

Multi-Stage Laptop Manufacturing System Simulation Study

Group 29

Jasper Zhou (s2477464), Chris Liu (s2156635) Shihan Wang (s2525282), Yiyang Cao (s2306471)

April 23, 2026

Summary

This report develops a discrete-event simulation model in Simul8 for a simplified multi-stage laptop manufacturing system operating on a 24/7 continuous schedule with a one-month results collection period. The baseline results show that the system does not perform well under the current settings, with an average cycle time of about 1575 minutes and the case manufacturing line identified as the main bottleneck. A series of five progressive experiments are then conducted to diagnose and address the layered bottleneck structure. The final improved model reduces the mean cycle time by 21.1%, increases the throughput from 2.82 to 7.14 laptops per hour, and lowers the rejection rate by 15.1 percentage points. These improvements are statistically significant ($p < 10^{-12}$, Welch t -test), and the report also discusses the main limitations of the model and suggests directions for further improvements.

1 Introduction

1.1 Background

The system under study is a simplified laptop manufacturing process with several interconnected stages. Each order may require multiple laptops upon arrival, and each laptop is produced through two parallel streams: motherboard manufacturing and case manufacturing, which are later integrated in the main assembly stage.

The system operates under a number of constraints. A capacity limit restricts the total number of laptops in the system, and different stages require specific types of workers and machines. Processing times vary across stages and follow a range of distributions, including uniform, triangular, and empirically observed data. In addition, features such as blocking, rework, and inventory control for key raw materials introduce further complexity and stochastic behaviour to the system.

Due to these characteristics, it is difficult to evaluate system performance analytically. The company is particularly interested in key performance measures such as the average cycle time, as well as in identifying the bottlenecks and inefficiencies. To address this, a discrete-event simulation model is developed to represent the system and evaluate its performance. The model is then used to explore possible improvements to how the system operates.

The key performance indicators (KPIs) used in this study include: (i) *mean cycle time*, the average time from order entry to boxing completion per laptop; (ii) *throughput*, the number of completed laptops per unit time; (iii) *order rejection rate*, the proportion of arriving orders turned away due to the capacity constraint; (iv) *queue waiting times*, which identify where delays are concentrated; and (v) *resource utilisation*, which reveals the tightest resource constraints. These KPIs are evaluated using 95% confidence intervals constructed from independent replications.

1.2 Objectives

The aim of this project is to develop a discrete-event simulation model to represent the laptop manufacturing system and to provide a basis for evaluating its operation and identifying potential improvements.

The analysis is carried out in four main steps:

- I. Statistical input analysis is conducted on the available data to identify suitable distributions for processing times across different stages and interarrival times. The resulting distributions are then used to modify the model, ensuring a realistic representation of the system.
- II. A simulation framework is constructed to capture the key processes of the system, including order arrivals, parallel production of motherboard and case components, assembly operations, resource constraints, and queueing mechanisms such as blocking and rework.
- III. The baseline system is evaluated through output analysis, including warm-up determination, replication planning, and diagnostic analysis based on queue statistics and resource utilisation.
- IV. The simulation results are then used to identify bottlenecks and propose targeted improvements, with the effectiveness of these changes validated through statistical testing.

1.3 Report Structure

In this section, we outline the structure of the report and summarise the main steps of the simulation analysis.

First, in Section 2, we conduct input analysis by fitting appropriate probability distributions to the observed data and validating key assumptions such as independence. The resulting distributions are used to modify the simulation model.

Next, in Section 3, we construct a discrete-event simulation model that captures the main processes of the system, including order arrivals, parallel production stages, resource constraints, and queueing mechanisms such as blocking and rework.

In Section 4, we evaluate baseline system performance using the simulation output. This includes warm-up analysis, determining the number of replications, and examining queue statistics, workstation utilisation, and resource usage. In Section 5, we diagnose system bottlenecks and explore improvements, comparing the baseline and improved configurations through a Welch two-sample t -test.

Finally, Section 6 discusses the limitations of the model, and Section 7 concludes the report. The overall analytical workflow is illustrated in Figure 1.




Figure 1: Analytical workflow of the study.

2 Input Analysis

The input analysis is based on fitting suitable distributions to the data. For each dataset, a set of candidate distributions is fitted to the data. These models are then compared using goodness-of-fit measures, including the Akaike Information Criterion (AIC), the Bayesian Information Criterion (BIC), and the Kolmogorov–Smirnov (KS) test, in order to select the most appropriate distribution. Once selected, the parameters of the chosen model are estimated using maximum likelihood estimation (MLE).

Four input datasets are analysed: *interarrival times*, *service times for the initial phase*, *service times for placing the keyboard and mouse*, and *service times for assembling the case (aluminium plates)*. Interarrival times are assumed to follow an exponential distribution, as specified in the system description, and this is checked against the data. For the remaining service time datasets, no prior distributional assumptions are imposed. Instead, an empirical modelling approach is adopted to identify the most appropriate probability distributions.

2.1 Candidate Distributions

Based on the characteristics of the data (e.g. positivity and skewness), several candidate distributions were considered, including exponential, gamma, Weibull, lognormal, normal, inverse Gaussian, and log-logistic distributions. For each dataset, histograms with overlaid fitted densities were examined to visually assess how well the distributions fit the data before applying formal goodness-of-fit tests. The visual inspection confirmed that the selected distributions align closely with the empirical data shapes, and no systematic departures were observed.

2.2 Independence Analysis

Independence is an important assumption for the analysis, since both the goodness-of-fit tests and the use of independent and identically distributed (i.i.d.) random variates in the simulation require that successive observations are not serially correlated. To check for this, the lag-1 autocorrelation r_1 was calculated for each dataset. Under the null hypothesis of independence, $\sqrt{n}r_1$ is approximately standard normal for large n , so the test statistic is $z = \sqrt{n}r_1$ and the two-sided p-value is obtained from the standard normal distribution.

Variable	Lag-1 ACF	Test Statistic (z)	p-value	Conclusion
Interarrival Times	0.0209	0.2958	0.7674	Independent
Initial Phase Service Times	-0.0222	-0.2958	0.7674	Independent
Keyboard and Mouse Service Times	0.0069	0.0977	0.9222	Independent
Case Assembly Service Times	0.1254	1.7750	0.0759	Independent

Table 1: Lag-1 autocorrelation tests for input variables

For all datasets, the p-values exceed the 5% significance level. The case assembly dataset shows a somewhat higher lag-1 autocorrelation ($r_1 = 0.1254$, $p = 0.076$), but this is still non-significant at the 5% level. Overall, there is no strong evidence against the independence assumption for any of the input variables.

