Under consideration for publication in Formal Aspects of Computing

Three Tokens in Herman's Algorithm[1]

Stefan Kiefer*, Andrzej S. Murawski‡, Joël Ouaknine*, Björn Wachter*, and James Worrell*

*Department of Computer Science, University of Oxford, UK

‡Department of Computer Science, University of Leicester, UK

Abstract. Herman's algorithm is a synchronous randomized protocol for achieving self-stabilization in a token ring consisting of N processes. The interaction of tokens makes the dynamics of the protocol very difficult to analyze. In this paper we study the distribution of the time to stabilization, assuming that there are three tokens in the initial configuration. We show for arbitrary N and for an arbitrary timeout t that the probability of stabilization within time t is minimized by choosing as the initial three-token configuration the configuration in which the tokens are placed equidistantly on the ring. Our result strengthens a corollary of a theorem of McIver and Morgan [13], which states that the *expected* stabilization time is minimized by the equidistant configuration.

1. Introduction

Self-stabilization is a concept of fault-tolerance in distributed computing. A system is self-stabilizing if, starting in an arbitrary state, it reaches a correct or *legitimate* state and remains in a legitimate state thereafter. Thus a self-stabilizing system is able to recover from *transient errors* such as state-corrupting faults. The study of self-stabilizing algorithms originated in an influential paper of Dijkstra [4]. By now there is a considerable body of work in the area, see [16, 5].

In this paper we consider self-stabilization in a classical context that was also treated in Dijkstra's original paper—a *token ring*, i.e., a ring of N identical processes, exactly one of which is meant to hold a token at any given time. If, through some error, the ring enters a configuration with multiple tokens, self-stabilization requires that the system be guaranteed to reach a configuration with only one token. In particular, we are interested in analyzing a self-stabilization algorithm proposed by Herman [8].

Herman's algorithm is a randomized procedure by which a ring of processes connected uni-directionally can achieve self-stabilization almost surely. The algorithm works by having each process synchronously execute the following action at each time step: if the process possesses a token then it passes the token to its clockwise neighbor with probability $1/2$ and keeps the token with probability $1/2$. If such a process decides to keep its token and if it receives a token from its neighbor then the two tokens are annihilated. Due to the way the algorithm is implemented we can assume that an error state always has an odd number of tokens, thus this process of pairwise annihilation eventually leads to a configuration with a single token.

It is easy to see that Herman's algorithm is almost surely self-stabilizing, but computing the time to termination is a challenging problem. This is characteristic of systems of interacting particles under random

[1] Research supported by EPSRC (EP/G069158/1). Stefan Kiefer is supported by a postdoctoral fellowship of the German Academic Exchange Service (DAAD).

Correspondence and offprint requests to: Stefan Kiefer, Department of Computer Science, University of Oxford, UK. e-mail: stefan.kiefer@cs.ox.ac.uk

motion, which are ubiquitous in the physical and medical sciences, including statistical mechanics, neural networks and epidemiology [12]. The analysis of such systems typically requires delicate combinatorial arguments [6]. Our case is no exception.

Given some initial configuration, let $\mathbf{T}$ be the time until the token ring stabilizes under Herman's algorithm. Previous analysis of $\mathbf{T}$ focused largely on $\mathbb{E}\mathbf{T}$, the expected value of the self-stabilization time. Herman's original paper [8] showed that $\mathbb{E}\mathbf{T} \leq (N^2 \log N)/2$ in the worst case (i.e., over all initial configurations with N processes). It also mentions an improved upper bound of $O(N^2)$ due to Dolev, Israeli, and Moran, without giving a proof or a further reference. In 2005, three papers [7, 13, 15] were published, largely independently, all of them giving improved $O(N^2)$ bounds. In [10] we studied various extensions of the protocol and improved the upper bound for an arbitrary initial configuration to $\mathbb{E}\mathbf{T} \leq 0.64N^2$.

McIver and Morgan [13] paid special attention to initial configurations with exactly three tokens. In this case, the protocol stabilizes as soon as two tokens meet. They found an explicit formula for computing $\mathbb{E}\mathbf{T}$ for any initial configuration with exactly three tokens: if $a, b, c \in \mathbb{N}$ with $a + b + c = N$ denote the initial token distances, then it holds that

$$\mathbb{E}\mathbf{T} = 4abc/N\,. \tag{1}$$

Assuming that N is a multiple of 3, we obtain $\mathbb{E}\mathbf{T} = \frac{4}{27}N^2$ for an *equidistant* configuration, which is a configuration with three equally spaced tokens. It follows from (1) that an equidistant configuration maximizes $\mathbb{E}\mathbf{T}$ over all three-token configurations. Furthermore, it was conjectured in [13] that this is the worst case among all starting configurations, including those with more than three tokens. This intriguing conjecture is supported by experimental evidence [2].

