

Distributed Self-Stabilizing Placement of Replicated Resources in Emerging Networks

Bong-Jun Ko, *Student Member, IEEE*, and Dan Rubenstein, *Member, IEEE*

Abstract—Emerging large scale distributed networking systems, such as P2P file sharing systems, sensor networks, and ad hoc wireless networks, require replication of content, functionality, or configuration to enact or optimize communication tasks. The placement of these replicated resources can significantly impact performance. We present a novel self-stabilizing, fully distributed, asynchronous, scalable protocol that can be used to place replicated resources such that each node is “close” to some copy of any object. We describe our protocol in the context of a graph with colored nodes, where a node’s color indicates the replica/task that it is assigned. Our combination of theoretical results and simulation prove stabilization of the protocol, and evaluate its performance in the context of convergence time, message transmissions, and color distance. Our results show that the protocol generates colorings that are close to the optimal under a set of metrics, making such a protocol ideal for emerging networking systems.

Index Terms—Replica placement, convergence, graph coloring.

I. INTRODUCTION

EMERGING large scale distributed networking systems, such as P2P file sharing systems, wireless ad hoc and sensor networks, and utility computing systems utilize replication of resources, information, and jobs at participating network nodes to increase the overall performance of the networked system. For example, P2P file sharing applications [6] replicate content at nodes throughout the network, reducing the expected retrieval time. In utility computing applications such as Grid Computing [5], a computing task replicated among multiple computing nodes not only further extends the distribution of the computation load, but also reduces the time spent transporting tasks to nodes. In multichannel wireless networks [2], [9], the capacity of the network is often increased by assigning channels to nodes such that the distance between nodes assigned to the same channel is maximized. In all these systems, the decision of where to replicate particular items, whether the items take the form of information, tasks, or resources, has a tremendous impact on the performance of the entire system.

While the use of replication in the different environments described above addresses a vastly different set of performance

issues, we make the observation that at an abstract level, the placement objectives are remarkably similar. In particular, it is important either to place different items in the vicinity of each node or to place the identical items as far away from each other as possible. We note that a placement that achieves one of these objectives also does it well for the other in many cases, i.e., if the distance between two identical resources is increased, it leaves room for different resources to be placed at nearby points, and *vice versa*.

There are several properties of emerging networks that complicate the design of a protocol to place or configure replicated resources.

- The number of participating nodes will often be in the thousands or millions.
- Participation will be dynamic, with nodes joining and leaving the network at unpredictable times, making it difficult to elect “leader” nodes that make decisions for large portions of the topology.
- Clocks will not be perfectly coordinated, making it difficult or impossible to perform synchronized actions among large collections of nodes. Unless carefully constructed, in such an environment, a protocol can enter a livelocked or deadlocked state.
- Transmissions across large distances will incur high bandwidth costs or will increase the time needed to decide where to place the replicas.

For these reasons, these networks cannot utilize protocols that implement centralized (leader-oriented) computation, require synchronized messaging, or require communications between distant nodes.

This paper describes a distributed, hierarchy-free, self-stabilizing protocol that performs placement of replicated resources. We address the fundamental problem of placing multiple copies of resources within the network, in which each node must hold some resource, such that an arbitrarily chosen resource is reachable over a short path from any starting point in the network, and that there are large distances between identical copies. We develop a protocol that exhibits several desirable properties:

- self-configuring: no third party is used to coordinate placement.
- fully distributed: all nodes are assigned an identical set of tasks to aid in the decision of the placement.
- scalable: information exchange is performed within localized areas.
- asynchronous: no clocks are required to sequence or coordinate the operations at nodes.

We explore this problem in the context of an arbitrary graph, where each node is colored with one of k colors, where each

Manuscript received February 17, 2004; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor E. Zegura.

The authors are with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: kobj@ee.columbia.edu; danr@ee.columbia.edu).

This work was supported in part by the National Science Foundation under Grant no. ANI-0411047, ITR ANI-0325495, and CAREER Award no. ANI-0133829, and by gifts from Lucent Technologies, Cisco Systems, and Intel Information Technology Research Council. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Digital Object Identifier 10.1109/TNET.2005.850196

color represents a resource class. The protocol operates by having each node continually change its color in a greedy manner to maximize its own distance to a node of the same color. To this, we add a novel mechanism that avoids race conditions to ensure that the protocol stabilizes in asynchronous communication environments indicative of these emerging networks.

We use both theoretical analysis and simulation to evaluate the performance of our protocol. Our theoretical analysis proves that the protocol always converges to a fixed coloring. This analysis also formally bounds the distance that must be traveled from a node to an arbitrarily chosen color within a factor of three of what can be achieved in an optimal coloring. The results of discrete-event simulation demonstrate that for randomly generated graphs, our protocol is scalable in that the convergence time and the number of protocol messages grow slowly with the number of nodes in the network, and that the coloring is on average much closer to the optimal allocation than the theoretical upper bound.

The rest of this paper is organized as follows. In Section II, we present the network model and formally state the coloring problem. Section III describes the distributed protocol, and Section IV discusses the issues involved in its extension and some potential applications. The convergence to a fixed coloring and theoretical bounds on performance are presented in Section V. Section VI presents performance evaluation results of the protocol observed through simulation. Section VII briefly reviews related work, and we conclude the paper in Section VIII.

II. MODEL AND PROBLEM FORMULATION

In this section, we introduce the network model and the optimization goals in the context of that model. Throughout the description of the model and the protocol, we assume the protocol messages exchanged between a pair of nodes are reliably delivered in the order in which they were transmitted, using some existing reliable transfer mechanism relevant to the particular network system (for example, hop-by-hop TCP connections).

We view a distributed networking system as an undirected, connected graph, $G = (N, E)$, where $N = \{1, 2, \dots, n = |N|\}$ is the set of (numbered) nodes of the graph and E is the set of edges. Each edge $e \in E$ is assigned an arbitrary, nonnegative weight. The **distance** between two nodes x and y , $d(x, y)$, is the sum of the weights of edges along a shortest path between x and y . Each node $n \in N$ is assigned one of k **colors**, c_1, \dots, c_k , where a node's color can change with time. Each color represents a specific set of resources being replicated. We assume that the value of k is determined in advance and remains fixed for the duration of the execution of the protocol. The **color** of node x at time t is denoted $C_t(x)$.

Node x 's **distance to color** c_i at time t , $d_t(x, c_i)$ is defined as the distance at time t between x and the closest node to x of color c_i i.e., $d_t(x, c_i) = \min_{1 \leq j \leq n} \{d(x, j) : C_t(j) = c_i\}$. Note that $d_t(x, C_t(x)) = 0$, and also that if no node is assigned to a particular color c_i , then $d_t(x, c_i) = \infty$. We define $\delta_t(x)$ to be the distance at time t between x and the closest other node of the same color, i.e., $\delta_t(x) = \min_{1 \leq j \leq n} \{d_t(x, j) : C_t(j) = C_t(x), j \neq x\}$.

In the context of this model, there are several different objective functions whose minimization or maximization would yield desirable placements of replicated resources depending on the design principle. For example, one could argue that the objective function is to:

- minimize $\max_x \max_i d_t(x, c_i)$;
- minimize $\sum_x \sum_i d_t(x, c_i)$; or
- maximize $\sum_x \delta_t(x)$.

The first two objective functions are used to place colors so that each node has a short distance through the graph (i.e., network) to reach any other color. Such placement is useful when it is desirable to have replicated contents or tasks (represented by colors) near to any arbitrarily chosen point in the network. With the last objective, we would like to place colors so that each node is far from nodes of the same color. Such placement is useful when nearby placement of identically-configured nodes causes a conflict, such as assigning two nodes to transmit upon the same wireless channels. Note that these two types of objectives fit naturally together: placing the same color far away makes room for other colors nearby.

Unfortunately, the above optimization problems are NP-hard as they can be reduced from 3-SAT problem [1]. Hence, we focus on developing a distributed protocol that computes approximations to these problems.

