



LOKAD

Quantitative Supply Chain



# Optimizing Scheduling in

# AEROSPACE

A Programmatic Approach to Managing  
Complexity in OEM and MRO Operations

[www.lokad.com](http://www.lokad.com)

# TABLE OF CONTENTS

Executive Summary .....3

Key Insights..... 4

1 Macro v Micro-level Planning..... 4

2 You Have a Bill of Materials Resources.....6

2.1 Managing Your Bill of Resources.....7

3 Probability Theory in Scheduling..... 8

3.1 “Exceptions” are the (Expensive) Norm.....9

3.2 Risk Management .....10

4 Why FIFO Fails at micro-level planning ..... 11

5 Solution to the Scheduling Problem ..... 12

5.1 Scheduling in “Emergencies” ..... 13

6 How Lokad Does It .....14

6.1 Scheduling Optimization Results .....14

7 Next Step..... 17

7.1 Special Thanks..... 17

7.2 About the Author ..... 17

Glossary.....18

## THEY TRUST US





“  
 You need a very granular plan to be generated every single day (e.g., the specific part numbers, specific tools, and exact amount of time people should be assigned to each sub-task).  
 ”

## EXECUTIVE SUMMARY

Effective scheduling is a computationally intensive problem. This problem is only exacerbated in aerospace, particularly for OEM and MRO players, as the combinatorial complexity involved is considerable. Though OEM and MRO companies differ when it comes to their respective bills of materials (for OEM, the bill of materials (BOM) is **known**; for MRO, it is **uncertain/probabilistic**), they share similar concerns when it comes to scheduling the manufacture or repair of parts.

It should be noted that these companies actually have a comprehensive **Bill of Resources** (i.e., required parts, tools, and people), rather than a purely physical bill of materials. This is critical as you need your full bill of resources to be **simultaneously available** to perform any given step in a manufacture or repair process. What’s more, you need to come up with an **optimized, feasible sequence of actions** (e.g., properly allocating people’s time among several aircraft competing for the same resources). This sequence of actions maximizes the efficiency of the process as a whole with respect to financial consequences. Even more challenging, the sequence needs to be revised each night and be ready for use the next morning. Unfortunately, if there is a change in the sequence (e.g., a part is missing, a tool is broken, or a technician is sick), the

generated sequence of manufacture/repair actions is now obsolete, as it was based on outdated information. Companies incorrectly think of these events as “emergencies”. Given the volume of parts, tools, and people involved in aerospace production and repair processes, there will (statistically speaking) be “an emergency” every day. These include situations where a single part of the bill of resources is missing, thus requiring the entire sequence of actions to be regenerated. As such, emergencies are not outliers; **they are the norm**. Companies have tried to solve the scheduling problem by combining **macro-level planning** (e.g., a Min-Max inventory policy) with simple heuristics to address production/repair needs (e.g., applying FIFO to a sequence of production or repair orders). These heuristics are ultimately very low-bandwidth fixes to what is a very complex problem; i.e., they ignore the granular problem of generating an **optimized sequence of actions**.

Historically, Lokad has optimized macro-level planning with probabilistic forecasts and stochastic optimization (thus replacing basic tools like Min-Max inventory policies). However, the more difficult challenge of deciding **exactly where and when to use resources** had proved elusive. This **micro-level planning** represents a completely different challenge, one in which the **optimal sequence of actions** is

generated for clients on a daily basis. Lokad does this, informing clients on a daily basis of exactly what parts, tools, and people should be sent to complete each step in a process. Further, this robust sequence of actions is generated to reflect the greatest financial return on investment for the client. When an element of the bill of resources is missing, which will happen almost every day, Lokad can regenerate the **sequence of actions** to remedy this problem in a handful of minutes. This avoids the wasted time, money, and bandwidth associated

with emergency staff meetings and sub-optimal uses of resources.

In short, you need a very granular plan to be generated every single day (e.g., the specific part numbers, specific tools, and exact amount of time people should be assigned to each sub-task). This is precisely what Lokad provides its aerospace clients. This micro-level planning, in conjunction with Lokad's macro-level inventory optimization, effectively solves the problem of scheduling in aerospace.

