

Rendezvous of Mobile Agents in Unknown Graphs with Faulty Links

J eremie Chalopin^{1,*}, Shantanu Das², and Nicola Santoro³

¹ LaBRI Universit  Bordeaux 1, Talence, France
chalopin@labri.fr

² School of Information Technology and Engineering, University of Ottawa, Canada
shantdas@site.uottawa.ca

³ School of Computer Science, Carleton University, Canada
santoro@scs.carleton.ca

Abstract. A group of identical mobile agents moving asynchronously among the nodes of an anonymous network have to gather together in a single node of the graph. This problem known as the (asynchronous anonymous multi-agent) rendezvous problem has been studied extensively but only for networks that are safe or fault-free. In this paper, we consider the case when some of the edges in the network are dangerous or faulty such that any agent travelling along one of these edges would be destroyed. The objective is to minimize the number of agents that are destroyed and achieve rendezvous of all the surviving agents. We determine under what conditions this is possible and present algorithms for achieving rendezvous in such cases. Our algorithms are for arbitrary networks with an arbitrary number of dangerous channels; thus our model is a generalization of the case where all the dangerous channels lead to single node, called the *Black Hole*. We do not assume prior knowledge of the network topology; In fact, we show that knowledge of only a “tight” bound on the network size is sufficient for solving the problem, whenever it is solvable.

1 Introduction

1.1 The Problem

Consider a networked environment, modelled as a simple connected graph, in which operate a set of mobile computational entities, called agents or robots. A central problem in such systems is the so called *rendezvous* (or *gathering*) problem which requires the agents to meet together in a single node of the network. This problem has been extensively studied in the literature (see, for example, [1,3,10,15,17,19,23]) under a variety of models with different assumptions on the identity of the network nodes and/or of the agents (anonymous or distinct labels), the existence of timing bounds on the agents’ actions (synchronous or asynchronous), the intercommunication mechanisms (whiteboards

* Supported in part by grant ANR-06-SETI-015-03 awarded by the Agence Nationale de la Recherche.

or pebbles/tokens), the amount and type of memory, etc. In spite of their widely different models, the existing studies on the rendezvous problems share the common assumption that the environment where the agents operate is safe.

Unlike previous studies on the rendezvous problem, we consider the case when the environment where the rendezvous must take place, is not safe. In our model, some of the edges in the graph are harmful for the agents; specifically, any agent that attempts to traverse any such an edge (from either direction) simply disappears, without leaving any trace. The location of the unsafe links are initially unknown to the agents; we only assume that the unsafe links do not disconnect the network. Notice that if all the edges incident to a node u are unsafe, then node u can never be reached by any agent and is equivalent to a *black hole*, i.e., a node that destroys any incoming agent (e.g., [8,9,11,12,13,18]). In other words, the black hole model is just a specific case of the model considered in this paper.

We investigate the problem in a very weak (and thus computationally difficult) setting: the network nodes do not have distinct identities (i.e., the network is anonymous), the agents are identical, and their actions (computations and movements) take a finite but otherwise unpredictable amount of time.

The only previous result for rendezvous in faulty networks was in the case of the ring network containing two unsafe links leading to a single node—the black hole [13]. Our investigation is thus a generalization of these studies to networks of arbitrary topology that contain faults at arbitrary locations.

1.2 Our Results

In this paper we provide a full characterization of the rendezvous problem of asynchronous anonymous agents in anonymous networks with unsafe links. Assuming that the safe part of the network is connected, any port on a safe node which leads to an unsafe part of the network, is called a faulty link. We present the following results in this paper:

- We first show that, if there are τ unsafe links in the network and k agents, then it is not possible, in general, for k' agents to rendezvous if $k' > k - \tau$.
- We then prove that the rendezvous of $k - \tau$ agents is deterministically possible only when the network is *covering minimal*. Even in this case, rendezvous is not possible if the agents do not know the size of the network or at least a tight upper bound. In fact, we prove that a loose upper bound $n \leq B_n \leq 2n$ is not sufficient.
- We then show that this result is tight. In fact, we present an algorithm, *RDV* that requires only the knowledge of a tight upper bound B on the number of nodes n , such that $n \leq B < 2n$. This algorithm allows rendezvous of $k - \tau$ agents in networks where such a rendezvous is possible; the rendezvous occurs with explicit termination for each surviving agent.
- The total number of moves made by the agents during the execution of algorithm *RDV* is $O(m(m + k))$ where m is the number of edges in G . We prove that this cost is optimal; in fact, we show that solving rendezvous of $k - \tau$ agents in networks where it is solvable, requires at least $\Omega(m(m + k))$ moves, even when the network topology is known a priori.

- Finally we show that, there exists no *effective* algorithm for *maximal* rendezvous, i.e. there does not exist an algorithm that when executed on any arbitrary network achieves the rendezvous of as many agents as deterministically possible on that network.

