Education of Chip Designers at a Large Scale: A Proposal

The renewed government interest in the semiconductor industry has recently fueled tremendous investments across the globe. President Biden's CHIPS Act has allocated US$53 billion to this industry, most of which has been earmarked for chip *fabrication*. The question is, Who will *design* the chips that must fill the capacity of these fabrication lines?

This article proposes a plan for educating future generations of chip designers such that they can enter the industry with sufficient knowledge and experience andbegin to develop products. While envisioned for a university environment, the plan can also be deployed in the industry with minor modifications. We demonstrate how 50 students were trained over the course of three quarters.

Typical Chip Development Process

Figure 1 illustrates the development stages that the industry follows for chips of moderate complexity. Beginning with product definition and performance specifications, engineers embark upon the design of the chip and, in about six to nine months, complete the first "draft." The result is in the form of tens or hundreds of "schematics." Next, the chip goes to "layout"; i.e., designers create geometries (mostly squares and rectangles) that must be built on the actual chip. Since layout introduces its own "parasitics," the result may not meet the performance targets, requiring iteration with the design. Anticipating that the package in which the chip will be housed will contribute additional parasitics, we account for them in this iteration phase as well. This cycle takes three to six months.

After the chip layout is completed and is verified to match the schematics, it is sent for fabrication, e.g., to Taiwan Semiconductor Manufacturing Co. Fabrication requires three to four months. Upon return, the chip is mounted in a suitable package and undergoes test and characterization; this phase takes another three to four months.

Chip Design Education: What to Teach?

Many undergraduate and graduate university curricula have offered courses in chip design for decades. Some also include design projects to give students hands-on experience. We can then ask, Does this experience sufficiently prepare the students for the industry? The answer is, generally, no.

As evident from Figure 1, a class design project spanning five to six weeks touches upon only a small part of the chip development process. Notably, the layout, fabrication, packaging, and test components of this flow are not explored. Even if chip layout is included, students do not find the opportunity to test, in the laboratory, the results of their labor. Consequently, they do not learn myriad practical issues that beset chip design.

One might argue that these missing components of chip design education can be learned after students enter the industry. However, product development cannot afford to place an inexperienced engineer in the critical path. Thus, such training in the industry occurs only slowly and haphazardly.

For these reasons, true design education must encompass all the phases depicted in Figure 1—perhaps at a smaller scale.

Digital Object Identifier 10.1109/MSSC.2024.3395073
Date of current version: 25 June 2024

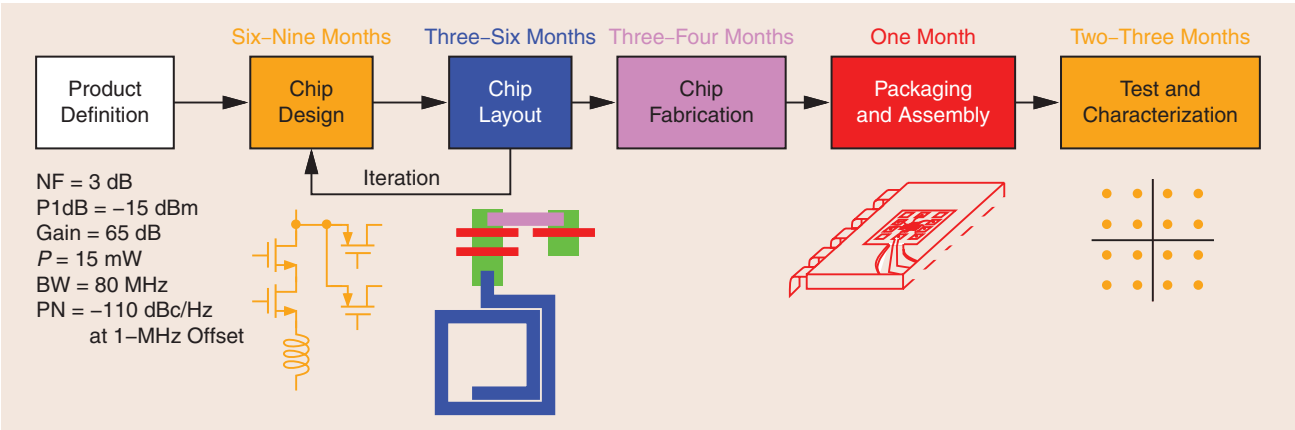

FIGURE 1: A typical chip development time frame.

Enter the "Tapeout" Class

In our community, a course or course sequence that practices the stages in Figure 1 is called a "tapeout" class. The idea dates back to the early 1980s, when John Newkirk, Rob Mathews, and Mark Horowitz pioneered such an effort at Stanford University. (I took that course in 1988, under Horowitz.)

In recent years, a number of universities have reignited the concept, including Carnegie Mellon, Columbia [1], Georgia Tech, Stanford, and the University of California, Berkeley [2]. Other activities have also been reported [3], [4].

When I was approached by Apple two years ago to create such a course at the University of California, Los Angeles (UCLA), I accepted with great excitement, especially because I saw the largest enrollment in our regular chip design courses and, hence, an opportunity to develop a scalable model for a tapeout class. The journey included numerous challenges and required numerous solutions, which I hope to share in this article.

