P.I.E.S.: An Engineer's "Do-It-Yourself" Knowledge System for Interpretation of Parametric Test Data

Jeff Yung-Choa Pan and Jay M. Tenenbaum

Schlumberger Palo Alto Research
3340 Hillview Avenue
Palo Alto, CA 94304

ABSTRACT

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, building the knowledge base, and then maintaining it to reflect process modifications and operating experience. The approach appears applicable to other process control and diagnosis tasks.

1. Introduction

This report summarizes our experience in building PIES (*Parametric Interpretation Expert System*), a knowledge-based system that diagnoses problems in semiconductor fabrication processes by analyzing parametric test data.

Parametric measurement, 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 process status: whether it is functioning correctly, and if not, what is the nature and cause of the abnormality. Currently this interpretation task is performed by a group of semiconductor specialists known as *failure analysis* or *yield enhancement* engineers. It routinely consumes a large proportion of their time. Moreover, 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 engineering domains: (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 AI professionals for this on-going support function. 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 may be too thick because the wafer was left in an oven too long or the oven temperature was too high. Finally, process problems are traced to root causes (e.g., the wafer was left in the oven too long because a timer broke).

The multi-level 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. 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, intermediate levels for hypothesization and verification. 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 more efficient to proceed step by step, hypothesizing and prioritizing or ruling out possibilities at the structural and process levels. In addition to being more efficient, this multi-level diagnosis leads to explanations that fabrication engineers find easy to comprehend.

A working knowledge-based system incorporating the above concepts was implemented in Franz Lisp on a VAX/Unix system at Schlumberger Palo Alto Research. 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.

2. Background

2.1. About Semiconductor Fabrication and Parametric Test

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 precise control of process equipment (e.g., epitaxy, oxidation, etching, 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 complete execution of the recipe. Hence, the need for parametric testing.

When abnormal measurements of some key parameters are detected, the wafer is rejected and 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 workload 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 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 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, semicoductor 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 to develop qualitative modeling and reasoning techniques that can supplement PIES' experience-oriented knowledge base.

2.2. Shallow-Level vs. Deep-Level Approach to Expert Systems

A conventional way to build an expert system for diagnosing process faults would be to rely on a knowledge engineer to capture the experience of fabrication engineers in the form of if-then or production rules [1]. An inference mechanism might then use a *forward chaining* inference process [2] 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 [3], 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 models of the (semiconductor) domain [4].

