HELIOS

Optimising airport performance

A careful balancing act

HELIOS ADVISER

Analysis & commentary for
decision makers in the transport industry

THE QUEEN'S AWARDS FOR ENTERPRISE 2009

Abstract

Airport operators, air navigation service providers, their customers and regulators have long known that optimising performance is a delicate balance of the key performance indicators. The main components of operational performance include efficiency; quality of service, mainly punctuality and delay; operational resilience; and environmental impact. Increasingly the regime is being broadened to encompass the entire passenger experience.

$$\textbf{Efficiency} = \frac{\text{Useful outputs}}{\text{Inputs}} \qquad \textbf{Utilisation} = \frac{\text{Capacity used}}{\text{Capacity available}}$$

Arrival punctuality = Actual arrival time − scheduled arrival time

Delay = Actual time taken − planned time taken

Resilience = ... the abilities to *withstand* or *recover* quickly from difficult conditions (Oxford English Dictionary)

This briefing paper describes a framework for operational performance assessment covering both efficiency and quality of service whilst being mindful of operational resilience. As many significant operational processes at airports and within air traffic management, e.g. security screening and air traffic control, are heavily dependent on human resources, application of the framework is described using these as real-life examples.

Although this framework has been developed with the air transport sector in mind, it is equally applicable to many other sectors.

Introduction

Whether they play the role of service provider, operator, economic regulator or indeed customer, optimisation of operational performance is becoming increasingly important to the protagonists along the air transport infrastructure value chain.

Performance improvement is a key objective for all involved in air transport

Clearly commercialised infrastructure providers – be they airports or air navigation service providers – desire to improve their performance in line with their strategic business objectives, whilst their customers wish to be assured that monopoly services are being delivered in an efficient and effective manner to meet their requirements. Equally, however, it is important that economic regulators have a good understanding of the performance of the regulated company in order to set realistic yet challenging performance improvement objectives. The main components of operational performance in airports and air traffic management (ATM) are, of course, efficiency, punctuality, operational resilience and environmental impact and the concept is increasingly being broadened to encompass the entire passenger experience.

Key operational performance areas include efficiency, operational resilience, punctuality, environmental impact and the passenger experience

Performance can be defined and assessed in a multiplicity of ways and at various levels: for the enterprise, for the business unit or for the service line. However, the most detailed picture is obtained by examining performance at the operational process level. Compiling a group of the most significant operational processes allows a picture to be drawn of the overall performance of the organisation, including areas of good or poor performance and improvement objectives.

A useful framework for operational performance analysis is based on the inputs (resources deployed), the outputs (the quantity delivered) and the outcomes (the quality delivered, associated price and environmental impact), but must also consider interactions with and between third parties, e.g. airlines, as well as the external drivers: e.g. complexity, regulatory environment, geography, etc. This framework is illustrated in the following figure.

Generic framework for operational performance analysis

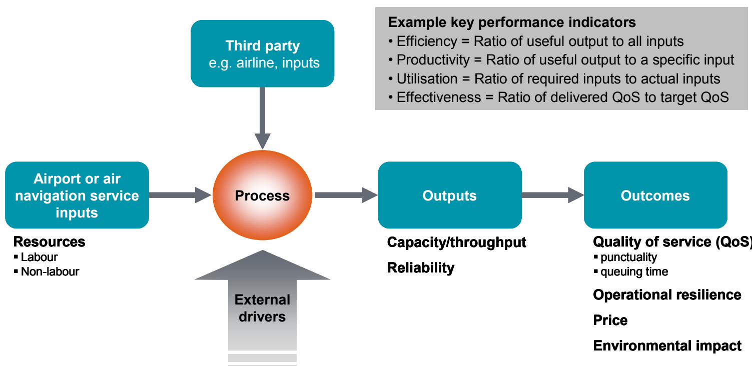

Measurement of operational performance must address both efficiency and service quality to paint a complete picture

The remainder of this brief document describes how this framework can be applied to assess efficiency and quality of service whilst being mindful of operational resilience based on real-life experience. As the most significant processes at airports are heavily dependent on human resources, e.g. air traffic control and security screening, examples discussed here are focused on this type of process.

