

Jeff Yung-Choa Pan and Jay M. Tenenbaum

PIES: An Engineer's Do-It-Yourself Knowledge System for Interpretation of Parametric Test Data

Introduction

This report summarizes our experience in building PIES, a knowledge-based system that diagnoses problems in semiconductor fabrication processes by analyzing parametric test data.

Parametric measurement, which is performed on test circuits at the end of a complicated semiconductor fabrication process, provides semiconductor engineers with early information to monitor the "health" of the overall fabrication process. Typically, hundreds of measurements are made on each wafer. The problem is to reduce the resulting ream of data to a concise summary of the process status: whether the process is functioning correctly and, if not, what the nature and cause of the abnormality is. Currently, this interpretation task is performed by a group of semiconductor specialists known as failure-analysis or yield-enhancement engineers and routinely consumes a large portion of their time. It is critical that problems be identified quickly to avoid a major operational loss.

For any knowledge system to be effective in this application, it must be able to deal with two common characteristics of an engineering domain: (1) knowledge about the domain matures progressively with experience, following a learning curve, and (2) the process sequence is subjected to continual modification. These characteristics entail ongoing maintenance of the knowledge base. Unfortunately, it is impractical to use highly trained artificial intelligence (AI) professionals for this continuing support function. The PIES approach to this problem is to provide a knowledge-acquisition environment that permits the failure-analysis engineers themselves to build up and maintain the actual contents of the knowledge base. The traditional AI knowledge engineering

task has been reduced to initially analyzing the domain and defining an appropriate structure for the knowledge base.

The structure of the knowledge base reflects the way fabrication engineers reason causally about semiconductor failures. First, measurement deviations are used to infer physical defects of wafer structure, such as the thickness or doping density of some layer being too high. These structural anomalies are then linked to problems in particular process steps; for example, a wafer layer might be too thick because the wafer was left in an oven too long or because the oven temperature was too high. Finally, process problems are traced to root causes; for example, the wafer was left in the oven too long because a timer broke.

The multilevel causal structure of the knowledge base permits fabrication engineers to codify their knowledge of and experience with failures of a fabrication process in a form they find natural: causal links that associate evidence at each level with hypotheses at the next level. Thus, there are associations linking deviated measurements to structural anomalies, anomalies to process problems, and process problems to root causes. A knowledge editor supports and enforces this conceptual structure.

The structure of the knowledge base also helps focus the diagnostic reasoning process by providing natural, inter-

Abstract The Parametric Interpretation Expert System (PIES) is a knowledge system for interpreting the parametric test data collected at the end of complex semiconductor fabrication processes. The system transforms hundreds of measurements into a concise statement of the overall health of the process and the nature and probable cause of any anomalies. A key feature of PIES is the structure of the knowledge base, which reflects the way fabrication engineers reason causally about semiconductor failures. This structure permits fabrication engineers to do their own knowledge engineering, to build the knowledge base, and then to maintain it to reflect process modifications and operating experience. The approach appears applicable to other process control and diagnosis tasks.

Jeff Yung-Choa Pan is a computer scientist at Schlumberger Palo Alto Research, 3340 Hillyview Avenue, Palo Alto, California 94304. Jay M. Tenenbaum is the director of the Artificial Intelligence Laboratory at Schlumberger Palo Alto Research.

mediate levels for hypothesizing and verifying. Usually, there are many root causes that could account for an observed set of parameter deviations. Instead of directly associating measurements with root causes, it is computationally efficient to proceed step by step, hypothesizing and prioritizing or ruling out possibilities at the structural and process levels. In addition to being efficient, this multilevel diagnosis leads to explanations that fabrication engineers find easy to comprehend.

A working knowledge-based system incorporating these concepts was implemented in Franz Lisp on a VAX/Unix system at Schlumberger Palo Alto Research (SPAR). This core system was then installed at Fairchild's fabrication facility in Puyallup, Washington, running on a VAX under VMS. The knowledge base was compiled and is maintained solely by failure-analysis engineers at the production site. Performance of the system is currently being evaluated.

