Designing Efficient Taxi Pick-up Operations at Airports

by
David Carvalho Teixeira da Costa
ANA - Aeroportos de Portugal, SA
Rua D, Edificio 120, Aeroporto de Lisboa
1700-008 Lisboa, Portugal
dccosta@ana.pt

Richard de Neufville
Engineering Systems Division
Massachusetts Institute of Technology
Cambridge, MA 02139 (U.S.A.)
ardent@mit.edu

Abstract

This paper provides a practical procedure for designing efficient taxi pick-up operations at airports. How to do this effectively is an open question. Solutions are not available. Practices differ widely, reflecting different approaches to, and lack of research on the subject. The solutions are often unsatisfactory. At many airports, passengers routinely suffer long waits outside, exposed to the elements, after a tiring journey. Such disagreeable experiences are avoidable.

Designing efficient taxi pick-up operations at airports is problematic. The peculiarities of this process preclude easy solutions. First, it involves queuing, so system performance is a highly non-linear function of the loads. Second, it features unstable transient situations as travelers typically arrive in bulk over short periods. Third, traffic is significantly differentiated, consisting of a wide variety of groups implying different service characteristics. Standard results from queuing theory thus do not apply usefully to this problem.

The design process therefore uses simulation based on detailed observation of the local practices. It involves 4 steps: (1) detailed local measurements of both the arrival of travelers, and the service rates provided by taxis in different queuing positions; (2) creation and validation of a simulation model sufficiently detailed to account for these realities; (3) exploration of design alternatives to estimate the characteristics of the service that they would provide; and (4) selection of a preferred design that properly balances efforts to minimize average and extreme wait times. The paper demonstrates this procedure through application to the Lisbon Airport.

Key Words: Airports, Taxis, Terminal Design, Simulation, Queuing Theory

Designing Efficient Taxi Pick-up Operations at Airports

by
David Carvalho Teixeira da Costa
ANA - Aeroportos de Portugal, SA
Rua D, Edificio 120, Aeroporto de Lisboa
1700-008 Lisboa, Portugal

Richard de Neufville
Engineering Systems Division
Massachusetts Institute of Technology
Cambridge, MA 02139, U.S.A.
ardent@mit.edu

The Problem

Taxi pick-up operations at airports are sources of frequent complaints. Passengers seeking to get taxis to complete their journeys routinely have long waits outside, exposed to the elements, after a tiring journey. For example, international passengers arriving in Boston wait outside, in eventual rain, wind and snow, for taxis to arrive sporadically from a holding lot about 2 kilometers away. Even when taxis are available, travelers cannot board them rapidly due to the design of the boarding area. Similarly, passengers arriving in Lisbon can easily queue up for 20 minutes even though the taxi pool is full – and only 100 meters away.

The problem of designing efficient taxi pick-up operations is important at the many airports where taxis service a significant fraction of the arriving passengers. Taxi operations are most important for airports close to their cities, particularly where public transit is either minimally available or provides poor service to travelers encumbered with bags. La Croix (1991) and Cardon (2007) have documented this logical conclusion. Boston and Lisbon are good examples of this situation, along with Bangkok, Chicago, Delhi, Frankfurt, Los Angeles, New York, Paris, San Francisco, Singapore, Sydney, Toronto and many others. Table 1 illustrates the importance of taxi operations at selected airports.

No coherent guidance is available to airport operators for how to design efficient taxi pick-up operations, despite the importance and prevalence of the issue. Essentially no information is available on the topic. As a practical matter, both the international standard design manual (IATA, 2004) and the major current textbooks (Ashford et al, 2011; Ashford and Moore, 1996; de Neufville and Odoni, 2003; Horonjett, 2010; Kazda and Caves, 2007) are silent on the issue. The references to taxi services are limited to a general view of curbside access and taxi availability, from nearby segregated "pools". These sources neither address the efficient design of taxi service, of passenger waiting areas and their interface, nor provide guidance on how to handle the mismatch between their distinct operational behaviors.

Table 1 - Taxis serve a significant portion of passengers at many airports

City/Airport	Taxi Share (%) of Passenger Traffic
Amsterdam	18
Brussels	20
Copenhagen	33
Frankfurt/Main	19
London/Heathrow	26
Madrid	40
Manchester (UK)	28
New York/Kennedy	42
Paris/de Gaulle	37
Paris/Orly	27
Rome/Fuimicino	32

Source: Kazda and Caves, (2007)

In this context, this paper proposes a comprehensive, coherent, and practical approach for designing efficient taxi operations. It comprehensively deals with the entire problem, from collecting data, through the analysis, to presenting results meaningfully for airport decision-makers. It is coherent in that it consistently deals with the reality of transient queuing processes. As demonstrated by the application to a major airport, the method is practical. Specifically, the method provides a process for:

- Gathering the kind of data needed the characterize taxi pick-up operations;
- Calibrating the data to peculiarities of local situation;
- Extrapolating this reality to alternative designs;
- Presenting the results showing the impact of alternative designs according to the range of significant metrics (such as average and maximal wait times, and the reliability of the process).

The proposed method thus provides airport managers with the kind of information they need to make informed decisions about how to best suit their taxi operations to local conditions.

Current Practice

Given the lack of detailed guidance, it is not surprising that airports have adopted a hodge-podge of designs and practices for their taxi operations. As regards the taxis, airports variously line up them in single, dual or even triple files for taxi queues in front of passengers. Single file arrangements are perhaps most common. Dual files, as practiced at Washington/Reagan and

Lisbon for example, have the advantage of offering the waiting passengers more taxis close to the head of their queue. Some airports, such as Singapore, use parallel loading, which permits taxis to exit their queue as soon as they are ready, without waiting for other taxis to complete their loading process. Figure 1 illustrates these alternatives.

Airport practices also differ in how they treat the passenger side of the taxi pick-up process. Some allow passengers to board from multiple points; others force them to pass through a single point of access. Some airports let passengers go directly to available taxis as they arrive, others -- such as Boston, constrain passengers to wait behind a barrier until a new batch of taxis is in place.

In short, there is no agreement on how to design and operate taxi pick-up operations. While we should expect that different solutions are preferable for different situations, it is safe to say that prevailing practices typically are historical legacies rather than the result of conscious analysis and consideration.
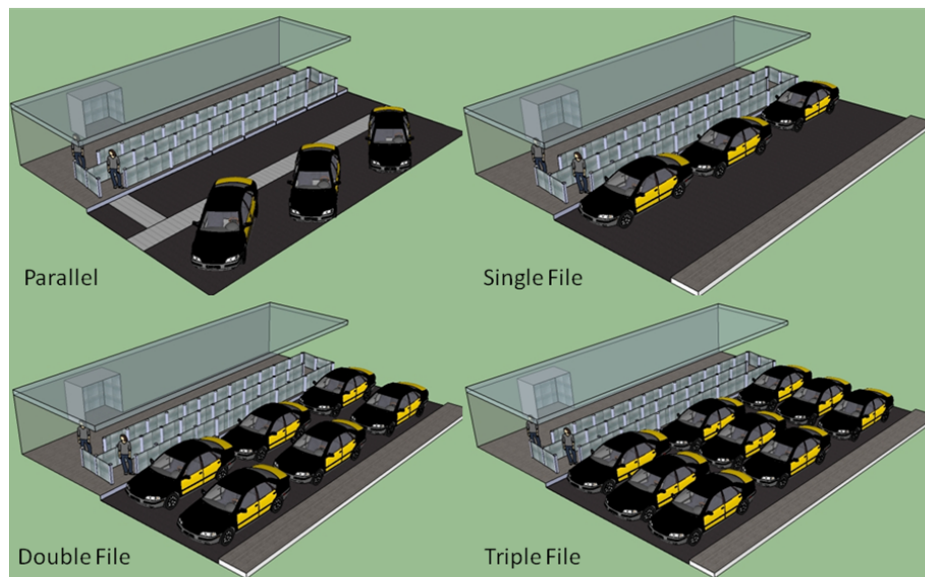

Figure 1 – Types of Configurations for Taxi Queues

Nature of the Problem

The problem of designing efficient taxi pick-up operations at airports is technically challenging. The analysis is conceptually difficult and the structure of the problem does not allow for closed form solutions. These facts go a long way to explaining why general solutions are not available. Understanding the nature of these difficulties points the way to developing an effective procedure for defining efficient designs.

