Irrigating Ad Hoc Networks in Constant Time

[Extended Abstract]

D. Dubhashi,
C. Johansson
Department of Computing Science
Chalmers University
SE 412 96 Göteborg, Sweden

O. Häggström
Department of Mathematical Statistics
Chalmers University
SE 412 96 Göteborg, Sweden

A. Panconesi, M. Sozio
University La Sapienza of Rome
Dipartimento di Informatica
via Salaria 113, 00198, Roma, Italy *

ABSTRACT

We propose very simple randomized algorithms to compute sparse overlay networks for geometric random graphs modelling wireless communication networks. The algorithms generate in constant time a sparse overlay network that, with high probability, is connected and spans the whole network. Moreover, by making use of the "power of choice" paradigm, the maximum degree can be made as small as $O(\log \log n)$, where n is the size of the network. We show the usefulness of this kind of overlays by giving a new protocol for the classical broadcast problem, where a source is to send a message to the whole network. Our experimental evaluation shows that our approach outperforms the well-known gossiping approach in all situations where the cost of a message can be charged to the pair (sender, receiver), i.e. to the edge connecting the two. This includes sensor networks.

Categories and Subject Descriptors

F.2 [**Analysis of Algorithms and Problem Complexity**]; G.2.2 [**Network Problems**]; G.3 [**Probabilistic Algorithms**]

General Terms

Algorithms, Experimentation, Performance, Theory

Keywords

Distributed Algorithms, Randomized Protocols, Wireless Networks, Overlay Networks, Ad Hoc Networks

*This work was funded by project WEB-MINDS of the Italian Ministry of University and Research under the FIRB program, and by the EC projects EYES and DELIS

SPAA'05, July 18–20, 2005, Las Vegas, Nevada, USA.

1. INTRODUCTION

In the study of wireless communication networks, the following setup is natural. A number n of nodes are distributed in some geographical region. Links can be established between any pair of nodes within a given distance r from each other.

> *What is a good strategy for connecting nodes to each other, in order to obtain good connectivity of the resulting global network?* (1.1)

We shall for simplicity take the geographical region to be the unit square $[0,1]^2$, and also make the standard probabilistic model assumption that the positions of the n nodes are random: independent and uniformly distributed on $[0,1]^2$.

An obvious answer to question (1.1) is to establish links between *all* pairs of nodes within distance r from each other. We shall refer to this as the *full visibility graph* and denote it as G_r^n, where n is the number of nodes. Computing G_r^n in a wireless environment may be very costly, so one would like to find a localized, distributed strategy (i.e., not requiring global optimization or coordination) which keeps the number of links incident to each node small. This problem arises for instance in the context of Bluetooth (BT) ad hoc networks. The BT standard requires network nodes to become aware of their neighbors (i.e. of nodes within transmission range). Once two neighbors become aware of each other a reliable communication channel between them can be established. This is the so-called *device discovery* phase the objective of which is to set up a globally connected network by means of these reliable links. As shown in [5], for each node to discover *all* neighbors, i.e. to set up G_r^n, is too time consuming. Fortunately, as shown in [20], reasonable time-outs, in the order of 10 seconds, give good statistical guarantees of connectivity. However, protocols based on absolute time limits are clearly not robust (i.e. independent neither of the technology nor of specific environmental conditions). Our first contribution is to give a robust local rule for generating connected BT networks. As a first approximation, BT device discovery is a randomized protocol that selects neighbours within range with uniform probability [6]. This motivates the following definition.

DEFINITION 1. *Fix $r > 0$ and a positive integer c. We take $G_{r,c}^n = (V_{r,c}^n, E_{r,c}^n)$ to denote the geometric random graph defined as follows.*

- *The vertex set $V^n_{r,c}$ consists of n points, picked independently according to the uniform distribution on* $[0,1]^2$.
- *Each node $v \in V^n_{r,c}$ connects to c nodes chosen uniformly at random among those within distance r. If the number of such nodes is less than c then v connects to all of them. This is done independently for all nodes v.*

The resulting graph $G^n_{r,c}$ is called the **irrigation graph** *with parameters r, c, and n.*

For this model we prove the following.

