

High-speed electrical testing of multichip ceramic modules

D. G. Manzer
J. P. Karidis
K. M. Wiley
D. C. Bruen
C. W. Cline
C. Hendricks
R. N. Wiggin
Y.-Y. Yu

This paper reports on the successful application of very-high-performance robotics in the electrical testing of multichip modules using only two probes, breaking with the old traditional array of probes as the primary test method. Complete production line tools include two high-speed Hummingbird® probing robots and precise x-y tables to carry them and a fast, accurate opens-shorts test. To ensure fast probe placement without damaging the part under test requires real-time control hardware and software to operate with extreme precision, flexibility, and programmability to accommodate any part. Finally, because a module can have nearly 100,000 points to be probed, computing an optimal path for the two probes to take for full testing of a part can greatly reduce test time.

Introduction

Work on what was to become Hummingbird* minipositioner-equipped two-point testers began in the mid-1980s to address testing issues in high-density glass-ceramic packaging especially for complex high-end IBM systems. The concept was to replace testing equipment which used a large cluster of probes ("bed of nails") with two single probes moving from point to point very quickly: the Hummingbirds. A single Hummingbird is shown above a ceramic multichip module (MCM) in **Figure 1**. While the research and technology were largely completed for the Hummingbird, the change at IBM from bipolar to CMOS transistor technology so changed the landscape for testing that the Hummingbirds were not employed. Nearly a decade later, the complexity of CMOS machines began to stress existing test technology again, and the Hummingbirds were successfully pressed into service for the glass-ceramic substrates of the S/390* zSeries* machines prior to the year 2000. Since then, the Hummingbirds have gained full acceptance in testing almost every part, from all of the IBM server product MCMs down to the smallest original equipment manufacturer (OEM) single-chip modules. For small modules, a "tray" of parts is placed in the tools. The machine vision and large work envelope of the tool can

test this collection of parts without human intervention until the entire tray is complete.

The first test employed was Time Domain Opens and Shorts (TDOS). A defect-free multichip module will have no net-to-net unintended electrically conductive shorts, nor will it have any point-to-point unintended electrical opens. Later, the Electronic Defect Detection System (EDDS) test, described elsewhere in this issue [1], was also included. The ultimate tools are equipped with both the TDOS and EDDS testers; these tools can conduct both tests by switching the appropriate tester in and out when the probes are on the appropriate pads.

Perhaps the most difficult class of parts tested to date is the complex substrate shown in **Figure 2**. Not only were there large numbers of contact points and total nets to be tested, but this substrate included six layers of organic insulator and thin-film wiring on top of the glass-ceramic. Even these nets were tested. The thin-film wiring is delicate and pliable. A hard probe impact or any scrubbing motion could plow up metal and ruin the part. Both the bare ceramic and the added thin-film layers had to be tested at various times during the manufacturing process.

The Hummingbird minipositioner

The Hummingbird minipositioner is a custom three-axis servomechanism which was designed to provide

©Copyright 2005 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/05/\$5.00 © 2005 IBM




Figure 1

A Hummingbird shown above an MCM.

extremely fast and accurate positioning of an electrical probe tip on or above a high-density electronic substrate. In serialized electrical testing of substrate wiring, high positioning speed is required because of the very large number of moves to be made. Excellent positioning accuracy and probe force control are required in order to avoid damaging the possibly delicate substrates. To provide the fastest and most accurate motion possible, the minipositioner workspace was limited to 13 mm by 13 mm by 1 mm.

This workspace is large enough to cover the planar region of interest in electrical test applications (e.g., some semiconductor chip footprints), yet small enough to permit high accelerations (>500 m/s/s). For applications requiring a larger work envelope, the total workspace is partitioned into an array of cells equal to or smaller than the minipositioner workspace, and the minipositioner itself is mounted on a suitable large-area positioner which moves the minipositioner between cells. In our probing tools, these larger-area positioners are highly accurate, commercial x - y tables. For coarse z motion, the part




Figure 2

Example of a complex substrate (G6 multichip ceramic module).

under test and its fixturing is on a programmable z -height stage. The z axis of the minipositioner is used to sense the position of the part under test, and even detects when it has been moved to the proper height.

x-y positioning mechanism and actuator design

The Hummingbird minipositioner provides x - y positioning with a five-bar kinematic linkage, as shown in **Figure 3**. The design allows the probe to be placed very close to external obstacles, such as another electrical probe from a second minipositioner, and allows for passive elimination of x - y reaction forces. The idea of high-precision “end-effectors” was pursued early and continues to be of interest to industry [2, 3]. The parallel-drive mechanism is kinematically very similar to those of larger robots [4–8] and smaller end-effectors¹ [9] that have been developed in the past. It differs significantly, however, in the details of its dynamically balanced and vertically symmetric design. As can be seen from the figure, two rotary moving-coil actuators are used to drive the two main links (link 1 and link 2) of a five-bar linkage around a common axis. Link 3 pivots around the end of link 1, while link 4 pivots around the end of link 2. A final rotary joint connects links 3 and 4. The last link of this degenerate five-bar mechanism is the center main-shaft itself, which has zero link length. Each of the aluminum-alloy links is designed to provide very high in-plane and