The analysis of taxi pick-up operations is conceptually difficult because it involves queuing. That is, people arrive looking for service [to board a taxi], they queue up behind the position(s) at which service is offered, and then exit the process [when they leave in a taxi]. The essential difficulty with a queuing problem is that its performance depends significantly and non-linearly on the state of the system.

In the first instance, the level of delays and the length of the queues depend greatly on the relative load on the system. Both the delays and the unreliability of the system in steady state increase rapidly as the rate of arrivals nears its nominal capacity (Figure 2a). Further, this non-linearity increases as the rate of arrivals at the point of service becomes more irregular (Figure 2b) and as defined by Equation 1 and 2 for a Poisson arrivals process.
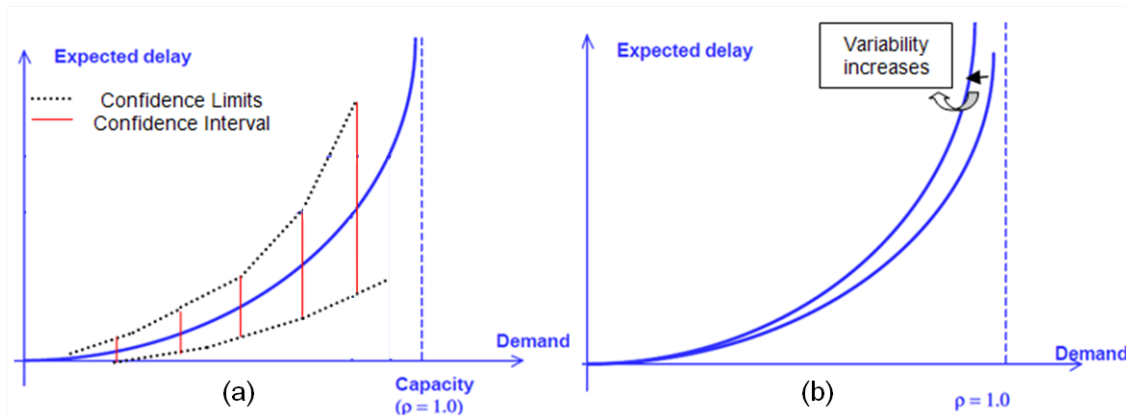

Figure 2 - Delay versus Utilization ratio ρ in steady state:(a) general curve and (b) as modified by unsteadiness of arrivals (Source: Odoni, n.d.)

Analytically, an expression for this non-linearity in steady state is

$$W_q = \frac{\lambda \; [(\frac{1}{\mu})^2 + \sigma_t^2]}{2(1-\rho)} = \frac{\lambda \; [E^2(t) + \sigma_t^2]}{2(1-\rho)}$$

Eqn.1

Where W_q is the expected waiting time in queue, λ is the average rate of arrivals, and μ is the average rate of service, that is the nominal capacity of the system. It follows that $(1/\mu)$ is the expected value of the service time, $E(t)$. In addition, σ_t is the standard deviation of the service time. Finally, ρ is the utilization ratio:

$$\rho = \lambda / \mu = \text{(average rate of arrivals) / (nominal capacity of process)} \qquad \text{Eqn.2}$$

As can be seen from Equation 1, when the rate of arrivals approaches the nominal capacity of the system, so that ρ comes close to 1, $(1-\rho)$ comes close to zero and delays increase rapidly.

This phenomenon leads counter-intuitively to the fact that the "practical capacity" of the system – that we can think of as the quantity of arrivals that could be served within acceptable delays and levels of reliability – is substantially less than the nominal capacity. Specialists understand this phenomenon, which provided the basis for the old practice of defining the capacity of runways in terms of their ability to handle arrivals with less than a stated number of minutes of delay (Horonjeff, 1962). Yet managers often miss this reality as they wrongly assume that a service system will operate satisfactorily up to its nominal capacity. Thus the designers for the automated baggage handling system for Denver International Airport failed to account for this and were dismayed to find that their system became too unreliable – and thus effectively at "practical capacity" -- whenever loads exceeded around 40% of its rated capacity.(de Neufville, 1994).

The analysis of taxi pick-up operations is further complicated because it involves transient loads, that is, a fluctuating rate of arrivals for service. The rate at which travelers arrive at taxi stands typically varies significantly over the course of any hour and any day, as various flights arrive. Periods of intense demand for taxis follow long slack periods. The resulting queuing process is thus rarely in steady state. There is then no closed-form description of the queuing process of the kind provided in Equation 1. Theoretical queuing analysis thus does not apply in useful detail.

Finally, the analysis is difficult because the stream of passengers arriving for taxi service is highly differentiated. Customers arrive in groups of various sizes and capabilities to deal with the process of boarding taxis. For example, a single business traveler coming to town for a day may hop into a taxi and move on quickly. In contrast, the family of four foreigners arriving for a vacation will require much more time as the load their many bags and try to explain their destination to the driver. Thus, the rate at which the taxi pick-up system can provide service is also a mixed bag of processing times. In this respect queues for taxi operations are quite distinct from many other queuing operations in airports, in which travelers come in groups but are processed as individuals, for example at security check points.

Additionally, the rate at which individual taxis can provide service depends on their position in the queue. Specifically, the rate at which the second or third taxi in a line can complete service [as by leaving the queue and leaving space for a replacement] usually depends on the speed at which the taxi(s) ahead of it provide service. This last point is particularly important for the design of taxi pick-up operations. Alternative designs for these processes position taxis differently, as Figure 1 shows. In order to evaluate the performance of design alternatives, we therefore need to have a method of analysis that is sensitive to the positions of the taxis in the alternative designs for the operations.

In summary, the nature of taxi pick-up operations at airports precludes theoretic, analytic solutions. Furthermore, as the processing rate of the pick-up operation depends greatly both on the kind of passenger group being served and the position of a taxi in the queue, we need a method of analysis that is sensitive to such details. A general, all-purpose model will not do. Analysts therefore depend on some form of detailed simulation carefully calibrated to the realities of a specific situation. This is the approach presented in this paper.

Review of Literature
This section reviews the elements relevant to the design of taxi pick-up operations at airports:
1. analyses of taxi pick-up operations;
2. practical guidance about managing queues;
3. approaches to dealing with transient queues; and
4. pertinent topics in simulation.

Analyses of Taxi Pick-up Operations:
To our knowledge, there are only four reports on the subject within the open domain. Notably, as indicated in the introduction, none of the standard textbooks or manuals on airport design and operations analyze taxi operations.

The first is that of Lee (1966) who presented a detailed analysis of the West London Air Terminal. His work stresses the necessity to examine carefully the behavior of arrivals and service time, in particular the distribution of groups of passengers that translate into single taxi hires and trips.

Curry and De Vany (1977) modeled the transient buildup and dissipation of taxi and bus queues at Dallas/Fort Worth Airport. They approximated the transient levels of arrivals as a series of steady-state cases, which they used to calculate expected wait times and other measures of performance. They did not suggest a general procedure for analyzing alternative ways to design the taxi operations.

Costa (2009) analyzed the taxi operations in detail at Lisbon International Airport. This paper reports on these results, as extended by his subsequent work as an employee of the operator of Lisbon and other Portuguese airports, ANA, Aeroportos de Portugal, SA.

Finally, based upon Costa's work, Ferreira, Pacheco and Teodoro (2010) tested a specific proposed alternative design of the taxi operations at Lisbon International Airport though means of closed-form formulaic approximations for important system characteristics.

Practical Guidance about Managing Queues

For designing taxi pick-up operations, perhaps the most important guidance from queuing theory is the idea that operators can improve overall service by providing differentiated service to customers with distinct service times. This is the principle behind providing express lanes for customers who can be served quickly at checkout counters for grocery stores and the like. These fast lanes expedite service and shorten the wait time and queue length for the overall process.