2.3 Distribution Fitting and Parameter Estimation

For each dataset, several candidate distributions were fitted and compared using three complementary goodness-of-fit measures: (i) the Kolmogorov–Smirnov (KS) test, which assesses the maximum distance between the empirical and fitted cumulative distribution functions (CDFs); (ii) the Akaike Information Criterion (AIC); and (iii) the Bayesian Information Criterion (BIC)¹. The KS test was performed with Monte Carlo adjustment to account for the fact that parameters are estimated from the same data. The AIC and BIC provide complementary model-selection criteria that penalise model complexity differently. The final distribution for each dataset was selected based on overall consistency across these measures. Parameters for the selected distribution were then estimated using maximum likelihood estimation (MLE).

Interarrival Times Several candidate distributions were fitted to the data, and the top three based on AIC, BIC, and KS test results are reported in Table 2.

Distribution	AIC	BIC	KS p-value
Exponential	1213.53	1220.13	0.321
Weibull	1214.04	1223.94	0.855
Gamma	1214.76	1224.66	0.686

Table 2: Goodness-of-fit results for interarrival times

¹Smaller AIC/BIC indicates a better fit for the same dataset. For the KS test, $p > 0.05$ indicates no rejection at the 5% level.

All these candidate distributions pass the KS test at the 5% level. The exponential model has the lowest AIC and BIC, while Weibull and gamma give slightly higher KS p-values.

Since the exponential distribution is also assumed in the system description and fits reasonably well, it is retained. Using MLE, the estimated rate parameter is $\hat{\lambda} \approx 0.1321$.

Service Times for the Initial Phase The chosen candidate distributions were fitted to the data. The three best-performing models according to AIC and BIC are summarised in Table 3.

Distribution	AIC	BIC	KS p-value
Weibull	717.92	727.82	0.2998
Normal	723.67	730.27	0.3052
Gamma	728.89	738.78	0.0635

Table 3: Goodness-of-fit results for initial phase service times

Three models are acceptable under the KS test at the 5% level. The Weibull distribution clearly outperforms the others in terms of AIC and BIC, whereas the normal distribution shows a marginally higher KS p-value.

Given this, the Weibull model is selected as it offers the most favourable overall fit to the data. The parameter estimates obtained via MLE are $\hat{\alpha} \approx 2.997$ and $\hat{\beta} \approx 4.414$.

Service Times for Placing the Keyboard and Mouse A set of candidate distributions was evaluated, and the three strongest models according to AIC and BIC are presented in Table 4.

Distribution	AIC	BIC	KS p-value
Exponential	1246.37	1252.97	0.5080
Gamma	1247.94	1257.84	0.7709
Weibull	1248.12	1258.01	0.7360

Table 4: Goodness-of-fit results for keyboard and mouse service times

All candidate distributions pass the KS test at the 5% level. While the exponential distribution achieves the lowest AIC, the differences across the three models are relatively small. In contrast, the gamma distribution produces a noticeably higher KS p-value, indicating a closer agreement with the empirical distribution.

Given this, the gamma distribution is selected as a more flexible model that better captures the observed variability in the data. Using MLE, the estimated parameters are $\hat{\alpha} \approx 0.9446$ and $\hat{\beta} \approx 8.6965$.

Service Times for Assembling the Case (Aluminium Plates) A number of candidate distributions were fitted, and the three leading models according to AIC and BIC are shown in Table 5.

Distribution	AIC	BIC	KS p-value
Lognormal	811.54	821.44	0.8623
Inverse Gaussian	811.64	821.54	0.9110
Loglogistic	813.52	823.41	0.9846

Table 5: Goodness-of-fit results for case assembly service times

These three distributions pass the KS test at the 5% level and provide a good fit to the data. The lognormal distribution achieves the lowest AIC and BIC values, although the inverse Gaussian and loglogistic distributions yield slightly higher KS p-values.

Given the very similar performance across the candidate models, the lognormal distribution is selected, as it provides a good balance between fit and model simplicity. Using MLE, the estimated parameters are $\hat{\mu} \approx 0.834$ and $\hat{\sigma} \approx 0.7869$.

2.4 Summary of Fitted Distributions

Table 6 summarises the selected distributions and their estimated parameters for all four input variables. Figure 2 shows the histograms of the observed data with the fitted probability density functions overlaid. In all cases, the fitted densities follow the shape of the empirical distributions closely, providing visual confirmation that the selected models are appropriate.

Variable	Distribution	Parameters (MLE)	Mean	KS p -value
Interarrival times	Exponential	$\hat{\lambda} = 0.1321$	7.57 min	0.321
Initial phase service	Weibull	$\hat{\alpha} = 2.997, \hat{\beta} = 4.414$	3.93 min	0.300
KB ² & Mouse service	Gamma	$\hat{\alpha} = 0.9446, \hat{\beta} = 8.6965$	8.21 min	0.771
Case assembly service	Lognormal	$\hat{\mu} = 0.834, \hat{\sigma} = 0.7869$	3.10 min	0.862

Table 6: Summary of fitted distributions and MLE parameter estimates




Figure 2: Histograms of the four input datasets with fitted probability density functions overlaid ($n = 200$ observations each). (a) Interarrival times: Exponential; (b) Initial phase: Weibull; (c) Keyboard & Mouse: Gamma; (d) Case assembly: Lognormal.

²KB = Keyboard.

3 Model Design

3.1 Model Overview

The simulation model was implemented in Simul8 to represent the laptop production system. Simul8 is a visual discrete-event simulation package that advances the simulation clock from one event to the next, updating system state variables (e.g. queue contents, resource availability, inventory levels) at each event transition. The model uses Simul8's built-in resource management to enforce worker and machine constraints, and Visual Logic scripts to implement non-standard features such as blocking, storage overflow, rework routing, and the (s, Q) inventory policies. The overall layout of the system is illustrated in Figures 3 and 4, covering the full process from order arrival to final dispatch.




Figure 3: Model layout (Part I): order arrival, admission check, order splitting, and parallel production streams converging at the Collect and Match station.




Figure 4: Model layout (Part II): Main Assembly, Place Into Case, Final Quality Check, and Boxing.