Measuring efficiency

Staff deployment, through rostering, is the key input to many operational processes

The starting point for efficiency measurement is the definition of the appropriate key performance indicators (KPIs) or metrics.

Using productivity as an example, in a human resource intensive operation or process the input is a measure of the staff deployed, e.g. the number of air traffic controllers or security agents on duty each hour. This is particularly dependent on the rostering process.

Traffic throughput, be it passengers or aircraft, is the key output of many operational processes

The output is the throughput of the service offered, e.g. flight hours controlled per hour for air traffic control or hourly passenger throughput for security screening. Care must be taken in defining the measure of output to ensure that it is representative of the task being undertaken, e.g. the number of flight hours controlled per hour is in many cases a better representation of the output of an air traffic controller's task than simply the number of flights controlled per hour.

Inputs and outputs are derived from operational and planning data at the relevant level of granularity, for example on an hourly basis. In principle, derivation of the KPI – in this case staff productivity – is a straightforward division of the output by the input as illustrated in the following figure. The KPIs can be described in terms of their statistical properties. Statistical probability distributions are easily obtainable from the sample, as illustrated in the next figure.

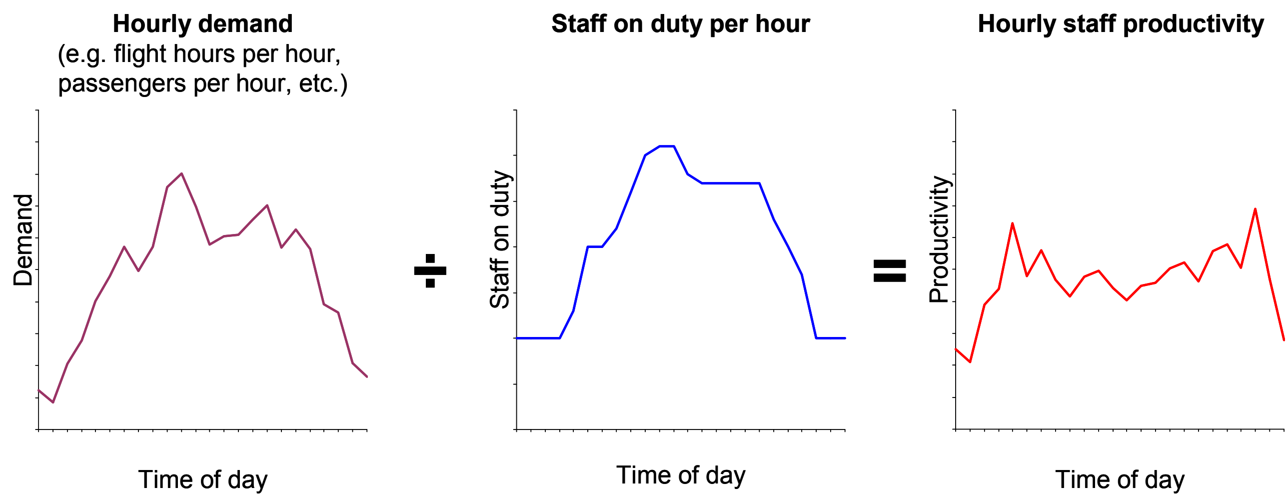

Productivity is the ratio of traffic throughput to staff deployment

In our experience the vast majority of efficiency data from various ATM and airport processes follow a normal distribution facilitating subsequent analysis.

The two most important parameters derived from the probability distribution are:

- the **mean**, which is the measure of average productivity (high mean = high productivity)
- the **standard deviation**, which gives an indication of how consistently the process matches supply with demand (low standard deviation = consistent match).

