

A multilevel approach to the control of a chemical–mechanical planarization process

Roland Telfeyan,^{a)} James Moyne,^{a)} and Nauman Chaudhry
*Electrical Engineering & Computer Science Department, Center for Display Technology & Manufacturing,
The University of Michigan, Ann Arbor, Michigan 48109-2108*

James Pugmire and Scott Shellman
Statistics Department, Brigham Young University, Provo, Utah 84602

Duane Boning and William Moyne
*Microsystems Technology Laboratories, Massachusetts Institute of Technology, Cambridge,
Massachusetts 02139*

Arnon Hurwitz
SEMATECH, Austin, Texas 78741

John Taylor
Compugenesis, Inc., Austin, Texas 78739

(Received 2 October 1995; accepted 11 March 1996)

A multilevel hierarchical control system has been designed and is being applied to chemical–mechanical planarization (CMP) process control. The current implementation of the control system incorporates closed-loop run-to-run (R2R) control and open-loop real-time monitoring, and can accommodate inter-cell control. The R2R control element is enabled via a generic cell controller (GCC) implementation that provides flexible automated control of the process and equipment, multiple control algorithm branches and fuzzy logic decision capability among the branches, simulation capabilities, hardware and software independence, and extensive GUI support for control and data analysis. The R2R element utilizes a linear approximation multivariate control algorithm (branch) that supports individual exponential weighted moving average (EWMA) modeling of advices (outputs), weighting of inputs, granularity, and input bounding. The real-time element of the control system utilizes a partial least squares (PLS) algorithm to identify real-time equipment input trace patterns and relate these patterns to alarming conditions. The entire control system is designed to provide multivariate control of CMP process removal rate and uniformity. As a result of extensive design of experiments and testing, the R2R control level has been demonstrated to achieve good control of removal rate and fair control of uniformity. The addition of the real-time element is expected to improve process control and reduce R2R process noise, thus leading to a more effective R2R control element. © 1996 American Vacuum Society.

I. INTRODUCTION

Control of semiconductor processes has been and continues to be a topic generating significant interest in both industry and academia. The study and implementation of control in the semiconductor manufacturing facility is usually broken down into three fundamental levels: real time (time critical), run-to-run (R2R), and factory level (intercell). The generally accepted vision is that these three levels of control together comprise a hierarchical control scheme for the entire facility as shown in Fig. 1.

In real-time process control, the product recipe is modified according to measurements taken *in situ* (during a process run). These measurements and recipe modifications are made continuously (or at high sampling rates); thus it is significant that they occur in a timely fashion. In run-to-run (R2R) control the measurements and product recipe modifications are carried out *ex situ*, i.e., between machine runs. (The term “run” means a single wafer processed individually

or a set of wafers processed simultaneously.) Factory-level or intercell control utilizes feedback and feed-forward mechanisms between processes to improve the quality of the entire fabrication process. The vertical cross section of the control scheme given in Fig. 1 depicts a number of significant research, development, and standardization problems at each of these control layers.^{1–3}

The effectiveness of control at any level is dependent upon the ability to sense numerous process and equipment parameters and to actuate equipment settings. Consequently, the inability to sense numerous process environment parameters of certain fundamental semiconductor manufacturing processes [like reactive ion etching (RIE) and chemical–mechanical planarization (CMP)] has impeded the development and implementation of real-time control.²

The realization of an effective and robust multilevel control system in semiconductor manufacturing necessitates important design requirements at each control level. At the R2R control level there are at least three such requirements. First, the controller must provide control over a wide range of