¹An end-effector is a tool or mechanism at the end of the robot arm of the minipositioner.

out-of-plane stiffness yet comprise a mass of only a few grams. Each of the five rotary joints in the linkage consists of two high-precision instrument bearings which are axially preloaded during assembly to completely eliminate all joint clearance and to provide high radial and axial bearing stiffness.

z-axis actuator design

To provide fast and accurate control of the vertical probe position and probe contact force, a miniature servo-controlled actuator was custom-developed to provide 1 mm of straight-line vertical travel with very high accelerations (up to 1,000 m/s/s) and a limited (1-mm) stroke. The basic actuator consists of two C-shaped stator cores separated by a permanent magnet, a coil that surrounds the inner legs of the stator cores, and an armature having two flat teeth which project toward and partially overlap the stator cores to provide a closure path for the magnetic flux, somewhat similar to that in [10]. The actuator operates somewhat similarly to one phase of a hybrid stepping motor, although with a different detailed tooth structure. It is not “stepped,” but rather controlled more like a dc motor. Armature motion away from the equilibrium center position is achieved by energizing the coil, which alters the magnetic flux distribution, thus generating a net force which is roughly proportional to the sign and magnitude of the coil current. A set of balls rolling in V-grooves act as separators and bearings for the armature. The attractive force which exists between the armature and the stator poles maintains the necessary preload on the bearings, even during high lateral accelerations.

Position sensing

The *z*-axis position is measured using an infrared light-emitting diode (LED) to differentially illuminate two phototransistors through an aperture mask which is mounted on the moving armature. The aperture mask is designed so that the analog differential signal from the two phototransistors is directly proportional to the position of the armature.

The *x*-*y* position of the probe is inferred from the orientations of the two main links measured using optical encoders mounted on the ends of links behind the drive actuators. The custom sensors consist of a large stationary grid plate and a compact, low-mass moving element which includes a reference phase plate, illumination, and optical sensors. The moving sensor contains two infrared LEDs and four photodetectors mounted on a common substrate and facing the same direction. Transmission-type encoder operation is achieved by providing a curved retro-reflector to provide backlighting of the grid plate, and the four photodetectors generate two differential quadrature




Figure 3

Actuator mechanism links (bearings removed): (a) Top view; (b) oblique view. (Voice coils are windings which together with a magnetic field comprise the direct drive actuators for the minipositioner linkages.)

signals in typical opto-encoder fashion. High-speed software interpolation of these sinusoidal signals having a spatial frequency of 10 cycles per degree yields a final link angle resolution of 0.001 degree, corresponding to an end-point resolution of approximately 0.4 μ m.

Minimization of reaction forces and torques

To minimize the effects of dynamic interaction between the minipositioner and the large-area positioner on which it is mounted, the minipositioner was designed to be low in mass (950 g) and to generate virtually no reaction forces or torques during *x*-*y* motion. It can be shown that the net *x*- and *y*-axis forces acting on the linkage




Figure 4

Hummingbird work envelope and 5,500- μm calibration grid (in μm).

assembly—and thus the net x - and y -axis reaction forces—must be zero at all times if the overall center of gravity (CG) of the linkage assembly does not move during planar motion. Therefore, the minipositioner linkage was very carefully designed to maintain a fixed CG. Static balance experiments with an actual linkage assembly demonstrated that the true CG of the system was well within 0.5 mm of the main-shaft center for all linkage orientations. Thus, for all practical purposes the system generates no net x - y reaction forces during planar motion.

However, just as a car tire can be statically balanced yet still generate out-of-plane reaction moments because of what is called *dynamic imbalance*, statically balancing the linkage alone is not sufficient to eliminate reaction moments. A key feature of the Hummingbird linkage design which distinguishes it from kinematically similar mechanisms [4, 5, 7, 8] is the nearly complete symmetry of all of the moving elements about the x - y plane passing through the middle of the assembly, as can be seen at least partially in Figure 3(b). Although the linkage assembly is not perfectly symmetric about the mid-plane because of slight asymmetries in the z actuator and sensors, a detailed CATIA model shows that one of the principal axes of inertia for the linkage is within approximately 2° of the z axis. For practical purposes, then, the Hummingbird is dynamically balanced and generates no reaction moments about the x or y axes during x - y motion. Many other fine details may be found in [11].

