

FLIGHT PLANNING, THE FORGOTTEN ASSET

Allen Dickinson

Qantas Airways Limited, Mascot NSW, Australia

Abstract

Flight planning is a complex, dynamic function central to the safe and efficient operation of commercial aviation. Despite its critical role, it is often undervalued by airline leadership as a routine cost rather than a strategic capability. Modern flight planning systems can, however, deliver substantial returns on investment, in some cases exceeding those made on other major capital assets such as aircraft. This paper examines the expanding role of advanced flight planning systems across the operational lifecycle, including aircraft evaluation, network and route planning, post-flight analysis, and regulatory compliance. It highlights the technological evolution of these systems, from manual methods using paper charts and slide rules to sophisticated, computer-based platforms incorporating advanced optimisation algorithms and, increasingly, artificial intelligence. A case study is presented of an operator that has leveraged such capabilities to develop highly complex and efficient network operations, demonstrating that these outcomes are contingent on investment in state-of-the-art flight planning technologies.

Keywords: flight planning, flight optimisation

1 Introduction

Airlines around the world are continually looking at ways to improve their operational efficiency. Given the very small operating margins of the industry this is not surprising, with IATA noting that industry wide, the net profit margin is forecast to be 3.9% in 2026. [1]

Fuel will typically account for 20 – 25% of an airline's operating expenses, although Low-Cost Carriers (LCCs), with their lower cost base, may report fuel expenses of up to 45% of the total. Fuel is frequently the second largest expense, after staff salaries and benefits, that an airline will incur.

With fuel costs being a prime driver of efficiency efforts, they are not the only factors at play.

American Airlines notes that: *"Our business has been and will continue to be materially affected by many changing economic, geopolitical, commercial, regulatory and other conditions beyond our control, including global events that affect travel behaviour"* [2] while Delta Air Lines noted in 2023 that: *"Significant damage to our reputation and brand, including as a result of significant adverse publicity or inability to achieve certain sustainability goals, could materially adversely affect our business and financial results".* [3]

Improvements in airline efficiency are therefore driven by multiple factors including:

- The need to control costs, specifically fuel related costs;
- Geopolitical forces;
- Regulatory changes;
- The need to be more sustainable; and,
- Societal expectations.

Esque, Fuchs and Riedel (2022) [4] noted that the aviation industry have made significant gains in terms of fuel efficiency, citing a 39 percent improvement in fuel consumption per passenger-kilometre, between 2005 and 2019. They stated that four factors that had driven this improvement, namely:

https://doi.org/10.71945/icas2026_1003

- Improvements achieved through fleet upgrades, with newer aircraft being considerably more efficient than their predecessors;
- Improvements and adoption of fuel efficiency programs such as engine-out taxi techniques and continuous descent approaches;
- Increased seat density on aircraft; and,
- Increased load factors.

They went on to state that two of these initiatives (seat density and load factors) were approaching their limits and that the industry will need to look at other areas in the coming years to achieve their aims of improved efficiency.

Of the other two initiatives highlighted in the work by Esque et al. (2022) both present difficulties.

Fleet upgrades are expensive and slow to implement. Airlines with relatively young fleets will not be motivated to replace them as the capital cost of their existing fleet may not have been fully amortised. Airlines with older fleets need to finance the capital cost of new aircraft. With net profit margins hovering around 4% financing the required investment becomes difficult with the International Air Transport Association (IATA) [5] noting that the Return on Invested Capital (ROIC) is expected to remain below the Weighted Average Cost of Capital (WACC) in 2026 (IATA, 2025) [6]. This is especially so considering the fact that fleet renewal programs typically occur over many years and that resilience to future geopolitical or other shocks such as the Global Financial Crisis, COVID and the Middle East conflict, may be weakened during periods of high capital investment and associated increased debt levels.

Improvements and adoption of fuel efficiency measures likewise face significant difficulties. Operational techniques such as reduced flap landings, engine out taxi, idle reverse thrust on landing and reducing the amount of discretionary fuel added by pilots are inherently culturally driven. Some operators, with very high levels of engagement from their pilots, are already achieving excellent results. These operators however have, or will soon, reach a natural limit where there is no further benefit to be gained, whilst those with poor levels of uptake face a large cultural challenge to improve performance with Kegan and Laskow Lahey (2011) [7] addressing some of the reasons people will not change, specifically the notion of a “competing commitment”. In the case of pilots who, at their core hold safety as paramount, there may be a perception that some of these techniques may reduce the level of safety. The maintenance of safety standards being their “competing commitment”.

Esque et al. (2022) concluded with a projection that overall fuel efficiency gains across the aviation industry will fall from approximately 3.4% per annum to 1.5% – 2% unless “carriers take more ambitious actions, such as accelerated fleet replacement, transformative infrastructure improvements, or other similarly bold measures”.

Given that there is a need for air carriers across the globe to improve efficiency, especially with respect to fuel consumption what is possible with today’s technology? This question is particularly pertinent given that there are a number of impediments to achieving high levels of identified improvement, including:

- the high costs of fleet replacement programs;
- the fact that some current programs have, or will reach, a natural limit; and,
- the significant cultural change necessary to embed some operational efficiency measures.

Improved flight planning techniques and practices may offer part of the solution yet are systemically overlooked by industry executives. Why is this so?

This paper will attempt to explore why Flight Planning is the Forgotten Asset.

2 What is Flight Planning?

Before exploring the benefits that a modern flight planning system can provide the operator, it is important to understand the scope of operations being considered. It is not a “one size fits all” discussion.

Not all operators need the same levels of complexity, and indeed introducing greater complexity may have a detrimental effect on other aspects of the operation being undertaken and may ultimately compromise safety.

So what is flight planning in the context of this discussion? The Australian Civil Aviation Safety Authority (CASA) [8] states:

> *“Flight planning involves consideration of the straightforward elements of a flight, such as how much fuel to carry, what radio frequencies to use—but it also involves anticipating the unusual and preparing a course of action should it occur.” (Flight planning, always thinking ahead, 2019)*

Aviation Services [9] defines flight planning as:

> *“Flight planning is the detailed process of creating a route for an aircraft before departure. This involves calculating various factors such as weather conditions, fuel requirements, altitude selection, and air traffic restrictions to determine the most efficient and safest path from departure to arrival.”*

These are just two of the numerous definitions that can be found with a simple search of the web.