Due to the limitations of space, the proofs of some lemmas and theorems have been omitted; These can be found in the full paper.

1.3 Related Work

The problem of *Rendezvous* has been extensively studied mostly using randomized methods (see [1] for a survey). Among deterministic solutions to rendezvous, Yu and Yung [23] and Dessmark et al. [10] presented algorithms for agents with distinct labels. In the anonymous setting, the problem has been studied under different models (synchronous or asynchronous), using either whiteboards [3] or, pebbles/tokens [19]. Some of the recent studies have focussed on minimizing the memory required by the agents for rendezvous([15,17]).

Most of these solutions are designed for anonymous graphs (i.e. graphs where nodes do not have distinct identities) which present the most challenging (i.e. computationally difficult) situations. The issue of computability in anonymous graphs, have been studied by many authors including Angluin [2], Yamashita and Kameda [22], Mazurkiewicz [20], Sakamoto [21], and Boldi and Vigna [4]. Most of these studies have concentrated on the problem of symmetry-breaking or leader election which is in fact, closely related (and sometimes equivalent [7]) to the rendezvous problem for mobile agents. However, all the above results are restricted to safe or fault-free networks.

Recently attention has focused on designing mobile agent protocols for networks which are faulty, in particular, where there is a black hole, that is a harmful network site that destroys any visiting agent. The research on such networks have concentrated on locating the black hole. In asynchronous systems, this has been studied under two different methods—using whiteboards [11,12] or using tokens [14] to mark edges. The objective here is minimizing the number of agents that fall into the black hole and the number of moves. In the case of synchronous agents, the objective is to minimize the time taken by the surviving agents to locate the black hole [9,18]. The general case of multiple black holes has been considered only by Cooper et al. [8]. All these problems assume that the team of agents start from the same node, i.e. they are co-located. When the agents start from distinct nodes, it is very difficult to gather the agents while avoiding the black hole nodes. This has been studied earlier only in the case of ring networks containing a single black hole, by Dobrev et al. [13], where the authors give solutions to rendezvous and near-gathering assuming the knowledge of topology and the size of the network.

2 The Model and Definitions

2.1 The Model

The environment is modelled by the tuple $(G, \xi, p, \lambda, \eta)$ where G is an undirected connected graph, ξ is a set of agents and p specifies the initial placement of the

agents in the graph G (i.e. $\forall A \in \xi, p(A) = v : v \in V(G)$). The number of nodes is denoted by $n = |V(G)|$ and the number of agents is denoted by $k = |\xi|$. The agents can move from one node to its adjacent node by traversing the edge connecting them. The edges incident to a node v are locally oriented i.e. they are labelled as $1, 2, \dots, d(v)$, where $d(v)$ is the degree of node v . Notice that each edge $e = (u, v)$ has two labels, one for the link or port at node u and another for the link at node v . The edge labelling of the graph G is specified by $\lambda = \{\lambda_v : v \in V\}$, where for each vertex u , $\lambda_u : \{(u, v) \in E : v \in V\} \rightarrow \{1, 2, 3, \dots, d(u)\}$ defines the labelling on its incident edges. For any edge (u, v) we use $\lambda(u, v)$ to denote the pair $(\lambda_u(u, v), \lambda_v(u, v))$.

The function $\eta : E(G) \rightarrow \{0, 1\}$ denotes which edges are safe/faulty. An edge $e \in E(G)$ is safe if $\eta(e) = 1$ and faulty otherwise. The faults are permanent, so any edge that is faulty at the start of the algorithm remains so until the end and no new faulty edge appears during the execution of the algorithm.

The node from where an agent A starts the algorithm (i.e. the initial location) is called the *homebase* of agent A . The agents are all identical (i.e. they do not have distinct names or labels) and they execute the same algorithm. An agent may enter the system at any time and at any location, and on entry, an agent immediately starts its individual execution of the algorithm. The system is totally asynchronous, such that every action performed by an agent takes a finite but otherwise unpredictable amount of time. As in previous papers on the subject, we assume that the agents communicate by reading and writing information on public whiteboards locally available at the nodes of the network. Thus, each node $v \in G$ has a whiteboard (which is a shared region of its memory) and any agent visiting node v can read or write to the whiteboard. Access to the whiteboard is restricted by fair mutual exclusion, so that, at most one agent can access the whiteboard of a node at the same time, and any requesting agent will be granted access within finite time. An agent that is granted access to the whiteboard at node v , is allowed to complete its activity at that node before relinquishing access to the whiteboard (i.e. access control is non preemptive).

