

Microeconomics of Advanced Process Window Control

Kevin M. Monahan, Xuemei Chen, Georges Falessi, Craig Garvin, Matt Hankinson, Amir Lev, Ady Levy, and Mike Slessor, KLA-Tencor Corporation

Fundamentally, advanced process control enables accelerated design-rule reduction, but simple microeconomic models that directly link the effects of advanced process control to profitability are rare or non-existent. In this work, we derive these links using a simplified model for the rate of profit generated by the semiconductor manufacturing process. We use it to explain why and how microprocessor manufacturers strive to avoid commoditization by producing only the number of dies required to satisfy the time-varying demand in each performance segment. This strategy is realized using the tactic known as speed binning—the deliberate creation of an unnatural distribution of microprocessor performance that varies according to market demand. We show that the ability of APC to achieve these economic objectives may be limited by variability in the larger manufacturing context, including measurement delays and process window variation.

Introduction

For the 90 nm technology node and beyond, we foresee a trend toward extended, generalized, factory-wide, model-based, advanced process control. The traditional microeconomic benefits of advanced process control (APC) include reduced process variation, accelerated shrinks, elimination of send-ahead wafers, fewer monitor wafers, shorter response times, reduced scrap, better tool matching, improved overall equipment effectiveness, faster yield ramps, lower parametric yield loss, better device performance, easier process transfer from pilot line to factory, and dramatically lower labor costs. These benefits, especially in the case of microprocessors, will continue to translate into higher gross margins for semiconductor products since they simultaneously lower manufacturing cost and increase average selling price.

In addition, we expect traditional APC architectures to be augmented at the 90 nm node to support the precursors of terahertz

transistor technology. At the 90 nm node, physical gate lengths are expected to reach 50 nm and decrease rapidly to 30, 20, and 15 nm in subsequent generations. Some microprocessor manufacturers may fabricate 50 nm gates at the 130 nm node. Such design rules will generate more rigorous process control requirements for rapidly shrinking, overlapping process windows. Ultimately APC architectures may be called upon to provide comprehensive control of transistor formation. To address the new requirements, APC architectures must incorporate the following:

- Infrastructure extensions that enable seamless root-cause determination and correction when the APC history analysis generates an action request. Examples are model-based process window monitors (PWM), advanced equipment control (AEC), fault detection/classification (FDC), integrated diagnostic monitors, and internet-based support.
- Factory-wide open APC frameworks that allow single-point MES integration and coordination of multiple APC applications running simultaneously in several process areas. Consider an L-effective controller that feeds forward CMP thickness data to correct CDs in litho, litho CD data to correct CDs in

etch, and etch CD data to correct L-effective during spacer formation or ion implantation.

- Business rule management that supports high-mix production and multiple APC strategies. The simple APC business rules developed initially in low-mix microprocessor and DRAM factories must be extended to address high-mix manufacturing contexts in foundries and to support multiple strategies that require different levels of metrology and process tool dedication for low and high ASP products.
- Multi-threaded APC applications that respond to more than one input and generate more than one output. Consider a litho area that exhibits systematic cell-to-cell, lot-to-lot, wafer-to-wafer, field-to-field, and die-to-die CD variation that might be corrected by feeding back dose offsets for each cell, lot, wafer, field, and die, respectively.
- High-bandwidth support for stand-alone, clustered, and integrated metrology tools. Stand-alone and 300 mm cluster metrology tools that generate data on a lot-to-lot basis will be joined by integrated metrology tools that generate data on a wafer-to-wafer basis. In some cases, bandwidth requirements may force on-tool distillation of correctables required for lot-based recipe modification.
- Physical model-based process control that can predict electrical characteristics. APC applications will evolve from simple offset corrections (CD control) and empirical response surfaces (overlay control) to physical models that predict electrical characteristics such as effective gate length (L-effective) and equivalent oxide thickness (EOT).
- Control of transistor formation using models that predict device performance. Model-based process control will evolve from simple models to more sophisticated physical models using measurements from multiple tools to predict device performance characteristics such as voltage threshold (V-threshold), off current (I-off), and drive current (I-drive).