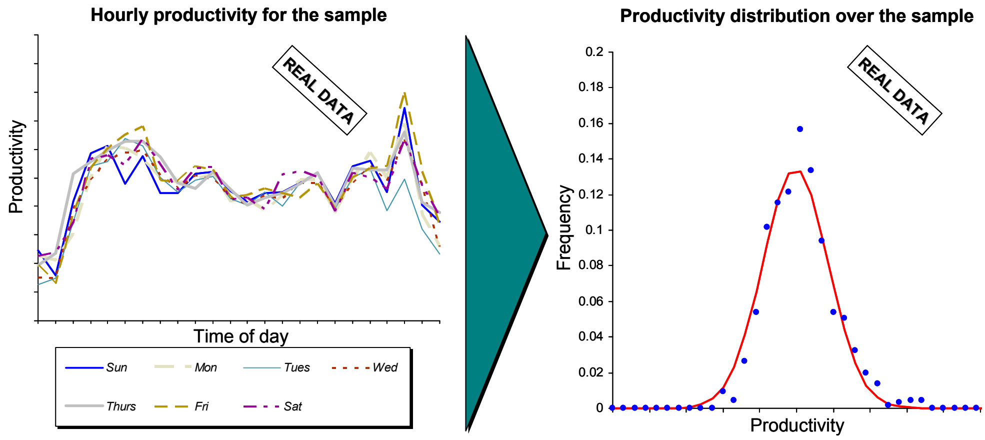

Efficiency KPIs are often best described in terms of their statistical properties

These statistical distributions can be used as the basis for **benchmarking**, that is comparing efficiency across a suitable range of comparators. The pros and cons of benchmarking are highlighted below.

These statistical descriptions can be used as the basis for both benchmarking and bottom-up analyses

In addition to benchmarking, the process can also be examined **bottom-up** in detail, for example in terms of the precise evolution of the productivity over time and detailed assessment of the match of supply and demand in terms of the target for the quality of the service. This document focuses on bottom-up techniques.

The evolution of the hourly productivity curve throughout the day often follows the general pattern illustrated in the following figure for both airport and ATM processes.

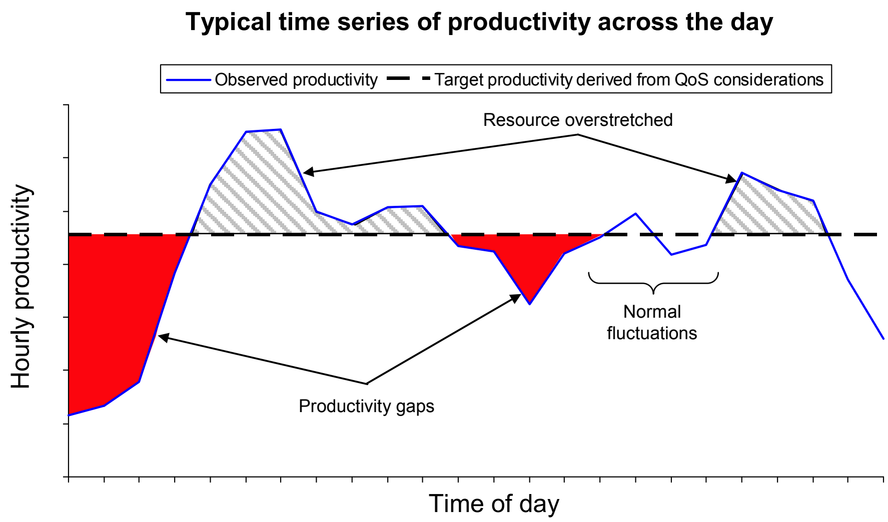

Periods of excessively low or high productivity are easily identified graphically from their time series

At some times the productivity is way above the target for extended periods and at others it is below the target. This implies that the system supporting the process oscillates between periods of being over-stretched (and possibly delivering poor quality and low resilience, e.g. long queues) and slack periods (and delivering inefficiently) where there are productivity gaps. It is clearly impractical to expect the productivity ever to match exactly the average but the oscillations should be small in amplitude, within a target upper and lower bound, and short in length. Note, as a bottom-up assessment, that this type of analysis gives little indication as to whether the average productivity is good, bad or indifferent – this comes from the comparative analysis of benchmarking.

Overly high productivity is likely to be associated with poor quality of service and low resilience

A similar picture can be derived from looking at another KPI – staff utilisation. Here the level of actual staff deployment is compared to that predicted as being required, based on analysis of the process and quality of service targets. This is illustrated in the figure below.

