Scaling Limits to AI Chip Manufacturing

Anonymous Authors[1]

Abstract

We analyze potential limitations to the growth of AI chip manufacturing capacity in the next decade. We find that bottlenecks in human capital imply that manufacturing capacity alone is unlikely to result in a cumulative compute capacity of more than 1e36 FLOP by 2034, which reduces the likelihood of some near-term scenarios for transformative AI. Nevertheless, compared to current levels of production, a significant expansion of manufacturing capacity seems plausible.

1. Introduction

1.1. Context

Scaling of computational resources for model training runs has resulted in significant improvements in AI performance, but could run into limitations over the next decade. One potential limitation to scaling is availability of cutting-edge GPUs: data centers for AI training require large numbers of AI accelerators, such as the Nvidia H100 and H200. Recent years have already seen chip shortages, and AI companies continue to scramble for GPUs. Efforts to expand manufacturing capacity are ongoing, but due to the technologically complex, internationally distributed supply chain for semiconductors, this is not a trivial task. Although the expected future economic value from AI could motivate unprecedented levels of investment in compute supply, a capacity expansion could not take place arbitrarily fast: production facilities need time to be built, personnel to be trained, and some input factors may remain scarce even at high levels of investment.

We study which potential bottlenecks to GPU manufacturing might occur under an extreme scaling scenarios, and derive plausible upper bounds to production over the next decade. We find that especially availability of high-skilled workers and power-generation capacity would pose hard limits to manufacturing capacity expansion, even assuming unprecedented investments in the semiconductor industry and strong political buy-in. Despite these upper bounds, we find that manufacturing capacity could likely be scaled significantly compared to status quo production levels. Our findings have implications for the speed of AI capability development.

1.2. Methodology

We focus our analysis on the most cutting-edge AI chips, at techology nodes below 7nm. We start by assessing existing production capacity for these nodes, directly, based on TSMC fab capacity, and indirectly, based on cumulative EUV exposure capacity by ASML machines. We then briefly present a scaling scenario based industry forecasts, which does not take into account an extreme increase in investments, but extrapolates from past developments and current company announcements.

We then model potential production bottlenecks using supply-demand comparisons across a chosen set of input factors: STEM workforce, training capacity, raw materials, electricity, water, and suitable sites. For each input factor, we calculate how much of this input factor would be required if we wanted to build a very large number of semiconductor fabs, and then compare this to how much of this input factor the world could supply, given a premise of extreme spending. Details and citations for these calculations are included in the appendix.

Instead of aiming for precise predictions, we aim to derive a robust upper bound to the space of possibilities. For this purpose, we calculate bottlenecks using assumptions that are heavily biased in the direction of more chip production being possible. By compounding layers of these biased assumptions, we derive a ceiling for production capacity that is robust, but well above any plausible true value.

2. Scaling under business-as-usual

We first look at the current cutting-edge AI chip production capacity, and at projections of how this is expected to develop in the next few years based on recent trends. Our results serve as a point of reference ("baseline"), which helps contextualize our later findings about more extreme

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

scenarios. Our baseline is also useful as an indication how many more chips could be produced if all currently suitable fabs used their entire capacity to produce AI chips.

We focus on TSMC and ASML, as they have technological and economic advantages over their competitors, resulting in a monopoly for ASML, and a quasi-monopoly for TSMC. TSMC is currently the only company that is able to produce the most cutting-edge nodes at a high volume. ASML is currently the only company that produces EUV machines needed for the production of cutting-edge nodes. Growth projections for both companies are based on a mixture of proprietary and publicly available data. (SEMI, 2023)[1]

2.1. TSMC baseline

As of Q4/2022, TSMC's production capacity across all nodes has grown to about 1.2 M wafers per month (WPM), which corresponds to a compound annual growth rate (CAGR) of 13% between 1990 and 2022. Until 2031, the SEMI World Fab Forecast projects that capacity will reach 2.1 M WPM, at a CAGR of 6% between 2023 and 2031. After 2031, there is no projection available.

TSMC had 12 below-7nm fabs in full production during Q1/2023. Between mid-2023 and early 2029, a further 17 below-7nm fabs are scheduled to reach full production capacity. This includes fabs that have reached "first silicon", i.e. starting to slowly ramp up production and calibrate equipment, or are currently in the construction or equipment installation phases, and fabs that have been planned or announced.

Using capacity numbers for each fab, we can calculate that TSMC's <7nm capacity at the beginning of 2023 was 0.4 M WPM, and that TSMC's projected capacity in 2031 will be 1.0 M WPM, representing an increase in production capacity by 2.3x compared to 2023. If we use TSMC's annual <7nm CAGR since 2023 to extrapolate until 2034, we arrive at 1.4 M WPM, which is an increase of 3.5x compared to 2023.

Most of the new 0.6M WPM <7nm wafer capacity planned from 2023 until 2031 will be for 3nm and 2nm nodes. Some additional capacity for 5nm and 6nm will be added, and production of <2nm nodes will slowly be rolled out.

2.2. ASML baseline

ASML's EUV systems are the most expensive single piece of equipment in a fab, tend to be the bottleneck of individual fabs, and are very complex to manufacture. Thus the supply of EUV machines could constrain the overall wafer production capacity.

[1]Sources for ASML numbers include (ASML, 2022a;c; Shilov, 2020; Cutress, 2020; Thoss, 2019; ASML, 2015; Verheyde, 2019; ASML, 2024; 2018; 2021b;a; 2022b)

As of Q3/2023, our estimates suggest that ASML has shipped 261 EUV lithography systems between 2013 and 2023. This includes all systems between the NXE:3300B and the latest-generation NXE:3600E. The numbers until 2023 are fairly reliable, as they are partly based on, or closely match, publicly available information from ASML.

ASML predicts in its Annual Report 2022 that 90 EUV systems will be shipped in 2025. Taking this prediction at face value, this would imply a CAGR of EUV systems of 21 % between 2020 and 2025. Using this CAGR to extrapolate growth suggests that 820 EUV systems will have been shipped by 2028, of which 44 will be in the high-NA category.

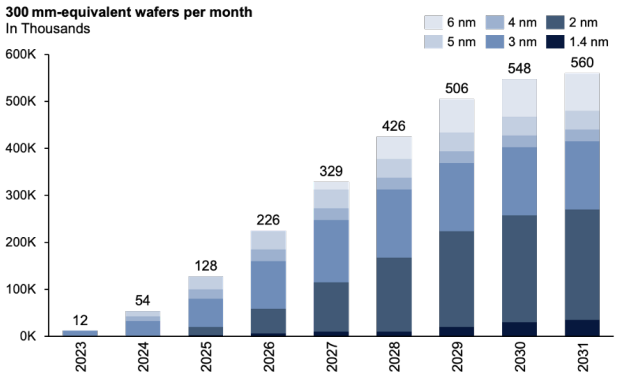

Figure 1. TSMC's total <7nm capacity per node until 2031 according to the SEMI World Fab Forecast.

Further projection until 2034 would be possible, but since we're only using CAGRs, and cannot rely on company announcements or high-fidelity inside information, this would be increasingly crude. Even the projection until 2028 is subject to significant uncertainty, especially when considering the cyclicality of the semiconductor market.

Using the cumulative number of shipped machines, we can calculate the worldwide maximum EUV lithography capacity in EUV-exposed wafers per month. For this, we multiply the number of machines with throughput per machine, measured in wafers per hour (WPH). The resulting capacity represents the theoretical maximum, as EUV machines usually don't run at full capacity 24/7.

Next, we can convert the maximum number of EUV exposures per hour into the number of completed wafers per month this would yield for a given technology node. This number depends on the number of EUV exposures necessary for that node. For this illustration, we will use TSMC's N3E process as a benchmark, as it is known to have 19 single-exposure EUV layers.

From these calculations, we conclude that by the end of 2023, the worldwide cumulative EUV capacity will be 33.2

M EUV exposures per month. Assuming the entire capacity is used for manufacturing TSMC's N3E node, this would result in a manufacturing capacity of 1.7 M wafers per month, assuming that all machines run 24/7 at 100% utilization.

Extrapolating EUV shipments and throughput until 2028 using CAGRs, we arrive at 153 M EUV exposures per month in 2028, of which 9 M from High-NA EUV. If we convert this to TSMC N3E-equivalent wafers, this would amount to 8.1 M wafers per month, but it is unclear how informative this benchmark is in 2028.

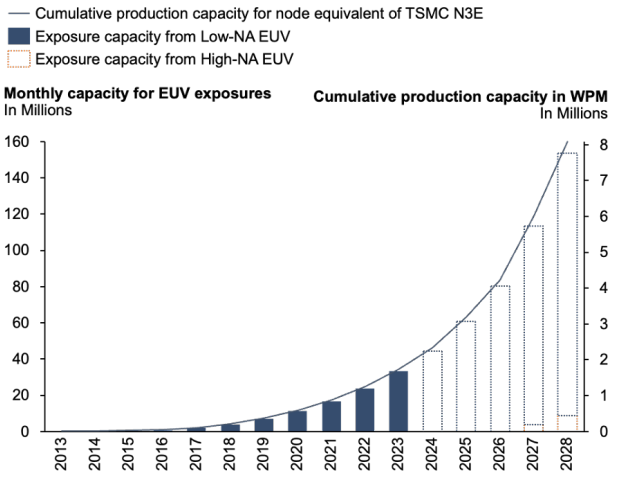

Figure 2. Cumulative monthly capacity for EUV exposures and corresponding production capacity in wafers per month (dashed lines indicate projections)

3. Potential Scaling Limits

3.1. Hypothetical scaling scenario

We explore the implications of a hypothetical scenario, with very large capital expenditures on chip production. For the sake of argument, we posit that this amount is equal to USD 105 T each year between 2024 and 2034, which is approximately equal to 2023 gross world product. Given this premise, we calculate how many fabs this could buy per year and which factors could become bottlenecks if this amount of fabs was built. This amount is chosen arbitrarily, to represent a large amount as a basis for calculations. It is not meant to be realistic. One of the newest 'greenfield' TSMC fab in Taiwan, "Fab 18 Phase 9" (3nm, 40kWPM), is projected to cost around USD 13.38 B This is the cheapest relevant fab we could find. Assuming a price of USD 13.38 B, the equivalent of 2023 GWP could buy 7,800 fabs.

3.2. Potential bottlenecks

3.2.1. Training capacity

A cutting-edge fab, requiring at least 2,400 highly skilled STEM professionals, faces significant expansion limits due to training bottlenecks. Assuming each worker needs at least two months of on-the-job training that cannot be expedited, and that each fab can train at most 33% of its existing STEM workforce annually (approximately 4,750 trainees for a fab with 2,400 relevant staff), growth is constrained. Even with fabs being constructible in two years and usable for training immediately upon completion (with training occurring during construction, potentially leading to a 22-month wait for the first cohort), the cumulative number of fabs is projected to be under 80 in 2026, below 440 by 2027, and less than 2,400 in 2028, starting from 20 available <7nm fabs in 2024. These optimistic projections hinge on on-job training not becoming massively scalable in the next decade, implying that limited training capacity will restrict fab construction to an average of less than 600 new fabs annually within the first six years of scaling.

3.2.2. STEM workforce

Constructing approximately 7,800 cutting-edge semiconductor fabs annually would necessitate recruiting around 18.2 million high-skilled STEM professionals each year. This figure is roughly 2.7 times the entire 2021 US semiconductor-relevant STEM workforce (6.7 million), nine times the estimated annual number of relevant STEM graduates worldwide (2 million), and about 90 times the number of such graduates in the US (200,000). These comparisons suggest that staffing such a large number of fabs per year until 2034 will be exceptionally difficult, even assuming each fab requires 2,400 professionals with degrees in fields like electrical engineering, computer science, chemistry, or physics, and that these roles have perfect background interchangeability. While the OECD produced approximately 1 million relevant graduates in 2021, the global limit for such specialized personnel makes this level of expansion highly challenging.

3.2.3. Power

Building approximately 7,800 cutting-edge fabs annually, each requiring 350 MW of power, would necessitate a yearly addition of electricity generation capacity (23,915 TWh) roughly equivalent to the entire world's electricity production in 2020. Although it's feasible to construct power plants alongside fabs (at an estimated additional cost of USD 0.4 billion per fab, or 3% of capital expenditure, and with gas plants being less complex than fabs), this would require an unprecedented rate of power plant construction, far exceeding any historical precedent of energy infrastructure expansion, even if the theoretical energy supply is not a con-

straint. This immense demand could not be met by existing grid capacity and would require a historically unmatched effort in building new power generation facilities.

3.2.4. Water

While constructing approximately 7,800 cutting-edge semiconductor fabs annually would lead to significant water consumption, equivalent to the usage of about 121 million average US households, this is not considered a relevant physical limit. A single fab (40,000 WPM with recycling) uses about 6.4 million cubic meters of water annually, comparable to 15,500 US households, with only 20% permanently lost. Even if building near ample fresh water isn't always possible, the substantial capital expenditure for fabs makes it feasible to incorporate large-scale water recycling facilities, desalination plants, and water pipelines (estimated at less than USD 0.5 billion per fab), making water scarcity a relatively inconsequential bottleneck compared to other factors.