What these definitions have in common is their emphasis on “a flight”, or, in other words, a single discrete event. It is the contention of this paper that to limit consideration to a discrete event, the true value of flight planning, and the latent value that can be extracted from the associated flight planning systems, is not fully realised.

Taking a step back, it is perhaps not surprising that the emphasis on the flight planning problem is limited to a single event. In the early days flight planning was a manual, time-consuming process, often completed by the pilot, who relied on paper charts, hard-copy weather reports and forecasts. Even today candidates sitting the Airline Transport Pilot Licence (ATPL) Flight Planning exam are tasked with preparing a single flight plan using manual techniques, again using hard copy data.

While having the requisite background knowledge to understand what is being presented on a flight plan, these manual techniques are a source of error and cannot, for practical reasons, consider multiple route options, variations from the “book” value of data and the myriad of other subtle influences that impact long range, heavy aircraft operations. It is for this reason that airlines have almost universally adopted some form of computerised flight planning system.

Yet having a computerised flight planning system does not guarantee either the maximisation of potential value, nor indeed does it guarantee optimisation. To ensure these two aims are met we need to consider the potential that flight planning systems offer and also the question; what constitutes optimisation?

2.1 Potential Use Cases for Flight Planning Systems

As stated there has been an historic emphasis on the role of flight planning and, by default, the systems associated with it, based on a single event. A modern flight planning system however has applications across the lifecycle of an operation commencing before an aircraft has even been ordered through to its ultimate disposal.

In considering the uses and capabilities that a modern flight planning system may offer an operator the following life cycle stages are considered:

- Aircraft evaluation
- Network Planning

- Pre-departure
- Day of Operation
- Post Flight

2.2 Aircraft Evaluation

Which aircraft to buy? This is a perennial problem and one that airline executives spend hours considering.

Having defined the mission/s that the aircraft will be called upon to undertake the critical question becomes which airframe engine combination will provide the greatest return on investment over the anticipated life of the asset. Essentially, from an operational perspective the following are some of the factors that need to be solved for:

- The fuel burn efficiency which is a major component of the operating cost.
- Payload capability which directly impacts potential revenue.
- The potential range capability to understand the potential for future markets.
- Flight time which impacts network schedule design.
- Long term consideration of aircraft and engine performance.

Each of these factors is essentially a flight planning problem albeit with different emphasis associated with each item. A modern system, with the ability to run large batches of flight plans using historical data, provides the capability to conduct large scale multi-scenario analysis that far exceeds the abilities of individual teams. Importantly, given the fact that at this stage of the lifecycle, the aircraft performance data will be provided by the original equipment manufacturer (OEM), it provides a baseline against which to measure the actual, in-service, performance of the chosen platform.

2.3 Network Planning

Network planning and management is widely considered the most critical function of an airline, because of its impact on an airline's financial performance (IATA, 2024). It is the strategic process of determining which cities to serve, flight frequencies, aircraft deployment, and schedule design — essentially shaping the airline's route network.

Again, the flight planning system provides the ability of the operator to determine if the network under consideration is technically flyable. On 24 February 2022, Russia invaded Ukraine creating significant impacts to civil aviation and resulting in the re-routing of aircraft, particularly between Asia and Europe (see Figure 1 below).

The flight planning system would have been used by those operators affected, to evaluate the flyability and profitability of the resulting route variations. An ability to predict the future operational impacts using historic data available in a modern system allows an operator to engage in advanced planning, particularly if a "pseudo" network can be inputted into the system. Such capabilities provide airline executives and network planners a level of insight into the future operation that is not possible using simpler, less sophisticated systems.

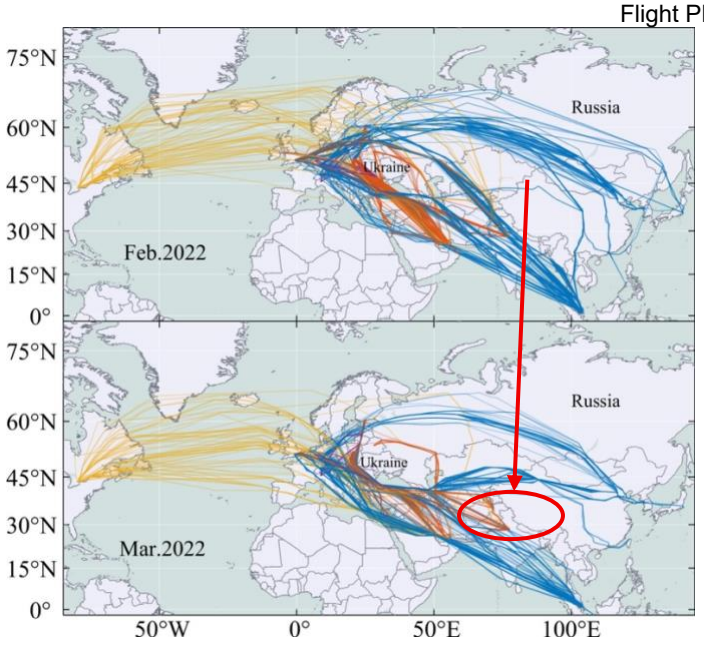


Figure 1 – Flight trajectory comparison before, and immediately after, the start of the Russia-Ukraine conflict [10]

The use of the flight planning system can assist operators to determine the ongoing viability of a particular route, determine the operational impact and pro-actively make informed decisions as to the allocation of assets.

2.4 Pre-departure

One of the great frustrations of passengers is to be told immediately before departure that their flight is not available and that they have been re-routed. Equally frustrating is to arrive at their destination and find that their baggage has not been on-loaded and is travelling on another flight. Often such issues are the result of seemingly minor changes in the forecast winds and/or temperatures along the route of the intended flight.

Equally frustrating, and costly, for the airline is the task of managing disruptions, particularly on flights that are operating at the limits of the aircraft's capability. While typically thought of as an issue for long-haul or ultra long-haul operations it is equally valid on shorter flights where the aircraft is operationally constrained.

