Strategic Uses of Machine Data

Timothy M Burke, PhD (Arch Systems Inc.)
Palo Alto, California

Originally presented at IPC-APEX Expo March 6-12, 2021

Abstract

EMS factories have collected and used machine data for many decades. Over that time, much of the value derived from machine data collection has come from three operational use cases: allowing fewer operators to simultaneously monitor more machines for errors, reducing common operational mistakes through programmatic interlocks, and maintaining traceability records in case of product recalls. There has been significantly less use of machine data for strategic optimization of factory operations, with the notable exception of asset utilization monitoring using simple calculations like Line Utilization and OEE. One of the historical reasons for the absence of large-scale analysis of machine data in the EMS industry has been that it was difficult to interpret machine data absent external context on what intended operation was being performed when the data was collected. More recently however, the advent of big data analysis techniques and machine learning algorithms has largely removed this traditional limitation. In this paper we discuss the difference between tactical and strategic data analysis approaches to the common EMS factory goals of lowering component attrition and increasing line utilization. We show how machine data can provide significant value at the strategic level if it is stored and analyzed in granular detail instead of being pre-aggregated into high level key performance indicators before being analyzed. As the EMS industry looks forward to Industry 4.0, we argue that one of the biggest areas of efficiency gain may come from such strategic data analysis.

Machine data can provide significant value at the strategic level if it is stored and analyzed in granular detail.

Introduction

Modern EMS factories are filled with machines that produce rich data about their activities. The exact data generated varies between machine type and vendor, but typically includes detailed information on operations performed, faults encountered during processing, material consumed, products produced, and recipes activated. This machine data is used for a wide variety of purposes inside of the factory.

One of these purposes is product traceability. The complexity of electronic products has long demanded that EMS manufacturers keep records of exactly which components were placed into each product. Depending on the industry and application, these records may be highly detailed and have retention periods measured in years or even decades. The general goal is to be able to react to a putative future report of a product failure or a known bad component lot and determine the minimal set of other products that may be affected and should be recalled or carefully inspected.

> Operational efficiency means increasing the number of defect-free products that can be built in a given time period with the same number of machines and workers.

For the purposes of this paper, we will not discuss product traceability any further. Instead, we will touch on the other uses of machine data inside the factory to improve operational efficiency. In this context, operational efficiency means increasing the number of defect-free products that can be built in a given time period with the same number of machines and workers.

Broadly speaking, there are three standard ways that machine data has been used to improve operational efficiency in an EMS factory:

1. **Realtime Operational Monitoring:** Machine performance data is collected and displayed at the line, area, or factory level in the form of monitors or dashboards. This can increase labor efficiency by allowing a smaller number of operators to monitor a larger number of machines. At the most basic level, it could consist of just allowing a remote operator to see the same display or Human Machine Interface (HMI) that would have otherwise been shown on the machine itself but without needing to walk to the machine. This permits an operator to sit in a centralized control room and observe the operation of

machines that are spread out across many lines or buildings. When a problem occurs, the operator may still need to walk to the machine to resolve the fault, but at least they no longer have to stand in front of the machine to know if there is a problem. More advanced versions of real-time monitoring include alert dispatching systems and interactive dashboards to analyze the current performance of different factory lines and compare them against standard targets. Real-time monitoring increases operational efficiency by reducing the time between a fault occurence and when it is resolved. Generally speaking, though, real-time monitoring does not decrease the total number of faults, just the mean time to resolve them.

2. **Process Interlocking:** As factories have begun to produce increasingly diverse mixes of products with high recipe complexity, there has been a need to ensure that normal, human setup mistakes are minimized. For example, if a circuit board requires 300 separate component feeders to be mounted in a pick-and-place machine and there is a 1 in 1,000 chance that any given component feeder is incorrectly installed, there is only a 74% joint probability that all 300 components are loaded correctly. The product can only be built successfully when all 300 component feeders are correctly installed. Even worse, if an incorrect setup is not detected and used to build products, those products will need to be scrapped: wasting material, production time and money. Hence, there is a need to reduce the odds of common mistakes so the overall probability that everything is set up correctly to build a product is closer to 100%. One of the ways to achieve this is to have the machine ask an external factory system to validate operator actions before the machine will accept them. For example, a pick-and-place machine could check the serial number of each component feeder and make sure it contains the expected component before starting production. Similarly, a machine could check a serial number printed on each circuit board before placing components on that circuit to ensure that it belongs to the expected work order for that factory line. There are many such examples of process interlocking, both simple and complex, that combine machine data with business rules defining allowed and disallowed conditions to reduce the chance of inadvertent mistakes.