3.2.5. Raw materials

The construction of semiconductor fabs, even at a massive scale, is not constrained by a relevant physical limit due to raw material shortages because all necessary elements are sufficiently abundant on Earth and are generally mined in adequate quantities. Even if chip production were to require unexpectedly large amounts of very rare elements—a scenario deemed unlikely as modern IC production uses a wide range of elements, with realistic estimates suggesting only about 0.015g of total material deposited per wafer, mostly non-rare—the resulting production limits would still be much higher than those imposed by other factors. Grossly overestimating demand by assuming each element constitutes 1% of a wafer's weight still shows that required amounts for most materials are small compared to global production, and reserves for even potentially scarce rare elements like Ruthenium and Iridium are virtually endless. In the improbable event of a raw material shortage, expanding supply by establishing new mines and refineries would be practical alongside the construction of thousands of fabs.

3.2.6. Suitable sites

The availability of physical space does not pose a relevant limit to semiconductor fab construction, as there are ample suitable building sites for hundreds of very large industrial zones, even under extremely conservative assumptions. These sites, typically required near population centers with well-educated talent pools and transport hubs, can accommodate the necessary infrastructure like roads, transmission lines, pipelines, and seismic isolation platforms. Even a conservative estimate considering OECD population centers with over 5 million inhabitants suggests enough space for approximately 6,700 fabs, with the true number of suitable sites likely being one to two orders of magnitude higher. Furthermore, megasite preparation and infrastructure development are relatively low-complexity and low-cost processes compared to the construction of the semiconductor fabs themselves, meaning space and site preparation do not present a bottleneck as significant as other factors.

3.3. Plausible Upper Bound until 2034

The previous section suggests that only (1) on-job training capacity and (2) availability of skilled STEM workers come close to being meaningful bottlenecks within the set of domains we investigated. Accepting the assumptions behind these bottlenecks, we derive an upper bound to annual production capacity within the next 10 years. **At most 24,000 fabs with 40kWPM capacity are likely to be operational in 2034, translating into a cumulative production capacity of 960 M WPM. Across 10 years, this would imply a cumulative compute capacity of < 1e36 FLOP.**

4. Potential Policy-relevant Implications

4.1. Chip manufacturing capacity is limited, but could be accelerated significantly

Accelerating compute growth through increases in manufacturing capacity alone will run into limitations. Assuming current levels of chip production efficiency will persist, the cumulative compute budget achieved through this method likely be lower than 1e36 FLOP after 10 years. The available budget for a 1-year training run will be even lower. This rules out some scenarios for transformative AI by 2034. (Ho, 2022) However, current trends in chip manufacturing capacity may not accurately predict future developments, as there appears to be substantial room for accelerated growth within the next decade, with the precise rate depending on the strength of incentives. If the perceived economic value of AI grows even further, many apparent bottlenecks on the current margin of investment could likely be addressed.

4.2. Human capital will be a particularly binding constraint to a rapid capacity expansion

Ambitious scale-ups of the chip manufacturing capacity require vast amounts of skilled labour. If scaling of manufacturing capacity remains a key lever for expanding compute supply, and in the absence of significant automation, countries with large pools of skilled labour, strong STEM education, efficient training programs, and attractive immigration policies will have an advantage. With significant automation of semiconductor manufacturing, countries with the best access to such automation would be advantaged. The degree to which AI allows for semiconductor manufacturing automation hence deserves more research attention.

Impact Statement

This paper presents work whose goal is to advance the field of techncial AI governance. One might argue that this paper contributes to an acceleration of AI capabilities, by analyzing bottlenecks in the compute supply. This might be a cause for concern for some. However, given the strong commercial incentives for acceleration, we doubt that this paper will have a strong counterfactual effect.

References

Arizona Technology Council. Tsmc announced its arizona microchip factory 3 years ago. here's what has happened since then. *Arizona Technology Council Blog*, May 2023. URL https://www.aztechcouncil.org/tsmc-announced-its-arizona-microchip-factory-3-years-ago-heres-what-has-happened-since-then/.

Asianometry. The insane engineering of the ps5, December 2020. URL https://www.youtube.com/watch?v=jlAWx_X5opA.

ASML. Asml posts record full-year 2014 sales of €5.86 billion. https://www.asml.com/en/news/press-releases/2015/asml-posts-record-full-year-2014-sales, January 2015.

ASML. Euv products and business opportunity. https://www.asml.com/-/media/asml/files/investors/investor-days/2018/euv-products-and-business-opportunity-christophe-foucquet.pdf, 2018.

ASML. Asml unveils euv technology training center in taiwan, July 2020. URL https://www.asml.com/en/news/press-releases/2020/asml-unveils-euv-technology-training-center-in-taiwan.

ASML. Asml euv products and business opportunity (investor day 2021 materials). https://www.asml.com/en/investors/investor-days, 2021a.

ASML. Asml euv products and business opportunity (investor day 2021 materials). https://www.asml.com/en/investors/investor-days, 2021b.

ASML. Asml annual report 2022. https://www.asml.com/en/investors/financial-results/annual-reports, 2022a.

ASML. Asml company strategy (investor day 2022 materials). https://www.asml.com/en/investors/investor-days, 2022b.

ASML. Asml investor day 2022 presentation materials. https://www.asml.com/en/investors/investor-days, 2022c.

ASML. Twinscan nxe:3400b - euv lithography systems. https://www.asml.com/en/products/euv-lithography-systems/twinscan-nxe3400b, 2024.

ASML. Financial overview, 2025. URL https://www.asml.com/en/investors/financial-overview.

Chang, C.-c. (tsmc's "newcomer training center" to create semiconductor talent, has trained more than 10,000 new recruits for the industry). *Yahoo News Taiwan / (Central News Agency)*, November 2023. URL https://tw.news.yahoo.com/%E5%8F%B0%E7%A9%8D%E9%9B%BB-%E6%96%B0%E4%BA%BA%E8%A8%93%E7%B7%B4%E4%B8%AD%E5%BF%83-%E6%89%93%E9%80%A0%E5%8D%8A%E5%B0%8E%E9%AB%94%E4%BA%BA%E6%89%8D-%E5%B7%B2%E5%9F%B9%E8%82%B2%E9%80%BE%E8%90%AC%E5%90%8D%E7%94%A2%E6%A5%AD%E7%94%9F%E5%8A%9B%E8%BB%8D-091614184.html.

CNA (Central News Agency). ntc (tsmc fab service center, taichung ntc opens). *CNA*, November 2021. URL https://www.cna.com.tw/news/afe/202111240314.aspx.

Cutress, I. Tsmc: We have 50% of all euv installations, 60% wafer capacity. https://www.anandtech.com/show/16042/tsmc-we-have-50-of-all-euv-installations-60-wafer-capacity, August 2020.

Data USA. Data usa, 2025. URL https://datausa.io/.

Dean, G. Chip maker tsmc needs to hire 4,500 workers for its new arizona plants. here's what it's looking for. *Yahoo Finance (via Business Insider)*, December 2022. URL https://finance.yahoo.com/news/chip-maker-tsmc-needs-hire-100000012.html.

fernstudieren.de. Abbrecherquote im fernstudium (dropout rate in distance learning), 2025. URL https://www.fernstudieren.de/haeufige-fragen/abbrecherquote-im-fernstudium/.

halbleiter.org. Chipfertigung (chip manufacturing), 2025. URL https://www.halbleiter.org/chipfertigung/.

Ho, A. Grokking "forecasting tai with biological anchors", 2022. URL https://epoch.ai/blog/grokking-bioanchors.

Intel. New mexico undergrad entry level process engineer job, 2025a. URL https://jobs.intel.com/en/job/albuquerque/new-mexico-undergrad-entry-level-process-engineer/41147/46684046176.

Intel. What does it take to build a fab? Technical report, 2025b. URL https://www.intel.com/content/dam/www/central-libraries/us/en/documents/what-does-it-take-to-build-a-fab.pdf.

Intel Newsroom. Technician training with ar/vr, March 2022. URL https://www.youtube.com/watch?v=mBKuUbihY8g.

International Monetary Fund. World economic outlook database, April 2025. URL https://www.imf.org/external/datamapper/NGDP_RPCH@WEO/OEMDC/ADVEC/WEOWORLD.

Jennings, R. and Söderberg, J. Tsmc prepares american engineers for its first u.s. fab. *AmCham Taiwan Topics*, August 2022. URL https://topics.amcham.com.tw/2022/08/tsmc-prepares-american-engineers-for-its-first-u-s-fab/.

Kerkmann, C. and Ströder, M. Arbeiten bei intel: So will der chiphersteller 3000 stellen in magdeburg besetzen (working at intel: This is how the chip manufacturer wants to fill 3000 positions in magdeburg). *Handelsblatt*, August 2023. URL https://www.handelsblatt.com/technik/it-internet/arbeiten-bei-intel-so-will-der-chiphersteller-3000-stellen-in-magdeburg-besetzen/29362636.html.

Khan, S. M., Mann, A., and Peterson, D. The chipmakers: U.s. strengths and priorities for semiconductor supply chains. Technical report, Center for Security and Emerging Technology (CSET), August 2020. URL https://cset.georgetown.edu/wp-content/uploads/CSET-The-Chipmakers.pdf.

Lee, Y. and Blanchard, B. Taiwan invests in next-generation talent with slew of chip schools. *Reuters*, March 2022. URL https://www.reuters.com/markets/funds/taiwan-invests-next-generation-talent-with-slew-chip-schools-2022-03-10/.

Liu, J. and Mozur, P. Tensions rise at tsmc's arizona factory over work culture clash. *The New York Times*, February 2023. URL https://www.nytimes.com/2023/02/22/technology/tsmc-arizona-factory-tensions.html.

Nam, H.-w. Korea to train 150,000 semiconductor workers in next decade. *The Korea Times*, July 2022. URL https://www.koreatimes.co.kr/www/tech/2023/08/129_333011.html.

NASA. Earth's energy balance. Fact Sheet FS-2008-03-042-GSFC(REV 1), March 2008. URL https://www.nasa.gov/pdf/135642main_balance_trifold21.pdf.

National Center for Science and Engineering Statistics (NCSES). International s&e higher education, 2019. URL https://ncses.nsf.gov/pubs/nsb20197/international-s-e-higher-education#u-s-position-in-global-s-e-higher-education.

National Science Foundation (NSF). Science and engineering indicators 2018 - appendix tables, 2018. URL https://www.nsf.gov/statistics/2018/nsb20181/data/appendix.

Phillips, N. How long will global uranium deposits last? *Scientific American*, October 2009.

PTS News Network. (tsmc's u.s. plant progress delayed, arizona union worries about impact on u.s. chip bill). *PTS News Network*, November 2022. URL https://news.pts.org.tw/article/609505.

Purdue University College of Engineering. Semiconductor degrees program, 2025. URL https://engineering.purdue.edu/semiconductors/degrees.

Ritchie, H., Roser, M., and Mathieu, E. Land use, March 2024.

SEMI. World fab forecast. https://www.semi.org/en/products-services/market-data/world-fab-forecast, 2023. URL https://www.semi.org/en/products-services/market-data/world-fab-forecast. Accessed: 2025-05-12.

Semiconductor Industry Association (SIA). The growing economic impact of the u.s. semiconductor industry. Technical report, May 2021. URL https://www.semiconductors.org/wp-content/uploads/2021/05/SIA-Impact_May2021-FINAL-May-19-2021_2.pdf.

Semiconductor Industry Association (SIA). U.s. semiconductor ecosystem map, 2025. URL https://www.semiconductors.org/u-s-semiconductor-ecosystem-map/.

Sharma, Y. Semiconductor talent shortage - a crisis or an opportunity? *University World News*, September 2022.

URL https://www.universityworldnews.com/post.php?story=20220903074733791.

Shepardson, D. and Alper, A. Biden to visit taiwan's tsmc chip plant in arizona, hail supply chain fixes. *Reuters*, December 2022. URL https://www.reuters.com/technology/biden-visit-taiwans-tsmc-chip-plant-arizona-hail-supply-chain-fixes-2022-12-06/.

Shilov, A. Asml ramps up euv scanners production: 35 in 2020, up to 50 in 2021. https://www.anandtech.com/show/15428/asml-ramps-up-euv-scanners-production-35-in-2020-45-50-in-2021, January 2020.

Taylor, N. A deeper dive into semiconductor fab requirements. *Deep Forest by Nathan Taylor (Substack)*, November 2023. URL https://deepforest.substack.com/p/a-deeper-dive-into-semiconductor.

Teague, A. Time, money and patience: The process of developing a megasite. *Cardinal News*, May 2023. URL https://cardinalnews.org/2023/05/01/time-money-and-patience-the-process-of-developing-a-megasite/.