Laboratory demonstrations show that the one remaining reaction moment resulting from x - y motion—the combined z -axis reaction moment generated by

the two link actuators—can be actively canceled using a third rotary actuator whose drive current is always proportional to the instantaneous sum of the currents in the two main drive actuators. In practice, however, this active torque cancellation has not been required owing to rigid mounting on the heavy x - y tables. Also, no effort was required to eliminate the vertical reaction force from the low-mass z actuator, since these forces are very small and generally not of practical concern. If one holds the minipositioner while it goes through its movements, one feels twisting moments as if the device body were trying to rotate about its center shaft. The moving masses and resulting inertias are small enough that the resulting torque does not punish the holder of the device.

Critical component life

There was initial concern over the reliability of the mechanical parts doing a heretofore unheard-of application. The flex cables connecting the probe and powering the z actuator, link coils, and sensing were tested successfully for several billion cycles under high-frequency (≤ 175 Hz) bending motion to exceed that experienced during the maximum linkage motion without electrical or mechanical fatigue failure. Life-tests of preloaded bearings under a variety of oscillating motion conditions with amplitudes of 0.25° – 19.0° , frequencies of 60–230 Hz, and inertial loads similar to the design loads also completed several billion cycles without substantial loss in bearing stiffness, which would have indicated incipient fatigue failure.

In practice on the test floor, the mechanical, electromechanical, and electronic durability of the Hummingbirds, controllers, and software have been spectacular. The tools have run 24 hours per day for weeks on end, stopping only for loading of a new part. The typical stop for failure is when the very hard tungsten carbide probe tips have been worn down to useless shorting nubs. In this case the impaired Hummingbird is swapped out for a good one, and the blunted probe tip is easily exchanged for a new one. After a quick z -actuator recalibration, the Hummingbird is ready for service again. The other rare failure mode is human—an operator catching a coat sleeve on a probe tip, or a part mounted incorrectly so that a disastrous crash results.

Hummingbird kinematics and calibration

The Hummingbird owes its high accuracy and precision not only to its electromechanical design, but also to a calibration methodology which aligns it with the x - y tables and adjusts the modeled kinematics to match that of the actual probe tip position to within a few microns. With a microscope and machine vision, an automated procedure can place the probe in the nine points (depicted in Figure 4) within the workspace, and any error in

position is nulled out with the x - y tables while viewing the position of the probe tip above a microscope. Using this nulling information, an optimal fitting method embedded in the control system software is run to remodel the kinematics. This method is so accurate that a 1- μm bend in the probe tip would be detected by the software as a difference in probe length, and the inverse kinematic model would be duly corrected. This procedure is executed whenever a new probe tip or complete Hummingbird assembly is installed on a tool, or if any part of the system has been otherwise disturbed. The method is a nonlinear optimization problem, solved by a Gauss-Newton least-squares method [12].

Hummingbird control

Kinematic calibration is essential, since the control system for x - y motion is not in Cartesian coordinates—the controlled variable of each axis of the linkage is its angle. To make matters worse, these angles are measured at the very back of the linkages, far from the probe tip.

For the link control problem, the measured values are the joint angles and the computed values are the currents to be injected into the coils, which are proportional to the torques generated. From here the rigid body dynamics of the Hummingbird linkage, defined in **Table 1**, determine the motion. These dynamics are identical to those of many large industrial robots, and similar control schemes can be adapted to control the Hummingbird [13]. Because the bearings are so nearly frictionless and the actuators are direct-drive and easily characterized, the Hummingbird is an ideal mechanical device for developing the necessary nonlinear control methods. However, the Hummingbird required exceptional digital hardware performance for its real-time control implementation.

Even with today's digital signal processors, it would be difficult to control the six high-speed axes of two Hummingbirds on one processor. The design choice of the day was to employ a processor for each axis, and in addition yet more processors to handle inverse kinematic calculations, sequencing of motion per axis, coordination between the two Hummingbirds, and communication with the tool controller, an industrial PC. The result was a nine-processor system, with each processor being an Inmos T805 transputer². The control processors for each axis run the control calculations at a strictly enforced sample period on 186 μs . A processor is dedicated to supervising and coordinating the motion of each Hummingbird, and the ninth processor, the root, is responsible for scheduling moves, coordinating the two Hummingbirds, and accepting commands and responding to the host PC. A schematic diagram of




Figure 5

Control system hardware; the compute node labels refer to the root processor, supervisors, long Hummingbird axis, short Hummingbird axis, and z axis.

