A Tighter Bound for the Self-Stabilization Time in Herman's Algorithm

Yuan Feng[a,1], Lijun Zhang[b,c,2]

[a] *University of Technology Sydney, Australia*
[b] *State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences*
[c] *Technical University of Denmark, DTU Compute, Denmark*

Abstract

We study the expected self-stabilization time of the Herman's algorithm. For N processors the lower bound is $\frac{4}{27}N^2$ ($0.148N^2$), and an upper bound of $0.64N^2$ is presented in [4]. In this paper we give a tighter upper bound $0.521N^2$.

We assume to have N processors numbered from 1 to N, clockwise, with N odd, organized in a ring topology. Each processor may or may not have a token. We denote the number of initial tokens by M, with M odd and $1 \leq M \leq N$. Let $z : \{1, \ldots, M\} \to \{1, \ldots, N\}$ be such that $z(1) < \cdots < z(M)$ and for all $i \in \{1, \ldots, M\}$, the processor $z(i)$ initially has a token; in other words, $z(i)$ is the position of the i-th token.

Herman's protocol in the traditional *synchronous variant* [2] works as follows: in each time step, each processor with a token passes its token to its clockwise neighbor with probability r (where $r \in (0, 1)$ is a fixed parameter), and keeps it with probability $1-r$. If a processor keeps its token and receives another one from its counterclockwise neighbor, then both of those tokens are annihilated. We refer to configurations with only one token as *legitimate* configurations. The protocol can be viewed as a finite state Markov chain with a single bottom SCC consisting of all of the legitimate configurations. So a legitimate configuration is reached with probability 1, regardless of the initial configuration. Hence, Herman's protocol is *self-stabilizing*.

An *asynchronous variant* of Herman's protocol is proposed in [4] which works similarly: each processor with a token passes the token to its clockwise neighbor with rate λ; i.e., a processor keeps its token for a time that is negative exponentially distributed with parameter λ, before passing the token to the neighbor. The asynchronous variant does not require processor synchronization.

Let $\mathbf{T}$ denote the time until only one token is left, i.e., until self-stabilization has occurred. We analyze its expectation $\mathbb{E}\mathbf{T}$ and provide a tighter upper bound. The following proposition provides a precise formula for $\mathbb{E}\mathbf{T}$ in the case $M = 3$.

Proposition 1 (cf. [6, 4]). *Let, for the synchronous protocol, $D = r(1 - r)$ with $r \in (0, 1)$; or, for the asynchronous protocol, $D = \lambda$ with $\lambda > 0$. Let $M = 3$, and let a, b, c denote the distances between neighboring tokens ($a + b + c = N$). Then $\mathbb{E}\mathbf{T} = \frac{abc}{DN}$.*

Proposition 1 is shown in [6] for the synchronous case with $r = \frac{1}{2}$, and is presented in [4] for the asynchronous case. For the case $r = \frac{1}{2}$, Proposition 1 implies $\mathbb{E}\mathbf{T} = \frac{4abc}{N}$: if N is a multiple of 3, one gets the lower bound $\mathbb{E}\mathbf{T} = \frac{4}{27}N^2$ with $a = b = c = \frac{N}{3}$. This is approximately $0.148N^2$. Indeed, McIver and Morgan [6] have conjectured that $0.148N^2$ is the worst case expected time among all configurations. Using the probabilistic model checker PRISM [5], the conjecture is validated for all rings with the size $N \leq 21$ that can be exhaustively analysed. Recently, [3] has studied the distribution of the self-stabilization time and shown that for an arbitrary t the probability of stabilization within time t is minimized under this configuration with $M = 3$.

In the synchronous case, Herman's original work [2] provides an upper bound $\mathbb{E}\mathbf{T} \leq N^2 \lceil \log N \rceil / 2$ for arbitrary configurations. In [1, 6], a concrete bound $2N^2$ is given for the case $r = \frac{1}{2}$. Then, the author of [7] applied the theory of coalescing random walks to Herman's protocol to obtain $\mathbb{E}\mathbf{T} \leq \left(\frac{\pi^2}{8} - 1\right) \cdot \frac{N^2}{r(1-r)}$, which is about $0.93N^2$ for the case $r = \frac{1}{2}$. By combining results from [7] and [6], the bound is further improved in [4] to $\left(\frac{\pi^2}{8} - \frac{29}{27}\right) \cdot \frac{N^2}{D}$, which is about $0.64N^2$ for the synchronous case $r = \frac{1}{2}$. The bound is shown to be valid for the asynchronous variant as well [4].

