# Abstract

Semiconductor manufacturers routinely test integrated circuits (IC's) while they are still part of the wafer. When the IC failure rate is high, a map of the test results typically reveals clusters of failures. Such 'failure patterns' evince circuit-to-process mismatches and processing problems, providing valuable information for design and process improvement.

We select five failure patterns that resemble simple geometrical 'objects'—Segments, Disks, Annuli, Bands, and Rings—and determine their uniform random distributions through use of the invariant pdf. Two to four geometrical variables govern an object's statistics by locating boundaries that divide the wafer map into zones. The geometry of these zones drives subsequent calculations, giving rise to eleven Shapes. One-zone shapes comprise Pass and Fail; two-zone shapes comprise Segments, A-Disks, B-Disks, A-Annuli, and B-Annuli; three-zone shapes comprise Bands, A-Rings, B-Rings, and C-Rings.

In addition to Shape, relevant features include Area, Location, and—if there are three zones—Orientation and Curve Direction. These four new 'feature variables' are real-valued functions of the geometrical variables, changing continuously as one pattern shape transforms into another.

This formulation enables construction of a synthetic wafer map generator, valuable for creating and studying classifiers. One study employs synthetic datasets with 160,000 maps and estimates the Bayes error of the synthetic population to be less than 1.5 percent. Other studies employ several classifiers, chief among which are nearest neighbor classifier N20 and prototype classifier kh, created with synthetic maps, and nearest neighbor classifier Xdata, created with industrial maps. These classifiers are evaluated using input sets of type `RAND`, containing synthetic maps, and `DATA`, containing industrial maps.

Classification of RAND reveals that N20 is the more accurate Shape classifier, scoring 95 percent, while kh is more accurate with Area and Location, at 96 and 94 percent, respectively. Xdata accuracies are 85, 91, and 83 percent. Classification of RAND plus bit noise demonstrates that N20 is affected less by noise than either kh or Xdata. Classification of DATA finds the kh classifier to be somewhat better than either N20 or Xdata with Shape, Area, and Location accuracies around 67, 77, and 75 percent, and with Xdata accuracy comparable to that of N20.

# Preface

This dissertation takes an old problem and looks at it in a new way. Yield modeling, diagnosis, and enhancement are problems that are as old as semiconductor manufacturing. The traditional approach is to use a variety of techniques from statistics. My approach has been to look for new constructs in probability with the goal of endowing a classifier with the ability to describe a wafer map failure pattern.

To illustrate the value of this approach, please allow me to tell you a story. I worked for almost five years at Hewlett-Packard/Agilent as a product engineer attached to a small silicon fab and its probe test facility. One of my products was a mixed signal communications IC that went into a transceiver module. This IC had to deliver more current than a specified minimum in order for the module to work, and the current it delivered was roughly proportional to the $\beta$ of the bipolar transistors.

Median $\beta$ was above 60 for all wafers manufactured with this bipolar process, but $\beta$ was not constant across any wafer. IC's at the center of the wafer had roughly a 20% smaller value of $\beta$ than IC's at the wafer's edge. This radius-dependent variation is a consequence of the physics of ion implantation. A wafer with median beta equal to 50, for example, would typically contain IC's with $\beta$ between 45 and 55. Wafers with central $\beta$ values of 40 or more were given a passing grade for $\beta$ and were sent to probe test.

In the case of my mixed signal IC, it could deliver the minimum current provided it was made with $\beta$ above 48 or so. During probe test, any IC that could not deliver the minimum current was marked with an ink dot, or 'inked', to prevent it from being packaged. Occasionally, when wafer-center $\beta$ would fall as low as 40, upwards of half the

IC's would be inked for insufficient output current. The resulting low yield would then be brought to my attention.

'Low' yield is relative to whatever happens to be average and does not necessarily imply anything beyond immediate revenue loss. However, some failure modes, such as those caused by shorts or leakage—either source-drain or emitter-collector—indicate the potential for more widespread problems. One of my responsibilities as product engineer was to hold wafers at probe test if there were any indications of more widespread problems. Whenever I saw a centered radial inking pattern associated with low current output, I could predict that the failures were caused by low $\beta$ and then reject or verify this hypothesis using process measurements. Alternatively, if the failing IC's were not grouped in the center, or if the process tests did not verify low $\beta$, then I would hold the lot for diagnosis or arrange for additional reliability testing of the finished modules.

While I was at Agilent, we instituted automatic sign-off procedures in order to conserve engineering resources. The implementation of automatic sign-off requires that the engineering team establish ship-or-scrap criteria for each product. Wafers that meet the criteria are sent to vendors to be sawn into dice, while those that fail are scrapped. Engineering judgement has no further role other than to review the criteria quarterly.

We were told that auto-sign-off criteria were to be based exclusively on the probe test statistics associated with each product and each failure category. In the case of this current test, the statistics dictated that wafers be scrapped when 35% of the dice failed for low current, yet wafers could be sawn up profitably even if 80% of the dice had failed. I had seen situations in which working IC's were desperately needed, low $\beta$ had caused failure rates above 35%, but I was able to prove that the wafers were usable. Under automatic sign-off, these wafers would be needlessly scrapped. Knowing the product as I did, I wanted the option to specify more detailed automatic criteria. I wanted to have a knowledge system that could evaluate the spatial distribution of the failing IC's and could compare probe-test measurements of output current with process measurements of $\beta$. The results of research described in this dissertation make it possible to construct such a system.