Terrill, M. Arizona legislators get up-close look at asu semiconductor facility. *ASU News*, March 2023. URL https://news.asu.edu/20230312-arizona-impact-arizona-legislators-get-upclose-look-asu-semiconductorfacility.

The Washington Post. Amid drought, arizona to study project that pipes in water from mexico. *The Washington Post*, 2025.

Thoss, A. How does the laser technology in euv lithography work? https://www.laserfocusworld.com/blogs/article/14039015/how-does-the-laser-technology-in-euv-lithography-work, August 2019.

TSMC. Tsmc 2022 sustainability report, 2022a. URL https://esg.tsmc.com/download/file/2022_sustainabilityReport/english/e-all.pdf.

TSMC. Tsmcnewcomer training with ar, November 2022b. URL https://www.youtube.com/watch?v=xUw7MygrsWE.

TSMC. Careers at tsmc, 2025. URL https://www.tsmc.com/static/english/careers/index.htm.

TSMC Careers. Supporting entry-level engineers, 2025a. URL https://tsmccareers.com/supporting-entry-level-engineers/.

TSMC Careers. Entry-level engineer training, 2025b. URL https://tsmccareers.com/entry-level-engineer-training/.

TSMC ESG. Newcomer training program to build a quality workforce, 2025. URL https://esg.tsmc.com/en/update/innovationAndService/caseStudy/44/index.html.

U.S. Bureau of Labor Statistics (BLS). Occupational employment statistics - naics 334400 semiconductor and other electronic component manufacturing, 2025. URL https://www.bls.gov/oes/current/naics4_334400.htm.

U.S. Energy Information Administration (EIA). How much carbon dioxide is produced when different fuels are burned?, September 2023. URL https://www.eia.gov/tools/faqs/faq.php?id=73&t=11.

Vanek, C. Tsmc starts training phoenix workers for arizona chip plant. *AZ Central*, May 2021. URL https://eu.azcentral.com/story/money/business/jobs/2021/05/12/taiwan-semiconductor-manufacturing-co-starts-training-phoenix-workers/5029771001/.

Verheyde, A. Amsl: 4.5 million cumulative euv wafers processed through 2018. https://www.tomshardware.com/news/amsl-45-million-cumulative-euv-wafers-processed-through-2018, December 2019.

Wealth Magazine Express. 6300 (tsmc newcomer training center exposed, more than 6,300 newcomers cultivated). *Yahoo News Taiwan*, November 2022. URL https://tw.stock.yahoo.com/news/%E5%8F%B0%E7%A9%8D%E9%9B%BB%E6%96%B0%E4%BA%BA%E8%A8%93%E7%B7%B4%E4%B8%AD%E5%BF%83%E6%9B%9D%E5%85%89-%E5%B7%B2%E5%9F%B9%E8%82%B2%E9%80%BE6300%E5%90%8D%E7%94%9F%E5%8A%9B%E8%BB%8D-071704811.html.

Wei, W.-S. (chip talent shortage! how does tsmc create a "newcomer training center" to solve the pain points of its own supply chain?). *(Manager Today)*, December 2022. URL https://www.managertoday.com.tw/articles/view/66014.

Wikipedia contributors. List of united states public university campuses by enrollment, 2025. URL https://en.wikipedia.org/wiki/List_of_United_States_public_university_campuses_by_enrollment.

Y, J. The big semiconductor water problem, 2025.

Yahoo Finance. Asml holding n.v. (asml.as) balance sheet, 2025. URL https://uk.finance.yahoo.com/quote/ASML.AS/balance-sheet?p=ASML.AS.

Yoon, J.-y. Chip industry talent plan sees s korea designate 20 universities. *University World News*, October 2021. URL https://www.universityworldnews.com/post.php?story=20211021144726611.

A. Details on Bottleneck Calculations

A.1. Training capacity

A.1.1. Summary

Professional training requirements for fab engineers at TSMC's Arizona fab suggest that there is a limit to how fast a firm like TMSC can expand its workforce. This is because capacity for on-job training can only expand at the rate at which fabs are built, which in turn depends on the speed at which suitable personnel can be trained.

Practical training, i.e. the type of training that takes place in a fab-like training center, is bottlenecked by the availability of sufficiently experienced instructors, although it seems possible that instruction capacity could be multiplied using VR. Theoretical training, i.e. the type of training that is focused on transmitting scientific knowledge on semiconductor production, could probably be massively parallelized through well-designed online training, which makes it unlikely to be a bottleneck.

This analysis purely considers the training requirements for engineers at TSMC, and does not take into account training of other types of employees or training at suppliers. Moreover, the potential on-job bottleneck relies on the assumption that some minimum amount of training time needs to be spent at the fab.

Production bottlenecks	Number of fabs possible per year	Key underlying assumptions	Comments
On-job training capacity	Cumulatively, at most ∼80 until 2026, ∼440 until 2027, and ∼2400 until 2028	There is no alternative to training new workers directly at fabs. Only two months of training is required for one engineer, fabs have large training capacity. Training is taking place continuously and without friction.	This seems to be a bottleneck even under pretty optimistic assumptions, e.g. training takes place continuously, without friction, at high throughput, there is no shortage of instructors.

Table 1. Production bottlenecks to annual construction of fabs resulting from limited workforce training capacity.

A.1.2. Approach

On the demand side, we focus on professional training requirements for fab engineers at TSMC's Arizona fab. By this we refer to training that is specific to semiconductor manufacturing, and goes beyond general engineering training at universities. We ignore training requirements of other fab employees, as no fine-grained data is available on this. We further assume that the labor requirements for each additional fab stay constant, assuming that no further automation - labor augmenting or substituting - takes place. On the supply side we extrapolate how much training TSMC (or TSMC-like companies) could provide if training capacity was scaled up massively. Within training we focus on three different types of training: (1) theoretical training, (2) practical training, and (3) on-job training.

We focus exclusively on the training requirements for skilled engineers, thereby neglecting the needs for other highly-skilled non-engineering workers. This narrow focus likely results in an overestimation of the number of fabs that could realistically be established. Furthermore, some uncertainty about the extent to which instruction capacity could be expanded, whether through dedicated training programs for instructors or the implementation of technologies like Virtual Reality (VR).

A.1.3. Demand

Before new engineering recruits can start working in the clean rooms, they need to undergo a few weeks to months of professional training. This is necessary to make them hit the ground running when they start taking over responsibility in the fab. Across different roles, at least three components of training are necessary:

- **Theoretical training:** getting recruits up to speed from a general engineering degree to specific knowledge about semiconductor manufacturing. For example: TSMC offers semiconductor programs at universities in Taiwan. These programs offer majors on components/integration, processes/modules, equipment engineering, advanced packaging, intelligent manufacturing, and advanced circuit design. In 2022, TSMC claimed that, since 2019, around 4,000 students had enrolled in these programs across universities in Taiwan. (TSMC, 2022a)
- **Practical training:** training recruits in a realistic fab-like environment and exposing them to real equipment. This takes place within a training center, to shorten training time in fabs and to avoid downtime and wear-out of the fab equipment.

For example: TSMC has a Newcomer Training Center in Taichung Science park in Taiwan, where new engineering recruits get 2 months of hands-on training with real equipment. (PTS News Network, 2022; Wealth Magazine Express, 2022)

- **On-job training:** bringing recruits to the fab to get them acquainted with the actual processes in the fab and teaching them on-site processes in a clean-room environment. For example: engineering graduates at TSMC's Kumamoto fab undergo 2-3 months of on-job training in a fab. (TSMC Careers, 2025a) Process engineering recruits at Intel's Albuquerque fab need to spend an estimated 6 months on-site during training. (Intel, 2025a)

Depending on the exact role, both training time and training content within each component will likely vary. An R&D engineer may spend more time on theory than an equipment engineer, whereas the latter may spend more time with machines.

How many fab engineers need to be trained?
TSMC claims that 3,000 of its 4,500 planned Arizona roles will require 4-year engineering degrees.(Terrill, 2023) The actual number of engineers may be higher, but in absence of better data, we will use this number and accept a potential overestimation of possible fabs. We then scale this down proportionally to 2,400 to account for a fab with a capacity of 40,000 WPM. Consequently, we will assume that building 10 fabs with 40 kWPM will require ∼24,000 engineers, building 100 fabs will require ∼240,000 engineers, and so on.

A.1.4. SUPPLY

Theoretical training Training tens or even hundreds of thousands of additional people in-person in universities would require a massive buildup of educational infrastructure. One way to provide this would be to build large higher education institutions that are entirely dedicated to semiconductor manufacturing. This is what a number of countries currently try to do, for example South Korea and Taiwan.

- Taiwan has opened four chip schools, each with capacity for 100 masters and PhD students. The schools are open year-round without breaks. (Lee & Blanchard, 2022) Around 4,800 students will benefit (Yoon, 2021) from the new Taiwanese chip schools (Sharma, 2022), at an investment of USD 340 M over 12 years.

- South Korea announced in 2022 that it wants to designate 20 universities as "specializing in semiconductors by 2026, with the goal of training 150,000 semiconductor workers over the next decade." (Wikipedia contributors, 2025)

Training 300,000 engineers each year could be achieved by building four semiconductor schools of size equivalent to Texas A&M university, which is the largest university in the US at ∼75,000 students. (Wikipedia contributors, 2025) Purely in terms of infrastructure, it seems possible to build 4 large universities within 10 years. The question is whether they can be built fast enough and whether enough qualified instructors can be trained within a few years. Whether this is possible or not, this approach does not seem like the most efficient way to train lots of engineers in a very short period of time. Instead, it seems more scalable to train people entirely through MOOCs, for example including live online instruction and online exams. This approach could massively speed up knowledge transmission, but it may come at a cost: close mentorship from experienced engineers may not be available and motivation to study may be lower. (fernstudieren.de, 2025) Still, at least for the theoretical part, it seems unlikely that these factors would be a hard bottleneck for scaling. Distance learning may be demotivating and may not provide the same degree of teacher-student interaction as an in-person degree, but it seems plausible that a well-designed MOOC, combined with strong intrinsic and extrinsic motivation for completion, would be sufficient to endow workers with the necessary knowledge. Purdue University - one of the top US engineering schools - already offers a comprehensive semiconductor degree program that is entirely online. (Purdue University College of Engineering, 2025) Based on these considerations, we expect that theoretical training in semiconductor fundamentals will likely not be a limiting factor, or at least not as limiting as other educational components.

Practical training In addition to theoretical training, fab engineers need to know how to handle and maintain the actual manufacturing equipment. At TSMC, this part of training is conducted at their Newcomer Training Centers, at which all new engineering recruits receive 6-8 weeks of training, for example in equipment, process, process integration, and manufacturing systems. (TSMC ESG, 2025) The training center is located within Fab 15A in Central Taiwan Science Park, and consists of several classrooms, each of which is fitted with a piece of real manufacturing equipment. (Wei, 2022)

Instructors are senior staff members from TSMC. (Wei, 2022) Between its inception in July 2021 and October 2022, the TSMC Newcomer Center trained around 6,300 workers. (Wealth Magazine Express, 2022)[2]

At the current estimated throughput, training 300,000 engineers would require approximately 50 new training centers, each of which endowed with skilled instructors and cutting-edge, expensive machines. Currently, the training center is located inside of Fab 15A, but likely it is possible to construct centers that are outside of fab buildings, so that the physical infrastructure would not be a major hurdle. In 2021, the training center was equipped with 20 machines that operate on the production line and 12 auxiliary machines (e.g. metrology). If we multiply this by 50, we arrive at 1,600 pieces of equipment across all training centers - not much compared to the equipment needed for a single fab. An Intel report estimates that one of their fabs requires ∼1,200 tools. (Intel, 2025b) The physical requirements for training capacity may decrease in the future, as training through augmented and virtual reality becomes more effective. Augmented reality and virtual reality (AR/VR) are already used extensively at both TSMC and Intel. In 2021, TSMC's newcomer engineer training also involved four weeks of augmented reality training, which cut short the overall training period from 6 to 4 months, while avoiding equipment attrition. (TSMC, 2022b) Augmented reality is also used at Intel (Intel Newsroom, 2022) for both training and regular operations. If AR/VR training becomes even better, then this may further shorten the training period on real equipment and hence the required training capacity at physical institutions. These considerations suggest that the amount of physical capital for a large-scale cluster of training centers would be small compared to the capital requirements of a large number of fabs.

