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 $\underline{e}$.*

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

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. □

4. Stabilization takes only $O(N^2)$ steps

We now establish the upper bound by showing that the expected time to stabilization is no more than a quadratic function of N. We begin with a technical lemma.

Lemma 4. *For $\underline{\underline{R}}$ as above we have* $\lim_{k\to\infty} \underline{\underline{R}}^k = \underline{\underline{[0]}}$.

Proof. From the proof of Lemma 3 we have $\lim_{k\to\infty} \underline{\underline{R}}^k \cdot \underline{[1]} = \underline{[0]}$, and the result then follows because all entries of $\underline{\underline{R}}$ are non-negative. □

The following lemma will be used to bound the stabilization complexity:

Lemma 5. *For "terminating" $\underline{\underline{R}}$ as above, suppose for some column vectors $\underline{e}, \underline{u}$ we have $\underline{\underline{R}} \cdot \underline{e} = \underline{e} - \underline{[1]}$ and $\underline{\underline{R}} \cdot \underline{u} \leqslant \underline{u} - \underline{[1]}$, where "$\leqslant$" is taken componentwise. Then $\underline{e} \leqslant \underline{u}$.*

Proof. We have immediately that $\underline{\underline{R}} \cdot (\underline{u} - \underline{e}) \leqslant (\underline{u} - \underline{e})$, whence by induction we obtain $\underline{\underline{R}}^k \cdot (\underline{u} - \underline{e}) \leqslant (\underline{u} - \underline{e})$ for all $k \geqslant 0$. Lemma 4 then gives

$$\underline{[0]} = \underline{\underline{[0]}} \cdot (\underline{u} - \underline{e}) = \lim_{k\to\infty} \underline{\underline{R}}^k \cdot (\underline{u} - \underline{e}) \leqslant \underline{u} - \underline{e}. \qquad \square$$

As a corollary we note that if in fact $\underline{\underline{R}} \cdot \underline{u} = \underline{u} - \underline{[1]}$ then $\underline{e} = \underline{u}$, so that (1) has at most one solution for terminating $\underline{\underline{R}}$.

Now from Lemmas 3 and 1 we know that the expected time $\underline{e}$ to stabilization satisfies $\underline{\underline{R}} \cdot \underline{e} = \underline{e} - \underline{[1]}$. From that and Lemma 5 we have our first result.

Lemma 6 (Herman's ring upper bound). *The expected time to stabilization of Herman's Ring is* $O(N^2)$.

Proof. Choose for an upper bound the column vector $\underline{u}$ of height #R whose r-entry is $2x_r(2N - x_r - 1)$ for each configuration r, where x_r is the *extent* of r, *the minimum length of any span of contiguous segments that includes all token-holding processors in* r. Note that since $x_r \leqslant N - 1$ for all configurations r, each entry of $\underline{u}$ is $O(N^2)$ as a function of N. Elementary (but detailed) calculation shows that $\underline{\underline{R}} \cdot \underline{u} \leqslant \underline{u} - \underline{[1]}$; see Appendix A.

Hence we have $\underline{e} \leqslant \underline{u}$, from Lemma 5, giving that each entry of $\underline{e}$ is $O(N^2)$. □

5. Exact stabilization for three initial tokens is quadratic

Now consider the special case in which exactly three processors have tokens initially. We give an exact value for the expected time to stabilization.

Lemma 7. *The expected time to stabilization of a ring with initially three tokens is $4abc/N$, where a, b, c are the initial separations of the tokens.* (*Note that $a + b + c = N$.*)

Proof. Let R_3 be the set of three-token configurations of the ring. Define column vector $\underline{e_3}$ over R_3 so that for r in R_3 the r-entry of $\underline{e_3}$ is $4a_r b_r c_r/N$, where a_r, b_r, c_r are the particular separations a, b, c in that configuration r. Let $\underline{\underline{R_3}}$ be the reduced transition matrix for R_3 only; we can make this restriction because the two-token case is impossible. Direct calculation (Appendix B) shows that $\underline{\underline{R_3}} \cdot \underline{e_3} = \underline{e_3} - \underline{[1]}$ and, from the corollary to Lemma 5, we see that $\underline{e_3}$ has that property uniquely.

