Magnifying the Engineer: Where AI Fits in Semiconductor Analytics

The AI buildout has made semiconductor manufacturing the constraint on the entire industry. Brad Hopper, VP of Vertical Markets at Spotfire, on what that pressure looks like inside the fab, and where AI helps engineers rather than replaces them.

THREE YEARS INTO THE GENERATIVE AI boom, the limiting factor on the industry is no longer demand. It is the capacity to manufacture. Fabs are running flat out, new ones are being built, and the engineers responsible for yield, quality, and throughput are working under conditions where mistakes are more expensive than they have been in years.

In this Q&A, Brad Hopper, VP of Vertical Markets at Spotfire, explains how analytics and AI are being applied to that problem. With a background in engineering and a passion for problem-solving, he has helped shape Spotfire's visual analytics technology into a powerful tool for industries ranging from semiconductors to oil and gas. Brad has a background in manufacturing, including degrees from UC Berkeley, and worked at AMD as a process and yield engineer. Both gave him unique insights into the potential applications of Spotfire's technology.

Brad Hopper

Q: How is the AI supercycle changing demand, and what do manufacturers need from their data?

Hopper: The last three years have been an incredible ride. Generative AI use cases began spiraling as soon as the first commercial tools launched at the end of 2022, though the semiconductor market itself did not bottom until early 2023. Since then, demand has grown continuously across the semiconductor lifecycle. Especially designers and manufacturers of CPU, GPU, memory, custom logic, networking and power chips have been selling out production for a year or more. The biggest limitation in the market is manufacturing capacity, so, of course, foundries and equipment providers are also growing as fabs are expanded and new ones are built. Historically, demand and capacity oscillate over a cycle of about four years. What makes this one a supercycle is that we are approaching that mark with no downturn in sight.

Manufacturers are operating under intense pressure to ramp new products faster, improve yield and quality, and increase throughput to meet their commitments. Spotfire is used in all of these pursuits. We are an industrial AI analytics company with a special focus on semiconductor manufacturing, and we have grown along with our customers as they use our technology more often and more deeply. They are especially taking advantage of our latest capabilities for high-capacity cloud computing, coupled with industry-native views and algorithms, and, of course, AI.

AI is truly an accelerant here because not only is AI demand driving growth for semiconductor devices in the market, but using AI internally helps manufacturers improve their processes to help them meet that demand. Spotfire's approach is to blend AI into the analysis process to magnify the expertise and creativity of engineers, not replace it; a message that resonates with semiconductor professionals.

> "Spotfire's approach is to blend AI into the analysis process to magnify the expertise and creativity of engineers, not replace it."

Q: Can you say something about who is primarily using this technology and why?

Hopper: First, a couple of words about our customers. The world's leading flash manufacturers, the largest semiconductor foundries, and the dominant high-bandwidth memory makers have been Spotfire customers for many years. These companies are running flat out, trying to produce chips for AI data centers and meet the insatiable demand for more and more powerful AI. As you could imagine, this means things are moving faster, the consequences

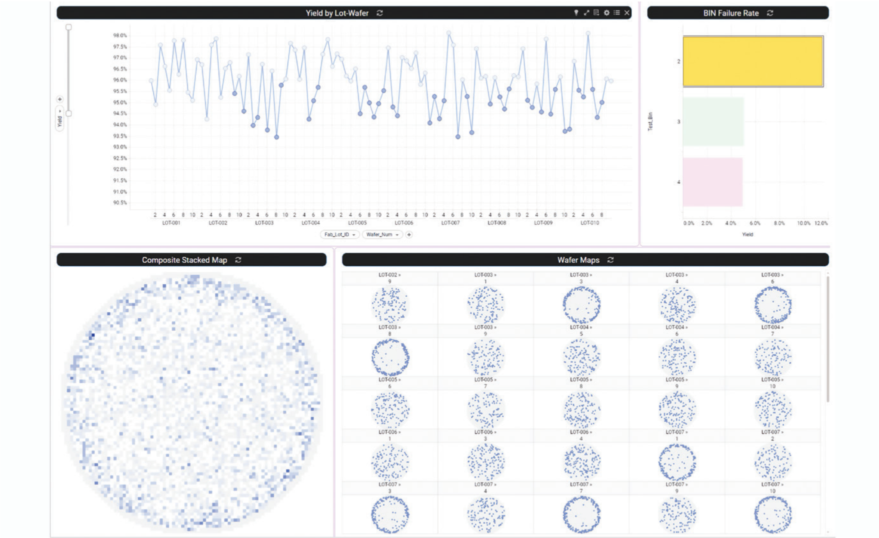

A Spotfire analysis combining yield trends, bin failure rates, and composite wafer maps gives fab engineers a unified view for spotting yield excursions and recurring defect patterns.

of mistakes are higher, and there is a strong desire to better understand and mitigate risk in decision-making.

At the heart of it all are three types of engineers: process engineers, product engineers and yield engineers. Each plays a pivotal role in bringing products from design and raw silicon to devices running in the data center.

Process engineers create the individual deposition, lithography, etch, implant and polish processes that, when sequenced hundreds of times, make up what's called a process technology. The elements of a process technology deeply depend on one another, so there is a complex balance of time-series sensor data and multi-faceted metrology data that have to be developed, characterized, and finely controlled to deliver a robust baseline for building semiconductor products. These engineers work for chip manufacturers, of course, but there are also a large number of process engineers working for equipment manufacturers, developing new processes alongside their customers.

Product engineers own the behavior of the devices themselves: What their specifications should be, and how they perform under different installed conditions such as the supported range of temperature, voltage, and frequency. So devices that need to operate in the cold vacuum of space, versus the hot and rough conditions in an automobile or military application, or that have to perform at ultra-low power in an IoT device out in the field for years, all will have different specifications to align with their use cases.