Ultimately, driven by compelling economic forces, today's APC, AEC, FDC, and PWM will evolve into a new species of process control that, for the purposes of this work, we refer to as Process Window Control (PWC). PWC opens a vast array of possibilities. Gate stack formation will take on new meaning as we force the overlap of multivariate process windows from deposition

Drive current is one of the most important performance characteristics in transistor design, particularly for microprocessors. It controls both the speed and fan-out capability of the device. Drive current¹ at saturation can be expressed as

$$I_{ds} \cong K \cdot \left(\frac{W}{2}\right) \cdot \left(\frac{\mu}{L}\right) \cdot \left(\frac{\epsilon}{T}\right) \cdot (V_{gs} - V_{th})^2$$

K is a function of multiple parametric errors, W is the gate width, L is the effective gate length, μ is the carrier mobility, T is the thickness of the gate dielectric, ϵ is the permittivity of the gate dielectric, V_{gs} is the gate voltage, and V_{th} is the threshold voltage at which switching begins to occur. In theory, a small error in any one parameter could affect performance substantially; but the true threat to drive current is in the complex interaction of these parameters in multiple process windows. The most thoroughly characterized parameter is L-effective, the optimization of which is constrained by the process windows for poly thickness, litho CD, etch CD, spacer profile, and implant dose and anneal conditions.

to alignment, exposure, etch, spacer formation, and ion implantation. Intra-field proximity effects in lithography may be compensated in the etch module by varying power, pressure, flow, and chemistry using feedback from embedded measurement tools. Cross-wafer non-uniformity induced in the etch module may be corrected in lithography using field-to-field and intra-field dose control. Factors that were once ignored in APC, such as lithographic defocus or spacer sidewall angle, may become an integral part of process correction as scatterometry proliferates. APC, AEC, FDC, and PWM will continue, but will function more robustly inside the context of process window control.

Microeconomic model

Fundamentally, advanced process control enables accelerated design-rule reduction, but simple microeconomic models that directly link the effects of advanced process control to profitability are rare or non-existent. In this work, we derive these links using a simplified model² for the rate of profit generated by the semiconductor manufacturing process. We have modified it to reflect the deliberate performance-binning used in

microprocessor factories. Thus, the rate of profit becomes

$$\dot{P} = -(\dot{W}_0 - \dot{W}) \cdot C_0 + \sum_j \dot{W}_j (d_j Y_j y_j p_j - C_j) \quad (2-1)$$

where \dot{W}_0 is the wafer output at full capacity, \dot{W} is the total number of wafer starts per week, C is the manufacturing cost per wafer, Y is the metrology-limited yield entitlement, y is the device yield expressed as a fraction of entitlement, d is the number of dies per wafer, p is the average selling price per die, i is the product index, and j is the binning index. Our simplified business model represents the gross rate of profit attributable to a factory. It does not include variable costs associated with packaging, marketing, or sales of the product; but the first term does take into account the "burn rate" associated with unused capacity. Unused capacity is modeled as a null product with zero average selling price and a manufacturing cost per wafer that is reduced by the incremental cost of maintenance and materials required for actual production (C_0).

Factory Burn Rate

Burn rate is due to fixed costs associated with capital investment, operation, and depreciation of a factory. For 90 nm manufacturing and beyond, huge capital investments may be required for 193 and 157 nm lithography tools, 300 mm wafer handling and automation, silicon-on-insulator substrate technology, atomic layer deposition for high-k gate dielectrics, SiGe epitaxy for raised source and drain, copper and low-k dielectric interconnect, and factory-wide metrology integration with comprehensive parametric APC and process window management systems. In the above model, 90 nm factories would be losing money before processing a single wafer.

The traditional strategy for minimizing the effect of fixed costs is to reduce manufacturing cycle time and operate near maximum capacity. Several tactics are employed to support this strategy. In a supply-limited environment, we fill the factory with high margin products to maximize profits. In a demand-limited environment, such as the current one, we load the factory with some lower margin products to offset the burn rate. If factories remain below capacity, we consolidate, shut down, or sell at depreciated value. The coming transition to 300 mm wafers will bring a 30% reduction in normalized die cost and could trigger these events for existing 200 mm factories.

Normalized Die Cost

Normalized die cost falls off sharply with larger wafer size and smaller design rules. We can estimate it by taking into account the manufacturing cost per wafer, wafer size, yield, and design-rule generation so that

$$c_j = \frac{C_j}{Y_j y_j d_j} = \frac{1}{d_0} \left[\frac{C_j}{D^2} \right] \left[\frac{S^2}{Y_j y_j} \right] \quad (2-2)$$

Here, S is the shrink ratio (e.g., 130 nm/90 nm), D is the wafer diameter ratio (e.g., 300 mm/200 mm), and d_0 is the initial number of dies per wafer.