3. **Historical Performance Tracking:** Factories strive to improve their operational efficiency over time by tracking key performance indicators whose numerical values capture some aspect of the factory's operations. One common set of such metrics for EMS factories is OEE (Overall Equipment Efficiency). OEE refers to a family of related statistics that attempt to quantify

what fraction of a factory's installed production capacity is being utilized. A 100% OEE score would mean that the factory was producing defect-free products as fast as physically possible whenever the factory intended to produce products. While OEE can be calculated manually, it is often computed using machine data to automatically capture the performance of each machine. Other statistics, such as first-pass yield or material attrition rates, can similarly be used to capture different aspects of factory operations and ensure that they either improve over time or, at least, remain constant.

Since factories are large, complex operations, there is a desire to simplify this complexity down to a few numbers that allow humans to quickly understand and compare overall performance without needing to reason about large sets of complex parameters.

Since factories are large, complex operations, there is a desire to simplify this complexity down to a few numbers that allow humans to quickly understand and compare overall performance.

The purpose of this paper is to discuss a fourth category of machine data collection: advanced machine data analytics. Like historical performance tracking, advanced machine data analytics, hereafter just called **advanced analytics** looks at how machine data values change over time in order to infer where problems are in the factory and which problems are most impactful to solve.

In this way, advanced analytics is similar to real-time operational monitoring in that its purpose is to uncover and resolve operational practices or situations that lead to inefficient factory operations. However, unlike real-time operational monitoring, its goal is not to find and quickly resolve minute-by-minute issues that occur and disappear rapidly, but rather to identify larger structural problems that reduce efficiency day-after-day or month-after-month.

The goal of real-time operational monitoring is to identify larger structural problems that reduce efficiency day-after-day or month-after-month.

A particular example of advanced analytics is predictive maintenance. Whereas real-time monitoring would accept that machines have a certain fault rate and endeavor to clear each fault as quickly as possible without necessarily reducing the number of faults, predictive maintenance complements this approach by predicting machine faults in advance and taking corrective action, such as recalibration, before production is impacted. Combining strategic (advanced analytics) and tactical (real-time monitoring) systems can reduce both the number of production faults and their mean time to correction.

The distinction between historical performance tracking and advanced analytics is that historical performance tracking visualizes and records coarse-grained, typically preaggregated data such as “OEE of each line per week”, whereas advanced analytics stores and analyzes as rich of a machine dataset as possible. For example, looking at component attrition, a historical performance tracking approach may store and compare the daily per-factory total attrition rate. An advanced analytics approach would instead store each individual mispick down to the feeder level along with the recipe that was active at the time and what the detailed error was. Obviously, you can recreate the historical performance tracking dataset from the advanced analytics dataset by aggregating the data appropriately, but it is generally not possible to go the other direction as there is significantly less detail inside of the historical performance tracking dataset.

One of the goals of this paper is to demonstrate that there is additional value in collecting and storing rich machine data for historical analysis because modern data analysis tools and methodologies have now made it possible to both efficiently collect and analyze this data. In additional to producing high-level performance tracking KPIs, the underlying rich machine data also more directly enables corrective action by pointing to specific problems that can be resolved to increase the performance metric. Typically, the high-level performance KPIs do not have the granular detail needed to understand why they are trending up or down, or what actions could be taken to improve them.

The following sections discuss the difference between these four approaches to using machine data in the context of two common EMS factory scenarios: reducing component attrition and increasing line utilization.

Component Attrition

The issue of material attrition is present in all factories, including EMS factories. For the purposes of this paper, material attrition is restricted to wasted electronic components loaded into a pick-and-place machine but not successfully placed onto a circuit board because of a mispick error during placement. It is generally not possible to achieve zero-component attrition, so the operational goal instead is to maintain attrition below a fixed level, such as 0.1% or 0.5%. An attrition rate of 0.1% would mean that one component was lost for every 999 components that were successfully placed onto a circuit board.

