# Characterization of Demand for Short Life-Cycle Technology Products


S. David Wu
Berrin Aytac
Lehigh University



Report No. 07T-005

# CHARACTERIZATION OF DEMAND FOR SHORT LIFE-CYCLE TECHNOLOGY PRODUCTS

S. David Wu and Berrin Aytac
Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA 18015
david.wu@lehigh.edu, bea4@lehigh.edu

**Abstract**

The increasing pace of product introductions in the high-tech markets leads to shortening product life cycles, which undergo rapid growth, maturity, and decline phases. The markets also face sudden changes, such as the decline in 2001 that followed the growth and expansion in the mid-to-late 1990's. In this environment, traditional time-series forecasting techniques that require long series of historical demand data with stable trend fail to generate accurate forecasts. As an alternative to these techniques, we study a demand-characterization method that uses cumulative *life-cycle growth models* to capture the changes in demand trend due to the short life-cycle nature of the products, and *leading-indicator products* to capture the changes in demand trend due to the changing market conditions: Demand patterns of the products are projected using several growth models. The combined estimate of these models are used to generate demand scenarios. With the advanced information from leading indicators, the projections of the models are (Bayesian) updated, and demand scenarios with smaller variance are obtained. Using data from three semiconductor manufacturing companies, we empirically validate that the leading indicator data replaces the actual data well and that the variance of the demand scenarios decreases.

## 1 Introduction

In the mid-to-late 1990s, high-tech industries, such as consumer electronics, telecommunications equipment, and semiconductors, grew rapidly. To reduce costs and cycle times, many firms developed and deployed supply-chain-management systems, but continued to rely on traditional demand planning in which the basis of a demand forecast is either an internal marketing judgement based on customer projections or short-to-medium term traditional time-series forecasting methods that rely heavily on historical data and require stable demand trend. Neither is effective in characterizing volatile, non-cyclic high-tech product demands and in responding to the fast changing conditions of the market place such as the succeeding industry decline in 2001.

In these challenging technology-driven market conditions, the firms rapidly innovate technology and introduce new products to maintain their competitive position. This leads to shortening technology life cycles for high-tech products, which are known to follow a general demand life cycle that starts with an initial growth (ramp up) followed by a period of stability and then a decline in sales when a new generation of products is introduced. In this study, we develop a medium-to-long term demand modeling approach that combines the insights from growth models and leading indicator products in order to characterize the demand life-cycles and advanced market signals in an insightful manner.

This study has been motivated by our extensive analysis of demand data of a major U.S. semiconductor manufacturing company. We conducted the analysis first in 1997-1999 and then in 2000- 2002. In the earlier study (Meixell and Wu, 2001), after analyzing demand data for some 3,500 products, we found that these products follow a few (approximately six) life-cycle patterns, and that the products can be grouped according to these patterns using statistical cluster analysis. More importantly, after doing correlation analysis on historical shipment data, we found that in each product group there exists a subset of leading indicator products that are likely to give advanced indication of changes in demand trends. In the latter study (Wu et al., 2006), we no longer group products based on similar life-cycle patterns but focus on exogenously defined product groups based on fab technology, market segments, and product families. We are again able to find strong leading indicators that behave similar to the product group at a certain period ahead of time. In the current study, we use leading indicator products as a means to to obtain advanced demand signals about the changing market conditions, and utilize this information to extend the available data set of the product group that they lead.

Besides leading indicators, the second key ingredient of our demand characterization analysis is the use of life-cycle growth models. Combining the concepts of product life cycles and forecasting, growth models provide attractive alternatives to the traditional time-series forecasting methods for short life-cycle (high-technology) products (Kurawarwala and Matsuo, 1996). Technological forecasting literature, which studies the diffusion of innovations in a population, suggests several growth models that we can use to describe a wide variety of life-cycle patterns observed in practice. Mahajan, Muller and Bass (1990) provide an extensive review of these models and study their applications to marketing. Meade and Islam (1998) collect 29 models in the technological forecasting literature and group them into classes. The classification is based on either the way the models incorporate randomness, or their shape that is characterized by the timing of the point of inflection, which corresponds to the point of the growth curve with the maximum rate of sales.

Availability of several models, however, makes it difficult to select the model that best describes the historical data of a given set of products and that generates accurate forecasts. As opposed to the technological forecasting literature, which mostly focuses on the performances of individual models, we use a large collection of models to characterize and predict demand growth patterns of the products. Furthermore, we describe demand in terms of scenarios that are obtained from the combined estimate of these models, rather than from the estimate of a single best model. However, one problem associated with the short life-cycle product environment is the scarcity of available data. The small number of data points would result in demand scenarios with large variation. In this paper, we extend the available data set of the product group with the use of several leading indicator products. The additional data from leading indicators streamline the characterization of the future demand patterns of the product group and generates demand scenarios with smaller variation.

There are other studies that use additional information from sources other than the historical data. Sultan et al. (1990) perform a meta-analysis of 213 applications of diffusion models from 15 articles. The impact of factors such as type of innovation, geographic effect, and the marketing variables on the value of parameters is determined. The results of the meta-analysis are used before any data becomes available. As new data arrives, the meta-analysis results are used as priors to Bayesian update the data-based estimates of the parameters. The Bayesian scheme appears to produce more robust results, particularly early in the product history. The inclusion of additional data is not through the direct use of the time series data of the product. It is a retrospective study, it does not take into account the future changes in the conditions of the market and the product. Islam et al. (2002) consider simultaneously occurring diffusion of innovation in several countries which have different starting times but similar dynamics of the diffusion process. Data from several countries are pooled using the linearized growth curve models. Significant parameter estimates, and more accurate forecasts with pooling when compared to a deterministic trend model. Similar to our approach, the use of the past data of a product in a country helps the prediction of the future of the same product in another country. However, Srinivasan and Mason (1986) state that the use of linearized version of the growth curves introduces a bias.

There are only a few studies that use Bayesian approach to forecast demand over product life cycles. Bewley et al (2001) compare the performances of the classical and the Bayesian approach using three variants of the logistic function, and conclude that the Bayesian approach produces more accurate forecasts and more reliable confidence intervals. Zhu et al (2004) propose an adaptive demand forecasting algorithm using Bass diffusion model. These studies use Bayesian approach to update the parameters of the growth models.

