

CONCURRENT PROBABILISTIC PROGRAMS, OR: HOW TO SCHEDULE IF YOU MUST*

SERGIU HART† AND MICHA SHARIR‡

Abstract. Consider a finite set of processes, such that each one may use randomizations in its course of execution; these processes are running concurrently, under a fair interleaving schedule. We analyze the worst-case probability of termination, i.e., program convergence to a specified set of goal states. Several methods for computing this probability are presented, and characterizations of the special case where it is identically 1 are derived. Specializations of these characterizations to the case of deterministic and nondeterministic programs, and to the case of programs with finite state spaces, are also discussed.

Key words. concurrent probabilistic program, scheduler, fairness, program termination, Markov chains

1. Introduction. This paper continues the study, begun in [HSP], of termination of concurrent probabilistic programs. The model that we assume is that of a finite set K of concurrent processes, each of which is allowed to use randomization, i.e., draw randomly according to probability distributions. These processes execute asynchronously, and we can thus consider each process $k \in K$ as a discrete Markov chain (with stationary transition probability matrix P^k) on the set I of common execution states. The overall execution behavior of these processes is described in terms of the interleaving pattern in which they are scheduled by some imaginary scheduler σ . Each process k scheduled at a state i can reach more than one subsequent state, so that to specify σ we need to consider all these transitions simultaneously. We may therefore represent σ as a tree (referred to as the *execution-tree* or the *transition-tree* induced by σ) each of whose nodes is labeled by a pair (i, k) , where $i \in I$ is the state reached at that node, and where $k \in K$ is the process to be scheduled there next. A node (i_1, k_1) will be a son of (i, k) in the tree if there exists a positive transition probability of reaching i_1 from i under a single execution step of process k , and if process k_1 will be next scheduled at i_1 , provided that this transition has indeed taken place.

Given such a σ , it induces in a standard manner a probability measure μ_σ on the space of all infinite sequences of states.

We consider here general schedules σ , with the sole restriction that they be *fair*, meaning that no process stops being scheduled; i.e., that the μ_σ -measure of the set of all tree paths on which each process $k \in K$ is scheduled infinitely often is 1.

This model is discussed and justified more fully in [HSP]. We note that it coincides with the model assumed by Lehmann and Rabin in [LR], and also with that used by Dubins and Savage [DS] in their study of optimal gambling strategies (with the essential exception that they do not require fairness). It does differ, though, from various other models used in the literature (cf. [Ra1], [Ra2], [RS1], [RS2]). The crucial distinction lies in the degree to which the imaginary scheduler can base its scheduling decisions on the outcome of random draws made by the processes, or, more generally, on their internal states. These more restrictive scheduling models usually correspond to situations in which the execution time of a single step of a process is independent of its current state and of the outcome of the random draws it has made. Our model is more general, and allows for such dependence, thereby being a more realistic model for

* Received by the editors September 15, 1982, and in revised form August 1, 1984.

† School of Mathematical Sciences, Tel Aviv University, Ramat Aviv, 69978 Tel Aviv, Israel.

‡ The work of this author was supported in part by the Bat-Sheva Fund at Tel Aviv University, and by the Office of Naval Research under grant N00014-75-C-0571 at the Courant Institute.

general concurrent or distributed probabilistic execution. Moreover, properties established for concurrent probabilistic programs under our model will continue to hold under the more restrictive models mentioned above, but not necessarily vice-versa (for example, Rabin's synchronization algorithm described in [Ra1] is shown in [HSP] to fail in our model).

In the preceding paper [HSP], we have analyzed *termination* of concurrent probabilistic programs having a *finite* state space. We have obtained there necessary and sufficient conditions for such a program to reach (with probability 1) a given set X of *goal states* from some initial state, under any fair schedule. These conditions can be checked mechanically, and are independent of the particular values of nonzero transition probabilities of the processes involved.

In this paper we generalize and extend these results to programs with *infinite* state spaces. As in the case of a single Markov chain, the analysis of program termination becomes much more complicated in the general case, and becomes dependent upon the actual values of the nonzero transition probabilities involved. The basic problem that we treat in this paper is the computation of the function φ on the set of states I , where, for each $i \in I$, $\varphi(i)$ is the minimum probability of program termination starting at state i , under any fair schedule. We establish various properties and characterizations of φ , and derive from them several techniques for the calculation of this function. This theory enables us to gain a better understanding of the structure of the (worst-case) convergence of the program towards termination. For example, one can interpret this convergence process as a game between the program and the scheduler, in which each move of the program requires the scheduler to schedule one of the processes and the scheduler responds by scheduling this process eventually, but only after scheduling some other processes prior to it, in a way which would hurt as much as possible the program's probability to terminate. We show that the optimal payoff for the program in this game is the function φ , provided that the game is long enough, where the length of such a game is measured by some (infinite) ordinal.

The various characterizations of φ are next used to obtain necessary and sufficient conditions for the special case $\varphi \equiv 1$ (i.e. for worst-case almost-sure termination from any initial state) to hold. Some of these conditions generalize similar conditions given in the preceding paper [HSP] for programs with finite state spaces. These characterizations of program termination are next specialized to the case in which the processes are deterministic or nondeterministic.¹ Some of these characterizations are shown to reduce to the conditions given by Lehmann, Pnueli and Stavi [LSP] for the termination of nondeterministic programs, while others are new. Finally, the special case of probabilistic programs with finite state spaces is reconsidered from the viewpoint of the general theory developed in this paper, enabling us to obtain the decomposition of the state space described in [HSP] in a different manner. The results of this paper are exemplified on several running example programs. The techniques developed in this paper can be immediately interpreted as (sound and complete) proof methods for almost sure program termination.

This paper is organized as follows: Section 2 presents the notations and terminology used in the paper, and begins the analysis of φ by establishing some more elementary properties of this function. Section 3 develops the main technical tools for the analysis

¹ A nondeterministic program is one where each execution step of any of its processes may lead from a state $i \in I$ to several succeeding states, but where there is no probability distribution associated with these states; instead, each of these succeeding states must be considered as being potentially the sole successor of i . Such a program is said to terminate if *every* execution sequence terminates.

and characterizations of φ , and obtains φ as the limit of a certain transfinite sequence of functions. Section 4 gives further characterizations of φ . Section 5 treats the special case $\varphi \equiv 1$ (i.e. of almost-sure worst-case termination), and derives various characterizations of this property. Section 6 specializes the preceding results to the case of deterministic and nondeterministic programs. The new characterization of termination of such programs is also given a direct proof. Section 7 treats the special case of probabilistic programs with finite state spaces. Some concluding remarks are presented in § 8.

2. Preliminaries. In this section we present our model of probabilistic concurrent programs in more precise terms, introduce some notations, and establish several preliminary properties of the worst-case termination probability of the program.

A *concurrent probabilistic program* consists of a finite set K of processes acting on a *state space* I ; each $i \in I$ is a common execution state of the processes, and is specified by the program location at each process, by the values of all variables—shared and private—etc. Each $k \in K$ can be regarded as a stationary discrete Markov chain on I . (This extra restriction of discreteness, which is quite adequate for actual programs, simplifies the analysis considerably, by avoiding the technical difficulties of treating non σ -additive measures, which would be otherwise necessary as in Dubins and Savage [DS].) Under this assumption, each process $k \in K$ is specified in terms of its transition probability matrix P^k , that is, for each $i, j \in I$, P_{ij}^k is the probability of reaching state j from state i in a single (indivisible) execution step of process k . The nonnegative matrix P^k is stochastic: for each i , $P_{ij}^k > 0$ for at most countably many j , and $\sum_{j \in I} P_{ij}^k = 1$.

As already stated, program execution is assumed to consist of *interleaving* execution steps of the processes, each executing in its turn one indivisible step. Let $i \in I$ be an initial execution state. Let $H(i)$ denote the set of all *finite execution histories* starting at i ; formally,

$$H(i) = \{i\} \times \left(\bigcup_{n=0}^{\infty} I^n \right).$$

An (infinite) *schedule* σ starting at i is simply defined as a function $\sigma: H(i) \rightarrow K$, that is, for each finite history $h \in H(i)$, $\sigma(h)$ is the next process to perform an execution step, given that execution has proceeded so far through the states in h . The set of all schedules starting at i will be denoted by $\Sigma(i)$. To each such schedule σ there corresponds an *execution tree*, defined inductively as follows. Each node of this tree is labelled by a pair (j, k) where j is the current execution state, and k the next process to be scheduled in this node. The root of the tree is labelled by $(i, \sigma(i))$. For each node ν in the tree, let $h \in H(i)$ be the sequence of states along the path from the root to ν , let j be the last state in h , and let $k = \sigma(h)$; then ν is labelled by (j, k) , and its sons are nodes labelled by $(j', \sigma(h, j'))$, (where (h, j') is the concatenation of j' to h) for $j' \in I$ such that $P_{jj'}^k > 0$.

Let $H^*(i)$ denote the set of all *infinite execution histories* starting at i , that is,

$$H^*(i) = \{i\} \times I^\infty \quad \left(\text{where } I^\infty = \prod_{n=1}^{\infty} I \right).$$

Each schedule $\sigma \in \Sigma(i)$ induces a probability measure μ_σ on the cylindrical σ -field on $H^*(i)$, such that for each cylinder $(i, i_1, i_2, \dots, i_n)$, consisting of all histories whose initial $n+1$ states are i, i_1, \dots, i_n ,

$$\mu_\sigma\{(i, i_1, \dots, i_n)\} = \prod_{s=0}^{n-1} P_{i_s, i_{s+1}}^{k_s},$$

where $i_0 = i$, $k_s = \sigma(i_0, i_1, \dots, i_s)$. Expectation with respect to μ_σ will be denoted by E_σ .

Let $H^* = \bigcup_{i \in I} H^*(i)$. Throughout the paper we will use the following notational convention: Elements of H^* —which we call paths or histories—will be denoted by π ; for each such π and each $n \geq 0$, the $(n+1)$ th state along π will be denoted by i_n , and the subpath consisting of the first $n+1$ states in π will be denoted by $\pi_n \equiv (i_0, i_1, \dots, i_n)$. A path π is a *fair path* with respect to a given schedule σ if each $k \in K$ appears infinitely often in the sequence $\{\sigma(\pi_n)\}_{n=0}^\infty$; the schedule σ is a *fair schedule* if $\mu_\sigma\{\pi: \pi \text{ is fair}\} = 1$. For each $i \in I$ we denote

$$\Sigma F(i) = \{\sigma \in \Sigma(i): \sigma \text{ is fair}\}.$$

Let $X \subset I$ be a given set of *goal states*, fixed henceforth. Our aim is to study the convergence of program execution to states in X ; we will therefore assume in the sequel, without loss of generality, that all states in X are *absorbing* for each $k \in K$; i.e., that $P_{i,i}^k = 1$ for each $i \in X$ and each $k \in K$.