Before discussing how machine data can reduce material attrition, we first need to describe the component placement process. Typically, electronic components are loaded into a component feeder, which is attached to the pick-and-place machine. The component feeder is an electromechanical device that presents components one at a time to a placement head for placement onto a circuit board from a carrier such as a reel or tray. Component feeders are durable machines that are typically individually serialized and require periodic maintenance and calibration.

Realtime Monitoring

The first way that machine data is used to reduce component attrition is by preventing scenarios where a misconfigured machine wastes several components at high speed before an operator can intervene. To prevent this, the machine can automatically stop if a given feeder mispicks more than N parts in a row. It is important to note though that there is a tradeoff involved in setting the limit N. Stopping the machine halts production on the factory line, which leads to a decrease in overall line performance. If the mispick limit is set too low, for example, stopping on every mispick, it could quickly trade one problem (elevated component attrition) for a worse problem (frequent line stoppage causing unacceptable performance slowdown). This tradeoff means there is an optimal choice of N that detects abnormal mispick sequences but does not have a high false positive rate.

Feeder Interlocking

Feeders are precision electromechanical machines that require periodic maintenance and calibration. Over time, an uncalibrated feeder is expected to gradually experience an increased mispick rate as its internal mechanisms slowly drift or desynchronize. Standard feeder maintenance programs address this problem by recalibrating each feeder periodically, such as every N days.

If a feeder is due for preventative maintenance, and hence would be expected to have an elevated mispick rate until such maintenance is performed, it can be blocked programmatically from being used to prevent an operator from skipping maintenance or as a signal to the factory operator that calibration must be performed. Feeder interlocks typically work by sending the barcode or serial number of a feeder to an external system at the time when it is attached to the pick-and-place machine. The external system can respond and say the feeder is blocked, which will halt the machine until the feeder is replaced.

Just like realtime monitoring, however, there is also a tradeoff involved in aggressively using feeder interlocks to reduce component attrition. Swapping an interlocked feeder that has been delivered to the factory line involves finding an alternative feeder, perhaps in a separate warehouse or storage area, then unloading the interlocked feeder and reloading the replacement feeder. This is a relatively slow process, especially considering that the interlocking happens during product changeovers, which are time-sensitive line-downtime events. If interlocks are used too aggressively, there could be a net negative impact on the overall factory performance by increasing average changeover times to reduce component attrition.

Historical Performance Tracking

> There is again a tradeoff here where the attrition rate needs to be high enough that it merits allocating limited expert resources to bring it back under control.

Generally speaking, there is not as much use for historical performance tracking to prevent component attrition other than if area or factory level attrition rates increase too much, such that it may trigger a more detailed review of feeder usage by an advanced manufacturing team until the rate lowers again. Since there are typically many more factory issues to resolve than advanced manufacturing experts to diagnose and solve them, there is again a tradeoff here where the attrition rate needs to be high enough that it merits allocating limited expert resources to bring it back under control.

Advanced Feeder Analytics

The previous sections demonstrated how the three standard uses of machine data together help control the component attrition rate in an EMS factory to an acceptable level. However, each of the three strategies comes with an associated tradeoff that prevents them from being used to prevent all mispicks:

- Realtime monitoring trades line stoppages for preventing excessive mispick rates.
- Feeder interlocking trades longer changeovers for preventing the use of feeders likely to have elevated mispick rates because they are due for maintenance.
- Historical attrition rate tracking ensures attrition rates never remain elevated for long periods of time, but requires the investment of finite engineering resources so the rates must be high enough, and for long enough, to justify assigning a team to resolve the issue.

If all component feeders were indistinguishable and had identical mispick rates, then the above three approaches would be optimal and there would not be additional value in strategic feeder monitoring or advanced analytics. However, it is generally not the case that all feeders perform and age identically. The follow sections show how advanced feeder analytics can improve on the attrition limits achievable by the three standard machine data uses by identifying with increasing specificity which exact feeders should be targeted for maintenance.

Usage Based Maintenance

The first way to decrease component attrition using advanced analytics is to recognize that different feeders will feed parts at different rates based on the exact component loaded into the feeder, and whether a given product needs many of those components or few of them. This difference means that feeders will accumulate different usage patterns over time. Like other electromechanical devices, one would expect that feeders with higher usage need more frequent maintenance. By tracking the number of placements performed on each feeder, you can change from time-based maintenance (every N days for all feeders) to usage-based maintenance (every P placements for all feeders).

