



A diffusion approximation to a single airport queue

David J. Lovell^{a,*}, Kleoniki Vlachou^a, Tarek Rabbani^b, Alexandre Bayen^b

^a Department of Civil and Environmental Engineering, Institute for Systems Research, University of Maryland, 1173 Glenn L. Martin Hall, College Park, MD 20742, USA

^b Department of Civil and Environmental Engineering, University of California at Berkeley, 760 Davis Hall, Berkeley, CA 94720, USA

ARTICLE INFO

Article history:

Received 17 February 2011

Received in revised form 1 March 2012

Accepted 11 April 2012

Keywords:

Queuing theory

Diffusion

Delay

Aviation system performance

Kolmogorov forward equation

Fokker–Planck equation

Finite element method

ABSTRACT

This paper illustrates a continuum approximation to queuing problems at a single airport, adapted from the well-known diffusion approximation, as encapsulated in the Kolmogorov forward equation of stochastic processes or the Fokker–Planck equation of physics. The continuum model is derived using special artifacts of the airport problem context, and a numerical solution scheme based on the finite element method is presented. The results are compared against known stationary results from the M/M/1 process, as well as against airport scenarios generated from real demand and supply data. In both cases, a Monte Carlo simulation is used to provide ground truth results against which to compare the diffusion model, and is shown that the results between the Monte Carlo and diffusion models match quite closely.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

Studies of queuing delays in the National Airspace System (NAS), and other large networks, for that matter, are typically conducted either in a Monte Carlo simulation environment, where a considerable amount of fidelity is available at the expense of computational efficiency, or with closed-form equilibrium queuing models fraught with distributional assumptions that are typically not very representative of real situations. A common example of the latter is the use of the Poisson process to represent arrival processes to queues, motivated by its mathematical tractability, even in the face of fairly compelling evidence that the system is not Markovian.

1.1. Existing queuing models

One well-known aviation queuing model is LMINET (Lee et al., 1997, 1998), in which a network of airport queues is represented by means of interconnected single-server queues. Each queue has a time-dependent Poisson arrival process, and an Erlang- k service process. One serious problem with this approach is that because the input process to each downstream node is Poisson, one cannot have independent control of its mean and variance. Thus, while the outputs from upstream nodes may have variances different than what the Poisson process would force, the model cannot interpret these properly. More importantly, any technologies or policies that might be adopted to reduce variance in the system (such as improved trajectory accuracy) cannot be modeled accurately. The goal of this paper is to provide a single-airport building block that might eventually be extended to a network environment, and that would allow for modeling of more complex and dependent interactions between aviation network nodes.

* Corresponding author. Tel.: +1 301 405 7995; fax: +1 301 405 2585.

E-mail address: lovell@umd.edu (D.J. Lovell).

Another single-airport queuing model commonly used in aviation is the DELAYS model developed at MIT, the methodology behind which is captured in Kivestu (1976), Horanjic (1990), and Malone (1995). This model uses a time-dependent Poisson arrival process and an Erlang- k service process, much the same as LMINET (both models have a common heritage). One major difference is that the DELAYS model was later adopted in a network structure that does not suffer from the same independence problems as outlined for LMINET. The Approximate Network Delays (AND) model was originally proposed in Malone (1995), but was not assembled into a working model until more recently (Malone and Odoni, 2001). The idea driving the AND model is that the DELAYS model, by itself, might produce excessively large estimates of delay, when fed purely scheduled arrival times. In reality, the network would not permit such large delays, as demands would be spread over time due to controller actions, metering by upstream queues, etc. Thus, the AND model iterates between the DELAYS model and a delay propagation algorithm, in an effort to find an estimate that more closely matches expectations. This is a heuristic approach, and it still suffers from the drawback this paper is intended to address, which is the strong dependence between arrival process mean and variance.

Another aviation queuing model is the National Airspace System Performance Analysis Capability (NASPAC), which was developed beginning in the 1980s by the Federal Aviation Administration (FAA) and Mitre Corporation. A good description of the original model can be found in Millner (1993). The model is now housed at the FAA, and continues to be developed (see for example Post et al., 2008). The model includes a number of detailed components, such as realistic fleet information, and fuel burn, but its queuing engine is quite rudimentary, consisting of a simple deterministic queue with scalar capacity values for the airports. The claimed path forward to dealing with real stochastic queuing effects is to incorporate Monte Carlo simulation (Post et al., 2008), which will seriously impact the computational complexity of the model, as described above and in the following paragraphs.

With the aviation system in mind, the idea behind this paper is to adapt a common continuous approximation technique known as the diffusion approximation to a queuing problem, with a specific interest in modeling arrival and departure delay statistics at an airport over the course of several hours or a day. The primary advantages of using the diffusion approximation for these purposes are that specific distributional assumptions can be relaxed in favor of an approximate description of the relevant stochastic processes by a small number of their time-dependent moments, that the full spectrum of probabilistic results can be obtained via a single run of the model, and that propagation of higher moments beyond the mean queue behavior can be captured.

