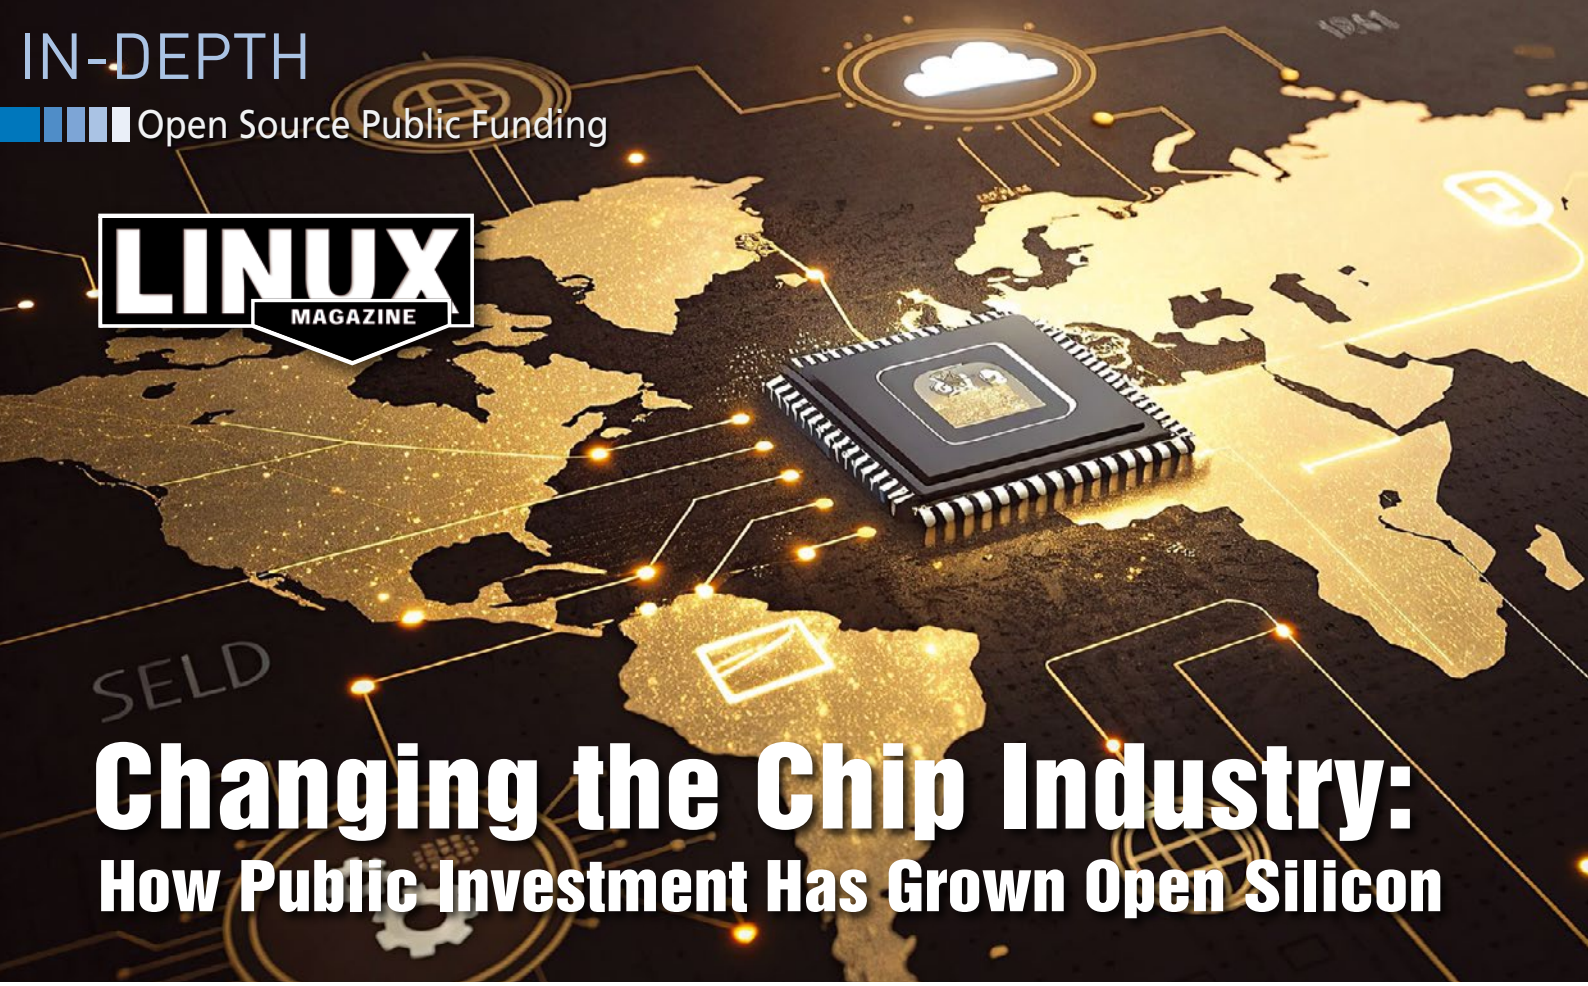

Changing the Chip Industry:
How Public Investment Has Grown Open Silicon

The first in a five-part series on how public funding of open source contributes to digital autonomy. *By Jeffrey A. McGuire*

In 2019, in a small restaurant in Paris, a man set a laptop on a dinner table for two other engineers to examine. What was on it should not have existed. They were looking at a Process Design Kit, or PDK: the complete set of physical rules and transistor models used by a chip factory in the United States to manufacture silicon. These documents were among the most fiercely guarded in the global technology industry; no engineer outside the factory's customer roster could see one without first signing a non-disclosure agreement (NDA). The fact that it was on a laptop in Paris, with no NDA attached, was the start of something.

Seven years later, that beginning has become a small constellation of open silicon projects. Together, they enable the design and manufacture of working chips with software anyone can read, modify, and share.

European Union public money has paid for a meaningful part of what now exists, and most Europeans have no idea. To see why it matters, you have to understand what designing a chip actually entails, and why almost none of that process was open before 2019.

How Chips Get Designed: EDA

Picture the work that ends in producing a microchip as a tall stack of software. At the top, humans sit at keyboards writing in hardware description languages such as SystemVerilog. Their code describes what the chip should do in terms of logic, not yet as physical shapes. Beneath that human-readable layer, a pile of automated tools translates the logic downward: Synthesis turns the description into gates, placement and routing decide where those gates sit on the silicon surface of the chip, verification checks the design against its specification, and sign-off tools confirm that the design meets the manufacturing rules and standards of the chip factory – called a "foundry" in the industry. At the bottom of the stack, the foundry uses finished designs to produce physical chips.

The stack of tools between the engineer and the foundry is all part of the category "Electronic Design Automation," or EDA. Most chip designers work "fabless," meaning they design entirely on their computers, send their finished files to a foundry that builds the chips, and the foundry ships the chips back. The whole supply chain is software-mediated, and most of that software comes from just three companies: Cadence, Synopsys, and Siemens EDA. Together, they control the bulk of the market – by most estimates, 75 to 90 percent.

The pricing of these tools is not public. People in the industry quote licensing fees in the million-euro range per engineer per year, with case-by-case custom discounts of 99 percent or more. That opacity is a kind of leverage: A vendor who decides who gets how big a discount also decides which customers, which sectors, and which countries can afford to participate. It is geopolitical influence dressed up as commercial pricing.

There is one more piece. When an engineer has finished with their chosen toolchain (top to bottom, logic to layout), foundries require a final check using a proprietary sign-off tool such as Siemens's Calibre before manufacturing begins. The open ecosystem has to interoperate with that gate; nobody can replace it without the foundry's blessing.