The concept of fast lanes opens up the possibility of better service in situations where operators could implement the equivalent of fast lanes. These might occur at airports with large traffic at where it might be sensible to have multiple pick-up points, separating destinations, as between central business areas and suburban destinations for example. This approach is common practice for regional bus service in many hub airports, notably in Asia at Tokyo/Narita and Seoul/Incheon. In these cases, passengers queue up in many different positions along the curb, each assigned to specific regional destinations.

It is also worth mention, because it analysts often forget, that small seemingly trivial changes in processing arrangements can have substantial effects on overall performance. Saving 15 seconds per operation might seem insignificant, but when these operations take two to three minutes to execute – as typical of many airport services and of taxi pick-ups in particular—such savings can increase total capacity by 10% or more.

Particularly relevant here is the practice of pre-positioning customers from a long queue in front of the place where they will receive service. At busy airports of entry, the US Immigration services thus now direct international arrivals, who start out in a single snake queue, into secondary queues just in front of the kiosk and officer who will process their documents. This practice minimizes the time between when one customer exits service and the next appears. It contrasts with the more general situation, as at check-in counters, where customers may not leave the head of the snake queue until they notice that one of the service channels is open. When the

open channel is 10 or more meters away, as can easily happen at busy airports, customers will take time both to notice that the service is available, and to move to the agent providing the service. The result is that this agent will have more time between customers, will provide lower capacity and the overall performance of the entire service process will be less than it could be. The lesson here is that the design of taxi pick-up operations should pay close attention to details that can expedite service and improve performance.

Newell's Approach to Dealing with Transient Queues

Newell (1982) developed and presented a practical way to estimate service quality provided by service systems under transient loads. Briefly, his approach assumes that as traffic builds up it is served without delay until arriving traffic equals the prevailing capacity of the service system. Further traffic would be served at the rate of this capacity, and customers in excess of this capacity would accumulate in a queue and would be served eventually as the number of customers arriving for service eventually became less then prevailing service capacity of the system. The concept leads to a simple graphical analysis. Figure 3, based on Costa's data from Lisbon illustrates the approach. As indicated, the horizontal distance between the arrivals and service defines the queue length, and the vertical distance approximates the wait time.
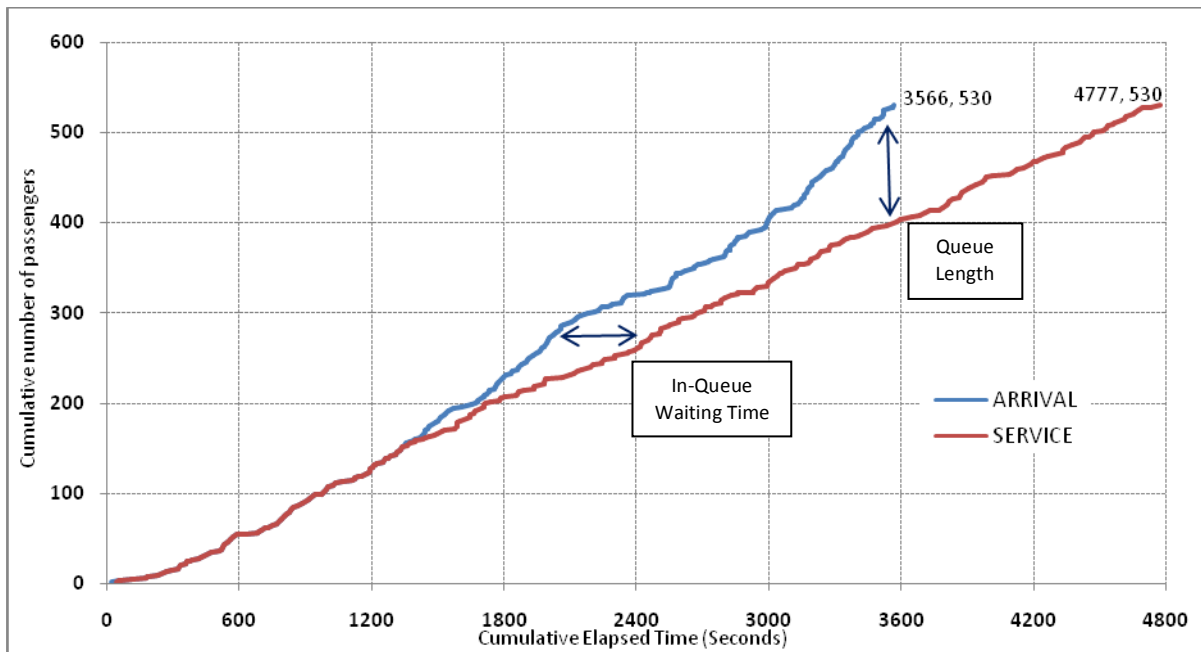

Figure 3 – Arrival and Exit Curves based for taxi pick-ups at Lisbon (Source: Costa, 2009)

Newell's approach is quick and convenient. It has the great advantage of providing a means to estimate the effect of changing service rates, for example by opening or closing check-in counters

or security inspection lanes. Its drawback is that it only works with a specific scenario. It thus does not provide guidance on distribution of wait times or the unreliability of the system. This approach thus does not help managers understand the "practical capacity" of a particular design of a service system, or give information about the frequency of unsatisfactory service that may define the overall acceptability of a design.

Overall, Newell's approach is perhaps most useful in providing insight into how managers of operations can affect average service by scheduling the shifts of personnel to provide service capacity at the right time and place. As airport operators and managers generally have little control over the availability of taxi service, Newell's approach does not seem applicable in that regard to the design of taxi pick-up operations.

Pertinent References in Simulation

Simulation is now widely invoked as a means to study and thus to improve airport operations of all kinds, both landside and airside. Analysts commonly use it to study various actual or potential bottlenecks in airport terminals, both in advance of construction and afterwards, once problems arise. Common applications include check-in processes and security inspection stations. Mumayiz and Schonfeld (1999) provide a collection of applications, which include those of Bender and Chang (1999) and Hathaway (1999) for roadway operations.

Generic computer-based simulations are easy to code. This means that there are many versions available for potential consumers. Generic simulations are however generally of much use in analyzing operations. This is because the output of plain vanilla simulations can be difficult to understand or appreciate. For this reason, providers of simulation services generally spend considerable time to prepare attractive visual images, typically in some form of video, to illustrate how their simulation works. These graphical capabilities differentiate the simulations in the marketplace. They are thus proprietary and unavailable for public discussion. Good graphics help sell simulations to customers. However, they do not themselves make the simulation analytically valuable.

The fundamental difficulty in applying simulations to people lies in describing the flows correctly. It is important to understand how these actually occur. They are typically quite complicated. People are not like parts in an assembly line that flow in straight paths from entry to exit. People wander around, between processes. Between check-in and the security inspection, they may buy a magazine or say good-bye to loved ones, for example. Understanding the complexity of these flows is a crucial part of creating effective simulations. Simulations prepared by persons who are

not deeply familiar with airport operations can thus easily be misleading, because they are naïve about how flows actually occur.

Complementarily, truly effective simulations depend on having good data about the details of the arrivals. As Lee (1966) stressed long ago, many processes are strongly affected by how customers arrive in groups. Are they single, or with family or friends? Are they frequent flyers that know their way around and can proceed expeditiously, or are they encumbered or proceeding with some difficulty? These factors can significantly affect the overall process. For this reason, aggregate data, as about the number of customers per hour or the distribution of arrival times, are not sufficient to define effective simulations of airport operations. This fact motivates the proposed process for analyzing taxi pick-up operations.

Proposed process

The proposed process for analyzing taxi pick-up operations consists of four steps:

1. Detailed measurement of arrivals (size and number of groups, interarrival times); and of service characteristics (the way loading times and exits from queue depend on configuration of queue and position in it);
2. Creation and calibration of simulation model to the specifics of the operation;
3. Exploration of alternatives, by applying the calibrated model to alternative designs; and
4. Choice of preferred design, which comes down to choosing between alternative distributions, thus involving multiple criteria – not just average wait time, but considering possible extremes.

