ORACLE

ORACLE

# Oracle Cloud Advanced ML Prognostics Innovations for Enterprise Computing Servers

## USPTO Tech Fair, May 19, 2022

**Guang Wang**

Principal Machine Learning Research Scientist, Oracle

 

## Safe Harbor

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Statements in this presentation relating to Oracle's future plans, expectations, beliefs, intentions and prospects are "forward-looking statements" and are subject to material risks and uncertainties. A detailed discussion of these factors and other risks that affect our business is contained in Oracle's Securities and Exchange Commission (SEC) filings, including our most recent reports on Form 10-K and Form 10-Q under the heading "Risk Factors." These filings are available on the SEC's website or on Oracle's website at http://www.oracle.com/investor. All information in this presentation is current as of September 2019 and Oracle undertakes no duty to update any statement in light of new information or future events.


# Oracle Advanced Machine Learning Prognostic Product Suite

*Background*

# Oracle AI/ML Platform for Prognostic Solutions

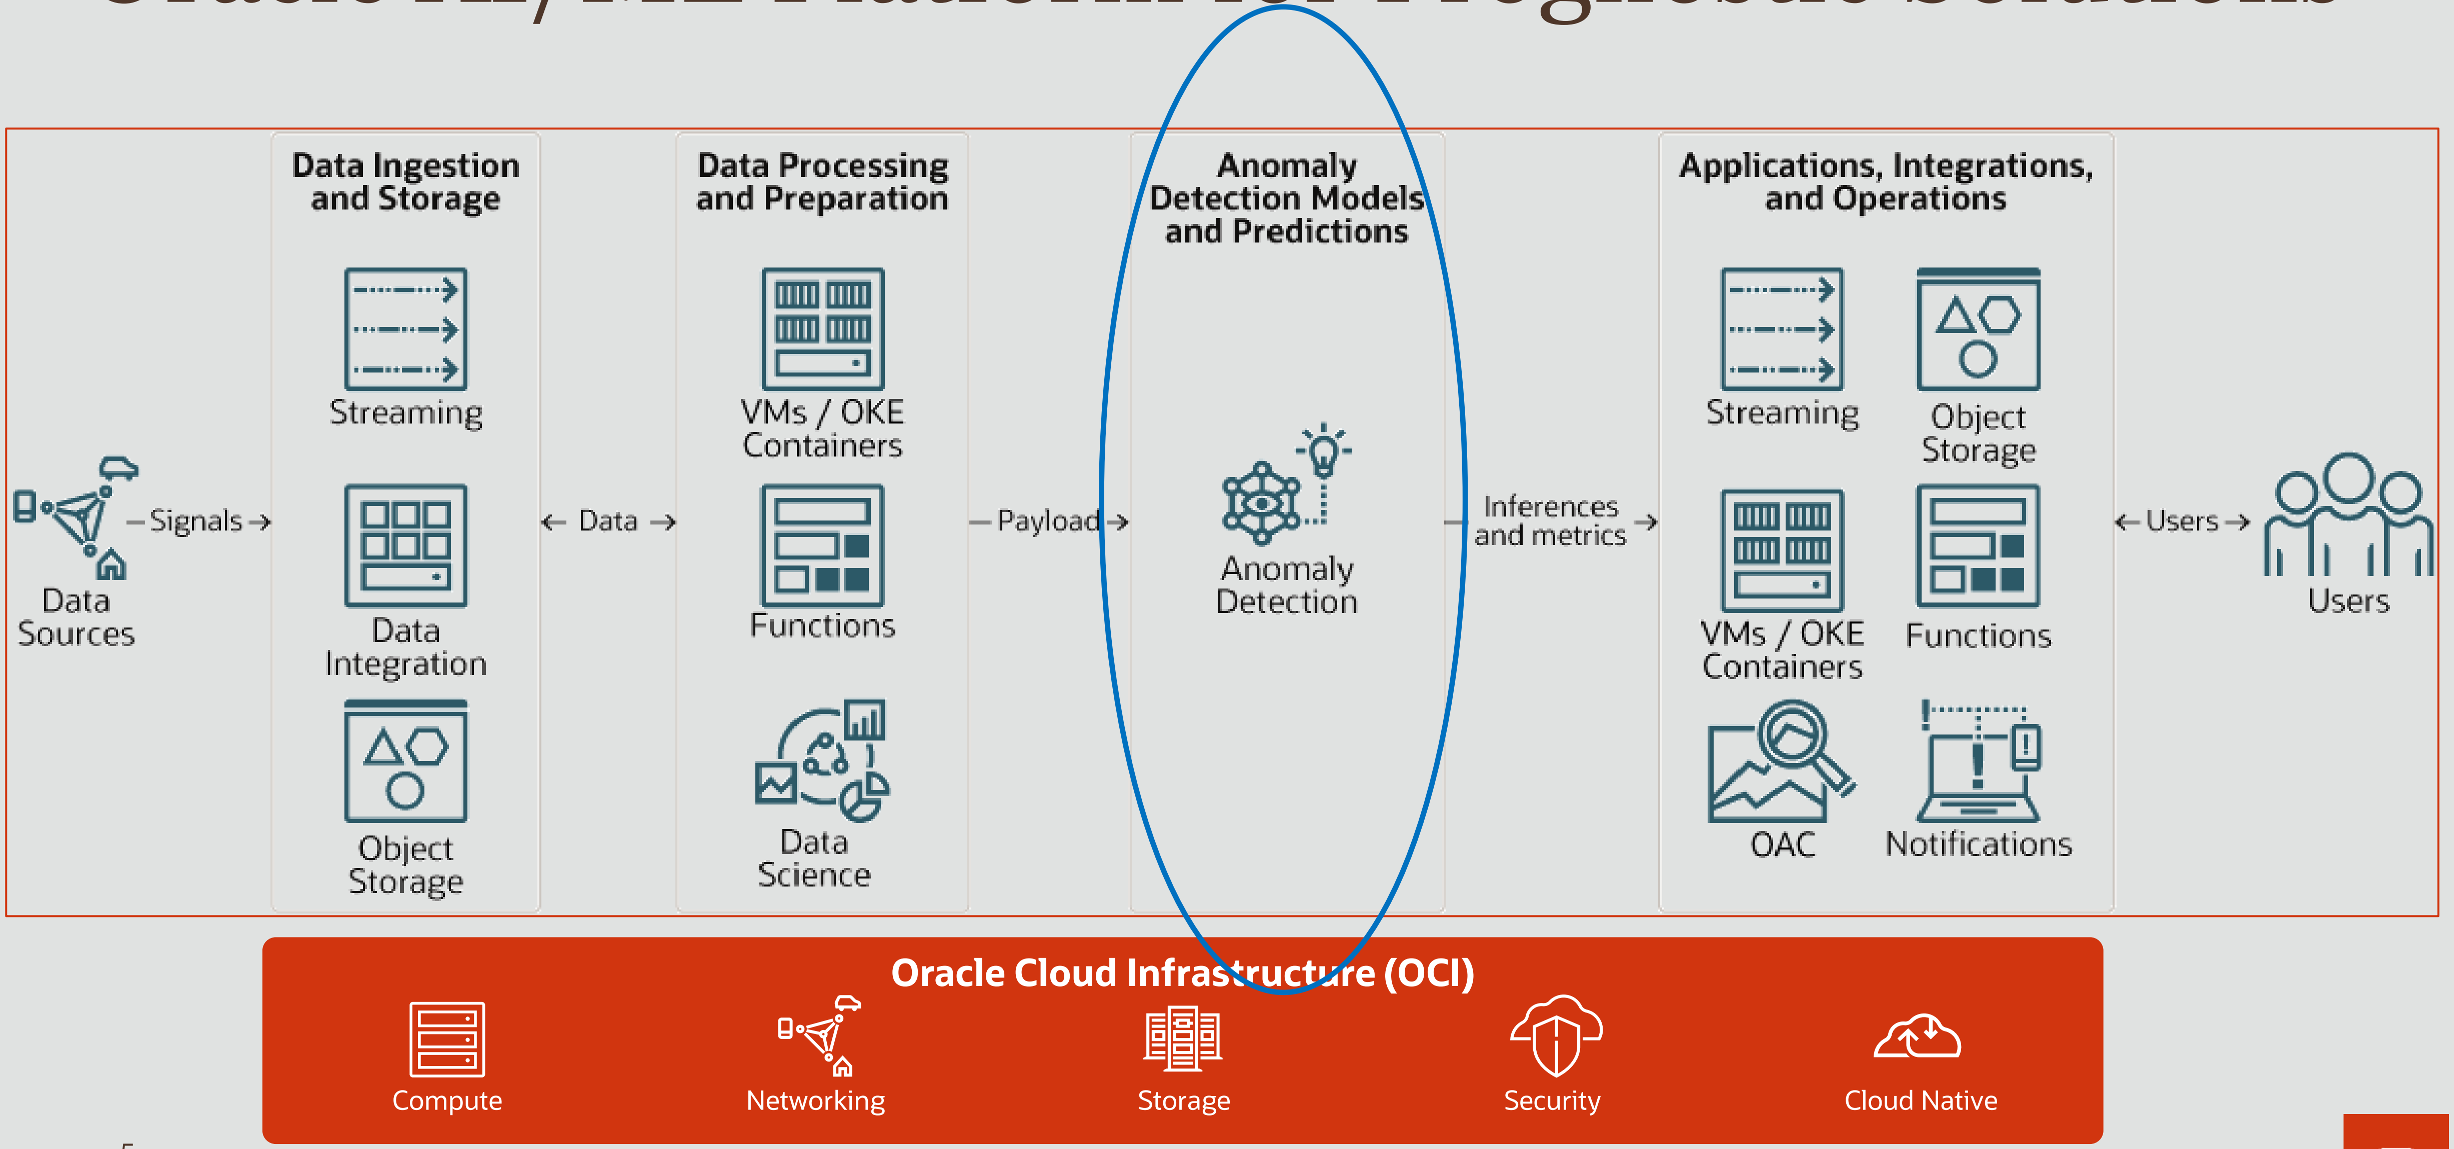

# Telemetry Signals in Enterprise Servers

Telemetry signals include many physical variables:

- voltages, currents, temperatures, fan speeds, and power levels

Telemetry signals correlate with system IO traffic, memory utilization, and system core utilizations

Prevalence of telemetry signals is growing exponentially

- One Oracle M6 server has 3400 sensors [same as a 1000 MW nuclear plant]
- One medium size data center now has 1M sensors

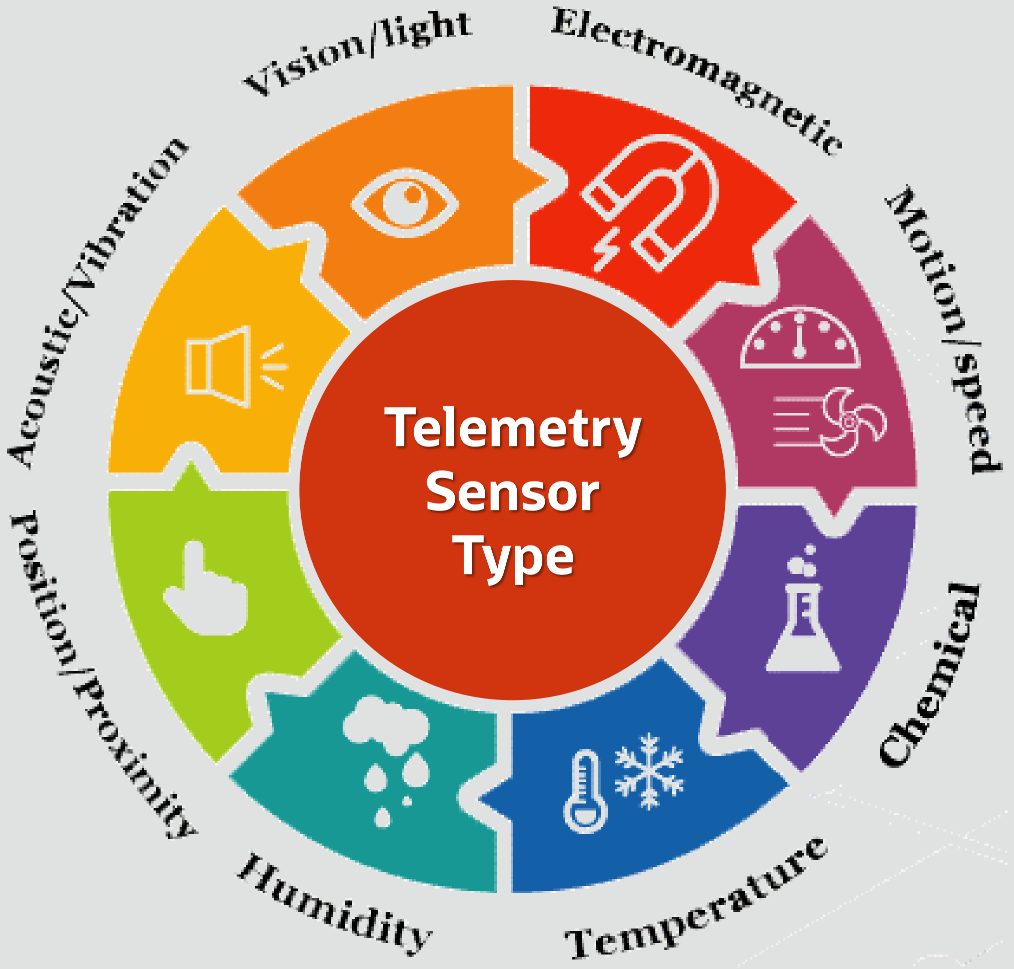

# Multivariate State Estimation Technique (MSET)

## Advanced Pattern Recognition Algorithm Using Telemetry Signals

- Nonlinear, nonparametric machine learning method for prognostic anomaly detection
- High sensitivity for detecting subtle anomalies in noisy or even chaotic time series metrics
- Ultra-low false-alarm and missed-alarm probabilities
- Ideal candidate ML algorithm for dense-sensor IoT applications