The Chokepoint: PDK

For most of the industry's history, almost none of the chip production software stack could practically be opened – and not because nobody tried. The bottom

layer, the foundry, kept the most important document a designer needed locked up: the Process Design Kit, or PDK.

A PDK is the bundle of files – transistor models, design rules, library cells, parameters – that a foundry hands its customers so their tools can produce designs the foundry can actually manufacture. Every foundry has its own PDK for each process it offers. Until 2019, all the commercially relevant PDKs in the world were under license and NDA.

Luca Alloatti, a silicon photonics researcher and one of the co-founders of the Free Silicon Foundation (F-Si), gave an analogy to help envision this situation. Imagine that Microsoft Word didn't have just one *.docx* format, but also *.docx1*, *.docx2*, *.docx3*, all the way through to *.docx1000000*, and each of those formats were individually under NDA. LibreOffice, in this world, would still want to read them. But if any LibreOffice contributor had ever signed an NDA on any *.docx* version and a parser for that version then appeared in the LibreOffice source tree, Microsoft could sue the whole project. Functionally, the entire ecosystem of free office software compatible with proprietary formats could not exist.

That is roughly the situation that produced closed EDA. A researcher's fix to make a free tool work with one foundry's PDK would live and die on their machine. There was no way for a public, compounding body of knowledge to form. Even when open chip-design tools such as the DARPA-funded project OpenROAD (started under Andrew B. Kahng at the University of California, San Diego) did connect to a real foundry after going through closed-PDK gymnastics, nobody could share those routines with other engineers.

Open source software cannot exist without a through line of openness: If your open code has to talk to secret code to function, you cannot legally publish it. The industry had arranged itself, for decades, to keep the fabrication layer secret.

The Unlock: The Paris Laptop

The man with the laptop in Paris was Jeff Carr, one of the co-founders of DigitalOcean, a US cloud computing company. Sometime around 2018, he became personally interested in the open silicon question. He drove out to a foundry called SkyWater Technology, parked outside, and walked in to ask how much it would cost to buy the factory. He did not buy it. But through some combination of his involvement and a push from Google, the SkyWater 130-nanometer PDK emerged: the first open Process Design Kit in the world.

The laptop in Paris a year later linked up the full EDA stack, from node (industry shorthand for a particular manufacturing process and its physical scale) to tapeout (the moment a finalized design is committed to a foundry for commercial production). A real foundry and a real manufacturing process, available to anyone, without any NDAs.

What followed Paris was a cascade. GlobalFoundries, a much larger US foundry, opened its GF180MCU process (a 180-nanometer node aimed at microcontrollers) in partnership with Google. Then came IHP Microelectronics in Germany, the country's national research foundry. OpenROAD's five-and-a-half-year DARPA program concluded at the end of 2023; the project has continued as a nonprofit, with paid commercial support offered through Precision Innovations. Since 2019, several hundred chips have been taped out using open tools. Google's open shuttle program alone has manufactured 360 designs, selected from more than 600 submissions across 19 countries.

In a "130-nanometer node," the smallest features are roughly 130 nanometers across. The cutting edge today is 3 nanometers and below, dominated by TSMC, the Taiwanese foundry that fabricates the world's most advanced chips. Open silicon is not competing at that frontier, but at mature nodes (130nm, 180nm, 90nm) where the bulk of real-world chip applications live: in automotive sensors, analog circuits, microcontrollers, home appliances, anything that does not need to push silicon to its physical limits. OpenROAD, the open source digital design tool, has been used for commercial tapeouts at nodes from 180 down to 12 nanometers.

Seven years ago, you could design openly but not fabricate openly. Today, open silicon projects have users whose work, built end-to-end with open tools, has reached industrial tapeout to produce chips at scale.

