

Brief Announcement: Broadcast via Mobile Agents in a Dynamic Network: Interplay of Graph Properties & Agents

William K. Moses Jr. ✉ 

Department of Computer Science, Durham University, UK

Amanda Redlich ✉ 

Department of Mathematics & Statistics, University of Massachusetts Lowell, MA, USA

Frederick Stock ✉ 

Department of Computer Science, University of Massachusetts Lowell, MA, USA

Abstract

In this paper, we revisit the problem of BROADCAST, introduced by Das, Giachoudis, Luccio, and Markou [OPODIS, 2020], where $k + 1$ agents are initially placed on an n node dynamic graph, where 1 agent has a message that must be broadcast to the remaining k ignorant agents. The original paper studied the relationship between the number of agents needed to solve the problem and the edge density of the graph. The paper presented strong evidence that edge density of a graph, or the number of redundant edges within the graph, may be the correct graph property to accurately differentiate whether $k = o(n)$ agents (low edge density) or $k = \Omega(n)$ agents (high edge density) are needed to solve the problem.

In this paper, we show that surprisingly, edge density is not in fact the differentiating property. The original paper presents graphs with edge density $1.1\bar{6}$ that require $\Omega(n)$ agents, however, we construct graphs with edge density $> 1.1\bar{6}$ and develop an algorithm to solve the problem on those graphs using only $o(n)$ agents. We further show that the relationship between edge density and number of agents is fairly weak by first constructing graphs with edge density tending to 1 from above that require $\Omega(n/f(n))$ agents to solve, for any function $f(n) \rightarrow \infty$ as $n \rightarrow \infty$. We then construct an infinite family of graphs with edge density $< \rho$ requiring exactly k ignorant agents to solve BROADCAST, for any $k > 0$ and $\rho > 1$.

Finally, we investigate different versions of connectivity as possible properties determining the number of agents. We show that edge connectivity is not related to the number of agents: it is possible for a graph to have low (constant) edge connectivity but require a high (linear) number of agents to solve BROADCAST. More generally, for any arbitrary λ , we construct a family of graphs on n nodes with edge connectivity $(n - 1)/\lambda$ requiring exactly $k = n - 2\lambda + 1$ agents. On the other hand, we show vertex connectivity *can* impact the required number of agents: for any graph with vertex connectivity ≥ 3 and minimum degree δ , $k \geq \delta - 1$ agents are required to solve BROADCAST. Finally, we show that for any graph containing a certain type of bond with m edges, BROADCAST is unsolvable when $k \leq m - 2$.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Distributed algorithms; Computing methodologies $\rightarrow$ Mobile agents

Keywords and phrases Mobile agents, mobile robots, broadcast, dynamic graph, dynamic network

Digital Object Identifier 10.4230/LIPIcs.SAND.2025.17

Related Version *Full Version:* <https://arxiv.org/abs/2504.05442>

1 Introduction

In this paper, we study the specific problem of BROADCAST in the context of mobile agents on a dynamic graph. This can be used to model situations in which actual robots are trying to move through constrained spaces (e.g., a building with rooms and corridors connecting



© William K. Moses Jr., Amanda Redlich, and Frederick Stock;
licensed under Creative Commons License CC-BY 4.0

4th Symposium on Algorithmic Foundations of Dynamic Networks (SAND 2025).

Editors: Kitty Meeks and Christian Scheideler; Article No. 17; pp. 17:1–17:5

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

those rooms). It can also be used to model situations on networks where each agent is a virtual agent that moves from computer to computer in order to spread some information or perform some task on a subset of computers. Considering mobile agents on *dynamic* graphs allows for broader applications in modeling the real world (e.g., earthquakes/fires changing the connectivity between rooms in a building) or a virtual one (e.g., two computers in a network become connected/disconnected because a cable is added/cut between them).

Within this paradigm, many problems have been studied, modeling an array of real world issues: gathering [3], scattering [2], pattern formation [6], dispersion [1], and convergence [4]. Here, we study the problem of BROADCAST as introduced in Das, Giachoudis, Luccio, and Markou [5]. In it, there is initially one agent with a message that is to be transmitted to a number of other agents located on the nodes of the graph.

The problem of BROADCAST, as studied in this paper, acts as an interesting bridge between agent-specific problems and traditional distributed computing. In traditional distributed computing, BROADCAST requires all nodes to participate in the broadcast of a message originating at one node. However, here, the number of agents is independent of the number of nodes and furthermore the agents can move around on the nodes. By studying such “bridge” problems, we may develop a deeper understanding of the links between the mobile agents on a graph model and standard distributed computing.

In the paper that introduced BROADCAST in this context, the focus was on the question of the minimum number of agents that allow the problem to be solvable. A number of graph classes of varying density were presented with lower and upper bounds on the necessary number of agents. Based on this, [5] observes the number of agents “apparently depends on the density of the underlying graph G or, the number of redundant edges in G ”.

However, this is not the case! Here we give infinite classes of graphs in which edge density and number of agents are unrelated. We further investigate edge connectivity, vertex connectivity, and other graph properties and analyze their potential relationships with BROADCAST.

