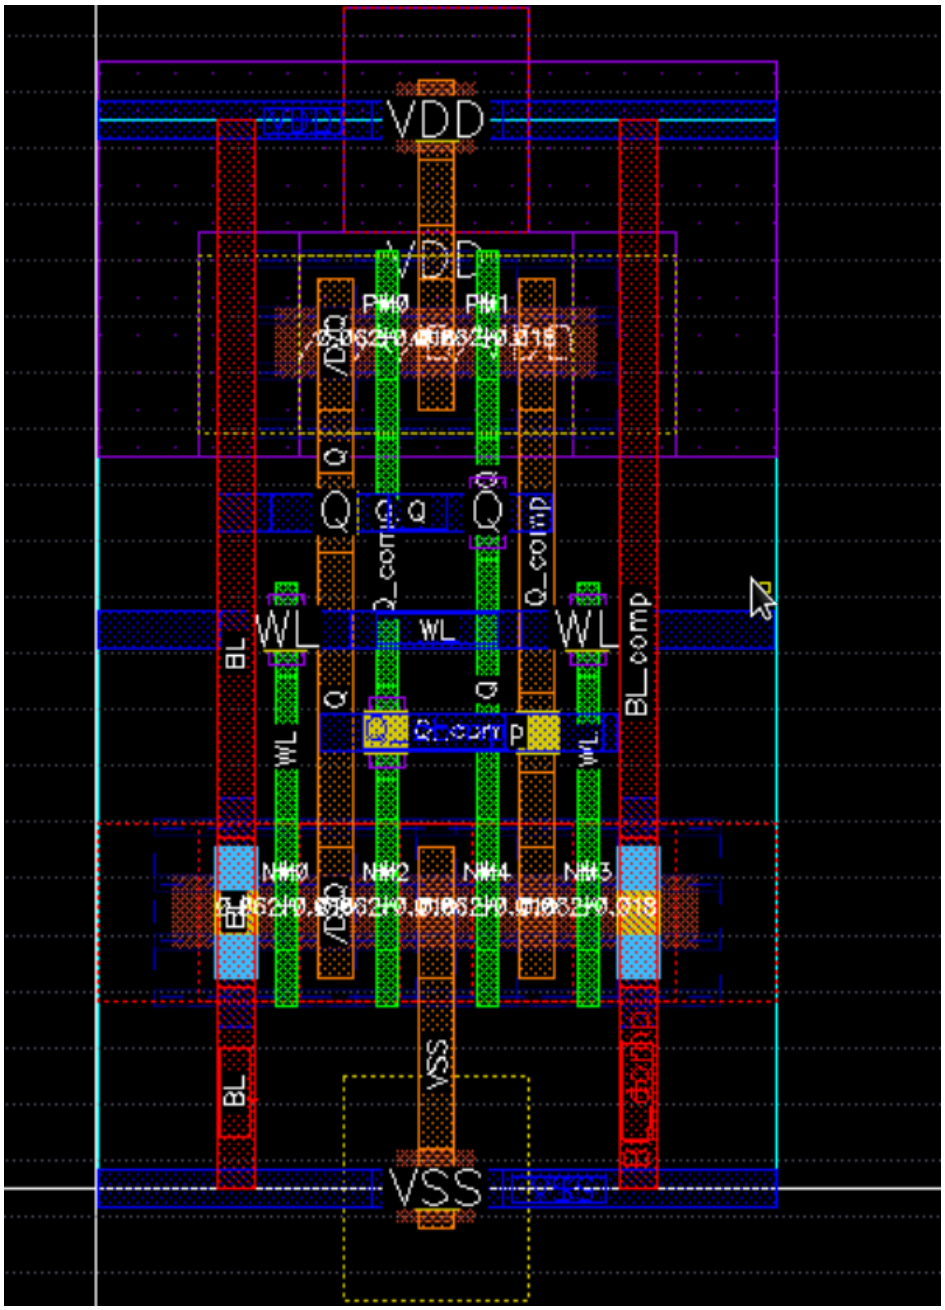


Figure 4: Bitcell layout (`proj1_lib/sram6t_v2/layout`), DRC and LVS clean. The cyan rectangle is the cell (PR) boundary, 0.578 µm × 0.912 µm.

Floorplan. The cell has two device rows. All four NMOS devices sit on one straight two-fin active strip near the bottom edge, in the order PG1, PD1, PD2, PG2, with one shared diffusion between each pair of neighbours: PG1 and PD1 share Q, PD1 and PD2 share the VSS source, and PD2 and PG2 share $\overline{Q}$. The two pull-ups sit on a second two-fin strip near the top edge, inside one NWell, and share their VDD source, which reaches the VDD rail through a short vertical

strap. All dummy polys were removed, so neighbouring gates are one gate pitch (86 nm) apart. Each inverter's gate poly runs straight from its pull-up to its pull-down, and each storage node is a vertical local-interconnect strap from the pull-up drain down to the shared PG/PD diffusion. The Q and $\overline{Q}$ cross-couple runs on two short horizontal M1 tracks between the rows, above and below the wordline, each landing on the opposite inverter's gate through a poly contact. The pass-gate polys extend up to the wordline, which contacts them through poly contacts. The shared PD source reaches the VSS rail through a short, straight strap, which keeps the pull-down source resistance low. The n-well tap and the substrate tap sit at the centre of the top and bottom edges, straddling the cell boundary.