EDA: The Philosophical Perspective

The Free Silicon Foundation, or F-Si, sits on the policy and philosophy side of this cluster. Among its co-founders is Thomas Kramer, an electrical engineer who studied at the *Eidgenössische Technische Hochschule Zürich* (ETH Zurich, the Swiss Federal Institute of Technology Zurich).

Thomas's path into open silicon started in a university lecture in the 2010s. As an electrical engineering undergraduate at ETH Zurich, he took a Very Large Scale Integration (VLSI) design class to learn how to arrange the millions or billions of transistors on a modern chip: The homework software required a signed NDA. That did not feel right to him, so he asked his lecturer if there were open source alternatives he could use at home – the lecturer did not know.

After Luca Alloatti completed his PhD in electrical engineering in Germany and a postdoctoral position at MIT – where he helped build the first microprocessor to use light rather than electricity for some of its internal connections – he found himself inside the "intellectual force field" of the Free Software Foundation. Richard Stallman, the FSF's founder, has spent more than 40 years arguing that software users deserve the freedom to run, study, modify, and share the code they depend on. "That contaminated me," Luca says with a smile. He went back to Europe to apply the same idea to silicon.

In 2018, when Luca brought the idea back to European academia, his colleagues called him quixotic, from Cervantes' deluded knight. But the windmills were not imaginary. Today, hundreds of chips have been taped out with open tools, the European Commission is funding open EDA development, and many of the same professors who shook their heads at Luca now describe themselves as long-standing supporters of open silicon. The knight turned out to have been looking at giants all along.

F-Si, part of the Next Generation Internet Zero (NGI0) Consortium, runs the annual Free Silicon Conference (FSiC) and publishes formal policy recommendations to the European Commission. They're also developing the Hardware

GPL (HGPL), a copyleft license for silicon chip designs that requires derivative works to carry the same terms. The HGPL will address a real, documented gap. The strongly reciprocal CERN Open Hardware Licence, CERN-OHL-S, is incompatible with the GPL, a clash of two strong copyleft licenses that both CERN and F-Si acknowledge. Because a complex silicon design is likely to incorporate or link GPL-licensed code, F-Si argues a GPL-compatible hardware license is needed to keep the free-hardware and free-software communities from fragmenting.

The foundation also maps where the field's knowledge actually lives: F-Si ran a language model over a September 2025 dump of the OpenAIRE thesis database and pulled out 13,272 related to chip design and EDA, each scored across 53 categories. The result is a public, searchable database of where the discipline has already done its thinking, useful both to a researcher entering a niche and to F-Si when they look for university groups to visit.

Underneath the policy and licensing work is a framing Luca returns to often. "Digital infrastructure is public infrastructure," just like roads, hospitals, and the power grid. "We all need it. We all rely on it." When public infrastructure goes private, the public pays twice: once with taxes to build it, again with fees to use it.

EDA: The Industrial Perspective

Naja, an open source EDA project that has received funding from NGI0 programs, sits on the industrial-adoption side of the EDA cluster. The company that built it, Keplertech, was founded by Christophe Alexandre (CEO) and Noam Cohen (CTO).

Christophe has spent his entire career in EDA. He did a PhD, co-founded a startup in 2005 and watched it get acquired by Mentor Graphics (originator of the Calibre sign-off tool), watched Mentor get acquired by Siemens, and decided after 15 years that he had had enough of corporate politics defending market position rather than improving products. Noam arrived from a parallel direction: electronic engineering studies in Israel, an internship at the US chip company Marvell, and a job at Synopsys in Tel Aviv that brought him to Paris through an internal transfer. He and Christophe met at Siemens. Both of them smile when they describe realizing they wanted to leave and build something else.

They had a specific frustration with the standard startup playbook in the EDA space. That playbook, Christophe explains, is: build a small plugin that solves one narrow problem inside one of the big three's toolchains, generate some traction, get acquired. "If you do just plugins," he says, "you're going to maybe replace one box in the stack, but you need to replace 50 boxes" to make the stack open. And that was not on offer.