Once again, a modern system, utilising weather models that provide operational weather data up to five days before the intended operation, can proactively predict a potential issue and allow planning to take place. Understanding that these models will not have the same level of fidelity as the operational data produced on the day of operation the system can, nevertheless, allow an operator to take a risk-based approach to managing the forecast payload. The operator may potentially decide to re-route passengers, perhaps giving them more choice, or offload freight to other services to free up capacity on the impacted flight.

Again, it is not a "one size fits all" approach as each individual flight may have different operational limitations due to the type of aircraft being utilised, the performance characteristics of individual aircraft (e.g. variations in the deviation from baseline performance) and technical limitations of an individual aircraft etc. A modern flight planning system will consider these, and other, variables and allow the operator to minimise the scale of any disruption and the cost that may be incurred.

2.5 Day of Operation

The planning of an individual flight is, on the one hand, conceptually the most easily understood of the flight planning systems functions yet on the other hand is the area that potentially offers the greatest opportunity for improvement. This is particularly so for smaller operators that do not have the resources

required to either acquire some of the capabilities available in a modern system, or do not have the experience or skills to maximise their effectiveness.

The central question to be asked on the day of operation is, what are we attempting to achieve operationally on this flight? Again, it is potentially not the same for every flight.

A flight that is operating to an airport with a defined curfew may on the one hand seek to minimise fuel burn on one flight where meeting the defined curfew is not an issue but seek to minimise time on another that may be at risk of arriving after the curfew time and be forced to divert.

A modern flight planning system should be able to plan multiple scenarios, offer them to the flight dispatch or operational control department and once a decision is made execute that decision with minimal further input.

Scenarios that may be useful to an operator could include:

- Minimum cost
- Minimum fuel
- Minimum time
- Minimum time within schedule
- Maximum payload
- Maximum tankering opportunity

Note that a minimum fuel plan and a minimum cost plan may not be the same. Fuel is only one of the costs incurred. Crew costs, airframe costs, overflight costs, potential diversion costs etc all contribute to the total and an advanced flight planning system should consider these variables to present solutions to the operator. The flight planning system therefore has a role as a cost control and decision-making system, features that are not required to meet the objectives when planning a single flight.

2.6 Post Flight Analysis

Vital to understanding the performance of any operation is the need to evaluate the relationship between the inputs made and the resulting outputs. In the context of a flight, the flight plan provides the baseline against which the subsequent operation can be measured.

In 2021 GE Digital announced that it had developed a new tool in collaboration with Qantas Airways Ltd. (Qantas) (GE, 2021) [11]. This app, provided to all pilots, allows them to individually analyse their own operation, compare their performance against their peers and review their operational decision making. It compares the flight plan profile to that actually flown and highlights variations, showing the impact of these variations. Reporting a 5.71 million kg avoidance in carbon emissions and a 15% increase in the use of fuel-saving techniques it serves as an example of the benefits of post flight analysis that would not be possible without the underlying flight planning system.

More broadly the flight planning system can, by extracting post flight data, refine the planning of flights to ensure greater accuracy based on actual inflight experience. This continual updating of the data results in a better understanding of the actual environment in which the aircraft operates, allows the operator to make better risk-based decisions, improves confidence in the system which results in lower amounts of discretionary fuel and ultimately lower required amounts of contingency fuel.

Ultimately, the flight planning system could conceptually take post flight data from every flight and build a unique cost model for the specific mission that is about to be flown. This cost model could conceivably take into account the time of day, vagaries that may be specific to a particular route, and the actual aircraft scheduled to fly the mission and build a truly one-off plan for that specific mission.

3 A Case Study – The Constellation Flight Planning System

3.1 Introduction

In October 2018 Qantas launched a new, custom-built flight planning system. Named “Constellation”, the introduction of the new system followed almost 10 years of study, design, development, testing and integration.

The following is a brief, high level description of the system.

Developing an in-house solution was not Qantas’s first choice and followed multiple evaluations of existing systems. What was discovered at the time was the fact that existing systems had core route optimisation algorithms that were conceived in the 1980s. Thirty years on, the understanding of optimisation techniques and the available computing power to support them meant that it was deemed to be an acceptable risk to undertake development of a new system. This decision has resulted in a system that is considered to be one of the, if not the, best systems in the world today.

The Flight Planning Engine is the brain of the system and was developed by researchers at the Australian Centre for Field Robotics (ACFR) at the University of Sydney in collaboration with Qantas. It represents a sophisticated automated system designed to generate optimal flight plans specifically for commercial aviation.

At its core, the system attempts to solve what is fundamentally an enormous and complex problem: given a departure airport, a destination airport, and a set of operational requirements, what is the best possible route for an aircraft to fly, and how much fuel does it need to carry to do so safely and economically?

To appreciate why this is a difficult problem, it helps to consider just how many variables are involved. The route an aircraft takes is not simply a straight line drawn on a map. It must account for wind conditions, temperature at altitude, the weight of the aircraft as fuel burns off during the flight, safety regulations that require the plane to always be within reach of an emergency airport, official airways and airspace restrictions, and the commercial priorities of the airline, which may wish to minimise fuel consumption, flight time, or total operating cost depending on the circumstances.

The ACFR system addresses all of these considerations simultaneously through a series of carefully designed algorithms that work together in a structured, layered fashion.

3.2 The Three-Layer Search Architecture

The system approaches the flight planning problem by breaking it down into three interconnected layers of search, each of which refines the solution produced by the one before it. These layers are known as the path search, the fuel search, and the speed search.

The outermost layer is the speed search, which is responsible for deciding a parameter called how fast the aircraft should fly. Aircraft speed in modern aviation is governed by the Cost Index (CI), which essentially represents a trade-off between the cost of fuel and the cost of time. A low CI instructs the aircraft to fly more slowly, conserving fuel at the expense of a longer flight. A high CI does the opposite, flying faster and burning more fuel but reducing crew and time-related costs. In practice the speed search tests approximately ten different speed settings and selects the one that produces the cheapest overall flight plan, noting however that there is theoretically no limit to the number of speed searches that can be made. The limit of ten represents the most common solutions while reducing computational time.

