Solution Brief

Manufacturing/Oil and Gas
AI-based Predictive Maintenance

Up To 3.1x Faster with Intel® Optimizations, SparkPredict Learns Faster to Deploy Quicker

Using multiple processes and Intel Distribution of Python, SparkPredict on 2nd Gen Intel Xeon® Scalable processors makes maintenance of industrial systems smarter

INTEL® AI BUILDERS MEMBER

sparkcognition™

Industry 4.0 modernization means more machinery will run manufacturing in the future. According to CB Insights, 76 percent of manufacturers reported in 2017 of having a smart factory initiative in the works. That means more equipment that will need periodic maintenance and repair.

Equipment downtime can severely impact manufacturing production and revenue. According to industry analysts, equipment failure causes 42 percent of unplanned downtime for industrial manufacturers and costs them an estimated $50 billion annually. Predictive maintenance can reduce downtime up to 50 percent and increase machine life up to 40 percent.

Often, before a major failure occurs, the equipment is generating signals that can be leveraged for a proactive response and help isolate issues before they become problems. Preventive maintenance operations can use these signals, when available, and artificial intelligence (AI)-enabled solutions can provide predictive alerts that help optimize production and minimize downtime.

SparkPredict from AI Builder member SparkCognition is a turnkey, machine learning-based solution that analyzes machine sensor data and returns actionable insights, flagging suboptimal operations and identifying impending failures before they occur.

AI-powered SparkPredict Optimizes Equipment Operations

SparkPredict uses data and machine learning to approximate the physics of industrial assets with a high degree of accuracy (Figure 1). The application:

- Uses machine learning insights and diagnostics to accelerate decision-making workflows.
- Predicts and understands a wide range of failures and anomalies.
- Includes human input (human-in-the-loop) integrated into the continuous learning solution.
- Provides a high-level visualization of what subsystems need attention.

With the flexibility of running either on-premise or in the cloud, SparkPredict identifies a broad spectrum of problematic asset behaviors, from critical, never-before-seen anomalies to commonly occurring events. With this data, businesses can better understand, predict, and aggregate failure insights across an entire fleet.

SparkPredict is unique in the scalable way it executes asset-centric models and makes insights usable to customers through data preparation, model-building and execution, and insight visualization. This end-to-end infrastructure provides the means to operationalize customer- or SparkCognition-developed models. With these capabilities, SparkPredict revolutionizes maintenance and delivers the following operational benefits:

- **Improve asset production:** Equip operations with predictive AI analytics that protect assets and keep them online.
- **Minimize maintenance costs:** Drive labor efficiencies during downtime with insights that inform repairs.
- **Future-proof operations:** Retain workforce knowledge with machine learning that codifies human expertise.
- **Exceed availability targets:** Predict more machine problems with less work and expand asset failure horizons.
- **Streamline work orders:** Take quick, informed repair actions with explainable failure indicators.
- **Operate with confidence:** Maintain predictive accuracy with automatic model retraining that improves models over time.

SparkPredict Workload Overview and Intel Optimizations

SparkPredict analyzes accumulated time-series samples from manufacturing equipment to predict potential machine issues. SparkPredict's workload comprises four steps:

1. Dimensionality reduction using Principal Component Analysis (PCA)
2. HDBSCAN clustering algorithm
3. Training the Classifier
4. Prediction

The HDBSCAN algorithm analyzes clusters of data, and the results are used to train the classifier to detect anomalies for predictive alerts of equipment malfunction. HDBSCAN is written in Python and uses several Python-based libraries, such as numpy and sklearn.