## KEY INSIGHTS

- Effective scheduling must consider a comprehensive **Bill of Resources (BOR)**-including parts, tools, and people-as missing any single resource can halt production or repair processes.
- Daily uncertainties like delayed parts, unavailable tools, or absent technicians **are the norm**. Any one of these can cause immediate and costly disruptions.
- Traditional methods like FIFO (First In, First Out) are simply insufficient policies for scheduling.
- Specialized algorithms and probabilistic forecasting are critical for efficient scheduling. These tools allow for the generation of sequences that reflect the client's bill of resources, and can produce intelligent revisions to the sequences when needed.

## 1 MACRO VS. MICRO-LEVEL PLANNING

To better understand why effective scheduling is so challenging, it's essential to examine the distinction between macro and micro-level planning. **Macro-level planning** can be understood as the long-term decision-making strategies (e.g., overall inventory policies for service level and staffing policies, etc.). This is making sure that, overall, you have everything you need most of the time to complete processes (with as high a certainty as possible). **Micro-level planning** can be understood as the day-to-day practical decisions you make to execute the processes, considering the resources that you have at that moment.

Below is a table detailing some of the most common (and critical) macro and micro-level questions when it comes to effective scheduling for OEM and MRO companies.

“

Macro-level planning can be understood as the long-term decision-making strategies. Micro-level planning can be understood as the day-to-day practical decisions you make to execute the processes, considering the resources that you have at that moment.

”



Resource	Macro-level Planning	Micro-level Planning
Parts	<ul style="list-style-type: none"> <li>• How much inventory do I require to cover the totality of my operations ?</li> <li>• Where should I store my inventory ?</li> <li>• How do I account for varying lead times (e.g., delayed inventory deliveries) at the warehouse level ?</li> <li>• What should I do with unused P/Ns ?</li> </ul>	<ul style="list-style-type: none"> <li>• What do I do if inventory is not available at the exact moment I need it in the manufacturing / repair facility ?</li> <li>• How do I reschedule my sequence of actions to reflect the absence of a critical resource ?</li> <li>• How does this absence financially impact my production/repair schedule ?</li> <li>• How does this absence impact my contractual commitments and client loyalty (especially in the case of MRO) ?</li> </ul>
Tools	<ul style="list-style-type: none"> <li>• What tools/machinery do I have access to ?</li> <li>• When are the tools going to be available for use ?</li> </ul>	<ul style="list-style-type: none"> <li>• What is the optimal order of actions to fully exploit the available tools and skills ?</li> <li>• Are all the necessary tools for each step of the process currently available ?</li> <li>• How do I reschedule actions if a specific tool is unavailable ?</li> </ul>
People	<ul style="list-style-type: none"> <li>• How many technicians do I have ?</li> <li>• How many can I realistically hire ?</li> <li>• When is each technician required onsite ?</li> <li>• Which skills (e.g., technicians with specific certifications, licenses, etc.) do I have access to ?</li> <li>• When are these skills going to be onsite ?</li> <li>• What should be the trainings / certifications schedule for each staff member in the upcoming year?</li> </ul>	<ul style="list-style-type: none"> <li>• Do I have the necessary number of people I need to complete a given process ?</li> <li>• Who is actually available on the day ?</li> <li>• How should I reschedule actions if a technician is unavailable ?</li> <li>• What should they do and in what order ?</li> <li>• Where and when should I send each person ?</li> <li>• Which technician utilizes their skills the most efficiently (e.g., using the fewest parts and least amount of time) ?</li> </ul>

Table 1. Outline of some macro and micro-level planning considerations for key resources in manufacturing and repair processes.

Importantly, most theories (and solutions) attempt to address **only** the macro-level questions, typically focusing exclusively on **physical, deterministic constraints** (e.g., the number of P/Ns to order). This is a deeply flawed perspective for a few key, predictable reasons.

Fundamentally, people tend to think in terms of **how things should be**, e.g., the PO should be here next week, Jennifer should be available for work tomorrow, etc. They think in terms of things following the plan. This is an easy but very simplistic way to process reality. By contrast, it is much harder to think



in terms of **how things could be**. This style of thinking involves enormous amounts of possibilities, contingencies, and conditionalities, and this becomes increasingly complex very quickly.

For example, the combinatorial complexity involved in managing P/Ns (ordering, reordering, allocating, liquidating, etc.) is enormous. This is especially true given BOM constraints (parts are needed in combination to perform an operation), thus the complexity goes far beyond the capabilities of simplistic solutions (e.g., human intuition, Excel, safety stock formulas, etc.).