^{a)}Electronic mail: moyne@umich.edu

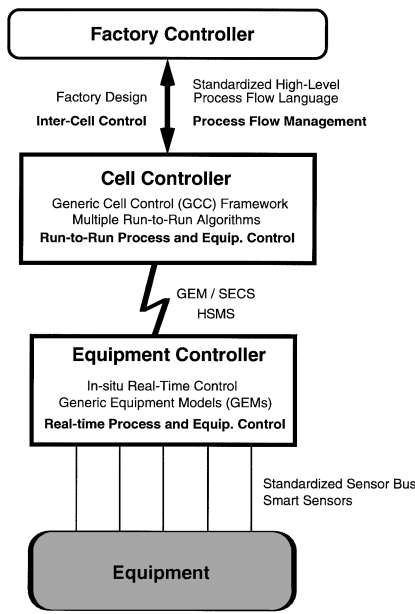


FIG. 1. Semiconductor manufacturing facility hierarchical control scheme.

process and equipment operating conditions. In order to do this, the controller generally must utilize a number of sequential control algorithms in a complementary fashion.^{4,5} Second, the controller must be able to provide control in the absence of real-time control and must be able to incrementally accommodate real-time control as technology advances. Third, the controller must be able to provide control as an integral part of the aforementioned hierarchical control scheme.⁶

We have designed a R2R control system that meets these requirements. The design is based on generic cell controller (GCC) concept. The GCC enabler provides a dynamic and reusable solution, provides for the complementary utilization of multiple sequential algorithms, provides control in tandem with as well as in the absence of real-time control, and can be configured to operate as an integral part of a hierarchical control scheme.

At the real-time control level, it is important to note that sensory or actuation technology for a process such as CMP is often insufficient to support this level of *in situ* control. However, as a precursor to control, the system design should provide for real-time monitoring of available equipment inputs so that they may be employed to provide a level of real-time diagnostics. This diagnostic capability must complement the R2R control capability; thus a requirement of the real-time monitor is that it provide a monitoring and diagnostic capability in the face of changing operating conditions (as suggested by the R2R controller). We have designed a real-time monitoring and diagnostics system that meets these requirements. With this design, equipment inputs are monitored in order to infer if the process has exceeded predefined specification limits, signifying that an exception condition has been detected. These exception conditions may be deduced by (1) an analysis of the current values of one or more equipment inputs, or (2) an analysis of the history or

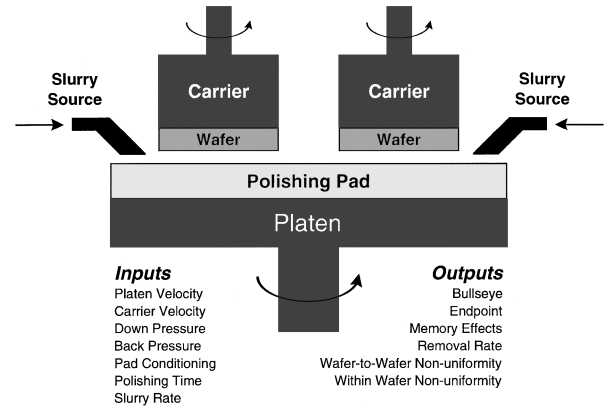


FIG. 2. Schematic of a typical CMP process.

“footprint” of one or more inputs. This information can then be utilized by a higher level (R2R) control system to pause, abort, or stop the process or, in more sophisticated systems, to alter the process to correct the problem.

II. BACKGROUND

A. Chemical–mechanical planarization

Chemical–mechanical planarization has become a widely accepted technology for multilevel interconnects. CMP of dielectric films is the planarization method of choice for 0.35 μm device technology. In addition to providing planarization, CMP has also been shown to reduce defect density and define vertical and horizontal wiring.⁷

CMP is basically a surface planarization method in which a wafer is affixed to a carrier and pressed face down on a rotating platen holding a polishing pad as shown in Fig. 2. A silica based alkaline slurry is applied during polishing thus providing a chemical and mechanical component to the polishing process. The general process goal is the preferential removal of high material across the wafer. Typical process