Lines and tiling. VDD and VSS are M1 rails along the top and bottom edges, and WL is an M1 line across the full cell width between the device rows; these three run edge to edge, so they join into continuous lines when cells abut in a row. BL and $\overline{\text{BL}}$ are vertical M2 lines running the full cell height just inside the left and right edges, so they join into continuous bitlines down a column, and each reaches its pass gate's outer diffusion through an M2–M1 via, an M1 pad and an M1–LiAct contact. Rows are intended to mirror about the shared supply rails, so the taps on the top and bottom edges are shared with the mirrored neighbour, and each row holds one VDD and one VSS rail between two rows of cells.

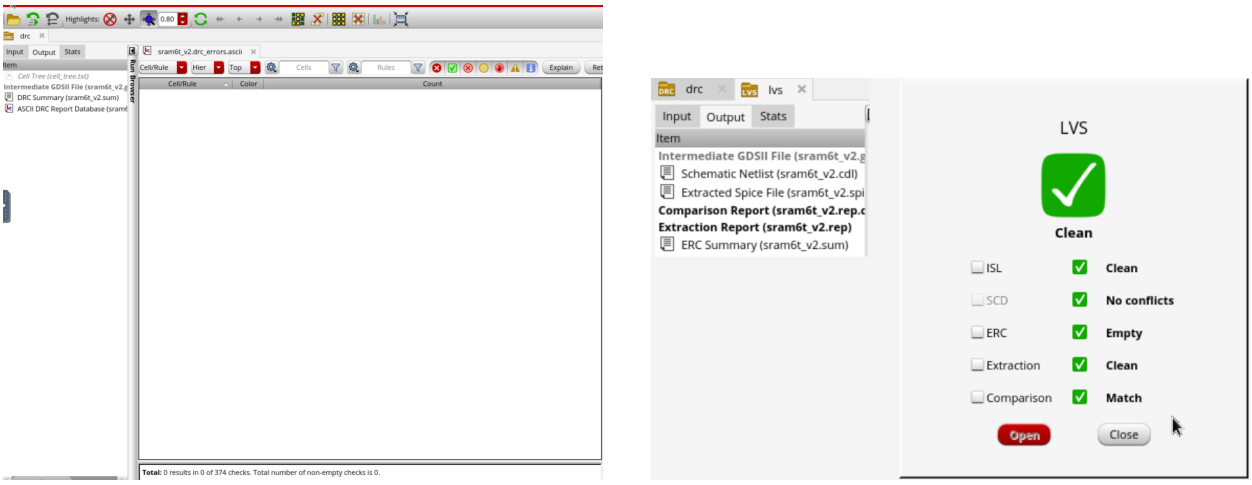


Figure 5: DRC (left) and LVS (right) summaries for the cell.

Quantity	Value
Cell width × height (µm)	0.578×0.912
Cell area (µm^2)	0.527
Cell area (F^2, $F = 10\,\text{nm}$, $F^2 = 1 \times 10^{-4}\,\mu\text{m}^2$)	5.27×10^3
$\sum_{i=1}^{6} W_i L_i$ (µm^2)	0.0067
Density factor $D = A_{cell} / \sum W_i L_i$	79

Table 7: Cell area, from the PR boundary (0.5271 µm^2). Every device is 2 fins, so $W_i = 0.062\,\mu\text{m}$ for all six and $\sum W_i L_i = 6 \times 0.062 \times 0.018 = 0.0067\,\mu\text{m}^2$. D is used for area estimates in Parts 2 and 3.

The cell is 0.578 µm wide, and 0.912 µm tall. Its width is set by the four NMOS gates on the

shared strip plus the bitline contacts and the M2 bitlines at the two side edges. Its height is set by the two device rows, the fin-grid spacing the OXIDE and CUTACTIVE rules force between separate actives, and the three M1 tracks (Q, WL, $\overline{Q}$) between the rows. Against the 2/4/6 layout ($2.650 \times 1.164\,\mu\text{m} = 3.08\,\mu\text{m}^2$) the cell is 4.6 times narrower and 1.3 times shorter, 5.9 times smaller overall, and D falls from 181 to 79. The width saving comes almost entirely from shared diffusions and removed dummy polys: in the 2/4/6 layout each device carried an edge poly and two dummy polys on each side. D remains large because a FinFET's drawn width is only the fin height (0.062 µm for two fins), so $\sum W_i L_i$ is tiny compared with the wiring and spacing the cell needs.

3 Parasitic extraction (1.3)

Extraction flow. We reran LVS using `cds_ff_mpt_pvs` with Quantus input data enabled, then extracted the cell in Quantus through the Pegasus interface (version 23.1). We used `cds_ff_mpt_qrc` at the `typical` corner and 25 °C, with coupled RC extraction and VSS as the reference. The resulting `av_extracted` view contains the six transistors, 244 parasitic capacitors, and 90 parasitic resistors across BL, $\overline{\text{BL}}$, WL, Q, $\overline{Q}$, VDD, and VSS. The earlier 2/4/6 cell had 538 capacitors and 168 resistors. After removing the dummy polys, the number of floating nets dropped from 42 to 1, and Quantus reported no warnings. Coupling capacitors below 0.01 fF were converted to grounded capacitances rather than removed. We selected `av_extracted` through a `config` view in the Hierarchy Editor for all post-layout simulations in Sections 1.4 and 1.5.