Companies also typically consider only the **physical elements** of the problem, such as P/Ns. However, this perspective overlooks the critical role **people** play in manufacturing and repair processes (see the micro-level questions listed above). People are subject to just as much variability as parts are, if not more so. If you ignore this fact, you are missing an enormous part of the problem - a problem that requires more creative solutions than trying to think in terms of **how things should be**.

For example, a human mind might be able to manage the macro-level questions if the scale is not very large (e.g., a very small number of P/Ns). In reality, even an enormous team of highly skilled humans would struggle (and likely fail miserably) to optimize the macro-level questions for any large-scale OEM or MRO company.

That said, there is no human alive (or team of humans) who could manually optimize the sequence of actions needed to answer the micro-level questions above. In fact, the enormous complexity involved explains why most mainstream practices **ignore the micro-level question entirely** and opt instead for very basic heuristics like FIFO (First In, First Out) and gut-instinct (e.g., who looks like they really need the P/N?).

Effectively answering the micro-level questions above requires algorithmic automation - the kind Lokad provides its clients. Anything less is an expensive exercise in wastefulness. Central to answering these micro-level questions is the **Bill of Resources** - a concept that is fundamental to achieving the kind of scheduling automation Lokad recommends.

## 2 YOU HAVE A BILL OF MATERIALS RESOURCES

Rather than a purely physical bill of materials (BOM), each OEM and MRO has to contend with an expansive **Bill of Resources (BOR)**. This refers to the three classes of resources involved in every single manufacturing/repair process. These classes are :

- 1) Parts (e.g., P/Ns)
- 2) Tools (e.g., machinery)
- 3) People (e.g., certified technicians)

Importantly, without the **simultaneous availability** of **all three of the resource classes**, one cannot complete a given process. For example, missing a single P/N, tool, or technician at a critical moment means the manufacturing or repair process halts completely - with immediate financial losses inflicted (through delays and additional costs to complete the process).

“  
 Rather than a purely physical bill of materials (BOM), each OEM and MRO has to contend with an expansive Bill of Resources (BOR). Importantly, without the simultaneous availability of all three of the resource classes, one cannot complete a given process.  
 ”

## 2.1 MANAGING YOUR BILL OF RESOURCES

Although the natures of BORs differ, OEMs and MROs share several challenges and requirements when it comes to managing their bills of resources. Below is a comparison highlighting these key variables :

Key Variables	OEM	MRO
Predictability of BOR	<b>Known</b> - Exact quantities and specifications are predefined for each product. For instance, manufacturing an A320 engine requires ~30,000 parts, ~200 specific tools, and ~150 skilled technicians & engineers.	<b>Unknown</b> - Resource requirements vary based on maintenance schedules and operational factors. Maintaining an A320 engine might involve ~30,000 parts, but exact P/Ns and quantities fluctuate with flight hours and conditions.
Resource Availability	Critical to have all materials, tooling, and skills <b>simultaneously available</b> in the precise order to avoid production delays. Any missing resource halts the manufacturing process, leading to immediate financial loss due to halted production.	Equally essential to ensure all necessary resources are <b>simultaneously available</b> when needed for maintenance tasks. Unavailability disrupts repair processes, causing aircraft downtime and revenue loss as the engine remains grounded.
Financial Implications	Delays in production (even if missing just one unit) directly translate to financial waste, as products like the A320 engine are valuable assets that generate revenue only when operational. Extended delays exacerbate financial losses.	Inability to complete maintenance promptly results in aircraft being out of service longer, increasing financial losses over time as the engine does not contribute to revenue-generating flights.
Resource Uncertainty	Even with a known BOR, uncertainties exist regarding the availability of resources at any given time, necessitating robust planning and probability assessments to mitigate risks.	The inherent unpredictability of maintenance needs requires MROs to understand and manage the probabilities of resource availability, ensuring readiness despite variable demands.
Micro-level Planning	Requires meticulous scheduling and inventory management to align all resources perfectly for efficient manufacturing.	Demands flexible & responsive planning to accommodate varying maintenance schedules & resource requirements based on real-time operational data.

