Invited Plenary Lecture, 1995 NIST Workshop on "Semiconductor Characterization: Present Status and Future Needs," pp. 459-66, W.M. Bullis, D.G. Seiler and A.C. Diebold Eds., AIP Press, Woodbury NY, 1995

In-Situ Process Control for Semiconductor Technology: A Contrast Between Research and Factory Perspectives

Or

Observations of a Researcher Who Has Seen the Light... At Dawn... Driving the Interstate to an Out-of-State Factory

John C. Bean
AT&T Bell Laboratories, Murray Hill NJ 07974

In the redirected 90's, many researchers have been cast into the brave new world of the factory. This happened to me, and it happened to the AT&T Bell Labs department I manage. In this paper I draw on that experience and suggest ways in which the researcher can increase his effectiveness in this new environment. I adopt a critical and irreverant factory viewpoint in the hope of better preparing researchers for the adventure they face. I focus on the likely disconnects in areas ranging from attitude to methodology, drawing from our experiences in an optoelectronics factory.

PROBLEM SELECTION

The bulk of this paper focuses on methodology. However, it is critical that we first address the subjects of problem selection and definition - critical because it is at this point that many researchers have already lost the battle to become effective factory contributors. Let's examine the classic research value system. Researchers are taught to pursue problems that offer the possibility of comprehensive and enduring understanding. The global nature of these attributes mean that research and researchers are naturally evaluated by communities extending well beyond a single institution or employer. This process presents a dilemma: how does the individual researcher achieve acknowledgment in a vast, essentially worldwide, community? The proven solution is to specialize and thereby define a community of more manageable dimensions. This process is further expedited by association with a specific problem or methodology, ideally of one's own invention.

However simplified, accept the above description and consider the behavior it stimulates. At the top of the value list is the desirability of knowledge. This knowledge should be complete and all-encompassing, it should be unambiguous, it should be fundamental, and it should be timeless. Researchers are the generators of this knowledge and they are evaluated almost exclusively on the basis of their expertise in a specific chosen field. This evaluation is, in turn, critically dependent on peer recognition. Recognition can be enhanced by tackling what are acknowledged to be the most difficult problems. Recognition can be enhanced by novelty, by invention, invention of analytical approaches, techniques and instrumentation. Further, a portfolio of small inventions is generally less recognizable than a single large invention with which the researcher will be irrevocably associated. This leads naturally to the researcher's role as champion and advocate of said tool or technique.

Contrast the above points with the needs of the factory engineer. A factory process is, of course, based on knowledge. However, that knowledge is largely acquired at an earlier development stage. That foundation should be largely complete by the time a process moves into a factory, at which point the emphasis turns to maintenance of a production status quo, that is, reproduction of a given product again and again with little or no deviation. Problems do, of course, arise but the emphasis will then be on expeditious solution, not on further extension of a knowledge base. This leads to one of the first disconnects I repeatedly observe in discussions between researchers and factory engineers. The engineer looks for a parameter that will define a successful process, or a tool to measure that parameter. The researcher dogmatically insists that a shift in that parameter can be produced by a range of phenomena and that a better parameter must be found. The discussion continues at cross purposes until frustration drives one of the parties to leave. The problem is that while the researcher may believe he is focusing on the engineer's need for a solution he is implicitly, and perhaps unconsciously, working another agenda: the need for unambiguous understanding. At the moment, however, the engineer does not need to resolve that ambiguity, he simply needs a handle, an identifier, a fingerprint of a successful process. In time he, or his colleagues, will attempt to resolve much of that ambiguity either as a part of improving the process or setting up the next product process. That, however, will occur at a future date in what is often a separate off-line investigation.