Note that it is not necessary for two agents A and B traversing the same edge $e = (u, v)$ of the graph, to arrive at node v in the same order in which they left node u . However, using the whiteboards at the nodes, it is easy to implement a first-in first-out (FIFO) strategy such that agents traversing an edge can be assumed to have reached their destination in order (i.e. an agent cannot overtake another while traversing an edge). For the rest of this paper, we shall assume this FIFO property; this will simplify the description of our algorithms.

2.2 Directed Graphs and Coverings

In this section, we present some definitions and results related to directed graphs and their coverings, which we use to characterize those network where rendezvous is possible. A directed graph (digraph) $D = (V(D), A(D), s_D, t_D)$ possibly having parallel arcs and self-loops, is defined by a set $V(D)$ of vertices, a set $A(D)$ of arcs and by two maps s_D and t_D that assign to each arc two elements of $V(D)$: a source and a target (in general, the subscripts will be omitted). A digraph D is

strongly connected if for all vertices $u, v \in V(D)$, there exists a path between u and v . A *symmetric* digraph D is a digraph endowed with a symmetry, that is, an involution $Sym : A(D) \rightarrow A(D)$ such that for every $a \in A(D)$, $s(a) = t(Sym(a))$. A bidirectional network can be represented by a strongly connected symmetric digraph, where each edge of the network is represented by a pair of symmetric arcs. In this paper, we consider digraphs where the vertices and the arcs are labelled with labels from a recursive label set L and such digraphs will be denoted by (D, μ_D) , where $\mu_D : V(D) \cup A(D) \rightarrow L$ is the labelling function. In general, the label on an arc a would be a pair (x, y) and the labelling μ_D should satisfy the property that if $\mu_D(a) = (x, y)$ then $\mu_D(Sym(a)) = (y, x)$, for every arc $a \in D$.

A digraph homomorphism γ between the digraph D and the digraph D' is a mapping $\gamma : V(D) \cup A(D) \rightarrow V(D') \cup A(D')$ such that if u, v are vertices of D and a is an arc such that $u = s(a)$ and $v = t(a)$ then $\gamma(u) = s(\gamma(a))$ and $\gamma(v) = t(\gamma(a))$. A homomorphism from (D, μ_D) to (D', μ'_D) is a digraph homomorphism from D to D' which preserves the labelling, i.e., such that $\mu'_D(\gamma(x)) = \mu_D(x)$ for every $x \in V(D) \cup A(D)$.

We now define the notion of graph coverings, borrowing the terminology of Boldi and Vigna[5]. A *covering projection* is a homomorphism φ from D to D' satisfying the following: (i) For each arc a' of $A(D')$ and for each vertex v of $V(D)$ such that $\varphi(v) = v' = t(a')$ there exists a unique arc a in $A(D)$ such that $t(a) = v$ and $\varphi(a) = a'$. (ii) For each arc a' of $A(D')$ and for each vertex v of $V(D)$ such that $\varphi(v) = v' = s(a')$ there exists a unique arc a in $A(D)$ such that $s(a) = v$ and $\varphi(a) = a'$.

The fibre over a vertex v' (resp. an arc a') of D' is the set $\varphi^{-1}(v')$ of vertices of D (resp. the set $\varphi^{-1}(a')$ of arcs of D).

If a covering projection $\varphi : D \rightarrow D'$ exists, D is said to be a *covering* of D' via φ and D' is called the base of φ . A symmetric digraph D is a *symmetric covering* of a symmetric digraph D' via a homomorphism φ if D is a covering of D' via φ such that $\forall a \in A(D), \varphi(Sym(a)) = Sym(\varphi(a))$. A digraph D is *symmetric-covering-minimal* if there does not exist any graph D' not isomorphic to D such that D is a symmetric covering of D' .

The notions of coverings extend to labelled digraphs in an obvious way: the homomorphisms must preserve the labelling. Given a labelled symmetric digraph (H, μ_H) , the minimum base of (H, μ_H) is defined to be the labelled digraph (D, μ_D) such that (i) (H, μ_H) is a symmetric covering of (D, μ_D) and (ii) (D, μ_D) is symmetric covering minimal.

The following results on digraph coverings were proved in [5].

Property 1. Given two non-empty strongly connected digraphs D, D' , each covering projection φ from D to D' is surjective; moreover, all the fibres have the same cardinality. This cardinality is called the number of sheets of the covering.

Property 2. If the digraph (H, μ_H) is a covering of (D, μ_D) via φ , then any execution of an algorithm \mathcal{P} on (D, μ_D) can be lifted up to an execution on

(H, μ_H) , such that at the end of the execution, for any $v \in V(H)$, v would be in the same state as $\varphi(v)$.

2.3 Definitions and Properties

Given any deterministic (distributed) algorithm \mathcal{P} and a network $(G, \xi, p, \lambda, \eta)$, the order in which the various actions are performed by the agents defines an *execution* of the algorithm on the network $(G, \xi, p, \lambda, \eta)$. We define the *synchronous* execution of an algorithm \mathcal{P} to be the particular execution where all agents start executing at exactly the same time and every action taken by any agent takes exactly one unit of time.