In this paper we analyze the distribution of $\mathbf{T}$ for three-token configurations in more detail. We show for an arbitrary timeout $t \in \mathbb{N}$ that any equidistant configuration minimizes the probability $\mathcal{P}(\mathbf{T} \leq t)$ among all three-token configurations. As $\mathbf{T}$ is an $\mathbb{N}$-valued random variable, we have $\mathbb{E}\mathbf{T} = \sum_{t=0}^{\infty} \mathcal{P}(\mathbf{T} > t)$, so our result strengthens the above-mentioned consequence of (1) that equidistant configurations maximize $\mathbb{E}\mathbf{T}$ over the three-token configurations. We illustrate our result with experimental data generated by the probabilistic model checker APEX [14, 11, 9].

In [10] we analyzed $\mathbb{E}\mathbf{T}$ by exploiting and adapting a technique that was developed by Balding [1] for a scenario from physical chemistry. The technique uses an application of the reflection principle to express the dynamics of the token interaction in terms of one-dimensional random walks with absorbing barriers. We reuse this technique in this paper to compute the distribution of $\mathbf{T}$ in the three-token case. Although notation considerably simplifies in the three-token case, the core of Balding's argument is preserved, so we hope that giving full details for the three-token case is worthwhile. The second ingredient of our analysis is an inductive argument about $\mathcal{P}(\mathbf{T} \leq t)$ over the timeout t. Using Balding's technique it suffices to give an inductive expression for the probability that by time t a one-dimensional random walk has hit a barrier.

Organization of the paper. After some preliminaries in Section 2, we prove in Section 3 our main result which states that equidistant configurations minimize the probability $\mathcal{P}(\mathbf{T} \leq t)$ among the three-token configurations. We obtain explicit expressions for $\mathcal{P}(\mathbf{T} \leq t)$ in Section 4. We also show how the APEX tool can be used to automatically compute these numbers, taking as input only an implementation of Herman's algorithm. We conclude in Section 5.

Happy birthday!

2. Preliminaries

We assume $N \geq 3$ processes organized in a ring topology, numbered from 1 to N, clockwise, according to their position in the ring. Each process may or may not have a token. Herman's protocol works as follows: in each time step, each process that has a token passes its token to its clockwise neighbor with probability $1/2$ and keeps it with probability $1/2$; if a process keeps its token and receives another token from its counterclockwise neighbor, then both of those tokens are annihilated. Notice that the number of tokens never increases, and can decrease only by even numbers. We assume in the following that initially exactly three processes have a token. Thus an initial configuration z is given by three numbers $1 \leq z_1 < z_2 < z_3 \leq N$ such that process z_i initially has a token; in other words, z_i is the position of the ith token. We are interested in the time $\mathbf{T}_z$ until the (only) annihilation takes place. We may drop the subscript if the initial configuration z is understood.

The protocol can be viewed as a Markov chain with two strongly connected components, a transient SCC

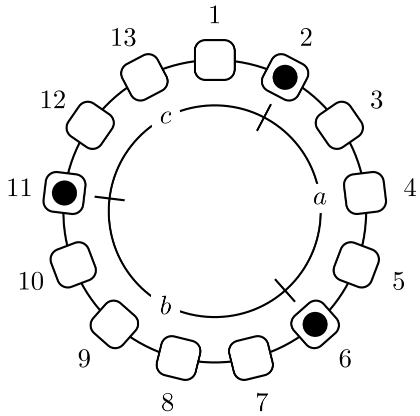

Fig. 1. An example of an equidistant configuration with $N = 13$ processes, token positions $(z_1, z_2, z_3) = (2, 6, 11)$, and token separations $(a, b, c) = (4, 5, 4)$.

containing the three-token configurations, and a recurrent SCC containing the one-token configurations. So a one-token configuration is almost surely reached, i.e., $\mathbf{T} < \infty$ holds with probability 1.

