Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

SUPPORTING FAB OPERATIONS USING MULTI-AGENT REINFORCEMENT LEARNING

Ishaan Sood[1], Abhinav Kaushik[1], Tom Bulgerin[1], P. Kumar[1], Abdelhak Khemiri[1], Jasper van Heugten[1], Johnny Chang[2], Sam Hsu[2], and Jeroen Bédorf[1]

[1]minds.ai Inc., Santa Cruz, CA, USA
[2]Micron Technology, Inc.,Taoyuan, TAIWAN

ABSTRACT

Over recent years semiconductor operations have grown in scope, dynamics, and complexity. Consequently, advanced scheduling algorithms and manufacturing engineers can no longer quickly estimate the best performing schedule. In this paper we present how machine learning, in real time, can be used to augment and support manufacturing engineers. The presented results, obtained from production deployments in Micron's fabs, show that AI assisted scheduling can improve Key Performance Indicators with little to no downside.

1 INTRODUCTION

Modern 300MM semiconductor fabrication facilities face challenges such as complex hierarchical structures, high financial stakes, and dynamic processes. These facilities contain thousands of tools with unique capabilities, resulting in intricate dispatching schedules. Frequent schedule updates are required due to maintenance, failures, and shifting priorities. This constant updating makes it impractical for Manufacturing Engineers (MEs) to manually optimize scheduling parameters.

Efforts have been made to use data science and machine learning to improve scheduling and operating margins (Lee et al. 2020; Lee et al. 2019a). This work demonstrates how reinforcement learning (RL) can aid manufacturing engineers (MEs) in scheduling decisions. A key goal is to improve MEs' well-being by mitigating time constraints from dynamic business priorities. Machine learning helps MEs by handling routine tasks, allowing them to focus on rare events. RL uses simulations and historical data to optimize scheduling based on current conditions, offering MEs suggestions to enhance Key Performance Indicators (KPIs) like throughput, tool idle time, and critical queue times.

2 METHOD

Most semiconductor job scheduling rules use advanced heuristics to improve KPIs, incorporating MEs' domain knowledge. The fabs used in this work use an Adaptive Dispatching Rule (ADR) from Li et al. (2013), where each idle tool assigns a score to waiting jobs based on weighted factors; the job with the highest score gets scheduled. MEs pre-assign weights for each factor, and factor values for each job are assigned dynamically by the scheduler. Due to the impracticality of assigning unique factor values to hundreds of tools, we developed a multi-agent RL solution that dynamically assigns factor weights based on each tool's Work-in-Progress (WIP) conditions to optimize fab performance. Compared to prior single-agent RL methods, our multi-agent solution offers finer control (Lee et al. 2019b).

MEs supply the KPI reward weights based on current demand and production plans, such as improving fab throughput or reducing product cycle time. Training uses nearly 2 months of real fab scheduling data, with the last 2 weeks for evaluation. The agent's performance is compared against the baseline weight values set manually by the ME.

3 RESULTS

An extensive User Acceptance Test (UAT) phase assesses solution quality, involving RL engineers, MEs, and the global operations team. Predefined acceptance criteria guides the agent, initially running in shadow mode to suggest changes and expected results. MEs decide shift-by-shift whether to follow these recommendations and eventually the agent runs autonomously. This paper presents results from an agent developed to increase the number of wafers processed by bottleneck tools in Micron's facilities. The agent is rewarded for improving tool efficiency without increasing queue time breaches. It adjusts actions per shift based on fab behavior, either pushing more work to tools or changing settings to reduce idle time. Figure 1 shows the per shift increase in tool throughput, with 0 indicating no improvement or negative effects, reverting to MEs' default settings.

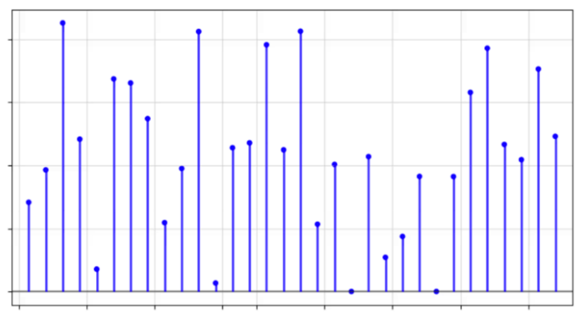

Fig. 1. Comparison between baseline and agent results. The x-axis shows the shift number (2 per day), the y-axis the percentage improvement of numbers of wafers processed by the bottleneck tools.

4 CONCLUSION

This work shows that advanced machine learning techniques, combined with the experience of MEs, has led to increased wafer productions and improved the quality of life of the engineers with virtually no downsides. At the same time, it allows MEs to focus directly on KPI settings instead of tool factors. The dynamic nature of the semiconductor fabrication process results in that the agents have to be monitored and updated. The updating process is part of a robust and automated Machine Learning Operations (MLOps) process and as such does not incur any additional work on the MEs. The MLOps process takes care of retraining triggers, rolling out updates, and automating the evaluation/analysis process for all the fabs in which the agent is deployed.

REFERENCES

J. -H. Lee, Y. Kim, Y. B. Kim, B. -H. Kim, G. -H. Chung and H. -J. Kim, "A Simulation-Based Sequential Search Method for Multi-Objective Scheduling Problems of Manufacturing Systems," 2020 Winter Simulation Conference (WSC), Orlando, FL, USA, 2020, pp. 1943-1953

J. -H. Lee, Y. Kim, Y. B. Kim, H. -J. Kim, B. -H. Kim and G. -H. Chung, "A Sequential Search Framework for Selecting Weights of Dispatching Rules in Manufacturing Systems," 2019 Winter Simulation Conference (WSC), National Harbor, MD, USA, 2019a, pp. 2201-2211

Li, L., Zijin, S., Jiacheng, N. and Fei, Q. Data-based scheduling framework and adaptive dispatching rule of complex manufacturing systems. Int J Adv Manuf Technol 66, 1891–1905 (2013). https://doi.org/10.1007/s00170-012-4468-6

W.J. Lee, B.H Kim, K. Ko, and H. Shin. Simulation based multi-objective fab scheduling by using reinforcement learning. In Proceedings of the 2019 Winter Simulation Conference. IEEE Press, 2019b, 2236–2247.