



Design-Technology Co-Optimisation at Gate-All-Around Nodes: Device, Cell, and Routing Considerations

Mathieu Tadesse^{1*}, Soren Alvarez², Ganesh Johansson³

¹Department of Electrical Electronics, Volta Institute for Advanced Electrical Electronics, Bahrain

²Department of Electrical Electronics, Volta Research Institute of Electrical Electronics, Argentina

³Department of Electrical Electronics, Netravati School of Electrical Electronics, Slovakia

Article Info

Volume 1, Issue 1, 2026

doi: [10.62715/ijvlsidsp.2026.v1.i1.pp8-14](https://doi.org/10.62715/ijvlsidsp.2026.v1.i1.pp8-14)

Received: 15 November 2025

Accepted: 04 January 2026

Published: 10 March 2026

Abstract

For most of the history of integrated circuits a new process node meant smaller pitches, and designers received the benefit through a fresh set of layout rules. That arrangement has ended. Gate and metal pitches now shrink slowly, and the density gained at each node comes largely from changes to the standard cell and to the way it is powered and wired, changes that must be negotiated between process engineers and designers before the technology is fixed. This practice is known as design-technology co-optimisation. The move from fin transistors to gate-all-around nanosheet devices makes it both more necessary and more intricate. This narrative review examines the considerations that arise at three levels. At device level it discusses sheet width as a continuous design variable, the number of stacked sheets, inner spacers and threshold voltage options. At cell level it covers track height reduction, the separation of n-type and p-type devices, forksheet and complementary stacked structures, and buried or backside power rails. At routing level it considers pin access, patterning restrictions and the rising resistance of narrow wires and vias. The review argues that these levels cannot be optimised separately and identifies questions that the open literature has not yet settled.

Keywords: *Design-technology co-optimisation, Gate-all-around transistors, Nanosheet devices, Standard cell architecture, Backside power delivery, Interconnect scaling*

© 2026 Mathieu Tadesse et al. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

1. Introduction

The scaling rules set out by Dennard *et al.* (1974) gave the semiconductor industry a simple contract. If every dimension and the supply voltage were reduced by the same factor, transistors became faster, denser and no more power-hungry per unit area, and the doubling of component counts anticipated by Moore (1965) could continue. Designers took part mainly as recipients: each node arrived with a rule deck, and a layout drawn for the old node could be shrunk to fit the new one. Voltage scaling stalled first, then the patterning of ever finer pitches became difficult and expensive, and the contract lapsed (Kuhn, 2012; Bohr and Young, 2017).

* Corresponding author: Mathieu Tadesse, Department of Electrical Electronics, Volta Institute for Advanced Electrical Electronics, Bahrain.
E-mail: mathieu.tadesse@volta.edu

© 2026 Mathieu Tadesse et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

What replaced it is a negotiated process. Before a node is defined, technologists and designers examine together which combinations of pitches, device structures and cell layouts give the best power, performance and area for representative circuits, and adjust both the process and the design rules in response. [Liebmann et al. \(2016\)](#) describe how this practice, design-technology co-optimisation or DTCO, grew out of the restrictions imposed by sub-resolution lithography. [Yeric \(2015\)](#) argued from the design side that with pitch scaling slowing, such co-optimisation would have to supply an increasing share of each node's benefit.

The introduction of gate-all-around (GAA) transistors, in which the gate surrounds horizontal sheets of silicon stacked one above another, gives the practice new material to work with. The nanosheet offers design freedoms that the fin did not, and it arrives together with changes in cell construction and power delivery that alter what the router sees. This review organises the published considerations by level: device, standard cell and routing.

2. Scope and Approach

This is a narrative review of the open literature: device and technology conference papers, journal articles, reference texts and the public roadmap. Foundry process details are confidential, and much of what is published comes from research consortia and from predictive design kits such as ASAP7 ([Clark et al., 2016](#)), so the account describes directions and trade-offs and does not attribute specific dimensions to any commercial node. Node names are treated as labels. The review does not cover memory-specific co-optimisation beyond a brief mention, nor two-dimensional channel materials, and it makes no claim to completeness.

