

A Distributed TDMA Slot Assignment Algorithm for Wireless Sensor Networks

Ted Herman¹ and Sébastien Tixeuil²

¹ University of Iowa, ted-herman@uiowa.edu

² LRI – CNRS UMR 8623 & INRIA Grand Large, Université Paris-Sud XI, France, tixeuil@lri.fr

Abstract. Wireless sensor networks benefit from communication protocols that reduce power requirements by avoiding frame collision. Time Division Media Access methods schedule transmission in slots to avoid collision, however these methods often lack scalability when implemented in *ad hoc* networks subject to node failures and dynamic topology. This paper reports a distributed algorithm for TDMA slot assignment that is self-stabilizing to transient faults and dynamic topology change. The expected local convergence time is $O(1)$ for any size network satisfying a constant bound on the size of a node neighborhood.

1 Introduction

Collision management and avoidance are fundamental issues in wireless network protocols. Networks now being imagined for sensors [24] and small devices [3] require energy conservation, scalability, tolerance to transient faults, and adaptivity to topology change. Time Division Media Access (TDMA) is a reasonable technique for managing wireless media access, however the priorities of scalability and fault tolerance are not emphasized by most previous research. Recent analysis [8] of radio transmission characteristics typical of sensor networks shows that TDMA may not substantially improve bandwidth when compared to randomized collision avoidance protocols, however fairness and energy conservation considerations remain important motivations. In applications with predictable communication patterns, a sensor may even power off the radio receiver during TDMA slots where no messages are expected; such timed approaches to power management are typical of the sensor regime.

Emerging models of *ad hoc* sensor networks are more constrained than general models of distributed systems, especially with respect to computational and communication resources. These constraints tend to favor simple algorithms that use limited memory. A few constraints of some sensor networks can be helpful: sensors may have access to geographic coordinates and a time base (such as GPS provides), and the density of sensors in an area can have a known, fixed upper bound. The question we ask in this paper is how systems can distributively obtain




Fig. 1. Two solutions to distance-two coloring

a TDMA assignment of slots to nodes, given the assumptions of synchronized clocks and a bounded density (where density is interpreted to be a fixed upper bound on the number of immediate neighbors in the communication range of any node). In practice, such a limit on the number of neighbors in range of a node has been achieved by dynamically attenuating transmission power on radios. Our answers to the question of distributively obtaining a TDMA schedule are partial: our results are not necessarily optimum, and although the algorithms we present are self-stabilizing, they are not optimally designed for all cases of minor disruptions or changes to a stabilized sensor network.

Before presenting our results, it may be helpful for the reader to consider the relation between TDMA scheduling and standard problems of graph coloring (since these topics often found in textbooks on network algorithms for spatial multiplexing). Algorithmic research on TDMA relates the problem of timeslot assignment to minimal graph coloring where the coloring constraint is typically that of ensuring that no two nodes within distance two have the same color (the constraint of distance two has a motivation akin to the well known hidden terminal problem in wireless networks). This simple reduction of TDMA timeslot assignment neglects some opportunities for time division: even a solution to minimum coloring does not necessarily give the best result for TDMA slot assignment. Consider the two colorings shown in Figure 1, which are minimum distance-two colorings of the same network. We can count, for each node p , the size of the set of colors used within its distance-two neighborhood (where this set includes p 's color); this is illustrated in Figure 2 for the respective colorings of Figure 1. We see that some of the nodes find more colors in their distance-two neighborhoods in the second coloring of Figure 1. The method of slot allocation in Section 6 allocates larger bandwidth share when the number of colors in distance-two neighborhoods is smaller. Intuitively, if some node p sees $k < \lambda$ colors in its distance-two neighborhood, then it should have at least a $1/(k+1)$ share of bandwidth, which is superior to assigning a $1/(\lambda+1)$ share to each color. Thus the problem of optimum TDMA slot assignment is, in some sense, harder than optimizing the global number of colors.

Contributions. The main issues for our research are dynamic network configurations, transient fault tolerance and scalability of TDMA slot assignment algorithms. Our approach to both dynamic network change and transient fault events is to use the paradigm of self-stabilization, which ensures the system




Fig. 2. Number of colors used within distance two

state converges to a valid TDMA assignment after any transient fault or topology change event. Our approach to scalability is to propose a randomized slot assignment algorithm with $O(1)$ expected *local* convergence time. The basis for our algorithm is, in essence, a probabilistically fast clustering technique (which could be exploited for other problems of sensor networks). The expected time for *all* nodes to have a valid TDMA assignment is not $O(1)$; our view is that stabilization over the entire network is an unreasonable metric for sensor network applications; we discuss this further in the paper's conclusion. Our approach guarantees that after stabilization, if nodes crash, TDMA collision may occur only locally (in the distance-three neighborhood of the faults).

