COMMUN. STATIST.—STOCHASTIC MODELS, 12(2), 215–243 (1996)

STOCHASTIC LIMIT LAWS FOR SCHEDULE MAKESPANS

E. G. COFFMAN, Jr.,[1] Leopold FLATTO[2] and Ward WHITT[3]
AT&T Bell Laboratories
Murray Hill, New Jersey 07974-0636

ABSTRACT

A basic multiprocessor version of the makespan scheduling problem requires that n tasks be scheduled on m identical processors so as to minimize the latest task finishing time. In the standard probability model considered here, the task durations are i.i.d. random variables with a general distribution F having finite mean. Our main objective is to estimate the distribution of the makespan as a function of m, n, and F, under the on-line greedy policy, i.e., where the tasks are put in sequence and assigned in order to processors whenever they become idle. Because of the difficulty of exact analysis, we concentrate on the asymptotic behavior as $n \to \infty$ or as both $m \to \infty$ and $n \to \infty$ with $m \leq n$. The focal point is the Markov chain giving the remaining processing times of the $m-1$ tasks still running at task completion epochs. The theory of stationary marked point processes is used to show that the stationary distribution of this Markov chain coincides with the order statistics of $m-1$ independent random variables having the equilibrium residual-life distribution associated with F. Convergence theory for general-state Markov chains is then applied to establish convergence results for the Markov chain of interest. Finally, central limit theorems are applied to show that what we can gain from a good list scheduling policy is asymptotically negligible compared to our degree of uncertainty about the makespan (i.e., its standard deviation).

Key Words: scheduling, multiprocessor scheduling, makespan, Markov chains, Harris-recurrent Markov chains, rates of convergence, small sets, stationary marked point processes, superposition

[1]Room 2D-155; email: egc@research.att.com
[2]Room 2C-359; email: leopold@research.att.com
[3]Room 2C-178; email: wow@research.att.com

1. Introduction

An integer $m \geq 2$ together with positive task running times $T_1, \ldots, T_n$ defines an instance of the *multiprocessor scheduling problem*: Schedule $T_1, \ldots, T_n$ on m identical processors $P_1, \ldots, P_m$ so as to minimize the latest task finishing time or *makespan*; i.e., partition the set $\{T_1, \ldots, T_n\}$ into subsets $P_1, \ldots, P_m$ so as to minimize the maximum subset sum

$$L_{m,n} \equiv \max_{1 \leq j \leq m} \sum_{\{i : T_i \in P_j\}} T_i \ .$$

To avoid trivialities, we assume that $n \geq m$ unless stated otherwise. This NP-complete problem finds application in operations research as a model of scheduling parallel machines in industrial job shops (see Lawler, et al. [19] and Blazewicz, et al. [4]). It has also had a prominent role in computer science, where the term multiprocessor originates. Along with a number of other fundamental NP-complete problems, it has served as a theoretical testbed for the development of new ideas in the design and analysis of algorithms (see Garey and Johnson [16]).

For the purposes of defining heuristic policies, it is convenient to assume that the tasks are presented in the form of a list $(T_1, \ldots, T_n)$. The on-line *greedy* policy is arguably the simplest (and fastest) heuristic for finding approximate solutions to the multiprocessor scheduling problem. This policy uses no advance information on the number or durations of tasks. The policy begins by assigning the first m tasks $T_1, \ldots, T_m$ to the m processors $P_1, \ldots, P_m$; the processors start running these tasks at time 0, while the remaining tasks wait. Thereafter, whenever a processor finishes its current task, the next waiting task, if any, is assigned to the idle processor. In queueing terminology the system operates as an m-server queue with a first-come first-served service discipline; n customers arrive to an empty system at time 0, and the latest of their departure times is the makespan. For our purposes the rule for resolving ties among processors is immaterial, so we leave it unspecified.

Understandably, the greedy policy was among the first policies studied when the probabilistic analysis of scheduling algorithms began some 15 years ago. In the standard probability model considered here, the task durations T_i are independent and identically distributed (i.i.d.) with a distribution $F(t) = P(T_i \leq t)$ having a finite first moment $E[T]$, where T denotes a generic task duration. The problem is to find the distribution of the makespan $L_{m,n}$ as a function of the number m of processors, the number n of tasks and the distribution F. The general aim is to bring out typical behavior rather than the highly unlikely worst-case behavior brought out by deterministic analysis. With explicit formulas in mind, probabilistic analysis is usually quite difficult, so research has often turned to large-n asymptotics.

In the past decade, several papers have been devoted to an asymptotic analysis of the greedy policy for general m; see Boxma [6], Bruno and Downey [7], Coffman and Gilbert [9], Han, Hong, and Leung [18] and Loulou [21] (the monograph by Coffman and Lueker [10] gives a general coverage). None of this work has led to limiting behavior as precise as that found for $m = 2$ in an early result of Feller [13], p. 208; instead, for $m \geq 3$, the analysis has resorted to various bounding techniques. However, Feller's result can in fact be generalized, as shown in Section 2.

The contributions of Sections 2–5 are outlined as follows. Fix $m \geq 2$, and extend the greedy process to the infinite time horizon; i.e., construct a greedy schedule from the infinite sequence $T_1, T_2, \ldots$. For $n \geq 1$, let C_n denote the n^{th} completion time; let R_i^n, $1 \leq i \leq m-1$, denote the residual times of those tasks still running at time C_n, ordered by increasing processor index; and let $X_1^n \leq X_2^n \leq \cdots \leq X_{m-1}^n$ denote the order statistics of the R_i^n. C_0 is defined to be 0, and $X^0 = (X_1^0, \ldots, X_{m-1}^0)$ gives the order statistics of $T_1, \ldots, T_{m-1}$. Figure 1 illustrates the definitions. Sections 2–4 study the Markov chain $\{X^n\}_{n \geq 0}$, $X^n = (X_1^n, \ldots, X_{m-1}^n)$, for a given initial state X^0. By an application of the theory of marked point processes, Section 2 identifies the

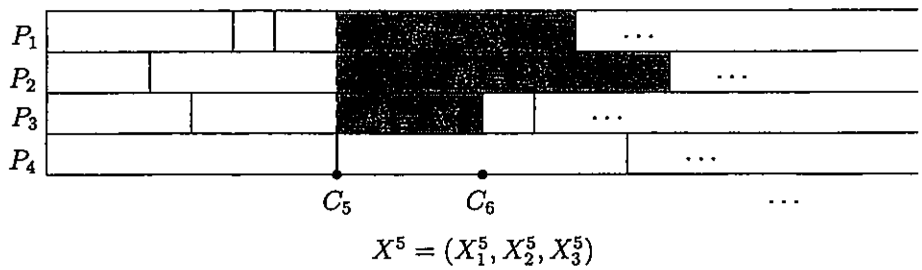

Figure 1: Illustration of definitions; $m = 4$.

invariant measure of $\{X^n\}$, thus extending Feller's result for $m = 2$ to general m. This result is then applied to asymptotics of the expected "error" $E[\alpha_{m,n}]$, where $\alpha_{m,n} = L_{m,n} - \frac{n}{m}E(T)$, and illustrated for specific distributions. Note that $\frac{n}{m}E[T]$ is an obvious lower bound to $E[L_{m,n}]$. Section 2 concludes with asymptotics in m that are derived as iterated limits $n \to \infty$, then $m \to \infty$. The results in Section 2 were reported informally in a recent survey by Coffman and Whitt [11].

The convergence issues related to the invariant measure of $\{X^n\}$ are examined in depth in Sections 3 and 4 under a broad class C of distributions F. The generality of C requires a comprehensive convergence theory of Markov chains; this theory is briefly reviewed in Section 3, using the recent text of Meyn and Tweedie [22] as the principal reference. The foundations presented in Section 3 are then used in Section 4 to prove several limit laws for the chain $\{X^n\}$.

Section 5 shifts to limit laws provided by central limit theorems, tools that apply naturally to the asymptotic analysis of schedule makespans. These tools are used in a policy-free formulation, i.e., limit theorems are proved which apply simultaneously to all scheduling policies. The paper concludes with final remarks in Section 6.

2. Asymptotics of the Greedy Policy

We begin by expressing $L_{m,n}$ in terms of the T_i and the residual running times X_i^n. Returning to Fig. 1, we see that, at time C_{n-m+1}, tasks $T_1, \dots, T_n$ have all started, $n-m+1$ of them have finished, and $m-1$ are still running with the ordered residual times X_i^{n-m+1}, $1 \le i \le m-1$. Then the tasks $T_1, \dots, T_n$ have a latest finishing time

$$L_{m,n} = C_{n-m+1} + X_{m-1}^{n-m+1} , \tag{2.1}$$

and a sum of running times that can be expressed as

$$\sum_{i=1}^{n} T_i = mC_{n-m+1} + \sum_{i=1}^{m-1} X_i^{n-m+1} . \tag{2.2}$$

Combining (2.1) and (2.2) gives

$$L_{m,n} = \frac{1}{m}\left[\sum_{i=1}^{n} T_i + mX_{m-1}^{n-m+1} - \sum_{i=1}^{m-1} X_i^{n-m+1}\right] . \tag{2.3}$$

To proceed, we need information on the random variables X_i^n. A direct approach analyzes the Markov chain $\{X^n\}$ using standard methods. To this end, assume for simplicity that F is absolutely continuous with a density f; we will return to more general F later. Under this condition, it is easy to verify that, for each n, the distribution of X^n will have a continuous density; we denote this density by $\pi_n(\mathbf{x})$, $\mathbf{x} = (x_1, \dots, x_{m-1})$, and let $\pi(\mathbf{x}) = \lim_{n\to\infty} \pi_n(\mathbf{x})$ be the stationary density, assuming that it exists.

