

Theoretical analysis of beaconless geocast protocols in 1D

Citation for published version (APA):

Gudmundssons, J., Kostitsyna, I., Löffler, M., Müller, T., Sacristán, V., & Silveira, R. I. (2018). Theoretical analysis of beaconless geocast protocols in 1D. In M. Nebel, & S. Wagner (Eds.), *2018 Proceedings of the 15th Workshop on Analytic Algorithmics and Combinatorics, ANALCO 2018* (pp. 62-76)
<https://doi.org/10.1137/1.9781611975062.7>

DOI:

[10.1137/1.9781611975062.7](https://doi.org/10.1137/1.9781611975062.7)

Document status and date:

Published: 01/01/2018

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Theoretical analysis of beaconless geocast protocols in 1D

Joachim Gudmundsson* Irina Kostitsyna[†] Maarten Löffler[‡] Tobias Müller[§]
Vera Sacristán[¶] Rodrigo I. Silveira[¶]

Abstract

Beaconless geocast protocols are routing protocols used to send messages in mobile ad-hoc wireless networks, in which the only information available to each node is its own location. Messages get routed in a distributed manner: each node uses local decision rules based on the message source and destination, and its own location. In this paper we analyze six different beaconless geocast protocols, focusing on two relevant 1D scenarios. The selection of protocols reflects the most relevant types of protocols proposed in the literature, including those evaluated in previous computer simulations. We present a formal and structured analysis of the maximum number of messages that a node can receive, for each protocol, in each of the two scenarios. This is a measure of the network load incurred by each protocol. Our analysis, that for some of the protocols requires an involved probabilistic analysis, confirms behaviors that had been observed only through simulations before.

1 Introduction

In mobile ad-hoc wireless networks there is no fixed infrastructure or global knowledge about the network topology. Nodes communicate on a peer-to-peer basis, using only local information. Thus messages between nodes that are not within range of each other must be sent through other nodes acting as relay stations. An important particular case of ad-hoc wireless networks are wireless sensor networks, in which a (usually large) number of autonomous sensor nodes collaborate to collectively gather information about a certain area.

Nodes are typically mobile devices whose location and availability may change frequently, resulting in a highly dynamic environment in which routing must be done on-the-fly. Typically, messages are not sent to a particular network address, but to some or all nodes within a geographic region. This is called *geocasting*. The main pieces of information used to send a message are the location of the source node, and that of the destination region (also referred to as *geocast region*), which is usually included in the actual message.¹

*School of Information Technologies, University of Sydney, joachim.gudmundsson@gmail.com

[†]Mathematics and Computer Science Dept., TU Eindhoven, i.kostitsyna@tue.nl, supported in part by the Netherlands Organisation for Scientific Research (NWO) under project no. 639.023.208 and by F.R.S.-FNRS.

[‡]Dept. of Information and Computing Science, Universiteit Utrecht, m.loffler@uu.nl, supported by the Netherlands Organisation for Scientific Research (NWO) under project no. 639.021.123 and 614.001.504.

[§]Johann Bernoulli Institute, Groningen University, tobias.muller@rug.nl, partially supported by the Netherlands Organisation for Scientific Research (NWO) under project no. 639.032.529 and 612.001.409

[¶]Dept. de Matemàtiques, Universitat Politècnica de Catalunya, {vera.sacristan,rodrigo.silveira}@upc.edu, partially supported by projects MINECO MTM2015-63791-R and Gen. Cat. DGR 2014SGR46. R.S. also supported by MINECO through the Ramón y Cajal program.

¹Some works use the term ‘packet’ to denote the indivisible unit of information sent between the nodes. In this paper we use the term ‘message’ instead as we are interested in counting the number of transmissions and not the higher level aspects of protocols.

Many geocast protocols have been proposed. In general, existing protocols can be divided into two groups: those that assume that each node also knows the location of its 1-hop neighbors (i.e., all nodes within range) and those that don't. In practice, the locations of neighbors can be obtained by regularly exchanging *beacon* messages in the neighborhood. Beacons imply a significant message overhead, which prevents these methods from scaling even to medium-size networks [3]: the problem is that in dense environments the number of messages *received* by each individual node, and thus the workload to decide whether and how to react to those messages, becomes prohibitive. For this reason, in this paper we are interested in the second group, the so-called *beaconless* geocast protocols.

Probably the most straightforward beaconless geocast protocol is simple flooding: each message is broadcasted to all neighbors, who in turn broadcast it to all their neighbors, and so on. Even though it is effective, the resulting message overhead is clearly unaffordable. From there on, there have been many improvements proposed. The ultimate goal is to reduce the message overhead as much as possible while still guaranteeing delivery. In the last few decades, many different geocast protocols have been proposed. Several of the most important protocols are the subject of this paper, and will be described in detail in the next section. For a thorough review of all existing beaconless geocast protocols we refer the reader to the surveys by Maihöfer [5] and Rührup [6].

Given the importance of geocast protocols and the many options available, there have been a few comparative studies that assessed the efficiency and efficacy of different methods under different scenarios, using low-level computer simulations. Maihöfer [5] presents simulations for four geocast protocols (flooding, two variants of flooding restricted to a forwarding zone, and a greedy-based protocol) for sparse networks with 100–1000 nodes that move randomly within a square region, and use random circular geocast regions as destinations. The results are analyzed in terms of total number of messages transmitted (network load) and success rate. As expected, the experiments show that flooding has the highest success rate, but does not scale well, while the other methods are much better in terms of network load, but suffer from lower success rate. Hall [1] evaluated four methods (M heuristic, T heuristic, CD and CD-P, and several combinations of them) in 14 different scenarios, many of them based on realistic training applications. The parameters studied were success rate and average latency. The simulation was done using a high fidelity simulator capable of modeling realistic MAC and queuing behavior, allowing to reflect the effects of high network load in the different protocols. The main conclusion in [1] is that CD-P performs best in terms of both success rate and latency. Interestingly, the experiments also show that for restricted flooding heuristics, an increase in the redundancy parameters does not always lead to higher success rate due to more collisions and medium contention.

In contrast with previous comparisons, in this paper we are interested in analyzing the behavior of beaconless geocast protocols from a theoretical perspective. To that end, we present a structured overview of six different protocols that represent the main existing protocols in the literature, and identify important quality criteria to analyze them mathematically.

The protocols analyzed are simple flooding, M heuristic [2], T heuristic [2], CD [1], CD-P [1], and delay-based protocols (that include, as particular cases, protocols like BLR [3], GeRaF [8] and GeDiR [7]). This selection of protocols reflects the main different types of beaconless geocast protocols, and includes those evaluated in previous computer simulations [1, 5].²

Several criteria can be taken into account when comparing the behavior of different protocols. The *success rate* measures the fraction of sent messages that actually reach the target. For those that arrive, the *hop count* indicates how many steps (forwards) are needed. In this paper we only

²Note that, in practice, protocols are often combined: a message is forwarded if at least one protocol requires it to do so. This increases success rates, at the cost of a higher total number of messages in the system.




Figure 1: Even though the 1D scenario serves as a first step towards the general question, it is already relevant in, e.g., narrow passages, where the local geometry is essentially 1-dimensional.

Protocol	Unbounded reach	Bounded reach (r)
Lower bound	$\Omega(k)$	$\Omega(k)$
Flooding	nk	$O(rk)$
M-heuristic	Mk	$\min\{Mk, 2rk\}$
T-heuristic	$\lceil \frac{n}{2T} \rceil k, \lceil \frac{n}{T} \rceil k$	$O(\frac{rk}{T})$
CD	$\begin{cases} \Theta(k^2 \log(\lceil n/k \rceil + 1)), & \text{if } k \leq n \\ \Theta(nk), & \text{if } k > n \end{cases}$	$O(k^{3/2})$
CD-P	$\Theta(k \log n)$	$\Theta(k)$
Delay-based	$\min\{2^k, n(1 + k - \log n)\}$	$O(\frac{nk}{r})$

Figure 2: Summary of the results for **RecMess** for the two settings and different protocols studied. The results for CD and CD-P hold with high probability (when n and/or k tend to infinity).

focus on what we consider to be the most significant measure within this context: the maximum number of messages that a node receives (**RecMess**). This parameter measures the work or energy consumption for a node, as well as the overall network load and therefore, its congestion. We note that network load is directly related to success rate, thus indirectly this aspect is also being considered, as done in previous comparisons [1, 5]; we do consider the theoretical success probability of the protocols (in case of no collisions) separately. We also note that, in most situations, **RecMess** is larger than the number of sent messages, because for intermediate nodes, the sending of a message occurs only as a consequence of receiving one before.

The behavior of a geocast protocol, in general, must be analyzed in the context of a particular geometric setting (i.e., a certain configuration of nodes and radio obstacles). In this paper, we focus on two fundamental geometric scenarios in 1D: unbounded range and bounded range. Even though it is clear that the full complexity of these protocols can only be appreciated in two dimensions, we show that the 1D scenarios considered, despite their apparent simplicity, already pose interesting challenges, and expose many of the essential differences between the protocols studied. Moreover, understanding 1D situations is useful for many 2D scenarios as well, in which there are local situations that behave essentially as one-dimensional (see Fig. 1).

For each scenario we analyze worst- and expected-case performance for each protocol. The results obtained corroborate many of the findings previously obtained by simulations, and provide new insights into the difficulties of the 2-dimensional case where, in addition, escaping from local optima is necessary and requires combining different techniques. In fact, we note that 1D settings have been used before to understand the behavior of the CD and CD-P heuristics [1]. A summary of our results is presented in Figure 2.