Related Work. The idea of self-stabilizing TDMA has been developed in [13, 14] for model that is more restricted than ours (a grid topology where each node knows its location). Algorithms for allocating TDMA time slots and FDMA frequencies are formulated as vertex coloring problems in a graph [17]. Let the set of vertex colors be the integers from the range $0..λ$. For FDMA the colors (f_v, f_w) of neighboring vertices (v, w) should satisfy $|f_v - f_w| > 1$ to avoid interference. The standard notation for this constraint is $L(\ell_1, \ell_2)$: for any pair of vertices at distance $i \in \{1, 2\}$, the colors differ by at least ℓ_i . The coloring problem for TDMA is: let $L'(\ell_1, \ell_2)$ be the constraint that for any pair of vertices at distance $i \in \{1, 2\}$, the colors differ by at least $\ell_i \bmod (\lambda + 1)$. (This constraint represents the fact that time slots wrap around, unlike frequencies.) The coloring constraint for TDMA is $L'(1, 1)$. Coloring problems with constraints $L(1, 0)$, $L(0, 1)$, $L(1, 1)$, and $L(2, 1)$ have been well-studied not only for general graphs but for many special types of graphs [2,11,18]; many such problems are NP-complete and although approximation algorithms have been proposed, such algorithms are typically not distributed. (The related problem finding a minimum dominating set has been shown to have a distributed approximation using constant time [12], though it is unclear if the techniques apply to self-stabilizing coloring.) Self-stabilizing algorithms for $L(1, 0)$ have been studied in [5,21,19,20, 7], and for $L(1, 1)$ in [6]. Our algorithms borrow from techniques of self-stabilizing coloring and renaming [6,7], which use techniques well-known in the literature of parallel algorithms on PRAM models [15]. To the extent that the sensor network model is synchronous, some of these techniques can be adapted; however working out details when messages collide, and the initial state is unknown, is not an entirely trivial task. This paper is novel in the sense that it composes self-stabilizing algorithms for renaming and coloring for a base model that has

only probabilistically correct communication, due to the possibility of collisions at the media access layer. Also, our coloring uses a constant number of colors for the $L(1, 1)$ problem, while the previous self-stabilizing solution to this problem uses n^2 colors. Due to space constraints, proofs are delegated to [9].

2 Wireless Network, Program Notation

The system is comprised of a set V of nodes in an *ad hoc* wireless network, and each node has a unique identifier. Communication between nodes uses a low-power radio. Each node p can communicate with a subset $N_p \subseteq V$ of nodes determined by the range of the radio signal; N_p is called the neighborhood of node p . In the wireless model, transmission is omnidirectional: each message sent by p is effectively broadcast to all nodes in N_p . We also assume that communication capability is bidirectional: $q \in N_p$ iff $p \in N_q$. Define $N_p^1 = N_p$ and for $i > 1$, $N_p^i = N_p^{i-1} \cup \{r \mid (\exists q : q \in N_p^{i-1} : r \in N_q)\}$ (call N_p^i the distance- i neighborhood of p). Distribution of nodes is sparse: there is some known constant δ such that for any node p , $|N_p| \leq \delta$. (Sensor networks can control density by powering off nodes in areas that are too dense, which is one aim of topology control algorithms.)

Each node has fine-grained, real-time clock hardware, and all node clocks are synchronized to a common, global time. Each node uses the same radio frequency (one frequency is shared spatially by all nodes in the network) and media access is managed by CSMA/CA: if node p has a message ready to transmit, but is receiving some signal, then p does not begin transmission until it detects the absence of signal; and before p transmits a message, it waits for some random period (as implemented, for instance, in [23]). We assume that the implementation of CSMA/CA satisfies the following: there exists a constant $\tau > 0$ such that the probability of a frame transmission without collision is at least τ (this corresponds to typical assumptions for multiaccess channels [1]; the independence of τ for different frame transmissions indicates our assumption of an underlying memoryless probability distribution in a Markov model).


Notation. We describe algorithms using the notation of guarded assignment statements: $G \rightarrow S$ represents a guarded assignment, where G is a predicate of the local variables of a node, and S is an assignment to local variables of the node. If predicate G (called the *guard*) holds, then assignment S is executed, otherwise S is skipped. Some guards can be event predicates that hold upon the event of receiving a message: we assume that all such guarded assignments execute atomically when a message is received. At any system state where a given guard G holds, we say that G is *enabled* at that state. The \square operator is the nondeterministic composition of guarded assignments; $(\square q : q \in M_p : G_q \rightarrow S_q)$ is a closed-form expression of $G_{q_1} \rightarrow S_{q_1} \square G_{q_2} \rightarrow S_{q_2} \square \dots \square G_{q_k} \rightarrow S_{q_k}$, where $M_p = \{q_1, q_2, \dots, q_k\}$.