Background

In the following two subsections, we present a brief discussion of semiconductor fabrication and parametric testing, and shallow-level versus deep-level approaches to expert systems.

Semiconductor Fabrication and Parametric Testing

Semiconductor devices are manufactured in two phases, as shown in figure 1: Wafers are first fabricated in batches (known as *lots*) in the controlled environment of a clean room; the wafers are then cut into "dice," which are individually packaged and tested. Parametric testing is performed on lots at the conclusion of the fabrication process, just before the wafers are cut.

The recipe for a modern semiconductor product typically contains more than 100 process steps. Each step is a chemical-physical interaction between a wafer and its environment under the precise control of process equipment, for example, epitaxy, oxidation, etching, and ion implantation. Although the result of each individual process step is monitored by a so-called in-process test (such as measuring the thickness of an oxide layer) to make sure that it is within tolerance, the combined effect of these process steps cannot be verified until the recipe has executed completely, hence the need for parametric testing.

When abnormal measurements of some key parameters are detected, the wafer is rejected and is sent for failure analysis accompanied by a complete test record of the lot. The job of the failure-analysis engineer is to diagnose the process step(s) responsible for the failure and take appropriate corrective action. The daily work load of a failure-analysis engineer thus depends on the number of rejected wafers during the previous day and the difficulties of those cases, each of which takes tens of minutes to hours to diagnose. A

knowledge-based system such as PIES can enhance the productivity of a failure-analysis engineer in two ways: first, it focuses an engineer's attention by reducing the flood of raw test data to a few likely failure candidates; second, it ensures an objective analysis by providing a complete and unbiased assessment of the situation.

Semiconductor fabrication was selected as a good experimental domain in which to pursue our long-term interest in applying AI technology to manufacturing. The choice was based on a number of considerations. First, there is high leverage: because of the high volume (millions of die a year), small-percentage increases in yield can result in considerable increases in profit. Second, the processes are not always well understood, so that actual operating experience is critical to achieving acceptable yields. It is important to be able to codify this experience so that it can be widely replicated and shared. Third, semiconductor fabrication is an ideal domain in which to pursue AI research on qualitative modeling and reasoning. Due to the ever-changing nature of fabrication technology, a knowledge system that is totally dependent on hand-coded, process-specific, task-specific, experiential knowledge is inefficient to maintain and difficult to generalize. Moreover, semiconductor engineers routinely invoke models of solid-state physics and silicon processing to explain a problem not encountered previously. To achieve the same level of competence as a human engineer, we set as a long-term goal the development of qualitative modeling and reasoning techniques that can supplement the PIES experience-oriented knowledge base.

A Shallow-Level Versus a Deep-Level Approach to Expert Systems

A conventional way to build an expert system for diagnosing process faults is to rely on a knowledge engineer to capture the experience of fabrication engineers in the form of if-then or production rules (Davis et al. 1975). An inference mechanism might then use a forward-chaining inference process (Winston 1984) to transform an input set of parametric symptoms into a set of possible faults. The approach so described is sometimes referred to as a shallow-level approach (Hart 1982), because its knowledge base records only aspects of experience acquired from human experts and not a model of the domain about which the system is supposed to be an expert. An alternative, deep-level approach would be to perform diagnosis by reasoning with causal models of the domain (Pan 1983).

A shallow-level approach is suitable when experience and not the exercise of theory plays the key role in performing a task. For a fixed problem, a shallow system can be built in a relatively short time and can be "tuned" to a high level of performance, as demonstrated by MYCIN (Shortliffe 1976). However, a shallow-level system will require re-engineering of its knowledge base whenever there is a change in the domain.

