SUCCESS STORY | LOCKHEED MARTIN

MINIMIZING DOWNTIME AND IMPROVING PRODUCTIVITY WITH AI-GUIDED PREDICTIVE MAINTENANCE

NVIDIA.

Lockheed Martin uses AI and NVIDIA DGX systems to distill insights from massive amounts of data and improve asset availability.

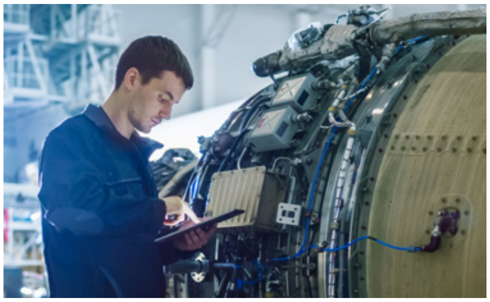

THE ESCALATING COST OF INDUSTRIAL UPTIME

In a study conducted by the U.S. Department of Commerce, manufacturing maintenance costs are estimated to be as high as 70 percent of the cost of goods produced. On top of that, approximately one-third of maintenance costs are unnecessary or improperly carried out.[1] Particularly in equipment-intensive industries, these costs are increasing and taking a disproportionate percentage of overall budget. One of the largest controllable manufacturing costs is in maintenance operations, and yet many maintenance systems are largely reactive, or time based and outdated. Unsurprisingly, 39 percent of manufacturing companies still use paper records to monitor maintenance in a global survey.[2] Furthermore, aging equipment (40%) and mechanical failure (24%) are the leading causes of unscheduled downtime, proving that a significant portion of this can be avoided with the right level of awareness.[2]

The U.S. government understands the impact of the cost of asset uptime, with 40 percent of the U.S. Department of Defense's fiscal year budget allocated to maintenance and operations, which is twice their procurement spend.[3] Lockheed Martin, the largest defense contractor in the world, manages and sustains fleets of aircraft for both government and commercial entities. The corporation is addressing the challenge of escalating maintenance costs in the defense sector by adopting AI for predictive maintenance to improve their client's asset availability, increase safety, and reduce costs.

CUSTOMER PROFILE

LOCKHEED MARTIN

Organization:	Industry:	Location:	Founded:	Website:
Lockheed Martin	Aerospace and defense	Bethesda, MD	1995	lockheedmartin.com

INFRASTRUCTURE

- **Training for small- to medium-sized projects:** DGX Station
- **Training and inference for the largest projects:** DGX server
- **Experimentation on smaller models:** Cloud instances with NVIDIA V100 Tensor Core GPU
- **AI workflow management software:** Domino Data Lab

SCHEDULED MAINTENANCE AVOIDS RANDOM FAILURES

Lockheed Martin's Data Analytics Innovations (DAI) Group collects high-frequency industrial sensor data from aircraft and helicopter fleets—which is very precise and fast data, as high as 10 Hz, or 1 data point every 0.1 seconds. This results in large-scale, time-series datasets. Sam Friedman, a senior data scientist in the DAI Group, has been leveraging AI models to find patterns in this data to conduct prognostics and health management (PHM) of their fleet, which includes usage-based lifing and condition-based maintenance. Usage-based lifing uses a regression approach to synthesize different sensor data streams to a single estimate of the remaining useful life of a part. Condition-based maintenance uses data about an asset—like where the aircraft has been and what conditions it has flown in—and uses AI risk modelling to classify the part into various categories that indicate its health and associated level of degradation. Once the degradation in one part is measured, costs can be further reduced, since degradation can also be estimated for many other parts without having to take the aircraft down. These two methods are very powerful in minimizing unplanned downtime by more accurately predicting when to take a part out of service for maintenance.

In the past, the DAI Group relied on traditional CPU-based systems to conduct PHM. With increasing amounts of sensor data being captured, along with the desire to use larger neural networks and models with more parameters, a development workstation with four NVIDIA TITAN X GPUs was installed. Needing more performance and a more optimized AI software stack to shorten the time to solution, DAI then deployed an NVIDIA DGX Station™ and two DGX™ servers. NVIDIA DGX systems are purpose-built to meet the demands of enterprise AI and data science, delivering the fastest start in AI development, effortless productivity, and revolutionary performance—for insights in hours instead of months.

KEEPING AIRCRAFT IN THE AIR WITH DEEP LEARNING

Another aspect of predictive maintenance involves analyzing maintenance records and parts replacement history. Using AI to select the right maintenance code helps improve data quality to increase equipment uptime. In the past, Lockheed Martin would require a significant expenditure of human effort to look at freeform text maintenance reports

> "We achieved a 10 percent boost in accuracy overnight because of the greater ability to train and tune parameters on the DGX."
>
> Sam Friedman
> Senior Data Scientist,
> Data Analytics
> Innovations Group
> Lockheed Martin

and determine failure codes manually. With some of their databases having over 8,000 distinct categories, this was no easy task.

The DAI Group used maintenance data, along with usage inputs, to develop a predictive model that determines the risk level of parts that could drive an aircraft-on-ground (AOG) event. The model even provides recommendations to the operator for what they need to do and a five-day warning so they have time to secure the part for replacement.