Execution Semantics. The life of computing at every node consists of the infinite repetition of finding a guard and executing its corresponding assignment or skipping the assignment if the guard is *false*. Generally, we suppose that when a node executes its program, all statements with *true* guards are executed in some constant time (done, for example, in round-robin order).

Shared Variable Propagation. A certain subset of the variables at any node are designated as *shared* variables. Nodes periodically transmit the values of their shared variables, based on a timed discipline. Given the discipline of repeated transmission of shared variables, each node can have a cached copy of the value of a shared variable for any neighbor. This cached copy is updated atomically upon receipt of a message carrying a new value for the shared variable.

Model Construction. Our goal is to provide an implementation of a general purpose, collision-free communication service. This service can be regarded as a transformation of the given model of Section 2 into a model without collisions. This service simplifies application programming and can reduce energy requirements for communication (messages do not need to be resent due to collisions). Let \mathcal{T} denote the task of transforming the model of Section 2 into a collision-free model.

To solve \mathcal{T} it suffices to assign each node a color and use node colors as the schedule for a TDMA approach to collision-free communication [17]. Even before colors are assigned, we use a schedule that partitions radio time into two parts: one part is for TDMA scheduling of application messages and the other part is reserved for the messages of the algorithm that assigns colors and time slots to nodes. The following diagram illustrates such a schedule, in which each TDMA part has five slots. Each overhead part is, in fact, a fixed-length slot in the TDMA schedule.



The programming model, including the technique for sharing variables, refers to message and computation activity in the overhead parts. Whereas CSMA/CA is used to manage collisions in the overhead slots, the remaining TDMA slots do not use random delay. During initialization or after a dynamic topology change, frames may collide in the TDMA slots, but after the slot assignment algorithm self-stabilizes, collisions do not occur in the TDMA slots.

With respect to any given node v , a solution \mathcal{T} is *locally stabilizing* with convergence time t if, for any initial system state, after at most t time units, every subsequent system state satisfies the property that any transmission by v during its assigned slot(s) is free from collision. Solution \mathcal{T} is *globally stabilizing* with

convergence time t if, for every initial state, after at most t time units, every subsequent system state has the property that all transmissions during assigned slots are free from collision. For randomized algorithms, these definitions are modified to specify expected convergence times (all stabilizing randomized algorithms we consider are probabilistically convergent in the Las Vegas sense). When the qualification (local or global) is omitted, convergence times for local stabilization are intended for the presented algorithms.

Several primitive services that are not part of the initial model can simplify the design and expression of \mathcal{T} 's implementation. All of these services need to be self-stabilizing. Briefly put, our plan is to develop a sequence of algorithms that enable TDMA implementation. These algorithms are: neighborhood-unique naming, maximal independent set, minimal coloring, and the assignment of time slots from colors. In addition, we rely on neighborhood services that update cached copies of shared variables.

Neighborhood Identification. We do not assume that a node p has built-in knowledge of its neighborhood N_p or its distance-three neighborhood N_p^3 . This is because the type of network under considering is *ad hoc*, and the topology dynamic. Therefore some algorithm is needed so that a node can refer to its neighbors. We describe first how a node p can learn of N_p^2 , since the technique can be extended to learn N_p^3 in a straightforward way.

Each node p can represent N_p^i for $i \in 1..3$ by a list of identifiers learned from messages received at p . However, because we do not make assumptions about the initial state of any node, such list representations can initially have arbitrary data. Let L be a data type for a list of up to δ items of the form $a : A$, where a is an identifier and A is a set of up to δ identifiers. Let sL_p be a shared variable of type L . Let message type mN with field of type L be the form of messages transmitted for sL_p . Let L_p be a private variable of a type that is an augmentation of L – it associates a real number with each item: $age(a : A)$ is a positive real value attached to the item.

Function $\text{update}(L_p, a : A)$ changes L_p to have new item information: if L_p already has some item whose first component is a , it is removed and replaced with $a : A$ (which then has age zero); if L_p has fewer than δ items and no item with a as first component, then $a : A$ is added to L_p ; if L_p has already δ items and no item with a as first component, then $a : A$ replaces some item with maximal age.

Let maxAge be some constant designed to be an upper limit on the possible age of items in L_p . Function $\text{neighbors}(L_p)$ returns the set

$$\{ q \mid q \neq p \wedge (\exists (a : A) : (a : A) \in L_p : a = q) \}$$

Given these variable definitions and functions, we present the algorithm for neighborhood identification.