Product engineers have to bridge from circuit simulation data to actual manufactured performance, adjusting designs to accommodate technology variation, define and monitor appropriate guardbands for new product bins, and optimize testing to reduce cost but ensure quality.

Yield engineers bridge all of the above, helping to troubleshoot when something goes wrong and finding opportunities to optimize yield across products, technologies, and processes to continuously improve productivity and quality. Yield analysis involves the widest array of data types, including facility conditions, process sensor data, metrology data such as lithography overlay and critical dimensions, defect density and distribution, inline electrical test, functional test, advanced packaging, packaged test, burn-in, and device reliability.

For all users, the key capabilities they rely on are blending data of disparate types, finding and exploring correlations, building predictive models and making improvement decisions. Especially important is aligning all of the various data spatially on dimensionally accurate, multi-layer wafer maps.

Q: So what's wrong with the tools they were previously using? Are they not powerful enough?

Hopper: Actually, it's the opposite problem. The semiconductor fab is filled with specialized tools that are very powerful, but it is that very power and specificity that makes them challenging to master and inflexible to bridge across heterogeneous data sources. These tools include Electronic Design Automation, Defect Management, Fault Detection & Classification, Lithography Overlay & CDs, General Metrology, Yield Management, and more. Because the most important problems in the factory usually involve bridging data of different types, engineers resort to spreadsheet exports to get around the specialized file formats and complex schemas of specialist tools. Spotfire is designed for exactly this challenge: To make it easier to access and blend heterogeneous data and reason over it to find insights. Spotfire is a decision layer that adds value to these specialist tools by bringing their data together and making it easier for engineers to solve cross-discipline problems.

Q: Where does AI fit into this picture? How is Spotfire + AI better than AI alone?

Hopper: AI is the catalyst that aligns engineering expertise and powerful analytics capabilities to accelerate decision-making. Current generations of AI are extremely powerful and are now capable of taking complex data and delivering the answer to sophisticated questions. AI can look into the data, write code to implement an algorithm, and draw a graph of a potential answer.

So why is this not sufficient? Typical engineering problem solving in the factory isn't about getting a single answer; instead, engineers must explore the space of answers and see where the trail leads, applying their own expertise along the way. In Spotfire AI, the native capabilities of Spotfire become tools that AI orchestrates and guides the user to understand as they perform analysis together. Here are a few before-and-after examples that might paint a picture.

Industry-native graphics

- **AI alone** can draw a graph of the results of a computation, but this graphing capability may not be optimized to industry-native standards.
- **Spotfire AI** has industry-native graphics built in that have heuristics and defaults that automatically render useful results without burning tokens to refine the visualization.

Live, not static

- **AI alone** produces result graphs that are stand-alone and static. Once drawn, they don't change or support further analysis.
- **Spotfire AI** can configure Spotfire visualizations that are "live." A user can see a subset of results in the visual, draw a complex shape around records of interest and ask, "What's driving this group?"

Certified computation

- **AI alone** is not optimized for computation. Typically, for data analysis, it will write uncertified Python code to execute a statistical function and then apply it to large data.
- **Spotfire AI** draws on certified statistical functions that already exist in the environment, helping the user choose the right one and parameterize it appropriately to the task.

Reusable methods

- **AI alone**, if the user spends time developing proprietary methods, produces work that is not easily reusable outside of that AI session.
- **Spotfire AI** allows building proprietary methods with AI and capturing them into reusable data functions that can be shared and managed centrally.

Verifiable results

- **AI alone** might show the user the top three most likely signals, but each is a stand-alone result, not linked to other analysis content and therefore difficult to verify.
- **Spotfire AI** can configure the analysis environment to link a variable importance table to drill-down visualizations that allow the user to explore the space of answers, each with its own additional visualizations.

Q: How would you sum this up for our readers?

Hopper: In my experience, the answer is usually something ordinary: A chamber that drifted after a parts change, two metrology tools calibrated slightly differently, a lot that took an unusual route through the line. What makes these hard is not the analysis itself; it is that the clues sit in different systems and nobody sees them together until someone goes to the trouble of putting them side by side. That is the part we try to make easier. AI helps by shortening the search and suggesting what to look at next, though the engineer is still the one who decides which thread to pull. If you are looking for somewhere to begin, the analysis your team rebuilds by hand every month is usually a good candidate.

Spotfire in Action

An analog and mixed-signal device maker connects distributions to position on the wafer. Marking an unusual group in a parameter histogram shows immediately where those die sit, and a ranked relationships table points to the process steps most associated with the variation. Following that trail, engineers isolated a single step where one of two tools was running at several times the median of the other, and traced it to a handler chuck that had missed a preventive maintenance interval. The whole sequence took an afternoon, against the weeks the same picture would take to assemble from separate systems.

A manufacturer of storage devices and the semiconductor read heads inside them maps every recorded defect onto the true physical geometry of the part, broken out by unit, by position, and by error code. Seen that way, particular error codes carried a distinct spatial signature that pointed to a chemical mechanical polishing step, and the trail ended at a single batch of slurry. A second template trends yield-relevant parameters week by week, sized by sample count so an engineer can judge significance at a glance. Both are standing analyses rather than one-off investigations, and both compress work that once ran to weeks into a matter of hours.

A fabless memory company validates devices across three foundry technologies at once. Bringing each foundry's test data into a single analysis turned validation cycles that had taken months into weeks, and it exposed an interaction with the process technology of one of them. Fixing a single optical proximity correction feature on one mask brought all three into robust margin.