1. Detailed measurement of arrivals and service characteristics

Passengers arrive at the taxi pick-up points singly and in groups. Further, the groups do not arrive regularly, but spread out over time. An obvious first step is thus to observe the sequence of persons arriving at the taxi stand. In this regard, note that this distribution of groups is almost certain to be different from the distribution of groups in other operations, such as security checkpoints. This is because people self-select for taxi service. Groups for example may preferentially choose taxis, because the cost per person is low relative to that of individual travelers, thus making taxi service a more attractive value proposition for groups. It is thus necessary to measure the arrival times and sizes of groups at the taxi stand. This observation should be straightforward.

Obtaining good data on the service process is more complicated. This is because each position in a queue of taxis picking up passengers will typically have its own service characteristics. These depend both on the distance between its position in the queue and the head of the queue

of passengers, which influences the time it takes to board its passenger(s), and on the its position in the queue, which determines how long it must wait for the taxis ahead of it to leave and thus permit it to complete the pick-up service. Furthermore, the service time depends on details of how the taxis can locate themselves in the queue, which depends on the way the lanes are marked off, the pattern of police or other enforcement, and so on. In general, the process of collecting data on the taxi-pick up process requires two phases. The first is a period of observation to identify the local particularities of the process to know what one should be measuring. That done, analysts can devise an effective plan for measuring the service process.

2. Creation and calibration of simulation model

Any simulation model is an assembly of basic modules, such as those depicting arrivals of customers for service into a queue, servicing processes, and departures from service to either some holding area or another queue. Creation of a suitable simulation model for taxi pick-up operations thus depends on having a good understanding of the nature of the flows at this location. This understanding is obtained through step 1. Once analysts have this, the creation of the simulation model is obvious.

The calibration of the model is tricky. As the measurements on the system include the determining elements of the queuing process (the arrivals and service rates), a correctly assembled simulation should reproduce observed results such as the queue lengths. However, this is often not the case. It may easily happen that the initial measurements in step 1 missed some factor that caused the distortion. Analysts should determine the cause of such discrepancies by observation and subsequent adjustment of the simulation until it leads to reasonably close representations of the observed behavior.

3. Exploration of alternatives

Given a simulation model satisfactorily calibrated to local behaviors of passengers and taxi operators, the analyst is now in position to explore and compare alternative designs of the taxi pick-up operations. This evaluation process should both consider the range of possible inputs and outputs.

It is important to consider the prospective performance of the pick-up system both with different levels of loads and patterns of demand. This is because queuing systems are non-linear, such that queue lengths and waiting times are likely to increase much faster than the level of loads (Figure 2). Furthermore, because taxi pick-up operations are subject to fluctuating loads, they rarely are in steady state: this means that they are sensitive to the precise pattern of arrivals for service. In short, the pick-up operations may exhibit widely different performance in terms of

queue lengths and wait times for the same average hourly or day level of traffic, depending on whether arrivals spread smoothly over the period or clump together at particular time.

Complementarily, the analysis should recognize that different designs produce different distributions of performance – some will be more reliable (have lower variance) than others. These differences matter. Customers are typically bothered both by delays and by their unpredictability. Therefore, the analyst should consider not only average performances of alternative designs, but also the range of performance. Note that the notion of a standard deviation in this context is not especially useful by itself. This is because performance is asymmetric: the good performance is limited to zero delays and queues, whereas bad performance can deviate greatly from the average. It is thus more appropriate to refer to maximum and minimum performances, especially for decision makers who may not appreciate the use of technical terms.

4. Choice of preferred design

In the context of uncertain outcomes, such as those resulting from queuing systems, it is generally best to refer to preferred, rather than best designs (de Neufville and Scholtes, 2011). This is because the measures of good performance and of cost reflect different, incomparable units. While we can agree that both delays and service unreliability are undesirable, we are unlikely to agree on how to compare these units – let alone give them a monetary value. Absent an agreed upon way to weight the different characteristics of any alternative, it is impossible to define an unambiguous best design. It is thus better to refer to preferred designs, those whose combinations of attributes stakeholders prefer.

The analysis must thus also recognize that the designs that the decision-makers prefer may differ substantially from those favored by the analyst. They may be more sensitive to costs, for example, or particularly responsive to unreliable service. This fact implies that the analyst should provide the decision-makers with the data on the range of performance for the several alternatives. The analyst might best do this by using a table as the following application indicates.

Case Application: Lisbon International Airport

Costa (2009) developed and applied the proposed process at the Lisbon International Airport, in cooperation with the airport operators, ANA. The case demonstrates the usefulness of the process and reports on the specific findings.

1. Detailed measurement of arrivals and service characteristics

The main taxi pick-up operation at Lisbon International Airport features two rows of taxis, each having two taxis providing service at any time. Taxis feed into these rows from a single file constantly replenished from taxi pool about 100m away. Customers line up in a snake queue whose head is at the top of the taxi row. Figure 4 shows the details.
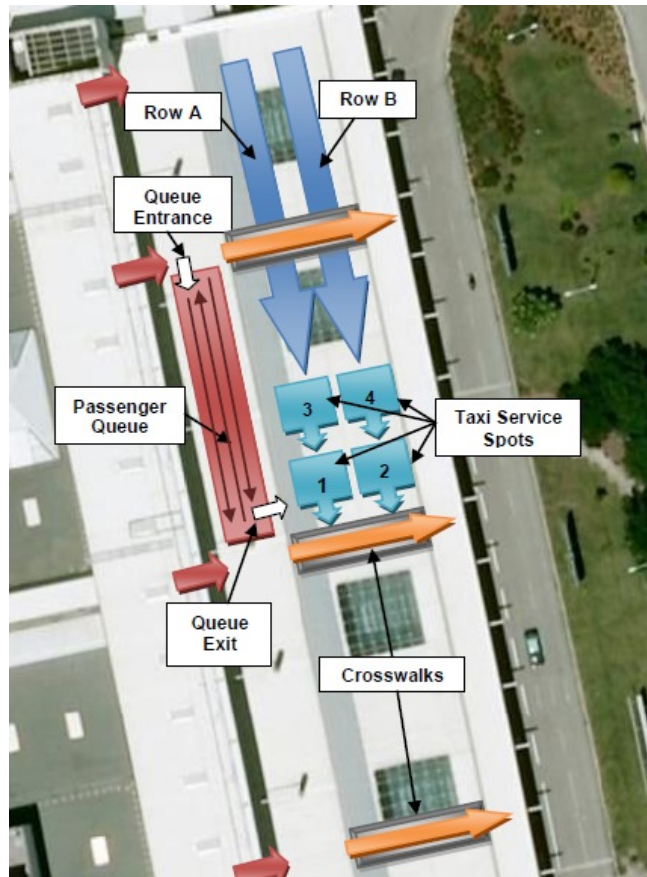
Figure 4. Configuration of the Lisbon Taxi Pick-up Operations at Terminal 1 Arrivals

Following extensive reconnaissance of the operations, Costa (2009) developed the data collection plan shown in Figure 5. It involves two observers. One monitored the performance of the loading process, that is, the service. The other observed the arrivals, noting both their arrival times and the extent to which the customers were part of a group. The observers took detailed measurements on different days.
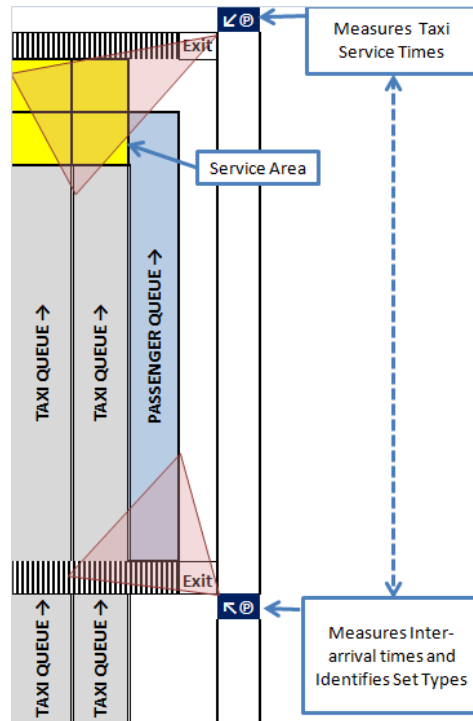
Figure 5 – Data Collection Plan