N0: *receive* $mN(a : A) \rightarrow \text{update}(L_p, a : A \setminus \{p\})$

- N1: $(\bigcap (a : A) \in L_p : \text{age}(a : A) > \text{maxAge} \rightarrow L_p := L_p \setminus (a : A))$
 N2: $\text{true} \rightarrow sL_p := (p : \text{neighbors}(L_p))$

The constant maxAge should be tuned to safely remove old or invalid neighbor data, yet to retain current neighbor information by receiving new mN messages before age expiration. This is an implementation issue beyond of the scope of this paper: our abstraction of the behavior of the communication layer is the assumption that, eventually for any node, the guard of N1 remains *false* for any $(a : A) \in L_p$ for which $a \in N_p$. By a similar argument, eventually each node p correctly has knowledge of N_p^2 and N_p^3 as well as N_p . In all subsequent sections, we use N_p^i for $i \in 1..3$ as constants in programs with the understanding that such neighborhood identification is actually obtained by the stabilizing protocol described above.

Building upon L_p , cached values of the shared variables of nodes in N_p^i , for $i \in 1..3$, can be maintained at p ; erroneous cache values not associated with any node can be discarded by the aging technique. We use the following notation in the rest of the paper: for node p and some shared variable var_q of node $q \in N_p^3$, let $\boxtimes \text{var}_q$ refer to the cached copy of var_q at p .

Problem Definition. Let \mathcal{T} denote the task of assigning TDMA slots so that each node has some assigned slot(s) for transmission, and this transmission is guaranteed to be collision-free. We seek a solution to \mathcal{T} that is distributed and self-stabilizing in the sense that, after some transient failure or reconfiguration, node states may not be consistent with the requirements of collision-free communication and collisions can occur; eventually the algorithm corrects node states to result in collision-free communication.

3 Neighborhood Unique Naming

An algorithm providing neighborhood-unique naming gives each node a name distinct from any of its N^3 -neighbors. This may seem odd considering that we already assume that nodes have unique identifiers, but when we try to use the identifiers for certain applications such as coloring, the potentially large namespace of identifiers can cause scalability problems. Therefore it can be useful to give nodes smaller names, from a constant space of names, in a way that ensures names are locally unique.

The problem of neighborhood unique naming can be considered as an N^3 -coloring algorithm and quickly suggests a solution to \mathcal{T} . Since neighborhood unique naming provides a solution to the problem of $L(1,1)$ coloring, it provides a schedule for TDMA. This solution would be especially wasteful if the space of unique identifiers is larger than $|V|$. It turns out that having unique identifiers within a neighborhood can be exploited by other algorithms to obtain a minimal N^2 -coloring, so we present a simple randomized algorithm for N^3 -naming.

Our neighborhood unique naming algorithm is roughly based on the randomized technique described in [6], and introduces some new features. Define $\Delta = \lceil \delta^t \rceil$ for some $t > 3$; the choice of t to fix constant Δ has two competing motivations discussed at the end of this section. We call Δ the *namespace*. Let shared variable Id_p have domain $0..\Delta$; variable Id_p is the *name* of node p . Another variable is used to collect the names of neighboring nodes: $Cids_p = \{ \boxtimes Id_q \mid q \in N_p^3 \setminus \{p\} \}$. Let $\text{random}(S)$ choose with uniform probability some element of set S . Node p uses the following function to compute Id_p :

$$\text{newld}(Id_p) = \begin{cases} Id_p & \text{if } Id_p \notin Cids_p \\ \text{random}(\Delta \setminus Cids_p) & \text{otherwise} \end{cases}$$

The algorithm for unique naming is the following.

$$\text{N3: } \text{true} \rightarrow Id_p := \text{newld}(Id_p)$$

Define $\text{Uniq}(p)$ to be the predicate that holds iff (i) no name mentioned in $Cids_p$ is equal to Id_p , (ii) for each $q \in N_p^3$, $q \neq p$, $Id_q \neq Id_p$, (iii) for each $q \in N_p^3$, one name in $Cids_q$ equals Id_p , (iv) for each $q \in N_p^3$, $q \neq p$, the equality $\boxtimes Id_p = Id_p$ holds at node q , and (v) no cache update message *en route* to p conveys a name that would update $Cids_p$ to have a name equal to Id_p . Predicate $\text{Uniq}(p)$ states that p 's name is known to all nodes in N_p^3 and does not conflict with any name of a node q within N_q^3 , nor is there a cached name liable to update $Cids_p$ that conflicts with p 's name. A key property of the algorithm is the following: $\text{Uniq}(p)$ is a stable property of the execution. This is because after $\text{Uniq}(p)$ holds, any node q in N_p^3 will not assign Id_q to equal p 's name, because N3 avoids names listed in the cache of distance-three neighborhood names – this stability property is not present in the randomized algorithm [6]. The property $(\forall r : r \in R : \text{Uniq}(r))$ is similarly stable for any subset R of nodes. In words, once a name becomes established as unique for all the neighborhoods it belongs to, it is stable. Therefore we can reason about a Markov model of executions by showing that the probability of a sequence of steps moving, from one stable set of ids to a larger stable set, is positive.