Figure 2 illustrates the possible one-step transitions from a state $\mathbf{x}'$ to a state $\mathbf{x}$ resulting from the assignment of a new task of duration t. These transitions have the following general forms:

$$(t + x_1, \dots, t + x_{m-1}) \to (x_1, \dots, x_{m-1}), \text{ if } 0 < t < x_1' ,$$
$$(t - x_1, t - x_1 + x_2, \dots, t - x_1 + x_{m-1}) \to (x_1, \dots, x_{m-1}), \text{ if } x_1' < t < x_2' ,$$
$$(t - x_i, t - x_i + x_1, \dots, t - x_i + x_{i-1}, t - x_i + x_{i+1}, \dots, t - x_i + x_{m-1}) \to (x_1, \dots, x_m),$$
$$\text{if } x_i' < t < x_{i+1}' \text{ and } 2 \le i \le m-2 ,$$
$$(t - x_{m-1}, t - x_{m-1} + x_1, \dots, t - x_{m-1} + x_{m-2}) \to (x_1, \dots, x_{m-1}), \text{ if } t > x_{m-1}' .$$

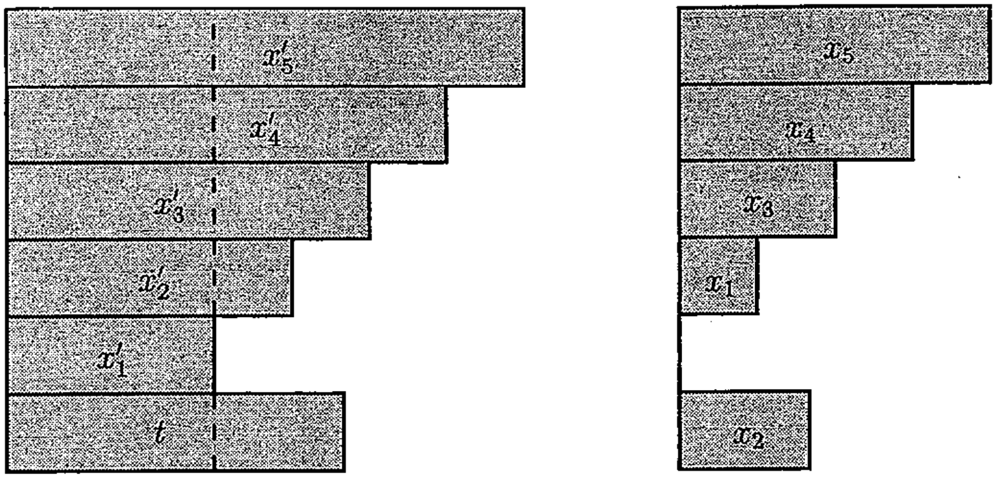

Figure 2: $m = 6$, $x'_2 < t < x'_3$. (In this and in any subsequent figures where the ordering of the processors is immaterial. the ordering is left unspecified.)

Thus, it is convenient to define

$$\begin{aligned} \mathbf{x}_0(y) &= (x_1 + y, \ldots, x_{m-1} + y) \\ \mathbf{x}_i(y) &= (y, x_1 + y, \ldots, x_{i-1} + y, x_{i+1} + y, \ldots, x_{m-1} + y), \quad 1 \le i \le m-1 \; , \end{aligned}$$

The state $\mathbf{x}$ can be obtained in one step from any state $\mathbf{x}_i(y)$, $0 \le i \le m-1$, provided the new task has duration $x_i + y$, where x_0 is defined to be 0. Hence,

$$\pi_{n+1}(\mathbf{x}) = \sum_{i=0}^{m-1} \int_0^\infty \pi_n(\mathbf{x}_i(y)) f(x_i + y) dy \; . \tag{2.4}$$

Sections 3 and 4 will prove convergence of π_n to a proper limit π independent of the starting state under conditions that properly include the stated density condition. In terms of random variables, we can write

$$(X_1^n, \ldots, X_{m-1}^n) \Rightarrow (X_1^*, \ldots, X_{m-1}^*) \quad \text{as} \quad n \to \infty \; , \tag{2.5}$$

where $\Rightarrow$ denotes convergence in distribution. Indeed, we prove in Section 4 a convergence of the probability measures in total variation as $n \to \infty$.

To obtain an explicit formula for the limiting distribution π, i.e., the distribution of $(X_1^*, \ldots, X_{m-1}^*)$ in (2.5), the Markov chain analysis now requires that we solve the stationary version of the rather awkward recurrence in (2.4). A key observation allows us to side-step this difficulty by applying the theory of stationary marked point processes, in an argument that makes no direct use of properties already established by the Markov chain approach. The observation is that the sequence $\{C_n\}$ generated by the greedy rule is equal stochastically to the superposition of m i.i.d. ordinary renewal processes; i.e., each is defined independently by F and each starts with a point at 0. The time-stationary version of each renewal process is the familiar equilibrium renewal process, in which the distance to the first point has the equilibrium residual-life distribution G with density $g(t) = [1 - F(t)]/E[T]$. Each original renewal process is the Palm (or synchronous) version of its time-stationary version.

Now consider the superposition of m i.i.d. copies of this time-stationary renewal process. This is a time-stationary point process with the distance to the first point from each component stream having distribution G. Since we want to look at the superposition process at completion times, we are interested in the Palm (synchronous) version of this stationary point process. Section 5.1 of Baccelli and Brémaud [2] characterizes this Palm version in terms of the Palm and stationary versions of the component processes. However, from this superposition process alone we cannot extract the stationary distribution of $(X_1^n, \ldots, X_{m-1}^n)$ directly. To do this, we mark the points of each component stream with the index of the processor on which it occurs, and then apply the corresponding superposition result for stationary marked point processes in Section 1.3.5 of Franken et al. [14]. This result shows that each of the $m - 1$ streams is equally likely to produce the current point (i.e., the m possible marks of the current point are equally likely), and that the residual time to the next point in each of the remaining $m-1$ streams has the distribution G. It follows that a stationary version of $(X_1^n, \ldots, X_{m-1}^n)$ at completion times coincides with the order statistics of $m-1$ i.i.d. random variables with distribution G. Thus, we have proved the following result (without any density assumption).

Theorem 1. *If F is any distribution with finite mean, then a stationary (invariant) distribution of $\{X^n\}$ is given by*

$$\pi(\mathbf{x}) = (m-1)! \prod_{i=1}^{m-1} g(x_i) \ . \tag{2.6}$$

For an alternative proof of Theorem 1, we can start by assuming that F has a density f and verify directly that (2.6) satisfies the stationary version of (2.4), with π_{n+1} and π_n replaced by π. If we remove constants, this is the same as verifying that $\widehat{\pi}$ solves (2.4), where

$$\widehat{\pi}(x_1, \ldots, x_{m-1}) = \prod_{i=1}^{m-1} [1 - F(x_i)] \ .$$

To show this, define $J(y) = \prod_{i=0}^{m-1} [1 - F(x_i + y)]$, where $x_0 = 0$. Then

$$J'(y) = -\sum_{i=0}^{m-1} \widehat{\pi}(\mathbf{x}(y)) f(x_i + y) \ ,$$

and hence

$$\begin{aligned} \sum_{i=0}^{m-1} \int_0^\infty \widehat{\pi}(\mathbf{x}_i(y)) f(x_i + y) dy &= \int_0^\infty \sum_{i=0}^{m-1} \widehat{\pi}(\mathbf{x}_i(y)) f(x_i + y) dy \\ &= -\int_0^\infty J'(y) dy = J(0) - J(\infty) \ . \end{aligned}$$

But $J(0) = \widehat{\pi}(x_1, \ldots, x_{m-1})$ and $J(\infty) = 0$, so we get

$$\sum_{i=0}^{m-1} \int_0^\infty \widehat{\pi}(\mathbf{x}_i(y)) f(x_i + y) dy = \widehat{\pi}(x_1, \ldots, x_{m-1}) \ ,$$

as desired.