2 Geocast protocols

We begin by describing the basic geocast framework operating at each node, following [1]. Each message is assumed to contain a unique ID, the location of the sender, and the destination (geocast region). When a node receives a message, it checks if it has already received a message with the same ID. If not, it creates a new record for the ID, and enqueues the message for potential later retransmission (possibly after storing some extra information about the message). When a message reaches the head of the transmission queue and is ready to be transmitted, a *heuristic check* is performed. If passed, the message is transmitted, otherwise discarded. The different protocols mostly differ on the heuristic check performed. In general terms, this check is a combination of local decision rules. Often, one of these rules is a location predicate to control the region where each message must travel (e.g., to guarantee that messages are only forwarded

within a certain zone that contains the geocast region).

In the remaining of the paper we focus on the following beaconless geocast protocol, which can be categorized as: simple flooding, restricted flooding, distance-based, and delay-based.

Simple flooding. The simple flooding protocol works as follows. When a node receives a message, first it checks if it has been broadcasted before. If not, then the message is broadcasted, and its ID stored in order to make sure it will not be broadcasted again.

This strategy is simple and robust, but it is non-scalable, as it produces an excessive and unnecessary network load. In the following, we describe several heuristics intended to reduce such flood load. Nevertheless, it is interesting to consider simple flooding not only for comparison purposes, but also because it is used as a building block in other protocols (e.g.,[4]).

2.1 Restricted flooding

In order to reduce the number of unnecessary transmissions of the same message, one can limit retransmissions in several ways [1]. The following two heuristics apply two different approaches for this: a direct limit in the number of retransmissions, or an implicit limit, by only making “far away” nodes retransmit already heard messages.

M heuristic [2]. The MinTrans (M) heuristic explicitly controls redundancy through a parameter M : A node broadcasts a received message if and only if the number of transmissions received for that ID is less than M . The redundant propagation allowed by the parameter M helps against problems such as message collisions and getting out from local optima.

T heuristic [2]. The Threshold (T) heuristic uses location information for spreading the geocast propagation outward: A node retransmits a received message if and only if the closest among all transmitters of messages with the same ID is at least a distance T from it.

2.2 Distance-based heuristics.

The previous heuristics are likely to have delivery failures in the presence of obstacles. The following protocols were designed to help solving this problem.

CD heuristic [1]. The Center-Distance (CD) heuristic relies on proximity: A node retransmits a received message if and only if its distance to the center of the geocast region is less than that of all originators of transmissions received for the message ID. This heuristic reduces some of the scalability problems of a previous method, Ko and Vaidya’s “Scheme 2” [4], which considered only the distance of the first transmitter heard of the same message instead of that of all of them.

CD-P heuristic [1]. This protocol uses priority queues in order to further reduce the scalability problems of the CD heuristic. It works as follows: Each time the node can transmit, it transmits any message that has not been transmitted at all yet (if any) or it retransmits, among all heard messages, the one whose transmission would give the largest reduction in distance to the center of the geocast region.

2.3 Delay-based heuristics

Some strategies to further reduce redundancy combine distances with retransmission delay.

BLR heuristic [3]. In the Beacon-Less Routing (BLR) protocol, each node determines when to retransmit a received message based on a dynamic forwarding delay function valued in $[0, \text{MD}]$, for MD a constant representing the maximum delay. The node retransmits the package after such delay, unless some other node does it before, in which case the retransmission is canceled. Three delay functions have been suggested in [3], based on the following parameters: r (transmission range), p (progress towards destination of the orthogonal projection of the current node onto the line connecting the previous node to the destination), and d (distance from current node to the source-destination line):

$$\begin{aligned} \text{delay}_1 &= \text{MD} \frac{r-p}{r}, \text{delay}_2 = \text{MD} \frac{p}{r}, \quad \text{and} \\ \text{delay}_3 &= \text{MD} \frac{e^{\sqrt{p^2+d^2}}}{e}. \end{aligned}$$

GeRaF heuristic [8]. Based on distance, the Geometric Random Forwarding (GeRaF) protocol divides logically the area around the destination of a message m into n_m areas $\mathcal{A}_1, \dots, \mathcal{A}_{n_m}$, where in \mathcal{A}_1 are all nodes closest to the destination, and so on. Once m is transmitted, up to n_m phases start, during which all nodes listen during a fixed amount of time. In the first phase, nodes in region \mathcal{A}_1 get to reply. If only one node replies, then that one will forward the message. If there are more, some collision resolution scheme must be used. If there is no reply, then it is the turn to reply for nodes in region \mathcal{A}_2 . This process continues until some node in the non-empty region closest to destination replies.

Greedy routing (beaconless version). Greedy routing consists in forwarding the message to the neighbor of the current node that is closest to the geocast region. Even though it does not guarantee delivery, greedy routing strategies are often used as building block of geocast protocols. For this reason we also consider greedy routing in our analysis. One example is Geographic Distance Routing (GeDiR) [7]. GeDir requires to know the position of all neighbors of a node. However, it can be made beaconless by using a delay function based on the following parameters: r (transmission range), d (distance from previous node to destinations), and x (distance from current node to destination).

$$\text{delay}_4 = \text{MD} \frac{x+r-d}{2r}.$$

This strategy tries to get out of local minima by sending the message to the best positioned neighbor, even if it is not closer to destination than the sender.

All protocols described in this section include a rule that states what to do if a node receives a message already in its queue. One option, like in BLR, is to always cancel the transmission of a message received twice. Another option also used in practice is to cancel only if the sender of the duplicate message is closer to the destination than the current node.

3 The 1D scenarios

In this paper we study two fundamental scenarios for a set of nodes in 1D, where the leftmost of $n+2$ nodes sends k messages to the rightmost node (i.e., the geocast region only contains the rightmost node). For simplicity, nodes are evenly spread at unit distance along the line. The two scenarios considered are nodes with unbounded range and nodes with bounded communication range (see Fig. 3). Each message contains the position of its last (re)sender and its destination. Each node stores all received messages in a queue that is managed in one way or another depending on the protocol used.

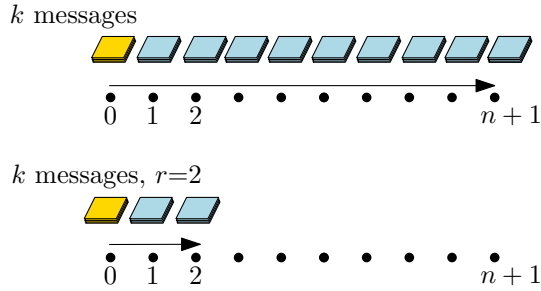


Figure 3: Illustration of the scenarios. Left: with unbounded reach, the k messages arrive immediately to all nodes, but that does not prevent intermediate nodes from forwarding the messages. Right: with range $r = 2$, the messages sent from node 0 only reach up to node 2, so forwards are necessary to reach the target, $n + 1$.

The n intermediate nodes form a dense bottleneck, a situation that can easily arise in practice (even for nodes in 2D). Once the transmissions start, collisions may happen. In order to cope with this problem, we assume fair medium access, i.e., the transmission is done in rounds, and in each round each node that has some message to transmit has the same probability to transmit it.

For each of the two settings and each protocol studied, we analyze the maximum number of messages that a node can receive, denoted **RecMess**.

4 Analysis: unbounded reach

In the *unbounded reach scenario*, all nodes are within the communication range of each other. This setting recreates a rather frequent situation where many messages must go through a high-density area. In this scenario we consider n nodes on a line with integer coordinates $\{1, \dots, n\}$, all within range of each other. We assume that, initially, all the nodes have the same k messages in their queues that were all received from the same node at coordinate 0, and are all directed to a node at coordinate $n + 1$. An example of the behavior of all protocols studied on an example sequence of random probes is shown in Figure 4.

4.1 Simple flooding.

Under this protocol, all the nodes will receive and retransmit all the messages: **RecMess** = nk .

4.2 M heuristic.

In the unbounded reach scenario this protocol guarantees that there are never more than M copies of each messages transmitted. Therefore, every node can receive every message at most M times, leading to **RecMess** = Mk .

4.3 T heuristic.

Due to the unit-distance distribution of the nodes and assuming $T \in \mathbb{N}$, when a node transmits a message, then this message is deleted from the queues of its T left and its T right neighbors. Thus: $\lceil \frac{n}{2T} \rceil k \leq \mathbf{RecMess} \leq \lceil \frac{n}{T} \rceil k$.

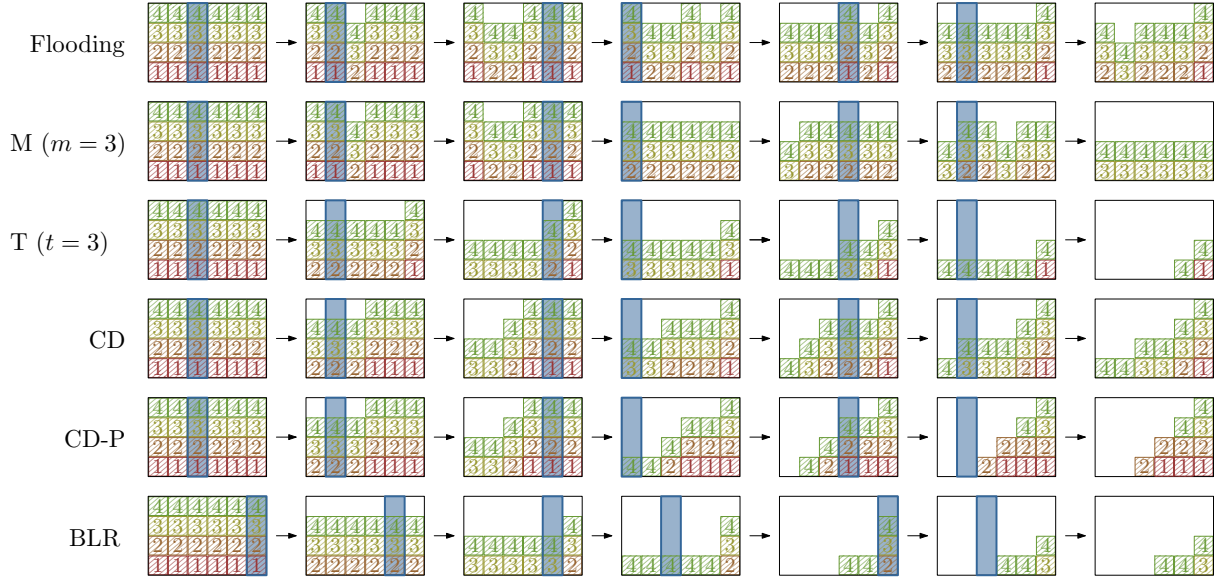


Figure 4: The different protocols run on an example data set with $k = 4$ and $n = 6$. Each row shows the first 7 steps of a different protocol. At each step, a matrix is shown. Columns represent the queues of the nodes, which all start containing the same four messages in the same order. A shaded column indicates the node that is about to forward a message in that step. The same random order is chosen everywhere except for BLR.