Utilisation is a measure of the staff deployed to the staff needed – again related to rostering

Assessment of staff utilisation

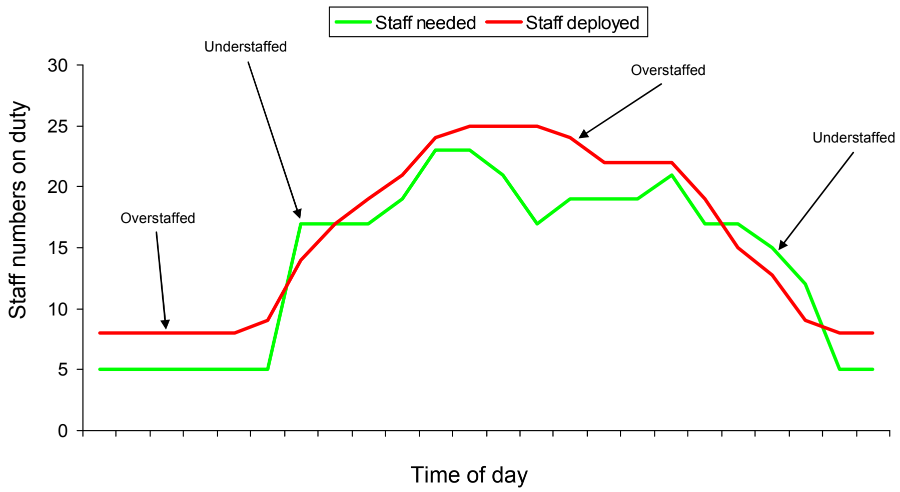

The figure shows periods where under and overstaffing occurs. These can be quantified in terms of parameters such as:

- the average utilisation (gaps between the value and 100% representing downtime)
- undersupply, defined as the periods where the minimum number of staff required exceeds the number on duty
- the converse, oversupply.

Again, periods of excessively low or high utilisation are easily identified graphically from their time series

Short and random periods of oversupply are very difficult to address but it is important that longer and systematic mismatches are addressed in the rostering process:

- undersupply can result in ineffective performance, e.g. delays, long queues and, importantly, poor resilience, that has a knock-on effect on customers, which can cause severe disruptions especially if occurring early in the day when the entire schedule can be affected
- oversupply implies inefficient use of resources with direct cost consequences.

Measuring quality of service

Queues

The amount of time spent queuing is a key indicator of quality of service

In both our examples of passenger security screening and air traffic control, the key measure of quality of service is the amount of time that the customer (ultimately the passenger) spends in a queue, be it waiting to pass through an archway metal detector (AMD) in a terminal building; sitting on the aircraft at the departure gate waiting for air traffic flow management clearance; holding on the taxiway; or in airborne holding patterns such as the stacks familiar to frequent travellers at busy airports. If queuing time is generically described as delay, then in all cases the quality of service measures are those of the delay experienced and the predictability of that delay. Effectiveness is defined in terms of the relationship between the realised delay and the target delay or queuing time.

Both the average queue length and its predictability are important - leading to a statistical analysis approach

As with efficiency, the delays are described by statistical distribution functions with the mean of the distribution being indicative of the average delay and the standard deviation being indicative of the predictability of the delay. The familiar queuing theory delay curve can also be derived from the delay distributions as illustrated in the following figure for stack holding typical of a busy airport. As capacity is a somewhat intangible quantity and is related to the acceptable level of delay, the delay curve below has been calibrated such that the capacity of the arrivals runways is set for an average 5-minute hold per inbound flight, as shown by the dashed lines.

Stack holding delay distribution and associated delay curve

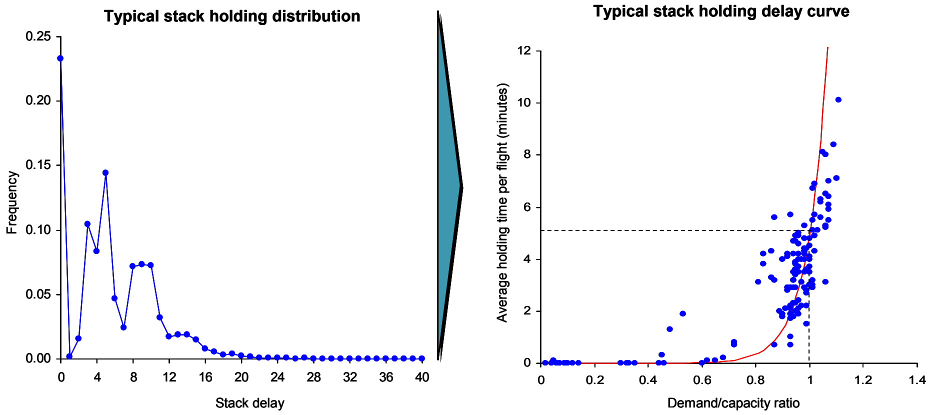

Source: Airline data

The capacity of the system is related to the maximum acceptable length of time spent queuing