Increasing wafer size results in lower die cost, but the benefit is partially offset by the opportunity cost of 300 mm facilities, equipment, and materials. Accelerated shrinks also result in lower die cost, but the benefit again is partially offset by the opportunity cost of advanced process tools, introduction of new materials, and potential reductions in yield.

The metrology-driven shrink is by far the most cost-effective alternative to the above tactics and can significantly extend the useful life of current process tools and factories. This solution requires relatively inexpensive APC and PWM applications supported by a factory-wide control framework. The investment needed for metrology-driven die-cost reduction is typically an order of magnitude less than that required for new process equipment and two orders of magnitude less than the cost of a 300 mm factory. Our microeconomic model, discussed in the next section, supports the use of parametric APC and PWM not only for improving device performance and die density but also for extending the life of existing 200 mm factories.

Revenue Potential

For the remainder of this work, we assume that consolidation tactics have reduced the burden of unused capacity to negligible levels. In this case, profitability can be expressed very simply as

$$\dot{P} = \sum_j \dot{W}_j d_j Y_j y_j (p_j - c_j) \quad (2-3)$$

The first term in the equation above represents the potential revenue from sale of products. In some cases, the semiconductor industry experiences significant cross-elasticity between types of products. For example, recent overproduction and subsequent commoditization in the DRAM business helped to maintain margins for microprocessors, since both share revenue generated by the personal computer industry. In the equipment

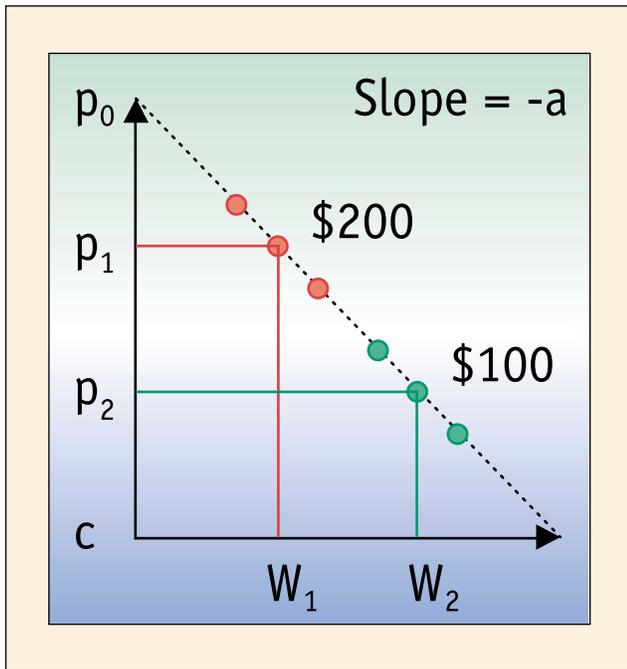


Figure 1. Elasticity of microprocessor pricing and wafer output. Profit is the area inside the boxes for the \$100 and \$200 parts. Drawing is exemplary and not to scale.

industry, the recent glut of depreciating process equipment created opportunities for yield management and process control companies that know how to extend the useful life of equipment and factories.

Next, we consider the microprocessor business, where price and output are also elastic in demand-limited markets. For the purposes of this discussion, we define a demand-limited market for a microprocessor product (*i*) in performance segment (*j*), as a market in which average selling price declines linearly with the rate of production (Figure 1) so that

$$p_y = -a_y \cdot \dot{W}_y + p_{0y} \quad 2-4$$

For a single product/performance segment, combining Equations 2-3 and 2-4 yields

$$\dot{P} = dYy(-a \cdot \dot{W}^2 + (p_0 - c) \cdot \dot{W}) \quad 2-5$$

Setting the derivative equal to zero and solving for \dot{W} , we have

$$\dot{W}_{max} = \frac{(p_0 - c)}{2 \cdot a} \quad 2-6$$

Evaluating Equation 2-5 for maximum profitability, we get

$$\dot{P}_{max} = \frac{dYy(p_0 - c)^2}{4 \cdot a} \quad 2-7$$

Clearly, microprocessor manufacturers can resist commoditization by producing only the number of dies required to satisfy the time-varying demand in each performance segment. This strategy is realized using the tactic known as speed binning, the deliberate creation of an unnatural distribution of microprocessor performance that varies according to market demand. Given constant elasticity, we can make the following observations:

- Narrow bands of profitability drive micro-segmentation of the market. For example, a hypothetical company with 80% market share, competing with a company that had 20% market share, would naturally produce both low-end and high-end microprocessors at multiple price/performance levels, denying any exclusive niche to its competitor.
- The microprocessor company with 80% market share would control global pricing. It could increase output, drop prices, and sacrifice near-term margins, depriving its competitor of the profit and capital required to invest in lower die-cost and in 300 mm factories. Economic forces would compel the smaller competitor to partner with other companies that can manufacture with lower die-cost or have more experience in 300 mm manufacturing.
- Maximum profitability is directly proportional to yield, density, and the square of the initial gross margin (Figure 2). Given the critical impact of initial gross margin, the first company to market enjoys substantial profitability due to market elasticity. Parametric process window control may be the key means to getting to market first. Gate CD control is an imminent example. In the next section, assuming constant yield, we establish the links between gate CD control, die-density, average selling price, and factory life extension.