Inside the speed search sits the fuel search, which is responsible for working out the correct take-off weight for the aircraft. This is important because the total weight of an aircraft at departure — including passengers, cargo, and fuel — directly affects how much fuel it will burn. The fuel search iteratively adjusts the take-off weight until it converges on a value that allows the aircraft to carry the required payload, known as the Zero Fuel Weight, while also satisfying all fuel reserve requirements. It uses a

well-established aviation formula called the Breguet fuel-range equation to efficiently estimate the correct take-off weight without requiring an excessive number of iterations.

At the innermost level is the path search, which finds the actual route — both the lateral path across the map and the vertical profile describing the aircraft's altitude throughout the flight. The path search operates on a fixed take-off weight and speed setting provided by the outer layers and seeks the route that minimises whatever cost function has been selected by the operator.

3.3 Building the Airmap

Before any route-finding can begin, the system must construct what it calls an airmap: a structured representation of all the possible paths the aircraft could legally and practically take between its departure and destination airports. This airmap is essentially a mathematical graph, where locations such as waypoints, navigation aids, and airports serve as nodes, and the permitted connections between them serve as edges.

The airmap is built by combining several different types of routes. The primary component is the conventional airway structure — the network of published flight paths that connect waypoints around the world, much like a road network connects towns and cities. Overlaid on this are free-flight regions, which are areas of airspace where aircraft are permitted to navigate more freely without being constrained to fixed airways. In these regions, the system generates a set of random sample points and connects nearby ones to one another, creating an artificial route network that gives the path search algorithm options to explore.

Additional route types are layered on top of this foundation. Organised Track Systems, such as the North Atlantic Tracks used by transatlantic flights, are pre-planned routes that change daily based on weather forecasts and are incorporated into the airmap when relevant. User Preferred Route directs, which are company-approved shortcuts between specific points, are also included. Where a flight is required to pass through specific geographical locations — known as via points — the airmap is constructed in segments, one for each leg between consecutive via points, and these segments are joined together.

To keep computation manageable, the airmap is intentionally limited in geographical scope. The system defines a flight region shaped like an ellipse centred on the departure and arrival airports, and only includes routes that fall within this boundary. Additionally, for twin-engine aircraft operating under Extended-range Twin-engine Operational Performance Standards rules — commonly known as ETOPS — the entire route must remain within a regulated flying time of suitable emergency airports, and any paths that venture outside this boundary are removed from the airmap during construction. Finally, routes that point significantly away from the destination are pruned from the airmap as they are unlikely to form part of an optimal plan.

The airmap is not static and is continually updated to ensure the most optimal path is available with changing environmental conditions.

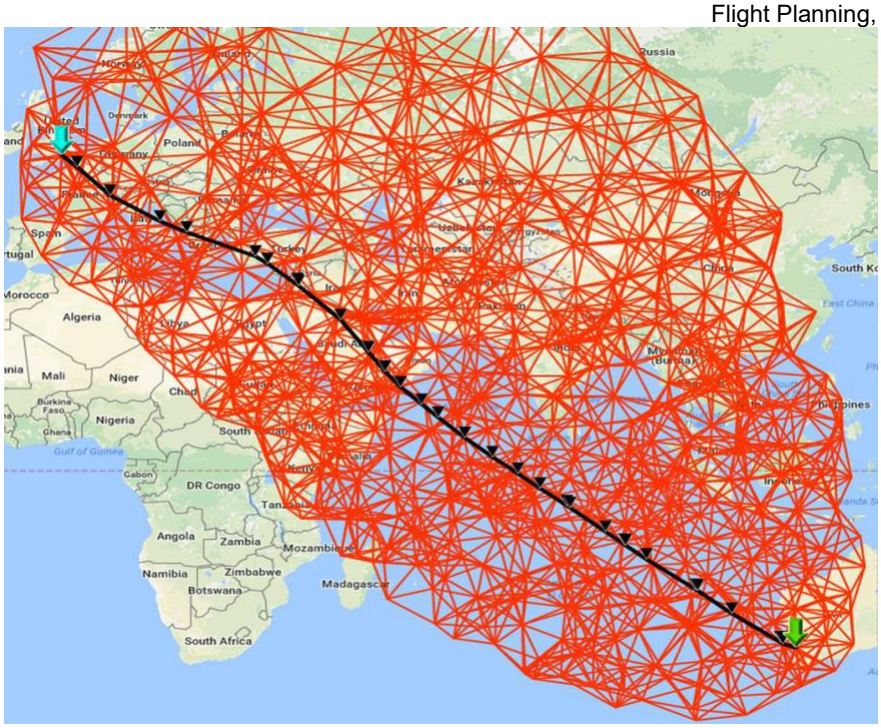


Figure 2 – Example of an airmap constructed for a flight between London and Perth

3.4 Finding the Optimal Route

With the airmap in place, the path search algorithm sets about finding the best route through it. This process is more complex than simply finding the shortest path on a map because each point along the route must be described not just by its geographical location but by a full description of the aircraft's state at that point — including its altitude, weight, the time it will arrive there, and the cost accumulated so far. These rich state descriptions allow the algorithm to account for factors such as altitude-dependent fuel efficiency, airspace closures that are only in effect at certain times of day, and weight-dependent aircraft performance.

The search itself is split into two complementary directions. A forward search starts at the departure airport and progressively builds up the climb and cruise portions of the flight, exploring paths through the airmap in order of increasing cost. Simultaneously, a reverse search works backwards from the destination, modelling the descent trajectory. This separation is valuable because it allows the complex process of finding the right point at which the aircraft begins its descent — the top of descent — to be handled efficiently. The two searches are then joined together in a process called stitching, which identifies points where the forward and reverse graphs meet and recalculates the descent profile with corrected weight and timing information.

