*Proceedings of the 2025 Winter Simulation Conference*
*E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.*

# A PRODUCT DIGITAL TWIN: HOW SIMULATION EVOLVED ALONGSIDE THE WORLD'S MOST ADVANCED GROCERY PICKING ROBOT

Robert Backhouse[1]

[1]Ocado Technology, Hatfield, UK

## ABSTRACT

Ocado Technology developed On-Grid Robotic Pick (OGRP) - a robotic picking solution for grocery fulfillment centers - with coordinated innovation across hardware, software, and system design. To support this development, we built a low-fidelity simulation model early and evolved it alongside the product. As OGRP matured, assumptions from each area were progressively integrated into the model, while simulation insight shaped design decisions - forming a "product digital twin". We describe how this approach of co-evolving the simulation supported exploration of a large and uncertain design space - aligning development streams, guiding rapid iteration, and enabling early development - ultimately driving faster delivery with lower risk. OGRP is now deployed around the world, picking over a million items per week and continuing to scale. This case study illustrates how evolving, system-level simulations can drive coordination in complex, fast-changing engineering projects.

## 1 INTRODUCTION: THE USE CASE FOR SIMULATION

Ocado Technology is the innovation factory behind the technology used by 13 online grocery operations around the world. Grocery fulfillment is a challenging market, with thin margins, large orders and tens of thousands of different products. To maximise efficiency, our fulfillment centers are built around the Hive, a proprietary automation system where a fleet of robots ("bots") move stock and customer orders around a dense storage grid. Precise coordination of hundreds of high-speed bots enables tens of thousands of items to be picked per hour, and simulation has always been crucial to understanding the complex emergent behavior of the system.

While the Hive offers industry-leading performance, the final step - manually picking items into customer bags - still accounts for a significant portion of the total labor cost. Automating this promised another step-change in warehouse efficiency. Robotic arms were a promising solution, but industry standard rigidly-programmed arms are poorly suited to the needs of grocery, with its highly varied product range, fragile items, and the need for dense but careful packing. So Ocado Technology developed OGRP - a robotic picking system specifically designed for the demands of grocery fulfillment.

OGRP combines developments in hardware (e.g. arm, end effector, supporting infrastructure), software (e.g. arm controllers, interaction with our bot orchestration system), AI (e.g. computer vision, behavior cloning) and process (e.g. site design, installation). These were all developed in parallel, but with each design heavily dependent on decisions in the other streams; so the development of OGRP can be characterized as exploring a complex, highly-coupled design space with high uncertainty. Efficient searching required rapid iteration, evaluating local decisions in the context of the overall system. Our simulation model provided this evaluation function, allowing the development teams to focus on the most promising directions, avoid dead ends, and move toward a better end product.

## 2 DEVELOPING A SIMULATION IN PARALLEL WITH THE PRODUCT

We already had a discrete-event simulation of the Hive system - written in Java using in-house libraries, in order to integrate with production control code and with our own tools for configuring and executing simulation experiments. So we began by extending our Hive simulation with a simple model of robotic picking. Despite a high degree of uncertainty, we assembled placeholder assumptions from each product stream into a consistent conceptual model of how the final solution might look. This early prototype allowed us to test the interaction between arms and bots, and confirm that the OGRP concept was viable - the arms could attain a reasonable throughput without excessive congestion for the bots.

From here, the simulation mirrored the evolving product, growing in fidelity and reflecting our best guess for OGRP end state. Individual assumptions were owned by the relevant product teams and combined to a central model of the system behavior. At any point, simulation could provide each stream with estimates for how their local design choices would affect overall system performance. This shared, evolving simulation acted as a “product digital twin” - a shared model of the intended end-state behavior that mirrored our latest understanding and enabled coherent, data-driven decisions.

## 3 BENEFITS OF THIS APPROACH

Early and iterative simulation led to valuable conceptual modeling conversations throughout the project. The initial modeling process helped to surface inconsistencies between product streams, identifying problems much earlier than they would have otherwise been discovered.

Once the initial prototype was in place, it provided consistent feedback on how local decisions affected the overall system behavior. Simulations highlighted when targets conflicted with other parts of the system or when the system failed to meet performance targets, avoiding development dead ends. Exploring a range of scenarios revealed which assumptions were the most impactful, allowing teams to focus on the most important performance drivers.

As the simulation gained fidelity alongside the maturing product, it became increasingly capable of answering specific, detailed questions. This naturally supported the shift toward lower-level decisions and finer trade-offs in the later stages of development.

Regularly simulating with a combined model revealed cross-cutting opportunities throughout the project. Outputs generated for one decision often exposed constraints or improvements elsewhere, which could then be assessed and prioritized immediately.

On top of the benefits around design space exploration, our “product digital twin” also allowed us to bring development forward. By testing algorithms and layout strategies against the evolving model, we could start development long before physical prototypes were available.

Finally, updating the simulation model acted as a visible point of commitment, reinforcing alignment across different development streams for each planned change.

## 4 OUTCOMES

OGRPs are now in active use, deployed in warehouses around the world. Arms and bots coordinate through the orchestration system developed and tested against the simulation. Our Luton fulfilment center was the first to be designed using the OGRP model, and the site is successfully handling the picking of over a million items per week by robotic arms. The OGRP technology and its deployment continue to evolve, and the simulation model is still being updated as the product itself changes. Ocado Intelligent Automation is starting to apply our robotics and orchestration systems to a range of non-grocery domains, bringing a stream of new use cases and optimization challenges for OGRP. The simulation remains a critical tool to support confident decisions, shape product development, and drive further optimizations for this highly adaptable robotic picking solution.