The intolerance of ambiguity is closely related to another chronic disconnect between researcher and factory engineer: the need for completeness. Research training dwells on those problems, such as the hydrogen atom, for which complete, closed, analytical solutions are possible. Many researchers seek to emulate such examples by choosing problems for which similar depth of understanding appears likely. This leads to the cliché of the scientist treating only very small systems, atoms or molecules, or extremely large systems, e.g. universes. Factory problems, of course, fall at neither extreme and complete models are not possible. To deal with this, modern factories are making increased use of statistics to define a history of common patterns. Statistics can thus be used to fill in the gaps in understanding. Statistical histories can be used to identify the most significant process parameters. Often, statistical histories can be used to define those areas that can be safely ignored, either because parameters are irrelevant or because intrinsic variation is small. Unfortunately, as researchers many of us have been taught that statistics are a poor man's substitute for complete knowledge and we will stubbornly resist not only using statistics but even learning the language (which for a Ph.D. is a relatively straight-forward process). However, in the absence of such supplemental information, the researcher is often left with either an intractable issue or compelled to define a problem domain that is so small as to be irrelevant in either return or time scale.

This leads to a discussion of specialization, invention and advocacy. Most researchers would characterize themselves as problem solvers. However, to the factory, researchers may appear more as solution providers. The distinction is important. In the factory, economic survival dictates that precedence be given to solution of a specific problem. The most appropriate solution is the solution that is fastest and cheapest. It does not matter if the solution is borrowed from someone or somewhere else. In fact this may the best way to ensure that it will be cheap and fast. An obsession with novelty, cleverness, invention and other research icons may in fact be counterproductive or even dangerous in this context. *Further, it is the problem that must take precedence, not the solution.* Again, these statements may seem obvious but I have repeatedly observed researchers ignore these principles. Our roles as expert, as inventor and advocate too often compel us to obsessively focus on our pet solution. Researchers repeatedly tune out aspects of a problem that do not fit their preconceived range of expertise and will in fact often try to turn discussion to new, more appropriate, problems! By surrendering our role as experts we may in fact undercut opportunities for recognition in an external community. It is ironic, however, that this surrender may allow us to use a much wider range of our expertise. Modern laboratory practices mean that the most successful researchers are in fact experts, of a sort, in a wide range of fields: chemistry, electrical engineering, computer programming, mechanical engineering, et cetera. World class experts, no, but experts and jacks-of-all-trades nevertheless. A focus on strictly pragmatic problem solving can allow the researcher unconstrained use of these skills and can, in fact, be surprisingly liberating and fun.

METHODOLOGY

It is time to get down to a more specific list of suggestions on how a researcher can increase factory impact, to define what I would call a path of least resistance. This will not be the only possible path but it is one relevant to our optoelectronic factory experience and one that may nevertheless offer insights to researchers in other production environments.

At the outset, I will lump possible activities into two categories: off-line and on-line. Off-line would include most of the essential work generally described as process or product characterization. These activities use many of the tools and methodologies with which the researcher is already familiar and comfortable. Because of that familiarity and comfort, and given the general thrust of this volume, I would like to focus on the other more foreign area of on-line investigation or what is generally known as process control. Process control has become increasingly important as virtually all semiconductor technologies move to the use of ever larger full wafers. Too many otherwise powerful off-line tools interfere with such processing or are fundamentally incompatible because of their sample destructive nature. Their use is thus delayed until processing is completed or restricted to secondary process monitoring samples. In either case the production process engineer is left with gaping blindspots or blind periods in which valuable product may be lost. Hence the compulsion to identify on-line supplements.

What should the researcher propose in such an on-line tool? First, preference should always be given to the tool that is **NON-INVASIVE**. Factory engineers are well-schooled experts on the myriad ways in which a process can go wrong. They will fiercely resist any proposal, however promising, if it may upset an established, successful process. The simplest way to protect against such a perturbation is by holding the actual processing chamber inviolate. This means that there is a built-in bias against tools that employ ion beams, electron beams, and to a lesser extent X-ray beams. Ion and electron beams have two shortcomings: they require vacuum and they entail introduction of foreign bodies in the process chamber, specifically guns and detectors. An X-ray apparatus can sidestep this problem by penetrating the wall of a process system. Unfortunately it as readily penetrates an operator and the bulky radiation shielding may require redesign of the processing system, the ultimate invasion. These considerations leave the door open for another alternative: optics. Because man is an optical animal, optical viewport materials are well characterized and they are most likely already in use on the processing system. Addition of optical viewports thus entails minimal risk and existing ports may, in fact, provide necessary access for what will then be a totally non-invasive optical tool.