III. ASYNCHRONOUS DISTRIBUTED COLORING PROTOCOL

In this section, we describe our Asynchronous Distributed Coloring (ADC) protocol and the procedure that each node performs. Our focus in this section is the development of a decentralized protocol with which each node learns of its nearby nodes' colors and how it selects and changes its own color without causing race conditions.

Since up-to-date color information about a node y takes time to reach node x , it will be important to distinguish between the perception that node x has about the current coloring and the actual coloring of the graph. We use $\hat{d}_t(x, c_i)$, and $\delta_t(x)$ to denote the distances to the closest node of color c_i and $C_t(x)$ respectively as **perceived by node** x at time t based on the most recent coloring information it has received. The **real** distances (i.e., as perceived by an oracle that knows every node's actual color at time t) remain denoted by $d_t(x, c_i)$ and $\delta_t(x)$.

We say that a node x is **locally stable** at time t if $\hat{d}_t(x, c_i) \leq \hat{\delta}_t(x)$ for all $1 \leq i \leq k$, i.e., x perceives that no color is further from it than the closest node whose color matches its own. Otherwise, the node is **locally unstable**. If $d_t(x, c_i) \leq \delta_t(x)$ for all $1 \leq i \leq k$, then at time, is said to be **stable**, and is otherwise **unstable**. A graph coloring is said to be **stable** when all nodes in the graph are stable.

A. Coloring Rule

We begin by describing the process used to adjust the color of a node. Each node is responsible for selecting its own color. We assume for now that, with the color change propagation protocol described in the next subsection, each node continually receives updated information of the colors of nearby nodes. More specifically, at any time t , a node x keeps $\hat{d}_t(x, c_i)$ for all colors




Fig. 1. x changes color from c_1 to c_2 . (a) Before: time t . (b) After: time t' .

c_i , $1 \leq i \leq k$ and $\hat{\delta}_t(x)$, and applies the following simple color changing rule.

The Color Change Rule: At any time t , a locally stable node keeps its color fixed. However, any locally unstable node x will try to change its color to a color c_i that is furthest from x , i.e., $\hat{d}_t(x, c_i) \geq \hat{d}_t(x, c_j)$ for all $1 \leq j \leq k$.¹

In other words, a node seeks to change its color to c_i if it perceives that the closest node of color c_i is further from it than the closest other node of its current color. Changing color in this manner is desirable in that it increases $\hat{\delta}_t(x)$, decreases $\sum_i \hat{d}_t(x, c_i)$ (i.e., average distance to a color), and does not increase but often decreases $\max_i \hat{d}_t(x, c_i)$ (i.e., distance to the furthest color).

Fig. 1 provides an illustrative example of a node x that starts as color c_1 at time t and completes its change to color c_2 at time $t' > t$. Nodes x and x' are initially colored c_1 and nodes z and z' are initially colored c_2 . In this example we use Euclidean distance to indicate the distance between nodes. In Fig. 1(a), x' is the closest node to x of color c_1 , and z is the closest node to x of color c_2 . Since $\delta_t(x) < d_t(x, c_2)$, x can increase its distance to a node of the same color by changing its color from c_1 to c_2 . The result is depicted in Fig. 1(b). We see that the shaded region around x (i.e., the minimum distance to the same color) has grown. Note that the color change by x is not necessarily favorable for other nodes in the graph. For instance, the change by x to color c_2 substantially shrinks the distance from node z to a node of its same color, c_2 . Because of this, it is not immediately obvious that repeated application of our coloring rule will lead to a stable coloring. This stability is proven in Section V.

An important characteristic of this coloring rule is that a node only needs the knowledge of the distance to the *closest* node of each color, i.e., the required information exchange is often available within a local region.

We now describe the messages exchanged by neighboring nodes to implement the ADC protocol. The messages exchanged between neighbors perform two essential functions. First, nearby nodes exchange information that describes the color configuration of nodes within a local area. Second, nodes seek approval from nearby neighbors before changing their colors to guard against simultaneous attempts to change colors which, if performed by nodes simultaneously, could livelock or deadlock the protocol before it reaches a stable configuration.

B. Color Propagation

A node needs to be continually updated about any changes in coloring at other nodes that affect its minimum distance

¹If there are more than one such furthest colors, x chooses any one of them arbitrarily.

to a color, since its own “desired” color is based on these values. We have devised a distributed method we call *Colored Bellman–Ford* or CBF for short, which exchanges the necessary information among nodes. Our description here assumes that the reader is familiar with the Bellman–Ford algorithm [4].

As in the Bellman–Ford algorithm, each node in CBF learns the distances to other nodes by exchanging a distance **update** message with each of its neighbor nodes. The difference is, while Bellman–Ford updates each node with the distances to all other nodes in the network, CBF updates information only about the distances to the closest nodes of all other colors, localizing the scope of the propagation of **update** messages.

In a preliminary design of CBF, there is a row per color and a column per neighbor, with entry (i, j) storing a pair consisting of the shortest known distance to color c_i along a path whose first hop passes through node j , and the identity of this closest node of color c_i . The distance is set to ∞ when no node of color c_i is known (or exists), and the identity is left as null. As in Bellman–Ford, a node chooses its shortest (known) distance to color c_i by selecting the column in row i that minimizes this distance. When a node running CBF receives an **update** from a neighboring node, as in Bellman–Ford, it updates its table in the respective column and forward **update** messages to all of its neighbors containing shortest distance information (length and node ID) for any colors to which its shortest distance has changed.

The one limitation of the above method is that a node x will often not learn about the closest other node of the same color, $C_t(x)$, when its neighbors all point to x as their closest node of that color. This problem is resolved by having each node store information pertaining to the *two* closest nodes of each color that can be reached through each neighbor. The information about these two closest nodes is also forwarded to neighbors when an update changes either member of this closest pair, and hence a node’s table entry for the same color must contain at least one node other than itself.

Clearly, CBF will continue to send updates as long as nodes continue to change their colors, since color changes affect the distances stored in the distance tables. We have proven the following result which will be used later to prove that our distributed protocol stabilizes to a coloring.

Claim 1: When nodes cease to change their colors, i.e., $\exists t, \forall x, \forall \tau > t, C_\tau(x) = C_t(x)$, CBF will terminate transmitting messages, and each node x will have accurate information about the distance to and the identity of the two closest nodes of each color $c_i \neq C(x)$, and of the closest node other than itself of color $C(x)$.

Proof: The proof is for a graph where edges have positive weights. We will prove the result for only the closest node of each color. The proof is trivially extended to cover the two closest nodes. We first show that “stale” information is removed such that CBF can only terminate when the nodes in row i of each node’s table are actually colored c_i . We then show that CBF only terminates when the node in this row is no further from x than any other node of color c_i .

To prove the first result, suppose x considers z to be its closest node of color c_i , yet z is no longer colored c_i . Let $n_\ell, n_{\ell-1}, \dots, n_0$ be any path along which the information

about z being color c_i had propagated to x . Clearly, $n_0 = x$ and $n_\ell = z$, where n_0 has the “wrong” information about z ’s color and n_ℓ has the “right” information, i.e., it knows that z is not color c_i . Thus, there is some minimum j where n_j has the “wrong” information and n_{j+1} has the “right” information. We can consider all paths over which the information about z propagated to x and choose the closest node to z that stores z as color c_i . Call this node y . y would then remove z from c_i unless it used information from a node that was further or equidistant from z than y . Hence, if y still “believes” that z is still of color c_i , then its observed distance to z grows. If it is known that nodes are not moving, y could simply ignore any reports that are not along the shortest path. Since y is the closest node to z that still “believes” z to be of color c_i , no node along the shortest path could be reporting z to be of color c_i , and y would have to choose some other node to be the closest node of color c_i .