3. From Fins to Nanosheets

A planar transistor loses control of its channel as the gate is shortened, because the drain field reaches under the gate. Wrapping the gate around more of the channel restores control, and the family of multiple-gate devices reviewed by [Colinge \(2004\)](#) follows from that principle; the underlying electrostatics are set out in standard texts ([Taur and Ning, 2009](#)). The FinFET of [Hisamoto et al. \(2000\)](#) placed the gate on both sides of a thin vertical fin, and the structure entered volume manufacture as the tri-gate transistor ([Auth et al., 2012](#)). Successive FinFET nodes made fins taller and thinner and reduced the number of fins per device.

That path ends when a device has one or two fins and the fin cannot be made usefully taller. The gate-all-around structure extends it. Silicon nanowire transistors with the gate on every side were demonstrated in research well before they were needed ([Singh et al., 2006](#)), and later work showed that vertically stacked wires could be built on bulk silicon in a replacement metal gate flow close to that of the FinFET ([Mertens et al., 2016](#)). Widening the wire into a sheet raises the drive current per footprint, and stacked nanosheets were presented as the successor to the fin by [Loubet et al. \(2017\)](#). A manufacturing technology built on such a multi-bridge-channel device followed ([Bae et al., 2018](#)).

4. Device-Level Considerations

The property of the nanosheet most often cited by designers is that its width is defined by lithography and can in principle take any value within a range. Fin devices are quantised: drive strength changes in whole fins, and a cell with two fins cannot be given a little more current. A sheet can be drawn wider for speed or narrower for low capacitance and leakage, which allows a library to offer high-performance and high-density variants on the same process ([Loubet et al., 2017](#)). The freedom is bounded. Wide sheets are harder to release and to fill uniformly with gate metal, and electrostatic control weakens somewhat as the sheet widens, so each process supports only a set of permitted widths.

The number of sheets in the stack is fixed by the process and not by the designer, yet it shapes every cell. More sheets give more current in the same footprint and also more gate and fringe capacitance, taller source and drain regions and a more demanding etch, so the gain per added sheet diminishes. Between the sheets, inner spacers separate the gate from the source and drain; their thickness trades parasitic

capacitance against the resistance of the path into the channel. The space between sheets also limits the thickness of the work-function metals that set the threshold voltage, which makes it harder to offer the several threshold options that low-power design relies on. Self-heating is a further concern, since the sheets are surrounded by materials that conduct heat poorly.

For co-optimisation, the lesson is that device choices are not judged by transistor figures of merit alone. Whether an extra sheet or a thinner spacer is worthwhile depends on the load the device drives, and in a dense logic block that load is increasingly the wiring. The evaluation therefore has to be made on circuits, typically ring oscillators with representative wire loads and then placed and routed blocks.


Static memory deserves a brief mention, because its bit cell is the densest layout on any chip and is co-optimised separately from logic. The cell needs a particular ratio of strengths among its pull-down, pass-gate and pull-up transistors, which fin quantisation made awkward to achieve. Adjustable sheet width restores some of that freedom, and the bit cell is consequently one of the first circuits on which a nanosheet process is assessed.

5. Cell-Level Considerations

The height of a standard cell is measured in tracks, the number of minimum-pitch horizontal metal lines that fit across it. Reducing the track count has been the main source of area scaling in recent nodes, and each reduction removes something: fins, routing tracks inside the cell, or room for the power rails. With nanosheets the devices can be kept reasonably wide while the cell shrinks, and what limits the height becomes the space between the n-type and p-type devices and the width taken by the rails. Table 1 lists the main scaling measures at each level with their benefits and costs, and Figure 1 shows how the levels depend on one another.