DEFINITION 2. *Let $s \in (0,1]$ be some fixed constant. An* s-giant component *of an undirected graph G is a connected subgraph of G containing at least ns of the vertices.*

PROPOSITION 3. *Fix $r > 0$ and $c \geq 2$. Then there exists a constant $s > 0$ such that,*

$$\lim_{n\to\infty} \mathbf{P}(G^n_{r,c}(u) \text{ is a } s\text{-giant component}) = 1 . \qquad (1.2)$$

The emergence of a giant component is interesting and potentially useful. But Proposition 3 is also a natural step toward proving the main result of this paper.

THEOREM 4. *Fix $r > 0$ and $c \geq 2$. Then*

$$\lim_{n\to\infty} \mathbf{P}(G^n_{r,c} \text{ is connected}) = 1 . \qquad (1.3)$$

Although we stated our results in terms of limits both statements hold with high probability. The probability of the complementary events goes to 0 as $\Theta(n^{-\epsilon})$ for some $\epsilon \in (0,1)$.

This theorem explains the above mentioned experimental results. If a BT device runs the device discovery protocol for a few seconds, it will have the possibility of discovering at least a couple of neighbours. Therefore true device discovery dominates G^n_{rc} (for $c = 2,3,4$) which, according to the theorem above, is connected with high probability. Note that Definition 1 also gives the desired robust local protocol for connectivity. The simulations presented in this paper show that the protocol generates connected networks not only "in the limit" but already realistic sized networks, for very small values of $c = 2,3,4$.

In general, Definition 1 embodies a very simple distributed algorithm to compute a sparse overlay (sub)network of the full visibility graph. In the synchronous, message-passing model of computation the running-time of the algorithm is constant. Also note that $G^n_{r,c}$ can be computed in a completely asynchronous fashion. In [19] it is shown that, for each r, there exists a c such G^n_{rc} is an expander with high probability (this result does not subsume the theorem above since here c is rather huge). All this makes the protocol useful in several contexts. Indeed, in [29] G^n_{rc} is used to give an efficient implementation of directed diffusion, a classical method used to compute routing tables. In [6] a variation of G^n_{rc} is used to solve with one stroke the above mentioned device discovery problem together with the scatternet formation problem in Bluetooth networks. The resulting protocol outperforms all previous existing solutions. A similar approach is pursued in the independent research of [3, 4].

In this paper we introduce another application of G^n_{rc} to the very basic *broadcast problem* where a source s is to send a message to every node in the network. Many broadcast (also called routing) protocols have been proposed in the literature for wireless networks. Some such as LAR [28] and DREAM [9] use GPS for position information whereas others such as AODV [31], ZRP [27] do not. All of these however use *flooding*: Starting from the source, every node that receives the message *for the first time* forwards it to its neighbours. Some heuristic optimizations are added on top of this basic scheme [26]. If the network is connected this process delivers the message to every node in the network. The communication cost of flooding however is typically too high and several alternative approaches have been proposed [23, 25]. A popular alternative is *vertex-based gossiping*: starting from the source, every node that receives the message for the first time forwards it to its neighbours with probability p. A related approach is *edge-based gossiping*: starting from the source, every node that receives the message for the first time forwards it to each neighbour with probability p. That is, when a node receives the message, a coin is flipped for every neighbour. Note that flooding is the limit case of gossiping when $p = 1$. Randomized gossip has become recognized as a component of building large scale distributed systems with robust communication (see, among others, [7, 32, 11, 12, 10, 13, 3, 4]). Hass, Halpern and Li [8] argue that randomized gossip can be used to significantly increase efficiency by reducing the number of messages sent by up to 35 %.

In this paper we propose the following alternative strategy which we dub *irrigating*: Flood through G^n_{rc}. We present empirical results demonstrating that irrigating is fairly superior to gossiping in every situation in which the cost of sending a message from a node u to another node v can be charged to the edge uv. This is the case for instance of sensor networks where the cost of sending a message is comparable to that of receiving it. The thorough empirical study of [18], based on the theoretical results and the preliminary empirical observations presented here, confirm that irrigating is a very effective approach for broadcasting in sensor networks.

In many applications, including Bluetooth, it is very important to generate sparse overlays with small maximum degree. It is not hard to see that the maximum degree of G^n_{rc} is $\Theta(\log n/\log\log n)$. By making use of the "power-of-choice" (see [2, 17]) we show that the maximum degree can be reduced substantially while still retaining a very simple protocol.