3.1.1 Order Arrival and System Entry

Orders arrive with exponential interarrival times, using the estimated rate $\hat{\lambda} = 0.1321$ (see Section 2). Each order consists of $X + 1$ laptops, where $X \sim \text{Binomial}(5, 0.6)$. A capacity constraint of 100 laptops is enforced at the Admission Check gate. If accepting an order would exceed this limit, the entire order is rejected. Accepted orders are split into individual laptop jobs before entering production.

3.1.2 Production Structure

Each laptop is processed through two main parallel streams. The *Motherboard stream* includes Initial Phase, Power Assembly, and Main Part Assembly. The *Case stream* begins with Cutting the Aluminium Plates, which produces two plates that feed two parallel sub-processes: Shaping the Screen $\rightarrow$ Placing the Screen (upper body), and Cutting $\rightarrow$ Placing Keyboard and Mouse (lower body). The three completed sub-assemblies—motherboard, upper body, and lower body—are synchronised at the Collect and Match station.

3.1.3 Assembly and Completion

Following synchronisation at the Collect and Match station, items proceed to Main Assembly, which consists of parallel Assemble HD (Hard Disk) and Assemble Case operations on shared workbenches. Subsequent stages include Place Into Case, Final Quality Check, and Boxing. A key feature of the system is that Boxing is performed at the order level, meaning all items within an order must complete processing before boxing can begin.

3.2 Processing Times and Resources

Processing times are assigned to each station based on either fitted distributions (see Section 2) or given specifications. Resource requirements include multiple worker types and specialised machines, with limited capacities that may create bottlenecks. A summary of processing time distributions and resource allocations is provided in Table 7.

Station	Processing Time	Workers	Machines
Initial Phase	Weibull($\hat{\alpha}, \hat{\beta}$) [†]	Type 1	—
Power Assembly	Uniform(2,6)	Type 1	—
Main Part Assembly	Triangular(5,8,15)	Type 2	—
Cutting Aluminum Plates	Fixed (8 minutes)	—	Roll & Cut Machine
Shaping the Screen	Uniform(3,5)	Type 3	Special Machine
Placing the Screen	Uniform(4,10)	Type 1	—
Cutting (Keyboard/Mouse)	Uniform(5,8) + Rework	Type 3	Special Machine
Placing Keyboard & Mouse	Gamma($\hat{\alpha}, \hat{\beta}$) [†]	Type 1	—
Assemble HD and Peripherals	$3X + 1, X \sim \text{Beta}(3, 2)$	Type 2	Workbench
Assemble the Case	Lognormal($\hat{\mu}, \hat{\sigma}$) [†]	Type 1	Workbench
Place the Motherboard	Uniform(4,8)	Type 1	Workbench
Quality Check	Triangular(8,14,16)	Type 4	—
Boxing	Uniform(4,8)	Type 5	—

Table 7: Processing time distributions and resource allocations. † indicates fitted distributions.

The resource capacities in the system are as follows: Type 1 workers (7), Type 2 workers (3), Type 3 workers (3), Type 4 workers (2), Type 5 workers (2), Roll-and-Cut machines (2), Special machines (3), and Workbenches (2). The factory operates continuously on a 24-hour schedule.

3.3 Additional System Features

Several additional modelling features are incorporated to capture key operational constraints and system dynamics more accurately.

Blocking and Inventory-Driven Suspension. The queue of the Power Assembly has a finite capacity of 3 items. When the queue is full, the Initial Phase is blocked and cannot release completed items until space becomes available. In Simul8, this is implemented by setting the queue capacity to 3 and enabling the *block* routing option on the upstream station³. A similar blocking mechanism is applied to

³Here the upstream station is Initial Phase, which feeds items into the Power Assembly queue.

inventory-dependent stations. When the on-hand stock of aluminium blocks or hard disks reaches zero, the corresponding production route (Cutting the Aluminium Plates or Assemble HD) is suspended via a Visual Logic script that sets the station to a *waiting* state. Production resumes automatically when the replenishment order arrives and the inventory variable is incremented above zero.

Storage Overflow. The queue of Main Part Assembly has a capacity of 2 items. Excess items are diverted to storage, with transfer times of 2 minutes to storage and 3 minutes back to the queue. This is modelled using a routing-out rule with a conditional branch: when the Main Part Assembly queue reaches capacity, items are redirected to a storage activity centre; a separate Visual Logic trigger monitors the queue count and moves items back (up to 2 at a time) when space becomes available.

Rework. In the cutting process for the keyboard and mouse, items may fail with probability 0.2 (first attempt) and 0.1 (second attempt). Failed items re-enter the back of the cutting queue for reprocessing. No item is processed more than three times. This is implemented using a label-based counter on each work item: after processing, a Visual Logic script increments the counter, checks the stage-dependent failure probability, and routes the item either back to the queue or forward to the next station.

Inventory Control. Aluminium blocks and hard disks follow (s, Q) inventory policies. When the on-hand inventory level falls below the reorder point s , an order of size Q is placed. The replenishment arrives after a fixed lead time of 15 hours. For aluminium, $s = 10$ and $Q = 50$ with an initial stock of 50 blocks, while for hard disks, $s = 10$ and $Q = 70$ with an initial stock of 70 units. In Simul8, the inventory levels are tracked using global variables that are decremented each time a unit is consumed and incremented upon delivery. A Visual Logic check at each consumption event triggers a replenishment order (scheduled as a future event) when the inventory falls below s . If the stock is zero at consumption time, the consuming station is suspended as described above until the replenishment arrives.

Order-Level Synchronisation (Boxing). The boxing process is triggered only after all items in an order complete quality checks. To implement this, an order-level label `OrderID` and a count label `OrderSize` are assigned to every laptop at the Order Splitting stage. Simul8’s *Collect* routing-in method is used at the Boxing queue: the queue collects work items that share the same `OrderID` value and releases them to the Boxing activity only when the number of collected items equals the recorded `OrderSize`. As a result, the last completed laptop in each order determines when boxing can begin, introducing order-size-dependent delays.

Component-Level Synchronisation (Collect and Match). Each laptop requires three sub-assemblies—motherboard, upper-body case, and lower-body case—to be completed before entering main assembly. A label-based matching scheme analogous to the order-level mechanism is used. At the Order Splitting stage, each laptop is assigned a unique `LaptopID`. This label is carried through all three parallel streams. The Collect and Match station uses a *Collect* routing-in rule that groups items by `LaptopID` and releases the matched triple to Main Assembly once all three components have arrived.