If y did not use the observation that its distance to z increased, then we have a situation where y is using stale information from some other node that believes z to be the closest node of color c_i . Hence, there is a cyclic dependency among a set of nodes, each believing the report of another that is supposedly closer to z . This is the same phenomenon that leads to the count-to-infinity phenomenon in the Bellman–Ford algorithm. Over time, y ’s distance to z would increase to infinity. Eventually, the distance to another node of color c_i would appear closer, and y would choose this node to be its closest node of color c_i .

To prove the second result, let n_0, \dots, n_ℓ represent the shortest path from y to x where y is a closest node of color c_i , and x ’s i th row contains another node z where $d(x, y) < d(x, z)$. Again, n_ℓ has the “wrong” information and n_0 has the “right” information. Thus, there is again some minimum j for which n_{j+1} has the wrong information and n_j has the right information. Clearly, under CBF, when n_j learned of this right information, it would update its tables and inform its neighbors of the change, including n_{j+1} . It follows via induction on j that eventually n_ℓ learns the right information. ■

CBF, coupled with the color-changing rule, allows nodes to change their colors and deliver their new color information to those nodes whose own color decisions are affected by that information. However, when nodes are actively changing their colors, the lag in time it takes for a node to learn about another node’s color change can lead to anomalous situations in which the procedure does not stabilize. A simple illustration of this phenomenon involves a graph to be colored with two colors, where all nodes are initially colored c_1 , and propagation delays between all pairs of nodes takes one unit of time. Since each node’s distance to c_2 is infinite, all nodes simultaneously change to c_2 and propagate their change information around the graph. Upon learning in the next time unit that their nearest neighbors are now colored c_2 , all nodes return to color c_1 , and the process repeats indefinitely.

Our goal here is to prevent this anomalous behavior, but we would like to permit as many locally unstable nodes to change their colors simultaneously when these simultaneous changes cause no such conflict.

C. Asynchronous Color Change Procedure

Coping with the asynchrony and handling possible anomalous behaviors is done by a 3-way handshake exchange using four types of protocol messages: **request**, **accept**, **reject**, and **decision**. These messages are used to ensure that the graph will converge to a stable coloring in which all nodes are stable. Before describing how these messages are used, we begin by looking at a monotonicity property, one of the critical properties that our protocol will satisfy to prove convergence.

1) *Monotonicity Property*: In Section V-A, we will show that the protocol produces a stable coloring as long as $\delta_t(x)$ increases whenever node x changes its color. Suppose a node x changes its color from c_1 to c_2 at time t . For $\delta_t(x)$ to increase, it is sufficient to ensure that $d_t(x, c_2) > \delta_t(x)$, which is satisfied when the following two conditions hold:

- There exists a node (other than x) at distance $\hat{\delta}_t(x)$ from x that is color c_1 at time t , ensuring that $\delta_t(x) \leq \hat{\delta}_t(x)$.
- There is no node of color c_2 within distance $\hat{\delta}_t(x)$ from x at time t , ensuring that $d_t(x, c_2) > \hat{\delta}_t(x)$.

Clearly, if these two conditions are satisfied, then so is $\delta_t(x) < d_t(x, c_2)$. Hence, by changing to color c_2 at time t , we have $\delta_{t'}(x) > \delta_t(x)$, where $\delta_{t'}(x)$ denotes the distance to the closest other node of color c_2 to x immediately after x changes its color to c_2 . In our distributed protocol, each node x changes its color only when the above two conditions are ensured. Note that x ’s decision uses only perceived distance information since this is the only information that is available to the node.

Now let us present our distributed protocol, beginning by describing at a high level how the messages are passed between nodes.

2) *Protocol Messages*: When x of color c_1 “perceives” (by CBF) that it can change its color to c_2 by the coloring rule, i.e., when $\hat{\delta}_t(x) < \hat{d}_t(x, c_2)$, it broadcasts a **request** message to all nodes whose distance from x are smaller than or equal to $\hat{\delta}_t(x)$ of x . We refer to this set of nodes as x ’s *time t disk*.² The request message states x ’s desire to change from color c_1 to color c_2 , and requests that one of the nodes in x ’s time t disk maintain color c_1 and none of these nodes change to color c_2 .

A node y that receives a **request** message from x responds either with an **accept** or **reject** message as follows. Node y immediately sends x a **reject** message if y satisfies either of the following properties:

- y is of color c_2 .
- y is the node that x “thinks” is of color c_1 , but y is in fact no longer of color c_1 .

Note that the former of the above two conditions, under which y rejects x ’s request, means that x ’s perception of having no nodes of color c_2 in its disk was incorrect. The latter condition means that the node that x depended on as the node of color c_1 is no longer of that color, hence c_1 may be further away from x than x had anticipated.

If neither of these properties hold, y sends x an **accept** message. If y accepts x ’s request to change from c_1 to c_2 , then y

²In Section III-C4, we will discuss how to scalably implement scoped broadcasting of messages within a node’s time t disk and the corresponding feedback mechanism.

is prohibited from changing from c_1 or changing to c_2 until it receives x 's **decision** message.

Now, if any node rejects x 's request, x aborts the color change. If all nodes accept x 's request, then x changes its color. Whether x changes its color or aborts, it sends a **decision** message to the same set of nodes to which it had sent a request, removing the color-change restriction that was placed on these other nodes.

3) *Handling Deadlock/Livelock*: Network propagation delays are problematic when trying to coordinate message requests within the network. In particular, we must address the possibility that two nodes, x , and y , each wishing to change color, request one another to hold their colors fixed until further notice. Nodes cannot simply accept the other's change requests when such conflict arises, as this could lead to a livelocking phenomenon, causing both to keep changing their colors. Both cannot simply reject the other—this leads to a livelock, preventing one another from changing to a “better” color. If both hold each other's request until they receive **accept** or **reject** to their own requests, a deadlock situation will arise. Also, it will often be the case that the distances over which x 's and y 's messages travel will differ. In particular, it could be the case that $\delta_t(x) < d(x, y) < \delta_t(y)$. Hence, even applying a simple tie-breaking rule between x and y is difficult, because one of the nodes may not even be aware of the conflict.

We address this dilemma by introducing a priority-based holding mechanism, in which conflicting requests from lower-priority nodes are held (without being responded) by higher-priority nodes until the latter complete their decision of whether or not to change.

More specifically, we have each node transition between two states: **STALLED** and **MOVING**. A **STALLED** node is precluded from transmitting requests to change its color. Locally stable nodes are by default always **STALLED**. A node is in **MOVING** state when it has issued its **request** to change color and has not yet issued its followup **decision**. Each node also maintains three sets that hold requests that it receives from other nodes: **stalling**, **master**, and **slave**. These messages are removed from a set only after the outcome of the request is known (i.e., the request was rejected, or the decision corresponding to an accepted request was received). Furthermore, we assign arbitrary, unique identifiers that can be used to impose a well-defined ordering among nodes. We write $i > j$ to indicate that node i is ordered before node j . Fig. 2 presents the state diagram relevant to these procedures, which are described as follows.

Initially, a node is in the **STALLED** state. Upon sending a request, a node transitions to the **MOVING** state, and transitions back to the **STALLED** state immediately after sending a decision message corresponding to the request that caused it to enter the **MOVING** state.

When in the **STALLED** state, a node x places in its stalling set any request that it accepts. x cannot transition to the **MOVING** state (i.e., cannot issue a change request) if its stalling set has an entry that conflicts with x 's desired color change. In particular, x cannot change from color c_1 to color c_2 if there is a request in its stalling set that asks to change to color c_2 , or that depends on x to be color c_1 .