# Prognostic Capability of MSET

- Creates a dynamic band around each sensor value
- Correlates all sensors simultaneously
- Earlier warning of potential sensor failures than thresholding based conventional ML solutions
- Allows users to specify both false alarm and missed alarm probabilities

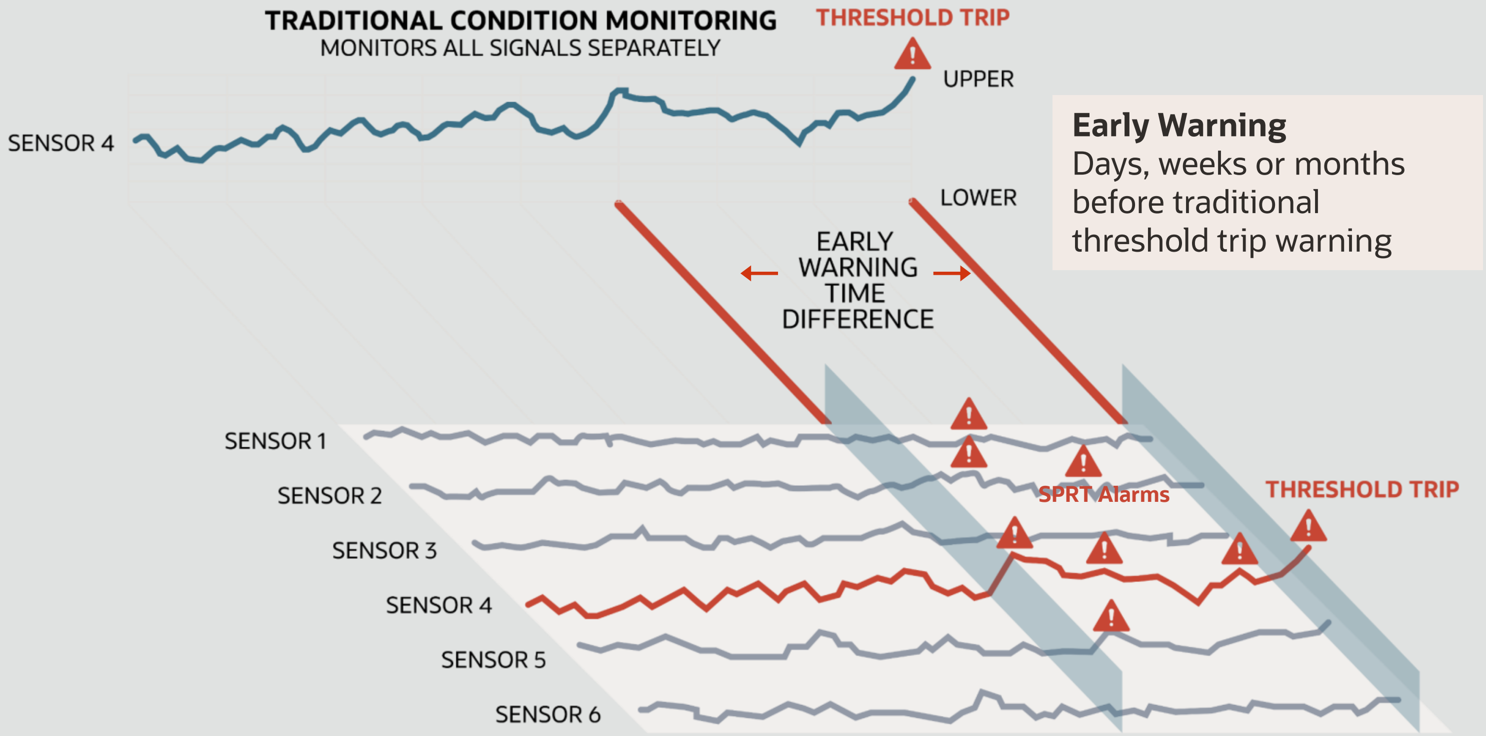

# MSET as a Product Suite for Prognostic

*Abraham Wald, (June 1943).*
*- Oracle uses modified SPRT*

Real world prognostic applications with suboptimal telemetry data

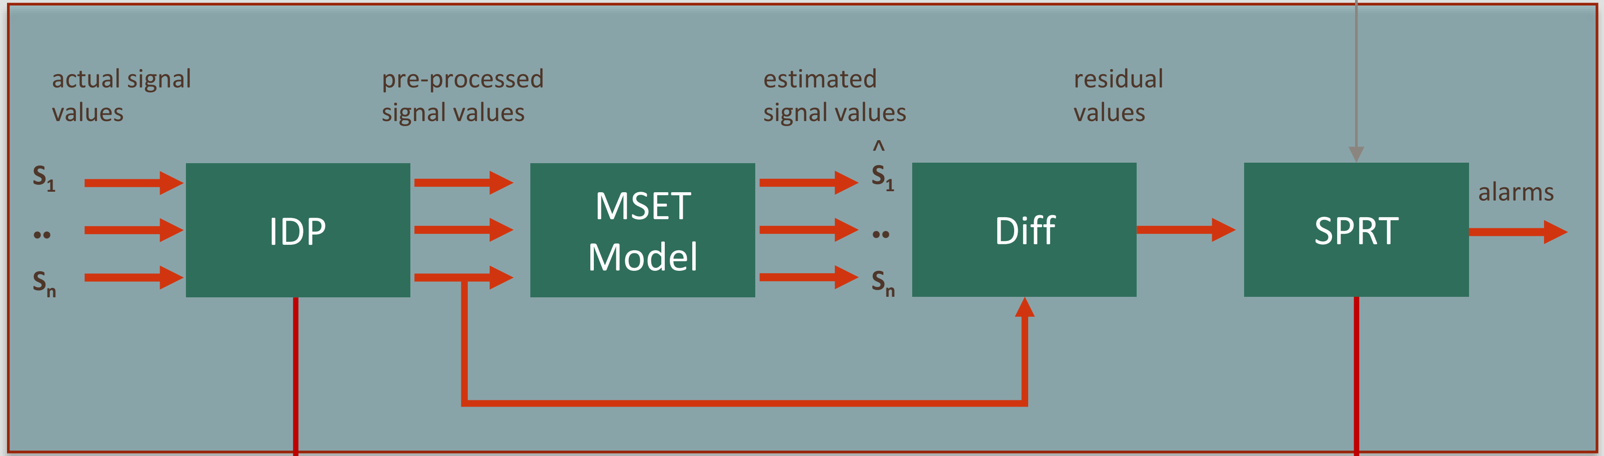


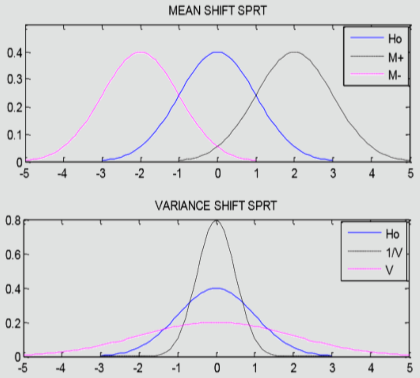


Sequential test for residual signals is based upon two hypotheses:

$H_0 : \mu = \mu_0 = 0$

$H_1 : \mu = \mu_1 = M$

**Data Historian Database**

**Telemetry Data Sensor Farm(s)**

<u>Intelligent Data Preprocessing</u>

- UnQuantize
- Analytical Resampling Process
- Missing Value Imputation
- Inferential Sensing
- Provenance Certification & Auditability
- TriPoint Clustering (TPC)
- Despiking
- Ambient Compensation
- Telemetry Parameter Synthesis System


Sequential Probability Ratio Test

- High sensitivity for subtle anomaly detection without increasing false alarm probability.
- Can accommodate any measurement noise, work with non-Gaussian noise signals.
- Sequential-binary hypothesis test compares reference distribution ($H_0$) vs degraded distribution ($H_1$)

# Top-5

# Intelligent Data Preprocessing Innovations

# #1

# Missing Values in Sensor Streams

*Optimal Missing Value Imputation*

# Background | Missing Values in Sensor Streams

- Common challenges in dense-sensor IoT time-series databases
- Often caused by saturated system bus or transmission errors
- Conventional solution is to fill in missing values through interpolations
- The preferred approach is to infer the missing values with the inter-correlations between the telemetry signals

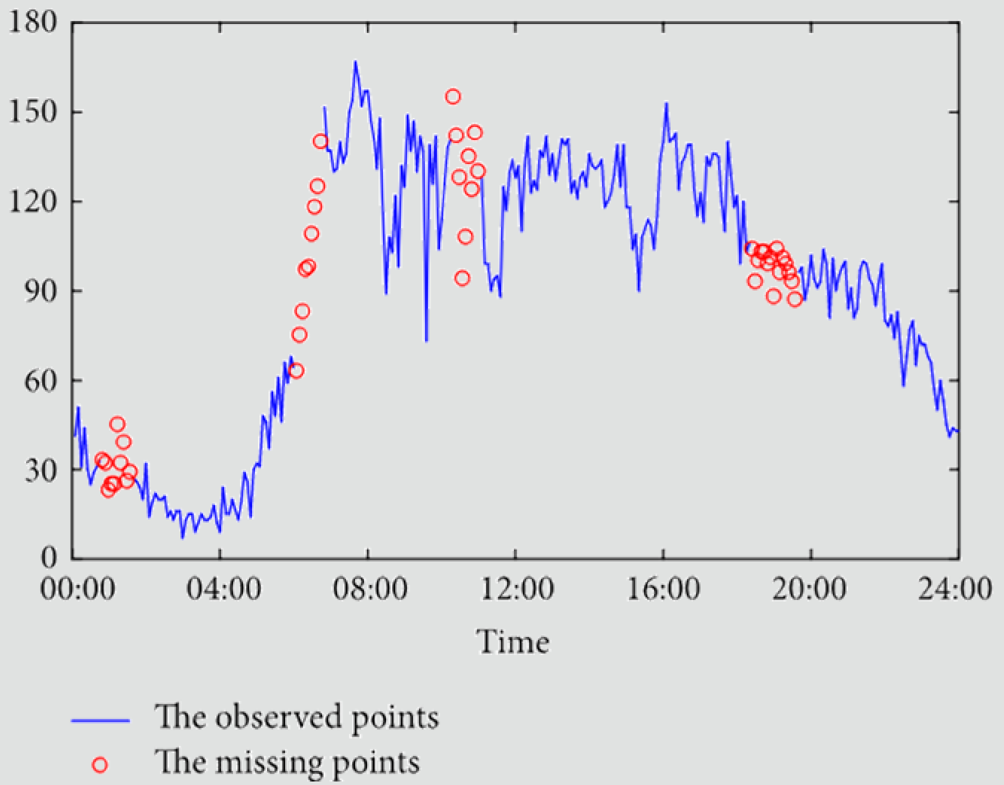

# Oracle's Missing Value Imputation[1] (MVI)

MVI Procedures:

- Pre-fill the missing observations with conventional interpolations
- Divide the dataset into two halves, A and B
- Train a MSET model using A, then apply the model to B to "update" the prior interpolated values in B
- Train a MSET model using B, then apply the model to A to update the prior interpolated values in A

Key differentiation value:

The inter-correlation between the signals is leveraged to fill in the missing values

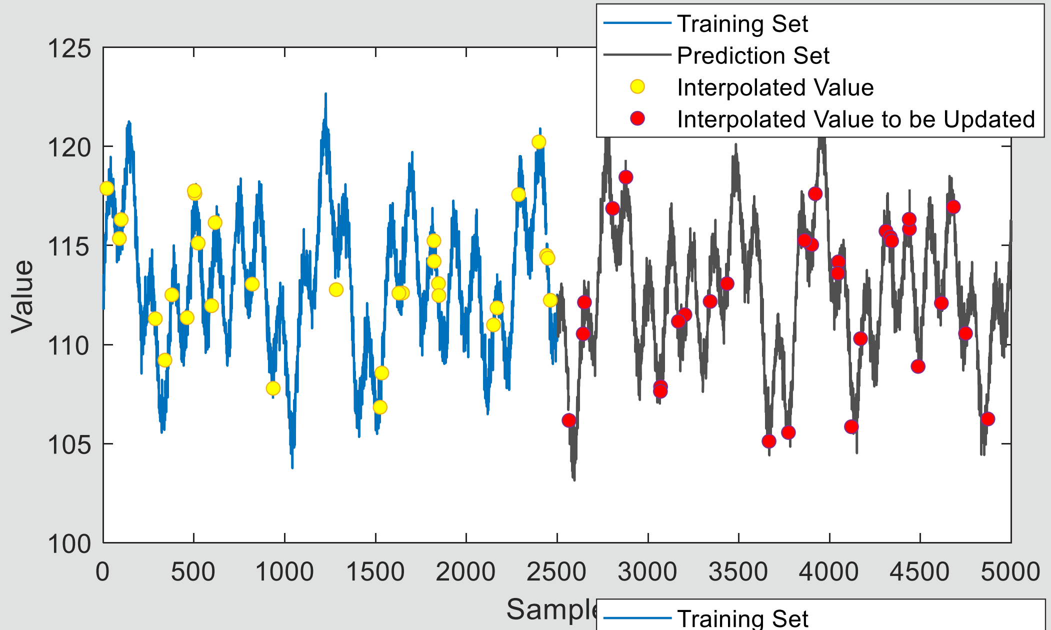


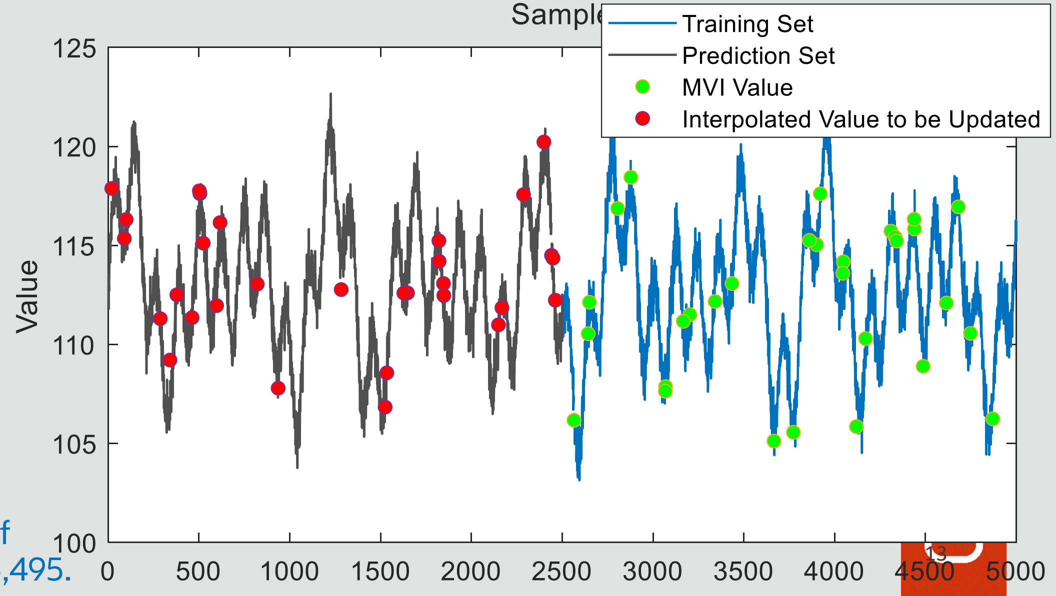


1. Wang, G.C., Gross, K.C. and Gawlick, D., Missing value imputation technique to facilitate prognostic analysis of time-series sensor data. U.S. Patent Application 16/005,495.

# Missing Value Imputation: Performance

Example application showing higher accuracy for MVI versus conventional linear interpolation

- The “valley” was missed by linear interpolation, which is likely indicative of a fault
- MVI was able to uncover the “valley” using the inter-correlations between signals, which can benefit both training and surveillance process

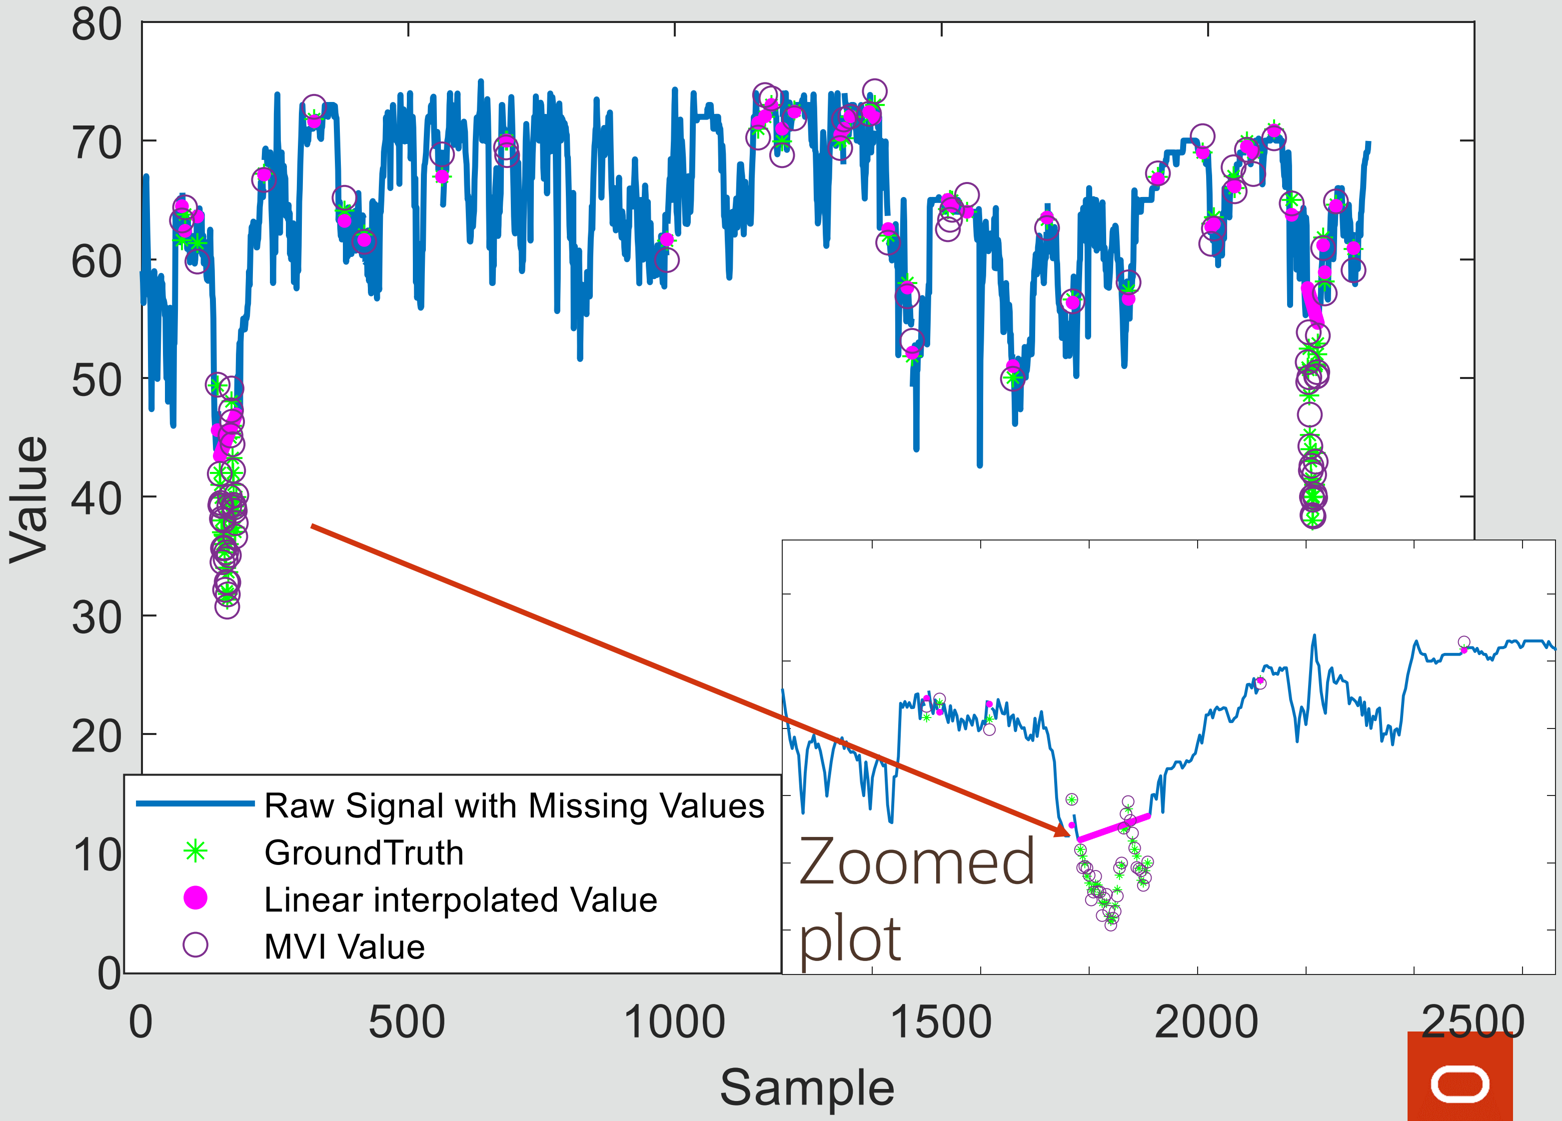

# #2

# Signal UnQuantization

*Turns low-resolution input signals into high-accuracy output signals*

# Background | Signal Quantization

- Many industries (including the enterprise computing industry) use 8-bit Analog/Digital conversion chips for physical sensors

- As a result of low-bit resolution, physical variables are severely quantized

- Machine Learning algorithms can't discern small variations in the quantized telemetry signals that could precede component degradation or system failure