The net impact of moving from time-based to usage-based preventative maintenance is typically a reduction in the number of maintenance events needed to achieve a given target mispick rate across the population of feeders since lightly used feeders are maintained less frequently than heavily used feeders.

Feeder Specific Analysis

The previous feeder maintenance programs have all assumed that feeders were indistinguishable objects drawn from a common population and it was not possible to track them individually. This means that time or usage-based maintenance rules need to be set for the "average feeder" and applied uniformly to all feeders.

If all feeders aged and performed identically, then this approach would be optimal. However, feeders, like other machines, are expected to follow a Pareto distribution instead, where a small number of feeders account for a large fraction of total mispicks. If the problematic hotspot feeders can be identified via advanced analysis techniques and taken out of service then the overall distribution of mispick rates shifts toward lower mispick rates and the attrition rate of a factory can be reduced below what is possible with the other approaches alone. This is shown schematically in Figure 1.

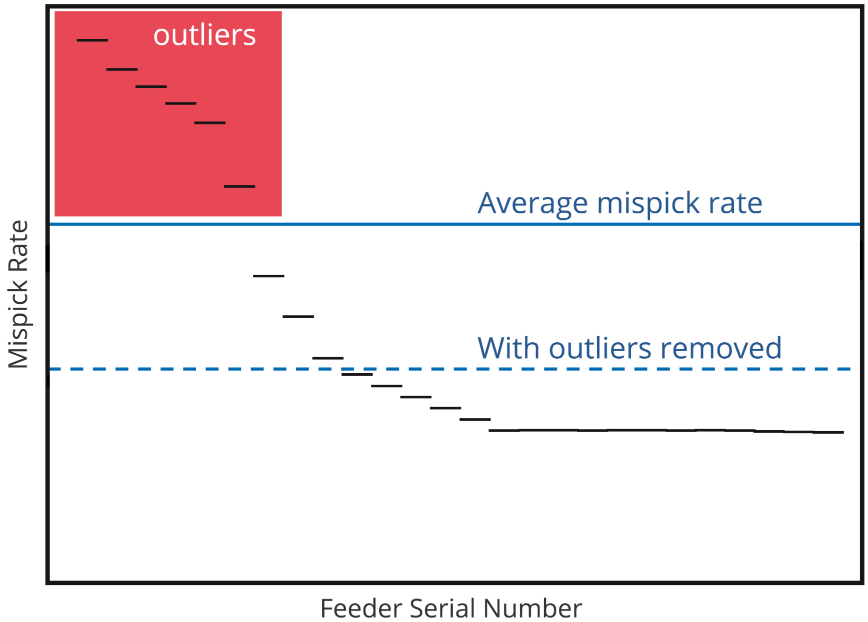

Figure 1: Achievable average mispick rate with and without high outliers identified and removed via advanced feeder analytics.

The unbalanced distribution of feeder errors, where a small number of feeders disproportionately increase the overall population or global mispick rate, means that it is possible to significantly improve on the attrition rate bounds possible with real-time monitoring, usage-based maintenance and interlocking alone.

Condition-based Maintenance

The simplest advanced feeder analytics system would be to implement a condition-based maintenance approach. Condition-based maintenance means that feeders are selected for maintenance based on their observed current performance. For example, you could define a monitoring window of 1 week. Every week, all feeders are ordered by the total number of mispicks and the worst 1% are selected for maintenance. The next week, this is repeated with a new set of top feeders selected for maintenance.

Condition-based monitoring will improve over usage-based monitoring when the following two criteria are satisfied:

- The monitoring period needs to be shorter than the standard maintenance interval since the gains from implementing condition-based maintenance will come from recalibrating feeders that would otherwise continue to experience an elevated mispick rate until their next scheduled maintenance.
- It must generally be the case that feeders which experience elevated mispick rates in a given monitoring period continue to experience elevated mispick rates in the next monitoring period as well. If this is not the case, then selecting feeders for maintenance based on past performance data will not predict the feeders most likely to have elevated mispick rates in the next period and a more complex predictive maintenance system is needed to show improvements.

Provided the above two conditions are met, condition-based feeder maintenance provides a straightforward way to reduce component attrition rates, often significantly, without needing either complex predictive algorithms or real-time intervention. One of the key points to emphasize about this approach is that there is no requirement that problematic feeders be removed immediately. Indeed, the analysis proceeds on the timescales of days or weeks and the goal is to identify those feeders that should be removed from service over the next few days or weeks, not necessarily immediately.