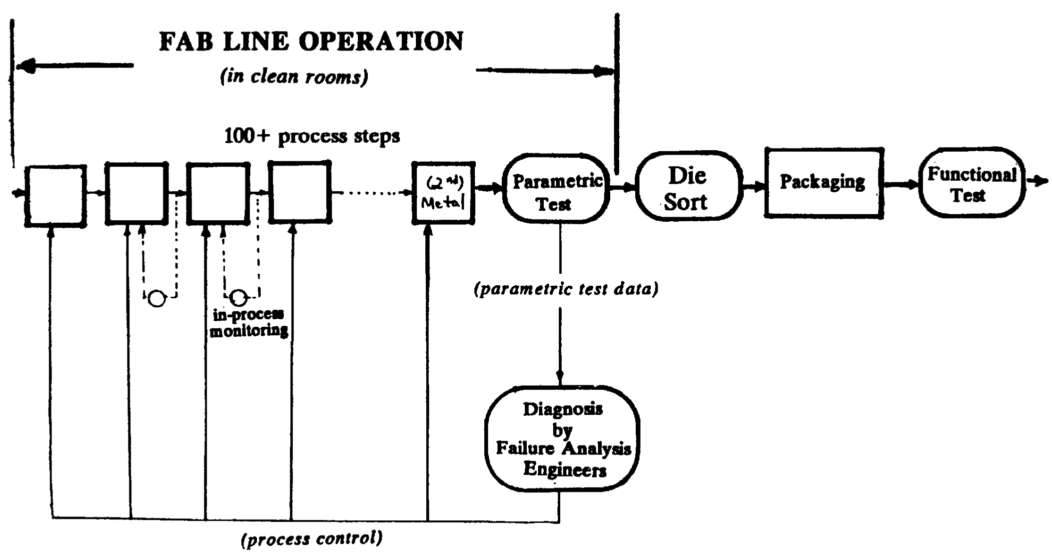

Figure 1 A Typical Semiconductor Manufacturing Process

A shallow-level approach is suitable when experience, 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 relative short time, and can be "tuned" to a high-level of performance, as demonstrated by MYCIN [5]. However, a shallow-level system will require re-engineering of its knowledge base whenever there is a change in the domain.

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 [6, 7, 8, 4] is promising, but the technique needs futher development before it can be useful in practice.

PIES' knowledge base approach falls between shallow and deep level approaches (*semi-deep*). 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. On the other hand, 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).

3. Approach

3.1. 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 a 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 that structure. PIES' diagnosis 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 --> rootcauses.*

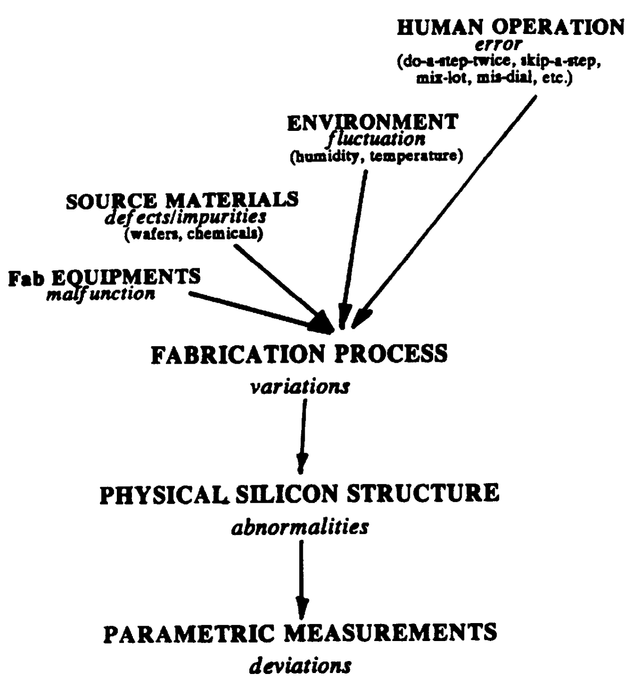

Figure 2 Multi-Level Propagation of Fabrication Failures

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 example, at the physical structure level, such modes would include incorrect thickness or doping density of particular wafer layers (e.g., the epitaxial layer). At the fabrication process level, the failure modes would include incorrect temperatures or gas densities during particular process steps (e.g., 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 (i.e., measurement deviations) and a set of possible causes (i.e., process failures).

Diagnosis proceeds as a multi-level hypothesis-verification process. Parametric measurements are first preprocessed to transform them from numeric values to qualitative ranges (e.g., 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 against 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 highest likelihood process failures. This iterated hypothesis-verification approach will identify the primary (i.e., most likely) failures. In many cases, it will also reveal multiple failures, that may 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 PIES' multi-level case-centered knowledge organization, thereby guiding an engineer to decompose his knowledge in a way that is both natural for him and required by PIES. Using the editor, an engineer can focus on the failure cases at any level. He 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, he can add it to the knowledge base, along with the expected symptoms and probable causes.

3.2. Knowledge Base

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

At each causal level, the knowledge base is decomposed into frame-like 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 PIES' current implementation 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 types of *associational link* which describe how this case is causally related to other types of failure. Other slots are used in conjunction with the knowledge base editor (see below) to group failure cases in ways that users find convenient.

