Selling Shovels in the Gold Rush:

My Journey From Particle Physics to the Semiconductor Sector

Tanmay Mudholkar
(PhD Physics, Carnegie Mellon)
FSPA Career Seminar, 2026-04-29

Outline

1. Overview of Career Paths
2. My Job Search
3. The Semiconductor Industry and my Current Job
4. Recommendations

Outline

1. Overview of Career Paths
2. My Job Search
3. The Semiconductor Industry and my Current Job
4. Recommendations

Career Path 1: Software/ML Engineering

- Typical Role Titles: “SDE”
- Skills Needed: Coding (python, also might depend on role, e.g. C/C++), “dev ops” (github workflows/automation/CI/CD, server maintenance etc.)
- Competition: CS PhDs/masters/undergrads, grads from other technical fields
- Hiring tends to be meritocratic and based on performance in coding interview
- Preparation:
  - Introduction to Algorithms AKA “CLRS” (Cormen, Leiserson, Rivest, Stein)
  - Cracking the Coding Interview (Gayle Laakman McDowell)
  - Leetcode Medium/Hard → Hard
- Your advantages: practical experience with large-scale code deployment (CMSSW!)
- Your disadvantages: no rigorous, structured acquisition of CS terminology or knowledge of data structures

Career Path 2: Quantitative Finance

- Typical role titles: “QR”, “Quant Dev”, “Quant Analyst”, “Quant Trader”
- Skills needed: statistical modeling/simulation, coding (typically Python)
- Competition: PhD Stats/ML and other technical fields
- “Academic Pedigree” counts, and interview performance matters too
- Preparation:
  - “Heard on the Street” (Timothy Falcon Crack)
  - “A Practical Guide to Quantitative Finance Interviews” (Xinfeng Zhou)
  - “Quant Job Interview Questions & Answers” (Joshi, Denson, Downes)
  - “Elements of Statistical Learning” (Hastie, Tibshirani, Friedman)
  - Leetcode Medium → Medium/Hard
- Your advantages: familiarity with prob/stats thinking and techniques
- Your disadvantages: particle physicists tend to have deep knowledge of one or two techniques; quant roles require broader knowledge

Career Path 3: Data Science

- Typical Role Titles: “Data Scientist”, “Data Analyst”, “Data Engineer”, “Research Scientist”
- Skills needed: statistical modeling/simulation, data pipelines (python, SQL), ML architectures (simple models + RNN, CNN, LSTM)
- Competition: masters/PhD students from technical fields
- Employers will pay attention to your CV
- Preparation:
  - Intro to Machine Learning by Andrew Ng on Coursera
  - “All of Statistics” (Larry Wasserman)
  - Leetcode Medium → Medium/Hard
- Your Advantages: automatic good match with the skillset in demand
- Your Disadvantages: unfamiliar nomenclature and “guesswork” required for datasets whose generative processes aren’t well-modeled

Outline

1. Overview of Career Paths
2. My Job Search
3. The Semiconductor Industry and my Current Job
4. Recommendations

My Background

- M.Sc. Physics (5-year Integrated Programme), Indian Institute of Technology, Kanpur
- Ph.D. Physics, Carnegie Mellon University. Thesis title: “Search for Stealth Supersymmetry in Proton-Proton Collisions with the CMS Experiment at the Large Hadron Collider”
- Advisor: Prof. Manfred Paulini
- Dissertation skills: how to estimate systematics for smoothed KDE, max likelihood inference, Cls technique, python/C++ coding, data flow automation, etc.
- Led DQM dev work for CMS ECAL for a few years. Skills: C++ coding, Github workflows/Jenkins, anomaly detection
- Helped with setting up DAQ and DQM on a testbeam at CERN. Skills: computer networks, basic firmware work, automation through rpi.
- Set up a T3 server at CMU and connected it to the grid. Skills: Linux system maintenance, bash scripting

Offer 1: Chem Eng Firm

- How I got in touch: job posting on LinkedIn
- Process: one introductory phone call (technical), one on-site day (one hour interview + 1.5 hours coding session + one hour HR interview)
- Reached out to me 2 weeks after on-site with an offer
- Politely declined

Offer 2: Quant Finance Firm

- How I got in touch: quant recruiter forwarded my details
- Process: one hour technical interview, then 4 hours on site (two 1.5 hour coding sessions, then one hour HR interview)
- Reached out to me a few days after on site
- Strongly considered offer, initially accepted it but declined much later (after talking to a couple of potential team members about potential WLB issues)