Measurement. To include both the extracted wiring capacitance and the transistor capacitance from the device models, we measured the BL and WL loads using AC simulation in `tb_cap` (Fig. 6). The test bench contains two extracted cells. For the BL measurement, we applied 0.8 V DC with a 1 V AC magnitude to BL, held WL at 0, and held $\overline{\text{BL}}$ at 0.8 V. This represents an unselected cell with precharged bitlines. For the WL measurement, we applied 0 V DC with a 1 V AC magnitude to WL and held both bitlines at 0.8 V. The other voltage sources had zero AC magnitude, making their connected pins AC grounds. Node sets were used to select the $Q = 0$ operating point in both cells. We calculated the capacitance from the AC current I through the driving source:

$$C = \frac{|\text{Im}\, I|}{2\pi f \cdot 1\,\text{V}}. \tag{11}$$

Over 1 kHz to 1 GHz the current magnitude rises in proportion to frequency (Fig. 6), so both loads are purely capacitive in this range; the values below are read at 1 MHz.

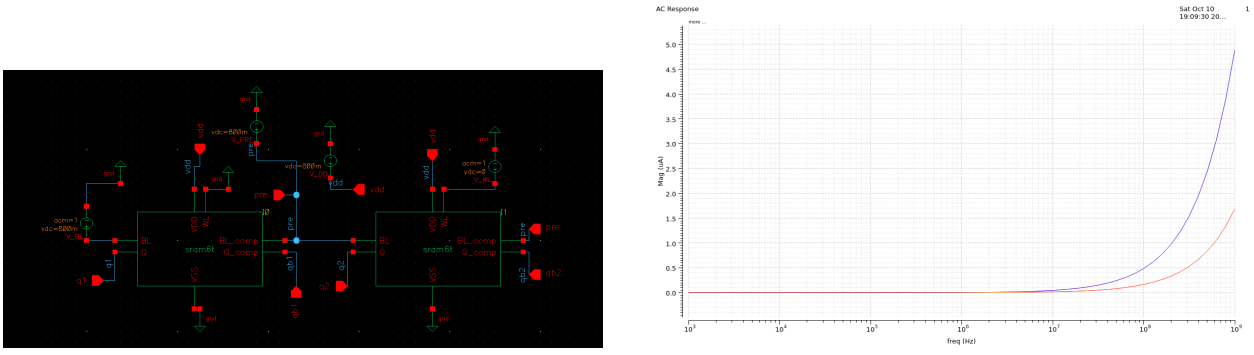


Figure 6: Left: capacitance test bench. Right: AC current drawn by the BL and WL drive sources (1 V AC). Both rise in proportion to f, as for a pure capacitance.

Node	C_{cell} (fF)	What it includes
Bitline (BL; $\overline{\text{BL}}$ by symmetry)	0.27	M2 bitline, M1 pad and vias, PG junction and overlap
Wordline (WL)	0.78	M1 wordline across the cell, two PG gates, coupling

Table 8: Extracted per-cell node capacitances, post-layout, `typical` corner, at 1 MHz (268.0 aF and 779.4 aF unrounded). These are the per-cell loads for Part 2.

Both loads are about 43–45% lower than in the 2/4/6 cell, which gave 0.47 fF and 1.41 fF. The two-fin pass gates halve the pass-gate junction on each bitline and the gate load the two pass gates put on the wordline, and the smaller cell (0.578 µm wide instead of 2.65 µm) shortens the wordline wire by a factor of 4.6. The wordline load is still about three times the bitline load, because the wordline carries both pass-gate gates and runs the full cell width on M1. For Part 2, a bitline in a column of 16 cells is loaded by about $16 \times 0.27 \approx 4.3$ fF before the precharge, write and sense circuits are added, and a wordline across 4 columns by about $4 \times 0.78 \approx 3.1$ fF before the driver. The 15 additional bitline loads in the functional check (1.5) add $15 \times 0.268 \approx 4.0$ fF.

4 Stability characterization (1.4)

All results in this section use the extracted view at $V_{DD} = 0.8$ V unless noted. SNM and V_{trip} are computed from exported DC sweeps with `p1tools.py` (Appendix A). Test bench schematics are in Appendix B.

The butterfly test bench (`tb_snm`, Appendix B) holds two extracted cells, selected through a `config` view. In each, one storage node is forced by an ideal source v_{in}, swept from 0 to V_{DD} in 1 mV steps, which breaks the feedback loop, and the other storage node is recorded. The first cell gives the $\overline{Q}$-side inverter and the second the Q-side inverter, so the two halves of the butterfly come from the two physically different halves of the layout. BL and $\overline{\text{BL}}$ are tied to V_{DD}. WL is a separate source, set to 0 for hold and to V_{DD} for read. Because the extracted halves are not identical, SNM_L and SNM_R differ slightly, and we report $\text{SNM} = \min(\text{SNM}_L, \text{SNM}_R)$.

4.1 Hold SNM (1.4a)

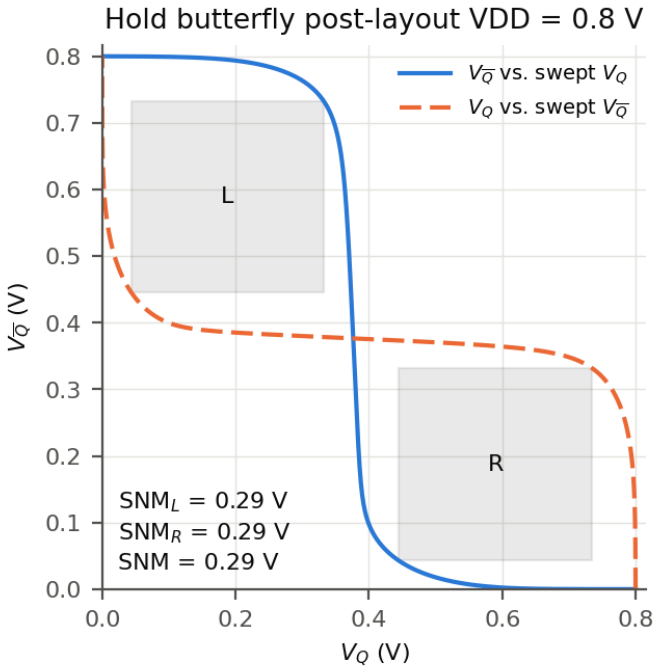


Figure 7: Post-layout hold butterfly, WL = 0, $V_{DD} = 0.8\,\mathrm{V}$. $\mathrm{SNM}_L = 0.288\,\mathrm{V}$, $\mathrm{SNM}_R = 0.289\,\mathrm{V}$, so $\mathrm{SNM}_{hold} = 0.29\,\mathrm{V}$ (36% of V_{DD}).

