Available online at www.sciencedirect.com
SCIENCE DIRECT®
Information Processing Letters 94 (2005) 79–84
Information Processing Letters
www.elsevier.com/locate/ipl

An elementary proof that Herman's Ring is $\Theta(N^2)$

Annabelle McIver [a], Carroll Morgan [b,*]

[a] *Department of Information and Computer Science, Macquarie University, NSW 2109, Australia*
[b] *Department of Computer Science and Engineering, University of New South Wales, NSW 2052, Australia*

Received 18 February 2004; received in revised form 6 October 2004
Available online 22 January 2005
Communicated by F.B. Schneider

Abstract

Herman's Ring [Inform. Process. Lett. 35 (1990) 63; http://www.cs.uiowa.edu/ftp/selfstab/H90.ps.gz] is an algorithm for self-stabilization of N identical processors connected uni-directionally in a synchronous ring; in its original form it has been shown to achieve stabilization, with probability one, in expected steps $O(N^2 \log N)$. We give an elementary proof that the original algorithm is in fact $O(N^2)$; and for the special case of three tokens initially we give an exact (quadratic) solution of $4abc/N$, where a, b, c are the tokens' initial separations. Thus the algorithm overall has worst-case expected running time of $\Theta(N^2)$. Although we use only simple matrix algebra in the proof, the approach was suggested by the general notions of *abstraction*, *nondeterminism* and *probabilistic variants* [A. McIver, C. Morgan, Refinement and Proof for Probabilistic Systems, Technical Monographs in Computer Science, Springer-Verlag, New York, 2004]. It is hoped they could also be useful for other, similar problems. We conclude with an open problem concerning the worst-case analysis.

Keywords: Randomized algorithms; Analysis of algorithms; Abstraction; Nondeterminism; Probabilistic variant

1. Introduction

Herman's Ring [1] comprises an odd number $N \geqslant 3$ of processors connected unidirectionally in a ring; at any moment each processor can hold either zero or one tokens. In each (synchronous) step of the *stabilization* algorithm, every token-holding processor decides independently with an unbiased coin-flip whether to *keep* its token (probability $1/2$) or to *pass* its token (also probability $1/2$) to the next processor downstream. If a *keeping* processor receives a token from its *passing* immediately-upstream neighbor, then the two tokens are annihilated.

Herman showed [1] that, from any initial state in which the number of tokens is odd, the system as a whole will with probability one eventually reach a *stable* state in which there is only one token; he has also shown that the expected number of synchronous steps until stabilization is $O(N^2 \log N)$.

* Corresponding author.
E-mail address: carrollm@cse.unsw.edu.au (C. Morgan).

0020-0190/$ – see front matter
doi:10.1016/j.ipl.2004.12.013

A number of researchers have described variations and improvements on the original algorithm, in some cases reducing the (time) complexity to $O(N^2)$ [3].

Here we show that Herman's *original* algorithm is $O(N^2)$;[1] and by giving an exact solution for the initial case of three tokens we show that in fact in the worst case it is $\Theta(N^2)$. The proof is given in elementary terms; the more general techniques that led to it are discussed in the conclusion.

2. Characterization of expected steps to stabilization

Let R (for ring) be the finite set of all ring configurations in which the number of tokens is odd and *more than* one. We write two-dimensional matrices, such as (Markov) transition matrices over R, with a double underline; column matrices, such as random variables over R, have a single underline; and if a matrix or vector has entries all the same scalar c then we write it $[c]$ with the appropriate number of underlines.

Let $\underline{\underline{R}}$ be the (#R)-by-(#R) transition matrix of probabilistic transitions determined by Herman's algorithm. It is *sub*-stochastic—its rows sum to no more than one—because only the "unstable" not-yet-terminated (i.e., more than one token) configurations are included in R.

The probability of *not* stabilizing on the very next step is $\underline{\underline{R}} \cdot \underline{[1]}$ (a column vector indexed by initial state); and thus in general $\underline{\underline{R}}^k \cdot \underline{[1]}$ gives the probabilities that stabilization will not occur within k steps. From elementary probability theory [5], the *expected time to stabilization* is then a column vector $\underline{e} = \sum_{k=0}^{\infty} \underline{\underline{R}}^k \cdot \underline{[1]}$ where this summation exists, provided stabilization occurs with probability one: each element of the vector gives the expected time from that initial state.

Where the summation does exist, matrix algebra shows that in fact we have $\underline{e} = \underline{[1]} + \underline{\underline{R}} \cdot \underline{e}$. We put these observations in a lemma:

Lemma 1. *If from every initial configuration r in* R *the expected steps $\underline{e}$ to stabilization is finite, then it satisfies*

$$\underline{\underline{R}} \cdot \underline{e} = \underline{e} - \underline{[1]}. \tag{1}$$

Conversely, if we have some $\underline{e}$ that satisfies (1) *uniquely then, provided we have established* (*by some other means*) *that the expected time to stabilization is everywhere finite, we will know it is given by that $\underline{e}$.*

3. Expected steps to stabilization is finite for Herman's Ring

We begin by showing that the ring's stabilization occurs "quickly" in the sense that the probability of not yet having stabilized decreases exponentially. We assume throughout that the ring size is fixed at N.

Lemma 2. *There are constants $c \geqslant 0$ and $0 \leqslant m < 1$ such that from any initial configuration r of the ring the probability $P_{k,r}$ that the ring will not yet have stabilized, after k steps, is no more than cm^k.*

Proof. Suppose at first that the number of steps is $(N-1)b$ for some b, i.e., that it comprises b "blocks" of $N-1$ steps each; select some fixed processor F. In each block of steps the chance of stabilization is no less than $\varepsilon = (1/2^N)^{N-1} > 0$, since that is a lower bound for the probability that in every one of the $N-1$ steps of the block only the nearest-downstream-to-F token is passed, while all others are kept.

The probability that stabilization does not occur in *any* of the b blocks is thus no more than $(1-\varepsilon)^b$, that is $P_{(N-1)b,r} \leqslant (1-\varepsilon)^b$. Writing $\lfloor \cdot \rfloor$ for the *floor* function, we therefore have for any k that

$$\begin{aligned} P_{k,r} &\leqslant P_{(N-1)\lfloor k/(N-1)\rfloor,r} \\ &\leqslant (1-\varepsilon)^{\lfloor k/(N-1)\rfloor} \\ &\leqslant cm^k, \end{aligned}$$

provided we set $c = 1/(1-\varepsilon)$ and $m = (1-\varepsilon)^{1/(N-1)}$. □

This quick stabilization gives us our finiteness result directly.

Lemma 3. *Stabilization occurs within a finite expected number of steps.*

Proof. Because the rth entry of column vector $\underline{\underline{R}}^k \cdot \underline{[1]}$ is just $P_{k,r}$, we have that Lemma 2 bounds $\sum_{k=0}^{\infty} \underline{\underline{R}}^k \cdot \underline{[1]}$ by $\sum_{k=0}^{\infty} \underline{[cm^k]}$, which converges. □

[1] Herman reports this result also [1], and notes that Dolev et al. have put it in the form of a game [4].