We denote by $a, b, c \in \mathbb{N}$ the initial token distances, i.e., $a := z_2 - z_1$, $b := z_3 - z_2$, and $c := z_1 - z_3 + N$. Observe that $a + b + c = N$. A configuration is called *equidistant* if $\lfloor N/3 \rfloor \leq a, b, c \leq \lceil N/3 \rceil$. Figure 1 illustrates an equidistant configuration for $N = 13$. For fixed N, all equidistant configurations are essentially the same (by rotational symmetry), so we sometimes speak about *the* equidistant configuration.

3. Main Result

In this section we prove our main result, which states that the equidistant configuration minimizes the probability of annihilation within time t.

Theorem 1. Let $N, t \in \mathbb{N}$ with $N \geq 3$. Let z^* denote the equidistant three-token configuration. Then

$$\mathcal{P}(\mathbf{T}_{z^*} \leq t) \quad \leq \quad \mathcal{P}(\mathbf{T}_z \leq t) \qquad \text{for all three-token configurations } z.$$

Fix N for the rest of the section. Proposition 2 below expresses $\mathcal{P}(\mathbf{T} \leq t)$ in terms of a discrete one-dimensional random walk on $\{0, \ldots, N\}$ with transitions

$$0 \xrightarrow{1} 0 \qquad \begin{array}{l} x \xrightarrow{1/4} x-1 \\ x \xrightarrow{1/2} x \\ x \xrightarrow{1/4} x+1 \end{array} \qquad N \xrightarrow{1} N \qquad \text{for } x \in \{1, \ldots, N-1\}. \tag{2}$$

For $t \in \mathbb{N}$, denote by $f_t : \{0, \ldots, N\} \to [0, 1]$ the function such that $f_t(x)$ is the probability that such a random walk started at x has hit the left absorbing barrier 0 by time t. Observe that

$$f_t(0) = 1 \quad \text{and} \quad f_t(N) = 0 \qquad \text{for all } t \in \mathbb{N}. \tag{3}$$

Proposition 2 (cf. [1, Theorem 2.1] and [10, Proposition 3]). Let $N, t \in \mathbb{N}$ with $N \geq 3$ and consider an initial three-token configuration with distances a, b, c. Then

$$\mathcal{P}(\mathbf{T} \leq t) = f_t(a) + f_t(b) + f_t(c) - f_t(a+b) - f_t(a+c) - f_t(b+c).$$

Proof. Fix a timeout t. We consider several events, i.e., sets of possible "trajectories" of the three tokens. We can formally think of a trajectory as a function $\omega : \{0, \ldots, t\} \times \{1, 2, 3\} \to \{1, \ldots, N\}$, assigning to a point in time and to a token its position on the ring. We have $\omega(0, i) = z_i$ for $i \in \{1, 2, 3\}$. Annihilations are disregarded, but given a trajectory ω it is clear, if, when and where tokens meet and would therefore be annihilated in Herman's algorithm.

Define D^0_{12} as the event that, by time t, tokens 1 and 2 have met, and no other pair of tokens has met before the first meeting of tokens 1 and 2. Note that in Herman's algorithm, the event D^0_{12} is that by time t tokens 1 and 2 have been annihilated. We define events D^0_{23} and D^0_{31} analogously for token pairs $(2,3)$ and $(3,1)$. Then we have

$$\text{“}\mathbf{T} \le t\text{”} \quad = \quad D^0_{12} \cup D^0_{23} \cup D^0_{31}\,, \tag{4}$$

where the unions are disjoint.

Define D_{12} as the event that by time t token 1 has "caught up" with token 2 in the clockwise direction (possibly after other collisions involving token 3). More formally, let $\Delta_\omega(s) := (\omega(s,2) - \omega(s,1)) \bmod N$ represent the clockwise distance from token 1 to token 2 at time point $s \le t$. Then D_{12} is the event that $\Delta_\omega(s)$ reaches 0 by time t and moreover that it first reaches 0 from value 1. The events D_{13}, D_{23}, D_{21}, D_{31} and D_{32} are defined analogously.