The basic problem studied in this paper is that of analyzing and computing the worst-case probability of the program to reach X (i.e., to *terminate*) when executed from a given initial state under a fair schedule. To formalize this notion, let χ_X be the characteristic function of X (defined on I); we extend this function to H^* by putting $\chi_X(\pi) = \lim_{n \rightarrow \infty} \chi_X(i_n)$ (recall that $\pi = (i_n)_{n \geq 0}$). Since X is absorbing, $\chi_X(\pi) = 1$ if X is ever reached along π , and 0 otherwise. The probability of reaching X under σ is then simply $E_\sigma(\chi_X)$. The following standard observation, which also establishes the measurability of the extended χ_X , will be quite useful in the sequel: For each $n \geq 0$ define a “truncated” extension $\chi_X^{(n)}$ of χ_X by putting $\chi_X^{(n)}(\pi) = 1$, if X is reached during the first n steps of π , and 0 otherwise. Then $E_\sigma(\chi_X^{(n)})$ is the probability of reaching X during the first n steps of σ , and we have

$$\lim_{n \rightarrow \infty} E_\sigma(\chi_X^{(n)}) = \sup_n E_\sigma(\chi_X^{(n)}) = E_\sigma(\chi_X).$$

The worst-case termination probability that we seek is defined, for each initial state $i \in I$, as

$$\varphi(i) = \inf_{\sigma \in \Sigma F(i)} E_\sigma(\chi_X).$$

We will shortly establish several preliminary properties of the function φ , but first we introduce additional notations concerning finite portions of program execution. Let \mathbb{N} denote the set of nonnegative integers, and put $\bar{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$. A *stopping time* N is a mapping from H^* into $\bar{\mathbb{N}}$ such that if $N(\pi) = m$ then $N(\pi') = m$ for each path π' which coincides with π at all steps up to, and including m . In other words, $N(\pi)$ may depend only on i_0, i_1, \dots, i_N —i.e., on states visited before this step, but not on future steps (i.e., on i_{N+1}, \dots). A *finite subschedule* at $i \in I$ is a pair $\tau = (\sigma, N)$ where $\sigma \in \Sigma(i)$, and where N is a stopping time on $H^*(i)$ such that $\mu_\sigma(N < \infty) = 1$ (this corresponds to the notion of “policy” of Dubins and Savage [DS]). The intuitive meaning of such a pair is the initial portion of σ up to, and including N ; in particular, the actual value of σ is relevant only up to the stopping time N . The set of all finite subschedules at i will be denoted by $T(i)$ (note that the empty subschedule—i.e., when $N \equiv 0$ —is included in $T(i)$).

In the sequel we will occasionally use the following standard decomposition of an infinite schedule $\sigma \in \Sigma(i)$: Let N be a stopping time, $\mu_\sigma(N < \infty) = 1$; then σ is equivalent to its initial portion $\tau = (\sigma, N)$, followed by the collection of continuation schedules; that is, for each $\pi \in H^*(i)$ (with $N < \infty$), the continuation $\sigma_{\pi_N} \in \Sigma(i_N)$ of σ after the *end state* i_N of τ . Note in particular that σ is fair if and only if each of the continuations σ_{π_N} is fair.

Let α be a real function on I . Then for each finite subschedule $\tau = (\sigma, N) \in T(i)$, the expectation of α with respect to τ is defined as

$$E_\tau(\alpha) = E_\sigma(\alpha(i_N)).$$

For example, let $\sigma \in \Sigma(i)$, and define, for each $n \geq 0$, $\tau_n = (\sigma, n) \in T(i)$ (i.e., with the stopping time $N \equiv n$). Then, as already noted, $E_\sigma(\chi_X) = \lim_{n \rightarrow \infty} E_{\tau_n}(\chi_X)$. Note also that $E_{\tau_0}(\alpha) = \alpha(i)$, and that $E_{\tau_1}(\alpha) = (P^k \alpha)(i)$, where $k = \sigma(i)$.

Having introduced all the required terminology, we begin by establishing a few elementary properties of the function φ .

PROPOSITION 2.1. (a) $\varphi \geq 0$; $\varphi|_X \equiv 1$.

(b) $\varphi(i) = \min_{k \in K} (P^k \varphi)(i)$ for each $i \in I$.

Proof. (a) is trivial, since X is absorbing; note also that $\varphi \leq 1$.

(b) To show that $\varphi(i) \leq (P^k \varphi)(i)$ for $k \in K$, $i \in I$, use a schedule $\sigma \in \Sigma F(i)$ which starts by scheduling k at i , and then continues so as to approximate φ at each of the resulting states. For the converse inequality, take a sequence of schedules σ_n in $\Sigma F(i)$ such that $E_{\sigma_n}(\chi_X)$ converges to $\varphi(i)$ and such that they all start by scheduling the same process $k \in K$ (since K is finite this is always possible); then it is easily seen that $(P^k \varphi)(i) \leq \varphi(i)$. More details can be found in [HS]. Q.E.D.

Extending standard notations in Markov chain theory, we say that a real function α on I is *subharmonic* if $\alpha \leq P^k \alpha$ for each $k \in K$. Similarly α will be called *min-harmonic* if $\alpha = \min_{k \in K} P^k \alpha$ (note that each min-harmonic function is subharmonic).

In the special case where K contains a single process k , the function φ is *harmonic* (i.e., $\varphi = P^k \varphi$). Moreover, it is well-known (cf. [SPH] for example) that φ is the smallest nonnegative harmonic function which is 1 on X . This might lead us to conjecture that for a general (finite) K , φ is also the smallest nonnegative min-harmonic function which is 1 on X . This, however, is not true in general, as can be seen from the following simple example: Let $I = \{0, 1\}$, $X = \{0\}$, and $K = \{1, 2\}$, with the nonzero transition probabilities $P_{1,0}^1 = P_{1,1}^2 = 1$. Obviously, any fair execution of this program brings it into X with certainty, so that $\varphi \equiv 1$, yet the function $\psi(0) = 1$, $\psi(1) = 0$ is a smaller nonnegative min-harmonic function which is 1 on X . The reason for this phenomenon is that fairness is not directly connected to the min-harmonicity of φ . Indeed, let us define a function ψ on I by

$$\psi(i) = \inf_{\sigma \in \Sigma(i)} E_\sigma(\chi_X), \quad i \in I.$$

(i.e., infimum over *all* schedules, not necessarily fair). Then it can be shown that

PROPOSITION 2.2. ψ is the smallest nonnegative min-harmonic function which is 1 on X .

Our next result is a strong form of a "zero-one law" for φ , which generalizes the zero-one law established in [HSP] for finite state spaces.

THEOREM 2.3 (zero-one law). $\inf_{i \in I} \varphi(i)$ is either 0 or 1. Moreover, for each $i \in I$ and $\sigma \in \Sigma F(i)$ define a sequence $\{f_n\}_{n \geq 0}$ of functions on $H^*(i)$ by putting $f_n(\pi) = \varphi(i_n)$, $\pi \in H^*(i)$, $n \geq 0$. Then $\{f_n\}$ converges μ_σ -a.s. to χ_X (extended to $H^*(i)$).

Proof. Let $i \in I$ and $\sigma \in \Sigma F(i)$ be given. The subharmonicity of φ implies that the sequence $\{f_n\}$ is a submartingale² on $H^*(i)$. Since $0 \leq f_n \leq 1$ for each $n \geq 0$, it follows from the (sub)martingale convergence theorem that $\{f_n\}$ converges μ_σ -a.s. to a limit f_∞ . Put $\tau_n = (\sigma, n)$, $n \geq 0$. Then

$$E_\sigma(\chi_X) = E_{\tau_n}(E_{\sigma_{\tau_n}}(\chi_X)) \geq E_{\tau_n}(\varphi) = E_\sigma(f_n)$$

² i.e., for all $n \geq 0$, $E_\sigma(f_{n+1} | \pi_n) \geq f_n$ where π_n is any history of length n with $\mu_\sigma(\pi_n) > 0$.

(since each σ_{π_n} is fair). Letting $n \rightarrow \infty$, we obtain

$$E_\sigma(\chi_X) \geq E_\sigma(f_\infty) \geq 1 \cdot \mu_\sigma\{\pi: f_\infty(\pi) = 1\}.$$

But for each $\pi \in H(i)$, if X is ever reached along π then $f_n(\pi) = \varphi(i_n) = 1$ for all sufficiently large n , so that $f_\infty(\pi) = 1$. Thus

$$\mu_\sigma\{\pi: f_\infty(\pi) = 1\} \geq \mu_\sigma\{\pi: X \text{ reached along } \pi\} = E_\sigma(\chi_X).$$

Therefore we must have equalities throughout; that is

$$E_\sigma(\chi_X) = E_\sigma(f_\infty) = \mu_\sigma\{\pi: f_\infty(\pi) = 1\}.$$

This, however, implies that f_∞ is almost everywhere either 0 or 1, and that $f_\infty(\pi) = 1$ if and only if X is ever reached along π . The zero-one law is now immediate, because if φ is not identically 1, take $i \in I$, $\sigma \in \Sigma F(i)$ such that $E_\sigma(\chi_X) = c < 1$. Then $\varphi(i_n) \rightarrow 0$ on a set of paths whose μ_σ -measure is $1 - c > 0$, thus there exists states with arbitrarily small φ , or $\inf_{i \in I} \varphi(i) = 0$. Q.E.D.

As a final preliminary note, we would like to point out that, unlike the case of a finite state space, the actual values of nonzero transition probabilities of the processes involved can have significant influence on the termination probabilities φ . This is indeed well known even for a single Markov chain. (Consider e.g. the case of a random walk on the nonnegative integers, where the "leftward" transition probability is p . Then the probability of converging towards 0 is identically 1 if $p \geq \frac{1}{2}$, and is exponentially decreasing otherwise; cf. [Ch] for details). Thus, for infinite state spaces there is no hope to obtain purely combinatorial analysis techniques (as have been developed in [HSP] for finite state spaces), and more complex techniques are needed. Development of such techniques is indeed the main purpose of the present paper.