A more difficult challenge may be the number of qualified instructors, as it seems less plausible that practical instruction could be massively parallelized through online education. Moreover, the instructors at TSMCs Newcomer Training Center are senior, with an average of more than 15 years of experience. (Chang, 2023) If this is a hard requirement for instruction, then training large numbers of instructors may be difficult. Unfortunately, we could not find data on the number of instructors at the TSMC Newcomer Centers, but we have data for a similar training center: the TSMC Factory Service Academy, at which engineers are trained who maintain the factory infrastructure (e.g. gas and electricity). The Factory Service Academy has 10 full-time lecturers with more than 15 years of experience, with class sizes capped at 12. (CNA (Central News Agency), 2021) Within 11 months, they trained ∼500 people, which we shall extrapolate to ∼550 in 12 months. (CNA (Central News Agency), 2021) The ASML EUV Technology Training Center in Tainan has 14 instructors and trains 360 engineers each year. (ASML, 2020) To the extent that instruction at these training centers are comparable, this would mean that we need between ∼550-830 experienced instructors to train 300,000 people. Through a joint effort by TSMC, Intel, and Samsung, it might be possible to recruit and train ∼550-830 experienced instructors within a relatively short time. TSMC alone has a workforce of >70,000. It depends, however, on the level of experience necessary for effective instruction, and the ease at which this experience can be transmitted through trainers. An ASML executive said that the firm spent several years codifying the tacit knowledge of experienced engineers in order to make it legible enough to be transmitted at scale. Based on these considerations, we find it likely that availability of experienced instructors for practical training will be a bottleneck, although further research is required. The necessary physical investment for large-scale practical training seems much less daunting by comparison, even negligible compared to what is needed to build a large number of fabs.

On-job training Next to theoretical and practical instruction, new engineers need to spend some time in the fab to become familiar with real production procedures. US engineers recruited for TSMC's Arizona fab need to spend 12-18 months in Taiwan. (Vanek, 2021; Jennings & Söderberg, 2022) Some fraction of this time is spent working alongside mentors at TMSC's most advanced fab in Taiwan. (TSMC Careers, 2025b) [3]

- If the numbers from TSMC's Kumamoto fab are representative, then new engineering recruits will spend 2-3 months with on-the-job training. (TSMC Careers, 2025a)
- Process engineering recruits at Intel's Albuquerque fab need to spend an estimated 6 months on-site during training. (Intel, 2025a)
- ASML says it takes around 18 months to train an EUV engineer. (ASML, 2020)

Assuming that in-fab training is strictly necessary, how many semiconductor engineers can be trained in a fab? There are three key factors here: (1) the number of fabs that are available for training people in any given year, (2) the number of

[2] It is unclear what kinds of employees are meant here.

[3] There is differing information on the time required for in-fab training, and it likely varies between different roles.For cutting-edge nodes, training likely needs to take place at the most cutting-edge fabs.

engineers that can be trained at a fab at any given time, and (3) the time engineering trainees need to spend at a fab. A simple calculation suggests that even under very optimistic assumptions, in-fab training capacity could become a bottleneck to fab scaling.

We make the following optimistic assumptions:

1. **Fab availability:**
   - 20 Arizona-fab-sized fabs at the most cutting-edge technology level are available for training in 2024 across TSMC, Samsung and Intel.
   - New fabs can be constructed in 2 years.
2. **Training capacity per fab:**
   - Around 2,400 engineers are on site in each fab.
   - Training capacity for engineers is ∼33% of the number of engineers on site.
   - Sufficiently many senior engineers are available to provide supervision, including transmission of tacit knowledge.
3. **Required training time:**
   - Trainees need to spend on average 2 months in training.
   - Training takes place continuously at maximum capacity, even as the new fabs in which trainees will work are still under construction.
4. **Other assumptions:**
   - There are no frictions from onboarding, workforce integration, and management.
   - Impediments to regular production from equipment strain and/or lack of senior engineering time are negligible.
   - Newly built fabs stay technologically relevant for at least 10 years, and can be used for training during this time.

Even under these optimistic assumptions, between year 0 and year 4, it would only be possible to build around 2400 fabs. The doubling rate of fabs would be limited during the first couple of years, as training of new generations of engineers depends on construction and training of the previous generations. After that, however, enough training capacity would be available for very fast growth. (**?**)[4]

Contrasting this calculation with an outside view also indicates that extreme growth seems unlikely. Let us again assume that we want to train 300,000 semiconductor engineers per year. Keep in mind that this is a lower estimate, as we are not even considering technical workers, facility support, and management. By today's standards, this is a very large number of people to train each year. For comparison:

- The entire US semiconductor industry employs around 270,000 people, of which only ∼25% are engineers. (Semiconductor Industry Association (SIA), 2021; U.S. Bureau of Labor Statistics (BLS), 2025)
- South Korea announced that it wants to train 150,000 semiconductor workers within the next 10 years. (Nam, 2022)
- In 2023, TSMC plans to hire 6,000 people, with a current headcount of 70,000 people.
- The US graduates ∼200,000 people in semiconductor-relevant subjects each year. (Khan et al., 2020)
- For comparison, TSMC's states in 2022 that >6,800 people have "benefited" from its university collaboration partnerships, with 4,000 people since 2019 who enrolled in its semiconductor university programs. (TSMC, 2022a)

[4] A few caveats to this calculation: (1) Training may become vastly more efficient and standardized as fabs become more advanced. (2) There may be economies of scale to training lots of people. Perhaps it would be worth opening a number of dedicated training fabs (especially if economic pressure does not play a role anymore). (3) Virtual reality may become so good that fab training is not necessary anymore. Or trained employees are not needed, as the AI will give instructions.

A.2. STEM workforce

A.2.1. Summary

A semiconductor fab with a capacity of 40,000 WPM requires at least ∼2,400 high-skilled professionals with the equivalent of a 4-year engineering degree, likely also encompassing people with other semiconductor-relevant university degrees in STEM. The most common degrees are in electrical engineering, computer science, computer engineering, mechanical engineering, chemistry, chemical engineering, materials science, and physics. Even if all known annual STEM graduates in these subjects in the world were recruited to work in fabs, this would still only result in slightly more than ∼800 fabs per year. If all relevant US graduates are recruited, this would result in ∼80 fabs per year.

The global workforce of STEM-educated professionals is much larger than the number of graduates. If significant numbers of people are recruited from this talent pool, then a much larger number of fabs would be possible, although still not orders of magnitude more. For example, if a cohort of working professionals with STEM backgrounds equivalent to all 6.7 M relevant workers in the US is recruited annually from the global STEM workforce, this still only results in ∼2,800 fabs per year. For our production bottleneck, we assume that we recruit 6.7 M relevant working professionals annually (∼2,800 fabs per year) and relevant global annual STEM graduates (∼800 fabs per year) resulting in ∼3,600 fabs per year. This is likely a vast overestimate, as it would imply that each year, a group of people of equivalent size to the entire US STEM workforce would be hired. Overall, the magnitude of the production bottlenecks varies across ∼2 OOM, depending on which restrictions one might want to put on the size of the talent pool (e.g. US only, OECD, world, all working professionals vs. new graduates).

All of these estimates purely compare demand for high-skilled professionals with supply of graduates or professionals in other industries, ignoring factors like individual motivation, immigration restrictions between countries, limits to talent integration, and general competence. Moreover, we ignore the additional labor requirements for staffing firms across the international supply chain. Both factors imply that these production bottlenecks are likely vast overestimates.

Production bottlenecks	**Number of fabs possible (total or per year)**	**Key underlying assumptions**	**Comments**
Availability of new STEM graduates per year in the US	∼80	Recruiting all semiconductor-relevant US STEM graduates in 2021. All semiconductor-relevant disciplines fungible.	This excludes STEM graduates from other countries.
Availability of new STEM graduates per year in the world	∼800	World STEM graduates (2014 or most recent year) recruited. All semiconductor-relevant disciplines fungible. Relies on extrapolation of degree shares from US share.	This includes all STEM graduates from countries for which education statistics are available.
Total of global STEM graduates & US working professionals	∼2,800	Recruiting the equivalent of 100% of the 2021 US STEM workforce each year. All semiconductor-relevant disciplines fungible. Relies on extrapolation of degree shares from US share. Includes bachelor degrees only for the US STEM workforce.	This scenario implicitly implies large-scale, world-wide poaching of engineers from other branches of industry.
Total of global STEM graduates & US working professionals	∼3,600	Recruiting the equivalent of 100% of the 2021 US STEM workforce each year. World STEM graduates (2014 or most recent year) recruited.	See previous comments.

Table 2. Production bottlenecks to annual construction of fabs resulting from limited STEM workforce availability.

A.2.2. Approach

On the demand side, we focus on STEM degrees that are directly relevant for semiconductor manufacturing. Based on a list by CSET (Khan et al., 2020), as well as the most common degree requirements for core semiconductor engineering openings at TSMC (TSMC, 2025), the following STEM degrees fall into this category:

- Electrical-, and computer engineering
- Computer science
- Mechanical engineering
- Chemistry and Chemical engineering
- Material science
- Physics

We focus on STEM workers for two reasons: (1) Workers with STEM degrees are likely in short supply relative to university-educated workers without STEM degrees or relative to unskilled workers. There is some evidence that skilled, non-university educated technicians may be in even shorter supply (Kerkmann & Ströder, 2023), but these seem harder to quantify, due to the multitude of different job requirements. (2) The data on STEM college graduates and workforce is reliable and precise (Asianometry, 2020) and semiconductor-relevant STEM disciplines are a good delimiter, as graduates from more unrelated subjects would be, on average, less able to transition into semiconductor manufacturing.

To quantify demand, we extrapolate from the publicly stated workforce requirements of TSMC's Arizona fab. We assume that the labor requirements for each additional fab stay constant, implicitly assuming that no automation or economies of scale take place. On the supply side, we quantify the size of global talent pools of relevant STEM graduates in the US, the OECD, and the wider world. A systematic overestimation of the available global workforce occurs because not all countries publish exact graduate statistics. At the same time, we likely underestimate workforce requirements as analyses often focus narrowly on publicly stated engineering needs for specific fabs, like TSMC's Arizona facility, thereby neglecting non-engineering STEM roles within the fab and the broader STEM demands across the supply chain. [5]

A.2.3. Demand

TSMC will hire 4,500 people for its two Arizona fabs. (Dean, 2022) They claim that 3,000 of its 4,500 planned Arizona roles will require 4-year engineering degrees. (Terrill, 2023) The actual number of semiconductor-relevant STEM workers may be higher, but we accept this, since we want to err on the side of overestimating potential for fab construction. The two Arizona fabs combined will have a capacity of 50,000 WPM. Assuming the workforce composition scales with constant proportions, this implies that a 40,000 WPM TSMC fab would require 2,400 workers with 4-year engineering degrees.

TSMC is not the only company that would require additional workers. Across the global supply chain, firms like ASML, Applied Materials, Tokyo Electron, LAM, and KLA would have to grow their workforce commensurate with the exploding demand from foundries. Hence, for each 40,000 WPM fab, large numbers of additional STEM workers would be required. Unlike with TSMC, in the supply chain it is much harder to rely directly on first-hand information on the number of STEM jobs needed. Some rough estimates suggest that 35% of new jobs will be at foundries and 65% at suppliers:[6]

In the absence of more precise numbers, we will just focus on labor requirements at TSMC directly, as we deem the numbers on these most reliable. The degree requirements for STEM workers at TSMC are not directly known. Our best proxy comes from looking at (1) an analysis of LinkedIn profiles at TSMC and (2) CSET data on green card applicants for semiconductor fabs between 2010-2018. (Khan et al., 2020) Both sources show that electrical engineers, material scientists, and chemists make up the bulk of workers. There is a mismatch on the relative importance of electrical engineers, but other than that, the splits look approximately similar. We assume that the actual skill distribution inside of an Arizona-type fab high-skilled workforce is the average of these two distributions.

This breakdown leaves open the question, to what extent certain fab roles can be performed by people with different STEM degree backgrounds. In some engineering roles, not much subject-specific background knowledge may be required, as most

[5] Furthermore, imprecision in these estimates arises from the difficulty in finding detailed subject-specific data outside the United States, with international overviews from sources like the OECD offering only broad categories such as "Information and Communication Technologies" or "Engineering, Manufacturing, and construction." Consequently, figures for individual subjects are often extrapolations based on US proportions, and ambiguities in how subjects like computer science are categorized (e.g., as natural sciences or engineering) and variations in degree quality across institutions can introduce further inaccuracies, potentially leading to either over- or underestimations, though the overall effect of this imprecision is considered not to be very strong.

[6] TSMC claims that its Arizona fab will create ∼9,500 high-tech jobs at regional suppliers. (Shepardson & Alper, 2022) Taken as representative, this would imply that 65% of overall (STEM and non-STEM) jobs will be at suppliers.

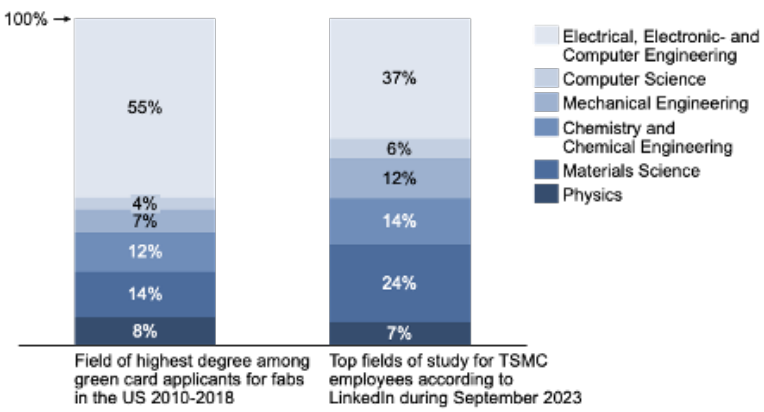

Figure 3. Degree backgrounds of (1) green card applicants planning to work at US semiconductor fabs 2010-2018 and (2) TSMC employees that were listed on LinkedIn in September 2023)

skills would be picked up on the job. Other roles may require slowly-acquired scientific niche-knowledge. Let us call the degree to which subjects are interchangeable the "background interchangeability".

