Simulation and Reinforcement Learning for Flexible Manufacturing Systems

Product: Generic Electronic products

KPI: production throughput

Geography: United States, California

Industry: Electronics

Project Time: 8 months

Customer: Undisclosed

The customer, who lives in the intersection of product development and material science and manufacturing, has the vision of understanding how the assembly of small form factor consumer and digital health devices may look like, as they integrate the latest set of technologies available within the industry 4.0 paradigm. In other words, they are interested in studying the mechanics and operational dynamics of flexible manufacturing systems.

Linear assembly lines are commonly underpinned by expensive human labor, and a fixed orientation process. Fixed lines are robust but inflexible and expensive to modify when the designs and the needs of businesses change. Some of the attributes that are essential components of a factory that takes advantage of the enabling technologies available, rendering it holistically efficient are:

- Adaptable
- Evolving
- Intelligent
- Flexible
- Efficient
- Connected

While the larger body of work takes all these into account, the interest is to understand how flexibility and adaptability under an intelligent and evolving framework can lead to greater efficiency.

Vision and Intent

Based on the facets mentioned above, the idea is to set out to architect a platform that would enable an automated assembly process. This means that it will be:

- Flexible: can run different orders of processes with minimal changeover.
- Adaptable: is able to adjust to the variety of processing needs and component requirements in short order using the same platform

Even though the assembly line is composed by a set of different manufacturing processes, Figure 1 shows the general architecture of one of the most important pieces of the system. Composed by pickheads, process heads and gantries.

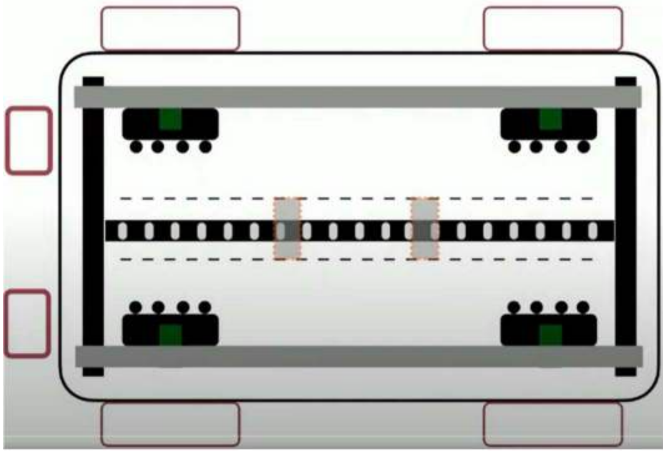

Figure 1. Architecture

Why Simulate?

While working on the architecture, it was clear that there were a finite number of rounding combinations, however an optimized solution for the relative motion paths for each agent was not intuitable. The complex interplay between resources is more apparent when one takes into consideration the effect of failures and errors within the system. The other aspect here is more mechanically themed, in reference to the large variety of products we would see coming down the line, it was important for the customer to have a test bed that would quantify the relationship between mechanical architecture and its ability to meet throughput requirements.

Additionally, a simulation allows to observe the effects a change in degree of freedom available or the placement position would have on throughput.

Description of the System

Figure 2 shows how the layout looks like in the simulation setting. In the center there is a conveyor that moves parts that are represented with different colors. Each part needs to go through a set of processes and the number of processes a part can undergo is variable.

Through the conveyor there are also 3 process heads, and each process head performs a different process on the parts. The number of process heads can also be variable depending on the type of part used.

Additionally, there are 4 pickhead structures that can also vary in number, size and can have different number of pickheads in it. On the image, only 4 pickhead structures with 4 pickheads each is shown. Each pickhead grabs components from bins and each pickhead has a unique type of component in it. Each pickhead also has a camera that verifies the correct positioning of a component on a part. The pickheads move horizontally through axes that can transport pickhead structure groups to the bins or to the conveyor for the different activities.