We define the *extended-view* of the network $(G, \xi, p, \lambda, \eta)$ as the labelled digraph (H, μ_H) such that, H consists of two disjoint vertex sets V_1 and V_2 and a set of arcs \mathcal{A} as defined below:

- $V_1 = V(G)$;
- $\mu_H(v) = |\{A \in \xi : p(A) = v\}|, \forall v \in V_1$;
- For every safe edge $e = (u, v) \in E(G)$, there are two arcs $a_1, a_2 \in \mathcal{A}$ such that $s(a_1) = t(a_2) = u$, $s(a_2) = t(a_1) = v$, and $\mu_H(a_1) = (\lambda_u(e), \lambda_v(e))$, $\mu_H(a_2) = (\lambda_v(e), \lambda_u(e))$.
- For every faulty edge $e = (u, v)$, there are vertices u' and $v' \in V_2$ with $\mu_H(u') = \mu_H(v') = -1$ and arcs $(u, u'), (u', u), (v, v')$ and $(v', v) \in \mathcal{A}$ with labels $(\lambda_e(u), 0)$, $(0, \lambda_e(u))$, $(\lambda_e(v), 0)$, and $(0, \lambda_e(v))$ respectively;

Here, the vertices in V_1 represent the (safe) nodes of the network and the vertices in V_2 represent (imaginary) Black-Holes. The label on a *safe* vertex v denotes the number of agents that started from the corresponding node, whereas the label on a *black-hole* vertex is always -1 . Intuitively, the *extended-view* can be thought of as a canonical representation of the network.

The following results follow from the definition of the *extended-view* of a network and the Properties 1 and 2.

Lemma 1. *For any deterministic algorithm \mathcal{P} , a synchronous execution of \mathcal{P} on the network $(G, \xi, p, \lambda, \eta)$ is equivalent to a synchronous execution of algorithm \mathcal{P} on the extended-view (H, μ_H) , such that the final state of any node in G is exactly same as the state of the corresponding vertex in H .*

Lemma 2. *If the extended view of two networks have same minimum-base (D, μ_D) then all nodes in the two networks which belong to the pre-image of a vertex $v \in D$ would always be in the same state, during a synchronous execution of any algorithm \mathcal{P} .*

3 Impossibility Results

In this section, we determine some necessary conditions for solution to the rendezvous problem. In the following, whenever the extended-view of a network is symmetric-covering minimal, we shall say the network is *minimal*.

Lemma 3. *In a network containing τ dangerous links and k dispersed agents, τ agents may die while executing any algorithm for rendezvous. Thus, it is not always possible to rendezvous more than $k - \tau$ agents even if the network topology is known to the agents.*

Lemma 4. *It is impossible to rendezvous $k - \tau$ agents in a network whose extended-view is not symmetric-covering minimal.*




Fig. 1. The network in (a) cannot be distinguished from the networks in (b) and (c) due to the slow edges. (The slow edges are marked by arrows and the dashed lines represent faulty edges).

In the next section, we show how to solve rendezvous of $k - \tau$ agents in any network that is minimal. Thus, we have a complete characterization of networks where rendezvous of $k - \tau$ agents is possible. However, solution to rendezvous requires at least some prior knowledge about the network, as we show below. Notice that since the network is asynchronous, an agent may take an arbitrarily long time to traverse some edge. Thus, a *slow* edge (i.e. one which the agents take a long time to traverse) is indistinguishable from a faulty edge. During the execution of an algorithm, the presence of slow edges may divide the network into two equal parts (for example, see Figure 1) and in such cases, it is not possible for the agents to terminate the algorithm unless an accurate estimate of network size is available.

Lemma 5. *It is impossible to solve rendezvous (with termination detection) of $k - \tau$ agents even in minimal networks if the agents know only an upper bound B on the number of nodes n such that $n \leq B \leq 2n$.*

The algorithm presented in this paper works only for networks which are minimal. We say that an algorithm \mathcal{P} is an universally *effective* algorithm for rendezvous of $w > 1$ agents if, when executed on any network where rendezvous of w agents is possible, algorithm \mathcal{P} always succeeds in achieving rendezvous within a finite time. We have the following negative result on the existence of such an algorithm.




Fig. 2. The networks in (a) and (b) have same topology but differ in the location of faulty edges (shown by dashed lines). (c) The minimum base for the network in (a).