Links to Gate CD Control

The variability of gate CD can be characterized using generalized ANOVA³, APC simulation, and process-window RSA to separate correctable systematic cell, lot, wafer, field, die, and site variation from lot-to-lot temporal variation. Benefits accruing from improved CD control are contingent upon a number of assumptions, including a price-performance premium, unsaturated

$$p \cong -a \cdot \dot{W} + p_0$$

$$\dot{P} = Yyd(-a \cdot \dot{W}^2 + (p_0 - c) \cdot \dot{W})$$

$$\dot{W}_{max} = \frac{(p_0 - c)}{2 \cdot a}$$

$$\dot{P}_{max} = \frac{Y \cdot y \cdot d \cdot (p_0 - c)^2}{4 \cdot a}$$

Figure 2. Dependence of maximum profitability on the yield entitlement, yield fraction, density, and the square of the initial gross margin.

market demand, scalability of interconnect speed, dimensional scalability of other critical layers, and the relatively small incremental cost of PWM and APC. The following calculations are meant to establish an entitlement for maximum incremental return, against which the actual ROI may be compared:

The entitlement for price increase due to improvement in L-effective control from σ_0 to σ is given approximately by

$$p - p_0 \cong \alpha \cdot \beta \cdot L_0 \cdot \left(1 - \frac{\sigma}{\sigma_0}\right) \tag{2-8}$$

α and β are, respectively, the incremental performance capability (~10 MHz/nm) and the price leverage (~0.2 \$/MHz). Analogously, the entitlement for increase in die density is approximately

$$d - d_0 \cong d_0 \cdot \left(\frac{\sigma_0^2}{\sigma^2} - 1\right) \tag{2-9}$$

The entitlement for factory life extension (t) in weeks can be estimated as

$$t - t_0 \cong \left(\frac{\tau}{\ln \chi}\right) \cdot \ln\left(\frac{\sigma}{\sigma_0}\right) \tag{2-10}$$

τ and χ are, respectively, the duration of the technology generation (~10⁴ weeks) and the scaling factor for each generation (~0.7).

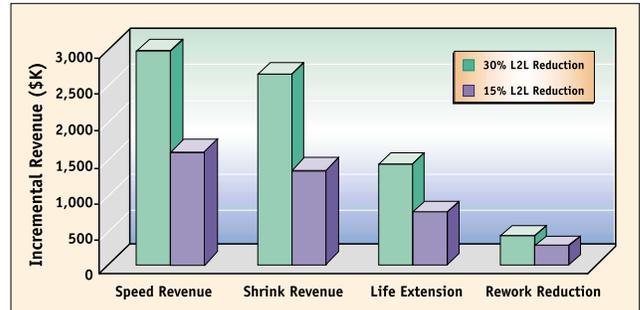


Figure 3. Pareto of predicted incremental revenue due to microprocessor speed improvement, increase of the density, factory life extension, and rework reduction after lot-to-lot (L2L) variation was decreased by 15% and 30%. Since the total variation model includes lot, wafer, field, and site components, the actual shrink ratios enabled were about 0.95 and 0.90, respectively. Actual incremental revenue is typically only a fraction of the entitlement, creating an opportunity for advanced process window control.

The Pareto chart in Figure 3 shows examples of incremental return due to microprocessor speed improvement, increase of die density, factory life extension, and rework reduction assuming APC produces a 15% to 30% decrease in lot-to-lot CD variation. Since the model for total CD variation includes lot, wafer, field, and site components, the actual shrink ratios are only about 0.95 and 0.90, respectively. Initial production is assumed to be 1000 wafers per week, with 150 dies per wafer at \$100 per die. These entitlements assume that gate CD is the critical speed and shrink bottleneck. We are not modeling the interactive incremental return due simultaneous increase of speed and density.