3. φ -iterates. Direct calculation of the function φ from its definition is rather complicated. The purpose of this section is to develop machinery needed for a simpler calculation and characterization of φ . Specifically, we will show that φ is the limit of a transfinite sequence of iterates of a certain operator. We will call these φ -iterates.

DEFINITION. We define an operator Q , and an auxiliary set of operators $\{Q^k\}_{k \in K}$, on the space of all bounded real functions on I , as follows: For each bounded real function α on I , each $i \in I$, and each $k \in K$, put

$$(Q^k \alpha)(i) = \inf_{\tau \in T(i, k)} E_\tau(\alpha),$$

where $T(i, k)$ is defined as

$$\{(\sigma, N+1): \sigma \in \Sigma(i), N \text{ an a.s. finite stopping time with } \sigma(\pi_N) = k\} \subset T(i);$$

i.e., $T(i, k)$ is the set of all subschedules which start at i , schedule k eventually almost surely, and stop right after scheduling k . Q is then defined as

$$(Q\alpha)(i) = \max_{k \in K} (Q^k \alpha)(i).$$

Let R be any of the operators Q^k or Q ; then plainly R is *monotone* (i.e., $\alpha_1 \leq \alpha_2$ implies $R\alpha_1 \leq R\alpha_2$), $R0 = 0$, and $R1 = 1$. The following lemma gives two characterizations of the operators Q^k , one of which is constructive while the other is not.

LEMMA 3.1. *Let α be a bounded real function on I . For each $k \in K$, $Q^k \alpha$ is the largest subharmonic function which does not exceed $P^k \alpha$ (i.e., (1) $Q^k \alpha \leq P^k \alpha$, (2) $Q^k \alpha$ is subharmonic, and (3) if $\beta \leq P^k \alpha$ is subharmonic, then $\beta \leq Q^k \alpha$). Furthermore, $Q^k \alpha$*

is the limit (or infimum) of the following nonincreasing sequence of functions:

$$\beta_1(i) = (P^k \alpha)(i), \quad i \in I,$$

$$\beta_{n+1}(i) = \min \left\{ \beta_n(i), \min_{l \in K} (P^l \beta_n)(i) \right\}, \quad i \in I, \quad n \geq 1.$$

Proof. Let β be the limit of the nonincreasing sequence $\{\beta_n\}$; then β is the largest subharmonic function $\leq P^k \alpha$. The rest follows by noting that, for each $n \geq 1$

$$\beta_n(i) = \inf_{\tau \in T_n(i, k)} E_\tau(\alpha),$$

where $T_n(i, k)$ consists of those subschedules in $T(i, k)$ which stop after at most n steps. Q.E.D.

LEMMA 3.2. For each subharmonic function α and each $k \in K$ we have

$$\alpha \leq Q^k \alpha \leq P^k \alpha.$$

Proof. By Lemma 3.1, $Q^k \alpha$ is the largest subharmonic function which is $\leq P^k \alpha$. Since α itself is subharmonic we have $\alpha \leq P^k \alpha$, so that $\alpha \leq Q^k \alpha$. Q.E.D.

DEFINITION. For each ordinal a we define on I real functions γ_a and γ_a^k , $k \in K$, by the following transfinite inductive process:

$$\gamma_0 = \gamma_0^k = \chi_X, \quad k \in K,$$

$$\gamma_a^k = \sup_{b < a} Q^k \gamma_b \quad \text{for each ordinal } a > 0 \text{ and } k \in K,$$

$$\gamma_a = \max_{k \in K} \gamma_a^k \quad \text{for each ordinal } a.$$

The functions γ_a^k and γ_a are called the φ -iterates of order a of the program (the reason for this terminology will be apparent at the end of this section).

Since X is absorbing, $Q^k \chi_X \geq \chi_X$, thus $\gamma_1^k \geq \gamma_0^k$ for each $k \in K$, hence $\gamma_1 \geq \gamma_0$. Also, by definition, $\gamma_a^k \geq \gamma_b^k$ for each pair of ordinals $a > b > 0$. Thus, for each $k \in K$ the transfinite sequence $\{\gamma_a^k\}_{a \geq 0}$ is nondecreasing, and so is the sequence $\{\gamma_a\}_{a \geq 0}$. From this it follows that $\gamma_{a+1}^k = Q^k \gamma_a^k$ for each ordinal a , and that $\gamma_a = \sup_{b < a} \gamma_b$ for limit ordinals a .

Since the transfinite sequence $\{\gamma_a\}_{a \geq 0}$ is nondecreasing, and each of its elements is obviously bounded between 0 and 1, this sequence must converge to a limit function γ , and there must exist an ordinal c such that $\gamma_c = \gamma$. (Indeed, for each $i \in I$ the transfinite sequence $\{\gamma_a(i)\}$ is a nondecreasing and bounded sequence of real numbers, and so must attain its supremum at some ordinal c_i ; the required ordinal c is simply $\sup_{i \in I} c_i$.) Obviously $Q\gamma = Q\gamma_c = \gamma_{c+1} = \gamma_c = \gamma$. Moreover, using standard fixpoint arguments, it is easily seen by transfinite induction that γ is the smallest fixpoint of Q which is $\geq \chi_X$.

Remarks. (1) To motivate these definitions, it is helpful to consider the following interpretation of the functions γ_a and γ_a^k : $\gamma_0(i)$ is just an indication whether $i \in X$. $\gamma_1^k(i)$ is the smallest probability of reaching X by any subschedule which starts execution at i , and is forced to schedule k eventually (a.s.). Thus $\gamma_1(i)$ is the smallest probability of reaching X that must be yielded by any subschedule starting at i which is forced to schedule any one of the processes at least once. Arguing inductively, $\gamma_n(i)$ is the smallest probability of reaching X that must be yielded by any subschedule starting at i which has to schedule any sequence of n processes one after the other. (Note that this sequence need not be specified in advance; rather the first process k_1

to be scheduled is specified, then the second process to be scheduled is specified, but it may depend on the state reached after scheduling k_1 , and so on.)

(2) $\gamma_n(i)$ can be viewed as the minmax value of a two-person zero-sum game $\Gamma_n(i)$. In this game, the aim of the first player, called "player X ," is to reach X during program execution with the highest possible probability, whereas the aim of the second player, called "the scheduler," is to prevent the program from reaching X as much as possible. The game $\Gamma_n(i)$ consists of n stages. Each stage starts at some state $j \in I$ (stage 1 starts at i). Player X chooses some $k \in K$, and then the scheduler chooses some $\tau \in T(j, k)$. The program is then run according to τ ; when it stops, the next stage is played. After n such stages, player X receives a payoff of one unit from the scheduler if a state in X has been reached, and zero otherwise.

This interpretation can be extended to higher-order ordinals. Specifically, for each ordinal α we define a collection of games $\Gamma_\alpha(i)$, for each $i \in I$, in the following transfinite inductive manner:

(i) $\Gamma_0(i)$ is the "empty" game; player X receives a payoff of 1 from the scheduler if $i \in X$, and zero otherwise.

(ii) If α is not a limit ordinal, say $\alpha = b + 1$, player X first chooses a process k and then the scheduler chooses a subschedule $\tau \in T(i, k)$, and the program is run according to τ ; for each end state j of τ , the game continues as $\Gamma_b(j)$.

(iii) If α is a limit ordinal, player X first chooses an ordinal $b < \alpha$, and then the game continues as $\Gamma_b(i)$.

The definitions (ii) and (iii) imply that after each stage, games with smaller ordinals are played; since every strictly decreasing sequence of ordinals is finite, it follows that every play of any of these games is finite, so that Γ_0 is reached eventually, and the payoff is therefore well defined. Moreover, by the definition of the sequence $\{\gamma_\alpha\}_{\alpha \geq 0}$, one easily obtains by transfinite induction, that $\gamma_\alpha(i)$ is precisely the value of $\Gamma_\alpha(i)$. Indeed, an ε -optimal strategy for player X is constructed as follows (for each $\varepsilon > 0$): If $\alpha = b + 1$, player X first chooses $k \in K$ for which $\gamma_\alpha(i) = \gamma_\alpha^k(i)$, and from each end state j of the subschedule $\tau \in T(i, k)$ subsequently chosen by the scheduler, he continues with an ε -optimal strategy of $\Gamma_j(b)$. If α is a limit ordinal, player X first chooses an ordinal $b < \alpha$ such that $\gamma_\alpha(i) - \gamma_b(i) < \varepsilon/2$, and then continues with an $\varepsilon/2$ -optimal strategy of $\Gamma_b(i)$. As for the scheduler, at the first ordinal $b + 1 \leq \alpha$ where he is called upon to move, he chooses $\tau \in T(i, k)$ such that $E_\tau(\gamma_b) - \varepsilon/2 < (Q^k \gamma_b)(i) = \gamma_{b+1}^k(i)$, and then he continues with an $\varepsilon/2$ -optimal strategy in the corresponding $\Gamma_b(j)$.

Furthermore, the ordinal c (at which $\gamma_{c+1} = \gamma_c$ is first obtained) is such that the expected payoff that player X can guarantee in the game $\Gamma_c(i)$ is the largest possible among all games $\{\Gamma_\alpha(i)\}_{\alpha \geq 0}$ —uniformly in the initial state i . As we shall see later in this section, this maximum payoff is exactly $\varphi(i)$.

(3) Note that if Q were σ -order continuous, i.e., if for any nondecreasing sequence $\{h_n\}$ of uniformly bounded functions we had

$$Q\left(\sup_n h_n\right) = \sup_n Qh_n,$$

then convergence of the γ_α 's would be attained at $c = \omega$ (the first infinite ordinal) or earlier. This is indeed so when I is a finite set, since then each such sequence $\{h_n\}$ converges uniformly to its supremum, in which case Q is clearly continuous. However, this does not hold in general, and so higher ordinals may be needed. (A similar phenomenon is noted by Lehmann, Pnueli and Stavi [LPS] concerning nondeterministic concurrent programs; see § 6 for a detailed comparison between their technique and ours.)