Lemma 6. *There does not exist any universally effective algorithm for rendezvous of $1 < w \leq k$ agents even if the network topology is known a-priori to the agents.*

Proof. Consider the two networks shown in Figure 2(a) and (b). Each network has the same topology and there are three faulty edges in each (but their locations are different). In the network of Figure 2(b) there are two edges (marked by arrows) which are very slow. Since slow edges can not be distinguished from faulty edges, the agents would not be able to determine whether they are in the first network or the second. Thus any rendezvous algorithm \mathcal{P} (that terminates within a finite time) must achieve the same result in both networks. Notice that the first network has an extended view which is not minimal (the minimum base is shown in Figure 2(c)). Since algorithm \mathcal{P} must fail to achieve rendezvous in the first network, it must also fail in the second one, even though it is possible to rendezvous in the second network. Thus algorithm \mathcal{P} is not *effective*.

4 Solution Protocol

In this section we present an algorithm for solving Rendezvous in faulty networks, using the knowledge of only an upper bound B on the network size, such that $n \leq B < 2n$. As shown in the previous section, there is no effective algorithm for Rendezvous in faulty networks. Our algorithm always works for any network whose extended-view is covering minimal, achieving the rendezvous of the maximum number of agents possible (i.e. $k - \tau$ agents). We also analyze the complexity of our algorithm and show that it is optimal in terms of the number of moves made by the agents.

Theorem 1. *For solving rendezvous of $(k - \tau)$ agents in an arbitrary network $(G, \xi, p, \lambda, \eta)$ without any knowledge other than the size of the network, the agents need to make at least $\Omega(m(m + k))$ moves in total.*

4.1 The Algorithm for Rendezvous

We can ensure that no more than one agent dies while traversing the same link, using the *cautious walk* technique as in [13]. At each node, all the incident edges are considered to be unexplored in the beginning. Whenever an agent A at a node u has to traverse an unexplored edge $e = (u, v)$, agent A first marks link $\lambda_u(e)$ as “Being Explored” and if it is able to reach the other end v successfully, it immediately returns to node u and re-marks the link $\lambda_u(e)$ as “safe”. During the algorithm we follow the rule that no agent ever traverses a link that is marked “Being Explored”. This ensures no more than τ agents may die during the algorithm.

We now briefly describe our algorithm for rendezvous (Algorithm RDV). Due to the space constraint, we present only an oversimplified version of the algorithm without the minor technical details. The complete pseudo-code for the algorithm can be found in the full paper.

At any stage of the algorithm, there are teams of agents, each team possessing a *territory* which is a connected acyclic subgraph of G (disjoint from other territories). Each team of agents tries to expand its territory until it spans a majority of the nodes. Once a team is able to acquire more than half the nodes of the network, it wins and agents from all other teams join the winning team to achieve rendezvous.

Initially each territory consists of only the starting node(homebase) of an agent (all agents that start from that node are in this team). Note that if an agent on start-up, finds that its homebase has already been acquired by some other team, it simply joins this team. The algorithm proceeds in a series of exploring and competing rounds. In an exploring round, the team of agents try to expand its territory by exploring new edges and acquiring new nodes. On the other hand, in the competing round a team tries to defeat another team and conquering their territory. The competition between two teams occurs by comparison using the tuple $(j, Code)$ where j is the round number and Code is an encoding of the territory(and its immediate neighborhood).

The territory of each team is a rooted tree and the information stored in the root defines the status of the agents in the team (e.g. whether they are in an exploring round or competing round). Every other node in the tree stores a pointer to its parent in the tree. The status of the root can either INIT-EXPLORE, INIT-COMPETE, COMPETE, LOST, or END.

Only one agent in a team can be competing state (this is called the active agent and all others are passive). When a team is in exploring round, the active agent initiates the exploration and thereafter every agent may participate in the exploration. Each agent explores one unexplored edge and if it succeeds in reaching the other side, it reports this to the root of the tree and the tree is updated accordingly. If the edge connects to a node that is already explored (by another agent or the same agent) then it is marked as tree edge (T-edge) and the new node becomes part of the tree. Otherwise it is marked as non-tree edge (NT-edge).

In the algorithm below, T refers to the tree representing the territory to which the agent belongs. The root of T is denoted by r and $\text{Root_Status}(T)$ is the status

of the root as written on the whiteboard of r (A copy of T is also stored on the whiteboard). A node u is called a *neighbor* of the tree T if there is a NT-edge between u and some node in T . Each agent starts in *active* state, but may become *passive* or *finished*. The following steps are executed by an agent A :

