Reservation Probabilities *

E. G. Coffman Jr.
Bell Labs, Lucent Technologies
Murray Hill, NJ 07974
egc@bell-labs.com

Predrag Jelenković†
Dept. of Electrical Engineering
Columbia University
New York, NY 10027
predrag@ee.columbia.edu

Bjorn Poonen
Department of Mathematics
University of California
Berkeley, CA 94720-3840
poonen@math.berkeley.edu

Abstract

Requests for a given resource arrive in a rate-λ Poisson stream, each specifying a future time interval to be reserved for its use of the resource. The advance notices and the requested intervals are represented by two independent i.i.d. sequences. A request is immediately accepted if it is for an interval that overlaps no currently reserved interval; otherwise, it is lost. We compute the *reservation probability*, which is the fraction of time the resource is booked. From the huge variety of potential applications covered by our model, relatively new applications in existing and proposed communication systems motivate the research here.

We begin by deriving explicit formulas for the reservation probability in several special, but interesting cases. For spread-out advance-notice distributions, our main result then establishes an intriguing connection between the reservation problem and the on-line interval packing problem. For example, if the advance-notice distribution is uniform on $[0, a]$ and all requested intervals have unit length, then large-a estimates of the reservation probability $p(a, \lambda)$ are available from the limit law

$$\lim_{a\to\infty} p(a,\lambda) = \int_0^{\lambda} e^{-2\int_0^v \frac{1-e^{-x}}{x}dx} dv,$$

where the constant on the right arises from connections with the classical car-parking problem. We prove a similar result for more general distributions, a result that extends earlier theorems for the on-line interval packing problem.

*`Advances in Performance Analysis, 1998, to appear`

†The research of this author was done while on staff at Bell Labs

1 Introduction

Requests for a resource arrive in a stationary Poisson stream at rate λ, each request specifying a future time interval, called a *reservation interval*, to be booked for its use of the resource. If the request is indeed approved, the interval is also called a *reserved* or *booked* interval. The *advance notices* (delays before reservation intervals are to begin) are independent and drawn from a distribution $\mathcal{A}$. The durations of reservation intervals are sampled from the distribution $\mathcal{B}$; these durations are independent of each other and the advance notices. A and B denote random variables with the distributions $\mathcal{A}$ and $\mathcal{B}$.

A request is immediately rejected (lost) if and only if the resource will not be available throughout its reservation interval, i.e., the resource has already been reserved for a time period overlapping the requested reservation interval. A sample path of this on-line (greedy) reservation process determines the fragmentation of available time into an alternating sequence of reserved and available (idle) intervals. Hereafter, we only consider stationary versions of reservation processes. (For all of the models that we study it is easily verified that the reservation process has a unique stationary regime.) We study the problem of computing a measure of the fragmentation, called the *reservation probability*, which is defined as the probability that the resource is in use at any time.

The variety of potential applications covered by models like the one above is huge, but it is the relatively new applications in existing and proposed communication systems, e.g., teleconferencing and video-on-demand systems, that have given a fresh impetus to research on reservation systems. In connection with implementations, it was argued in [4] that a slight modification of the proposed Internet resource reservation protocol (RSVP) can provide a framework for creating advance reservation services in the Internet environment. The existence of a commercial resource reservation service offered by AT&T was reported in [8].

We have chosen a baseline model because of its tractability and because we seek new insights supported by an exact analysis of mathematical models. The details of more elaborate systems, either existing or proposed, can be found in the literature [8]; the final section (Section 5) reviews a number of features in connection with promising open problems.

Most previous work in the communications field is quite recent and focuses as much on engineering issues as mathematical foundations. Past engineering research has dealt with the implementation issues of incorporating distributed advance-notice reservation protocols in current networks [4, 6, 10, 15]. In the combinatorial analysis of mathematical models, typical problems have been formulated as follows: Given k identical copies of a resource that can accommodate requests one at a time, and a set of requests each defined by a start and holding time, determine optimal schedules according to standard objectives such as the fraction of requests scheduled, resource utilization, etc. For results on problems of this type, see [1, 9, 11, 14]. In the stochastic model of [8], the authors propose an efficient admission control algorithm that allows sharing of resources between immediate requests (no advance notice) and advance reservations. The work of Virtamo[14] is closer to ours in that he also examines the reservation problem from the point of view of interval packing. However, his analysis of the single resource case is less general and lacks our limit law establishing the underlying connection between the two problems.

Subsequent sections show that, in spite of the simple structure of our reservation model, an exact analysis appears to be quite difficult without further simplification. In the next section, we analyze a number of simplified models, beginning with an elementary slotted system in which reservation intervals have unit durations starting at integer times. Interference among requests is greatly simplified in that a reservation interval not covering time t can not conflict with (overlap) one that does. We then study a "short notice" system where every advance notice is shorter than all reserved intervals. With this assumption, there can be at most one outstanding reservation at any time, but the analysis still requires some effort.

In our last simplified model, we consider bimodal advance-notice distributions where $A = 0$ or a for some fixed a. A successful analysis relies on our ability to handle the book-ahead ($A = a$) requests separately from the immediate ($A = 0$) requests. Recall that the bimodal model is also studied in [8].

Section 4 presents our most general result, in the form of a limit law. Specifically, we obtain asymptotic reservation probabilities as the advance-notice distribution becomes progressively more spread out (e.g., for a sequence of distributions with densities, the suprema of the densities tend to 0). Our main result establishes an intriguing connection between the reservation problem and a stochastic version of *on-line interval packing* [3]. In this problem, intervals of varying lengths arrive at random times and are packed in a given 'containing' interval $[0, x]$ if they are subintervals of $[0, x]$ and do not overlap intervals previously packed. Section 3 studies the stochastic packing problem in isolation, and extends the model in [3] to account for general advance-notice distributions with finite support. The first part of Section 4 extends usual interval packing rule to one with indeterminate, or "fuzzy" behavior at the boundaries 0 and x. The latter rule is used in the second part of Section 4 in a proof of the limit law. The third and last part of Section 4, for the special case of unit intervals, contains a stronger sample path relationship between the regular and fuzzy-boundary packing rules. Section 5 concludes the paper with a brief discussion of open problems.

2 Tractable Models

2.1 Slotted time

A reservation interval is a single time slot which can be any unit-duration interval beginning at an integer time. The advance-notice distribution is discretized and denoted by $\{a_i, i \geq 0\}$, where a_i is the probability that an arrival at time t wants a copy of the resource during $[\lceil t\rceil + i, \lceil t\rceil + i + 1)$. Let N_i be the number of arrivals in the slot $[-(i+1), -i)$, $i \geq 0$, so that $\mathrm{E}N_i = \lambda$. The reservation probability p is computed as the probability that at time 0 the interval $[0,1)$ is reserved. The arrival process need not be Poisson, but we assume that the N_i are i.i.d. random variables.

Let an attempt to reserve $[0,1)$ be called a 'hit', and let H be the number of hits by time 0. Write $H = \sum_{i\geq 0} H_i$, where H_i is the number of hits from arrivals in the interval $[-(i+1), -i)$, $i \geq 0$. The H_i are i.i.d. random variables with means λa_i, $i = 0, 1, \ldots$. Then, in terms of generating functions, $\mathrm{E}z^H = \prod_{i\geq 0} \mathrm{E}z^{H_i}$, where

$$\begin{aligned}
\mathrm{E}z^{H_i} &= \mathrm{E}[\mathrm{E}[z^{H_i}|N_i]] = \mathrm{E}\left[\sum_{j=0}^{N_i} \binom{N_i}{j} a_i^j (1-a_i)^{N_i-j} z^j\right] \\
&= \mathrm{E}(1 - a_i + z a_i)^{N_i},
\end{aligned}$$

so if we let $q_n = \mathbf{P}(N_i = n)$, then $\mathrm{E}z^{N_i} = \sum_{n=0}^{\infty}(1 - a_i + za_i)^n q_n$, and hence

$$\mathrm{E}z^H = \prod_{i=0}^{\infty} \sum_{n=0}^{\infty} (1 - a_i + za_i)^n q_n \tag{1}$$

In principle, we can obtain p, the probability of at least one hit, by subtracting from (1) the first coefficient in its expansion in powers of z.

Example. For Poisson arrivals we have $q_n = e^{-\lambda}\lambda^n/n!$, so

$$\begin{aligned}
\mathrm{E}z^H &= \prod_{i=0}^{\infty} \sum_{n=0}^{\infty} (1 - a_i + za_i)^n \frac{\lambda^n}{n!} e^{-\lambda} = \prod_{i=0}^{\infty} e^{-\lambda a_i (1-z)} \\
&= e^{-\lambda(1-z)}.
\end{aligned}$$

Thus, the number of hits in any interval is Poisson with mean λ, *independent of the advance notice distribution.* We have proved

Proposition 1 *In a slotted system with Poisson arrivals at rate λ and any advance notice distribution, the reservation probability is $1 - e^{-\lambda}$.*

∎

Remarks Note that unit reservation intervals without restriction to equally spaced starting times is a difficult problem, since it is no longer true that a time t is covered by a reserved interval if and only if there has been at least 1 request for a reservation interval covering t; in a slotted system all such intervals must be the same interval, but not in a system without the slot restriction where intervals not covering t can cause intervals that do cover t to be rejected.

The slotted model is one of the few models that is easy to generalize to $k \geq 1$ copies of the resource. In this case, a request, say for $[i, i+1)$, is rejected if and only if there are already k outstanding requests for $[i, i+1)$; and p is the probability that, at time 0, all k copies of the resource are reserved. Then the reservation probability is computed as 1 minus the sum of the first k terms in the expansion of (1) in powers of z. For the example of Poisson arrivals, we obtain $p = \sum_{i \geq k} e^{-\lambda} \lambda^i / i!$. ∎

Effects of the advance-notice distribution. We give examples below for which p depends on the advance-notice distribution.