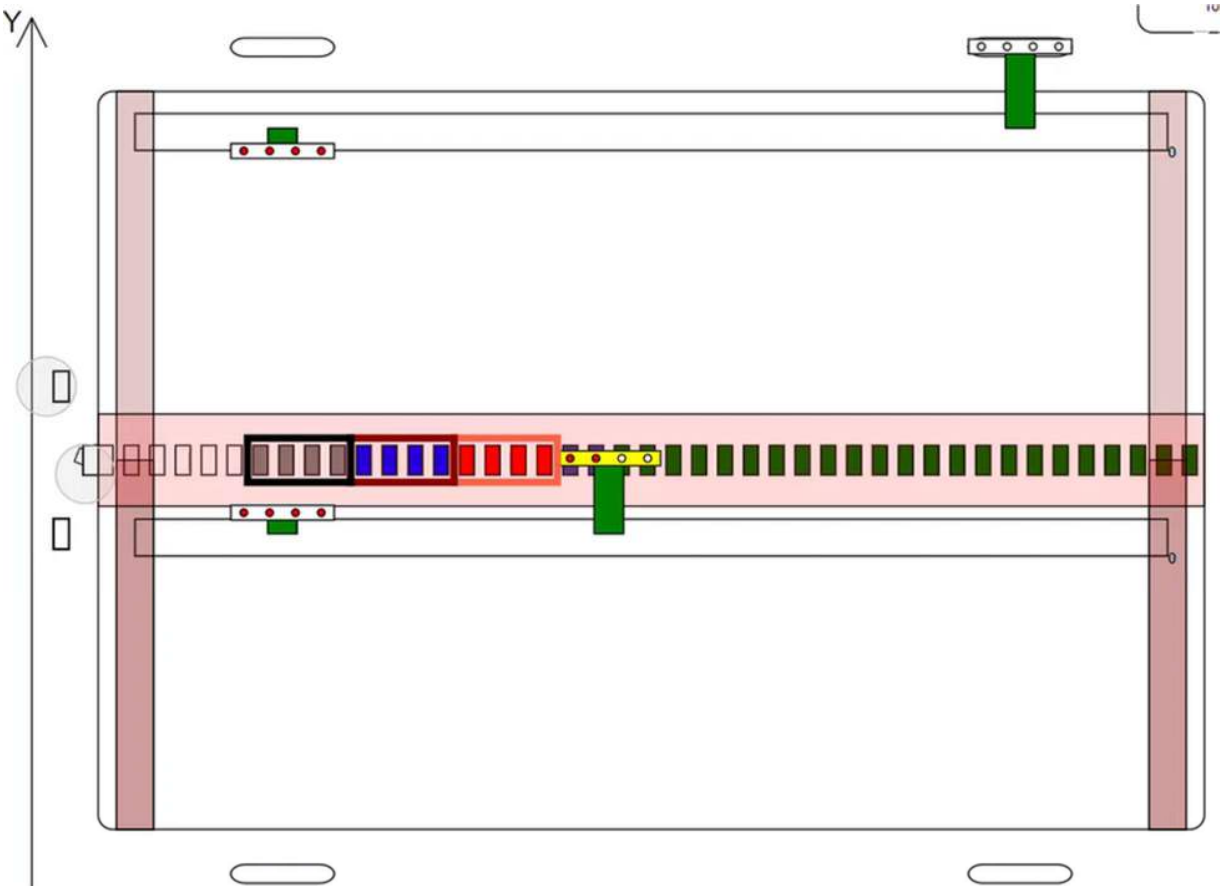

Figure 2. Simulation Model

This robotic system is just one of many other machines that are present in the assembly line, but it's the one that can produce bottlenecks and hence is the most important one to improve throughput.

It was important for the customer to be able to modify the layout for testing, and figure 3 shows an example of the things that the simulation was able to modify from the physical characteristics of the machine in order to test:

- Pickhead diameter
- Distance between pickheads
- Structure lengths
- Distance to axes
- Etc.