Lemma 1. *Starting from any state, there is a constant, positive probability that $\text{Uniq}(p)$ holds within constant time.*

Corollary 1. *The algorithm self-stabilizes with probability 1 and has constant expected local convergence time.*

Using the names assigned by N3 is a solution to $L(1,1)$ coloring, however using Δ colors is not the basis for an efficient TDMA schedule. The naming obtained by the algorithm does have a useful property. Let P be a path of t distinct nodes, that is, $P = p_1, p_2, \dots, p_t$. Define predicate $Up(P)$ to hold if $id_{p_i} < id_{p_j}$ for each $i < j$. In words, $Up(P)$ holds if the names along the path P increase.

Lemma 2. *Every path P satisfying $Up(P)$ has fewer than $\Delta + 1$ nodes.*

This lemma shows that the simple coloring algorithm gives us a property that node identifiers do not have: the path length of any increasing sequence of names is bounded by a constant. Henceforth, we suppose that node identifiers have this property, that is, we treat N_p^i as if the node identifiers are drawn from the namespace¹ of size Δ .

4 Leaders via Maximal Independent Set

Simple distance two coloring algorithms may use a number of colors that is wastefully large. Our objective is to find an algorithm that uses a reasonable number of colors and completes, with high probability, in constant time. We observe that an assignment to satisfy distance two coloring can be done in constant time given a set of neighborhood leader nodes distributed in the network. The leaders dictate coloring for nearby nodes. The coloring enabled by this method is minimal (not minimum, which is an NP-hard problem). An algorithm selecting a maximal independent set is our basis for selecting the leader nodes.

Let each node p have a boolean shared variable ℓ_p . In an initial state, the value of ℓ_p is arbitrary. A legitimate state for the algorithm satisfies $(\forall p : p \in V : \mathcal{L}_p)$, where

$$\mathcal{L}_p \equiv (\ell_p \Rightarrow (\forall q : q \in N_p : \neg \ell_q)) \wedge (\neg \ell_p \Rightarrow (\exists q : q \in N_p : \ell_q))$$

Thus the algorithm should elect one leader (identified by the ℓ -variable) for each neighborhood. As in previous sections, $\boxtimes \ell_p$ denotes the cached copy of the shared variable ℓ_p .

$$\text{R1: } (\forall q : q \in N_p : q > p) \rightarrow \ell_p := \text{true}$$

$$\text{R2: } (\exists q : q \in N_p : \boxtimes \ell_q \wedge q < p) \rightarrow \ell_p := \text{false}$$

$$\text{R3: } (\exists q : q \in N_p : q < p) \wedge (\forall q : q \in N_p \wedge (q > p \vee \neg \boxtimes \ell_q)) \rightarrow \ell_p := \text{true}$$

Although the algorithm does not use randomization, its convergence technically remains probabilistic because our underlying model of communication uses CSMA/CA based on random delay. The algorithm's progress is therefore guaranteed with probability 1 rather than by deterministic means.

Lemma 3. *With probability 1 the algorithm R1-R3 converges to a solution of maximal independent set; the convergence time is $O(1)$ if each timed variable propagation completes in $O(1)$ time.*

¹ There are two competing motivations for tuning the parameter t in $\Delta = \delta^t$. First, t should be large enough to ensure that the choice made by `newld` is unique with high probability. In the worst case, $|N_p^3| = \delta^3 + 1$, and each node's cache can contain δ^3 names, so a choosing $t \approx 6$ could be satisfactory. Generally, larger values for t decrease the expected convergence time of the neighborhood unique naming algorithm. On the other hand, smaller values of t will reduce the constant Δ , which will reduce the convergence time for algorithms described in subsequent sections.

5 Leader Assigned Coloring

Our method of distance-two coloring is simple: colors are assigned by the leader nodes given by maximal independent set output. The following variables are introduced for each node p :

$color_p$ is a number representing the color for node p .