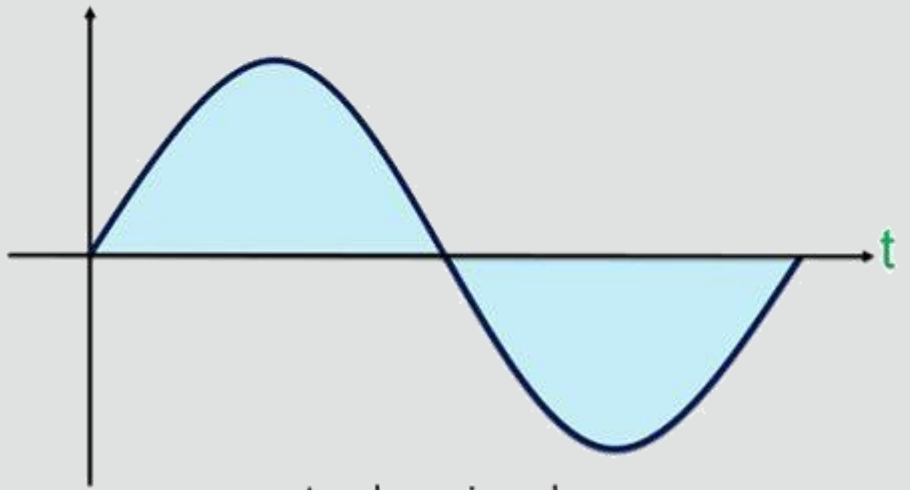


Analog signal

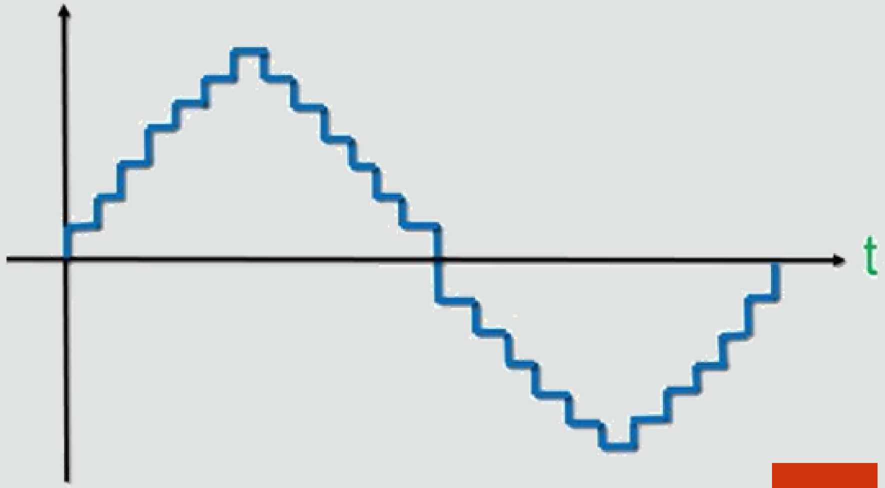


Quantized signal

# Signal Quantization Examples

- Three examples of quantized signals with 8-bit A/D chips (blue) vs. true signals (i.e., high-resolution) with 16-bit A/D chips (red)

  (a): a voltage signal with level-2 quantization consisting of 2 oscillating values

  (b & c): temperature and voltage signals with leve-4 quantization

- The non-linear trends and dynamics were absent in the quantized signals

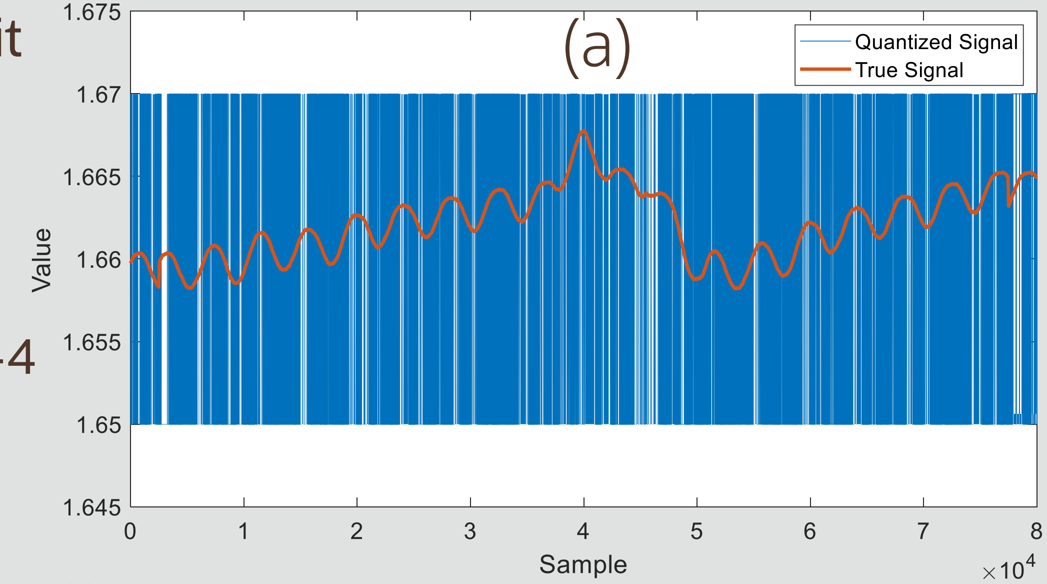


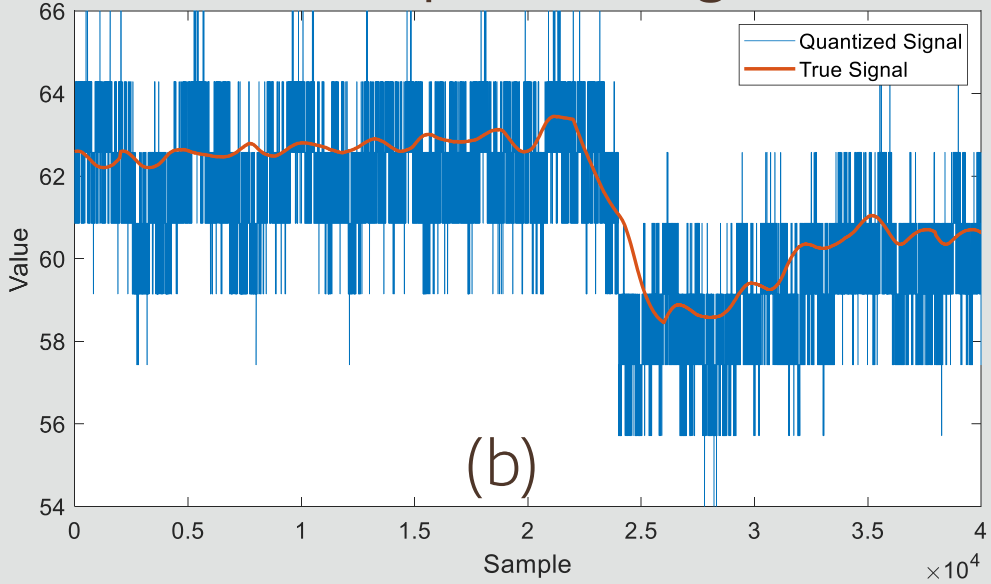


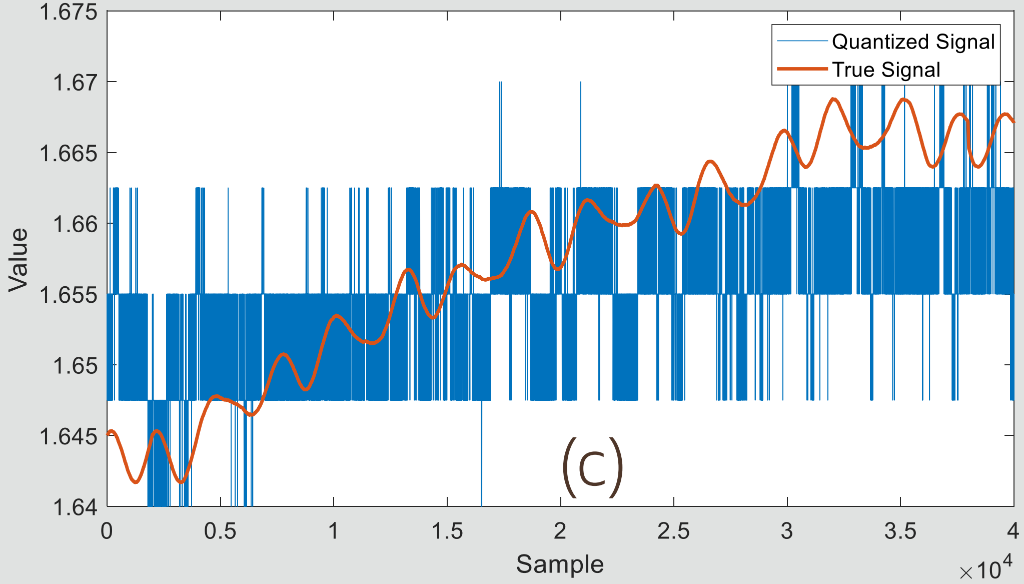

# Oracle's UnQuantize Techniques

Two approaches:

- Bin Density
  - Implements a sliding window over the quantized signal
  - Continuously update Bin Density values as a function of time
  - Applies to lower levels (below 4) of quantization

- Fourier Decomposition
  - Compute Fourier transform and pick N number of modes from Power Spectral Density
  - Reconstruct smooth signals with inverse Fourier transform
  - Applies to higher levels (4 and above) of quantization

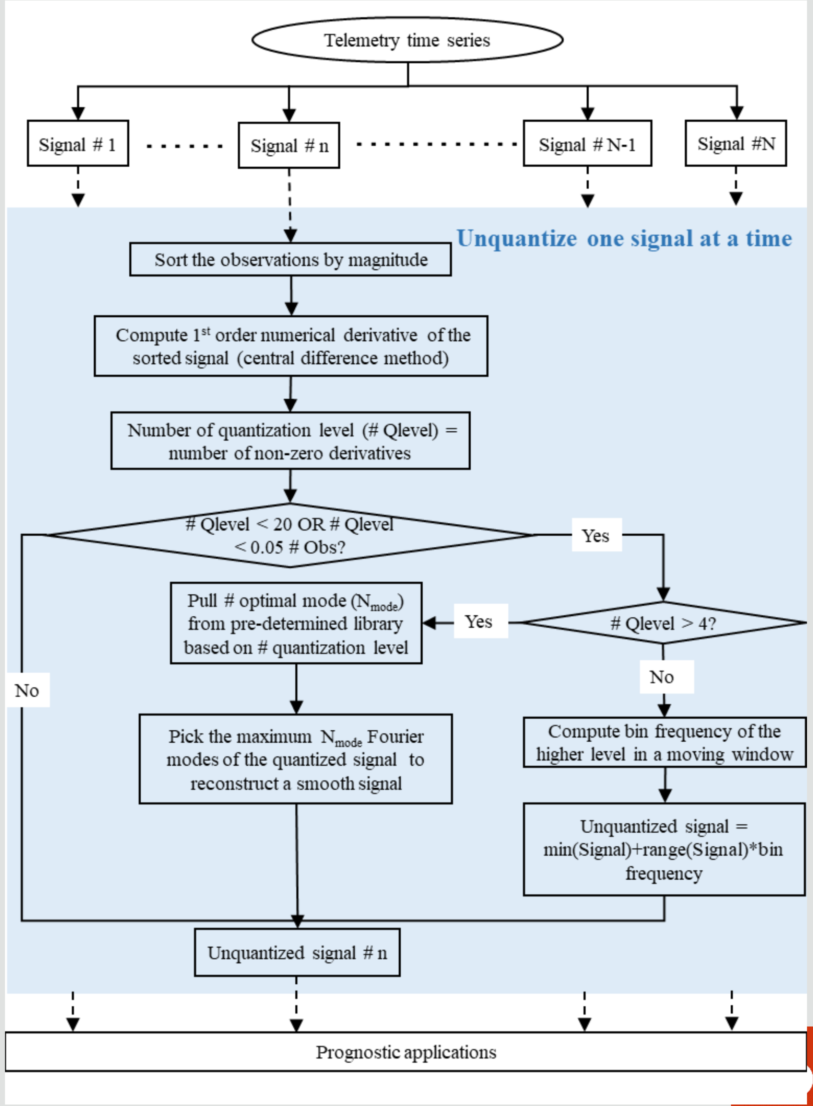

# UnQuantize: Performance

- Drastic difference between the quantized sensor readings (blue) and the genuine signal characteristics (red) on both signals
- The global trend and local dynamics were missing in the quantized signals
- The UnQuantize algorithm reveals the true characteristics of the two signals, benefiting the downstream anomaly detection tasks