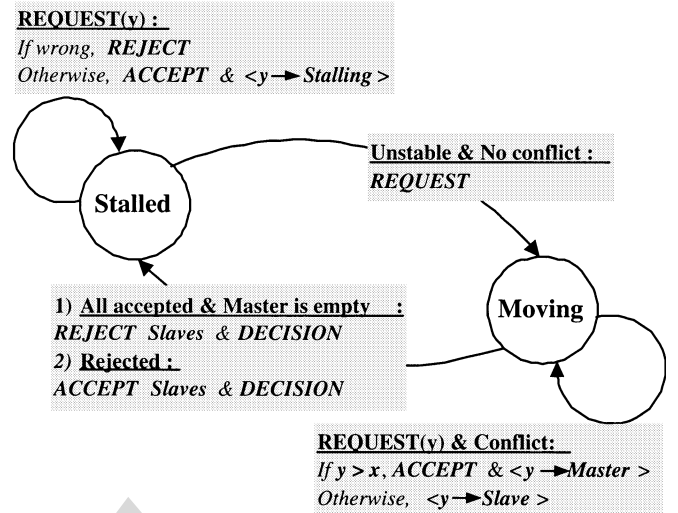


Fig. 2. State diagram for handling protocol messages.

When in the **MOVING** state (x has issued a request), x places a node y 's request in its stalling set if the request does not conflict with its own intended color change. However, if y 's request does conflict, it differentiates the request based on the order of x and y as follows. If $x > y$, then x places y 's request in its slave set without accepting or rejecting the request. On the other hand, if $x < y$, then x places y 's request in its master set and immediately accepts the request. x will not change its color while its master set contains accepted requests. A message in x 's master set is deleted when x receives the decision regarding the message.

If x 's request is accepted by all nodes and its master set is empty, it rejects (and removes) all messages in the slave set, sends a decision message, changes its color, and transitions to the **STALLED** state. If x 's request is rejected, it accepts and removes all messages in its slave set, transfers all messages in the master set into the stalling set, and transitions to the **STALLED** state without changing its color.

This prioritized **STALLED/MOVING** procedure prevents race conditions by having requests from lower-priority nodes “frozen” until nodes of higher order complete their decision of whether or not to change. The protocol does not deadlock because the node with the highest order that is in the **MOVING** state will never be “frozen” by another node. A formal proof of this claim is presented in Section V-A.

4) *Implementing Scoped Broadcast and Feedback*: We have assumed that x has the ability to broadcast messages to nodes in its time t disk and scalably receive feedback from all of these nodes. There are several ways to scalably implement this feedback mechanism, and choosing the right one depends heavily on the underlying routing mechanism of the network. If a shortest path multicast routing protocol is enabled in the network, a node can broadcast a message by broadcasting it down this shortest path tree, where at each hop, nodes on the path track the distance the message has traveled from the source thus far. If a node receives a message that has traveled farther than the radius of the disk, that node need not forward the message any further. If no such protocol is available, a simple solution is to have each node

contact each of its neighbors. Again, by tracking the distance traveled thus far, all copies of the message eventually die out, but not until all nodes within the disk have been reached (often several times). Less costly solutions can be implemented, but due to lack of space, we demonstrate mere feasibility.

In both broadcast frameworks discussed above, a feedback mechanism can be implemented to return accept or reject messages back to the original requester by sending the response to each request received through a neighboring node back to that neighboring node. When a node finds that it is not in conflict with a request, it should not send back an accept immediately, but instead wait first for accepts from all neighbors to which it forwarded a request. Once all those neighbors have responded with accepts, it too can forward its accept. If any neighbor sends a reject message, that reject message can be forwarded immediately to the node from which it was directly received from. For instance, in a shortest path multicast tree, accepts propagate back up the tree from which they came.

To summarize the ADC protocol, nodes use the distributed CBF method to exchange updates with neighbor nodes of nearest colorings. When a node wishes to change its color, it broadcasts and receives messages about this color change that respectively travel no further than the distance to the closest node perceived to be of the same color. Since these procedures occur within localized regions and have a small set of restrictions that prevent two nodes from simultaneously changing color, many of the color changes can occur simultaneously. In large-scale networks, this parallelism significantly reduces the time to reach a stable coloring in comparison to a mechanism that permits only a single color change at a time.

IV. DISCUSSION

In this section, we present some variations and extensions of the graph coloring problem that may be of use in practice. For each extension we can show that the ADC protocol will still converge. Next, we discuss practical applications, for which our protocol can offer useful solutions either directly or via one of those extensions.

A. Extensions and Variations

1) *Hardwired Nodes:* We can relax the assumption that all nodes participate in the ADC protocol and allow some nodes to fix their colors during the execution of the protocol. This is useful, for instance, in networks where there is an original source of information from which replicas are copied, and the point of origin of the information retains the original copy.

2) *Assigning Multiple Colors to a Node:* Each node can be assigned an arbitrary number of colors that it can store. This can be done by mapping the graph to another graph where a node that can store ℓ colors in the original graph is mapped to an ℓ -clique of nodes that can each store one color, where the edges between nodes in the ℓ -clique have length zero.³ Our coloring protocol can be immediately used without modification if each

node with ℓ colors in the original graph performs the procedure for ℓ independent nodes in the converted graph. The proof and all supporting results follow trivially.

This can be of use in practice, when some node has larger capacity than the others, for example, nodes with big storage, a wireless node supporting multiple channels, etc.

3) *Variable Density Coloring:* Each color c_i can be assigned a density, β_i , such that a node y of color c_1 is considered to be “closer” to x than a node z of color c_2 whenever $\beta_1 d_t(y, x) < \beta_2 d_t(z, x)$. Nodes of colors with larger density will more densely populate the graph. For instance, consider coloring a given graph with two colors c_1 of density ϵ and c_2 of density $1/\epsilon$ as $\epsilon \rightarrow 0$. Having a single node x of color c_1 in the graph is sufficient, since no matter how physically far this node is from another node y in the graph, $\beta_1 d(x, y)$ is infinitesimally small.

B. Applications

1) *Wireless Channel Allocation:* In multichannel wireless networks, transmissions can still conflict under conventional channel allocation mechanisms that assign different channels to nodes that share a common communication neighbor. This is because the interference range of a node is typically greater than its transmission range. Moreover, nodes’ mobility in such networks gives rise to a need to reconfigure the channel assignment as the topology changes.

Our graph coloring protocol creates long distances between nodes utilizing the same channel (represented by color) given a fixed number of channels. This reduces both the interference level and the frequency of channel reconfiguration. Even when the allocation by our coloring needs to be reconfigured, this can be done quickly as the simulation results in Section VI indicate. However, a more thorough evaluation in highly dynamic environments is yet to be done.

2) *Distributed Leader Election:* Leader election algorithms are useful building blocks in network protocols that impose a hierarchical structure, e.g., in hierarchical routing in large scale ad hoc networks. One challenge in those protocols is to distribute the burden of being leader. Using the method for assigning multiple colors to a node, our coloring protocol can be applied in this environment such that each node can find a node of every other color in some of its neighbors. Once the protocol stabilizes, the leader re-election problem simplifies to having each node take its turn for being the leader.

3) *Distributed Resource Directory Service:* In large scale distributed networks that require decentralized directory services, the cost of using the directory is reduced if the directory information is accessible at nearby nodes. Our protocol can be used to map each index (or set of indices) of the resource (or content) to a specific color using a predetermined method such as a hash function, having the index to be stored at a nearby node. An additional benefit of our protocol is that, while determining the location of resources, the direction that a search should take to reach the closest copy of the index is computed at each hop (by CBF). This makes it trivial to design a search that heads directly toward the closest copy of the index.

³To ensure the convergence of CBF, a very small positive value ϵ can be assigned to the edges in ℓ -clique.

V. ANALYSIS

In this section, we prove the convergence of our ADC protocol, and provide formal results about its performance in comparison to “optimal” colorings.

A. Convergence

We begin by proving that the ADC protocol always converges to a stable coloring, i.e., nodes all reach a point where they cease changing their color and therefore cease sending messages.

Claim 2: If node x is in the STALLED state and node y has an accurate view of node colors (in particular, $\hat{\delta}_t(y) = \delta_t(y)$ and $\hat{d}_t(y, c_i) = d_t(y, c_i)$ for all $1 \leq i \leq k$), then x will not prevent y from changing color (i.e., x does not “freeze” y ’s messages).