With the wordline off, the pass gates carry no current and each inverter swings fully, from 0.80 V to 0.1 mV. With equal pull-up and pull-down fin counts the inverters are close to balanced, and the transfer curves switch sharply at $V_M \approx 0.376\,\mathrm{V}$, near $V_{DD}/2$. Both lobes are large and almost equal, $\mathrm{SNM}_L = 0.288\,\mathrm{V}$ and $\mathrm{SNM}_R = 0.289\,\mathrm{V}$, giving $\mathrm{SNM}_{hold} = 0.288\,\mathrm{V}$, against 0.251 V for the 2/4/6 cell, whose weaker pull-up put its switching point lower. An independent brute-force search on a 0.7 mV raster gives 0.287 V and 0.289 V.

4.2 Read SNM (1.4b)

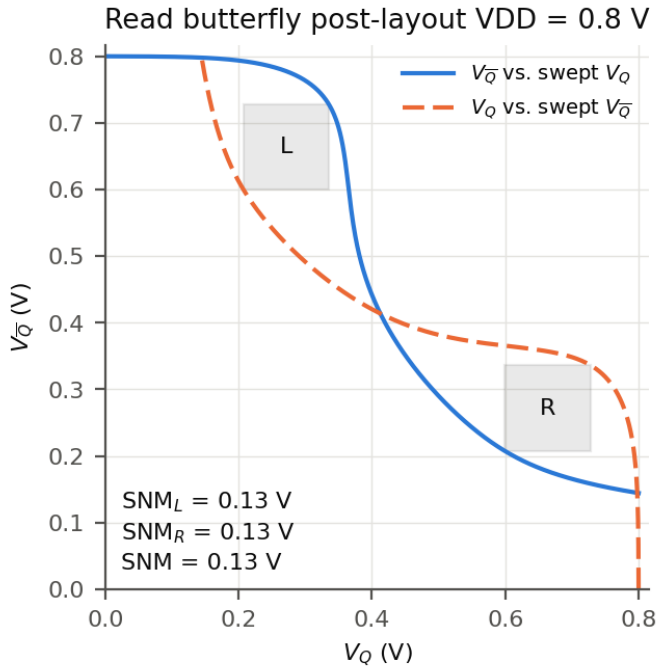


Figure 8: Post-layout read butterfly, WL = BL = $\overline{\text{BL}} = V_{DD} = 0.8\,\text{V}$ (ideal sources). $\text{SNM}_L = 0.1280\,\text{V}$, $\text{SNM}_R = 0.1286\,\text{V}$, so $\text{SNM}_{read} = 0.13\,\text{V}$ (16% of V_{DD}).

Turning the wordline on lets the pass gate pull the "0" node up against the pull-down. The low end of each transfer curve rises from 0.1 mV to 0.145 V, the switching point moves up to 0.413 V, and the curve becomes much less steep after it switches. Both lobes shrink, to $\text{SNM}_{read} = 0.128\,\text{V}$ (raster check 0.128 V), less than half the hold value, so read stability is still the smallest static margin of the cell. It is nevertheless 21 mV larger than the 2/4/6 cell's post-layout 0.107 V: the two-fin pull-up raises the switching point by about 40 mV, which more than offsets the higher "0" node that $\beta = 1$ brings.

The post-layout value equals the schematic-level 0.127 V of 1.1 within 1 mV. Layout parasitics raise the "0" node by 13 mV, from 0.131 V to 0.145 V, but they also raise the switching point slightly, and the SNM is unchanged. In the 2/4/6 cell the same comparison cost 11 mV of SNM and 27 mV to 32 mV on the "0" node. The difference comes from the pull-down source path: the read current is about half as large here (21.8 µA, Section 1.4d), and the shared PD source reaches the VSS rail through a short, straight strap, so the IR drop across the source wiring is smaller (13 mV at 21.8 µA corresponds to about 0.6 kΩ). The two halves are equal to within 0.1 mV, because the layout is mirror-symmetric.

	"0" node during read (V)
Hand estimate (1.1, $\beta = 1$)	0.16
Schematic simulation, final cell	0.131
Post-layout simulation, $Q = 0$ / $Q = 1$	0.145 / 0.145

Table 9: Read-disturb voltage on the "0" node at $V_{DD} = 0.8\,\mathrm{V}$. Simulated values are the stable intersections of the read butterfly, which are the closed-loop DC solutions of the cell.

The hand estimate, $V_Q = V_{ov}\left(1 - \sqrt{\beta/(1+\beta)}\right) = 0.16\,\mathrm{V}$, is 22% above the schematic value of 0.131 V, the same fraction as in the β sweep. The main reason is the body effect, which the hand analysis neglects: the pass gate's source is the "0" node, so as that node rises the pass gate's threshold rises and its current falls, which holds the node lower than the square-law balance predicts. The square law also overstates the pass gate's current at this short gate length, where velocity saturation limits it. Wiring resistance pushes the post-layout value back up by 13 mV, to 0.145 V, which is still 10% below the hand estimate: the hand analysis is conservative for this cell.

