

Self-Stabilizing Ring Orientation Using Constant Space

Jaap-Henk Hoepman*

CWI, Amsterdam, The Netherlands

E-mail: jhh@cwi.nl

The ring-orientation problem requires all processors on an anonymous ring to reach agreement on a direction along the ring. A self-stabilizing ring-orientation protocol eventually ensures that all processors on the ring agree on a direction, regardless of the initial states of the processors on which the protocol is started. In this paper we present two uniform deterministic self-stabilizing ring-orientation protocols for rings with an odd number of processors using only a constant number of states per processor. The first protocol operates in the link-register model under the distributed daemon, and the second protocol operates in the state-reading model under the central daemon. Both protocols do not assume an upper bound on the length of the ring and are therefore applicable to dynamic rings. As an application of our techniques we are able to prove that under the central daemon on an odd-length ring, the link-register model and the state-reading model are equivalent in the sense that any self-stabilizing protocol for the one model can be transformed to an equivalent, self-stabilizing protocol in the other model. © 1998 Academic Press

1. INTRODUCTION

On oriented rings, processors agree on a direction along the ring. Distributed algorithms on rings are more easily derived if it is known that the ring is oriented (cf. [ASW88]) and may be more efficient than similar algorithms for unoriented rings (cf. [San84]). To orient a ring the processors must choose a left and right neighbour consistently around the ring, such that the left neighbour of each processor considers that processor to be its right neighbour. Processors can distinguish between a first and second neighbour, but to exclude trivial solutions the processors are required to be otherwise identical.

A comprehensive study on uniform rings in the asynchronous message passing model has been published by Attiya *et al.* [ASW88]. They showed that there is no

* Partially supported by the Dutch foundation for scientific research (NWO) through NFI Project ALADDIN, under Contract NF 62-376. A preliminary version of this paper appeared in the Int. Workshop on Distributed Algorithms 1994 [Hoe94].

deterministic protocol to orient even-length rings, and that there also cannot exist a protocol to orient rings of arbitrary length, if the protocols are required to terminate. Syrotiuk and Pachl [SP87] presented a simple asynchronous ring-orientation protocol using message passing that is only guaranteed to work for rings whose lengths are odd and bounded. These papers do not address self-stabilization.

Self-stabilizing protocols are protocols that will eventually satisfy their specification, regardless of the initial state they were started in. Self-stabilization was introduced by Dijkstra [Dij74, Dij82] and is a framework in which one can derive fault-tolerant protocols capable of recovering from transient errors. This type of error can change the state of certain processors, but leaves the processors themselves in working order. Now consider a self-stabilizing protocol running on a set of processors and consider the state just after the last error. This could just as well have been the initial state of the protocol, so the protocol must attempt to recover from this error. As the protocol is self-stabilizing it will be able to do so, provided the next error does not occur too soon. Therefore if transient errors are infrequent enough, self-stabilizing protocols keep the system in a correct state most of the time.

Our interest in self-stabilizing ring-orientation protocols is threefold. First of all, several self-stabilizing protocols that run on oriented rings have been published. For instance Burns and Pachl [BP89] have shown that deterministic self-stabilizing protocols can break symmetry on oriented rings of prime size. Recently, Itkis *et al.* [ILS95] constructed a constant space protocol for the same problem. Our results imply that the ring does *not* have to be oriented to achieve these results. Second, if a ring can be oriented deterministically, we are interested in the necessary cost of doing so. Finally, our self-stabilizing ring-orientation protocols allow us to show that two models of interprocessor communication frequently used in the literature on self-stabilization are in fact equivalent on odd-length rings.

Several models for interprocessor communication and processor scheduling in self-stabilizing systems have been proposed in the literature. In the *state-reading* model [Dij74] processors communicate by reading each others' state. Second, in the *link-register* model [DIM93] a processor uses separate shared registers to communicate with each of its neighbours. Third, in the *message-passing* model processors communicate by sending messages (a)synchronously to other processors, where each message may incur an unbounded but finite delay.

The scheduling of processor steps is also an important issue. The *central daemon* [Dij74] schedules one processor at a time. This processor then performs one step in which it reads the information it needs, does some local processing, and finally writes its new state before returning control to the daemon. This model corresponds to systems with a high granularity of atomicity or multitasking systems where a single processor activates each process in turn. The *distributed daemon* [Bur87] may schedule several processors concurrently, but it is assumed that all scheduled processors first read the information they need, before any of them is allowed to write. Systems that stabilize under the distributed daemon will stabilize on synchronous systems (where processors proceed in lock-step) as well as on multitasking systems where some (not necessarily one) processors activate the processes in the self-stabilizing system. Finally, the *read/write daemon* [DIM93] allows arbitrary interleaving of

processor steps, much like the interleaving semantics considered for wait-free shared memory constructions.

Self-stabilizing ring-orientation protocols were studied by Israeli and Jalfon [IJ93]. They prove that no uniform deterministic self-stabilizing ring-orientation protocols exist in (a) the link-register model under the distributed daemon for even-length rings, (b) the state-reading model for either (b1) even-length rings under the central daemon or (b2) arbitrary rings under the distributed daemon. This leads them to construct a randomized self-stabilizing ring-orientation protocol in the link-register model under the distributed daemon for arbitrary rings. To complement their impossibility results, they also present a uniform deterministic self-stabilizing ring-orientation protocol to orient odd-length rings in the link-register model under the distributed daemon. This protocol assumes knowledge of an upper bound on the length of the ring; it uses a nonconstant number of states per processor.

Recently, Tsai and Huang [TH95] also studied self-stabilizing ring orientation. Their main contribution is a deterministic protocol that will orient any (also even) ring under the central daemon in a model where the neighbour of a node can see whether the orientation of that node points towards or instead points away from it. This model is stronger than the state-reading model: a node can convey different information to each of its neighbours, which makes it essentially equivalent to the link-register model. In Section 4.1 we show, however, that deterministic ring orientation of arbitrary rings in the link-register model under the central daemon is a trivial corollary of the Israeli–Jalfon protocol.

We present two uniform deterministic self-stabilizing ring-orientation protocols for odd-length rings, both using only a constant number of states per processor. The first protocol operates in the link-register model under the distributed daemon. This protocol is an adaption of the general randomized Israeli–Jalfon protocol. Contrary to the deterministic Israeli–Jalfon protocol for odd-length rings, our protocol does not depend on the length of the ring. This implies that our protocol can be used on dynamic rings on which the number of processors may change over time (provided that the length of the ring is odd in between these changes). The second protocol operates in the state-reading model under the central daemon and complements the impossibility results of Israeli and Jalfon. Note that these results do not contradict the impossibility results of Attiya *et al.* [ASW88], as self-stabilizing protocols can never be required to terminate. Nor do these results contradict the impossibility of uniform self-stabilizing mutual-exclusion on rings of nonprime length proven by Dijkstra [Dij82], as a cyclic symmetrical configuration is not necessarily unoriented.

The number of models encountered in the literature on distributed algorithms is overwhelmingly large. Some of these models are clearly distinct, other models may only differ significantly for certain classes of problems. It is a major challenge to explore these different models and to find conditions or problem areas such that using either model yields the same result. These models are then called equivalent. We prove that under the central daemon for a class of graphs, including oriented rings, the state-reading and link-register model are equivalent in the sense that any self-stabilizing protocol for the one model can be transformed to an equivalent, self-stabilizing protocol in the other model. Using our second protocol this proves that

the link-register model and the state-reading model are equivalent, in the same sense, under the central daemon on odd-length rings. Our results extend those of Gouda *et al.* [GHR90] to system models often used in the literature on self-stabilization.

The structure of this paper is as follows. Section 2 describes the model of a distributed system assumed throughout this paper. Several formal definitions of self-stabilization and their ramifications are discussed in Section 3. Then in Section 4 we start with a brief description of the Israeli–Jalfon protocol and continue presenting our uniform deterministic self-stabilizing ring-orientation protocol in the link-register model under the distributed daemon for odd-length rings. Section 5 describes the uniform deterministic self-stabilizing ring-orientation protocol in the state-reading model under the central daemon for odd-length rings. We conclude this paper by showing the equivalence of the link-register and the state-reading model on odd-length rings under the central daemon in Section 6 and presenting some interesting directions of further research in Section 7.