3.4 Simulation Settings

The simulation clock is configured for continuous 24-hour, 7-day operation. Each replication runs for a total of 60,480 minutes (around 1 month after excluding the warm-up period), comprising a warm-up period of $t_w = 20,160$ min (14 days) and a results collection period of 40,320 min (28 days, approximately one month). See Section 4 for the warm-up determination.

The baseline model is evaluated over 100 independent replications to ensure statistical reliability, while intermediate configurations are tested with 20 replications each and the final improved model with 50 replications.

4 Baseline Performance Evaluation and Output Analysis

4.1 Steady-State Setting and Warm-up Analysis

The laptop manufacturing system is modelled as a steady-state system, with orders arriving continuously and operations running on a stable 24-hour schedule. However, each simulation run begins with an empty system, meaning that the initial observations are affected by start-up effects and are not representative of steady-state behaviour. To address this, an appropriate warm-up period t_w is determined using the *Moving Average Method*.

Specifically, the cycle-time observations $X_1, X_2, \dots, X_n$ are collected from a single long run. Given a window size w , the moving average is computed as

$$X_i(w) = \frac{1}{w} \sum_{k=i}^{i+w-1} X_k, \quad i = 1, 2, \dots, n - w + 1.$$

Figure 5 is used as a graphical diagnostic for initialisation basis. The warm-up point t_w is taken as the earliest point where the moving average appears stable, remaining around a constant level with no clear trend.




Figure 5: Moving-average warm-up analysis using a single long-run cycle-time trace. Each point shows the cycle time of a laptop at the Final Quality Check station. Dashed lines show three moving-average windows ($w = 100, 500, 1000$); the vertical line marks the chosen warm-up cutoff $t_w = 20,160$ min (14 days).

The moving averages stabilise by approximately 20,000 minutes. We therefore set $t_w = 20,160$ min (exactly 14 days) as the warm-up period and discard all observations collected before this time in every replication. Each replication runs for a total of 60,480 minutes (= 42 days), so the results collection period is $60,480 - 20,160 = 40,320$ min (≈ 1 month). All reported performance statistics are based on the post-warm-up portion of each replication only.

4.2 Replications and Cycle-Time Estimation

The main output measure is the average cycle time per laptop, measured as the time from order entry to boxing completion. To assess whether the number of replications is sufficient, we apply the *two-stage procedure* with a target half-width $\varepsilon = 1$ minute for a 95% confidence interval.

Initial Analysis. We run $m_0 = 100$ independent replications. The initial results give a mean cycle time of $\bar{Y} = 1574.79$ min and a standard deviation of $S = 312.24$ min.

Replication Requirement and Practical Choice. The two-stage procedure estimates the required number of replications as

$$m^* = \left\lceil \left(\frac{t_{\alpha/2, m_0-1} \cdot S}{\varepsilon} \right)^2 \right\rceil = \left\lceil \left(\frac{1.984 \times 312.24}{1} \right)^2 \right\rceil \approx 383,844.$$

This is computationally infeasible for this model. We therefore report results based on the available 100 replications and acknowledge this as a limitation.

The resulting 95% confidence interval for the mean cycle time is (1513.59, 1635.98) min, with a half-width of 61.20 min. Although this does not meet the 1-minute target, it remains sufficiently precise for bottleneck identification and for comparing alternative system designs. While the cycle-time estimate gives an overall measure of the system performance, it does not by itself indicate whether the system is capacity constrained, so we next examine the throughput and rejection.

4.3 Throughput and Order Rejection

In the baseline model (100 replications), the system produces about 1893 completed laptops per collection period (approximately 2.82 laptops per hour through Final Quality Check), of which about 257 boxed orders are dispatched. In each run, around 5328 orders enter the system, of which about 538 are accepted (split into individual laptop jobs) and about 4790 are rejected because the system is already full (maximum 100 laptops in process at a time). This gives an estimated rejection rate of

$$\text{Rejection rate} = \frac{4790.1}{5328.0} \approx 89.9\%.$$

The 95% confidence intervals for the key throughput measures are: FQC completions 1893.2 ± 26.4 , boxed orders 257.4 ± 3.6 , and rejection rate $89.9\% \pm 0.5\%$. This result shows that the current setup cannot handle demand effectively. Orders arrive faster than they can be processed, so most of them are turned away. From the results above, it can be concluded that the system is under severe pressure, but they do not yet indicate where the delay is concentrated. This is then examined through queueing performance.

4.4 Queueing Performance

The queue statistics show where waiting time is concentrated. Table 8 reports the top queues in the baseline system (sorted by average waiting time), rather than the full queue list. Queue names follow the Simul8 labels; in particular, “Ready for Collecting A” (Collection A) is the queue where completed cases wait before entering main assembly, and “Queue of Case” is the queue immediately before case manufacturing.

Queue	Avg wait (min)	95% CI ⁴	Avg size	Max size
Ready for Collecting A	1052.61	(973.16, 1132.07)	48.78	87.8
Queue of Case	1015.95	(936.50, 1095.39)	47.13	85.8
Queue of Main Assembly	291.46	(287.62, 295.30)	13.48	46.1
Queue of Boxing	31.74	(31.12, 32.35)	1.50	6.0
Queue of Quality Check	27.52	(27.17, 27.87)	1.29	14.1
Assemble Case Finished	21.11	(21.06, 21.16)	0.99	2.0
Queue of Cutting	4.71	(4.61, 4.82)	0.27	7.9
Queue of Screen	2.60	(2.54, 2.65)	0.12	3.0
Main Assembly Queue	1.94	(1.92, 1.95)	0.09	2.0

Table 8: Top internal queues in the baseline system (sorted by average waiting time; not the full list)

The two dominant queues, *Ready for Collecting A* (≈ 1053 min) and *Queue of Case* (≈ 1016 min), both belong to the case line and together account for most of the cycle-time delay, indicating that case production is the primary bottleneck. The *Queue of Main Assembly* (≈ 291 min) indicates that main assembly stage is frequently delayed, likely because cases do not arrive quickly enough to match the motherboard and peripheral streams. Among the remaining queues, *Assemble Case Finished* (≈ 21 min) suggests a secondary timing mismatch between the hard-disk and case-assembly sub-lines.

Figure 6 provides a process-flow view to clarify the queue names and their relationships.