$minl_p$ is meaningful only for p such that $\neg \ell_p$ holds: it is intended to satisfy

$$minl_p = \min \{ q \mid q \in N_p \wedge \boxtimes \ell_q \}$$

In words, $minl_p$ is the smallest id of any neighbor that is a leader. Due to the uniqueness of names in N_p^3 , the value $minl_p$ stabilizes to a unique node. $spectrum_p$ is a set of pairs (c, r) where c is a color and r is an id. Pertaining only to nonleader nodes, $spectrum_p$ should contain $(color_p, minl_p)$ and $(\boxtimes color_q, \boxtimes minl_q)$ for each $q \in N_p$.

$setcol_p$ is meaningful only for p such that ℓ_p holds. It is an array of colors indexed by identifier: $setcol_p[q]$ is p 's preferred color for $q \in N_p$. We consider $color_p$ and $setcol_p[p]$ to be synonyms for the same variable. In the algorithm we use the notation $setcol_p[A] := B$ to denote the parallel assignment of a set of colors B based on a set of indices A . To make this assignment deterministic, we suppose that A can be represented by a sorted list for purposes of the assignment; B is similarly structured as a list.

dom_p for leader p is computed to be the nodes to which p can give a preferred color; these include any $q \in N_p$ such that $minl_q = p$. We say for $q \in dom_p$ that p dominates q .

f is a function used by each leader p to compute a set of unused colors to assign to the nodes in dom_p . The set of *used* colors for p is

$$\{ c \mid (\exists q, r : q \in N_p \wedge (c, r) \in spectrum_q \wedge r < p) \}$$

That is, used colors with respect to p are those colors in N_p^2 that are already assigned by leaders with smaller identifiers than p . The complement of the *used* set is the range of possible colors that p may prefer for nodes it dominates. Let f be the function to minimize the number of colors preferred for the nodes of dom_p , ensuring that the colors for dom_p are distinct, and assigning smaller color indices (as close to 0 as possible) preferentially. Function f returns a list of colors to match the deterministic list of dom_p in the assignment of R5.

R4: $\ell_p \rightarrow dom_p := \{ p \} \cup \{ q \mid q \in N_p \wedge \boxtimes minl_q = p \}$

R5: $\ell_p \rightarrow setcol_p[dom_p] := f(\{ c \mid \exists q : q \in N_p \wedge r < p \wedge (c, r) \in \boxtimes spectrum_q \})$

R6: $true \rightarrow minl_p := \min \{ q \mid q \in N_p \cup \{ p \} \wedge \boxtimes \ell_q \}$

R7: $\neg \ell_p \rightarrow color_p := \boxtimes setcol_r[p]$, where $r = minl_p$

R8: $\neg \ell_p \rightarrow spectrum_p := (color_p, minl_p) \cup \bigcup \{ (c, r) \mid (\exists q, c, r : q \in N_p : c = \boxtimes color_q \wedge r = \boxtimes minl_q) \}$

Lemma 4. *The algorithm R4-R8 converges to a distance-two coloring, with probability 1; the convergence time is $O(1)$ if each timed variable propagation completes in $O(1)$ time.*

Due to space restrictions, we omit the proof that the resulting coloring is minimal (which follows from the construction of f to be locally minimum, and the essentially sequential assignment of colors along paths of increasing names).

6 Assigning Time Slots from Colors

Given a distance-two coloring of the network nodes, the next task is to derive time slot assignments for each node for TDMA scheduling. Our starting assumption is that each node has equal priority for assigning time slots, *ie*, we are using an unweighted model in allocating bandwidth. Before presenting an algorithm, we have two motivating observations.

First, the algorithms that provide coloring are local in the sense that the actual number of colors assigned is not available in any global variable. Therefore to assign time slots consistently to all nodes apparently requires some additional computation. In the first solution of Figure 1, both leftmost and rightmost nodes have color 1, however only at the leftmost node is it clear that color 1 should be allocated one ninth of the time slots. Local information available at the rightmost node might imply that color 1 should have one third of the allocated slots.

The second observation is that each node p should have about as much bandwidth as any other node in N_p^2 . This follows from our assumption that all nodes have equal priority. Consider the N_p^2 sizes shown in Figure 2 that correspond to the colorings of 1. The rightmost node p in the first coloring has three colors in its two-neighborhood, but has a neighbor q with four colors in its two-neighborhood. It follows that q shares bandwidth with four nodes: q 's share of the bandwidth is at most $1/4$, whereas p 's share is at most $1/3$. It does not violate fairness to allow p to use $1/3$ of the slot allocation if these slots would otherwise be wasted. Our algorithm therefore allocates slots in order from most constrained (least bandwidth share) to least constrained, so that extra slots can be used where available.

To describe the algorithm that allocates media access time for node p , we introduce these shared variables and local functions.