In general, we believe it should be possible to represent a network of queues using methodology similar to the methods herein, although the results to date apply only to a single queue with a general arrival and general service process. The presentation of the approach will continue as follows: in Section 2, we show the derivation of the foundational partial differential equation that represents the system dynamics. That is followed by the development of the continuous equations necessary to establish the boundary and initial conditions that assure the meaningful solution of a meaningful problem. In Section 3, we show a numerical approximation scheme that is based on the finite element method (FEM), and that is used to solve the problem by computer. In Section 4 we show some results illustrating the use of the model. The results are compared against known stationary results from the M/M/1 process, as well as against airport scenarios generated from real demand and supply data. In both cases, a Monte Carlo simulation is used to provide ground truth results against which to compare the diffusion model. The paper closes in Section 5 with conclusions and some suggestions for further work.

2. Model development

In this section we introduce the modeling assumptions that lead to the particular continuum approximation for queuing systems known as the diffusion approximation. This consists of a governing differential equation, which is presented first, and which represents the primary dynamics of the system. This equation is valid for a closed subset of the real numbers representing all realistic values of the system state, but some boundary conditions must be imposed to prevent physically meaningless results outside of this interval. We also describe the set of initial conditions required to represent any particular queuing problem for which a solution is sought.

2.1. Governing differential equation

Diffusion methods have been applied to queuing problems in a variety of domains, including road transportation (Newell, 1971), computer networks (Kobayashi, 1974), and more general queuing systems (Gaver, 1968; Kimura, 1983). No significant use of them in an aviation setting is recorded in the literature. The development of the model shown in the following pages borrows very heavily from the exposition of Kimura (1964), which develops the diffusion approximation in the context of a very different application, that of population genetics. The reason for following the template of that paper, however, is that the treatment is very thorough but also accessible to readers without prior experience in diffusion methods, and it can be adapted readily to the aviation context.

Suppose we model the arrival process to an airport as a single-server queue. This is admittedly an abstraction, because there are frequently multiple cornerpost entry points to the airport, often the possibility of multiple arrival runways, and incoming aircraft do not physically line up in queue in the same manner as customers at a grocery store, or even vehicles at a traffic signal. Nevertheless, it is common to model the competition amongst multiple arriving aircraft for the capacitated

resource (the arrival runway system) as a queue, with the interpretation that the delays thereby imparted are assigned and incurred at en route locations farther away from the airport.

Let $Q(t)$ represent the time-dependent random variable describing the length of the (virtual) queue for arrival aircraft at time t . While beyond the scope of this paper, the ultimate goal of this endeavor is to model more complicated aviation networks. In that context, one could use the airport node being described here to model an arrival or departure resource like a runway, a gate, or an esoteric en route node intended to represent a capacity constraint in the airspace itself.

The first assumption necessary for consideration of continuum models is that of continuity; i.e., that the queue length measurement at any given time need not be an integer. Because aircraft only come in discrete units, this is obviously an artificial construct. However, we are mostly interested in using queue length measurements as preliminaries to computing delay statistics, so they will be averaged over a large time domain. This is a stochastic queuing system, and the probability density function for the queue length x at time t is denoted $f(x; t)$. A graphical example of f is shown in Fig. 1. In this notional example, the queue density transitions over the time interval $[0, 10]$, with a mean that increases and then decreases again, and a variance that changes similarly.

We also define the probability density transition function $g(\delta x, x; \delta t, t)$ as the probability density associated with a change in queue length from x to $x + \delta x$ in the time interval $[t, t + \delta t]$. An example of $g(\cdot)$ for a single choice of t and δt is shown in Fig. 2.




Fig. 1. Queue length probability density function $f(x; t)$.




Fig. 2. State transition probability function $g(\delta x, x; \delta t, t)$.

The density function for the queue length at some future time $t + \delta t$ can be expressed using the continuous Kolmogorov–Chapman equation:

$$f(x; t + \delta t) = \int f(x - \delta x; t)g(\delta x, x - \delta x; \delta t, t)d(\delta x) \quad (1)$$

This equation encapsulates conditioning over all of the possible queue states $x - \delta x$ at time t from which a transition to the state x at time $t + \delta t$ is possible. The necessary assumption to use this equation is that the transition probabilities of the state of the queue can be described entirely by the function g , regardless of the history of the prior queue states.

If we use the condensed notation $fg = f(x; t)g(\delta x, x; \delta t, t)$, then we can expand the integrand of (1) as a Taylor series around the point x as follows:

$$f(x - \delta x; t)g(\delta x, x - \delta x; \delta t, t) = fg - \delta x \frac{\partial}{\partial x}(fg) + \frac{(\delta x)^2}{2!} \frac{\partial^2}{\partial x^2}(fg) - \frac{(\delta x)^3}{3!} \frac{\partial^3}{\partial x^3}(fg) + \dots \quad (2)$$

We then substitute (2) back into (1), and interchange integration and differentiation. This presumes, of course that our functions are well-behaved (i.e., bounded) in this sense.

$$f(x, t + \delta t) = f \int g d(\delta x) - \frac{\partial}{\partial x} \left\{ f \int (\delta x) g d(\delta x) \right\} + \frac{1}{2} \frac{\partial^2}{\partial x^2} \left\{ f \int (\delta x)^2 g d(\delta x) \right\} - \dots \quad (3)$$