In the next section, we introduce the characterization of the demand patterns of the products using cumulative life-cycle growth models. We model the uncertainty in the projections of these models in Section 3. Using Bayesian approach, we incorporate the additional information from leading indicators to update the projections and the probabilities of the models that are obtained with the available data only. In Section 4, we consider the combined estimate of the models, from which we generate demand scenarios in Section 5. The implementation of our approach on a real-life data set is presented in Section 6, and concluding remarks follow in Section 7.

## 2 Life-Cycle Demand Projections

Growth models are often used to produce medium-to-long-term forecasts. Since the use of non-cumulative growth models requires the precise prediction of the patterns of a given time-series data, we use cumulative life-cycle models, which is more efficient for medium-to-long-term forecasting. For products with short life cycles, the life-cycle sales typically go through one phase of rapid growth, maturity, and decline. That is, they follow a single-modal curve. Therefore, an S-shaped curve characterizes the projection of the cumulative proportion of the total expected life-cycle demand that realizes over time.

With this new characterization of demand with S-shaped growth curves, forecasting the demand in each time period requires the projection of the cumulative life-cycle demand patterns of the products into future. We use several cumulative life-cycle growth models, each of which exhibits different features, to describe the demand pattern of a given set of products (Figure 1). Given the projection of the life-cycle demand of the products up to any time T, we forecast the proportion of the projected demand that will be met by any time T

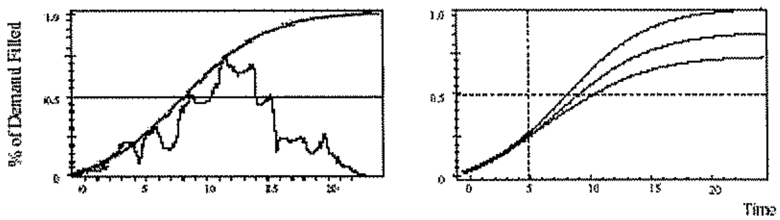


Figure 1: Cumulative Demand and Projections of Different Curves

+ M, $M > 0$: Let X(t) be the proportion of the life-cycle demand that has been met by time t ($0 \leq X(t) \leq 1, t > 0$), and let $\widehat{X}_k(t|\Theta_T)$ represent the estimate of the cumulative growth model k, $k = 1, ..., K$, given data up to time T, i.e., $\Theta_T = \{X(1), ..., X(T)\}$, such that:

$$X(t) = \widehat{X}_k(t|\Theta_T) + \epsilon_k(t|\Theta_T), \tag{1}$$

where $\epsilon_k(t|\Theta_T)$ represents the estimation error. We assume that it is normally distributed with mean zero and a constant known variance $\sigma^2_{\epsilon_k}$. Therefore, the estimate of a cumulative life-cycle growth model for time T+M is assumed to generate an unbiased forecast for the actual proportion of the total life-cycle demand that will be met by time T+M.

# 3 Uncertainty in the Life-Cycle Demand Projections

Given the importance of associating a measure of uncertainty with a point forecast, we express the uncertainty in the estimates of the growth models. For this purpose, we use Bayesian approach, which treats the estimate of each growth model as a random variable and describes each random variable with a probability distribution. This allows a complete characterization of the uncertainty associated with the estimates.

Bayesian approach starts with a prior distribution that summarizes the information available. With the observation of new data, the prior distribution is updated, which results in the posterior distribution. In the light of this information, we use the prior projections of the growth models that are obtained using the data available up to time $T$ to derive the prior distributions of the estimates of the models (Section 3.1.). We then use several leading indicators to represent the actual data that realizes after time $T$. The estimate of a growth model over the data set extended with each leading indicator provides a sample life-cycle projection for the true projection of the model that would be obtained with actual data only (Section 3.2.). Bayesian updating the prior distribution with the information from leading indicators, we obtain posterior distributions that characterize the posterior life-cycle projections as the new estimates (Section 3.3.).

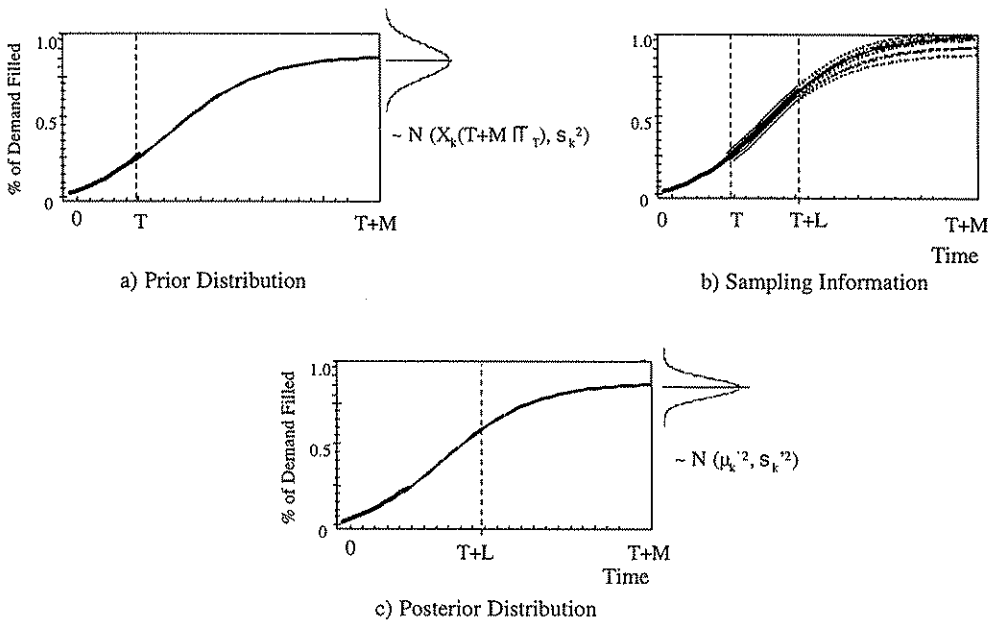


Figure 2: Probability Distribution of the Estimate of a Growth Curve

## 3.1 Prior Life-Cycle Projections

Remembering that the estimate of a model with the data available up to time T, i.e., $\widehat{X}_k(T+M|\Theta_T)$, k=1,...,m, is an unbiased estimator of the actual data at time $T+M$, i.e., $X(T+M)$, we use $\widehat{X}_k(T+M|\Theta_T)$ to obtain prior information for $X(T+M)$: In particular, we assume that the actual data is normally distributed around the prior mean with a certain variance, $\sigma_k^2$:

$$\widetilde{X}_k(T+M) \sim N\left(\widehat{X}_k(T+M|\Theta_T), \sigma_k^2\right)$$

where $\widetilde{X}_k(T+M)$ denote the random variable that corresponds to the estimate of model k for actual data at time $T+M$.