4.4 CD heuristic.

The number of messages received in the CD heuristic depends heavily on the order in which the nodes are chosen. In the best case, the last node is chosen consecutively k times, and every node receives only k messages, since all the nodes delete the message as soon as they hear it from a node farther ahead. In the worst case, the nodes are chosen from left to right, and all nodes receive all messages n times, as bad as in simple flooding. So, $k \leq \mathbf{RecMess} \leq nk$.

However, assuming fair medium access, which dictates a uniform distribution for the chosen node at each time step on the set of all nodes with non-empty queues, we can use probabilistic analysis to make a more precise prediction of the number of received messages in practice. We can show that if k is significantly smaller than n , the number of messages depends only logarithmically on n , at the cost of a quadratic dependence on k .

Theorem 1. *In the unbounded reach scenario under the CD heuristic, with high probability,*

$$\mathbf{RecMess} \in \begin{cases} \Theta(k^2 \log(\lceil n/k \rceil + 1)), & \text{if } k \leq n, \\ \Theta(kn), & \text{if } k > n. \end{cases}$$

To prove the theorem, we model the problem as an initially empty $n \times k$ grid of squares (n columns of height k), as shown in Fig. 4. At each time step a random column is selected among those columns that are not yet completely full. Having chosen a column, we fill the highest empty square in that column, together with all the squares to its left.³

We are interested in the (random) time T when all squares have become filled. Note that this is the same as asking for the time when the n -th column has been hit k times. What makes it tricky to analyze is that while initially we have chance $1/n$ of hitting the n -th column, this

³Clearly, this formulation is equivalent to starting with a full grid, and removing squares (messages) when a random non-empty column gets chosen.

probability will increase as time goes on and more columns get filled. We consider the case where the number of columns n or the number of rows k goes to infinity (or both).

In order to prove Theorem 1, we prove the following, more specific, result.

Theorem 2. *There exist constants $0 < c_L < c_U < \infty$ such that:*

1. *If $k = k(n) \leq n/1000$ then*

$$\mathbb{P}\left(c_L \cdot k^2 \log\left(\frac{n}{k}\right) \leq T \leq c_U \cdot k^2 \log\left(\frac{n}{k}\right)\right) \rightarrow 1,$$

as $n \rightarrow \infty$;

2. *If $k = k(n) \geq n/1000$ then*

$$\mathbb{P}(c_L \cdot kn \leq T \leq c_U \cdot kn) \rightarrow 1,$$

as $n \rightarrow \infty$.

3. *If n is constant then*

$$\mathbb{P}(c_L \cdot kn \leq T \leq c_U \cdot kn) \rightarrow 1,$$

as $k \rightarrow \infty$.

The random variable S . We will first consider the time S until at least one column becomes full. Of course $S \leq T$. What simplifies the analysis of S as opposed to T is that until time S in each time step we are choosing a column uniformly at random from all n columns.

Lemma 3. *For all $k \leq n$ and all t we have that:*

$$\mathbb{P}(S < t) \leq \left(\frac{2e^2 t}{k^2}\right)^k.$$

Proof. Let X_1, X_2, \dots denote the indices of the columns taken in each time step. Since we are only interested in S we can and will just pretend that nothing changes in the way we choose columns after one or more have filled up. In other words, from now on in this proof we suppose we are given an infinite sequence X_1, X_2, \dots of random variables (column indices) that are chosen independently and uniformly at random from $\{1, \dots, n\}$.

We now make the following crucial observation : There is a full column at time t if and only if there is a sequence $t_1 < \dots < t_k$ of length k such that $t_k \leq t$ and $X_{t_1} \geq X_{t_2} \geq \dots \geq X_{t_k}$. That is, the sequence X_1, \dots, X_t must contain a non-increasing subsequence of length k .⁴ To complete the proof, we give a very crude upper bound on the probability that there exists such a non-increasing subsequence in X_1, \dots, X_t . We will use the first moment method, i.e., we will just count—or rather bound—the expected number of all such subsequences.

First we choose the k time steps t_1, \dots, t_k on which a non-increasing subsequence is to occur. This can be done in

$$\binom{t}{k} \leq \left(\frac{et}{k}\right)^k,$$

ways (using a standard bound on binomial coefficients).

Assuming that we know the time steps t_1, \dots, t_k on which a non-increasing subsequence is to be formed, we wish to count the number of ways we can choose values for X_{t_1}, \dots, X_{t_k}

⁴This is loosely connected to the famous longest increasing subsequence problem for random permutations. Also there we need a permutation of length $\text{const} \times n^2$ before we can expect to find an increasing subsequence of length n .

to make them non-increasing. To specify a non-increasing sequence of length k with elements $\in \{1, \dots, n\}$ it is enough to specify how many times each number is hit. That is, we have $a_1, \dots, a_n \in \{0, \dots, k\}$ such that $a_1 + \dots + a_n = k$, where a_i is the number of times we use number i in the sequence. The sequence will start with a_n times the number n , followed by a_{n-1} times the number $n-1$, and so on. The number of ways of choosing values for X_{t_1}, \dots, X_{t_k} is thus:

$$\binom{k+n-1}{k} \leq \left(\frac{e(n+k)}{k}\right)^k \leq \left(\frac{2en}{k}\right)^k.$$

Next, we remark that for any given t_1, \dots, t_k and a_1, \dots, a_k the probability that X_{t_1}, \dots, X_{t_k} take the values specified (implicitly) by a_1, \dots, a_n simply equals n^{-k} .

Putting it all together we find that

$$\mathbb{P}(S < t) \leq (et/k)^k \cdot (2en/k)^k \cdot n^{-k} = (2e^2t/k^2)^k,$$

as required. \square

Lemma 4. *There exists a universal constant $c > 0$ such that for all $k \geq n$ and all $t \leq cnk$:*

$$\mathbb{P}(S < t) \leq \left(\frac{16t}{nk}\right)^{k/2}.$$

Proof. The argument is similar to that in the proof of Lemma 3, except that now we do not choose the times t_1, \dots, t_k . Instead we observe that if there exists a non-increasing subsequence that corresponds to a_1, \dots, a_n (as above, with $a_1 + \dots + a_n = k$, and $a_i \in \{0, \dots, k\}$ the number of times we use i in the sequence) then one can construct (possibly another) such sequence by first waiting until the first time t_n when we will have seen a_n occurrences of n , then waiting until the first time $t_{n-1} > t_n$ such that we have seen a_{n-1} occurrences of $n-1$ in the interval $(t_n, t_{n-1}]$ and so on. The sequence we obtain via this procedure will be the first to occur among all sequences corresponding to the same vector $\bar{a} = (a_1, \dots, a_n)$. For a given such vector \bar{a} we can therefore compute the probability that a subsequence corresponding to \bar{a} occurs before time t as follows:

$$\mathbb{P}\left(\begin{array}{c} \text{there exists a} \\ \text{subsequence} \\ \text{corresponding} \\ \text{to } \bar{a} \text{ in } X_1, \dots, X_t \end{array}\right) = \mathbb{P}(Y \leq t), \quad (1)$$

where $Y = Y_1 + \dots + Y_k$ is a sum of i.i.d. geometrically distributed random variables with common success probability $p = 1/n$. To see this, think of first “listening to the n -th column” until it has been selected a_n times, then “listening to the $(n-1)$ -st column”, and so on. At each time step the desired column is selected with probability $1/n$, independently of the previous history of the process.

In the full version we provide a generalized Chernoff bound on geometric series, which gives us

$$\begin{aligned} \mathbb{P}(Y \leq t) &\leq \exp[-(kt/\mathbb{E}(Y)) \cdot H(\mathbb{E}(Y)/t)] \\ &\leq \exp[-(t/n) \cdot H(nk/t)]. \end{aligned}$$

(Here $H(x) := x \ln x - x + 1$, and we have used $\mathbb{E}(Y) = nk$.) Now notice that $H(x) \geq \frac{1}{2} \cdot x \ln x$ for all x sufficiently large. That is, for all $x \geq x_0$ with x_0 a suitable constant—for instance, taking $x_0 := e^2$ is sufficient. Thus, we can choose c such that $(t/n) \cdot H(nk/t) \geq (k/2) \ln(nk/t)$ for all $t \leq cnk$.