Figures 6 to 8 provide the data on the arrivals. Figure 6 shows that about 2/3 of the customers arrive reasonably close together, and the rest spread out over time. This is consistent with a Poisson distribution, which might be expected. Figures 7 and 8 indicate that about 3/4 of the groups consist either of single customers or a couple. Note that while larger groups are relatively infrequent in this case, together they account for about 40% of the total customers.
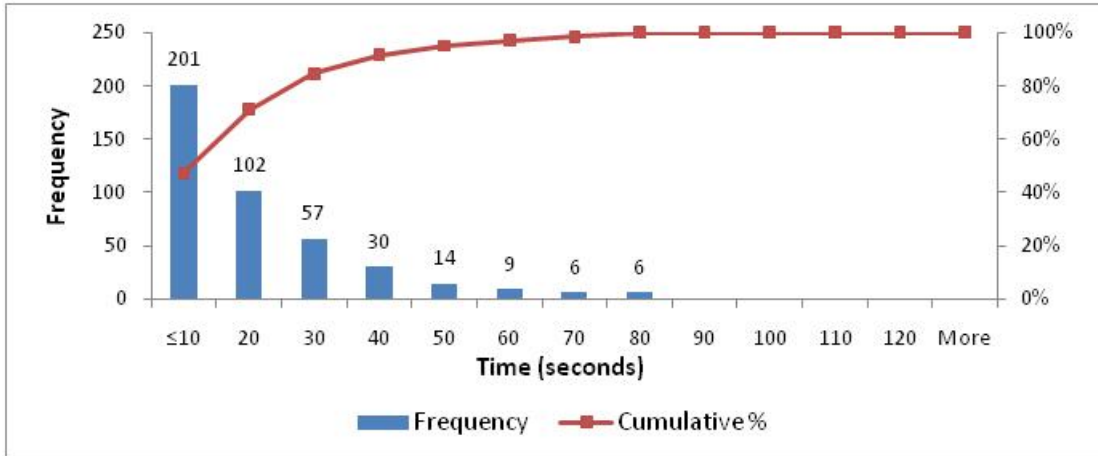
Figure 6. Histogram for Inter-Arrival Times for Groups
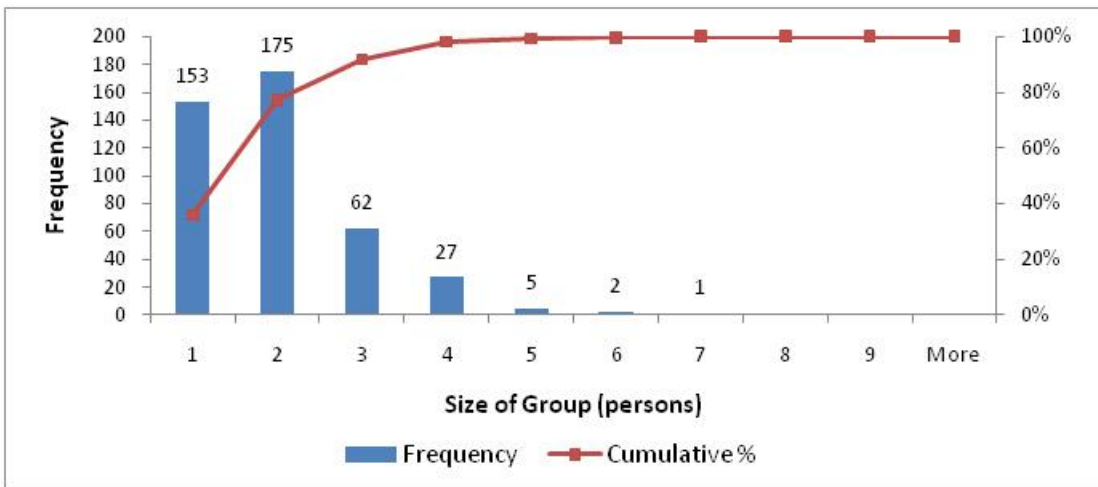
Figure 7. Histogram for Group Size
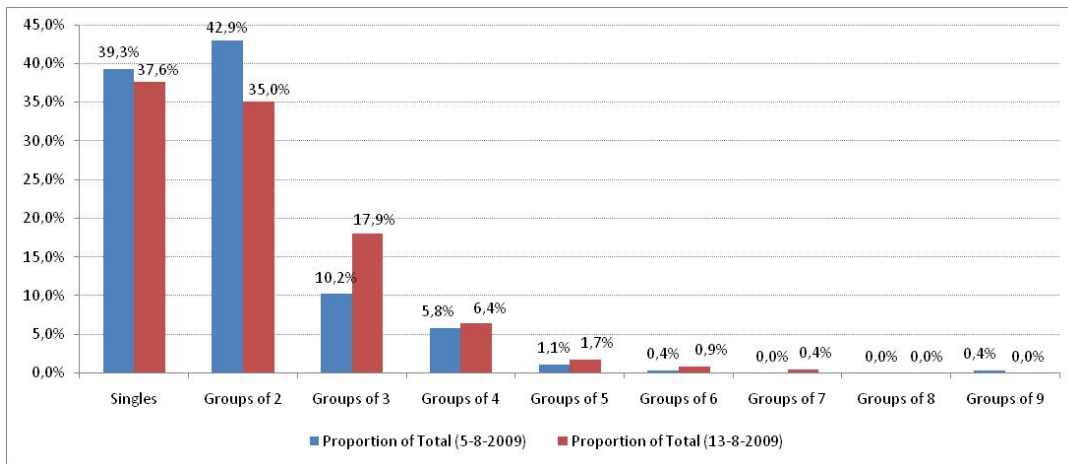
Figure 8. Comparison between Group Size Proportions from the different measurements

Figures 9 and 10 indicate the service times for the taxi pick-up operations. The four histograms in Figure 9 give the service times associated with each of the four positions indicated in Figure 4. Servers 1 and 2 are at the front of the two rows, and servers 1 and 3 are closest to the curb and thus to the customers waiting in the queue. Figure 10 provides the averages and standard deviations on these data. As expected, the taxis second in line and further out take longer times – both because their passenger take longer to get to these taxis, especially if they have to cross a lane of taxis, and because the taxis cannot complete their service and move out until the taxi in front of them has also completed its service. Figure 11 illustrates the conflicts between taxis and pedestrians that lead to these degradations in performance.
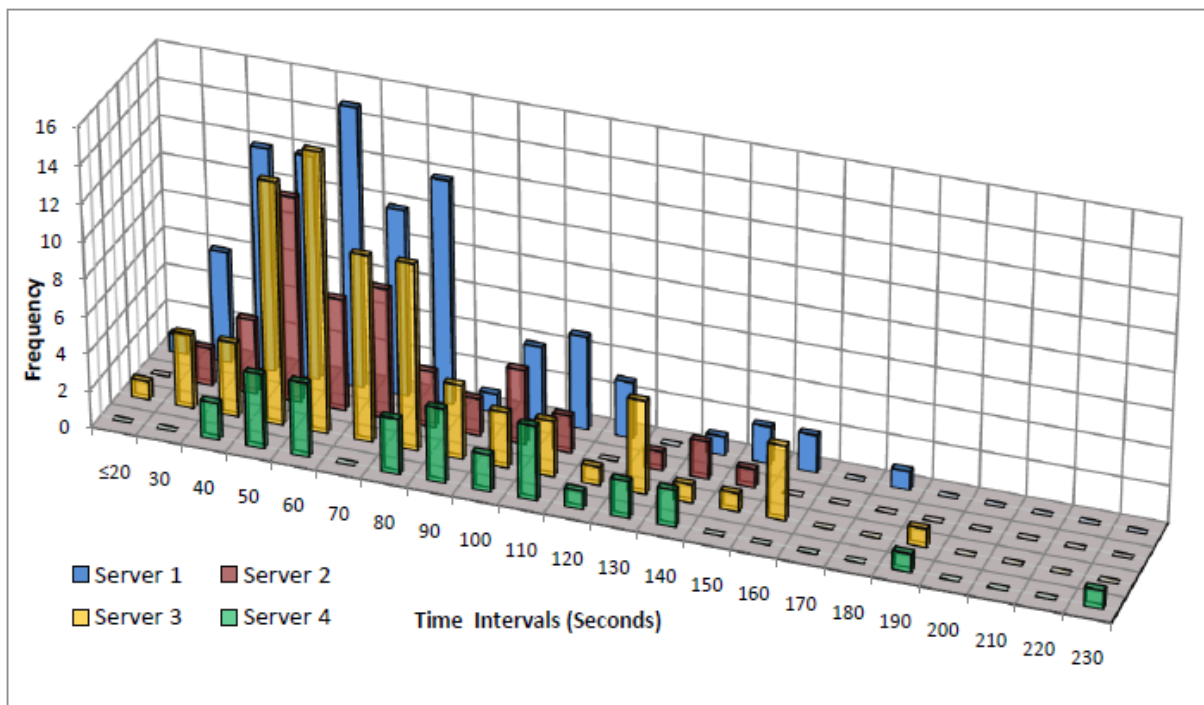
Figure 9. Histograms for Service Times