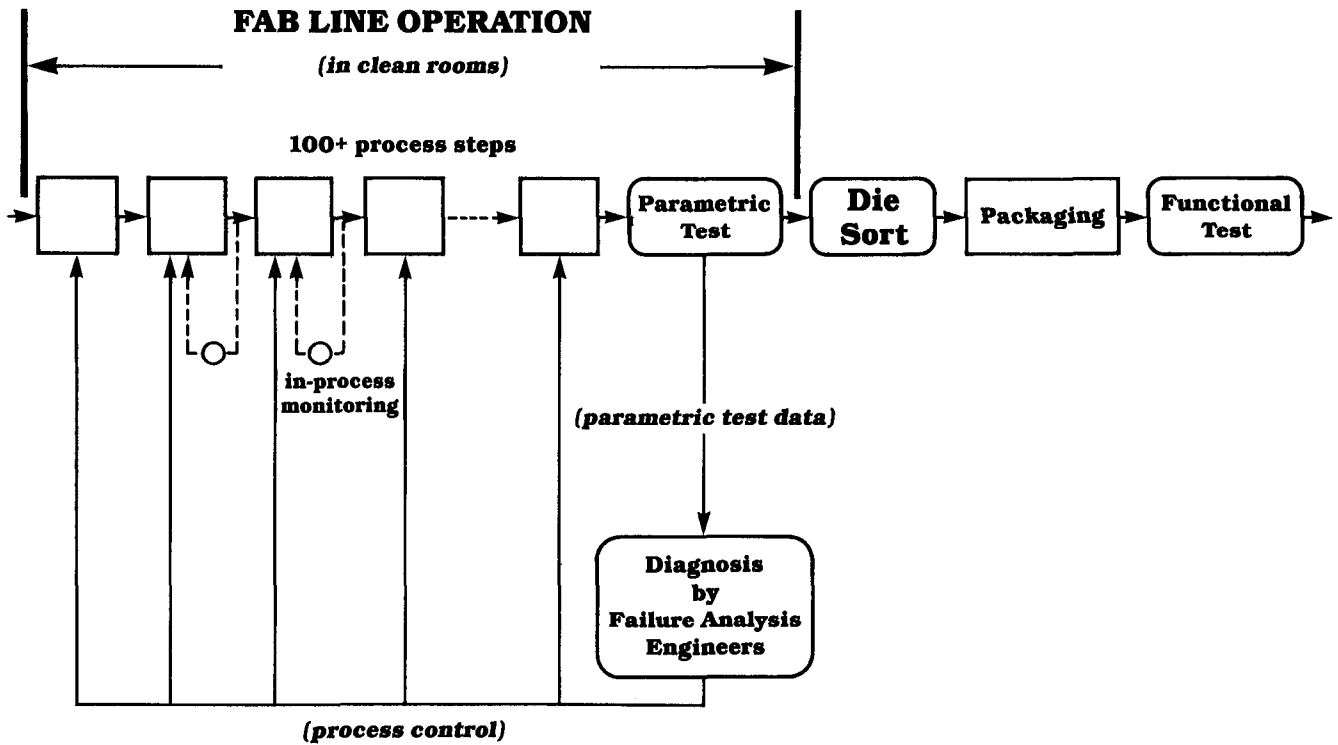


Figure 1. A Typical Semiconductor Manufacturing Process.

The deep-level approach complements the weakness of the shallow-level system because of its potential to derive solutions for unanticipated situations from the underlying principles of the domain. It is particularly advantageous in engineering-oriented domains where a complete or partial domain theory already exists. The progress made in the direction of qualitative modeling and reasoning (Forbus 1984, DeKleer 1979, Kuipers and Kassirer 1983, Pan 1983) is promising, but the technique needs further development before it can be useful in practice.

The PIES knowledge base approach falls between the shallow and deep-level approaches (semideep). It is similar to a shallow-level system in that it attempts to help domain experts in formalizing their experience and to apply the knowledge so acquired in diagnosis. However, it explicitly represents the structure of the domain in terms of multiple causal levels and uses such conceptual levels to communicate naturally with domain experts (in both knowledge acquisition and diagnostic reporting).

Approach

Overview

Figure 2 shows the causal chain through which fabrication failures originate and propagate. The root cause is either a malfunction in some fabrication equipment, contamination

in the source materials or clean-room environment, or human error. Any of these causes will result in variations in the fabrication process, which, in turn, will produce physical abnormalities in the wafer structure and corresponding deviations in parametric measurements associated with the structure. The PIES diagnostic approach is to isolate the possible causes of observed symptoms by reversing this causal chain level by level, following the sequence of measurement deviations, physical-structure abnormalities, process variations, and root causes.