Since g is a proper density function, then for any choices x , t , and δt , it must be the case that $\int g d(\delta x) = 1$. Hence we simplify the first term on the RHS of (3), and then subtract f from both sides and divide by δt :

$$\frac{f(x, t + \delta t) - f(x, t)}{\delta t} = - \frac{\partial}{\partial x} \left\{ f(x, t) \frac{1}{\delta t} \int (\delta x) g d(\delta x) \right\} + \frac{1}{2} \frac{\partial^2}{\partial x^2} \left\{ f(x, t) \frac{1}{\delta t} \int (\delta x)^2 g d(\delta x) \right\} - \dots \quad (4)$$

The limits of two of the elements contained in the RHS of (4) are frequently called the “infinitesimal” mean and variance, respectively:

$$\lim_{\delta t \rightarrow 0} \frac{1}{\delta t} \int (\delta x) g(\delta x, x, \delta t, t) d(\delta x) \equiv M(x; t) \quad \forall x, t \quad (5)$$

$$\lim_{\delta t \rightarrow 0} \frac{1}{\delta t} \int (\delta x)^2 g(\delta x, x, \delta t, t) d(\delta x) \equiv V(x, t) \quad \forall x, t \quad (6)$$

We now make a second assumption, which is that all of the important information about the transition density function g can be captured adequately in its first and second moments, as in (5) and (6), respectively. This is not a severe limitation; for situations where this is not the case, additional infinitesimal moments can be defined, and the analyst is then responsible for providing that information as well. In fact, in aviation applications, the best contemporary network models, such as LMINET (Lee et al., 1997, 1998) only deal with the propagation of average behavior, and usually with independent Poisson processes at each downstream node. Thus, including $V(x; t)$ is already a step forward. For the present case, assuming that the first two moments suffice, this is tantamount to the assumption:

$$\lim_{\delta t \rightarrow 0} \frac{1}{\delta t} \int (\delta x)^n g(\delta x, x; \delta t, t) d(\delta x) = 0 \quad n \geq 3 \quad \forall x, t \quad (7)$$

Then, taking the limit of (4) as $\delta t \rightarrow 0$ and substituting (5) and (6) yields:

$$\frac{\partial f(x; t)}{\partial t} = \frac{1}{2} \frac{\partial^2}{\partial x^2} V(x; t) f(x; t) - \frac{\partial}{\partial x} M(x; t) f(x; t) \quad (8)$$

Eq. (8) is commonly called the *Kolmogorov forward equation* in the stochastic processes literature, or the *Fokker–Planck equation* in the physics literature. In the second case, the term $M(x; t)$ is referred to as *drift*, while the term $V(x; t)$ is called *diffusion*. Eq. (8) is the governing differential equation (GDE) for our queuing system.

2.2. Boundary conditions

In this section, we develop the boundary conditions that prevent the model from generating non-zero probabilities for states that are not physically possible, including negative values of the queue length. A similar constraint can be imposed to prevent the possibility of what might be considered unnaturally large queue lengths. The upper bound is more difficult to specify precisely, but it is necessary from a pragmatic standpoint in the numerical scheme because the solution space must be bounded, as will be seen in Section 3. Using schedule data for a particular day, one could establish an absolute upper bound on the queue length by presuming that all of the aircraft had arrived, and none had been served, which is obviously an extreme case.

Because the random variable $Q(t)$ represents a queue length, it makes no sense for it to be negative. Thus, we want to apply an auxiliary condition that can guarantee that

$$f(x; t) = 0 \quad x < 0, \quad \forall t \quad (9)$$

This cannot be accomplished by simply saying that (9) must be true; an additional differential equation must be specified that follows the same temporal evolution as (8), and whose effect is to guarantee that (9) holds. Assuming that the initial conditions obey (9) (as they should), a way to do this is to guarantee that the “net probability flux” (what would be thought of as the mass flux if this were a problem in physics) across the point $x = 0$ is always zero.

We fix a point x in one dimension and consider the probability flux across that point in both directions. By integrating all possible increasing transitions that cross this barrier, and subsequently all possible decreasing transitions that cross the same barrier, and then adding them together, we arrive at the following requirement that the net probability flux be zero. This constraint is referred to in the physics or stochastic processes literature as a *reflecting barrier*.

$$f(0; t)M(0; t) - \frac{1}{2} \frac{\partial}{\partial x} f(x, t)V(x; t) \Big|_{x=0} = 0, \quad t > 0 \quad (10)$$

At all times, f must also be a proper density function:

$$f(x; t) \geq 0 \quad \forall x, t \quad (11)$$

$$\int f(x; t) dx = 1 \quad \forall t \quad (12)$$

These last two conditions are notoriously difficult to enforce in a numerical solution scheme (Kumar et al., 2006). We discuss this further in Section 3.