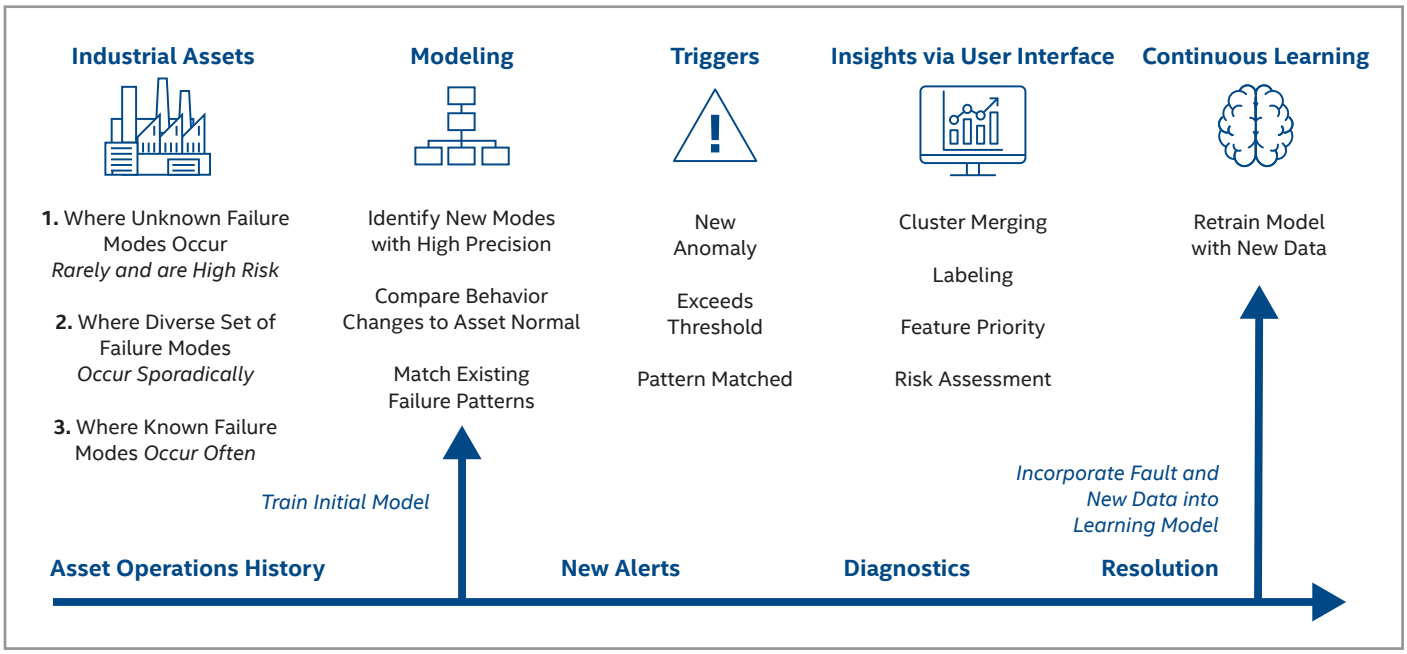

Figure 1. SparkPredict Architecture Overview.

Optimization

Initial preliminary experiments suggested that the knee of the performance curve is equal to the number of cores on the hardware (Figure 2). This suggests that the workload uses an unoptimized HDBSCAN algorithm, which is entirely single-threaded, allowing for performance improvements through parallelism.

Using multiple processes, each of the sample groups could be run in a single process, launched by a main thread. Each process would run all four steps and then return results to the main thread. The number of maximum processes allowed was calculated using the following:

Number of processes = Max (20, Number of cores – (1 core for the main thread + 1 core for the OS))

Engineers determined that beyond 20 processes, memory stalls take over and prevent further gains on the selected hardware platform.

Additionally, since the algorithm was written in Python, developers could use the Intel® Distribution of Python, which integrates the Intel Math Kernel Library for Deep Neural Networks (Intel MKL-DNN). With the Intel Distribution of Python, SparkPredict could directly leverage the technologies built into the 2nd Gen Intel Xeon® Scalable processors to accelerate performance.

After implementing these improvements and running benchmark tests, engineers were able to achieve speedup of up to 3.1x faster using multiple processes compared to a single process (Figure 3).

Conclusion

SparkPredict helps business operations reduce downtime, improve production efficiencies, and save costs through machine learning-enabled maintenance and predictive alerting. With multi-process and library optimizations, Intel was able to achieve up to 3.1x faster performance of SparkPredict than before optimizations. These optimizations mean faster time to deployment and faster relearning when in the field.

Figure 2. Execution Time Evaluation Test Results.

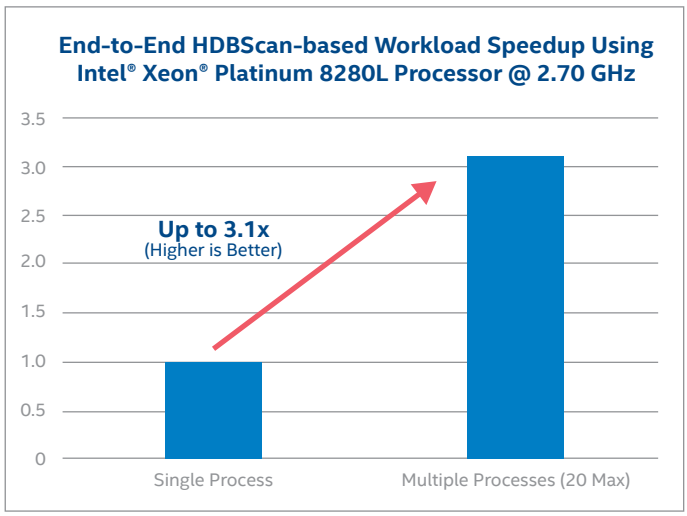

Figure 3. Post-optimization speedup tests achieved of up to 3.1x faster performance.[1]

For more information about SparkPredict, visit **https://www.sparkcognition.com/products/sparkpredict**

Learn more about the Intel AI Builders program at **https://builders.intel.com/ai/membership**.

INTEL® AI BUILDERS MEMBER

SparkCognition builds artificial intelligence solutions for applications in energy, oil and gas, manufacturing, finance, aerospace, defense, and security.

[1] System Configuration: Two-socket server configured with 2 x Intel® Xeon® Platinum processor 8180 (2.50 GHz, 28 cores), 192 GB DDR4@2666MHz memory (12 x 16 GB DIMMS), 1.5 TB Intel® SSD (SC2BX01), CentOS Linux* release 7.5.1804 (Core) (3.10.0-862.9.1.el7.x86_64);

Baseline software application: TensorFlow-Serving r1.9 (https://github.com/tensorflow/serving); Intel

Optimized software application: TensorFlow-Serving r1.9 + Intel MKL-DNN (https://mirror.bazel.build/github.com/intel/mkl-dnn/archive/0c1cf54b63732e5a723c5670f66f6dfb19b64d20.tar.gz) + optimizations (availability of optimizations expected in TensorFlow-Serving release 1.10).

[2] https://www.cbinsights.com/research/future-factory-manufacturing-tech-trends/

[3] https://partners.wsj.com/emerson/unlocking-performance/how-manufacturers-can-achieve-top-quartile-performance/

[4] https://www.mckinsey.com/business-functions/operations/our-insights/manufacturing-analytics-unleashes-productivity-and-profitability#

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit: http://www.intel.com/performance

 0920/AU/HBD/PDF ♻ Please Recycle 344189-001US

intel®