4.3 Write margin (1.4c)

As written, the specification stores $Q = 1$, holds BL at V_{DD} and sweeps $\overline{\mathrm{BL}}$ down. With $Q = 1$ stored, $\overline{Q}$ is already 0, so pulling $\overline{\mathrm{BL}}$ low only reinforces the stored state and the cell can never flip. We therefore swept the bitline on the side that stores the 1: test bench `tb_wm` (Appendix B) stores $Q = 1$ through a node set, holds $\mathrm{WL} = \overline{\mathrm{BL}} = V_{DD}$, and sweeps BL from V_{DD} down to 0 in 1 mV steps on the extracted cell. This is the same measurement the specification intends, since by symmetry it is equivalent to storing $Q = 0$ and sweeping $\overline{\mathrm{BL}}$.

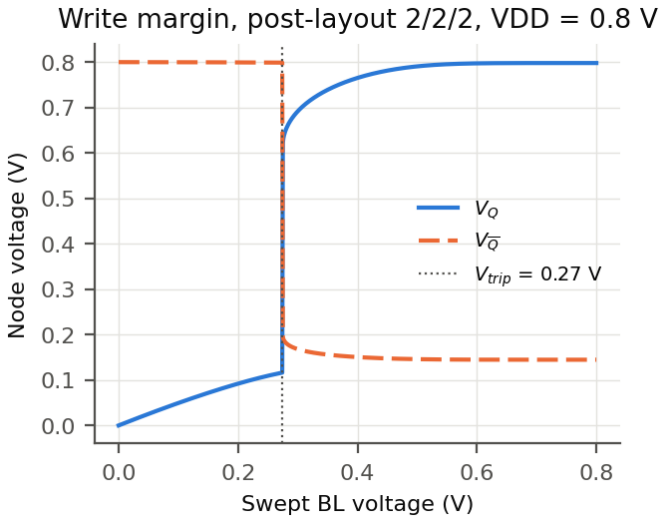


Figure 9: Post-layout write-margin sweep at $V_{DD} = 0.8\,\mathrm{V}$, $\mathrm{WL} = \overline{\mathrm{BL}} = V_{DD}$, BL swept downward. The cell flips at $V_{trip} = 0.275\,\mathrm{V}$, so $\mathrm{WM}_{BL} = 0.27\,\mathrm{V}$ (34% of V_{DD}).

Reading the sweep from the right: at $\mathrm{BL} = V_{DD}$ the cell holds $Q = 0.80\,\mathrm{V}$, with the read disturb of 0.145 V on $\overline{Q}$. As BL falls, the pass gate pulls Q down against the pull-up, to 0.62 V just before the trip. At $\mathrm{BL} = 0.275\,\mathrm{V}$ the falling Q crosses the switching threshold of the opposite inverter, the loop regenerates, and Q drops to 0.12 V while $\overline{Q}$ rises to 0.80 V. Below the trip point

Q follows BL down through the pass gate. The write margin is therefore $\mathrm{WM}_{BL} = V_{trip} = 0.27\,\mathrm{V}$: a write driver that pulls BL to ground writes with 0.27 V to spare.

This is 55 mV less than the 2/4/6 cell's 0.33 V, the expected cost of $\gamma = 1$: the pass gate now fights a pull-up as strong as itself rather than half as strong. The hand analysis of 1.1 predicted that a $\gamma = 1$ cell still writes, with the "1" node pulled to about 0.28 V at BL = 0, 0.13 V below the switching point. The simulation agrees: the cell flips well before BL reaches 0, so the conservative hand estimate understates the margin.

4.4 Cell read current (1.4d)

I_{cell} is the DC current a selected cell sinks from the bitline on its "0" side, with WL = BL = $\overline{\mathrm{BL}}$ = $V_{DD} = 0.8\,\mathrm{V}$ held by ideal sources. Test bench `tb_icell` (Appendix B) runs a DC operating point on the extracted cell, with a node set choosing the stored value, and records the current into the BL and $\overline{\mathrm{BL}}$ pins. We ran it once with $Q = 0$ and once with $Q = 1$, to measure both halves of the cell.

Stored value	BL current (µA)	$\overline{\mathrm{BL}}$ current (µA)
$Q = 0$	21.77	≈ 0
$Q = 1$	≈ 0	21.77

Table 10: Bitline currents of the selected extracted cell at $V_{DD} = 0.8\,\mathrm{V}$ (21.7723 µA and 21.7738 µA unrounded). Only the bitline facing the "0" node conducts.

The current flows from the precharged bitline through the pass gate and the pull-down into VSS; the bitline on the "1" side carries no current. The two halves agree to within 1.5 nA (0.007%), against a 10% difference in the 2/4/6 cell, because the 2/2/2 layout is mirror-symmetric about its centre line and both halves see the same wiring. We take $I_{cell} = 21.8\,\mu\mathrm{A}$. This is 57% of the 2/4/6 cell's 38 µA: the pass gate has half the fins, and with $\beta = 1$ the pull-down no longer holds the "0" node as low, so the pass gate's gate overdrive and drain-source voltage are both slightly lower. The bitline load fell by almost the same factor (Section 1.3), so a 16-cell bitline of $16 \times 0.268\,\mathrm{fF} = 4.29\,\mathrm{fF}$ discharges at $21.8\,\mu\mathrm{A}/4.29\,\mathrm{fF} \approx 5.1\,\mathrm{mV/ps}$, and a 100 mV split develops in about 20 ps, the same as for the larger cell, before the sense and precharge circuits of Part 2 add their loading.