After the graph search identifies an initial optimal path, the system applies a process of continuous optimisation to fine-tune the positions of the free-flight waypoints. This uses a well-established mathematical technique that iteratively adjusts the waypoints in the direction that most reduces the cost of the route, in a manner analogous to finding the lowest point in a hilly landscape by always stepping downhill. Once this local optimisation is complete, the refined waypoints are added back into the airmap and the graph search is run again to ensure the final path respects all airspace and vertical profile constraints.

A further refinement step searches for opportunities to perform step climbs and descents — moving the aircraft to a higher or lower altitude partway through a segment rather than waiting until a waypoint — which can yield additional fuel savings by allowing the aircraft to cruise at a more efficient altitude either as a result of it becoming lighter or a change in the environmental conditions that may warrant a change in altitude subject to the aircraft's capabilities and any defined restrictions.

3.5 Safety Contingencies and Emergency Fuel

A significant portion of the system's complexity is devoted to ensuring that the aircraft always carries enough fuel to handle emergencies. Three specific types of contingency planning are incorporated into the flight plan, each addressing a different scenario.

The first, known as DPD, or Designated Point Depressurised, ensures that from any point along the route, the aircraft has enough fuel to divert to at least one suitable airport if a depressurisation of the aircraft occurs.

The second, DP1, or Designated Point One Engine Inoperative, addresses the scenario of an engine failure, ensuring the aircraft can always reach a safe airport even while descending from cruise altitude on a reduced power setting.

The third, CPE or Critical Point ETOPS contingency, applies specifically to twin-engine aircraft flying over remote areas and requires that sufficient fuel be carried to reach an ETOPS-approved airport at any point where the aircraft is more than sixty minutes from such an airport.

In the case of DPD and DP1, the system identifies what is called the critical point — the location along the route where diversion to an emergency airport would require the same amount of fuel and ensures that this worst-case fuel requirement is covered. In the case of CPE, an equi-time point is calculated.

Finding these critical points involves searching along the planned route for the point at which the cheapest available diversion airport changes from one airport to another, indicating an equal-cost point between two options. Binary search techniques are used to locate these transition points precisely.

The system also plans a separate alternate flight plan, calculating the fuel and route required to fly from the arrival airport to a nominated alternate airport in the event that landing at the destination is not possible.

3.6 Fuel Policy and Reversion

Once the path and contingency calculations are complete, the fuel search applies a detailed fuel policy that specifies all the fuel that must be loaded onto the aircraft. This includes the trip fuel needed to fly the planned route, a fixed fuel reserve that must always be maintained, a contingency allowance representing a percentage of the trip fuel, the fuel required to divert to the alternate airport, any special holding fuel for anticipated air traffic delays, and planned additional fuel requested by the operator. The total of all these requirements determines the take-off weight, which feeds back into the path search in the next iteration.

There are occasions when the sum of all fuel requirements exceeds the capacity of the aircraft's fuel tanks, or when carrying all planned fuel would result in a take-off weight that is too heavy to be safe. In these situations, the system applies a process called reversion. Reversion involves systematically removing fuel from certain permitted categories — called revertible fuel buckets — until the plan becomes feasible. Some categories of fuel, such as mandatory fixed reserves, cannot be reduced regardless of circumstances and are termed non-revertible. The reversion algorithm must carefully track which reductions are permissible and ensure that removing fuel from one category does not inadvertently make the plan non-compliant with requirements in another.

A related concept is fuel tankering, which applies when fuel prices differ significantly between airports. If fuel is substantially cheaper at the departure airport, the system can plan to load extra fuel at departure to avoid purchasing the more expensive fuel at the destination, within the constraints imposed by weight limits and tank capacity.

3.7 Speed Selection and Final Output

Once the fuel search has converged on a valid take-off weight for a given speed setting, the outer speed search compares the total cost of flight plans produced at each of the tested Cost Index values and selects the one that best satisfies the operator's objectives. Where two speed settings produce equivalent results, the higher Cost Index — representing a faster flight — is preferred. Crucially, if a

particular speed setting results in the aircraft being unable to carry the required payload, that option is rejected regardless of its cost advantage, as preserving payload capacity takes precedence over cost optimisation.

The final output of the system is a complete flight plan specifying the optimal route, altitude profile, take-off weight, speed schedule, and all associated fuel quantities. This plan is the product of many hundreds of individual calculations working in concert, each informed by real-world data including weather forecasts, aircraft performance tables, airspace restrictions, and commercial fuel policies.

3.8 Conclusion

The ACFR Flight Planning Engine represents a carefully engineered solution to one of aviation's most computationally demanding problems. By decomposing the flight planning challenge into structured layers — route finding, fuel calculation, and speed selection — and by employing a combination of graph search, continuous mathematical optimisation, and iterative refinement, the system is capable of producing high-quality, fully compliant flight plans that balance safety, regulatory requirements, and commercial efficiency.

The depth of the system's design reflects both the complexity of modern commercial aviation operations and the sophistication of the algorithmic techniques that computer science and applied mathematics have made available to solve them.

4 The Results

The comments in this section are, by necessity, general in nature due to the potentially sensitive nature of the results. Accordingly, a few case studies are used to illustrate some of the advantages that have materialised.

At the time of implementation Qantas expected to realise an efficiency gain of between 0.6% and 1% which would have equated to a fuel saving of up to $40M per annum (Hatch, 2018) [14]. System performance has met or exceeded expectations.

Since 2018 a number of direct comparison studies with another flight planning system in common use has consistently shown that the Qantas system is able to find more efficient flight paths as a result of its search methodology. Figure 3 shows a flight from Sydney to Honolulu. Constellation initially tracked on a more easterly heading to take advantage of the increased tailwind. In this example the Constellation system saved 32 minutes of flight time and had a fuel saving of 8.7%.