2.3. Initial conditions

The functions $M(x; t)$ and $V(x; t)$ represent the first and second moments, respectively, of the rate at which the length of the queue is changing at time t , given that its current state is x . In a queuing system where the arrival process is independent of the service process, then with the possible exception of $x = 0$ and an upper reflecting barrier, there is no reason to suspect that these functions should change across x . In such situations, it is only necessary to specify how these functions change over time. For most aviation applications, for example, one would expect $M(x; t)$ to be positive at the beginning of the day, negative at the end of the day, and perhaps with some additional cycles in between. One would expect $V(x; t)$ to be small (approaching zero) at the beginning and end of the day and something larger in between, and of course never negative. If this construction were extended to a queuing network, these functions could be derived entirely from the outputs $\{f_i(x; t)\}$ of upstream queues i , with some time lags and with some rules for mixing them together.

Although we explicitly prevent negative queues, it also makes sense to preclude initial conditions that would seem in conflict with this goal. Thus, we require that

$$M(0; t) \geq 0 \quad \forall t \quad (13)$$

At any node to which this method is applied, one can imagine that $M(x; t)$ will be computed as the differential of the difference between the arrival rate, which we might get from the outputs of upstream processes, and the departure rate, which is related to the capacity of the airport or other resource. This being the case, (13) simply prevents an airport from serving traffic that does not exist.

At some airports, however, the rate of queue growth might depend on its current state. For example, if the total capacity of the airport is divided between arrivals and departures, and the airport has some control over that split, then in cases when there is an excess of arrivals, the airport might choose to emphasize arrivals over departures to ameliorate this queue. This is tantamount to a temporary increase in the arrival capacity of the airport. If this were repeatable and quantifiable behavior, that could be captured in differences in $M(x; t)$ across different values of x .

We must specify an initial queue length distribution. For real airport problems, the queue is empty at the beginning of the day, so one might require:

$$f(x; 0) = \delta(x) \quad (14)$$

where $\delta(\cdot)$ is the Dirac delta function. Alternatively, one might consider analyzing a problem starting at some other point in the middle of the day, in which case the restriction (13) is not required.

3. Numeric scheme

In order to solve a system including partial differential equations and their associated boundary and initial conditions, a numerical scheme is necessary to convert that continuum problem into some discrete form appropriate for solution by computer (Pepper and Heinrich, 1992). In this paper we present a discretization method based on the well-known finite element

method (FEM) that is appropriate for our problem. Certainly a host of other schemes could be attempted, including finite difference methods.

The FEM scheme developed for this problem consists of transforming the governing differential equation with its boundary and initial conditions into linear algebraic equations that can be solved at every time step. This transformation is possible by constructing a discrete approximation to the queue length density function $f(x; t)$ using the N Lagrange basis functions ϕ_1, \dots, ϕ_N . Each basis function has a triangular shape; the collection of them is illustrated in Fig. 3 for $N = 4$. Mathematically, the basis functions can be represented as follows:

$$\phi_1(x) = \left(\frac{l_2 - x}{l_2 - l_1}\right)_+$$

$$\phi_j(x) = \left(\min\left\{\frac{x - l_{j-1}}{l_j - l_{j-1}}, \frac{l_{j+1} - x}{l_{j+1} - l_j}\right\}\right)_+, \quad j = 2, \dots, N - 1$$

$$\phi_N(x) = \left(\frac{x - l_{N-1}}{l_N - l_{N-1}}\right)_+$$

The approximation for f can then be expressed using these basis functions as:

$$f^{L,N}(x; t) = \sum_{j=1}^N a_j(L\Delta t) \phi_j(x) \tag{15}$$

where L is the number of time steps Δt , N is the number of Lagrange basis functions, and $\{a_j\}$ are the parameters of the approximation. Using the finite element method, the “solution” of the problem essentially amounts to determining the values $\{a_j\}$.

The left hand side of the PDE (8) can now be approximated by:

$$\frac{\partial f(x, t)}{\partial t} \approx \frac{f^{L+1} - f^L}{\Delta t}$$

and the dynamics can be re-written as:

$$\frac{f^{L+1} - f^L}{\Delta t} = \frac{1}{2} \frac{d^2}{dx^2} (V^{L+1} f^{L+1}) - \frac{d}{dx} (M^{L+1} f^{L+1}) \tag{16}$$

We enforce (16) by defining the residue r , which is essentially the difference between the LHS and RHS of (16),

$$r = \frac{1}{2} \frac{d^2}{dx^2} (V^{L+1} f^{L+1}) - \frac{d}{dx} (M^{L+1} f^{L+1}) - \frac{f^{L+1} - f^L}{\Delta t}$$

We force that residue to zero by using a test function $w(x)$. We equate all of the projections of the residue on w to be zero; i.e., $\int_{\Omega} r w dx = 0$, where Ω is the domain of interest in x and $\partial\Omega$ its boundary. Integrating by parts yields:

$$\frac{1}{2} \int_{\Omega} \frac{d}{dx} (V^{L+1} f^{L+1}) \frac{dw}{dx} dx - \int_{\Omega} M^{L+1} f^{L+1} \frac{dw}{dx} dx + \int_{\Omega} \frac{f^{L+1}}{\Delta t} w dx = \int_{\Omega} \frac{f^L}{\Delta t} w dx + \left[\frac{1}{2} \frac{d}{dx} (V^{L+1} f^{L+1}) - M^{L+1} f^{L+1} \right] w \Big|_{\partial\Omega} \tag{17}$$

where the last term on the RHS depends on the boundary conditions.