4.5 Supply dependence (1.4e)

V_{DD} (V)	$\mathrm{SNM}_{\mathrm{read}}$ (V)	"0" node (V)	WM_{BL} (V)	Bistable in read	Flips in write
0.8	0.13	0.14	0.27	yes	yes
0.7	0.11	0.12	0.24	yes	yes
0.5	0.081	0.079	0.16	yes	yes

Table 11: Read SNM and write margin against supply, post-layout. "0" node is the higher of the two read-disturb levels. Generated by `p1tools.py table` from the exported sweeps.

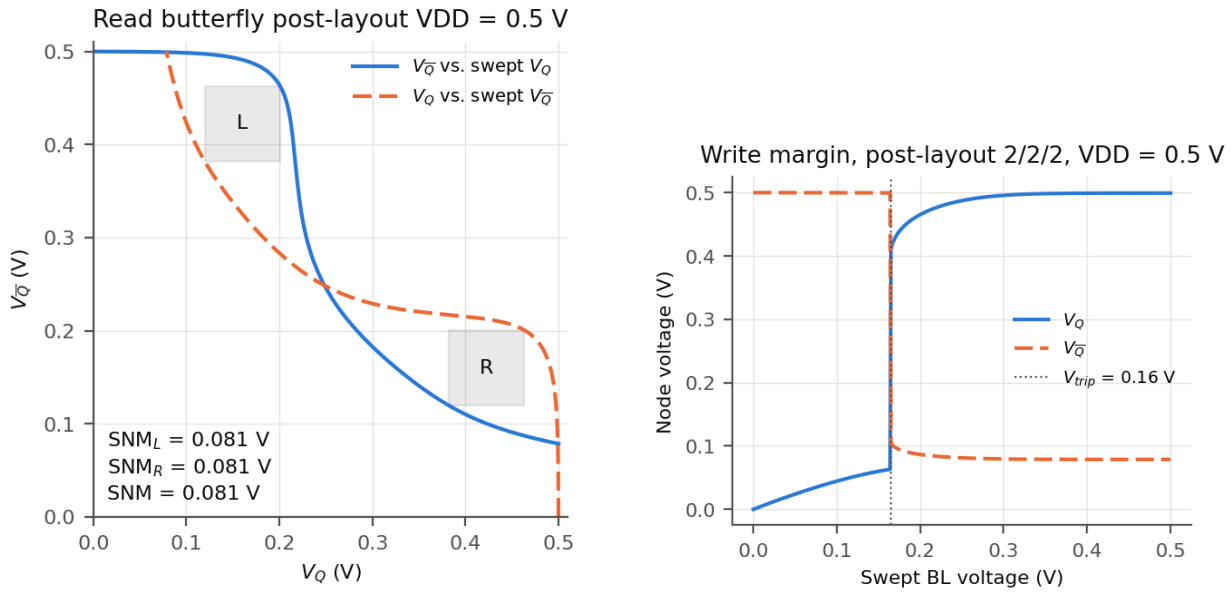


Figure 10: Post-layout behaviour at $V_{DD} = 0.5\,\mathrm{V}$. Left: read butterfly, $\mathrm{SNM}_{read} = 0.081\,\mathrm{V}$. Right: write sweep, $V_{trip} = 0.165\,\mathrm{V}$. The 0.7 V plots are in `fig/butterfly_read_0p7_222.png` and `fig/write_margin_0p7_222.png`.

Both margins shrink with the supply but stay nearly constant as a fraction of it. Read SNM goes from 0.128 V at 0.8 V to 0.114 V at 0.7 V and 0.081 V at 0.5 V, which is 16% of V_{DD} at every supply. The write trip point goes from 0.275 V to 0.236 V and 0.165 V, which is 34%, 34% and 33% of V_{DD}. At every supply the read butterfly has three intersections and two open lobes, and the write sweep flips the cell. **The cell therefore still reads and writes at** 0.5 V, with margins of 81 mV (read) and 165 mV (write). This is a nominal, `tt`-corner result: device mismatch and slow corners, which this analysis does not include, would erode the read margin first.

5 Functional check (1.5)

Test bench `tb_tran` (Appendix B) simulates one extracted cell, selected through a `config` view, as one cell of a 16-cell column. Each bitline carries a lumped capacitor of $15 \times C_{BL} = 15 \times 0.268\,\mathrm{fF} = 4.02\,\mathrm{fF}$, the bitline load of the other 15 cells from Section 1.3, on top of the cell's own extracted load. All control signals are ideal pulse sources with 20 ps edges at $V_{DD} = 0.8\,\mathrm{V}$. WL is driven directly. Each bitline has a precharge PMOS (`p1svt`, 2 fins) to VDD, gated by an active-low precharge pulse, and a write-driver NMOS (`n1svt`, 2 fins) to ground, gated by its own write pulse. The cell starts from an initial condition $Q = 1$, so the first write must actually flip it. The bitlines are precharged whenever no operation is in progress.

Phase	Time window (ns)	WL high (ns)	Precharge	BL / $\overline{\text{BL}}$ drive
Write 0	1–2	1.2–1.8	off	BL pulled to 0, $\overline{\text{BL}}$ floating
Read	3–4	3.2–3.8	off	both floating (precharged)
Write 1	5–6	5.2–5.8	off	$\overline{\text{BL}}$ pulled to 0, BL floating
Read	7–8	7.2–7.8	off	both floating (precharged)

Table 12: Stimulus for the functional check. Between the phases (0–1, 2–3, 4–5 and 6–7 ns) WL is low and both bitlines are precharged to V_{DD}; the write pulse turns on with the precharge off, 0.2 ns before WL rises.