Beyond being non-invasive, an ideal process control tool should be **NON-PERTURBING**. What do I mean by this? Consider the following. In many process systems, samples are only very loosely positioned. This may help to accommodate thermal cycling or it may be implicit in the mechanics of transferring samples through load-locked systems such as cluster tools, MOCVD or MBE systems. In addition, samples may well rest on opaque stages, e.g. heating, cooling and or electrical biasing assemblies. Finally, in many epitaxy and coating systems, samples rotate to promote deposition uniformity. Combine this with loose positioning and the result is wobble. This wobble can be quite large and I have measured values as large as $\pm 3°$. In principle, all of these shortcomings can be engineered out of a process system and the promise of enhanced process control tools may drive this process. However, for the moment, these shortcomings exist and a process control tool that requires their elimination will be hugely perturbing. For instance, an otherwise non-invasive optical tool such as an ellipsometer may have to be bypassed because it implicitly requires sample alignment to a fraction of a degree, a condition that simply does not exist in many current pieces of equipment. An ideal optical tool should instead accommodate existing equipment shortcomings. This can be accomplished though scanning, phase locking or other synchronization methods. An even more robust tool can be based on what a colleague christened "the lucky photon principle," that is, the design of a fixed optical illumination and collection system with acceptance angles so wide that regardless of the sample dance, at least the occasional photon makes it through the optical path.

As a third attribute, the ideal process control should be truly **ON-LINE rather than IN-SITU**. The distinction here parallels the previous paragraph and may be best explained by a specific illustration. In the last ten years the technique of RHEED oscillations[1] has received intense attention in the research community from which it was proposed as process control tool capable of unambiguously measuring growth of single atomic layers. In this context, the specifics of the acronym and technique are irrelevant. Suffice it to say that it uses an electron beam to detect the roughening of a crystalline surface as a new atomic plane is nucleated and is indeed in-situ and capable of detecting a single such plane. It is, however, off-line in one critical sense: it requires growth on a precisely oriented crystal surface. Most research uses such crystals. Most production does not. Discovering this fact, many research advocates suggested that one could nevertheless grow on one such crystal each day and thereby calibrate the system for later off-orientation growth. Again, this is possible but in effect it takes a very expensive growth apparatus off-line and forces it to serve the role of a comparatively inexpensive thickness measurement tool. This perverts the whole idea of effective capital utilization. Other popular optical techniques, such as Reflection Difference Spectroscopy[2],

---

[1] J.H. Neave, B.A. Joyce, P.J. Dobson and N. Norton, *Appl. Phys.* **A31**, 1 (1983)

[2] D.E. Aspnes, J.P. Harbison, A.A. Studna and L.T. Florez, *Phys. Rev. Lett.* **59**, 1687 (1987)

may suffer from the same need to operate a production tool in a non-production mode.

As a final generalized attribute, the ideal process control tool should **FOCUS NARROWLY ON CRITICAL PROCESS PARAMETERS**. This focus is dictated by the fact that a broader analysis almost always entails additional expense and complexity, in both instrumentation and interpretation. Again, an illustration is useful. In a transparent material, the optical path is the product of the layer thickness and the layer refractive index. While a large variety of relatively simple optical tools can measure optical path, very few can unambiguously determine the two underlying parameters. Spectral ellipsometry is one such technique. Fueled by the researcher's compulsion for completeness and resolution of ambiguity, ellipsometry has become hugely popular accounting for at least half of the 150 odd papers I browsed in preparing this article. However, if one is only trying to control a process, the ambiguous measure of simple optical path may be more than adequate! While it is possible that there may be compensating errors in thickness and refractive index, it is extremely unlikely, especially when one considers that for semiconductors, refractive index is limited to a rather small range of values and thickness is not. Production is probabilistic process. If the ambiguity latent in a simple measure of optical path means that one will miss a small subset of possible process deviations, so be it, especially, if the tool for measuring path (or another comparably general parameter) is an order of magnitude less expensive or accommodating of process constraints than the full-blown ambiguity-resolving instrument. As I will describe in the final case study, many such opportunities exist. To paraphrase another factory colleague "shoot for the fat rabbits."

IMPLEMENTATION

I would now like to offer a laundry list of specific suggestions on both hardware and software. I will build on the example of an on-line optically based process control tool. This entails a loss of generality, but I again believe at least the flavor of the comments will have broader relevance.

