Technical Disclosure Commons

Defensive Publications Series

November 2020

SYSTEM AND METHOD FOR PREDICTING HARDWARE FAILURES OF ELECTRONIC DEVICES USING ONBOARD SENSORS AND DEVICE LIFECYCLE DATA

Dmitry Goloubew

Shi-Jie Wen

James Turman

Sumit Parimoo

SYSTEM AND METHOD FOR PREDICTING HARDWARE FAILURES OF ELECTRONIC DEVICES USING ONBOARD SENSORS AND DEVICE LIFECYCLE DATA

AUTHORS:
Dmitry Goloubew
Shi-Jie Wen
James Turman
Sumit Parimoo

ABSTRACT

Techniques herein provide a capability to predict failures of hardware by using onboard sensors and provide for the ability to move from detection to prediction for hardware failures. In turn, such techniques can help to reduce downtime due to marginal hardware and improves network availability. The techniques can also help to reduce unnecessary maintenance and changes related to replacing hardware that has not failed, which can lead to business efficiencies for both customers and vendors and may become even more important in a Network-As-A-Service (NAAS) context in which both sides are paid by the vendor.

DETAILED DESCRIPTION

Every company operating a large fleet of electronic devices has to deal with hardware failures, repairs, and replacement. While such issues may be an expected part of doing business, some aspects of electronic device lifecycle can be very impactful to both vendor and customers businesses such as, for example:

- Some failures of marginal hardware are non-obvious. Devices may 'appear to work' such that even diagnostics are at times passing, but some device functions may be impaired. Such failures, at times, can lead to serious outages and, thus, every improvement in detecting and eliminating such failures is important.
- In some cases due to manufacturing processes, variations for a significant batch of devices may receive a component that does not have its corresponding designed life span. At the time of manufacturing, such devices and components may pass all tests but weeks/months afterward may fail in the field due to the

premature/rapid aging of the marginal component. On the technical side, this often leads to failures as described above. On the business side, this can lead to massive recall / replacement / rework campaigns when more parts have to be recalled from the field than necessary (due to an inability to find devices that have marginal components).

- For customer-managed hardware replacement, hardware that is not failing may be replaced, which means that the expense and operational risks of uninstalling, shipping, and inspecting hardware are unnecessary. The impacts of such replacements are shared by both customer and vendor.
- In a Network-As-A-Service (NAAS) context, all impacts from the aforementioned issues are transferred entirely on a vendor, thereby placing even more emphasis on balancing/preventing outages caused by hardware failures while replacing only the devices with failing hardware.

At a technical level, many failures of the type explained above are caused by semiconductor aging. There are 4 mechanisms related to semiconductor reliability, degradation, and failures, including Electromigration (EM), Time-Dependent Dielectric Breakdown (TDDB), Hot Carrier Injection (HCI), and Negative Bias Temperature Instability (NBTI). Modern semiconductors are designed taking into account the understanding of these mechanisms.

Correctly manufactured components typically have a life expectation spanning over a decade. Should, however, some undetected process variations occur during component manufacturing, such a component may still be functional and pass functional tests, but have much shorter lifespan (e.g., often weeks/months instead of decades). Because high-level functional tests are passed, it is difficult to detect the presence of such marginal components before they fail.

Presented herein are techniques to address such issues by providing a system to predict semiconductor (transistor-level) failures by using data from embedded sensors. This system is focusing on intrinsic sensors. Examples of such sensors may include frequency and voltage sensors. Through continuous monitoring, it is possible to detect the onset of rapid aging when device is still functional, but is moving fast towards a failure.

The predictive nature of such detection allows proactive action to be taken, avoiding possible downtime or outage.

Previously, such an approach was not practical as developing a monitoring system for each component individually was uneconomical. Techniques herein address this issue by building a general pipeline that adapts to individual characteristics of the components using machine learning.

During operation utilizing the system of this proposal, sensor data is combined with 'healthy population' data collected during the manufacturing stage of device lifecycle and 'aged/weak population' data from stress-testing and other sources.

Combining data collection using telemetry and processing using machine learning with healthy and weak population examples provides for the ability to follow the health of individual components. As representations of specific components may shift to a low-density space or to a vicinity of weak population, the system can calculate risks for components involved. Figure 1, below, provides a high-level overview illustrating the core concepts of the system.