We partition D_{12} into three disjoint events $D_{12} = D^0_{12} \cup D^{31}_{12} \cup D^{23}_{12}$, where D^0_{12} was defined above; $D^{31}_{12} := D^0_{31} \cap D_{12}$ is the event that the first collision is between tokens 1 and 3, before token 1 eventually catches up with token 2; $D^{23}_{12} := D^0_{23} \cap D_{12}$ is the event that the first collision is between tokens 2 and 3, before token 1 eventually catches up with token 2. Events $D^{12}_{13}, D^{23}_{13}, D^{12}_{23}, D^{31}_{23}, D^{23}_{21}, D^{31}_{21}, D^{23}_{31}, D^{12}_{31}, D^{31}_{32}, D^{12}_{32}$ are defined analogously, so that we have

$$\begin{aligned} D_{12} &= D^0_{12} \cup D^{31}_{12} \cup D^{23}_{12} & D_{21} &= D^{23}_{21} \cup D^{31}_{21} \\ D_{23} &= D^0_{23} \cup D^{12}_{23} \cup D^{31}_{23} & D_{32} &= D^{31}_{32} \cup D^{12}_{32} \\ D_{31} &= D^0_{31} \cup D^{23}_{31} \cup D^{12}_{31} & D_{13} &= D^{12}_{13} \cup D^{23}_{13}\,, \end{aligned} \tag{5}$$

where the unions are disjoint. Observe that token 2 cannot catch up with token 1 before token 1 or token 2 meets token 3, so there are no events D^0_{21}, D^0_{32} or D^0_{13}.

By the reflection principle, the events D^{31}_{12} and D^{31}_{32} have the same probability. To show this in detail, we establish a bijection π between D^{31}_{12} and D^{31}_{32} as follows. Given $\omega \in D^{31}_{12}$ we define the trajectory $\pi(\omega)$ by "switching" the movements of tokens 1 and 3 after their first collision. Formally, let $t' < t$ be the time of the first meeting of tokens 1 and 3, i.e., $\omega(t',1) = \omega(t',3)$. Define

$$\begin{aligned} \pi(\omega)(t'',1) &:= \begin{cases} \omega(t'',1) & t'' \le t' \\ \omega(t'',3) & t'' \ge t' \end{cases} \\ \pi(\omega)(t'',2) &:= \omega(t'',2) \\ \pi(\omega)(t'',3) &:= \begin{cases} \omega(t'',3) & t'' \le t' \\ \omega(t'',1) & t'' \ge t'\,, \end{cases} \end{aligned}$$

Notice that the trajectories ω and $\pi(\omega)$ have the same probability and that π indeed defines a bijection between D^{31}_{12} and D^{31}_{32}. Similarly, one can show

$$\begin{matrix} \mathcal{P}\left(D^{31}_{12}\right) = \mathcal{P}\left(D^{31}_{32}\right) & \mathcal{P}\left(D^{12}_{23}\right) = \mathcal{P}\left(D^{12}_{13}\right) & \mathcal{P}\left(D^{23}_{31}\right) = \mathcal{P}\left(D^{23}_{21}\right) \\ \mathcal{P}\left(D^{23}_{12}\right) = \mathcal{P}\left(D^{23}_{13}\right) & \mathcal{P}\left(D^{31}_{23}\right) = \mathcal{P}\left(D^{31}_{21}\right) & \mathcal{P}\left(D^{12}_{31}\right) = \mathcal{P}\left(D^{12}_{32}\right)\,. \end{matrix} \tag{6}$$

Combining all these observations, we obtain

$$\begin{aligned} \mathcal{P}\left(\text{“}\mathbf{T} \le t\text{”}\right) &\overset{(4)}{=} \mathcal{P}\left(D^0_{12}\right) + \mathcal{P}\left(D^0_{23}\right) + \mathcal{P}\left(D^0_{31}\right) \\ &\overset{(6)}{=} \left(\mathcal{P}\left(D^0_{12}\right) + \mathcal{P}\left(D^{31}_{12}\right) + \mathcal{P}\left(D^{23}_{12}\right)\right) + \left(\mathcal{P}\left(D^0_{23}\right) + \mathcal{P}\left(D^{12}_{23}\right) + \mathcal{P}\left(D^{31}_{23}\right)\right) \\ &\quad + \left(\mathcal{P}\left(D^0_{31}\right) + \mathcal{P}\left(D^{23}_{31}\right) + \mathcal{P}\left(D^{12}_{31}\right)\right) - \left(\mathcal{P}\left(D^{23}_{21}\right) + \mathcal{P}\left(D^{31}_{21}\right)\right) \\ &\quad - \left(\mathcal{P}\left(D^{31}_{32}\right) + \mathcal{P}\left(D^{12}_{32}\right)\right) - \left(\mathcal{P}\left(D^{12}_{13}\right) + \mathcal{P}\left(D^{23}_{13}\right)\right) \\ &\overset{(5)}{=} \mathcal{P}\left(D_{12}\right) + \mathcal{P}\left(D_{23}\right) + \mathcal{P}\left(D_{31}\right) - \mathcal{P}\left(D_{21}\right) - \mathcal{P}\left(D_{32}\right) - \mathcal{P}\left(D_{13}\right) \\ &= f_t(a) + f_t(b) + f_t(c) - f_t(a+b) - f_t(a+c) - f_t(b+c)\,. \end{aligned}$$