Figure 6: Simplified process-flow diagram for interpreting the key queues.

Overall, the queueing performance shows that the case line is a main source of the delay. To check whether the interpretation is consistent with the station behaviour, we then examine the workstation and resource utilisation.

4.5 Workstation and Resource Utilisation

Table 9 presents the utilisation breakdown for the main workstations. The key patterns are summarised below.

Workstation	Waiting %	Working %	Blocked %	Completed
Cutting the Al Plates	3.35	18.65	78.00	940
Place Into Case	85.82	14.06	0.00	944
Manufacturing the Motherboard	2.70	97.30	—	—
Boxing 2	96.17	3.83	—	257
Final Quality Check ⁵	0.00	0.00	0.00	1893
Main Part Assembly ⁵	0.00	0.00	0.00	1869

Table 9: Workstation utilisation in the baseline system (selected stations)

The Cutting station is blocked for 78.00% of the time; under the current setting, this is mainly driven by aluminium raw-material shortages. Conversely, Place Into Case is idle for 85.82% of the time, waiting for inputs to arrive. Boxing also shows a very high waiting proportion (96.17%), because packaging can start only after all laptops in the same order are completed; delays in earlier stages therefore create a knock-on effect at the packing stage. Taken together, these observations indicate that throughput in the case-manufacturing chain is severely constrained, with the aluminium inventory policy $(s, Q) = (10, 50)$ likely to be a major limiting factor. Manufacturing the Motherboard operates at 97.30% utilisation, confirming that the motherboard stream is continuously busy.

Table 10 shows the utilisation of the shared worker resources.

Worker Type	Capacity	Average Utilisation (%)	95% CI
Type 1 Worker	7	17.06	(16.93, 17.19)
Type 2 Worker	3	18.78	(18.63, 18.94)
Type 3 Worker	3	18.52	(18.38, 18.65)
Type 4 Worker	2	29.72	(29.50, 29.93)
Type 5 Worker	2	3.83	(3.80, 3.86)
Workbench	2	71.61% working, 28.39% waiting	
Roll & Cut Machine	2	18.65	(18.51, 18.79)
Special Machine	3	18.52	(18.38, 18.65)

Table 10: Utilisation of worker resources in the baseline model

Type 4 workers have the highest utilisation (29.72%), as they serve the Final Quality Check station through which every completed laptop passes. Overall, worker utilisation is low: all worker types are below 30%. This suggests that labour is not the primary constraint; instead, the low rates mainly reflect upstream material shortage⁶ rather than excess staffing. The Workbench is the most utilised shared equipment (71.61% working) and could become a constraint if throughput increases substantially.

5 Improvements

5.1 Bottleneck Diagnosis

This section presents a systematic bottleneck analysis using four complementary indicators from the baseline results: (i) queue waiting times, (ii) station blocking and shortage percentages, (iii) worker and machine utilisation, and (iv) throughput counts. The analysis reveals a layered bottleneck structure, where the primary constraint obscures secondary and tertiary constraints that only become apparent once it is relieved.

5.1.1 Primary Bottleneck: Material Shortage in the Case Line

The queue analysis in Section 4 identifies two queues with dramatically larger waiting times than the rest of the system:

⁵These stations have near-zero processing times (assembly/joining points), so their utilisation percentages round to 0%.

⁶Primarily the aluminium supply for the case line; see Section 5.1.

- **Queue of Case** (≈ 1016 min): laptops wait for case manufacturing to begin.
- **Ready for Collecting A** (≈ 1053 min): finished case sub-assemblies wait for the motherboard and peripheral streams before entering main assembly.

These two queues together account for the majority of the cycle-time delay, while no other internal queue exceeds 291 minutes. The utilisation data supports this: Cutting the Aluminium Plates is *blocked* for 78.00% of the time, meaning that the station frequently cannot proceed with production. Simultaneously, *Place Into Case*⁷ is *idle* for 85.82%, due to a lack of inputs from the congested case line.

The root cause is the $(s, Q) = (10, 50)$ aluminium inventory policy with a 15-hour lead time. Given an average demand of roughly 1893 laptops per collection period (each consuming one aluminium block), the total aluminium demand is approximately 1893 blocks over 40,320 minutes. With $Q = 50$ and a lead time of 900 minutes, each replenishment cycle covers only about 50 blocks while demand accumulates. As a result, cutting machines are frequently idle while waiting for resupply, leading to prolonged queues downstream⁸.

5.1.2 Cascading Effects Across the System

The case-line bottleneck spreads through the entire system in three ways:

1. **Main assembly synchronisation delay.** Because the motherboard line completes much faster than the case line, motherboards accumulate at the synchronisation queue (Queue of Main Assembly, ≈ 291 min), waiting for cases to arrive.
2. **Admission overload.** The slow case line causes in-progress laptop counts to remain near the 100-unit admission limit for extended periods, resulting in an order rejection rate of 89.9%. Most arriving demand is simply turned away.
3. **Downstream under-utilisation.**⁹ Worker and machine utilisation rates are uniformly low (all workers below 30%, as shown in Table 10), not because the workforce is oversized, but because material shortage limits the volume of work entering the system.

5.1.3 Secondary Bottleneck: Hard Disk Supply Shortage

To isolate the effect of the aluminium inventory fix, an intermediate experiment (“AL-fix”) was carried out in which *only* the aluminium policy was changed to $(s, Q) = (90, 150)$, while the hard disk policy and all other settings remained at baseline values. The results (Table 13, column 2) reveal a dramatic bottleneck shift.

The aluminium fix essentially eliminates the case-line shortage: Queue of Case drops from 1016 to 6.4 min, and Ready for Collecting A from 1053 to 6.6 min. Cutting is no longer blocked (baseline blocked 78.0% $\rightarrow$ 0%). However, a *new* bottleneck immediately appears at Main Assembly: its queue time explodes from 291 to 874 min ($\times 3.0$), with an average queue size of 48.1 items. Cases now arrive at the synchronisation point much faster, but hard disks cannot keep pace. The Workbench reaches 97.35% utilisation (up from 71.61%), operating close to full capacity. Meanwhile, the Cutting station is now *waiting* 78.08% of the time: the 100-unit WIP (work-in-progress) cap combined with Main Assembly congestion limits how many new orders can enter the system, leaving Cutting idle despite ample raw material. As a result, overall throughput improves only modestly: FQC completions rise from 1893 to 2252 (+19%), and the rejection rate remains high at 88.1%.