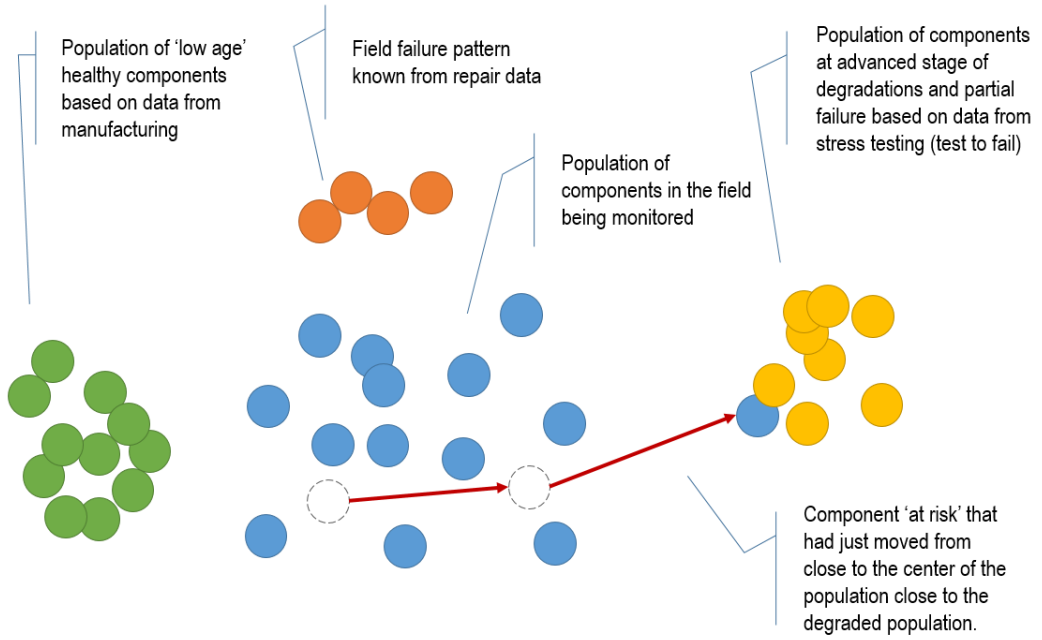

Figure 1: System Concepts

Utilizing the system, risk of failure is evaluated as proximity changes to weak or known failure pattern groups. The representation shown in Figure 1, above, is a product

of combining an unsupervised processing pipeline creating representations and use of known good and known degraded populations as the supervision signal.

As illustrated in the following Figures, the system of this proposal consists of 3 parts: training, accelerated inference, and automated analysis. Figure 2, below, illustrates example details associated with system training operations that may be performed.

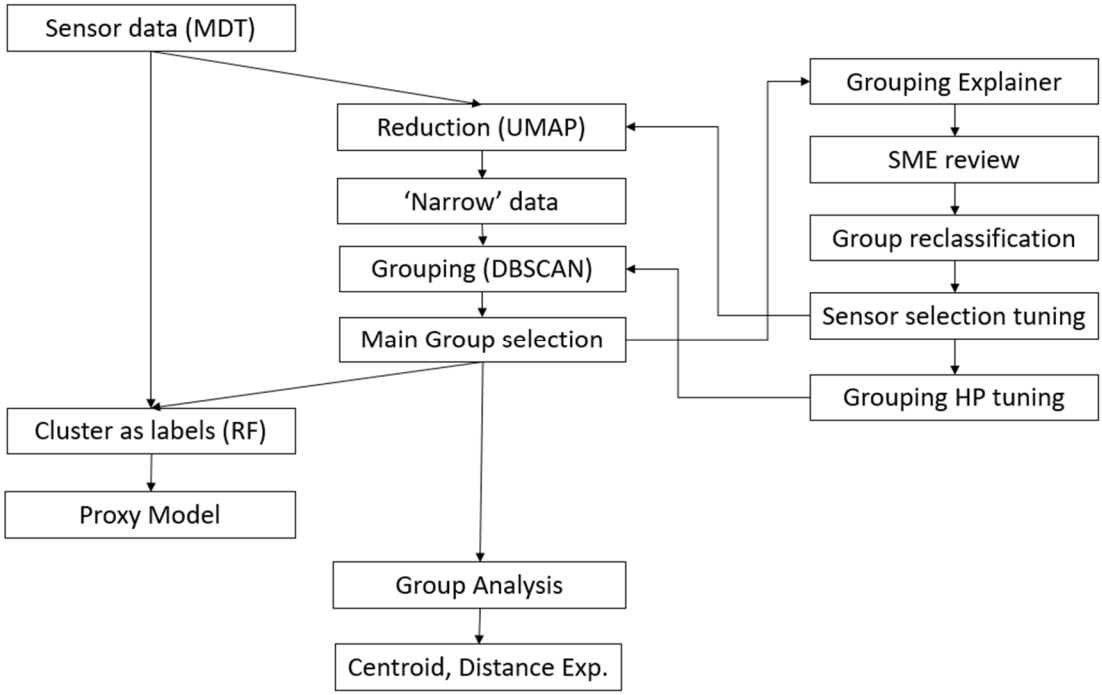

Figure 2: System Training Operations

Training operations may utilize a Uniform Manifold Approximation and Projection (UMAP) algorithm for dimensionality reduction from 'wide' space (equal to number of sensors) to a narrow space - which can be quickly examined visually. Subsequently low-dimensional representation is segmented utilizing a Density-based Spatial Clustering of Applications with Noise (DBSCAN) algorithm with silhouette analysis driving the hyper-parameters of the DBSCAN algorithm.