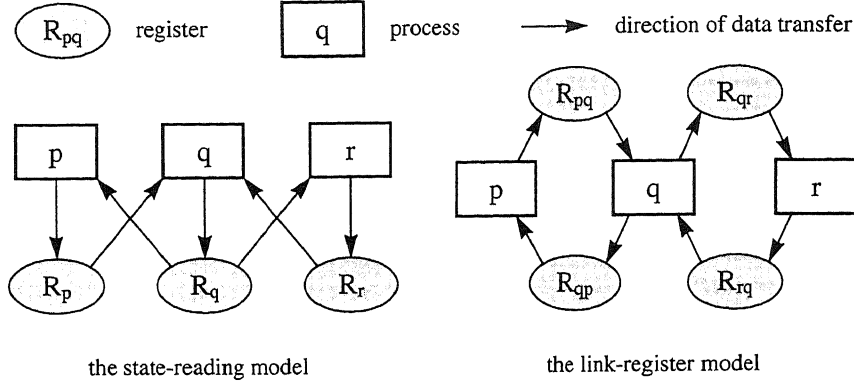
2. THE MODEL

We consider an *anonymous ring* $R = (V_R, E_R)$ consisting of an odd number n of clockwise numbered *nodes* $q \in V_R = \{0, \dots, n-1\}$ and *clockwise edges* $e \in E_R = \{pq \mid p, q \in V_R \wedge q = (p+1) \bmod n\}$. The nodes of the ring are numbered purely for notational convenience: as the ring is anonymous no node has access to its number. Two nodes p, q are *neighbours* if either pq or qp is an element of E_R . Neighbouring nodes can communicate with each other directly. Each node can distinguish a first and second neighbour. For neighbours p of q we define $port_q(p)$ to equal 1 if p is the first neighbour of q and 2 if p is the second neighbour of q .¹ In the remainder of this section, let nodes p and r be neighbours of a node q , with $port_q(p) = 1$.

In this paper two models of communication are considered (see Fig. 1). In the *link-register model*, q will communicate with p and r using separate registers: R_{qp} is written by q and read by p , whereas R_{qr} is written by q and read by r . In the *state-reading model* q stores its state in a register R_q readable by both p and r . The state-reading model is weaker than the link-register model, since in the state-reading model a processor q cannot introduce asymmetry in the states observed by p and r , whereas it can do so in the link-register model by writing different values to R_{qp} and R_{qr} .

The *state* s_q of a node q is comprised of its internal state and the contents of the registers it writes. The *configuration* C of the system is the Cartesian product $\prod_{q \in V_R} s_q$ over the states of all nodes in the ring. We write $C[q]$ for the state of node q in configuration C , and similarly $C[R_q]$ for the value of register R_q in configuration C . Node q can update its state according to its *program* δ_q . Each *step* of node q in configuration C changes q 's state to $\delta_q(C)$. In the link-register model, $\delta_q(C)$ is defined as $\delta_q(C[R_{pq}], C[q], C[R_{rq}])$. In the state-reading model, $\delta_q(C)$ is defined as $\delta_q(C[R_p], C[q], C[R_r])$. A *protocol* consists of a program δ_q for every node $q \in V_R$. A protocol is *uniform* if for all $p, q \in V_R$ we have $\delta_p = \delta_q$.

¹ Of course node q is not necessarily the first neighbour of node p if $pq \in E_R$.

FIG. 1. Registers used by q to communicate with its neighbours.

A *schedule* is an infinite sequence $(P_i)_{i \geq 0}$ of *activations* $P_i \subseteq V_R$. A schedule is *fair* if each node $q \in V_R$ occurs in infinitely many activations P_i . Define the sequence $(t_i)_{i \geq 0}$ for a given schedule $(P_i)_{i \geq 0}$ setting $t_0 = 0$ and, for all $j > 0$, setting t_j to the minimal t such that $\bigcup_{i=t_{j-1}}^{t-1} P_i = V_R$. This sequence is unique and partitions the schedule into *rounds* i , starting at t_i and ending at t_{i+1} , such that in each round each node is activated at least once. Thus a fair schedule is partitioned into infinitely many rounds. Under the *central daemon* one node is activated at a time to execute exactly one step: hence the central daemon only *allows* P_i that consist of exactly one node. Under the *distributed daemon* a set of nodes is simultaneously activated to concurrently execute exactly one step each. It allows arbitrary P_i . All nodes executing a step must have read the values in neighbouring registers before any node can be allowed to write the new value. Under both daemons activating P_i in configuration C yields a configuration C' , denoted $C \rightarrow_{P_i} C'$, such that for all $q \notin P_i$ we have $C'[q] = C[q]$, whereas for all $q \in P_i$ we have $C'[q] = \delta_q(C)$.

A schedule $(P_i)_{i \geq 0}$ and an *initial configuration* C_0 induce an *execution* $E = (C_i)_{i \geq 0}$ such that for all $i \geq 0$ we have $C_i \rightarrow_{P_i} C_{i+1}$. We write $E(i)$ for the i th configuration in execution E , and if $i \leq j$ we write $E(i) \Rightarrow_E E(j)$. An execution is fair if it is induced by a fair schedule.² We write \mathcal{E} for the set of all fair executions allowed by the daemon under consideration. Let the schedule be partitioned into rounds as above. Then in execution E , round i starts in configuration $E(t_i)$, for which we write $\hat{E}(i)$.

We use some additional definitions in this paper. A property X is called *stable* in configuration C (of execution E) if X holds in all configurations C' with $C \Rightarrow_E C'$. A property X is called *stable from configuration* C_f *up to configuration* C_l (of execution E) if this property holds in all configurations C' with $C_f \Rightarrow_E C' \Rightarrow_E C_l$. A *clockwise chain* is a sequence of nodes $q_0 \cdots q_k$, not necessarily $k < n$, such that for all i with $0 \leq i < k$, $q_i q_{i+1} \in E_R$. An *anticlockwise chain* is a sequence of nodes $q_0 \cdots q_k$ such that for all i with $0 \leq i < k$, $q_{i+1} q_i \in E_R$. A *chain* is either a clockwise or an anticlockwise chain.

² Note that we do allow stuttering (i.e., transitions $C \rightarrow_{\emptyset} C$) to occur in executions. In fact we need stuttering to make all executions infinite.

3. ABOUT SELF-STABILIZATION

A self-stabilizing protocol is a protocol that, when started in an arbitrary initial configuration, will eventually behave according to its specification. If we want to give a formal definition of self-stabilization, we first have to formalize what we mean by the specification of a protocol. In the early papers on self-stabilization the specification was viewed as describing the set of configurations \mathcal{L} , called the *legitimate configurations*, the protocol should be in. A mutual exclusion protocol, for instance, should always be in a configuration in which at most one node is executing its critical section. In this setting a protocol is *self-stabilizing* to a specification \mathcal{L} if for every execution E the protocol will eventually reach a legitimate configuration $E(i)$, and once the configuration is legitimate it will remain legitimate forever: $(\forall E \in \mathcal{E}, \exists i \geq 0 :: E(i) \in \mathcal{L})$ and $(\forall C \in \mathcal{L}, \forall P :: C \rightarrow_P C' \Rightarrow C' \in \mathcal{L})$.

One drawback of configuration-based specifications becomes apparent if we again consider mutual exclusion. Usually it is required that the privilege—the node allowed to execute its critical section—is passed fairly among all nodes competing. This fairness-property cannot be expressed in a configuration-based specification and is therefore not captured in the above definition of a self-stabilizing protocol.

Another, more general, way to view the specification of a protocol is as describing the set of behaviours $\mathcal{B} \subseteq \mathcal{E}$ the protocol should abide: i.e., any execution of the protocol should belong to its specification. In this setting a protocol is *self-stabilizing* to specification \mathcal{B} if for all executions E of the protocol there exists an $i \geq 0$ such that all executions of the protocol starting in configuration $E(i)$ belong to \mathcal{B} :

$$(\forall E \in \mathcal{E}, \exists i \geq 0, \forall F \in \mathcal{E} : F(0) = E(i) :: F \in \mathcal{B}).$$

In other words, a self-stabilizing protocol eventually cannot violate its specification.