Table 1 Nonlinear ordinary differential equations (ODE) of motion of the Hummingbird.

$$\begin{pmatrix} \ddot{\theta}_1 \\ \ddot{\theta}_2 \end{pmatrix} = -A^{-1}B \begin{pmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \end{pmatrix} + A^{-1} \begin{pmatrix} \tau_1 \\ \tau_2 \end{pmatrix},$$

where

$$A = \begin{pmatrix} a & b \cos(\theta_1 + \theta_2) \\ b \cos(\theta_1 + \theta_2) & c \end{pmatrix},$$

$$B = \begin{pmatrix} 0 & b \sin(\theta_1 + \theta_2) \\ b \sin(\theta_1 + \theta_2) & 0 \end{pmatrix}.$$

- a , b , and c are constants having units of inertia.
- τ_1 and τ_2 are the torque inputs.
- θ_1 and θ_2 are the joint angles.

this control hardware is shown in **Figure 5**. The interconnections are the actual serial links used by T805s for interprocessor and interprocess communication, four 20-Mb/s links per processor. The root and supervisor processors run a real-time operating system kernel, while the controls for the links—defined as long and short—and z axis each run their processors using the native hardware task scheduling for their multiple tasks. Their highest-priority task is running the servo control algorithms, which include exchanging information of positions and torques over the serial links, reading the sensing electronics, and commanding the current amplifiers. There are many low-priority tasks, such as communication, which are run and are scheduled by the processor hardware. Even with today's digital signal

²The T805 transputer was a product of Inmos, a member of the SGS-Thompson group of companies.




Figure 6

A 500-count (~4 mm) short link move.

processing (DSP) and real-time operating system technology, it is difficult to replicate the performance of this flexible architecture.

Figure 6 shows the motion of one axis of the linkage in response to the commanded current generated by the control laws. The bottom of the figure shows the actual current in the actuator, the control law being such that a near “bang-bang” motion for much of the move, roughly 4 mm, is commanded. The final settle-out of the move is the designed response of the nonlinear control laws, calculated to cancel out all nonlinearity seen in the equations of Table 1 and to decouple the interactions between the two links [13]. The commanded currents for the large motion are filtered to avoid producing square waves, since this would place unnecessary stress on the linkage structure and bearings. Since this move went from a relatively low-inertia state to a higher-inertia one, a larger deceleration force was required, taking the control system nearly to its designed limit of 5 A. The probe tip and the z actuator experienced accelerations in excess of 500 m/s/s. The position error graph shown in Figure 6 is valid only when the servo mode has changed from the large “seek” to a final settle-out. The unit count translates




Figure 7

Two Hummingbirds mounted on x - y tables working on an MCM.

to roughly $0.4 \mu\text{m}$ of probe tip motion in the linkage orientation of this move. Thus, a fairly large move is accomplished to submicron precision in well under 20 ms. To our knowledge, this performance is still unchallenged by that of any other robotic manipulator. An early version of the system and its performance may be seen in [14].

Finally, typically two Hummingbirds are employed on a tester, along with the control electronics, software, and commercial x - y tables. **Figure 7** is a photograph of such a configuration at work on an MCM.

Time domain opens and shorts measurement system for flexible testing

Flexible electrical test methods that are independent of a product footprint are critical to the packaging business. The Time Domain Opens and Shorts (TDOS) measurement system is a defect detection method for use on flexible probers that is functionally equivalent to the “bed-of-nails” method of detecting opens and shorts in substrates. While the bed-of-nails method uses resistive measurement techniques in combination with voltage stress, most flexible prober measurement systems either forgo the voltage stress or have voltage stress but only run in a pass/fail mode at a certain resistive cutoff level. TDOS provides voltage stress equivalent to the bed-of-nails method and can provide a resistive reading on high-resistance shorts at or above $100 \text{ k}\Omega$ at a speed faster than that supplied by a standard capacitance meter.

The TDOS measurement system is a collection of three distinct tests. The first is a single-point stimulus and response method in which external discrete electrical components are placed in series with the net under test

and a voltage pulse is applied. The output waveform is analyzed to determine whether a short or open exists and if so, to determine its characteristics. A second part of TDOS, resistive opens, measures the resistance of a net directly using two probes and a constant current. A third piece, shorts isolation, measures any shorts called by the single-point measurement against other suspect nets and the voltage plane to pinpoint the nets involved in the short and to provide a more exact measurement of the shorted resistance. The use of all three tests on the same prober helps to give an accurate picture of the product under test.

TDOS measurement system hardware

The TDOS measurement system electronics consist of a voltage pulse generator, a large capacitor and resistor, buffer amplifiers, relays, analog-to-digital converters, and a digital input-output (DIO) card to communicate with the prober.

The first piece of the measurement system is a single-point opens and shorts detector in which external resistor/capacitor components are placed in series with the net under test [15]. The stimulus is a square wave of known voltage V_{in} . The choice of these components and the input voltage used is dependent on the product specifications. Systems for testing current packaging products typically use a pulse of 100 V for 300 μ s, an external capacitance C_0 of approximately 1,000 pF, and an external resistance R_0 of around 1 G Ω . The resulting output waveform is then analyzed. **Figure 8** shows the circuit for this method: C_{net} is the net under test, and R_{leak} is the leakage resistance between C_{net} and a possible shorted net $C_{shorted}$.