To start, I strongly suggest the use of **OFF-THE-SHELF EQUIPMENT**. From the factory perspective, this has many advantages. Among the less obvious, but most important, is the fact that someone else is responsible for maintaining the expertise implicit in the design and manufacture of such an item. To put it another way, a $50,000 instrument becomes radically less attractive if its guaranteed operation requires one to pay the $200,000 plus annual loaded salary of a staff Ph.D. The factory may be willing to pay your salary to build the instrument but it will be strongly averse to an open ended service contract.

One should similarly aim for **MODULAR COMPONENTS,** which is another way of saying that off-the-shelf items should mate directly with off-the-shelf items. Same argument as above. In the context of my hypothetical optical process control tool these recommendations lead me to highlight the tremendous opportunities now offered by optical fibers complete with shielding and coax-like end connectors. These connectorized fibers mate with a large and ever-growing array of standard optical components. In a single issue of an optical industry trade-rag, I found ads for light sources (broadband, LED, laser, both CW and modulated), detectors (tunable, miniature), modulators, stabilizers, power meters, rotators, switches, splitters, combiners, isolators, circulators and if that were not enough, there were offers to design and fabricate custom but fully packaged fiber compatible items. One can also purchase fiber compatible subsystems. A dramatic example is provided by the modular spectrometers now being offered by several vendors. These units have no moving parts (not even an on-off switch), they fit in the palm of the hand, and, by borrowing CCD detection elements from copying machine technology, they are available with a price tags under $2000.

To be fair, the classical optical table still has a role in the research lab where it offers a wider range of options than will perhaps ever be available in connectorized modules. That does not, however, reopen the door to its use in the factory when well-engineered alternatives, such as that illustrated in figure 1, exist. Unlike the research-derived optical bench with its singular emphasis on versatility, the system of figure 1 places equal emphasis on ultimately locking in a design in a robust manner that will *not* be easily perturbed and will *not* require repeated expert alignment.

Discussion of alignment brings me to the next suggestion: use **NON-VARIABLE COMPONENTS**. This may seem like an almost trivial suggestion, but it is one at the heart of the research vs. production conflict. A researcher designs for the unexpected and thus incorporates in as much flexibility as possible. On the other hand, the factory engineer knows that an adjustment provides one way to optimize a process but a near infinite number of ways to louse it up. If your job is to maintain a production status quo, you want to design out variability. Again, let me provide an optics illustration. In the lab, one modulates light intensity with a neutral density wheel. However, when the proper intensity is determined and a move to the factory is contemplated, it is time to throw out the wheel. One option is to select a connectorized fiber with a core diameter chosen to pass the optimum light intensity. Yes, this will entail purchase of a selection of fibers and some swapping to find the one that is exactly right. However, when that process is complete, your job is done and the need for your expertise is removed. Anyone can note the part number of that fiber, order it and replace it.