Problem of Export Control

Shortly after I began to work on the tapeout class, I realized that one critical component—fabrication—might not be possible! The U.S. Department of Commerce has designated most fabrication technologies of various chip manufacturers as "export-controlled"; i.e., they cannot be accessed by citizens of certain countries.

I then asked other universities how they handled this restriction. One university would have such citizens work with FPGAs, in essence replacing chip design with coding. Another would ask the course teaching assistant to serve as a mediator between the export-controlled technology and students of those countries. Since we cannot discriminate among our students at UCLA, these remedies were not feasible for us.

I then learned that a few technologies were exempt from export control, but some entailed long fabrication cycle times (six months), which would not align with our time frame (described in the next section). I finally selected SkyWater's 130-nm process, as it could deliver chips in about four months. This decision was met with great skepticism by others, as a number of tapeout classes using SkyWater had seen complete failure in their chips. Nevertheless, I decided to take the risk.

Course Structure and Timeline

In a quarter system, such as ours, the course can be structured as illustrated in Figure 2. We offer conventional analog and digital chip design courses in the fall quarter, which most circuits-oriented graduate and some undergraduate students take.

With a proper foundation thus established, students begin the tapeout class (called "ECE209" internally) in the winter quarter, with a two-unit course. This is followed by a four-unit course in the spring and another two-unit course the next fall. Only the one in the spring counts toward students' degree requirements. Such a structure amounts to one year of training, allowing the instructor to impart a great deal of practical knowledge.

First Course

The objectives of the first course are threefold: 1) to teach the students the design of certain building blocks that are likely to appear in their final chips; 2) to help them construct the layout of these circuits, extract the parasitics, and quantify how much the performance degrades; and 3) to guide them toward their final project topic.

With 70 students in the first course, and a potentially wide range of project topics, it was necessary to implement scalability in the affairs, from guiding the students through the learning process to reviewing their work and answering their questions. Specifically, two issues arose. First, how would 70 students learn the exact mechanics of the design, layout, and verification of the building blocks in SkyWater's technology? For this and other training purposes, I asked my Ph.D. students to create several videos. Second, how would I check the 70 students' work for each building block? To this end, they would bring their results to class on a flash drive; I would randomly select different students to present their design and the layout to the class. I would then provide detailed feedback and suggest modifications. While some students might not have a chance to present their work, they still learned from others in class, as they all worked on the same building block in a given week.

The building blocks covered in this course included ring oscillators, latches, flip-flops, frequency dividers, the StrongArm comparator,

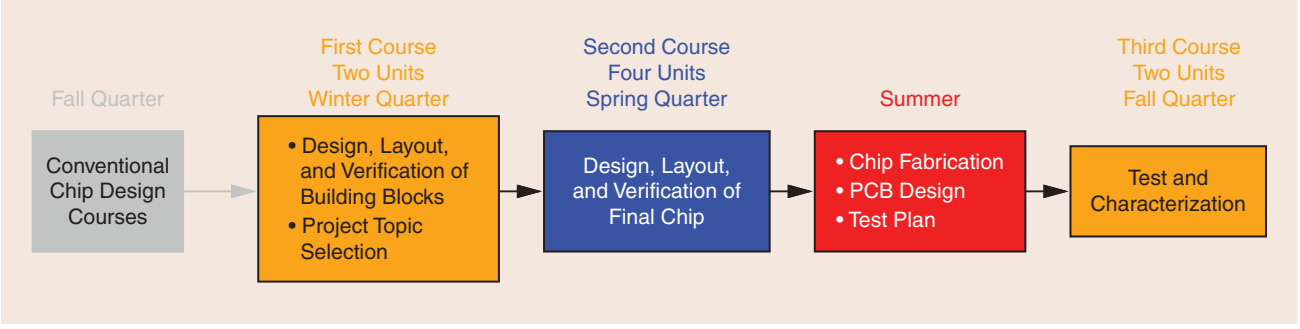

FIGURE 2: The structure of our tapeout course.

registers, and *LC* oscillators. Spanning six homework assignments over 10 weeks and gradually increasing in complexity, these circuits helped the students gain a detailed understanding of the entire design flow for a broad range of applications. Besides the in-class presentations, these assignments were not evaluated; in our second offering of the course this year, we have employed a teaching assistant for grading the homework.

In this course, I also taught practical issues, such as device and interconnect parasitics, capacitor and inductor structures, critical layout considerations, I/O techniques, ESD devices, and package parasitics.

This course was accompanied by a timeline stipulating certain action items for each week. These included finding a partner, selecting a project topic, reading books or papers about that topic, simulating at least one building block of the selected project, and developing a tentative plan for the design and simulation of the final chip. The partners would need to specify the building blocks that each of them owned and would present their parts separately in the design reviews in the second course (as explained below.)

Of the 70 students in this course, 50 continued to the next one. Those who stopped at this juncture received only a pass grade but learned a great deal.