A domain expert's knowledge about possible causal connections between two types of failure is represented in PIES by associational links. A link may be one of two types: *causes* or *caused-by*, and is further distinguished between *intra-level* and *inter-level*, 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, 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. PIES representation for this problem, known as *BASE-DISTRIBUTION-deep*, is shown (in its pretty-print form) as the following:

Knowledge about a case of *physical structure* defect:
BASE DISTRIBUTION deep
****

Possible effects at measurement level --
1: ((parametric-measurement WE10BETA low) very-likely)
2: ((parametric-measurement RB1 low) probably)
3: ((parametric-measurement RB2 low) very-likely)
4: ((parametric-measurement WE10-CBO low) probably)
5: ((parametric-measurement SOT2-CBO low) probably)
6: ((parametric-measurement SOT-B-SU very-low) probably)
7: ((parametric-measurement SOTBETAF low) probably)

Possible causes at process level --
1: ((BASE-IMPLANT ENERGY high) likely)
2: ((BASE-DRIVE FURNACE-TEMPERATURE high) likely)
3: ((BASE-DRIVE DIFFUSION-TIME long) likely)

Possible causes at SAME physical-structure level --
1: ((BASE-OXIDE THICKNESS low) likely)

****

In this example, the failure type of BASE DISTRIBUTION deep is said to be causally related to other types of failure at the process level, measurement level, and the physical structure level itself. As indicated, if it occurs, it may result in seven types of measurement deviation, but some of them are more likely to manifest (e.g., WE10BETA) than others (e.g., RB1).

3.3. 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: it relieves AI specialists from on-site visits and lengthy *knowledge engineering* sessions with domain experts; it avoids misunderstanding, and thus mistranslation of knowledge from domain experts to AI specialists; and it allows domain experts to incorporate new experience quickly into the knowledge base. This last feature makes the system more suitable than the traditional expert system approach in dealing with a changing domain.

The primary function of the knowledge editor is to guide domain experts in codifying their knowledge and expertise in a form consistent with the PIES knowledge base. During a knowledge engineering session, the knowledge editor first allows the domain expert to focus his attention on one of PIES causal levels. Within that particular level, the knowledge editor allows the user to maintain his own hierarchy of failure concepts. For example, at the physical structure level, he may wish to group together all failures associated with the same wafer layer, and within any one layer, all failures of a particular type (e.g., doping problems). This support of concept hierarchies helps the expert to organize the many types of failure known to the knowledge base. The knowledge editor provides its users with easy commands to create and traverse his hierarchy, to define new failure cases, and subsequently to fill in or modify the contents (slots) of a failure case. In summary, the PIES' knowledge editor guides a domain expert to decompose his failure-related expertise into the structure required by PIES' knowledge base. It ensures that the knowledge that is codified is both syntactically and semantically correct.

