Maximizing Product Quality and Production Yield in Silicon Wafer Manufacturing

About the Customer

As a leading silicon wafer manufacturer, this company produces wafers in a variety of sizes, including specialty wafers with specific features to meet the particular needs of different customer types. It has a global customer base made up primarily of electronics and semiconductor manufacturers.

The Customer's Need

A silicon wafer is a thin slice of crystalline silicon used in electronics for the fabrication of integrated circuits and in photovoltaics for conventional, wafer-based solar cells. The wafer serves as the substrate for microelectronic devices built in and over the wafer and undergoes many microfabrication process steps such as doping or ion implantation, etching, deposition of various materials, and photolithographic patterning.

Given the minute, delicate, and sophisticated nature of circuits, microchips and the like, the quality of the silicon wafer is critical to the success of those companies which use them as inputs to their own products. So, this manufacturer invests heavily in ensuring its silicon wafers are of the highest possible quality.

At the same time, the manufacturer must be mindful of its production quantities, and in particular its production yield - that is, how many wafers it is able to produce in a given production facility and from a given number of silicon ingots, the form silicon takes before being sliced into wafers.

The Challenge

Skyrocketing demand and quality requirements increase pressure

While the need to maximize product quality and production yield has always existed for silicon wafer producers, industry forces have recently made these issues particularly important. The global supply of silicon wafers is lagging as the growing markets for electronics and solar power drive demand for more semiconductors, in which the wafers are a necessary component. There is even speculation that the semiconductor industry has entered a "supercycle," a period longer than its usual three- to five-year market cycle. If so, wafer supplies could be squeezed further.

This demand is not easily satisfied, however, as it comes with stringent quality and reliability requirements from customers. The widespread adoption of electronic devices in our daily lives - part of the “internet of things” - has increased the production volumes on one hand, while on the other hand has made the semiconductor wafer fabrication process increasingly complex. With billions of parts being shipped by the semiconductor manufacturers of high dependencies, quality requirements are now specified in terms of “defect per billion” instead of the “defect per million” of the past.

And yet, for this manufacturer, the answer is not as straightforward as simply building more silicon wafer production facilities. A semiconductor wafer fabrication plant is one of the most capital-intensive production facilities that exists. Typically, a modern fabrication plant costs more than $3 billion these days. If the production of a new facility is mistimed, it could come online just after demand has peaked and has begun to decline. At the same time, there is stiff competition in the industry, meaning competitive moves, such as on pricing, could impact what demand a particular manufacturer experiences.

These factors make product quality and production yield more important than ever before. Long before the manufacturer can consider adding more fabrication capacity to its business, it must increase production by decreasing loss due to defects and increasing the yield from its existing production facilities.

The Solution

Applying data science to the wafer production process

There are countless opportunities for the manufacturer to use data science to improve product quality and production yield. Many of them involve predicting with great precision exactly how a step in the production process should be executed, to maximize the desired results.

A wafer fabrication facility is replete with extremely sophisticated production equipment that generates large volumes of data that can be analyzed. The equipment also can be adjusted very precisely to change the duration, distance, pressure, temperature, etc., of practically any aspect of the wafer production process. This creates an opportunity for data on past production efforts, and the outcome of those efforts, to be used to adjust future production efforts to create better quality and yield results.

Predicting wafer polishing duration

An example of this is the wafer polishing step. This is typically one of the last processes undertaken before the wafer is ready to be delivered to the customer.

Most prime grade silicon wafers go through 2-3 stages of polishing, using progressively finer slurries or polishing compounds. Wafers are polished either on the front side or double side polished for the largest wafers (300mm). Polishing produces a mirror finish. This surface must be free of topography, micro-cracks, scratches, and residual work damage.

The polishing process occurs in two steps, stock removal and final chemical mechanical polish. The mechanism for material removal is similar to that which occurs in glass polishing: chemical degradation or weakening of the surface film followed by abrasion by mechanical interaction with the particles of a polishing pad. The material removal process removes a very thin layer of silicon and is necessary to produce a wafer surface that is damage-free. The final polish does not remove any material but gives the wafer a mirror finish.

The more precisely the manufacturer can determine how much abrasion will be necessary to achieve the desired result, the less waste of silicon will occur (from overpolishing) and the less likelihood of a defect occurring (from underpolishing).

The manufacturer is using RapidMiner to build machine learning models that predict exactly how much polishing will be needed for each wafer. These models take into account all the data available that might be used for the prediction such as:

- **Data about the wafer:** doping levels (what else besides silicon is in the wafer, and how much), crystal orientation within the wafer, and the desired properties of the final product ("abrasion type")
- **Data about the equipment being used:** the machine type, the polishing pad type, the polish type
- **Data about previous polishing job:** the abrasion type last performed, and other differences between the last job and this job

The RapidMiner team also supported the manufacturer in developing an analytical approach to this problem, as there were a number of intricacies that needed to be worked through. For example, traditional approaches to these types of analyses are based on the study of the effect of a single influence factor. But in this case, the nature of the problem was such that different attributes depend on each other and effect the result. So, with RapidMiner's help, the manufacturer was able to overcome the challenges of a univariate approach and do true multivariate analysis to predict the amount of polish each wafer required.

Also, wafers are produced in batches, where subsequent batches share common attributes. This means the production process yields strong dependencies between individual results, which needed to be considered in the analysis. The RapidMiner team helped the manufacturer build a validation scheme which factored in these issues, and enables the client to get accurate performance estimation, and thus proper hyper parameter tuning.

And RapidMiner was able to help the manufacturer understand the impact of same or different production run configurations on its machine learning models. Since the modelling approach incorporates data on the runtime of the previous batch of produced wafers, the prediction process is easier when the pervious run had the same configuration and the previous runtime could be used as a proxy. Otherwise, it's harder to predict the correct settings - but the yield of machine learning model is also higher.

Other uses of RapidMiner in wafer production

The manufacturer has also started to use data science and RapidMiner with other aspects of maximizing product quality and production yield.

For example, a wafer can have defects attributed to a variety of causes. Some defective wafers can be repaired using a rework procedure, which brings the wafer up to adequate quality levels. The manufacturer has started to use RapidMiner to predict which defective wafers can be repaired, versus which are lost causes. This saves the manufacturer from wasting time reworking wafers that will never meet quality standards.

The Results

RapidMiner improves the manufacturer's wafer production results

By applying data science to the silicon wafer production product through the use of RapidMiner, the manufacturer has been able to increase product quality and production yield. This has helped the manufacturer to stay competitive and achieve better financial results amidst a fast-moving and rapidly-changing marketplace for silicon wafers.