	Unit	Row A		Row B	
Server		Server 1	Server 3	Server 2	Server 4
Average	(seconds)	66,8	73,5	67,6	90,4
Standard Deviation		33,2	35,5	30,0	43,2
Global Row average		70,0		76,7	
Global Row Standard Deviation		34,4		37,3	

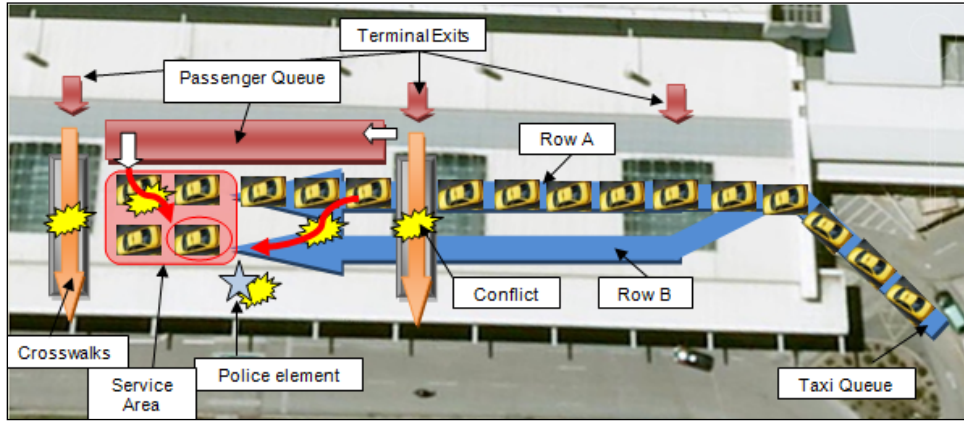Figure 10. Service Time averages and standard deviations

Figure 11. Main conflicts degrading taxi service under observed conditions.
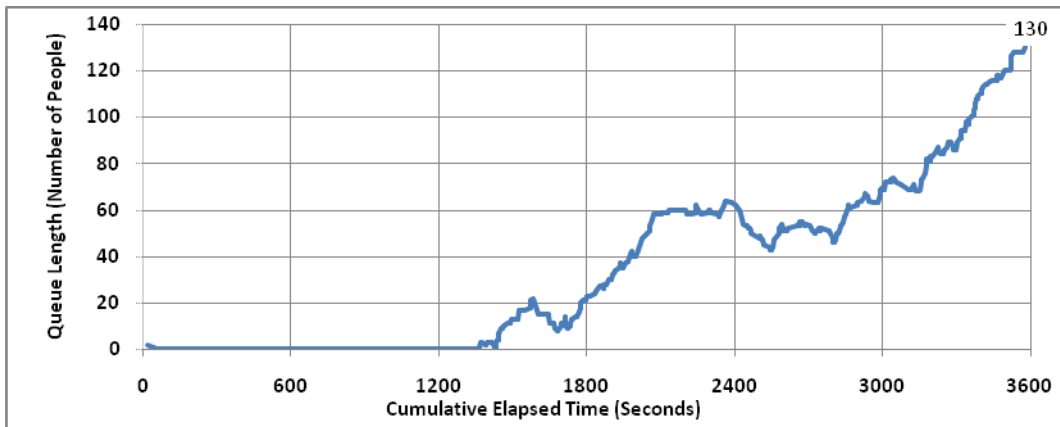
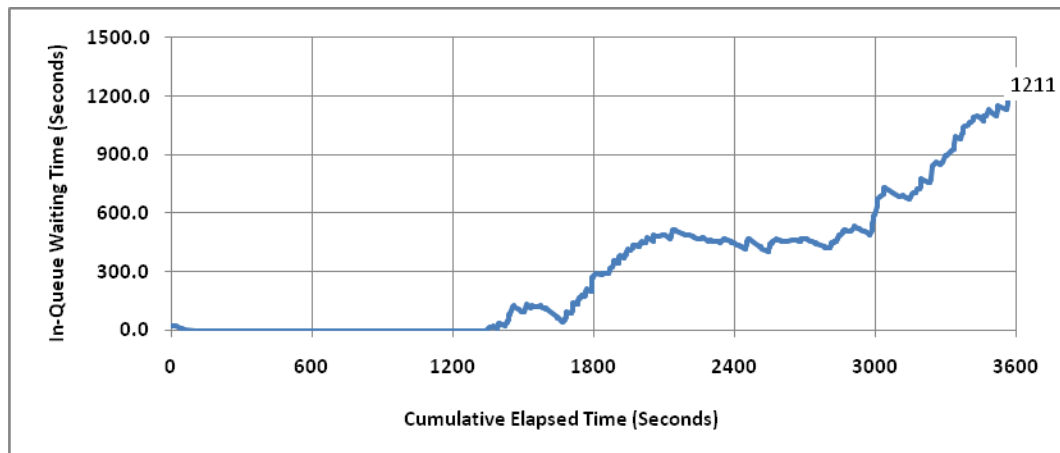

Figure 12. Queue Length evolution

Figure 13. In-Queue Waiting Time evolution

Figures 12 and 13 describe the overall performance of the taxi pick-up operations observed at Lisbon International Airport. Once the delays start during a busy period, they increase rapidly, leading both to long queues (as many as 130 in a queue were observed) and to long wait times (more than 20 minutes). In this connection, it should be stressed that the delays were not due to a lack of available taxis – many were consistently available in the taxi pool about 100m away. The delays can be attributed to the unsatisfactory design of the pick-up operations. This phenomenon stimulated the interest in the issue.