To illustrate the possible discontinuity of Q (and hence the need for higher ordinals), consider the following example (in which both processes involved are actually deterministic).

Example 1. Let $K = \{1, 2\}$, and let $I = I_1 \cup I_2$, where $I_1 = \mathbb{N} \times \{1\}$, $I_2 = \mathbb{N} \times \{2\}$, and $X = \{(0, 1)\}$. The nonzero transition probabilities are

$$P_{(n,1),(n-1,1)}^1 = P_{(n,1),(n,1)}^2 = 1, \quad n > 0,$$

$$P_{(n,2),(n+1,2)}^1 = P_{(n,2),(n,1)}^2 = 1, \quad n \geq 0.$$

These transitions are displayed in the following diagram:



It is easily seen that

$$\gamma_n(i, 1) = \begin{cases} 0, & i \geq n, \\ 1, & i < n, \end{cases} \quad i \in \mathbb{N}.$$

By definition of γ_ω we thus have

$$\gamma_\omega(i, 1) = 1, \quad i \in \mathbb{N}.$$

On the other hand, $\gamma_n(i, 2) = 0$ for each $i, n \in \mathbb{N}$ (to obtain $(Q^2 \gamma_n)(i, 2)$, schedule process 1 sufficiently many times so as to reach a state $(j, 2)$ with $j \geq n$, and then schedule process 2). Thus

$$\gamma_\omega(i, 2) = 0, \quad i \in \mathbb{N}.$$

But $\gamma_{\omega+1} = Q\gamma_\omega > \gamma_\omega$. Indeed, for each $(i, 2) \in I_2$ we have

$$\gamma_{\omega+1}(i, 2) = (Q\gamma_\omega)(i, 2) = (Q^2 \gamma_\omega)(i, 2) = \gamma_\omega(j, 1) = 1 \quad (\text{where } j \geq i).$$

Thus $\gamma_{\omega+1} \equiv 1$, and convergence of the φ -iterates is attained at the ordinal $\omega + 1$.

Remarks. (1) In the game-theoretic terminology established earlier, player X cannot achieve a nonzero payoff in any of the games $\Gamma_n(i, 2)$, $n \in \mathbb{N}$, or even $\Gamma_\omega(i, 2)$, because if the number of rounds n is fixed in advance, the scheduler will initially schedule process 1, $n+1$ times, and this will prevent player X from reaching X in n moves. On the other hand, a payoff of 1 is guaranteed in $\Gamma_{\omega+1}(i, 2)$ as follows: Player X first chooses process 2; no matter what subschedule in $T((i, 2), 2)$ will be chosen by the scheduler, it will end at some state $(j, 1)$ in I_1 , and the game continues from there as $\Gamma_\omega(j, 1)$. Now player X chooses the ordinal $j < \omega$, and this guarantees its entry into X after j additional moves, by requiring to schedule process 1 in each of these moves.

(2) One can easily obtain along similar lines examples where higher and higher ordinals are needed to attain convergence.

(3) If we take in Example 1, $P_{(1,1),(0,1)}^1 = \frac{1}{2}$ (instead of 1) and $P_{(1,1),(0,2)}^1 = \frac{1}{2}$ (instead of 0), it can be verified that the first ordinal c where $\gamma_c = \varphi \equiv 1$ is $c = \omega^2$.

The main purpose of this section is to prove that $\gamma = \varphi$. The proof of this assertion is quite involved and will be split into proving both inequalities $\gamma \leq \varphi$ and $\gamma \geq \varphi$. It consists of the following sequence of lemmata.

LEMMA 3.3. *Each of the φ -iterates γ_a^k and γ_a is subharmonic.*

Proof. Lemma 3.1 and the fact that the supremum of subharmonic functions is subharmonic. Q.E.D.

LEMMA 3.4. *For each $k \in K$, $Q^k \varphi \leq \varphi$.*

Proof. For each $\sigma \in \Sigma F(i)$ let N be the first time k is scheduled, and let $\tau = (\sigma, N+1) \in T(i, k)$. Then, if $\sigma_{\pi_{N+1}}$ denotes the continuation of σ after the end of τ , we have

$$\begin{aligned} E_\sigma(\chi_X) &= E_\tau[E_{\sigma_{\pi_{N+1}}}(\chi_X)] \quad (\text{because } N+1 \text{ is a stopping time}) \\ &\geq E_\tau[\varphi(i_{N+1})] \quad (\text{because } \sigma_{\pi_{N+1}} \in \Sigma F(i_{N+1})) \\ &\geq (Q^k \varphi)(i) \quad (\text{by definition of } Q^k). \end{aligned}$$

Since this holds for each $\sigma \in \Sigma F(i)$, we have $\varphi(i) \geq (Q^k \varphi)(i)$. Q.E.D.

PROPOSITION 3.5. $\varphi = Q\varphi = Q^k \varphi$, for each $k \in K$.

Proof. By the preceding lemma, $Q\varphi = \max_{k \in K} Q^k \varphi \leq \varphi$. On the other hand, for each $k \in K$, $Q^k \varphi \geq \varphi$ by Lemma 3.2, since φ is subharmonic by Proposition 2.1(b). Q.E.D.

LEMMA 3.6. $\gamma \leq \varphi$.

Proof. We will show, using transfinite induction, that $\gamma_a \leq \varphi$ for each ordinal a . For $a=0$, $\varphi \geq \chi_X = \gamma_0$ (see Proposition 2.1(a)). Assume $\gamma_b \leq \varphi$ for each $b < a$; then $Q^k \gamma_b \leq Q^k \varphi = \varphi$ by the preceding proposition, thus $\gamma_a^k \leq \varphi$ for each $k \in K$, so that $\gamma_a \leq \varphi$. Thus $\gamma = \gamma_c \leq \varphi$. Q.E.D.

LEMMA 3.7. $\varphi \leq \gamma$.

Proof. Note that, since $\gamma = \max_{k \in K} Q^k \gamma$, we have

$$(*) \quad \gamma(i) \geq \inf_{\tau \in T(i, k)} E_\tau(\gamma), \quad i \in I, \quad k \in K,$$

(actually, with equality holding for at least one k , although we will not make use of this fact). Let $i \in I$ be given. Choose $\varepsilon > 0$ and a sequence $\varepsilon_n \downarrow 0$ such that $\sum_n \varepsilon_n = \varepsilon$. Let $\{k_n\}_{n \geq 1}$ be a fixed sequence of processes in which each $k \in K$ appears infinitely many times. We will use (*) to construct a fair schedule σ starting at i by building it layer-by-layer from subschedules, as follows: Suppose that the first n layers of σ have already been constructed, the union of which being some subschedule τ_n starting at i (initially, τ_0 is "empty"). The $(n+1)$ th layer of σ is defined by appending to τ_n at each of its end nodes j a subschedule $\bar{\rho}_j \in T(j, k_{n+1})$ such that

$$\gamma(j) \geq E_{\bar{\rho}_j}(\gamma) - \varepsilon_n$$

(such a subschedule exists by (*)). Repeating this process inductively, we obtain the required (infinite) schedule σ , which is fair by our choice of the sequence $\{k_n\}_{n \geq 1}$.

Let $\{N_n\}_{n \geq 0}$ be the increasing sequence of stopping times defined by our construction; namely—the n th layer (i.e., τ_n) ends at N_n (in particular $N_0 \equiv 0$). For each $n \geq 0$ define the function

$$g_n(\pi) = \gamma(\pi_{N_n}), \quad \pi \in H^*(i);$$

in particular, $g_0 \equiv \gamma(i)$. By the choice of the subschedules ρ_j we have

$$(**) \quad g_n \geq E_\sigma(g_{n+1} | \pi_{N_n}) - \varepsilon_{n+1}, \quad n \geq 0.$$

Hence, the sequence of functions $\{g'_n\}_{n \geq 0}$ given by

$$g'_n \equiv g_n - \sum_{m=1}^n \varepsilon_m, \quad n \geq 0$$

forms a supermartingale, which is bounded between 1 and $-\varepsilon$. Hence it converges almost surely to a limit g'_∞ , so that $\{g_n\}$ converges almost surely to the function

$$g_\infty \equiv g'_\infty + \sum_{m=1}^{\infty} \varepsilon_m = g'_\infty + \varepsilon.$$

Note that $\gamma|_X \equiv 1$; thus, if X is reached along π , then $g_\infty(\pi) = 1$, because for all sufficiently large n we will have $g_n(\pi) = 1$. Hence, by (**),

$$\begin{aligned} \gamma(i) &= g'_0 \geq E_\sigma(g'_\infty) = E_\sigma(g_\infty) - \varepsilon \\ &\geq \mu_\sigma(g_\infty = 1) - \varepsilon \\ &\geq \mu_\sigma(X \text{ is reached}) - \varepsilon \\ &= E_\sigma(\chi_X) - \varepsilon \geq \varphi(i) - \varepsilon. \end{aligned}$$

Since ε was arbitrary, the proof is complete. Q.E.D.

Thus we have shown

THEOREM 3.8. $\varphi = \gamma$.

Next, we give an example of explicit calculation of φ as the limit of the φ -iterates.

Example 2. Let $I = \mathbb{N}$, $X = \{0\}$, $K = \{1, 2\}$ such that each process is a random walk on I (with X absorbing). It turns out that a fair interaction of two random walks, under the worst kind of schedules, yields essentially the same absorption probabilities as those yielded by the "worse" of the two walks alone. We exhibit here one simple case:

$$\begin{aligned} P_{i,i-1}^1 &= \frac{1}{3}, & P_{i,i+1}^1 &= \frac{2}{3}, & i &\geq 1, \\ P_{i,i-1}^2 &= 1, & i &\geq 1. \end{aligned}$$

It can be inductively shown that the φ -iterates for this program are

$$\gamma_{n-1}(i) = \begin{cases} \xi_{n-i} / \xi_n, & 0 \leq i \leq n, \\ 0, & i > n, \end{cases} \quad n \geq 1,$$

where $\xi_i = 2^i - 1$, $i \geq 0$, and

$$\gamma_\omega(i) = \gamma_{\omega+1}(i) = \varphi(i) = \frac{1}{2^i}, \quad i \geq 0.$$

Comparison with the iterates $\hat{\gamma}_n$ and their limit $\hat{\gamma}_\omega$ for the case in which only process 1 is activated shows that $\gamma_\omega = \hat{\gamma}_\omega$ but $\gamma_n > \hat{\gamma}_n$ for each finite n . Thus the fair interleaving of process 2 with process 1 increases the probability of convergence under any finite number of fairness constraints, but does not affect the overall (worst-case) convergence probability.