Predictive Maintenance

In situations where it is justified, one can go further than condition-based feeder maintenance and implement a predictive maintenance program on top of the same historical machine data. Condition-based maintenance can be thought of as a very simple form of predictive maintenance where the prediction is that machines will continue to perform identically in the future to how they performed in the past. In general, this is usually not the case. For example, feeder mispick rates tend to increase over time as feeders age, not stay the same. However, it may still be the case that for the shorter periods in question, feeders can be assumed to generally perform identically from one monitoring period to the next and no additional complexity needs to be introduced.

When necessary, though, it is possible to substitute more complicated algorithms for predicting future feeder performance based on past feeder data. For example, regression, statistical inference, or even advanced techniques such as neural networks could be used to predict – not which feeders were the worst performers in the previous monitoring period (which has already passed), but which will become the worst performers in the next monitoring period.

Regression, statistical inference, or even advanced techniques such as neural networks could be used to predict – not which feeders were the worst performers in the previous monitoring period, but which will become the worst performers in the next monitoring period.

Line Utilization

In addition to minimizing component attrition, another operational need for an EMS factory is maximizing Line Utilization (LU), which can be defined simply as the number of components placed on each line in 24 hours divided by the rated placement capacity of the installed pick-and-place machines on that line. Unlike OEE, for example, LU makes no correction for planned downtime or factory closures. Maximizing LU is important because SMT machines are expensive assets and there is therefore a rational desire to maximize the return on that investment by producing as many products as possible over the operational lifetime of each machine.

Since LU in EMS factories is typically measured directly as the number of placements performed by the pick-and-place machines on each line, machine data collection from those machines can be used to calculate LU automatically.

While there generally is not a role for interlocking in increasing LU, there are roles for both real-time monitoring and historical performance tracking. Real-time monitoring can be used to increase LU by comparing hour-by-hour line performance against a standard target performance that is believed to be achievable and taking immediate corrective action whenever the real-time output of the line drops below the target.

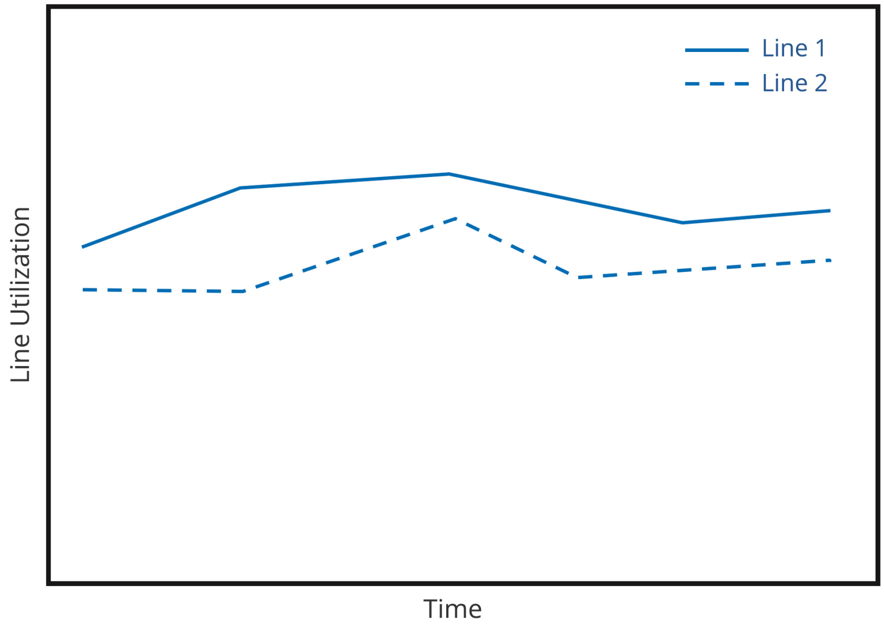

Figure 2: A typical plot of Line Utilization over time comparing 2 hypothetical lines in a factory. Note how difficult it is to interpret the root cause of the observed LU differences by looking at the plot.

Historical performance tracking looks for trends in LU over time by comparing different lines to identify opportunities for increasing LU by focusing finite engineering resources on the lowest performing lines. However, there is rarely enough information in the LU metric itself to directly enable corrective action. To understand why this is the case, a typical plot of Line Utilization versus time for two different factory lines is shown schematically in Figure 2.