We assume that the interval is closed, and that at the right boundary $x = 1$, we would like the net probability flux to be 0. For some large l , the probability density function will approach 0 for all $x > l$. This will make the net probability flux approach zero at $x = l$, although it cannot be absolutely guaranteed. This is discussed more in the conclusions. Together with Eq. (10), we conclude:

$$\left[\frac{1}{2} \frac{d}{dx} (V^{L+1} f^{L+1}) - M^{L+1} f^{L+1} \right] w \Big|_{\partial\Omega} = 0$$




Fig. 3. Lagrange basis functions for the finite element method.




Fig. 4. M/M/1 queue with 1000 replications, $\rho = \mu/\lambda = 0.375$.

We parameterize the test function w with the Lagrange basis functions $\{\phi_i\}$ and parameters $\{b_i\}$:

$$w(x) = \sum_{i=1}^N b_i \phi_i(x) \tag{18}$$

We use the Lagrange approximations of w and f to obtain:

$$\sum_{i=1}^N b_i \left[\sum_{j=1}^N a_j^{L+1} K_{ij} - R_i \right] = 0 \quad \forall \{b_i\} \tag{19}$$

where

$$K_{ij} = \frac{1}{2} \int_{\Omega} V^{L+1} \phi_j' \phi_i' dx - \int_{\Omega} M^{L+1} \phi_j \phi_i' dx + \frac{1}{\Delta t} \int_{\Omega} \phi_j \phi_i dx$$

$$R_i = \frac{1}{\Delta t} \int_{\Omega} \phi_i \left(\sum_{j=1}^N a_j^L \phi_j \right) dx$$

In the last two equations, we denote $a_j^L = a_j(L\Delta t)$ and suppress the dependence of the basis functions $\{\phi_i\}$ on x for the sake of clarity. As mentioned before, we have also assumed that the function $V(x; t)$ is constant in x .

Since the set $\{b_i\}$ is arbitrary, (19) is equivalent to solving the linear algebraic equations:

$$\sum_{j=1}^N a_j^{L+1} K_{ij} = R_i \quad \text{for } i = 1, 2, \dots, N \tag{20}$$

The solution of (20) is the set of parameters $\{a_j\}$ which define $f(x; t)$ according to (15). One of the advantages of the finite element method is the ability to solve these algebraic equations element by element. The N Lagrange basis function approximation defines $N - 1$ elements, which makes it possible to solve $N - 1$ independent algebraic equations.

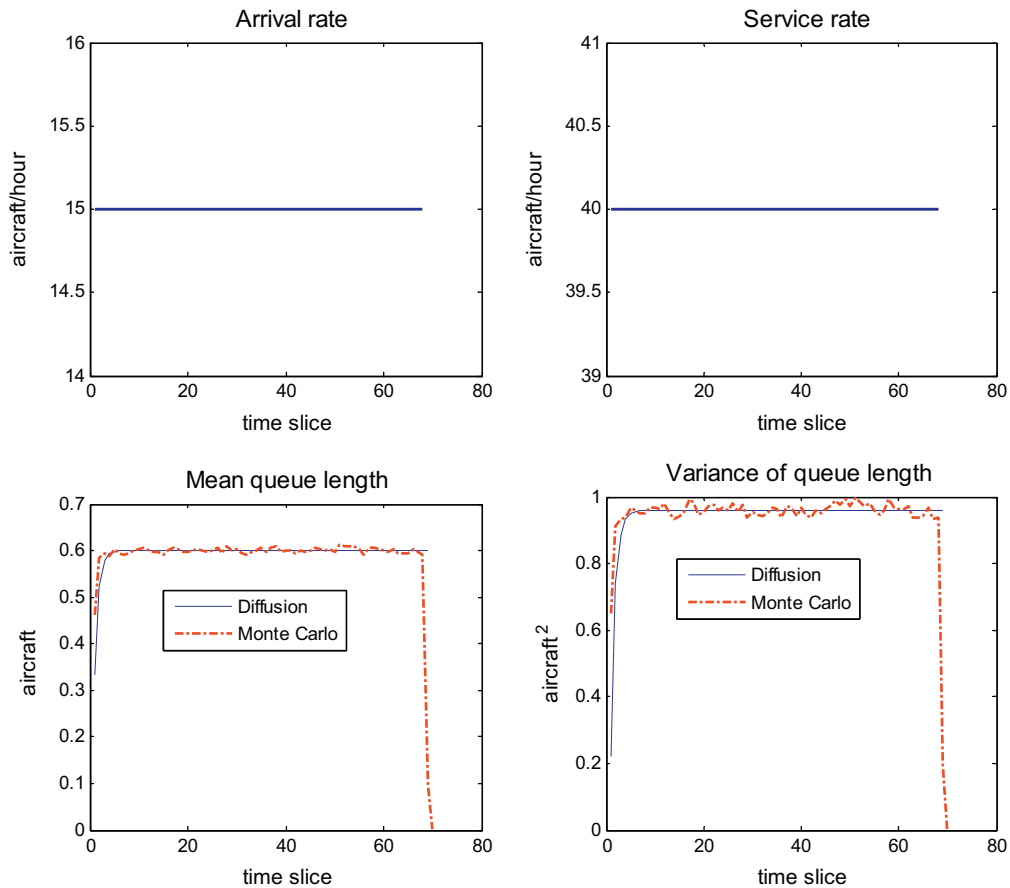


Fig. 5. M/M/1 queue with 10,000 replications, $\rho = \mu/\lambda = 0.375$.