Using this as a starting point, Burns *et al.* [BGM93] called a protocol *pseudo-stabilizing* to specification \mathcal{B} if for all executions $E = (C_j)_{j \geq 0}$ of the protocol there exists an $i \geq 0$ such that $(E(j))_{j \geq i} \in \mathcal{B}$; i.e., if

$$(\forall E \in \mathcal{E}, \exists i \geq 0 :: (E(j))_{j \geq i} \in \mathcal{B}).$$

In other words, a pseudo-stabilizing protocol eventually will not violate its specification. Pseudo-stabilizing protocols are weaker than self-stabilizing protocols. The first *will not* violate its specification, whereas the other *cannot* violate its specification. As the second statement depends on the assumption that the system will not be disrupted by another transient error, the difference seems artificial in practice. As Burns *et al.* [BGM93] observed that it is much easier to make pseudo-stabilizing protocols than to make self-stabilizing protocols, one might favour the first. However, pseudo-stabilizing protocols are not required to stabilize within a certain amount of time. In fact, if \mathcal{B} is closed under taking suffixes, one easily sees that if one can prove an upper bound on the stabilization time of a pseudo-stabilizing protocol, then this protocol is also self-stabilizing³ [Tel94]. Thus pseudo-stabilizing protocols that are not self-stabilizing actually do not guarantee that the system will be legitimate even once.

³ Observe that if i is bounded, $(\forall E \in \mathcal{E}, \exists i : 0 \leq i < k :: (E(j))_{j \geq i} \in \mathcal{B})$, so if \mathcal{B} is closed under taking suffixes, $(\forall E \in \mathcal{E} :: (E(j))_{j \geq k} \in \mathcal{B})$.

One can build self-stabilizing protocols from scratch, or one can try to combine previous results to obtain more generally applicable protocols. Dolev *et al.* [DIM93] introduced *fair protocol combination* as a useful tool to construct a self-stabilizing protocol from two, simpler, self-stabilizing protocols. Informally speaking, fair protocol combination combines a master-protocol PM —that stabilizes to a certain specification provided certain external conditions hold—with a slave-protocol PS that stabilizes to executions in which exactly these conditions hold. The resulting protocol is a self-stabilizing version of PM without requiring those external conditions. In the combined protocol the steps of both protocols are taken alternately; the states of both protocols are merged so that the master protocol can read the state of the slave protocol. For more details we refer to Dolev *et al.* [DIM93]. Their construction of a self-stabilizing mutual exclusion protocol for arbitrary graphs is a good example of the use of fair protocol combination. It combines a self-stabilizing mutual exclusion protocol that only operates on tree-shaped graphs (the master protocol) with a self-stabilizing spanning tree protocol for arbitrary graphs (the slave protocol). The slave protocol ensures that the master protocol is eventually run a tree-shaped graph, which in turn guarantees that the combined protocol will eventually satisfy the mutual exclusion requirements.

4. RING ORIENTATION IN THE LINK-REGISTER MODEL

In the ring-orientation problem it is required that all nodes agree on an orientation. That is, all nodes should choose a left and right neighbour such that the left neighbour of each node considers that node to be its right neighbour. Thus a self-stabilizing ring-orientation protocol must stabilize to the set of executions $\mathcal{R}_O = \{E = (C_0 C_1 \dots)\}$ where

$$(\forall q \in V_R, \exists p \in V_R :: \text{left}(q) = p \wedge \text{right}(p) = q \text{ is stable in } C_0 \text{ of } E).$$

To obtain a uniform deterministic self-stabilizing ring-orientation protocol for odd-length rings, we use the randomized ring-orientation protocol of Israeli and Jalfon [IJ93]. This protocol is composed of two layers, combined using fair protocol combination, both operating in the link-register model under the distributed daemon. The lower layer is (what we will call) a randomized self-stabilizing *neighbour-ordering* protocol that will stabilize to a state in which any two neighbours agree on the order between them. The second layer is a deterministic self-stabilizing ring-orientation protocol assuming that all neighbours are mutually ordered, using only a constant number of states per node.

Intuitively the second layer of the Israeli and Jalfon protocol operates as follows. Initially certain nodes may hold a token, while others may be about to create tokens. If the ring is not oriented initially, at least one token is present or about to be created. However, once a node has created a token, it will never create a new token.⁴ Each token has a fixed direction in which it travels around the ring. Whenever a node passes

⁴ This may seem to conflict with self-stabilization, but is easily achieved if one makes sure that no configuration in which a token may be generated by a node q can be the result of a step of q .

a token, this token directs the node. If two tokens meet, one of them will be eliminated based on the neighbour-ordering between the two nodes. Due to the elimination of opposite tokens, eventually all tokens will travel in the same direction around the ring, and thus the ring will eventually be oriented. For details we refer to [IJ93].

As an alternative for the lower layer we present a uniform deterministic self-stabilizing neighbour-ordering protocol for odd-length rings, using only a constant number of states per node, in the link-register model under the distributed daemon. Combining this with the second layer of the Israeli and Jalfon protocol yields the desired self-stabilizing ring-orientation protocol for odd-length rings using a constant number of states per node.

4.1. Neighbour Ordering in the Link-Register Model

In the neighbour-ordering problem it is required that eventually any two neighbours p and q agree on an order $<$ between them and that once $p < q$ holds, $p < q$ remains to hold forever. That is, a self-stabilizing neighbour-ordering protocol must stabilize to the set of executions $\mathcal{B}_{NO} = \{E = (C_0 C_1 \dots)\}$ where

$$\forall pq \in E_R :: (p < q \text{ stable in } C_0 \text{ of } E) \vee (p > q \text{ stable in } C_0 \text{ of } E).$$

In this section we develop a uniform deterministic self-stabilizing neighbour-ordering protocol operating in the link-register model under the distributed daemon.⁵ It is based on certain properties of odd-length rings that we derive in the next few paragraphs.

Define for neighbours $p, q \in V_R$ the relations \ll and \equiv by

$$\begin{aligned} p \ll q & \quad \text{if } \text{port}_p(q) < \text{port}_q(p), \text{ and} \\ p \equiv q & \quad \text{if } \text{port}_p(q) = \text{port}_q(p). \end{aligned}$$

Label the edges $pq \in E_R$ (recall that E_R only contains clockwise edges) with $O \in \{\ll, \equiv, \gg\}$ such that pOq if and only if pq is labelled O (see Fig. 2). If we consider chains $q_0 \dots q_{k+1}$, then

$$q_0 \gg q_1 \equiv \dots \equiv q_k \ll q_{k+1} \text{ implies } k \text{ is even, and} \quad (1)$$

$$q_0 \gg q_1 \equiv \dots \equiv q_k \gg q_{k+1} \text{ implies } k \text{ is odd.} \quad (2)$$


This leads us to the following claim.

CLAIM 4.1. *For any ring R with edges labelled as defined above, the number of edges labelled \equiv is even.*

⁵ Under the *central daemon* neighbour ordering is easily achieved deterministically. Let each register R_{pq} contain a value in the set $\{0, 1\}$, and define $p < q$ if $R_{pq} < R_{qp}$. For any edge pq run the following protocol on p (and q , with p and q interchanged):

$$\text{if } R_{pq} = R_{qp} \text{ then invert } R_{pq}.$$

This protocol will stabilize in exactly 1 round.

FIG. 2. Labelling edges with \ll , \equiv , and \gg .

Proof. Let R be an arbitrary ring, and consider neighbouring nodes p, q, r , and s (in that order). If $q \ll r$ it is easily seen that $\text{port}_q(r) = \text{port}_r(s) = 1$ and $\text{port}_r(q) = \text{port}_q(p) = 2$. Thus we can remove all edges not labelled \equiv (merging their endpoints) and the remaining ring has the same number of \equiv -labelled edges as R has. But obviously if all edges in the remaining ring are labelled \equiv , the ring must have an even number of edges. ■

From this claim it follows that any odd-length ring exhibits a certain asymmetry. The construction of our protocol will take this asymmetry as point of departure.

COROLLARY 4.2. *Let R be an odd-length ring, whose edges are labelled as defined above. Then there exists at least one edge labelled with \ll or \gg , and the parity of the number of edges labelled \gg is unequal to the parity of the number of edges labelled \ll .*

Now the straightforward approach to construct a self-stabilizing protocol for ring orientation of odd-length rings might be one in which one generates tokens on \ll - or \gg -labelled edges, letting them travel in the direction of the \ll or \gg . This will not work, however, because the parity-difference between \ll - and \gg -labelled edges on which it is based may be destroyed by two causes: initially extra tokens may be present on \equiv -labelled links, and each \ll -labelled edge has to generate infinitely many tokens, because it can never know it generated one. Apparently our previous corollary alone will not do: we are in need of an additional property of odd-length rings.