$base_p$ stabilizes to the number of colors in N_p^2 . The value $base_p^{-1} = 1/base_p$ is used as a constraint on the share of bandwidth required by p in the TDMA slot assignment.

$itvl_p$ is a set of intervals of the form $[x, y)$ where $0 \leq x < y \leq 1$. For allocation, each unit of time is divided into intervals and $itvl_p$ is the set of intervals that node p can use to transmit messages. The expression $|[x, y)|$ denotes the time-length of an interval.

$g(b, S)$ is a function to assign intervals, where S is a set of subintervals of $[0, 1)$. Function $g(b, S)$ returns a maximal set T of subintervals of $[0, 1)$ that are disjoint and also disjoint from any element of S such that $(\sum_{a \in T} |a|) \leq b$.

To simplify the presentation, we introduce S_p as a private (nonshared) variable.

$$\begin{aligned} \text{R9: } & \text{true} \rightarrow \text{base}_p := | \{ \boxtimes \text{color}_q \mid q \in N_p^2 \} | \\ \text{R10: } & \text{true} \rightarrow S_p := \bigcup \{ \boxtimes \text{itvl}_q \mid q \in N_p^2 \wedge \\ & (\boxtimes \text{base}_q > \text{base}_p \vee \\ & (\boxtimes \text{base}_q = \text{base}_p \wedge \boxtimes \text{color}_q < \text{color}_p)) \} \\ \text{R11: } & \text{true} \rightarrow \text{itvl}_p := g(\text{base}_p^{-1}, S_p) \end{aligned}$$

Lemma 5. *With probability 1 the algorithm R9–R11 converges to an allocation of time intervals such that no two nodes within distance two have conflicting time intervals, and the interval lengths for each node p sum to $|\{ \text{color}_q \mid q \in N_p^2 \}|^{-1}$; the expected convergence time of R9–R11 is $O(1)$ starting from any state with stable and valid coloring.*

It can be verified of R9–R11 that, at a fixed point, no node $q \in N_p^2$ is assigned a time that overlaps with interval(s) assigned to p ; also, all available time is assigned (there are no leftover intervals). A remaining practical issue is the conversion from intervals to a time slot schedule: a discrete TDMA slot schedule will approximate the intervals calculated by g .

7 Assembly

Given the component algorithms of Sections 2–6, the concluding statement of our result follows.

Theorem 1. *The composition of N0–N3 and R1–R11 is a probabilistically self-stabilizing solution to \mathcal{T} with $O(1)$ expected local convergence time.*

8 Conclusion

Sensor networks differ in characteristics and in typical applications from other large scale networks such as the Internet. Sensor networks of extreme scale (hundreds of thousands to millions of nodes) have been imagined [10], motivating scalability concerns for such networks. The current generation of sensor networks emphasizes the *sensing* aspect of the nodes, so services that aggregate data and report data have been emphasized. Future generations of sensor networks will have significant actuation capabilities. In the context of large scale

sensor/actuator networks, end-to-end services can be less important than regional and local services. Therefore we emphasize local stabilization time rather than global stabilization time in this paper, as the local stabilization time is likely to be more important for scalability of TDMA than global stabilization time. Nonetheless, the question of global stabilization time is neglected in previous sections. We speculate that global stabilization time will be sublinear in the diameter of the network (which could be a different type of argument for scalability of our constructions, considering that end-to-end latency would be linear in the network diameter even after stabilization). Some justification for our speculation is the following: if the expected local time for convergence is $O(1)$ and underlying probability assumptions are derived from Bernoulli (random name selection) and Poisson (wireless CSMA/CA) distributions, then these distributions can be approximately bounded by exponential distributions with constant means. Exponential distributions define half-lives for populations of convergent processes (given asymptotically large populations), which is to say that within some constant time γ , the expected population of processes that have not converged is halved; it would follow that global convergence is $O(\lg n)$.

We close by mentioning two important open problems. Because sensor networks can be deployed in an *ad hoc* manner, new sensor nodes can be dynamically thrown into a network, and mobility is also possible, the TDMA algorithm we propose could have a serious disadvantage: introduction of just one new node could disrupt the TDMA schedules of a sizable part of a network before the system stabilizes. Even if the stabilization time is expected to be $O(1)$, it may be that better algorithms could contain the effects of small topology changes with less impact than our proposed construction. One can exploit normal notifications of topology change as suggested in [4], for example. Another interesting question is whether the assumption of globally synchronized clocks (often casually defended by citing GPS availability in literature of wireless networks) is really needed for self-stabilizing TDMA construction; we have no proof at present that global synchronization is necessary.