DEFINITION 5. *Fix $r > 0$ and positive integers c and d. We take $G^n_{r,c,d} = (V^n_{r,c,d}, E^n_{r,c,d})$ to denote the geometric random graph defined as follows.*

- *The vertex set $V^n_{r,c,d}$ consists of n points, picked independently according to the uniform distribution on* $[0,1]^2$.
- *Each node $v \in V^n_{r,c,d}$ connects to c vertices. Each time v connects to a vertex it chooses d vertices independently and uniformly at random among those within distance r. It then connect only with the lowest degree vertex among those d vertices. This is done independently for all nodes v.*

The following theorem shows that while still achieving connectivity, the "power–of–choice" paradigm achieves a near-exponential decrease in the maximum degree and also a qualitative amelioration of the entire degree sequence profile.

THEOREM 6. *Fix $r > 0$ and positive integers $c \geq 2$ and $d > 0$. Then*

(a)

$$\lim_{n\to\infty} \mathbf{P}(G^n_{r,c,d} \text{ is connected}) = 1\,. \tag{1.4}$$

(b) The maximum degree of a vertex in $G^n_{r,c,d}$ is $\Theta(\log\log n)$ with high probability.

(c) The degree sequence of $G^n_{r,c,d}$ is stochastically majorized by the degree sequence of $G^n_{r,c,d'}$ for $d' \leq d$.

Intuitively the third statement means that the degree sequence is more concentrated around the mean and hence more regular. For instance, the highest degree is smaller, but the least degree is higher.

Assuming that each processor makes its "two choices" atomically, it is possible to show that the distributed algorithm embodied in Definition 5 can be implemented in $O(\log n)$ steps per processor. For lack of space the discussion is omitted from this extended abstract.

We shall consider one more model, which captures the natural situation in which links may fail.

DEFINITION 7. *Fix $r > 0$, $c \geq 2$, and $p \in (0, 1)$. The graph $G^n_{r,c,p} = (V^n_{r,c,p}, E^n_{r,c,p})$ is defined as the geometric random graph obtained by the following procedure.*

- *First, generate a Bluetooth graph $G^n_{r,c} = (V^n_{r,c}, E^n_{r,c})$ as in Definition 1.*
- *Then, for each edge $e \in E^n_{r,c}$ independently, delete e with probability $1-p$ (thus keeping it with probability p).*

The resulting graph $G^n_{r,c,p}$ is called a **thinned irrigation graph** *with parameters r, c, p and n.*

Our main result for thinned irrigation graphs is the following, which shows that, provided p is not too small, the giant component result (Proposition 3) extends to the thinned irrigation model. However, there is a striking qualitative difference in that the connectedness result (Theorem 4) does *not* obtain:

THEOREM 8. *Fix $r > 0$, $c \geq 2$, and $p > \frac{1}{c}$. Then there exists a constant $s > 0$ such that*

$$\lim_{n\to\infty} \mathbf{P}(G^n_{r,c,p} \text{ has an } s\text{-giant component}) = 1\,. \tag{1.5}$$

On the other hand,

$$\lim_{n\to\infty} \mathbf{P}(G^n_{r,c,p} \text{ is connected}) = 0\,. \tag{1.6}$$

We remark that an easy branching process comparison shows that the emergence of a giant component *fails* when p is sufficiently small. In fact, our main technique for proving the results stated above is also comparisons with branching processes, but these are somewhat more sophisticated. This difference comes from the fact that for the small p result, the comparison works by showing that a certain branching process *dominates* the connected component containing a given vertex, while for our main results, the comparison goes the other way.

The use of branching process comparisons to establish connectivity properties is quite standard in percolation theory (see, e.g., Meester and Roy [16]), but deserves to become better known in the study of wireless networks.

In the next section, we prove Proposition 3, while Sections 3 and 4 are devoted to the proofs of Theorems 4 and 8 (Equation 1.6), respectively. For lack of space, the proof of part (a) of Theorem 6 is omitted altogether. It follows from an argument similar to the proof of Theorem 4. Section 5 presents our experimental evaluation of irrigating and gossiping for the broadcast problem.