For example, in a knowledge-engineering session to build the knowledge base for diagnosing failures in Fairchild's ISO-Z bipolar process, our collaborator at Fairchild/Puyallup site chose to focus his attention on the physical structure level. PIES' editor helped him to organize known cases of physical structure failures into a hierarchy, and allowed him to traverse the hierarchy to a particular case of interest: BASE-DISTRIBUTION-deep, as shown in figure 3. To organize what he knew about the failure, the expert conceptualized relevant causalities centered around BASE-DISTRIBUTION-deep, as shown in figure 4. The knowledge editor allowed him to establish associational links from BASE-DISTRIBUTION-deep to other known failure cases at *effect level* (measurement level), *cause level* (process level), and *self* (physical structure level). The editor allows him to *add, delete,* or *replace* associational links, as necessary.

```
$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$
          MULTI-LEVEL CASE EDITOR   <SPAR PIES-1.0>
             last updated: February 1985
$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$

  Welcome to Fairchild/Schlumberger Parametric Interpretation Expert System

Available Commands are:
 <HELP> show new-perspective up top Mark-Case fill edit write
 write-and-quit exit abort expert beginner reset reset!
 display-highlight-on display-highlight-off

Enter Command (<CR> for listing of Case-Path) =>> newp
 1: measurement  2: physical-structure  3: process  4: root-causes


Enter selection (0 for redisplay, <CR> for physical-structure) => 2

Case-Library last modified on Wed Jun 12 11:07:18 1985
Total of 82 cases from file <physical.cas> loaded!!

You are now referring to the TOP of physical-structure
 1: COLLECTOR  2: EPI  3: ETCHED-SILICON  4: ISO-OX  5: SINK
 6: EMITTER  7: BASE  8: FIELD  9: SILICON  10: METAL-1
 11: ISO-ISLAND  12: VIA  13: BASE-OXIDE  14: LVCEO-RESISTOR
 15: GROUND-TAP  16: GUARD-RING  17: SIDEWALL  18: METAL-2

Enter Command (<CR> for listing of Case-Path) =>> 7

You are now referring to (BASE) of physical-structure
 1: EXTRINSIC-Q  2: DISTRIBUTION  3: INTRINSIC-Q

Enter Command (<CR> for listing of Case-Path) =>> 2

You are now referring to (BASE DISTRIBUTION) of physical-structure
 1: deep  2: shallow

Enter Command (<CR> for listing of Case-Path) =>> 1
You are now referring to (BASE DISTRIBUTION deep) of physical-structure

    ====================================

CASE-NAME: BASE-DISTRIBUTION-deep at physical-structure level

           ------
  Following symptoms at measurement level are to be resulted from this case:
    1: ((parametric-measurement WE10BETA low) very-likely)
    2: ((parametric-measurement RB1 low) probably)
    3: ((parametric-measurement RB2 low) very-likely)
    4: ((parametric-measurement WE10-CBO low) probably)
    5: ((parametric-measurement SOT2-CBO low) probably)
    6: ((parametric-measurement SOT-B-SU very-low) probably)
    7: ((parametric-measurement SOTBETAF low) probably)


           ------
  Following causes at process level is to result in this case:
    1: ((BASE-IMPLANT ENERGY high) likely)
    2: ((BASE-DRIVE FURNACE-TEMPERATURE high) likely)
    3: ((BASE-DRIVE DIFFUSION-TIME long) likely)


           ------
  Following causes at CURRENT physical-structure level can result in this case:
    1: ((BASE-OXIDE THICKNESS low) likely)

    ====================================
```

Figure 3 A Sample Knowledge Engineering Session under PIES' Knowledge Editor

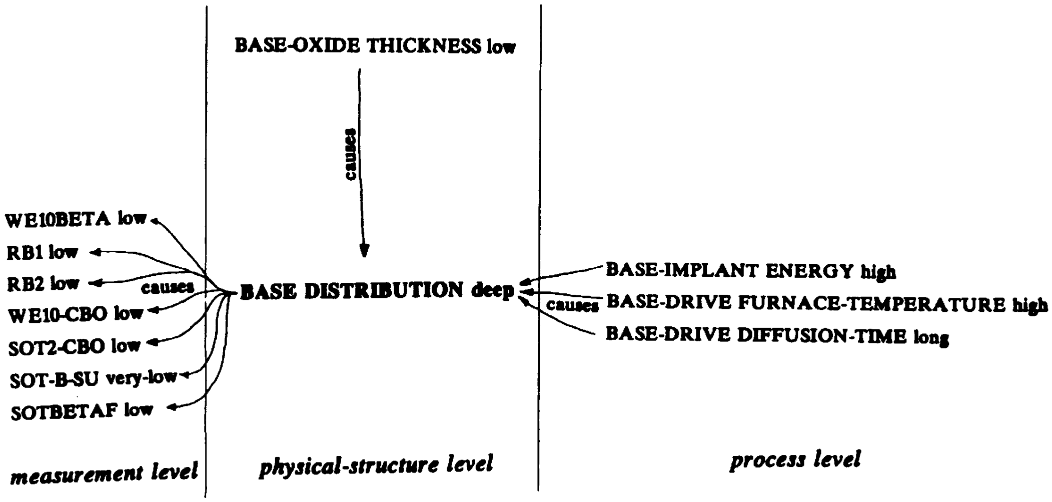

Figure 4 Organization of Concepts causally related to BASE-DISTRIBUTION-deep

From our experience, failure analysis engineers with no AI background were capable of mastering PIES's knowledge editor after a brief (less than an hour) tutorial session.