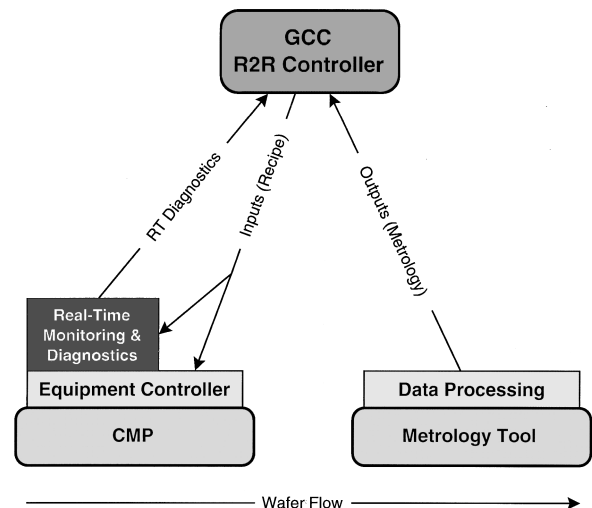


FIG. 3. Data and product flow in a typical CMP process.

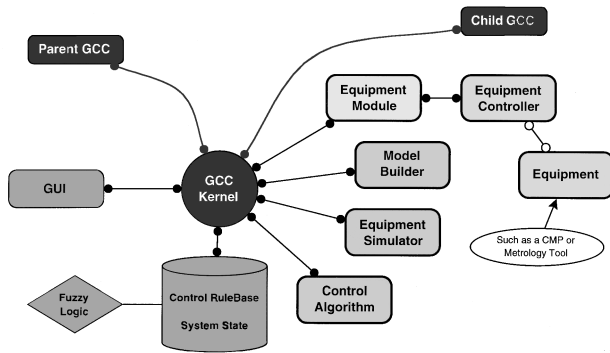


FIG. 4. Schematic of a typical GCC system.

metrics include removal rate and within-wafer uniformity. Equipment and process parameters that are typically utilized to control the process include polish time, pressure, rotation speed, and parameters that impact the conditioning of the polishing pad such as conditioning profile. A multilevel control system utilizing these metrics and parameters is shown in Fig. 3.

There are a number of characteristics of CMP that make it an ideal candidate for the development, implementation and test of multilevel control. First, the process is not well understood. This combined with factors such as inconsistency and degradation of consumables, and lack of sensors and actuators makes CMP a challenging candidate for control. Second, as there is a lack of *in situ* sensors for CMP, *in situ* control is not yet feasible; thus a system that features R2R control combined with real-time equipment monitoring and diagnostics appears to be the appropriate form of control that can be applied to CMP at this time. Third, as wafer surface roughness directly impacts the CMP process and results from upstream processes, CMP is an ideal candidate for future integration of intercell control.

The CMP process has been described in much greater detail elsewhere in the literature.^{7,8} The CMP R2R control problem is detailed in Refs. 6 and 9. A summary of the state of the art of CMP utilization including a discussion of limitations of the process and its control is presented in Ref. 10. A second effort focused on the development of CMP R2R control is described in Refs. 11 and 12.

B. The generic cell controller

The GCC is a discrete control mechanism that utilizes a relational database as opposed to procedural code to store sequential control information. The theory of operation of the GCC is documented in the literature.^{13,14} The main feature of the GCC that makes it an attractive R2R control enabler is that the database schema is tailored for the storage of event driven sequences that dictate how the system will respond to events; these sequences of events and responses are stored as data in the database. For this reason the GCC is capable of enabling complex and dynamic control scenarios that are characteristic of many R2R control systems. Further, due to GCC database schema and software interaction speci-

fications, a very high degree of modularity is established with GCC applications. This results in both high portability and transferability of software, and a capability to easily incorporate commercially available software components into the system.