Points to note in the figure are the peaks in the holding time distribution corresponding to the time taken to complete a circuit in the stack and the exponential form of the delay curve. Delay curves for other types of queue, e.g. air traffic flow management, show different forms as queuing relationships are dependent on the type and behaviour of the queue, principally whether the inputs and outputs are ordered or random in nature.

Different types of operation show different types of queue behaviour

The delay curve derived from historical data can be used to predict the impact of changes in demand (scheduling changes) and operational enhancements on delay without the need for complex simulations as long as the changes are not too large. Thus the statistical/queuing theory approach is ideally suited to investigating multiple what-if scenarios.

The shape of the delay curve at the point of operations also gives a good indication of the operational resilience for the system. If most operations occur at the steep part of the curve, small fluctuations in demand and/or capacity will have a large impact on delay and hence the system is likely to suffer from low resilience.

The point on the delay curve at which the system is operating gives an indication of its resilience

Punctuality

The length of time spent in queues is one measure of the quality of service that the system is delivering. However, more relevant in the passenger's perception is punctuality - measured as the difference between the actual

Punctuality is the KPI that is most tangible in the passenger's perception

time of an event and its scheduled time, most notably arrival at the destination airport. Punctuality can also be described through statistical distribution functions. Experience in analysing punctuality indicates that a variety of distributions occur as shown in the following figure.

Different types of punctuality distribution are indicative of different types of behaviour

Illustration of some typical types of punctuality distribution

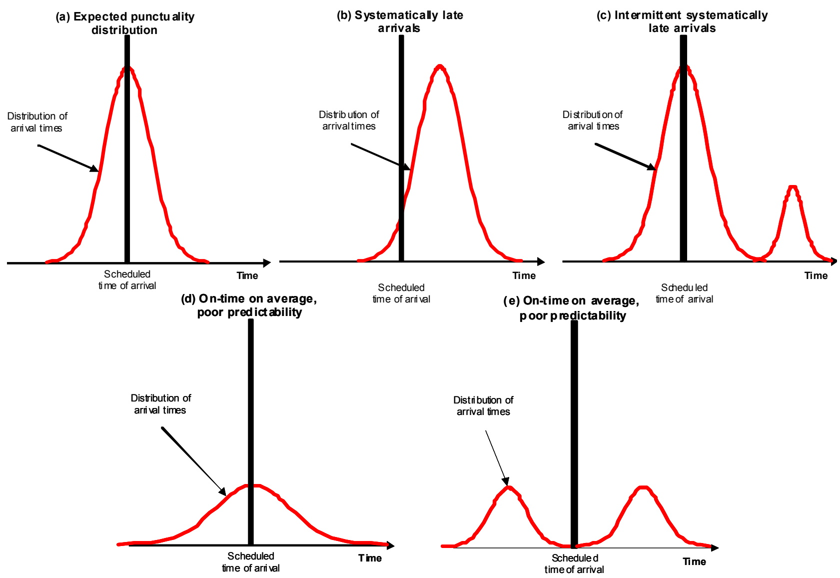

These distributions show why the standard deviation of the distribution is important as well as the average. Figure (d) shows a distribution that has on-time average punctuality but with a low, broad distribution indicating a very broad range of arrival times around the schedule where very late arrivals are cancelled out by very early arrivals. Figure (e) shows a more extreme case where arrivals are either early or late (with none actually on time) but again with late arrivals being cancelled out in the average by early arrivals.