References

1. D. Bertsekas, R. Gallager. *Data Networks*, Prentice-Hall, 1987.
2. H. L. Bodlaender, T. Kloks, R. B. Tan, J. van Leeuwen. Approximations for λ -coloring of graphs. University of Utrecht, Department of Computer Science, Technical Report 2000-25, 2000 (25 pages).
3. D. E. Culler, J. Hill, P. Buonadonna, R. Szewczyk, A. Woo. A network-centric approach to embedded software for tiny devices. In *Proceedings of Embedded Software, First International Workshop EMSOFT 2001*, Springer LNCS 2211, pp. 114-130, 2001.
4. S Dolev and T Herman. Superstabilizing protocols for dynamic distributed systems. *Chicago Journal of Theoretical Computer Science*, 3(4), 1997.
5. S. Ghosh, M. H. Karaata. A self-stabilizing algorithm for coloring planar graphs. *Distributed Computing* 7:55-59, 1993.

6. M. Gradinariu, C. Johnen. Self-stabilizing neighborhood unique naming under unfair scheduler. In *Euro-Par'01 Parallel Processing, Proceedings*, Springer LNCS 2150, 2001, pp. 458-465.
7. M. Gradinariu, S. Tixeuil. Self-stabilizing vertex coloration of arbitrary graphs. In *4th International Conference On Principles Of Distributed Systems, OPODIS'2000*, 2000, pp. 55-70.
8. M. Haengi, X. Liu. Fundamental throughput limits in Rayleigh fading sensor networks. *in submission*, 2003.
9. T Herman, S Tixeuil A distributed TDMA slot assignment algorithm for wireless sensor networks. Technical Report, University of Iowa Department of Computer Science, 2004. CoRR Archive Number cs.DC/0405042.
10. J. Kahn, R. Katz, and K. Pister. Next century challenges: mobile networking for "smart dust". In *Proceedings of the Fifth Annual International Conference on Mobile Computing and Networking (MOBICOM '99)*, 1999.
11. S. O. Krumke, M. V. Marathe, S. S. Ravi. Models and approximation algorithms for channel assignment in radio networks. *Wireless Networks* 7(6 2001):575-584.
12. F. Kuhn, R. Wattenhofer. Constant-time distributed dominating set approximation. In *Proceedings of the Twenty-Second ACM Symposium on Principles of Distributed Computing*, (PODC 2003), pp. 25-32, 2003.
13. S. S. Kulkarni, U. Arumugam. Collision-free communication in sensor networks. In *Proceedings of Self-Stabilizing Systems, 6th International Symposium*, Springer LNCS 2704, 2003, pp. 17-31.
14. S. S. Kulkarni and U. Arumugam. Transformations for Write-All-With-Collision Model. In *Proceedings of the 7th International Conference on Principles of Distributed Systems (OPODIS)*, Springer LNCS, 12/03. (Martinique, French West Indies, France).
15. M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing* 15:1036-1053, 1986.
16. M. Mizuno, M. Nesterenko. A transformation of self-stabilizing serial model programs for asynchronous parallel computing environments. *Information Processing Letters* 66 (6 1998):285-290.
17. S. Ramanathan. A unified framework and algorithm for channel assignment in wireless networks. *Wireless Networks* 5(2 1999):81-94.
18. S. Ramanathan, E. L. Lloyd. Scheduling algorithms for multi-hop radio networks. *IEEE/ACM Transactions on Networking*, 1(2 1993):166-177.
19. S. Shukla, D. Rosenkrantz, and S. Ravi. Developing self-stabilizing coloring algorithms via systematic randomization. In *Proceedings of the International Workshop on Parallel Processing*, pages 668-673, Bangalore, India, 1994. Tata-McGrawhill, New Delhi.
20. S. Shukla, D. Rosenkrantz, and S. Ravi. Observations on self-stabilizing graph algorithms for anonymous networks. In *Proceedings of the Second Workshop on Self-stabilizing Systems (WSS'95)*, pages 7.1-7.15, 1995.
21. S. Sur and P. K. Srimani. A self-stabilizing algorithm for coloring bipartite graphs. *Information Sciences*, 69:219-227, 1993.
22. G. Tel. *Introduction to Distributed Algorithms*, Cambridge University Press, 1994.
23. A. Woo, D. Culler. A transmission control scheme for media access in sensor networks. In *Proceedings of the Seventh International Conference on Mobile Computing and Networking (Mobicom 2001)*, pp. 221-235, 2001.
24. F. Zhao, L. Guibas (Editors). *Proceedings of Information Processing in Sensor Networks, Second International Workshop, IPSN 2003*, Springer LNCS 2634. April, 2003.