Naja's choice of Apache 2.0 follows directly. A permissive license, it lets industrial users incorporate the code into their own products without being required to release their work in turn. F-Si pushes copyleft for a different reason, to preserve the public-good character of the work downstream. The contrast between the two licenses is tactical rather than philosophical. Both want adoption and openness: They are aiming at different doorways.

Naja's foundation is a set of EDA tool-building libraries written in C++. These building blocks are accessible through Naja EDA, a thin Python wrapper that can be installed with a single command (*pip install najaeda*). On top of Naja EDA sits Kepler Formal, an equivalence checker – a tool that confirms two versions of a chip design are functionally identical even after someone has rewritten or optimized one of them. OpenROAD has adopted Kepler Formal to verify that its design transformations don't change chip functionality between iterations.

Naja started near the bottom of the human-relevant part of the stack, in a step called post-synthesis, and is now climbing upward toward the engineers who write the original code. Christophe describes the trajectory as moving "closer to the humans." A SystemVerilog front end is integrated; logic equivalence checking at the human-readable level is the next step.

Kepler Formal is also about to ship as a GitHub Action, an automation that runs whenever someone proposes a change to an open chip design project hosted on GitHub. If a contributor submits a pull request, a GitHub runner can invoke Kepler Formal and confirm that the change has not silently broken the design's logic. Continuous integration like this is unremarkable in software development, but it does not exist anywhere else in EDA today.

A note in defense of the industry Naja is trying to "chip" away at, which Noam and Christophe both volunteered without prompting: Chips cannot be patched. A bad chip is a recall, a lawsuit, or a generation of products written off. The cost of an EDA license, even at the proprietary list price, is dwarfed by just the electricity bill for the servers that run the tools – never mind the cost of materials and the other consequences of failure. The industry's conservatism is rational, even when the lock-in it creates is not.

Different Doors, Same Building

F-Si and Naja can look like opposites: a philosophical, copyleft-pushing, policy-focused European foundation on one side and a permissive-licensed, industrial-minded startup on the other. The people doing the work do not see it that way.

Luca puts the point plainly: "I would say there are not different approaches; it's one big one." Within the open ecosystem, between F-Si and Naja, OpenROAD, and dozens of others, there is no real conflict. People share tools, attend each other's conferences, and contribute to each other's projects. Christophe has presented Naja at F-Si's own Free Silicon Conference every year from 2022 to 2025. The choice of license is a tactical decision about which doors a project wants to walk through, not a wall between rival camps.

The cooperation runs deeper than conference invitations. Naja's front end relies on Verilator and the Slang parser, both of which originated in the US open source ecosystem. OpenROAD, which began life at the University of California, San Diego, is integrating Kepler Formal into its own workflow. The users span continents, and the dependencies cross borders; the ecosystem is international by design.

In this industry, "international by design" cuts to the heart of what digital autonomy can and cannot mean. A

modern chip is the product of a global supply chain that nobody, anywhere, fully controls. The foundries that build leading-edge silicon are concentrated in Taiwan, with smaller capacity in Korea, the US, and (increasingly) Germany. The machines those foundries use to print transistors at the nanometer scale come from a single Dutch company, ASML, with no real competitors at the leading edge. End markets are everywhere. No region is anywhere close to being able to do the whole thing on its own.

The Next-Generation Talent Question

Today, an electronics engineering student in Europe will mostly learn on proprietary tools. That is what their professor has set up. That is what the university's licensing arrangement covers, usually under an academic discount of 99 percent or more from one of the three big EDA vendors. And that is what students want to learn, of course, because that is what companies hire for.

Jason Cong, a computer scientist at UCLA, has argued for opening chip design to engineers trained in software. In a 2024 conference keynote on "democratizing" chip design, he described teaching undergraduates with no hardware background to design a working chip accelerator in about a week and a half. European academia is not yet structured around that insight; chip design programs tend to treat software as a tributary rather than as a peer discipline.