4. Characterizations of φ . This section contains the main results of the paper. Using the machinery developed in §§ 2 and 3, we will derive several characterizations of φ , which provide a variety of rather simple techniques for its calculation, or for deriving various properties of this function. Obviously, the most important such

property is whether $\varphi \equiv 1$ (i.e., whether the program terminates almost surely from any initial state). Relaxation of the characterizations of φ given here will enable us to derive necessary and sufficient conditions for program termination, and these conditions are presented in § 5.

THEOREM 4.1. (a) φ is the smallest fixpoint of the equation

$$\varphi = Q\varphi$$

which is $\geq \chi_X$.

(b) φ is the smallest simultaneous solution of the equations

$$\varphi = Q^k \varphi \quad \text{for each } k \in K,$$

which is $\geq \chi_X$.

Proof. By Propositions 3.5 and 2.1(a), $\varphi = Q\varphi = Q^k \varphi$ for each $k \in K$, and $\varphi \geq \chi_X$. To prove (a) we repeat the argument used in the proof of Lemma 3.6. That is, let $\psi \geq \chi_X$ be such that $\psi = Q\psi$. Then $\psi \geq \gamma_0$, thus $\psi = Q\psi \geq Q\gamma_0 = \gamma_1$, and by transfinite induction $\psi \geq \gamma_a$ for each ordinal a , thus $\psi \geq \gamma = \varphi$. As for (b), note that $\psi = Q^k \psi$ for all $k \in K$ implies $\psi = Q\psi$, and then use (a). Q.E.D.

Next we restate the second assertion of Theorem 4.1 in a manner which makes it more convenient for actual calculation of φ .

DEFINITION. Let α be a real function on I . We say that α has *property (A)* if the following are satisfied:

(A.1) $\alpha|_X \equiv 1$;

(A.2) α is subharmonic;

(A.3) for each $k \in K$ the only subharmonic function lying between α and $P^k \alpha$ is α itself.

(Note that the constant function 1 has property (A).)

THEOREM 4.2. φ is the smallest nonnegative function on I having property (A) (i.e., if $\alpha \geq 0$ satisfies (A), then $\alpha(i) \geq \varphi(i)$ for each $i \in I$).

Proof. By Lemmata 3.1 and 3.2, (A.2) and (A.3) imply $\alpha = Q^k \alpha$ for all $k \in K$, or $\alpha = Q\alpha$. We now use Theorem 4.1. Q.E.D.

Theorem 4.2 suggests the following procedure for computing φ : Take any nonnegative subharmonic function $\alpha \geq \chi_X$. For each $k \in K$ compute the largest subharmonic function which is $\geq P^k \alpha$, and require that it coincide with α . Find the general solution of these constraints, and obtain φ as the smallest such solution. Later on in this section we will use this procedure to compute φ for several exemplary programs, and show that this technique is quite feasible in practice.

Put

(A.2') α is min-harmonic,

and let property (A') be defined as the conjunction of (A.1), (A.2') and (A.3). Then we also have

COROLLARY 4.3. φ is the smallest nonnegative function having property (A').

Proof. Immediate, since φ itself is min-harmonic, by Proposition 2.1(b), and every min-harmonic function is also subharmonic. Q.E.D.

Remark. In carrying out the calculations of the procedure just outlined, it may sometimes be more convenient to employ the "1-complement" version of Theorem 4.2; that is, instead of computing φ we compute the function $\psi \equiv 1 - \varphi$, which is then the largest function ≤ 1 which is a fixpoint of the equation

$$\psi(i) = \min_{k \in K} \sup_{\tau \in T(i,k)} E_\tau(\psi)$$

or, alternatively, is the largest function $\beta \leq 1$ having *property (B)*, defined as

$$(B.1) \quad \beta|_X \equiv 0;$$

$$(B.2) \quad \beta \text{ is superharmonic, i.e., } \beta \geq P^k \beta \text{ for each } k \in K;$$

(B.3) for each $k \in K$, the only superharmonic function between $P^k \beta$ and β is β itself.

(Again, we can replace (B.2) by (B.2'), namely require that β be *max-harmonic*, that is $\beta = \max_{k \in K} P^k \beta$.)

The usefulness of this complementation lies in the fact that property (B) is positively homogeneous (i.e., β satisfies (B) implies $\lambda \beta$ satisfies (B) for every $\lambda > 0$, where $(\lambda \beta)(i) \equiv \lambda \cdot \beta(i)$); note that (A) was not such (due to (A.1)). For example, we obtain

COROLLARY 4.4. $\varphi \equiv 1$ if and only if no bounded function having some positive entries has property (B).

Proof. Assume β satisfies (B) so that $\lambda \equiv \sup_{i \in I} \beta(i) < \infty$ and is positive. Then $(1/\lambda)\beta$ also satisfies (B) and is ≤ 1 . Q.E.D.

We can also give now a second short proof of the Zero-One Law for φ ; namely, that $\inf_{i \in I} \varphi(i)$ is either 0 or 1 (see Theorem 2.3; however, the original proof is more elementary).

Second proof of the zero-one law (Theorem 2.3). Let $\psi = 1 - \varphi$ and put $\lambda = \sup_{i \in I} \psi(i)$. If $0 < \lambda < 1$, then the function $\psi' = (1/\lambda)\psi$ is larger than ψ , satisfies (B), and is ≤ 1 —contradicting the fact that ψ is the largest such function.

Examples. We will now apply the techniques presented in this section to several programs, to compute the function φ for each of these programs. These examples include two programs with finite state spaces (which had already been analyzed in a preceding paper [HSP] by different special techniques developed there for finite-state programs), and another program having an infinite state space.

Example 3. Let $K = \{1, 2\}$. The following program arises in an analysis of freedom from lockout in a simple synchronization protocol (cf. [HSP, Example 1] for details). Using a notation slightly different from that of [HSP], we have $I = X \cup \{i_1, i_2, i_3, i_4\}$, with nonzero transition probabilities

$$P_{i_1, X}^1 = P_{i_1, i_1}^2 = 1,$$

$$P_{i_2, i_1}^1 = P_{i_2, i_4}^1 = P_{i_2, i_1}^2 = P_{i_2, i_4}^2 = \frac{1}{2},$$

$$P_{i_3, i_3}^1 = P_{i_3, i_2}^2 = 1,$$

$$P_{i_4, i_4}^1 = P_{i_4, i_3}^2 = 1.$$

To compute φ , we first write down the form of the general subharmonic function which is 1 on X . Such a function $\alpha = (\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ (where α_t is a shorthand for $\alpha(i_t)$, $1 \leq t \leq 4$) must satisfy

$$\alpha_1 \leq 1,$$

$$\alpha_2 \leq \frac{1}{2}\alpha_1 + \frac{1}{2}\alpha_4,$$

$$\alpha_3 \leq \alpha_2,$$

$$\alpha_4 \leq \alpha_3.$$

Next, we spell out condition (A.3) for such an α : First consider $k = 1$. It is easily checked that the function

$$P^1 \alpha = (1, \frac{1}{2}\alpha_1 + \frac{1}{2}\alpha_4, \alpha_3, \alpha_4)$$

is also subharmonic. Hence we must have $\alpha = P^1\alpha$, i.e.,

$$\alpha_1 = 1, \quad \alpha_2 = \frac{1}{2}\alpha_1 + \frac{1}{2}\alpha_4 = \frac{1}{2} + \frac{1}{2}\alpha_4.$$

Similarly, for $k=2$ we have

$$P^2\alpha = (\alpha_1, \frac{1}{2}\alpha_1 + \frac{1}{2}\alpha_4, \alpha_2, \alpha_3),$$

which is also seen to be subharmonic. Hence $\alpha = P^2\alpha$, i.e.,

$$\alpha_2 = \alpha_3, \quad \alpha_3 = \alpha_4.$$

Thus we have $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 1$. That is, the only—and thus, the smallest—function satisfying (A) is $\varphi \equiv 1$.

Example 4. This example is also taken from [HSP], and arises in the analysis of another synchronization protocol. Here $K = \{1, 2\}$, $I = X \cup \{i_1, \dots, i_6\}$, and the transition probability matrices are

	P^1							P^2						
	i_1	i_2	i_3	i_4	i_5	i_6	X	i_1	i_2	i_3	i_4	i_5	i_6	X
i_1	$\frac{1}{2}$						$\frac{1}{2}$	$\frac{1}{2}$				$\frac{1}{2}$		
i_2		$\frac{1}{2}$		$\frac{1}{2}$				$\frac{1}{2}$				$\frac{1}{2}$		
i_3	$\frac{1}{2}$						$\frac{1}{2}$			$\frac{1}{2}$	$\frac{1}{2}$			
i_4		$\frac{1}{2}$		$\frac{1}{2}$						$\frac{1}{2}$	$\frac{1}{2}$			
i_5					$\frac{1}{2}$	$\frac{1}{2}$		1						
i_6					$\frac{1}{2}$	$\frac{1}{2}$				1				

It is straightforward to check that a general subharmonic function $\alpha = (\alpha_1, \dots, \alpha_6)$ which is 1 on X must satisfy

$$\alpha_1 = \alpha_2 = \dots = \alpha_6 \leq 1.$$

It now follows that (A.3) holds for each such function α , because any function constant on $I - X$ and lying between α and $P^1\alpha$ (resp. $P^2\alpha$) must coincide with α (since $P^1\alpha$ (resp. $P^2\alpha$) coincides with α at some of these states). Thus φ , which is the smallest nonnegative such function, is χ_X .