Finally, we treat a general F with finite mean by representing it as a limit of distributions F_n, each having a density f_n and finite mean $E[T_n]$. In particular, we construct F_n so that, as $n \to \infty$, $F_n(t) \to F(t)$ at each continuity point of F

and $E[T_n] \to E[T]$; such constructions are always possible. Then as $n \to \infty$, the associated invariant densities π^n converge pointwise almost everywhere with respect to Lebesgue measure on $\mathbb{R}^{m-1}$ to the candidate invariant density π. By Scheffé's theorem, p. 224 of Billingsley [3], the associated invariant distributions on $\mathbb{R}^{m-1}$ also converge in total variation and thus weakly. To conclude the argument we need a continuity property of the sequence of Markov chains initialized by their invariant distribution. For each n, the Markov chain is stationary, so it suffices to show that the distributions converge after one transition, and for this, it suffices to show that $P^{(n)}(x_n, \cdot) \Rightarrow P(x, \cdot)$ whenever $x_n \to x$, where $\Rightarrow$ denotes weak convergence of probability measures and $P^{(n)}(x, B)$ is the transition kernel of chain n. But this desired continuity property is easily proved, as shown in Lemma 1 of Section 4.

It is important to note that, in general, the invariant distribution in (2.4) need not be unique. For example, this is the case with deterministic running times. We establish uniqueness and convergence in Section 4 under an extra condition. If uniqueness and convergence do hold, then by (2.3), (2.5), and Theorem 1, we have the limit

$$L_{m,n} - \frac{1}{m}\sum_{i=1}^{n} T_i \Rightarrow \alpha_m \equiv X^*_{m-1} - \frac{1}{m}\sum_{i=1}^{m-1} X^*_i \quad \text{as} \quad n \to \infty\ , \tag{2.7}$$

where $X^*_1, \ldots, X^*_{m-1}$ are the order statistics of the i.i.d. random variables $R^*_1, \ldots, R^*_{m-1}$ with distribution G. Hence, in (2.7) we may replace $\sum_{i=1}^{m-1} X^*_i$ by $\sum_{i=1}^{m-1} R^*_i$.

We now consider expected values. We assume from now on that $E[T^2] < \infty$. Then the residual-life distribution G can be shown to have the mean $(\sigma^2(T) + (E[T])^2)/2E[T] = E[T](v^2(T) + 1)/2$ where $\sigma(T)$ is the standard deviation of F and $v(T) = \sigma(T)/E[T]$ is the coefficient of variation of F. From (2.3) and (2.7) we obtain

$$E(\alpha_{m,n}) = E(X^{n-m+1}_{m-1}) - \frac{1}{m}\sum_{i=1}^{m-1} E(X^{n-m+1}_i) \tag{2.8}$$

and from (2.7),

$$\begin{aligned} E(\alpha_m) &= E(X^*_{m-1}) - \frac{1}{m}\sum_{i=1}^{m-1} E(X^*_i) = E(X^*_{m-1}) - \frac{1}{m}\sum_{i=1}^{m-1} E(R^*_i) \\ &= \int_0^\infty [1-G^{m-1}(x)]dx - \left(\frac{m-1}{m}\right)\left(\frac{v^2(T)+1}{2}\right) E(T) \ . \end{aligned} \tag{2.9}$$

These observations yield the following corollary to Theorem 1.

Corollary 1. *If* $\lim_{n\to\infty} E(X^n_i) = E(X^*_i)$ *for* $1 \le i \le m-1$, *then*

$$E(\alpha_{m,n}) = \int_0^\infty [1-G^{m-1}(x)]dx - \left(\frac{m-1}{m}\right)\left(\frac{v^2(T)+1}{2}\right) E(T) + o(1) \quad as\ n\to\infty \ . \tag{2.10}$$

In Section 4, we prove that $E(X^n_i)$ approaches $E(X^*_i)$ exponentially fast, provided F has an exponential tail. Hence, for these cases, the term $o(1)$ in (2.10) can be improved to $O(\rho^n)$ for some $0 < \rho < 1$.

The integral appearing on the right side of (2.9) can seldom be evaluated in closed form. We mention two cases where it can.

(i) Consider the uniform distribution $F(t) = t$, $0 \le t \le 1$. Then $g(t) = 2(1-t)$, $0 \le t \le 1$. From (2.9) we obtain

$$E(\alpha_m) = \int_0^1 [1-(2t-t^2)^{m-1}]dt - \frac{m-1}{3m} \ . \tag{2.11}$$

After the change of variables $u = 2t - t^2$, the integral in (2.11) can be evaluated in closed form. We find that

$$E(\alpha_m) = \frac{2m+1}{3m} - \frac{2^{2m-2}}{2m-1} \Big/ \binom{2m-2}{m-1} \ . \tag{2.12}$$

(ii) Suppose we have the exponential distribution $F(t) = 1 - e^{-\kappa t}$, $t \ge 0$, for a given rate parameter $\kappa > 0$. In this case the C_n, $n \ge m$, are the epochs of a Poisson process at rate $m\kappa$, so that for all i and $n \ge m$, the X^n_i are the order statistics of $m-1$ i.i.d. random variables with the distribution $G = F$. Then (2.3) and (2.7) give

$$E[\alpha_{m,n}] = \frac{1}{\kappa}[H_m - 1] \quad \text{for all } m \text{ and } n \ge m \ , \tag{2.13}$$

where $H_m = \sum_{j=1}^{m} 1/j$ (see also Coffman and Gilbert [9]).

Because of the increased use of massively parallel computers, it is natural to consider asymptotics as $m \to \infty$. From expressions like (2.12) and (2.13), large-m asymptotics for the mean of the time-stationary random variable α_m can be obtained directly. For example, (2.12) and Stirling's formula give

$$E[\alpha_m] = \frac{2}{3} - \frac{\sqrt{\pi}}{2\sqrt{m}} + O\left(\frac{1}{m^{3/2}}\right) \quad \text{as} \quad m \to \infty \tag{2.14}$$

when F is the uniform distribution on $[0, 1]$; similarly, when F is the exponential distribution, (2.13) and asymptotics for H_m give

$$E[\alpha_m] = \frac{1}{\kappa}(\ln m - 1 + \gamma) + O\left(\frac{1}{m}\right) \quad \text{as} \quad m \to \infty \, , \tag{2.15}$$

where γ is Euler's constant $(0.5772\ldots)$.

More generally, we can obtain asymptotic properties of $E[\alpha_m]$ from (2.10). When F has support $[0, b]$, (2.7) and the strong law of large numbers implies that

$$\alpha_m \to \alpha \equiv b - E[T]\frac{v^2(T) + 1}{2} \quad \text{w.p.1} \quad \text{as} \quad m \to \infty \, . \tag{2.16}$$

From (2.16) we can see how F influences the asymptotic error α. For a given bound b, α *decreases* in $E[T]$ and $v^2(T)$. For given b and $E[T]$, the lowest value of α is $b/2$, which is approached by the two-point distribution with mass $E[T]/b$ on b and mass $(b - E[T])/b$ on 0; e.g., see p. 120 of Whitt [26].

It is interesting that for this extremal two-point distribution the greedy policy is optimal for all m and n; i.e., there is a distribution with finite positive variance for which greedy gives the minimum expected error. The optimality of the greedy policy in this case is trivial to see because the makespan is the same as for a random number of tasks, each with a constant running time b. In this case, all work-conserving policies (in which no processor is idle when there is a task that has not started) are obviously optimal. This two-point distribution is not in the class of absolutely continuous distributions, but it is approached by such distributions.

The term X^*_{m-1} in formula (2.7) for α_m obviously becomes more important when F does not have finite support. The asymptotic behavior of X^*_{m-1} as $m \to \infty$ is described by classical extreme-value theory; see Leadbetter, Lindgren and Rootzén [20] and Reiss [23]. This extreme value theory applies to the iterated limit as first $n \to \infty$ and then $m \to \infty$ provided (2.5) is still valid.

Since the superposition of m i.i.d. renewal processes, appropriately scaled, converges to a Poisson process as $m \to \infty$, see e.g., Çinlar [8], one might expect that the general formula for $E[\alpha_m]$ in (2.9) and (2.10) would in some sense approach the formulas for the exponential distribution in (2.13) and (2.15), but this is *not* the case. For the question here, the superposition limit theorem does not apply. The superposition limit theorem implies that the distribution of X^*_1 is asymptotically exponential as m gets large, but in (2.9) we focus on X^*_{m-1} and $\sum_{i=1}^{m-1} X^*_i$.

An interesting open problem is the joint limiting behavior as $m \to \infty$ and $n \to \infty$. Above, we considered only the iterated limit in which first $n \to \infty$ and then $m \to \infty$. If $m = n$, then the extreme-value theory for i.i.d. random variables with distribution F describes the makespan. It would be interesting to develop different asymptotics in intermediate cases.