Operationally, flights that arrive very early are as disruptive as those that arrive very late

In addition, although from the passengers' perspective, arriving early is (usually) beneficial, from the airport's and other operators' perspectives, early arrivals can be as disruptive as late arrivals. The optimum, therefore, is for flights to arrive within a narrow window centred on their scheduled

Punctuality distribution

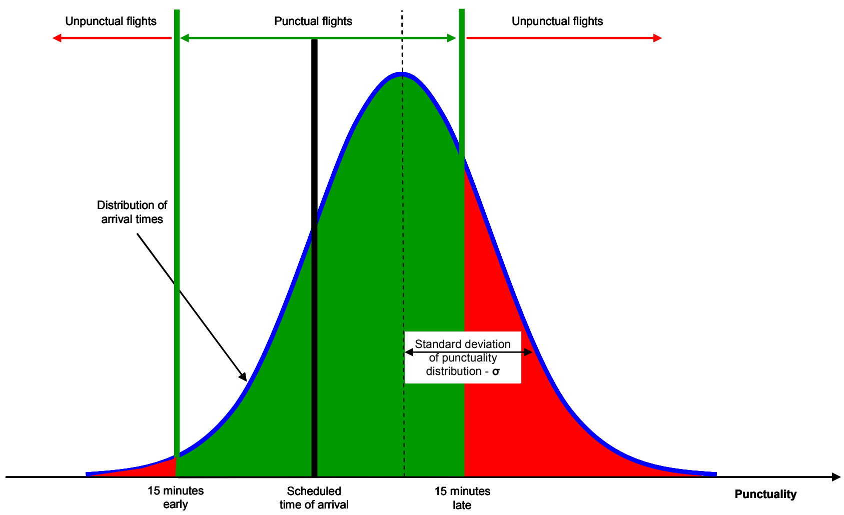

arrival time. This window is illustrated in terms of the punctuality distribution of a particular flight in the following figure for an arbitrary ±15 minute window. Those parts of the distribution highlighted in green comply with the punctuality target and those in red do not.

From this it is possible to derive a confidence level, based in historical data that the flight will arrive within the punctuality target based on the mean and standard deviation of the distribution. Various confidence levels are shown in the following figure where, if the mean/standard deviation combination lie in:

Statistical analysis gives an indication of the level of confidence or expectation that a flight will be punctual

- the blue zone, then the punctuality confidence level is 68% or better
- the yellow zone, then the punctuality confidence level is 95% or better
- the green zone, then the punctuality confidence level is 99% or better.

Punctuality confidence levels

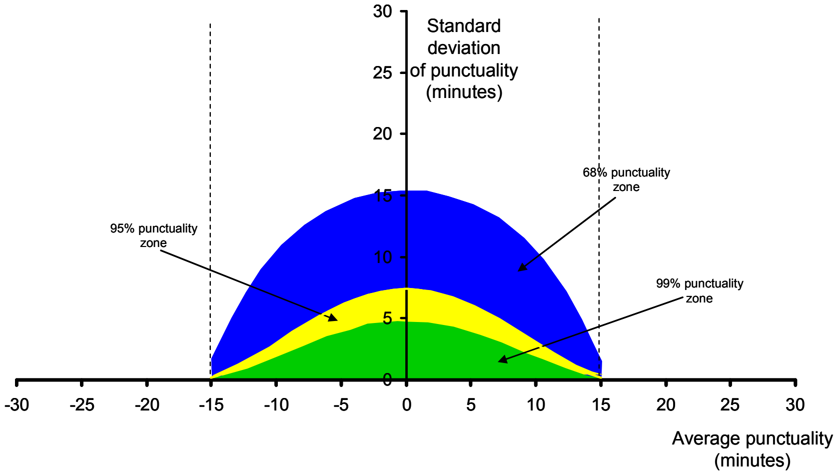

Good punctuality at a high confidence level is very challenging

Experience in analysing real punctuality data indicates that meeting a target of 15 minutes to a high level of confidence is extremely challenging.

Bringing it all together

Our simplified examples of efficiency and quality of service measurement above show that, in human resource intensive processes, both are related to the ratio of the demand to the number of staff deployed. The key performance indicators for efficiency and quality of service can also indicate when the system is over-stretched and hence are proxy indicators for operational resilience.