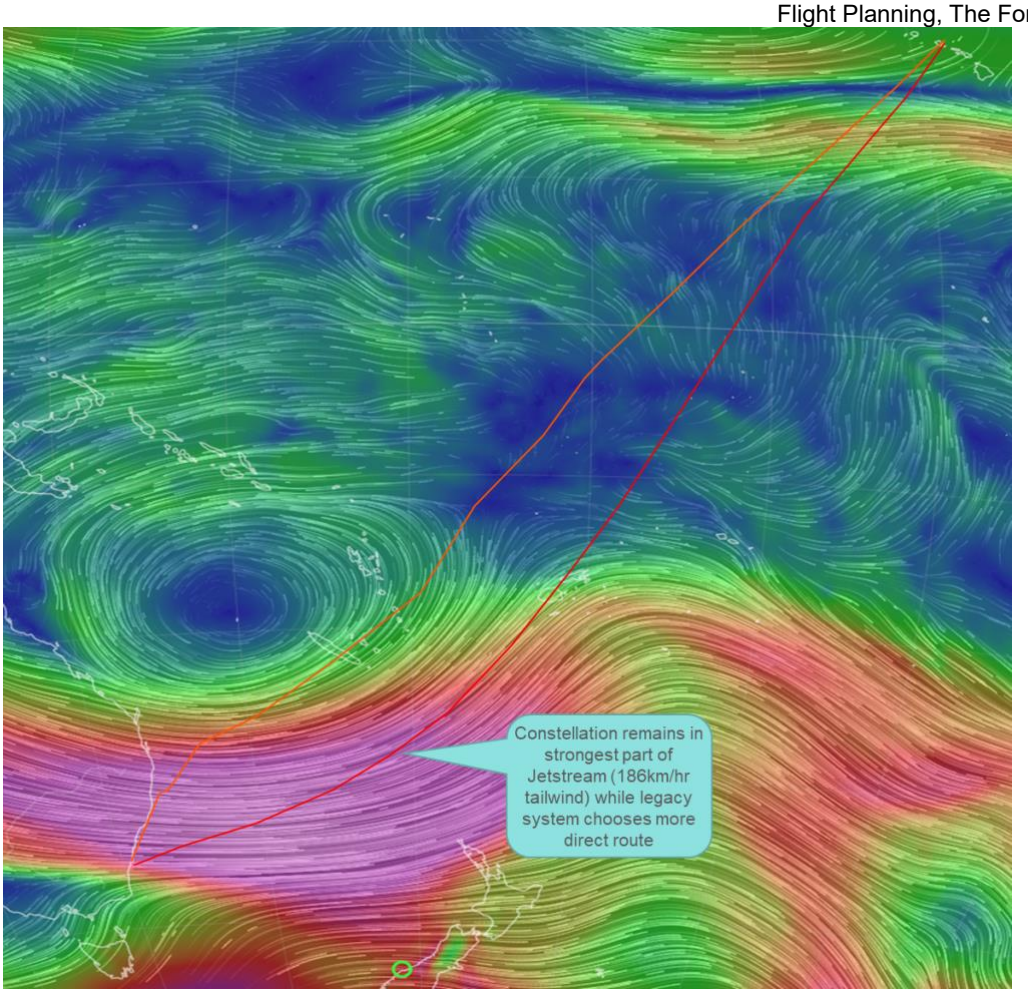


Figure 3 – Sydney to Honolulu flight path comparison

While the above example may appear extreme it was not the only example that demonstrated this scale of saving. Although this example is an outlier, the savings demonstrated on this route were typically in the range of 2.5% – 3.5%.

In one study of 1,400 flights, the Constellation system demonstrated potential fuel savings of approximately 2% across the network when compared to the other system.

Compared to the previous flight planning system used by Qantas, Constellation provided what may initially appear to be counter-intuitive results. In one cited case (Hatch, 2018), a flight to Johannesburg was planned to fly 160 nautical miles (300 kilometres) further than it would normally, but in doing so cut the headwinds it experienced by two-thirds. The aircraft arrived only three minutes later than scheduled and saved more than a tonne of fuel. The question often asked is: “How can an aircraft fly for longer yet use less fuel”? There are a number of reasons why this may be so, but in the case of the Constellation system Qantas deliberately designed the system so that the planning system matched the actual behaviour of the aircraft. In the case of the FMC the Mach Number flown for a given CI is impacted by the environmental conditions experienced. This directly impacts the fuel flow. Constellation models this behaviour and is therefore able to accurately trade time for fuel or vice versa.

From a flight dispatch point of view significant savings were realised due to the level of automation built into the system. Using the previous flight planning system, which lacked the automation of Constellation, the Perth to London Flight would take a dispatcher approximately 5 hours to plan. Constellation now plans the same flight in less than 15 minutes with greater accuracy and efficiency.

From a Qantas perspective the system has been a success albeit not without difficulties during its development. The development of the Constellation system importantly demonstrates that a modern, technically advanced flight planning system can deliver significant benefits to an airline.

5 Does it Make Financial Sense?

Does it make financial sense to invest in a new flight planning system? The answer is very much: it depends!

The building of a new system such as Constellation is a massive undertaking, filled with risk and, for most operators, this risk could not be justified.

In the case of Qantas, it had a very specific set of requirements (Riegler, 2020) [17]. Among these Qantas sought to:

- Maintain its unique fuel policy that allows it to undertake ultra long-haul operations. No flight planning system available offered this capability.
- Qantas wanted the system to be future-proofed and capable of adopting new technologies and planning practices including:
  - The introduction of System Wide Information Management (SWIM);
  - Flight and Flow Information in a Collaborative Environment (FFICE);
  - Trajectory Based Operations (TBO);
  - The ability to directly exchange data with ANSPs, hence the adoption of an AIXM database; and,
  - Use of cloud-based computing to improve speed, allow scalability and increase redundancy.

Few if any operators will have a similar set of requirements that would justify building their own system. That does not mean however that they should accept the status quo and not seek to improve the performance of their flight planning system either through replacement or upgrade.

Let's consider a hypothetical, extremely simplified, scenario for a moment and consider an airline that is looking to introduce a new fleet of 50 A321s as replacements for an existing aircraft type (It could be any other type but the A321 is used purely as an example).

The list price for an A321XLR is approximately $140M USD (ACC Aviation, 2025) [12] but let's assume that the airline can obtain them for $70M. Total capital cost is therefore $3.5 Billion. These aircraft have an efficiency premium of approximately 35% over older aircraft that it is replacing. The purchase of a new aircraft is, in this example, gaining significant efficiency but at the considerable cost of approximately $100M per percentage point of efficiency improvement. Importantly, the efficiency gain is only available to that fleet of aircraft.