Algorithm RDV

```

< T, rNum > := Initialize;
While (Status(A) ≠ finished) {
  Case(Status(A) is active and Root_Status(T) is INIT-EXPLORE) {
    If (|T| > B/2) terminate;
    InitExplore;
    If (Root_Status(T) = LOST) become passive and exit Case;
    Set Root_Status(T) to EXPLORE and become Passive;
  }
  Case( Status(A) is passive ) {
    While(Root_Status(T) ∉ { EXPLORE, INIT-EXPLORE, END } )
      Sleep until woken up;
    If(Root_Status(T) is END) become finished and exit;
    Explore an unexplored edge and go back to  $r$  to update  $T$ ;
    If(Root_Status(T) is EXPLORE) // i.e. no active agents
      become active and set Root_Status(T) to INIT-COMPETE
  }
  Case(Status(A) is active and Root_Status(T) is INIT-COMPETE) {
    do{  $T_{OLD} := T$  ;
      InitCompete( $rNum$ );
      If (Root_Status(T) = LOST) become passive and exit Case;
      Result = Compete( $rNum$ );  $rNum := rNum + 1$ ;
    }While(  $T_{OLD} \neq T$  );
    If(Root_Status(T) ≠ LOST) Root_Status(T):= INIT-EXPLORE;
  }
}

```

PROCEDURE *Initialize*: If the homebase of the agent is already part of a tree, then the agent joins this team in passive state. Otherwise, it initializes the tree with only the homebase node and then starts the algorithm in active state with $rNum = 1$.

PROCEDURE *InitExplore*: The agent initiates a new exploring round by traversing T writing “EXPLORE” on every node and waking up as many agents as needed for exploration (more precisely, it wakes-up x agents where x is the number of unexplored edges currently incident to T , unless less than x agents are available in T —in which case it wakes-up everyone).

PROCEDURE *InitCompete*(j): The agent initiates the competing round j by traversing T , writing ($COMPETE, j$) on each node also assigning labels to the nodes in T . Next, it reads the labels written on the neighboring nodes and constructs an encoding of T and its neighbors, called CODE (this is used for comparisons). Finally it writes the CODE on every node in T .

PROCEDURE *Compete(j)*: During this procedure, agent A competes with each node u that is neighbor of T , until it wins or loses. If u is marked END, agent A terminates after writing END on every node of T , waking up any sleeping agents and merging with the tree containing u . Else, if u is in a bigger competing round or in the same competing round with a larger Code, then agent A loses (i.e. it becomes passive and goes back to its root to sleep). Otherwise if u is in exploring round or in a smaller competing round or same round but has smaller Code, then $\text{Invade}(u)$ is invoked to determine if the result of the competition is win or loss.

PROCEDURE *Invade(u)*: The agent attempts to acquire the tree containing node u . The agent follows the father-links from u to reach the current root r_u of u . At node r_u it uses the usual comparison criteria and if r_u is bigger or equal, then it loses (i.e. it becomes passive and goes back to its root to sleep). Otherwise it wins and it acquires the Tree rooted at r_u , by reversing the father-links in the path from r_u to v where v is the node in the agent's territory from where it started the *Invade* procedure.

PROCEDURE *terminate*: The agent writes END on all nodes in T and then writes $\text{Root_Status}(T)=\text{END}$ at the root of T ; The agent now becomes finished and locally terminates.

4.2 Analysis of the Algorithm

Theorem 2. *Algorithm RDV correctly solves Rendezvous for $k - \tau$ agents in any network whose extended view is symmetric-covering minimal.*

This result follows from the following lemmas:

Lemma 7. *During the algorithm RDV, the following holds: (i) Each territory is a Tree. (ii) The territories are disjoint. (iii) There is at most one active agent in each territory.*

Proof. (i) Initially each territory is a tree by construction; addition of a new edge does not create cycles because only tree edges are added. When two trees are merged, one is 'larger' than other and the active agent of the larger tree performs the merging by changing the father-links and updating the root. (Notice that there can not be two active agents in a territory.) Thus the merged territory is a tree. (ii) Each tree is a rooted tree and every node contains a pointer to its parent in the tree. Thus, a node cannot belong to two trees. (iii) Initially an agent becomes active only when it successfully constructs T and writes it at the root node. Thus, due to the mutual exclusion property of the whiteboards, there is only one active agent initially in the tree (all other agents in the same tree must start in passive state). A passive agent can become active only when there are no active agents in the tree.

Lemma 8. *There is no deadlock in the algorithm RDV.*

Proof. Notice that as long as there is some active agent, the algorithm progresses. In each tree T , there is at least one active agent unless root-status is EXPLORE. Suppose the root-status of every tree is EXPLORE, then each agent is currently exploring some unexplored edge. There are only τ faulty links in the network and $k > \tau$, so at least one of the exploring agents must return safely from the exploration. This agent now becomes active.