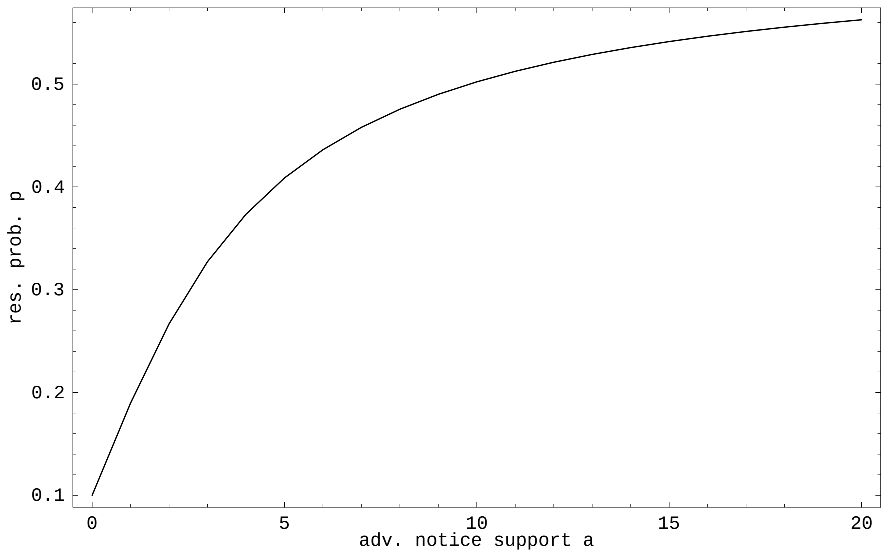

Figure 1: Illustration for example 1

Example 1 First, we demonstrate how the presence of the advance-notice distribution can increase the reservation probability and the utilization of the system. Informally, one would expect this to occur when there are spikes in the arrival distribution. This is because the randomization in the advance-notice distribution smooths out the arrival process and increases the chances of reservation. For example, if we take $\mathbf{P}(N_i = 10) = 1 - \mathbf{P}(N_i = 0) = 0.1$, and $\mathbf{P}(A = j) = 1/a, 0 \leq j \leq a$, then the dependence of the reservation probability on the advance-notice support parameter a is as shown in Figure 1. From the figure we can see that, for $a = 20$, the reservation probability is almost six times that for zero advance notice, $a = 0$. In general, the increase of the reservation probability and the utilization of the system can be made arbitrarily large by increasing the peak value of N_i.

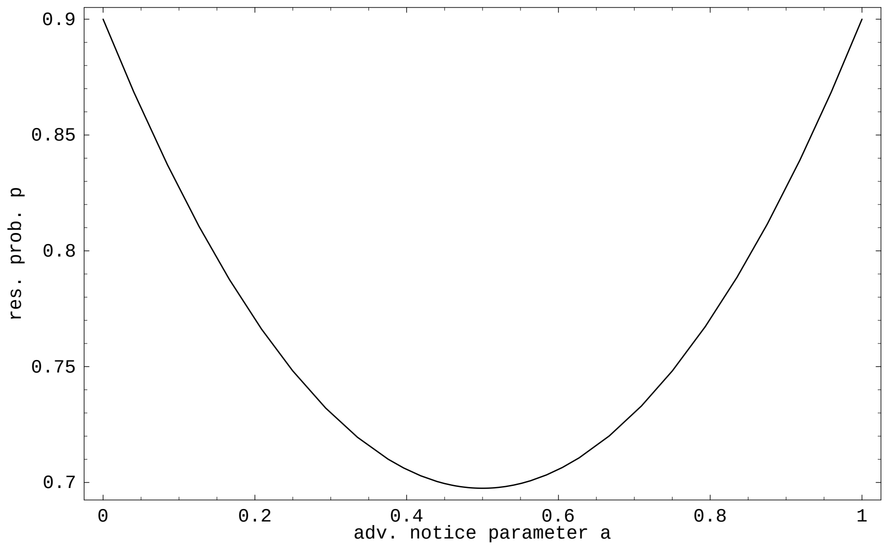

Figure 2: Illustration for example 2

Example 2 Contrary to the previous example, there are situations when the presence of advance notice can decrease the reservation probability (and the utilization of the system). To see this consider a sequence of Bernoulli arrivals $\mathbf{P}(N_i = 1) = 1 - \mathbf{P}(N_i = 0) = 0.9$. Clearly, without the advance notice distribution the probability that the request is rejected is equal to zero, and the reservation probability is $p = 0.9$. However, with the presence of advance notices the rejection probability is positive and the reservation probability is decreased. For the choice $\mathbf{P}(A = 10) = 1 - \mathbf{P}(A = 0) = a$, the dependence of the reservation probability on a is plotted in Figure 2.

2.2 Short notice

For the remainder of the paper, we return to the continuous-time model with Poisson arrivals. The calculation of p is still not difficult if we assume that the support of the advance-notice distribution is to the left of the support of the distribution of the durations of reservation intervals, i.e., the longest advance notice is shorter than the shortest reservation interval. This system is much simplified in that at any time t there can be at most one reserved interval starting at a time $t' > t$. (To see this, note that the existence of two such intervals would imply that the later one had an advance notice exceeding the duration of the earlier one.)

To compute the reservation probability p, we first compute the expected duration of idle periods when the resource is not being used. We focus on a reserved interval, let I denote the duration of the next idle period, and define $f(t)$ as the expected value of I given that time t remains in the current reserved interval. If no time remains, then I is the time to the next arrival plus an advance notice A, and so we have the initial condition $f(0) = 1/\lambda + \mathrm{E}A$. Note also that $f(t)$ must be monotone nonincreasing in t and constant for $t \geq a$, where a is the supremum of the support of $\mathcal{A}(x)$.

For $0 < t \leq a$, we describe the behavior of $f(t)$ in a small interval $[t - \Delta t, t]$ as follows:

$$f(t) = (1-\lambda\Delta t)f(t-\Delta t) + \lambda\Delta t \mathrm{E}[(A-t)\mathbf{1}_{A>t}] + \lambda\Delta t P(A \leq t)f(t) + o(\Delta t). \quad (2)$$

The first and third terms reflect the fact that there is no change in $[t - \Delta t, t]$ if no arrival occurs, or if an arrival occurs but is rejected (it specifies an advance notice less than the remaining time of the current reserved interval). The second term corresponds to an accepted arrival with an advance notice $A > t$; the next idle period will have duration $A - t$.

Now (2) yields the first-order differential equation

$$\begin{aligned} 0 &= -\lambda f(t) - f'(t) + \lambda \mathrm{E}[(A-t)\mathbf{1}_{A>t}] + \lambda P(A \leq t) f(t) \\ &= -f'(t) - \lambda P(A > t) f(t) + \lambda \mathrm{E}[(A-t)\mathbf{1}_{A>t}], \end{aligned} \quad (3)$$

which has the general solution

$$f(t) = e^{-\int_0^t \lambda P(A>u)du} \left[\int_0^t \lambda \mathrm{E}[(A-x)\mathbf{1}_{A>t}] e^{\int_0^x \lambda P(A>u)du} dx + f(0) \right], \quad (4)$$

where $f(0) = 1/\lambda + \mathrm{E}A$. Now introduce $\int_0^x P(A > u)du = \mathrm{E}A - \int_x^\infty P(A > u)du$ and expand (4) to get

$$f(t) = e^{\int_t^\infty \lambda P(A>u)du} \left[\int_0^t dx \int_x^\infty \cdot\lambda(u-x) d\mathcal{A}(u) e^{-\int_x^\infty \lambda P(A>u)} + f(0) e^{-\lambda \mathrm{E}A} \right]. \quad (5)$$

Integrating by parts gives $g(x) := \int_x^\infty \lambda P(A > u)du = \int_x^\infty \lambda(u-x)d\mathcal{A}(u)$, so (5) reduces to

$$f(t) = e^{g(t)} \int_0^t g(x)e^{-g(x)}dx + (1/\lambda + \mathrm{E}A)e^{-\lambda \mathrm{E}A}. \tag{6}$$

Finally, $g(t) \to 0$ as $t \to a$, so

$$f(a) = \int_0^a g(t)e^{-g(t)}dt + (1/\lambda + \mathrm{E}A)e^{-\lambda \mathrm{E}A}, \tag{7}$$

is the unconditional expected duration of an idle period. This and the Law of Large Numbers for stationary processes gives

Proposition 2 *In the short notice system, the reservation probability is* $p = \mathrm{E}B/[f(a) + \mathrm{E}B]$, *where* $f(a)$ *is given by (7).*

2.3 Bimodal advance notice

In this section, as another tractable special case of the advance-notice distribution, the bimodal distribution is considered, where $P(A = a) = 1 - P(A = 0) = q$. For simplicity, we also assume unit reservation intervals. We can assume that $a > 1$, for otherwise, we would just have another instance of the short-notice problem of the previous section.

Note that, according to the Poisson decomposition theorem, those requests that occur with advance notice a compose a Poisson process of rate $q\lambda$. Similarly, the requests that arrive with zero advance notice ($A = 0$) form a Poisson process of rate $(1-q)\lambda$. In addition, these two processes are independent of each other.

The main observations that make the bimodal set-up tractable are the following: first, since $a > 1$, the requests that arrive with advance reservation time a can not collide with the requests that have zero advance notice. Hence, the time between two successive reserved intervals that were requested with advance notice a is exponentialy distributed with parameter $q\lambda$. We term this random variable T_a.

Second, the reservation requests with zero advance notice can only fill the empty gaps of durations T_a between the reservations made with advance notice a. To compute the length of an idle interval in this reservation process, let T_0 be a random variable independent of T_a and exponentially distributed with parameter $(1-q)\lambda$. Then, by the strong Markov property, the duration of an idle interval is equal in distribution to a random variable I with $I = T_a$ if $T_0 + 1 > T_a$ and $I = T_0$ otherwise. Then an easy calculation gives

$$\mathbf{E}I = \frac{1 - e^{-q\lambda}}{q\lambda} - (1-q)e^{-q\lambda} + \frac{e^{-q\lambda}}{\lambda}, \tag{8}$$

and hence the following result.

Proposition 3 *In the case of bimodal advance-notice distributions with $a > 1$ and unit reservation intervals, we have $p = 1/[1+\mathbf{E}I]$, where $\mathbf{E}I$ is given by (8).*

Remark. As a partial check of our results, a straightforward evaluation of (7) shows that $\lim_{a\uparrow 1} f(a) = \mathbf{E}I$, where $\mathbf{E}I$ is the same as in (8). ∎

We have established that, for bimodal advance-notice distributions, explicit formulas for p are available for all values of λ, a, and q, $0 \le q \le 1$. This remains true under several extensions of the model. For example, it is useful to have durations for immediate requests that differ from those booking ahead. Hofri[1] has pointed out that the following such model is easily analyzed.