The knowledge base in PIES consists of four levels that correspond directly to those in figure 2. At each level, we enumerate observed failure modes. For instance, at the physical-structure level, such modes include incorrect thickness or doping density of particular wafer layers, such as the epitaxial layer. At the fabrication process level, the failure modes include incorrect temperatures or gas densities during particular process steps, for example, oxidation or ion implantation. Rules provided by the fabrication engineer link failure modes at adjacent levels. Thus, EPI-thickness-high is associated with abnormally high temperature during the epitaxial process stage.

Fabrication engineers often find it convenient to organize their knowledge around specific failure cases, each corresponding to an observed or expected anomaly in physical structure. Associated with each such structural anomaly are a set of expected symptoms (that is, measurement deviations) and a set of possible causes (that is, process failures).

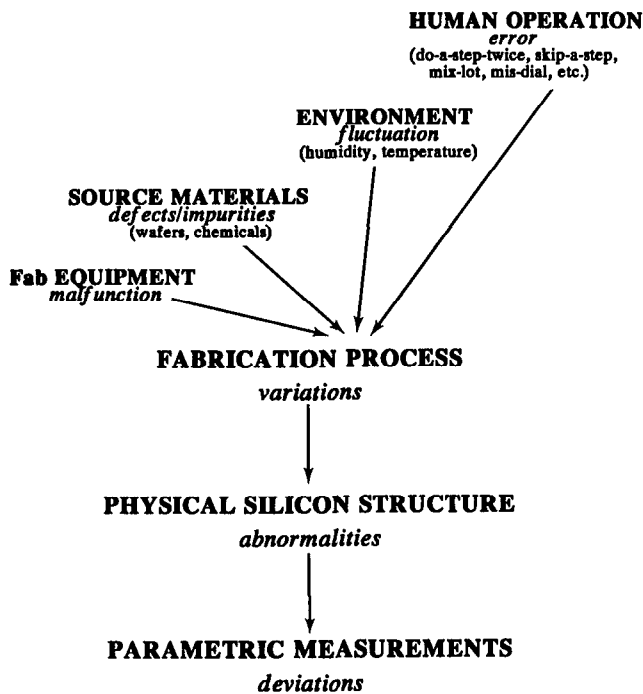


Figure 2. Multilevel Causal Structure of Fabrication Diagnostic Knowledge.

Diagnosis proceeds as a multilevel hypothesis-verification process. Parametric measurements are first pre-processed to transform them from numeric values to qualitative ranges, for example, normal, high, very high. Each measurement that is abnormal implicates one or more physical-structure problems. The expected symptoms associated with each of these hypothesized physical-structure problems are compared with the complete set of abnormal measurements. A score is assigned corresponding to how well the expected symptoms match the observed ones. The scores are compared, and hypotheses with significantly lower scores are eliminated from consideration. The same hypothesis-verification process is then used to select the most probable process failures based on the surviving structural problems. Finally, the root causes are selected that best explain the most probable process failures. This iterated hypothesis-verification approach identifies the primary, that is, most likely, failures. In many cases, it also reveals multiple failures that might be independent of, or causally related to, the primary failure.

The PIES knowledge editor makes it possible for a fabrication engineer without AI training to build and maintain the knowledge base. It does this by directly supporting the PIES multilevel case-centered knowledge organization, thereby guiding an engineer to decompose knowledge in a way that is both natural and required by PIES. Using the editor, an engineer can focus on the failure cases at any level and can create or delete cases as well as their associational links to other cases at the same or adjacent levels in the causal chain. For

example, having discovered a new type of physical-structure failure, the engineer can add the failure to the knowledge base along with the expected symptoms and probable causes.

Knowledge Base

The top level of the PIES knowledge base is organized into four explicit causal levels: measurement, physical structure, process, and root cause. As part of the representational mechanism in PIES, the causal sequence among these four levels is described by a set of symbolic links that are used by both the knowledge editor and the diagnostic reasoner.

At each causal level, the knowledge base is decomposed into framelike structures, called *failure cases* or *cases* for short, each encoding knowledge about a type of failure at that level.