From the figure, it's clear that Line 1 has higher LU than line 2, but other than that, there is not much else that can be learned from the figure. In particular, it cannot be inferred that Line 2 is being operated worse than Line 1 or which line is a better candidate for operational improvement. For example, it could be the case that Line 2 is having significant operational issues but it could also be the case that Line 2 is building a slow, complex product that requires slower machine cycles and it is actually being operated closer to its ideal throughput than Line 1 given the constraints of the products that each is making.

The same challenges occur when comparing other historical monitoring metrics between lines such as OEE, products produced, attrition rates, etc. Top-level monitoring KPIs are too coarse to capture any of the nuances about what each line is doing in order to provide insight into the reasons why one has higher LU than the other.

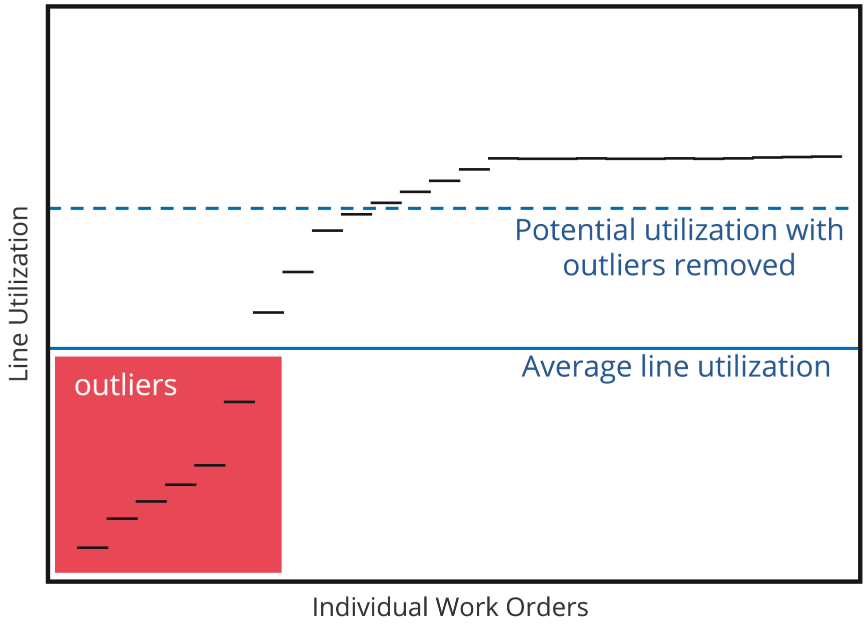

Figure 3: Disaggregated line utilization showing the individual performance of each work order follow a pareto-like distribution where specific, outlier work orders have a significant impact on the overall utilization of the line.

This is the same as the situation with historical performance tracking of component attrition rates. In one such case, we were able to uncover actionable metrics by disaggregating all of the feeders and tracking each one individually. In the case of

LU, the corresponding individual object to track is each work order on each line, where a work order defines a single contiguous period of time where N units of a single product are being made on a single line.

Figure 3 shows that tracking the individual performance of work orders on each line, instead of the aggregate impact of those work orders over time, provides an actionable path for increasing LU. Figure 3 is drawn in explicit parallel to Figure 1 to emphasize the point that both follow the same approach of serialized tracking of individual objects which are then analyzed using advanced techniques to determine which individual object should be investigated.

Once the outlier work orders and associated product recipes are identified using advance machine data analytics, there are several possible corrective steps that could be taken. For example:

- Ranking the work orders by LU loss and assigning an expert team to resolve the underlying problems.
- Performing further analysis on the lowest performing work orders by looking for commonalities in machine characteristics that distinguish low performing work orders from high performing work orders.

Conclusion

In this paper we have reviewed the common uses of machine data for operational improvement in EMS factories: real-time monitoring, performance tracking and process interlocking. We then discussed a fourth, emerging use of machine data for advanced analytics, which uses large, highly detailed datasets collected automatically from machines over long periods of time to surface specific opportunities for operational improvement that are not directly achievable using the other 3 methods. We showed two specific use cases for advanced analytics to reduce component attrition and improve LU. Looking forward to the continued adoption of Industry 4.0 applications in the EMS industry, we expect that advanced machine data analytics will become a standard, indispensable tool for factory operational excellence just like the other three uses of machine data.