The number of sequences \bar{a} is at most

$$\binom{n+k-1}{k} \leq \binom{2k}{k} \leq 4^k.$$

Putting these bounds together we find that, for $t \leq cnk$:

$$\mathbb{P}(S < t) \leq 4^k \left(\frac{t}{nk}\right)^{k/2} = \left(\frac{16t}{nk}\right)^{k/2},$$

as required. \square

Combining Lemmas 3 and 4 we see that:

Corollary 5. *There exists a universal constant $c > 0$ such that $\mathbb{P}(S < cnk) \rightarrow 0$ as $n \rightarrow \infty$ for any sequence $k = k(n)$ satisfying $k \geq n/1000$.*

Lemma 6. *There exists an absolute constant $C > 0$ such that for all n, k and all $t \geq Ck^2$ we have*

$$\mathbb{P}(S > t) \leq e^{-t}.$$

Proof Sketch. We can assume that without loss of generality $k < n/1000$ (otherwise, if $k \geq n/1000$ then $S \leq 1000k^2$ is deterministically true since in every time step a square gets filled and there are $nk \leq 1000k^2$ squares in total). First we wait until the first time one of the last $\lfloor n/k \rfloor$ columns are hit, then we wait until the next time after that the next $\lfloor n/k \rfloor$ columns are hit, and so on. The probability that at time t , we have not had k successes in the above process equals the probability that $X > t$ where X is the sum of k independent $\text{geo}(p)$ random variables with $p = \lfloor n/k \rfloor/n = \Theta(1/k)$. We can thus apply a generalized Chernoff bound (see full version for details). Note that $\mathbb{E}(X) = \Theta(k^2)$ and that $H(x) = x \ln x - x + 1 \geq x$ for all $x \geq x_0$ (where x_0 is of course a suitable constant, $x_0 := e^2$ will do). Thus, we can indeed choose $C > 0$ such that $t/\mathbb{E}(X) \geq x_0$ for all $t \geq Ck^2$, giving

$$\mathbb{P}(S > t) \leq \mathbb{P}(X > t) \leq e^{-\mathbb{E}(X) \cdot H(t/\mathbb{E}(X))} \leq e^{-t},$$

for all k, n, t satisfying the conditions of the lemma. \square

Upper bound. We will consider a modification of the process that will certainly take at least as long as the original version (and is easier to analyze). We distinguish different “phases” of the process. In the first phase, starting from the empty $k \times n$ grid, we throw columns until at least one column with index $\geq n/2$ is full. While doing this, we ignore the height constraint on columns with index $\leq n/2$. That is, columns with index $\leq n/2$ are allowed to have height $\geq k$ and we will still select columns uniformly from all n columns. As soon as a column with index $> n/2$ obtains height k , the first phase ends. We now throw away the left half of the grid, and empty the right half. So we now have an empty $(k/2) \times n$ grid. We now repeat the about modified process, i.e. randomly select columns and ignore the height constraints of the left half until we get some full column in the right half (which corresponds to the rightmost quarter of the original grid) and repeat. We keep going on until we start a phase starts when the number of remaining columns is $\leq 1000k$. In the last phase we simply carry out the process as usual, waiting until all the last column is filled. (Observe that there will be $\Theta(\log(n/k))$ phases in total.)

Proof of the upper bound. Let C as provided by Lemma 6. Notice that if we restrict attention to the right half of the columns that have not yet been filled at the start of the phase (i.e. in phase i these are columns with indices $n2^{-(i+1)} + 1, \dots, n$) then the number of times this set of columns is hit in a period of t time steps since the start of the phase behaves like a $\text{Bi}(t, \frac{1}{2})$ random variable (lets call it X) and the probability that the phase lasts longer than t is precisely

the probability that $S_{k, n2^{-(i+1)}} > X$. We see that for all $i \leq I$ (where $I := \lceil \log(n/1000k) \rceil$ is the least i such that $n2^{-i} \leq 1000k$) and all $t \geq 100Ck^2$:

$$\begin{aligned} \mathbb{P}(\text{phase } i \text{ takes longer than } t) &\leq \mathbb{P}(\text{Bi}(t, \frac{1}{2}) < t/100) \\ &\quad + \mathbb{P}(S_{k, n2^{-(i+1)}} > t/100) \quad (2) \\ &\leq e^{-(t/2) \cdot H(2/100)} + e^{-t/100} \\ &\leq 2e^{-t/100}, \end{aligned}$$

where we've applied the vanilla Chernoff bound and Lemma 6. and that $\frac{1}{2}H(2/100) > 1/100$.

We distinguish two cases. First, we suppose that $k \geq \sqrt{n}$. In that case (2) gives that

$$\mathbb{P}(\text{phase } i \text{ takes longer than } 100Ck^2) = \exp[-\Omega(n)].$$

Hence, we see that if we write $K := 100Ck^2I + 1000k^2$ ($= \Theta(k^2 \log(n/k))$) then assuming $k \geq \sqrt{n}$ we have

$$\begin{aligned} \mathbb{P}(\text{the process takes longer than } K \text{ to complete}) \\ \leq \log n \cdot e^{-\Omega(n)} = o(1). \end{aligned}$$

It remains to consider the case when $k \leq \sqrt{n}$. Note that in this case we have $I \rightarrow \infty$. Let L_i denote the duration of phase i . By Lemma 3 we have that

$$\mathbb{E}(L_i) \geq \frac{k^2}{4e^2} \cdot \mathbb{P}(L_i \geq \frac{k^2}{4e^2}) \geq \frac{k^2}{4e^2} \cdot \frac{1}{2} = \Omega(k^2). \quad (3)$$

On the other hand, applying Lemma 6 and a standard formula for the expectation of non-negative random variables we have that

$$\begin{aligned} \mathbb{E}(L_i) &= \int_0^\infty \mathbb{P}(L_i > t) dt \\ &= \int_0^\infty \mathbb{P}(L_i > t) dt \\ &\leq 100Ck^2 + 2 \int_{100Ck^2}^\infty e^{-t/100} dt \\ &\leq 100Ck^2 + 2 \int_0^\infty e^{-t/100} dt \\ &= O(k^2). \end{aligned} \quad (4)$$

(Here it is important to note that the constants hidden inside the Ω, O notation in (3) and (4) are universal. That is, the same constants work for all n, k in the range considered.) Similarly to (4) we also have that:

$$\begin{aligned} \mathbb{E}(L_i^2) &= \int_0^\infty \mathbb{P}(L_i^2 > t) dt \\ &= \int_0^\infty \mathbb{P}(L_i > \sqrt{t}) dt \\ &\leq 10^4 C^2 k^4 + 2 \int_{10^4 C^2 k^4}^\infty e^{-\sqrt{t}/10} dt \\ &= O(k^2), \end{aligned}$$

where the multiplicative constant hidden inside the $O(k^4)$ is universal.

Writing $L := \sum_{i=1}^I L_i$, we have that $\mathbb{E}(L) = \Theta(Ik^2)$ and $\text{Var } L = O(Ik^4)$. Thus, we can apply Chebyshev's inequality to see that

$$\mathbb{P}(L \leq \mathbb{E}(L)/2) \leq \frac{4 \text{Var } L}{(\mathbb{E}(L))^2} = O(1/I) = o(1),$$

Since $I = \Theta(\log(n/k))$ tends to infinity under the assumption that $k \leq \sqrt{n}$. As L is the time until the process completes, this completes the proof of the upper bound. \square

Lower bound. Again we will modify the process, this time so as to obtain a process that will terminate quicker (and is easier to analyse). We will choose constants $0 < \alpha, \beta, \gamma, \delta, \varepsilon < 1$ satisfying some additional demands (which include $\alpha \ll \beta, \beta < \gamma^2$) that will be specified as we go along with the proof.

- We divide time into “rounds” (periods) of length $t = \delta k^2$.
- We divide the columns into groups. Group i consists of $n_i := n\alpha^i$ consecutive columns. (Group 1 contains columns $1, \dots, n_1$, group two columns $n_1 + 1, \dots, n_1 + n_2$ and so on.) We shall be choosing α small enough so that in fact $\sum n_i < n$. That the sum is strictly less than n which will not be a problem for our argument.⁵
- At the start of each round, the columns of groups $1, \dots, i$ will already have been “filled”, while none of the columns in groups $i + 1, \dots$ will be full. (The first round is round 0.) At the end of the period we will fill one or more additional groups of columns completely, according to the following rules:
 - If at the start of the current round group i is the last filled group as above, and during the round in some group $i + j$ a non-increasing sequence of length $k\gamma^j$ gets created only taking into account what happened in the current period to the columns in that group, then we fill all columns up to and including that entire group, and move on to the next period. (We do not drop any more columns during the rest of round. That is, as soon as non-increasing sequence of length $k\gamma^j$ gets created in group $i + j$ the round ends.)
 - Otherwise, if by the end of a round no group obtained a non-increasing sequence of the desired length (among the selections in that period), then we fill up group i , and move on to the next period.
- The process terminates once a group with index $\geq I := \varepsilon \log(n/k)$ has been filled.

This process works because in a round when groups $1, \dots, i$ have been filled but group $i + 1$ has not, then in group $i + j$ the longest non-increasing sequence that could possibly exist in that group (now taking in to account all periods) is no longer than

$$\begin{aligned} \sum_{0 \leq k \leq i} k\gamma^{i+j-k} &= \sum_{\ell=j}^{i+j-1} k\gamma^\ell \\ &\leq k\gamma^j / (1 - \gamma). \end{aligned}$$

Thus, even if all the longest non-increasing sequences of the groups $i + 1, i + 2, \dots$ can be combined into one non-increasing sequence, then this sequence would still have length at most

$$\sum_{j=1}^{\infty} k\gamma^j / (1 - \gamma) = \frac{\gamma k}{(1 - \gamma)^2}.$$

This also gives that even if in the current round some group achieves the goal we have set for it (namely a sequence of length $k\gamma^j$ for group $i + j$) then the longest sequence will still be no longer than $k\gamma(1 + 1/(1 - \gamma)^2) < k$ (provided we chose the constants γ sufficiently small – which we can assume without loss of generality). In particular we never fill columns completely, except at the end of a round when we fill one or more entire groups of columns.

Next, we need a bound on the probability that group $i + j$ “fires” in the period when groups $1, \dots, i$ have been filled by group $i + 1$ has not.