The cases have *slots* for encoding attributes that describe a particular type of failure. Examples of such attributes in the current implementation of PIES are the “popular” name commonly used by domain experts to refer to a failure case, comments from fabrication engineers about the failure, and most significantly, four associational link types that describe how this case is causally related to other failure types. Other slots are used in conjunction with the knowledge base editor (described below) to group failure cases in ways that users find convenient.

A domain expert’s knowledge about possible causal connections between two failure types is represented in PIES by associational links. A link can be one of two types: causes or caused-by, which are further distinguished between intra-level and interlevel, depending on whether the other failure case it refers to is at the same or a different causal level. Each associational link has an associational strength, which is a heuristic estimation of the strength of the causal relationship, and can be one of five quantized states: must, very-likely, likely, probably, and maybe.

As an example, a common failure in a bipolar ISO-Z process at the physical-structure level occurs when an ion-implantation problem alters the distribution of doping in the base region of a transistor. The PIES representation for this problem, known as BASE DISTRIBUTION deep, is shown (in its pretty print form) in figure 3.

In this example, the failure type of BASE DISTRIBUTION deep is said to be causally related to other failure types at the process level, the measurement level, and the physical-structure level itself. If this failure it occurs, it can result in seven types of measurement deviation, some of which are more likely to manifest themselves. For example, WE-10BETA is more likely than RB1.

Knowledge Editor

The knowledge editor enables domain experts to build and maintain the PIES knowledge base without on-site help from AI specialists. Acquiring knowledge directly from domain experts has several advantages in practice: (1) it relieves AI




Figure 5. Organization of Concepts Causally Related to BASE DISTRIBUTION deep.

(process level), and self level (physical- structure level). The editor allowed our collaborator to add, delete, or replace associational links as necessary. Our experience showed that failure-analysis engineers with no AI background were capable of mastering the PIES knowledge editor after a brief (less than an hour) tutorial session.

Diagnostic Reasoner

The PIES diagnostic-reasoning mechanism exploits the multiple causal-level structure of the knowledge base to diagnose the root cause of failure from a given set of parametric test data. Before actually starting the diagnostic process, symbolic “symptoms” have to be abstracted from raw test data (in this experiment, the raw data were recorded by Fairchild’s Keithley tester). The symptom abstraction process follows two steps: first, noisy data points (due to bad test-probe contact or random failure) are removed from the data set by a statistical method; then, a statistical average and a standard deviation are computed for each parametric measurement over all wafers in a given lot. This information is compared with expert-provided limits to produce a qualitative estimation of the measurement, for example, EPI-R very-low. The resulting qualitized measurements form the initial symptom set.

The diagnostic process is performed by progressing level by level through a sequence of hypothesization and confirmation steps, as explained in the overview. At each level, a set of probable failures is filtered from initial hypo-

theses suggested by the likely faults isolated at the previous stage of reasoning (or the initial symptom set). The level-to-level isolation cycle repeats itself, following the inverted causal chain, until it reaches a final diagnostic conclusion at the root cause level.

Let us follow through an example of this reasoning chain. EPI-R is a measurement of electric resistivity from a test structure within a layer of epitaxial material. (It is designed to monitor the result of the epitaxial process.) One possible explanation for an observed low EPI-R measurement, which readily follows a basic principle of semiconductor physics, is that the EPI layer was too thick—a physical-structure failure directly confirmable by other expensive, time-consuming material-analysis techniques. Tracing further back along the causal chain, we find a thick EPI layer can result from, among other factors, an abnormally high temperature during the EPI process. The final step is to identify possible root causes of this failure, which leads to, among others, a faulty thermostat—an equipment failure—that resulted in a higher than normal EPI process temperature.

At each stage of the level-to-level diagnosis, the isolation of failures from hypotheses at the previous level is achieved in four steps: hypothesization, implication, confirmation, and thresholding.

The hypothesization step is designed to heuristically retrieve from among all known types of failures a *suspect set* that includes only those failure cases which are “reasonably” implicated by given symptoms. The *sensitivity*, that is,

how strong the evidence has to be for a hypothesis to be included in the suspect set, is an adjustable threshold.