2. GIANT COMPONENT

For the purpose of proving Proposition 3, and also later, Lemma 9 below will be useful. Fix an integer k such that

$$k > \frac{\sqrt{5}}{r} \tag{2.7}$$

and partition $[0, 1]^2$ into k^2 subsquares of size $\frac{1}{k} \times \frac{1}{k}$ in the obvious way. (One point of this choice of k is that it ensures that any two points sitting in adjacent subsquares are within distance r from each other; this will be needed in the proof of Theorem 4.)

Let $\mathcal{A}$ be the event *"each of the k^2 subsquares contains at least $\frac{n}{2k^2}$ points"*.

LEMMA 9. *For any fixed k, we have*

$$\lim_{n\to\infty} \mathbf{P}_n(\mathcal{A}) = 1\,. \tag{2.8}$$

Proof: Fix a square S and let X denote the number of points in S. Then, $\mu := EX = \frac{n}{k^2}$ and, by the Chernoff-Hoeffding bound,

$$\mathbf{P}_n(X < \frac{n}{2k^2}) \leq e^{-n/8k^2}.$$

Thus, the probability that some square has less than the required number of points is at most $k^2 e^{-n/8k^2}$. □

To investigate connected components of the irrigation graph, we shall employ the following method, which we will call the **sequential discovery procedure** (this is simply a breadth-first exploration). First, select a node v_0 at random (among all n nodes). Then consider the c edges chosen by v_0 (in the device discovery procedure of Definition 1), and denote the endpoints (other than v_0) of these edges by $v_1, \ldots, v_c$. Then continue with the edges chosen by v_1, and so on, in a breadth first search manner. Each time a new node is encountered, the node reached by it is included in our list of nodes, and the choice is deemed a **success**. Sometimes, the edge leads to a node already seen in this procedure, in which case the choice is said to be a **failure**. At any point of this search procedure, we may stop, and those vertices encountered whose outgoing edges have not been investigated (yet), are called **fresh** nodes.

At various stages of our arguments, we will invoke a comparison between the sequential discovery procedure and a (Galton–Watson) **branching process**. Such a branching process (see, e.g., Harris [14] or Asmussen and Hering [1]) begins with m_0 individuals. Each of these begets, independently of the others, a number of offspring, which has some given distribution f on the non-negative integers. Each of these children then has a number of children for themselves, again independent with distribution f. And so on, again in a BFS manner. One of two things will happen: either the branching process dies out after a finite number of generations, or it survives (forever). Excluding the trivial case where f puts unit mass on 1, it is well known that the branching process has positive probability of surviving if and only if f's first moment is strictly greater than 1.

We shall be particularly concerned with a branching process whose offspring distribution is the binomial distribution $\mathrm{Bin}(2, p)$. This can be compared to the sequential discovery procedure for $c = 2$ in the following way. Suppose that we can show that up until some given stage S of the sequential discovery procedure, each choice of a new node to connect to has probability at least p (conditionally on everything seen so far) of being a success. Then we can make a joint construction (a so-called coupling; see [15]) of the sequential discovery procedure and the $\mathrm{Bin}(2, p)$ branching process in such a way that each individual in the branching process corresponds to a particular node (not shared by any of the other individuals of the branching process) of the sequential procedure, up until the given stage S. We say in this case that the sequential procedure up until stage S **stochastically dominates** the branching processes (for stochastic domination see for example [15]). We will show that the sequential discovery procedure first generates almost surely a set of $\log n$ points and that from then on each point u begets offsprings with distribution $\mathrm{Bin}(2, p_u)$, with $p_u \geq \frac{3}{4}$. It follows from a standard application of stochastic domination that the survival probability of the sequential discovery procedure is at least that of $\log n$ independent branching processes with distribution $\mathrm{Bin}(2, \frac{3}{4})$.

Proof of Proposition 3: We prove the result for $c = 2$ only, which is obviously enough since adding edges is not going to destroy a giant component.

Run the sequential discovery procedure until the outgoing edges of $\log(n)$ nodes are investigated (or until there are no more fresh nodes, in which case we are stuck).