Suppose that immediate requests are for b time units, and book-ahead requests are for one time unit, with the other parameters as before. To avoid the short-notice set-up, and to keep immediate requests from interfering with book-ahead requests, we need $a > \max(b, 1)$. Sample paths of the reservation process consist of advance reservations (unit intervals booked ahead) alternating with sequences of 0 or more length-b intervals requested immediately. It is easily verified that the stationary distribution of the number V of accepted immediate requests between consecutive advance reservations is geometric with mean

$$\mathrm{E}V = \frac{(1-q)e^{-\lambda q b}}{1-(1-q)e^{-\lambda q b}}.$$

The expected value of the time C between the beginnings of consecutive advance reservations is $\mathrm{E}C = 1 + 1/(\lambda q)$, so we can compute the reservation probability from $p = [1 + b\mathrm{E}V]/\mathrm{E}C$.

It is worth mentioning that the analysis can also be generalized to n-modal ($n \ge 2$) advance-notice distributions $\mathbf{P}(A = a_i) = q_i, 0 \le i \le n$, with $a_i - a_{i-1} > 1, i > 0$, and unit reservation intervals. In this case the reservation requests that arrive with the largest advance notice a_n can not overlap the other requests with smaller advance notices. The same reasoning applies to the requests with advance notices a_i, $i \le n-1$. Thus, an equivalent recursive procedure for constructing the resulting reservation process is as follows. First, construct a reservation process that is obtained from Poisson arrivals of requests at rate $q_n\lambda$ and advance notice a_n. Then fill the gaps (idle intervals) in this process from the Poisson arrivals of requests at rate $q_{n-1}\lambda$ and advance notice a_{n-1}. Next, take the process resulting from this construction and fill its gaps with the Poisson request arrivals at rate $q_{n-2}\lambda$. Continue this procedure recursively,

[1] private communication

by using Poisson rates $q_{n-3}\lambda, \ldots, q_0\lambda$, respectively. Unfortunately, although the construction of the reservation process appears to be straightforward, its analysis, even for $n = 3$, appears to be complicated. Further analysis of the n-modal system is left as an open problem.

3 On-line Interval Packing

The results of the next section relate asymptotics of the *interval packing problem* to asymptotics of our reservation problem. In the interval packing problem, intervals arrive randomly in $\mathbf{R}_+^2$ according to a Poisson process in the two dimensions representing arrival times t and the left endpoints of the arriving intervals. Interval lengths are i.i.d. with distribution $\mathcal{B}$. The intensity is 1, i.e., an average of one interval arrives per unit time per unit distance. For a given $x > 0$, an arriving interval is packed (or accepted) if and only if it is a subinterval of $[0, x]$ and it does not overlap an interval already accepted. The problem is to find, or at least estimate, the function $K(t, x)$, which is the expected length of the intervals accepted by time $t > 0$, assuming that none has yet been accepted by time 0 ($[0, x]$ is initially empty).

To relate the reservation and interval packing problems, suppose we restrict ourselves to advance notice distributions uniform on $[0, y]$; we generalize this assumption in the next section. Consider the reservation problems defined by an arrival rate λ, a duration distribution $\mathcal{B}$, and advance-notice distributions uniform on $[0, y_n]$, $n \geq 1$, where the y_n form an increasing sequence. The next section shows that the asymptotic large-n behavior of these reservation problems is the same as the asymptotic large-x behavior at time λ in the interval packing problem with $\mathcal{B}$ the interval-length distribution.

The remainder of this section derives an estimate for $K(t, x)$, assuming that the support of $\mathcal{B}$ is contained in $[\delta, d], \delta > 0$; for simplicity we fix $\delta = 1$ as the time unit. The next section will then make use of this estimate in an asymptotic analysis of the reservation problem. Our analysis generalizes that in [3] for the case of unit-length intervals ($d = 1$), but leads to a result expressed in terms of functions that, in general, can be computed only numerically. Following the analysis below, special cases where explicit results are possible will be illustrated by an analysis of the case $B \in \{1, 2\}$ with $\mathbf{P}(B = 1) = p = 1 - \mathbf{P}(B = 2) \geq 0$.

Let $\{B_i, i \geq 1\}$ be a sequence of i.i.d. random variables independent of the arrival process, with B_i denoting the length of the i-th arriving interval. Let $b_1 := \mathbf{E}B_i$, $b_2 := \mathbf{E}B_i^2$ denote the first two moments.

3.1 The transform of $K(t, x)$

Denote by $L_x(t)$ the total length of the intervals packed at time t in $[0, x]$, and let $K(t, x) := \mathbf{E}L_x(t)$. Consider the behavior of K during $[t, t + \Delta t]$ in the interval $[0, x]$, with $x \geq d$. Arriving intervals that overlap (x, ∞) are rejected, so if no interval contained in $[0, x]$ arrives at some time in $[0, \Delta t]$, then nothing happens and $K(t + \Delta t, x) = K(t, x)$; the probability of this happening is $\Delta t(x - b_1)$. On the other hand, if an interval with length in $[z, z + \Delta z]$, $1 \leq z < d$, arrives at position $[y, y + \Delta y]$ with $0 \leq y \leq x - z$, then $K(t, t + \Delta t)$ will be the sum of $K(t, y)$ and $K(t, x - y - z)$ plus z for the interval packed. The probability of this event is $\Delta t \Delta \mathcal{B}(z) \Delta y$, so we have, as $\Delta t \to 0$,

$$\begin{aligned} K(t + \Delta t, x) &= [1 - \Delta t(x - b_1)]K(t, x) \\ &+ \Delta t \int_1^d d\mathcal{B}(z) \int_0^{x-z} \mathbf{E}[L_y(t) + L_{x-y-z}(t) + z]dy + o(\Delta t). \quad (9) \end{aligned}$$

Simplifying the integral, dividing by Δt, and letting $\Delta t \to 0$ gives

$$\frac{\partial K(t, x)}{\partial t} = -(x - b_1)K(t, x) + b_1 x - b_2 + 2 \int_1^d d\mathcal{B}(z) \int_0^{x-z} K(t, y)dy. \quad (10)$$

Next, define the transforms $\mathcal{K}(t, x) := \int_d^\infty e^{-ux} K(t, x)dx$, and $\mathcal{B}^*(u) := \int_1^d e^{-uz} d\mathcal{B}(z)$. The transform of (10) is then

$$\begin{aligned} \frac{\partial \mathcal{K}(t, u)}{\partial t} &= \frac{[b_1 + u(b_1 d - b_2)]e^{-du}}{u^2} + \frac{\partial \mathcal{K}(t, u)}{\partial u} + b_1 \mathcal{K}(t, u) \\ &\quad +2 \int_1^d d\mathcal{B}(z) \int_d^\infty e^{-ux} dx \int_0^{x-z} K(t, y)dy. \quad (11) \end{aligned}$$

To simplify the last term in (11), define the functionals

$$\mathcal{H}_0(t, u) := \int_0^d e^{-uz} K(t, z)dz \quad (12)$$

and

$$\mathcal{H}_1(t, u) := u \int_1^d d\mathcal{B}(z) \int_z^d e^{-ux} dx \int_0^{x-z} K(t, y)dy, \quad (13)$$

determined by $K(t, z)$ at the "boundary" $1 \leq z \leq d$. A more convenient representation of $\mathcal{H}_1(t, u)$ follows from an integration by parts,

$$\mathcal{H}_1(t, u) = \int_1^d d\mathcal{B}(z) \left(-e^{-ud} \int_0^{d-z} K(t, y)dy + e^{-uz} \int_0^{d-z} e^{-uy} K(t, y)dy \right). \quad (14)$$

Then the last term of (11) can be written

$$\begin{aligned}
2\int_1^d d\mathcal{B}(z)\int_d^\infty e^{-ux}dx\int_0^{x-z}K(t,y)dy \\
&= 2\int_1^d d\mathcal{B}(z)\left(\int_z^\infty - \int_z^d\right)e^{-ux}dx\int_0^{x-z}K(t,y)dy \\
&= 2\int_1^d d\mathcal{B}(z)\int_0^\infty K(t,y)dy\int_{z+y}^\infty e^{-ux}dx - \frac{2\mathcal{H}_1(t,u)}{u} \\
&= \frac{2\mathcal{B}^*(u)}{u}\mathcal{K}(t,u) + \frac{2\mathcal{B}^*(u)}{u}\mathcal{H}_0(t,u) - \frac{2\mathcal{H}_1(t,u)}{u} \qquad (15)
\end{aligned}$$

which when substituted into (11) yields

$$\frac{\partial\mathcal{K}(t,u)}{\partial t} = \frac{\partial\mathcal{K}(t,u)}{\partial u} + \left(b_1 + \frac{2\mathcal{B}^*(u)}{u}\right)\mathcal{K}(t,u) + \frac{\mathcal{C}(t,u)}{u^2}, \qquad (16)$$

where

$$\mathcal{C}(t,u) := [b_1 + u(b_1 d - b_2)]e^{-du} + 2u\mathcal{B}^*(u)\mathcal{H}_0(t,u) - 2u\mathcal{H}_1(t,u). \qquad (17)$$

To solve (16), it is convenient to define $\mathcal{M}(v,s) := \mathcal{K}(v,s-v)$, so that

$$\frac{\partial\mathcal{M}(v,s)}{\partial v} = \frac{\partial\mathcal{K}(t,u)}{\partial t} - \frac{\partial\mathcal{K}(t,u)}{\partial u}\bigg|_{t=v,u=s-v.} \qquad (18)$$