The two remaining boundary conditions to enforce on the solution are (11) and (12). As described in Kumar et al. (2006), we can enforce (12) by scaling the solution appropriately. The non-negativity constraint is harder to enforce. One possible solution is the partition of unity finite element method (PUFEM), described in Kumar et al. (2006). For the time being, however, the problems that we have solved always result in positively valued density functions, and they solve very quickly, so a more complex solution method is not justified unless that situation changes.

4. Model validation and results

In this section, we show some results of applying the modeling with different input data sets. In order to validate our model, we used Monte Carlo simulation as ground truth. We run the simulation a number of iterations (1000 and 10,000) and averaged over this number in order to get the mean and the variance of the queue length. The first set of experiments involve comparisons against the steady state M/M/1 queue. This is not a very useful system for modeling airports, but because the stationary moments are known, we can demonstrate that the diffusion model converges to the proper equilibrium solution.

Fig. 4 shows how the results from the diffusion model compare to the results from the Monte Carlo model. The latter results are for an M/M/1 queue with arrival rate $\mu = 15$ aircraft/h, and a service rate of $\lambda = 40$ aircraft/h. We are assuming here that the airport is only handling arrivals, and therefore the service rate is a proxy for the runway capacity in that condition. The traffic intensity is thus $\rho = \mu/\lambda = 0.375$. The equilibrium queue length is then given by:

$$\bar{Q} = \frac{\rho}{1 - \rho} = 0.6$$

and the equilibrium variance by:

$$\text{Var}(Q) = \frac{\rho}{(1 - \rho)^2} = 0.96$$

Both the Monte Carlo and the diffusion results obviously converge to these values, although the diffusion model does so much more smoothly. That is because in this figure, only 1000 replications of the Monte Carlo simulation were conducted, hence a certain amount of noise around the equilibrium values. Fig. 5 shows similar results for Monte Carlo runs with 10,000 replications instead.

Observe that as the number of replications for the Monte Carlo simulation increases, it follows much better the diffusion solution and the equilibrium solution. One important advantage of the diffusion model is the solution time. The Monte Carlo simulation required 10.86 s for 1000 runs and 106.9 s for 10,000 runs. The diffusion model completes in one iteration, which takes about 8.2 s.

In Fig. 6, we show results for some realistic airport demand and supply profiles. In this case, the demand profile is from the published (OAG) schedule for Miami International Airport, from a peak day in 2007. The capacity profile is a single cluster from a k -means cluster analysis on airport arrival rates (AARs), generated using the methodology shown in Liu et al. (2008). These demand and capacity data have been used for previous studies on queuing, see for example Hansen et al. (2009). The arrival data show considerable fluctuation over the course of the day, while the capacity profile is nearly flat. We modeled the arrival process as a non-stationary Poisson process, and the service process as a non-stationary Erlang- k process, with $k = 7$. This was done for two reasons, first to show that the diffusion model produces good results with different distribution assumptions, and second because this has been shown to be a reasonable model for a single airport server process in other literature (see for example Malone and Odoni, 2001). The Monte Carlo results include 1000 replications.

From observation of the figure, one can tell that the diffusion model replicates the Monte Carlo ground truth quite well, in both the first and second moments. This is a very uncongested day, so the mean queue length remains quite low over the entire day.

The final set of results comes from a peak day at Chicago O'Hare International Airport. These results are shown in Fig. 7. The demand profile is very oscillatory, and it frequently surpasses the capacity over the first three quarters of the day. Thus, larger mean queue lengths are to be expected. The demand subsides towards the end of the day. The capacity profile appears to oscillate smoothly between 22 and 23 aircraft/h. This is just an artifact of specifying integer 15-min service rates derived