⁵We are ignoring rounding for now. This can be settled by noting we can assume n, k and $\alpha, \beta, \gamma, \delta$ are all powers of two (for n, k this can be shown using the monotonicity of T in n, k – worsening the constant c_L a bit.)

Lemma 7. *Provided $\alpha, \beta, \gamma, \delta, \varepsilon$ are chosen appropriately, and i, j are such that $i + j \leq I$ (the number of groups) the following holds. Consider the situation in which we start a round in which groups $1, \dots, i$ have been filled and group $i + 1, i + 2, \dots$ not. Let $E_{i,j}$ be the event that group $i + j$ achieves a sequence of length $k\gamma^j$ during this period. Then*

$$\mathbb{P}(E_{i,j}) \leq 10^{-j}.$$

Moreover, if $k \geq \sqrt{n}$ then we even have

$$\mathbb{P}(E_{i,j}) = \exp[-\Omega(n^{1/4})],$$

for all $1 \leq i \leq I - 1$ and $1 \leq j \leq I - i$.

Proof. Note that if $E_{i,j}$ holds then one of the following two possibilities must have occurred. Either **(a)** group j has received $\geq k^2\beta^j$ columns (in the current period), or **(b)** it has received less than $k^2\beta^j$ columns but nonetheless a sequence of length $k\gamma^j$ was created. By Lemma 3 (applied with $k' = \max(k\gamma^j, 1), t' := k^2\beta^j, n' := n\alpha^j$) we see that

$$\begin{aligned} \mathbb{P}(\mathbf{(b)}) &\leq \left(\frac{2e^2 t'}{k'^2}\right)^{k'} \\ &= \left(\frac{2e^2 \beta^j}{\gamma^{2j}}\right)^{\max(k\gamma^j, 1)} \\ &\leq 2e^2 (\beta/\gamma^2)^j. \end{aligned} \tag{5}$$

To deal with the probability that **(a)** holds we distinguish two further possibilities: **(a-1)** we have $k^2\beta^j \geq 1$, or **(a-2)** we have $k^2\beta^j < 1$. We find that under assumption **(a-1)** we can apply the vanilla Chernoff bound to obtain

$$\begin{aligned} \mathbb{P}(\mathbf{(a)}) &= \mathbb{P}(\text{Bi}(\delta k^2, \alpha^j) \geq k^2\beta^j) \\ &\leq \exp[-\delta k^2 \alpha^j H((\beta/\alpha)^j)] \\ &\leq \exp[-\delta k^2 \beta^j \ln((\beta/\alpha)^j)/2] \\ &\leq (\alpha/\beta)^{j/2}. \end{aligned} \tag{6}$$

where in the third line we have used that $H(x) \geq \frac{1}{2}x \ln x$ for $x \geq x_0$ with x_0 sufficiently large and that we can choose α, β such that $\beta/\alpha \geq x_0$. Under the assumption **(a-2)** we see that

$$\begin{aligned} \mathbb{P}(\mathbf{(a)}) &= \mathbb{P}(\text{Bi}(\delta k^2, \alpha^j) \geq 1) \\ &\leq \delta k^2 \alpha^j \\ &= (\delta k^2 \beta^j) \cdot (\alpha/\beta)^j \\ &\leq (\alpha/\beta)^j, \end{aligned} \tag{7}$$

using assumption **(a-2)** for the third line.

Combining the above bounds we see that, provided α, β, γ were chosen appropriately, we have

$$\mathbb{P}(E_{i,j}) \leq 2e^2 (\beta/\gamma^2)^j + (\alpha/\beta)^{j/2} < 10^{-j},$$

the last inequality again holding if we chose—as we may assume we have— α, β, γ appropriately.

We now consider the situation where $k \geq \sqrt{n}$. Note that in this case, having chosen the constants appropriately,

$$k\gamma^j \geq \sqrt{n\gamma^{2I}} \geq \sqrt{n\gamma^{4\varepsilon \log(n)}} = \Omega(n^{1/4}).$$

Thus, the second line of (5) give in fact that

$$\mathbb{P}(\mathbf{(b)}) = \exp\left[-\Omega\left(n^{1/4}\right)\right]$$

Similarly, having chosen the constants correctly, we have

$$k^2 \beta^j > n \beta^I \geq n^{1+\varepsilon \log(\beta)} = \Omega(\sqrt{n}).$$

Thus, the case **(a-2)** does not apply, and the third line of (6) gives that

$$\mathbb{P}(\mathbf{a}) = \exp[-\Omega(\sqrt{n})].$$

This shows that indeed $\mathbb{P}(E_{i,j}) = \exp[-\Omega(n^{1/4})]$ in the case when $k \geq \sqrt{n}$. \square

This last lemma is all we need to prove part 1 of Theorem 2, for the case when $k \geq \sqrt{n}$.

Proof of Theorem 2. First assume $k \geq \sqrt{n}$. By the previous lemma, the probability that during any round of the process some group “fires” is at most

$$\begin{aligned} \mathbb{P}(\text{fires}) &\leq \sum_{i+j \leq I} \mathbb{P}(E_{i,j}) \\ &= O\left((\log n)^2\right) \cdot \exp[-\Omega(n^{1/4})] = o(1). \end{aligned}$$

This means that, with probability $1 - o(1)$, there will be a total of $I = \Theta(\log(n/k))$ rounds and in each round we spend the full amount of time δk^2 . So indeed, with probability $1 - o(1)$, the process does not complete before $\delta k^2 I = \Omega(k^2 \log(n/k))$.

In the rest of the proof we will thus be assuming that $k \leq \sqrt{n}$. Note that this in particular implies that $I = \Theta(\log(n/k)) \rightarrow \infty$. We will denote J_i be the number of groups that get filled in round i .

Appealing to Lemma 7, we see that sequence J_1, J_2, \dots is stochastically dominated by an i.i.d. sequence J'_1, J'_2, \dots where $\mathbb{P}(J'_1 = j) = 10^{-j}$ for $j \geq 2$ and $\mathbb{P}(J'_1 = 1) = \frac{89}{90}$. It is easily checked that $\mathbb{E}(J'_1)^2 < \infty$. Let us set $I' := \varepsilon I$ for some suitable chosen $0 < \varepsilon < 1$. By the law of large number (which holds for i.i.d. sequence with a finite second moment) we have that

$$\mathbb{P}(J'_1 + \dots + J'_{I'} > 2I' \mathbb{E}(J'_1)) = o(1).$$

Thus, choosing ε appropriately (namely so that $2\varepsilon \mathbb{E}(J'_1) < 1$), we we have

$$\mathbb{P}(J_1 + \dots + J_{I'} > I) \leq \mathbb{P}(J'_1 + \dots + J'_{I'} > I) = o(1).$$

In other words, with probability $1 - o(1)$, there are at least I' rounds.

Now let Y_i be a $\{0, 1\}$ -valued variable that equals one if in round i no group fired. Observe that

$$\mathbb{P}(Y_i = 1) \geq 1 - \sum_{j=1}^{I-i} \mathbb{P}(E_{i,j}) \geq 1 - \sum_{j=1}^{\infty} 10^{-j} = \frac{8}{9}.$$

Observe that Y_1, Y_2, \dots stochastically dominates an i.i.d sequence Y'_1, Y'_2, \dots of Bernoulli($\frac{8}{9}$) random variables. Again by the law of large numbers we have that

$$\begin{aligned} \mathbb{P}(Y_1 + \dots + Y_{I'} > \frac{4}{9} I') &\geq \mathbb{P}(Y'_1 + \dots + Y'_{I'} > \frac{4}{9} I') \\ &= 1 - o(1). \end{aligned}$$

So this shows that with probability $1 - o(1)$ there were at least $\frac{4}{9} I' = \Omega(\log(n/k))$ rounds in which no group fired, which means the round lasted the full amount of δk^2 time. Hence, with probability $1 - o(1)$, at time $\delta k^2 I' = \Omega(k^2 \log(n/k))$ the process has not yet completed.

Finally, to prove parts 2 and 3 of Theorem 2, note that in each step of the process at least one square of the grid gets filled. Thus $T \leq nk$ (with probability one). Since $T \geq S$ the lower bound in 2 follows immediately from Corollary 5 and the lower bound in 3 from Lemma 3. \square

4.5 CD-P heuristic.

Similarly to the CD heuristic, the difference between the best and worst case is large: we again have $k \leq \mathbf{RecMess} \leq nk$.

For this protocol, however, the number of messages reduces more quickly than in the case of CD. We show that the difference is significant and speeds up the process by roughly a factor k .

Theorem 8. *In the unbounded reach scenario under the CD-P heuristic, with high probability, $\mathbf{RecMess} \in \Theta(k \log n)$.*

Under the CD-P heuristic nodes prioritize messages in the queue by the distance to their senders. Thus, when a node i is chosen to send a message, it will choose a message that is currently present in the maximum number of nodes. Then, all the nodes that are further away from the destination, i.e., all the nodes with coordinates less than i will delete this message from their queues. We will show that, with high probability, it takes $\Theta(k \log n)$ steps until all nodes have empty queues.

To do so we represent the number of nodes in which a message is present as random variables and will bound them from above by continuous random variables with a simpler distribution at every iteration of the protocol.

We let $\mathcal{R}^t = (R_1^t, R_2^t, \dots, R_k^t)$ where R_i^t denotes the number of nodes that contain message i at time t . Thus, at time $t = 0$ we have $\mathcal{R}^0 = (n, \dots, n)$. One iteration of the CD-P protocol corresponds to the largest value in \mathcal{R} being randomly reduced. In other words, we find an index i such that $R_i^t = \max_{j=1, \dots, k} R_j^t$; we set $R_j^{t+1} := R_j^t$ for all $j \neq i$; and we choose R_i^{t+1} uniformly at random from $\{0, \dots, R_i^t - 1\}$.

We let T denote the first t for which $R_1^t = \dots = R_k^t = 0$. We shall always assume that $n \geq 2$, because if $n = 1$ then trivially $T = n$.