StarCat Core Voltage and Temperature Signals

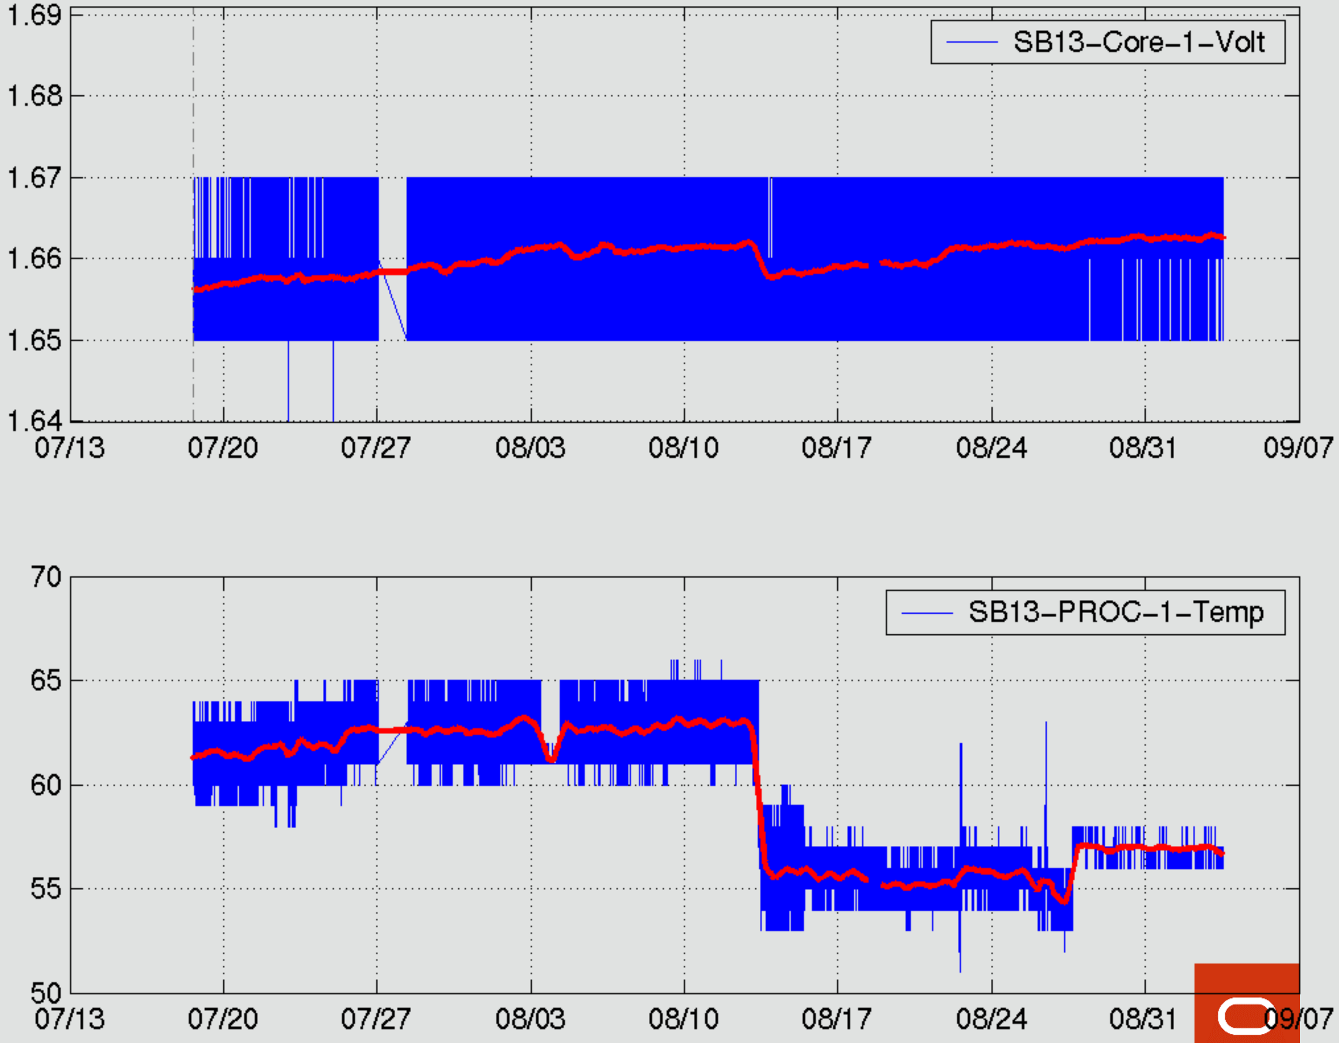

# UnQuantize: State-of-the-art

**Challenge Addressed**: Most enterprise computers, and many IoT industries use 8-bit A/D conversion chips for physical sensors. As a result, physical variables (e.g. voltages, currents, temperatures, fan speeds) can be severely quantized. This causes large uncertainties in prognostics analyses of large time-series databases, and for real-time IoT applications causes jitter inefficiencies in feedback/control loops.


- Li, M. and Gross, K.C., Oracle International Corp, 2019. Dequantizing low-resolution IoT signals to produce high-accuracy prognostic indicators. U.S. Patent 10,496,084.
- Gerdes, M.T., Gross, K. and Wang, G.C., 2021. Unquantize: Overcoming Signal Quantization Effects in IoT Time Series Databases. In Advances in Security, Networks, and Internet of Things (pp. 621-636). Springer, Cham.

# #3

# Analytical Synchronization Process

*Essential for multi-sensor prognostics*

# Background | Out-of-Sync Telemetry Signals

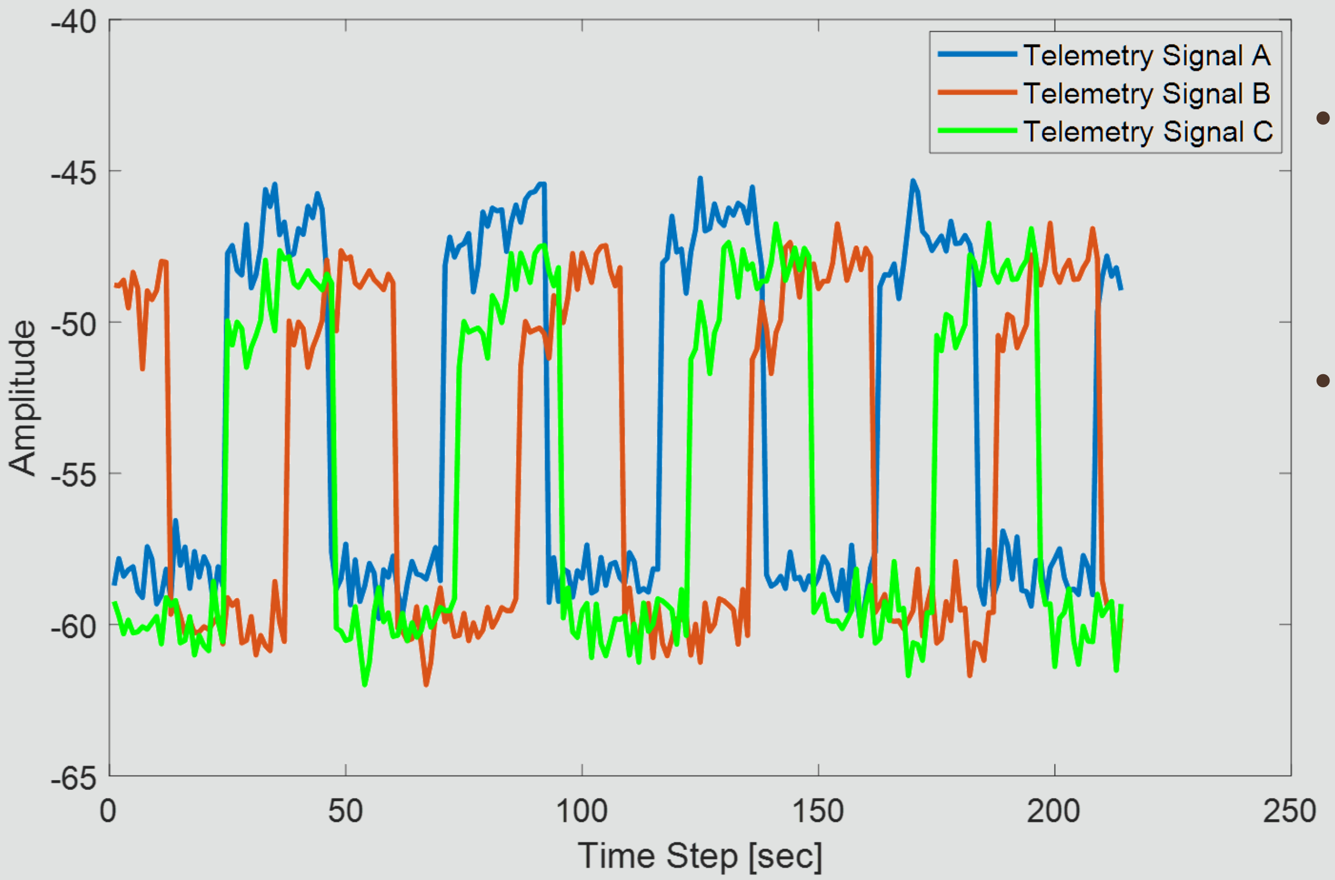


- Out-of-sync telemetry signals originate with distributed data-acquisition modules with out-of-sync clocks
- Clock mismatch issues will cause almost all time-series ML algorithms to fail
  - ➢ Clock mismatch “blurs” the patterns of correlations among signals
  - ➢ Normal signals with clock skews can be falsely identified as anomalous

# Oracle's Analytical Synchronization Process[1]

## Essential for Multi-Signal Prognostics

- Examines telemetry signals and autonomously uses one of 3 methods to align the out-of-sync signals
- Considers the trade-off between the highest possible alignment accuracy and overhead compute cost
- Benefits the downstream training and/or detection tasks

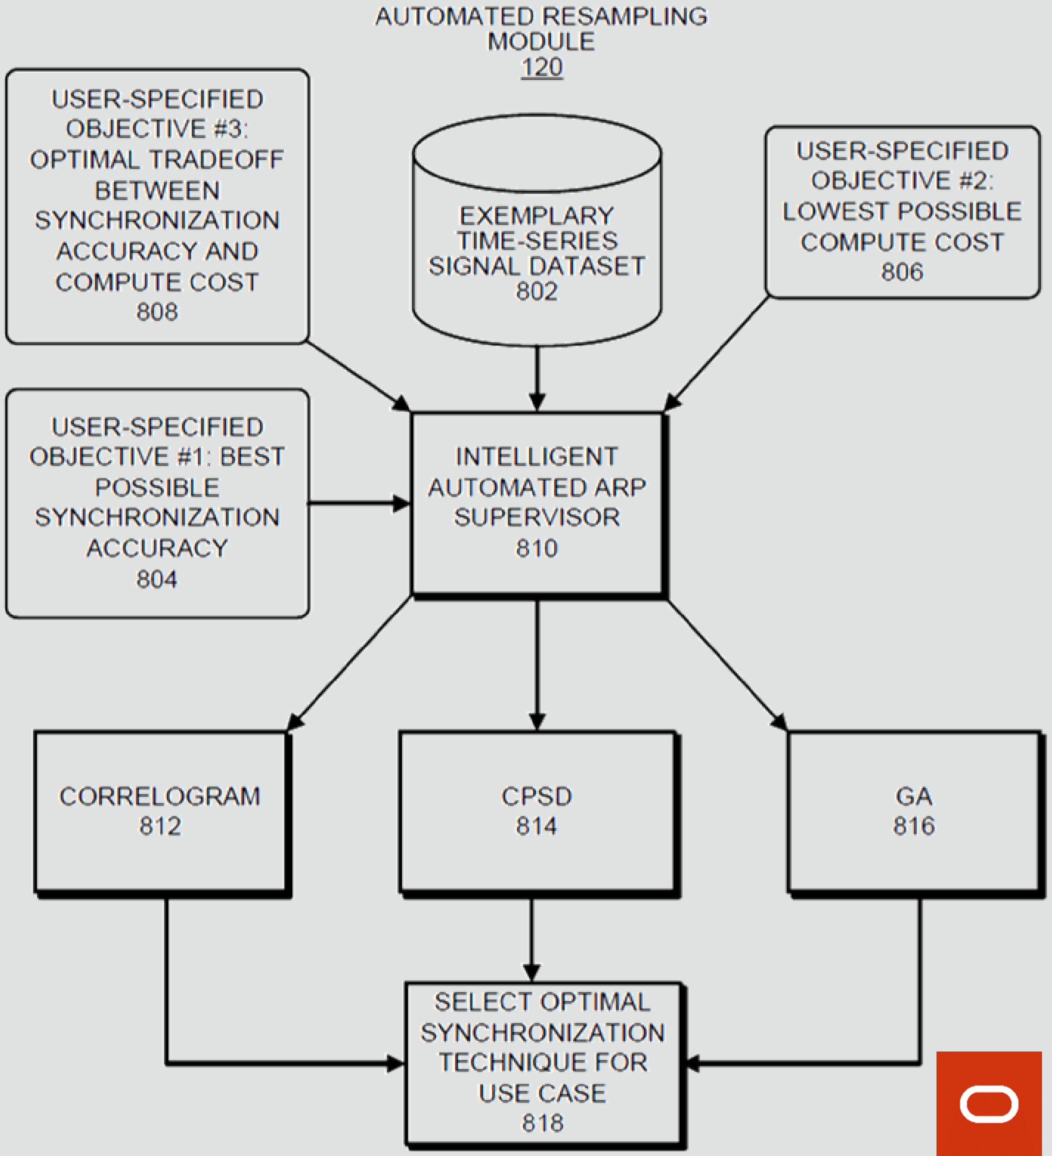


1. Automated Analytic Resampling Process for Optimally Synchronizing Time-series Signals, K. C. Gross and G. C. Wang, U.S. Patent Application 16/168,193.