Actual incremental return on investment is typically a small fraction of entitlement. This return can be increased dramatically by controlling the other components of CD variation at the site, field, and wafer levels (Figures 4 and 5). The extension of APC to more critical steps and layers could remove speed and shrink bottlenecks in other parts of the process. Finally, the effectiveness of the APC algorithm itself could be improved by making changes in the manufacturing context that minimize measurement delay and the contribution of unexplained process variation, as we discuss in the next section.

Controlling gate CD

The two generic means of controlling gate CD in state-of-the-art factories are advanced process control and process window management. APC seeks to adjust process parameters inside process windows. PWM seeks to align the process windows, primarily by

140 nm GL		Litho		Etch	
3-Sigma	SYS	RDM	SYS	RDM	
LOT	3.50	5.00			10.00
WAFER	0.70	2.00	1.00		3.50
FIELD	2.50	2.50	3.50		3.50
SITE	3.50	7.00	3.50		7.00
TOTAL	11				14

50 nm GL		Litho		Etch	
3-Sigma	SYS	RDM	SYS	RDM	
LOT	1.30	1.80			3.50
WAFER	0.25	0.70	0.35		1.30
FIELD	1.00	1.00	1.30		1.30
SITE	1.30	2.50	1.30		2.50
TOTAL	4				5

Figure 4. Estimates of generalized ANOVA for 140 nm gates (top) and scaled estimates for 50 nm gates (bottom). Stabilizing and correcting lower-level nested systematic components constitutes an opportunity for PWM and APC. The apparently random lot variation in etch is due to lack of chamber tracking and lack of litho-etch tool dedication.

enabling recurrent equipment matching. Both APC and PWM are evolutionary next-steps in the “copy exactly” philosophy of semiconductor manufacturing.

Advanced Process Control - APC

Advanced process control is a closed-loop control system in which process correctables are fed back to the preceding module or fed forward to the next module to

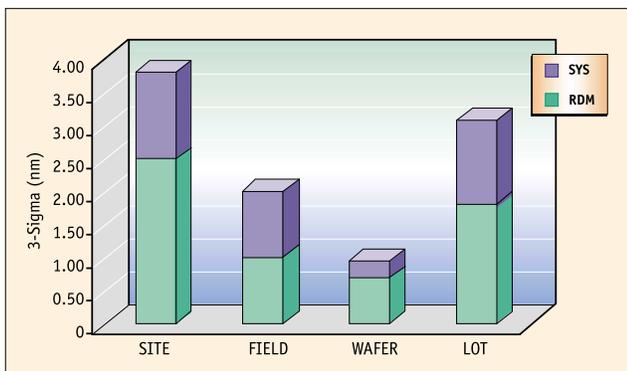


Figure 5. Graphical representation of the 50 nm lithography data. The largest single component of variation is associated with intra-field sites, suggesting that the order of economic value for APC is site, lot, field, and then wafer. The composite estimates support the idea of intra-field dose control and raise some questions about the value of wafer-to-wafer adjustments. Note that much of the “random variation” is due to multiple systematic variations confounded by sparse sample plans.

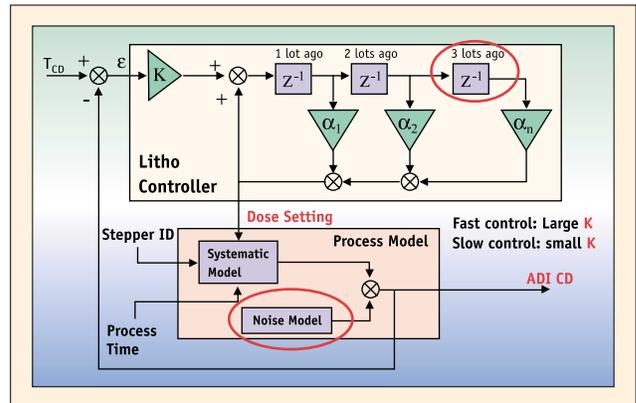


Figure 6. APC requires a control model (e.g., litho controller) and a process model (e.g., CD versus dose). The control model can be more effective if CD measurement data is both accurate (e.g., error less than 1 nm) and current (e.g., delivered in less than 1 hour). The process model is more useful when the systematic model variation is large relative to the process noise contribution. Noise may derive from process window variation that is outside the context of the process model and therefore not correctable by APC.