For the final equality, notice that $\mathcal{P}\left(D_{12}\right) = f_t(a)$, because the event D_{12} (where token 1 catches up with token 2, disregarding token 3) is in a bijection with the event that the one-dimensional random walk started at $a = z_2 - z_1$ has been absorbed at 0 by time t: the random walk models the distance between tokens 1 and 2. Similarly, we have $\mathcal{P}(D_{21}) = f_t(b+c)$ etc. $\square$

For $t \in \mathbb{N}$ and $y \in \{0, \ldots, N\}$, define a function $g_{t,y} : \{0, \ldots, N-y\} \to [0,1]$ by

$$g_{t,y}(x) := f_t(x) + f_t(y) + f_t(N-x-y) - f_t(x+y) - f_t(N-y) - f_t(N-x)\,.$$

Observe that we have $\mathcal{P}(\mathbf{T} \le t) = g_{t,b}(a)$ by Proposition 2, which establishes the range $[0,1]$ of $g_{t,y}$. If the initial positions of tokens 2 and 3 are fixed (thus determining their distance b), the function $g_{t,b}$ describes how $\mathcal{P}(\mathbf{T} \le t)$ depends on the initial position of token 1 (which determines a). The following lemma states that $g_{t,b}$ is minimized by placing token 1 "halfway" between tokens 2 and 3.

Lemma 3. Let $t \in \mathbb{N}$ and $y \in \{0, \ldots, N\}$. The function $g_{t,y}$ is minimized by $x = \lfloor (N-y)/2 \rfloor$ and $x = \lceil (N-y)/2 \rceil$.

Proof. Define $G_{t,y} : \{0, \ldots, N-y\} \to \mathbb{R}$ by

$$G_{t,y}(x) := f_t(x) + f_t(N-x-y) - f_t(x+y) - f_t(N-x)\,. \tag{7}$$

Notice that $G_{t,y}$ differs from $g_{t,y}$ only by a constant $f_t(y) - f_t(N-y)$. The definition of the one-dimensional random walk (2) implies

$$f_t(x) = \frac{1}{4} f_{t-1}(x-1) + \frac{1}{2} f_{t-1}(x) + \frac{1}{4} f_{t-1}(x+1) \qquad \text{for } t \ge 1 \text{ and } 1 \le x \le N-1, \tag{8}$$

hence we have for $t \ge 1$ and $y \in \{0, \ldots, N\}$ and $1 \le x \le N-y-1$:

$$\begin{aligned}
G_{t,y}(x) &\overset{(7)}{=} f_t(x) + f_t(N-x-y) - f_t(x+y) - f_t(N-x) \\
&\overset{(8)}{=} \frac{1}{4} f_{t-1}(x-1) + \frac{1}{2} f_{t-1}(x) + \frac{1}{4} f_{t-1}(x+1) \\
&\quad + \frac{1}{4} f_t(N-x+1-y) + \frac{1}{2} f_t(N-x-y) + \frac{1}{4} f_t(N-x-1-y) \\
&\quad - \frac{1}{4} f_{t-1}(x-1+y) - \frac{1}{2} f_{t-1}(x+y) - \frac{1}{4} f_{t-1}(x+1+y) \\
&\quad - \frac{1}{4} f_{t-1}(N-x+1) - \frac{1}{2} f_{t-1}(N-x) - \frac{1}{4} f_{t-1}(N-x-1) \\
&\overset{(7)}{=} \frac{1}{4} G_{t-1,y}(x-1) + \frac{1}{2} G_{t-1,y}(x) + \frac{1}{4} G_{t-1,y}(x+1)\,.
\end{aligned} \tag{9}$$