Implementations may utilize similar algorithms such as Hierarchical DBSCAN (H-DBSCAN) or Ordering Points To Identify the Clustering Structure (OPTICS), hierarchical clustering, and/or the like. As dimensionality reduction is computationally intensive, a proxy model (e.g., utilizing a Random Forest algorithm) can be utilized to allow fast

inference, which permits making predictions using very little resources and, thus, allows the system to scale. It should be noted that alternative algorithms can be utilized for different implementations such as Gradient Boosting Machine (GBM), neural networks, and/or the like. The probability of errors introduced by bypassing dimensionality reduction is compensated by batch-reconfirming predictions on a periodic basis, as discussed in further detail below with reference to periodic analysis operations that may be performed by the system.

Figure 3, below, illustrates example details associated with accelerated inference operations that may be performed.

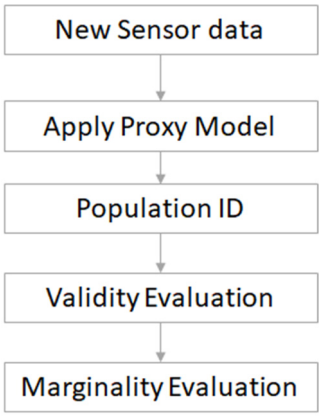

Figure 3: Accelerated Inference Operations

The inference step, as illustrated in Figure 3, above, applies the proxy model to a data point (e.g., a sensor reading) to establish whether the given data point belongs to common population. Additionally, distances to centers of interest (e.g., healthy, aged, known issues, etc.) are evaluated. Together with prior history, this provides for the ability to produce a marginality score for new data points very quickly. Predictions are validated versus pre-calculated distance distributions for validity. For example, if distances fall out of the expected distributions, a prediction is discarded and for batch workflows a valid prediction will be made during periodic workflow. For flows requiring near-real time prediction, a data point prediction is recalculated using the full workflow, as illustrated above via the training operations.

Figure 4, below, illustrates example details associated with periodic analysis operations that may be performed.

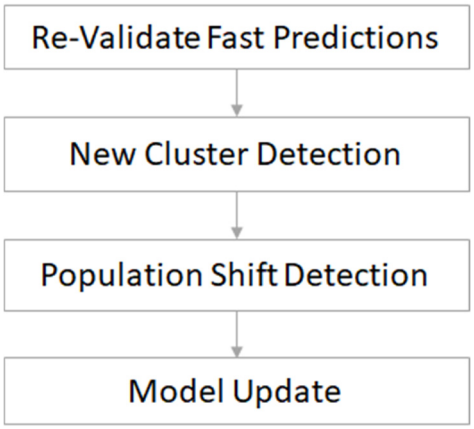

Figure 4: Periodic Analysis Operations

Periodic analysis may serve several purposes. For example, periodic analysis may provide for the ability to revalidate predictions in a batch, thus detecting/correcting any potential imprecision introduced by proxy modelling. However, a main purpose of periodic analysis may be to detect large changes that are not specifically related to a particular device/component, but rather to a population as a whole. Such issues can be detected by spotting formations of new clusters in a low-dimensional representation. Even earlier dynamics can be seen by observing the shifts within a cluster such as, for example, when a subset of devices starts to move toward the fringes of a cluster.

Thus, provided herein are a system and techniques to utilize intrinsic sensor readings for evaluation of marginality and risk using the representation of multidimensional sensor data. Data from various stages of device lifecycle such as manufacturing and accelerated stress-testing can be utilized to complement representations produced through unsupervised and self-supervised techniques and enable predictions for products before instances of products reach old age.

Additionally, techniques herein may provide for establishing hyper-parameters for grouping (e.g., utilizing DBSCAN or other density-based algorithm) using silhouette analysis and may also facilitate automated centroid splitting for complex shaped clusters to balance accuracy and computational load during inference. Further, techniques herein may provide cluster membership explanations using a Shapley algorithm for analysis and interpretation. Further, techniques herein may facilitate the acceleration of inference utilizing proxy modelling such that inference can be performed in very resource-lean manner, for example, on a device, and also later through slow-batch confirmation.

Moreover, techniques herein may facilitate population-level discovery via new cluster formation detection and population shift detection.

In summary, techniques herein provide a capability to predict failures of hardware by using onboard sensors and provide for the ability to move from detection to prediction for hardware failures. In turn, this helps to reduce downtime due to marginal hardware and improves network availability. The techniques also help to reduce unnecessary maintenance and changes related to replacing hardware that has not failed, which can lead to business efficiencies for both customers and vendors.