# Analytical Synchronization Process: Example

Correlogram

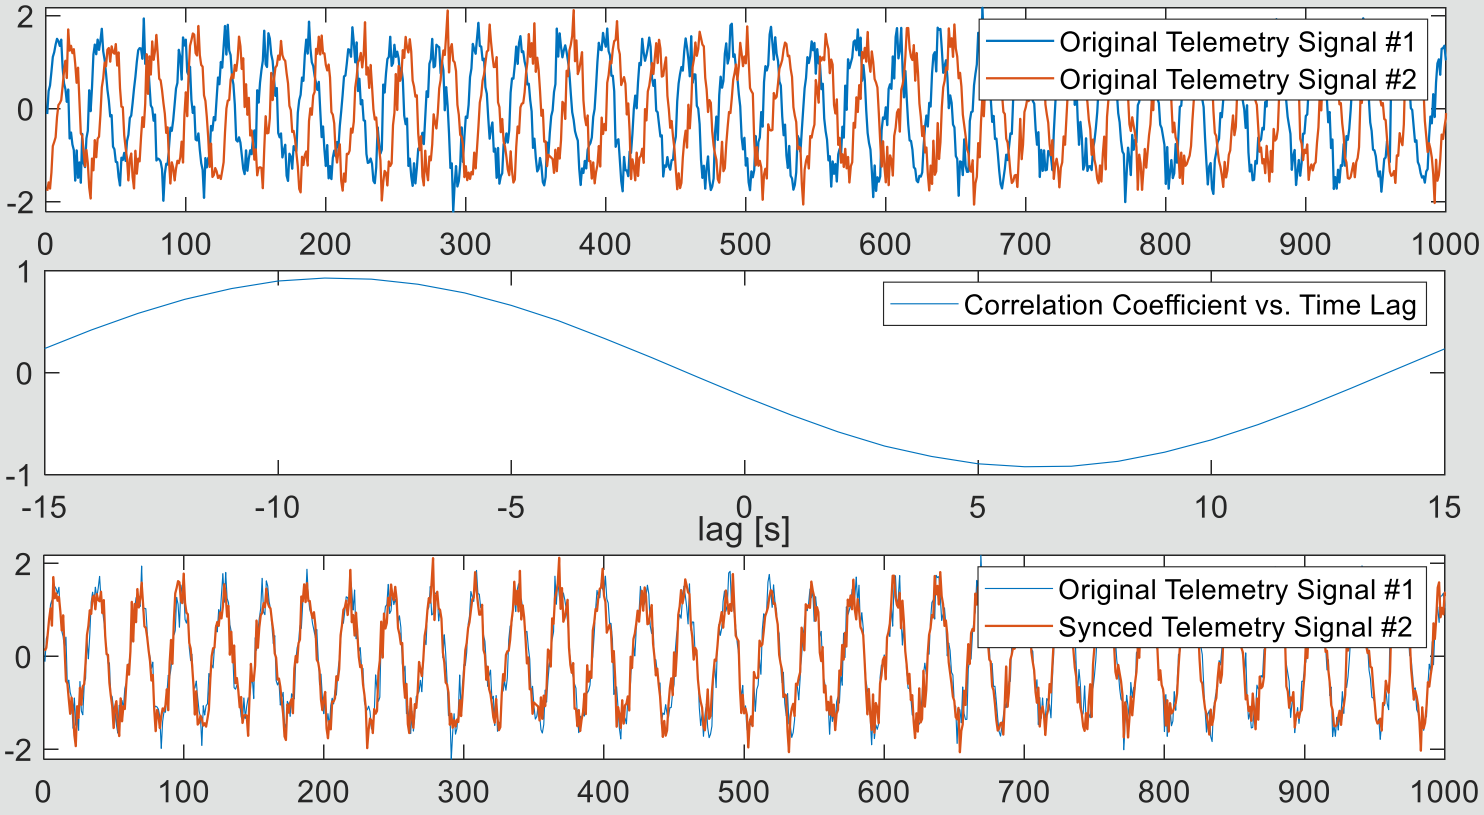


Cross Power Spectral Density

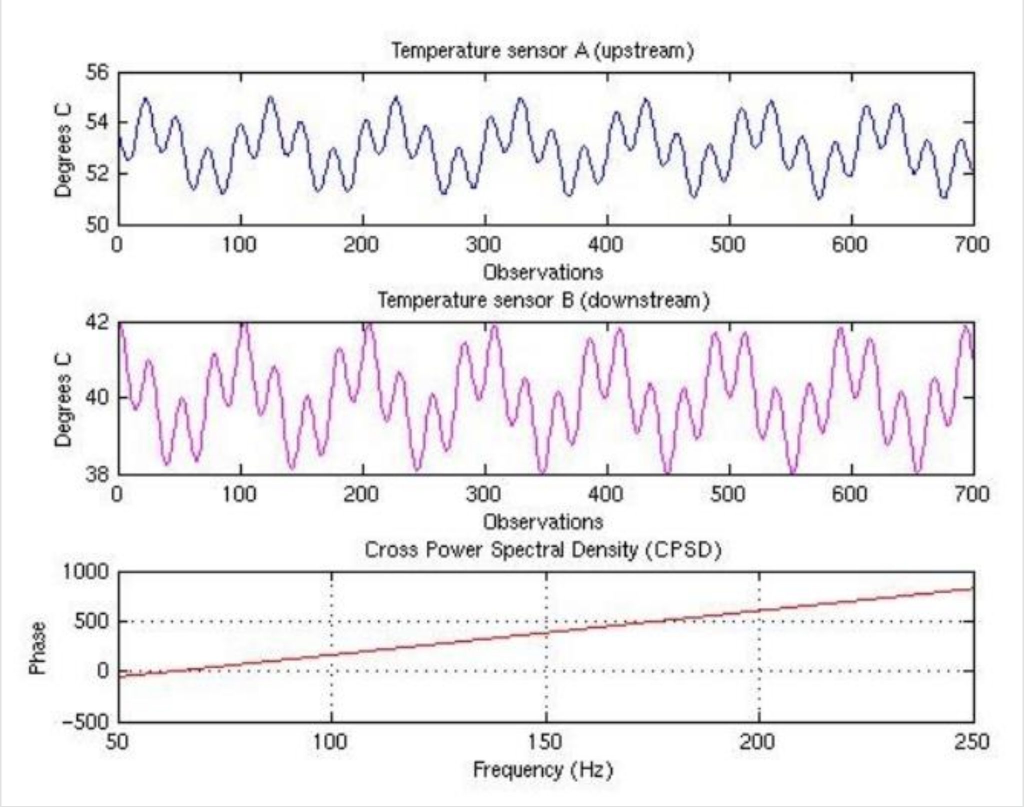

# Analytical Synchronization Process: State-of-the-art

**Competitive Differentiation**: Oracle's Analytical Synchronization Process assure optimal machine learning prognostics for all types of variable sampling rate, variable clock-skew challenges across all IoT industries.

.

- "High-Accuracy Synchronization of Signals from Computer Systems," K. C. Gross and K. Vaidyanathan, Case ID SUN080852, U.S. Patent 8,214,682.
- "Synchronizing Signals Related to Real-Time Prognostics of Enterprise Computer Systems," K. C. Gross and K. Vaidyanathan, Case ID SUN080126, U.S. Patent 8,365,003.
- "Automated Analytic Resampling Process for Optimally Synchronizing Time-series Signals," K. C. Gross and G. C. Wang, U.S. Patent Application 16/168,193.
- "Real Time Empirical Synchronization of IoT Signals for Improved AI Prognostics," Wang, G.C. and Gross, K., 2018, December. In 2018 International Conference on Computational Science and Computational Intelligence (CSCI) (pp. 954-959). IEEE.
- "Compression/dilation Technique for Synchronizing Time-series Signals Used to Detect Unwanted Electronic Components in Critical Assets Based on EMI Fingerprints, " G.C. Wang and K.C. Gross, Oracle International Corp, 2021. U.S. Patent 11,210,400.