Recall from (3) that $f_t(0) = 1$ and $f_t(N) = 0$, hence $g_{t,y}(0) = 1$. We now prove that the function $g_{t,y}$ is monotonically decreasing in $\{0, \ldots, \lfloor (N-y)/2 \rfloor\}$, which suffices to prove the statement of the lemma, as $g_{t,y}$ is, by its definition, symmetric around $(N-y)/2$. We proceed by induction on t. For $t = 0$ and $y = 0$ we have $g_{0,0}(x) = 1$ for $x \ge 0$. For $t = 0$ and $y > 0$ we have $g_{0,y}(0) = 1 > 0 = g_{0,y}(x)$ for $x \in \{0, \ldots, \lfloor (N-y)/2 \rfloor\}$. Let $t \ge 1$. We show $g_{t,y}(x-1) \ge g_{t,y}(x)$ for $x \in \{1, \ldots, \lfloor (N-y)/2 \rfloor\}$.

- Let $x = 1$. Then $g_{t,y}(0) = 1 \ge g_{t,y}(1)$.
- Let $2 \le x < (N-y)/2$. Then

$$\begin{aligned}
G_{t,y}(x-1) &\overset{(9)}{=} \frac{1}{4} G_{t-1,y}(x-2) + \frac{1}{2} G_{t-1,y}(x-1) + \frac{1}{4} G_{t-1,y}(x) \\
&\overset{\text{IH}}{\ge} \frac{1}{4} G_{t-1,y}(x-1) + \frac{1}{2} G_{t-1,y}(x) + \frac{1}{4} G_{t-1,y}(x+1) \\
&\overset{(9)}{=} G_{t,y}(x)\,,
\end{aligned}$$

where the inequality marked with "IH" holds by the induction hypothesis. This implies $g_{t,y}(x-1) \ge g_{t,y}(x)$, as $G_{t,y}$ and $g_{t,y}$ differ only by a constant.

- Let $x = (N-y)/2$. (This case only occurs if $N-y$ is even.) Then

$$\begin{aligned}
G_{t,y}(x-1) &\overset{(9)}{=} \frac{1}{4}G_{t-1,y}(x-2) + \frac{1}{2}G_{t-1,y}(x-1) + \frac{1}{4}G_{t-1,y}(x) \\
&\overset{\text{S}}{=} \frac{1}{4}G_{t-1,y}(x-2) + \frac{1}{4}G_{t-1,y}(x-1) + \frac{1}{4}G_{t-1,y}(x) + \frac{1}{4}G_{t-1,y}(x+1) \\
&\overset{\text{IH}}{\geq} \frac{1}{4}G_{t-1,y}(x-1) + \frac{1}{2}G_{t-1,y}(x) + \frac{1}{4}G_{t-1,y}(x+1) \\
&\overset{(9)}{=} G_{t,y}(x) \;,
\end{aligned}$$

where the inequalities marked with "S" and "IH" hold by symmetry around x and by the induction hypothesis, respectively. This implies $g_{t,y}(x-1) \geq g_{t,y}(x)$, as in the previous case.

□

Now we can prove Theorem 1.

Proof of Theorem 1. Suppose that z is a three-token configuration with minimum probability (among three-token configurations) to stabilize by time t. Write a, b, c for the token distances in z and recall that $\mathcal{P}(\mathbf{T} \leq t) = g_{t,b}(a)$. The function $g_{t,b}$ is minimized at a, thus, applying Lemma 3, we have that $\lfloor (a+c)/2 \rfloor \leq a \leq \lceil (a+c)/2 \rceil$. We conclude that $|a-c| \leq 1$ and, by symmetry, we likewise have that $|a-b|, |b-c| \leq 1$. This implies $\lfloor N/3 \rfloor \leq a, b, c \leq \lceil N/3 \rceil$, i.e., z is equidistant. □

4. Computing Self-Stabilization Probabilities

For a concrete expression for $\mathcal{P}(\mathbf{T} \leq t)$, we only need an expression for f_t and to apply Proposition 2. By [3, Section 2.2, Equation (25)], if we define

$$u(x,j) := \frac{\sin\frac{j\pi}{N}}{1-\cos\frac{j\pi}{N}} \cdot \sin\frac{j\pi x}{N} \qquad \text{and} \qquad v(j) := 1 - \frac{1}{2}\left(1 - \cos\frac{j\pi}{N}\right) \;,$$

we get

$$f_t(x) = 1 - \frac{x}{N} - \frac{1}{N} \cdot \sum_{j=1}^{N-1} u(x,j)v(j)^t \qquad \text{for } t \in \mathbb{N} \text{ and } x \in \{0, \ldots, N\}. \tag{10}$$

We apply (10) to compute $\mathcal{P}(\mathbf{T} \leq t)$ for the three-token configurations in a ring with $N = 9$ processes. Figure 2 shows all (up to rotational symmetry) such configurations.[2] We thus obtain the numbers reported in Figure 3. Observe that the equidistant configuration "J" indeed minimizes $\mathcal{P}(\mathbf{T} \leq t)$.