The GCC provides for intelligent routing between the various software components, or modules, involved in the R2R control task. These modules may include commercially available software such as communications drivers and controller algorithms. The block diagram in Fig. 4 depicts the various components of an example GCC system. The GCC provides an environment for the comparative evaluation of optimization and control algorithms, as it can incorporate any number of these algorithms and provide sequences for their selective utilization. The algorithm evaluation can include the investigation of paradigms for the complementary utilization of a number of control algorithms to achieve more robust control.^{4,14}

Related results associated with the development and deployment of the GCC implementation described in this article have also been documented in literature. Specifically, practical issues associated with the utilization of a R2R control algorithm are discussed in Ref. 15. Issues of R2R control that relate specifically to CMP are discussed in Ref. 9. A detailed discussion of the design and operation of the current GCC in providing CMP process R2R control is provided in Ref. 6.

III. RUN-TO-RUN CONTROL

At the level of R2R control, we are integrating into the GCC multiple optimization and control software modules, a linear-approximation SPC controller and an artificial neural net optimizer and controller. We have designed and implemented a multiple-control branch-selection algorithm which allows the GCC to employ these multiple control algorithms in a complementary fashion.

A. Multiple control branches and branch selection algorithm

In recent years, various algorithms have been developed for R2R optimization and control, e.g., Ultramax¹⁶ (Ultramax is a trademark of the Ultramax Corporation, Cincinnati, Ohio), and a linear approximation-based control algorithm.¹⁷ There does not appear to exist any single algorithm that may be used throughout the entire range of R2R control and optimization.⁴ Fortunately, the ranges of applicability of these algorithms can be roughly expressed in terms of the process being near or far from its optimum. The process engineer in the manufacturing facility in general will have many “rules of thumb” for process control. It is these sources of knowledge, even if vague or imprecise, which we must exploit in order to carry out the selection of control algorithms.⁵

An effective controller design should, therefore, provide a mechanism for utilizing rules of all these types in making the control decision. Toward this end, we have used fuzzy set theory to develop an algorithm selection module (ASM) and

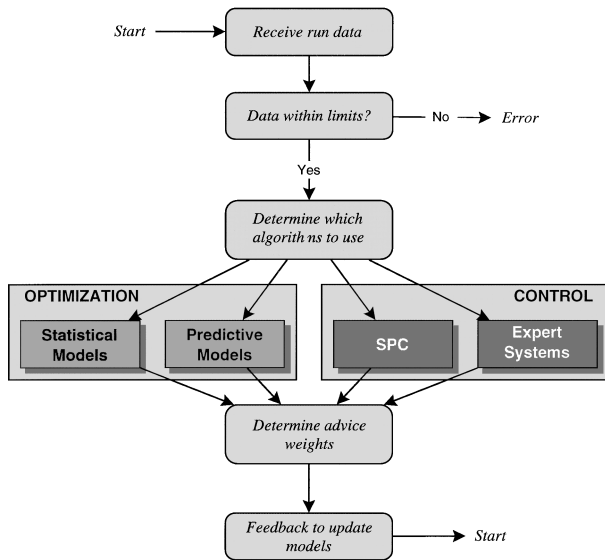


FIG. 5. Selection of optimization and control algorithms.

have incorporated it into the GCC system. The ASM utilizes a fuzzy rule-based decision-making process to select from multiple algorithms. The fuzzy rules express knowledge about the suitability of the different control algorithms based on the process conditions. For a detailed description of the algorithm selection module and its integration with the GCC, see Refs. 5 and 18. The basic functionality of the ASM is depicted in Fig. 5.