The solution to this circuit can be simplified by making certain choices in the values of the external components C_0 and R_0 . The simplified equation of the circuit in Figure 8 when $C_0 \gg C_{net}$, $R_0 C_0 \gg R_{leak} C_{net}$, $R_0 C_0 \gg R_{leak} C_{shorted}$ is

$$V_{out} = V_{in} \left\{ \frac{C_{net} + C_{shorted}}{C_0} \cdot \exp[-t/R_0(C_{net} + C_{shorted} + C_0)] - \frac{C_{shorted}}{C_0} \cdot \exp(-t/R_{leak} C_{shorted}) \right\}. \quad (1)$$

Inspecting the voltage response of the output waveform at the rising edge of the pulse determines the characteristics of the net under test. A flat output pulse indicates a good net, a dead short, or an open, because the circuit acts like a capacitive voltage divider in these cases. For a good net or an open net, the simplified equation is

$$V_{out} = V_{in} \frac{C_{net}}{C_{net} + C_0}. \quad (2)$$

For a dead short, the simplified equation of the circuit is




Figure 8

Single-point opens and shorts circuit.

$$V_{out} = V_{in} \frac{C_{net} + C_{shorted}}{C_{net} + C_{shorted} + C_0}. \quad (3)$$

As with traditional capacitance measurements, the TDOS single-point technique requires a database of known good or “learned” values from several parts to determine a nominal value for that net. If the measured value falls close to the learned value, the net is good. A value below the learned value by a certain percentage indicates an open due to reduced net capacitance. A measured value that exceeds the learned value by a certain percentage is called a short because of the addition of $C_{shorted}$ into Equation (3). However, typical capacitance techniques can detect only low-resistance shorts.

The TDOS system also has the ability to detect and call out shorts above 100 k Ω , a near short, by using information about the slope of the top of the response waveform. Beyond the leading edge of the pulse, an increase in voltage indicates that another net is being charged. This implies that a high-resistance short exists. The slope of the response, (dV/dt) , in Equation (4) is inversely proportional to the near-short resistance R_{leak} as long as the conditions of Equation (1) are still met:

$$\frac{dV}{dt} = V_{in} \frac{C_0}{(C_0 + C_{net})^2 R_{leak}}. \quad (4)$$

By using a very-fast-rise-time pulse generator as the input to this circuit and an analog-to-digital converter (ADC) with a sample rate greater than 1 MHz, the slope of the response can be measured. By measuring the slope of the output waveform and calculating the net capacitance C_{net} from the rising edge of the pulse, Equation (4) gives the leakage resistance directly. If the waveform shows that it is fully charged, we can also compute the combined capacitance of the shorted network by using the voltage level at the end of the pulse

in Equation (2). Thus, all of the information required to characterize the net under test can be directly measured or computed from the directly measured values.

Voltage stress capability of TDOS

Voltage stress is important in detecting shorts because leakage resistance between two nets can decrease significantly under high voltage. Some near shorts or latent shorts can be detected only under high voltage. When only one probe is used for detecting shorts, as is the case with TDOS, the voltage is applied between the net under test and a reference plane; the reference plane can be a voltage plane within the part or external to the part. There is a constant voltage between the net and voltage plane for the duration of the pulse, but the net-to-net stress voltage is transient. The net-to-net voltage is given by

$$V_{\text{net}} = V_{\text{in}} \{ \exp[-(t_{\text{stress}} + t_{\text{measure}})/(R_{\text{leak}} C_{\text{shorted}})] \}. \quad (5)$$

Specifications for voltage stress consist of a voltage applied for a certain minimum time, and since the stress voltage in TDOS is a transient voltage that decays exponentially, a voltage greater than that of the product electrical specification must be applied in order to obtain the required stress for the required time period. On the basis of Equation (5), in order to guarantee the minimum stress time and minimum stress voltage, the following inequality must hold:

$$R_{\text{leak}} C_{\text{shorted}} > \frac{t_{\text{stress}}}{\ln\left(\frac{V_{\text{in}}}{V_{\text{stress}}}\right)}. \quad (6)$$

The worst case is a short to the smallest net C_{shorted} . By setting the R_{leak} threshold high enough, it is easy to fulfill the requirements of high leakage sensitivity and short measurement time [16].

TDOS resistive opens detection

The TDOS resistive opens detection is unique not for its circuit, which is simply a constant-current source in parallel with the net under test fed into a buffer amplifier, but because this circuit can be left in the path of the single-point TDOS measurement with no detrimental effects to either of the two tests, thereby eliminating the need for isolation relays and the time delays associated with switching them. The current used is determined by the electrical specification for the product; however, accuracy is lost if the current is too low. This circuit has been successfully used to provide near-parametric-resistance measurements for some packaging products.