The suspect set so derived is by no means exhaustive; a potential failure might not be included because the symptoms stipulated for hypothesizing the failure are not observable from the given test circuit. A reasoning step, known as implication, expands the original suspect set by including additional hypotheses that are implicated by any failure case already included in the suspect set. Such implication is based on the intralevel causalities coded in the knowledge base. For example, one intralevel causal link coded in the ISO-Z knowledge base indicates that the physical-structure failure BASE-OXIDE THICKNESS low is a potential cause of another physical structure failure, BASE DISTRIBUTION deep (as shown in figure 5). However, base-oxide thickness is not directly monitored by any ISO-Z test structure; therefore, BASE-OXIDE THICKNESS low can only be included in the suspect set through the implication step after a failure it might cause, (for example, BASE DISTRIBUTION deep) has been hypothesized.

In the confirmation step, expected symptoms of each failure case in the suspect set are matched against the failure hypotheses concluded thus far in the diagnostic process. The matching process computes a "score" for each failure case, indicating how close the case's expected symptoms match the conclusion derived from the given measurement data.

Following the confirmation step, the failure cases in the suspect set are sorted according to their matching scores. Thresholding is done to exclude those failure cases which have relatively low scores. The remaining suspect set serves as the system's diagnostic conclusion for the current level and is passed on to the next stage of reasoning.

Results of the PIES Experiment

The PIES experiment was conducted in three stages: knowledge base construction, system tuning, and performance evaluation.

With the PIES knowledge editor installed in the Fairchild Puyallup production environment, a knowledge base for diagnosing the Fairchild ISO-Z bipolar process was constructed by failure-analysis engineers on site. In the resulting ISO-Z knowledge base, 342 types of failure cases were identified: 101 failure types are associated with the measurement level, 82 with the physical-structure level, and 159 with the process level. The knowledge base also encodes about 600 associational links among the identified cases. Today, the knowledge base is competently maintained by Fairchild's failure-analysis engineers.

The performance of PIES was evaluated by analyzing parametric test data from problem lots that represent a fair sample of challenging cases encountered and recorded during the production history of the ISO-Z process. For each case of lot data tested, the PIES diagnostic result was compared with the recorded conclusion reached by failure-analysis engineers at the time of its occurrence.

Initially, diagnostic results from only 10 of the 25 cases tested were judged to be satisfactory by experts. The major reason for these unsuccessful diagnoses was, not surprisingly, missing knowledge in the PIES knowledge base. The problems were subsequently corrected by Fairchild engineers with a modification of the knowledge base using the PIES knowledge editor. After this initial system tuning, correct diagnosis was achieved on each of the 25 cases in the original set. At the next phase, our Fairchild collaborators tested the updated system against test data from another 18 randomly selected problem lots. Twelve achieved satisfactory diagnostic results, and according to the process engineers, some of these were more objective (that is, contained a more exhaustive set of possible causes) than the original diagnoses. Again, missing knowledge accounted for the misdiagnoses.

Conclusions and Future Research

The experience at Puyallup with the Fairchild ISO-Z process suggests that with continued tuning PIES can become an effective productivity enhancement tool for failure-analysis engineers. More importantly, the Puyallup experiment demonstrates the feasibility of transferring responsibility for building and maintaining the knowledge base of an expert system from AI specialists to the people who possess first-hand knowledge of a domain. We believe that this transfer is inevitable if expert systems are to become practical in continually evolving domains such as engineering and manufacturing. The experiment also confirms the expected weakness of any shallow-level approach; that is, a system which relies solely on coded experiential knowledge must be expected to fail when encountering a processing failure not previously seen.

In addition to its primary role in process diagnosis, the PIES knowledge base is also valuable as a knowledge carrier to document, propagate, and replicate engineering experience. In the semiconductor industry, a new process is usually developed in an R&D environment and then transferred to manufacturing facilities in different geographical locations. In the transfer, precious operating experience is lost, and it is often necessary to physically transfer personnel along with the process to regain acceptable yields. PIES can be used to document the diagnostic experience acquired during a process-development phase and then pass that experience to manufacturing engineers at remote sites without the need to move people.