Proof: Since x is STALLED, its slave set is empty. Hence, it does not “freeze” messages. Since y has an accurate view of the network coloring and has issued a request, no node will reject the request. x will therefore immediately accept the request. ■

Claim 3: If the graph is unstable, then there exists a node that is or will be in the MOVING state.

Proof: Our proof is by contradiction. Assume the graph is unstable, yet all nodes remain in the STALLED state indefinitely. Since nodes cannot issue requests to change color in the STALLED state, the coloring ceases to change. Since CBF is used to perform color updates, we have from Claim 1 that all nodes will eventually obtain the correct coloring of the graph. Thus, any node y that is unstable will eventually become locally unstable. In such a case, the only event that keeps y in the STALLED state is a node, x , whose request conflicts with the request that y would like to issue. If x has a pending request, it must be in the MOVING state, contradicting our assumption. On the other hand, if x is in the STALLED state, it follows from Claim 2 that x will not prevent y from changing color, hence no node stops y from going to the MOVING state, again leading to a contradiction. ■

Lemma 1 (Liveness): If the graph is unstable, then some node will eventually successfully change its color.

Proof: Our proof is by contradiction. Assume the graph is unstable but no node ever changes color. By Claim 1, all nodes eventually have correct view of graph coloring. Since the graph is unstable, by Claim 3, there exists a node that is or will be in the MOVING state. Let v be the node of highest order among those in the MOVING state. Since v ’s order is higher than any other nodes in the MOVING state, v is not in the slave set of any other node (only nodes in the MOVING state may have a nonempty slave set, and all such nodes have lower order than v). Furthermore, since v has the correct view of graph coloring, v ’s request cannot be rejected. Hence nothing can stop v from changing its color, which contradicts the assumption that no node ever changes color. ■

Our final lemma, before proving our first main result, is similar in flavor to a lemma we presented in [12] that proves the convergence of the coloring rule on a graph whose edges have identical length. The proof in this lemma requires that we consider a vector $V(t) = \langle \delta_t(n_1), \delta_t(n_2), \dots, \delta_t(n_{|N|}) \rangle$, where $\{n_1, n_2, \dots, n_{|N|}\}$ are all the nodes of the graph, ordered such that $\delta_t(n_i) \leq \delta_t(n_{i+1})$ for all $1 \leq i < |N|$, i.e.,

$V(t)$ is an ordered $|N|$ -component vector whose components are the $\delta_t(x)$ for all nodes x . We say vector $V(t) < V(t')$ when $V(t)$ is *lexicographically smaller than* $V(t')$, i.e., if $V(t) = \langle v_1, v_2, \dots, v_{|N|} \rangle$ and $V(t') = \langle w_1, w_2, \dots, w_{|N|} \rangle$ then there is some $0 < i \leq |N|$ where $v_j = w_j$ for all $j < i$ and $v_i < w_i$.

Claim 4: If x is the only node that changes color between time t and t' , then $\delta_{t'}(x) > \delta_t(x)$.

Proof: Let x change from color c_1 to color c_2 at time τ , $t < \tau < t'$.⁴ By the 3-way handshaking protocol, x can change its color only when its request to change from c_1 to c_2 has been accepted by all nodes within distance $\hat{\delta}_t(x)$, which means that at time t , there is at least one node of color c_1 at distance $\hat{\delta}_t(x)$ and no node of color c_2 in x ’s time t disk. Hence x changes color after ensuring that $\hat{\delta}_t(x) = \delta_t(x)$ and $d_t(x, c_2) > \hat{\delta}_t(x)$. Since no other node in x ’s time t disk changes color between times t and t' , $\delta_{t'}(x) > \hat{\delta}_t(x)$, and hence $\delta_{t'}(x) > \delta_t(x)$. ■

Lemma 2 (Monotonicity): When one node changes color between times t and t' , then $V(t) < V(t')$.

Proof: Let x be the one node that changes from color c_1 to color c_2 . We have $\delta_{t'}(x) > \delta_t(x)$ from Claim 4. Consider any node y where $\delta_{t'}(y) < \delta_t(y)$. For this decrease to occur, some node that at time t was not colored $C_t(y) = C_{t'}(y)$ must have changed to $C_{t'}(y)$ by time t' . Since only x changed its color during this time interval, the node must be x and $C_{t'}(y) = c_2$. Since x is no closer than $\delta_{t'}(x)$ to a node of the same color at time t' and x is now the closest node of color c_2 to y , we have that $\delta_{t'}(y) = d(y, x) \geq \delta_{t'}(x)$. Thus, $\delta_t(x) < \delta_{t'}(x) \leq \delta_{t'}(y)$, i.e., some component of $V(t)$ increased to yield $V(t')$ (in this case, the component contributed by node x), and any decreasing components remain larger than the final value of the increased component (for instance, y ’s component). This is clearly a lexicographic increase, hence $V(t) < V(t')$. ■

Theorem 1 (Convergence): ADC Protocol converges to a stable coloring.

Proof: If the graph is unstable, by Lemma 1, there will always be at least one node which changes its color, and whenever a node changes its color, by Lemma 2, there is a lexicographic increase in $V(t)$. Since each component is either bounded by the maximum diameter of the graph or else equals ∞ , we must reach a time T where $V(t)$ is fixed for all $T > t$. This means that nodes cease changing their colors, and therefore the coloring is stable. ■

B. Formal Bounds

Now we prove some theoretical bounds on color distances generated by our distributed protocol. For the purpose of evaluating the performance of a stable coloring, we define the super-optimal distance of a node as follows.

Definition 1: The **super-optimal distance** of a node x in a graph that is colored with k colors, $d_{opt}(x)$, is the distance to the k th closest node to x including x .

The super-optimal distance of a node is the minimum distance for which there exists a coloring in which the node can reach all

⁴If two nodes change at exactly the same time, clearly the result is the same as in a system where the difference in time is positive but bounded below any ϵ . We can then choose t and t' such that their difference is less than ϵ .




Fig. 3. Example of stable coloring: $\max d(x, c_i) = 3d_{\text{opt}}(x)$.

k colors (including its own color). We coin the term as “super-optimal” because for a large number of graphs, there is no single coloring in which each node can reach all colors within its super-optimal distance.

Theorem 2: If the graph is stable, then for each node x , $d_t(x, c_i) \leq 3d_{\text{opt}}(x)$ for all $i = 1, \dots, k$. This is a tight bound.

Proof: Consider an arbitrary coloring of a graph in which there is a node x and a color c_i where $d(x, c_i) > 3d_{\text{opt}}(x)$. Let y be the closest node to x of color c_i . Since there are at least k nodes within distance $d_{\text{opt}}(x)$ from x , and none of which is color c_i , there are two nodes w and z within this distance that are both of the same color, c_j . Since both of these nodes are within distance $d_{\text{opt}}(x)$ from x , they are at most distance $2d_{\text{opt}}(x)$ from one another, and more than distance $2d_{\text{opt}}(x)$ from any node of color c_i . It follows that neither w nor z is stable, hence the graph is not stable. ■

We note that this bound is tight, i.e., that there can exist a stable coloring in which $d(x, c_i) = 3d_{\text{opt}}(x)$ for some color c_i . Fig. 3 demonstrates this, where all edges in the figure have weight 1 and each node is labeled by one of four colors. The node colored Y has super-optimal distance of 1, the coloring of the graph is stable, and a distance of 3 is needed to reach color W .

Our next result extends the k -hop bound Lemma for unit-length edge graphs in [12] to graphs whose edge lengths are arbitrary.

Lemma 3: Let $D_k(x)$ be the minimum distance from node x to a node that is k hops from x . Then for any color $c_i \neq C_t(x)$, $d_t(x, c_i) < D_k(x)$ in a stable coloring.