The same numbers can be obtained using the probabilistic model checker APEX [14, 11, 9]. The tool APEX needs as input only an implementation of Herman's algorithm. In Figure 4 we show such an implementation for $N = 9$. Running APEX on this code produces the probability $\mathcal{P}(\mathbf{T} \leq 5)$ for the equidistant configuration J. The other probabilities from Figure 3 are obtained by straightforward modifications of the the code. Thus APEX provides a way of computing these numbers without a specially tailored analysis of Herman's algorithm. The implementation in Figure 4 uses the classical encoding of configurations, where each process holds a single bit: a process has a token if and only if the process's bit coincides with the bit of the counterclockwise neighbor process. See [11] for more details on APEX and its application on Herman's algorithm.

5. Conclusions

We have studied the distribution of the self-stabilization time $\mathbf{T}$ for three-token configurations. Our analysis relies on the reflection principle which allows one to reduce the interactions of the tokens to a few one-dimensional random walks. Arguing by induction on t then suffices to prove our main result, which states

[2] We do not consider mirror symmetry because the APEX implementation below is not mirror symmetric.

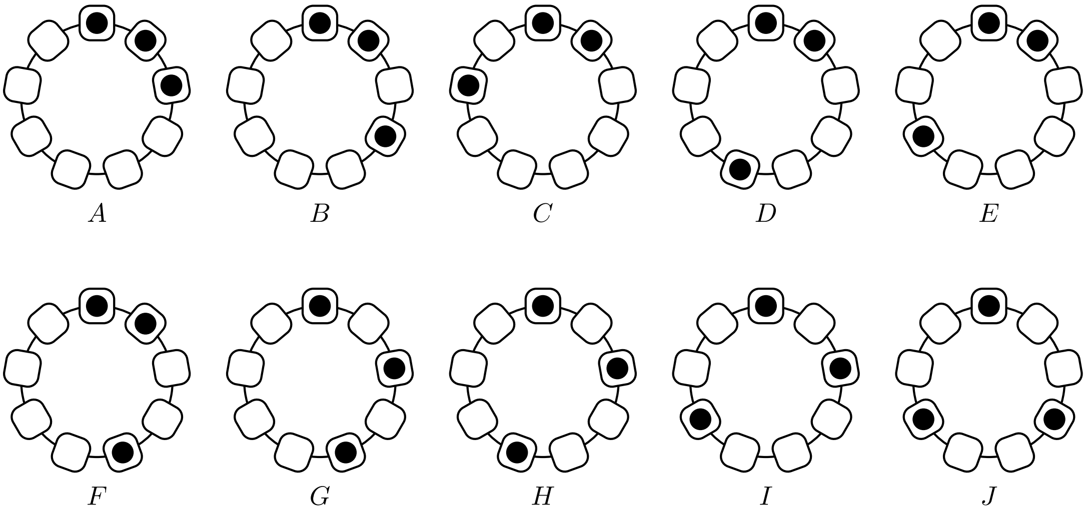

Fig. 2. All three-token configurations for $N = 9$.

	$t=1$	$t=2$	$t=3$	$t=4$	$t=5$	$t=6$	$t=7$	$t=8$	$t=9$	$t=10$
A	.50000	.68750	.78125	.83594	.87109	.89526	.91278	.92606	.93650	.94499
B	.25000	.43750	.56250	.64844	.70996	.75586	.79144	.81996	.84344	.86317
C	.25000	.43750	.56250	.64844	.70996	.75586	.79144	.81996	.84344	.86317
D	.25000	.37500	.45313	.51563	.57031	.61914	.66278	.70168	.73624	.76690
E	.25000	.37500	.46875	.54297	.60352	.65405	.69696	.73384	.76582	.79370
F	.25000	.37500	.46875	.54297	.60352	.65405	.69696	.73384	.76582	.79370
G	0	.12500	.25000	.35547	.44238	.51465	.57562	.62775	.67275	.71189
H	0	.06250	.14062	.22266	.30273	.37793	.44696	.50946	.56556	.61562
I	0	.06250	.14062	.22266	.30273	.37793	.44696	.50946	.56556	.61562
J	0	0	.04687	.11719	.19629	.27612	.35248	.42334	.48793	.54615