B. The MIT run-by-run controller

The MIT run-by-run (MIT RbR) controller is the main control algorithm that is in use in the current GCC implementation. It provides sequential control advices through a multivariate adaptive model-based approach. At the heart of the controller is a linear model of the system; outputs are expressed as a linear combination of the control inputs (with constant gain), plus a constant offset term:

$$y = Ax + c. \quad (1)$$

Based on a R2R measurement, the controller adapts the offset terms c , while the gain matrix A remains fixed. The model is updated recursively by an exponentially weighted moving average (EWMA) update of the offset term based on the error between model prediction and measurement:

$$c_t = \alpha(y_t - Ax_t) + (1 - \alpha)c_{t-1}. \quad (2)$$

Here c_{t-1} is the offset term used on the previous run. The selection of the weight coefficient α is based on consideration of noise, drifts, shifts, and model error. An especially important consideration is the selection of these coefficients in a closed loop control scenario. In general, the system will compensate for drifts or shifts in the system more rapidly and closely with larger weights. On the other hand, larger weights will also cause the controller to react more aggressively to noise in the outputs, which can increase the overall noise through the system. Methods for optimal selection of

weight coefficients are discussed in Ref. 15. Current research is investigating methods for dynamic selection of the weight based on observation of the system itself (e.g., to decrease the weight if one observes that the noise in the system increases over time).

Once the model has been updated, the model is inverted to solve for the plant controls that best achieve the desired targets for the output parameters of interest. Around this basic structure of R2R control, a number of important practical issues are addressed by the controller algorithms.^{9,15} First, appropriate normalization of inputs and outputs is performed, so that the relative magnitude of individual parameters in the multivariate solution does not bias the solution. In conjunction with this normalization, user supplied weights may be provided (1) to weight the relative importance of achieving each of several multiple targets, and (2) to weight the relative willingness to move or change control inputs. Second, simple weighted least squares solutions are used to either achieve the closest possible fit to the desired set of targets (if these cannot be satisfied exactly due to bound constraints on the inputs) or to achieve the desired targets with the smallest weighted change to the current recipe. Finally, the controller explicitly handles quantization of input parameters (e.g., speed only capable being adjusted in integer units) and finds an acceptable near-optimal process recipe subject to quantization. A particularly important opportunity for improvement to the MIT RbR controller is the utilization of *in situ* machine sensor information to guide the selection of process recipe changes, or to motivate and guide adaptation of model gain terms. A description of the MIT RbR controller software implementation can be found in Ref. 19 and, in its incorporation as a module of the GCC, in Ref. 6.

C. An artificial neural network branch

The MIT RbR controller is successful in detecting small shifts and drifts and noise. But it performs poorly for a process with large shifts and drifts. Artificial neural network (ANN) technology is a promising tool for not only control, but also optimization, of manufacturing processes. ANNs are massively parallel and dynamic in nature and are very efficient in modeling nonlinear processes. They can handle many inputs simultaneously. An ANN basically consists of three layers: an input layer, an output layer, and one or more hidden layers linked to each other through connections. ANNs are predicted to perform better for nonlinear processes and also to successfully detect large shifts and drifts.

All the advantages of ANN technology can be incorporated into the GCC by using an additional ANN branch alongside the MIT RbR controller. The fuzzy branch determination technique explained in Sec. III A provides the GCC with the flexibility to employ this set of two control algorithms in a complementary fashion.

We have used a commercial ANN based software tool to develop our CMP process model and extract the controller model from it. Simulation models of complex processes can be quickly and easily built directly from historical (DOE)

data. Once a model has been created, the software tool can predict the future behavior of the process, validate sensor information, and control the process.

IV. REAL-TIME MONITORING

As mentioned above, there are very few sensor technologies that are capable of directly monitoring CMP process variables in real time. However many CMP equipment inputs may be monitored in real time. We have designed a real-time CMP monitoring system that utilizes this information to provide real-time process and equipment diagnostics; the system utilizes historical data to derive profiles or footprints of the system, and then compares real-time data against these footprints to provide system diagnostics. The control structure is illustrated in Fig. 3.