Selection of Project Topics

In addition to covering good design and layout practices, the first course also helped the students converge on the type of circuit that they wished to implement as their final chip. I provided a long list of possible topics, e.g., ADCs, DACs, crystal oscillators, analog or digital PLLs, DLLs, RF receivers, RF PAs, baseband analog filters for RF receivers, distributed amplifiers, clock and data recovery circuits, DFEs, high-speed PRBS generators, energy harvesting systems, chopper-stabilized amplifiers, TDCs, digital-to-time converters, digital thermometers, dc–dc converters, digital multipliers, digital adders, digital accumulators, digital filters, digital $\Delta\Sigma$ modulators, and FFT engines. Students could also propose their own topic.

The selection of a topic was governed by three constraints:

1) The chip had to be designed, laid out, and verified in about nine weeks so that it could be sent for fabrication in early June. Students could even begin their designs in the winter quarter so as to give themselves more time.
2) The chip had to lend itself to development by two partners. This number is somewhat flexible, but I felt that it was a good compromise between the chip complexity and the possibility of chip failure due to miscommunication among partners.
3) The fabricated chip had to be tested using the equipment available to this class. As the first course proceeded, I continued to compile a list of necessary equipment, as explained below.

By the end of the first course, students had finalized their 25 project topics. We had eight PLLs, eight ADCs, two DACs, three RF receivers, one dc–dc converter, and three digital multipliers.

Choice of Design Software

For their tapeout classes, some universities have relied on "open source" software, i.e., tools available for free to both academia and industry. While serving a broad community, such tools present several shortcomings in comparison to the commercial tool Cadence. First, the lack of cohesion among their constituent parts makes them difficult to use, hence the semiconductor industry's preference for Cadence. This means that the use of Cadence in a tapeout class prepares the students better for industry. Second, open-source tools can contain serious errors and inconsistencies, thus frustrating the students and even leading to chip failure. I thus decided on Cadence, especially because it is available to most universities for free. As open-source tools continue to improve, they may lend themselves to robust integration in one platform and prove practical for tapeout classes.

Industry Mentors

As part of my effort toward making the tapeout class scalable, I envisioned that the 25 teams could benefit from regular interaction with engineers in the industry. With the project topics defined, I wrote to my contacts at various companies and asked for experts who could mentor the students. The response was overwhelmingly positive.

I matched 25 experts from Apple, Boston Scientific, Broadcom, Intel, Samsung, and Texas Instruments with the students. The mentors' time commitment consisted of 1 h per week (by Zoom) for nine weeks. Students confirmed later that they learned a great deal from their mentors. This relationship also opened doors to the students for future employment.

Second Course

The bulk of the design took place in this course for about nine weeks. To ensure that all the teams moved forward at a proper pace, I decided to conduct a design review with each team every week. I emphasized that, as the students had seen in the first course, layout, extraction, and design iteration would consume a great deal of time. The teams were then required to complete their schematic designs in the first five weeks and layout in the next four.

I allocated three afternoons per week to the design reviews, giving each team 10 to 15 min. The first afternoon was dedicated to PLLs, the second to Nyquist-rate ADCs, and the third to the others. I also realized that the review proceeded more efficiently by Zoom, as students could readily go back and forth between their slides and their Cadence database. To ensure that both partners moved along well, I posed questions to them separately and created action items for

each. The students also found that attending an entire session proved beneficial because they could learn from my feedback to other teams as well. As mentioned in the previous section, the students also met with their industry mentors once a week.

Most teams completed their schematic simulations in four weeks. I then went through their work carefully to ensure a sound design or to suggest a last round of modifications. I found it more rigorous to examine their actual Cadence files rather than learn only from their slides.

Most teams finished their layout by the seventh week. Again, I examined each chip layout myself and provided feedback.

As the last step before the tapeout, the teams were asked to extract their entire chip and perform a "pad-to-pad" simulation so as to capture the effect of I/O circuits and include an estimate of package parasitics.

By the end of the ninth week, all 25 chips were ready to go to fabrication.

Inductors, I/O Pads, and ESD Devices

The SkyWater process design kit provides various active and passive devices but not spiral inductors. The eight PLLs and three RF receivers, on the other hand, did rely on such components. Realizing that the tight time frame did not allow the class to learn inductor modeling and design, I asked one of my Ph.D. students to create a library with values ranging from 1 to 8 nH. He designed the inductors in SkyWater's process, simulated them using Cadence's EMX tool, and provided both the layout and the models. The class then chose only from this library.

The I/O pads originally available from SkyWater included a great deal of circuitry and large ESD protection devices, limiting the speed to only tens of megahertz. To accommodate the gigahertz speeds in our projects, I asked the project aggregator, Efabless, to create simple pads with moderate ESD structures. The actual I/O circuits sensing or driving these pads were designed by the teams themselves. Figure 3 presents our I/O circuit examples: a self-biased inverter for receiving off-chip clocks or LO waveforms and an open drain PMOS device for driving off-chip instrumentation at high speeds.