Table 2. Comparative overview of key BOR variables between OEM and MRO contexts

Crucially, whether your bill of resources is known or unknown, you don't know with 100% certainty what resources will be available **at any given time**. This is a fundamental aspect of answering the macro and

micro-level questions mentioned earlier, particularly if you want to do so in a financially efficient way. To answer these questions, though, we need to understand the probabilities of each resource's availability.



### 3 PROBABILITY THEORY IN SCHEDULING



Imagine an OEM wants to produce a compressor rotor blade for an APU (auxiliary power unit). This process requires ~20 P/Ns. Now, imagine an MRO wants to perform minor repairs on that same APU. For this, they might need 100 individual P/Ns. Imagine both companies have set macro-level inventory policies where they always have 99% service level for all P/Ns (or 99% chance of parts arriving on time, subject to lead times). This is a very expensive inventory policy, especially in MRO as their demand patterns are less predictable, but both companies feel that this will cover needs in the vast majority of situations.

To calculate (simply) the likelihood that the OEM and MRO will have all parts available at the same time, the desired service level should be raised to the power of the number of parts: (service level)<sup>number of parts needed</sup>. For the OEM, this is  $(0.99)^{20}$ . For the MRO, the likelihood is  $(0.99)^{100}$ .

Under these conditions :

- There is an **81.79%** chance that the OEM will have all 20 parts available at the same time.

- There is a **36.6%** chance that the MRO has all 100 parts available at the same time.

The impact of this is clear.

- There is almost a **20%** probability that the OEM will encounter expensive delays in their manufacturing process and the scheduled sequence of actions will have to be regenerated.
- There is almost a **65%** probability that the MRO will encounter expensive delays in their repair process and the scheduled sequence of actions will have to be regenerated.

Bear in mind that no sensible company would have a 99% service level for every single part (given the extreme costs associated with such a policy). Fortunately, Lokad has already published extensive resources discussing this problem in aviation. For more, please review Joannès Vermorel's **public lecture** on the topic.



### 3.1 “EXCEPTIONS” ARE THE (EXPENSIVE) NORM

In aerospace “exceptions” are not black swan events; they are the norm. In other words, resource unavailability is ordinary, not unusual. Although global emergencies (such as pandemics disrupting lead times) are relatively uncommon, deviations from the expected (e.g., delayed deliveries and unavailable technicians) are a reality of day-to-day business. Given the size of the bill of resources (parts, tools, and people) involved in OEM and MRO processes, the likelihood that something will be absent at any particular moment is significant (see an example in the previous chapter) - and it will result in the scheduled sequence of actions needing to be regenerated. This, however, takes time, and during that time, the process is halted.

Specific financial information for the cost of downtimes in aerospace is scarce as the expected cost can only be estimated with (very) wide ranges - and sometimes cannot even be accurately estimated due to the sheer volume of variables.<sup>1</sup> However, our own estimates show that downtimes are at least as costly as downtimes in similarly complex and high-value industries. For example, the financial impact of downtimes in the automotive industry (comparably complex to aerospace) can be over **\$2,000,000<sup>2</sup> per hour**, and continues to rise as “the

growing complexity and interdependence of auto-production systems means downtime in one process can halt production across a big part of a plant”.<sup>3</sup> Aerospace operations can be expected to experience costs at least as high (if not higher) given the level of complexity involved in OEM and MRO operations.

Aerospace companies may opt for generous safety stock policies to protect themselves against costly stockouts, but this, fundamentally, **does not address the scheduling problem** presented by an absence. If a technician is sick, or a crucial tool is unavailable, the sequence of actions needs to be regenerated. Trying to manually resolve these exceptions (e.g., unplanned downtimes) each time they arise is an incredibly costly and wasteful approach.

Simply, an event that could realistically happen every day should not be considered an exception - it is a feature of the system and not a bug. As such, it should be **proactively understood** that parts of the bill of resources could be missing on any given day. Thus, your software solution should be able to adapt to the situation and generate a new feasible sequence of actions as quickly as possible. This is only possible if one’s macro and micro-level planning factors the critical role of probability theory.