Bringing the results of all of the analyses together adds insight into identification of root causes and the impacts of potential improvement programmes

Examination of the correlation of punctuality performance with efficiency, queue length and operational resilience is also useful in identifying and fixing the root causes of poor punctuality. Notwithstanding the small variations in efficiency that must occur between individual security agents or air traffic controllers and the fact that capacity comes in large lumps (new security lanes opened, new sectors opened, etc), combination of the analysis techniques allows for identification of shortcomings in productivity or utilisation due to, for example, the rostering system, and also the impact that this would have in the queue length and punctuality. Conversely, the techniques could be used as a planning tool to identify the levels of staff needed to meet a target queue length.

Adding benchmarking

Benchmarking gives insights through comparison that are not possible with bottom-up analysis

The techniques outlined above could be described as bottom-up assessment in that they use detailed operational data to derive KPIs and assess these in detail internally to the organisation to understand and address any fluctuations. They do not give any indication as to whether the achieved efficiency and quality of service is good, bad or indifferent. This latter assessment can be made through benchmarking, where the KPIs and associated operating practices for a particular organisation's process are compared externally with the KPIs and operating practices describing the same (or similar processes) in other organisations.

Benchmarking and bottom-up assessment are complementary techniques that bring added insights when applied together

The difficulties and 'apples-compared-to-oranges' shortcomings of benchmarking are well-known and are not repeated here. However, the following table compares benchmarking and bottom-up assessment techniques, and shows that they are complementary. Separately they both give useful insights but used together, they can give the complete picture.

Pros and cons of complementary analysis techniques

Benchmarking	Bottom-up assessment
Outward looking	**Inward looking**
1. Black-box	1. Looks inside the box
2. Data light	2. Detailed data needed
3. Relatively quick analysis	3. Analysis can be complex
4. Sample dependent but sample gives point of reference	4. No external data needed but gives no point of reference
5. Needs adjustment to account for different external drivers across sample	5. No adjustment needed for self-consistent data
6. Gives relative position within the sample showing whether performance is reasonable without scope for improvement if best in class	6. Always scope for improving the process but gives no indication whether achieved results are reasonable compared to peers
7. Based on statistical descriptions of KPIs	7. Based on actual KPI values and their evolution
8. Best practices can be identified – can identify whether process should be changed	8. Best practices not identified but gives indication of how process can be improved

Concluding remarks

Performance improvement must be approached holistically to balance the impact on the KPIs

Measuring and improving operational performance must consider the efficiency, quality of service and operational resilience elements of a process, principally due to the conflicting effects that the improvement levers can have. Furthermore, our experience of measuring performance in airports and air traffic management shows that benchmarking and bottom-up techniques are complementary. Separately they give useful insights but used together, they can give the complete picture.

Statistical techniques are very useful to analyse multiple what-if scenarios quickly and quantitatively

Finally, the techniques described are based on statistical and analytical approaches rather than large-scale simulations. These techniques, although based on historical data, can provide very useful insight and are extremely flexible and efficient in assessing multiple what-if scenarios where the simulation alternative offers more detail but is cumbersome. They also offer the advantage that they are easy to understand and sensitivities to the key variables can be identified and tested.

Although the examples are based on airports and air traffic control, the techniques are easily transferable to other industries.

About Helios

Helios is a management and technology consultancy specialising in airports, air traffic management and air transport more generally. We support central and local government, regulators, airport operators, development agencies, master planners, engineering companies and investment banks.

We help our customers to:

- measure, understand and improve airport operational performance and resilience, both airside and landside
- determine the capacity of current and proposed airport infrastructure, both landside and airside, as well as understand the capacity implications of growth in demand and schedule changes
- understand the potential of new technologies and procedures to reduce environmental impact
- develop strategies for route retention and development.

Our success has been recognised through two Queen's Awards for Enterprise (in 2004 and 2009).

Get in touch…

For further information please contact:

Dr Michael Fairbanks

Helios
29 Hercules Way
Aerospace Boulevard
AeroPark
Farnborough
Hampshire
GU14 6UU
UK

E mike.fairbanks@askhelios.com
T +44 1252 451 651
F +44 1252 451 652
W www.askhelios.com

HELIOS

Helios
29 Hercules Way
Aerospace Boulevard | AeroPark
Farnborough | Hampshire | GU14 6UU | UK

T +44 1252 451 651
F +44 1252 451 652
E info@askhelios.com
W www.askhelios.com