

Distance-2 Collision-Free Broadcast Scheduling in Wireless Networks

Valentin Pollet, Vincent Boudet, Jean-Claude König
UM/CNRS
LIRMM - UMR 5506 - CC 477
France
Email: {pollet,boudet,konig}@lirmm.fr

Abstract—In this paper, we study the distance-2 broadcast scheduling problem in synchronous wireless networks of known topology. Two constraints are taken under consideration: the schedule must be collision-free and the nodes at distance 2 must be informed by nodes at distance 1. In general graphs, a tight bound of $\mathcal{O}(\log(n)^2)$ slots to complete the broadcast is known, n being the number of nodes at distance 2. We improve this bound to $\mathcal{O}(\log(n))$ in unit disk graphs, and to $\mathcal{O}(1)$ when the neighbourhoods of the nodes are circular intervals.

I. INTRODUCTION

WE consider the communication model proposed by [1] in which nodes communicate at synchronous slots, using only one frequency. When a node transmits at a given slot, all the nodes connected to it may receive the message. Collisions occur when a node receives several messages at the same slot. Collisions cannot be detected. When exactly one of its neighbours transmits at a given slot, a node is said to be informed.

We study the distance-2 broadcast scheduling problem (D2B) when the topology of the graph is known in advance. Given a graph and a source node, one must schedule the roles of the nodes over several slots in order to inform every node at distance 2 from the source. We assume that nodes at distance 2 from the source cannot be used to inform other nodes. We quantify the quality of a scheduling by the number of slots it uses, the less the better. The number of slots used is an intuitive measure of the time taken to complete the broadcast. This problem is motivated by the fact that knowledge at distance 2 is often assumed when designing communication protocols.

A. Related work

This work is primarily inspired by [2], [3], dealing with D2B in general graphs. They give polynomial algorithms to schedule distance-2 broadcasts using $\mathcal{O}(\log(n)^2)$ slots, where n is the number of nodes at distance 2 from the source. Their result provides a tight upper bound on the number of slots needed to complete broadcast since there exists a family of graphs of diameter 2 requiring a logarithmic number of slots to complete broadcast.

The global broadcast problem, in which a source must flood the whole network, has been widely studied in general graphs. See [4] for a survey.

On a theoretical point of view, solving D2B under the collision model is linked to the exact cover problem. Authors

in [5] give a polynomial algorithm to solve the weighted covering problem for sets of pseudo-disks in the plane. Their algorithm can be adapted to decide the existence of a 1-slot solution to D2B on unit disk graphs.

B. Our results

Our work focuses on restricted classes of graph, and show that the upper bounds in those cases are strictly lower.

First, we exhibit a family of instances over unit disk graphs such that solving D2B requires exactly $\log(n) + 1$ slots, n being the number of nodes at distance 1 from the source. We further prove that it is always possible to complete broadcast using $\mathcal{O}(\log(n))$ slots, tightening the bound.

Then we consider a more restrictive case: when the neighbourhoods of the nodes at distance 1 are circular arcs of the nodes at distance 2. We give a simple greedy algorithm yielding solutions using at most 3 slots. This algorithm proves a constant upper bound on the number of slots needed to complete distance-2 broadcasts in these graphs.

C. Notations

An instance of the distance-2 broadcast scheduling problem (D2B) is given by a graph $G = (V, E)$ and a node x_e of this graph. The immediate neighbours X of x_e are called the *source nodes*. We call *target nodes* the set of nodes Y at distance 2 from x_e . This terminology is motivated by the fact that nodes from Y cannot be scheduled in the broadcast. Assuming so, solving D2B then consists in scheduling the transmission of some nodes in X such that every node in Y gets informed at some slot. A solution to D2B is a set of subsets of X , $\{X_1, \dots, X_k\}$, k being the number of slots the solution uses.

It is clear that X can be fully informed in one slot when x_e acts as the sole transmitter, the bounds we prove in the following do not count this trivial step.