By Lemma 9, we may assume that the event in (2.8) happens, and condition on that event. By the choice (2.7) of k, this means that each time a node selects another node to connect to, there are at least $\frac{n}{2k^2}$ nodes to choose from. And each time, there are at most $2\log(n)$ nodes that have already been seen, so each edge has probability at most

$$\frac{4k^2 \log(n)}{n} \tag{2.9}$$

of hitting a node that has already been seen. Hence, the probability that *at least one* of the $2\log(n)$ choices is a failure, is at most

$$2\log(n)\frac{4k^2\log(n)}{n} = \frac{8k^2(\log(n))^2}{n}, \tag{2.10}$$

which tends to 0 as $n \to \infty$. Hence, the probability that *all* choices, up until the outgoing edges of $\log(n)$ nodes are investigated, are successful, tends to 1 as $n \to \infty$.

Hence, we have shown that with probability approaching 1 as $n \to \infty$, we get a connected component with at least $2\log(n)$ nodes. But this is not enough to prove Proposition 3, which asserts a component whose size is *linear* in n.

We can, however, continue the sequential discovery procedure from the $\log(n)$ fresh nodes that we have (assuming that all choices so far have been successful). Let us continue the sequential procedure until the stage S when either a total of $\frac{n}{8k^2}$ nodes have been found (or no fresh nodes remain). Before stage S, each new discovery has, by an analogous argument as that used to establish (2.9), probability at most

$$\frac{n/8k^2}{n/2k^2} = \frac{1}{4}$$

of not being successful. It follows that the sequential discovery procedure starting from the $\log(n)$ nodes until stage S stochastically dominates a $\mathrm{Bin}(2, \frac{3}{4})$ branching process with the same initial number of individuals. Let $\mathcal{A}$ be the event *"the sequential procedure fails to survive until $\frac{n}{8k^2}$ nodes are found"* , let $\mathcal{B}$ be the event *"a $\mathrm{Bin}(2, \frac{3}{4})$ branching process starting with $\log(n)$ individuals dies out"* and let $\mathcal{C}$ be the event *"a $\mathrm{Bin}(2, \frac{3}{4})$ branching process starting with 1 individual dies out"*. We therefore get, conditionally on no failures associated with the first $2\log(n)$ nodes,

$$\mathbf{P}(\mathcal{A}) \leq \mathbf{P}(\mathcal{B}) = \mathbf{P}(\mathcal{C})^{\log(n)} = (1-\alpha)^{\log n} \tag{2.11}$$

where $\alpha > 0$ is the survival probability of a $\mathrm{Bin}(2, \frac{3}{4})$ branching process starting from a single individual (an easy calculation shows that $\alpha = \frac{8}{9}$, but we only need the fact that $\alpha > 0$, which follows from the fact that the offspring distribution has expectation $\frac{3}{2} > 1$). The sum of (2.10) and (2.11) tends to 0 as $n \to \infty$, whence (1.2) follows with $s = \frac{1}{8k^2}$, and we are done. □

Note that by Lemma 9, (2.10) and (2.11) the probability of not having a giant component is $\Theta(n^{-\epsilon})$ for $\epsilon > 0$.

3. CONNECTIVITY

In this section we go on to prove the connectedness of $G^n_{r,c}$ asserted in Theorem 4. We begin by proving the following strengthening of Proposition 3.

Let $\mathcal{D}$ be the event *"every node of $G^n_{r,c}$ is in some s-giant component"*.

PROPOSITION 10. *Fix $r > 0$ and $c \geq 2$. Then there exists a constant $s > 0$ such that*

$$\lim_{n\to\infty} \mathbf{P}(\mathcal{D}) = 1 .$$

Proof: Again, it suffices to consider $c = 2$. As in the previous section, let α denote the survival probability of a $\mathrm{Bin}(2, \frac{3}{4})$ branching process starting from a single individual.