3. Convergence of $\{X^n\}$: Preliminaries

In this and the next section, we study the conditions under which the distributions π_n, $n \geq 1$, defined in (2.4) converge to an invariant measure for all initial measures, and the stronger conditions under which this convergence is geometrically fast. (The limit necessarily must be π in (2.6).) For this purpose, we cite various theorems in Meyn and Tweedie [22], hereafter referred to simply as MT, and then find conditions on F guaranteeing that the Markov chain $\{X^n\}$ satisfies these conditions. Roughly speaking, the conditions are of two kinds: A condition of the first kind specifies a property of the chain directly, e.g., irreducibility, aperiodicity, etc. A condition of the second kind invokes the existence of a nonnegative function on the state space with certain desired properties related to the chain. These conditions translate into two kinds of conditions on F. The first kind involves smoothness properties of F, while the second involves the behavior of F at infinity. The remainder of this section briefly reviews the basic theory. The next section applies the theory to establish convergence properties of $\{X^n\}$ under an appropriate condition on F.

Let $\mathbf{X}$ be a general state space and let $\mathcal{B} = \mathcal{B}(\mathbf{X})$ be a σ-field of subsets of $\mathbf{X}$. For the chain $\{X^n\}_{n\geq 0}$ analyzed in the next section, $\mathbf{X}$ will be the set of vectors $\mathbf{x} = (x_1, \ldots, x_{m-1})$ with $0 \leq x_1 \leq \cdots \leq x_{m-1}$, and $\mathcal{B}$ will be the Borel σ-field of sets in $\mathbf{X}$. Let $\{\Phi^n\}_{n\geq 0}$ be a generic Markov chain on $\mathbf{X}$ governed by an initial probability measure μ_0 and a one-step transition probability $P(\mathbf{x}, B)$, $\mathbf{x} \in \mathbf{X}$, $B \in \mathcal{B}$. The n-step transition probabilities are denoted by $P^n(\mathbf{x}, B)$, so that $P(\mathbf{x}, B) = P^1(\mathbf{x}, B)$. We use $P^n_{\mathbf{x}}(\cdot)$ for the probability of events in n-step transitions from an initial state $\mathbf{x}$, i.e., $\mu_0(\Phi^0 = \mathbf{x}) = 1$. A hitting probability is denoted by $H(\mathbf{x}, B)$, the probability that $\Phi^n \in B$ for some $n > 0$, starting in $\mathbf{x}$.

In the sequel, all measures are assumed to be nontrivial, i.e., not identically zero. The standard notation $\|\cdot\|$ will be used for the total variation norm. We continue to use π to denote an invariant measure, i.e., a measure π for which $\pi(A) = \int P(\mathbf{y}, A)\pi(d\mathbf{y})$ for all $A \in \mathcal{B}$. In what follows, we avoid the concept of "petite" sets emphasized by MT as it is not necessary for our work. Instead, we use the more restricted notion of "small" sets as described below; see p. 106 of MT.

Let ν be a measure on $\mathcal{B}$ and let k be a positive integer. Set A is called (ν, k)-*small* (or simply *small* with the existence of ν, k understood) if $P^k(\mathbf{x}, B) \geq \nu(B)$ for all $\mathbf{x} \in A$, and for all $B \in \mathcal{B}$. If $\nu(B) > 0$, then the above condition implies that B can be reached in k steps from anywhere in A with a positive probability that is independent of the starting point. Intuitively, this means that all points $\mathbf{x} \in A$ are not "too far" from B, and in this sense, A is small. For the chain $\{X^n\}$, we will impose a condition on F guaranteeing that all compact sets are small.

The following result applies only to the case when the entire state space $\mathbf{X}$ is small, and is the basis (in Theorem 6) for proving the desired convergence of $\{X^n\}$ under the class of distributions supported on a finite interval, which includes the uniform distributions illustrated in Section 2.

Theorem 2. (Theorem 16.2.4 of MT). *Suppose* $\mathbf{X}$ *is* (ν, k)*-small under a measure* ν *and a positive integer* k. *Then there exists a unique invariant measure* π *such that*

$$\|P^n(\mathbf{x}, \cdot) - \pi\| \leq \theta^n$$

for all $n > 0$, $\mathbf{x} \in \mathbf{X}$, *where* $\theta = [1 - \nu(\mathbf{X})]^{1/k}$.

The norm $\|\cdot\|$ appearing in Theorem 2 is the total variation norm. Note that, since $0 < \nu(\mathbf{X}) \leq P^k(\mathbf{x}, \mathbf{X}) = 1$, we have $0 \leq 1 - \nu(\mathbf{X}) < 1$. Thus, Theorem 2 ensures geometric convergence.

To obtain a similar result in greater generality, we introduce extensions of the familiar notions of irreducibility, aperiodicity, and recurrence for Markov chains on countable state spaces to more general state spaces. First, suppose there exists a measure ψ on $\mathcal{B}$ such that for any set $B \in \mathcal{B}$, $\psi(B) > 0$ implies and is implied by $H(\mathbf{x}, B) > 0$. Then $\{\Phi^n\}$ is said to be *ψ-irreducible.* Thus, the sets of positive ψ measure are precisely those sets that are hit with positive probability from all starting points $\mathbf{x}$. Typically, to prove ψ-irreducibility, one first finds a measure ψ' such that $\psi'(B) > 0$ implies $H(\mathbf{x}, B) > 0$ for all $\mathbf{x} \in \mathbf{X}$. Then one computes a maximal irreducibility measure ψ from ψ' for which the reverse implication also holds (see Section 4.2 of MT). This approach can be used in the proof of the ψ-irreducibility result of Corollary 2 in Section 4.

If in addition to the ψ-irreducibility of $\{\Phi^n\}$, we have that $\psi(B) > 0$ implies

$$P_{\mathbf{x}}(\Phi^n \in B \text{ for infinitely many } n) = 1 \text{ for all } \mathbf{x} \in \mathbf{X} ,$$

then $\{\Phi^n\}$ is called *Harris recurrent*; Section 9.1 of MT. Finally, if $\{\Phi^n\}$ is ψ-irreducible and there exists a set $A \in \mathcal{B}$ such that A is (ν, k)-small for some $k \geq 1$, $\nu(A) > 0$, and such that

$$g.c.d.\{r \geq 1 : \ A \text{ is } (\delta_r \nu, r)\text{-small for some } \delta_r > 0\} = 1 \ ,$$

then $\{\Phi^n\}$ is *aperiodic* (MT, Section 5.4).

In what follows $V(\mathbf{x})$ denotes a nonnegative function on $\mathbf{X}$, to be called a *potential function.* The function ΔV denotes the *drift operator* on V and is defined by (MT, p. 174)

$$\Delta V = \int P(\mathbf{x}, d\mathbf{y}) V(\mathbf{y}) - V(\mathbf{x}) \ .$$

The function V is said to be *unbounded off small sets* if $\{\mathbf{x} : V(\mathbf{x}) > r\}$ is small for all $r \geq 0$. For the chain $\{X^n\}$, we will verify the above condition by demanding that $V(\mathbf{x})$ be continuous for all $\mathbf{x} \in \mathbf{X}$ and that $\lim\limits_{|\mathbf{x}| \to \infty} V(\mathbf{x}) = \infty$, where $|\mathbf{x}| = \max(|x_1|, \ldots, |x_{m-1}|)$. These conditions imply that $\{\mathbf{x} : V(\mathbf{x}) \leq r\}$ is compact, and our condition on F implies that compact sets are small.

We conclude this section with two theorems that also apply when $\mathbf{X}$ is not small; the first gives conditions for convergence, and the second gives conditions for geometric convergence.

Theorem 3. (Theorems 9.18 and 13.33 in MT). *Let $\{\Phi^n\}$ be ψ-irreducible and suppose there exists a potential function $V(\mathbf{x})$ such that $\Delta V \leq 0$, except possibly on some small set, and such that V is unbounded off small sets. Then $\{\Phi^n\}$ is Harris recurrent. If in addition, $\{\Phi^n\}$ is aperiodic and has an invariant measure π, then*

$$\left\| \int \mu_0(d\mathbf{x}) P^n(\mathbf{x}, \cdot) - \pi \right\| \to 0 \quad \text{as} \quad n \to \infty$$

for all initial probability measures μ_0 on $\mathbf{X}$.

Theorem 4. (Theorem 15.01 in MT). *Let $\{\Phi^n\}$ be ψ-irreducible and aperi-*

odic, and suppose there exist a potential function $V \geq 1$, constants $c, \beta > 0$ and a small set A with $\Delta V \leq -\beta V + c \cdot 1_A$ for all $\mathbf{x} \in \mathbf{X}$. Then for some $\eta > 0$ and ρ, $0 < \rho < 1$, we have that

$$\sup_{|\varphi| \leq V} \left| \int P^n(\mathbf{x}, d\mathbf{y})\varphi(y) - \int \pi(d\mathbf{y})\varphi(y) \right| \leq \eta V(\mathbf{x})\rho^n \ .$$

In particular, for φ the indicator function of some set B, $|P^n(\mathbf{x}, B) - \pi(B)| \leq \eta V(x)\rho^n$, so that $\|P^n(\mathbf{x}, \cdot) - \pi\| \leq 2\eta V(x)\rho^n$.

We remark that, as a trivial consequence of Theorems 2, 3 and 4, the invariant measures in these theorems are unique.