We quantify the uncertainty in $\widetilde{X}_k(T+M)$ by its variance, $\sigma_k^2$. Under the assumption that the estimate of a growth model and the error between the actual data and the estimate are independent (1), $\sigma_k^2$ can be expressed as the sum of the variance of the estimate of the growth model, $\omega_k^2$, and the variance of the the error, $\sigma_{\epsilon_k}^2$ :

$$\begin{aligned} \widehat{var}\left(\widetilde{X}_k(T+M)|\Theta_T\right) &= \widehat{var}(\hat{X}_k(T+M|\Theta_T)) + \widehat{var}(\epsilon_k(T+M|\Theta_T)) \\ \sigma_k^2 &= \omega_k^2 + \sigma_{\epsilon_k}^2. \end{aligned}$$

The main sources of the uncertainty in the estimate of a growth model are known to be the estimation error in its parameters and the inadequacy of the selected model. To estimate the variance of this quantity, Meade and Islam (1995) propose three approaches: The first approach, the approximated variance approach, uses the variance of the linearized Taylor series expansion of the model at the final estimates of the parameters and their (asymptotic) covariance matrix. The second approach, the explicit density approach, derives the estimate of the variance by explicitly modeling the uncertainties in the parameter estimates. The third approach, bootstrapping, is a non-parametric approach that develops the distribution of the estimate of the growth curve by resampling the residuals.

## 3.2 Samples of Life-Cycle Projections

The aim of using leading indicator products is to obtain advanced indications about the demand patterns of a group of other products. Given data up to time $T$, we determine a candidate leading indicator product based on the absolute value of the correlation coefficient between the demand time series of the candidate product and that of the product group shifted by the time lag of the leading indicator. Once we determine a leading indicator product which leads the demand pattern of the other products at a time lag L, by regressing the time series of the leading indicator over $[1, T-L]$ against the time series of the product group over $[L+1, T]$, we determine the relationship that transforms the leading-indicator-based data into the data of the product group.

At time period $T$, we use the $L$-period data from a leading indicator to substitute the actual data of the product group through periods $T+1$ to $T+L$, which are unknown at time $T$. As illustrated in Figure 2b, this advances our position on the projected curve from time $T$ to time $T+L$. In order to increase the likelihood that the added data correctly replaces the actual data, we consider several leading-indicator products. With each leading-indicator product, we obtain the projection of each model, each of which gives a sample life-cycle projection for the true projection of the model that would be obtained if actual data was used over $[1, T+L]$.

We assume that the data from leading indicators provide unbiased samples for the actual data over T+1 to T+L; therefore, the estimate of a model with a leading indicator also provides an unbiased sample for the true estimate of the model that would be obtained with actual data over 1 to T+L. We use a sample of m leading indicators to derive the sampling information about the actual data at time $T+M$.

The estimate of model k with leading indicator j, $j = 1, ..., m$, is obtained over the estimation period extended with the leading-indicator-j-based data, $\Theta^j_{T+L}$. We can express our assumption that the estimate of a model with each leading indicator provides an unbiased sample for the true estimate of the model with actual data, i.e. over $\Theta_{T+L}$, as:

$$\widehat{X}_{kj}(T+M|\Theta^j_{T+L}) = \widehat{X}_k(T+M|\Theta_{T+L}) + \varepsilon_{kj} \tag{2}$$

where $\varepsilon_{kj}$ represents sampling noise. We assume that $\varepsilon_{kj}$ is normally distributed with mean zero and constant known variance $\varphi^2_k$. Integrating our assumption that the estimates of the models are unbiased estimators of the actual data gives that the estimates of the models with the leading indicators are unbiased estimators of the actual data. As a result, under the assumption that the leading indicators are independent, it is obtained that $\{\widehat{X}_{kj}(T+M|\Theta^j_{T+L}), j = 1, ..., m\}$ are independently and identically distributed observations from a normal distribution with mean $X(T+M)$ and variance $\tau^2_k = \sigma'^2_{\epsilon_k} + \varphi^2_k$, where $\sigma'^2_{\epsilon_k}$ represents the variance of the estimation error at time $T+L$. We can summarize the sampling information as:

$$\widetilde{X}_k(T+M) \sim N\left(\frac{1}{m}\sum_{j=1}^{m}\widehat{X}_{kj}(T+M|\Theta^j_{T+L}), \frac{\tau^2_k}{m}\right)$$

By the strong law of large numbers, with probability one the sample mean provides an unbiased estimate for the actual data at time $T+M$ with variance $\tau^2_k/m$, and the confidence in the sample mean increases with the number of leading indicators.

**Biasedness of Leading-Indicator Data**

There is a possibility that the data from a leading indicator is not an unbiased estimator of the actual data over $[T+1, T+L]$. Thus, the estimate of a model with this leading indicator will not be an unbiased estimator of the true estimate of the model obtained with actual data over $[1, T]$, either. If we describe the biasedness as a random variable and characterize it as normally distributed with mean $b_{kj}$ and variance $\theta_{kj}^2$, it follows that:

$$\begin{aligned} E[\widehat{X}_{kj}(T+M|\Theta_{T+L}^j)] &= E[\widehat{X}_k(T+M|\Theta_{T+L}) + \beta_{kj}] \\ &= X(T+M) + b_{kj}. \end{aligned}$$

The mean and the variance of the sampling distribution changes to:

$$E\left[\frac{1}{m}\sum_{j=1}^{m}\widehat{X}_{kj}(T+M|\Theta_{T+L}^j)\right] = X(T+M) + \frac{1}{m}\sum_{j=1}^{m} b_{kj}$$

$$Var\left[\frac{1}{m}\sum_{j=1}^{m}\widehat{X}_{kj}(T+M|\Theta_{T+L}^j)\right] = \frac{\tau_k^2}{m} + \frac{\sum_{j=1}^{m}\theta_{kj}^2}{m^2}.$$

The results indicate that biasedness increases the variance of the sampling estimate; the mean of the estimate may shift depending on the relative biasedness of the leading indicators.