From Lemma 3 we know that the expected time to stabilization is finite. Thus by Lemma 1, the r-entry of $\underline{e_3}$ gives the expected time to stabilization for each r in R_3. □

6. Stabilization takes $\Theta(N^2)$ steps

Our main result follows directly from the bounds proved above.

Theorem 8. *Herman's Ring takes expected $\Theta(N^2)$ steps to stabilization in the worst case.*

Proof. From Lemma 6 Herman's Ring is $O(N^2)$. Its worst-case lower-bound complexity is the worst over all possible initial configurations, including the three-token configurations where $a + b + c = N$; from Lemma 7 it is thus $\Omega(N^2)$, since $abc \approx N^3/27$ when $a \approx b \approx c \approx N/3$. □

7. Conclusion

The idea of the "extent" arises from using probabilistic variants to show termination. The general tech-

"ideal" probabilities →	$(\ldots 0, 0, 1/4$,	$1/2$,	$1/4, 0, 0, \ldots)$
	↑	↑	↑
outcomes for $x_{r'}$ →	$x_{r'} = x_r - 1$	$x_{r'} = x_r$	$x_{r'} = x_r + 1$

Fig. 1. Ideal outcome if others tokens' effects are ignored.

nique was described by Hart et al. [6], and was later embedded in the quantitative program logic of Morgan et al. [2,7,8]; it seems simpler than most other methods for showing termination with probability one.

The extent and associated "ideal transitions" (Fig. 1) are similar to the "inter-token distance" measure used by Herman and others [1,9]. Like Duflot et al. [9] we use abstraction, but we abstract the whole ring—Herman, and subsequently Duflot et al., abstract from the two-token case only; the subsequent decomposition of the original system into a succession of halves then introduces the unnecessary factor of $\log N$.

For inspiration we relied on a model of probabilistic programming more extensive than the Markov setting [2,7]; its general application to analysis of expected complexity is described by Celiku and McIver [10]. Other treatments of expected complexity include Dolev et al.'s scheduler-luck games [4].

We are not the first to note the $\mathrm{O}(N^2)$ upper bound for Herman's Ring: Rosaz's asynchronous leader-election algorithm specializes to a synchronous version which is close to Herman's, and which also is shown to have that $\mathrm{O}(N^2)$ complexity [3] (although not using an elementary proof). However we believe the lower bound is new: we borrowed it from a gambling-game puzzle [11, p. 103]. Intriguingly the numbers our lower bound yields for the *maximum* initial separation agree to six places with the PRISM [12] probabilistic model-checker by Kwiatkowska and her colleagues for *all* initial configurations.

Is the worst-case initial configuration for Herman's Ring therefore just three maximally-separated tokens, for all odd ring sizes N?

Appendix A. Calculations for Lemma 6

We show that $\underline{\underline{R}} \cdot \underline{u} \leqslant \underline{u} - \underline{[1]}$; the calculations are elementary, if intricate.

In summary we argue as follows. The extent x_r of a configuration r is the minimal number of contiguous segments containing all tokens; it behaves roughly as a random walk with absorbing barrier at 0 and reflecting at $N - 1$.

Pick some configuration r with extent x_r; focus on a particular (minimal) span in r of that length x_r; and observe that single step from r to some other r' will have three possible outcomes with respect to the new extent $x_{r'}$.

(1) With probability $1/4 = 1/2 \times 1/2$ the leading token of the span is kept but the trailing token is passed; in this case the extent $x_{r'}$ after the step satisfies $x_{r'} > x_r$.
(2) With probability $1/2 = 1/4 + 1/4$ both tokens are passed or both kept; in this case we have $x_{r'} \leqslant x_r$.
(3) With probability $1/4$ the trailing token is kept but the leading token is passed; in this case we have $x_{r'} \leqslant x_r + 1$.