- If the background interchangeability is small, this means that each role given to an electrical engineer also requires the particular skills of an electrical engineer, and not that of some other type of engineer. If this is true, then the distribution of study backgrounds is the result of the particular skill requirements of each role in the fab.
- If the background interchangeability is large, then most roles could be performed by any type of person with a STEM background. The distribution of study backgrounds may then reflect supply of degrees.
- There may be an option in between, where the skill overlap between some roles is large, but not between others.

Most of the current open engineering roles at TSMC and Intel are open to graduates with different engineering backgrounds, plus graduates from "related fields". (**?**) This suggests that background interchangeability is large, and that many important skills can be successfully taught on the job. Nevertheless, for the purpose of calculating production bottlenecks, we will also consider a scenario with a smaller background overlap. Moreover, we do not differentiate between roles for PhDs, masters, and so on, which likely results in an overestimation of how many fabs can be fully staffed, as some roles may not be suitable for bachelor graduates.

A.2.4. Supply

Across all countries that publish education statistics (National Center for Science and Engineering Statistics (NCSES), 2019; National Science Foundation (NSF), 2018), approximately 2 M people graduate each year with a first degree that is related to semiconductor manufacturing. In the OECD there are approximately 1 M graduates per year, of which ∼200,000 are from the US. (Data USA, 2025)

In an aggressive scaling scenario, recent university graduates would not be the only available types of workers. Strong enough incentives could motivate workers from the existing world-wide STEM workforce to re-train and switch careers. Many more workers would then be available - a calculation shows that in the US, approximately 6.7 M workers with at least first degrees in semiconductor-relevant subjects are available. (Data USA, 2025) The world-wide workforce is likely much larger than the US workforce, but difficult to quantify.

A.3. Raw materials

A.3.1. Summary

No material used for AI chip fabrication is scarce enough to be a likely bottleneck for production, given strong economic incentives to re-allocate material flows or to increase output where necessary, as assumed in extreme scenarios. (halbleiter.org, 2025)

Semiconductor-relevant subjects	US STEM graduates in 2021 (bachelor/master/PhD)	Estimate for OECD 2021 (bachelor/master/PhD)	Estimate for the world 2021 (first degree only)
Electrical, Electronic, and Computer Engineering	46,410	313,296	573,894
Computer Science	56,186	176,615	132,776
Mechanical Engineering	46,015	310,629	728,642
Chemistry and Chemical Engineering	33,080	143,808	392,853
Materials Science	4,069	24,877	56,971
Physics	12,240	34,430	68,445
Total	**198,000**	**1,003,656**	**1,953,581**

Table 3. Number of graduates in semiconductor-relevant STEM subjects in 2021. (Data USA, 2025; National Center for Science and Engineering Statistics (NCSES), 2019; National Science Foundation (NSF), 2018)

US workforce in semiconductor-relevant subjects 2021 (first degrees only)	Number of people (estimate)
Electrical, Electronic, and Computer Engineering	1,927,275
Computer Science	1,697,509
Mechanical Engineering	1,258,377
Chemistry and Chemical Engineering	1,349,454
Materials Science	87,355
Physics	426,835
Total	**6,746,805**

Table 4. Number of people in the US workforce with first degrees in semiconductor-relevant subjects. (Data USA, 2025)

Production bottlenecks	Number of fabs possible per year	Key underlying assumptions	Comments
Raw materials	No relevant limit	Mines and refineries are relatively low complexity facilities, so they can be built as fast or faster than other input supply chain components. Demand is small in absolute terms, i.e. wafers only require small quantities of materials. Available mineral supply and reserves are very large relative to that, even in extreme scenarios.	

Table 5. Production bottlenecks to annual construction of fabs resulting from limited raw material supply.

A.3.2. Approach

Modern IC production makes use of almost the entire periodic table of elements for various very specialized applications, which also are potentially vastly different between e.g. a photosensor, a memory chip, a logic chip, or a display. It is hard to estimate the amounts of various elements used for a single logic chip wafer. Therefore we choose to grossly overestimate demand for each element, and compare it to global supply of this element, based on global mining and production data.

A.3.3. Demand

We start with a demand estimate for each element by assuming it is equal to 1% of wafer weight (=1.25 g). This is a vast overestimate of material needs, as most of a wafer's 0.5 mm thickness is silicon, with the actual layers making up the transistors are just a few nanometers thick. If we roughly assume 100x of today's < 7 nm production, at 1.25 g/wafer around 600 metric tons per year (t/y) of each element would be needed. Despite our vast overestimate, 600 t/y is still an

exceedingly small amount compared to global production of most elements. For example, 600 t/y is roughly a fifth of annual gold production.

Ruthenium is the only element we found that could have a remote chance of being a bottleneck. The metal is used as a capping layer in EUV masks. These layers are only 2.5nm thick though, requiring around 0.7 mg of material for a mask that is used for many wafers. This amount of material use would mean Ruthenium supplies would be virtually endless even in the most extreme scaling scenarios.

Across all production steps, approximately 0.015 g of partially non-abundant materials are deposited on each wafer. (halbleiter.org, 2025) With this figure a 100x TSMC scenario would require a total of 7.2 t/y of various materials for all deposition steps in wafer production, which is just above the global production volume of the exceedingly rare metal Iridium alone. Assuming that - for the sake of argument - a quarter of the material is Iridium and that it takes one year to double Iridium production, this would result in a limit of around 4,000 fabs/y. But this is an exceedingly conservative estimate yet again, so we cannot take this as a production bottleneck. Instead, this indicates that no relevant limits will result from raw material supply shortages.

Going through the elements that are produced in quantities smaller than 600 t/y, none seems to be overly relevant for logic chip production, or to have viable reserves that are not many multiples of current yearly production, not to mention unproven reserves.

A.3.4. SUPPLY

Z	Element	World Prod. (t/y)	Estimated Usage in AI Chips	Source(s)
76	Osmium	<1	Most likely none	Osmium World Council
37	Rubidium	6	Very low; 100k t reserves	USGS, Mordor Intel.
77	Iridium	7	Reserves »1k t; 16.8% electronics	Statista
81	Thallium	10	Very low, most likely	Wikiwand
21	Scandium	20	Very low; Likely no SCI app.	MCS 2022, SciDirect, SputterTargets
45	Rhodium	21	Very low; 0.57% electronics	Statista
44	Ruthenium	31	Capping layer EUV photomasks; 5k t reserves;	Statista
55	Caesium	45	Very low, most likely	EarthMag
75	Rhenium	59	Very low; Tiny amounts in some SME/experimental	MCS 2022, SEMI
72	Hafnium	67	»1M t reserves; 2-3% in Zr ores	CRM Alliance, USGS
32	Germanium	140	>200k t reserves	MCS 2022
78	Platinum	180	Low; 5-7% demand from electronics	MCS 2022, BullionVault
46	Palladium	200	Low; 82% demand from catalytic converters	MCS 2022, Statista
4	Beryllium	260	Low; World res. >100k t; 60% in U.S.	MCS 2022, USGS
(1)	Deuterium	323	Very low	WorldBank, Isowater
31	Gallium	430	>100k t res.	MCS 2022
52	Tellurium	580	Very low; Potential future transistor material	MCS 2022, Nature
49	Indium	920	Usage info not specified	MCS 2022
73	Tantalum	2,100	Usage info not specified	MCS 2022
80	Mercury	2,200	Usage info not specified	Statista
79	Gold	3,100	Usage info not specified	Statista
34	Selenium	3,330	Usage info not specified	Statista
39	Yttrium	10,000	Low; Used in SME coatings	MCS 2022, SAMaterials
14	Silicon	16,780	Wafers (300mm) in tons	Statista, WolframAlpha
83	Bismuth	19,000	Usage info not specified	MCS 2022
48	Cadmium	24,000	Usage info not specified	USGS MCS 2023
47	Silver	26,000	Usage info not specified	Statista
53	Iodine	33,000	Usage info not specified	Statista
33	Arsenic	40,539	Usage info not specified	MCS 2022, WolframAlpha
92	Uranium	48,888	Usage info not specified	Statista

Table 6. Overview of the 34 least-produced elements relevant for semiconductor manufacturing. Production in metric tons/year.

A.4. Electricity

A.4.1. Summary

A modern 40 kWPM fab requires 350 MW of electrical power. A "gigafab" comprising 6 such fabs requires about 2100 MW, as much as the output of a typical nuclear power plant with two reactors. US industrial electricity would suffice for 326 fabs, the world's for 3094 fabs. This is equivalent to 33x (US) 309x (world) of TSMC's current < 7 nm capacity. Extreme scaling of chip production would require the construction of additional power plants around fab locations, the speed and scale of which would rival or exceed the periods of fastest power addition in human history.

Production bottlenecks	**Number of fabs possible per year**	**Key underlying assumptions**	**Comments**
US industry	326	Rerouting electricity from all industrial consumers.	It would be more practical to build more power plants, due to scarce transmission capacity.
World industry	3,094		
Uranium	No relevant limit	10 y at commercially viable supplies (2009 data)	The quantity of potentially available mineral supply is generally much larger.
Gas	No relevant limit	10 y at commercially viable supplies (2020 data)	
Solar	No relevant limit	1% surface area Earth, 20% efficiency	
Historical comparison	Growth speed at historical extreme is 856/y	Highest historical single year growth rate for the entire world. Chinese additions taken as 65% efficient modern plants. All power plants for fabs.	Low detail BOTEC

Table 7. Production bottlenecks to annual construction of fabs resulting from limited electricity supply.

A.4.2. Approach

Fabs require a significant amount of electricity, for example to maintain precise environmental conditions using air conditioning, for powering equipment like EUV machines (1 kW each) or for operating ultra-pure water systems.

Still, in comparison to other industries, electricity consumption is not very high in absolute terms, as no major material conversion occurs. A single electric arc furnace for steel recycling requires as much power as an entire fab. Aluminum plants may require 5 times as much power per facility area.

A straightforward estimate for this potential bottleneck is to compare electricity consumption per fab with US (or world) electricity production to estimate how many fabs could be powered with current supply. We will do this below for orientation, but we think it is more informative to follow an approach that considers the construction of new power plants rather than a rerouting of electricity from the existing grid. This also means that comparisons to US or world electricity consumption can only be taken as bottlenecks if for some reason rerouting power would be possible, but construction would not be possible, which seems implausible.

A.4.3. Demand

The first phase of the Arizona plant (Fab 21 Phase 1, hereafter F21P1) will have around 200 MW of electrical power made available. We are assuming 100% utilization, as fabs are operated 24/7, so this is constant 200 MW for 8,760 h in a year, i.e. 1.752 terawatt hours (TWh). The project in Arizona is a multi-billion dollar investment for TSMC, but small on the scales considered here. It requires "only" relocation of transmission lines, not the construction of new power plants. However, they also need to build out an existing substation and add a second one. In contrast, new fabs in Taiwan seem to necessitate new power generation in some cases. However, although it is difficult to find out exactly, the two new 2 nm fabs in Taichung with 40 kWPM each apparently require 350 MW each, which fits well with the higher capacity and more advanced node. In this case yearly electricity use per fab would be 3.066 TWh.

A.4.4. SUPPLY

With its industrial energy supply of around 1000 TWh, the US can power 326 fabs, equal to to ∼13,040 kWPM or ∼33x current < 7 nm capacity of TSMC. The worldwide industrial energy supply of around 9492 TWh could power 3,094 fabs, or ∼123,760 kWPM ∼309x current < 7 nm capacity of TSMC.

Calculating with total energy supply, the the annual US electricity production of ∼4,243 TWh could power 1,383 fabs (40 kWPM, 2 nm) and the world electricity production of 2020 could power 7786 fabs.

Global energy supply would not be a bottleneck, as theoretically available sources for power far exceed any reasonable demand.[7]

So in conclusion, electricity needs for large-scale fab construction are substantive in absolute terms, but can be met with power plant construction that is not subject to meaningful physical limits. Construction of power infrastructure at such scales would be a massive undertaking, but most likely less complex than fabs.

A.5. Water

A.5.1. SUMMARY

A 40,000 WPM fab with recycling capacity requires 6.4 Mm^3 of water annually, with only 20% of that being lost in the case of F21 P1 in Arizona. This is an amount similar to the collective water use of 15,500 average US households and only a small fraction of local use in the vicinity of the fab. Many options for providing water to fabs exist, given that one could buy water usage rights, deplete local aquifers, build pipelines, or simply build in non-desert locations. Water use is therefore unlikely to be a limit to scaling under extreme scaling assumptions.