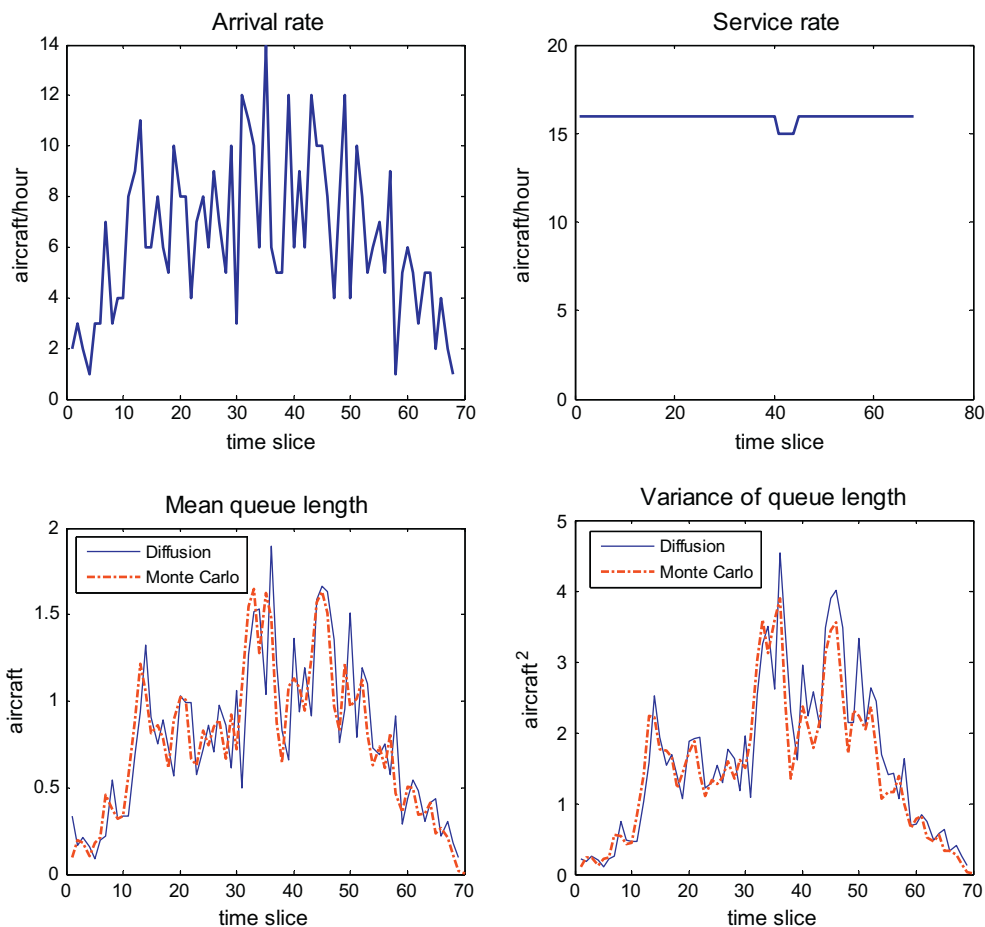


Fig. 6. Diffusion and Monte Carlo queuing results for Miami International Airport.

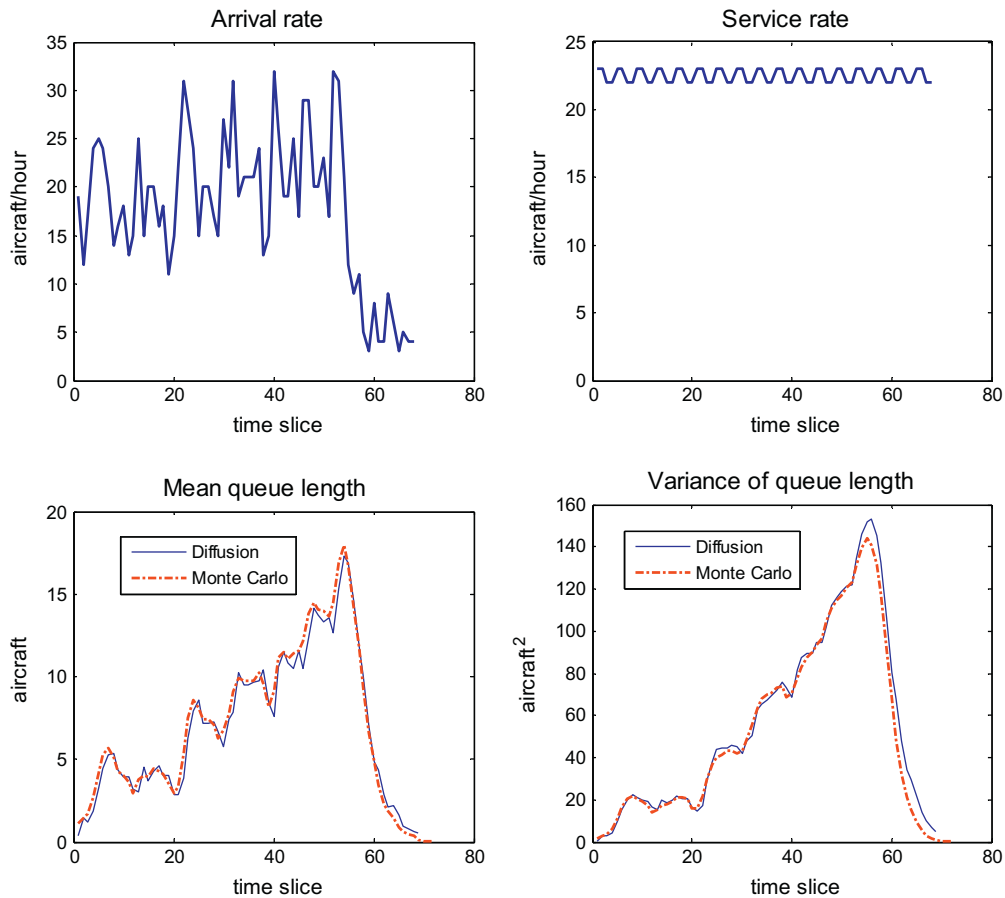


Fig. 7. Diffusion and Monte Carlo queuing results for Chicago O'Hare International Airport.

from hourly rates that happen not to be multiples of 4. Again, the profiles of the first and second moments of the queue length over time match quite closely between the diffusion and the Monte Carlo models.