Table 1: Principal co-optimisation measures at gate-all-around nodes, with the benefit sought and the cost incurred			
Level	Measure	Benefit sought	Cost or risk
Device	Variable sheet width	Drive strength tuned per cell	Limited set of legal widths; fill uniformity
Device	More stacked sheets	Current per footprint	Added capacitance; taller, harder etch
Cell	Fewer tracks	Smaller cell area	Less internal routing; pin access suffers
Cell	Forksheet n-p wall	Tighter n-p spacing	Gate no longer fully surrounds the sheet
Cell	Stacked complementary devices	n and p share one footprint	Complex integration; vertical contacts
Cell and routing	Buried or backside power rails	Signal tracks freed; lower supply drop	Wafer thinning and bonding; heat removal
Routing	Alternative wire metals	Lower resistance in narrow lines	New integration and reliability questions

The separation between n and p devices is set by patterning margins for the work-function metals and the gate cut. The forksheet device addresses it directly: a dielectric wall is placed between the n and p sheet stacks, so that they can sit much closer, at the price of a gate that wraps three sides of each sheet and not four. [Weckx et al. \(2017\)](#) examined the structure for static memory cells as an instance of device and design being optimised together. The complementary FET goes further and stacks the n-type device on top of the p-type device, or the reverse, removing the lateral separation altogether ([Ryckaert et al., 2018](#)). Its cell can be very low, but the connections to the lower device and the choice of which signals leave the cell at which level become the central layout problem.



Power rails are the other consumer of cell height. In a conventional cell they run in the lowest metal layers along the top and bottom edges and must be wide to carry current. Moving them out of the signal stack frees tracks. A buried power rail is a metal line placed below the devices, in the isolation between them; with a backside power delivery network, the rails are fed through the thinned wafer from a dedicated grid on its rear face. Prasad *et al.* (2019) analysed both options on a processor design and showed that the benefit appears as reduced supply voltage drop and relieved routing congestion as well as cell area. The costs lie in wafer bonding and thinning, in alignment between the two faces, and in removing heat from devices that no longer sit on a thick silicon substrate.

Smaller measures add up alongside these. Placing the gate contact over the active region, replacing the dummy gates between cells with a single diffusion break, and self-aligning gate cuts and contacts each recover a fraction of a track or a gate pitch. Each also tightens the coupling between layout and process, because the layout patterns permitted are those the process module was tuned to print.

6. Routing-Level Considerations

A cell that is small but cannot be reached is of no use. As track counts fall, the pins of a cell have fewer points at which a wire from the routing layers can land, and neighbouring cells compete for the same tracks. Pin accessibility has therefore become a criterion in cell design, checked by placing and routing trial blocks during technology definition and not after it. The algorithms involved are those of conventional physical design (Kahng *et al.*, 2011); what has changed is that their outcome feeds back into the cell architecture and even into the choice of metal pitches.

Patterning shapes the routing fabric. The lower metal layers of advanced nodes are drawn as unidirectional lines on a fixed grid, with line ends defined by separate cut patterns, because that is what multiple patterning and extreme ultraviolet exposure can print with adequate margin; the optical reasons are covered in lithography texts (Mack, 2007). The designer loses the freedom to jog a wire within a layer, and every change of direction costs a via. Rules on cut spacing and on the minimum length of a line segment decide whether two pins in adjacent cells can both be connected, so they belong to the co-optimisation as much as the cell layout does.

Resistance is the third issue. When the width of a copper line approaches the mean free path of its electrons, scattering at surfaces and grain boundaries raises the resistivity, and the barrier and liner layers take a growing share of the cross-section. Gall (2016) showed that the product of bulk resistivity and mean free path predicts which metals suffer least from this size effect, which is the basis for the interest in cobalt, ruthenium and molybdenum for the narrowest lines and vias. For the circuit, high wire and via

resistance means that the delay of a short net may be dominated by the wire and not by the transistor that drives it. This weakens the return on device improvements and strengthens the case for evaluating every technology option on routed designs.

7. Interactions Across Levels

The three levels cannot be closed one at a time. A wider sheet raises drive current but may force a taller cell; a lower cell saves area but reduces pin access, and the router then needs more length and more vias, which adds resistance that the wider sheet was meant to overcome. Moving power to the backside relaxes the cell and the lower metal layers at once and so changes the best choice of track height and metal pitch. The published studies handle these loops by building a provisional process design kit and cell library for each candidate, running a small set of benchmark designs through synthesis, placement and routing, and comparing the outcomes.