“  
In aerospace “exceptions” are not black swan events; they are the norm. An event that could realistically happen every day should not be considered an exception - it is a feature of the system and not a bug.  
”

<sup>1</sup>For this reason, relevant information in comparably complex industries was used as a benchmark for our own internal figures (and those of experts I contacted). Feel free to take these figures with a pinch of salt, and I invite any expert to contact with me a more accurate range at [c.doherty@lokad.com](mailto:c.doherty@lokad.com)

<sup>2</sup>Siemens, The True Cost of Downtime 2022, (2023).

<sup>3</sup>Idem.

## 3.2 RISK MANAGEMENT

In aerospace, almost everything is probabilistic, including risks. Lead times vary, as do the financial implications of making decisions that result in critical resources being absent. As discussed in the previous section, the consequences of these poor decisions can be both swift and brutal :

For an OEM, failing to manufacture a compressor rotor blade on time not only affects the immediate assembly of the APU but could impact other areas of aircraft production, e.g., the overall production of engines. This knock-on effect could cost the OEM millions in missed deadlines & contractual penalties.

For an MRO, failing to repair an APU on time could result in a plane being grounded (AOG event). This single AOG event could cost hundreds of thousands of dollars (per day, for the duration of a delay), once the direct and indirect costs are factored.

It is worth noting that the examples discussed

above consider only the physical materials (P/Ns). In reality, effective scheduling factors the availability of the entire bill of resources (materials, tools, and skills). Companies need to consider their bill of resources every day and at scale in order to protect themselves from the economic risks of missing a vital resource. Doing this requires a software solution that can natively operate with probabilities and optimize decisions and sequences in the presence of all of this complexity.

Thinking probabilistically allows companies to make financially sensible decisions that reflect the variety of scenarios (with their respective probabilities) that they may encounter in their manufacturing/repair processes. This reduces the financial waste associated with excess inventory and creates a much more efficient allocation of one's bill of resources.

Alternatively, companies can fall back on simple heuristics that completely ignore this complexity and expose them to unnecessary financial risks.



## 4 WHY FIFO FAILS AT MICRO-LEVEL PLANNING

The most common mistake when trying to address micro-level planning in aerospace is for OEMs and MROs to deploy simplistic policies that ignore the complexity of the task entirely. This is understandable given the **sheer scale of complexity** involved in finding an optimized sequence of actions to produce or repair components. Factoring all the constraints associated with one's bill of resources, and then producing a sequence of actions that is ranked in terms of urgency and profitability, is a sophisticated and computationally-intensive task. In the face of this complexity, companies often feel there is no good option and thus deploy FIFO as a simple solution. However, FIFO is a low-bandwidth solution for a high-dimensional problem.

Consider a situation where an MRO receives two engines for repair. **Engine A** and **Engine B** arrive in that order, with both having the same expected repair time. In reality, **Engine A** (which arrived first) actually requires more parts than **Engine B**. Below is a table outlining how a simple heuristic like FIFO would adversely impact operations when compared to a more sophisticated scheduling optimization.



“

FIFO is a low-bandwidth solution for a high-dimensional problem. The distinctions between each set of choices are not always clear to the human eye, but a sophisticated algorithm can identify and process all the different variables and generate an optimal schedule.

”

Key Variables	FIFO	Optimized Scheduling
Repair sequence	Engine A is allocated parts first because it arrived earlier.	Engine B is prioritized for parts as it can be fully repaired immediately, even though it didn't arrive first.
Bill of Resources	Engine A is missing other critical parts, thus preventing repairs from being fully completed at this time.	Engine B has all necessary parts available, allowing for prompt repair and return to service.
Repair Outcome	Despite receiving parts first, Engine A remains out of service longer due to the other missing critical parts, thus delaying revenue generation. (Engine B also remains out-of-service in this scenario.)	Engine B is repaired quickly and returned to operation, ensuring continuous revenue flow and client satisfaction. Engine A is fully repaired once all necessary parts are available.
Client Impact	Unlike the owner of Engine A, Engine B belongs to a client with no spare engines. This client thus faces significant financial penalties for delays.	Engine B is repaired promptly, avoiding financial penalties and maintaining a strong client relationship.



Key Variables	FIFO	Optimized Scheduling
Overall Efficiency	FIFO overlooks context and urgency, causing inefficiency and financial losses for the MRO and its clients. For instance, aircraft often need simultaneous repairs across multiple modules, making FIFO an impractical approach that leads to poor financial decisions.	Optimized scheduling minimizes downtime and financial losses by considering urgency, resource availability, and overall impact. E.g., it can factor the need for simultaneous repairs across multiple aircraft modules, ensuring more efficient resource allocation and better financial outcomes.