4. Convergence Theorems for Greedy Schedules

We first define a class of distribution F and then apply the results of Section 3 to the chain $\{X^n\}$, $X^n = (X_1^n, \ldots, X_{m-1}^n)$.

Definition of Class $\mathcal{C}$. *Let $\mathcal{C}$ denote the class of distributions F for which there exists a "positive-density" interval $[a, b]$, $0 < a < b$, on which $F(x)$ has a continuous, strictly positive derivative $F'(x)$. For $F \in \mathcal{C}$ and a, b given, let $\xi \equiv \xi(F) > 0$ be such that $F'(x) > \xi$, $a \leq x \leq b$.*

Remark. The class $\mathcal{C}$ seems natural and general. It is slightly smaller than the class of *spread-out* distributions, which in turn is slightly smaller than the class of *nonlattice* distributions; see p. 140 of Asmussen [1]. The results here can easily be extended to spread-out distributions, because F is spread-out if and only if the m-fold convolution F^{m*} belongs to $\mathcal{C}$ for some m. ■

We first show that, if $F \in \mathcal{C}$, then $\{X^n\}$ has the desired properties, viz., those of the first kind mentioned at the beginning of Section 3. To do this we need a little more notation.

Let $Q_t = \{\mathbf{x} = (x_1, \ldots, x_{m-1}) : 0 \leq x_1 \leq \cdots \leq x_{m-1} \leq t\} = \{\mathbf{x} \in \mathbf{X} : |\mathbf{x}| \leq t\}$, so that $Q \equiv Q_\infty$ is the entire state space $\mathbf{X}$. For any c, $a < c \leq b$, we define $A_c = \{\mathbf{x} : a \leq x_1 \leq \cdots \leq x_{m-1} \leq c\}$. Finally, in what follows, λ denotes $(m-1)$-dimensional Lebesgue measure, and k always denotes a positive integer.

Theorem 5. *Let $F \in \mathcal{C}$ with positive-density interval $[a, b]$. For any fixed c, $a < c < b$, there exists a function $N_t > 0$ defined for $t > 0$, and a function $D(n, t) > 0$ defined for $n > N_t$, $t > 0$, such that*

$$P^n(\mathbf{x}, B) \geq D(n,t)\lambda(B) \quad \textit{for} \quad t > 0,\ n > N_t,\ \mathbf{x} \in Q_t,\ B \subseteq A_c\ .$$

The following is a direct consequence of Theorem 5; we omit the proof as it amounts to a straightforward checking of definitions.

Corollary 2. *If $F \in \mathcal{C}$, then $\{X^n\}$ is ψ-irreducible and aperiodic. In addition, compact sets are small.*

We break up the proof of Theorem 5 into two parts. After stating a routine continuity result, the first part computes a lower bound on the probability of transitions from a state $\mathbf{x}$ to a neighborhood of the origin. Then the second part lower bounds the probability of transitions from a neighborhood of the origin to subsets of A_c, these transitions being effected by running $m-1$ additional tasks.

The first result follows from the observations made in deriving (2.4) along with a trivial induction argument. Proof details are omitted.

Lemma 1. *For each n, X^n is a continuous (deterministic) function of the variables*
$x_1, \ldots, x_{m-1}, t_1, \ldots, t_n$, where $(x_1, \ldots, x_{m-1})$ is a given initial state and $t_1, \ldots, t_n$ is a given sequence of the task running times.

Lemma 2. *With c fixed as in Theorem 5, let $K_t = (m+1)\max\left\{\frac{am}{c-a}, \frac{t}{c-a}\right\}$. Then there exists a $\delta_{\epsilon,t,k} > 0$ such that $P^k(\mathbf{x}, Q_\epsilon) > \delta_{\epsilon,t,k}$ for $\epsilon, t > 0$, $k > K_t$, $\mathbf{x} \in Q_t$.*

Proof. It is convenient to consider $P^k(\mathbf{x}, Q_\epsilon)$ as a function of $T_1, \ldots, T_k$, i.e., $P^k(\mathbf{x}, Q_\epsilon) = P(X^k(\mathbf{x}; T_1, \ldots, T_k) \in Q_\epsilon)$. We first produce for given $t > 0$, $k > K_t$, $\mathbf{x} \in Q_t$ a deterministic list of running times $t_1, \ldots, t_k$, depending on $\mathbf{x}$, which are contained in $[a, c]$ and lead to $X^k(\mathbf{x}; t_1, \ldots, t_k) = 0$. Then we perturb the t_i's to obtain the lemma.

Let $k = mq + r$, where q, r are integers with $0 \le r \le m - 1$. Consider k tasks with running times $\frac{qc}{q+r}$, $c - \frac{x_1}{q}, \ldots, c - \frac{x_{m-1}}{q}$, the value $\frac{qc}{q+r}$ assumed $q + r$ times, and each of the other $m - 1$ values assumed q times. These values are contained in $[a, c]$ for all $\mathbf{x} \in Q_t$, provided $\frac{qc}{q+r} \ge a$ and $c - \frac{t}{q} \ge a$, which is equivalent to $q \ge \max\left[\frac{ar}{c-a}, \frac{t}{c-a}\right]$. Since $q > \frac{k}{m+1}$ and $r < m$, this is guaranteed by the condition

$$k > K_t = (m+1) \max\left[\frac{am}{c-a}, \frac{t}{c-a}\right] ,$$

as stated in the lemma.

The t_i's are a certain permutation of the values listed above. First, let the processors $P_1, \ldots, P_m$ be indexed so that P_{i+1} carries a task with residual running time x_i, where $0 = x_0 \le \cdots \le x_{m-1}$. The k tasks are run as follows. At time 0, P_1 starts running in succession the $q + r$ tasks with running times $\frac{qc}{a+r}$, and for $1 \le i \le m - 1$, P_{i+1} starts running in succession at time x_i the q tasks with running times $c - \frac{x_i}{q}$. For the above set of tasks, let C_i be the i^{th} completion time, with $C_0 = 0$. The task running times $t_1, \ldots, t_k$ are defined to be the respective running times of the tasks started at times $C_0, C_1, \ldots, C_{k-1}$. Figure 3 illustrates the definitions for $k = 7$, $m = 3$, and hence $q = 2$, $r = 1$.

The m processors finish their tasks at time $C_k = qc$. Thus, for $t > 0$, $k > K_t$, $\mathbf{x} \in Q_t$, we have task durations t_i such that $a \le t_1, \ldots, t_k \le c$ and $X^k(\mathbf{x}; t_1, \ldots, t_k) = 0$. The set $\{(\mathbf{x}; T_1, \ldots, T_k) : \mathbf{x} \in Q_t,\ a \le T_1, \ldots, T_k \le b\}$ is a compact subset of $\mathbb{R}^{k+m-1}$, and by Lemma 1, $X^k(\mathbf{x}; T_1, \ldots, T_k)$ is uniformly continuous on such sets. We may therefore choose $\delta' = \delta'_{\epsilon,t,k} < b - c$ so that

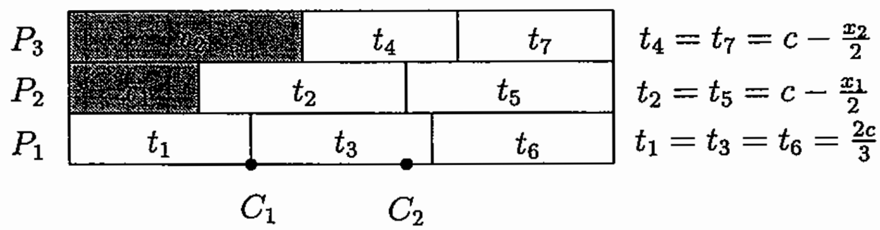

Figure 3: Illustration for Lemma 2, $m = 3$.

$$X^k(\mathbf{x}; T_1, \ldots, T_k) \in Q_\epsilon \quad \text{whenever} \quad \mathbf{x} \in Q_t \quad \text{and} \quad 0 < T_1 - t_1, \ldots, T_k - t_k < \delta' . \tag{4.1}$$

Since $[t_i, t_i + \delta'] \subset [a, b]$, $1 \leq i \leq k$, and since $F'(x) > \xi$, $a \leq x \leq b$, we obtain

$$P(0 < T_i - t_i < \delta') > \xi\delta', \ 1 \leq i \leq k .$$

These observations along with the independence of the T_i's shows that, for $\epsilon, t > 0$, $k > K_t$, $\mathbf{x} \in Q_t$,

$$\begin{aligned} P(X^k(\mathbf{x}, T_1, \ldots, T_k) \in Q_\epsilon) &\geq P(0 < T_1 - t_1, \ldots, T_k - t_k < \delta') \\ &= P(t_1 < T_1 < t_1 + \delta') \cdots P(t_k < T_k < t_k + \delta') > (\xi\delta')^k . \end{aligned} \tag{4.2}$$