DEFINITION 4.3. Mark an arbitrary, nonzero number of edges originally labelled \equiv with \bowtie instead. Then a clockwise chain $p_{l+1}p_l \cdots p_0q_0 \cdots q_rq_{r+1}$ is called a \bowtie -delimited chain K^\bowtie if the edges $p_{l+1}p_l$, p_0q_0 and q_rq_{r+1} are the only edges marked \bowtie in K^\bowtie .

Note that this definition also captures the case in which only one edge is marked \bowtie , because chains are allowed to span the ring more than once. This definition is used in the following lemma to expose yet another source of asymmetry on odd-length rings.

LEMMA 4.4. *Let R be an odd-length ring, whose edges are labelled as defined above. Mark an arbitrary number of edges with \bowtie according to the previous definition. Then there exists a \bowtie -delimited chain $p_{l+1}p_l \cdots p_0q_0 \cdots q_rq_{r+1}$ such that the parity of the number of edges between p_l and p_0 labelled \gg is unequal to the parity of the number of edges between q_0 and q_r labelled \ll .*

Proof. Let λ be the total number of edges labelled \gg and let ρ be the total number of edges labelled \ll . By Corollary 4.2 we know that λ is odd (even) whereas ρ is even (odd). Given a ring R marked with \bowtie , consider the set $\{K_i^\bowtie\}$ of all \bowtie -delimited chains along R . For each K_i^\bowtie let l_i be the number of \gg -labelled edges left of the middle \bowtie , and let r_i be the number of \ll -labelled edges right of the middle \bowtie .

As all edges marked \bowtie were labelled with \equiv , and as all other edges occur exactly once left and exactly once right of the middle \bowtie of some K_j^\bowtie , we see that $\sum l_i = \lambda$ and $\sum r_i = \rho$. Assume λ odd and ρ even (the other case leads to the same result). Then the number of odd l_i must be odd, whereas the number of odd r_i must be even. Therefore there must be a chain K_i^\bowtie for which the parity of l_i does not equal the parity of r_i . ■

Now we are ready to give an informal description of the neighbour-ordering protocol. Each register R_{pq} holds a field $dir \in \{0, 1\}$ such that $R_{pq}.dir < R_{qp}.dir$ if and only if $p < q$. If an edge pq is unordered, i.e., if neither $p < q$ nor $p > q$, then we write $p = q$. If $p \neq q$, neither p nor q will try to order the edge pq . If $p = q$, node q can try to order pq by inverting $R_{qp}.dir$, but under the distributed daemon both p and q might try to order pq simultaneously. In that case the net result will be zero, and if the daemon always schedules p and q simultaneously, pq will never become ordered.

To avoid the above livelock schedule we propose the following. First we make sure that for any p and q with $p \gg q$, only q tries to order pq . By Corollary 4.2 at least one such pq exists, which means that eventually one pair of nodes will be ordered. To order the remaining edges pq with $p \equiv q$, we only allow nodes q that are already ordered with respect to their other neighbour r (i.e., nodes q with $q \neq r$) to try to order pq .

Now there are two cases to consider

1. $o = p = q \neq r$: p will not try to order pq so q must be allowed to do so, or
2. $o \neq p = q \neq r$: the situation for p and q is symmetric and livelock might still occur.

To enable q to tell case 1 and 2 apart, R_{pq} will store whether $o = p$ in a field $ord \in \{\neq, =\}$. To break the symmetry in case 2 we apply Lemma 4.4: the parity of \gg -labelled edges left of p must be unequal to the parity of \ll -labelled edges right of q for at least one edge pq in case 2. We let each node maintain the parity of \gg -labelled edges coming in from the other side in a field $parity \in \{0, 1\}$, and only allow q to try to order pq if the parity it holds equals 1.

4.1.1. The protocol. Let p and r be the neighbours of node q . In the neighbour-ordering protocol each register R_{qp} has the following fields:

- $youare \in \{1, 2\}$, to encode the ordering \ll between p and q . This field is always set such that $R_{qp}.youare = port_q(p)$.
- $dir \in \{0, 1\}$, to encode the desired node-ordering $<$ between p and q .

- $ord \in \{ \neq, = \}$, to tell p whether $q = r$ or not.
- $parity \in \{0, 1\}$, holding the parity of the number of \ll -labelled edges coming in from the other side (i.e., through r).

Thus the protocol uses $2^4 \times 2^4$ states per node. Define for neighbouring nodes p and q

$$\begin{aligned}
 p \ll' q & \quad \text{if } R_{pq}.youare < R_{qp}.youare, \\
 p \equiv' q & \quad \text{if } R_{pq}.youare = R_{qp}.youare, \\
 p < q & \quad \text{if } R_{pq}.dir < R_{qp}.dir, \quad \text{and} \\
 p = q & \quad \text{if } R_{pq}.dir = R_{qp}.dir.
 \end{aligned}$$

These predicates can be evaluated locally at both p and q .

In the neighbour-ordering protocol all nodes run the same program, consisting of two subroutines. The subroutine for a node q with neighbours p and r to update the register used to communicate with p (i.e., R_{qp}) is presented in Protocol 4.1. Node q runs a similar subroutine to update R_{qr} (obtained by swapping p and r in the subroutine code). The subroutine consists of a set of guarded commands, denoted by **if**-statements. Whenever node q is scheduled by the distributed daemon to take a step, all commands (both for updating R_{qp} and R_{qr}) whose guards are true are executed. At the start of each step, q reads R_{pq} and R_{rq} once, and q will write the new contents for R_{qp} and R_{qr} once at the end of each step.

```

1   $R_{qp}.youare := port_q(p)$                                 (* force  $\gg'$  to correspond to  $\gg$  *)
2  if  $q = r$  then  $R_{qp}.parity := 0$                             (* anchor the parity-chain *)
3  if  $q \ll' r \wedge q \neq r$  then  $R_{qp}.parity := (R_{rq}.parity + 1) \bmod 2$ 
                                                                    (* increase parity for incoming  $\ll'$  from other side *)
4  if  $q \not\ll' r \wedge q \neq r$  then  $R_{qp}.parity := R_{rq}.parity$ 
                                                                    (* pass on parity to other side *)
5  if  $q = r$  then  $R_{qp}.ord := '='$  else  $R_{qp}.ord := '\neq'$ 
                                                                    (* pass ordering to other side *)
6  if  $p = q$                                                     (* only try to order not yet ordered pairs *)
7  then if  $p \gg' q$  then invert  $R_{qp}.dir$ 
                                                                    (* note that only the head of  $\gg'$  directs the arc *)
8  else if  $p \equiv' q \wedge q \neq r \wedge R_{pq}.ord = '='$ 
      then invert  $R_{qp}.dir$                                        (* case  $\dots o \approx p = q \neq r \dots$  *)
9  else if  $p \equiv' q \wedge q \neq r \wedge R_{pq}.ord = '\neq' \wedge R_{qp}.parity = 1$ 
      then invert  $R_{qp}.dir$                                        (* case  $\dots o \not\approx p = q \neq r \dots$  *)

```

Protocol 4.1. Neighbour-ordering subroutine for node q to adjust R_{qp} .

4.1.2. *Proof of correctness.* Throughout the proof, let $(P_i)_{i \geq 0}$ be an arbitrary fair schedule under the distributed daemon, let C_0 be an arbitrary initial configuration, and let E be the execution they induce.

CLAIM 4.5. If $p < q$ holds in a configuration C of E , then $p < q$ is stable in C of E .

Proof. The only steps of p or q that can change $R_{pq}.dir$ or $R_{qp}.dir$ (and therefore the ordering between p and q) are those on line 7, 8, and 9. But these three steps are guarded by the condition $p = q$ (see line 6), so the result follows. ■

Thus it remains to be shown that there exists an i such that for all $pq \in E_R$ we have $p \neq q$ in configuration C_i of execution E . We prove this by exhibiting an upper bound on the number of rounds of E after which a configuration satisfying this condition is guaranteed to be reached. This also provides us with an upper bound on the stabilization time, measured in rounds, of the protocol. Recall that $\hat{E}(i)$ is the configuration at the start of round i in execution E .