Generalizations

The same multilevel knowledge structure discussed in this article can be used to interpret parametric test data for any semiconductor fabrication process. Currently, Fairchild engineers at several sites are building PIES knowledge bases for their latest processes. In a broader sense, PIES can be applied to many other diagnostic problems in which a se-

quence of causal levels can clearly be identified. Underlying PIES is an explicitly defined "shell" that can easily be reconfigured to reflect the appropriate causal structure. The extensibility of PIES has already been demonstrated by applying it to diagnose problems in a photolithographic process. This knowledge base, constructed by a photolithographic expert at Fairchild's Research Center, encodes causal connections between visually acquired symptoms—for example, the exposed pattern on a wafer is out of focus along only one axis—and its causes—for example, the stepper stage control gain is too high. Many other applications to in-process monitoring and controlling are under consideration. The ability to do one's own knowledge engineering is a very powerful incentive, luring engineers to try new applications.

Toward a Deeper Knowledge System

We argued previously that in engineering applications there is a continuing need to update the knowledge base to reflect changes in the domain. PIES addresses this problem by transferring responsibility for knowledge base maintenance to the domain experts. An alternative, which is based on current AI research at SPAR and other laboratories, is to provide the computer with "deeper" models that enable it to account for observed symptoms using fundamental engineering theories of the domain. In the case of semiconductor fabrication, knowledge of device physics and process technology can be used to create models that show how fabrication processes affect wafer structure and how changes in structure affect the electrical behavior of test circuits. These models can be used to derive explanations for fabrication problems not previously encountered (Mohammed and Simmons 1986). They can also be used to automatically update the knowledge base when the process recipe or a test circuit changes. Finally, the models can be used to validate the completeness and correctness of knowledge contributed by domain experts; for example, are there any alternative explanations that could account for an observed symptom. In the near future, we hope to integrate PIES with a system based on causal process models to realize these advantages.

Acknowledgments

The authors would like to express their sincere thanks to Dr. Harry G. Barrow, chief scientist of Schlumberger Palo Alto Research, whose broad vision helped to shape the PIES project. Harry directly contributed many ideas discussed in this paper and personally implemented part of the PIES system. Mike Slama and Les Smith from the high-speed memory logic division of Fairchild at Puyallup, Washington, originally suggested the topic of parametric test interpretation, served as the domain experts to construct and maintain the PIES knowledge base for this experiment, and evaluated the system's performance. Dr. Harold Brown of the Knowledge Systems Laboratory, Stanford University, served as a consultant for this project and contributed valuable ideas. Dr.

Tony Crossley of Fairchild's Palo Alto Research Center helped to make this paper readable for semiconductor engineers.

References

- Davis, R., Buchanan B., and Shortliffe, E. 1975 Production Rules as a Representation for a Knowledge-Based Consultation Program. Technical Report, STAN-CS-75-519, Department of Computer Science, Stanford Univ
- DeKleer, J. 1979 Causal and Teleological Reasoning in Circuit Recognition Ph D diss., Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology
- Forbus, K. 1984 Qualitative Process Theory Ph D diss., Dept of Electrical Engineering and Computer Science, Massachusetts Institute of Technology. Hart, P. 1982 Directions for AI in the Eighties Technical Report (612), Fairchild Camera and Instrument Corp.
- Kuipers, B., and Kassirei, J. 1983 How to Discover a Knowledge Representation for Causal Reasoning by Studying an Expert Physician In: *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, 49-56
- Mohammed, J., and Simmons, R. 1986 Qualitative Simulation of Semiconductor Fabrication In: *Proceedings of the Fifth National Conference on Artificial Intelligence*
- Pan, Y. 1983. Qualitative Reasonings with Deep Level Mechanism Models for Diagnoses of Dependent Failures. Ph D. diss. CSL Report T-132, Department of Computer Science, University of Illinois, Urbana/Champaign.
- Shortliffe, E. 1976 *Computer-Based Medical Consultation MYCIN* New York: American Elsevier
- Winston, P. 1984 *Artificial Intelligence* Reading, Mass: Addison Wesley