## 3.3 Posterior Life-Cycle Projections

The Bayesian update of the prior distribution, which is based on the estimate of the models with actual data available up to current time $T$, with the sampling information, which approximates the true estimate of the models with several leading indicators, gives the posterior distribution of the actual data at time $T+M$ :

**Proposition 1** *When the m leading-indicator-based sampling estimates follow a normal distribution with mean $X(T+M)$ and variance $\tau_k^2$, and when $X(T+M)$ has prior normal distribution with mean $\widehat{X}_k(T+M|\Theta_T)$ and variance $\sigma_k^2$, the posterior estimate is also normally distributed:*

$$\widetilde{X}_k(T+M) \sim N(\mu_k', \sigma_k'^2)$$

*where*

$$\begin{aligned} \mu_k' &= \frac{1/\sigma_k^2}{1/\sigma_k^2 + m/\tau_k^2}\,\widehat{X}_k(T+M|\Theta_T) + \frac{m/\tau_k^2}{1/\sigma_k^2 + m/\tau_k^2}\,\frac{1}{m}\sum_{j=1}^{m}\widehat{X}_{kj}(T+M|\Theta_{T+L}^j) \\ \sigma_k'^2 &= \frac{\sigma_k^2\tau_k^2}{m\sigma_k^2 + \tau_k^2} \end{aligned}$$

The derivation of the posterior distribution is given in Appendix A.3. The updated life-cycle demand projections correspond to the means of the posterior distributions, which are the weighted averages of the prior projections and the sampling means, with the weights being inversely proportional to the variances of the prior distribution and the sampling distribution, respectively.

The use of additional data from leading indicators leads to the following conclusions:

**Corrolary 1** *Bayesian updating the projections of the growth models with the additional data from leading indicators reduces the variance of the estimates of the models, i.e.,* $\sigma_k'^2 \leq \sigma_k^2$.

**Corrolary 2** *prior variance* $\geq$ *posterior variance with some biased leading indicators* $\geq$ *posterior variance with unbiased leading indicators.*

**Corrolary 3** *The variance of the estimates asymptotically reduces to zero as the number of leading indicators increases.*

The assumption of normal distribution is not restrictive. The reason for making this assumption can be expressed as follows: Normal distribution provides a reasonable characterization of the sampling and the estimation errors. In our case, this results in a normal sampling distribution. The corresponding natural conjugate prior distribution, which is selected for mathematical convenience, is also normal distribution. The use of natural conjugate prior distributions has the advantage that the induced posterior distribution is in the same family as the prior distribution, and therefore allows for closed-form expressions. The analysis is still valid if this assumption is released. In particular, when the variance is unknown, sample variance can be used, in which case t-distribution is a more accurate representation of the uncertainty in the sampling estimates.

# 4 Combined Estimate of the Life-Cycle Projections

In this section, we consider the combination of the estimates of the projections of the growth models and show that the variance of the combined estimate reduces with the additional data from leading indicators.

The topic of forecast combination has been widely studied since the seminal work of Bates and Granger (1969). The motivation for combining forecasts has been to avoid the risk associated with the choice of the "best" single forecasting model and to aggregate information from different models. Several empirical studies show that combining forecasts improves the forecasting accuracy and reduces the variance of forecasting errors.

Among the many different combination methods, we consider the one that minimizes the variance of the combined forecast (Dickinson 1973): Under the assumption that the estimates of different models are independent, this method assigns each model a weight, $p_k$, that is inversely proportional to the variance of its estimation error. For the prior distribution of the estimates,which are normally distributed with mean $\hat{X}_k(T+M|\Theta_T)$ and variance $\sigma_k^2$, the weights become:

$$p_k = \frac{1/\sigma_k^2}{\sum_{j=1}^{K} 1/\sigma_j^2}$$

.

Furthermore, consider the combined forecast, $Y$, defined by $Y = \sum_{k=1}^{K} p_k \widetilde{X}_k(T+M)$. From the Reproductive Property of the Normal Distribution, which states that the linear combination of independent normal random variables is normally distributed, it follows that

$$Y = \sum_{k=1}^{K} p_k \widetilde{X}_k(T+M) \sim N\left(\sum_{k=1}^{K} p_k \hat{X}_k(T+M|\Theta_T), \sum_{k=1}^{K} p_k^2 \sigma_k^2\right),$$

where the variance of the combined forecast is :

$$\sigma_c^2 = \sum_{k=1}^{K} p_k^2 \sigma_k^2 = \frac{1}{\sum_{k=1}^{K} 1/\sigma_k^2} \tag{3}$$

As it can be noticed, the variance of the combined forecast is no greater than the variance of the estimate of any model. The same results also hold for the posterior distribution of the estimates, which are normally distributed with mean $\hat{\mu}_k$ and variance $\sigma_k^{\prime 2}$. For this case, we denote by $\sigma_c^{\prime 2}$ the variance of the combined forecast. Figure 3 illustrates the distribution of the estimates of individual models and combined forecast at any future period T+M.

We can reach the following conclusion:

**Proposition 2** *Reduced variances of the estimates of the models with more data points from leading indicators result in reduced variance of the combined forecast, i.e.,* $\sigma_c^{\prime 2} \leq \sigma_c^2$.

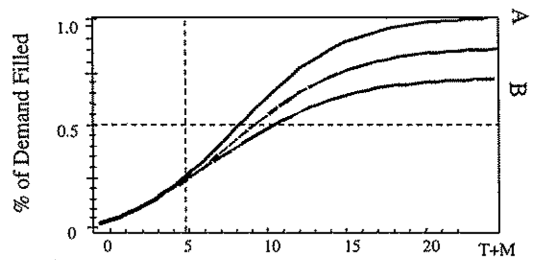


(a) Prediction with growth models

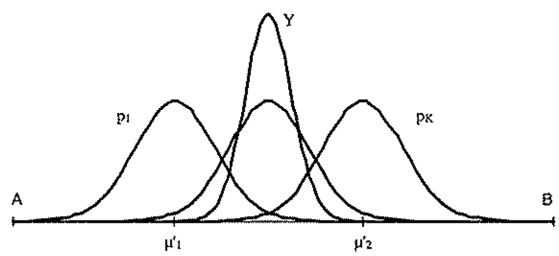


(b) The probability distribution of the estimate of the growth models and the combined forecast at any time T+M

Figure 3: Distribution of the Estimates of the Growth Models and the Combined Forecast for time T+M

# 5 Life-Cycle Demand Scenarios and the Variation of the Scenarios

We combined the projections of several growth models as a weighted average of the estimates of the models that minimizes the variance of the combined forecast. This allowed us to consider the relative accuracy of the models and to obtain a combined forecast with smaller variance than the variance of its component model forecasts. We showed that the additional data from leading indicators reduce the variance of the estimates of the individual models and therefore the variance of the combined forecast. The main thesis of this paper is that using the resulting probability distribution of the combined forecast with smaller variance, demand scenarios with smaller variance are obtained (Figure 4).