The inequalities are because of "complicating effects" due to the precise arrangement of tokens in configuration r over which the extent x_r has been measured. For example, since there can be several shortest spans, the value of x can decrease even though case (2) or (3) was the outcome for the span we chose to observe: that will happen if some *other* minimum span shortens, even though this one did not. Similarly, we may find that the leader catches up to the trailer in case (3) when $x_r = N - 1$, or in case (1) the leader-minus-one catches up to the leader. In both of those outcomes the resulting annihilation of the colliding tokens might decrease x by more than one.

The distribution of $x_{r'}$ after the step from r to r' can be written as row-vector of width $(N - 1)$; and the above shows that the row can be thought of as taking the "ideal" outcome as in Fig. 1 and then introducing the inequalities by post-multiplying with a lower-triangular matrix $\underline{\underline{L}}$, with $(\leqslant 1)$-summing rows, that shifts probabilistic weights some distance (possibly zero) to the left towards lower values of $x_{r'}$. Fig. 2 illustrates the case of a five-place three-token ring $(N - 1 = 4)$ of extent 3 in configuration [•·○•·], where the black tokens • indicate the span chosen.

$$\underset{\substack{\uparrow\\ x=2}}{\overset{\textit{ideal}\text{ outcome}}{(0,\tfrac14,\tfrac12,\tfrac14)}}\cdot\begin{pmatrix}0&0&0&0\\0&\frac12&0&0\\0&0&\frac34&0\\0&\frac12&\frac12&0\end{pmatrix}=\underset{\substack{\uparrow\\ x=4}}{\overset{\textit{actual}\text{ outcome}}{(0,\tfrac14,\tfrac12,0)}}\leftarrow \textit{some row } r$$

Fig. 2. Accounting for the effects of other tokens.

Consider, for example, the second column of $\underline{\underline{L}}$. Its second entry $1/2$ "throws away" $(1-1/2)$ of position $x=2$'s value to annihilation; but its fourth entry $1/2$ "steals" $1/2$ of position $x=4$'s value $1/4$, i.e., $1/8$, and gives it to position $x=2$; the same applies to position $x=3$, which loses $1/8$ to annihilation but again gains $1/8$ from position $x=4$. Position $x=4$ however, having lost half its value to each of $x=2$ and $x=3$, becomes zero.

We see below that the idealized system will have the exact solution $\underline{u}$ for expected time to termination [13, Ch. XIV, Sec. 3[2]]; and we show that left-shifting as illustrated in Fig. 2 can only decrease that expectation (i.e., speed up its termination). That will give us the result.

This is the detailed argument. Consider an idealized system $\mathsf{X}=\{1..N-1\}$ of extents, whose $(N-1)$-by-$(N-1)$ sub-stochastic transition matrix $\underline{\underline{X}}$ is determined by the following rules for a single step in which x,x' are the row, column indices, respectively:

when $x=1$—set $x'=x$ with probability $1/2$ and $x'=x+1$ with probability $1/4$ (the probability $1/4$ transition to zero is implicit);
when $1<x<N-1$—set $x'=x-1$ with probability $1/4$, $x'=x$ with probability $1/2$ and $x'=x+1$ with probability $1/4$;
when $x=N-1$—set $x'=x-1$ with probability $1/4$ and $x'=x$ with probability $3/4$.