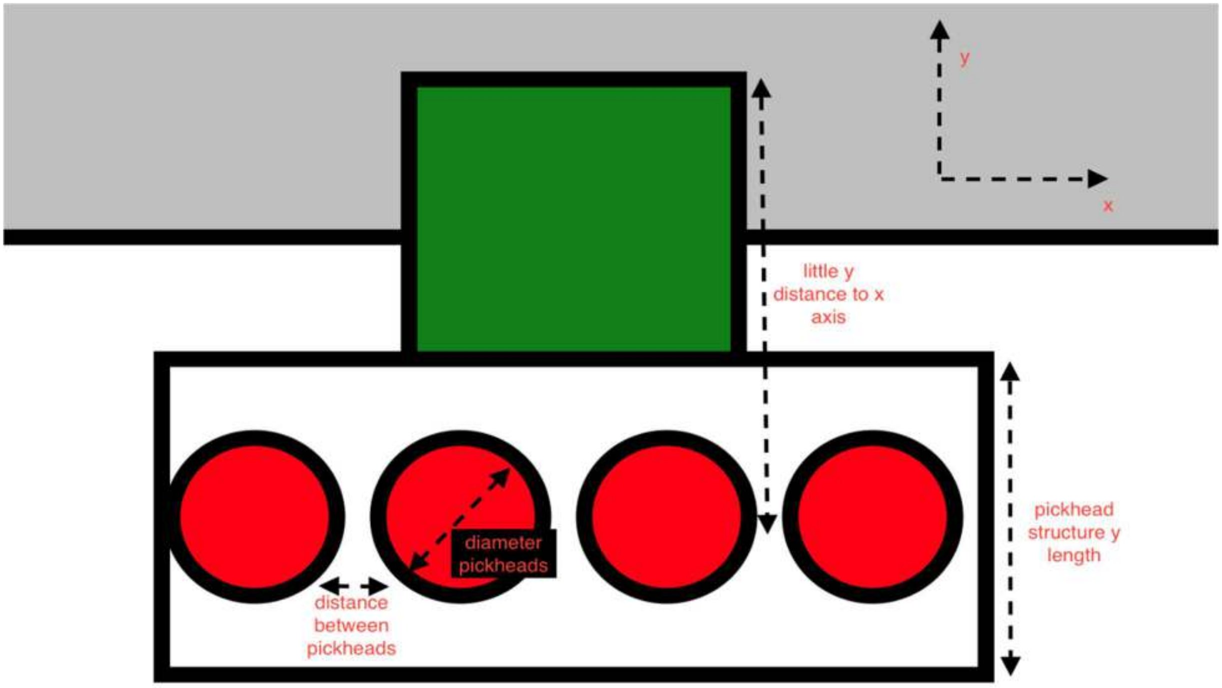

Figure 3 – Pickhead structure

Throughput results

The baseline for the system behavior is a set of rules mimicking the behavior of a PLC that would guide all the system agents into completing all parts effectively. These algorithms allowed a throughput of 518 units per hour.

With this baseline in mind, we developed for a few months different reinforcement learning solutions that would improve the coordination of the agents in order to maximize the throughput getting a result of 870 units per hour, which is a 68% improvement. Considering the revenue for an individual part, this is translated into additional revenue of around $100,000 USD per day or more than 30 million dollars per year.

Many different architectural and mechanical structures were tested during this analysis, that lead to improvements up to an additional 1% on the artificial intelligent solution with a hybrid between reinforcement learning and expert rules.

Errors and failures

In the final phase of this project, we added errors and failures that can occur during the process of the parts. There are 3 types of errors:

- Placing the component incorrectly
- Having a defective component
- The part becomes defective.

The addition of failures and errors required a complete redesign of the reinforcement learning solution due to the several additional complexities associated to it, and reduced the throughput to 740 parts per hour, which is still 40% better than the base case.

How was RL applied?

This is a complex simulation, and to implement reinforcement learning properly, we needed 3 key ingredients:

- Framing the problem as a multi-RL agent problem: the dynamic nature and uniqueness of each element of the manufacturing process makes this necessary
- Filter out illegal or invalid actions: some agents need to wait till it’s their turn so some manual restrictions during Reinforcement learning was necessary
- Mix and match global and individual rewards to teach the RL policy desired behavior.
    - Maximize throughput.
    - Minimize unfinished products.

There were 10 different agents for the RL process: 3 process heads move left or right, 4 pickhead structures can move left or right, 2 axes can move up and down and the conveyor can move or stop.

In general, to obtain a RL policy, the simulation was run 10,000 times in about 2 hours.

Conclusions

The work developed using simulations to improve automated flexible manufacturing systems and robotics has shown to be beneficial not only for the revenue of the customer, but also for the advancement of the technology, the opportunity to generate better value for the clients and the learning that leads to an incentive to apply these techniques in a multiplicity of similar endeavors. In the beginning of this work, it was apparent that expert-rule systems were going to give the necessary results, but it was easy to discover that by using artificial intelligence, the improvement of the system holistically as well as the improvement of the individual agents that compose this system is limitless.