Production bottlenecks	Number of fabs possible per year	Key underlying assumptions	Comments
Water	No relevant limit	Desalination plants with water pipelines where necessary.	Current local political issues in e.g. Arizona or Taiwan around these topics are on small scales relative to construction activity in an extreme scaling scenario.

Table 8. Production bottlenecks to annual construction of fabs resulting from limited water supply.

Ultra-pure water is used to clean wafers between fabrication steps, such as etching, baking, deposition, lithography, and so on. The requirements for the number of particles in the water are extremely strict, because even very small particles can ruin a chip at any point in the process, decreasing yield. (Y, 2025) Most of the water can be recycled, which is increasingly common practice. Once a fabs water reservoir is filled up, it only needs to be replenished from outside supply at a rate of a few percent per year. Usually, recycling water internally is only done when water isn't abundant. (**?**Arizona Technology Council, 2023)

There is enough water globally to supply a very large number of fabs, and yet water usage is often cited as a problem for fabs today, especially in dry places like Arizona. To estimate whether this could become a problem in extreme cases, we look at the water usage of TSMC's Arizona fab and compare it to local supply in Phoenix. We also consider how fab water use compares to other large consumers, and whether the local supply could be expanded.

A.5.2. DEMAND

TSMC's F21 P1 fab in Arizona requires 12 Mm^3/y assuming no recycling, which approximately matches the larger Taichung P1 fab's requirement of 18 Mm^3/y: Both yield around 4.7 Mm^3/y per 10 kWPM wafer production capacity. While some sources say that 98%+ can be recycled, current systems as planned recycle about 65% of the water (e.g. TSMC F21 P1). With this efficiency, the necessary water to replenish the fab's water inventory is 4.3 Mm^3/y. Assuming similar efficiencies

[7]Solar radiation: 4.4e16 W equivalent to 2.5e5 fabs, with 20% efficiency and 1% of earth surface area. (NASA, 2008) Uranium: 8e12 W for 10 y with known commercially viable reserves in 2009 and some reactor improvements, equivalent to ∼23,000 fabs. (Phillips, 2009) Many OOM more with Uranium extraction from seawater. Natural gas: 16e12 W for 10 y with known commercially viable reserves in 2020 and 66% efficiency, equivalent to ∼46,000 fabs. (U.S. Energy Information Administration (EIA), 2023)

for a typical 40 kWPM fab like in Taichung yields 6.4 Mm^3/y.

A.5.3. SUPPLY

Comparing Arizona water use and relevant other numbers indicates that the 6.4 Mm^3/y water demand for one fab is relatively small when we look at larger scales of a few fabs per city or metropolitan area. The largest US nuclear power plant happens to be located in Arizona. Most nuclear power plants are cooled with water from rivers or seawater. The Palo Verde Nuclear Generating Station instead evaporates 100 Mm^3/y for waste heat rejection. In an extreme scaling scenario, it seems likely that at least as much water would be available for fabs, sufficient for 15 facilities, even if this would mean depleting the local aquifers, as is the case for the Palo Verde plant. The Phoenix area around F21P1, including agriculture, consumes 2847 Mm^3/y, enough for 445 fabs in that part of the Sonoran desert around Phoenix, Arizona alone (if all water rights were reallocated to fabs). 6.4 Mm^3/y is 0.22% of that. A USD 5.4 B proposal envisions a 322 km long pipeline from a desalination plant at the coast in Mexico to Arizona. (The Washington Post, 2025) It would have an initial capacity of 370 Mm^3/y in the first phase and up to three times as much in potential later phases. Electricity requirements for desalination of the amount of water required for a 40 kWPM fab with 65% recycling are 2.6 MW per fab. This is a negligible fraction (0.8%) of the total electricity demand.

A.6. Suitable sites

A.6.1. SUMMARY

Based on a very conservative estimate more than 66,635 km^2 is available at suitable sites, which is sufficient for 6,664 fabs. Given the extremely conservative assumptions (see below), which allow us to easily add an OOM or two more fabs when we relax one or few of them, it seems clear that space is not an issue. Hence, neither sites nor logistics seem like a meaningful limit for large-scale fab construction. Megasite preparation is currently a slow, political process, but physically, it consists of low-complexity work, like grading and road building. Hence, under the assumptions of this analysis, the effort to prepare building sites seems minor compared to other tasks. Even under extremely conservative assumptions, there are enough suitable building sites for hundreds of very large industrial zones close to population centers. This proximity in turn makes the placement near major roads, rail lines, seaports, or airports easy.

Production bottlenecks	Number of fabs possible per year	Key underlying assumptions	Comments
Suitable sites	No relevant limit	Over 10 years. Only >5 M OECD cities. Other extremely conservative assumptions. Relaxing them would yield 1-2 OOM.	5 M is larger than many cities in Taiwan that host most of the world's advanced fabs, like Tainan (1.86 M), Hsinchu (0.45 M), Taichung (2.83 M). Only 1% of Earth's habitable surface area is urban or other areas covered by buildings, roads, etc.

Table 9. Production bottlenecks to annual construction of fabs resulting from limited site availability.

A.6.2. APPROACH

For current fab projects, it is often challenging to find a site that can support all necessary requirements at sufficiently low cost. Flattening land, building roads, transmission lines, pipelines, etc. costs money, and land close to well-educated talent pools may be scarce or hard to acquire.[8] Moreover, seismic requirements may render sites unusable that are located close to highways, airports, or rail lines. (Taylor, 2023) Under current circumstances, megasite preparation projects undertaken by local and national governments often take many years to bring to completion. (Teague, 2023)

Preparing sites for construction is mostly not very complex compared to fab construction itself, as it involves things like utility connections for water and electricity, roads, and earthworks. The inputs for a fab are expensive and specialized, but they are not particularly voluminous or heavy, so as long as a major road or railway is adjacent to the site, logistical roadblocks seem minor. A more complicated requirement could be seismic isolation: even small vibrations can disturb the precise processing steps in semiconductor production. However, many fabs in Taiwan and the US are located in areas with

[8] Source: expert interview.

high, or at least occasional, seismic activity. (Semiconductor Industry Association (SIA), 2025) Moreover, both TSMC's Fab 21 and ASML's Veldhoven site are located next to large roads.

Given that both physical site preparation and seismic isolation are already done with relative ease under current economic circumstances, it seems likely that they could be done as a regular step of the construction of fabs. To us, it seems unlikely enough to be a serious bottleneck. What makes suitable sites rare and slow to prepare today, apart from economics, is mostly the political side of things. Environmental reviews, assembling a large enough plot of land from various owners, and getting the buy-in to spend millions on something that might only yield benefits in a decade (or not at all) is much more likely to be the cause of delays.

A.6.3. Demand

How much building space needs to be available for fabs and associated buildings? A simple calculation is based on the F21P1 area including surrounding tech park and mixed use areas, which would yield 15 km^2 for a gigafab. We quadruple this number to account for other buildings to be more conservative, which yields 60 km^2.

A.6.4. Supply

How much suitable space is available for building fabs? Only 1% of the earth's habitable land area is cities or other built up area, while around 46% are being used for agriculture. (Ritchie et al., 2024) We therefore construct an estimate of available agricultural or other land around cities with the following extremely conservative assumptions:

- Only cities with $>$ 5 M inhabitants are eligible, which reduces the number of cities to 81.[9]
- Only OECD cities have the necessary talent pools, which reduces the number of cities further to 14 (by scaling by the fraction of OECD population vs. the world).
- All cities have the same population density as the most dense of the 81, i.e. Dhaka, Bangladesh. This yields 167 km^2 on average.
- All cities are at the sea or have some other geographical feature around their perimeter that makes those areas unsuitable, so they are roughly shaped like a half-circle. This yields a radius of 10 km and a perimeter of 53 km. It should also justify the assumption that logistics is not a problem if a few 100 km of six lane roads or equivalents are added.
- The area around this shape within a 30 km distance from the half-circle shaped city, summed over all 14 cities is 66,635 km^2.

A.7. Capital expenditure

A.7.1. Summary

One of the newest 'greenfield' TSMC fab in Taiwan, "Fab 18 Phase 9" (3 nm 40 kWPM), is projected to cost around USD 13.38 B (SEMI, 2023), whicih is comparativelt cheap. In contrast, TSMC's new Arizona fabs cost a sum equivalent to USD 32 B per 40 kWPM fab (SEMI, 2023). World GWP in 2023 was $\sim$USD 105 T.(International Monetary Fund, 2025) Assuming a price of USD 13.38 B, the equivalent of 2023 GWP could buy $\sim$7,800 fabs.

This approach obviously ignores many economic considerations. It is not to be seen as a prediction, but merely as an initial order-of-magnitude guess for how many fabs might realistically be built in any given year, to be used as a basis for bottleneck calculations.

A.7.2. Approach

For the demand side, we select the cheapest relevant TSMC fabrication plant (fab) from the SEMI World Fab Forecast and utilize its capital cost as the lowest cost estimate. For the supply side, we employ arbitrarily chosen points of reference, specifically ranging from 10% to 100% of the 2023 Gross World Product (GWP).

[9] This excludes Phoenix, Arizona by a thin margin (4.85 M when using the metropolitan area figures, instead of the city proper), and many cities in Taiwan that actually have most of the world's advanced fabs, like Tainan (1.86 M), Hsinchu (0.45 M), Taichung (2.83 M). The entire province of Noord-Brabant, where ASML's EUV fabrication is located, has a population of 2.56 M people. For more on suitable cities with skilled talent pools, see this analysis: (**?**)

Table 10. Production bottlenecks to annual construction of fabs resulting from limited availability of capital.

Production bottlenecks	Number of fabs possible (per year)	Key underlying assumptions	Comments
10% 2023 GWP	∼780 fabs	USD 13.38 B per fab, 10% GWP	No change over time modeled
100% 2023 GWP	∼7,800 fabs	USD 13.38 B per fab, 100% GWP	No change over time modeled

The approach systematically underestimates costs by using the lowest available fab CAPEX and does not exhaustively cover all supply chain, operational, or surrounding economic expenditures, which could be significant. Furthermore, the reliance on external data sources like the World Fab Forecast means potential inaccuracies cannot be ruled out, as costs are not independently modeled. Finally, the assumption of constant costs over time simplifies dynamics where prices could either decrease with efficiency or increase with demand. As such, this estimate is intended as a reasonable starting point for analysis rather than a precise prediction.

A.7.3. Demand

While TSMC's most recent 2-3 nm 40 kWPM fabs in Taiwan cost around USD 13.38 B and have cost the same in the past fairly consistently, they announced a cost of USD 40 B for F21 P1-2 in Arizona (Liu & Mozur, 2023), which together account for just 50 kWPM, implying a cost of USD 32 B per 40 kWPM fab. Among other things, they cite 4.5 times higher construction costs as a reason. TSMC Fab 18 Phase 9 (F18 P9) serves as our lower bound for costs. It cost USD 13.38 B in total, which includes USD 2.3 B for construction and USD 11.08 B for equipment. We chose this fab as lower bound for several reasons: (1) It produces chips in a relevant node (3 nm), at a capacity of 40 kWPM, (2) It has a greenfield construction, meaning that costs for preparing the site, providing water and power infrastructure are included (among other things) and (3) It is comparatively cheap (partly due to its position in Taiwan, the home country of TSMC). This is helpful for erring on the side of underestimating costs. There are several CAPEX needs beyond the needs for fabs. Modeling these in detail is out of scope for this analysis. [10]

A.7.4. Supply

Supply of CAPEX depends on various actors' willingness and ability to spend. For this report, we assume an extreme willingness to spend equalling 100% of 2023 GWP, which is roughly USD 105 T. (International Monetary Fund, 2025)

B. Suppliers (ASML Case Study)

B.1. Overview

Extreme scaling of AI chip production would have to go hand-in-hand with extreme scaling of the entire semiconductor supply chain, which includes Tier 1 suppliers like ASML, Applied Materials, KLA, LAM, and Tokyo Electron, Tier 2 suppliers like Carl Zeiss and Trumpf, and thousands of other Tier N companies. Long lead times for equipment are already a common complaint in the semiconductor industry, and some key suppliers have struggled to deliver despite strong monetary incentives. It seems likely that an extreme scaling scenario would multiply and prolong such delays.

How would an extreme scaling scenario impact suppliers' ability to deliver? To understand this, we interviewed 2 former and 1 current ASML senior executive, and asked them if it was possible to scale up production by 10x or more within 10 years. We find it plausible that this inside view is representative for wider parts of the supply chain, as it contains many companies similar to ASML, in the sense that they produce niche high-tech products that require a high degree of company-wide technical expertise.

[10] For example, CAPEX is needed for Tier 2 to Tier N suppliers. ASML has a total gross value of USD 7 B for property, plant and equipment as of end of 2022(Yahoo Finance, 2025) and Linde plans a plant for USD 600 M in Arizona for F21 Phase 1 & 2.