As part of my effort toward making the tapeout class scalable, I envisioned that the 25 teams could benefit from regular interaction with engineers in the industry.

The 25 chips would need to be assembled into larger blocks before submission for fabrication. To this end, we created the pad frame in Figure 4, where ESD devices and some supply and ground pads are included and the remaining pads can serve as inputs, outputs, or dc lines.

The chips were then arranged in 4 × 4-mm blocks, as depicted in Figure 5. The second block contains additional test circuits developed by my Ph.D. students for characterizing the 130-nm process up to tens of gigahertz.

As can be seen in Figure 5, the spacing between different projects is not sufficient to allow dicing without damaging some circuits. In other words, one set of horizontal and vertical cuts would keep only some of the circuits intact, requiring a different set for rescuing the others. Since we expected to receive about 100 samples of each block, this issue was not deemed serious.

Summer

While on internships or engaged in research during the summer, the teams also pursued a few activities for the tapeout class. First, they created a test plan describing how exactly they would characterize their chips, what parameters they wished to measure, and what type of equipment they would require.

Second, they designed their PCBs based on a tutorial video produced by one of my Ph.D. students. As practiced in my research lab, I decided to have the chips directly

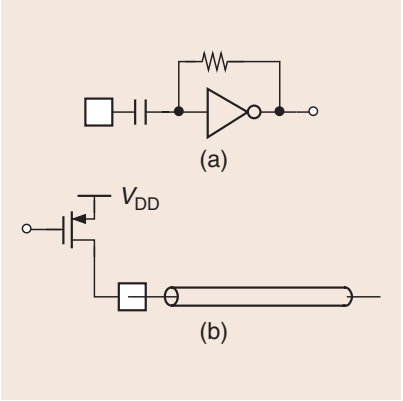

FIGURE 3: The (a) self-biased inverter for receiving off-chip clock or LO waveforms and (b) open-drain PMOS device for driving off-chip 50-Ω instrumentation.

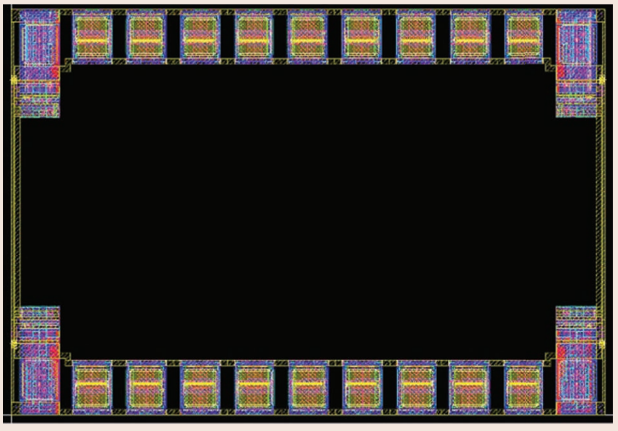

FIGURE 4: The pad frame used by all the projects.

bonded to the PCBs so as to avoid packaging and minimize parasitics. I held two PCB design reviews with each team to teach them about package and test issues. As part of this task, they also created a bonding diagram showing how the chip connected to the PCB and what off-chip components were necessary. Figure 6 gives examples of PCB design, a bonding diagram, and a test setup. The PCBs were submitted for fabrication in August and returned in September.

Test Equipment

In addition to the cost of chip fabrication, the test equipment also presents a heavy financial burden. After receiving the teams' test plans, I set out to create a lab infrastructure that could serve, at a given point in time, about half of the class. The plan was to spread the teams over several lab sessions per week.

Besides power supplies and moderate-speed oscilloscopes found in most labs, the necessary equipment fell into four categories:

1) "analog" input generators, i.e., instruments that deliver sinusoids for testing ADCs or serve as references for PLLs, as local oscillators or inputs for RF receivers, and as clocks for ADCs; the self-biased inverter of Figure 3(a) was used on each chip to convert the sine waves to square waves for use as clocks
2) digital input generators for characterization of DACs and digital multipliers
3) spectrum analyzers for capturing the analog outputs of PLLs, RF receivers, and DACs
4) logic analyzers or FPGAs for capturing the digital outputs of ADCs and digital multipliers.

Fortunately, the advent of low-cost USB-based instrumentation has provided myriad solutions. Table 1 summarizes examples that proved economically viable for our purpose.

Tens of other parts, such as subminiature version A (SMA) cables, bias tees, baluns, dc blocks, ESD protection straps, and tools necessary for lab work were also acquired.

Fabricated Chips

The chips returned from fabrication in SkyWater's 130-nm technology early October. Figure 7 provides die photographs of the two blocks. As a quick initial test of functionality, I used the probe station in my lab to check the operation of a ring oscillator and an *LC* oscillator. Both functioned as expected, indicating that fabrication had proceeded properly.

An Unforeseen Challenge

Our in-house facility diced the blocks and bonded each circuit to its respective PCB. The students were eager to embark upon testing, but soon, team after team reported that most of their pads appeared to be shorted to one another, a phenomenon that I had not encountered in my 35 years of chip development.