Now let $\underline{v}$ be the column vector of height $N-1$ abstracted from $\underline{u}$, so that its x-entry is $2x(2N-x-1)$. Elementary algebra then shows that $\underline{\underline{X}}\cdot\underline{v}=\underline{v}-\underline{[1]}$. We set out the $1<x<N-1$ case as an example: element x of $\underline{\underline{X}}\cdot\underline{v}$ equals

$$\begin{aligned}&1/4\times 2(x-1)\big(2N-(x-1)-1\big)\\&\quad+1/2\times 2x(2N-x-1)\\&\quad+1/4\times 2(x+1)\big(2N-(x+1)-1\big)\\&=xN-x^2/2-N+x/2\\&\quad+2xN-x^2-x\\&\quad+xN-x^2/2+N-3x/2-1\end{aligned}$$

which is just $2x(2N-x-1)-1$, element x of $\underline{v}-\underline{[1]}$ as required. In the other two cases we rely on $x=1$ and $x=N-1$, respectively; the result is the same.

We now make the connection between the ideal X and the actual R systems. Let the $(\#\mathsf{R})$-by-$(N-1)$ matrix $\underline{\underline{A}}$ (for abstraction) contain value one in row r column x just when x is the extent of r, and zero otherwise. As a result we have $\underline{u}=\underline{\underline{A}}\cdot\underline{v}$ immediately.

Now the $(\#\mathsf{R})$-by-$(N-1)$ matrix $\underline{\underline{R}}\cdot\underline{\underline{A}}$ gives in its row r, at position x, the probability that one step from r in the *actual* system R will result in a new configuration of extent x. Matrix $\underline{\underline{A}}\cdot\underline{\underline{X}}$ has the same size, but its row r gives at each x the probability that first converting initial r to its extent and then taking one step, in the *idealized* system $\underline{\underline{X}}$, will give final extent x. Thus, as Fig. 2 showed, we must have a "shifting" relationship between corresponding rows: we can write for some lower-triangular row-$(\leqslant 1)$-summing $\underline{\underline{L_r}}$ the equality $(\underline{\underline{A}}\cdot\underline{\underline{X}})_{\langle r\rangle}\cdot\underline{\underline{L_r}}=(\underline{\underline{R}}\cdot\underline{\underline{A}})_{\langle r\rangle}$, where $\langle r\rangle$ selects row r and $\underline{\underline{L_r}}$ may depend on r.

Now inspection of column vector $\underline{v}$ shows it is monotonic in x. (The entries $2x(2N-x-1)$ of $\underline{v}$ are non-decreasing as x varies from 1 to $N-1$.) Because of $\underline{\underline{L_r}}$'s special properties, and that monotonicity, we have $\underline{\underline{L_r}}\cdot\underline{v}\leqslant\underline{v}$, and so

$$(\underline{\underline{R}}\cdot\underline{\underline{A}})_{\langle r\rangle}\cdot\underline{v}=(\underline{\underline{A}}\cdot\underline{\underline{X}})_{\langle r\rangle}\cdot\underline{\underline{L_r}}\cdot\underline{v}\leqslant(\underline{\underline{A}}\cdot\underline{\underline{X}})_{\langle r\rangle}\cdot\underline{v},$$

for all r. Thus taking all rows at once gives in fact $\underline{\underline{R}}\cdot\underline{\underline{A}}\cdot\underline{v}\leqslant\underline{\underline{A}}\cdot\underline{\underline{X}}\cdot\underline{v}$, and we can now conclude our argument with simple matrix algebra. We have

[2] We adapt the expected duration $z(a-z)$ of the $1/2$, $1/2$ random walk: replace the walker's position z by the extent n; replace the upper barrier a by $2N-1$ because our upper barrier is reflecting; and multiply by 2 because our probabilities are $1/4$, $1/4$.

Token movements	Final values		
	a'	b'	c'
No token passed	a	b	c
	$a-1$	b	$c+1$
One token passed	$a+1$	$b-1$	c
	a	$b+1$	$c-1$
	a	$b-1$	$c+1$
Two token passed	$a+1$	b	$c-1$
	$a-1$	$b+1$	c
All token passed	a	b	c

Fig. 3. Effects of one step on token separations a, b, c. The separations a, b, c are named from upstream to downstream; each outcome occurs with probability $1/8$.

$$\begin{aligned}\underline{\underline{R}}\cdot \underline{u} &= \underline{\underline{R}}\cdot \underline{\underline{A}}\cdot \underline{v} \leqslant \underline{\underline{A}}\cdot \underline{\underline{X}}\cdot \underline{v} = \underline{\underline{A}}\cdot (\underline{v} - \underline{[1]}) \\ &= \underline{u}\cdot \underline{\underline{A}}\cdot \underline{[1]} = \underline{u} - \underline{[1]},\end{aligned}$$

which is the inequality we sought.