B.2. Could ASML scale EUV machine production fast? (according to former senior ASML executives)

B.2.1. Context

The section below presents the state of evidence on extreme ASML scaling after 5 hours of expert interviews. All information that is directly attributable to a source is marked by a footnote directly in the sentence (i.e. there are no overarching footnotes which include multiple sub-bullets, like in the section above). Other sentences are inferences and extrapolations from our side.

The information gathered in this section was largely gathered by prompting experts with a hypothetical 10x scaling scenario, and asking them what would happen if this scenario came true. It was also clear from context that we are interested in a rapid scaling scenario, say, within 10 years. For the purpose of this section, I will call this scenario *10 y / 10x*. We mostly did not push for more extreme scenarios, as it seemed like the expert's intuition might break down at a certain point, and the information gathered for the 10x scenario was likely informative for more extreme scenarios.

This section mostly considers the case in which more manufacturing facilities would be built, not a case in which existing capacity is shifted, or a case in which fabs are repurposed. This is partly because the experts did not touch on this topic much, and partly because we believe that these scenarios are less relevant based on evidence from other sources. Nevertheless, it seems appropriate to spend some additional time investigating the relevance of these scenarios elsewhere.

B.2.2. Summary

- Expanding geographically has challenges (e.g. limited space near HQ, strains on local labor/housing) but seems feasible through strategies like expanding to the wider region, hiring internationally, and getting government support.
- Expanding ASML's workforce requires hiring/training many additional skilled technical workers and engineers. Talent seems transferable from other industries and ASML is an attractive employer, but training capacity could be a bottleneck.
- Key suppliers like Zeiss and Trumpf already struggled to meet much lower production increases, so supply chain scaling seems challenging from ASML's perspective. A lack of visibility deep into the supply chain poses additional hurdles.
- Organizational change management has proven difficult, but focused executive attention and extensive resources could enable adaptations needed for scaling production.
- Overall, geographic, human capital, supplier, and organizational challenges exist, but none seem fully insurmountable given a concerted global effort. The finely-tuned manufacturing process reliant on ad-hoc fixes seems like the most compelling potential barrier. Both are weak impressions that are entirely dependent on evidence from three former senior ASML employees.

B.2.3. Evaluating the business move

Before companies would consider scaling production, they would carefully consider whether doing so would be a good business move. ASML in particular is known as a methodical player that carefully and deliberately plans its next steps.[11] The expert interviews hinted at two general classes of business-related roadblocks associated with a *10 y / 10x* scenario:

- There may still be a perceived residual risk that demand doesn't materialize
  - **Avoiding crushing fixed cost:** If capacity is built and demand does not materialize, then ASML and many suppliers may fail to amortize the massive fixed cost associated with the expansion.[12] This could have catastrophic consequences, and depending on the scale of expansion and the severity of the demand slowdown, it could plausibly bring companies close to bankruptcy.
  - **Getting enough trust:** Investors and financiers would have to mobilize large sums for the necessary investments, and depending on their perceived risk of demand instability, they might choose to cap their funding at a lower bar than required for a *10 y / 10x* scenario.

[11] Expert 1 [16.05.23]
[12] Expert 2 [17.05.23]

- Suppliers may be unwilling to serve extreme demand:
  - **Avoiding brain drain:** Going for a mass production strategy could alienate innovation-driven engineers and may result in a brain drain away from ASML.[13]
  - **Maintaining quality:** Companies may be risk-averse to be able to maintain quality and reliability. High pressure for speed may result in corner-cutting and increased risk-taking during production.[14]
  - **Staying agile:** ASML would generally be worried about becoming too large and unwieldy.[15]

One might argue that the first class of risk is ruled out by the scenario we assume: a sustained and predictable 10x demand increase for AI chips. We agree that this would be true in a scenario in which demand would be 100% guaranteed, e.g. through government backing. In a scenario with less-than-100% demand guarantee, even small risks of demand collapse may, in expectation, be hugely detrimental for companies and risk for financiers. Our interpretation of our scenario is that we are not yet in a war economy, and there may still be a small risk of an unexpected AI winter. Therefore, it seems reasonable to include the first class of risk.

B.2.4. Expanding geographically

A 10x expansion plan would require large capital investments in new ASML production facilities. Increasing the capacity by shifting existing production capacity and repurposing existing fabs would likely not be enough. Assuming that financing would not be an issue, how would ASML go about expanding?

Ideally, ASML would like to expand within its comfort zone, that is in the vicinity of its Veldhoven headquarters. This has several advantages.

- **"Finely-tuned" production process:** EUV machines are not automated mass products, so volume production requires the physical presence of engineers and suppliers in order to fix problems as they come up. Replicating this arrangement abroad would be hard.[16]
- **Collaborative culture:** it would help foster the kind of trust and collaboration that evolves from being together in a room and fixing things by trial and error.[17]
- **Employee preferences:** it would avoid a situation in which regular long-distance travel would be required by many of ASML's engineers,[18] or in which engineers would have to move their entire families to another country.[19]
- **IP protection:** geographical proximity makes it easier to protect IP from espionage and leaks.[20]
- **Geopolitical constraints:** ASML would want to shield its operations from geopolitical tensions.[21] This suggests that they would strongly prefer an expansion within the Netherlands, or at least within the EU.

However, expanding the Veldhoven facility would face a number of problems.

- **Limited space in Veldhoven:** the ASML headquarter is jammed between a large highway and the city of Veldhoven. Any attempt at expansion would involve crossing the highway and/or buying up property in the surrounding areas. Both may exacerbate already existing frictions between local inhabitants and ASML.[22]

[13] Expert 3 [12.05.23]
[14] Expert 2 [17.05.23]
[15] Expert 1 [16.05.23]
[16] Expert 2 [11.05.23]
[17] Expert 2 [11.05.23]
[18] Expert 1 [10.05.23]
[19] Expert 3 [12.05.23]
[20] Expert 1 [10.05.23]
[21] Expert 3 [12.05.23]
[22] Expert 1 [10.05.23]

- **Collaborative environment:** the Netherlands is a collaborative environment, so many stakeholders would have a say in expansion plans, for example the suppliers, the unions, the government, the local population, and the municipality.[23] This makes it more likely that extreme expansion plans would be limited or delayed.

- **Strains on the local housing and labor market:** ASML already puts a significant strain on the local labor market. Local car manufacturers in the Eindhoven area find it difficult to find talent because people prefer to work in ASML's clean rooms.[24] A 10x expansion would exacerbate this problem, so that both ASML and suppliers may not be able to find enough talent locally. If talent is sourced from abroad, then housing could become a problem.[25,26]

Still, overall it seems like many of the barriers to a fast geographical expansion could be solved with enough determination and resources, using the following strategies:

- **Expanding close to home:** even if ASML cannot expand in Veldhoven directly, it seems plausible that an expansion close to home would be possible. Expanding in the wider Eindhoven area could be possible, and there is still the option to expand into the northern Netherlands, or into close parts of Germany and Belgium. This would not address potential problems with employee travel and company culture, but international expansions by high-tech companies happen frequently, so it is unlikely that this would be a hard limiting factor.

- **Hiring internationally:** if the (southern-) Dutch labor market is exhausted, ASML and its local suppliers would have to double down on hiring internationally. This may be a challenge for company culture and the local communities, as employees from various cultural backgrounds would have to be integrated into company processes, and the surrounding cities and villages would have to absorb an inflow of newcomers. Again, this would be a challenge, but does seem solvable through foresight and determination.

- **Getting support by the government:** it seems likely that the Dutch government would go to lengths to retain the geopolitical leverage provided by ASML. This could result in administrative fast tracks for ASML and would generally put opponents of an extreme expansion at a disadvantage. Moreover, it seems likely that many Dutch municipal governments would be delighted to accommodate ASML, and the prosperity this would bring. The Eindhoven region in particular has aspirations to be a "European Silicon Valley", so it would be especially accommodating to a high-tech industrial expansion.[27]

Overall, a geographical expansion would lead to significant challenges to ASML. We find it plausible that none of these challenges would pose a roadblock so severe that it would make a *10 y / 10x* expansion extremely unlikely. Most of the counterarguments seem to be of the form "this would be very cumbersome and involve extra efforts and require extra planning and care" but none of them seem like clear showstoppers.

To us, the finely-tuned system of supply and assembly – which requires lots of ad-hoc attention and problem fixing – seems like the most compelling argument why a *10 y / 10x* expansion could be unlikely. It is unclear how easily this type of close-knit manufacturing arrangement can be replicated ad-hoc elsewhere, as it may be based on tacit knowledge distilled from years of iteration between the local companies involved. Still, this point is somewhat opaque to us, and we find it plausible that we're overestimating how important it really is.

On balance, we're left with a weak, preliminary view that, despite the significant challenges associated with it, a geographical expansion would not be the most critical bottleneck to a *10 y / 10x* scenario.

B.2.5. Expanding the workforce

How many more people would be needed? Manufacturing of EUV machines is a labor-intensive process. It is not mostly automated like the car industry, but much more similar to other low-volume high-tech industries, like the aerospace industry. Hence, a *10 y / 10x* expansion would – at least initially – result in a massive hiring campaign across various departments.

[23] Expert 1 [10.05.23]

[24] Expert 2 [11.05.23]

[25] In 2018, the Veldhoven location had 12,000 employees. If the location only grows 5x, then ASML alone would make up >25 % of Eindhoven's population.

[26] Expert 3 [12.05.23]

[27] Conversation with ASML insider

Not all of ASML's 40 k employees would have to be doubled for each doubling of production capacity. The manufacturing division would likely see rapid growth, especially for the people who assemble machines in the clean rooms. In contrast, the product engineering team (i.e. the people who design the machines) would likely not have to grow by much, as their designs are scalable. Here is a non-exhaustive and rough expert assessment of what would have to scale and what would not have to.

Roles which would have to scale proportionally Production supervisors/operators: more production requires more people that execute assembly steps (at least in the absence of automation).[28]

Roles which would have to scale significantly, but less than proportionally Manufacturing engineers: more manufacturing engineers would be needed to 1) plan and design production facilities and 2) streamline production by reducing the number of steps and outsourcing processes to suppliers. However, there would be economies of scale, as some of these activities would be scalable across production facilities.[29] Service engineers[30]: for each additional machine sold, there would have to be additional service engineers taking care of the machine at the client site. Proportional scaling would not be necessary, as multiple machines at the same client could be served by the same team of service engineers.[31] Others: supplier quality management, testing, and program management.[32]

Roles which would scale somewhat Product engineering: scaling up production would not require significantly more effort to improve machine designs.[33,34] Finance/Marketing/other support functions[35]: more people would be required to steer the necessary organizational changes, but there would be economies of scale.

To summarize, the rough expert assessment is that FTE increases would largely be driven by 1) non-engineer, yet highly-skilled technical workers 2) engineers in manufacturing and service. Note that there are three implicit assumptions about production efficiency behind this assessment.

1. The first is that production will not be automated, or at least that automation will play an insignificant role. If automation does become important, then this would probably result in a hiring surge for automation engineers over the first couple of years, followed by a ramp-down of highly skilled manufacturing employees after manual labor is obviated by machines.

2. The second is that machine designs will not be developed further, but that the existing "2023 design" would simply be copied as much as possible. If this is not the case, then demand for engineers manufacturing, and service would be even higher than in the alternative scenario, as lower product maturity would require a stronger engineering presence in the manufacturing process. Increased R&D demands would also require additional hiring in the product design department.

3. The third is that the operating model of the company would not be oriented towards pre-qualification, a setup in which responsibility for functioning machine parts is shifted much more to suppliers relative to the current setup, potentially leading to an overall reduction of manufacturing engineers and operators, although this is very speculative.

How realistic are each of these assumptions?

Automation Automation could be possible in principle, but is not done due to the low production volume of EUV machines.[36] The underlying reason is unclear, but it could be that the product is not yet mature enough for assembly processes to be executed automatically without frequent interruptions for ad-hoc problem-solving. It could also be that the massive capital investment required for automation is not economically viable at the current margin. Both could be true simultaneously. There is anecdotal evidence that ASML failed at automating and scaling a complex internal process, despite spending large amounts of money on it. The process involved multiple company functions, such as finance, sales,

[28] Expert 1 [15.05.23]
[29] Expert 1 [15.05.23]
[30] Expert 1 [15.05.23]
[31] Expert 1 [15.05.23]
[32] Expert 1 [15.05.23]
[33] Expert 2 [11.05.23]
[34] Expert 3 [12.05.23]
[35] Expert 1 [15.05.23]
[36] Expert 3 [12.05.23]

engineering, shipping, compliance, and more. The reason for failure is unclear, but one problem was that reconfiguring the SAP setup turned out to be very complicated.[37] Another corporate finance automation project was scheduled over a horizon of 5 years, and after 3 years, it was unclear whether anything had changed at all.[38] These examples are flawed, as they are largely unrelated to manufacturing, but they hint at the general complexity associated with automation.