LEMMA 4.6. For all $pq \in E_R$, if $p \neq q$ then $p \neq q$ in $\hat{E}(2)$.

Proof. First observe that after round 0 each node has taken step 1 at least once, so we have $p \equiv q \Leftrightarrow p \equiv' q$, and similarly $p \gg q \Leftrightarrow p \gg' q$ and $p \ll q \Leftrightarrow p \ll' q$ for all $pq \in E_R$. Clearly these properties are stable in $\hat{E}(1)$. Take an arbitrary edge $pq \in E_R$ with $p \neq q$. Assume $p \gg q$ (the case $p \ll q$ is handled similarly). Then $p \gg' q$ during round 1, which means that p cannot apply steps 7, 8, or 9 on R_{pq} during round 1. If $p \neq q$ in $\hat{E}(1)$, then we are done according to Claim 4.5. Otherwise, q will apply step 7 on R_{qp} in round 1, setting $p \neq q$. ■

This leaves edges pq with $p \equiv q$. We first show that every node q faithfully conveys its order-relation with neighbour r to the other neighbour p .

CLAIM 4.7. Let p and r be the neighbours of q . If $q \neq r$ is stable from $\hat{E}(i)$ to $\hat{E}(j)$, then $R_{qp}.ord = \neq$ is stable from $\hat{E}(i+1)$ to $\hat{E}(j)$. Similarly, if $q = r$ is stable from $\hat{E}(i)$ to $\hat{E}(j)$, then $R_{qp}.ord = =$ is stable from $\hat{E}(i+1)$ to $\hat{E}(j)$.

Proof. If $q \neq r$ is stable from $\hat{E}(i)$ to $\hat{E}(j)$, then q will apply step 5 in rounds i through $j-1$, setting $R_{qp}.ord = \neq$. Thus $R_{qp}.ord = \neq$ is stable from $\hat{E}(i+1)$ to $\hat{E}(j)$. The case for $q = r$ is handled similarly. ■

The next claim shows that if for a certain edge rp we have $r = p$ for l rounds, then all nodes q with distance $k < l$ from p (without another equal pair of nodes in between) correctly store the parity of \gg -labelled edges (i.e., those pointing away from p) between p and q .

CLAIM 4.8. Let $rp_0p_1 \dots p_k p_{k+1}$ be a chain along the ring such that $p_{i-1} \neq p_i$ for all i with $1 \leq i \leq k$. Let $i > 2$ be an arbitrary round and suppose $r = p_0$ is stable from $\hat{E}(i)$ to $\hat{E}(i+n)$. Define $\sigma_k = |\{p_i \mid 1 \leq i \leq k \wedge p_{i-1} \gg p_i\}|$, i.e., the number of \gg -labelled edges between p_0 and p_k . Then $R_{p_k p_{k+1}}.parity = \sigma_k \bmod 2$ is stable from $\hat{E}(i+k+1)$ up to $\hat{E}(i+n)$.

Proof. Proof by induction on k . In the base case $k=0$ we have $\sigma_0=0$. Since $r=p_0$ holds during round i through $i+n-1$, p_0 applies step 2 in these rounds, setting $R_{p_0 p_1}.parity=0$ as required. Thus this property is stable from $\hat{E}(i+1)$ up to $\hat{E}(i+n)$.

Assume the induction hypothesis holds for k , and let $p_k \neq p_{k+1}$ (hence $k < n-1$). Then $R_{p_k p_{k+1}}.parity = \sigma_k \bmod 2$ from $\hat{E}(i+k+1)$ up to $\hat{E}(i+n)$. Then there are two cases to consider. If $p_k \gg p_{k+1}$, then $p_k \gg' p_{k+1}$ stable in $\hat{E}(i+k+1)$ (as $i \geq 2$) and $\sigma_{k+1} = 1 + \sigma_k$. Now p_{k+1} takes step 3 in rounds $i+k+1$ through $i+n-1$, setting $R_{p_{k+1} p_{k+2}}.parity$ to $(R_{p_k p_{k+1}}.parity + 1) \bmod 2$ in these rounds as required. If $p_k \ggg p_{k+1}$, then $p_k \ggg' p_{k+1}$ stable in $\hat{E}(i+k+1)$ (as $i \geq 2$) and $\sigma_{k+1} = \sigma_k$. So p_{k+1} takes step 4 in rounds $i+k+1$ through $i+n-1$, setting $R_{p_{k+1} p_{k+2}}.parity$ to $R_{p_k p_{k+1}}.parity$ in these rounds as required. ■

THEOREM 4.9. *The neighbour-ordering protocol stabilizes on an odd-length ring, under the distributed daemon, to a configuration in which for any two neighbours p, q , $p \neq q$. Furthermore, once $p < q$, then $p < q$ holds forever. The system stabilizes in at most $2 + n^2$ rounds.*

Proof. We prove that in every sequence of n rounds, the number of edges pq with $p=q$ decreases by at least 1, until none are left. Then after at most n^2 rounds, $p \neq q$ holds for all edges $pq \in E_R$. By Claim 4.5 the theorem follows. Towards a contradiction, let $i > 2$ be an arbitrary round, and assume that for all edges pq with $p=q$, $p=q$ is stable from $\hat{E}(i)$ to $\hat{E}(i+n)$. There are two cases to consider.

1. Suppose that in configuration $\hat{E}(i)$ there exist triplets o, p, q with $o = p = q$. By Corollary 4.2 and Lemma 4.6 for at least one we actually have the quartet $o = p = q \neq r$. By Claim 4.7, and by assumption that $o = p$ is stable, then $R_{pq}.ord = "="$ and $R_{qp}.ord = "\neq"$ during round $i+1$. Consequently, in round $i+1$ node p cannot apply any of the steps 7, 8, and 9. On the other hand, as $p=q$ during round $i+1$ by assumption, node q will apply step 8 in round $i+1$, setting $p \neq q$ contrary to assumption.

2. Suppose that in configuration $\hat{E}(i)$ for all edges pq with $p=q$ and $op \in E_R$ and $qr \in E_R$ we have $o \neq p = q \neq r$ (the only other possible case). Now mark all edges $pq \in E_R$ with $p=q$ (and hence $p \equiv q$) with \bowtie . By Lemma 4.4 there exists at least one \bowtie -delimited chain $p_{l+1} = p_l \cdots p_0 = q_0 \cdots q_r = q_{r+1}$ such that the parity of the number of edges between p_l and p_0 labelled \gg is unequal to the parity of the number of edges between q_0 and q_r labelled \ll . Note that for all edges pq between p_l and p_0 or between q_0 and q_r we have $p \neq q$.

Consider the above chain. By assumption $p_{l+1} = p_l$, $p_0 = q_0$, and $q_r = q_{r+1}$ are stable from $\hat{E}(i)$ to $\hat{E}(i+n)$. Then using Claim 4.8, $R_{p_0 q_0}.parity \neq R_{q_0 p_0}.parity$ during round $i+n-1$. By Claim 4.7 also $R_{p_0 q_0}.ord = "\neq"$ and $R_{q_0 p_0}.ord = "\neq"$ during round $i+n-1$. Then in round $i+n-1$ either p_0 or q_0 (but not both) will apply step 9 setting $p \neq q$ contrary to assumption that $p=q$ holds up to $\hat{E}(i+n)$.

This completes the proof. ■

5. RING-ORIENTATION IN THE STATE-READING MODEL

In this section we present a uniform deterministic self-stabilizing ring-orientation protocol for odd-length rings operating in the state-reading model under the central daemon, using only a constant number of states per node. This is the best we can hope for, as Israeli and Jalfon [IJ93] have already shown that such a protocol is impossible under the distributed daemon and under the central daemon if the length of the ring is even. Most ring-orientation protocols operate by forwarding a token with a fixed direction around the ring. But in the state-reading model both neighbours of a node read the same state. If that node were to hold a token with a direction, it is not immediately obvious how to forward that token to the one neighbour it points to without possibly forwarding it to the other neighbour as well. In view of this observation it is perhaps surprising that it is possible at all to orient a ring in the state reading model.