The scope of the exercise is also widening. The roadmap now treats system-technology co-optimisation, which adds three-dimensional stacking, packaging and the partitioning of a product across several dice, as the continuation of DTCO (IEEE, 2022). Salahuddin et al. (2018) make a related argument: that future gains will come from combining device, memory and architectural changes and not from any one of them. On that view, the nanosheet transistor is one move in a longer game whose pieces are increasingly above the device.

8. Gaps and Open Questions

Several matters are not resolved in the open literature. Public evidence on variability in nanosheet devices, in particular variation of sheet thickness and its effect on threshold voltage across a library of different widths, is thin compared with what was published for fins. The thermal behaviour of blocks with backside power delivery has been modelled more than it has been measured. Predictive design kits, on which academic work depends, lag the industry: the most widely used one models a fin process, and openly available kits for nanosheets with backside power are only beginning to appear. As a result, published place-and-route studies of such technologies are difficult to reproduce.

There is also a methodological gap. Co-optimisation studies report results for a handful of benchmark blocks, and how far the conclusions carry over to designs with different mixes of logic, memory and wiring is seldom tested. Automating the loop, so that cell libraries are generated and assessed for many candidate rule sets without manual layout, is an active area and a natural meeting point for process engineers and the design automation community.

9. Conclusion

Gate-all-around transistors do not restore the old pattern in which a better device translated directly into a better chip. Their value is realised through the standard cell, the power delivery network and the routing fabric, and each of those constrains the device in return. The literature reviewed here points consistently to the same working method: candidate technologies are judged on placed and routed circuits, with wire resistance and pin access counted from the start. The questions that remain open, on variability, heat and the generality of benchmark results, are ones on which openly published data would be particularly valuable.

Conflict of interest

To be declared by the author(s) before publication.

Funding

To be declared by the author(s) before publication.