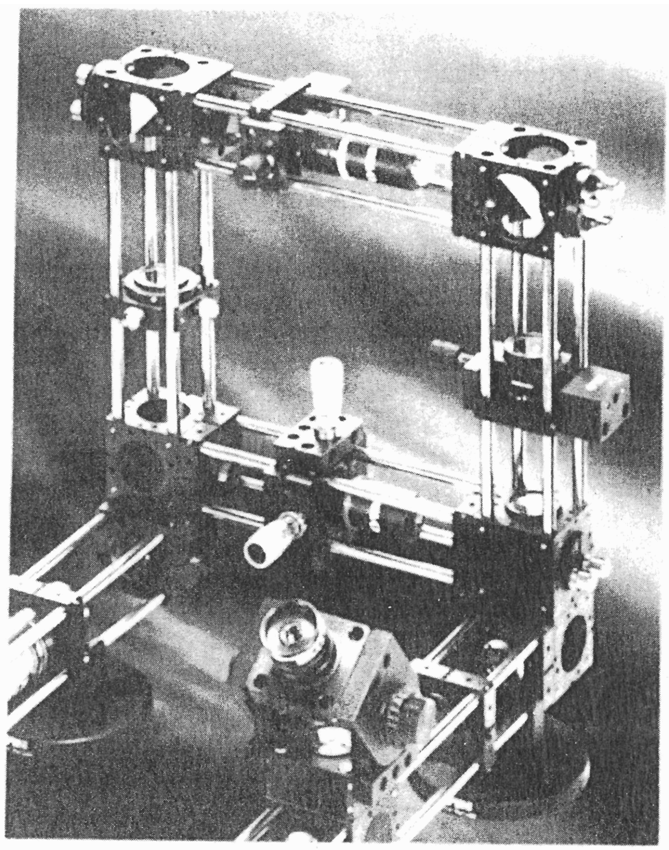

Fig. 1 Rail based modular optical bench system of Spindler-Hoyer Inc, Milford Massachusetts. Among other applications these units have been employed in optical systems built to withstand the vibration and impact loads of launch aboard the U.S. Space Shuttle (by permission of Spindler-Hoyer).

Beyond these hardware suggestions, if you are designing an instrument for the factory, also tailor the data presentation to the factory. As the instrument should focus narrowly on the most relevant process parameters, so the data should **HIGHLIGHT PROCESS DEVIATIONS**. Use ratios, offsets or differences *from the production norm.* Absolute numerical values are in a sense irrelevant; it is the deviation that is critical in an established production process. The absolute values were determined by other instruments at other stages. A simple idea, again, but I have seen the implementation of this idea crystallize the solution of a production problem that had been festering in a factory for years.

Beyond highlighting deviations, **PRESENT DATA FOR OPERATORS NOT SCIENTISTS**. I am not downplaying the competence of operators, I am simply acknowledging that our favored forms of presentation build on our lust to maximize information, even at the expense of clarity or relevance. The presentation of on-line control information should remove such extraneous information.

As data presentation will now almost certainly occur on a personal computer screen, **USE BROADLY ACCEPTED SOFTWARE PLATFORMS**. There is no longer an excuse for compelling a operator to face a DOS prompt or the manual for a unique single purpose software package. Virtually all factory denizens are familiar with graphical interfaces as implemented in the increasingly indistinguishable Mac or Windows format. Use of such a format, guarantees that most operators will already be familiar and comfortable with the mundane file and hardcopy functions that make up much of every program. Further, the prevalence of these standards means that most instrument vendors are being compelled to offer libraries (in Window's parlance: DLL's) that take care of the detailed hardware interface code. While this is still a somewhat painful work-in-progress, the existence of such libraries will soon make the programmer's job far simpler, and make it possible for the instument originator, who knows the equipment best, to drive that process.

This leads to a final software suggestion, **USE A GRAPHICAL SOFTWARE INTERFACE**. The above packages are intrinsically image based and build on the fact that the human brain is a superb image processor. Mathematical curves may incorporate precision but they have nowhere near the impact or immediacy of graphics. This is especially true for those less comfortable with mathematics, a group that includes virtually all non-scientists except mathematicians. Rather than illustrating this single point, I now offer a case study that illustrates a variety of the above suggestions.

CASE STUDY: CONTROL OF EPITAXIAL GROWTH BY INTERFEROMETRIC REFLECTOMETRY

This study is a prime example of borrowing and building upon another's ideas. In this case I am the borrower, and the borrowees are Kevin Killeen and Bill Breiland of the Sandia National Laboratories. The problem is that of determining when one has correctly grown the desired complex layer sequence of semiconductor materials. The proposed solution is based on interferometric reflectometry.[3] Use is made of penetrating wavelengths such that the reflected signal will be a complex sum of interfering reflections from interfaces throughout the structure. This is one instance in which complexity is a plus because it contributes to the uniqueness of the signal which, when presented properly, makes it easier to detect minute process deviations.