We start by making the observation that the (distribution of the) completion time would be the same if we changed the process as follows: first we repeatedly keep setting R_1^{t+1} to a random value in $\{0, \dots, R_1^t - 1\}$ until we reach a situation where $R_1^t = 0$; next we do the same for R_2 , then for R_3 and so on. The reason that this gives the same (distribution of the) completion time is that for every time step t of the original process, every R_i^t that is nonzero is sure to become the maximum value in some future round $t' \geq t$.

Let us thus write Z for the (random number) of rounds it takes for R_1^t to reach zero, starting from $R_1^0 = n$. Then $T \stackrel{d}{=} Z_1 + \dots + Z_k$, where Z_1, \dots, Z_k are i.i.d. distributed like Z .

Let U_1, U_2, \dots be i.i.d. uniform on $[0, 1]$ and let us denote

$$\begin{aligned} Z_{\text{upper}} &:= \min\{t : n \cdot U_1 \cdot \dots \cdot U_t < 1\}, \\ Z_{\text{lower}} &:= \min\{t : n \cdot U_1 \cdot \dots \cdot U_t < t + 1\}. \end{aligned}$$

We have:

Lemma 9. *We have $Z_{\text{lower}} \leq_{st} Z \leq_{st} Z_{\text{upper}}$.*

Proof. For notational convenience we write $V_i := nU_1 \dots U_i$, $W_i := V_i - i$. We start with the (stochastic) upper bound. It is enough to show that $R_1^i \leq_{st} V_i$ for all i , since that will prove that

$$\begin{aligned} \mathbb{P}(Z > t) &= \mathbb{P}(R_1^{\lfloor t \rfloor} \geq 1) \\ &\leq \mathbb{P}(V_{\lfloor t \rfloor} \geq 1) = \mathbb{P}(Z_{\text{upper}} > t). \end{aligned} \tag{8}$$

We will use induction to prove $R_1^i \leq_{st} V_i$ for all i . Clearly $R_1^0 = n = V_0$ so the base case holds. Let us thus suppose that $R_1^i \leq_{st} V_i$ for some i . By Strassen's theorem, there is a coupling for (R_1^i, V_i) so that $R_1^i \leq V_i$ with probability one. If under this coupling, we condition on the

event that $V_i = x, R_1^i = \ell$ (where necessarily $\ell \leq \lfloor x \rfloor$), then $V_{i+1} = V_i U_{i+1}$ is chosen uniformly at random on the interval $[0, x]$, while R_1^{i+1} is chosen uniformly from the set $\{0, \dots, \ell - 1\}$. Thus, for all $0 \leq y \leq x$ we have

$$\mathbb{P}(V_{i+1} \geq y | V_i = x, R_1^i = \ell) = \frac{x - y}{x} = 1 - \frac{y}{x},$$

and

$$\begin{aligned} \mathbb{P}(R_1^{i+1} \geq y | V_i = x, R_1^i = \ell) &= \frac{(\ell - 1) - \lceil y \rceil}{\ell} \\ &= 1 - \frac{\lceil y \rceil + 1}{\ell} \leq 1 - \frac{y}{x}. \end{aligned}$$

Summing over ℓ , this simplifies to

$$\begin{aligned} \mathbb{P}(V_{i+1} \geq y | V_i = x) &= \sum_{\ell} \mathbb{P}(V_{i+1} \geq y | V_i = x, R_1^i = \ell) \mathbb{P}(R_1^i = \ell | V_i = x) \\ &= 1 - \frac{y}{x}, \end{aligned} \tag{9}$$

and similarly

$$\mathbb{P}(R_1^{i+1} \geq y | V_i = x) \leq 1 - \frac{y}{x}. \tag{10}$$

Hence

$$\begin{aligned} \mathbb{P}(V_{i+1} \geq y) &= \int_y^n \mathbb{P}(V_{i+1} \geq y | V_i = x) f_{V_i}(x) dx \\ &\geq \int_y^n \mathbb{P}(R_1^{i+1} \geq y | V_i = x) f_{V_i}(x) dx \\ &= \mathbb{P}(R_1^{i+1} \geq y), \end{aligned} \tag{11}$$

so that the upper bound is proved.

For the lower bound, we remark it is sufficient to show $W_i \leq_{\text{st}} R_1^i$ for all i , by an argument analogous to (8). Again we use induction. The details are similar to the previous case and left for out for space reasons; refer to the full version for a complete proof. \square

Corollary 10. $\mathbb{P}(Z_1 > t) \leq n(1/2)^t$.

Proof. This follows by Markov's inequality since $\mathbb{E}(nU_1 \dots U_t) = n(1/2)^t$. \square

Corollary 11. *There exists a universal constant $c > 0$ such that*

$$\mathbb{P}(Z < c \log n) \leq \frac{1}{\log n} \quad (\text{for all } n \geq 2)$$

(Note that the distribution of Z implicitly depends on n .)

The proof of this corollary makes use of the observation that if U is uniform on $[0, 1]$ and $W := \ln(1/U)$ then $\mathbb{E}W = \text{Var } W = 1$. For completeness we spell out the straightforward computations.

Lemma 12. *Let U be uniform on $[0, 1]$ and let $W := \ln(1/U)$. Then $\mathbb{E}W = \text{Var } W = 1$.*

Proof. We have that

$$\mathbb{E}W = \int_0^1 -\ln u \, du = [-u \ln u + u]_0^1 = 1,$$

and

$$\mathbb{E}W^2 = \int_0^1 \ln^2 u \, du = [u \ln^2 u - 2u \ln u + 2u]_0^1 = 2.$$

So indeed $\text{Var } W = \mathbb{E}W^2 - (\mathbb{E}W)^2 = 1$. \square

Proof of Corollary 11. Since $\mathbb{P}(Z = 0) = 0$, we can assume that $n \geq n_0$ for some sufficiently large constant n_0 (at the price of possibly having to lower the universal constant c a bit). We put $t := \lfloor \frac{\ln n}{1000} \rfloor$ and we note that

$$\begin{aligned} \mathbb{P}[nU_1 \dots U_t < t + 1] \\ = \mathbb{P}[\ln(1/U_1) + \dots + \ln(1/U_t) > \ln n - \ln(t + 1)]. \end{aligned}$$

Writing $S := \ln(1/U_1) + \dots + \ln(1/U_t)$ for convenience, we see that $\mathbb{E}S = \text{Var } S = t$ by Lemma 12 above, and hence

$$\begin{aligned} \mathbb{P}[S > \ln n - \ln(t + 1)] \\ &= \mathbb{P}[S - \mathbb{E}S > \ln n - \ln(t + 1) - \mathbb{E}S] \\ &\leq \mathbb{P}[S - \mathbb{E}S > \ln n - t - \mathbb{E}S] \\ &\leq \frac{t}{(\ln n - 2t)^2} \\ &\leq \frac{1}{\ln n}, \end{aligned}$$

having used $t \geq \ln(t + 1)$ in the second line; Chebyshev in the third line; and that n is sufficiently large in the fourth line. \square

Corollary 13. *There exist universal constants $c, C > 0$ such that, if $n \rightarrow \infty$ and $k = k(n)$ satisfies $1 \leq k = o(\log n)$ then*

$$\mathbb{P}(ck \log n < T < Ck \log n) \rightarrow 1.$$

Proof. For the upper bound, we notice that

$$\begin{aligned} \mathbb{P}(T > Ck \log n) &\leq k\mathbb{P}(Z > C \log n) \\ &\leq kn(1/2)^{C \log n} \rightarrow 0, \end{aligned}$$

using Corollary 10 and where the limit follows assuming C is sufficiently large. For the lower bound, we argue similarly noting that

$$\mathbb{P}(T < ck \log n) \leq k\mathbb{P}(Z_1 < c \log n) \leq \frac{k}{\log n} \rightarrow 0,$$

using Corollary 11 (assuming c was chosen smaller than the constant provided there) and using that $k = o(\log n)$. \square

This last corollary proves our main result in the case when $n \rightarrow \infty$ and k is either constant or does not go to infinity too fast. To complete the proof of the main result, it now suffices to consider the case when $k \rightarrow \infty$ and $n = n(k) \geq 2$ is arbitrary. Before we tackle this case we need one more observation.

Corollary 14. *There exist universal constants $c, C > 0$ such that*

$$\begin{aligned} c \log n &\leq \mathbb{E}Z \leq C \log n, \\ c \log^2 n &\leq \mathbb{E}Z^2 \leq C \log^2 n, \end{aligned} \quad (\text{for all } n \geq 2)$$

(Again we remind the reader that Z depends implicitly on n .)

Proof. It follows from Corollary 11 that

$$\mathbb{E}Z \geq c \log n \cdot \mathbb{P}(Z \geq c \log n) = \Omega(\log n),$$

if c is as in Corollary 11. The same argument also gives $\mathbb{E}Z^2 \geq (\mathbb{E}Z)^2 = \Omega(\log^2 n)$.