Lemma 9. (i) *If the network is minimal, then exactly one node has Root-Status(T) = END.* (ii) *When one node has Root-Status(T) = END, every alive agent in G eventually joins this tree.* (iii) *$(k - \tau)$ agents eventually become finished and reach the node having Root-Status(T) = END.*

Proof. (i) First notice that due to the constraint on the bound B , only one Tree T can have a size greater than $B/2$. Due to Lemma 7, the trees are disjoint and each has a unique root. Thus, at most one node may have Root-Status(T)=END. We now show that at least one node eventually reaches Root-Status(T)=END. A team of agents in tree T stops expanding its Tree only when either Root-Status(T)=END or, all the trees neighboring T have the same CODE and round number. In the later case, the team starts an exploring round. Notice that there must exist a Tree T which has some non-faulty edges incident to it that are either unexplored or being explored (otherwise the network is either disconnected or, not minimal). So one of these edges would be added to T . Thus the size of T would keep increasing until it contains more than $B/2$ nodes and Root-Status(T) = END.

(ii) Every agent either dies or reaches a node labelled “END”. All nodes that are labelled “END” are part of the same tree.

(iii) Due to the use of *cautious walk*, at most τ agents may die. Thus each of the surviving agent eventually reaches a node v labelled “END” and goes to the root of this tree which is the node having Root-Status(T)=END.

We now analyze the cost of algorithm RDV. We first count the number of competing rounds and exploring rounds performed during the algorithm.

Lemma 10. *There are $O(m + k)$ competing rounds. The number of exploring rounds is at most the number of competing rounds.*

Proof. A new competing round is started whenever T is expanded, i.e. the new territory T contains one more edge or one more agent than the previous territory T_{OLD} . Thus there can be at most $(m+k)$ competing rounds. After every exploring round completes, there is one competing round. So, the number of exploring rounds can not be more than the number of competing rounds.

Lemma 11. *The agents make at most $O(n(m + k))$ moves in the all the exploring rounds combined.*

Proof. The Procedure InitExplore() makes $|T|$ moves in a tree T . Only one agent (the active one) in every tree executes this procedure. So, this accounts for $O(n)$

moves per exploring round and $O(n(m+k))$ moves in total. Other than that, every passive agent that is woken up makes $O(n)$ moves to go to an unexplored edge, explore it and report it to the root. Since each unexplored edge will be explored once (or at most twice), this cost can be counted per newly explored edge. Thus, this accounts for $O(n.m)$ moves.

Lemma 12. *The agents make at most $O(m(m+k))$ moves in all the competing rounds combined.*

Proof. Only the active agent in a tree participates in the competing round. During procedure InitCompete every edge in G is traversed a constant number of times; this accounts for $O(m)$ moves per round. Similarly $O(m)$ moves are made per round during procedure Compete, for the comparisons with each neighbor. Each execution of Invade() takes $O(n)$ moves, because only tree-edges are traversed and no edge is traversed twice. Whenever an agent execute Invade(), it either wins or loses. A losing agent never competes again, so the total contribution from losing agents is $O(k.n)$. Edges traversed by the winning agents are disjoint, so this accounts for $O(n)$ moves per round.

Due to the above lemmas and Theorem 1, we have the following result:

Theorem 3. *The moves complexity of algorithm RDV is $O(m(m+k))$. Thus algorithm RDV is optimal.*

5 Conclusions

We considered the problem of rendezvous of mobile agent in a faulty network and showed that it is possible to rendezvous at most $k - \tau$ in any network containing k dispersed agents and τ faulty links. We determined the condition under which this is possible and gave an algorithm for solving the problem under this condition. The algorithm we presented is optimal in terms of the total number of moves made by the agents and requires no prior information about the network topology (except the size). Moreover, we showed that it is impossible to have an *effective* algorithm for rendezvous, one that always achieves the rendezvous of as many agents as possible in any given network.

Notice that the only information needed by our algorithm is a strict upper bound on the number of nodes. We assumed that the faulty links do not disconnect the network. In the case of a disconnected network, we can still rendezvous the agents in a connected component if we know a good bound on its size. For example, if the component contains a majority of the nodes (i.e. more than half of them), then original network size can be used as the bound. In this case, if τ is equal to the number of outgoing edges from the component then we can rendezvous $k' - \tau$ agents where k' agents are initially located in this component. For networks containing a single black hole (which does not disconnect the network), our algorithm can be used to rendezvous $k - d_{BH}$ agents where d_{BH} is the degree of the black-hole.

For the results in this paper, we considered two optimization criteria—minimizing the number of agents that are destroyed and the number of moves taken by the surviving agents to rendezvous. It would be interesting to also consider the optimization of whiteboard memory or agent memory when solving the rendezvous problem in faulty networks.