Table 3. A comparison of FIFO and optimized scheduling for engine repairs in an MRO.

The example above is black-and-white, but every real-life scheduling choice is a more granular (and less obvious) version of that scenario. The distinctions between each set of choices are not always clear to the human eye, but a sophisticated algorithm can identify and process all the different variables (e.g., criticality of repairs, required bill of

resources, availability of each resource, financial implications of each choice, etc.) and generate an optimal schedule. Crucially, such an algorithm is able to do this **at scale every single day**.

This kind of algorithmic solution is, for Lokad, the only practical solution to overcome the shortcomings of simplistic approaches like FIFO.

## 5 SOLUTION TO THE SCHEDULING PROBLEM

Given all the constraints and bottlenecks associated with effective scheduling, companies need a solution that is capable of:

- Proactively optimizing POs, theoretical allocation, and staffing plans (macro-level).
- Quickly addressing exceptions as they arise, such as real-time allocation to address a missing resource (micro-level).
- Conveniently visualizing (and white-boxing) the process with dashboards.

Lokad's **Supply Chain Scientists** use a specially created programming language (named **Envision**) to design a unique and adaptive algorithm to address each client's situation and problem(s). These algorithms are called "numerical recipes".

These numerical recipes generate the decisions Lokad provides its clients - for example, an **optimized sequence of actions** for the repair of an engine (including optimized macro-level decisions). This means the client knows:

- ✓ How many parts to order and when to order them
- ✓ What orders (if any) should be expedited to avoid late delivery
- ✓ Where to store parts and when to allocate them

- ✓ What parts, equipment, and skills (people) are needed to complete a sequence of operations
- ✓ Which resources should be allocated at what time and in what place (at both macro and micro-level)
- ✓ When each step should start and how long each step should take

“  
Lokad's Supply Chain Scientists use a specially created programming language (named Envision) to design a unique and adaptive algorithm to address each client's situation and problem(s).  
”

## 5.1 SCHEDULING IN “EMERGENCIES”

No matter how effectively one prepares, there is always the chance that a resource will be unavailable when it is needed. Sometimes inventory arrives late, machinery breaks down, or technicians are sick. When these events strike, the optimized sequence of actions that was originally generated is completely useless.

For example, imagine a process contains 20 steps, and the company requires a **specific technician** to be available to complete Step 10 at a **very specific time**. The value and utility of Steps 1 to 9 and Steps 11 to 20 are all contingent upon Step 10 being completed **by that specific technician at a specific time**. If the technician is absent, the entire 20-step process is ruined. Thus, if that single step cannot be completed, the entire sequence needs to be redesigned.

Redesigning a sequence of actions at the last minute presents companies with a difficult choice :

- Generate a new sequence slowly - feasible but less optimal.
- Generate a new sequence quickly - feasible and more optimal.

The slow option will ultimately generate a high-quality sequence of actions but will take a significant amount of time (approximately as long as the original sequence needed, e.g., several hours). During this downtime, production/repair is halted, and the company loses time and money. Meanwhile, the fast option will ultimately generate a good sequence of actions - less optimal than the original sequence, obviously - **in a matter of minutes**.

In aerospace, where time is an extraordinarily expensive asset, the fast option is preferable to the slow one. To this end, Lokad’s Supply Chain Scientists are capable of generating optimal sequences of actions each day for clients and, importantly, can create new sequences in emergency situations **within a handful of minutes**.

The net result is clients begin each day with a financially optimized set of instructions for the efficient execution of a given process. If a resource is suddenly missing, clients are presented with a brand new sequence to address the emergency in as little as 6 minutes (the time it takes to brew a coffee).

