DMBSI-CARE: A Causal Process-Mining Framework for Root-Cause Analysis and Bottleneck Attribution in Wafer Fabrication

Mohammad Sharifur Rahman [1,*], M M Manjurul Islam [1], Karl McCreadie [1], Saugat Bhattacharyya [1], Cormac McAteer [2], Bryan John Baker [3], Nuala Parker [2] and Girijesh Prasad [1]

[1] Intelligent Systems Research Centre, Ulster University, Londonderry BT48 7JL, UK
[2] Seagate Technology, Londonderry BT48 0BF, UK
[3] Seagate Technology, One Disc Dr, Minnesota 55435, USA
* Correspondence: rahman-m11@ulster.ac.uk

ABSTRACT

Cycle-time root-cause analysis in wafer fabrication still relies on descriptive severity rankings, yet severity does not equate to causal impact, does not reveal how delays propagate, and cannot explain why a particular lot was slow. DMBSI-CARE closes these gaps by seeding a structural causal model with a process-mining Directly-Follows Graph, estimating per-step effects with double machine learning, decomposing each effect through severity-aligned mediators, and aggregating them into a per-lot causal cycle-time bill of materials. On a log of 1,406 lots and 7.9 million events, it surfaces leverage-step severity-ranking misses and attributes 73–99% of the indirect effect to processing variability rather than queueing.

KEYWORDS: root-cause analysis; causal inference; process mining; cycle time; semiconductor manufacturing

1. INTRODUCTION

In a modern wafer fabrication facility, many wafers undergo more than 200 processing steps over several months, and most of that time is spent waiting rather than processing. Cycle time drives work-in-progress capital, customer lead time and buffer inventory, so reducing it is among the most valuable improvements a fab can pursue [1]. Root-cause analysis has historically relied on severity ranking [2]: a multi-axis composite score is computed for each step, weights are calibrated to measured cycle time, and the highest-ranked steps become intervention targets; the Dynamic Multi-Criteria Bottleneck Severity Index (DMBSI) [3] is a recent example.

Severity indices are descriptive, which limits causal decision-making in three ways: severity tells where time accumulates, not whether intervening on a step would reduce cycle time; it is a single number, silent on whether delay propagates through queueing, processing-time variability or rework; and it is fab-wide, so it cannot explain why one specific lot was slow. This work presents DMBSI-CARE (DMBSI — Causal Attribution and Root-cause Explanation), which supplies the missing causal estimate, mechanism decomposition and per-lot attribution.

2. AIM AND OBJECTIVES

The aim of this work is to turn descriptive severity ranking into causal accounting for wafer-fab cycle time. Four objectives support this aim: (i) to construct, for each candidate bottleneck step, a structural causal model [4] whose graph topology is constrained by a Directly-Follows Graph mined from the event log [5], so that a valid backdoor adjustment set is obtained by graph traversal at constant cost per step; (ii) to decompose each per-step effect through mediators aligned with the DMBSI sub-scores, making every indirect effect interpretable per severity dimension; (iii) to decompose each individual lot's measured cycle time into per-step causal contributions — a causal cycle-time bill of materials; and (iv) to establish sufficient conditions under which the construction recovers the true causal structure and yields a √N-consistent estimate of each effect.

3. METHODOLOGY

DMBSI-CARE is a five-stage pipeline. Stage 1 reduces the raw event log — 7,916,117 work-in-progress transactions covering 1,406 lots, 231 distinct steps and 650 production days from a high-volume fab — to a 153,379-row step-lot panel. For each lot–step pair the panel holds the treatment (cumulative wait at the step, in hours), the outcome (the lot's total cycle time), three context covariates (priority class,