We proceed using the sequential discovery procedure as in the proof of Proposition 3, with the following modification. Instead of initially running it until the outgoing edges of $\log(n)$ nodes have been checked, run it until the outgoing edges of $a\log(n)$ nodes have been checked, where a is a fixed number chosen so that

$$a > \log\left(\frac{1}{1-\alpha}\right) .$$

The estimate in (2.10) then becomes replaced by $\frac{8k^2a^2(\log(n))^2}{n}$. However, since the result we are trying to prove concerns all n points simultaneously, we need to improve on this estimate (which, when multiplied by n, fails to approach 0). To do this, note we can afford to have *one* failed edge during the discovery of the outgoing edges of the first $a\log(n)$ nodes without very much damage (there will still be $a\log(n)$ fresh edges at the end of this search). To estimate the probability that at least two choices fail, note that there are less than $\frac{(a\log(n))^2}{2}$ pairs of times during the procedure at which the choices can fail, and for each such pair the probability of failure in both is at most $\left(\frac{2a\log(n)}{n/2k^2}\right)^2$ (assuming as before the event in Lemma 9). The probability that at least two of the $2a\log(n)$ choices are failures is therefore at most

$$\frac{(a\log(n))^2}{2}\left(\frac{2a\log(n)}{n/2k^2}\right)^2 = \frac{8k^4a^4(\log(n))^4}{n^2}, \quad (3.12)$$

which tends to 0 at a rate which (as we shall see) is fast enough for our purposes.

Again imitating the proof of Proposition 3, we go on to run the sequential discovery procedure until a total of $\frac{n}{8k^2}$ nodes have been found. Let $\mathcal{E}$ be the event *"the sequential procedure fails to survive until $\frac{n}{8k^2}$ nodes are found"*. Since we begin with $a\log(n)$ fresh nodes, the analogue of (2.11) becomes

$$\mathbf{P}(\mathcal{E}) \leq (1-\alpha)^{a\log(n)} = n^{-b} \quad (3.13)$$

where $b = -a\log(1-\alpha)$, and $b > 1$ by the choice of a.

On the event in Lemma 9 (whose probability tends to 1), we can bound the probability that *some* node fails to sit in an s-giant component (with $s = \frac{1}{8k^2}$) by adding the estimates in (3.12) and (3.13) and multiplying by the number of nodes n. This yields

$$\frac{8k^4a^4(\log(n))^4}{n} + n^{1-b} \quad (3.14)$$

which still tends to 0 as $n \to \infty$, so the proof is complete. □

Proof of Theorem 4: As usual, we need only consider the $c = 2$ case. Note that in view of Proposition 10 with the estimate

$$s = \frac{1}{8k^2} \quad (3.15)$$

that comes out of its proof, the only thing that can cause (1.3) to go wrong is if there exists an $\varepsilon > 0$ such that

$$\limsup_{n\to\infty} \mathbf{P}(\mathcal{F}) \geq \varepsilon\,. \quad (3.16)$$

where $\mathcal{F}$ is the event *"$G^n_{r,c}$ contains at least two distinct $\frac{1}{8k^2}$-giant components"*.

Now consider the experiment of generating $G^n_{r,c}$ and then picking two of its nodes at random; let A denote the event that these two nodes end up in the same connected component. By conditioning on the first of these nodes, we see that (3.16) implies that

$$\limsup_{n\to\infty} \mathbf{P}(\neg A) \;\geq\; \frac{\varepsilon}{8k^2}\,.$$

In order to prove the theorem, it therefore suffices to show that

$$\lim_{n\to\infty} \mathbf{P}(\neg A) \;=\; 0\,. \quad (3.17)$$

Let us denote the two nodes chosen at random by v_0 and v_1. By Proposition 10 and the estimate (3.15), we may assume that v_0 is in a connected component of at size least $\frac{n}{8k^2}$. Then, by the pigeonhole principle, at least one of the k^2 subsquares of $[0,1]^2$ introduced in Section 2 contains at least $\frac{n}{8k^4}$ nodes of that connected component. Let us pick such a subsquare and denote it by B.

Next, fix an integer m, and run the sequential discovery procedure starting from the other node v_1, with the following restriction. As soon as an edge fails to lead to a new node, we give up. Assuming this does not happen, we run the procedure until the outgoing edges of exactly $m-1$ nodes have been investigated; this leaves us with exactly m fresh nodes. Let $w_1, \ldots w_m$ denote the fresh nodes after having checked the outgoing edges of $m-1$ nodes in the sequential procedure. Pick one of these vertices, w_i, and denote the subsquare it sits in by $B_{i,0}$. We can then find a sequence of subsquares $B_{i,1}, B_{i,2}, \ldots, B_{i,\ell}$, $\ell \leq 2k$, such that

(i) for each $j \in \{0, 1, \ldots, \ell-1\}$, the subsquares $B_{i,j}$ and $B_{i,j+1}$ are adjacent, and

(ii) $B_{i,\ell} = B$.