Continuous R&D If ASML chooses to "freeze" and mass-produce the EUV design of 2023, this would result in increasing product maturity, and thereby decrease the amount of engineering attention necessary for the construction of each machine.[39] This may take time: DUV took roughly 10-15 years to mature, and EUV may take up to 20 years to mature.[40] At the same time, it seems unrealistic that there would be no changes to the machine at all. Customers would at least demand certain design enhancements, like increased power and efficiency.[41] Hence, even in the copy-paste scenario, there would be some significant R&D effort by ASML.

Pre-qualification In the current setup, ASML takes on much of the burden of ensuring that individual parts of the machine work as intended, from testing and fixing during assembly, up to client-site service. Under pre-qualification, the suppliers would certify that individual parts work as intended, so that ASML's engineering investment could be reduced.[42] If this is a viable option, then it could have implications for workforce composition, but at this point, it is very unclear to me how far away this system is from implementation.

Overall, it is very unclear how important each of the three points will be, as they all feel pretty speculative. A weak conclusion might be: If we assume – contrary to the baseline "manual copy-paste without pre-qualification" scenario – that there will be some automation, some continuous R&D, and some degree of pre-qualification, then we should expect somewhat more demand for manufacturing and product design engineers, somewhat less demand for manufacturing supervisors and operators, and somewhat less demand for supplier quality control people.

Hiring ability As mentioned earlier, ASML is already putting a strain on the local labor market in Eindhoven, so it would have to hire internationally to be able to mount a *10 y/10x* expansion. The expert interviews did not go into detail enough to give a quantitative answer as to whether global labor markets for engineers and highly-skilled technical workers would be large enough to sustain a *10 y/10x* expansion of ASML, especially given that the entire supply chain would be competing for the same talent.

At least there is a qualitative argument why ASML would be able to draw talent from adjacent industries:

- **Talent is transferable from other industries:** for various ASML roles, there are similar roles in other industries. Supplier quality engineers could be drawn from the automotive, aerospace, and medical industries. Manufacturing supervisors and operators could be hired from other industries with complicated manufacturing procedures, especially if these procedures resemble those in semiconductor manufacturing.[43] For example, as the automotive industry moves towards electric vehicles, its manufacturing procedures may become more similar to the semi industry, thereby increasing the global talent pool available to ASML.[44]
- **ASML is a popular company:** ASML is a well-known and popular company, so they would do well in a talent war, relative to many competitors from other industries.[45]

Training capacity Even if ASML manages to hire enough employees for a *10 y/10x* expansion, it still needs to get these workers up to speed in their various domains. Three factors suggest that this could be done surprisingly quickly:

- **ASML already has the ability to train lots of people:** ASML has grown from 20,000 employees in 2018 to 36 k

[37] Expert 3 [12.05.23]
[38] Expert 3 [12.05.23]
[39] Expert 1 [15.05.23]
[40] Expert 3 [12.05.23]
[41] Expert 1 [15.05.23]
[42] Expert 1 [10.05.23]
[43] Expert 2 [11.05.23]
[44] Expert 2 [11.05.23]
[45] Expert 1 [15.05.23]

employees in 2022.(ASML, 2025) They have gotten very good at large-scale training and are consistently improving their training services.[46] Remote training is available for many parts of the organization, including engineering.[47]

- **Many relevant employees can be trained within 1 year:** Many of the employees most relevant for a *10 y / 10x* expansion could be trained in 1 year at most, for example engineers in service manufacturing and supplier quality, as well as manufacturing supervisors and operators.[48]

- **Training-intensive employees would not have to scale as much:** It seems plausible that product design engineers take the longest time to train, as improving the actual machine design seems like the most complicated task at ASML. Perhaps it takes more than 1 year to get a product design engineer up to speed, but as mentioned earlier, a *10 y / 10x* scenario may not necessitate a significant expansion of product design department.

Despite these considerations, we find it plausible that training capacity could be a bottleneck. There is necessarily an upper limit to how fast a workforce can be trained[49], and it is hard to predict how close ASML would come to this limit in a *10 y / 10x* scenario. The limit could be reached when training capacity starts to trade off noticeably against production capacity,[50] as engineers and others have to spend more and more of their time on training people, relative to spending time in production or development. There is also a limit to how much money can be spent to fix this problem – the maximum rate of knowledge transmission for an individual human reaches a natural limit.

Overall, the experts agree that a human capital expansion seems like a significant challenge, but it looks like many of the key difficulties aren't insurmountable. And this assessment doesn't even take into account that, in our scaling scenario, ASML would be able to offer eye-watering salary packages to everyone, from the engineering level to the factory floor. The biggest uncertainty seems to be around training capacity. On first sight, it seems like training should not take that long and that ASML is well-prepared for training many people, but there is a residual uncertainty how much training would really be required, and whether the relevant employees within the company would be able to provide the necessary capacity without undermining regular production.

B.2.6. Supply chain scaling

Key suppliers could fail Between 2016 and 2021, Carl Zeiss and Trumpf SMT were consistently the biggest bottlenecks to ramping up EUV production. During this time, the goal was going from 24 to 40 machines a year.[51] Both suppliers seemed to struggle with quality and repeatability, as well as with hiring and training people.[52] There were also more company-specific reasons for the delay:

- **Carl Zeiss failed to expand capacity:** generally, Carl Zeiss mirrors are known as one of the longest-lead-time investment items.[53] In 2016, ASML went as far as acquiring a stake in Zeiss in order to make sure the planned ASML investments would be possible and to demonstrate its commitment.[54] The exact source of the Carl Zeiss delay is unclear, but it may have to do with product maturity. The manufacturing process involves much process control, and everytime a deviation is measured, a very slow process starts to correct it.[55] Only a few dozen mirrors are produced each year, so there may have not been enough learning cycles to achieve maturity.[56]

- **Trumpf struggled with adapting the laser design:** in this case, the delay was likely caused by Trumpf's inability to keep up with the evolving design specification for the laser.[57]

[46] Expert 2 [11.05.23]
[47] Expert 2 [11.05.23]
[48] Expert 1 [15.05.23]
[49] Expert 3 [12.05.23]
[50] Expert 3 [12.05.23]
[51] Expert 1 [10.05.23]
[52] Expert 1 [10.05.23]
[53] Expert 1 [10.05.23]
[54] Expert 1 [10.05.23]
[55] Expert 2 [17.05.23]
[56] Expert 2 [17.05.23]
[57] Expert 1 [10.05.23]

In the case of a massive expansion, both suppliers may struggle to reconcile the strong demand for their products with their preference to keep production close to home in Germany.[58] They would face similar problems as ASML, in that they would exhaust local and regional labor markets.[59] On the other hand, they may benefit from increasing product maturity, which could enable less meticulous process control for Zeiss and less need for continuous R&D for Trumpf.

Overall, the expert evidence on suppliers does not seem conclusive so far. We find it plausible that key suppliers could be an important bottleneck, but more research is required to settle this point.

Lack of supply chain transparency Even if key suppliers can be brought on board with a *10 y / 10x* scaling scenario, there are some supply chain factors that are currently outside of ASML's control.

- **ASML cannot monitor its entire supply chain:** there are >300,000 components in an EUV machine and ASML has each of them listed in its bill of materials, but it would not be feasible to audit every single source.[60] Moreover, ASML does not have complete visibility of its level 2 to level N suppliers.[61] As a result, some bottlenecks in deeper layers are only detected years after a capacity expansion is initialized and dealt with ad-hoc with high pressure[62], as was the case with a producer of specialized aluminum.[63] After a bottleneck is detected, some investigations into supply chain issues take years to resolve.[64]
- **Suppliers may not know that they are crucial to ASML:** there are many suppliers deep down in the supply chain that produce things like electrical connectors, gas piping, valves, ceramics, and silicium blocks.[65] Many of these players may not have a good overview of how the semi market is developing, as they are way down in the supply chain and their products are being used across industries.[66] Due to its low production volume, ASML may be a minor customer to these suppliers next to companies with high throughput. This may lead to bad coordination between ASML and its lower-layer suppliers. For example, a supplier may phase out, or "end-of-life" a product without telling ASML.[67]

B.2.7. Scaling the organization

A *10 y / 10x* scenario would be accompanied by sweeping organizational changes geared at making ASML's production more scalable. Two pieces of evidence suggest that this may not be trivial:

- **Change management is difficult:** implementing change in an organization is hard and different to merely optimizing existing operations:[68] "A messy scientist may be good at producing output, but bad at scaling his process".[69] In general, there will always be a degree of resistance to changes in ASML's operating model.[70] For example, there have been worries in ASML's leadership team and engineering team about losing ASML's "historical operating mode" due to hypergrowth.[71]
- **Existing change management programs failed:** there are examples of previous change management programs that failed or got sidetracked, for example the previously-mentioned corporate finance automation project.[72]

On the other hand, ASML invests much attention into change management. Very experienced executives are focused on fixing internal processes. For example, the executive VP of an important business line is now fully focused on process

[58] Expert 1 [10.05.23]
[59] Expert 1 [10.05.23]
[60] Expert 2 [11.05.23]
[61] Expert 2 [11.05.23]
[62] Expert 3 [12.05.23]
[63] Expert 2 [17.05.23]
[64] Expert 3 [12.05.23]
[65] Expert 2 [11.05.23]
[66] Expert 2 [11.05.23]
[67] Expert 2 [17.05.23]
[68] Expert 3 [12.05.23]
[69] Expert 3 [12.05.23]
[70] Expert 3 [12.05.23]
[71] Expert 1 [16.05.23]
[72] Expert 3 [12.05.23]

improvement. There is a new department which is employing "business architects" whose full-time job it is to figure out how to scale the organization.[73]

Overall, although it certainly seems possible, we do not expect that change management would be a significant bottleneck in an extreme scaling case. It seems solvable by spending enough money on internal and external experts, who could support the organization in streamlining its processes and preserving its company culture.

Glossary

Advanced Semiconductor Materials Lithography (ASML): The only supplier of advanced EUV lithography machines required to produce advanced logic chips.

Aquifer: An underground geological formation containing groundwater.

Capital expenditure (CAPEX): Investments in assets such as equipment, infrastructure, etc.

CHIPS Act: 'Creating Helpful Incentives to Produce Semiconductors' — one of three components under the 2022 passed U.S. federal statute called the CHIPS and Science Act.

Desalination: The process of removing salt from seawater to produce freshwater.

Enhanced N3 node (N3E): A "3 nm" chip process produced by TSMC with various kinds of updates from the N3 process. Note that all modern process technology nodes do not have relevant dimensions of the size their name implies.

Extreme Ultraviolet Lithography (EUV or EUVL): An optical lithography technology used in semiconductor device fabrication to make integrated circuits (ICs).

Fab: A semiconductor fabrication plant (commonly called a fab; sometimes foundry) is a factory for semiconductor device fabrication. See previous chapter for a more detailed definition.

Gross world product (GWP): The total economic output of the world, including all goods and services produced globally. Also referred to as "global GDP".

Megawatt (MW): A unit of power measurement equal to one million watts, commonly used in the context of industrial power generation and consumption.

Nanometre (nm): A nanometer is one billionth of a meter (M), also expressed as 0.000000001 or 10^{-9} meters.

Node: The technology node (also process node, process technology or simply node) refers to a specific semiconductor manufacturing process and its design rules. Different nodes often imply different architectures.

NXE:3300B: A type of EUV lithography system developed by ASML to enable chip manufacturing at the 7 and 5 nm nodes.

NXE:3600E: A new generation of EUV lithography systems developed by ASML to enable chip manufacturing at the 3 nm node and beyond.

Organization for Economic Co-operation and Development (OECD): An intergovernmental organization with 38 member countries, founded in 1961 to stimulate economic progress and world trade.

Production bottleneck: The maximum possible limit of a parameter, such as performance or speed, constraining the fabrication process.

Production capacity: The maximum quantity of microchips or semiconductor devices a manufacturing facility can produce within a specified time frame.

Science, Technology, Engineering, and Mathematics (STEM): Core disciplines that collectively cover a range of scientific and technical knowledge.

[73] Expert 3 [12.05.23]

Seismic activity: The occurrence of ground movements, including earthquakes and tremors, resulting from the shifting of Earth's tectonic plates.

SEMI: An industry association comprising companies involved in the electronics design and manufacturing supply chain.

Taiwan Semiconductor Manufacturing Company (TSMC): A Taiwanese multinational semiconductor contract manufacturing company with $\sim$90 % market share in $<$ 7 nm chips.

Transmission capacity: The ability of an electrical power grid to transmit electric power.

Wafer: A crystalline silicon disk serving as the base of integrated circuits.

Wafers per Month (WPM): Measure of production capacity; 1 kWPM = 1,000 wafers per month, 1 MWPM = 1 million wafers per month. Note that this refers to wafer starts per month, i.e. it counts wafers that do not make it through the process due to defects. This difference does not matter for most cases in our report, so we omit this qualifier.