The lemma then follows by letting $\delta_{\epsilon,t,k} = (\xi\delta')^k$. ■

Proof of Theorem 5. Let $Z^k = (Z_1^k, \ldots, Z_{m-1}^k) = (X_{m-1}^k, X_{m-1}^k - X_1^k, \ldots, X_{m-1}^k - X_{m-2}^k)$ and suppose that $T_{k+i} = Z_i^k + U_i^k$, $1 \leq i \leq m-1$, where $0 \leq U_1^k \leq \cdots \leq U_{m-1}^k$. As illustrated in Fig. 4, we have $X^{k+m-1} = U^k = (U_1^k, \ldots, U_{m-1}^k)$. Let $Z^k + B$, $B \in \mathcal{B}(Q)$, denote the set B translated by Z^k, i.e., $Z^k + B$ contains just those vectors $(Z_1^k + x_1, \ldots, Z_{m-1}^k + x_{m-1})$ with $(x_1, \ldots, x_{m-1}) \in B$. Then

$$\{(T_{k+1}, \ldots, T_{k+m-1}) \in Z^k + B\} \subseteq \{X^{k+m-1} \in B\}, \quad B \in \mathcal{B}(Q) . \tag{4.3}$$

Let $t_1, \ldots, t_k$, δ', and K_t be as defined in the proof of Lemma 2, and let $P^{(1)}$, $P^{(2)}$ be the respective probability distributions of $(T_1, \ldots, T_k)$ and $(T_{k+1}, \ldots, T_n)$, $n = k + m - 1$. By the independence of the T_i's, we obtain from (4.3)

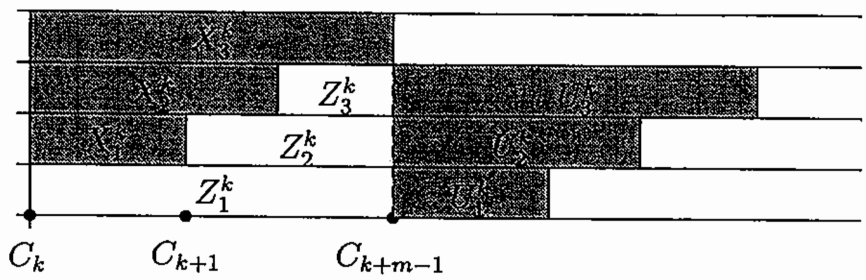

Figure 4: Illustration for Theorem 5, $m = 4$.

$$
\begin{aligned}
P(X^n \in B) &\geq P((T_{k+1}, \ldots, T_n) \in Z^k + B) \\
&\geq P(0 < T_1 - t_1, \ldots, T_k - t_k < \delta', (T_{k+1}, \ldots, T_n) \in Z^k + B) \\
&= \int_{\{0 < T_1 - t_1, \ldots, T_k - t_k < \delta'\}} P^{(2)}((T_{k+1}, \ldots, T_n) \in Z^k + B) dP^{(1)} \, .
\end{aligned}
\tag{4.4}
$$

As in Lemma 2, consider $t > 0$, $k > K_t$, $\mathbf{x} \in Q_t$, $B \subset A_c$, and put $\epsilon = b - c$. Suppose that $0 < T_1 - t_1, \ldots, T_k - t_k < \delta'$ and hence by (4.1), $X^k \in Q_\epsilon$. This is equivalent to $Z^k \in Q_\epsilon$, so $Z^k + B \subset A_b$. Thus, since $F'(x) > \xi$, $a \leq x \leq b$,

$$
P^{(2)}((T_{k+1}, \ldots, T_n) \in Z^k + B) \geq \xi^{m-1} \lambda(Z^k + B) = \xi^{m-1} \lambda(B) \ , \tag{4.5}
$$

the inequality following from the translation invariance of Lebesgue measure. We conclude from (4.2), (4.4), and (4.5) that $P(X^n \in B) \geq \xi^{m-1} \delta_{\epsilon,t,k} \lambda(B)$. Theorem 5 follows by choosing $N_t = K_t + m - 1$ and $D(n,t) = \xi^{m-1} \delta_{\epsilon,t,n-m+1}$. ∎

With Theorem 1, Section 2, and Theorem 5 as the foundation, we proceed to prove three limit laws for the chain $\{X^n\}$. These results correspond to Theorems 2–4. *We assume in each theorem that* $F \in \mathcal{C}$. By the remark at the end of Section 3, the invariant measure in each of the following theorems is unique, and hence is the one given by formula (2.6).

Theorem 6. *Let* $F \in \mathcal{C}$ *and let* F *have compact support, say* $F(x) = 1$ *for* $x \geq x_0$. *Then, for any* $s > x_0$, *there exists a* $\theta(s)$, $0 < \theta(s) < 1$, *such that*

$\|P(\mathbf{x},\cdot) - \pi\| \le \theta^n(s)$ *for all* $\mathbf{x} \in Q_s$ *and* $n > 0$.

Proof. From the assumptions, it follows that, with probability 1, the residual running times at any time are at most s. Thus, Q_s may be taken as the state space of $\{X^n\}$. This state space is compact, so by Corollary 2 it is small. The theorem thus follows from Theorems 1 and 2. ■

Theorem 7. *If* $F \in \mathcal{C}$ *and* $E[T] < \infty$, *then for any initial distribution* μ_0,

$$\left\| \int \mu_0(d\mathbf{x}) P^n(\mathbf{x},\cdot) - \pi \right\| \to 0 \quad \text{as} \quad n \to \infty\ .$$

Proof. By Corollary 2, $\{X^n\}$ is ψ-irreducible and aperiodic, so in view of Theorems 1 and 3, we will obtain the desired result if we can exhibit a potential function V with the two properties in Theorem 3. We choose $V(\mathbf{x}) = mx_{m-1} - (x_1 + \cdots + x_{m-1})$ as illustrated in Fig. 5(i).

Since $x_1 \le \cdots \le x_{m-1}$, we obtain $V \ge mx_{m-1} - (m-1)x_{m-1} = x_{m-1}$, so $\lim_{|\mathbf{x}|\to\infty} V(\mathbf{x}) = \infty$. Thus, V is unbounded off small sets, which verifies the second of the two properties in Theorem 3.

Let the new task assigned in state $\mathbf{x}$ have duration t and cause a transition from state $\mathbf{x}$ to state $\mathbf{y}$. As illustrated in Fig. 5(ii), (iii), we have

$$V(\mathbf{y}) = mx_{m-1} - (x_1 + \cdots + x_{m-1} + t) = V(\mathbf{x}) - t, \text{ if } t \le x_{m-1} \tag{4.6}$$

$$V(\mathbf{y}) = mt - (x_1 + \cdots + x_{m-1} + t) = V(\mathbf{x}) + (m-1)t - mx_{m-1}, \text{ if } t > x_{m-1} \tag{4.7}$$

Thus,

$$\begin{aligned}
\Delta V &= \int [V(\mathbf{y}) - V(\mathbf{x})] P(\mathbf{x}, d\mathbf{y}) \\
&= -\int_0^{x_{m-1}} t dF(t) + \int_{x_{m-1}}^{\infty} [(m-1)t - mx_{m-1}] dF(t) \\
&= -\int_0^{\infty} t dF(t) + \int_{x_{m-1}}^{\infty} m(t - x_{m-1}) dF(t) \\
&\le -\int_0^{\infty} t dF(t) + m \int_{x_{m-1}}^{\infty} t dF(t)\ .
\end{aligned}$$

Since $E[T] = \int_0^\infty t dF(t) < \infty$, we have that $\lim_{x_{m-1}\to\infty} \int_{x_{m-1}}^{\infty} t dF(t) = 0$. We

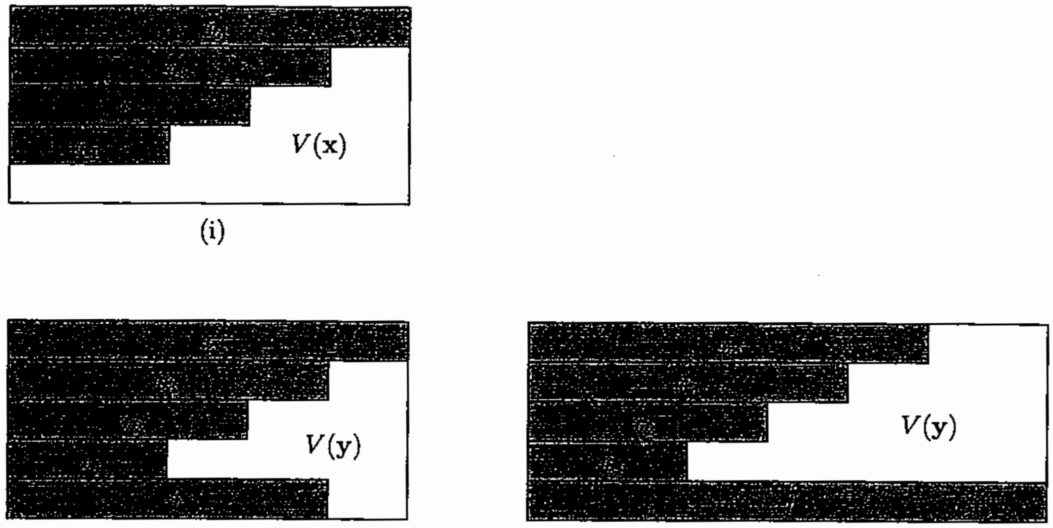

(i)

(ii) $t \leq x_4$ (iii) $t > x_4$

Figure 5: Illustration for Theorem 7, $m = 5$.

conclude that $\Delta V < 0$ for x_{m-1} sufficiently large, say $x_{m-1} > s_1$. Then $\Delta V < 0$, except on the compact set Q_{s_1}, so the first and only remaining condition in Theorem 3 is satisfied. ■

Theorem 8. *Let $F^c(x) \equiv 1 - F(x) = O(e^{-\kappa x})$ for some $\kappa > 0$. Then there exist positive constants η, ζ, and ρ, with $\rho < 1$, such that*

$$\|P^n(\mathbf{x}, \cdot) - \pi\| \leq \eta e^{\zeta(mx_{m-1} - x_1 - \cdots - x_{m-1})} \rho^n$$

for all $\mathbf{x} = (x_1, \ldots, x_{m-1}) \in Q$ and $n > 0$.