# #4

# Ambient Variable Compensation

***Automatic removal ambient variations from load-dependent variations***

# Background | Ambient Temperature Variation

External (ambient) parameter variations “superimpose” on internal sensor signals

Example shows internal thermal sensors for Utility Transformers

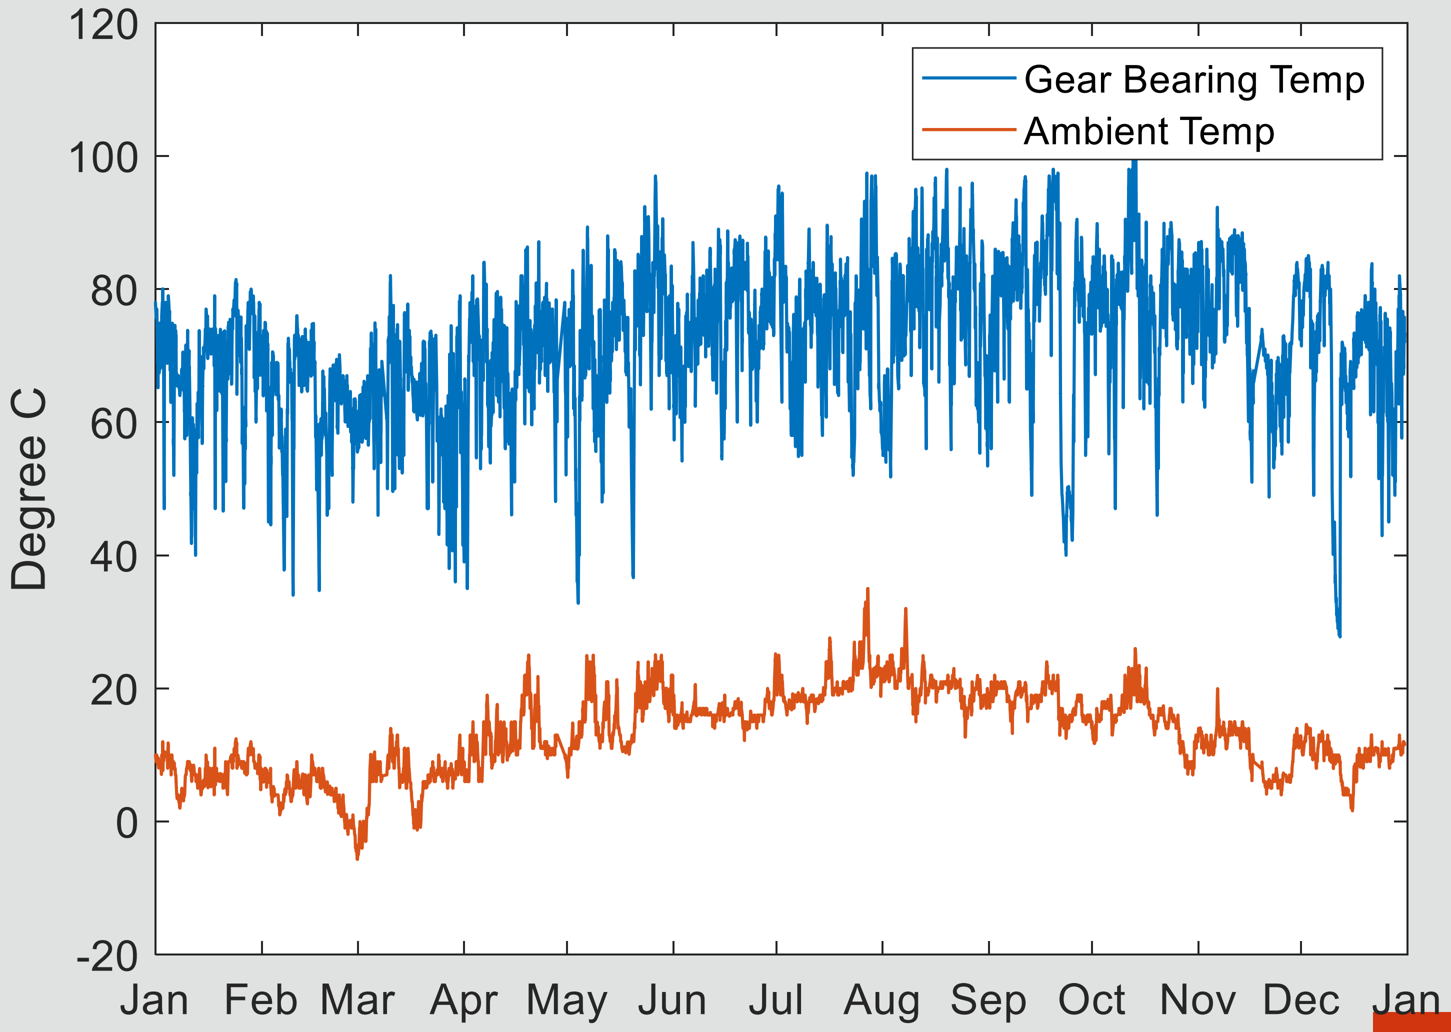

# Impact of Ambient Temperature Variation

## Real world wind turbine prognostic use case

Ambient temperature variation causing high false alarms


# Oracle's Ambient Compensation Technique[1]

- Infers ambient thermal dynamics from the data consisting of multiple temperature sensors
- Removes the environmental dependencies from all temperature related sensor readings
- Benefits the downstream training and detection tasks

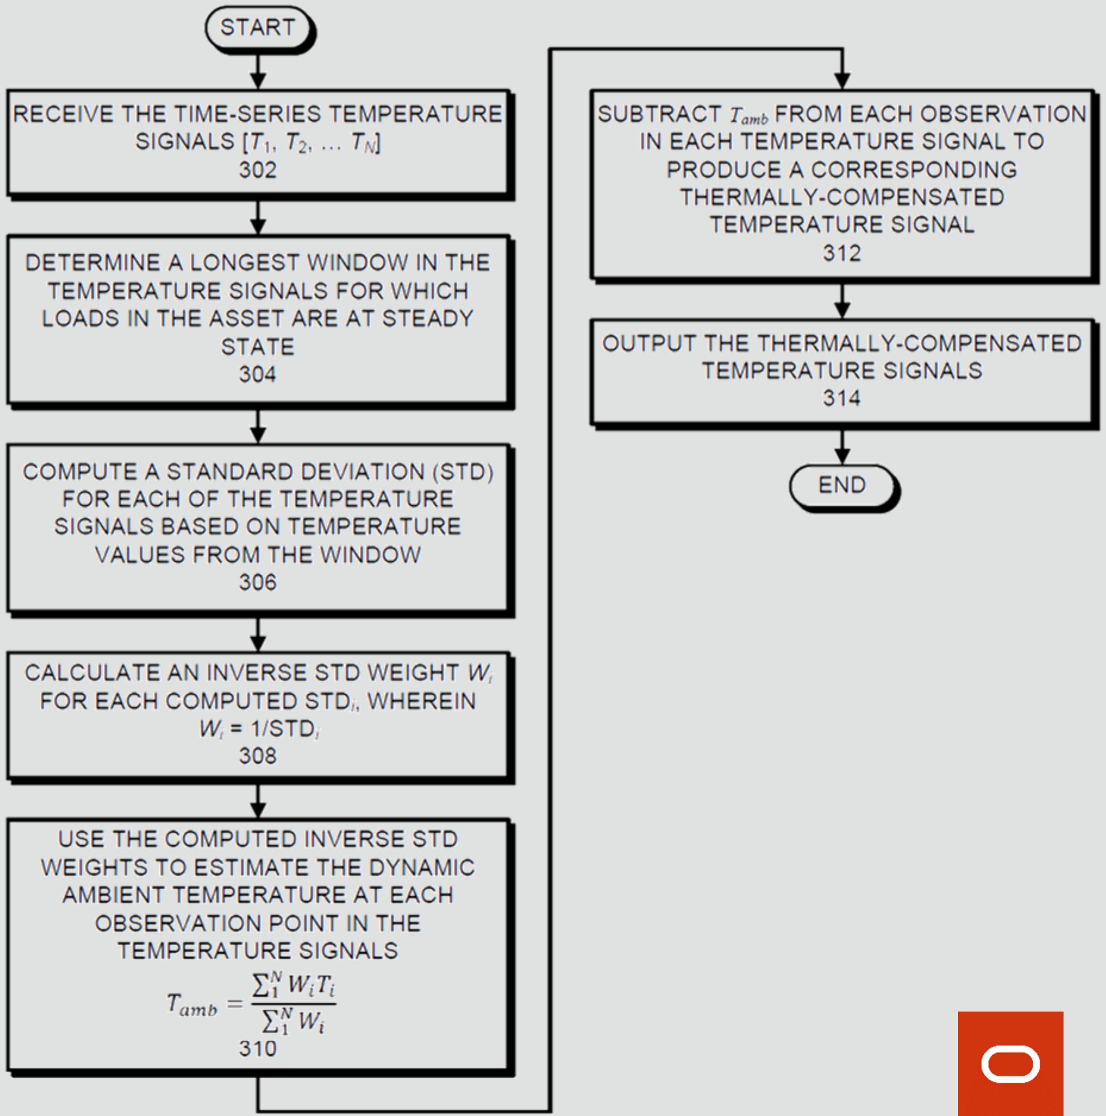


1. Gross, K.C., Wang, G.C. and Wetherbee, E.R., Oracle International Corp, 2021. *Thermally-compensated prognostic-surveillance technique for critical assets in outdoor environments*. U.S. Patent 10,929,776.

# #5

# Digital Twin

*A virtual copy of a real-world asset, augmented by real-time data and analytics*

# Background | Oracle's Digital Twin Innovation

Innovation for Advanced Prognostics of Complex Engineering Assets (since 2003)