If we substitute (18) into (16), then for any fixed s, (16) becomes a first-degree, ordinary differential equation

$$\frac{\partial\mathcal{M}(v,s)}{\partial v} = \left(b_1 + 2\frac{\mathcal{B}^*(s-v)}{s-v}\right)\mathcal{M}(v,s) + \frac{\mathcal{C}(v,s-v)}{(s-v)^2}, \qquad (19)$$

with the boundary condition $\mathcal{M}(0,s) = \mathcal{K}(0,s) = 0$. The solution is

$$\mathcal{M}(v,s) = \int_0^v \frac{\mathcal{C}(\xi,s-\xi)}{(s-\xi)^2}\exp\left(\int_\xi^v\left(b_1 + \frac{2\mathcal{B}^*(s-x)}{s-x}\right)dx\right)d\xi. \qquad (20)$$

Finally, by using $\mathcal{K}(t,u) = \mathcal{M}(t,t+u)$ and replacing v,s with $t,t+u$ in (20), we obtain

$$\mathcal{K}(t,u) = \int_0^t \frac{\mathcal{C}(\xi,u+t-\xi)}{(u+t-\xi)^2}e^{b_1(t-\xi)}e^{2\int_\xi^t \frac{\mathcal{B}^*(u+t-x)}{u+t-x}dx}d\xi, \qquad (21)$$

which gives, after changes of variables and easy manipulations,

$$\mathcal{K}(t,u) = \frac{1}{u^2}\int_0^t \mathcal{C}(t-z,u+z)e^{b_1 z}e^{-2\int_u^{u+z}\frac{1-\mathcal{B}^*(y)}{y}dy}dz. \qquad (22)$$

It remains to compute the unknown functions $\mathcal{H}_0(t,u)$ and $\mathcal{H}_1(t,u)$, or equivalently the function $K(t,x)$ with x at the boundary $1 \leq x \leq d$. It is convenient to do this by computing a recurrence other than the one in (9). Instead of basing the recurrence on what happens in the first Δt time units, as in (9), we base it on the time of the first arrival. Observe first that $\exp\left(-t\int_0^x (x-w)d\mathcal{B}(w)\right)$ is the probability that there are no intervals packed in $[0,x]$ during $[0,t]$. Thus, using $\exp[-v\int_1^x (x-w)d\mathcal{B}(w)]dvdyd\mathcal{B}(z)$ as the probability that the first arrival occurs during $[v, v+dv]$ with left endpoint in $[y, y+dy]$ and length in $[z, z+dz]$, we compute, for $x \geq 1$,

$$\begin{aligned} K(t,x) &= \int_0^t \exp\left(-(t-v)\int_1^x (x-w)d\mathcal{B}(w)\right) dv \int_1^x d\mathcal{B}(z) \\ &\qquad \int_1^{x-z} (z + K(v,y) + K(v, x-y-z))dy \\ &= \frac{\int_1^x z(x-z)d\mathcal{B}(z)}{\int_1^x (x-w)d\mathcal{B}(w)} \left(1 - e^{-t\int_1^x (x-w)d\mathcal{B}(w)}\right) \\ &\qquad +2\int_0^t e^{-(t-v)\int_1^x (x-w)d\mathcal{B}(w)} dv \int_1^x d\mathcal{B}(z) \int_0^{x-z} K(v,y)dy, \quad (23) \end{aligned}$$

which for $1 \leq x < 2$, reduces to

$$K(t,x) = \frac{\int_1^x z(x-z)d\mathcal{B}(z)}{\int_1^x (x-w)d\mathcal{B}(w)} \left(1 - e^{-t\int_1^x (x-w)d\mathcal{B}(w)}\right). \tag{24}$$

Thus, in general, $K(t,x)$, $1 \leq x \leq d$, is evaluated inductively from (23) and (24), and then $\mathcal{H}_0$ and $\mathcal{H}_1$ are evaluated from (12) and (13).

3.2 Estimates of $K(t,x)$

Leading-term asymptotics in x are given by the following result, in which

$$\alpha(t) := \int_0^t \mathcal{C}(t-v, v)e^{b_1 v} \exp\left(-2\int_0^v \frac{1-\mathcal{B}^*(y)}{y}dy\right) dv, \tag{25}$$

where $\mathcal{C}$ is given by (17).

Theorem 4 *For any $t > 0$, we have $K(t,x) \sim \alpha(t)x \quad \text{as} \quad x \to \infty$.*

Proof: From (22) and (25) it follows easily that $\mathcal{K}(t,u) \sim \alpha(t)/u^2$ as $u \downarrow 0$. This means that, since $\int_0^x K(t,z)dz$ is monotonically increasing in x, Karamata's Tauberian theorem gives $\int_0^x K(t,z)dz \sim \alpha(t)x^2/2$ as $x \to \infty$. Next, since x is the expected number of arrivals per unit time in $[0,x]$, we have $\frac{\partial}{\partial t}K(t,x) \leq x$, which implies in particular that $\frac{\partial}{\partial t}K(t,x)/x^2 \to 0$, as $x \to \infty$. Now divide (10) by x^2, let $x \to \infty$, and substitute the above asymptotic limits to obtain the result of the theorem. ∎

In order to obtain higher-order error terms, we use complex analysis and the Cauchy residue theorem to evaluate directly the Laplace-transform inversion formula

$$K(t,x) = \frac{1}{2\pi i}\int_{\sigma - i\infty}^{\sigma + i\infty} \mathcal{K}(t,u)e^{xu}du; \quad \sigma > 0, x > 0. \tag{26}$$

Hereafter, u should always be taken as complex. The main result needed for the limit law of Section 4 is

Theorem 5 *Fix $\xi, T > 0$ with $\xi > 2T > 0$. Then there exists a constant $\gamma(t)$, such that*[2]

$$\sup_{0 \le t \le T} |K(t,x) - (\alpha(t)x + \gamma(t))| = O(e^{-\xi x}).$$

Remark: By optimizing the hidden multiplicative constant in $O(e^{-\xi x})$, the error bound is easily improved to $O(e^{-\xi x \log x})$ for every $\xi \in (0,1)$. To avoid distracting clutter, we do not prove this stronger result; it is not needed for the limit law of the next section. For the case $d = 1$, the tighter error bound can be found in [3]. ∎

The proof of Theorem 5 is based on the following three lemmas; the proofs of these lemmas are deferred to the next subsection. The first lemma determines analytic properties of $\mathcal{K}$ and gives leading terms of its Laurent expansion.

Lemma 6 *For fixed t, $\mathcal{K}(t,u)$ is analytic for all $u \neq 0$; at $u = 0$, it has a second order pole and the expansion*

$$\mathcal{K}(t,u) = \frac{\alpha(t)}{u^2} + \frac{\gamma(t)}{u} + \cdots, \tag{27}$$

where $\alpha(t)$ is the same as in (25) and $\gamma(t)$ is the constant given by the integral (22) with the integrand replaced by the second order term of its expansion at $u = 0$.

The next lemma gives a growth estimate for $|\mathcal{K}|$.

[2] Here, and in the remaining results of this section, the hidden multiplicative constants of the $O(.)$ estimates depend on ξ and T, unless stated otherwise.

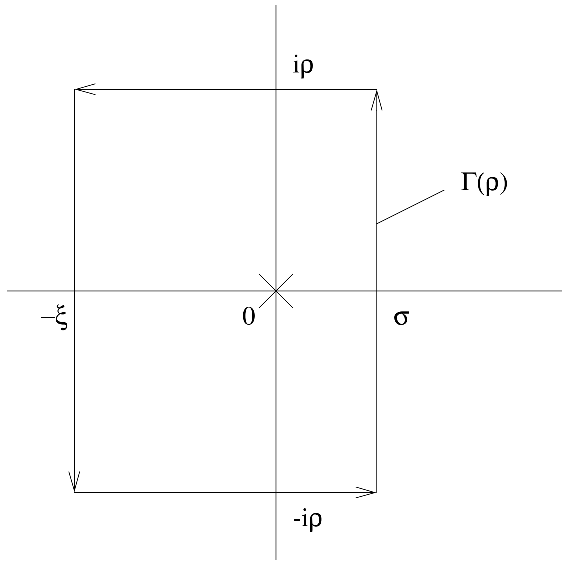

Figure 3: Rectangular contour of integration.

Lemma 7 *Fix* ξ, T *with* $\xi > 2T > 0$. *Then*

$$\sup_{0 \leq t \leq T} |\mathcal{K}(t,u)| = O\left(\frac{1}{|u|}\right), \quad |u| > 2T, \ \Re u \geq -\xi.$$

The following bound completes our preparation for the residue theorem applied in the proof of Theorem 5.

Lemma 8 *If we fix* ξ, T *with* $\xi > 2T > 0$, *then*

$$\sup_{0 \leq t \leq T} \left| \int_{-\xi - i\infty}^{-\xi + i\infty} e^{xu} \mathcal{K}(t,u) du \right| = O(e^{-\xi x}).$$

Proof of Theorem 5: First observe that, by Lemma 6, $\mathcal{K}(u)e^{xu}$ is analytic for all u, except for a pole at $u = 0$. From (27) and the power series expansion $e^{xu} = 1 + xu + \cdots$, we conclude that

$$\mathcal{K}(u)e^{xu} = \frac{\alpha(t)}{u^2} + \frac{\alpha(t)x + \gamma(t)}{u} + \cdots, \quad u \neq 0. \tag{28}$$

For any $\sigma > 2T > 0$, apply the residue theorem to the rectangular contour $\Gamma(\rho)$ sketched in Figure 3, where ρ and ξ are any two reals satisfying $\min(\rho, \xi) > 2T$.

As $\mathcal{K}(u)$ is analytic on and inside $\Gamma(\rho)$, except for a pole at $u=0$, we get from (28)

$$\frac{1}{2\pi i}\int_{\Gamma(\rho)}\mathcal{K}(u)e^{xu}du=\alpha(t)x+\gamma(t). \tag{29}$$

By Lemma 7, the total contribution of the integrals along the horizontal sides of $\Gamma(\rho)$ tends to 0 as $\rho\to\infty$, so (29) becomes

$$\frac{1}{2\pi i}\int_{\sigma-i\infty}^{\sigma+i\infty}\mathcal{K}(u)e^{xu}du=\alpha(t)x+\gamma(t)+\frac{1}{2\pi i}\int_{-\xi-i\infty}^{-\xi+i\infty}\mathcal{K}(u)e^{xu}du. \tag{30}$$

But by Lemma 8, we have

$$\sup_{0\le t\le T}\left|\frac{1}{2\pi i}\int_{-\xi-i\infty}^{-\xi+i\infty}\mathcal{K}(u)e^{xu}du\right|=O(e^{-\xi x}), \tag{31}$$

which together with (30) implies $\sup_{0\le t\le T}|K(t,x)-(\alpha(t)x+\gamma(t))|=O(e^{-\xi x})$, as desired. ∎

We illustrate the results with two examples.