Noam reframes the choice from the student's perspective: "If you want a job, learn proprietary tools." Industrial recruitment is downstream of the tools on a graduate's resume. "If you want to start a startup, learn open." New companies cannot launch without paying a million-euro license bill before they have even built a prototype. These are two paths into the same industry, and the second path only exists thanks to the open ecosystem.

One more piece of the puzzle, from Luca: Many of the same professors who could add open tools to the curriculum have an interest in keeping their relationships with proprietary vendors. Discounts of 99 percent or more are contingent on vendor goodwill, so a professor who agitates publicly for open alternatives may find their next discount renegotiation harder. That dynamic is hard to see from outside academia, but it almost certainly shapes what gets taught.

This is the same pattern that locks European public administrations into Microsoft and Amazon Web Services at a different layer of the technology stack. The argument is not that the proprietary tools are bad: Having no real alternative is. Digital autonomy means having options available.

What Public Money Buys

The first EU Chips Act, passed in 2023, put roughly 43 billion euros behind one goal: doubling Europe's share of global chip manufacturing by 2030. The proposal for the Chips Act 2.0, adopted this June, aims to position the EU within the broader ecosystem by complementing newer initiatives such as the Cloud and AI Development Act and the Open Source Strategy.

The goal of reducing dependency has sometimes seemed out of reach, and the European Court of Auditors has pointed to the gap between European research and European fabrication as the core weakness – the second version of the Act explicitly focuses on increasing demand and innovation. The open EDA community has a sharper critique. A meaningful share of that public money has gone straight back out the door as license fees, paid to the same three proprietary vendors the investment was meant to make Europe less dependent upon. European startups working inside EU-funded design centers build their chips on cloud-hosted Cadence and Synopsys: software the public paid for and the vendors still own.

EuroCDP, the EU Chips Design Platform, is one potential answer. Coordinated by the Belgian research institute imec with 12 European research partners, it was set to onboard its first startups and SMEs in early 2026. EuroCDP gives European chip designers access to commercial EDA at negotiated discounts and to open source EDA tools, plus fabrication, training, and startup support. Within the platform, open source is treated as a pillar in its own right, with the same standing as the commercial offering. Whether the open pillar grows in proportion to the commercial one is the big "open" question.

What makes open silicon possible today exists in part because the European Commission and member-state funders have backed it for years, through programs few Europeans know exist. Increasing the technological sovereignty Europe says it wants comes down to sustained investment of public money in building tools, rather than renting them.

The Chips Act 2.0 broadens the agenda to include packaging, materials, and cross-border investment in fabrication, and it runs directly into this choice. Money spent on licenses is money not spent on engineers, designs, or training, so the same public grant goes further when the toolchain underneath it is free and open – the same is true when it comes to industrial design and fabrication. The second wave of European semiconductor legislation introduces "Grand Challenges" and "Demand Accelerators" to offer tangible resource support at the policy level, bringing public procurement money into the open silicon space.

This article was made possible by support from NLnet Foundation through Linux New Media's Topic Subsidy Program. NLnet Foundation is the coordinator of the NGI Zero Programs which are funded by the European Commission.

Links for further reading:

- *EU Open Source Strategy Plays Key Role in Tech Sovereignty Package*
- *Google's Open MPW shuttle program announcement*
- *OpenROAD Initiative News and Updates*
- *EuroCDP*
- *European Commission Policy Page: Chips Act 2.0*
- *European Parliament Legislative Train Schedule for tracking progress*
- *Open EDA ecosystem supported by NGI Zero*
- *NGI Zero Programs*

Links for your toolkit:

- *F-Si EDA toolchain directory (GitHub)*
- *F-Si EDA and chip-design thesis database*
- *SkyWater SKY130 open PDK (GitHub)*
- *IHP Microelectronics open PDK (SG13G2)*
- *Naja (GitHub)*
- *Kepler Formal (GitHub)*