A. Real-time CMP data

In a recent series of articles^{20–23} methods were developed to monitor process behavior in batch chemical processes using real-time process data. The methods presented in these papers compare parameter data from current batches to typical or average profiles based on historical production. Multivariate statistics model the correlation structure and reduce the high-dimensional and highly correlated parameter data to a few latent variables whose values are monitored on a multivariate control chart. Although the chemical processes discussed in Refs. 20–23 differ widely in function and purpose from CMP, the data structure and objectives of CMP match well enough to benefit from these ideas.

Developing a real-time process monitor for CMP in the R2R control context adds another layer of complexity beyond that seen in the chemical processes of Refs. 20–23. In R2R control, process inputs change frequently, directly influencing both parameter profiles and product measurements. As a result, the profiles with which current data should be compared depends on the recipe and other initial conditions. Behavior monitoring becomes a matter of comparing incoming parameter readings against expected profiles rather than average profiles. The following sections describe the data and situation more rigorously, present functions that should be performed by a real-time monitor for CMP, and discuss the need to develop a method beyond those found in the literature.

B. Data used in real-time monitoring

Each run provides a single observation on each variable in an input and initial condition matrix \mathbf{X} , a real-time parameter array \mathbf{Y} , and a product-related data matrix \mathbf{Z} .

The \mathbf{X} matrix plots run versus input variables (down pressure, carrier velocity) and initial condition variables (environmental data, summary data from the previous runs parameters, or initial readings on parameters at the beginning of the current run).

The \mathbf{Y} array has three dimensions: run number, real-time sample number, and real-time parameter value. That is, for each run number, it shows the real-time parameter value for

each real-time sample within the run. Again, a row in the \mathbf{Y} array contains all the parameter data for a single run.

The \mathbf{Z} matrix plots run number versus product-related data (removal rate, post mean thickness, and removal non-uniformity). Although their use will not be emphasized here, these variables can be modeled and predicted from the real-time parameters and initial condition variables as suggested in Refs. 20–23.

In CMP, a relationship of the following form exists between a row vector in the \mathbf{X} matrix and a row matrix in the \mathbf{Y} array.

$$y_i = F(x_i) + e_i. \quad (3)$$

In this expression, y_i represents a matrix of parameter data from one run (a single row of the \mathbf{Y} array); x_i represents a vector of input and initial condition variables from one run (a single row of the \mathbf{X} matrix); F is a matrix-valued function; and e_i represents a matrix of random errors. The ability of a real-time monitor to generate expected profiles against which to compare incoming data relies on estimating the function $F(x_i)$ based on historical runs.

C. Functions to be performed by a real-time monitor

Considering the R2R control context, a real-time monitor for CMP must perform a series of functions.

- (1) The monitor must create a model based on historical production data that can predict parameter profiles from values in the x_i vector. This model must be applicable to recipes not seen in the historical data set and versatile enough to handle a variety of linear and nonlinear relationships between elements of x_i and elements of y_i .
- (2) The monitor must reduce the dimensionality of the parameter data to just a few summary variables (components) that can be monitored on a R2R multivariate control chart.
- (3) A trade-off must be made between simplicity and sensitivity. Too much dimension reduction may ignore information and decrease sensitivity. On the other hand, not enough dimension reduction may not adequately simplify monitoring charts.
- (4) The monitor must be capable of tracing abnormal or out-of-control signals in the multi-variate control chart back to the features in the real-time parameter profiles that most likely caused them. This provides a helpful diagnostic tool, without which a real-time monitor would be much less useful. This tool should be able to determine which parameter at what time interval contributed the most to the signal.
- (5) The monitor should be able to predict the values of the summary variables each time a new real-time parameter observation is received. Because polishing times tend to be short in CMP (on the order of minutes), this feature is not as critical to CMP as it may be to some longer chemical processes. However, it will still serve as an early warning tool and provide an additional way of pinpointing the time in the process at which abnormalities occur.

- (6) The monitor should create a model based on historical production data that can predict product-related data from values in the x_i vector and y_i matrix. Although this is a valuable use of real-time data, the methods suggested in Ref. 21 effectively perform this function (see below) and need no further development for application to CMP.