One way to generate life-cycle demand scenarios is simply to use random sampling

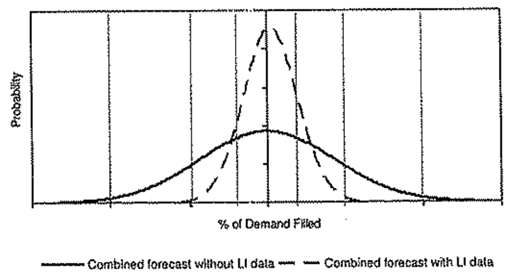


Figure 4: Sampling Life-Cycle Demand Scenarios from the Distribution of the Combined Forecasts

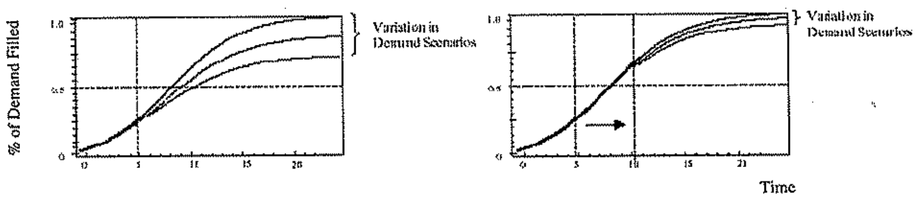


Figure 5: Life-Cycle Demand Scenarios

from the probability distribution of the combined forecast. As an alternative to random sampling, there are several sampling techniques that the literature suggests as appropriate for generating demand scenarios more systematically and obtaining demand scenarios with smaller variance as compared to those generated with random sampling. One such technique, for example, is a specific case of stratified sampling, which provides considerable intuitive appeal: The range of probable demand values are divided into n segments of equal probability, where n is the number of scenarios to be generated (Figure 4). Then, a random demand is generated from each range as a possible demand scenario(Figure 5).

As illustrated in Figure 5, through a Bayesian update, the additional data from leading indicators can reduce the variation on the projected demand growth patterns (i.e., the life-cycle scenarios) by a significant margin. One of the reasons is that a small movement on the time axis might correspond to drastic change on the curve, especially when the point of inflection is included in the movement.

# 6 Empirical Analysis

## 6.1 Experimental Data

For the purpose of empirically demonstrating the impact of additional information from leading-indicator products on the characterization of life-cycle patterns of the products, we use disguised real data sets from three semiconductor manufacturing companies, which we will refer to as company A, company B, and company C. The latter two companies are members of the Semiconductor Research Corporation.

Company A provided us with the weekly sales data of around 3000 products from January 2002 to December 2003. Each product has attributes such as strategic business unit, business entity, fab process group, tester group, and package type. The data from company B is the weekly sales data of 228 products from April 2004 to March 2006. All the products are in the same strategic business unit. The data from company C is the weekly sales data of over 2000 products from February 2005 to March 2006. We know the business unit, wafer-fab process group, package group and technology code of each product. In order to lessen the impact of the short-term fluctuations that carry less information in medium-to-long term forecasting, we transformed the data into monthly figures. We normalized the monthly data by taking into consideration that some fiscal months are of 4 weeks and others are of 5 weeks.

For each company, we demonstrate the characterization of the life-cycle demand pattern of an exogenously determined product group, as in Wu et al. (2006). The product group we consider for company $A$ consists of 120 products that share the same fab capacity and belong to the same market segment. For company $B$, we consider all the 228 products in the given strategic business unit. For company $C$, we study a group of 969 products in a strategic business unit.

## 6.2 Experimental Design

Meade and Islam (1998) study 29 models from the technological forecasting literature using both simulated and real data sets, and determine seven well-performing models in terms of fitting and forecasting. In order to practically perform our analysis, we use these models, namely, simple logistic, Mansfield, Gompertz, Floyd, Weibull, extended logistic, and cumulative log-normal. Simple logistic and Mansfield are symmetric models about the point of inflexion which is fixed when value of cumulative proportion is 0.5. Gompertz and Floyd are asymmetric models with a point of inflexion fixed at a value of cumulative fraction less than 0.5. The point of inflexion for Weibull, extended logistic, and cumulative log-normal is flexible and can take a range of values that includes 0.5 depending on the parameter values. Different shapes of these models allow the characterization of a wide range of life-cycle patterns observed in practice.

Growth models describe the projections of the the cumulative proportion of the total life-cycle demand that is met over time. Hence, the given time-series data of the products needs to be transformed into the scale that represents the cumulative proportion. This requires an estimate of the expected total life-cycle sales, also known as the market potential. This quantity can either be estimated as a parameter from the given data set, or its market estimate, which can be determined through market surveys or management judgement, can be provided as an input to the model. We use the market estimates, since the estimation procedure results in unrealistic estimates of this parameter in order to increase the fitting performance. Heeler and Hustad (1980) and Tigert and Farivar (1981) are among the empirical studies that compare these two approaches and report that the accuracy of the forecasts increases significantly when the estimate of this quantity is an input.

We proceed as follows: There are 24-month data for company $A$ and company $B$, and 14-month data for company $C$. We first use the initial 9-month period ($T = 9$) as the *estimation period*, in which the prior life-cycle projections of the individual models and the combined forecast are obtained. Next, we use the $L$-period demand signals from $m$ leading-indicator products. This extends the estimation period to month $T + L$, the first $T$ months from the actual data and the following $L$ months from the leading-indicator-based data. The prior projections of the models and the combined forecast are updated with this additional information. We use the remainder of the data, i.e., $[T + L + 1, 24]$ as the *validation period* to compare the performances of the estimates with and without additional data.