Proof. By Corollary 2, $\{X^n\}$ is ψ-irreducible and aperiodic, and compact sets are small. Hence, by Theorem 4, it suffices to produce a potential function V satisfying $V \geq 1$ and, for some $\beta, c > 0$ and some compact set A,

$$\Delta V = -\beta V + c \cdot 1_A \ . \tag{4.8}$$

We choose $V = e^{\zeta W}$, where $0 < \zeta < \frac{\kappa}{m-1}$ and $W(\mathbf{x}) = mx_{m-1} - x_1 - \cdots - x_{m-1}$, i.e., W is the function V appearing in the proof of Theorem 7. Since $W \geq 0$, we have $V \geq 1$, so it remains to verify (4.8). For this, we use (4.6) and (4.7), replacing V by W.

Note first that integration by parts together with $F^c(x) = O(e^{-\kappa x})$ shows that

$$\int_0^\infty e^{\zeta(m-1)x} dF(x) = -\int_0^\infty e^{\zeta(m-1)x} dF^c(x) < 0$$

for $0 < \zeta < \frac{\kappa}{m-1}$. Then (4.6), (4.7) give

$$\begin{aligned}
\int V(\mathbf{y})P(\mathbf{x}, d\mathbf{y}) &= \int_0^{x_{m-1}} e^{\zeta[W(\mathbf{x})-t]} dF(t) + \int_{x_{m-1}}^\infty e^{\zeta[W(\mathbf{x})+(m-1)t-mx_{m-1}]} dF(t) \\
&= V(\mathbf{x}) \left[\int_0^{x_{m-1}} e^{-\zeta t} dF(t) + \int_{x_{m-1}}^\infty e^{\zeta(m-1)t-mx_{m-1}} dF(t) \right] \qquad (4.9) \\
&\leq V(\mathbf{x}) \left[\int_0^\infty e^{-\zeta t} dF(t) + \int_{x_{m-1}}^\infty e^{\zeta(m-1)t} dF(t) \right] .
\end{aligned}$$

Now choose $\epsilon > 0$ so that $F(\epsilon) < 1$. Then

$$\begin{aligned}
\int_0^\infty e^{-\zeta t} dF(t) &= \int_0^\epsilon e^{-\zeta t} dF(t) + \int_\epsilon^\infty e^{-\zeta t} dF(t) \\
&\leq \int_0^\epsilon dF(t) + e^{-\zeta\epsilon} \int_\epsilon^\infty dF(t) = 1 - (1 - e^{-\zeta\epsilon})(1 - F(\epsilon)) \qquad (4.10)
\end{aligned}$$

Integration by parts along with $F^c(x) = O(e^{-\kappa x})$ gives

$$\begin{aligned}
\int_{x_{m-1}}^\infty e^{\zeta(m-1)t} dF(t) &= -\int_{x_{m-1}}^\infty e^{\zeta(m-1)t} dF^c(t) \\
&= e^{\zeta(m-1)x_{m-1}} F^c(x_{m-1}) + \zeta(m-1) \int_0^\infty F^c(t) e^{\zeta(m-1)t} dt \qquad (4.11) \\
&= O(e^{-[\kappa - \zeta(m-1)]x_{m-1}}) .
\end{aligned}$$

From (4.9)–(4.11) we get for x_{m-1} sufficiently large, say $x_{m-1} > s_2$,

$$\int V(\mathbf{y})P(\mathbf{x}, d\mathbf{y}) \leq (1-\beta)V(\mathbf{x}), \quad \beta = 1 - \frac{1 - e^{-\zeta\epsilon}}{2}(1 - F(\epsilon)) . \qquad (4.12)$$

For $\mathbf{x} \in Q_{s_2}$, we have $V(\mathbf{x}) \leq e^{\zeta m s_2}$. Hence, by (4.9)

$$\begin{aligned}
\Delta V &= V(\mathbf{x}) \left[\int_0^\infty e^{-\zeta t} dF(t) + \int_{x_{m-1}}^\infty e^{\zeta(m-1)t} dF(t) - 1 \right] \\
&\leq e^{\zeta s_2} \int_0^\infty e^{\zeta(m-1)t} dF(t) . \qquad (4.13)
\end{aligned}$$

Finally, (4.8) follows from (4.12), (4.13) by choosing $A = Q_{s_2}$, β as in (4.12), and

$$c = e^{\zeta s_2} \int_0^\infty e^{\zeta(m-1)t} dF(t) .$$

■

We conclude this section with a limit law for the expectation $E[\alpha_{m,n}]$.

Theorem 9. *Let $1 - F(x) = O(e^{-\kappa x})$ for some $\kappa > 0$. Then for some ρ, $0 < \rho < 1$, we have $E[\alpha_{m,n}] = E[\alpha_m] + O(\rho^n)$ as $n \to \infty$.*

Proof. Let η, ζ, and ρ be the constants appearing in Theorem 8. From (2.8) and (2.9) we conclude that

$$E[\alpha_{m,n}] - E[\alpha_m] = E[X_{m-1}^{n-m+1}] - E[X_{m-1}^*] - \frac{1}{m} \sum_{i=1}^{m-1} (E[X_i^{n-m+1}] - E[X_i^*]) , \tag{4.14}$$

where

$$E[X_i^{n-m+1}] - E[X_i^*] = \int y_i [P^{n-m+1}(dy) - \pi(dy)], \quad 1 \le i \le m-1 . \tag{4.15}$$

Let $V(\mathbf{x})$ be defined as in the proof of Theorem 8. Then

$$V = e^{\zeta[m x_{m-1} - x_1 - \cdots - x_{m-1}]} \ge e^{\zeta x_{m-1}} \ge \zeta x_i, \quad 1 \le i \le m-1 . \tag{4.16}$$

From (4.15), (4.16), and Theorem 4, we conclude that

$$\left| E[X_i^{n-m+1}] - E[X_i^*] \right| \le \frac{\eta}{\zeta} e^{\zeta[m x_{m-1} - x_1 - \cdots - x_{m-1}]} \rho^{n-m+1}, \quad 1 \le i \le m-1 . \tag{4.17}$$

The theorem follows from (4.14) and (4.17). ■

5. Policy-Free Error Asymptotics

From (2.7) it is clear that the relative size of the error $\alpha_{m,n}$ compared to the makespan $L_{m,n}$ itself is asymptotically negligible as $n \to \infty$. For large n, obviously the dominant part of $L_{n,m}$ for any policy is the sum of all the processing times divided by m. In this section we establish a stronger result. We show that *for any policy the limiting behavior of the error $\alpha_{m,n}$ as $n \to \infty$*

is independent of the policy. In particular, the mean of the error $\alpha_{m,n} = L_{m,n} - \frac{n}{m}E[T]$ for a given policy is asymptotically negligible compared to the standard deviation of the makespan (which is the same as the standard deviation of the error). In a probabilistic setting, what we can gain from a good policy is asymptotically negligible as $n \to \infty$ compared to our degree of uncertainty about the makespan.

Central limit theorems (CLTs) and functional central limit theorems (FCLTs) provide asymptotics that exhibit this property for general distributions and a policy-free set-up, i.e., a model yielding results simultaneously valid for all policies. The policies to be considered in the illustrations below are those in the class of *list scheduling* (LS) policies. Such a policy begins by computing a permutation $\tau_n = (\tau(1), \ldots, \tau(n))$ of the integers $1, \ldots, n$, and then schedules the ordered list $(T_{\tau(1)}, \ldots, T_{\tau(n)})$ by the greedy rule. For any given sequence $T_1, T_2, \ldots$, an LS policy defines a sequence of permutations $\{\tau_n, n \geq 1\}$.