reduce parametric variation on product wafers. The success of APC is based on two assumptions: the existence of an effective control model and the existence of a valid process model (Figure 6). For the control model to be effective, measurement data must be supplied in a timely manner so that high control gains may be used without the danger of overcorrecting (Figure 7). For the process model to be valid, it must account for a useable fraction of the observed variation and leave a minimum of unexplained variation in the form of residuals. Unexplained variation can have both systematic “lack-of-fit” and random “repeatability” components. Quantum effects notwithstanding, if we considered enough factors, nearly all the parametric variation in a

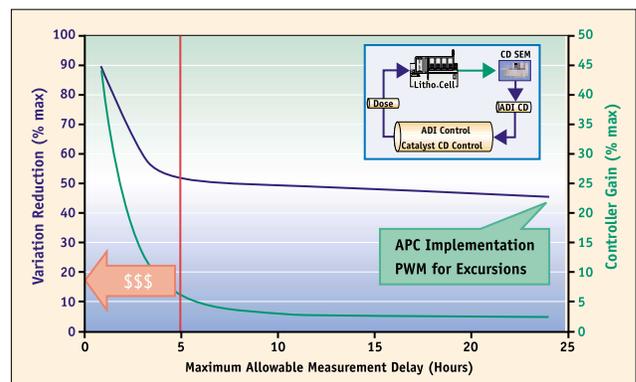


Figure 7. Estimate of lot APC effectiveness versus time delay. About 90% of value is achieved for delays less than 1 hour.

semiconductor factory would be classified as systematic. The postulate of APC, however, is that control can be accomplished by adjustment of a few critical parameters inside otherwise stable, multivariate process windows.

The effectiveness of APC can be compromised by delays in either processing or measurement. Figure 7 shows simulated reduction of lot-level variation as a function of maximum allowable measurement delay. These control curves are similar regardless of the origin of the factory data used for simulation. In some cases, the effectiveness of CD APC can double by taking delays down to about 1 hour from typical values in the range of 5-10 hours. This behavior is also observed for APC in overlay applications. For lot-level control, measurement delays can be reduced by increasing stand-alone metrology capacity to reduce queuing time, by clustering metrology to reduce delays in lot transfer, and by integrating metrology into the process tools. Given the relatively low wafer-to-wafer variation shown in Figure 5, delay time reduction, faster excursion response, cumulative statistical analysis, and direct input-output correlation may be the real opportunity for integrated metrology and compensation for incremental cost on every tool.

Alternatively, in the case of stand-alone CD and overlay metrology, we have found that dramatic improvements in APC performance can be achieved by optimizing capacity and sample plans⁴. For example, increasing an overlay sample plan from 4 arbitrarily selected fields to 10 statistically optimized fields may nearly double the potential effectiveness of OL APC. In addition, the larger sample plans collect more cross-wafer and intra-field data and usually provide more statistical leverage for correction of lower-level, nested spatial variation. The requisite metrology capacity pays for itself by accelerating shrinks and enhancing profitability.

Process Window Management – PWM

Process window management is an open-loop control system in which multiple parameters are monitored for the purpose of detecting excursions and matching process equipment using monitor wafers. PWM is a powerful tool for dealing with unexplained intra-module or inter-module process variation that can reduce performance of APC, particularly, since performance gaps commonly arise from collapse and/or displacement of process windows. As an example, Figure 8 shows litho and etch CD windows as a function of litho defocus. The optimal litho and etch windows are displaced by 0.2 micrometers, creating a 30% etch-CD risk and the

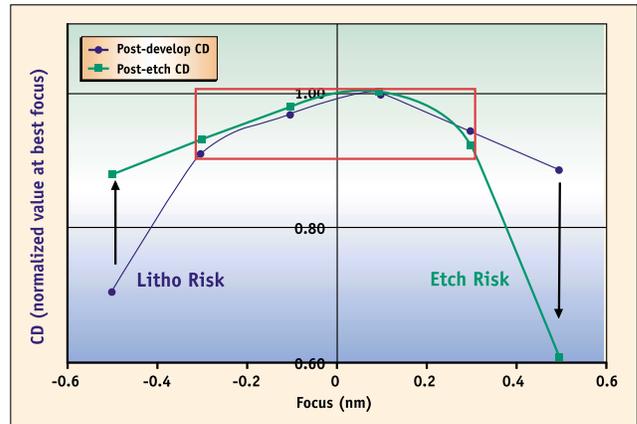


Figure 8. Example of a litho-etch focus window interaction that can produce “unexplained variation” in CD APC applications.

potential for catastrophic yield loss. Most CD APC applications are blind to this class of inter-module variation.