Intuitively the protocol operates as follows. Let each node have a *colour* taken from the set $\{0, 1\}$. Try to give neighbouring nodes alternating colours by inverting the colour of a node if it has the same colour as both its neighbours.⁶ Of course on an odd-length ring this is never completely possible, so we are bound to end up with patterns like 001 and 110 around the ring. Call such patterns tokens. It is worth noting that Herman [Her90] used the same idea to achieve probabilistic self-stabilization to a single token on odd length, but oriented, rings.

Now the idea is to make these tokens travel around the ring, orienting the nodes they visit. This requires us to impose a direction on tokens, for which we let each node store a *phase* taken from the set $\{+, -\}$. A token is directed if the nodes with equal colour have opposite phase: then the direction of the token points from the node with phase $+$ to the node with phase $-$. Undirected tokens can be directed by letting the middle node in a token invert its phase.

Let us write 0_- for a node with colour 0 and phase $-$, and consider the pattern $0_+ 0_- 1_- 0_-$ around the ring. If we allow the second node to change its state to 1_+ we get the pattern $0_- 1_+ 1_- 0_-$: the token has moved one step in its direction. Now let a node also maintain its orientation, taken from $\{\leftarrow, \rightarrow\}$ —for instance by storing the port of the neighbour the head of the arrow points to. If a token moves one step, the node changing colour is oriented into the direction of the token. Our protocol depends on the fact that a token always keeps the same direction, until it is eliminated. Then in the situation $0_+ 0_- 1_+ 0_-$ it is unwise to allow the second node to change its state to 1_+ as this would yield the situation $0_+ 1_+ 1_+ 0_-$. Here the token becomes undirected and thus may decide to invert its direction: $0_+ 1_- 1_+ 0_-$. Therefore, in situations like $0_+ 0_- 1_+ 0_-$ we must wait until the third node sets its phase to $-$.

We have already seen that in almost all configurations (except for the special case in which all nodes have the same colour) at least one token exists. If we can make sure that eventually all tokens travel in the same direction around the ring, the ring will eventually become oriented. Now consider what happens when two tokens with opposite direction meet. This situation is depicted in Fig. 3 (where the steps are

⁶This trick cannot be applied immediately under the distributed daemon, because the distributed daemon is free to schedule all nodes simultaneously. If all nodes in the ring have the same colour, then under such a schedule all nodes will simultaneously invert their colour.

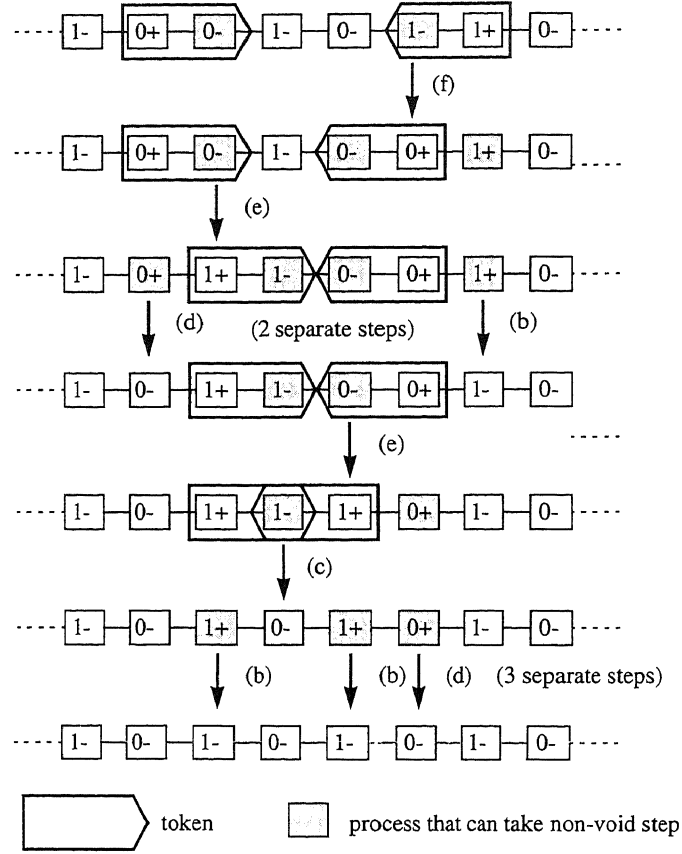


FIG. 3. Example of two tokens with opposite direction meeting.

taken from the complete description of the protocol in Protocol 5.1). Other nodes in this example may take a step instead, but this leads to essentially the same situation. We see that both tokens are eliminated altogether.

5.1. The Protocol

Let p and r be the neighbours of q . Each node q stores its whole state (i.e., its colour, phase, and orientation) in register R_q . In the ring-orientation protocol, all nodes q run the same program, presented in Protocol 5.1 as a state-transition function δ . The tables list the applicable steps for a node q with neighbours p and r , depending on the state of p , q , and r (i.e., the contents of R_p , R_q , and R_r). The entries under q' list the new state of q after applying δ . To reduce the size, symmetric steps (with p and r interchanged) are not shown. So steps (e) through (j) actually represented two steps. If only some components of a state are specified, the other components may have an arbitrary value. If in the new state the orientation is not specified, it should equal the direction of q in the old state. In cases (e) and (f), the direction of q in the new state should be read as pointing from p to r . Node q can always determine this orientation because in these steps p and r hold opposite colours. The port of p and r is used to encode the direction.

δ	p	q	r	q'	δ	p	q	r	q'	δ	p	q	r	q'
(a)	0	0	0	1 ₋	(e)	0 ₊	0 ₋	1 ₋	$\overrightarrow{1_+}$	(g)	0 ₋	0 ₋	1	0 ₊
(b)	0	1	0	1 ₋	(f)	1 ₊	1 ₋	0 ₋	$\overrightarrow{0_+}$	(h)	0 ₊	0 ₊	1	0 ₋
(c)	1	1	1	0 ₋						(i)	1 ₋	1 ₋	0	1 ₊
(d)	1	0	1	0 ₋						(j)	1 ₊	1 ₊	0	1 ₋

Protocol 5.1. Ring-orientation program for node q .

Each table contains related steps. The first table lists the steps that try to colour the ring alternately and to make sure that alternately coloured nodes have phase $-$. The second table lists the steps responsible for forwarding the token, and the third table lists the steps that break the symmetry in possible tokens by imposing one direction on them.

5.2. Proof of Correctness

Throughout the proof, let $(P_i)_{i \geq 0}$ be an arbitrary fair schedule under the central daemon, let C_0 be an arbitrary initial configuration, and let E be the execution of the protocol they induce. If $n = 1$, the protocol is trivially correct, so in the remainder of the proof we assume $n \geq 3$. We prove correctness of the algorithm in stages. In each stage we define a set of configurations and show that the protocol will converge to this set, when started in a configuration from the set of the previous stage. All sets are shown to be closed under transitions of the protocol. In the final stage the set of configurations will contain exactly those that are oriented. To describe the configurations in a set we use regular expressions.

DEFINITION 5.1. A configuration *matches* a regular expression L if the concatenation of the states of some chain (clockwise or anticlockwise) of length n matches the regular expression. A regular expression L is *closed* under steps of the protocol if for all configurations C that match L , C' with $C \rightarrow_P C'$ for arbitrary allowed P matches L as well.

In other words, a configuration matches a regular expression if we can cut the ring between a pair of nodes and if the resulting chain, or string, of states matches the regular expression.

Define the regular expression L_1 by

$$\begin{aligned}
 L_1 &= (O_1 I_1)^*, \\
 O_1 &= 0 \mid 0_+ 0_- 0_+ \mid 0_- 0_+ \mid 0_+ 0_-, \quad \text{and} \\
 I_1 &= 1 \mid 1_+ 1_- 1_+ \mid 1_- 1_+ \mid 1_+ 1_-.
 \end{aligned}$$

LEMMA 5.2. *Starting in an arbitrary configuration, we eventually reach a configuration matching L_1 . Furthermore, L_1 is closed.*

Proof. Consider 0-chains, i.e., chains of nodes coloured only with 0 that are delimited at both ends by 1-coloured nodes. If all nodes in the ring are coloured 0,

only step (a) will be applicable by all nodes creating one 0-chain after the first node takes a step. For 0-chains define the width w as

$$\begin{aligned} w(0_+) &= 1.1, \\ w(0_-) &= 1, \quad \text{and} \\ w(0_1 \cdots 0_k) &= k + p(0_1 0_2) + p(0_k 0_{k-1}), \end{aligned}$$

where the penalty p for the endpoints of the chain is given by

$$p(ab) = \begin{cases} 1.5 & \text{if } a = 0_-, \\ 1.6 & \text{if } ab = 0_+ 0_+, \\ 0 & \text{otherwise.} \end{cases} \quad \text{and}$$

1-chains and their associated penalty and width are defined similarly. It is easily checked that all 0-chains (1-chains) K matching O_1 (I_1) are exactly those⁷ for which $w(K) \leq 3.5$. So it remains to show that for any chain K with $w(K) > 3.5$, every nonvoid step taken by a node on that chain will decrease its width and will not create new chains with width greater than 3.5.