A conventional reflectometry spectrum is presented in figure 2. It is a good example of how not to present data in the factory. The spectrum has plenty of wiggles and without a computer only an idiot-savant would stand a chance of relating such wiggles back to product characteristics. In fact, even the savant would fail because these wiggles are a measure of optical paths, which, as discussed above, do not map unambiguously back to physical characteristics. Now, imagine that one proposes using reflectometry as the real-time monitor of a growth or etching process. Figure 3 presents a subset of the data produced: a superposition of spectra, one for each layer of a complex device (in this case a vertical cavity surface emitting laser or a resonant cavity photodiode). Things have gone from bad to worse. At this point, I believe Killeen and Brieland made a breakthrough by going graphical. Rather than superimposing such curves in a marginally more intelligible manner, they made each curve into a single scan line of a CRT image. Peaks were converted to white pixels, valleys to black, intermediate points to appropriate gray values. Consecutive curves were thus translated into consecutive scan lines to produce an image such as that in figure 4. What does this convoluted "fingerprint" mean? *It means that if you reproduce it you have almost certainly grown the same structure I grew.* No more, no less. It does not give you fundamental information on the structure. It does not eliminate all ambiguity and in fact the very careful design of an alternate structure could reproduce this image. However, by random process variations, such a false reproduction would be exceedingly unlikely and the fingerprint thus serves as a very effective process monitor.

I now invoke the rule on highlighting process deviations. In figure 4, I superimpose the fingerprint of a growing structure on the background fingerprint of an accepted standard. It is relatively easy to see that the fingerprint of the growing structure (below the horizontal white demarcation line) is offset slightly to the right of the background reference. Because I am measuring optical path, the degree of offset is a measure of the drift in physical dimension: right corresponds to longer or larger. In fact, this information is not necessary. It is only important that the deviation of the fingerprint scale in some smooth way with the error in the process. A statistical history and correlation with product yields would then establish the limits of acceptable alignment.

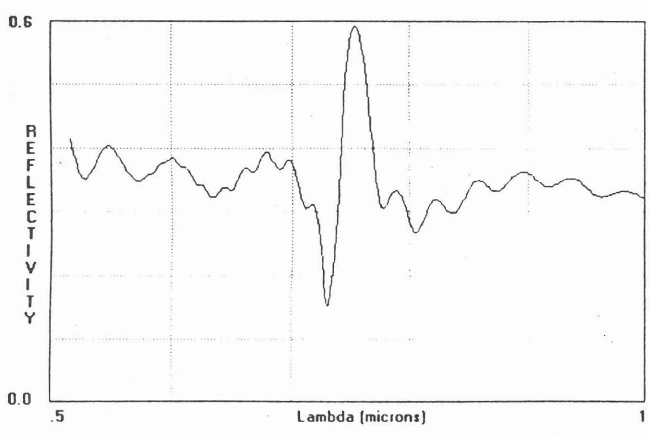

Fig. 2 Conventional spectrum of visible light reflected back from complex semiconductor layer structure.

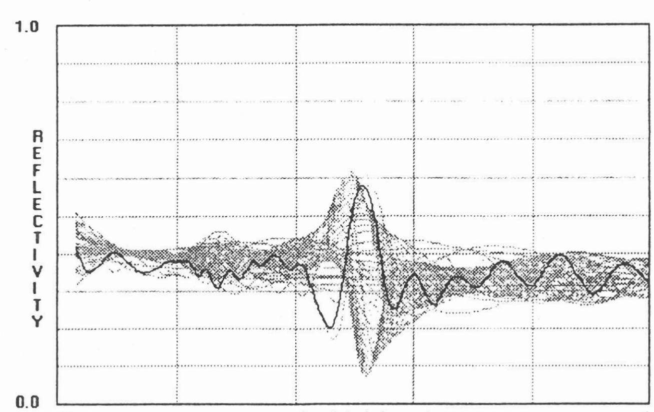

Fig. 3 Superposition of spectra such as those in fig. 3 illustrating nearly incomprehensible build-up of information obtained during growth of structure.

[3]K.P. Killeen and W.G. Breiland, *J. Elec. Materials* **23**, 179 (1994)