Example 5 ("The Two Combs"). Let $K = \{1, 2\}$, $I = X \cup \mathbb{Z}$ (where \mathbb{Z} denotes the set of signed integers); the nonzero transition probabilities are

$$\begin{aligned} P_{n,n+1}^1 &= p_n, & P_{n,X}^1 &= p'_n = 1 - p_n, \\ P_{n,n-1}^2 &= q_n, & P_{n,X}^2 &= q'_n = 1 - q_n, \end{aligned} \quad n \in \mathbb{Z}.$$

To avoid degeneracy, we assume that $0 < p_n q_{n+1} < 1$ for each $n \in \mathbb{Z}$. Denote, for $n \in \mathbb{Z}$,

$$P_n = \prod_{m=n}^{\infty} p_m, \quad Q_n = \prod_{m=-\infty}^n q_m.$$

Denote by (C^+) the condition

$$(C^+) \quad \prod_{n>0} p_n > 0 \quad \text{and} \quad \limsup_{n \rightarrow \infty} q_n = 1,$$

and by (C^-) the condition

$$(C^-) \quad \prod_{n<0} q_n > 0 \quad \text{and} \quad \limsup_{n \rightarrow -\infty} p_n = 1.$$

- PROPOSITION 4.5. (a) If neither (C^+) nor (C^-) hold, then $\varphi \equiv 1$.
 (b) If (C^+) holds but (C^-) does not hold, then $\varphi_n = 1 - P_n$, $n \in \mathbb{Z}$.
 (c) If (C^-) holds but (C^+) does not hold, then $\varphi_n = 1 - Q_n$, $n \in \mathbb{Z}$.
 (d) If both (C^+) and (C^-) hold, then $\varphi_n = 1 - \max\{P_n, Q_n\}$, $n \in \mathbb{Z}$.

Proof. It will be more convenient to work in "1-complement" mode, calculating $\psi \equiv 1 - \varphi$, and using property (B). The calculation of ψ proceeds through the following steps (for details, see [HS]).

- (1) If $\psi_n = 0$ for some $n \in \mathbb{Z}$, then $\psi \equiv 0$.
- (2) Put $\psi^1 = P^1\psi$, $\psi^2 = P^2\psi$; if $\psi \equiv 0$ then it is impossible to have for some $n \in \mathbb{Z}$, $\psi_n^1 = \psi_n$ and $\psi_{n+1}^2 = \psi_{n+1}$.
- (3) $\psi_n > \psi_n^1 \Rightarrow \psi_m > \psi_m^1$ for each $m \leq n$, and $\psi_n > \psi_n^2 \Rightarrow \psi_m > \psi_m^2$ for each $m \geq n$.
- (4) Thus only the following four cases are possible:
 - (a) $\psi = \psi^1 = \psi^2 \equiv 0$;
 - (b) $\psi_n = \psi_n^1 > \psi_n^2$ for each $n \in \mathbb{Z}$;
 - (c) $\psi_n = \psi_n^2 > \psi_n^1$ for each $n \in \mathbb{Z}$;
 - (d) there exists $n_0 \in \mathbb{Z}$ such that $\psi_n = \psi_n^1 > \psi_n^2$ for each $n > n_0$, and $\psi_n = \psi_n^2 < \psi_n^1$ for each $n < n_0$.
- (5) Suppose $\psi > 0$. If, for some $n_0 \in \mathbb{Z}$, $\psi_n = \psi_n^1$ for each $n > n_0$, then $\prod_{n > n_0} p_n > 0$. Similarly, if $\psi_n = \psi_n^2$ for each $n < n_0$, then $\prod_{n < n_0} q_n > 0$.
- (6) In particular, if $\prod_{n > 0} p_n = \prod_{n < 0} q_n = 0$, then $\psi \equiv 0$.
- (7) Suppose $\psi > 0$. If, for some $n_0 \in \mathbb{Z}$, $\psi_n = \psi_n^1$ for each $n > n_0$, then $\limsup_{n \rightarrow \infty} p_n q_{n+1} = 1$. Similarly, if $\psi_n = \psi_n^2$ for each $n < n_0$, then $\limsup_{n \rightarrow -\infty} p_n q_{n+1} = 1$.
- (8) The following is a partial converse to (7): Let $\psi > 0$ be any max-harmonic function satisfying $\psi_n = \psi_n^1$ for each $n > n_0$, and suppose $\limsup_{n \rightarrow \infty} p_n q_{n+1} = 1$. Then the unique superharmonic function lying between ψ and ψ^2 is ψ itself. A similar statement holds if $\psi_n = \psi_n^2$ for each $n < n_0$ and $\limsup_{n \rightarrow -\infty} p_n q_{n+1} = 1$.
- (9) In case (b) condition (C^+) holds; in case (c) condition (C^-) holds; and in case (d) both conditions (C^+) and (C^-) hold.
- (10) Conversely, if (C^+) holds but (C^-) does not, then case (b) must occur. Similarly, if (C^-) holds but not (C^+) , then case (c) must occur.
- (11) Finally, if both (C^+) and (C^-) hold, then case (d) must occur. Q.E.D.

5. Verification of program termination. The results developed in the two preceding sections provide us with methods for calculating the function φ for any concurrent probabilistic program. However, in many applications the only question of interest concerning φ is whether $\varphi \equiv 1$, i.e., whether the program terminates almost surely from any initial state under any fair schedule. In this section we will present several characterizations of program termination, the first two of which are straightforward specializations of the general results of the preceding sections, while the third involves a somewhat different approach, generalizing that used in [HSP] for finite state spaces.

PROPOSITION 5.1. $\varphi \equiv 1$ if and only if no min-harmonic function smaller than 1 has property (A).

Proof. See Corollary 4.3.

PROPOSITION 5.2. $\varphi \equiv 1$ if and only if there exist an ordinal c and transfinite sequences of functions $\{\delta_a^k\}_{a \leq c}$, $k \in K$, and $\{\delta_a\}_{a \leq c}$ having the following properties:

- (1) $\delta_0 = \delta_0^k = \chi_X$, $k \in K$;
- (2) δ_a^k is subharmonic for each $a \leq c$ and each $k \in K$;
- (3) $\delta_a \leq \max_{k \in K} \delta_a^k$, $a \leq c$;
- (4) $\delta_{a+1}^k \leq P^k \delta_a^k$, $k \in K$, $a < c$;
- (5) $\delta_a^k \leq \sup_{b < a} \delta_b^k$, for limit ordinals a , and $k \in K$;
- (6) $\inf_{i \in I} \delta_c(i) > 0$.

Proof. If $\varphi \equiv 1$ then the φ -iterates can be taken as the δ 's. Conversely, if such sequences of functions are given, then by transfinite induction $\delta_a \leq \varphi$ for each ordinal a . In particular $\delta_c \leq \varphi$, so that $\inf_{i \in I} \varphi(i) > 0$, and by the zero-one law (Theorem 2.3) we must have $\varphi \equiv 1$. Q.E.D.

Our next characterization of program termination generalizes one of the characterizations given in [HSP] for finite-state programs. Intuitively speaking, if the program does not always terminate, then there must exist some "ergodic structure" of nonterminating states, through which an "adversary" fair scheduler can iterate forever without reaching X . Unlike the case of a finite state space, where such a structure was a single " K -ergodic" set, ergodicity in general state spaces is a much more complex notion, and is defined as follows.

DEFINITION. A K -ergodic chain is a nonincreasing sequence $\{E_n\}_{n \geq 1}$ of nonempty subsets of $X^c \equiv I - X$ such that

$$\lim_{n \rightarrow \infty} \sup_{\substack{i \in E_n \\ m \geq 1}} (Q\chi_{E_m^c})(i) = 0.$$

In other words, let $n \geq 1$, $i \in E_n$, $m \geq 1$ and $k \in K$ be given. Then there exists a subschedule in $T(i, k)$ which reaches E_m with probability tending to 1 uniformly as $n \rightarrow \infty$. That is, without losing too much probability, we can reach any of the sets E_m from any state in E_n after scheduling any required process.

THEOREM 5.3. $\varphi \equiv 1$ if and only if $I - X$ does not contain any K -ergodic chain.

Before proving this theorem, we need two lemmata.

LEMMA 5.4. Let $\delta > 0$, and define $D = \{i \in I: \varphi(i) \geq \delta\}$. Then $\varphi \geq Q\chi_D$.

Proof. Let $i \in I$, $k \in K$, and $\sigma \in \Sigma F(i)$. For each $n \geq 1$ define a stopping time N_n on $H^*(i)$ so that $N_n(\pi)$ is the n th time k has been scheduled along π ; note that $\{N_n\}_{n \geq 1}$ is an increasing sequence of μ_0 -a.s. finite stopping times, whose limit is $+\infty$. For each $n \geq 1$ the subschedule $\tau_n = (\sigma, N_n) \in T(i, k)$, so that

$$(Q^k\chi_D)(i) \leq E_{\tau_n}(\chi_D) = \mu_\sigma\{\varphi(i_{N_n}) \geq \delta\}.$$

Consider the sequence of functions $\{f_m\}_{m \geq 0}$, defined by $f_m(\pi) = \varphi(i_m)$, $m \geq 0$, $\pi \in H^*(i)$. By Theorem 2.3 $\{f_m\}$ converges a.s. to a limit f_∞ , such that $f_\infty(\pi)$ is 1 if X is reached, and is otherwise 0. Therefore we also have $\varphi(i_{N_n}) \rightarrow f_\infty$ a.s. as $n \rightarrow \infty$, so that

$$\mu_\sigma\{f_\infty \geq \delta\} \geq \overline{\lim}_{n \rightarrow \infty} \mu_\sigma\{\varphi(i_{N_n}) \geq \delta\} \geq (Q^k\chi_D)(i).$$

Since $\delta > 0$, we have $f_\infty(\pi) \geq \delta$ if and only if $f_\infty(\pi) = 1$, or, alternatively, if and only if X is reached along π . Thus

$$E_\sigma(\chi_X) = \mu_\sigma\{f_\infty \geq \delta\} \geq (Q^k\chi_D)(i),$$

from which our assertion follows. Q.E.D.