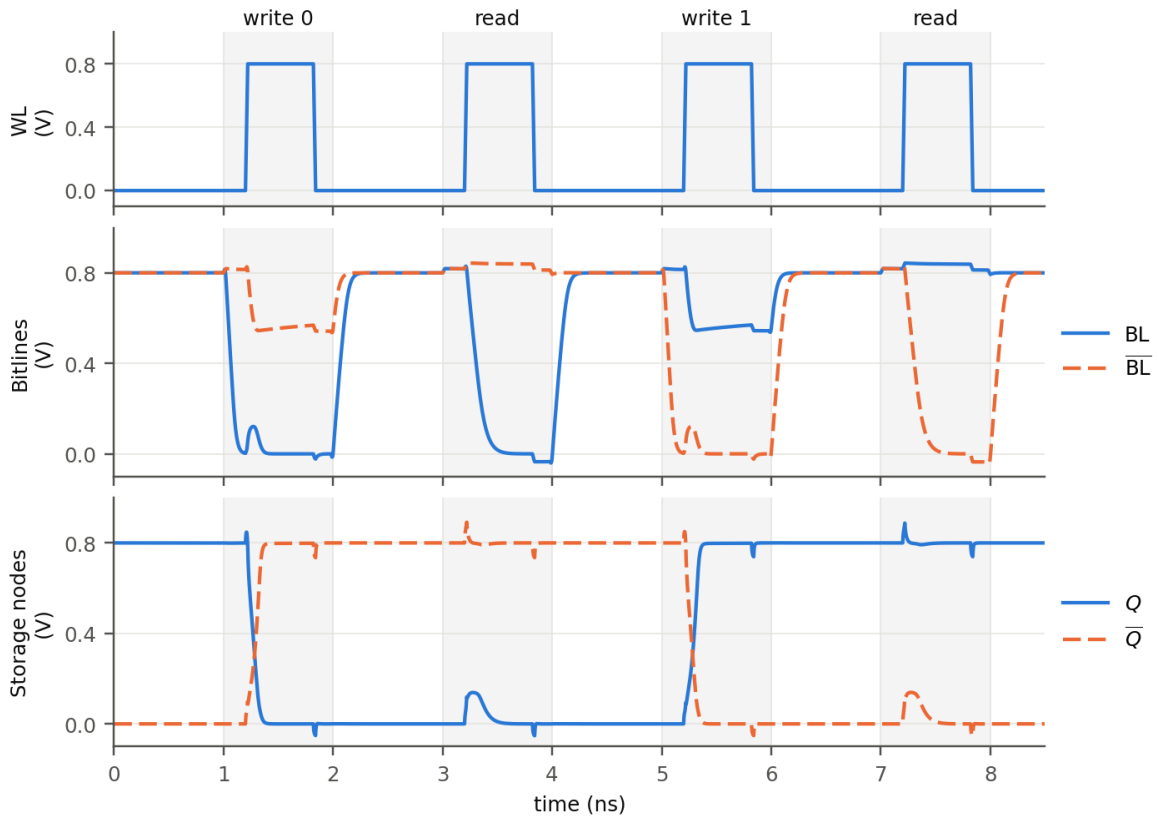


Figure 11: Functional check on the extracted cell with a 16-cell bitline load: write 0, read, write 1, read. Shaded bands are the four operation windows.

Measurement	First (0 stored / written)	Second (1 stored / written)
Write: both nodes cross $V_{DD}/2$ (ps after WL 50%)	88	88
Write: both nodes within 10% of the rails (ps)	131	130
Read: peak "0"-node bump (V)	0.139	0.139
Read: lowest "1"-node level (V)	0.75	0.75
Read: 100 mV bitline split (ps after WL 50%)	26	24
Read: bitline differential at WL fall (V)	0.85	0.85

Table 13: Measured from the exported transient (`tran_222.csv`). "First" is the write-0 phase and the first read (which reads $Q = 0$ on BL); "second" is the write-1 phase and the second read (which reads $\overline{Q} = 0$ on $\overline{\text{BL}}$).

Writes complete. In the write-0 phase BL is already at 15 mV when WL rises. Q falls through $V_{DD}/2$ 59 ps after the WL edge, $\overline{Q}$ rises through it 88 ps after, and both nodes are within 10% of the rails after 131 ps, well inside the 0.6 ns wordline pulse. The write-1 phase is the mirror image (88 ps and 130 ps). The writes are about 25% faster than in the 2/4/6 cell (174 ps), because the storage nodes carry half the transistor capacitance and much shorter wires, even though $\gamma = 1$ makes the static write margin smaller. The bitline that is not driven is floating during a write: when WL opens, the cell's "0" node on that side discharges it, to 0.54 V, until the cell flips. The write still completes, but a real write driver should hold the other bitline at V_{DD}.

Reads do not disturb. During the first read the "0" node Q peaks at 0.139 V, 64 ps after the wordline rises, and then decays as BL discharges and the pass-gate current falls. This peak is just below the DC read-disturb level of 1.4b (0.145 V) and far below the switching point of the opposite inverter (about 0.41 V), so the cell does not flip, and the "1" node never drops below 0.75 V. The second read is identical (0.139 V), as the symmetric layout predicts. The bitlines split by 100 mV within 24 ps to 26 ps of the wordline edge, close to the estimate of about 20 ps from I_{cell}/C_{BL} in 1.4d, and by the end of the wordline pulse the selected bitline has discharged completely, a differential of 0.85 V. The precharge PMOS restores a discharged bitline to 90% of V_{DD} within 0.15 ns. The short spikes on Q, $\overline{Q}$ and the bitlines at each wordline edge (up to 0.89 V and down to -0.05 V) are capacitive coupling of the 20 ps wordline edges through the pass gates; they recover within about 0.1 ns and do not affect the stored data.