Example 3: Consider the case where all intervals are of unit length, i.e., $d=1$. Here, $\mathcal{H}_0(t,u)=\mathcal{H}_1(t,u)\equiv 0$, $\mathcal{B}^*(v)=e^{-v}$, which when substituted into $\mathcal{C}$ and (25) yields

$$\alpha(t)=\int_0^t e^{-2\int_0^v\frac{1-e^{-y}}{y}dy}dv, \tag{32}$$

the incomplete Rényi constant [3]. As noted earlier, the $d=1$ case was analyzed extensively in [3], where the result above first appeared. Note that, as $t\to\infty$, the interval packing process in $[0,x]$ tends towards an absorbing state in which no more intervals can be packed (all gaps are less than 1 in length). This absorbing state is stochastically identical to the absorbing state in the classical car parking problem on $[0,x]$. Rényi [13] showed that the expected occupied part of (the expected number parked in) $[0,x]$ in these absorbing states is asymptotically $\alpha(\infty)x+O(1)$ as $x\to\infty$. It was proved in [3] that the time to absorption in the interval packing problem is finite almost surely, but its mean is infinite for all $x>2$. For related results and various refinements, see the discussion in [3].

Example 4: Consider the case where interval lengths have only the values 1 or 2, with probabilities $\mathbf{P}(\mathcal{B}_i=1)=1-\mathbf{P}(B_i=2)=p$. From (24), we have for $1\le x<2$, $K(t,x)=1-e^{-pt(x-1)}$, and the transform

$$\mathcal{H}_0(t,u)=\frac{e^{-u}-e^{-2u}}{u}-\frac{e^{-u}-e^{-2u-pt}}{u+pt},\quad \mathcal{H}_1(t,u)\equiv 0.$$

Substituting $\mathcal{H}_0(t,u)$, $\mathcal{H}_1(t,u)$ and $\mathcal{B}^*(v) = pe^{-v} + (1-p)e^{-2v}$ into $\mathcal{C}$, and then $\mathcal{C}$ and $\mathcal{B}^*$ into (25) yields an explicit, albeit complicated, formula for $\alpha(t)$, and hence the asymptotics of Theorem 4. ∎

3.3 Proofs of lemmas

Proof of Lemma 6: Observe that $\int_u^{u+v}[1-\mathcal{B}^*(y)]dy/y$ is an entire function of u, as are $\mathcal{H}_0$ and $\mathcal{H}_1$ and hence $\mathcal{C}$ by (12), (13), and (17). Then by (22), for any fixed t, $\mathcal{K}(t,u)$ is analytic for all $u \neq 0$. At $u = 0$, it has a second order pole, and by (22), it has the Laurent expansion given in (27) for some constant $\gamma(t)$. ∎

Proof of Lemma 7: Consider the following elementary inequalities

$$|\mathcal{B}^*(u)| \leq e^{\xi d}, \quad \Re u \geq -\xi, \tag{33}$$

$$\frac{1}{|u+z|} \leq \frac{2}{|u|} \leq \frac{1}{T}, \quad |u| > 2T,\ 0 \leq z \leq T, \tag{34}$$

which together imply

$$\left| e^{2\int_u^{u+z} \frac{\mathcal{B}^*(y)}{y} dy} \right| \leq e^{2e^{\xi d}}, \quad \Re u \geq -\xi,\ |u| > 2T,\ 0 \leq z \leq T. \tag{35}$$

It is easy to verify that $|\mathcal{H}_i(t-z, u+z)| \leq 2d^2 e^{\xi d}$, $\Re u \geq -\xi$, for $i = 0, 1$, so in conjunction with (17), (33), and (34), we find that

$$\frac{|\mathcal{C}(t-z, u+z)|}{|u+z|^2} = O\left(\frac{1}{|u|}\right), \quad \Re u \geq -\xi,\ |u| > 2T,\ 0 \leq z \leq t \leq T. \tag{36}$$

Finally, after making a change of variables in (21) and then substituting (35) and (36), we complete the proof of the lemma. ∎

Proof of Lemma 8: To simplify the proof, we break $\mathcal{K}$ into two parts $\mathcal{K}_0, \mathcal{K}_1$ and prove the lemma for each part individually. The parts of $\mathcal{K} = \mathcal{K}_0 + \mathcal{K}_1$ are determined in turn by dividing $\mathcal{C}$ into two parts $\mathcal{C}_0, \mathcal{C}_1$. To define these parts, recall the expression for $\mathcal{H}_1$ in (14). We associate the first of the two double integrals with $\mathcal{C}_0$ and the second with $\mathcal{C}_1$ and then divide the remaining terms of (17) so that $\mathcal{C} = \mathcal{C}_0 + \mathcal{C}_1$ with

$$\mathcal{C}_0(t,u) = u(b_1 d - b_2)e^{-ud} + 2ue^{-ud} \int_1^d d\mathcal{B}(z) \int_0^{d-z} K(t,y)dy, \tag{37}$$

$$\mathcal{C}_1(t,u) = b_1 e^{-du} + 2u\mathcal{B}^*(u)\mathcal{H}_0(t,u) - 2u\int_1^d d\mathcal{B}(z)e^{-uz}\int_0^{d-z} e^{-uy}K(t,y)dy. \tag{38}$$

We take care of $\mathcal{K}_1$ first, i.e., we prove Lemma 8 with $\mathcal{K}$ replaced by $\mathcal{K}_1$, where $\mathcal{K}_1$ is given by (21) with $\mathcal{C}$ replaced by $\mathcal{C}_1$ above. This proof will be easy once we prove a $O(1/|u|)$ upper bound on the suprema over $0 \le z \le t \le T$ of $|\mathcal{H}_0(t-z, u+z)|$ and the modulus of the double integral in (38) with t, u replaced by $t-z, u+z$.

To prove the desired bound on $|\mathcal{H}_0|$, we will integrate (12) by parts and bound each of the resulting terms. But to justify that step, we need the following properties of the derivative of $K(t,x)$ with respect to x.

Lemma 9 *$\partial K(t,x)/\partial x$ exists for $1 \le x \le d$ and $t \ge 0$. Furthermore, for any $T > 0$,*

$$\sup_{0\le t\le T,\ 1\le x\le d} \frac{\partial K(t,x)}{\partial x} = O(1), \tag{39}$$

where the hidden constant depends only on T.

Proof: The existence of $\partial K(t,x)/\partial x$ follows easily from the integral representation of $K(t,x)$, $1 \le x \le d$, given by formulas (23) and (24). The bound in (39) follows by differentiating (23) and (24), and by observing that all of the terms in the derivative are bounded functions over $1 \le x \le d$, $0 \le t \le T$. ∎

We are now ready to prove

Lemma 10 *Fix ξ, T, with $\xi > 2T > 0$, let $\Re u = -\xi$and let $\mathcal{I}_1(t,u)$ denote the double integral in (38). Then*

$$\sup_{0\le z\le t\le T} \max(|\mathcal{H}_0(t-z,u+z)|, |\mathcal{I}_1(t-z,u+z)|) = O\left(\frac{1}{|u|}\right).$$

Proof: To prove the bound for $|\mathcal{H}_0(t-z,u+z)|$, use the first part of Lemma 9, integrate by parts the expression for $\mathcal{H}_0(t-z,u+z)$ (see (12)) and obtain

$$\begin{aligned}\int_0^d e^{-(u+z)y}K(t-z,y)dy = \quad & - \quad \frac{e^{-(u+z)d}K(t-z,d)}{u+z} \\ & + \quad \frac{1}{u+z}\int_0^d e^{-(u+z)y}\frac{\partial K(t-z,y)}{\partial y}dy.\end{aligned}$$

Then, by applying the bound in Lemma 9, the inequalities in (34), and the estimates $K(t,y) \le y \le d$ and $|e^{-(u+z)y}| \le e^{\xi d}$ for $y \le d$, we obtain the desired estimate

$$\sup_{0\le z\le t\le T} |\mathcal{H}_0(t-z,u+z)| = O\left(\frac{1}{|u|}\right). \tag{40}$$

Using exactly the same arguments, we can prove a similar estimate for $|\mathcal{I}_1(t-z,u+z)|$; we skip the details. ∎

We now claim that, for fixed ξ, T with $\xi > 2T > 0$, if we let $\Re u = -\xi$, then

$$\sup_{0\le z\le T} |\mathcal{K}_1(t,u)| = O\left(\frac{1}{|u|^2}\right). \tag{41}$$

A proof based on Lemma 10 and the estimates (33)-(35) is completely analogous to the proof of Lemma 7, so we omit the details. From (41) we can conclude that, as desired,

$$\sup_{0\le t\le T} \left|\int_{-\xi-i\infty}^{-\xi+i\infty} e^{xu}\mathcal{K}_1(t,u)du\right| = O(e^{-\xi x}). \tag{42}$$

We prove Lemma 8 with $\mathcal{K}$ replaced by $\mathcal{K}_0$, where $\mathcal{K}_0$ is given by (21) with $\mathcal{C}$ replaced by $\mathcal{C}_0$ in (37); changing variables in (21), this gives

$$\mathcal{K}_0(t,u) = \int_0^t \frac{\mathcal{C}_0(t-z,u+z)}{(u+z)^2} e^{b_1 z} e^{2\int_u^{u+z}\frac{\mathcal{B}^*(y)}{y}dy} dz. \tag{43}$$

Using (37), it is helpful to put (43) in the form

$$\begin{aligned}\mathcal{K}_0(t,u) &= \int_0^t \frac{\mathcal{C}_*(t,z)e^{-ud}}{(u+z)} e^{2\int_u^{u+z}\frac{\mathcal{B}^*(y)}{y}dy} dz \\ &= \int_0^t \frac{\mathcal{C}_*(t,z)e^{-ud}}{(u+z)}\left(e^{2\int_u^{u+z}\frac{\mathcal{B}^*(y)}{y}dy} - 1\right)dz + \int_0^t \frac{\mathcal{C}_*(t,z)e^{-ud}}{(u+z)}dz, \end{aligned}\tag{44}$$

where

$$\mathcal{C}_*(t,z) := e^{-(d-b_1)z}\left[b_1 d - b_2 + 2\int_1^d d\mathcal{B}(z)\int_0^{d-z} K(t-z,y)dy\right]. \tag{45}$$

We treat the two integrals in (44) separately beginning with the second.