In the theorem below we improve this bound further to $0.521N^2$, and propose a similar bound for the asynchronous protocol version:

Theorem 1. *Let, for the synchronous protocol, $D = r(1-r)$ with $r \in (0, 1)$; or, for the asynchronous protocol, $D = \lambda$ with $\lambda > 0$. Then, for all N and for all initial configura-*

[1] Supported partially by Australian ARC grants DP110103473 and FT100100218.

[2] Supported by the EU FP7 Programme under grant agreement no. TREsPASS (318003), by IDEA4CPS and MT-LAB (a VKR Centre of Excellence).

tions, we have

$$\mathbb{E}\mathbf{T} \leq \left(\frac{\pi^2}{8} - \frac{7 \cdot \zeta(3)}{4} + 1\right) \cdot \frac{N^2}{D}$$

where ζ is the Riemann zeta function. Hence, $\mathbb{E}\mathbf{T} \leq 0.521N^2$ in the synchronous case with $r = \frac{1}{2}$.

Proof The proof is based on [7] for the synchronous case, and [4] for the asynchronous case. For $M \geq 3$, let τ_M denote the maximal expected time for a configuration with M tokens to reach a configuration with fewer than M tokens, where the maximum is taken over all M-token configurations. It is shown that

$$\tau_M \leq \frac{1}{M^2} \cdot \frac{N^2}{D} \tag{1}$$

in [7] for the synchronous case, and [4] for the asynchronous case, respectively. Nakata [7] shows then $\mathbb{E}\mathbf{T} \leq \tau_3 + \tau_5 + \tau_7 + \cdots$. Further, since $\frac{1}{1^2} + \frac{1}{3^2} + \frac{1}{5^2} + \cdots = \frac{\pi^2}{8}$, it follows that $\mathbb{E}\mathbf{T} \leq \left(\frac{\pi^2}{8} - 1\right) \cdot \frac{N^2}{D}$. The improvement in [4] is obtained by replacing the bound $\tau_3 \leq \frac{1}{9} \cdot \frac{N^2}{D}$ with $\tau_3 \leq \frac{1}{27} \cdot \frac{N^2}{D}$, which follows from Proposition 1.

Our improvement is based on a tighter bound on τ_M for all $M \geq 5$. Observe that there must exist three neighboring tokens such that their distances are at most $\lfloor 2N/M \rfloor$. Without loss of generality, we assume $z_{12} + z_{23} \leq \lfloor 2N/M \rfloor$, where we write z_{uv} for $z(v) - z(u)$. Now consider these three tokens at $z(1)$, $z(2)$ and $z(3)$. With the same argument as in [7] it holds that the expected time until two of these three tokens are annihilated is an upper bound on τ_M. By applying Proposition 1, it thus holds

$$\tau_M \leq \frac{z_{12} z_{23} (N - z_{12} - z_{23})}{DN} =: f(z_{12}, z_{23})$$

with constraints $z_{12}, z_{23} \geq 1$ and $z_{12} + z_{23} \leq \lfloor 2N/M \rfloor$. We rewrite $f(z_{12}, z_{23})$ as follows:

$$\begin{aligned} f(z_{12}, z_{23}) &= \frac{z_{12}}{DN}(-z_{23}^2 + (N - z_{12}) z_{23}) \\ &= \frac{z_{12}}{DN}\left[-\left(z_{23} - \frac{N - z_{12}}{2}\right)^2 + \frac{(N - z_{12})^2}{4}\right] \end{aligned}$$