TDOS isolation of shorts

The shorts isolation portion of the TDOS measurement system uses the results of the single-point shorts test to determine the location and value of the short. This is important, because some substrates can be repaired. There are two shorts isolation circuits in the TDOS system. The first can be used for most nets and voltage planes with low to medium capacitance; the second is for use when the first cannot sufficiently charge the net within the time allocated to this test because of the large capacitance of the net under test. The challenge in shorts isolation is the large range of resistances that may be encountered and the need for stress voltage. The shorts isolation circuit can be solved to come up with the following equation for the leakage resistance R as a function of the output voltage measured (A and B are constants that depend on the specific circuit components chosen and are determined by the range of resistances to be measured):

$$R_{\text{leak}} = \frac{A}{B + V_{\text{measured}}}. \quad (7)$$

For a range of 10 k Ω to 333 m Ω , circuit components are chosen such that A is approximately 4,000,000 and B is approximately 500,000.

Occasionally the second shorts isolation circuit is needed. This circuit technique also uses the slope of the response to determine the resistance. A positive slope on the output indicates a resistance above a preset cutoff level, and a negative slope indicates a resistance below the cutoff level. The value of the resistance can be estimated using the actual slope value compared to a calibrated standard. The cutoff level can be adjusted higher or lower according to the product specification by changing some component values within the circuit [17].

The TDOS measurement, combining three distinct tests, can nearly fully characterize the package under test. Measurement times for opens and shorts detection are of the order of 6 ms to 15 ms depending on the ADC and processing hardware used. This makes the TDOS system fast and flexible enough to benefit from a high-speed probing system.

Hummingbirds and global test path optimization

Given that testing involves tens of thousands of probing contacts per MCM and that each MCM might be tested multiple times during its construction, some thought has been given to optimizing the motions of the tester. Having nothing to do with control, this is a logistical problem in predetermining the sequence of desired probe touchdowns on the part. The Hummingbirds are so fast that little can be gained in total test time by optimizing their individual position schedules. Much of the total test time is taken by the relatively slow x - y table moves. Here, a globally

optimized solution can affect total test time. The normal algorithm used to determine where to move the tables is a heuristic, locally greedy one; the software makes decisions with a narrow vision of future consequences. A casual observer watching the tool in this mode of operation can notice some table moves that may appear odd or wasteful. However, a global solution is known to be “hard” in the theoretical sense. The problem is made more difficult by the fact that the tables can easily collide if a poor schedule is requested, and some probe orientations are impossible because of the interference of both the x - y tables and the Hummingbirds being in the same workspace. What follows is a tractable method to address this problem.

The module is divided into 12-mm partitions for the purpose of testing. Optimizing these partitions for an optimal covering is yet another problem to be solved, but the methodology which follows can produce a reasonable covering. Once these partitions are defined, the movement of the x - y tables becomes, with some added complexity, akin to the traveling salesman problem. While there are many numerical routines that solve this problem well and rapidly, our problem is easy to formulate and “solve” using a simulated annealing methodology [18–20].

Traveling salesman and simulated annealing problems on MCMs

There is a cost associated with moving the tables. The formal cost function calculates the cost as the length of a complete path to cover the whole module in microns; it is the sum of the longest axis moves for each table move in the path. On a per-point basis, the cost of moving the probes from one pair of points to the next is determined by the current placement of both probes as a function of the destination placement. All costs are associated with the need to move one or both Hummingbirds into a new partition on the module and so require the tables to move. If the movement does not cause either table to move, there is no associated cost. In this case, both new points are in the same partitions as the previous points, and only the minipositioners have to move. The following section defines the criteria of costs associated with moving the tables.

Collision avoidance costs

1. The destination points are in the same partition. This forces the partition to be split into two separate partitions.
2. The destination points must be swapped AND the table movements must be sequenced.
3. The destination points must be swapped.
4. The table moves must be sequenced.

The cost of moving the tables from one partition to another is determined by the table that has to make the

longest move in any one direction; all table axes are moved asynchronously. It is assumed that all axes move at approximately the same speed. The length of the move is measured in terms of the number of microns the axis must travel in order to proceed from its current position to the destination position.

Collision detection

Collisions occur when the movement of the tables from one partition to another results in the tables colliding while in motion. This happens either when the tables cross paths or when a “race” condition exists: They are moving in the same direction and the trailing table catches up with the leading table and collides with it. A *collision-avoidance algorithm* is encoded in the cost function to detect collisions and take action to avoid them. The corrective action involves either swapping the points between the two probes and/or “sequencing” the table moves. Sequencing implies that a single movement of both tables is split into two separate moves. The first move of the sequence will move only one or more selected axes; the second, which completes the total move, is not started until the first move of the sequence is completed.

Solution path

Since there is no cost associated with moving to points that do not require the tables to move, computing time can be saved by removing these points from the path. The annealing algorithm runs only on the points that require the tables to move. This is accomplished through a pre-processing step before annealing begins. After the annealing is completed, the points that were removed from the path are re-inserted into the final solution in a post-processing step.

Results of simulated annealing and the MCM problem

The software developed was tested against several known completely solved traveling salesman problems. The implementation of the algorithm achieved solutions very close to the known global minimums. Using this method achieves better than a 30% improvement in test path length over the current locally greedy algorithm on the testers. This can help in allocating fewer testers for a particular product.