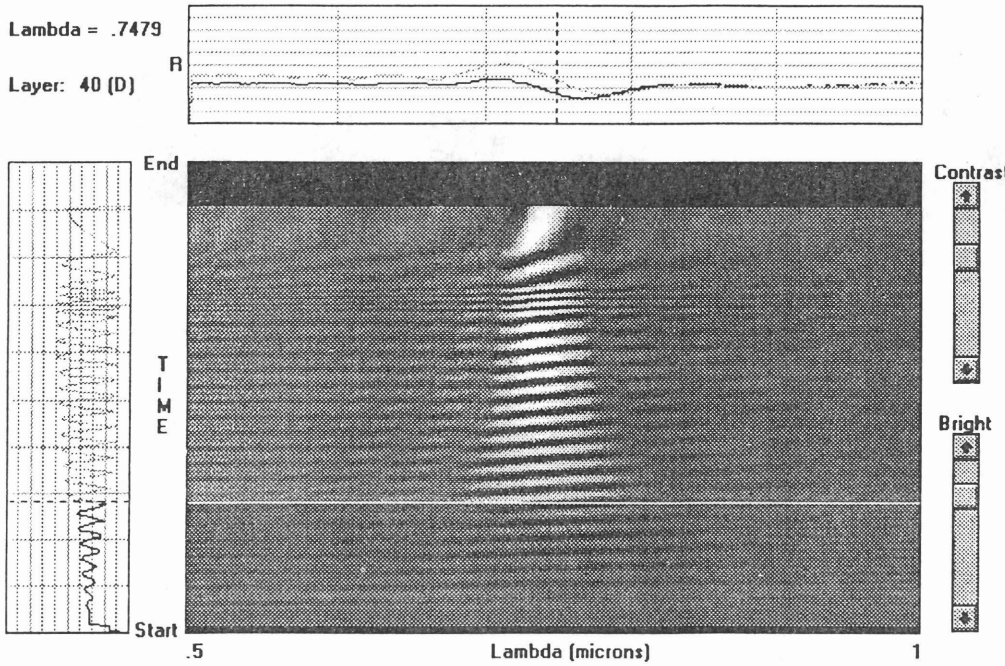

Fig. 4 Translation of conventional spectra into time-dependent gray-scale "fingerprint" after the work of Kevin Killeen and Bill Breiland of the Sandia National Laboratories. Lower section of fingerprint taken from actual growing sample. This is superimposed on reference background fingerprint of known "good" sample. The slight left/right mis-alignment is indicative of process deviation. This deviation highlighted in fingerprint cross-sections in top and side boxes.