Now we consider the right-hand side expression as a function in z_{23} with respect to the previous constraints. Observe that the constraint $z_{12} + z_{23} \leq \lfloor 2N/M \rfloor$ implies that $z_{23} < \frac{N - z_{12}}{2}$. Thus for any valid fixed z_{12}, the value $f(z_{12}, z_{23})$ is increasing in $z_{23} \in [1, \lfloor 2N/M \rfloor - z_{12}]$. It follows then that $f(z_{12}, z_{23})$ takes its maximal value if $z_{12} + z_{23} = \lfloor 2N/M \rfloor$. In this case we reformulate $f(z_{12}, z_{23})$ by

$$\begin{aligned} f(z_{12}, z_{23}) &= \frac{1}{DN} z_{12} \cdot (\lfloor 2N/M \rfloor - z_{12}) \cdot (N - \lfloor 2N/M \rfloor) \\ &\leq \frac{1}{4DN} (\lfloor 2N/M \rfloor)^2 \cdot (N - \lfloor 2N/M \rfloor) \end{aligned}$$

Note that the function $x^2 \cdot (N - x)$ is increasing between 1 and $2N/3$, and $\lfloor 2N/M \rfloor \leq 2N/M < 2N/3$ for $M \geq 5$. We further derive that

$$\begin{aligned} \tau_M &\leq f(z_{12}, z_{23}) \\ &\leq \frac{1}{4DN} (2N/M)^2 \cdot (N - 2N/M) \\ &= \frac{1}{M^2} \cdot \frac{N^2}{D} (1 - \frac{2}{M}). \end{aligned}$$

Comparing with the Nakata bound in Eq.(1), our bound has been improved by $\frac{2}{M^3} \cdot \frac{N^2}{D}$.

With the new bound, and noting that $\frac{1}{1^3} + \frac{1}{3^3} + \frac{1}{5^3} + \cdots = \frac{7}{8}\zeta(3)$, we have

$$\begin{aligned} \mathbb{E}\mathbf{T} &\leq \sum_{M=3,5,7\ldots} \frac{1}{M^2} \cdot \frac{N^2}{D} (1 - \frac{2}{M}) \\ &= \left(\frac{\pi^2}{8} - \frac{7\zeta(3)}{4} + 1\right) \cdot \frac{N^2}{D}. \end{aligned}$$

Note that $\zeta(3) \approx 1.2021$. We have $\mathbb{E}\mathbf{T} \leq 0.521N^2$ in the synchronous case with $r = \frac{1}{2}$. □

Essentially, our proof exploits the precise solution for the three token configurations in Proposition 1, whereas the strategy in [7] uses the precise solution for the random walk. We conclude by noting that our approach can be exploited to improve bounds for $M = 5$ and $z_{12} = 1$. For this case we can derive a tighter bound for τ_5 using $z_{12} = 1$ in the proof of Theorem 1, and have $\mathbb{E}\mathbf{T} \leq \frac{1}{27}\frac{N^2}{D} + \frac{N}{4D}$. If $z_{23} = 1$ holds in addition, we have $\mathbb{E}\mathbf{T} \leq \frac{1}{27}\frac{N^2}{D} + \frac{1}{D} - \frac{2}{DN}$.

References

[1] L. Fribourg, S. Messika, and C. Picaronny. Coupling and self-stabilization. *Distributed Computing*, 18:221–232, 2005.

[2] T. Herman. Probabilistic self-stabilization. *Information Processing Letters*, 35(2):63–67, 1990.

[3] S. Kiefer, A. S. Murawski, J. Ouaknine, B. Wachter, and J. Worrell. Three tokens in Herman's algorithm. *Formal Asp. Comput.*, 24(4-6):671–678, 2012.

[4] S. Kiefer, A. S. Murawski, J. Ouaknine, J. Worrell, and L. Zhang. On stabilization in Herman's algorithm. In *ICALP (2)*, pages 466–477, 2011.

[5] M. Z. Kwiatkowska, G. Norman, and D. Parker. Probabilistic verification of Herman's self-stabilisation algorithm. *Formal Asp. Comput.*, 24(4-6):661–670, 2012.

[6] A. McIver and C. Morgan. An elementary proof that Herman's ring is $\Theta(n^2)$. *Inf. Process. Lett.*, 94(2):79–84, 2005.

[7] T. Nakata. On the expected time for Herman's probabilistic self-stabilizing algorithm. *Theoretical Computer Science*, 349(3):475–483, 2005.