It logically follows that smaller operators may choose not to invest in an advanced flight planning system due to the scale of operation. Let's take an example of an operator with 10 aircraft. Using the same numbers the cost per percentage point of efficiency is:

10 aircraft × $70M = $700M cost of fleet. Divided by 35% efficiency = $20M per percentage point of efficiency. In this case the operator may choose not to explore the opportunity.

If the operator can invest in a new, or upgraded, flight planning system and achieve a defined percentage increase in efficiency at a cost less than the cost of the same percentage gain per aircraft then it is worthy of consideration when measured solely against the purchase of a new fleet.

The flight planning system however has additional advantages over the purchase of an aircraft, including:

- The efficiency gain achieved by a flight planning system is available to all aircraft in the fleet, old and new.
- The efficiency gains achieved across the fleet may allow an operator to meet additional Environmental, Social and Governance (ESG) benefits.
- The run costs of a new flight planning system will be substantially lower than that of an aircraft.
- Once established there are few ancillary costs such as the cost of simulators, spare parts etc.

- Given that there would no doubt be greater automation with a new system there are likely manpower savings available.
- A new system has the potential to deliver operational benefits as previously described, again these benefits are available to all aircraft in the fleet.

The above is a hypothetical example designed more than anything to start a conversation. Clearly any business case would be very complex, and no two operators would be the same, but it does serve to illustrate that funding the flight planning system can be a worthwhile endeavour and deserving of consideration.

6 Why is Flight Planning the Forgotten Asset?

The title of this paper is: “Flight Planning, The Forgotten Asset”. Why this paper, why this time?

With backlog orders of more than 15,000 Airbus and Boeing aircraft (Holve, 2026) [15], it is clear that airlines are undergoing a significant re-fleeting program, the value of which will be in the tens of billions of dollars. These aircraft will no doubt offer considerable benefits in terms of greater fuel efficiency, with associated lower carbon emissions. These benefits will take decades to deliver.

These same operators will, in many cases, continue to use the same flight planning systems with little consideration as to what efficiency benefits could be gained from a comparatively modest investment. Conceptually operators could combine forces to either develop a new system or drive change to their existing systems yet there seems to be little imperative to do so. Given that airlines are demanding greater efficiency to reduce costs, become more environmentally sustainable, meet new and emerging societal expectations, and improve safety it seems counter-intuitive that a system that can deliver these benefits does not gain more attention.

What makes senior executives reluctant to drive change with respect to flight planning?

Not surprisingly there is little, if any literature or studies that examine this issue and as a result the contents of this section offer some potential thoughts rather than any definitive statement of fact.

- Flight planning is hard and few people fully comprehend the process of flight planning (Aviation 360, 2023) [13]. Even within a typical airline there may be no one that understands the complexity of the flight planning system or the inputs that are going into producing a valid plan. Almost inevitably there will be no one with a sound understanding of flight path optimisation, this almost exclusively being the domain of academics, the aircraft manufacturers and NASA. It is not surprising therefore that senior executives faced with a request to fund enhancements to a flight planning system do not have the knowledge to adequately evaluate the request. Worse, they do not have people available who can articulate the problem in simple terms.
- The true value of the flight planning system is not appreciated, and it is viewed as a cost rather than a strategic asset. All too often an airline executive will ask, “How much does the flight plan cost?” If one were to answer, “12 dollars”, then it is likely to be countered by a statement that their planning system only costs $10 per plan and therefore they are not interested. The true cost of a flight plan is more accurately the cost of production PLUS the lost opportunity. The $10 plan, if it produces a plan that will cost 500kg of fuel more than the alternative system is costing approximately $510, far more expensive than the alternative. This lack of understanding is another reason flight planning upgrades face difficulty gaining funding.
- Associated with the fact that executives focus on the cost rather than the benefits is that any business case will inevitably have a negative Net Present Value (NPV). If associated benefits, particularly with respect to fuel savings were included the NPV would almost certainly be positive and may actually exceed the NPV of other capital expenditure being considered.
- There is confusion between accuracy and optimisation. Supposedly intelligent, skilled, operationally savvy people in an airline will state that their current system is producing good results as the actual fuel burn of the aircraft matches the plan. This is accuracy and is the entry level criteria of any flight planning system. Optimisation is different and involves the best possible

outcome. Faced with a chorus of operational staff extolling the accuracy of a current system senior executives will inevitably fail to understand that there may be better alternatives.

- Flight planning is seen as the domain of the flight dispatch department. Analogous to pilots who can operate an aircraft with extreme levels of skill the flight dispatchers are very skilled at their core tasks. Just as a pilot might design an aircraft with 12 engines (you can't be too careful!) flight dispatchers will define a set of requirements that may make their working life easier but do not actually improve efficiency. Just as aircraft purchases are the domain of the finance department so too the purchase or upgrade of the flight planning system should be based on sound financial decisions and organisational need.



Figure 4 – An aircraft designed by pilots?

- Flight planning systems are boring! Many of us grew up building model aircraft. We like looking at a bright shiny aircraft and senior executives can speak glowingly of the business class flatbed seat or the latest French champagne that will be served to customers on the new aircraft. Given the choice between selling the vision or selling the latest computer screen the aircraft wins out!

It is not the fault of the airline executive that they do not appreciate the value of a modern flight planning system. Very few, if any, people in their organisation would truly understand the potential benefits that could be unlocked with system improvements.

7 Conclusion

The purpose of this paper was not to propose that every airline needs to build or purchase a new flight planning system. Rather it was intended to illustrate that there are significant operational advantages to be gained from having a system that is truly fit for purpose and that is designed to maximise the operational efficiency.

The example of Qantas is perhaps an extreme one. An operator, driven by specific needs, building an entirely new system inhouse. It was a high-risk strategy but one that is showing huge benefits and is achieving efficiencies that would probably not have been achieved had an alternative path been taken. It serves as an example of what can be achieved.

For airline executives there is a clear message that the flight planning system that your organisation uses can be either an incredibly powerful strategic asset or a system that is failing to deliver the most optimal solution for your operation. It is probably time to take a closer look at the system and not let it languish in the flight dispatch department.