3.4. Diagnostic Reasoner

PIES diagnostic reasoning mechanism exploits the multiple causal level structure of the knowledge base to diagnose rootcause 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 was recorded by Puyallup's Keithley tester). The symptom abstraction process follows two steps: first, noisy data points (due to bad probe contact or random failure) are removed from the data set by a statistical method; then a statistical average and standard deviation is 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 (e.g., 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 hypotheses suggested by 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 rootcause level.

Let us follow through an example of this reasoning chain. EPI-R is a measurement of electrical 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 more expensive, time-consuming material analysis techniques. Tracing further back along the causal chain, 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 - which resulted in 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 three 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 - while the "sensitivity" (i.e., 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 may not have been included because the symptoms stipulated for hypothesizing that 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 *intra-level causalities* coded in the knowledge base. For example, one intra-level 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 4). 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 may cause (e.g., 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 by the diagnosis process so far. The matching process will compute a "score" for each failure case, indicating how close the case's expected symptoms match against 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 the reasoning.

4. 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, among which, 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, and is today competently maintained by Puyallup's failure analysis engineers.

The performance of PIES was evaluated by analyzing parametric test data from problem lots which 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, 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 those unsuccessful diagnoses was, not surprisingly, missing knowledge in PIES' knowledge base. The problems were subsequently corrected by Puyallup 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 Puyallup collaborators tested the updated system against test data from another 18 randomly-selected problem lots. Among those, 12 achieved satisfactory diagnostic results, and according to the fab engineers, some even "outperformed" the original diagnoses. Again, missing knowledge accounted for the misdiagnoses.

5. Conclusions and Future Research

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, namely, a system that 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 moving people.

5.1. Generalizations

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

5.2. Toward a Deeper Knowledge System

We have 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, 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 electrical behavior of test circuits. These models can be used to derive explanations for fabrication problems not previously encountered [9]. They can also be used to update automatically the knowledge base when the process recipe or test circuits change. Finally, they can be used to validate knowledge contributed by domain experts for completeness and correctness (e.g., 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.

ACKNOWLEDGEMENT

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 and 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 System Laboratory, Stanford University, served as a consultant for this project and contributed valuable ideas. Dr. Tony Crossley of Fairchild's Palo Alto Reserach Center helped to make this paper readable by semiconductor engineers.

REFERENCES

[1] Davis, R., et al. (1975). Production Rules as a Representation for a Knowledge-Based Consultation Program. Stanford University. STAN-CS-75-519.

[2] Winston, P. (1984). Artificial Intelligence. Addison Wesley, Pub.

[3] Hart, P. (1982). Directions for AI in the Eighties. Fairchild Technical Report No. 612.

[4] Pan, Y. (1983). Qualitative Reasonings with Deep-Level Mechanism Models for Diagnoses of Dependent Failures. Ph.D. Dissertation. University of Illinois, Urbana/Champaign. CSL Report T-132.

[5] Shortliffe, E. (1976). Computer-Based Medical Consultation: MYCIN. American Elsevier, Pub.

[6] Forbus, K. (1984). Qualitative Process Theory. Ph.D. Dissertation. M.I.T. AI-TR-789.

[7] DeKleer, J. (1979). Causal and Teleological Reasonings in Circuit Recognition. Ph.D. Dissertation. M.I.T. AI-TR-529.

[8] Kuipers, B. and Kassirer, J. (1983). How to Discover a Knowledge Representation for Causal Reasoning by Studying an Expert Physician. *Proceedings, 8th International Joint Conference on Artificial Intelligence.*

[9] Mohammed, J. and Simmons, R. (1986). Qualitative Simulation of Semiconductor Fabrication. *Proceedings, AAAI Conference, 1986.*