1.1 Preliminaries & Definitions

Graph Definitions

Initially, we are given a graph $G(V, E)$, with node set V and edge set E , which serves as the base graph. In each round i , the adversary chooses some (possibly empty) subset of edges E' from G to remove for that round such that the graph $G_i(V, E \setminus E')$ is still connected.


The *edge density* ρ of a graph with m edges and n nodes is m/n .

A *generalized theta graph*, denoted by $\theta_{d_1, d_2, \dots, d_\ell}$, consists of two fixed nodes joined by ℓ paths, where each path i contains $d_i \geq 1$ nodes. See Figure 1 for an example illustration.

A *cut* in a graph is a set of edges whose deletion renders the graph disconnected. The *edge connectivity* of a graph is the size of the smallest cut of the graph. The *vertex connectivity* of a graph is the smallest number of nodes whose deletion renders the graph disconnected. A *bond* of a graph is a set of edges such that if all the edges in the set are removed the graph becomes disconnected, but if all but any one edge is removed the graph is still connected. In other words, a bond is a minimal cut set of edges.

Agent Description

We consider that initially there is one agent with the message $\mathcal{M}$ and k agents without the message. To keep our analysis consistent with the original context of [5], we maintain the same assumptions on agents used in that paper. We assume that agents have unique IDs,



■ **Figure 1** An example theta graph $\theta_{7,4,2,3,4,3,1,7}$.

have some local memory for computation, and each agent starts at a unique node of G . For each $i \geq 0$, each agent has global visibility of the graph G_i , including any edges the adversary has deleted, and can see which agents occupy which nodes of the graph, including whether the agents have knowledge of $\mathcal{M}$ (call them *source agents*) or not (call them *ignorant agents*). As stated in [5], these assumptions can be justified by considering that the adversary is very strong and so the agents need fairly strong capabilities in order to be able to solve the problem.

Time

Time proceeds in synchronous rounds. In each round i , the following occurs in the given order.

1. Adversary chooses a (possibly empty) subset of edges E' and removes them from G to create a connected subgraph $G'(V, E \setminus E')$.
2. Agents perform local computation and/or communicate with co-located agents and choose an edge to move through (or possibly decide to stay at the same node).
3. Agents move through chosen edges.

Problem Statement

The problem statement, as stated in [5], is given below.

► **Definition 1** (The BROADCAST Problem). *Given a constantly connected dynamic network $\mathcal{G}$ based on an underlying graph G consisting of $n \geq 2$ nodes, a source agent that has a message $\mathcal{M}$ and $k \geq 1$ ignorant agents that are initially located at distinct nodes of the network, can the message $\mathcal{M}$ be broadcast to all the agents?*

2 Our Contributions

In this paper, we focus on a fine-grained analysis of the edge density of a graph with n nodes and its relationship to the number of agents needed to solve BROADCAST on that graph. The work of Das, Giachoudis, Luccio, and Markou [5] used edge density to divide the classes of graphs into sparse and dense graphs and showed that there exist sparse graphs (rings) for which BROADCAST can be solved using $o(n)$ agents and more dense graphs (grids) for which BROADCAST requires at least $\Omega(n)$ agents. In particular, they show that for rings,

with edge density 1, $k = 2$ ignorant agents suffice to solve BROADCAST and for $2 \times L$ grids, including a 2×3 grid with edge density $1.1\bar{6}$, $k \geq L = \Omega(n)$ ignorant agents are required to solve BROADCAST.

One might naturally assume that for graphs with edge density greater than $1.1\bar{6}$, $k = \Omega(n)$ agents are required to solve BROADCAST. However, this is not the case. We show that there exists a class of graphs, commonly known as generalized theta graphs, that may have edge density larger than $1.1\bar{6}$ that in fact require only $k = o(n)$ ignorant agents to solve BROADCAST. This surprising result shows that edge density, by itself, may not be the dividing line between graphs that require $k = o(n)$ and $k = \Omega(n)$ ignorant agents. An example of such a construction is as follows. Consider a generalized theta graph with ℓ paths between two nodes, such that each path is of length d . The number of nodes n is $\ell d + 2$ and the edge density is $1 + (\ell - 2)/(\ell d + 2)$, which can be rewritten as $1 + (n - 2 - 2d)/(nd)$. The number of ignorant agents needed to solve BROADCAST is ℓ . When $d = \omega(1)$, we see that $\ell = o(n)$. Let $n = 100$ and $d = \log^* n = 4$. Then the given graph requires $o(n)$ ignorant agents and has edge density $1.225 > 1.1\bar{6}$.

A complementary question one might then ask is, whether in the range of edge density $\rho \in (1, 1.1\bar{6})$, do graphs require very few agents, i.e., $k \ll n$? We answer this question in the negative by constructing, for any arbitrary $f(n) \rightarrow \infty$, a family of graphs with edge density tending to 1 from above (and therefore in particular $< \rho$) requiring $\Omega(n/f(n))$ agents. For example, taking $f(n) = \log n$ gives a graph with edge density tending to 1 but agents tending to ∞ .