At first, given that probing the oscillators had confirmed proper operation, I surmised that the on-chip ESD protection devices might have been damaged during bonding. But I then examined a bonded chip closely (Figure 8), noting that parts of the bond wires had landed on the chip's passivation. Recalling that we had run a ground line and a supply line on the chip in parallel with the pads, I wondered whether the bond wires had broken through the passivation and shorted themselves to either of these lines. Ordinarily, I would have dismissed this possibility, as I knew from past experience that typical passivation is in fact similar to hard glass and difficult to break. However, when I attempted to scratch SkyWater's passivation

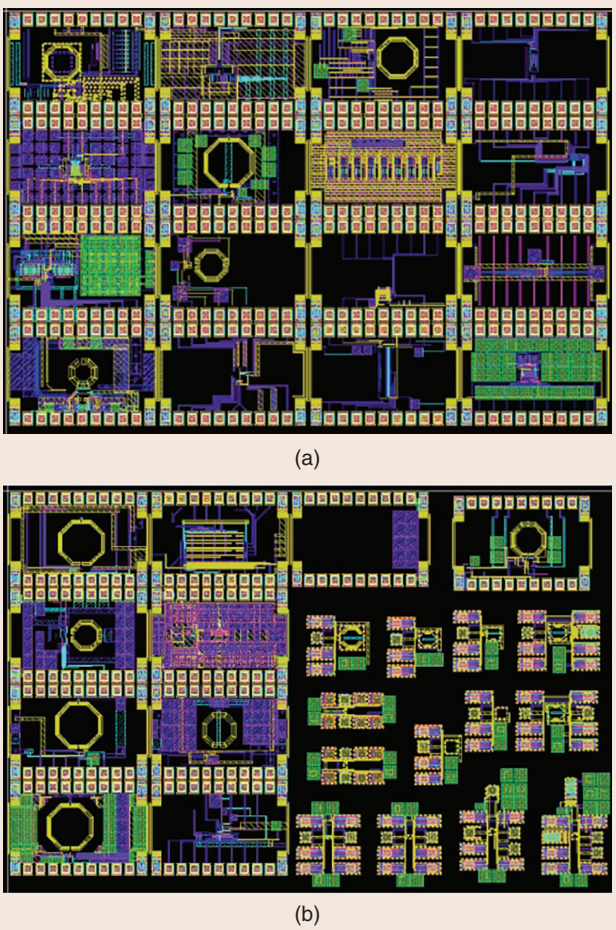

FIGURE 5: The two blocks containing the 25 projects along with other test circuits. (a) One block containing 16 projects, and (b) the other block containing nine projects and additional test circuits.

with a probe, I noticed that it was soft and brittle, a point confirmed by Efabless.

At this juncture, it appeared that we had no recourse, especially because it was not completely certain that the shorts occurred due to the penetration of bond wires through the passivation. To complicate matters further, we faced a race against time, as some of the students in the class planned to graduate soon and would need their grades. Nonetheless, I pursued two possible solutions. First, I decided to switch from our