By a standard formula for the expectation of nonnegative random variables, we have

$$\begin{aligned} \mathbb{E}Z &= \int_0^\infty \mathbb{P}(Z \geq t) dt \\ &\leq C \log n + \int_{C \log n}^\infty n(1/2)^t dt \\ &= O(\log n), \end{aligned}$$

provided C is chosen sufficiently large. Similarly

$$\begin{aligned} \mathbb{E}Z^2 &= \int_0^\infty \mathbb{P}(Z \geq \sqrt{t}) dt \\ &\leq C \log^2 n + \int_{C \log^2 n}^\infty n(1/2)^{\sqrt{t}} dt \\ &= O(\log^2 n), \end{aligned}$$

using the substitution $z := \ln(2)\sqrt{t}$ in the third line, and that $z \leq e^{z/2}$ for all sufficiently large z in the third line. \square

Corollary 15. *There exist universal constants c, C such that, if $k \rightarrow \infty$ and $n = n(k) \geq 2$ is arbitrary then*

$$\mathbb{P}(ck \log n \leq T \leq Ck \log n) \rightarrow 1.$$

Proof. Since $T = Z_1 + \dots + Z_k$, the previous corollary shows that $\mathbb{E}T = k\mathbb{E}Z = \Theta(k \log n)$ and $\text{Var } T = k \text{Var } Z = \Theta(k \log^2 n)$. It follows that

$$\mathbb{P}(|T - \mathbb{E}T| > \frac{1}{2}\mathbb{E}T) \leq \frac{4 \text{Var } T}{(\mathbb{E}T)^2} = O\left(\frac{1}{k}\right) = o(1),$$

using Chebyshev and Corollary 14. To conclude, we observe that, provided we chose c small enough and C big enough, we have $\{T < ck \log n\}, \{T > Ck \log n\} \subseteq \{|T - \mathbb{E}T| > \frac{1}{2}\mathbb{E}T\}$. \square

Proof of Theorem 8. We remark that Corollaries 13 and 15 together prove the theorem, taking the minimum of the two c 's and the maximum of the two C 's provided by these corollaries for the full result. \square

4.6 Delay-based heuristic.

We assume that the delay is chosen such that it increases by exactly one time step per node; that is, $\text{MD} = r$. The nodes delete messages from their queues when they receive them for the second time. Thus, every message is retransmitted only once, no matter which delay function is used. Therefore, $\mathbf{RecMess} = k$.

In the variant in which the nodes delete messages from their queues only when they receive a duplicate from a node that is closer to the destination, we get:

$$\mathbf{RecMess} = \begin{cases} 2^k & \text{if } k < \log n, \\ n + n(k - \log n) & \text{if } k \geq \log n. \end{cases}$$

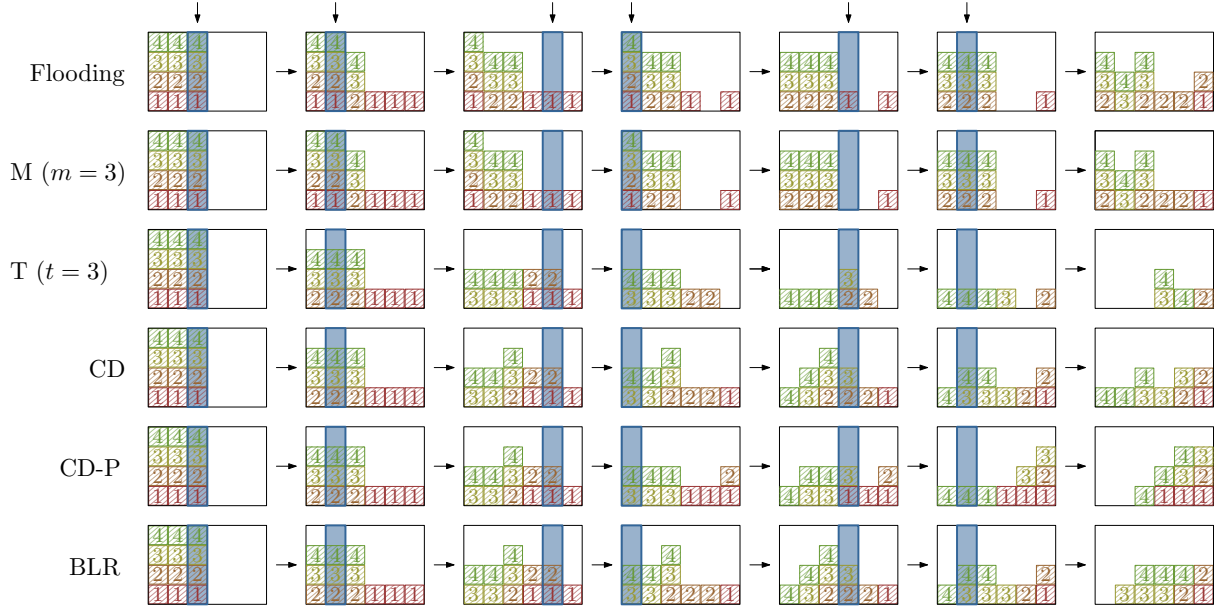


Figure 5: The different protocols run on the same data as Figure 4, but with bounded radius $r = 3$.

5 Analysis: bounded reach

In the *bounded reach scenario* each node can communicate with r neighbors to its left and r to its right, for r a parameter. This scenario generalizes the previous one, allowing to evaluate the effect that node density (indirectly related to r) has on the different protocols. An example of the behavior of all protocols is shown in Figure 5.

5.1 Simple flooding.

Under this protocol, every node will receive each message from at most $2r$ of its neighbors: $\mathbf{RecMess} = O(rk)$.

5.2 M heuristic.

If $2r \leq M$ this protocol is equivalent to the previous one. If $2r > M$ it is equivalent to the M heuristic for the unbounded reach scenario. Therefore $\mathbf{RecMess} = O(\min(M, 2r)k)$.

5.3 T heuristic.

If $T \geq r$, no message will ever be forwarded. If $T < r$, then each node u can receive a message from at most $2r$ nodes. Each time it receives one, at least T and at most $2T$ of the nodes within reach of u delete the message from their queues. Thus: $\mathbf{RecMess} = O(\frac{rk}{T})$.

5.4 CD heuristic.

For the CD heuristic, in the best case, each message is transmitted only by the node closest to the target. Then each node receives each message only twice: once from the left and once from the right. In the worst case, messages are retransmitted by the node farthest from the target, causing all nodes to retransmit every message. However, because of the bounded reach, the flooding effect is somewhat mitigated: each node has only $2r$ neighbors and receives each message $2r$ times. So, $2k \leq \mathbf{RecMess} \leq 2rk$. As before, we can make a more precise probabilistic

statement assuming fair medium access. The proofs of the following theorem and the next one are deferred to the next section.

Theorem 16. *In the bounded reach scenario under the CD heuristic, w.h.p., $\mathbf{RecMess} \in O(k^{3/2})$.*

Consider as before an $n \times k$ table, where columns represent message queues of the nodes, and rows represent the messages that are in the queues. All messages start in the leftmost r columns. Whenever a node retransmits a message, it gets deleted from all nodes to its left and added to the $r - 1$ nodes to its right. So, each message is always present in exactly r consecutive nodes (except at the end of the process), that we will refer to as a *train* of r messages.

In CD, when a node gets to retransmit, it picks the message with the lowest id. This means that the trains of messages are always ordered by id: the lower the id, the closer the train is to the destination. Notice that the trains cannot “overtake” each other. Thus, when a node chooses a message to retransmit, it retransmits the message that corresponds to the train most ahead of the others. On the contrary, in the CD-P protocol, when a node sends a message, it chooses it from the train that is most behind.

From this we can already see that the average speed at which the trains progress towards the destination is smaller than $r/2$ nodes per move for CD, and greater than $r/2$ nodes per move for CD-P. If a train of messages does not overlap with the other trains, its expected progress during the next move is $r/2$ nodes. If a train does overlap with other trains, if it is chosen to progress under the CD heuristic, it moves by less than $r/2$ nodes, and under CD-P heuristic, by more than $r/2$ nodes.

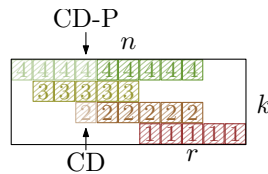


Figure 6: CD vs CD-P in bounded reach scenario.

Lower bound. In the best case any algorithm will need at least $\frac{kn}{r}$ steps, as a message cannot progress by more than r nodes at a time.

Upper bounds. We first prove Theorem 17, since it is simpler. Then we focus on Theorem 16.

Proof of Theorem 17. The expected progress for a train of r messages at each move is greater than $\frac{r}{2}$. Therefore, after $\frac{2kn}{r}$ steps, the expected number of messages left in the queues will be 0. The total number of messages sent is $O(\frac{kn}{r})$, and every node receives a fraction $O(\frac{r}{n})$ of all the messages, therefore, $\mathbf{RecMess} = O(\frac{kn}{r} \cdot \frac{r}{n}) = O(k)$. \square

5.5 CD-P heuristic.

As before, the best and worst cases are the same as for the CD heuristic, that is, $2k \leq \mathbf{RecMess} \leq 2rk$.

However, we can prove that with high probability and for $r \ll n$, the actual number of messages received in this scenario is only a constant factor larger than the best case; that is, linear in k .