The root cause mirrors the aluminium problem: the hard disk policy $(s, Q) = (10, 70)$ with a 15-hour (900 min) lead time cannot sustain the Workbench at higher throughput. When the hard disk policy is

⁷Place Into Case is located after Main Assembly; it receives finished assemblies from the Workbench.

⁸Downstream of Cutting: Shaping the Screen, Placing the Screen, Cutting (KB/Mouse), and Placing KB & Mouse, which all depend on the aluminium plates produced by the Cutting station.

⁹When case production is delayed, Main Assembly, Place Into Case, Quality Check, and Boxing all receive fewer items.

also expanded to (70, 100) (the “AL+HD” configuration, Table 13 column 3), FQC completions jump from 2252 to 3324 (+47.6%), confirming that the hard disk shortage was the limiting secondary constraint.

5.1.4 Testing Quality Check Capacity

Having resolved the two material-supply bottlenecks, we hypothesised that the Final Quality Check might become a constraint, since it processes every completed laptop using only 2 Type 4 workers. An intermediate experiment (“AL+HD+QC”) was therefore conducted, adding a third Type 4 worker while keeping all other settings identical to AL+HD.

The results are striking: *every throughput and queue metric is identical between AL+HD and AL+HD+QC* (Table 13, columns 3–4). FQC completions, boxing counts, cycle time, and all queue statistics are unchanged. The *only* difference is per-worker utilisation: Type 4 Worker utilisation drops from 52.19% to 34.80%, indicating that the additional worker merely shares the existing load rather than relieving a limiting constraint. This confirms that Quality Check was *not* a bottleneck at the AL+HD throughput level (≈ 3324 FQC completions per collection period).

5.1.5 Tertiary Constraint: Workbench Capacity

With both inventory policies expanded and QC confirmed as a non-bottleneck, the data points to the Workbench as the next limiting constraint. In the AL+HD configuration, Workbench utilisation is 83.87% (up from 71.61% at baseline), and the Queue of Main Assembly remains substantial (335 min average waiting). The Workbench processes three operations sequentially — Assemble HD, Assemble Case, and Place Into Case — and any throughput increase adds pressure on this shared resource.

In the final improved configuration, which increases Workbenches from 2 to 4 (along with expanded inventories, WIP limit increase to 200, and the addition of one Type 4 worker), Workbench utilisation reaches 99.74%. Even with the doubling of Workbenches, the resource operates close to full capacity, confirming the Workbench as the tightest remaining constraint. The improved model achieves 4799 FQC completions but at the cost of elevated queue times at Main Assembly (701 min) and Quality Check (201 min), as the much higher throughput now places additional pressure on these downstream stations¹⁰

5.1.6 Summary

Table 11 summarises the layered bottleneck structure and the corresponding improvements.

Layer	Constraint	Key Evidence	Intervention
1	Aluminium supply	Cutting blocked 78%; case queues > 1000 min	Increase (s, Q): (10, 50) $\rightarrow$ (90, 150)
2	Hard disk supply	AL-fix: Main assemble queue 291 $\rightarrow$ 874 min; WB 97.35%	Increase (s, Q): (10, 70) $\rightarrow$ (70, 100)
3	Quality Check (non-bottleneck)	AL+HD+QC = AL+HD in all throughput metrics	(confirmed not limiting)
4	Workbench (WB) capacity	83.87% in AL+HD; 99.74% with 4 WBs in improved	Increase from 2 to 4 Workbenches

Table 11: Layered bottleneck structure and corresponding interventions

5.2 Changes in the Improved Model

Based on the layered bottleneck diagnosis, changes were applied incrementally to verify that each intervention addressed the identified constraint. Table 12 summarises the parameter modifications across all five configurations tested.

¹⁰After the Collect-and-Match stage, each completed laptop goes through Main Assembly and Quality Check before Boxing.

Parameter	Type	Baseline	AL-fix	AL+HD	AL+HD+QC	Improved
Al blocks (s, Q)	Inventory	(10, 50)	(90, 150)	(90, 150)	(90, 150)	(120, 200)
Hard disks (s, Q)	Inventory	(10, 70)	(10, 70)	(70, 100)	(70, 100)	(80, 120)
Type 4 Workers	Resource	2	2	2	3	3
WIP Limit	System	100	100	100	100	200
Workbenches	Resource	2	2	2	2	4

Table 12: Summary of parameter configurations across all five experiments

Step 1: Aluminium inventory expansion (AL-fix). The (s, Q) policy for aluminium blocks was increased from (10, 50) to (90, 150), raising both the reorder point and order quantity to ensure that cutting machines are supplied with raw material more continuously. This change nearly eliminated the case-line shortage (case queue times dropped from ~ 1016 to ~ 6 min), but total throughput improved only modestly: FQC completions rose from 1893 to 2252 per collection period (+19%). The Workbench reached 97.35% utilisation, confirming that the hard disk supply became the new limiting constraint (Section 5.1.3).

Step 2: Hard disk inventory expansion (AL+HD). The hard disk policy was expanded from (10, 70) to (70, 100) on top of the aluminium change. This produced the largest single-step throughput gain: FQC completions jumped from 2252 to 3324 (+47.6%), and the rejection rate dropped from 88.1% to 82.1%. Workbench utilisation fell from 97.35% to 83.87%, confirming that the hard disk supply constraint was now partially relieved. Cycle time also decreased dramatically, from 1169 to 614 min.

Step 3: Quality Check capacity test (AL+HD+QC). Adding a third Type 4 worker to the AL+HD configuration produced *no change* in throughput, cycle time, or any queue metric. The only effect was reduced per-worker utilisation at Quality Check (52.19% $\rightarrow$ 34.80%). This confirmed that Quality Check was not a bottleneck at the AL+HD throughput level (see Section 5.1.4).

Step 4: Combined improvements (Improved). The final improved configuration simultaneously applies more generous inventory policies (Al = (120, 200); HD = (80, 120)), one additional Type 4 worker, an increased WIP limit of 200, and 4 Workbenches (up from 2). This achieves the highest throughput: 4799 FQC completions (+153.5% vs baseline) and 649 boxed orders, with the rejection rate dropping from 89.9% to 74.8%. Mean cycle time is 1242 min, a 21.1% reduction from baseline. Notably, the cycle time is *higher* than in the AL+HD configuration (614 min), because the WIP limit increase to 200 allows many more orders to enter the system simultaneously, increasing congestion despite the greater throughput.