The main thesis of this study is that the variance of the combined forecast reduces with the additional data. In order to verify if the reduction in variance is significant, we construct a paired comparison experimental design, which compares the variance of the combined forecast with and without additional information over the same sets of data. The variances of the combined forecast for each period in the validation period over all the companies provides a block of observations for the paired comparison design. The null hypothesis that variance remain the same $H_0 : \sigma_C^2 = \sigma_C'^2$ is tested against the one-sided alternative hypothesis $H_1 : \sigma_C^2 > \sigma_C'^2$. The hypothesis is evaluated using a t-test statistic based on the standardized average difference between the variance of the prior estimates and the variance of the posterior estimates over all the blocks .

We identify three factors that can affect the results : (i) number of leading-indicator products, (ii) number of data points from a leading-indicator product, and (iii) the biasedness of leading-indicator products. From Corrolary 3, it follows that variances of the estimates decrease at a diminishing rate with the number of leading indicators in case all the leading indicators are unbiased. Furthermore, by the law of large numbers, as the number of leading indicators goes to infinity, the variance of the estimates goes to zero. However, it is likely that data from some of the leading indicators are biased, and therefore the results are distorted. We want to measure if there exist biased leading indicators, how significantly they affect the results. In addition, leading-indicator-based data are used to replace the actual data. As the number of data points from a leading indicator increases, the estimation quality of the leading indicators may be affected, which in turn has an impact on the estimation quality of the growth models. To analyze the impact of these factors, we consider different number of leading indicators ($m = 5, 10, 15, 20$) that provide data up to a time lag of 5 months ($L = 1, ..., 5$). We present the experimental results in the next section.

## 6.3 Experimental Results

We start our analysis by illustrating the impact of a single leading indicator on the estimates of the models. Figure 6a plots the prior projections of the models obtained over the estimation period with 9 months of actual data. The demand signals from a leading indicator with a time lag 5 are used to extend the estimation period to month 14, and the projections of the models change as in Figure 6*b*.

The updated projection of each model with this leading indicator provides a sampling instance for the true projection of the model that would have been obtained over the estimation period with 14 months of actual data. Similarly projecting with several leading indicators give the sampling distribution of the estimate of each model. Combining the sampling distribution with the prior distribution determined at time 9 gives the posterior distribution.

We perform the analysis using different combinations of $m$ and $L$. Our goal in

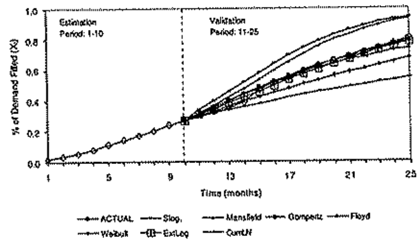


a. Prediction over the first 10 months of actual data (without a leading indicator)

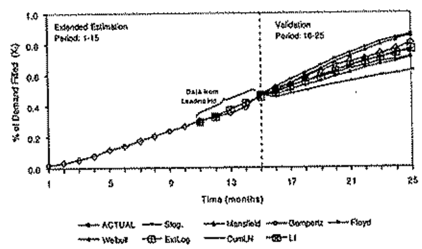


b. Prediction over the 15 months of data (with 5-month demand signal from a leading indicator)

Figure 6: Projection of Cumulative Demand with Life-Cycle Growth Models

performing this computational testing is to study: (i) How closely do the sampling means approximate the true estimates with actual data? (ii) How significant is the reduction in the variance of combined forecast with respect to the number of leading indicators used? (iii) Does the number of data points from each leading indicator have a notable effect on the results? (iv) Are there any biased leading indicators? If there are, what is the impact of biasedness on the mean and variance of the combined forecast?

*Approximation Quality of the Sampling Estimates*

Our objective in using Bayesian forecasting is to capture the future demand signals in the highly volatile market environment with limited historical data. We aim to capture the future demand signals via leading indicators. It is inevitable that the true prediction of leading indicators of the actual demand data to be realized in future is critical to the success of our approach. In order to assess the impact of the prediction power of the leading indicators on the given data sets, we compare the mean estimates of the sampling distributions that are the average of the estimates of several leading indicators against the true estimates that would have been obtained with actual data.

Figure 7 depicts the mean estimate of the sampling distribution under different

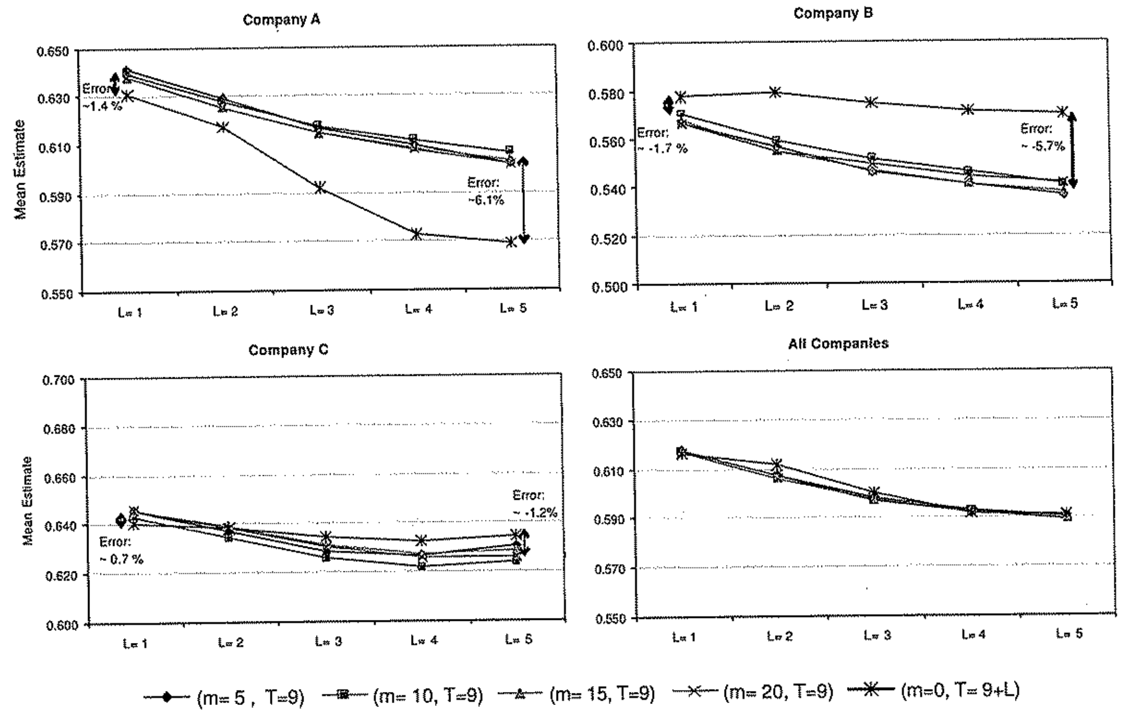


Figure 7: Comparison of sampling means against the true estimates with actual data.

combinations of $m$ and $L$ across all the companies. The mean estimate for each combination of $(m, L)$ is the average of the sampling means of all the seven models for all the periods in the validation period using $T$-period actual data and $L$-period leading-indicator-based data as the estimation period. The true estimates are the averages of the estimates of the models over the validation period if actual data were used over $[1, T + L]$, and shown with the line that corresponds to the tag $(m = 0, T = 9 + L)$.