One might then wonder if there is any relationship between the exact edge density of a graph and the exact number of ignorant agents needed to solve for that edge density. From [5], we see that for trees, that have edge density < 1 , $k = 1$ suffices. For rings, that have edge density exactly 1, $k = 2$ suffices. For a complete graph, that has edge density $(n - 1)/2$, $k \geq n - 2$ ignorant agents are needed. The question then becomes: are there specific values ρ and k such that any graph with edge density $< \rho$ requires $< k$ ignorant agents for some constant k ? We answer this in the negative. For any $k > 0$ and any $\rho > 1$, we construct an infinite family of graphs with edge density $< \rho$ requiring exactly k ignorant agents to solve BROADCAST. In other words, there is no threshold ρ and number of ignorant agents k for which the statement “any graph whose edge density is below ρ requires $< k$ ignorant agents” is correct.

A natural next line of inquiry is to look towards other graph properties and ask about their relationship with the number of ignorant agents needed to solve BROADCAST. We begin with edge connectivity, that is the defining constraint on the adversary. We show that it is possible for a graph to have low (constant) edge connectivity but require high (linear) number of agents to achieve BROADCAST. (Contrast this, for example, with the ring, which has edge connectivity 2 and also requires 2 agents.) In fact, for arbitrary λ , we construct a family of graphs on n vertices with edge connectivity $(n - 1)/\lambda$ but requiring exactly $n - 2\lambda + 1$ ignorant agents for BROADCAST.

As edge connectivity is thus irrelevant, we next investigate vertex connectivity. In this context we develop an adversary strategy for any graph with vertex-connectivity at least 3. In particular, for all $\beta \geq 3$, all graphs with vertex connectivity β and minimum degree δ require at least $k \geq \delta - 1$ ignorant agents to solve BROADCAST. (Note that for specific classes of graphs we have improved on this bound using earlier arguments; however, this is a more broadly-applicable lower bound.)

We also develop an adversary strategy for any graph which has a specific type of bond with m edges, showing BROADCAST is unsolvable for $k \leq m - 2$ ignorant agents.

2.1 Technical Overview & Challenges

While the problem of BROADCAST seems relatively simple, solving it requires involved algorithms with many small but crucial details. For example, in the process of solving BROADCAST the number of ignorant and source agents is constantly changing as ignorant agents become source agents. Therefore any winning strategy must adjust for the changing ratio between these two classes of agents.

Although [5] observes a potential link between edge density (i.e. “redundant” edges) and number of agents needed, we prove that is not always the case. This is surprising, as more redundant edges means the adversary has more potential strategies available. In fact one of the main technical challenges was to design graphs and BROADCAST strategies that are robust against the adversary’s many options. This edge-density condition inspired our choice of generalized theta graphs; by carefully choosing functions for internal path length and multiplicity, we are able to fully analyze all potential strategies and to generate families of graphs with low edge densities and a large number of required agents. Other constructions combined well-understood subgraphs (cliques, paths) using structural graph techniques to maintain local strategies’ effectiveness in the larger graph. For example, we combine carefully-selected instances of paths and cliques by identifying the two graphs at a single vertex; this allows us to manipulate edge density without affecting BROADCAST strategies.

The edge and vertex connectivity results required new approaches. As connectivity is a global (not local) property, it was not possible to merge subgraphs and use their average values. However, it is possible to combine multiple identical cliques in such a way to preserve their edge-connectivity. Then the strategies for adversary and BROADCAST on a larger graph are exactly generated by the strategies for each subclique. Another approach is to use other graph theoretic properties such as bonds and degrees, and the high connectivity of complete graphs, to construct algorithms and counterexamples. In particular, the adversary is able to define a strategy based on a cut set of vertices or edges, and construct a spanning tree of the remaining connected graph that prevents BROADCAST.

References

- 1 John Augustine and William K. Moses Jr. Dispersion of mobile robots: A study of memory-time trade-offs. *CoRR*, abs/1707.05629, [v4] 2018 (a preliminary version appeared in ICDCN’18).
- 2 L. Barriere, P. Flocchini, E. Mesa-Barrameda, and N. Santoro. Uniform scattering of autonomous mobile robots in a grid. In *IPDPS*, pages 1–8, 2009.
- 3 Mark Cieliebak, Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. Distributed computing by mobile robots: Gathering. *SIAM J. Comput.*, 41(4):829–879, 2012. doi:10.1137/100796534.
- 4 Reuven Cohen and David Peleg. Robot convergence via center-of-gravity algorithms. In *Proc. of the 11th International Colloquium on Structural Information and Communication Complexity, (SIROCCO)*, pages 79–88, 2004. doi:10.1007/978-3-540-27796-5_8.
- 5 Shantanu Das, Nikos Giachoudis, Flaminia L. Luccio, and Euripides Markou. Broadcasting with mobile agents in dynamic networks. In *24th International Conference on Principles of Distributed Systems (OPODIS 2020)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.
- 6 Ichiro Suzuki and Masafumi Yamashita. Distributed anonymous mobile robots: Formation of geometric patterns. *SIAM J. Comput.*, 28(4):1347–1363, 1999. doi:10.1137/S009753979628292X.