D. Need for CMP process monitor development

To meet these criteria, and to be able to be applied to CMP and R2R control, the methods presented in Refs. 20–23 must be extended or changed. The existing approaches perform functions 2–4, but the processes for which they were developed did not merit the development of function 1. A more detailed description of a method described in Ref. 21 clarifies the need for additional consideration of the profile modeling function.

Three separate methods for monitoring batch processes are suggested by Ref. 21, the method of choice depending on the types of data available. One of these methods makes use of input, initial condition, and product measurement data and is based on multiblock multiway PLS (partial least squares or projection to latent structures).

In terms of the notation presented here, PLS models the product-related variables in the \mathbf{Z} matrix in terms of the \mathbf{X} matrix and \mathbf{Y} array (note the different notation in Refs. 20–23). The functional relationship can be expressed as

$$z_i = G(x_i, y_i) + f_i. \quad (4)$$

In this expression, z_i represents a vector of product data (a single row of the \mathbf{Z} matrix); \mathbf{G} is a vector-valued function; x_i represents a vector of input and initial condition data (a single row of the \mathbf{X} matrix); y_i represents a matrix of parameter data (a single row of the \mathbf{Y} array); and f_i represents a vector of random errors.

Multiblock multiway PLS estimates the function G , retaining the first few latent variables related to the \mathbf{Y} array for behavior monitoring. Incoming parameter data is then compared to average profiles based on the information retained in these latent variables. Deviations from these average profiles represent random error, giving rise to the following relationship between average profiles and incoming profiles:

$$y_i = \bar{y} + e_{i*}. \quad (5)$$

In this expression, y_i again refers to a matrix of parameter data; \bar{y} represents a matrix containing average parameter profiles; and e_{i*} represents a matrix of random errors.

Comparison between Eqs. (3) and (5) highlights the difference between methods applicable to CMP process monitoring (under R2R control) and those that work for the batch processes discussed in Refs. 20–23. Equation (3) suggests comparison of current parameter profiles to expected profiles, while Eq. (5) suggests comparison with average profiles. The need to address the model in Eq. (5) arises only because of R2R control. For wafers polished at a fixed recipe and constant initial conditions, Eq. (3) simplifies to Eq. (5), and the methods presented in Refs. 20–23 can be applied.

However, ignoring the relationship in Eq. (3) when it exists (as in CMP) dumps all this explainable variation into error, inflating the variance of latent variables and the multivariate control limits.

V. CONCLUSION

In this article we have presented a design for a multilevel control system that is applicable to semiconductor manufacturing. The design incorporates real-time and run-to-run (R2R) control elements and provides a comprehensive, robust and dynamic control solution. Elements of this design have been developed and applied to the control of a CMP process. Specifically, a R2R control system has been developed and deployed that provides R2R multivariate control of material removed and uniformity by analyzing and improving the process recipe. A real-time monitoring and diagnostics system has also been developed and deployed that analyzes equipment inputs (including R2R recipe variations) and input histories and deduces exception conditions from input value combinations and input footprints.

In the near future we see the potential for research and development at each of the three levels of control applied to the CMP process. At the real-time level, open-loop diagnostics will be further integrated with the R2R controller, with the long term goal being the development of sensing technology and corresponding control algorithms to achieve closed-loop real-time control. At the R2R level, efforts are continuing on developing additional control branches and integrating them into the multibranch control scheme so as to enhance the robustness of the controller. Finally, at the intercell level efforts are focused on implementing and testing enabling mechanisms that can carry out intercell feedback and feed-forward control.³

ACKNOWLEDGMENTS

The authors wish to express their appreciation to SEMATECH, Compugenesis, Inc., and the University of Michigan Center for Display Technology & Manufacturing who have funded this research.