References

- Auth, C., Allen, C., Blattner, A., Bergstrom, D., Brazier, M., Bost, M. et al. (2012). [A 22nm high performance and low-power CMOS technology featuring fully-depleted tri-gate transistors, self-aligned contacts and high density MIM capacitors. 2012 Symposium on VLSI Technology](#), 131-132. <https://doi.org/10.1109/VLSIT.2012.6242496>
- Bae, G., Bae, D.-I., Kang, M., Hwang, S.M., Kim, S.S., Seo, B. et al. (2018). [3nm GAA technology featuring multi-bridge-channel FET for low power and high performance applications. 2018 IEEE International Electron Devices Meeting](#). <https://doi.org/10.1109/IEDM.2018.8614629>
- Bohr, M.T. and Young, I.A. (2017). [CMOS scaling trends and beyond. IEEE Micro](#), 37(6): 20-29. <https://doi.org/10.1109/MM.2017.4241347>
- Clark, L.T., Vashishtha, V., Shifren, L., Gujja, A., Sinha, S., Cline, B., Ramamurthy, C. and Yeric, G. (2016). [ASAP7: A 7-nm finFET predictive process design kit. Microelectronics Journal](#), 53: 105-115. <https://doi.org/10.1016/j.mejo.2016.04.006>
- Colinge, J.-P. (2004). [Multiple-gate SOI MOSFETs. Solid-State Electronics](#), 48(6): 897-905. <https://doi.org/10.1016/j.sse.2003.12.020>
- Dennard, R.H., Gaensslen, F.H., Yu, H.-N., Rideout, V.L., Bassous, E. and LeBlanc, A.R. (1974). [Design of ion-implanted MOSFET's with very small physical dimensions. IEEE Journal of Solid-State Circuits](#), 9(5): 256-268. <https://doi.org/10.1109/JSSC.1974.1050511>
- Gall, D. (2016). [Electron mean free path in elemental metals. Journal of Applied Physics](#), 119(8): 085101. <https://doi.org/10.1063/1.4942216>
- Hisamoto, D., Lee, W.-C., Kedzierski, J., Takeuchi, H., Asano, K., Kuo, C., Anderson, E., King, T.-J., Bokor, J. and Hu, C. (2000). [FinFET: A self-aligned double-gate MOSFET scalable to 20 nm. IEEE Transactions on Electron Devices](#), 47(12): 2320-2325. <https://doi.org/10.1109/16.887014>
- IEEE (2022). [International Roadmap for Devices and Systems, 2022 Edition: More Moore](#). Institute of Electrical and Electronics Engineers, Piscataway, NJ.
- Kahng, A.B., Lienig, J., Markov, I.L. and Hu, J. (2011). [VLSI Physical Design: From Graph Partitioning to Timing Closure](#). Springer, Dordrecht. <https://doi.org/10.1007/978-90-481-9591-6>
- Kuhn, K.J. (2012). [Considerations for ultimate CMOS scaling. IEEE Transactions on Electron Devices](#), 59(7): 1813-1828. <https://doi.org/10.1109/TED.2012.2193129>
- Liebmann, L.W., Vaidyanathan, K. and Pileggi, L. (2016). [Design Technology Co-Optimization in the Era of Sub-Resolution IC Scaling](#). SPIE Press, Bellingham, WA.
- Loubet, N., Hook, T., Montanini, P., Yeung, C.-W., Kanakasabapathy, S., Guillorn, M. et al. (2017). [Stacked nanosheet gate-all-around transistor to enable scaling beyond FinFET. 2017 Symposium on VLSI Technology](#), T230-T231. <https://doi.org/10.23919/VLSIT.2017.7998183>
- Mack, C.A. (2007). [Fundamental Principles of Optical Lithography: The Science of Microfabrication](#). John Wiley & Sons, Chichester.
- Mertens, H., Ritzenthaler, R., Hikavy, A., Kim, M.S., Tao, Z., Wostyn, K. et al. (2016). [Gate-all-around MOSFETs based on vertically stacked horizontal Si nanowires in a replacement metal gate process on bulk Si substrates. 2016 IEEE Symposium on VLSI Technology](#).
- Moore, G.E. (1965). [Cramming more components onto integrated circuits. Electronics](#), 38(8): 114-117.
- Prasad, D., Nibhanupudi, S.S.T., Das, S., Zografos, O., Chehab, B., Sarkar, S. et al. (2019). [Buried power rails and back-side power grids: Arm CPU power delivery network design beyond 5nm. 2019 IEEE International Electron Devices Meeting](#). <https://doi.org/10.1109/IEDM19573.2019.8993617>
- Ryckaert, J., Schuddinck, P., Weckx, P., Bouche, G., Vincent, B., Smith, J. et al. (2018). [The complementary FET \(CFET\) for CMOS scaling beyond N3. 2018 IEEE Symposium on VLSI Technology](#), 141-142. <https://doi.org/10.1109/VLSIT.2018.8510618>
- Salahuddin, S., Ni, K. and Datta, S. (2018). [The era of hyper-scaling in electronics. Nature Electronics](#), 1:

442-450. <https://doi.org/10.1038/s41928-018-0117-x>

Singh, N., Agarwal, A., Bera, L.K., Liow, T.Y., Yang, R., Rustagi, S.C. et al. (2006). [High-performance fully depleted silicon nanowire \(diameter \$\leq 5\$ nm\) gate-all-around CMOS devices](#). *IEEE Electron Device Letters*, 27(5): 383-386. <https://doi.org/10.1109/LED.2006.873381>

Taur, Y. and Ning, T.H. (2009). [Fundamentals of Modern VLSI Devices](#), 2nd edition. Cambridge University Press, Cambridge.

Weckx, P., Ryckaert, J., Putcha, V., De Keersgieter, A., Boemmels, J., Schuddinck, P. et al. (2017). [Stacked nanosheet fork architecture for SRAM design and device co-optimization toward 3nm](#). *2017 IEEE International Electron Devices Meeting*.

Yeric, G. (2015). [Moore's law at 50: Are we planning for retirement?](#) *2015 IEEE International Electron Devices Meeting*. <https://doi.org/10.1109/IEDM.2015.7409607>