As shown in the figure, the estimation performance of the sampling mean tends to degrade with $L$. The sampling mean is within 2% range of the true estimates with actual data when the available data is extended with one-month demand signals from leading indicators. The range increases up to around 6% as the estimation period is extended with more data points from leading indicators. These changes are most notable for company A and company B. A possible reason is that data from leading indicators estimates the actual data over $[T + 1, T + L]$. As $L$ increases, that is as the number of data points to estimate increases, the approximation quality of the leading indicators decreases. This, in turn, affects the estimation quality of the models. The figure also suggests that the impact of the number of leading indicators on the estimate of the sampling distribution is not significant.

*Diminishing Value of Leading Indicators*

The empirical analysis that compares the variances of the combined forecast over all the time horizons in the validation period across the data sets from the three companies rejects the the null hypothesis that the variance of combined forecast obtained from prior and posterior projections is equal. That is, the reduction in variance is significant. The test's $p$-value is $< 0.01$.

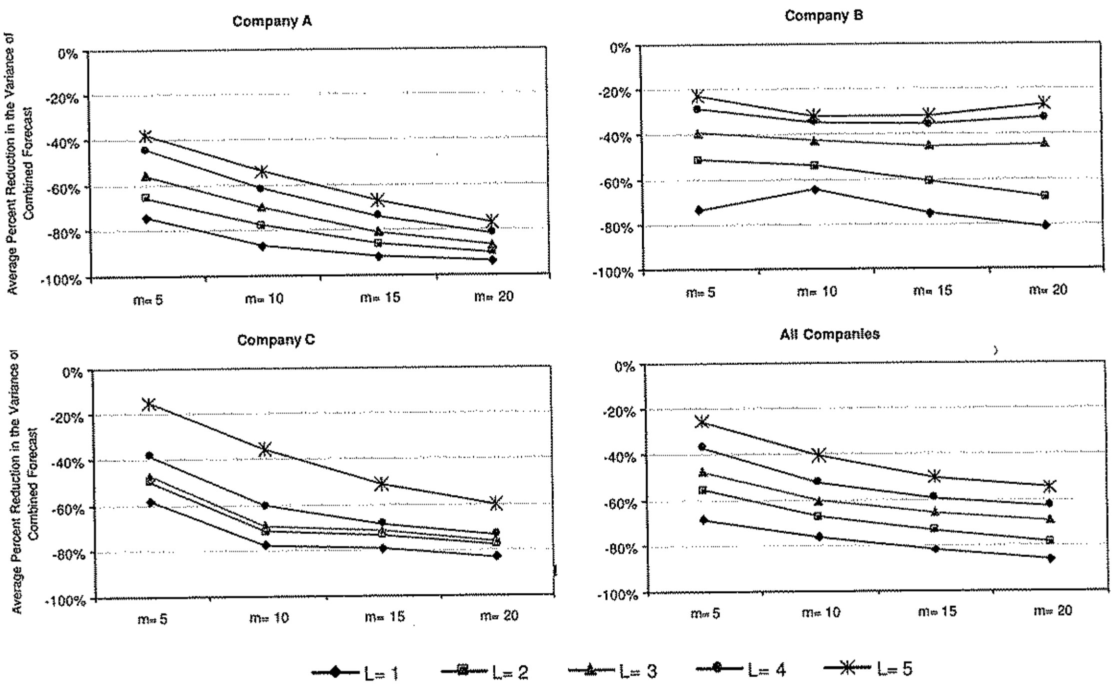


Figure 8: Percent reduction in the variance of the combined forecast with $m$ leading indicators, each of which provides $L$ periods of data.

Figure 8 displays the percent reduction in the variance of prior combined forecast under different combinations of $(m, L)$ across the data from all the companies. The value shown for each combination of $(m, L)$ is the average reduction over all the time horizons in the validation period. When all the leading indicators are unbiased, from Corrolary 3, it follows that the marginal value of the additional data (as measured by the reduction in the variance) diminishes when more leading indicators are considered. This is most apparent in the results for company $A$. However, the results for company $B$ follows a conflicting pattern, which is an indication of existence of bias in the estimates of leading indicators.

To measure if there is any systematic bias in the L-period data provided by a leading indicator, we use a test that evaluates the null hypothesis that the errors have zero mean, for which a t-test is appropriate. The test results indicate that all of the 20 leading indicators for company $A$ and company $C$ provide unbiased data at 0.01 level of significance. However, 8 of the 20 leading indicators for company $B$ provide biased data at 0.01 level of significance. This confirms that smaller reduction with $m$ is due to the existence of biased leading indicators. This reveals that there is a possible trade-off in increasing $m$. Leading indicators are selected based on their relative ranking in terms of the absolute value of correlation coefficient between their demand time-series data and the shifted time-series data of the product group. As $m$ increases, leading indicators with low absolute value of correlation coefficients are included. This increases the possibility of inclusion of biased leading indicators. Therefore, increasing $m$ has also a negative impact on the reduction in variance unlike the case where all leading indicators are unbiased.

Another notable effect that is apparent in the results is that the rate of change in the reduction of variance tends to be less sensitive to $m$ for small values of $L$, when all leading indicators are unbiased. The possible reason is that as $L$ decreases the approximation quality of the leading indicators are likely to be better and the variation between their estimates be smaller. As a result, the impact of $m$ is smaller. Overall, the results indicate that it is possible to obtain $20 - 80\%$ reduction with the additional data across the data sets of all the companies.

*Impact of Time-Lag of the Additional Information*

Figure 8 reveals that reduction in variance tends to be less for larger values of $L$ due to the relatively poor estimation quality of the leading indicators. In addition, in case of unbiased leading indicators, for large values of $m$, the rate of change is less sensitive to $L$. This may suggest the use of as many leading indicators as possible; however, due to the increasing possibility of having biased leading indicators, this is not a valid suggestion.