¹B. Rashap, M. Elta, H. Etemad, J. Fournier, J. Freudenberg, M. Giles, J. Grizzle, P. Kabamba, P. Khargonekar, S. Lafortune, J. Moyne, D. Teneketzis, and F. Terry, Jr., IEEE Trans. Semicond. Manuf. (1995).

²J. Moyne, R. Telfeyan, A. Hurwitz, and J. Taylor, *Proceedings of the Sixth Annual SEMI/IEEE ASMC, Boston, November 1995* (SEMI, Mountain View, CA, 1995).

³N. Chaudhry, J. Moyne, and E. Rundensteiner, 42nd National Symposium of the American Vacuum Society, Minneapolis, MN, October 1995.

⁴J. Moyne, H. Etemad, and M. Elta, Proc. SPIE 2091, 379 (1994).

⁵J. Moyne, N. Chaudhry, and R. Telfeyan, J. Vac. Sci. Technol. A 13, 1787 (1995).

⁶R. Telfeyan, J. Moyne, A. Hurwitz, and J. Taylor, *187th Meeting of The Electrochemical Society, May 1995* (The Electrochemical Society, Pennington, NJ, 1995).

⁷R. Jairath, J. Farkas, C. Huang, M. Stell, and S. Tzeng, Solid State Technol. 71 (1994).

⁸P. Singer, Semicond. Int. 48 (1994).

⁹D. Boning, A. Hurwitz, J. Moyne, W. Moyne, S. Shellman, T. Smith, J. Taylor, and R. Telfeyan, in IEEE Proceedings of the 17th International Electronic Manufacturing Technology Symposium, October 1995 (unpublished).

- ¹⁰M. Martinez, *Solid State Technol.* **26** (1994).
- ¹¹A. Hu, X. Zhang, E. Sachs, and P. Renteln, in *IEEE Proceedings of the 15th International Electronic Manufacturing Technology Symposium*, October 1993 (unpublished).
- ¹²A. Hu, H. Du, S. Wong, P. Renteln, and E. Sachs, in *IEEE Proceedings of the 16th International Electronic Manufacturing Technology Symposium*, October 1994 (unpublished).
- ¹³J. Moyne and L. McAfee, Jr., *IEEE Trans. Semi. Manuf.* **5**, 77 (1992).
- ¹⁴J. Moyne, U.S. Patent No. 5,469,361 (November 1995).
- ¹⁵D. Boning, W. Moyne, T. Smith, J. Moyne, and A. Hurwitz, in Ref. 2.
- ¹⁶C. W. Moreno, *ISA Proc.* (June 1986).
- ¹⁷E. Sachs, A. Hu, and A. Ingolfsson, *IEEE Trans. Semicond. Manuf.* **81**, 26 (1995).
- ¹⁸N. Chaudhry, J. Moyne and E. Rundensteiner, University of Michigan, CSE Division, Technical Report, CSE-TR-224-94, November 1994.
- ¹⁹W. Moyne, SM thesis, MIT EECS, May 1995.
- ²⁰P. Nomikos and J. F. MacGregor, *AIChE J.* **40**, 1361 (1994).
- ²¹T. Kourti, P. Nomikos, and J. F. MacGregor, *J. Process Contr.* (1995).
- ²²J. F. MacGregor, and T. Kourti, *Control Eng. Practice* **3**, 403 (1995).
- ²³P. Nomikos and J. F. MacGregor, *Technometrics* **37**, 41 (1995).
- ²⁴S. W. Butler and J. Stefani, *Proceedings of the American Control Conference*, June 1993 (unpublished).
- ²⁵U. Dayal, E. Hanson and J. Widom, in *Modern Database Systems: The Object Model, Interoperability and Beyond*, edited by W. Kim (ACM, 1995).
- ²⁶N. Gehani and H. Jagadish, *Proceedings of the 17th International Conference on Very Large Data Bases, September 1991* (Morgan Kaufman, Los Altos, CA, 1991), pp. 327–336.