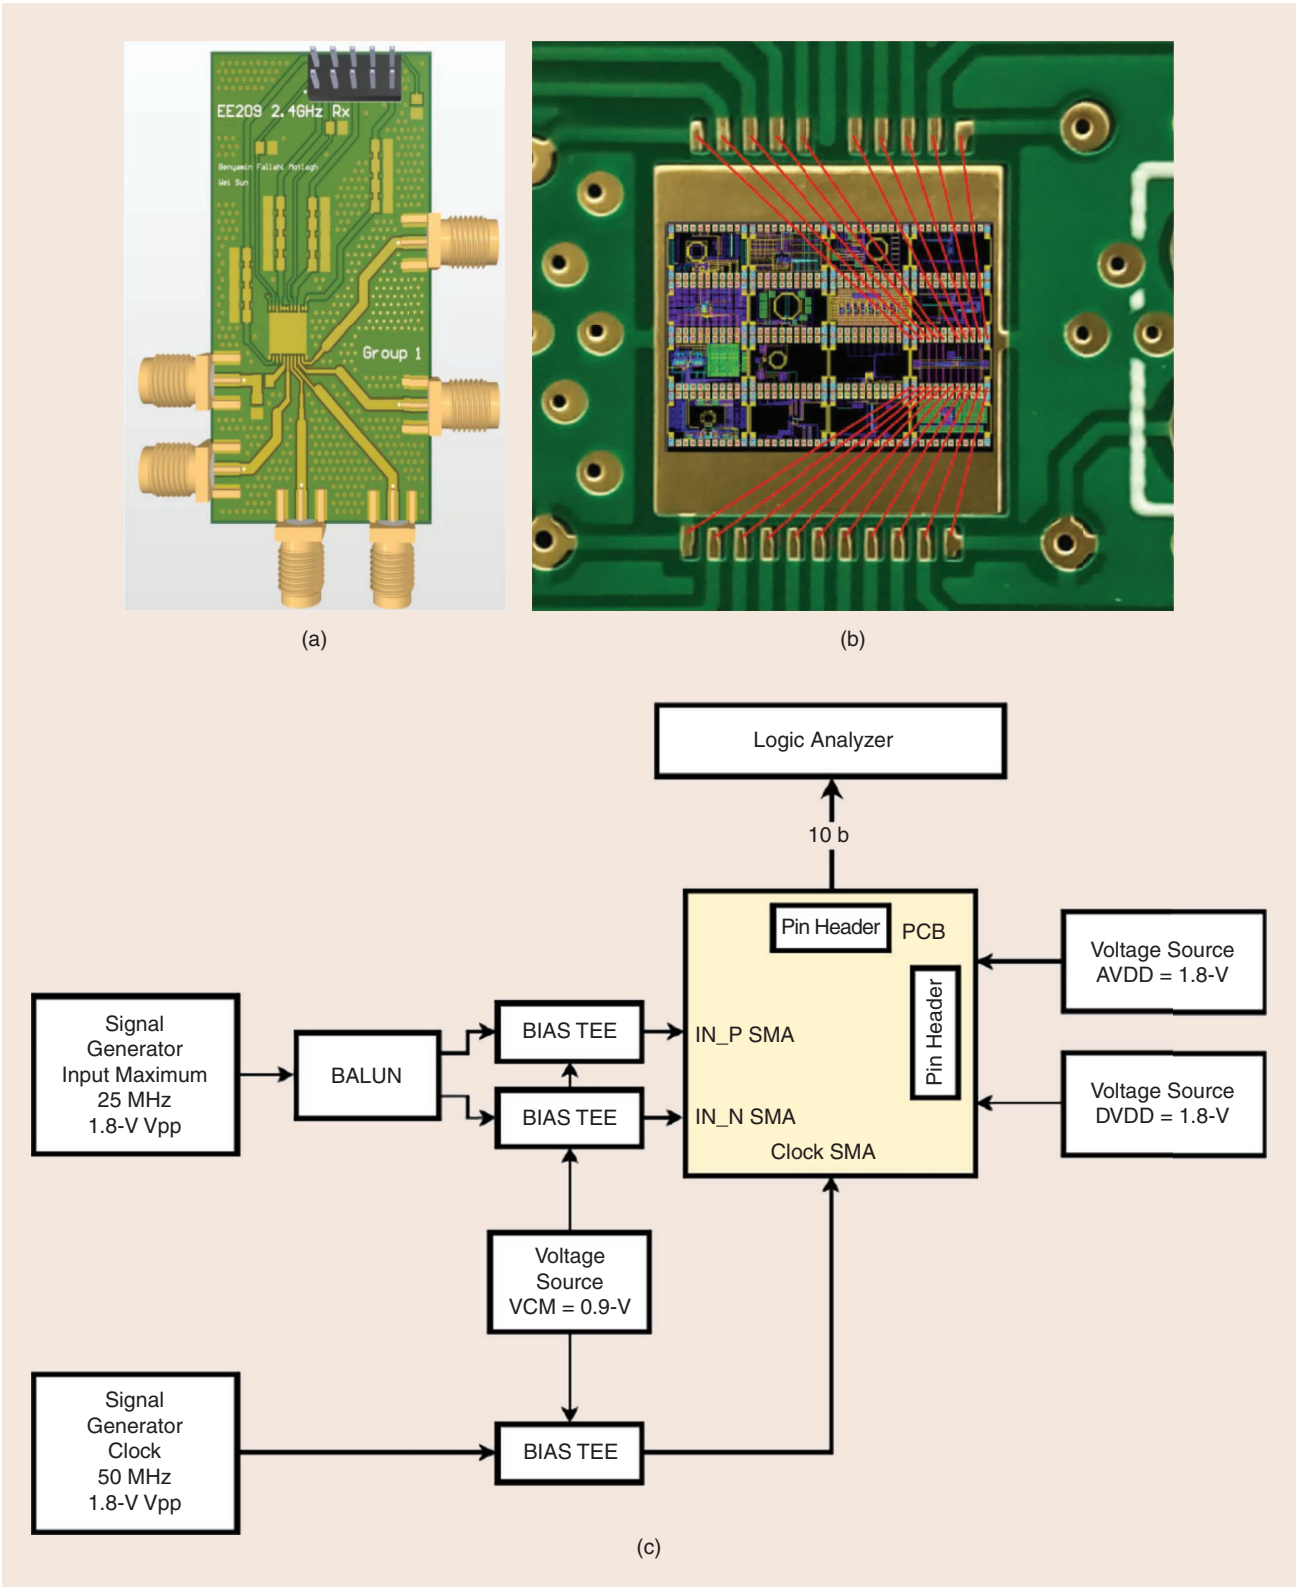

FIGURE 6: Examples of (a) PCB design, (b) a bonding diagram, and (c) a test plan. SMA: subminiature version A.

bonding method, called "wedge bonding," to another called "ball bonding." Depicted in Figure 9, the two differ in that ball bonding can confine the bond wire to the pad area, thus avoiding passivation breakage.