© The Author(s) 2026
D. Nikam et al. (eds.), *Proceedings of the Global Symposium on Adaptive Manufacturing (GSAM 2026)*, Advances in Engineering Research 308,
https://doi.org/10.2991/978-94-6239-813-9_13

route family and product code) and three mediators that double as DMBSI sub-scores: queue depth (S2), within-step processing variability (S3) and rework hours (S5). Ten candidate treatment steps were drawn from the intersection of the DMBSI top twenty and the steps with sufficient graph out-degree. Stage 2 builds a per-step directed acyclic graph whose edges are filtered by the Directly-Follows Graph at threshold $\tau = 50$; traversal of this graph returns the structural backdoor set {PRIORITY, ROUTE, PROD} for every candidate step. Stage 3 estimates the per-step average treatment effect (ATE) by double machine learning with cross-fitting [6]. Stage 4 splits each ATE into direct and indirect effects through S2, S3, and S5. Stage 5 multiplies each step's estimated effect by the lot's actual wait to obtain per-lot attribution, then aggregates it by product. E-value sensitivity bounds, negative-control placebo tests and a stability analysis over $\tau \in \{25, 50, 100, 250\}$ run alongside the estimation stages.

4. RESULTS

Severity rank does not track causal impact. The two largest cleanly identified effects — 18.55 h of cycle time per hour of step wait (95% CI [13.23, 23.87]) and 18.23 h/h ([14.04, 22.41]) — belong to steps ranked only eighth and ninth by DMBSI severity, an order of magnitude above the top five cluster, whose effects span 1.67 to 2.90 h/h. A one-hour wait reduction at the eighth-ranked step is worth roughly six one-hour reductions at the median top five step, so a severity-only pipeline would have prioritised the wrong steps. Across the DMBSI top five, 73–99% of the indirect effect routes through processing variability (S3); queue depth accounts for only 6–15%. This qualifies the queueing-theoretic intuition that queue length drives cycle time in re-entrant flows, and it carries a direct policy implication: variance-reduction scheduling should be prioritised over queue-depth control at these steps. Separately, the step with the highest severity contribution and the highest graph centrality returned an estimate of 177.8 h/h whose confidence interval spans [−264.6, +620.2]; too many downstream steps depend on its wait for any observational adjustment to disentangle its effect, and the framework flags it as non-identifiable rather than reporting an unusable number. Finally, the per-lot bill of materials reconstructs the measured cycle time of the five slowest lots (up to 10,687 h) to within 5%, with the dominant contributing step varying by product.

5. CONCLUSIONS

DMBSI-CARE turns a descriptive bottleneck ranking into a causal account of wafer-fab cycle time: which steps cause it, through which mechanism, and how each individual lot's hours are built up step by step. On the studied fab log, the framework located leverage points, the severity ranking missed, identified processing variability rather than queueing as the dominant mechanism, and flagged a high-severity hub whose effect cannot be measured observationally. An identification theorem certifies the construction under explicit, checkable conditions. Cross-fab replication and real-time deployment within a cyber-physical fab architecture are the natural next steps. Funding support from the UKRI Strength in Places Fund project (81801), Smart Nano-Manufacturing Corridor, is gratefully acknowledged.

REFERENCES

1. Mönch, L.; Fowler, J.W.; Mason, S.J. Production Planning and Control for Semiconductor Wafer Fabrication Facilities; Springer: New York, NY, USA, 2013.
2. Roser, C.; Nakano, M.; Tanaka, M. Shifting bottleneck detection. In Proceedings of the Winter Simulation Conference, San Diego, CA, USA, 8–11 December 2002; pp. 1079–1086.
3. Rahman, M.S.; McCreadie, K.; Bhattacharyya, S.; Islam, M.M.M.; McAteer, C.; Baker, B.J.; Parker, N.; Prasad, G. Dynamic multi-criteria bottleneck severity index (DMBSI) for semiconductor wafer manufacturing: A genetically optimised framework for reentrant production systems. arXiv 2026, arXiv:2607.24819. https://doi.org/10.48550/arXiv.2607.24819
4. Pearl, J. Causal inference in statistics: an overview. Statist. Surv. 2009, 3, 96–146.
5. van der Aalst, W.M.P. Process Mining: Data Science in Action, 2nd ed.; Springer: Berlin, Germany, 2016.
6. Chernozhukov, V.; Chetverikov, D.; Demirer, M.; Duflo, E.; Hansen, C.; Newey, W.; Robins, J. Double/debiased machine learning for treatment and structural parameters. Econom. J. 2018, 21, C1–C68.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