Fig. 3. $\mathcal{P}(\mathbf{T} \le t)$ for the configurations from Figure 2, rounded to 5 decimal places.

that the equidistant configuration minimizes $\mathcal{P}(\mathbf{T} \le t)$ for any timeout t. We also have numerically confirmed our results by computing $\mathcal{P}(\mathbf{T} \le t)$ both using our tailored analysis and the general-purpose APEX tool.

A conjecture by McIver and Morgan [13] states that the three-token equidistant configuration maximizes $\mathbb{E}\mathbf{T}$ among all configurations, even among those with more tokens. Our result may give evidence for a stronger conjecture: Does the three-token equidistant configuration minimize $\mathcal{P}(\mathbf{T} \le t)$ among all configurations and for all t? Numerical experiments with 5 and more tokens seem to support this conjecture.

References

[1] D. Balding. Diffusion-reaction in one dimension. *J. Appl. Prob.*, 25:733–743, 1988.

[2] PRISM case studies. Randomised self-stabilising algorithms. `http://www.prismmodelchecker.org/casestudies/self-stabilisat`

[3] D. Cox and H. Miller. *The theory of stochastic processes*. Chapman & Hall/CRC, 2001.

[4] E. W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Commun. ACM*, 17(11):643–644, 1974.

[5] S. Dolev. *Self-Stabilization*. MIT Press, 2000.

[6] R. Durrett and H. Kesten (eds). *Random Walks, Brownian Motion and Interacting Particle Systems*. Birkhauser Verlag AG, 1991.

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

[8] T. Herman. Probabilistic self-stabilization. *Information Processing Letters*, 35(2):63–67, 1990. Technical Report at `ftp://ftp.math.uiowa.edu/pub/selfstab/H90.html`.

[9] S. Kiefer, A.S. Murawski, J. Ouaknine, B. Wachter, and J. Worrell. Language equivalence for probabilistic automata. In *Proceedings of CAV*, volume 6806 of *LNCS*, pages 526–540. Springer, 2011.

```
const N:=9;
const STEPS:=5;
var%2 x[N];
var%2 z;
var%3 token;
var%N i;

var%(STEPS+1) counter;

token:=2;

// set bits for configuration ''J''
x[1]:=1;
x[4]:=1;
x[7]:=1;

while((not (token=1)) and (counter<STEPS)) do
  {
    counter:= counter + 1;
    token := 0;
    i := 0;
    z := x[0];
    while (succ(i)) do
    {
      if (x[i] = x[succ(i)]) then x[i]:=coin
      else x[i] := x[succ(i)];
      if i then if (x[i] = x[i-1]) then
        token:=case(token)[1,2,2];
      i:=succ(i)
    };
    if (x[i] = z) then x[i]:=coin else x[i]:=z;
    if (x[i] = x[i-1]) then token :=case(token)[1,2,2];
    if (x[0] = x[i]) then token :=case(token)[1,2,2];
  };
if (token=1) then skip else diverge
:com
```

Fig. 4. APEX code for Herman's algorithm.

[10] S. Kiefer, A.S. Murawski, J. Ouaknine, J. Worrell, and L. Zhang. On stabilization in Herman's algorithm. In *Proceedings of ICALP*, volume 6756 of *LNCS*, pages 466–477. Springer, 2011.

[11] A. Legay, A. S. Murawski, J. Ouaknine, and J. Worrell. On automated verification of probabilistic programs. In *Proceedings of TACAS*, volume 4963 of *LNCS*, pages 173–187. Springer, 2008.

[12] T.M. Liggett. *Interacting particle systems*. Springer, 2005.

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

[14] A. S. Murawski and J. Ouaknine. On probabilistic program equivalence and refinement. In *Proceedings of CONCUR*, volume 3653 of *LNCS*, pages 156–170. Springer, 2005.

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

[16] M. Schneider. Self-stabilization. *ACM Comput. Surv.*, 25(1):45–67, 1993.