Such an endeavor would need to be delegated to external assembly houses, but it faced yet another difficulty. The exposed copper on a PCB is typically covered by tin or silver to avoid oxidation. However, ball bonding can attach wires only to a finish called "electroless nickel electroless palladium immersion gold" (ENEPIG). Our modification therefore required both another round of PCB fabrication with this finish and ball bonding by an assembly house. This process took about five weeks.

In parallel with this development, I also explored the possibility of testing the circuits on a probe station. The principal issue here was that the 25 chips had 25 different pinouts, disallowing any custom-designed probe cards. But we still ordered two probe cards with simple needles that would land on the two sides of the chip. This solution did in fact prove useful for one team that had made a fatal mistake in its PCB design.

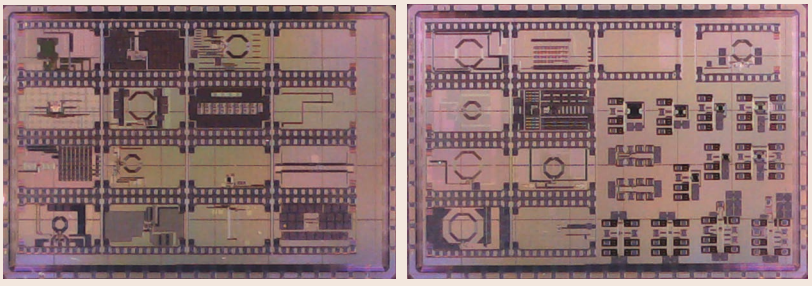

FIGURE 7: Die photographs of the two blocks. (a) One block containing 16 projects, and (b) the other block containing nine projects and additional test circuits.

TABLE 1. LIST OF EQUIPMENT

RF GENERATORS	FREQUENCY RANGE	UNIT PRICE	NUMBER OF UNITS
Windfreak SynthUSB3	12.5 MHz–6.4 GHz	US$350	20
Vaunix LMS-602D	1.5–6 GHz	US$2,000	1
LOGIC ANALYZER AND PATTERN GENERATORS	**MAXIMUM FREQUENCY**	**PRICE PER UNIT**	**NUMBER OF UNITS**
Digital Discovery logic analyzer	800 MS/s	US$325	2
Saleae Logic 8 logic analyzer	100 MS/s	US$500	2
Saleae Logic Pro 8 logic analyzer	500 MS/s	US$1,000	1
TUL 1M1-M0000127 FPGA	100 MS/s	US$130	10
DSLogic U3Pro32	250 MHz	US$400	4
SPECTRUM ANALYZERS	**FREQUENCY RANGE**	**UNIT PRICE**	**NUMBER OF UNITS**
SeeSii TinySA	100 kHz–5.3 GHz	US$210	10
Signal Hound SA124B	100 kHz–12.4 GHz	US$2,600	1
HIGH-SPEED TEST EQUIPMENT	**MAXIMUM FREQUENCY**	**UNIT PRICE**	**NUMBER OF UNITS**
Optellent BER tester	10 Gb/s	US$5,000	1
Agilent oscilloscope and sampling head	20 GHz	US$5,600	1

Third Course

With ball-bonded chips, testing moved briskly forward, and the teams began to see encouraging results for the last few weeks of the quarter. Interaction with individual groups allowed me to help them overcome testing difficulties (or mistakes). Of the 25 chips, 20 were tested successfully.

The other five partially or completely failed. An RF receiver did downconvert the 2.4-GHz input signal but provided no gain. A PLL did have proper VCO operation, but the second stage in its divider chain failed because of insufficient swings delivered by the first stage. A DAC did not work smoothly because its input register did not update properly. And another PLL's VCO simply did not oscillate, possibly due to the team's incompetence or negligence.

Discussion

Another scalability consideration relates to how the equipment and the lab can remain organized with 50 students frequenting the place. The issue was compounded by the fact that the lab space was also used by other courses at other times. To maintain order, I hired a lab assistant (who happened to be one of the students in the tapeout class) and created an equipment checkout form that the students needed to fill out before they could receive what they needed.

From the feedback that I received, the students found the tapeout class enormously useful, both as a learning experience and as a key stage in their careers.

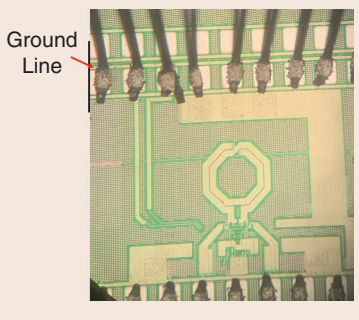

FIGURE 8: The problem of wedge bonding with soft passivation.