Theorem 17. *In the bounded reach scenario under the CD-P heuristic, w.h.p., $\mathbf{RecMess} \in \Theta(k)$.*

Proof. The expected progress for a train of r messages is trickier to calculate. In the case where all the trains overlap, it is

$$\begin{aligned}\mathbb{E}(\text{prog.}) &= \sum_1^{k-1} \frac{(R_{i+1} - R_i)^2}{2(r + R_k - R_1)} + \frac{r^2}{2(r + R_k - R_1)} \\ &= f(R_1, R_2, \dots, R_k),\end{aligned}$$

where $R_1 \leq R_2 \leq \dots \leq R_k$, and R_i is the furthest node from the destination that still has message i , or, in other words, the tail of the train. Finding the minimum of this function, we get:

$$\begin{aligned}\frac{\partial f}{\partial R_1} &= -\frac{R_2 - R_1}{r + R_k - R_1} + \sum_1^{k-1} \frac{(R_{i+1} - R_i)^2}{2(r + R_k - R_1)^2} \\ &\quad + \frac{r^2}{2(r + R_k - R_1)^2} = 0, \\ \frac{\partial f}{\partial R_i} &= \frac{R_i - R_{i-1}}{r + R_k - R_1} - \frac{R_{i+1} - R_i}{r + R_k - R_1} = 0 \text{ for } 1 < i < k, \\ \frac{\partial f}{\partial R_k} &= \frac{R_k - R_{k-1}}{r + R_k - R_1} - \sum_1^{k-1} \frac{(R_{i+1} - R_i)^2}{2(r + R_k - R_1)^2} \\ &\quad - \frac{r^2}{2(r + R_k - R_1)^2} = 0.\end{aligned}$$

From the middle equations we get that $R_i = \frac{R_{i-1} + R_{i+1}}{2}$ when the expected shift is minimized. Thus, the expected shift is

$$\begin{aligned}\mathbb{E}(\text{prog.}) &\geq \sum_1^{k-1} \frac{(\frac{R_k - R_1}{k-1})^2}{2(r + R_k - R_1)} + \frac{r^2}{2(r + R_k - R_1)} \\ &= \frac{(R_k - R_1)^2}{2(k-1)(r + R_k - R_1)} + \frac{r^2}{2(r + R_k - R_1)}.\end{aligned}$$

Let $R_k - R_1 = x$, and $\mathbb{E}(\text{progress}) = f(x)$. Again find the minimum of the function:

$$\begin{aligned}\frac{df}{dx} &= \frac{x}{(k-1)(r+x)} - \frac{x^2}{2(k-1)(r+x)^2} \\ &\quad - \frac{r^2}{2(r+x)^2} = 0.\end{aligned}$$

We get equality

$$x^2 + 2xr + r^2 - kr^2 = 0$$

from which we get that the expected shift is minimized when $x = r(\sqrt{k} - 1)$. Thus,

$$\mathbb{E}(\text{prog.}) > \frac{r}{\sqrt{k} + 1}.$$

If the trains of messages do not overlap, the expected progress per move is even higher. Therefore, after $\frac{(\sqrt{k+1})kn}{r} = O(\frac{k^{3/2}n}{r})$ steps, the expected number of messages left in the queues will be 0. The total number of messages sent is $O(\frac{k^{3/2}n}{r})$, and every node receives a fraction $O(\frac{r}{n})$ of all the messages, therefore, $\mathbf{RecMess} = O(\frac{k^{3/2}n}{r} \cdot \frac{r}{n}) = O(k^{3/2})$. \square

5.6 Delay-based heuristic.

Since all messages get deleted from the queue when heard by a node for the second time, $\mathbf{RecMess} = O(\frac{nk}{r})$.

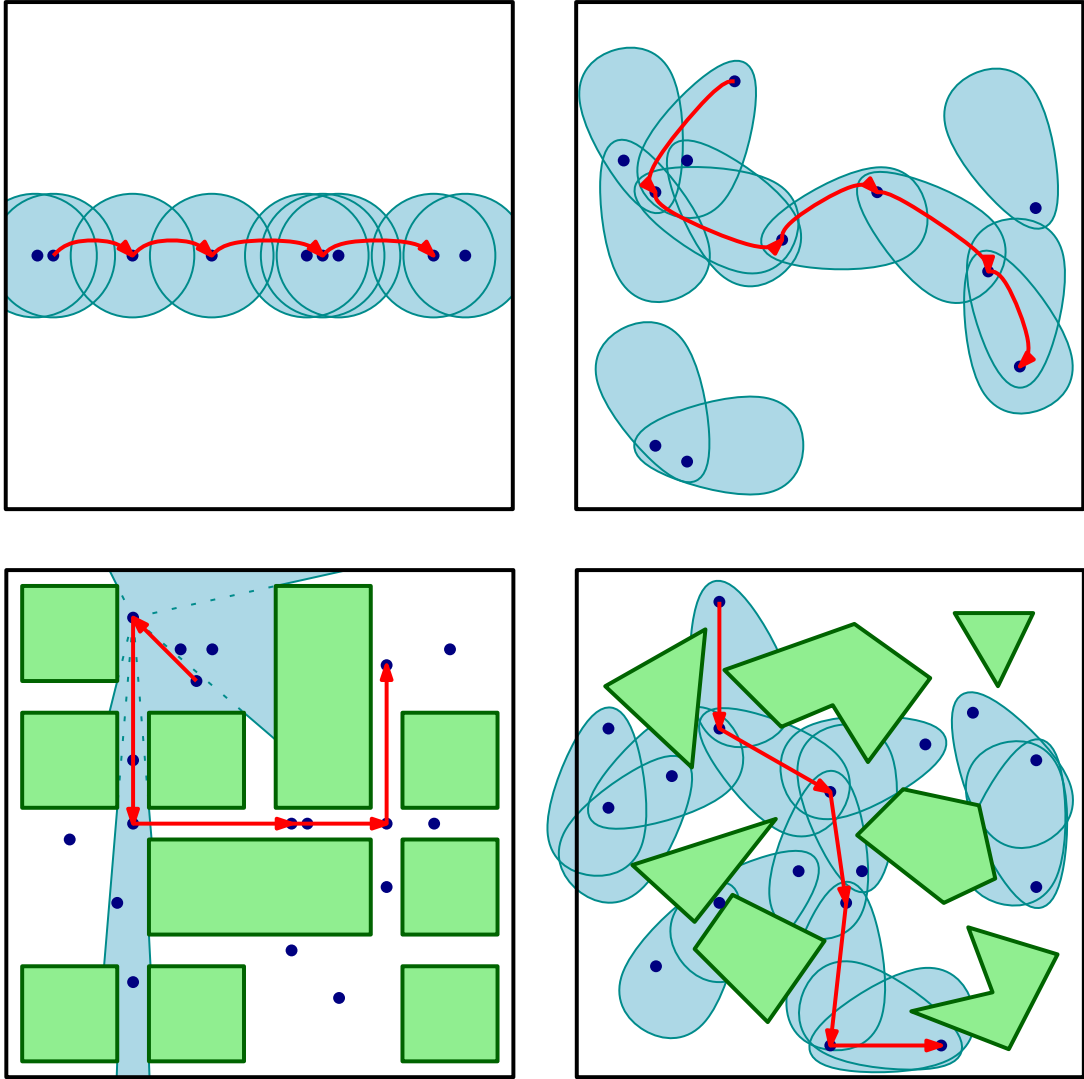


Figure 7: Scenarios of increasing complexity. (i) 1D scenario. (ii) 2D scenario; non-uniform ranges. (iii) Regular obstacles; visibility. (iv) Irregular obstacles; combining visibility and non-uniform ranges.

6 Discussion

Beaconless geocast protocols are used in practice in 2D scenarios. They differ from the 1D ones in a few but important characteristics, most notably that obstacles (like buildings, which cannot be traversed by the transmission signal) need to be surrounded, and that local optimization strategies fail to guarantee delivery. Therefore, combinations of different strategies need to be used in order to achieve delivery guarantees and, at the same time, keep the network load within reasonable bounds. The network load analysis in this cases is difficult, and almost only experimental results exist. This motivated studying the 1D case. We have shown that the rigorous analysis of geocast protocols even in simple 1D scenarios can be interesting—and challenging. Indeed, all protocols give rise to different load bounds, with CD and CD-P being particularly subtle to analyze due to the fact that an involved probabilistic analysis is required.

Our theoretical analysis confirms behaviors that had been observed before only through simulations. The different expression for **RecMess**, summarized in Figure 2, make evident the differences between flooding and the restricted flooding (M- and T-heuristics) and delay-based

protocols, as well as the advantage of CD over them for small values of k . Moreover, our analysis gives theoretical support to the performance claims of CD-P: Hall [1] showed through simulations that CD-P scales better and improves its efficiency over CD considerably. Our results indeed corroborate this. In the unbounded reach scenario, CD has serious scalability problems when the number of messages (k) approaches the number of nodes (n), making it behave like flooding, while CD-P's dependence on k remains always linear, and, more importantly, only logarithmic on n . For the bounded reach scenario, CD again shows an overhead in k (this time of only $O(\sqrt{k})$), while CD-P achieves the optimal asymptotic performance of $\Theta(k)$.

The results in this paper are a first step towards analyzing geocast protocols in generic geometric settings. This paper has focused on two relatively simple scenarios. Between these basic scenarios and the final intricacies of real-world situations, several abstractions of the geocast problem of varying complexity can be imagined, which are definitely worth studying next (see Figure 7).

References

- [1] R. J. Hall. An improved geocast for mobile ad hoc networks. *IEEE Transactions on Mobile Computing*, 10(2):254–266, 2011.
- [2] R. J. Hall and J. Auzins. A tiered geocast protocol for long range mobile ad hoc networking. In *IEEE Military Communications conference (MILCOM)*, pages 1–8, 2006.
- [3] M. Heissenbüttel, T. Braun, T. Bernoulli, and M. Wälchli. BLR: Beacon-less routing algorithm for mobile ad-hoc networks. *Computer Communications*, 27(11):1076–1086, 2004.
- [4] Y. Ko and N. H. Vaidya. Geocasting in mobile ad hoc networks: Location-based multicast algorithms. In *2nd Workshop on Mobile Computing Systems and Applications (WMCSA)*, pages 101–110, 1999.
- [5] C. Maihöfer. A survey of geocast routing protocols. *IEEE Communications Surveys and Tutorials*, 6(1-4):32–42, 2004.
- [6] S. Rührup. Theory and practice of geographic routing. In *Ad Hoc and Sensor Wireless Networks: Architectures, Algorithms and Protocols*. Bentham Science, 2009.
- [7] I. Stojmenovic and X. Lin. Loop-free hybrid single-path/flooding routing algorithms with guaranteed delivery for wireless networks. *IEEE Transactions on Parallel and Distributed Systems*, 12(10):1023–1032, 2001.
- [8] M. Zorzi. A new contention-based mac protocol for geographic forwarding in ad hoc and sensor networks. In *IEEE Int. Conference on Communications (ICC)*, pages 3481–3485, 2004.