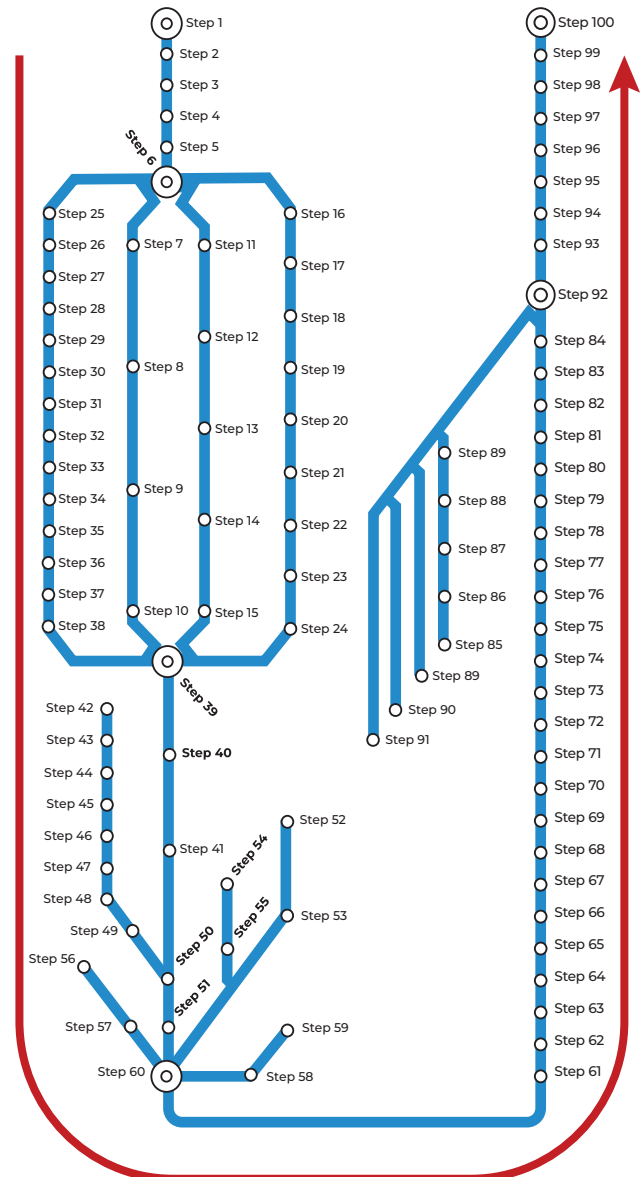


Figure 1. “Subway” PERT map visualizing the intricate and interconnected nature of a 100-step production process.

“

If a resource is suddenly missing, clients are presented with a brand new sequence to address the emergency in as little as 6 minutes (the time it takes to brew a coffee).

”



## 6 HOW LOKAD DOES IT


Lokad leverages several unique supply chain paradigms to address the problems discussed in this article. Firstly, instead of a general purpose programming language, Lokad uses a domain-specific language (DSL) named Envision. This DSL was specifically designed for the optimization of supply chain decision-making - even under extreme time pressure.

Lokad also combines probabilistic forecasting technology to identify the full range of possible scenarios (e.g., demand for a P/N) and the associated likelihood of each scenario. Using this information, in combination with differentiable programming and AI Pilots, our Supply Chain Scientists code solutions that optimize the decisions our clients make. Examples include prioritized lists of inventory purchases, allocation recommendations, and optimal production/repair sequences. Importantly, these lists are **generated**

**automatically every day**, without the need of manual intervention.

Each solution is white-boxed for our clients using **custom dashboards and cockpits**. This allows end users to unpack and understand the information in various degrees of granularity-ranging from top-level summary to fine-grained detail.

### Supply Chain Scientists



### 6.1 SCHEDULING OPTIMIZATION RESULTS

The ideas described here are not theoretical, nor are they impossible. Lokad utilizes the technology described here with large-scale aerospace clients, covering the macro and micro-level planning concerns mentioned earlier (see **Macro vs Micro-Level Planning**).

For example, Lokad can improve the scheduling and assignment of production line workers in an aeronautics manufacturing plant. Such a project is exceptionally challenging given the scale of operations this entails and the enormous bill of resources required. **Table 4** outlines the key optimization constraints Lokad typically considers when providing such a solution.

Optimization Factor	BOR Resource	MRO
Shift Scheduling	People	Align task assignments with workers' shifts and production line operating hours.
Competencies & Certifications	People	Match tasks with workers' skills and certifications.
Task Management	All	Consider task complexity, dependencies, duration, and maximum number of workers allowed simultaneously.