LEMMA 5.5. Let $\{G_n\}_{n \geq 1}$ be a nondecreasing sequence of subsets of I , all of which contain X , and let $\{\varepsilon_n\}_{n \geq 1}$ be a sequence of positive numbers converging to 0. Suppose that

$$Q\chi_{G_m}(i) \leq \varepsilon_n$$

for each $m, n \geq 1$ and each $i \in G_n^c$. Then

$$\varphi \leq \sup_m Q\chi_{G_m}.$$

Proof. Put $\beta \equiv \sup_m Q\chi_{G_m}$. The above assumption concerning $\{G_n\}$ can be restated as

$$Q\chi_{G_m} \leq \varepsilon_n \cdot \chi_{G_n^c} + 1 \cdot \chi_{G_n} = \varepsilon_n + (1 - \varepsilon_n)\chi_{G_n},$$

for each $m, n \geq 1$. This implies

$$\beta \leq \varepsilon_n + (1 - \varepsilon_n)\chi_{G_n}, \quad n \geq 1,$$

and thus

$$Q\beta \leq Q(\varepsilon_n + (1 - \varepsilon_n)\chi_{G_n}), \quad n \geq 1.$$

However, it is easily checked that for any scalars $a, b > 0$ and any nonnegative function α we have

$$Q(a + b\alpha) = a + bQ\alpha.$$

Hence,

$$Q\beta \leq \varepsilon_n + (1 - \varepsilon_n)Q\chi_{G_n} \leq \varepsilon_n + (1 - \varepsilon_n)\beta,$$

for each $n \geq 1$. Letting $n \rightarrow \infty$, we obtain $Q\beta \leq \beta$. But β is subharmonic (as a supremum of subharmonic functions), thus $\beta = Q\beta$ (see Lemmata 3.1 and 3.2), implying $\varphi \leq \beta$ by Theorem 4.1. Q.E.D.

Proof of Theorem 5.3. The theorem is now an easy consequence of the last two lemmata. For example, if φ is not identically 1, then, by Lemma 5.4, the collection $\{E_n\}_{n \geq 1}$ is a K -ergodic chain, where

$$E_n = \left\{ i \in I : \varphi(i) < \frac{1}{n} \right\}.$$

The converse statement follows similarly from Lemma 5.5 (for more details, see [HS]). Q.E.D.

Example 5 revisited. Consider the three cases in the example of "the two combs" in which $\varphi < 1$. It is easily verified that in case (b) the chain $E_n^+ = \{i : i \geq n\}$, $n \geq 1$, is K -ergodic; similarly, in case (c) the chain $E_n^- = \{i : i \leq -n\}$, $n \geq 1$, is K -ergodic; and in case (d) both these chains are K -ergodic.

Remark. In the case of a finite state space, ergodicity is manifested in a single set (see [HSP]). In an analogous manner, we could have considered here the set $E = \{i \in I : \varphi(i) = 0\}$ as a natural candidate for being K -ergodic (that is, consider the constant chain $E_n = E$, $n \geq 1$). There are two problems, however, with this approach, which make it infeasible for general state spaces. One problem is that E may be empty (as is indeed the case in Example 5 just considered). Moreover, even if E is not empty, it may happen that, starting from some $i \in E$, we never reach E again, but instead reach states j at which $\varphi(j)$ is arbitrarily small, but positive. Thus, for general state spaces ergodicity must be defined in terms of an infinite chain of sets rather than in terms of a single set. (In the finite case, though, any K -ergodic chain must reduce to a constant set from a certain index on.) Note that this phenomenon occurs in Markov chains as well.

We conclude this section with a further property of φ .

PROPOSITION 5.6. *There exists a nondecreasing sequence $\{D_n\}_{n=1}^{\infty}$ of subsets of I such that*

- (1) $\varphi|_{D^c} \equiv 0$, where $D = \bigcup_{n=1}^{\infty} D_n$,
- (2) $\varphi = \lim_{n \rightarrow \infty} Q\chi_{D_n}$.

Proof. Put $D_n = \{i : \varphi(i) \geq 1/n\}$. Q.E.D.

Note that in the case $\varphi \equiv 1$ we can take $D_n = I$ for all n . Moreover,

COROLLARY 5.7. *If I is a finite set, then there exists $D \subset I$ such that*

- (1) $\varphi|_{D^c} \equiv 0$,
- (2) $\varphi = Q\chi_D$.

6. Comparison with the deterministic and the nondeterministic cases. In this section we consider the special case in which each process is deterministic; programs with nondeterministic processes are also included, since any such program can be simulated by a deterministic one involving additional processes. (For example, suppose that one of the processes k_1 makes a nondeterministic choice from some set A of alternatives; the same behavior can be achieved by introducing a new shared variable ν which k_1 sets to some value in A prior to making the choice, and by introducing another process k_2 whose only action is to iterate ν over the set A . k_1 then makes its choice deterministically, depending on the current value of ν . Thus the nondeterminism is now transferred to the scheduler—the final choice depends on how many times k_2 has been scheduled in between.) Thus, by specializing the various equivalent criteria for program termination developed so far in this paper to the deterministic case, we can obtain similar criteria for the termination of deterministic (or nondeterministic) concurrent programs. As it turns out, the criterion obtained in this way from the characterization of φ as the limit of the φ -iterates (Theorem 3.8, Proposition 5.2) essentially coincides with the known criterion of Lehmann, Pnueli and Stavi [LPS]. On the other hand, specialization of Theorem 4.2 leads to a new characterization for deterministic and nondeterministic termination. So as not to make this characterization appear too deep, we provide a direct nonprobabilistic proof of its validity.

We begin by observing that in the deterministic case all transitions have probability 0 or 1, so that each of the operators Q^k , $k \in K$, and Q , when applied to a function which takes only the values 0, 1 (i.e., a characteristic function of some subset of I) yields a similar function. Hence each of the φ -iterates γ_a^k , $k \in K$, (resp. γ_a) is a characteristic function of the form $\chi_{G_a^k}$ (resp. χ_{G_a}). Note also that a characteristic function χ_A is subharmonic if and only if for each $k \in K$ and each $i \in A$ the (unique) k -transition from i is to a state in A , i.e., there are no transitions from states in A to states outside A . Hence, spelling out the conditions in Proposition 5.2 in terms of the subsets of I corresponding to the functions appearing there, we obtain the following.

COROLLARY 6.1. *A deterministic program terminates if and only if there exist transfinite (increasing) sequences $\{G_a^k\}_{a \geq 0}$, $k \in K$, and $\{G_a\}_{a \geq 0}$ of subsets of I having the following properties:*

- (1) $G_0 = G_0^k = X$, $k \in K$;
- (2) *there are no transitions from states in G_a^k to states outside G_a^k , for each ordinal a and each $k \in K$;*
- (3) $G_a = \bigcup_{k \in K} G_a^k$, for each ordinal a ;
- (4) *for each $k \in K$ and each ordinal a , all k -transitions from states in G_{a+1}^k are to states in G_a ;*
- (5) $G_a^k = \bigcup_{b < a} G_b^k$, for each limit ordinal a and each $k \in K$;
- (6) *there exists an ordinal c such that $G_c = I$.*

These conditions, however, are merely a rephrasing of the characterization for termination of “just” programs given by Lehmann, Pnueli and Stavi in [LPS]. To see this, define a function ρ from I to the ordinals by

$$\rho(i) = \min \{a : i \in G_a\}, \quad i \in I,$$

and a function $h : I \rightarrow K$ which maps each $i \in G_{\rho(i)}$ to some $k \in K$ such that $i \in G_{\rho(i)}^k$. Then it is easily checked that these functions satisfy the conditions in [LPS] for just termination, i.e.: the “ranking” map ρ never increases during execution; activating process $h(i)$ at state i always strictly decreases the value of ρ ; and h remains unchanged

as long as ρ does not decrease. By the remark made in the beginning of this section, it is easily seen that Corollary 6.1 also applies to nondeterministic (nonprobabilistic) programs. Thus our Proposition 5.2 generalizes the results of [LPS] to the probabilistic case.

Next we specialize Theorem 4.2 to the deterministic case. As noted above, we can identify subharmonic functions with subsets A of I such that there are no transitions from A to A^c . This notion is formalized in the following.

DEFINITION. (a) A *cut* (I_0, I_1) is a partition of I into two disjoint subsets I_0, I_1 with $I_0 \cup I_1 = I$, such that $X \subset I_1$ and such that there are no transitions from I_1 to I_0 .

(b) For each cut (I_0, I_1) and each $k \in K$ put

$$I_0^k = \{i \in I_0 : \text{the } k\text{-transition from } i \text{ is into } I_0\}.$$

(Note that χ_{I_1} is subharmonic and $\cong \chi_X$, and that $P^k \chi_{I_0} = \chi_{I_0^k}$.) Using these notations, Theorem 4.2 translates into the following theorem (which merely states that (A.3) does not hold for any subharmonic function < 1).

THEOREM 6.2. *A deterministic program terminates if and only if for each nontrivial cut (I_0, I_1) (i.e., $I_0 \neq \emptyset$) there exists $k \in K$ and another cut (J_0, J_1) such that*

$$I_0^k \subseteq J_0 \subsetneq I_0.$$

Proof. As this result (and its appropriate generalization to the nondeterministic case) is new, and may be of interest in its own right, we provide here a direct proof of this characterization, which does not use the probabilistic techniques developed in this paper.

Assume first that the condition of the theorem does not hold, i.e., that there exists a nontrivial cut (I_0, I_1) such that for each $k \in K$ and each set $I_0^k \subseteq J \subsetneq I_0$, the pair (J, J^c) is not a cut, that is, there exist transitions from J^c to J (these transitions can only be from states in $I_0 - J$). Let $i \in I_0$, and let $F(i)$ denote the set of all states in I_0 (including i) reachable from i by some finite sequence of process activations. We claim that for each $k \in K$, $F(i)$ intersects I_0^k . For otherwise, put $J = I_0 - F(i)$, so that $I_0^k \subset J \subsetneq I_0$. By our assumption there exist transitions from $I_0 - J = F(i)$ into J , which contradicts the definition of $F(i)$. This implies that I_0 is ergodic, i.e., that there exists a fair schedule σ which can keep the program in I_0 forever. To prove this it suffices to show that for each $i \in I_0$ and each $k \in K$ there exists a finite scheduling sequence starting at i and ending by scheduling k and reaching a state in I_0 . Since $F(i) \cap I_0^k \neq \emptyset$, take a finite sequence of process activations which takes the program from i into some state in $F(i) \cap I_0^k$, and then schedule k , thereby reaching a state in I_0 .

Next suppose that the condition of the theorem does hold. We will construct a "ranking" function ρ from I to the ordinals and an "assistance" function $h: I \rightarrow K$ which will satisfy the conditions of [LPS] for program termination. These functions will be constructed in the following transfinite inductive manner. Put initially $\rho|_X = 0$, and define $h|_X$ in an arbitrary manner. Suppose inductively that ρ and h have already been defined on some subset M of I such that (L, M) is a cut (where $L = M^c$). By the above condition, there exists $k \in K$ and another cut (J, J^c) such that $L^k \subseteq J \subsetneq L$. Put $H = L - J \neq \emptyset$, and define $\rho|_H = 1 + \sup \rho|_M$, and $h|_H = k$. Note that since (J, J^c) is a cut, there are no transitions from H into J , and since H is disjoint from L^k , each k -transition from H is into M . Repeating this construction transfinitely, we obtain everywhere-defined functions ρ and h whose properties imply by [LPS] that the program terminates. Q.E.D.