The root-cause of etch window displacement may not be the CDs themselves, but rather the resist sidewall angles as they change through focus and erode during etch. Figure 9 shows a cross-section SEM micrograph of a gate stack, upon which we have superimposed profiles that were generated using library-based, spectroscopic scatterometry (SCD). The SCD profiles characterize erosion of the resist during etch and accurately measure the slope of the underlying silicon, all with sub-nanometer repeatability. We expect that PWM will be most effective when embedded in smart metrology tools that provide large amounts of unique information.

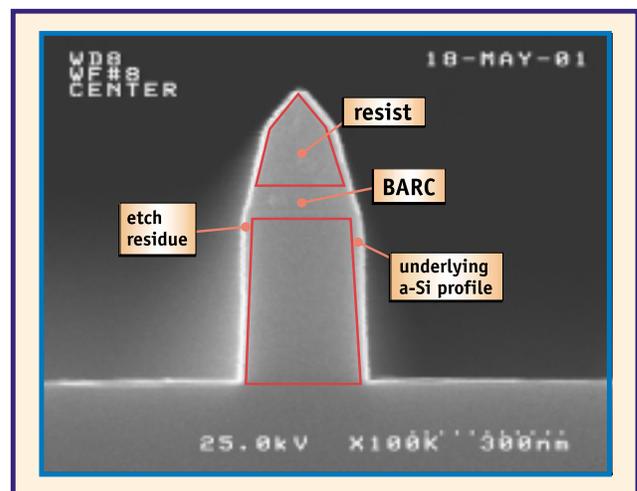


Figure 9. A potential root-cause of litho-etch interaction is resist sidewall angle (not necessarily the CD itself). Profiles were generated using library-based scatterometry.

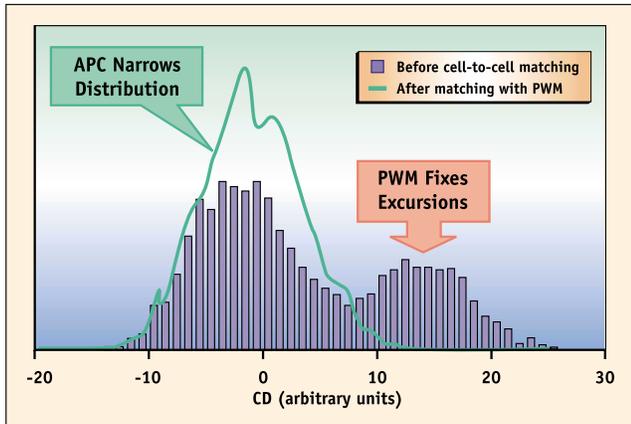


Figure 10. Use-cases for process window management (PWM) and advanced process control (APC). PWM can be used to predict and correct excursions and to stabilize process windows while APC narrows “in-control” distributions. APC can access PWM monitor data as part of the manufacturing context.

In addition to its role in supporting APC, PWM can be used by itself as a fully automated, web-based monitor for process tool stabilization and tool matching across entire areas. In lithography, CD PWM is implemented using multiple focus-exposure matrices measured on a SEM or SCD metrology tool. Automated analysis of these matrices can generate dose, focus, and depth-of-focus metrics for each lithography cell and each critical layer in a factory. For example, Figure 10 shows a process window excursion (histogram chart) that was corrected (line chart) using exposure dose information from a CD SEM-PWM system. Figure 11 illustrates a case where focus information was used to monitor focal plane tilt over five sites across a scanner field. As shown, focal plane tilt shrinks the cross-field overlap

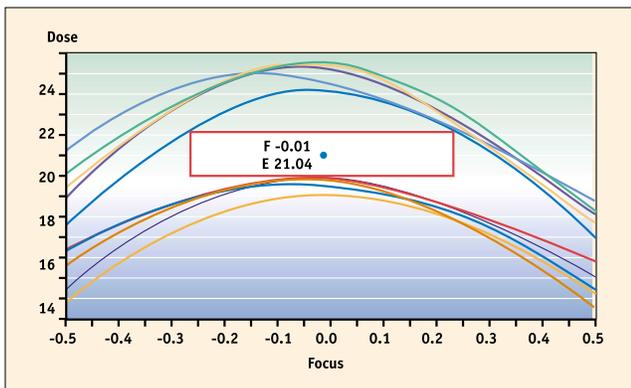


Figure 11. A large overlap window is critical for manufacturing. Here, PWM assists in determining the common lithographic focus-exposure window over five positions in the scan field. Overlap windows can be spatial, temporal, or tool-to-tool.