Conclusions

This work began with a simple question posed in an informal meeting: “How fast can two probes go?” The resulting answer, after years of work, is that two probes can support 40 to 100 probing cycles per second, not counting the time to do a measurement or any coarse table moves. Since some measurements can be very rapid, this solution has proved viable and has been the mainstay

in testing high-density ceramic parts. The Hummingbird is used as an exemplar of engineering and robotic discipline in courses at Johns Hopkins University, the University of Michigan, and McGill University. Its role in testing continues to evolve.

Acknowledgments

Moving two probes required the invention of new test technology such as TDOS. Also, implementing this solution required the work of a great many more individuals than the authors of this paper. The ingenuity involved and risk-taking incurred spanned many disciplines, business organizations, and continents. We thank many individuals in IBM Japan, IBM East Fishkill, and especially IBM Research in Yorktown Heights, New York, for an unfailing dedication to building a system that could barely be imagined, let alone created.

*Trademark or registered trademark of International Business Machines Corporation.

References

1. E. J. Yarmchuk, C. W. Cline, and D. C. Bruen, "Latent Defect Screening for High-Reliability Glass-Ceramic Multichip Module Copper Interconnects," *IBM J. Res. & Dev.* **49**, No. 4/5, 677–685 (2005, this issue).
2. A. Sharon and D. Hardt, "Enhancement of Robot Accuracy Using Endpoint Feedback and a Macro-Micro Manipulator System," *Proceedings of the American Control Conference*, San Diego, June 1984, pp. 1836–1842.
3. R. L. Hollis, "Design for a Planar XY Robotic Fine Positioning Device," *Proceedings of the ASME 1985 Winter Annual Meeting*, Miami, November 1985, pp. 291–298.
4. H. Asada and K. Youcef-Toumi, "Analysis and Design of a Direct-Drive Arm with a Five-Bar-Link Parallel Drive Mechanism," *J. Dynamic Syst., Measurement, & Control* **106**, 225–230 (September 1984).
5. T. Kokkinis, L. Leka, S. Telschow, J. Wilson, and T. Miller, "Kinematics and Design of a Direct-Drive Robotic Arm for Microelectronics Applications," *Proceedings of the IASTED International Symposium*, AGFA Press, B-1640 Mortsel, Belgium, May 1987.
6. Z. C. Yan, D. X. Guo, T. Takagami, and M. Takano, "A Study on Vibration of Flexible Robot Arm—Use of Viscous Dampers and Structural Damping to Control of a 5-Bar Linkage Robot," *J. Faculty Eng. Univ. Tokyo* **XXXIX**, No. 2 (1987).
7. H. Kazerooni, "Design and Analysis of the Statically Balanced Direct-Drive Robot Manipulator," *Robotics & Computer-Integrated Manuf.* **6**, No. 4, 287–293 (1989).
8. K. Youcef-Toumi and A. T. Y. Kuo, "Design and Control of a High-Speed Direct-Drive Manipulator," *Int. J. Prod. Res.* **27**, No. 3, 375–394 (1989).
9. H. Kazerooni, "Direct-Drive Active Compliant End Effector (Active RCC)," *IEEE J. Robotics & Automation* **4**, No. 3, 758–766 (June 1988).
10. B. A. Sawyer, "Magnetic Positioning Device," U.S. Patent 3,457,482, July 22, 1969.
11. J. P. Karidis, G. McVicker, J. P. Pawletko, L. C. Zai, M. Goldowsky, and R. E. Brown, "The Hummingbird MiniPositioner—Providing Three-Axis Motion at 50 G's with Low Reactions," *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, Nice, France, May 1992, pp. 685–692.
12. D. G. Manzer and L. C. Zai, "Kinematic Calibration by Nonlinear Estimation," *Research Report RC-17176*, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, September 1991.
13. L. C. Zai, L. F. Durfee, D. G. Manzer, J. P. Karidis, M. P. Mastro, L. W. Landerman, and K. C. Ho, "Control of a Hummingbird Minipositioner with a Multi-Transputer MARC Controller," *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, Nice, France, May 1992, pp. 534–541.
14. J. P. Karidis, L. C. Zai, D. G. Manzer, and L. F. Durfee, "The Hummingbird Minipositioner—Robotics Too Fast To See," *Video Proceedings of the 1992 IEEE International Conference on Robotics and Automation Conference*, Nice, France, May 1992.
15. K. Woo, "Apparatus and Method for Resistive Detection and Waveform Analysis of Interconnection Networks," U.S. Patent 5,266,901, 1993.
16. J. Craig and K. Woo, "Voltage-Stressing and Testing of Networks Using Moving Probes," U.S. Patent 5,438,272, 1995.
17. H. Bhatia, D. Long, and K. Wiley, "Method and Circuit for Electrical Testing of Isolation Resistance of Large Capacitance Networks," U.S. Patent 6,573,728, 2003.
18. E. H. L. Aarts and J. H. M. Korst, *Simulated Annealing and Boltzmann Machines: a Stochastic Approach to Combinatorial Optimization and Neural Computing*, Wiley-InterScience Series in Discrete Mathematics and Optimization, Wiley-Interscience Press, New York, 1989.
19. M. E. Johnson, *Simulated Annealing and Optimization*, American Science Press, Columbus, OH, 1988.
20. S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, "Optimization by Simulated Annealing," *Science* **220**, 671–681 (1983).