Remark. To obtain the nondeterministic version of Theorem 7.2, define the sets I_0^k by

$$I_0^k = \{i \in I_0 : \text{there exists a } k\text{-transition from } i \text{ into } I_0\},$$

and leave all other definitions and assertions unchanged.

Example 1 revisited. Let us apply Theorem 6.2 to the deterministic program given in Example 1, § 3. As is easily checked, a cut (I_0, I_1) of I must be one of the following three types:

- (a) $I_0 = [0, n] \times \{2\}$, for some $n \in \mathbb{N}$;
- (b) $I_0 = \mathbb{N} \times \{2\}$;
- (c) $I_0 = (\mathbb{N} \times \{2\}) \cup ([n, \infty] \times \{1\})$ for some $n \geq 1$.

In cases (a) and (b), $I_0^2 = \emptyset$, so that $(I_0^2, (I_0^2)^c) = (\emptyset, I)$ is a cut satisfying the condition of Theorem 6.2. In case (c), $I_0^1 = (\mathbb{N} \times \{2\}) \cup ([n+1, \infty] \times \{1\})$, so that $(I_0^1, (I_0^1)^c)$ is the required cut. Thus program termination is ensured by Theorem 6.2.

7. Programs with finite state space. In a preceding paper [HSP], the special case of concurrent probabilistic programs with finite state spaces has been analyzed, and a characterization of almost sure program termination in terms of the existence of a certain decomposition of the state space has been obtained. In this section we show how to obtain this characterization from the general theory developed so far in this paper.

Let us now assume that I is finite, and that $\varphi \equiv 1$. We will obtain a decomposition of I into (finitely many) disjoint sets $\{I_m\}_{m \geq 0}$ such that the following properties hold (here we use the notation $P_{i,E}^k \equiv \sum_{j \in E} P_{ij}^k$):

- (a) $I_0 = X$;
- (b) for each $m \geq 1$, each $i \in I_m$ and each $k \in K$, if $P_{i,J_m}^k = 0$ then $P_{i,I_m}^k = 1$ (where $J_m = \bigcup_{s < m} I_s$);
- (c) for each $m \geq 1$ there exists $k = k(m) \in K$ such that for each $i \in I_m$, $P_{i,J_m}^k > 0$. (This is the decomposition obtained in [HSP].)

To obtain it, we proceed inductively. Initially put $I_0 = X$. Suppose I_0, \dots, I_{m-1} have already been constructed. Let $J_m = \bigcup_{s < m} I_s$ and $H_m = J_m^c$. If $H_m = \emptyset$, the decomposition is complete. Otherwise, $\chi_{J_m} \not\equiv 1 = \varphi$, thus we cannot have $\varphi \leq \chi_{J_m}$, therefore $Q\chi_{J_m} \not\leq \chi_{J_m}$ (indeed, for any $\alpha \geq \chi_X$, if $Q\alpha \leq \alpha$ then $\varphi \leq Q\alpha \leq \alpha$). Thus there exists $k \in K$ such that $\delta \equiv Q_{J_m}^k \not\leq \chi_{J_m}$. Define

$$c = \max_{j \in H_m} \delta(j) \quad \text{and} \quad I_m = \{i \in H_m : \delta(i) = c\}.$$

Note that $c > 0$, for otherwise $\delta|_{H_m} \equiv 0$, so that $\delta \leq \chi_{J_m}$, which is impossible.

It can be now seen that conditions (b) and (c) hold for I_m .

Remarks. (1) The converse statement, namely that the existence of such a decomposition implies that $\varphi \equiv 1$, is also easy to establish, e.g., by proving that $\min_{i \in I} \varphi(i) > 0$, and using the zero-one law (cf. [HSP] for a detailed proof).

(2) Once the existence of such a decomposition has been established, it can also be obtained in the following different manner (for details, see [HS]).

Define an equivalence relation on I so that $i, j \in I$ are equivalent if and only if $\alpha(i) = \alpha(j)$ for every subharmonic function α . The sets I_m , $m \geq 0$, are then simply the equivalence classes of this relation. Furthermore, to assign indices to these sets, look for a subharmonic function α which assumes distinct values on each of these sets, and order the equivalence classes in decreasing order of the values of α on them. Such a separating subharmonic function always exists; moreover the decomposition will satisfy condition (c) if and only if $\varphi \equiv 1$. (Thus, in the finite case there always exists a

decomposition of I into an ordered sequence of sets I_m , $m \geq 0$, which satisfy conditions (a) and (b); each of these sets either satisfies conditions (c) or else is K -ergodic.)

An open problem is whether the existence of a similar decomposition is equivalent to $\varphi \equiv 1$ in the general (discrete) case as well.

8. Conclusions. In this paper we have analyzed termination of concurrent probabilistic programs having discrete infinite state spaces. Our aim has been to calculate the worst-case probability of such a program to reach a given set of terminating states under an arbitrary but fair scheduling of its processes. We have obtained several characterizations of the required probability function φ , which yielded useful techniques for the calculation of this function. Specializing to the case of deterministic (or nondeterministic) programs, our techniques have been shown to generalize known techniques for proving termination of such programs, and also to yield new such techniques. From the point of view of the theory of probability, our results extend the classical theory of optimal gambling strategies by Dubins and Savage [DS] to the case where such strategies must be "fair."

The model that we have introduced in this paper and in the preceding one [HSP] for the (fair) execution of concurrent probabilistic programs is very general, natural, and easy to work with, and we believe that it should serve as a standard model for execution of such programs. A more detailed discussion concerning this model can be found in [HSP].

The techniques developed in the present paper can be immediately interpreted as *sound* and *complete* proof methods for termination of concurrent probabilistic programs. It would be interesting to generalize these techniques to proof methods for additional properties of such programs, or, alternatively, to develop temporal probabilistic logics, based upon our techniques, for reasoning about such programs (see, e.g., [HS2]). One such generalization can be achieved as follows: Let α be a subharmonic function defined on I . For each schedule σ define $E_\sigma(\alpha)$ as $\lim_{n \rightarrow \infty} E_{(\sigma, n)}(\alpha)$ (which always exists, by the subharmonicity of α). Then we want to compute

$$\varphi_\alpha(i) = \inf_{\alpha \in \Sigma F(i)} E_\sigma(\alpha),$$

which generalizes the function $\varphi \equiv \varphi_{xx}$ studied in this paper. Intuitively, $\varphi_\alpha(i)$ is the smallest "long-term" expected value of α under a fair schedule starting at i . Most of the theory developed in §§ 3 and 4 can be generalized to the case of a general subharmonic α .

(An interesting choice for α is where $\alpha|_{I-X} \equiv 0$, $\alpha|_X \geq 0$; then φ_α gives the smallest expected value of α upon termination under fair execution. Thus, appropriate adaptation of the techniques developed in this paper will enable us to derive lower bounds for the expected value of such functions upon termination; compare with [SPH]).

A final corollary of the results developed in this paper concerns *bounded waiting time* (cf. [Ra1] for example). Let us define a *round* of execution as a portion of the execution during which each process has been scheduled at least once. We then say that the program has the (*local*) *bounded waiting time property* at some $i \in I$ if for each $\varepsilon > 0$ there exists an integer $N = N(i, \varepsilon)$ such that the probability of reaching X from i after N rounds under any fair schedule is at least $1 - \varepsilon$. The program has the *global bounded waiting time property* if the above holds for all $i \in I$ and if N is independent of i . A simple application of Theorem 3.8 implies the next corollary.

COROLLARY 8.1. *The program has the local bounded waiting time property at $i \in I$ if and only if $\gamma_\omega(i) = 1$. It has the global property if and only if $\gamma_\omega \equiv 1$ (i.e., if the convergence ordinal of the program is $\leq \omega$) and $\gamma_n \uparrow 1$ uniformly on I as $n \uparrow \omega$.*

REFERENCES

- [Ch] K. L. CHUNG, *Markov Chains with Stationary Transition Probabilities*, Springer-Verlag, New York, 1967.
- [DS] L. E. DUBINS AND L. J. SAVAGE, *Inequalities for Stochastic Processes; How to Gamble if You Must*, Dover, New York, 1976.
- [HSP] S. HART, M. SHARIR AND A. PNUELI, *Termination of concurrent probabilistic programs*, ACM Trans. Prog. Languages and Systems, 5 (1983), pp. 356–380.
- [HS] S. HART AND M. SHARIR, *Concurrent probabilistic programs, or how to schedule if you must*, Tech. Rept., School of Mathematical Sciences, Tel Aviv University, May 1982.
- [HS2] ———, *Probabilistic propositional temporal logics*, Proc. 16th Symposium on Theory of Computing, April 1984, pp. 1–13.
- [LPS] D. LEHMANN, A. PNUELI AND J. STAVI, *Impartiality, justice, fairness: the ethics of concurrent termination*, Proc. 8th ICALP Conference, 1981, pp. 264–277.
- [LR] D. LEHMANN AND M. O. RABIN, *On the advantages of free choice: a symmetric and fully distributed solution to the dining philosophers' problem*, Proc. 8th Symposium on Principles of Programming Languages, 1981, pp. 133–138.
- [Ra1] M. O. RABIN, *N process synchronization by a $4 \log_2 N$ -valued shared variable*, J. Comp. Syst. Sciences, 25 (1982), pp. 66–75.
- [Ra2] ———, *The choice coordination problem*, Acta Inform., 17 (1982), pp. 121–134.
- [RS1] J. REIF AND P. SPIRAKIS, *Distributed algorithms for synchronizing interprocess communication within real time*, Proc. 13th Symposium on Theory of Computing, 1981, pp. 133–145.
- [RS2] ———, *Unbounded speed variability in distributed communication systems*, Proc. 9th Symposium on Principles of Programming Languages, 1982, pp. 46–56.
- [SPH] M. SHARIR, A. PNUELI AND S. HART, *Verification of probabilistic programs*, this Journal, 13 (1984), pp. 292–314.