Lemma 11 *Fix ξ, T satisfying $\xi > 2T > 0$, and let $\Re u = -\xi$. Then*

$$\left|\int_{-\xi-i\infty}^{-\xi+i\infty} e^{ux}du \int_0^t \frac{\mathcal{C}_*(t,z)e^{-ud}}{(u+z)}dz\right| = O(e^{-\xi x}). \tag{46}$$

Proof: We have $|u+z| \geq |u|/2$ for $|u| \geq 2T,\ 0 \leq z \leq T$, and so

$$\left|\frac{1}{u+z} - \frac{1}{u}\right| \leq \frac{2T}{|u|^2}. \tag{47}$$

Also, an integration by parts shows that

$$\left|\int_{-\xi-i\infty}^{-\xi+i\infty} \frac{e^{u(x-d)}}{u} du\right| = O(e^{-\xi x}). \tag{48}$$

Then by (47), (48), and the fact that $\sup_{0\leq z\leq T} \mathcal{C}_*(t,z) < \infty$, the left-hand side of (46) has the upper bound

$$\left|\int_{-\xi-i\infty}^{-\xi+i\infty} e^{u(x-d)} du \int_0^t \mathcal{C}_*(t,z) \left(\frac{1}{u+z} - \frac{1}{u}\right) dz\right| + \left|\int_0^t \mathcal{C}_*(t,z) dz \int_{-\xi-i\infty}^{-\xi+i\infty} \frac{e^{u(x-d)}}{u} du\right|$$
$$\leq e^{-\xi(x-d)} \left|\int_0^t C_*(t,z)dz\right| \int_{-\infty}^{\infty} \frac{2T}{\omega^2+\xi^2} d\omega + O(e^{-\xi x}) = O(e^{-\xi x}),$$

as claimed. ∎

Turning to the first integral in (44), we start with a bound on its second factor.

Lemma 12 *Fix ξ, T with $\xi > 2T > 0$, and let $\Re u = -\xi$. Then*

$$\sup_{0\leq z\leq T} \left|\exp\left(2\int_u^{u+z} \frac{\mathcal{B}^*(y)}{y} dy\right) - 1\right| = O\left(\frac{1}{|u|}\right). \tag{49}$$

Proof: Note that, by (33) and (34),

$$\left|\int_u^{u+z} \frac{\mathcal{B}^*(y)}{y} dy\right| \leq \frac{2T e^{\xi d}}{|u|}, \quad \Re u = -\xi,\ \xi > 2T. \tag{50}$$

The lemma follows by combining (50) with the elementary inequalities $|e^z - 1| \leq e^{|z|} - 1 \leq |z|$, which hold for any complex z. ∎

Letting $\mathcal{I}_0$ denote the first integral in (44), we see that, by Lemma 12, the inequalities of (33), and the boundedness of $\mathcal{C}_*$,

$$\left|\int_{-\xi-i\infty}^{-\xi+i\infty} e^{xu} \mathcal{I}_0(t,u) du\right| = O(e^{-\xi x}), \tag{51}$$

which together with (46) for the second integral yields

$$\sup_{0\leq t\leq T}\left|\int_{-\xi-i\infty}^{-\xi+i\infty} e^{xu}\mathcal{K}_0(t,u)du\right| = O(e^{-\xi x}). \tag{52}$$

Lemma 8 is immediate from (42) and (52), since $\mathcal{K} = \mathcal{K}_0 + \mathcal{K}_1$. ∎

4 Limit Law

The asymptotic analysis of this section begins in Section 4.1 with preliminary results that apply to any interval-length distribution on $[1, d]$, $1 \leq d < \infty$. In Section 4.2 we apply the results of Section 4.1 to prove our main result, the asymptotic reservation probability as the advance-notice distribution becomes progressively more spread out, in a sense to be formalized in Section 4.2.

The special case $d = 1$ is treated separately in Section 4.3; the preliminary results for this case are much stronger and of independent interest; we will prove sample-path monotonicity properties that do not hold for arbitrary $d > 1$.

To help guide the reader, we give in rough outline the intuition behind the limit law for the case where advance notices are uniform on $[0, a]$, intervals have unit length, and request arrivals are Poisson at rate λ. At some time later than x, consider the (unit) intervals that were reserved in $[0, x]$, assuming that x is large and that a is large relative to x. Reservations in $[0, x]$ could not have been made earlier than $-a$, since all advance notices are no larger than a. It is easy to see that, between $-a$ and x, the process of making reservations in $[0, x]$ is approximately the same as the process of packing unit intervals in $[0, x]$; i.e., the reserved intervals are scattered (approximately) uniformly at random throughout $[0, x]$. The comparison fails to be exact for two reasons.
(i) During $[-a, -a + x]$ and $[0, x]$ the intervals reserved in $[0, x]$ are not located uniformly over $[0, x]$, as in interval packing.
(ii) The reservation process allows reserved intervals to overlap 0 and x, but the interval packing process does not.
But as x becomes large, and a becomes large relative to x, one expects the influence of these edge effects to be negligible. The remainder of this section will prove that this is indeed so, and in a much more general setting.

4.1 Fuzzy packing processes

Consider the two-dimensional Poisson arrival process of Section 3, with rate 1 per unit time per unit length. We define an *interval packing process on* $[0, x]$

with fuzzy boundary at x (or simply a *fuzzy process $\mathcal{P}$*) to be a process that packs disjoint intervals with left endpoints in $[0, x]$ as follows: If such an interval overlaps any interval already packed, it is rejected. Otherwise, if it is wholly contained in $[0, x]$, then it is packed, but if it overlaps x, then its acceptance is decided by some given, but arbitrary (possibly randomized) on-line rule. In our application to reservation problems, we will use this model to handle cases in which the acceptance behavior at the boundary is unknown because it depends on events happening outside the window of attention. Let $L_{\mathcal{P},x}(t)$ be the total length of all the packed intervals during time $[0, t]$ whose left endpoints are in $[0, x]$ (including the length of the interval, if any, that covers x); and let $K_{\mathcal{P}}(t, x) := \mathbf{E}L_{\mathcal{P},x}(t)$. The following lemma proves the intuitively appealing fact that $K_{\mathcal{P}}(t, x)$ and $K(t, x)$ can not differ by more than a constant.

Lemma 13 *Let $\mathcal{P}$ be a fuzzy packing process on $[0, x]$ with fuzzy boundary at x. Then for any fixed $t > 0$, we have $|K_{\mathcal{P}}(t, x) - K(t, x)| = O(1)$.*

Proof: Consider the regular process (no overlapping of x is allowed) and the fuzzy process $\mathcal{P}$ operating on the same realization of the Poisson arrival process. Assuming that $\mathcal{P}$ accepts an interval, say I, overlapping x, we let $\tau \equiv \tau_{\mathcal{P}}, 0 \le \tau \le t$, denote the time when it is packed. If no such interval exists, we set $\tau = t$. Let y be the position of the left endpoint of I, and let ℓ be its length. Observe that up to time τ both the regular and fuzzy processes have packed the same intervals. If we let z be the position of the right endpoint of the rightmost interval already packed by both the regular and fuzzy processes, then this observation implies

$$\begin{aligned}\mathbf{E}[(L_{\mathcal{P},x}(t) - L_x(t))|\tau, y, z] &= \mathbf{E}L_{y-z}(t-\tau) + \ell - \mathbf{E}L_{x-z}(t-\tau) \\ &\le \mathbf{E}L_{y-z}(t-\tau) + d - \mathbf{E}L_{x-z}(t-\tau) \qquad (53)\end{aligned}$$

Theorem 5 implies that, for any $t > 0$, there exists a positive constant $c = c(t)$ such that, for all $0 \le \tau \le t$ and all $x \ge 0$,

$$\alpha(t-\tau)x - c \le \mathbf{E}L_x(t-\tau) \le \alpha(t-\tau)x + c, \qquad (54)$$

so the combination of (53) and (54) yields

$$\begin{aligned}\mathbf{E}[(L_{\mathcal{P},x}(t) - L_x(t))|\tau, y, z] &\le d + [\alpha(t-\tau)(y-z) + c] - [\alpha(t-\tau)(x-z) - c] \\ &\le d + 2c - \alpha(t-\tau)(x-y) \le d + 2c. \qquad (55)\end{aligned}$$

Removing the conditioning in (55), we compute the upper bound

$$K_{\mathcal{P}}(t, x) - K(t, x) \le d + 2c. \qquad (56)$$

Similarly, we obtain the lower bound

$$\begin{aligned}
\mathbf{E}[(L_{\mathcal{P},x}(t) - L_x(t))|\tau, y, z] &= \mathbf{E}L_{y-z}(t-\tau) + l - \mathbf{E}L_{x-z}(t-\tau) \\
&\geq \alpha(t-\tau)(y-z) - c - [\alpha(t-\tau)(x-z) + c] \\
&\geq -2c - \alpha(t-\tau)(x-y) \geq -(2c+d), \qquad (57)
\end{aligned}$$

since $\alpha(.) \leq 1$ and $x - y \leq d$. After removing the conditioning, we obtain a lower bound that, in conjunction with (56), proves the lemma. ∎

By symmetry, we can also define a fuzzy boundary at 0. We say that $\mathcal{P}$ is a fuzzy packing process on $[0, x]$ with a fuzzy boundary at zero if it packs disjoint intervals with *right* endpoints in $[0, x]$ as follows: If such an interval overlaps an interval already packed, it is rejected. Otherwise, if it is wholly within $[0, x]$ it is packed, but if it overlaps 0, its acceptance is decided by a given but arbitrary on-line rule. Let $L_{\mathcal{P},x}(t)$ be the total length of intervals packed by $\mathcal{P}$ during $[0, t]$ whose left endpoints are in $[0, x]$ (note that this *excludes* the length of the packed interval, if any, that covers 0); and let $K_{\mathcal{P},x}(t) := \mathbf{E}L_{\mathcal{P},x}(t)$.