Optimization Factor	BOR Resource	MRO
Machine Tools Usage	Tools	Efficient allocation of machinery for tasks.
Safety Constraints	All	Ensure all safety protocols are met.
Component Requirements	Parts	Manage necessary components for each task.
Supplier Deliveries	People	Schedule based on planned supplier deliveries.
Unforeseen Events	All	Handle unplanned staff absences, delivery delays, execution anomalies requiring the redoing of certain tasks, or administrative delays.

Table 4. Key optimization constraints for a production line scheduling project at a typical large-scale aeronautics company.

To ensure operational agility and high user adoption, Lokad dedicates considerable time to the modular design of the project’s scheduling optimizer. This is a critical tool in the overall success of a typical initiative, and its key features are summarized in **Table 5**.

Scheduler Features	Consideration(s)
Modular design	Facilitate easy modification and addition of constraints without the need for re-coding, allowing for adaptable and scalable solutions.
Flexibility	Adapt to new business requirements through different project iterations with the operational team, ensuring the scheduler remains relevant and effective.
Ease of Implementation	Smooth integration into existing production systems, reducing downtime and enhancing user adoption.

Table 5. Key features for a typical production scheduler.

Lokad applies several performance optimization techniques to its scheduling systems. The ultimate aim is maximizing speed and reliability of outputs while minimizing overheads. See Table 6 for key insights in this regard.

Performance Optimization	Consideration(s)
Parallelization	Enhance computational performance by running multiple processes simultaneously, reducing overall processing time.
Pre-Calculations	Execute "identified pre-calculations" during off-hours (e.g., at night) to save daytime resources and ensure readiness for daily operations.



Performance Optimization	Consideration(s)
Responsive Execution	Combine high computing power for optimal results with quick recalculations to handle daytime changes efficiently, ensuring the scheduler remains responsive to dynamic conditions.

Table 6. Typical performance optimization strategies applied to scheduling systems.

Ultimately, Lokad’s optimized scheduling system provides time savings, enables quick responses to unforeseen events, and provides greater operational insight (up to 2 weeks ahead) into the consequences of decision-making. The full impact of a typical scheduling optimization on production line management is summarized in **Table 7**.

Production impact	Explained
Time Savings	Reduces the time required for planning and allocation, allowing managers to focus on strategic tasks.
Holistic View	Provides a comprehensive overview of the entire production line, enabling better monitoring and management.
Progress Visualization	Allows visualization of planned progress across different time horizons, aiding in tracking and forecasting.
Impact Analysis	Helps to understand the effects of strategic decisions, especially when dealing with unforeseen events, facilitating proactive management.
Informed Decision-Making	Enables quick and informed decisions with full awareness of short and medium-term impacts (up to 2 weeks), enhancing operational agility and effectiveness.

Table 7. Typical impact of a scheduling optimization project on a large-scale production line in aerospace.





## 7 NEXT STEP

To learn more about Lokad's scheduling solution, you can request a discovery call at [contact@lokad.com](mailto:contact@lokad.com).

For more information pertaining to the ideas discussed here, please visit [www.lokad.com](http://www.lokad.com).

### 7.1 SPECIAL THANKS

This paper was produced in close collaboration with several talented and experienced supply chain experts, including Simon Schalit, Luciano Lisiotti, Fabian Hoehner, Tristan Oualid, Rémi Quentin, Victor Noisette, and Alexey Tikhonov. An additional thanks to Timothy Russo and Joshua Bradshaw for generously providing external review.

### 7.2 ABOUT THE AUTHOR

**Conor E. Doherty** is Head of Communication at Lokad, as well as a philosophy and rhetoric instructor at Sciences Po (Paris). Conor has published on a wide array of supply chain topics, including best inventory and forecasting practices. To contact the author, email [c.doherty@lokad.com](mailto:c.doherty@lokad.com) or message through LinkedIn.



## GLOSSARY

AI	Artificial Intelligence
AOG	Aircraft on Ground
APU	Auxiliary Power Unit
BOM	Bill of Materials
BOR	Bill of Resources
DSL	Domain-Specific Language
FIFO	First in, First Out
MRO	Maintenance, Repair, and Overhaul (sometimes “Operations”)
OEM	Original Equipment Manufacturer
PERT	Process Evaluation and Review Technique
PO	Purchase Order
P/N	Part Number
SCS	Supply Chain Scientist





**LOKAD**

Quantitative Supply Chain

[www.lokad.com](http://www.lokad.com)