Proof: The proof is by contradiction. Suppose, for some node x and for some color $c_i \neq C_t(x)$, $d_t(x, c_i) \geq D_k(x)$ in a stable coloring. Since there exist at least k nodes (including x) along any shortest path from x leading to a closest node of color c_i , two nodes along the path must be of the same color. Consider any shortest path from x to a node of color c_i . Let y and z be two nodes of the same color c_j on this k -hop path, where, without loss of generality, y is the closer node to x along this path (y may in fact be x , in which case $c_j = C_t(x)$). It is easy to see from the properties of a shortest path that y is closer to z than to any node of color c_i — otherwise x would be closer to this node of color c_i than to z , violating the assumption that no node of color c_i is within k hops from x . Hence, y is unstable, leading to a contradiction. ■

Corollary 1: Let w_{max} be the weight of the longest edge in the graph and k be the number of colors. In any stable coloring, $d_t(x, c_i) < kw_{\text{max}}$ for all colors $c_i, i = 1, \dots, k$, and all nodes x .

Corollary 2: If all edges in the graph (V, E) have the same edge weights, then for all nodes x and all colors $c_i, i = 1, \dots, k$, x can reach c_i within $k - 1$ hops when the graph has converged to a stable coloring.

The above bounds are tight when applied to the general class of connected graphs, since the equality holds for the case of a chain containing k nodes and $k - 1$ unit-length edges. Last, we present a Lemma that lower-bounds the distance between nodes of the same color.

Lemma 4: $\delta_t(x) \geq d_{\text{opt}}(x)$ in a stable coloring.

Proof: Consider a coloring where $\delta_t(x) < d_{\text{opt}}(x)$. Since fewer than k nodes are within distance $\delta_t(x)$ of x , there is some color $c_i \neq C_t(x)$ for which $d_t(x, c_i) > \delta_t(x)$. Hence, the coloring is not stable. ■

VI. PERFORMANCE EVALUATION

In this section, we evaluate the performance of our protocol via discrete event-driven simulation, measuring its transient behavior (time to converge to a stable coloring, number of messages per node) as well as the quality of the colorings produced in the context of the distances to colors. Each simulation run is performed on a connected, undirected graph, which is generated randomly as follows. Given a set of nodes, edges are added to the graph between pairs of nodes chosen uniformly, and edge weights are selected uniformly at random within the interval [1], [10]. We vary the number of edges that are added and only select graphs that are connected (we add additional edges at random when the graph is not connected). We then classify graphs by the average degree of the nodes. Results presented here, unless explicitly stated otherwise, are based on graphs whose average degrees round to 5, i.e., are between 4.5 and 5.5.

We vary the number of nodes, n , that form the graph and generate 100 different random graphs for each value of n considered. For each graph, we vary k , the number of colors that are used to color nodes in the graph, and run a separate simulation for each value of k and each graph. At the start of the simulation, a node is assigned a color (from the k) uniformly at random and all nodes transmit CBF update messages. The simulation terminates when there is no outstanding message from any node, meaning the graph is stable. The message propagation delays along the edges are proportional to the weight of the edge. We choose proportional delays as a starting point for the simulations since often the edge weights are simply a measure of delay. We reiterate that the protocol will converge regardless of what temporal or spatial dependence exists between edge delays and edge weights.

A. Transient Behavior

We evaluate the transient behavior of the protocol by measuring convergence time, the number of color changes and messages sent or forwarded per node. The convergence time is measured as the time elapsed until the graph enters a stable state, starting from the instance that every node is assigned a randomly chosen color.

Fig. 4 depicts the average convergence time with 95% confidence intervals from the simulations as we vary the number of nodes in the graph and the number of colors. In Fig. 4(a), we vary the number of nodes along the x axis, with each curve fixing the number of colors used in that set of simulations. We see that as n grows large, the increase in time grows in a logarithmic manner. Fig. 4(b) replots these points but varies the

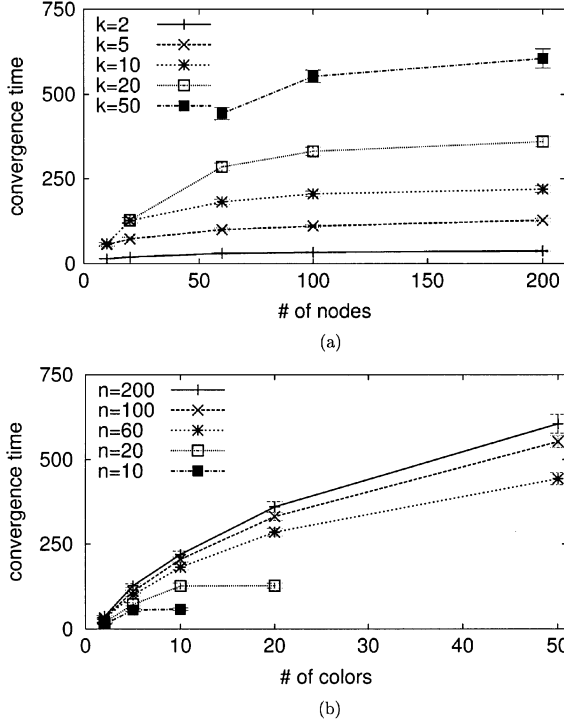


Fig. 4. Convergence time. (a) Convergence time versus graph size. (b) Convergence time versus number of colors.

number of colors along the x axis with each curve fixing the number of nodes used in that simulation. We see that the convergence time grows sublinearly with a change in the number of colors, but clearly the time is more sensitive to the number of colors than it is to the number of nodes. These results confirm the intuition that localized regions quickly fill with all possible colors, such that colorings in areas of the graph distant from a node have little impact on its coloring decision. Intuitively, increasing the number of nodes merely increases the number of these regions, having little effect on their individual convergence times, while increasing the number of colors increases the area of each region (and hence the super-optimal distance of each node), resulting in more time taken for the protocol to converge.

Fig. 5 plots the convergence time of the graph versus the average super-optimal distance of nodes in the graph over all simulation runs where $n = 100$, irrespective of the number of colors or average degree of the graph. The plot suggests that convergence time is roughly proportional to the average super-optimal distance of a graph. This supports our intuition described above, in that the super-optimal distance is a rough indicator of the size of the region that a node must observe to determine its own color.

Next, we turn our attention to the number of times a node changes its color. Fig. 6(a) and (b) plots the average number of color changes and of color change requests per node, respectively, as a function of the number of nodes in the graph. We see the average number of changes per node is between 0 and 2, while the number of the attempts to change color (including those that subsequently cancel) is roughly double the number of changes. These preliminary results suggest that the number of color changes, while somewhat sensitive to the number of

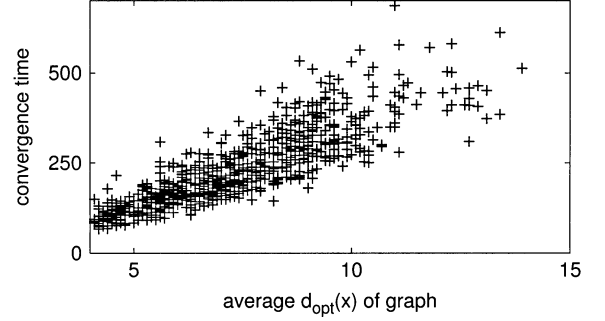


Fig. 5. Convergence time versus super-optimal distance.

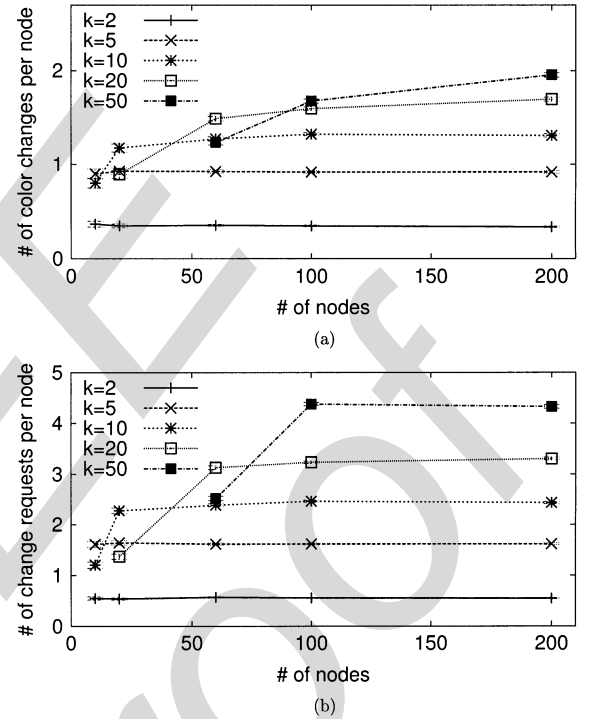


Fig. 6. Number of color changes/requests per node. (a) Number of color changes per node. (b) Number of color change requests per node.

colors, appears to grow slowly in both the number of nodes and colors.