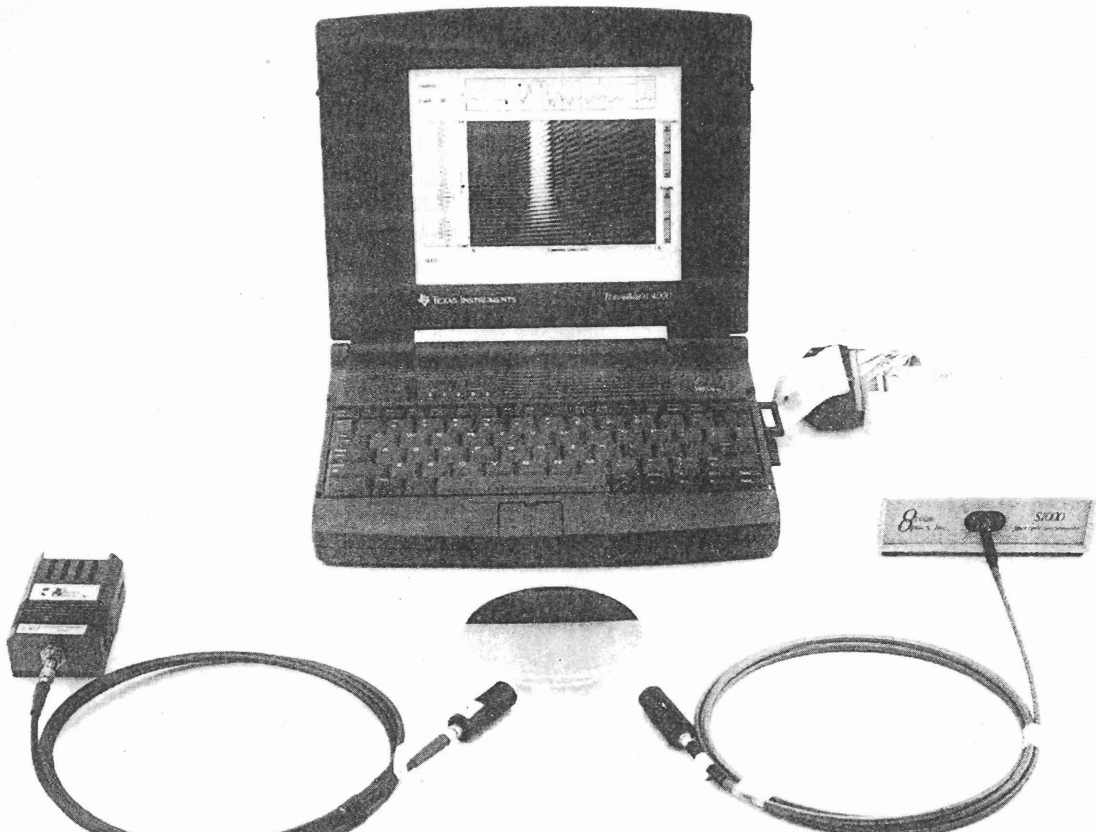

Fig. 5 Completely modular, off-the-shelf, sample tolerant, implementation of interferometric reflectometer. Appended to a Molecular Beam Epitaxy growth system, this simplified unit produced the data of the above figures.

In this superimposed image, I have added the ability to take cross sections along particular planes of the fingerprints. The top box shows cross sections immediately above and below the demarcation; the side box shows a cross section at the wavelength printed at upper left or indicated by the dotted line in the top box. These are called up by simply clicking on the fingerprint at a point of interest. The particulars are unimportant. The significance is that, in a graphical programming environment (here Microsoft Visual Basic and Windows), a very wide range of graphical options is available and easily implemented. The right format is the one that proves most successful at highlighting process deviations.

Now let's turn to the hardware. In presenting this paper I used an ad entitled "Reflectance Measurements Made Simple." For copyright reasons I'll simply describe it here. It pictures a lamp housing with a beam directed out at a chopper assembly through a beam-splitter to the sample and back to the beam splitter to a focal length matcher appended to a mechanical spectrometer appended to a detector cabled to a digital lock-in amplifier to a PC. Their disconnected nature suggests that all of these components were to be secured to an optical bench. This may be a simple implementation by research standards, but it is entirely inconsistent with the factory principles I have attempted to illustrate. It is too complex; it is too susceptible to misalignment; and it is entirely incapable of accommodating a real-life wobbling wafer.

I did not see how Killeen and Breiland implemented their spectrometer and was compelled to improvise in accordance with the above guidelines. My apparatus is shown in figure 5. Starting from the top, it avoids the use of exotic workstations or languages, custom microprocessors or ROMs. It is driven by an off-the-shelf DOS laptop with a single PCMCIA input/output card running a Visual Basic program, under Windows. The spectrometer, at right, is the simple no-adjustment module described above (with a price less than one half that of the laptop). The illumination source is a light bulb in the small module at the left. Input and output light paths are fixed by the coiled optical fibers, which were selected to achieve the optimum illumination intensity and output collection efficiency. At the ends of these fibers, fixed simple lenses spread the light path out to a 3.5° angle. At the actual 0.6 meter sample distance, this divergence is slightly larger than the sample wobble, ensuring that there will always be a "lucky photon" making it back through the collection path. The software automatically smoothes and re-normalizes the data for each run. Further, the software offers the operator complete default setup options which include automatic start of data collection upon a signal from the process crystal growth apparatus. With the exception of the software, every single item in this system comes straight out of one of several catalogs.

SUMMARY

The instrument above is still a work-in-progress and although it is moving onto our development lines I cannot yet say that it is successfully monitoring product. However, I can say that is has been accepted, indeed embraced, for trial use in a way that simply would not have occurred three years ago. The difference lies in many of the points above. To start with, for this and a number of other instruments, we took the time first to learn the culture, the priorities and the statistical language of our factory customers. We used that knowledge not only to build credibility but to focus on their most important problems, not just our favorite solutions. We learned that complete and unambiguous information was not always necessary and, liberated by that, we were able to focus on much simpler and direct tools. Instead of just inventing, we endeavored to borrow as much as possible, from other scientists and engineers and from catalogs of existing equipment. We tried to deliver non-invasive little systems that offered minimal risk, expense or inconvenience. Finally, we tried to construct tools that had no invisible strings leading back to us, to our expertise, to our loaded salary, systems that could be operated and maintained by anyone, in the hope that such value would ultimately create a much stronger bond.