II. UNIT DISK GRAPHS

In this section, D_u denotes the unit disk centred at the origin. We can assume without loss of generality that the source node x_e is at the origin. When a disk is referred to as D_i for some label i , then C_i will be the corresponding circle *i.e.* the border of D_i .




Fig. 1. Three disks (D_1, D_2, D_3). The gray areas (partitioning the union of the disks) are the target areas induced by these disks. Each target area contains its signature

Let us assume that the nodes lie on a plane and that each source node has a communication disk - its *source disk*. Target nodes must be covered using these disks. An instance of D2B is now a couple (\mathcal{D}, Y) , \mathcal{D} being a set of source disks and Y a set of targets. Let $\mathcal{D} = \{D_1, \dots, D_m\}$ be a set of labelled source disks. Given $\mathcal{D}' \subset \mathcal{D}$, the proper intersection of all the disks in \mathcal{D}' is $\cap_{x \in \mathcal{D}'} x \setminus \bigcup_{y \in \mathcal{D} \setminus \mathcal{D}'} y$. If said proper intersection is non-empty, we say that \mathcal{D}' induces a *target area*. We define the *signature* of a target area as the set containing all the labels of the disks inducing this area. In Figure 1, three disks D_1, D_2, D_3 intersect each other. The gray areas are target areas. Each area has a signature. For instance the target area of signature (1, 2) is the intersection between D_1 and D_2 minus D_3 . Note that the target areas partition the union of all the disks.

A. Unit disk graphs requiring a logarithmic number of slots

Theorem II.1. *Let $m \in \mathbb{N}$. Let $\mathcal{D} = (D_1, \dots, D_m)$ be a set of distinct source disks which centres lie on the same radius of the unit disk. There exists a set Y of points such that solving D2B over (\mathcal{D}, Y) requires $1 + \lfloor \log_2(|Y|) \rfloor$ slots.*

Proof. We can assume that \mathcal{D} is ordered by the disks' centres' distance to the origin. For all D_i in \mathcal{D} , x_i is the centre of D_i . The *signature matrix* \mathcal{M} given in Figure 2 contains the signatures of all the non-empty target areas induced by \mathcal{D} .

Now we build an instance of D2B by placing a target node in each target area. Denote n the number of target nodes ($n = \frac{m(m+1)}{2}$) and t_k the number of slots needed to inform k targets for all $k \in \{1, \dots, n\}$. Clearly $t_1 = 1$.

For some i , let us remove from \mathcal{M} all the target areas which signatures contain i (see the gray part in Figure 3). Note that the remaining areas to cover may form two disjoint instances—*residual instances*—of the problem on smaller graphs. One

(1)		
(1, 2)	(2)	
\vdots	\vdots	\ddots
(1, ..., m)	(2, ..., m)	... (m)

Fig. 2. Signature matrix \mathcal{M} . Each entry is the signature of a target area induced by \mathcal{D} . Note that an entry (i, j) corresponds to the signature (i, \dots, j) .

instance uses disks of indexes 1 to $i - 1$ and the other uses disks of indexes $i + 1$ to m . Note that if $i = 1$ or $i = m$, there is only one residual instance.

(1)			
\vdots		\ddots	
(1, ..., i)	...	(i)	
\vdots	\vdots	\vdots	\ddots
(1, ..., m)	...	(i, ..., m)	... (m)

Fig. 3. When node i is the sole emitter at a given slot, the gray rectangle is covered. The remaining uncovered parts of the matrix are disjoint and can be processed in parallel without risking to interfere at an uncovered area.

Observe that the signature matrices of the residual instances have the same form as the original instance. Whichever i is chosen to emit at first slot, at least one of the residual matrices has size at least $\lceil \frac{m-1}{2} \rceil$. Since the two residual instances are disjoint, they can be processed in parallel. The following recurrence on t_n is now verified.

$$t_n \geq t_{\lceil \frac{n-1}{2} \rceil} + 1 \quad (1)$$

$$\geq 1 + \lfloor \log_2(n) \rfloor \quad (2)$$

Now let us prove that the recurrence equation can be tight. The two following properties are verified.