Fig. 7 depicts the average number of messages transmitted by a node, including CBF updates as well as request, decision, accept and reject messages. In the simulation, request, decision, accept, and reject messages are forwarded within the time t disk along the shortest-path tree rooted at x . CBF update messages received at x from a neighbor y regarding the color change of node z are disregarded (and not forwarded further) when y is not on the shortest path from x to z .

Fig. 7(a) plots the average number of messages transmitted over all types of messages, and Fig. 7(b) restricts the plots to only those messages that are not CBF updates. We note that the CBF update messages contribute to roughly $2/3$ the message overhead.

We again observe a logarithmic increase in the growth as a function of n , suggesting again that the rate of increase in the number of messages transmitted by a node as a function of n becomes negligible as n grows large. While the total number

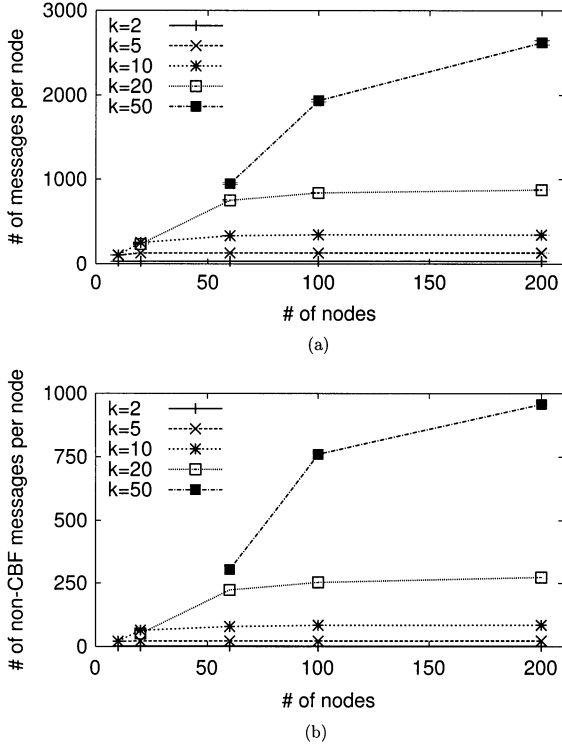


Fig. 7. Number of messages per node. (a) All messages. (b) Non-CBF messages.

of messages at first glance seems high for large values of k , we suspect the number is negligible in comparison to the number of messages that these same environments will require as part of the routing protocol, or in comparison to the amounts of data that will be transmitted. In addition, we have made no attempts to optimize the number of messages and suspect that aggregation and randomized suppression techniques can significantly lower the number of messages transmitted (for example, we could “piggyback” some update messages in decision messages).

B. Coloring Distance

We evaluate the performance of our protocol by comparing the color distances in stable graphs generated by the protocol to the super-optimal distances and to distances in graphs where each node chooses its color uniformly at random.

Fig. 8(a) depicts the distance from each node to its furthest color within the stable coloring, i.e., $\max_{1 \leq i \leq k} d(x, c_i)$ averaged over all nodes x . As expected, this distance remains constant as n is varied, i.e., proximity to the furthest color is insensitive to the size of the graph. Given the expected exponential increase in the number of nodes that can be reached as one increases the travelling distance linearly, it follows our intuition that this distance appears to be on the order of $\log k$ for large n (note that each curve roughly doubles the number of colors and are evenly spaced apart).

Fig. 8(b) plots the node’s distance to its furthest color normalized by its super-optimal distance in the graph, i.e., $(\max_{1 \leq i \leq k} d(x, c_i)) / (d_{\text{opt}}(x))$. While we have shown that this value can be as high but no higher than 3, in practice we find the value no higher than 1.2 on average. This is quite impressive when one considers the fact that this value cannot

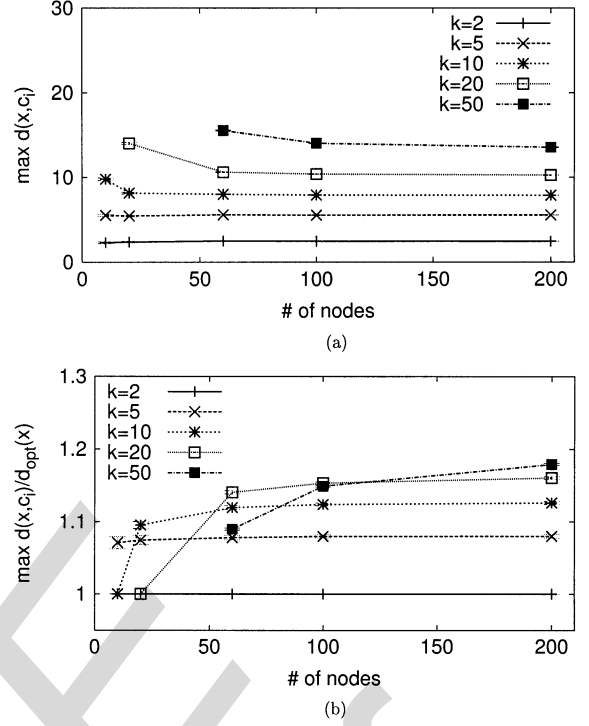


Fig. 8. Distance to furthest color. (a) $\max d(x, c_i)$. (b) $(\max d(x, c_i)) / (d_{\text{opt}}(x))$.

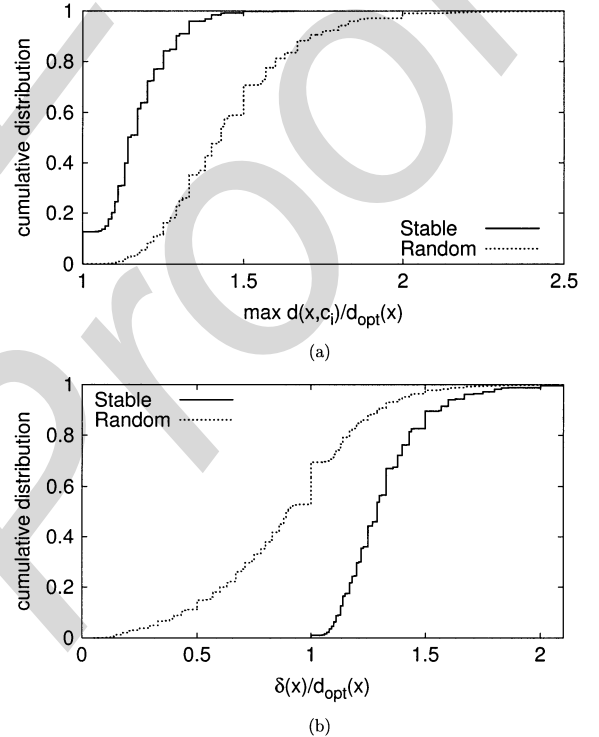


Fig. 9. Color distance comparison: Stable Coloring versus Random coloring. (a) $(\max d(x, c_i)) / (d_{\text{opt}}(x))$. (b) $(\delta(x)) / (d_{\text{opt}}(x))$.

be less than 1 and often there is no coloring for which this value is 1 simultaneously for all nodes.