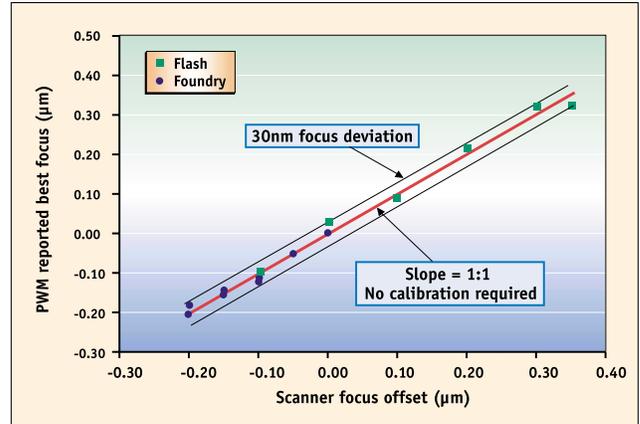


Figure 12. Performance of the PWM system must exceed that of available in-situ metrology on process equipment. Flash and foundry data from two different companies indicate maximum error of about 30 nm without calibration (~12 nm with calibration).

window relative to the individual process windows. Finally, Figure 12 demonstrates system-level performance for PWM that exceeds the capability of in-situ metrology available on scanners. Combining flash memory and foundry data from two different factories, we observed a maximum focus tracking error of 30 nm without calibration. Our development work indicates that, with calibration, the noise floor for CD SEM-PWM system is less than 12 nm-3σ.

Conclusion

Accelerated semiconductor technology roadmaps, fueled by powerful economic forces, have created the phenomenon of rapidly shrinking process windows. This trend is driving us toward ubiquitous control and forcing convergence of today’s APC, AEC, FDC, and PWM technologies into a more generalized Process Window Control (PWC). In this work, we have discussed the APC and PWM components of PWC and have shown how they can be combined with measurement tools in a strategy to increase the profitability of semiconductor manufacturers, primarily by accelerating shrinks, improving performance, and lowering costs. Specifically, we conclude that

- Economic considerations will drive demand for more granularity in parametric control, particularly intra-field and field-to-field dose control in lithography.
- Control of module-to-module parametric window overlap will become a necessity as gate CD control merges with L-effective control and ultimately with electrical parametric control.

- Far more effective and complex APC applications will appear as measurement and process delay times are reduced and better models are developed to minimize unexplained variation.
- Further reduction of unexplained variation will be accomplished using process window management, smart metrology tools, and statistically optimized sample plans.
- Further reduction of delays will be accomplished by increasing the level of metrology integration, clustering, and dedication.

Finally, in difficult economic times, the metrology-driven shrink may be the most cost-effective alternative to massive investment in new process equipment, or factories, since it both extends the life of current assets and provides the foundation for higher returns during economic recovery. Advanced process window control is the key enabler of this strategy.

Acknowledgements

The authors would like to acknowledge the support and encouragement of Gary Dickerson, Rod Browning, Tom Long, Ellis Chang, Ed Charrier, Dennis Fortino, Brian Trafas, and Avi Cohen, all of KLA-Tencor

Corporation. We also acknowledge fruitful discussions with Julia Castillo, Arun Chatterjee, Jim Hamilton, Fabio Pintchovski, Viral Hazari, Kwame Eason, Andy Kirkpatrick, and Satya Akella, also of KLA-Tencor Corporation. In addition, we reserve special thanks for Robert Routh of Atmel Corporation and others who helped us refine some of the concepts presented in this work.

References

1. C. Meade and L. Conway. Introduction to VLSI Systems. Menlo Park: Addison-Wesley, 1980, pp 1-5. Also, K. Itoh. VLSI Memory Chip Design. New York: Springer-Verlag, 2001, pp. 1-99.
2. K. M. Monahan, S. Ashkenaz, X. Chen, P. Lord, M. Merrill, R. Quattrini, and J. Wiley, "Accelerated yield learning in aggressive lithography," Proc. SPIE, Vol. 3998, p. 492 (2000).
3. R. C. Elliott, R. K. Nurani, S.-J. Lee, L. Ortiz, M. Preil, J. G. Shanthikumar, T. Riley, and G. Goodwin, "Sampling plan optimization for detection of lithography and etch CD process excursions," Proc. SPIE 3998, pp. 527 (2000).
4. Metrology sample plan optimization is a value-added service provided by KLA-Tencor Corporation.

A version of this article was originally published in the 2002 SPIE Microlithography proceedings 4692, SPIE Microlithography Conference, March 3-8, 2002, Santa Clara, California, USA.