- 1) There exists a target area of signature $(1, \dots, m)$. In any solution (X_1, \dots, X_k) there is thus a X_i of cardinality 1.
- 2) Given a solution (X_1, \dots, X_k) , permuting two sets in the solution does not change it.

Taking these two facts under account, we can assume that in any solution the first slot is occupied by a single source node *i.e.* $|X_1| = 1$.

Now if we choose the central node $i = \lceil \frac{m-1}{2} \rceil$ to transmit at each recursion step, then the solution produced is necessarily optimal. Indeed, in each residual instance the target area with the biggest signature has to be covered. By choosing the central node as sole transmitter for the first step, we ensure that said area is covered, and that the residual instances have roughly the same size. We then have $t_n = 1 + \lfloor \log_2(n) \rfloor$. \square

B. A logarithmic number of slots always suffices

We call *angular region* any part of the plane delimited by two half-lines with a common extremity being the origin. An

angular region is of angle α if the angle formed by the half-lines delimiting it is α .

Lemma II.1. *Let \mathcal{D} be a set of source disks, and Z be an angular region of angle $\frac{\pi}{2}$. For all D_i, D_j in \mathcal{D} , C_i and C_j intersect at most once in $Z \setminus D_u$.*




Fig. 4. Circles of centres A and B intersecting at points P and Q. Using the angle at center theorem on γ and β , we show the angle α to be at least $\frac{\pi}{2}$.

Proof. Let D_1 and D_2 be two disks of centres A and B, containing the origin O. Let P and Q be the intersection points between C_1 and C_2 . Let C be the intersection between the line (OA) and D_1 . We can assume without loss of generality that A and B lie in a half of the unit disk centred in O. We can also assume that A is closer to O than B is. Now let α be the angle (\vec{OP}, \vec{OQ}) , β the angle (\vec{AQ}, \vec{AP}) and γ the angle (\vec{CP}, \vec{CQ}) (see Figure 4). Since β and γ intercept the same arc of C_1 , the angle at center theorem states that $\gamma = \frac{\beta}{2}$. By construction, since the triangle OPQ is contained in the triangle CPQ, we have $\alpha \geq \gamma$. Then, $\beta > \pi$ because if $\beta = \pi$ then PQ is a diameter of both C_1 and C_2 meaning that these two circles are equal. We thus have $\alpha > \frac{\pi}{2}$ and therefore C_1 and C_2 cannot intersect more than once in $Z \setminus D_u$. \square

Given a set of source disks \mathcal{D} and a disk $D \in \mathcal{D}$ labelled i we say that D induces a *proper area* if there exists a target area induced by \mathcal{D} of signature (i) . Now consider an angular region Z of angle $\frac{\pi}{2}$, we define $\mathcal{D}_{|Z}$ as the set of all disks in \mathcal{D} intersecting Z non-emptily. We then remove from $\mathcal{D}_{|Z}$ all the disks not inducing a proper area in Z . Now order the disks in $\mathcal{D}_{|Z}$ according to the angle to the proper area they induce (see Figure 5). After this ordering $\mathcal{D}_{|Z} = (D_1, \dots, D_m)$ if S_i and S_j are proper areas induced by D_i and D_j with $i < j$ then $\forall P \in S_i, \forall Q \in S_j, (\vec{OP}, \vec{OQ}) \in]0, 2\pi[$.

We will next assume that $\mathcal{D}_{|Z}$ does not contain three circles intersecting in one point in Z . It is not hard to work around it, but simplifies the proof of the following lemma.




Fig. 5. Ordering disks w.r.t the proper areas in a quadrant. D_1 's proper area appears first then D_2 's then D_3 's.