Friedman's team used deep learning to automate this task with higher levels of confidence than their CPU solution. "We achieved a 10 percent boost in accuracy overnight because of the greater ability to train and tune parameters on the DGX. On the CPU server, we were running an XGBoost model. Using the DGX server, we were able to raise overall classification accuracy by leveraging an ensemble method, combining inputs from our original model and from neural network models built with intensive embedding layers and from neural networks that were more fit to task at hand, including the sequence-to-sequence (seq2seq) model," Friedman said. Today, nine out of ten maintenance records can be automatically classified without a subject matter expert, with over 95 percent accuracy. When you're dealing with hundreds of thousands of records, that translates to significant cost savings.

"Before, our models were limited in size and expressivity as they were CPU bound," said Friedman. "Once we moved to GPU-based models on DGX systems, we could leverage more complex neural networks, including the newer recurrent neural networks (RNNs) and transformer models."

A NEW ERA IN MAINTENANCE

"The reliance on traditional analytics techniques and a legacy CPU-only compute infrastructure that lacks the needed processing power to analyze the volume and variety of data fast enough has impacted our ability to deliver scalable and robust predictive maintenance solutions. With the adoption of GPU-accelerated AI algorithms and DGX systems, we're able to improve failure prediction performance while reducing false alarms by building context-aware AI models at a scale, which wasn't possible before. Today, whether we're training our own models or grabbing a trained model off the shelf like LSTM, RNN, or 1DCNN for time-series data or BERT for

“The powerful capabilities of DGX systems enabled us to take advantage of the full gamut of deep learning methods to solve our day-to-day problems.”

Sam Friedman
Senior Data Scientist,
Data Analytics
Innovations Group
Lockheed Martin

textual data, we can easily approach the problem knowing we have enough compute on the DGX,” said Friedman.

With the large amounts of data handled by Friedman’s team, the ability to train neural networks in parallel and split out layers across multiple GPUs within the DGX is important. The NVIDIA® NVLink® architecture within the DGX enables ultra-high bandwidth and low-latency connections between all the GPUs, so training runs much faster. And the team can speed up hyperparameter tuning by running jobs on multiple instances on different GPUs at the same time. According to Friedman, “By spreading tasks across powerful GPUs in the DGX, we no longer have to downsample our data or aggregate it and reduce its expressiveness.” Although Friedman’s team also uses cloud instances, DGX systems remain their sole resource for GPU compute, as it’s harder to do model and data parallelism across cloud instances.

Friedman’s team was once running several CPU-based models in the cloud, on the development box, or on high-performance computing (HPC) clusters. Since they moved to GPU-accelerated models on DGX systems, they saw massive speedups. For instance, training a corpus of a few million documents was at least 18X faster on a DGX than their CPU-based server. Additionally, in moving from the development box to the DGX Station, Friedman’s team saw an almost 2X speedup in training time with no change to architecture or code. “It was a complete change,” said Friedman. “Speedups from the DGX system were like night and day.”

Today, Friedman’s team uses DGX servers for production, where they have models already trained and hyperparameters tuned and optimized, and for retraining the models on a regular basis. Complementary to this is the DGX Station, which they use for experimentation and model development.

TAKING PREDICTIVE ANALYTICS TO THE NEXT LEVEL

“Before we had access to GPUs in DGX systems, a lot of natural language processing (NLP) models were simply not an option to us. We found that we couldn’t use a lot of the more complex and advanced neural networks—including time-series forecasting using RNNs and sequencing modeling—on CPUs or weaker GPUs. Today, we’re able to pre-process hundreds of thousands of handwritten maintenance records, which are then used along with other sources to build a robust dataset for predictive

maintenance. Innovation is no longer stifled. Our developers can engage in one-off projects without needing to worry about waiting for days and days for the results to come back," said Friedman. The approaches developed in commercial space have already been applied to the sustainment of land, air, sea, and space platforms of the Department of Defense (DOD), such the Sikorsky S-92 helicopter, the F-22 and F-35 stealth aircrafts, the CH-53E King Stallion helicopter, and the U-2 Dragon Lady jet. Lockheed Martin has already developed over one hundred NLP algorithms, which they've provided to the government.

Beyond predictive maintenance of fleets, AI-infused sustainment with usage-based lifing and condition-based maintenance is broadly applicable across industries: Food and beverage companies are faced with strict regulatory standards where broken equipment can lead to major health issues or food spoiling. Hospitals need to maximize service time of medical devices like ultrasound and MRI systems and keep these systems reliable and safe. Banks need to ensure ATM machines remain up and running and don't break down in the middle of a transaction. Manufacturing companies—such as railway infrastructure, computer numerical control (CNC) machines in factories, paper mills, and fleets of buses and automobiles—need to ensure early fault detection and keep production lines up. And just about every industry has data centers and IT hardware that can benefit from predictive maintenance to improve service availability and minimize data loss. Unplanned downtime is crippling—to production, to safety, even to brand. As roughly $647 billion is lost annually on a global basis by manufacturers across all industries,[4] using AI to gain insights into asset health can have a significant impact, leading to tremendous cost savings, improved operational efficiencies and service availability, increased safety, and more satisfied customers.

To learn more about NVIDIA DGX systems, visit: **www.nvidia.com/dgx**

www.nvidia.com

1 National Institute of Standards and Technology, U.S. Department of Commerce. The Costs and Benefits of Advanced Maintenance in Manufacturing. April 2018.
2 Plant Engineering. 2019 Maintenance Study. March 2019.
3 U.S. Department of Defense. FY 2019 Defense Budget.
4 Robert Rice and Richard Bontatibus. International Society of Automation. Predictive Maintenance Embraces Analytics.