Appendix B. Calculations for Lemma 7

We must show that $\underline{\underline{R_3}}\cdot \underline{e_3} = \underline{e_3} - \underline{[1]}$.

In a three-token system there are eight equiprobable outcomes for a single step, ranging from “all tokens kept” to “all tokens passed”. Their effects on a, b, c are tabulated in Fig. 3. Direct calculation of the rth entry in $\underline{\underline{R_3}}\cdot \underline{e_3}$ gives

$$1/8 \times 4/N \times \begin{cases} abc + (a-1)b(c+1) \\ +(a+1)(b-1)c + a(b+1)(c-1) \\ +a(b-1)(c+1) + (a+1)b(c-1) \\ +(a-1)(b+1)c + abc, \end{cases}$$

where a, b, c are the separations in configuration r. Via $a+b+c=N$ that expression simplifies to $4abc/N - 1$, which is the rth entry of $\underline{e_3} - \underline{[1]}$, as required.

References

[1] T. Herman, Probabilistic self-stabilization, Inform. Process. Lett. 35 (1990) 63–67, also http://www.cs.uiowa.edu/ftp/selfstab/H90.ps.gz.

[2] A. McIver, C. Morgan, Abstraction, Refinement and Proof for Probabilistic Systems, Technical Monographs in Computer Science, Springer-Verlag, New York, 2004.

[3] L. Rosaz, Self-stabilizing token circulation on asynchronous uniform unidirectional rings, in: PODC 2000, ACM, New York, 2000, pp. 249–258.

[4] S. Dolev, A. Israeli, S. Moran, Analysing expected time by scheduler-luck games, IEEE Trans. Software Engrg. 21 (1995) 429–439.

[5] G. Grimmett, D. Stirzaker, Probability and Random Processes, second ed., Oxford Science Publications, Oxford, 1992.

[6] S. Hart, M. Sharir, A. Pnueli, Termination of probabilistic concurrent programs, ACM Trans. Program. Lang. Syst. 5 (1983) 356–380.

[7] C. Morgan, A. McIver, K. Seidel, Probabilistic predicate transformers, ACM Trans. Program. Lang. Syst. 18 (3) (1996) 325–353.

[8] C. Morgan, Proof rules for probabilistic loops, in: H. Jifeng, J. Cooke, P. Wallis (Eds.), Proc. 7th BCS-FACS Refinement Workshop, Springer-Verlag, Berlin, 1996; http://ewic.bcs.org/conferences/1996/refinement/papers/paper10.htm.

[9] M. Duflot, L. Fribourg, C. Picaronny, Randomized finite-state distributed algorithms as Markov chains, in: PODC 2001, in: Lecture Notes in Comput. Sci., vol. 2180, Springer-Verlag, Berlin, 2001, pp. 240–254.

[10] O. Celiku, A. McIver, Cost-based analysis of programs mechanised in HOL, http://web.comlab.ox.ac.uk/oucl/research/areas/probs/bibliography.html.

[11] R. Honsberger, Mathematical Diamonds, Dolcani Mathematical Expositions, vol. 26, Mathematical Society of America, Washington, DC, 2003.

[12] PRISM, Probabilistic symbolic model checker, http://www.cs.bham.ac.uk/~dxp/prism/casestudies/self-stabilisation.html.

[13] W. Feller, An Introduction to Probability Theory and its Applications, vol. 1, second ed., Wiley, New York, 1959.