Lemma II.2. *Let S be a target area induced by $\mathcal{D}_{|Z}$ such that $S \cap Z \neq \emptyset$ then the signature of S is of the form $\sigma(S) = (i, i+1, \dots, i+p)$ for some i, p .*

Proof. Let us assume it is not the case, thus there is a gap in the signature of S : $\sigma(S) = (i, \dots, k, k+j, \dots, i+p)$ with $j > 1$. Consider the set of source disks $D' = \{D_k, D_{k+1}, D_{k+j}\}$ which we can renumber $D' = \{D_1, D_2, D_3\}$ to clarify things. Now there exists an area of signature $(1, 3)$. Using Lemma II.1 we know that there is a unique intersection point P between C_1 and C_2 in Z . Let Z^- (resp. Z^+) be the part of Z located counterclockwise before (resp. after) the line (OP). Let $D_i^+ := D_i \cap Z^+$ and $D_i^- := D_i \cap Z^-$ for $i = 1, 2, 3$. Since C_1 intersects C_2 only once, we know that $D_2^- \subset D_1^-$ and $D_1^+ \subset D_2^+$. As a consequence, the area of signature $(1, 3)$ cannot lie in Z^+ , it has to lie in Z^- .

- 1) Assume that C_1 and C_3 do not intersect in Z^- , then $D_3^- \subsetneq D_1^-$ because otherwise the proper area induced by D_3 appears before the proper area induced by D_1 which is absurd. Note that C_2 and C_3 have to intersect in Z^+ because if not
 - a) either $D_2^+ \subsetneq D_3^+$ and then D_2 does not induce a proper area in Z since $D_2^- \subsetneq D_1^-$ which is absurd;
 - b) or $D_3^+ \subsetneq D_2^+$ and then the proper area induced by D_3 appears in Z^- thus before the one induced by D_2 which lies in Z^+ . That is absurd as well.

Thus, C_2 intersects C_3 in Z^+ and we know they do not intersect in Z^- using Lemma II.1. Under these assumptions $D_3^- \subsetneq D_2^-$ because otherwise $D_2^- \subsetneq D_3^-$ and $P \in D_3^-$, thus $D_3^- \cap D_1^- \neq \emptyset$ or $D_3^- \subsetneq D_1^-$, in both cases it is absurd. Thus, we have $D_3^- \subsetneq D_2^-$, and in that case the area of signature $(1, 3)$ cannot exist, which is absurd since we assumed it did exist.

- 2) Finally, assume that C_1 and C_3 do intersect in Z^- . Then C_3 intersects C_1 before P (because we assume that three circles never intersect in one point) and the proper area induced by C_3 appears before the proper area induced by C_2 , absurd again.

Both cases lead to contradictions, the area of signature (1,3) cannot exist. The signature of S is thus of the form $(i, i+1, \dots, i+p)$, thus the lemma. \square

Theorem II.2. *Let \mathcal{D} be a set of source disks, then all the target areas induced by \mathcal{D} can be covered using $\mathcal{O}(\log_2(n))$ slots.*

Proof. Partition the plane in four angular sectors of angle $\frac{\pi}{2}$ Z_1, Z_2, Z_3, Z_4 . First, consider Z_1 then using Lemma II.2 any target area appearing in Z_1 has signature of form $(i, i+1, \dots, i+j)$ for some i, j (after removing the disks not inducing proper areas and reordering $D|_{Z_1}$). We can thus build a signature matrix for $D|_{Z_1}$ having the same form as the matrices considered in subsection II-A. Since we showed that instances with such signature matrices can be dealt with using $\mathcal{O}(\log_2(n))$ slots, we can cover any target area appearing in Z_1 using $\mathcal{O}(\log_2(n))$ slots. We can then do the same for Z_2, Z_3 and Z_4 sequentially and cover any target area induced by D in $\mathcal{O}(\log_2(n))$ slots. \square

III. CIRCULAR ARC NEIGHBOURHOODS

We now suppose that the set Y of target nodes can be placed on a circle, and the neighbourhoods of X can be represented by proper arcs on that circle - that is a target node y can be reached by a source node x if y lies on the arc representing x 's neighbourhood. For instance, on Figure 6, x_1 has neighbourhood $\{y_1, y_2, y_3, y_4\}$.