Received September 15, 2004; accepted for publication March 9, 2005; Internet publication August 17, 2005

Dennis G. Manzer *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (manzer@us.ibm.com).* Dr. Manzer is a Research Staff Member at the IBM Thomas J. Watson Research Center. He received a B.S. degree in engineering and applied science at Yale University in 1978 and M.S. and Ph.D. degrees at Cornell University in 1980 and 1983, respectively. He has worked in robotics, printing, electrical testing, high-speed image processing, picosecond imaging analyses of working microprocessors, and control in disk drives.

John P. Karidis *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (karidis@us.ibm.com).* Dr. Karidis is an IBM Distinguished Engineer with more than 45 issued patents and more than 20 years' experience in advanced technology and product development. He received his B.S., M.S., and Ph.D. degrees in mechanical engineering from Pennsylvania State University in 1980, 1982, and 1983, respectively, subsequently joining IBM Research. After spending ten years developing advanced hardware technologies for printing and robotics applications, he spent several years in various IBM product divisions developing unique products such as the "butterfly" keyboard on the ThinkPad* 701C and the ThinkPad TransNote* portfolio computer. Currently, he is working in the Research Division, with a focus on cooling technologies.

Kathleen M. Wiley *IBM Systems and Technology Group, Mid Hudson Research Park, 1580 Route 52, East Fishkill, New York 12533 (kmwiley@us.ibm.com).* Ms. Wiley is a Development Engineer at IBM in East Fishkill. She received a B.S. degree in electrical engineering from Lehigh University in 1989, subsequently joining IBM, where she has worked on the development of inspection and test equipment for packaging products and is an author or coauthor of four patents. Ms. Wiley is a member of the Institute of Electrical and Electronics Engineers.

Dominic C. Bruen *IBM Systems and Technology Group, Mid Hudson Research Park, 1580 Route 52, East Fishkill, New York 12533 (bruen@us.ibm.com).* Mr. Bruen is an Associate Engineer in the MLC Test Engineering Department at the IBM facility in East Fishkill, New York. He received an A.A.S. degree in electrical technology from DeVry Institute of Technology in 1990, a B.A. degree in computer science from Mount Saint Mary College in 1995, and an M.S. degree in computer science from Marist College in 2000.

Christopher W. Cline *IBM Systems and Technology Group, Mid Hudson Research Park, 1580 Route 52, East Fishkill, New York 12533 (cline@us.ibm.com).* Mr. Cline is an Advisory Engineer in the MLC Test Engineering Department at the IBM facility in East Fishkill, New York. He received a B.S. degree in 1983 and an M.S. degree in 1984, both in mechanical engineering, from Columbia University. Mr. Cline is the author or coauthor of two patents and four technical papers.

Charles Hendricks *IBM Systems and Technology Group, Mid Hudson Research Park, 1580 Route 52, East Fishkill, New York 12533 (chendric@us.ibm.com).* Dr. Hendricks joined IBM in 1974 after receiving a Ph.D. degree in mechanical engineering from Pennsylvania State University. In his IBM career he has held engineering and management positions in semiconductor manufacturing engineering and development, semiconductor

program management, and multilayer ceramic (MLC) packaging manufacturing engineering. He is currently manager of MLC New Products Manufacturing Engineering and Diagnostics in the East Fishkill MLC Manufacturing organization.

Robert N. Wiggin *IBM Systems and Technology Group, Mid Hudson Research Park, 1580 Route 52, East Fishkill, New York 12533 (Robert_N_Wiggin@us.ibm.com).* Mr. Wiggin joined IBM in 1979 after receiving a B.S. degree in electrical and computer engineering from Clarkson College. He leads the development and application of cost-effective electrical test assurance methodologies and repair solutions at the Multilayer Ceramic manufacturing facility and manages the MLC Test/Repair Engineering Department.

Yuet-Ying Yu *IBM Systems and Technology Group, 2455 South Road, Poughkeepsie, New York 12601 (yuy@us.ibm.com).* Ms. Yu is currently an Equipment/Process Engineer in MCM Burn-in and Test at the IBM Poughkeepsie facility. She has worked in the MLC Test Engineering Department for the past 15 years as an Equipment Engineer. She received her B.S. degree in mechanical engineering from City College of New York in 1984 and her M.S. degree in mechanical engineering from Rensselaer Polytechnic Institute in 1989. Ms. Yu is an author or co-author of 18 U.S. patents.