*Impact of Biasedness of Leading Indicators*

Leading indicators are unbiased estimators of the actual data over the estimation period $[1, T]$. However, it is possible that the data they provide over $[T+1, T+L]$ is biased. As stated earlier, 8 of the 20 leading indicators for company $B$ provide biased demand signals at 0.01 level of significance. The theoretical results imply that with biasedness the variance of the estimates increases and the mean of the estimates may shift depending on the relative biasedness of the leading indicators. To empirically support these claims, we replaced the 8 leading indicators diagnosed as biased with unbiased ones.

The average of the changes over all the periods in the validation period and over all the companies are summarized in Table 1 for all combinations of $(m, L)$. For instance, the variance with biased leading indicators is 0.43% more when 20 leading indicators with a time lag 5 are considered. Also, the mean estimate of the combined forecast with biased leading indicators is 0.003 less than the mean estimate with all the leading indicators unbiased. This corresponds to around 0.5% change in the mean estimate. Overall, the change in percent variance reduction is less than 5%, and the change in the mean is less than 0.012, which is about 2.5% of the mean with unbiased leading indicators.

| | Change in % Variance Reduction | | | | Change in Mean | | | |
|---|---|---|---|---|---|---|---|---|
| | m=5 | m=10 | m=15 | m=20 | m=5 | m=10 | m=15 | m=20 |
| L= 1 | -0.15% | 0.42% | 0.16% | -0.02% | -0.005 | -0.002 | -0.003 | -0.003 |
| L= 2 | 2.31% | 2.68% | 0.79% | 0.58% | -0.008 | -0.010 | -0.005 | -0.005 |
| L= 3 | -1.32% | 1.87% | 0.39% | 0.67% | -0.006 | -0.007 | -0.004 | -0.004 |
| L= 4 | 4.19% | 3.68% | 1.16% | 0.47% | -0.012 | -0.009 | -0.005 | -0.003 |
| L= 5 | 1.34% | 2.91% | 1.04% | 0.43% | -0.006 | -0.008 | -0.005 | -0.003 |

Table 1: Impact of biasedness of leading indicators on the percent reduction in variance and on the mean estimate of the combined forecast (in terms of deviation from the estimates with unbiased leading indicators at 0.01 level of significance).

In general, there is no discernable pattern in the change in both variance and mean forecast with respect to $L$. In terms of $m$, we observe that the changes with biased leading indicators tend to be smaller when number of leading indicators is 15 or 20 when compared to 5 and 10. That is, the impact of biasedness is less with more leading indicators. There are a few negative signs in the change in variance reduction which indicate the contradictory result that there is smaller reduction in variance with unbiased leading indicators. Since there is no exact method to diagnose the biasedness, the possible reason is related to the method we choose the unbiased leading indicators.

# 7 Conclusions

In this paper, we propose a demand characterization approach for technology products that have short life cycles and high demand volatility. Our approach is based on life-cycle growth modeling and Bayesian forecasting. We characterize the life-cycle patterns of the products with technological growth models in order to capture the volatility in demand due to the short life-cycle nature of the products. With Bayesian forecasting, we quantify the uncertainty associated with the forecasts and produce distributional estimates for the random demand. Our proposed approach uses leading indicator products to learn from the future. Bayesian forecasting updates the distributional forecasts with the information obtained from leading indicators, and hence captures the volatility in demand due to the market conditions. From the distributional forecast of the combined estimate of the life-cycle models, we develop demand scenarios.

Given the observed data set and several growth models characterizing the technology life cycle, the approach allows us to develop a streamlined way to modeling demand scenarios for a particular market segment in a technology driven market. The scenarios can be integrated into the supply-demand planning decision systems of the companies. Inclusion of the leading indicator demand reduces the variability in future demand scenarios. This, in turn, has a potential to improve the decision-making activities of the companies in terms of both computational time and the expected total operational costs.

One drawback of this approach is, however, that leading indicators might provide biased information. In order to alleviate this affect, we use several leading indicators, the biasedness of each can cancel out, but at the expense of increased variance of the estimates.

Computational testing on real-world data provided by three semiconductor manufacturing companies suggests that the proposed approach is effective in capturing the short life-cycle nature of the products and early demand signals, and capable of reducing the uncertainty in the demand forecasts by more than 20%.

# Acknowledgement

Our research is supported by the Semiconductor Research Corporation (SRC) grant 2004-OJ-1223.

## References

J. M. Bates and W. J. Granger. The combination of forecasts. *Operational Research Quarterly*, 20(4):451–468, 1969.

J. P. Dickinson. Some statistical results in the combination of forecasts. *Operational Research Quarterly*, 24(2):253–260, 1973.

R. M. Heeler and T. P. Hustad. Problems in predicting new product growth for consumer durables. *Management Science*, 26(10):1007–1020, 1980.

A. A. Kurawarwala and H. Matsuo. Product growth models for medium-term forecasting of short life cycle products. *Technological Forecasting and Social Change*, 57:169–196, 1998.

V. Mahajan, E. Muller, and F.M. Bass. New product diffusion models in marketing: A review and directions for future research. *Journal of Marketing*, 54:1–26, 1990.

N. Meade and T. Islam. Prediction intervals for growth curve forecasts. *Journal of Forecasting*, 14:413–430, 1995.

N. Meade and T. Islam. Technological forecasting– model selection, model stability, and combining models. *Management Science*, 44(8):1115–1130, 1998.

M. J. Meixell and S.D. Wu. Scenario analysis of demand in a technology market using leading indicators. *IEEE Transactions on Semiconductor Manufacturing*, 14:65–75, 2001.

S. J. Press. *Subjective and Objective Bayesian Statistics: Principles, Models, and Applications*. J. Wiley and Sons, Inc., New Jersey, 2003.

V. Srinivasan and C. H. Mason. Nonlinear least squares estimation of new product diffusion models. *Marketing Science*, 5(2):169–178, 1986.

D. Tigert and B. Farivar. The bass new product growth model: A sensitivity analysis for a high technology product. *Journal of Marketing*, 45:81–90, 1981.

S. D. Wu, B. Aytac, R. T. Berger, and C. A. Armbruster. Managing short life-cycle technology products for agere systems. *Interfaces*, 36:234–247, 2006.