Table 13 presents the full progressive comparison across all five configurations.

	Baseline	AL-fix	AL+HD	AL+HD+QC	Improved
<i>Configuration</i>					
AI (s, Q)	(10, 50)	(90, 150)	(90, 150)	(90, 150)	(120, 200)
HD (s, Q)	(10, 70)	(10, 70)	(70, 100)	(70, 100)	(80, 120)
Type 4 Workers	2	2	2	3	3
WIP Limit	100	100	100	100	200
Workbenches	2	2	2	2	4
<i>Throughput</i>					
FQC completed	1893	2252	3324	3324	4799
Boxed orders	257	309	465	465	649
Throughput (laptops/hr)	2.82	3.35	4.95	4.95	7.14
Rejection rate (%)	89.9	88.1	82.1	82.1	74.8
<i>Key queues (avg wait, min)</i>					
Ready for Collecting A	1053	6.6	7.3	7.3	130
Queue of Case	1016	6.4	6.7	6.7	133
Queue of Main Assembly	291	874	335	335	701
<i>Utilisation (%)</i>					
Workbench working	71.61	97.35	83.87	83.87	99.74
Cutting working	18.65	21.92	32.42	32.42	94.61
Cutting blocked	78.00	0.00	1.36	1.36	0.00
Mean cycle time (min)	1575	1169	614	614	1242
<i>Cycle time 95% CI (min)</i>					
Half-width	61.20	101.78	52.58	52.58	55.44

Table 13: Progressive experiment results across five configurations (baseline: 100 reps; AL-fix, AL+HD, AL+HD+QC: 20 reps each; improved: 50 reps). Columns 2–4 each add one change to the previous configuration; the Improved column combines all changes with further parameter adjustments. AL+HD+QC is identical to AL+HD in all metrics except per-worker utilisation, confirming QC is not a bottleneck.

Figure 7 provides a visual summary of the three main KPIs across all five configurations, clearly illustrating the layered improvement pattern.




Figure 7: Progressive comparison of (a) throughput, (b) mean cycle time, and (c) order rejection rate across the five experimental configurations. AL+HD+QC is identical to AL+HD in all three metrics.

5.3 Statistical Validation

To confirm that the observed cycle-time reduction is statistically significant, a Welch two-sample t -test is applied. The baseline model ($n_1 = 100$, $\bar{X}_1 = 1574.79$ min, $S_1 = 312.24$ min) and the improved model ($n_2 = 50$, $\bar{X}_2 = 1242.30$ min, $S_2 = 200.01$ min) were run with independent random number streams, so

an independent-samples test is appropriate.

Under $H_0: \mu_{\text{baseline}} = \mu_{\text{improved}}$, the Welch t -statistic is

$$t = \frac{1574.79 - 1242.30}{\sqrt{312.24^2/100 + 200.01^2/50}} = 7.89,$$

with degrees of freedom $\nu \approx 139.0$ and $p = 8.00 \times 10^{-13}$. The null hypothesis is rejected at any conventional significance level. A 95% confidence interval for the difference in means is (249.19, 415.79) min, corresponding to a reduction of approximately 21.1%.

The improvement is both statistically significant and practically meaningful. Although the two models have unequal sample sizes, Welch’s test remains valid; the smaller improved-model sample mainly widens its own confidence interval rather than affecting the validity of the comparison.

We note that, since the baseline and improved models were run with *independent* random number streams in Simul8, a paired- t test (which requires common random numbers) cannot be applied here. The Welch independent-samples t -test is therefore the appropriate method. As discussed in Section 6, using common random variables would reduce variance and strengthen the comparison, but was not feasible in this implementation.

5.4 Performance of the Improved Model

Table 14 compares the key performance indicators between the two models.

Metric	Baseline	Improved	Change
Mean cycle time (min)	1574.79	1242.30	−21.1%
95% CI for mean (min)	(1513.59, 1635.98)	(1186.86, 1297.74)	—
95% CI half-width (min)	61.20	55.44	—
Rejection rate (%)	89.9	74.8	−15.1 pp ¹¹
FQC completions	1893	4799	+153.5%
Boxed orders	257	649	+152.5%
Throughput (laptops/hr)	≈ 2.82	≈ 7.14	+153%

Table 14: Performance comparison between baseline and improved models

Queue improvements. Table 15 shows how the most important queues changed.

Queue	Avg wait (min)		Avg size	
	Baseline	Improved	Baseline	Improved
Ready for Collecting A	1052.61	129.92	48.78	15.14
Queue of Case	1015.95	132.93	47.13	15.49
Queue of Main Assembly	291.46	701.31	13.48	83.18
Queue of Quality Check	27.52	201.28	1.29	23.94
Queue of Boxing	31.74	12.39	1.50	1.48
Assemble Case Finished	21.11	23.25	0.99	2.77

Table 15: Queue comparison between baseline and improved models

The case line queues dropped dramatically (Queue of Case from 1016 to 133 min), confirming that the inventory policy change addressed the root cause of aluminium shortage. Queue of Boxing also improved from 31.74 to 12.39 min. However, Queue of Main Assembly increased from 291 to 701 min and Queue of Quality Check rose from 27.52 to 201 min, reflecting the much higher throughput (≈ 4799 vs.

¹¹pp = percentage points.

≈ 1893 FQC completions) placing greater pressure on downstream stages¹². The Assemble Case Finished queue remained similar (21→23 min), indicating that the hard disk–case synchronisation delay was not substantially altered by these changes.

Workstation and resource changes. Table 16 summarises how key resources responded.

Resource	Baseline	Improved
Type 1 Worker utilisation	17.06%	43.26% ↑
Type 2 Worker utilisation	18.78%	47.59% ↑
Type 3 Worker utilisation	18.52%	47.05% ↑
Type 4 Worker utilisation	29.72%	50.21% ↑
Type 5 Worker utilisation	3.83%	9.64% ↑
Workbench working	71.61%	99.74% ↑
Roll & Cut Machine util.	18.65%	47.31% ↑
Cutting blocked	78.00%	0.00% ↓
Place Into Case waiting	85.82%	81.22% ↓

Table 16: Resource utilisation comparison between baseline and improved models

All worker utilisation rates increased substantially (e.g. Type 1: 17.06% → 43.26%; Type 4: 29.72% → 50.21%). Cutting blocked time dropped from 78.00% to 0.00%, confirming that the aluminium shortage has been completely eliminated. The most striking change is Workbench utilisation reaching **99.74%**, identifying it as the new limiting constraint in the improved model. Despite having 4 Workbenches (up from 2), the greatly increased throughput saturates all available capacity. Figure 8 visualises the utilisation changes across all resources.