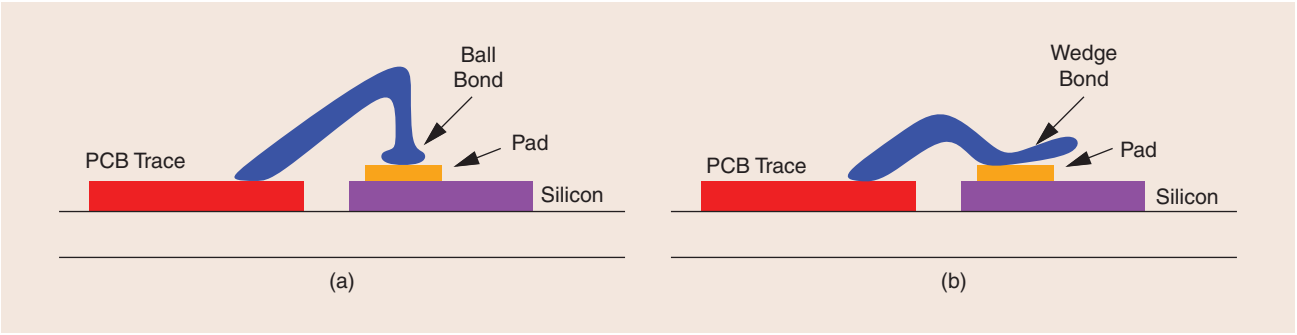

FIGURE 9: (a) Ball bonding and (b) wedge bonding.

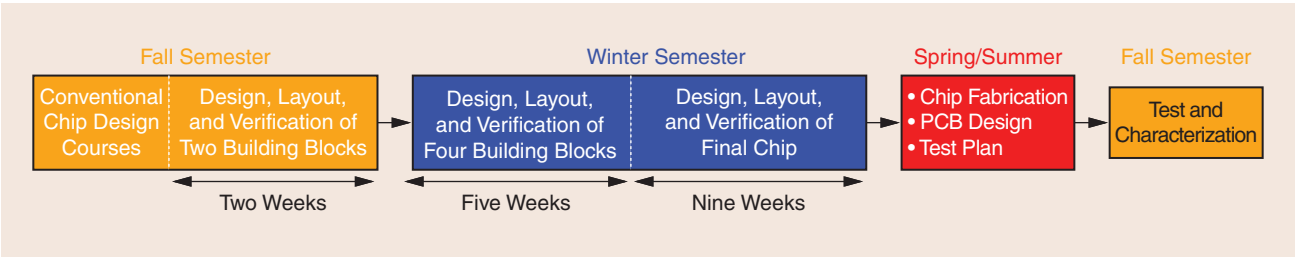

FIGURE 10: The proposed course structure for semester systems.

For a large tapeout class, one might recruit a teaching assistant as well. I did not have one. One concern is that it may be difficult to find a person who is familiar with the broad range of topics adopted for design in this course. Nonetheless, a teaching assistant can still help with tool setup and homework grading.

Extension to Semester Systems

The proposed blueprint can be applied to semester systems with some modifications. Appearing in Figure 10 is a possible plan. Assuming that the academic year begins with conventional chip design courses, one can allocate the last two weeks of the semester to the design, layout, and verification of two building blocks, e.g., a ring oscillator and a frequency divider. In the winter semester, another four blocks can be pursued while the students decide on their main project topic. This plan allows more time for chip fabrication.

Conclusion

Our first attempt at creating a large-scale tapeout class was fairly successful, motivating us to offer this course sequence once a year. If a dozen universities train students at this scale, we can deliver hundreds of chip designers to industry every year.

If a dozen universities train students at this scale, we can deliver hundreds of chip designers to industry every year.

Critical to such an endeavor is scalability, primarily in terms of 1) checking the students' building blocks in the first course, 2) closely monitoring the students' design and layout in the second course, and 3) lab equipment and organization in the third course. Recruiting industry mentors proves particularly helpful for the second matter.

The instructor does need to have an in-depth knowledge of the project topics as well as extensive experience in chip design and measurement. If left to their own devices, even highly intelligent students may develop faulty chips, thereby resenting their experience.

—Behzad Razavi

Acknowledgment

The author wishes to thank Ajith Amerasekera, Jeff DiCorpo, Tim Edwards, Matias Jara, Mohamed Kassem, Ken Mai, Boris Murmann, Borivoje Nikolic, Priyanka Raina, Arijit Raychowdhury, Jafar Savoj, Kshitiz Tyagi, Andrew Wright, and Jared Zerbe.

References

[1] P. Kinget, "Teaching IC design: From concepts to testing a fabricated custom chip," *IEEE Solid-State Circuits Mag.*, vol. 15, no. 3, pp. 87–93, Summer 2023, doi: 10.1109/MSSC.2023.3283976.

[2] D. C. Burnett et al., "Tapeout class: Taking students from schematic to silicon in one semester," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, May 2018, pp. 1–5, doi: 10.1109/ISCAS.2018.8351506.

[3] Y. Tsividis and D. A. Antoniadis, "A multiproject chip approach to the teaching of analog MOS LSI and VLSI," *IEEE Trans. Educ.*, vol. 25, no. 2, pp. 48–53, May 1982, doi: 10.1109/TE.1982.4321539.

[4] B. Murmann and M. Kassem, "Democratizing IC design: The story of a new movement and the launch of the SSCS PICO program," *IEEE Solid-State Circuits Mag.*, vol. 13, no. 2, pp. 123–130, Fall 2021, doi: 10.1109/MSSC.2021.3072273.