Lemma 14 *Let $\mathcal{P}$ be a fuzzy packing process on $[0, x]$ with fuzzy boundary at 0. Then for any fixed $t > 0$, we have $|K_{\mathcal{P}}(t, x) - K(t, x)| = O(1)$.*

Proof: Define τ to be the time when an interval overlapping 0 is packed; let y be the position of the right endpoint of that interval, and let z be the position of the left endpoint of the leftmost interval packed before τ. In terms of these quantities, the proof proceeds in analogy with the proof of Lemma 13; we omit the details. ∎

We will also wish to consider packing processes with fuzzy boundaries at *both* ends of $[0, x]$. This packing process is obtained simply by combining the definitions for the fuzzy processes at 0 and x.

Corollary 15 *Let $\mathcal{P}$ be a fuzzy packing process on $[0, x]$ with fuzzy boundary at 0 and x. Then for any fixed $t > 0$, we have $|K_{\mathcal{P}}(t, x) - K(t, x)| = O(1)$.*

Proof: The proof is analogous to the proofs of Lemmas 13 and 14, except that here we define τ to be the first time when one of the points 0 or x is covered. Then, after this time, the fuzzy process becomes fuzzy at only one of the boundaries and we can use the preceding lemmas. Again, we omit the details. ∎

Corollary 16 *Let $\mathcal{P}$ be any packing process on $[0, x]$ that is fuzzy at 0 and x. Then the expected number of intervals packed by $\mathcal{P}$ is $K_{\mathcal{P},x}(t) = \alpha(t)x + O(1)$.*

Proof: Combine Corollary 15 and Theorem 5. ∎

4.2 Advanced-notice limit law

For any fixed advance-notice distribution $\mathcal{A}$ with compact support, and for any $\lambda, x \geq 0$, we define $G_{\mathcal{A}}(\lambda, x)$ to be the limit as T approaches infinity of the expected length of reserved intervals with left endpoint in $[T, T+x)$, assuming that requests are arriving with total rate λ and with advance-notice distribution given by $\mathcal{A}$. Clearly $G_{\mathcal{A}}(\lambda, x)$ is nondecreasing as a function of x, and it satisfies $G_{\mathcal{A}}(\lambda, x+y) = G_{\mathcal{A}}(\lambda, x) + G_{\mathcal{A}}(\lambda, y)$. It follows that $G_{\mathcal{A}}(\lambda, x) = p_{\mathcal{A}}(\lambda)x$ for some nonnegative number $p_{\mathcal{A}}(\lambda)$, which represents the probability that the resource is booked at a randomly chosen time.

We say that an advance-notice distribution $\mathcal{A}$ has flatness bound (L, ϵ) if for any real z, $|\mathcal{A}(z+L) - \mathcal{A}(z)| \leq \epsilon$. In this section, we are interested in the behavior of $p_{\mathcal{A}}(\lambda)$ as the advance-notice distribution $\mathcal{A}$ becomes "very flat" in the sense that L becomes large and ϵ becomes small.

Theorem 17 *Suppose that for $i = 1, 2, \ldots$, we have a distribution $\mathcal{A}_i$ with compact support, and suppose that $\mathcal{A}_i$ has flatness bound (L_i, ϵ_i). If $\lim_{i\to\infty} L_i = \infty$ and $\lim_{i\to\infty} \epsilon_i = 0$, then*

$$\lim_{i\to\infty} p_{\mathcal{A}_i}(\lambda) = \alpha(\lambda),$$

where α is given by (25).

Proof: To ease notation we temporarily fix one of the distributions $\mathcal{A} = \mathcal{A}_i$, with flatness bound (L, ϵ). The proof will depend on the choice of certain positive parameters $x \in \mathbf{R}$ and $k \in \mathbf{Z}$. As it will turn out, we will need $x \equiv x_i$ and $k \equiv k_i$ to satisfy the conditions $1 \ll x < L$, $\epsilon \ll 1/k$, $k\epsilon x \ll 1$, $x^2/k \ll 1$, where "$f \ll g$" means that the ratio f/g of the quantities f and g associated with $\mathcal{A}_i$ tends to zero as i tends to infinity. These conditions can be satisfied by choosing $k := \lfloor \epsilon^{-1/2} \rfloor$ and $x := \min(L-1, \epsilon^{-1/6})$ for each $\mathcal{A}$.

Because of the flatness bound, it is possible to choose $0 = a_0 < a_1 < \cdots < a_k$ such that the support of $\mathcal{A}$ is contained in $[a_0, a_k]$ and such that

$$|\mathbf{P}(A \in [a_0, a_j]) - j/k| \leq \epsilon, \quad 1 \leq j \leq k,$$

and $|a_j - a_{j-1}| > 2L$. For $j = 1, 2, \ldots, k$, define

$$p_j := \mathbf{P}(A \in [a_{j-1} + x, a_j - x]).$$

It then follows from the flatness bound that $p_j = 1/k + O(\epsilon)$, if $0 < x < L$.

We will say that a request for a reservation with left endpoint in $[T, T+x)$ occurs during episode j if its advance notice is in $I_j := [a_{j-1}+x, a_j - x]$. All requests during episode j are made between times $T - a_j + x$ and $T - a_{j-1}$. In particular all requests in epside j come after all requests in episode $j+1$. Let $W = [a_0, a_k] \setminus \bigcup_j I_j$. Then the expected number of requests for reservations with left endpoint in $[T, T+x)$ which do not fall into any episode is

$$\int_{-\infty}^{\infty} \mathbf{P}\left(A \in W \text{ and } A + t \in [T, T+x)\right) \lambda\, dt. \tag{58}$$

(The integration is with respect to t, which represents the time at which the request is made. Since $\mathcal{A}$ has compact support, we could replace the limits of integration by finite numbers if desired.) By writing the probability $\mathbf{P}(.)$ in (58) as an integral of a distribution, and interchanging the order of integration, we find that this expected number of non-episode requests is $\mathbf{P}\left(A \in W\right) \lambda x$, which by the flatness bound is at most $2k\epsilon\lambda x$. In particular, the probability that there exists a non-episode request at all is bounded by $2k\epsilon\lambda x$, which by assumption is negligible. (Any event occurring with probability tending to zero as $\mathcal{A}$ flattens out may be ignored in our analysis, since its contribution to the expected number of reserved intervals will be $o(x)$, where here and for the rest of the proof $o(x)$ denotes a function that when divided by x tends to zero as $\mathcal{A}$ flattens out.)

Similarly we see that the expected number of episode j requests with left endpoint in $[T, T+x)$ equals $p_j \lambda x$. The analogous statement for any subinterval of $[T, T+x)$ is true (with the x in the assertion replaced by the length of the subinterval), and the numbers of episode j requests in disjoint subintervals are independent, so it follows that the episode j requests will have left endpoints uniformly distributed in $[T, T+x)$, and that their number is a discrete Poisson random variable with mean $p_j \lambda x = \lambda x(1/k + O(\epsilon)) = O(\lambda x/k)$. The probability of two or more such requests occurring during episode j is then $O\left((\lambda x/k)^2\right)$, so the probability of there being *some* episode with two or more requests is at most $k \cdot O\left((\lambda x/k)^2\right) = O\left((\lambda x)^2/k\right)$, which by assumption is negligible. If we discard sample paths with non-episode requests or with two or more requests in some episode, then the behavior of our reservation process in $[T, T+x)$ exactly mimics the fuzzy packing process with time parameter $p_1\lambda + p_2\lambda + \cdots + p_k\lambda$, provided that we also discard in the latter those sample paths in which two packing requests are made in a single episode, where the episodes in the process are defined as the successive time intervals of length $p_1\lambda$, $p_2\lambda$, $\ldots$, $p_k\lambda$ whose concatenation forms the entire time interval over which the packing occurs. The fraction of sample paths so discarded in the packing problem is at most $\sum_{j=1}^{k}(p_j\lambda x)^2$, which is again negligible. Thus the expected length $G_{\mathcal{A}}(\lambda, x)$ of reserved intervals with left endpoint in $[T, T+x)$ differs by at most $o(x)$ from the expected length of intervals packed with left endpoint in $[0, x)$ in a fuzzy packing process (i.e., in $[0, x]$ with fuzzy boundaries at both ends). By Corollary 16, the expected length of packed intervals is $\alpha(\sum_{1 \leq j \leq k} p_j \lambda)x + O(1)$. But

$$1 - o(1) = 1 - 2k\epsilon \leq \sum_{j=1}^{k} p_j \leq 1,$$

so by the continuity of α, the expected length of packed intervals is $\alpha(\lambda)x + o(x)$, as $x \to \infty$. Thus, $G_{\mathcal{A}}(\lambda, x) = \alpha(\lambda)x + o(x)$ also. Dividing by x, and taking the limit as x goes to infinity, we find $p_{\mathcal{A}}(\lambda) = \alpha(\lambda) + o(1)$. We complete the proof by applying this to $\mathcal{A}_i$ with i tending to infinity. ∎

Remark: Let $[0, a_i]$ be the support of $\mathcal{A}_i$. Assume that the reservation process is in its stationary regime and that $t = 0$ represents the present time of the reservation process. Then, it is of interest to compute the fraction of time $p_{\mathcal{A}_i}(\eta, \lambda)$ that in some future interval $(\eta a_i, \eta a_i + x), 0 \leq \eta < 1$, the resource is booked; x is the same as in the proof of the preceding theorem. This is a more user-oriented performance measure which gives an indication of the chances of booking ahead if the advance-notice is roughly ηa_i. Thus, if in addition to the assumptions of Theorem 17, we assume that for any $0 \leq \eta < 1$, $\lim_{i \to \infty} \mathbf{P}(A_i \geq \eta a_i) = \theta(\eta)$, then the proof of the previous theorem can be easily adapted to obtain

$$\lim_{i \to \infty} p_{\mathcal{A}_i}(\eta, \lambda) = \alpha(\theta(\eta)\lambda),$$

where α is given by (25). ∎

4.3 Fuzzy packing processes with $d = 1$

In this subsection, we specialize to the case of unit intervals ($d = 1$) and establish a much stronger coupling between regular and fuzzy packing processes. The results of this section are of independent interest, though they are not needed elsewhere in the paper.