Figure 8: Resource utilisation comparison between the baseline and improved models. The Workbench reaches near-full capacity (99.74%) in the improved model, identifying it as the new limiting constraint.

5.5 Discussion and Recommendations

The improved model reduces mean cycle time by 21.1% and increases throughput by 153%, as confirmed by the Welch t -test ($p = 8.00 \times 10^{-13}$). The rejection rate decreased by 15.1 percentage points (89.9% → 74.8%), and the aluminium shortage has been completely eliminated (cutting blocked: 78.00% → 0.00%).

¹²Main Assembly → Place Into Case → Quality Check → Boxing.

Throughput–cycle-time trade-off. Interestingly, the improved model’s cycle time (1242 min) is substantially higher than that of the intermediate AL+HD configuration (614 min), despite achieving much higher throughput. This is because the WIP limit increase from 100 to 200 allows many more orders to enter the system simultaneously, increasing work-in-progress congestion and thus queueing time per order. The trade-off between throughput and per-order cycle time is a classic feature of systems operating near capacity.

Quality Check capacity. A key finding of this study is that adding a third Type 4 worker produced *zero* change in throughput, cycle time, or queue metrics (AL+HD+QC vs. AL+HD). Quality Check was not a bottleneck at the current throughput level, despite appearing to be a candidate based on worker utilisation alone. This underscores the importance of systematic bottleneck decomposition experiments, rather than relying solely on utilisation metrics.

Remaining issues. The Workbench has emerged as the new limiting constraint: its utilisation reached 99.74% in the improved model, even after doubling from 2 to 4 units. Queue of Main Assembly increased from 291 to 701 min, and Queue of Quality Check rose from 28 to 201 min, reflecting the substantially higher throughput. Adding further Workbench capacity, or reducing processing time at the Workbench station, would be the logical next intervention.

Recommendation. We recommend implementing the inventory policy changes and Workbench expansion as a first priority, as they address the two primary supply-chain bottlenecks and deliver a large, statistically validated improvement. In a second phase, Workbench capacity should be further expanded (or processing time reduced) to relieve the 99.74% utilisation constraint and enable additional throughput gains.

6 Limitations

- **Confidence interval precision.** As reported in Section 4.2, the baseline model still produces a 95% confidence-interval half-width of 61.20 minutes for the mean cycle time, which is much higher than the target of the project of 1 minute. This means that the model is suitable for broad diagnosis and comparative analysis, but not for highly precise point estimation. The improved model has a half-width of 55.44 min based on 50 replications. The two-stage procedure estimates that approximately 383,844 replications are needed to achieve a half-width of 1 minute for the mean cycle time. Running over 383,000 replications was computationally infeasible for this model.
- **Warm-up determination.** The warm-up analysis is based on the moving average method, which is widely used and provides intuitive diagnostic results. However, the selection of warm-up points remains to some extent subjective. Any subjective judgement may result in slight variations in the cut-off values. In this model, the selected 20,160 minutes (14 days) is an appropriate figure, but it is not the only possible value; it is a figure arrived at through subjective judgement. A more formal approach (e.g. MSER-5 or the batch-means method) could provide a more rigorous cut-off point.
- **Limited intermediate-model replications.** The intermediate experiments (AL-fix, AL+HD, AL+HD+QC) were each evaluated with only 20 replications, yielding wider confidence intervals compared to the baseline (100 reps) and improved (50 reps) models. Additional replications would improve precision for the intermediate configurations.
- **Simplified operational assumptions.** The model assumes continuous 24-hour operation with no shift changes, machine breakdowns, or maintenance downtime. Real manufacturing systems would exhibit additional variability from these factors.

- **Limited exploration scope.** Although five configurations were tested in a progressive bottleneck decomposition (aluminium, hard disk, QC capacity, WIP/Workbench), only one or two target values were tested for each parameter. A more systematic exploration (e.g. factorial experiments or simulation-based optimisation) could identify better configurations.
- **Absence of financial analysis.** The proposed improvements (larger inventory orders, additional Workbenches and workers) would incur additional costs in practice, including capital expenditure, holding costs, and labour costs. However, no cost data were available for this system, so a cost-benefit analysis could not be performed. In a real-world setting, the financial feasibility of each intervention should be evaluated alongside the operational performance gains reported here.

7 Conclusion

This study developed a discrete-event simulation model of a multi-stage laptop manufacturing system using Simul8. Through input analysis, appropriate probability distributions were fitted to four process-time datasets and validated using goodness-of-fit tests.

The baseline model revealed an average cycle time of approximately 1575 minutes, with the case manufacturing line identified as the primary bottleneck: the two dominant queues (*Ready for Collecting A* and *Queue of Case*) together account for over 80% of the total waiting time, and the Cutting station is blocked 78.00% of the time due to aluminium shortages.

A series of five progressive experiments provided a systematic bottleneck decomposition. Fixing only the aluminium inventory policy eliminated the case-line shortage but exposed a secondary hard disk supply bottleneck, with the Main Assembly queue rising from 291 to 874 min and Workbench utilisation reaching 97.35%. Adding the hard disk inventory fix produced the largest single-step throughput gain: FQC completions jumped from 2252 to 3324 (+47.6%). A dedicated experiment confirmed that Quality Check is *not* a bottleneck: adding a third Type 4 worker produced zero change in any throughput or queue metric. The Workbench was identified as the true tertiary constraint, reaching 99.74% utilisation in the final improved model.

The final improved model achieves a 21.1% reduction in mean cycle time (to ≈ 1242 min), a 153% increase in throughput (2.82 to 7.14 laptops/hr), and a 15.1-percentage-point reduction in the order rejection rate. These improvements were confirmed as statistically significant by a Welch two-sample t -test ($p = 8.00 \times 10^{-13}$).

The study also identified the Workbench as the emerging limiting constraint in the improved system, with working utilisation at 99.74%. Future work should explore expanding Workbench capacity or reducing Workbench processing time, and applying variance reduction techniques to improve the precision of performance comparisons.