X and Y can be arbitrarily ordered as follows. Pick one of the source nodes x to be the first one in the X order ($x = x_1$), then number $\{y_1, \dots, y_k\}$ its neighbours. Now number x_2 the node which arc starts after x_1 , and its neighbours can be numbered following x_1 's neighbours. One can do so until every node in X and every node in Y have been ordered (see Figure 6).




Fig. 6. Dots are target nodes, arcs represent source nodes

Theorem III.1. *3 slots are always enough to solve D2B in the circular arc neighbourhoods case.*

Proof. The following procedure produces a solution to D2B using at most 3 slots.

- 1) order X and Y as previously explained
- 2) let y be the first target in Y not yet covered

- 3) add to the solution the arc covering y ending the furthest possible after y
- 4) go to 2 while Y is not entirely covered

Denote x_1, \dots, x_k the nodes selected during the procedure. Then $N(x_i) \cap N(x_{i+2}) = \emptyset$ for all $i = \{1, \dots, k-3\}$. Indeed, if $N(x_i) \cap N(x_{i+2}) \neq \emptyset$ then consider the target y that forced the algorithm to pick x_{i+1} , and the target y' that made it pick x_{i+2} . We know that $y \notin N(x_i)$, and that the arc x_{i+1} ends after the arc x_{i+2} (otherwise x_{i+2} would have been picked instead since $N(x_i) \cap N(x_{i+2}) \neq \emptyset$), but then y' being the first target after x_{i+2} would be contained in $N(x_{i+1})$ and x_{i+2} would not have been picked, that is absurd thus $N(x_i) \cap N(x_{i+2}) = \emptyset$.

Possibly $N(x_{k-2}) \cap N(x_k) \neq \emptyset$. Now, let $X_1 = \{x_i, i < k \text{ is odd}\}$, $X_2 = \{x_i, i < k \text{ is even}\}$ and $X_3 = \{x_k\}$. $\{X_1, X_2, X_3\}$ is thus a broadcast using three slots. \square

IV. CONCLUSION

We studied the distance-2 broadcast problem in two specific cases. In general graphs, there is a tight $\mathcal{O}(\log(n)^2)$ upper bound on the number of slots needed to solve the problem. We improved this bound in unit disk graphs to $\mathcal{O}(\log(n))$. In the case the neighbourhoods can be represented as circular intervals, we proved the bound to be even lower: $\mathcal{O}(1)$.

In the future we would like to extend our research to other classes of graph, or more accurate communication models. Another interesting problem is the maximal cover in a fixed number of slots.

REFERENCES

- [1] R. Bar-Yehuda, O. Goldreich, and A. Itai, "On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization," *Journal of Computer and System Sciences*, vol. 45, no. 1, pp. 104–126, 1992, ISSN: 0022-0000. DOI: 10.1016/0022-0000(92)90042-H.
- [2] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg, "A lower bound for radio broadcast," *Journal of Computer and System Sciences*, vol. 43, no. 2, pp. 290–298, 1991, ISSN: 0022-0000. DOI: 10.1016/0022-0000(91)90015-W.
- [3] O. Cogis, B. Darties, S. Durand, J. König, and G. Simonet, "The mv-decomposition: Definition and application to the distance-2 broadcast problem in multi-hops radio networks," in *Fifth IFIP-TCS, 2008*, 2008, pp. 115–126. DOI: 10.1007/978-0-387-09680-3_8.
- [4] D. Peleg, "Time-efficient broadcasting in radio networks: A review," in *ICDCIT 2007, Bangalore, India, December 17-20. Proceedings*, T. Janowski and H. Mohanty, Eds. Springer Berlin Heidelberg, 2007, pp. 1–18, ISBN: 978-3-540-77115-9. DOI: 10.1007/978-3-540-77115-9_1.
- [5] T. Chan and E. Grant, "Exact algorithms and apx-hardness results for geometric packing and covering problems," *Computational Geometry*, vol. 47, no. 2, pp. 112–124, 2014, Special Issue: 23rd Canadian Conference on Computational Geometry (CCCG11), ISSN: 0925-7721. DOI: 10.1016/j.comgeo.2012.04.001.