Let S_n and M_n be the sum and maximum of $T_1, \ldots, T_n$, and note that both quantities are invariant under permutations of $T_1, \ldots, T_n$. Let the number m_n of processors be a nondecreasing function of n, and denote the makespan and error under permutation τ_n by $L_n^{\tau_n}$ and $\alpha_n^{\tau_n}$. From (2.3) we obtain the basic inequality

$$|S_n - m_n L_n^{\tau_n}| \leq m_n M_n \quad \text{for all} \quad \tau_n \ , \tag{5.1}$$

from which we see that the limiting behavior of $L_n^{\tau_n}$, $\alpha_{m,n}^{\tau_n}$ is determined by the asymptotics of (S_n, M_n). Typically, when a CLT holds for S_n, M_n is asymptotically negligible compared to S_n. (See §4.5 of Resnick [24] for further discussion of the asymptotic behavior of (S_n, M_n).) In our case, we have the CLT

$$n^{-1/2}(S_n - nE[T]) \Rightarrow N(0, \sigma^2(T)) \quad \text{as} \quad n \to \infty \ , \tag{5.2}$$

where $\Rightarrow$ denotes convergence in distribution and $N(a, b)$ denotes a normally distributed random variable with mean a and variance b. It then follows from (5.1) and Theorem 4.1 of Billingsley [3] that, if

$$m_n n^{-1/2} M_n \Rightarrow 0 \quad \text{as} \quad n \to \infty ,$$

then for any sequence of permutations $\{\tau_n, n \geq 1\}$,

$$\begin{aligned} n^{-1/2} m_n \alpha_n^{\tau_n} &= n^{-1/2}(S_n - nE(T)) + n^{-1/2}(m_n L_n^{\tau_n} - S_n) \\ &\Rightarrow N(0, \sigma^2(T)) \quad \text{as} \quad n \to \infty . \end{aligned} \tag{5.3}$$

For example, suppose T is exponentially distributed. Since $M_n / \ln n \Rightarrow E[T]$, $n \to \infty$ (e.g. see Leadbetter, Lindgren, and Rootzén [20]), then (5.3) holds if $m_n = o(n^{1/2} / \ln n)$, $n \to \infty$.

For a fixed number $m_n = m$, $n \geq 1$, of processors, no explicit assumption about M_n needs to be made. This can be seen in the general setting of the following FCLT for S_n. In terms of the usual diffusion-limit scalings, define the normalized processes

$$\mathbf{S}_n \equiv \mathbf{S}_n(t) = \frac{S_{\lfloor nt \rfloor} - E[T]nt}{n^{1/2}} , \quad t \geq 0$$

$$\boldsymbol{\alpha}_n^{\tau_n} \equiv \boldsymbol{\alpha}_n^{\tau_n}(t) = \frac{L_{\lfloor nt \rfloor}^{\tau_n} - (E[T]/m)nt}{n^{1/2}} , \quad t \geq 0 .$$

If $\mathbf{B}$ denotes standard (zero drift, unit diffusion) Brownian motion, then we have

$$\mathbf{S}_n \Rightarrow \sigma(T)\mathbf{B} \quad \text{as} \quad n \to \infty , \tag{5.4}$$

where $\Rightarrow$ denotes weak convergence in the Skorohod space $D \equiv D([0,1], \mathbf{R})$ (see Ethier and Kurtz [12]). By the continuous mapping theorem with the maximum jump functional, we deduce from (5.4) that $n^{-1/2} M_n \Rightarrow 0$ as $n \to \infty$. Hence, for any sequence of permutations $\{\tau_n, n \geq 1\}$,

$$\boldsymbol{\alpha}_{m,n}^{\tau_n} \Rightarrow \frac{\sigma(T)}{m}\mathbf{B} \ \text{in} \ D \quad \text{as} \quad n \to \infty . \tag{5.5}$$

This gives the approximation

$$L_{m,n}^{\tau_n} \approx \frac{nE[T]}{m} + n^{1/2} N(0, \sigma^2(T)/m^2) , \tag{5.6}$$

in which τ_n does not appear, since the effect of the permutation τ_n is of order

M_n, which is asymptotically negligible compared to $n^{1/2}$.

We remark that the setting for the above limit laws can be broadened considerably, covering interesting cases where the independence assumption or the identical-distribution assumption does not hold.

Acknowledgment. We thank Donald P. Gaver, Alan Weiss and Paul E. Wright for helpful discussions.

References

[1] Asmussen, S., *Applied Probability and Queues*, Wiley, New York, 1987.

[2] Baccelli, F. and Brémaud, P., *Elements of Queueing Theory*, Springer, New York, 1984.

[3] Billingsley, P., *Convergence of Probability Measures*, Wiley, New York, 1968.

[4] Blazewicz, J., Ecker, K., Schmidt, G., and Weglarz, J., *Scheduling in Computer and Manufacturing Systems*, Springer-Verlag, Berlin, 1993.

[5] Boxma, O. J., A Probabilistic Analysis of the LPT Scheduling Rule, in E. Gelenbe, editor, *Performançe '84: Proceedings of the Tenth International Symposium on Models of Computer System Performance*, North-Holland, 1984.

[6] Boxma, O. J., A Probabilistic Analysis of Multiprocessor List Scheduling: the Erlang Case, *Stochastic Models*, 1 (1985) 209–220.

[7] Bruno, J. L. and Downey, P. J., Probabilistic Bounds on the Performance of List Scheduling, *SIAM J. Comput.*, 15 (1986) 409–417.

[8] Çinlar, E., Superposition of Point Processes, in *Stochastic Point Processes: Statistical Analysis, Theory and Applications*, P. A. W. Lewis (ed.), Wiley, New York, 1972, 549–606.

[9] Coffman, E. G., Jr. and Gilbert, E. N., On the Expected Relative Performance of List Scheduling, *Oper. Res.*, 33 (1985) 548–561.

[10] Coffman, E. G., Jr. and Lueker, G. S., *Probabilistic Analysis of Packing and Related Partitioning Problems*, Wiley, New York, 1991.

[11] Coffman, E. G., Jr. and Whitt, W., Recent Asymptotic Results in the Probabilistic Analysis of Makespan Scheduling, in *Scheduling Theory and Its Applications*, P. Chrétienne, E. G. Coffman, Jr., J. K. Lenstra and Z. Liu (eds). Wiley, New York, 1995, 15–31.

[12] Ethier, S. N. and Kurtz, T. G., *Markov Processes: Characterization and Convergence*, Wiley, New York, 1986.

[13] Feller, W., *An Introduction to Probability Theory and Its Applications*, Volume II, second edition, John Wiley & Sons, New York, 1971.

[14] Franken, P., König, D., Arndt, U., and Schmidt, V., *Queues and Point Processes*, Wiley, New York, 1982.

[15] Frenk, J. B. G. and Rinnooy Kan, A. H. G., The Asymptotic Optimality of the LPT Rule, *Math. Oper. Res.*, 12 (1987) 241–254.

[16] Garey, M. R. and Johnson, D. S., *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, New York, 1979.

[17] Graham, R. L., Bounds for Certain Multiprocessing Anomalies, *Bell. Sys. Tech. J.*, 45 (1966) 1563–1581.

[18] Han, S., Hong, D., and Leung, J. Y.-T., On the Asymptotic Optimality of Multiprocessor Scheduling Heuristic Algorithms, Tech. Rep., Computer Science Dept., University of Nebraska, Lincoln, NE 68588-0115, 1992.

[19] Lawler, E. L. Lenstra, J. K., Rinnooy Kan, A. H. G., and Shmoys, D. B., Sequencing and Scheduling: Algorithms and Complexity, Rep. BS-R8909,

Centre for mathematics and Computer Science, 1009AB Amsterdam, 1992. (To appear in *Handbook of Operations Research and Management Science*, North-Holland, Amsterdam.)

[20] Leadbetter, M. R., Lindgren, G., and Rootzén, H., *Extremes and Related Properties of Random Sequences and Processes*, Springer, New York, 1983.

[21] Loulou, R., Tight Bounds and Probabilistic Analysis Two Heuristics for Parallel Processor Scheduling, *Math. Oper. Res.*, 9 (1984) 142–150.

[22] Meyn, S. P. and Tweedie, R. L., *Markov Chains and Stochastic Stability*, Springer-Verlag, London, 1993.

[23] Reiss, R.-D., *Approximate Distributions of Order Statistics*, Springer-Verlag, New York, 1989.

[24] Resnick, S. I., Point Processes, Regular Variation and Weak Convergence, *Adv. Appl. Prob.*, 18 (1986) 66–138.

[25] Rinnooy Kan, A. H. G. and Frenk, J. B. G., On the Rate of Convergence to Optimality of the LPT Rule, *Discrete Applied Mathematics*, 14 (1986) 187–198.

[26] Whitt, W., On Approximations for Queues, I: Extremal Distributions, *AT&T Bell. Lab. Tech. J.*, 63 (1984) 115–138.

Received: 1/3/1995
Accepted: 9/15/1995

Recommended by Sidney I. Resnick, Editor