2. Creation and calibration of simulation model

Figure 14 shows the basic configuration of the simulation model for the taxi pick-up operations for the system as observed. Its notable feature is that it includes separate service modules for each of the taxi positions. This contrasts with the naïve alternative of assuming that the service rate is the same for each taxi, and that the overall system can be reasonably considered to be a single queue. Costa implemented the model using the modular capabilities of the Simul8 package. (Shalliker, et al., 2002)
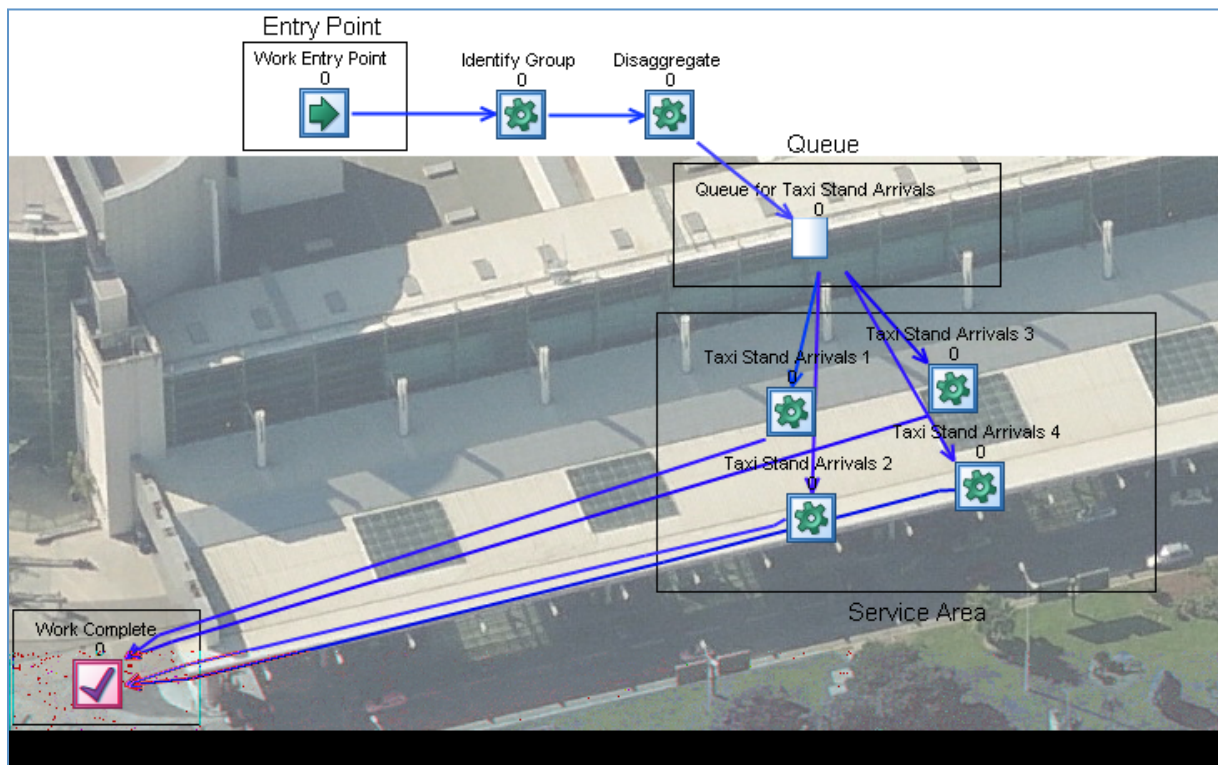

Figure 14. Configuration of Simulation Model for Taxi Pick-up Operations

Figure 15 illustrates the process of calibrating the simulation model. It compares the queue length observed with that predicted by the simulation. In this case, the simulation estimated the maximum length of queue and the pattern of its evolution reasonably accurately, although it did estimate that the queue would build up earlier. A more detailed analysis, beyond the scope of this study, could investigate such phenomena. The calibration is satisfactory for the purposes of this demonstration of the process.
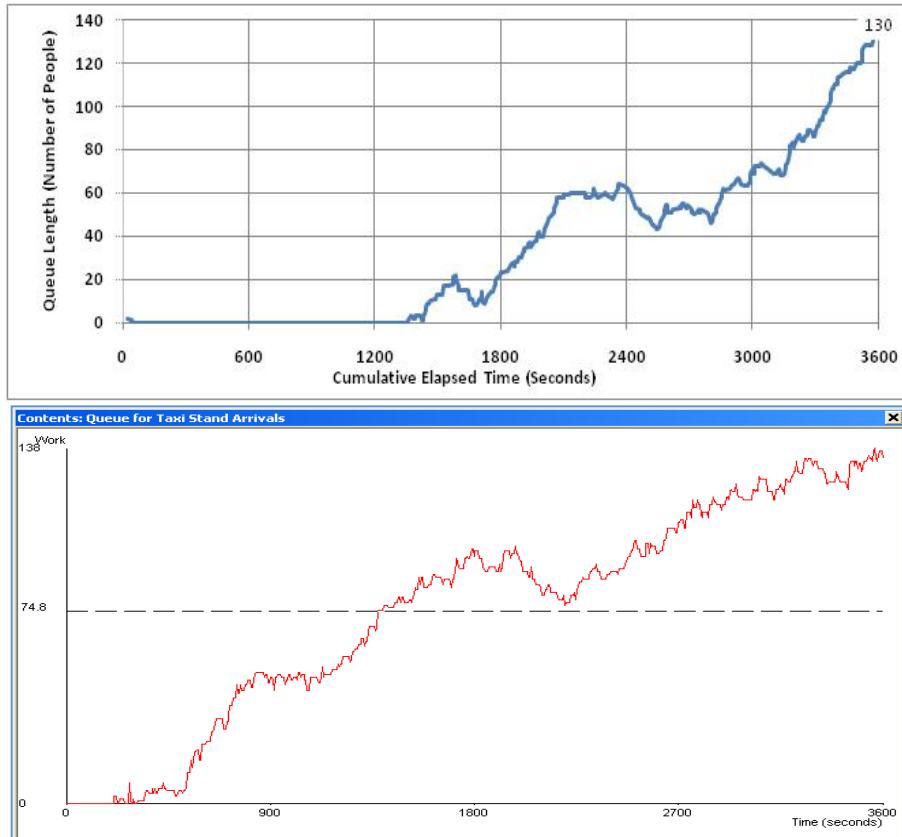

Figure 15. Queue Evolution for the Empirical (above) and Simulation (below) Methods

3. Exploration of alternatives

This application built four alternative scenarios and tested them for performance:

- Scenario I, *Increase the number of service lanes/rows and servers*, consisted in enlarging the inner-curbside area to allow for three lanes of taxis to serve simultaneously at peak hours, instead of the current two. The modeling of this situation consisted in adding two extra servers to the service area, with the assumption that they follow service time distributions that are similar to those of the remaining servers.

- Scenarios IIA and IIB, *Different service area configuration (single lane, multiple servers)*, feature an *inner* service lane, loading several taxis at a time, and a *free* lane, allowing taxis to safely and quickly bypass front colleagues. Scenario II.A had three service positions and Scenario II.B had four.
- Scenario III, *Multiple queues and service segmentation,* explored the possibility of creating a secondary and smaller queue (and dedicated road access) with two service positions near the exit of the inner-curbside road, on the far left end of the Terminal facade, simultaneously keeping the current existing taxi service area functioning.