References

1. Alpern, S., Gal, S.: *The Theory of Search Games and Rendezvous*. Kluwer, Dordrecht (2003)
2. Angluin, D.: Local and global properties in networks of processors. In: *STOC '80. Proc. 12th ACM Symp. on Theory of Computing*, pp. 82–93. ACM Press, New York (1980)
3. Barrière, L., Flocchini, P., Fraigniaud, P., Santoro, N.: Rendezvous and Election of Mobile Agents: Impact of Sense of Direction. *Theory of Computing Systems* 40(2), 143–162 (2007)
4. Boldi, P., Vigna, S.: An effective characterization of computability in anonymous networks. In: Welch, J.L. (ed.) *DISC 2001. LNCS*, vol. 2180, pp. 33–47. Springer, Heidelberg (2001)
5. Boldi, P., Vigna, S.: Fibrations of graphs. *Discrete Math.* 243, 21–66 (2002)
6. Chalopin, J., Das, S., Santoro, N.: *Rendezvous of Mobile Agents in Anonymous Networks with Faulty Channels*. Technical Report TR-2007-02, University of Ottawa (2007)
7. Chalopin, J., Godard, E., Métivier, Y., Ossamy, R.: Mobile agents algorithms versus message passing algorithms. In: Shvartsman, A.A. (ed.) *OPODIS 2006. LNCS*, vol. 4305. Springer, Heidelberg (2006)
8. Cooper, C., Klasing, R., Radzik, T.: Searching for black-hole faults in a network using multiple agents. In: Shvartsman, A.A. (ed.) *OPODIS 2006. LNCS*, vol. 4305, pp. 320–332. Springer, Heidelberg (2006)
9. Czyzowicz, J., Kowalski, D., Markou, E., Pelc, A.: Searching for a black hole in synchronous tree networks. *Combinatorics, Probability and Computing* (to appear, 2007)
10. Dessmark, A., Fraigniaud, P., Kowalski, D., Pelc, A.: Deterministic rendezvous in graphs. *Algorithmica* 46, 69–96 (2006)
11. Dobrev, S., Flocchini, P., Prencipe, G., Santoro, N.: Mobile agents searching for a black hole in an anonymous ring. *Algorithmica* (to appear, 2007)
12. Dobrev, S., Flocchini, P., Prencipe, G., Santoro, N.: Finding a black hole in an arbitrary network: optimal mobile agents protocols. *Distributed Computing* (to appear, 2007)
13. Dobrev, S., Flocchini, P., Prencipe, G., Santoro, N.: Multiple agents rendezvous in a ring in spite of a black hole. In: Papatriantafyllou, M., Huel, P. (eds.) *OPODIS 2003. LNCS*, vol. 3144, pp. 34–46. Springer, Heidelberg (2004)
14. Dobrev, S., Flocchini, P., Kralovic, R., Santoro, N.: Exploring a dangerous unknown graph using tokens. In: *Proc. of 5th IFIP International Conference on Theoretical Computer Science (TCS'06) (2006)*
15. Gasieniec, L., Kranakis, E., Krizanc, D., Zhang, X.: Optimal memory rendezvous of anonymous mobile agents in a unidirectional ring. In: Wiedermann, J., Tel, G., Pokorný, J., Bieliková, M., Štuller, J. (eds.) *SOFSEM 2006. LNCS*, vol. 3831, pp. 282–292. Springer, Heidelberg (2006)

16. Godard, E., Métivier, Y., Muscholl, A.: Characterization of classes of graphs recognizable by local computations. *Theory of Computing Systems* 37(2), 249–293 (2004)
17. Klasing, R., Markou, E., Pelc, A.: Gathering asynchronous oblivious mobile robots in a ring. In: Asano, T. (ed.) *ISAAC 2006*. LNCS, vol. 4288, pp. 744–753. Springer, Heidelberg (2006)
18. Klasing, R., Markou, E., Radzik, T., Sarracco, F.: Hardness and approximation results for black hole search in arbitrary networks. *Theoretical Computer Science* (to appear, 2007)
19. Kranakis, E., Krizanc, D., Markou, E.: Mobile agent rendezvous in a synchronous torus. In: Correa, J.R., Hevia, A., Kiwi, M. (eds.) *LATIN 2006*. LNCS, vol. 3887, pp. 653–664. Springer, Heidelberg (2006)
20. Mazurkiewicz, A.: Distributed enumeration. *Inf. Processing Letters* 61(5), 233–239 (1997)
21. Sakamoto, N.: Comparison of initial conditions for distributed algorithms on anonymous networks. In: *PODC '99*. Proc. 18th ACM Symposium on Principles of Distributed Computing, pp. 173–179. ACM Press, New York (1999)
22. Yamashita, M., Kameda, T.: Computing on anonymous networks: Parts I and II. *IEEE Trans. Parallel and Distributed Systems* 7(1), 69–96 (1996)
23. Yu, X., Yung, M.: Agent rendezvous: A dynamic symmetry-breaking problem. In: Meyer auf der Heide, F., Monien, B. (eds.) *ICALP 1996*. LNCS, vol. 1099, pp. 610–621. Springer, Heidelberg (1996)