(a) Creates two 0-chains, with at least 2 0-coloured nodes less than the original chain, whereas the maximal penalty of the new endpoint for both chains is at most 1.6. Hence the width of both chains is less than the original.

(b) Does not change any 0-chains.

(c) Creates a 0-chain with width 1.

(d) If (d) is nonvoid, it changes 0_+ with width 1.1, to 0_- with width 1.

(e) Decreases the length of the chain by 1, but the endpoint may change from 000_- to $0_+ 0_+$ resulting in an increase of penalty of at most 0.1.

(f) If 0_- in (f) is a complete 0-chain on its own, (f) changes this to chain $0_+ 0_-$ with width 3.5. Otherwise, it increases the length of the 0-chain (with endpoint 0_- and hence penalty 1.5) by 1 to a chain with endpoint $0_+ 0_-$ and hence penalty 0. So the width decreases by 0.5.

(g) Decreases the penalty of the endpoint from 1.5 to 0.

(h) Decreases the penalty of the endpoint from 1.6 to 1.5.

(i), (j) Do not change any 0-chains.

This completes the proof. ■

All configurations matching the regular expression L_2 defined below only contain tokens in one direction.

$$L_2 = (O_2 I_2)^*, \quad O_2 = 0 \mid 0_+ 0_-, \quad \text{and} \quad I_2 = 1 \mid 1_+ 1_-.$$

⁷ All chains of 4 or more 0s have width exceeding 4, $w(0_- 00) \geq 4.5$, $w(0_+ 0_+) = 5.2$, $w(0_- 0_+) = 3.5$, and $w(0_- 0_-) = 5$.

LEMMA 5.3. *From a configuration matching L_1 , we eventually reach a configuration matching L_2 . Furthermore, L_2 is closed.*

Proof. We prove convergence and closure similar to the proof of Lemma 5.2. To prove that the system eventually reaches a configuration matching L_2 , consider a configuration matching L_1 . This configuration may fail to match L_2 for two reasons. First of all it may contain patterns like $0_+0_-0_+$ and $1_+1_-1_+$, and secondly it may contain tokens (0_+0_- and 1_+1_-) in opposite directions.

If there are tokens with opposite directions on the ring, uniquely match pairs of opposite tokens in the following manner (the result is similar to the matching of braces in expressions). Let K be a chain that does not contain any tokens. Then the opposing tokens a_+a_- and b_-b_+ in the chain $a_+a_-Kb_-b_+$ are defined to match. Inductively, the same holds if K contains only pairs of matching opposing tokens and no single unmatched tokens. The distance between two matching opposing tokens equals 1 plus the length of the intermediate chain K .

For 0-chains matching L_2 define the entropy e as follows

$$\begin{aligned} e(0) &= 0, \\ e(0_+0_-0_+) &= 1/n, \\ e(0_-0_+) &= e(0_+0_-) = \text{the distance to its matching token,} \\ &\quad 0 \text{ if it does not exist.} \end{aligned}$$

The entropy of 1-chains is defined similarly. The total entropy of a configuration is defined to be the sum of the entropies of its 0 and 1 chains. Observe that configurations that match L_2 are exactly those for which the corresponding entropy equals 0. The number of matching opposing tokens is at most $n/4$, with maximum entropy $n-3$ each. Therefore the maximum entropy is bounded by n^2 . We consider the effect of each step of Protocol 5.1 on the entropy.

(a), (c) Decrease the entropy from $1/n$ to 0 ($10_+0_-0_+1$ becomes $10_+1_-0_+1$; three chains with entropy 0 are created).

(b), (d) Do not change the entropy.

(e) If the token matches with an opposing token, with distance larger than 1, this distance decreases by 1. $10_+0_-1_-0$ changes to $10_+1_+1_-0$ moving the token one step into its direction. If the distance equals 1, then $10_+0_-1_-1_+0$ changes to $10_+1_+1_-1_+0$. Thus the entropy decreases from 2 (1 for each token) to $1/n$. If the token does not match with an opposing token, then the entropy does not change. However, for as long as there are opposing tokens, in each round one of them will decrease its entropy. Refer also to Fig. 3 for an example.

(f) Similar to (e).

(g), (h), (i), (j) Do not apply to chains matching L_2 .

This completes the proof. ■

All configurations matching the regular expression L_3 defined below are oriented.

$$L_3 = (O_3 I_3)^*, \quad O_3 = \vec{0} \mid \vec{0_+} \vec{0_-}, \quad \text{and} \quad I_3 = \vec{1} \mid \vec{1_+} \vec{1_-}.$$

LEMMA 5.4. *From a configuration matching L_2 , we eventually reach a configuration matching L_3 . Furthermore, L_3 is closed.*

Proof. Closure is proven as in Lemma 5.3, noting that (e) and (f) do not change the orientation. To prove that the system eventually reaches a configuration matching L_3 , note that because the ring has odd length, at least one token must exist, and that for at least one token we have $0_+ 0_- 10$ or $1_+ 1_- 01$. Thus the head of at least one token can eventually take a step, moving the token one position into the direction of the token and directing the former head as well. Eventually one token must have travelled around the ring completely, at which time the ring is oriented. ■

From Lemmas 5.2, 5.3, and 5.4 we easily obtain the following theorem.

THEOREM 5.5. *The ring-orientation protocol stabilizes, under the central daemon, to a configuration in which the ring is oriented, provided the length of the ring is odd.*

This theorem has a curious consequence. Dijkstra [Dij82] showed that no uniform deterministic self-stabilizing mutual exclusion protocol exists for rings of nonprime size. Burns and Pachl [BP89] complemented this impossibility result with a uniform self-stabilizing mutual exclusion protocol for *oriented* rings of prime size, operating in the state-reading model under the central daemon. Combining the protocol of Burns and Pachl with our second ring-orientation protocol using fair protocol combination (cf. Section 3) proves the following theorem.

THEOREM 5.6. *On unoriented rings of prime size, self-stabilizing mutual exclusion can be achieved under the central daemon using a uniform protocol. Moreover, using the protocol of Itkis et al. [ILS95] instead, the total number of states per processor can even be kept constant.*

6. ON THE EQUIVALENCE OF SELF-STABILIZING SYSTEM MODELS

An overwhelming amount of models for distributed systems can be found in the literature, some of which only differ on seemingly minor points. This diversity is caused by the fact that slight alterations to a model may have a huge effect on the (im)possibility or (in)efficiency of certain protocols.⁸ A major challenge is to explore these different models and find conditions or application areas under which these models are equivalent. In the area of self-stabilization, research in this area has already been started by Gouda *et al.* [GHR90]. In this section we will show that on oriented rings the link-register and state-reading model are equivalent. Using the results of the previous sections, we also show that, as a consequence, the link-register model and the state-reading are eventually equivalent on odd-length rings under the central daemon.

What do we mean by equivalence between models? Several definitions seem to be appropriate (cf. [GHR90, LV93]). Intuitively two models are equivalent if they are of equal strength: whatever is possible in one model is also possible in the other, and vice versa. We adopt the following formal definitions.

⁸ As exemplified by the large body of literature on how to reach agreement in the presence of faults.

DEFINITION 6.1. Protocol P_1 *simulates* protocol P_2 if there exists a recursive mapping μ such that for all executions E_1 of P_1 , $\mu(E_1)$ is an execution of P_2 . Protocol P_1 *eventually simulates* protocol P_2 if there exists a mapping μ such that for all executions E_1 of P_1 , a suffix of $\mu(E_1)$ is an execution of P_2 . Protocols P_1 and P_2 are (eventually) *equivalent* if both (eventually) simulate each other. System model M_1 (eventually) *simulates* system model M_2 if for all deterministic protocols P_2 on M_2 there exists a deterministic protocol P_1 on M_1 that (eventually) simulates P_2 . System models M_1 and M_2 are (eventually) *equivalent* if both (eventually) simulate each other.