Consider a fuzzy packing process $\mathcal{P}$ on $[0, x]$ with a fuzzy boundary at x. For a given sample path, define $R(x)$ to be a time-ordered list of the requests whose left endpoints are contained in $[0, x]$. (With probability one, $R(x)$ is finite.) Note the following equivalent way of defining an interval packing process on $[0, x]$ with fuzzy boundary at x: Begin the process by removing from $R(x)$ some subset of the intervals that contain x according to a given, but arbitrary rule (possibly involving further randomness). The process then attempts to pack the remaining intervals of $R(x)$ by the greedy rule, i.e., in order, discarding any that overlap previously packed intervals. If $x \leq y \leq x+1$, and we consider the process $\mathcal{P}_{x,y}$ that removes from $R(x)$ all intervals not entirely contained in $[0,y]$, then we obtain the regular (non-fuzzy) packing process, but on $[0,y]$. For a given fuzzy process $\mathcal{P}$, let $N_{\mathcal{P}}$ be the random variable that counts the number of intervals in $R(x)$ that actually get packed in $[0,x+1]$ by $\mathcal{P}$. (Note that, since $d=1$, $N_{\mathcal{P}} \equiv N_{\mathcal{P}}(t)$ is the same as $L_{\mathcal{P},x}(t)$ from the previous section.) We define $N_x = N_{\mathcal{P}_{x,x}}$ to be the number of intervals in $R(x)$ that get packed by the regular process $\mathcal{P}_{x,x}$ in $[0,x]$ that considers only intervals entirely contained in $[0,x]$.

Lemma 18 *Let $\mathcal{P}$ be a packing process on $[0,x]$ with fuzzy boundary at x. Then $N_x \leq N_{\mathcal{P}} \leq N_x + 1$.*

Remark: It is easy to construct a counterexample showing that this result does not hold for $\mathrm{E}L_x$ when $d > 1$. ∎

Proof: We prove the result for all x simultaneously, by induction on $\#R(x)$, the cardinality of $R(x)$. For a sample path in which $\#R(x) = 0$, we have $N_x = N_{\mathcal{P}} = 0$. More generally, if $\mathcal{P}$ packs no intervals at all, then neither does $\mathcal{P}_{x,x}$, so $N_x = N_{\mathcal{P}} = 0$. Otherwise consider the first interval $[y-1,y]$ packed by $\mathcal{P}$. If it is contained in $[0,x]$, it is also the first interval packed by $\mathcal{P}_{x,x}$, and from then on, the processes are identical to the left of $[y-1,y]$ and the numbers of intervals packed to the right of $[y-1,y]$ by the two process can be compared by the inductive hypothesis (with x replaced by $x-y$ and with $R(x)$ replaced by the subcollection of intervals entirely contained within $[y,x+1]$, all shifted y to the left).

It remains to consider the case in which the first interval $[y-1,y]$ packed by $\mathcal{P}$ straddles x (and hence is not packed by $\mathcal{P}_{x,x}$). From then on, the number of intervals packed by $\mathcal{P}$ equals the number of intervals N_{y-1} in $R((y-1)+1)$ packed by $\mathcal{P}_{y-1,y-1}$, bringing the total to $N_{\mathcal{P}} = N_{y-1}+1$. On the other hand, N_x also equals $N_{\mathcal{P}_{y-1,x}}$. The inductive hypothesis gives $N_{y-1} \leq N_{\mathcal{P}_{y-1,x}} \leq N_{y-1}+1$, which is equivalent to $N_{\mathcal{P}_{y-1,x}} \leq N_{y-1}+1 \leq N_{\mathcal{P}_{y-1,x}}+1$, which is $N_x \leq N_{\mathcal{P}} \leq N_x+1$. ∎

Corollary 19 *As a function of $x \geq 0$, $K(t,x)$ is nondecreasing.*

Proof: It suffices to prove $K(t,x) \leq K(t,y)$ when $0 \leq x \leq y \leq x+1$. Applying the left half of Lemma 18 with $\mathcal{P} = \mathcal{P}_{x,y}$ shows that we have the desired

λ	$\alpha(\lambda)$	a	$\hat{P}(a,\lambda)$
2	.593	5	.599
		10	.597
		20	.596
		40	.595
		80	.595
10	.716	5	.735
		10	.723
		20	.720
		40	.718
		80	.718
50	.741	5	.837
		10	.791
		20	.758
		40	.745
		80	.744

Table 1: Reservation probabilities for $\mathcal{A}$ uniform on $[0, a]$ and unit intervals. The $\hat{P}(a, \lambda)$ are values from simulations.

inequality sample by sample. ∎

Remark: As in the preceding section, one can consider a fuzzy boundary at 0 instead of x, and fuzzy boundaries at both 0 and x. By using the same combinatorial arguments as in Lemma 18 and adapting the notation in the obvious way, one can easily obtain the respective inequalities $N_x \leq N_{\mathcal{P}} \leq N_x + 1$, and $N_x \leq N_{\mathcal{P}} \leq N_x + 2$ for these cases. ∎

5 Final Remarks

There are several avenues for further research in reservation systems. Expanding the class of models leading to exact results, or even good estimates, is one such avenue. However, even incremental extensions to the models of Sections 2 can lead to difficult problems. We give just two examples where this seems to be the case. First, suppose we extend the slotted system to include reservation intervals covering 2 slots as well as 1. The analysis must then account for the fact that patterns of reservations in one period of time can effect the patterns in much later periods. One way to avoid this dependence is to impose an 'alignment' on the length-2 intervals; e.g., suppose requests are such that length-2 reservation intervals begin only at even integer times. Then the reservation patterns in $[i, i+2]$, i even, are independent of the reservation patterns elsewhere.

Assuming $k = 2$ copies of the resource is another important incremental extension. With multiple copies, the protocol is as given in Section 2: a request is denied if and only if, at some time, its reservation interval overlaps two currently reserved intervals. We have seen that this extension is easy to handle in the slotted model of Section 2, but it is not nearly so easy in the other models.

Further investigation into other types of asymptotic behavior may be more rewarding. A central limit theorem for 'wasted' time (the time the resource goes unreserved) may be an especially worthwhile goal in view of a similar result for the interval packing problem [3], which extended a central limit theorem by Dvoretzky and Robbins [5] on the parking problem. Asymptotic results can often be useful approximations over wide ranges of parameter values. Table 1 illustrates this fact via the limit law of Section 5 for advance notice distributions uniform over $[0, a]$ and unit interval lengths. The numbers suggest that, for λ moderately small, $\alpha(\lambda)$ approximates the true reservation probability to within a few percent even for small a. As λ becomes large, a good approximation relies on a becoming large as well.

The reservation probability is a relatively simple, system oriented performance measure analogous to the probability of a busy server in single-server queueing theory. Customer oriented performance measures would also be of interest, prime examples being the conditional probabilities that a request is accepted given the advance notice or given the duration of the interval requested.

Algorithmics is yet another obvious avenue of research. The greedy rule is a theme of this paper, but in general it is unlikely to be an optimal policy in the sense of maximizing the reservation probability. As a simple example, consider unit reservation intervals and advance notices that are uniform over some large interval. An intriguing question is: When, if ever, is it better to reject a request even though the interval it requests is available?

Acknowledgments The authors are pleased to acknowledge many helpful discussions with L. Flatto, E. N. Gilbert, and M. Hofri.

REFERENCES

[1] E. M. Arkin and E. B. Silverberg. Scheduling jobs with fixed start and end times. *Discrete Appl. Math.*, 18:1–8, 1987.

[2] N. H. Bingham, C. M. Goldie, and J. L. Teugels. *Regular Variation.* Cambridge University Press, Cambridge, 1987.

[3] E. G. Coffman, Jr., Leopold Flatto, Predrag Jelenković, and Bjorn Poonen. Packing random intervals on-line. *Algorithmica* (to appear). Bell Labs, Lucent Technologies, Murray Hill, NJ 07974, 1997.

[4] M. Degermark, T. Köhler, S. Pink, and O. Schelén. Advance reservations for predictive service. In *NOSSDAV'95*, pages 3–15, 1995.

[5] A. Dvoretzky and H. Robbins. On the "parking" problem. *MTA Mat. Kut. Int. Kžl.*, 9:209–225, 196 4.

[6] D. Ferrari, A. Gupta, and G. Ventre. Distributed advance reservations for real time con nections. In *NOSSDAV'95*, pages 18–27, 1995.

[7] Juan Garay, Inder Gopal, Shay Kutten, and Yishay Mansour. Efficient on-line call control algorithms. *J. Algorithms*, 23:180-194, 1997.

[8] A. Greenberg, R. Srikant, and W. Whitt. Resource sharing for book-ahead and instantaneous-request calls. In *ITC 15*, pages 539–548, Washington, D.C., USA, June 1997.

[9] J. J. Harms and J. W. Wong. Performance modeling of a channel reservation service. *Computer Net. ISDN Sys.*, 27:1487–1497, 1995.

[10] K-Q. Liao and J. W. Roberts. Videoconference traffic and network design. IEEE Trans. Comm., COM-35:275-282, 1997.

[11] R. J. Lipton and Andrew Tompkins. On-line interval scheduling. *Proc. 5th Ann. ACM-SIAM Symp. Disc. Alg.*, 302-311, 1994.

[12] P. Ney. A random interval filling problem. *Annals of Math. Stat.*, 33(33):702–718, 1962.

[13] A. Rényi. On a one-dimensional random space-filling problem. *MTA Mat. Kut. Int. Kžl.*, 3:109–127, 195 8.

[14] J. T. Virtamo. A model of reservation systems. *IEEE Trans. Comm. Sys.*, 40:109–118, 1992.

[15] L. C. Wolf, L. Delgrossi, R. Steinmetz, S. Schaller, and H. Witting. Issues of reserving resources in advance. In *NOSSDAV'95*, pages 28–38, 1995.