6 Discussion

Choosing the minimum-area cell (PU = PG = PD = 2 fins) bought a 5.9× smaller cell (0.53 µm^2), 43–45% less bitline and wordline capacitance per cell, and faster writes, at the cost of a smaller write margin (0.27 V against 0.33 V) and a smaller read current (21.8 µA against 38 µA). It did not cost read stability: the two-fin pull-up raised the switching point enough that the read SNM grew from 0.107 V to 0.128 V, and because the bitline capacitance fell with the read current, a 100 mV bitline split still takes about 25 ps. Read stability remains the smallest static margin, at 16% of V_{DD} at every supply down to 0.5 V, with the write margin next at 33–34%. For Part 2 we would keep this cell and move the well and substrate taps into tap cells along the array, which removes the tap space from every cell. The periphery should hold the undriven bitline at V_{DD} during writes, and the sense amplifier can fire within a few tens of picoseconds of the wordline edge, long before

the bitline fully discharges. If the write margin proves too small across corners, a write assist such as a slightly negative bitline during writes would recover it without changing the cell.

A SNM and write-margin extraction method

Each butterfly uses two DC sweeps with the feedback broken: $V_{\overline{Q}}$ while V_Q is swept, and V_Q while $V_{\overline{Q}}$ is swept. Both curves are plotted on the same V_Q–$V_{\overline{Q}}$ axes. Following Seevinck et al. [1], the axes are rotated by 45°, so each curve becomes a single-valued function. Within each lobe, the largest vertical gap between the rotated curves is the diagonal of the largest inscribed square, so the side is that gap divided by $\sqrt{2}$. We report SNM $= \min(\mathrm{SNM}_L, \mathrm{SNM}_R)$. A brute-force raster search (`p1tools.py snm --check`) gives the same values within the grid resolution. V_{trip} is the swept bitline voltage at which $V_Q - V_{\overline{Q}}$ first changes sign, found by linear interpolation.

B Test bench schematics

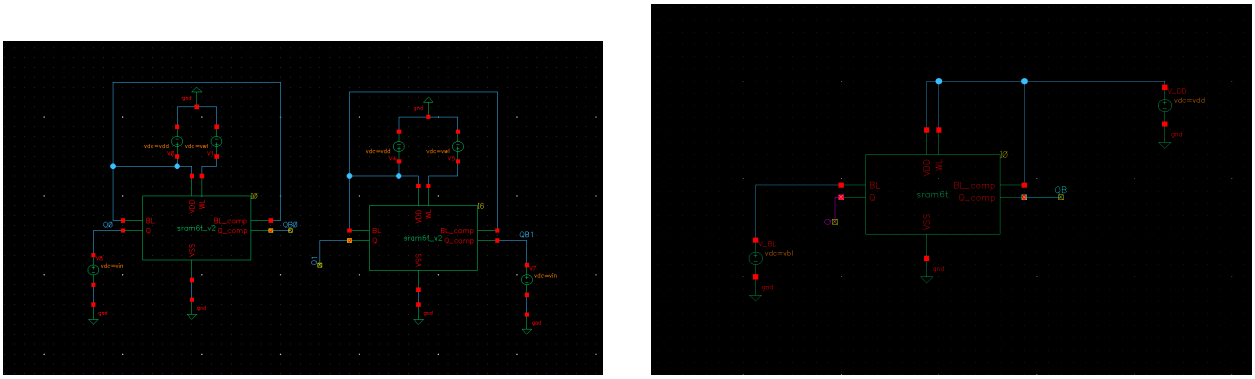


Figure 12: DC test benches for 1.4a–c.

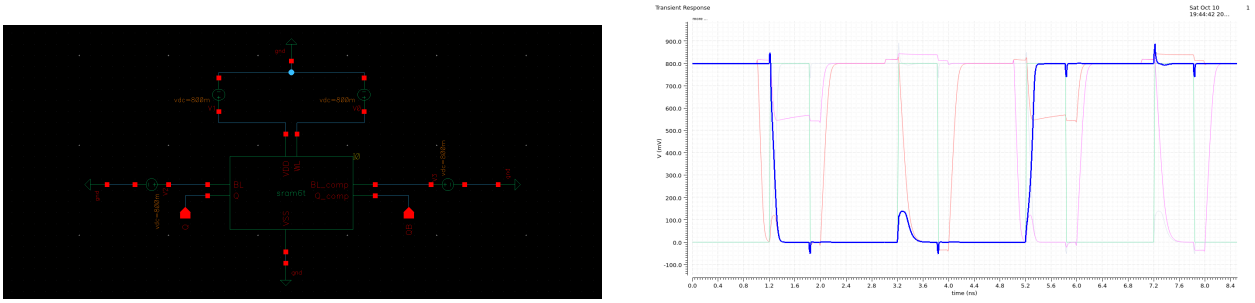


Figure 13: Test benches for 1.4d and 1.5.

References

[1] E. Seevinck, F. J. List, and J. Lohstroh, "Static-noise margin analysis of MOS SRAM cells," *IEEE J. Solid-State Circuits*, vol. 22, no. 5, pp. 748–754, Oct. 1987.