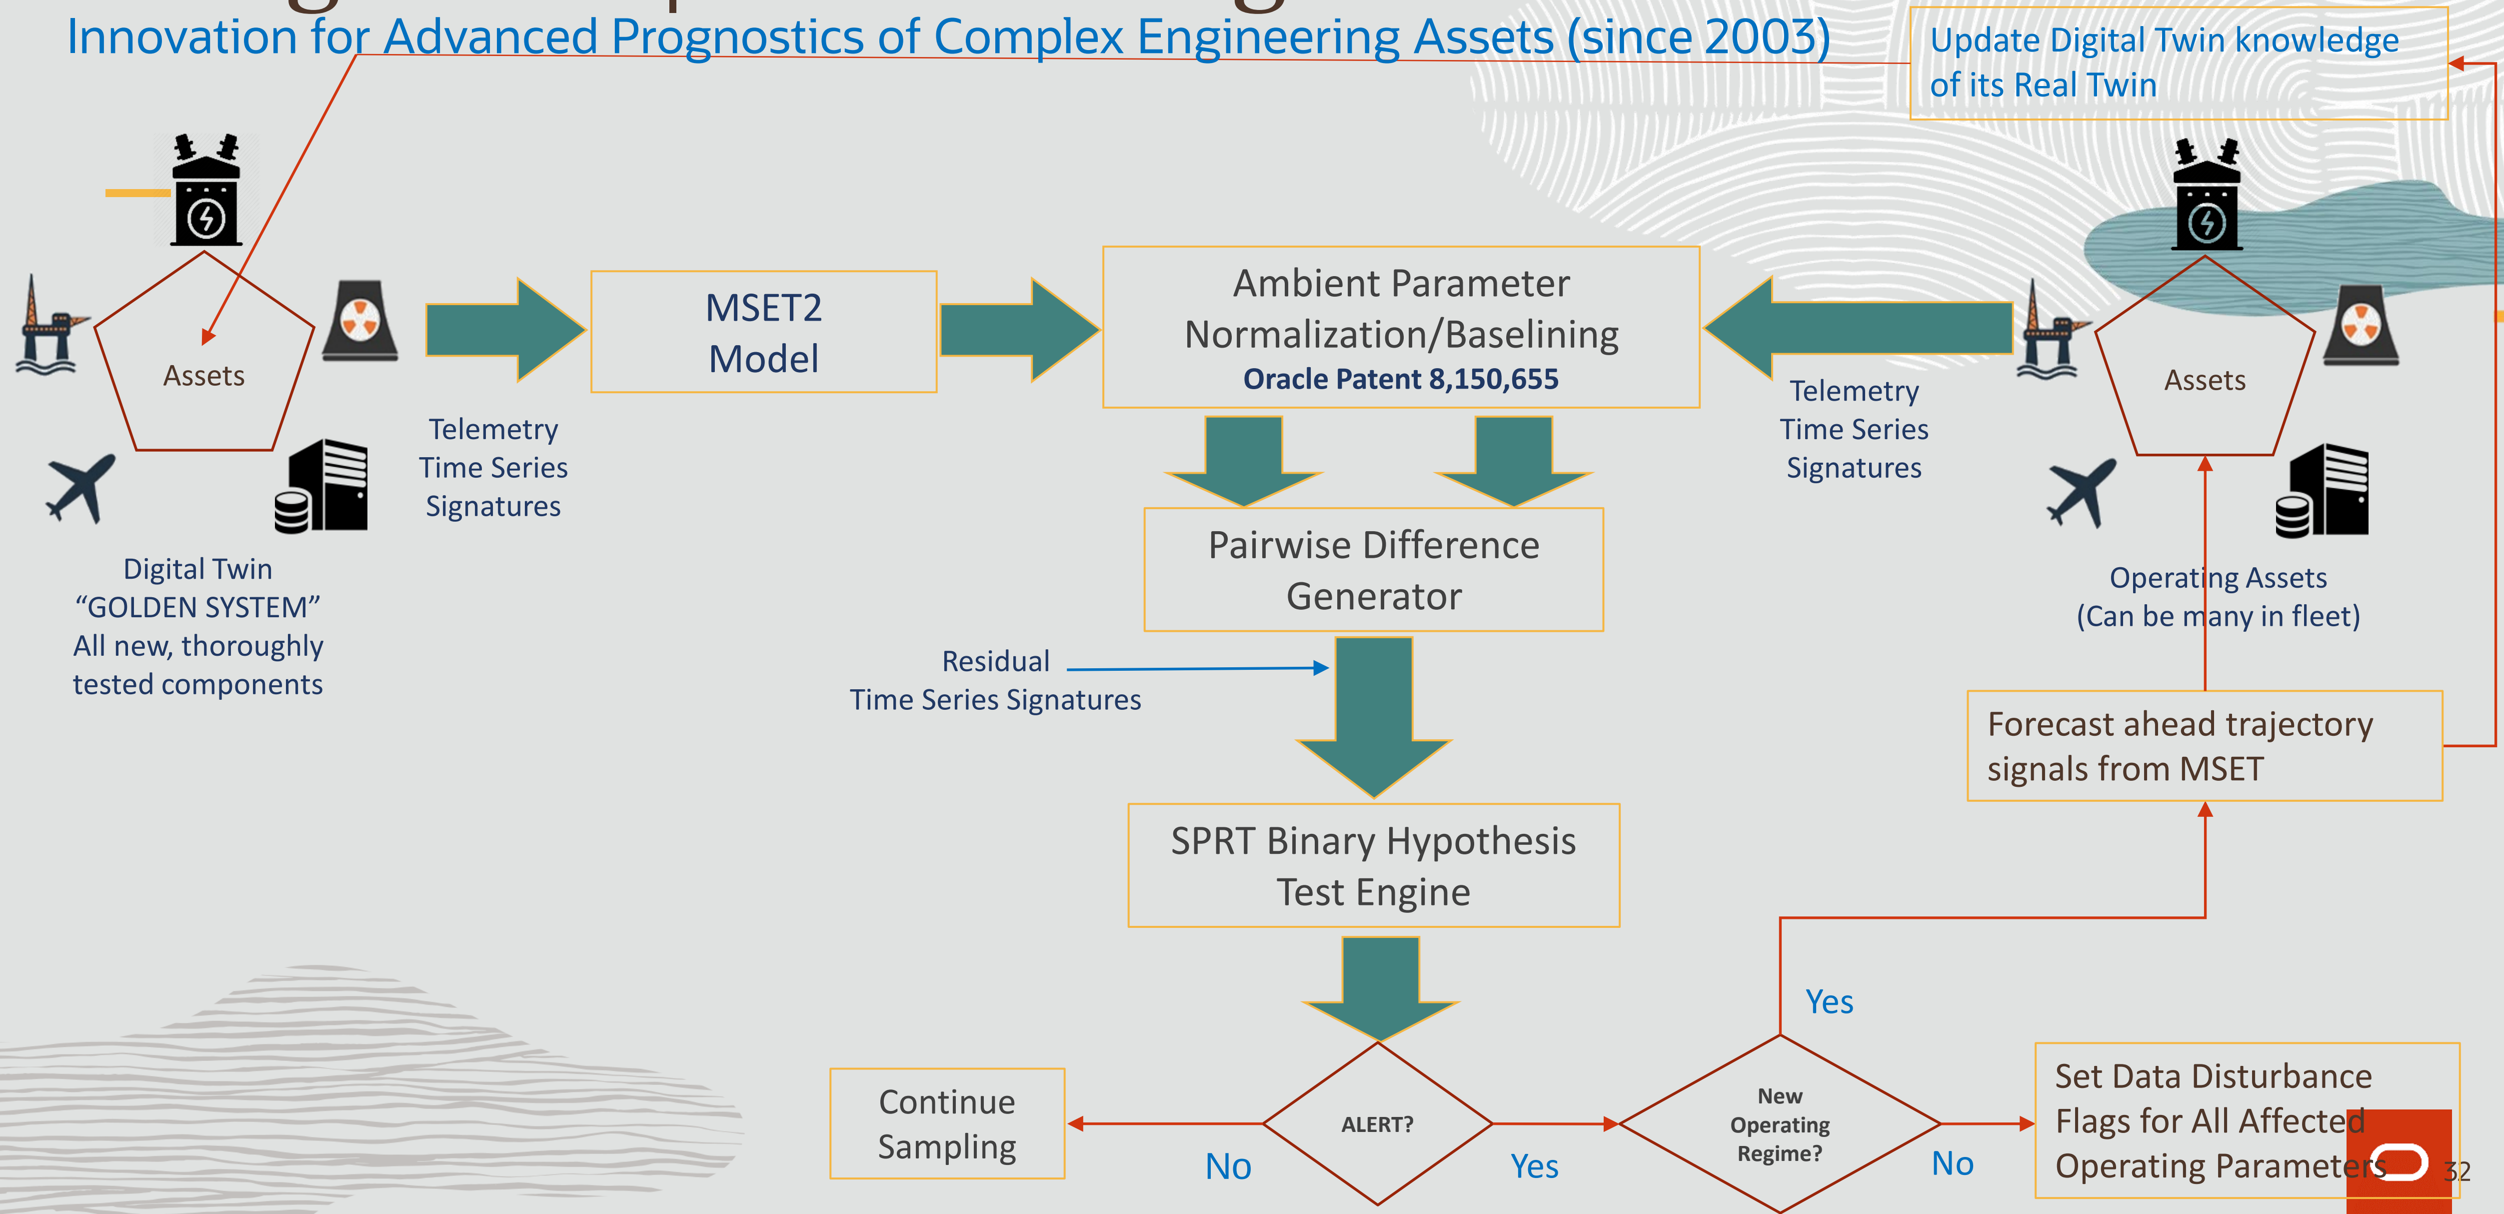

# The Extension of Digital Twin: EMI Fingerprints

For detection of Counterfeit Components in Server Electronic System

- Train on One Golden Asset (certified to have no counterfeits)
- Scan any number of assets in the field, or at loading docks, or ports of entry

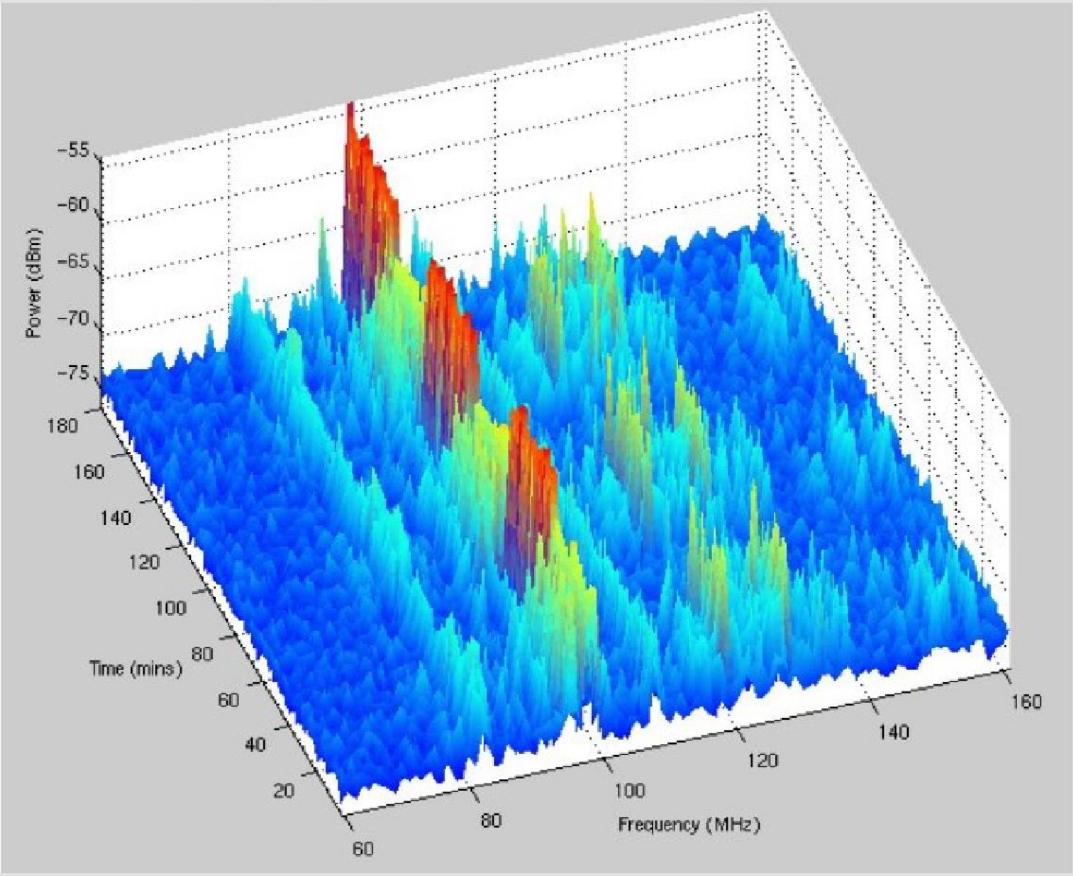


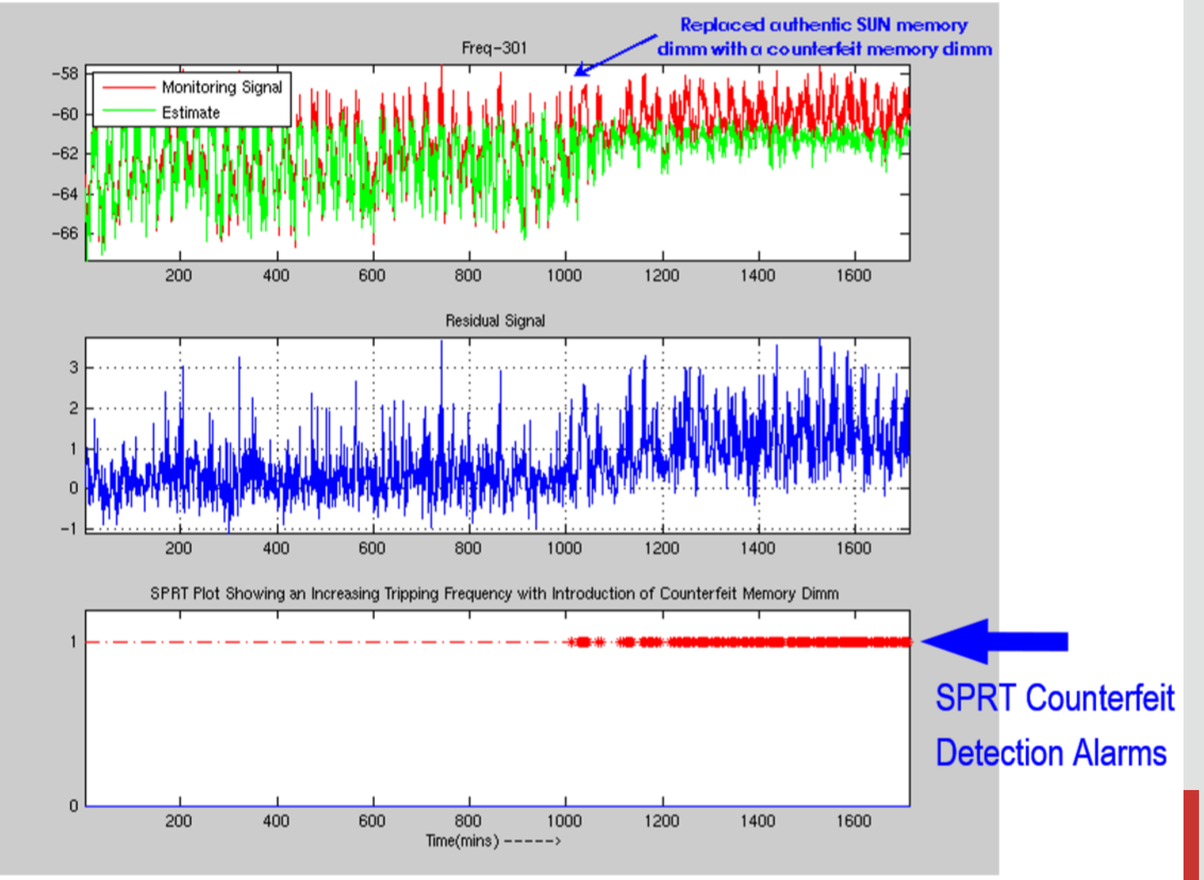

# Top-5 Innovations Summary

# Oracle Prognostics-Centric Innovations

## Intelligent Data Preprocessing

**Common Issues**

- Missing values in IoT sensor data
- Variable sampling rates of IoT data
- Asynchronous signals (Clock mismatch)
- Low resolution sensor signals
- Sensor drifting out of calibration

**MVI**
Optimal Missing Value Imputation

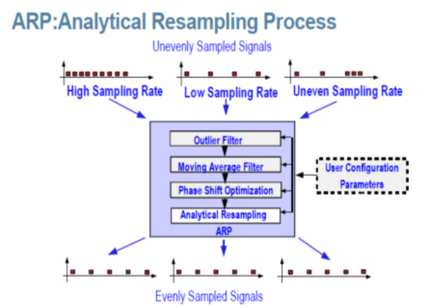


**UnQuantize**
Low resolution input signals into high accuracy output signals upstream

**Inferential Sensing**

- Optimal Server Fan Control
- Digital Twin for Prognostics
- Energy Waste Reduction
- Intelligent Power Monitoring
- Remaining Useful Life

# Thank You

Technical Questions?

Guang Wang: guang.wang@oracle.com

Kenny Gross: kenny.gross@oracle.com

Resources:

Anomaly Detection Documentation: https://docs.oracle.com/en-us/iaas/Content/anomaly/using/home.htm

Oracle MSET2 Blog: https://blogs.oracle.com/bigdata/real-time-machine-learning-use-case

Interested in Trying Out MSET Based Anomaly Detection Service?

Viji Krishnamurthy (PM):

viji.krishnamurthy@oracle.com

ORACLE