Figures 16 and 17 show the simulation results for the several tested scenarios, in terms of queue size and queuing time indicators, respectively.
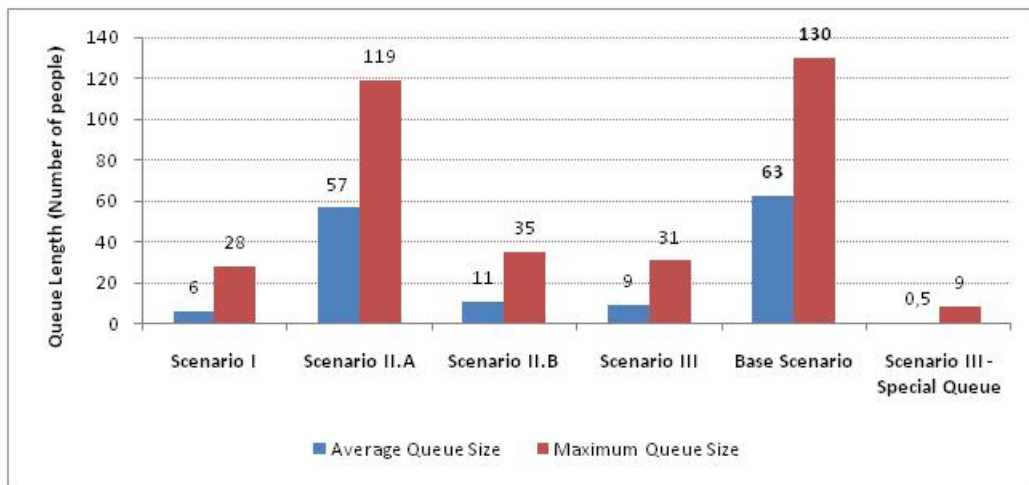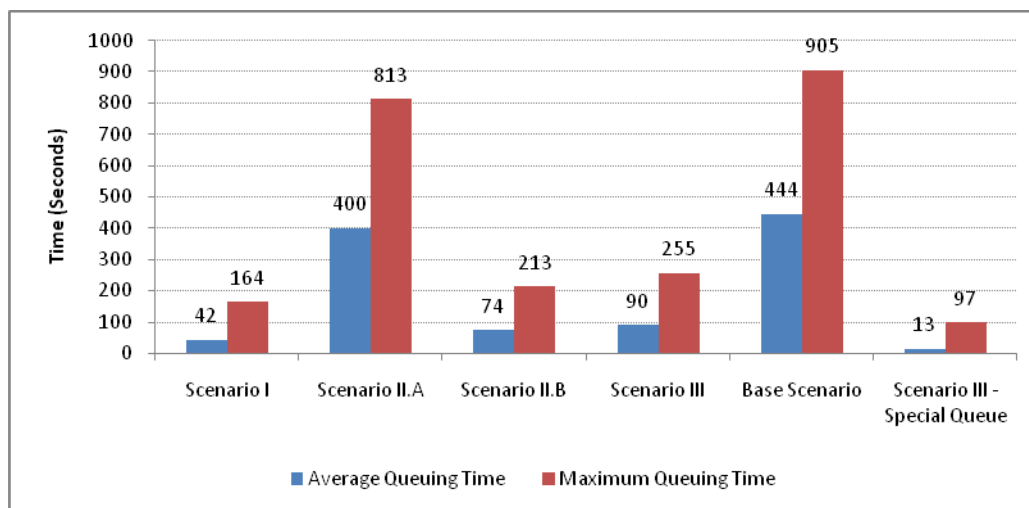

Figure 16. Results for the main Queue Size indicators

Figure 17. Results for the main Queuing Time indicators

4. Choice of preferred design

Any coherent and responsible evaluation of alternatives must take into account not only the main average indicators of service and queuing, but also extreme values as seen in Figure 18. The combination of these two evaluation perspectives allows the analyst to weight and show the peak-hour system behavior better.

Other interesting and important indicators can become important, especially in space or cost-constrained contexts, and decisively influence the final choice of preferred design. In the case of Lisbon, although Scenario I was consistently the best in terms of results, the airport operator perceived it as a less safe, space demanding and expensive alternative. The decision-makers considered Scenario IIB the preferred design, because it implied better levels of safety and affordability at the cost of only slightly less performance.

		Average Queue Size	Maximum Queue Size	Average Queuing Time (s)	St. Dev. Of Queuing Time	Maximum Queuing Time (s)	Average % Server Utilization	Extra Space Consumption	Safety	Implementation Cost
Base Scenario		63	130	444	264	905	**98,6**	🟢	🟡	🟢
Scenario I		**6**	**28**	**42**	**43**	**164**	86,6	🔴	🔴	🔴
Scenario II.A		57	119	400	236	813	**98,6**	🟢	🟢	🟢
Scenario II.B		11	35	74	58	213	92,0	🟢	🟢	🟢
Scenario III	Common Queue	9	31	90	70	255	90,9	🔴	🟡	🔴
	Special Queue	1	9	13	23	97	53,8			

Figure 18. Summary of the results for the main system performance indicators

Further Application: JetBlue Terminal at New York/Kennedy

Based on the successful application in Lisbon, we have been invited to collaborate on a similar application to the new JetBlue Terminal at New York/Kennedy (Kamga et al, 2011). This process appears to be working well. As of the date of submission of this paper (August 1, 2011), these results are unavailable for release. If appropriate, we would hope to be able to include them in the presentation version of this paper.

Conclusion

This paper proposes and demonstrates a practical solution to a long-standing issue in the design of ground transportation services for airports. It provides airport managers with the means to analyze taxi operations quantitatively, which has previously not been available. Based on this approach, they should be able to investigate alternative designs for their taxi pick-up sites, identify preferred solutions, and improve these operations that so often have been deficient.

Acknowledgements

References

Ashford, N., Mumayiz, S., and Wright, P., Airport Engineering: Planning, Design and Development of 21st Century Airports, 4th ed., John Wiley and Sons, New York, NY, 2011.

Ashford, N., Stanton, H. and Moore, C., Airport Operations, 2nd ed., McGraw-Hill, New York, NY, 1996.

Bender, G. and Chang, K., "Roadway and Curbside Traffic Simulation for Las Vegas McCarran International Airport," pp. 28-35 in Mumayiz, S. and Schonfeld, Paul, eds., "Airport Modeling and Simulation," Proceedings of the International Conference on Airport Modeling and Simulation, American Society of Civil Engineers, NY, 1999.

Cardon, N., "The place of the taxi through urban mobility: its practices, positioning & potential for expansion," City on the move, Institut de la Ville en Mouvement, presentation at the Taxi Festival, Lisbon, 2007.

Costa, D., "Performance and Design of Taxi Services at Airport Passenger Terminals," Masters Thesis, Instituto Superior Técnico, Universidade Técnica de Lisboa, Lisbon, 2009.

Curry, G., De Vany, A. and Feldman, R., "A queuing model of airport passenger departures by taxi: Competition with a public transportation mode," Transportation Research, Vol. 12, pp. 115-120, 1977.

Ferreira F., Pacheco A. and Teodoro, F., "Evaluation of taxi services provision on airport terminals curbside for picking up passengers," 74th European Study Group with Industry, University of Aveiro, Aveiro, Portugal, 2010.

de Neufville, R., "The Baggage System at Denver: Prospects and Lessons," Journal of Air Transport Management, Vol.4, no. l, pp. 229-236, December 1994.

de Neufville, R. and Odoni, A., Airport Systems: Planning, Design and Management, McGraw-Hill, New York, NY., 2003.

de Neufville, R. and Scholtes, S., Flexibility in Engineering Design, MIT Press, Cambridge, MA, 2011.

Hathaway, D., "Landside Simulation of Washington National Airport," pp.253-264 in Mumayiz, S. and Schonfeld, Paul, eds., "Airport Modeling and Simulation," Proceedings of the International Conference on Airport Modeling and Simulation, American Society of Civil Engineers, NY, 1999.

Horonjeff, R., The Planning and Design of Airports, McGraw-Hill, New York, NY, 1962.

Horonjeff, R., McKelvey, F. X., Sproule, W., and Young, S.. Planning and Design of Airports, 5th edition, McGraw-Hill, New York, NY, 2010.

IATA International Air Transportation Association, The Airport Development Reference Manual, 9th edition, IATA, Montreal, Canada, 2004.

Kamga, C., Conway, A., Yazici, A., Snghal, A., and Aslam, N., "Strategies to improve taxi service for jetBlue customers at JFK terminal 5," Final Report draft, University Transportation Research Center, Region 2, New York, June, 2011.

Kazda, A. and Caves, R., Airport Design and Operation, 2nd ed., Elsevier, Oxford, 2007.

La Croix, S., Mak, J. and Miklius, W., "Evaluation of alternative arrangements for the provision of airport taxi service," Working Paper 199110, Manoa, University of Hawaii, Department of Economics, 1991.

Lee, A., Applied Queuing Theory, Macmillan, London, UK, 1966.

Mumayiz, S. and Schonfeld, P., eds., "Airport Modeling and Simulation," Proceedings of the International Conference on Airport Modeling and Simulation, American Society of Civil Engineers, NY, 1999.

Newell, G., Applications of Queuing Theory, Chapman and Hall, London, UK, 1982.

Odoni, A., "Lecture notes, Airport Systems Planning and Design, MIT," n.d.

Shalliker J. and Ricketts, C., "An Introduction to Simul8 Release Nine," University of Plymouth, Plymouth, UK, 2002.