Finally, we compare the coloring distance in the stable graphs to that of random colorings in which nodes select their colors uniformly at random. Fig. 9(a) plots the fraction of nodes over

all simulation runs whose distance to the closest node of the furthest color is less than x times the super-optimal distance. The curves plot results for both colorings generated by our protocol and those generated at random, where x is varied along the x axis. The results are taken from all simulations with $n = 200$ and $k = 20$. The gain obtained by using our mechanism is clear. For example, almost 80% of nodes in colorings generated by our protocol have the distances to their furthest color within a factor of 1.25 of their super-optimal distances, while only 10% of nodes are within this factor in a random coloring. Similarly, almost all nodes in a coloring generated by our protocol have the distances to their furthest color within a factor of 1.5 of their super-optimal distance in a coloring generated by our protocol while only 60% are within this factor in a random coloring.

Fig. 9(b) plots the fraction of nodes whose distance to the closest node of the same color is within a factor of x of the super-optimal distance. As proven in Lemma 4, this factor never falls below one, whereas in a random coloring, this factor is less than one for more than 50% of the nodes. This significant improvement is observed in all the other sets of simulations with different parameters.

VII. RELATED WORK

In [12], we presented a very preliminary centralized version of the greedy coloring protocol used here. The work makes multiple simplifying assumptions that preclude its use in practical networking environments. First, it was assumed that only one node could change its color at a time—a difficult property to ensure in a distributed, asynchronous medium. Second, it was assumed that each node could obtain from an oracle the current coloring of all other nodes within the rest of the graph. Last, all graph edges were assumed to be unit length. We relaxed all of these assumptions in [13], which we extend in this paper by providing additional theoretical and experimental results.

Our goal of optimizing distance to the nearby copy of resource shares many similarities with the optimal facility location problems. There have been numerous results in this area, such as polynomial-time, constant-factor approximation algorithms [3], [7], [15], and greedy methods [8] that further improve upon previous approximations. Ref. [10] is probably the most closely related to our work in the sense that it solves the problem of placing replicated objects in arbitrary nodes in the network, but it uses an off-line, centralized algorithm, and to our best knowledge, there has been no exploration of simple distributed approaches that address the problem in our context.

A large body of work has also looked at the problem of placing replicas within peer-to-peer overlays [11], [14]. However, these works differ from ours in that the focus is on identifying content that is popular and ensuring that this content is placed nearby. Our goal, in contrast, is to find a placement strategy that minimizes the distance that needs to be searched to find an arbitrarily chosen piece of content.

VIII. CONCLUSION

We have presented a decentralized, fully distributed, scalable protocol that places replicated resources in a network of arbitrary

topology such that the furthest distance one must travel to find a particular copy of a resource is only slightly larger than optimal, and the distance between identical copies is large. Simulation results suggest that running time, messages, and the number of color changes for a node scale logarithmically in the number of nodes in the graph and colors used to color the graph. These properties make the protocol ideal for assigning locations of replicated resources that are needed in emerging distributed networking environments such as ad-hoc, sensor, and overlay networks.

There are a number of interesting open problems of both a theoretical and practical nature. First, it remains an open question as to whether the convergence time of our distributed protocol is polynomial in the number of colors and nodes of the graph. Simulation results suggest this to be the case, but we have not yet been successful in demonstrating this formally. Another challenge is dealing with frequent changes in topology, either due to node movement, node failure, or joining and leaving of participants, which is a common characteristic of these emerging networking environments. Since our protocol converges quickly starting from a random coloring, we suspect it will also quickly adapt to such changes. However, shifting of replicated resources may be expensive and so the practical details of such shifts must be considered.

ACKNOWLEDGMENT

The authors would like to thank the reviewers for their helpful and constructive comments.

REFERENCES

- [1] M. Adler, private communication.
- [2] A. A. Bertossi and M. A. Bonuccelli, "Code assignment for hidden terminal interference avoidance in multihop packet radio networks," *IEEE/ACM Trans. Networking*, vol. 3, no. 4, pp. 441–449, Aug. 1995.
- [3] M. Charikar, S. Guha, E. Tardos, and D. B. Shmoys, "A constant-factor approximation algorithm for the k -median problem," *J. Comput. Syst. Sci.*, vol. 65, no. 1, pp. 129–149, Aug. 2002.
- [4] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. Cambridge, MA: MIT Press, 2001.
- [5] I. Foster and C. Kesselman, Eds., *The Grid: Blueprint for a New Computing Infrastructure*. San Mateo, CA: Morgan Kaufmann, 1999.
- [6] Gnutella [Online]. Available: <http://www.gnutella.co.uk/>
- [7] A. Goel, P. Indyk, and K. R. Varadarajan, "Reductions among high dimensional proximity problems," in *Proc. Symp. Discrete Algorithms (SODA)*, Washington, DC, 2001, pp. 769–778.
- [8] <PLEASE PROVIDE PAGE NUMBERS.> S. Guha and S. Khuller, "Greedy strikes back: Improved facility location algorithms," in *Proc. Symp. Discrete Algorithms (SODA)*, San Francisco, CA, Jan. 1998.
- [9] L. Hu, "Distributed code assignments for cdma packet radio network," in *IEEE/ACM Trans. Networking*, vol. 1, Dec. 1993, pp. 668–677.
- [10] <PLEASE PROVIDE PAGE NUMBERS.> J. Roberts, J. Kangasharju, and K. Ross, "Object replication strategies in content distribution networks," in *Proc. Web Caching and Content Distribution Workshop (WCW'01)*, Berlin, Germany, Jun. 2001.
- [11] <HOW MAY THIS REFERENCE BE ACCESSED? URL?> J. Kangasharju, K. W. Ross, and D. A. Turner, "Optimal content replication in P2P communities," 2003. manuscript.
- [12] <PLEASE PROVIDE PAGE NUMBERS.> B.-J. Ko and D. Rubenstein, "A greedy approach to replicated content placement using graph coloring," in *Proc. SPIE ITCOM Conf. Scalability and Traffic Control in IP Networks II*, Jul. 2002.
- [13] <PLEASE PROVIDE PAGE NUMBERS.> B.-J. Ko and D. Rubenstein, "Distributed, self-stabilizing placement of replicated resources in emerging networks," in *Proc. 11th IEEE Int. Conf. Network Protocols (ICNP 2003)*, Nov. 2003.

- [14] <PLEASE PROVIDE PAGE NUMBERS.> <2001 OR 2002?> L. Qiu, V. Padmanabham, and G. Voelker, "Replication strategies in unstructured peer-to-peer networks," in *Proc. ACM SIGCOMM'01*, Aug. 2002.
- [15] D. B. Shmoys, E. Tardos, and K. Aardal, "Approximation algorithms for facility location problems (extended abstract)," in *ACM Symp. Theory of Computing (STOC)*, El Paso, TX, May 1997, pp. 265–274.



Bong-Jun Ko (S'02) received the B.S. and M.S. degrees in electrical engineering from Seoul National University, Seoul, Korea, and is currently working toward the Ph.D. degree in electrical engineering at Columbia University, New York, NY.

He worked as a Research Engineer at LG Electronics in Korea before joining Columbia University, and has interned at IBM T. J. Watson Research Laboratory, Yorktown Heights, NY. His current research interests include large-scale distributed systems and mobile and wireless networks.

Mr. Ko is the recipient of the 2003 IEEE ICNP Best Paper Award.

Dan Rubenstein (M'00) received the B.S. degree in mathematics from the Massachusetts Institute of Technology, Cambridge, the M.A. degree in mathematics from the University of California at Los Angeles, and the Ph.D. degree in computer science from the University of Massachusetts, Amherst.

He has been an Assistant Professor of Electrical Engineering and Computer Science at Columbia University, New York, since 2000. His research interests are in network technologies, applications, and performance analysis, with a recent emphasis on resilient and secure networking, distributed communication algorithms, and overlay technologies.

Dr. Rubenstein has received a National Science Foundation CAREER Award, the Best Student Paper Award from the ACM SIGMETRICS 2000 Conference, and a Best Paper Award from the IEEE ICNP 2003 Conference. He has been a member of the ACM since 2000.

IEEE
PROOF