Offer 3: AMD

- How I got in touch: CMU physics alum's former boss was having difficulty filling an open req, reached out to alum, alum forwarded my CV
- Process: one hour introductory technical interview + three one hour technical interviews all on separate days, all online
- Reached out to me a few days after final interview with offer
- At my request, AMD hosted me in Austin after I already had the offer and I had the opportunity to talk to future team members
- Accepted offer around the same time as Offer 2, joined AMD in October 2022

Outline

1. Overview of Career Paths
2. My Job Search
3. The Semiconductor Industry and my Current Job
4. Recommendations

The Semiconductor Industry

Software

- Drivers, OS kernels etc.
- Test suites
- EDA tools for chip design
- Data Center Infra
- Compilers/toolkits (ROCm, CUDA etc.)
- Data engineering and ”yield analytics”

Hardware

- Process Tool Ownership
- Chip Architecture
- Thermal/mechanical characterization
- Collaboration with external foundries/OSATs
- “Qualifying” specific technologies through DoE

Heterogeneous Integration

Typical Projects at Work

- Yield Analytics: build dashboard that shows “yields” by product, workweek
- Statistical Inference: if fraction f chips fail in condition 1, what’s the expected distribution of f in condition 2?
- Data Engineering: set up a pipeline to share only the aggregated data we want with external suppliers
- Data Science: build predictive model to guess whether defect rate for new design would be sustainable

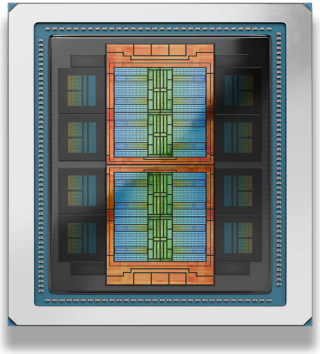

AMD MI350

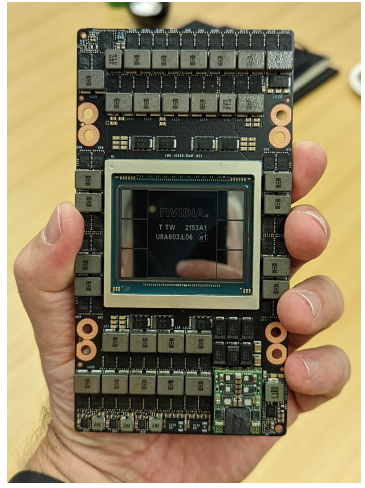

NVIDIA H100

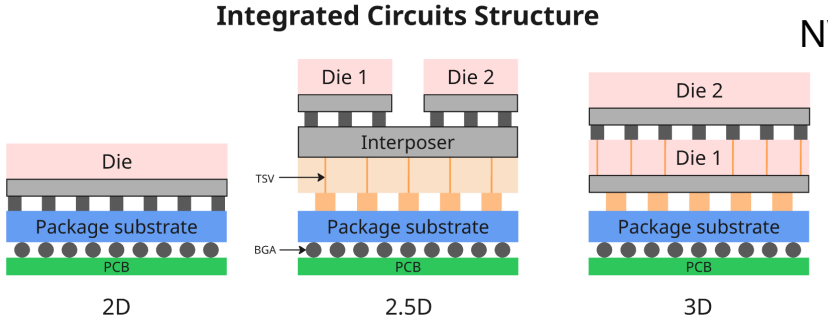

Outline

1. Overview of Career Paths
2. My Job Search
3. The Semiconductor Industry and my Current Job
4. Recommendations

General Recommendations (Technical)

- Use industry standard open-source libraries like `pandas`.
  - Avoid using `ROOT`. Use minimal script to filter data from `MiniAOD`, analyze/plot with industry-standard tools (`pandas+bokeh/matplotlib`).
  - If performance is an issue, use `polars`, numba…
- Prepare leetcode half a level higher than your interview target. (expect medium → prepare medium/hard)
- Familiarize yourself with common statistical methods and learn their names. (“Principal Component Analysis” is often used to estimate systematics.)

General Recommendations (non-technical)

- $<N_{evts}> = \sigma \times \int \mathcal{L}\ \mathrm{d}t$ -- application quality and quantity both help!
- Try to apply through as many means as possible
  - Don't hesitate to reach out to long-lost connections in your network
- Prepare thoroughly for "behavioral questions"
  - Prepare answers to all common questions (including ones outside your direct experience) and practice them with friends
- Prepare multiple versions of your resume, include and order skills/achievements/projects based on job description (LaTeX makes this easy)
- The best negotiating lever is to have more than one job offer
- Do not prematurely decline an offer unless you're sure. There is no upper limit on the number of offers you are allowed to accept at a time
- Follow glassdoor / levels.fyi to get salary information and compare against offer

Thank You!