For both of these last two sets of results, the Monte Carlo runs complete in about 31 s, and the diffusion runs in about 8 s. The time required for the Monte Carlo runs is directly proportional to the number of replications, so if more precision were required, for example 10,000 runs, then the run time would be closer to 310 s. The diffusion model is immune to these considerations.

5. Conclusions

The paper has presented the mathematical construction of a continuum approximation to a queuing system that might represent a single congested resource in the National Airspace System, such as an airport, a runway, or some en route resource. The result is derived from the diffusion approximation. A numeric solution scheme based on the finite element method is also shown.

The use of this type of approximation requires one to be comfortable with some of the assumptions made in the paper, such as the willingness to consider non-integer queue lengths. That notwithstanding, the method has seen considerable application in other areas of queuing theory that also deal with countable objects, so this assumption is not unique to the aviation context. The power of the approach is its ability to deal with a wide range of possible distributions, without having to specify them to any more detail than their first few moments.

This result is a stepping stone in what will hopefully be a larger system of inquiry into the use of such continuum approximations to study systems of aviation queues. In particular, the ability to model the propagation of both the mean and the variance of delay statistics through a connected network would mark a major leap forward in the performance analysis of the aviation system.

In this study we conducted a validation effort of our model. We were able to replicate the known steady-state results from that small set of queuing systems for which equilibrium results are known in closed form. The results in such cases showed that the diffusion approximation gives the same results, but very quickly. Furthermore, a Monte Carlo exercise

was also conducted for a number of other cases whose solutions cannot be found analytically. Again, the diffusion model seemed to perform very well, and it is much faster than running large numbers of Monte Carlo simulations.

References

- Gaver Jr., D.P., 1968. Diffusion approximations and models for certain congestion problems. *Journal of Applied Probability* 5 (3), 607–623.
- Hansen, M., Nikolieris, T., Lovell, D., Vlachou, K., Odoni, A., 2009. Use of queuing models to estimate delay savings from 4D trajectory precision. In: *Proceedings of the 8th USA/Europe Air Traffic Management Research and Development Seminar*, Napa Valley, CA.
- Horanjic, B.R., 1990. Some Queuing Models of Airport Delays. S.M. Thesis. Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA.
- Kimura, M., 1964. Diffusion models in population genetics. *Journal of Applied Probability* 1 (2), 177–232.
- Kimura, T., 1983. Diffusion approximation for an M/G/m queue. *Operations Research* 31 (2), 304–321.
- Kivestu, P., 1976. Alternative Methods of Investigating the Time-Dependent M/G/k Queue. S.M. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Kobayashi, H., 1974. Application of the diffusion approximation to queuing networks II: Nonequilibrium distributions and applications to computer modeling. *Journal of the ACM* 21 (3), 459–469.
- Kumar, M., Singla, P., Chakravorty, S., Junkins, J.L., 2006. The partition of unity finite element approach to the stationary Fokker–Planck equation. In: *Proceedings of the AIAA/AAS Astrodynamics Specialist Conference*, Keystone, CO.
- Lee, D.A., Kostiuk, P.F., Kaplan, B.J., Escobar, M., Odoni, A., Malone, B., 1997. Technical and Economic Analysis of Air Transportation Management Issues Related to Free Flight. Logistics Management Institute, NS501T1.
- Lee, D.A., Long, D., Etheridge, M.R., Plugge, J.R., Johnson, J.P., Kostiuk, P.F., 1998. A Method for Making Cross-Comparable Estimates of the Benefits of Decision Support Technologies for Air Traffic Management. Logistics Management Institute, NS710S1.
- Liu, P.-c.B., Hansen, M., Mukherjee, A., 2008. Scenario-based air traffic flow management: from theory to practice. *Transportation Research Part B* 42, 685–702.
- Malone, K., 1995. Dynamic Queuing Systems: Behavior and Approximations for Individual Queues and Networks. Ph.D. Thesis, Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA.
- Malone, K., Odoni, A., 2001. The Approximate Network Delays Model, Working Paper. Operations Research Center, MIT.
- Millner, D.C., 1993. Design of the NASPAC Simulation Modeling System. The MITRE Corporation, MTR 92W0000135, FAA, Contract No. DTFEA01-93-C-00001.
- Newell, G.F., 1971. *Applications of Queueing Theory*. Chapman-Hall.
- Pepper, D.W., Heinrich, J.C., 1992. *The Finite Element Method: Basic Concepts and Applications*. Taylor & Francis.
- Post, J., Gulding, J., Noonan, K., Murphy, D., Bonn, J., Graham, M., 2008. The modernized National Airspace System Performance Analysis Capability (NASPAC). In: *Proceedings of the 26th Congress of the International Council of the Aeronautical Sciences (ICAS)*, Anchorage, AK.