We choose to show simulation of M_2 by M_1 by giving a general method to convert a protocol in M_2 to a simulating protocol in M_1 and describing the mapping μ from executions in M_1 to M_2 . As in the definition, both models are equivalent if we can show that both simulate each other.

THEOREM 6.2. Let $G=(V, E)$ be an undirected graph, where each node $p \in V$ labels its edges $pq \in E$ with $lab_p(q)$. Suppose there exists a function f such that for all $pq \in E$ we have $lab_p(q) = f(lab_q(p))$, and suppose that for all $pq, qr \in E$ with $pq \neq pr$ we have $lab_p(q) \neq lab_p(r)$. Then on G , the state-reading model and the link-register model are equivalent.

Proof. Consider an arbitrary node q in G . A protocol in the state-reading model is easily transformed into an equivalent protocol in the link-register model, by changing all writes to R_q into writes of the same value to all registers R_{qp} with $qp \in E$ and changing all reads from R_p , for some $qp \in E$, to reads from R_{pq} . Initially, for all $qp \in E$, the contents of R_{qp} should equal the contents of R_q . Then μ maps all executions in the link-register model to executions in the state-reading model by mapping the contents of R_{qp} for all $qp \in E$ (that by construction always hold the same value) to the contents of R_q .

A protocol in the link-register model is transformed into an equivalent protocol in the state-reading model as follows. Split, in the state-reading model, the register R_q into as many fields $R_q.to[\cdot]$ as there are edges $qp \in E$. Replace a write to R_{qp} by a write of the same value to $R_q.to[lab_q(p)]$. Change reads from R_{qp} to reads from $R_q.to[f(lab_p(q))]$. Then μ maps all executions in the state-reading model to executions in the link-register model by mapping the contents of $R_q.to[lab_q(p)]$ to the contents of R_{qp} . The result of applying μ is indeed an execution in the link register model because (i) if $qp \neq qr$ then $lab_q(p) \neq lab_q(r)$ so no write is overwritten by a wrong write, and (ii) $lab_p(q) = f(lab_q(p))$ so a value read from R_{qp} equals $R_q.to[f(lab_p(q))] = R_q.to[lab_q(p)]$ which equals the value written by a write to R_{qp} . ■

Oriented rings and also cliques and hypercubes with sense of direction have a labelling as in the above theorem. From the self-stabilizing ring-orientation protocol presented in the previous section, we easily obtain the following corollary.

COROLLARY 6.3. For odd-length rings, the link-register and state-reading model are eventually equivalent under the central daemon.

Simply combine the simulation in the proof of Theorem 6.2 with the ring-orientation protocol for the state-reading model. Note that the simulation of the state-reading protocol by the link-register protocol ensures that the contents of R_{qp} and R_{qr} are equal. A transient error may disturb this invariant, but equality will be re-established as soon as q takes its first step after the error.

Observe that the ring-orientation protocol for the state-reading model requires a central daemon, so Corollary 6.3 only holds under the central daemon. A similar corollary cannot be obtained for odd-length rings under the distributed daemon, because Israeli and Jalfon already showed that no ring in the state-reading model can be oriented deterministically under the distributed daemon.

7. CONCLUSIONS AND FURTHER RESEARCH

We have shown that there exist uniform deterministic self-stabilizing ring-orientation protocols using only a constant number of states per node for odd-length rings, both in the link-register model under the distributed daemon and in the state-reading model under the central daemon. Further research might be directed at deriving similar protocols for other graphs with a regular structure, such as cliques or hypercubes.

We have also shown that the link-register model and the state-reading model are eventually equivalent on odd-length rings under the central daemon, showing that a self-stabilizing protocol designed for the one model can be transformed to an equivalent, self-stabilizing protocol for the other model. We are unaware of similar theorems exploring the relation between the central daemon and the distributed daemon.

ACKNOWLEDGMENTS

The author thanks Marina Papatriantafyllou and Philippas Tsigas for the discussions on this problem while they were visiting the CWI in Amsterdam. The anonymous referees are gratefully acknowledged for the valuable suggestions that helped improve the paper.

Received July 5, 1995; final manuscript received October 23, 1997

REFERENCES

- [ASW88] Attiya, H., Snir, M., and Warmuth, M. K. (1988), Computing on an anonymous ring, *J. Assoc. Comput. Mach.* **35**(4), 845–875.
- [Bur87] Burns, J. E. (1987), Self-stabilizing rings without demons, Tech. Rep. GIT-ICS-87/36, School of Information and Computer Science, Georgia Institute of Technology, Atlanta, Georgia.
- [BGM93] Burns, J. E., Gouda, M. G., and Miller, R. E. (1993), Stabilization and pseudo-stabilization, *Distrib. Comput.* **7**(1), 35–42.
- [BP89] Burns, J. E., and Pachl, J. (1989), Uniform self-stabilizing rings, *ACM Trans. Program. Languages Systems* **11**(2), 330–344.
- [Dij74] Dijkstra, E. W. (1974), Self-stabilizing systems in spite of distributed control, *Comm. Assoc. Comput. Mach.* **17**(11), 643–644.
- [Dij82] Dijkstra, E. W. (1982), Self-stabilization in spite of distributed control, in “Selected Writings on Computing: A Personal Perspective,” pp. 41–46, Springer-Verlag, New York.

- [DIM93] Dolev, S., Israeli, A., and Moran, S. (1993), Self-stabilization of dynamic systems assuming only read/write atomicity, *Distrib. Comput.* **7**(1), 3–16.
- [GHR90] Gouda, M. G., Howell, R. R., and Rosier, L. E. (1990), The instability of self-stabilization, *Acta Inform.* **27**(8), 697–724.
- [Her90] Herman, T. (1990), Probabilistic self-stabilization, *Inform. Process. Lett.* **35**(35), 63–67.
- [Hoe94] Hoepman, J.-H. (1994), Uniform deterministic self-stabilizing ring-orientation on odd-length rings, in “8th Int. Workshop on Distributed Algorithms, Terschelling, The Netherlands” (G. Tel and P. M. B. Vitányi, Eds.), Lecture Notes in Computer Science, Vol. 857, pp. 265–279, Springer-Verlag, Berlin/New York.
- [IJ93] Israeli, A., and Jalfón, M. (1993), Uniform self-stabilizing ring orientation, *Inform. and Comput.* **104**(2), 175–196.
- [ILS95] Itkis, G., Lin, C., and Simon, J. (1995), Deterministic, constant space, self-stabilizing leader election on uniform rings, in “9th Int. Workshop on Distributed Algorithms, Le Mont-Saint-Michel, France” (J.-M. Hélary and M. Raynal, Eds.), Lecture Notes in Computer Science, Vol. 972, pp. 288–302, Springer-Verlag, Berlin/New York.
- [LV93] Lynch, N. A., and Vaandrager, F. W. (1993), Forward and backward simulations, Part I: Untimed systems, Tech. Rep. CS-R9313, Stichting Mathematisch Centrum (CWI), Amsterdam.
- [San84] Santoro, N. (1984), Sense of direction, topological awareness and communication complexity, *ACM SIGACT News* **16**(2), 50–56.
- [SP87] Syrotiuk, V., and Pachl, J. (1987), A distributed ring orientation algorithm, in “2nd Int. Workshop on Distributed Algorithms, Amsterdam, The Netherlands” (J. van Leeuwen, Ed.), Lecture Notes in Computer Science, Vol. 312, pp. 332–336, Springer-Verlag, Berlin/New York.
- [Tel94] Tel, G. (1994), “Introduction to Distributed Algorithms,” Cambridge Univ. Press, Cambridge, UK.
- [TH95] Tsai, M.-S., and Huang, S.-T. (1995), Self-stabilizing ring orientation protocols, in “2nd Workshop on Self-Stabilizing Systems, Las Vegas, NV,” Dept. of Comp. Science, University of Nevada, Las Vegas, Box 454019 Las Vegas, NV 89154-4019, pp. 16.1–16.14.