Whilst building a new system may not be practical, having active discussions as to the features needed to maximise efficiency and provide the level of functionality that one needs to operate their business can lead to significant operational improvements.

Beyond the individual airline, what does the future hold for flight planning systems?

Professor Salah Sukkarieh from the Australian Centre for Field Robotics provided the following thoughts:

The next step beyond single-flight optimisation is network-wide optimisation. Today most systems optimise one flight at a time. The frontier is treating the airline as a coupled system: fleet, crew, slots, maintenance, and trajectories solved together, because a locally optimal flight plan can easily be globally suboptimal once you account for the rest of the operation.

Three further shifts are coming. Stochastic optimisation, where the planner explicitly accounts for forecast uncertainty rather than assuming a single weather scenario. Real-time in-flight reoptimisation, where the plan adapts continuously as conditions evolve, rather than being locked at dispatch. And environmental optimisation, including contrail avoidance and non-CO2 effects, which is moving from research into regulatory expectation faster than most executives realise.

There is also Sustainable Aviation Fuel. The point here is not SAF itself, but what SAF does to the flight planning problem. Today's optimisers largely minimise a single thing, fuel against time, subject to constraints. SAF converts that into a genuinely multi-objective problem that has to be solved as one: fuel cost, SAF cost and availability, tankering trade-offs, performance penalties from lower energy density, and emissions trade-offs across CO2, NOx, particulates, and contrails. None of these move in the same direction. A plan that is optimal for cost is rarely optimal for emissions, and a plan that is optimal for CO2 may be worse for contrails or NOx.

The inputs are also genuinely uncertain. SAF availability at destination airports is patchy, credit schemes mean the actual blend onboard may not be known until shortly before departure, and the science on contrail formation under SAF is still unsettled. This is exactly the kind of problem that spreadsheets and rules of thumb cannot solve, and exactly where a proper optimisation platform shows its value. As environmental regulation tightens and SAF moves from voluntary to mandated, the airlines that can solve this multi-objective problem operationally, not just in policy documents, will hold a measurable cost and compliance advantage. That is the next leg of the investment case.

Underneath all of these, the common thread is the same. Better models, better search, better decisions under uncertainty.

Contact Author Email Address

The contact author email address is: allendickinson76@gmail.com

Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The author also confirms that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] Airline profitability stabilizes with 3.9% net margin expected in 2026, 2025, *Press Release No. 54, International Air Transport Association (IATA)*, viewed 11 May 2026, https://www.iata.org/en/pressroom/2025-releases/2025-12-09-01/

[2] American Airlines Group Inc. *American Airlines Group Annual Report, Year Ended December 31, 2024*. Author, Fort Worth.

[3] Delta Air Lines Inc. *Delta Air Lines Annual Report, Year Ended December 31, 2023*. Author, Atlanta.

[4] Esque A, Fuchs G and Riedel R. Fuel efficiency: why airlines need to switch to more ambitious measures. *McKinsey & Company*, 2022, https://www.mckinsey.com/industries/aerospace-and-defense/our-insights/future-air-mobility-blog/fuel-efficiency-why-airlines-need-to-switch-to-more-ambitious-measures

[5] IATA. How to work in airline network planning: a critical role in airline management. 2024, viewed 28 May 2026, https://www.iata.org/en/publications/newsletters/iata-knowledge-hub/how-to-work-in-airline-network-planning/

[6] IATA. Airline profitability stabilises with 3.9% net margin expected in 2026. *IATA Press Release No. 54*, 9 December 2025, viewed 15 May 2026, https://www.iata.org/en/pressroom/2025-releases/2025-12-09-01/

[7] Kegan R and Laskow Lahey L. The real reason people won't change. *Harvard Business Review's 10 Must Reads on Change Management*, Harvard Business School Publishing Corporation, Boston, pp 119-136, 2011.

[8] Civil Aviation Safety Authority. Flight planning, always thinking ahead. 2019, https://www.casa.gov.au/sites/default/files/2022-10/flight-planning-always-think-ahead-booklet-vfr.pdf

[9] Aviation Services. Flight planning: overview and basic terminology. Viewed 11 May 2026, https://an.aero/flight-planning-overview-and-basic-terminology/

[10] Airspace closure challenges: exploring the impact of the Russia-Ukraine conflict on flight operations and pathways to solutions. *Transportation Interdisciplinary Perspectives*, Vol. 31, May 2025, https://www.sciencedirect.com/science/article/pii/S2590198225000752

[11] GE. GE Digital's FlightPulse mobile app provides airline pilots insight from flight prep to flight assessment. *GE Press Release*, 09 March 2021, viewed 01 June 2026, https://www.ge.com/news/press-releases/ge-digitals-flightpulse-mobile-app-provides-airline-pilots-insight-from-flight-prep

[12] ACC Aviation. Airbus A321XLR estimated pricing (2023-2025). https://accaviation.com/about-us/

[13] Aviation 360. Airlines turn to AI to attain ultra-long flight destinations. 23 April 2023.

[14] Hatch P. How subtle changes to flight paths will save Qantas millions a year. *The Sydney Morning Herald*, 10 December 2018, https://www.smh.com.au/business/companies/how-subtle-changes-to-flight-paths-will-save-qantas-millions-a-year-20181126-p50icy.html

[15] Holve G. Airbus and Boeing report May 2026 commercial aircraft orders and deliveries. *Flightplan, Forecast International*, 10 June 2026, https://flightplan.forecastinternational.com/2026/06/10/airbus-and-boeing-report-may-2026-commercial-aircraft-orders-and-deliveries/

[16] Qantas Airways Ltd. *Qantas Annual Report 2025*. Author, Sydney.

[17] Riegler M. Case study: building innovative route optimizing flight planning platform at Qantas. *Aircraft IT OPS, March/April 2020*, https://www.aircraftit.com/articles/case-study-building-innovative-route-optimizing-flight-planning-platform-at-qantas/

[18] Ryanair Group. *Ryanair Group Annual Report 2025, Year Ended March 2025*. Author, Swords.