Fix such a sequence, and consider the "naked-branch" sequential discovery procedure starting from w_i, and denote the nodes found along this branch by $w_{i,1}, w_{i,2}, \ldots, w_{i,\ell}$. Given the event in Lemma 9 (which we may assume happens), the probability that $w_{i,1}$ ends up in $B_{i,1}$ is at least $\frac{n/2k^2}{n} = \frac{1}{2k^2}$ (due to our choice (2.7) of k). Given that, the conditional probability that $w_{i,2}$ ends up in $B_{i,2}$ is at least $\frac{1}{2k^2}$. And so on. Finally, given that $w_{i,\ell-1}$ is in $B_{i,\ell-1}$, the conditional probability that $w_{i,\ell}$ is in the connected component of v_1, is at least $\frac{n/8k^4}{n} = \frac{1}{8k^4}$. Multiplying these conditional probabilities yields that $w_{i,\ell}$ has probability at least

$$\left(\frac{1}{2k^2}\right)^{\ell-1}\frac{1}{8k^4} \geq \left(\frac{1}{2k^2}\right)^{2k-1}\frac{1}{8k^4} \quad (3.18)$$

of being in the connected component of v_0.

On the event that no checked edges result in failures (which we assume), the m different "naked branches" move independently, so (3.18) implies that the probability that none of them hit the connected component of v_0 is at most

$$\left(1 - \left(\frac{1}{2k-1}\right)^{k-1}\frac{1}{8k^4}\right)^m.$$

We have thus shown that

$$\limsup_{n\to\infty} \mathbf{P}(\neg A) \leq \left(1 - \left(\frac{1}{2k^2}\right)^{2k-1}\frac{1}{8k^4}\right)^m. \quad (3.19)$$

Now, m was arbitrary, and the right hand side of (3.19) can be made as small as we wish by picking m large. Hence (3.17) is established and the proof is complete. □

A careful examination of the estimates of failure probabilities in the proof above show that the probability of not having connectivity is at most $\Theta(n^{-\epsilon})$, for $\epsilon > 0$.

4. THINNED IRRIGATION

In this section, we go on to consider the thinned irrigation graph, and in particular to prove Theorem 8. We begin with the easy part.

Proof of Theorem 8, part (1.6): First construct the (non-thinned) irrigation graph $G^n_{r,c}$. This graph contains at most cn edges, so the nodes have on average degree at most $2c$. It follows that at least $\frac{n}{2}$ nodes have degree at most $4c$. Among $\frac{n}{2}$ such nodes, we can find a subset A of size at least $\frac{n}{2(4c+1)}$, such that no two of the nodes in A share an edge.

When we now delete edges from $G^n_{r,c}$ to obtain the thinned irrigation graph $G^n_{r,c,p}$, each node in A gets all its edges deleted with probability at least $(1-p)^{4c}$, independently for different nodes in A. Hence, the probability that *at least one* node in A becomes isolated in $G^n_{r,c,p}$ is at least

$$1 - \left(1 - (1-p)^{4c}\right)^{\frac{n}{2(4c+1)}} ,$$

which tends to 1 as $n \to \infty$, and (1.6) follows. □

Proof of Theorem 8, part (1.5): Omitted from this extended abstract. □

5. EXPERIMENTS

In this section we perform an empirical evaluation of two algorithms for the broadcast problem, (vertex and edge-based) gossiping and our new irrigating approach. Recall that in the broadcast problem a source is to send a message to all nodes in the network. In the gossip protocol each node that receives the message *for the first time* forwards it to its neighbours with probability p, while irrigating consists of flooding through G^n_{rc}. It is not hard to see that, for any fixed p and r, vertex gossip does deliver the message to all nodes for large n with high probability. An outline of the proof is as follows. For large n every little box of the subdivision of the unit square will contain $\Theta(n)$ nodes. Thus the source will deliver the message to $\Theta(n)$ nodes. Thus every adjacent box will contain $\Theta(pn) = \Theta(n)$ nodes that will try to gossip. With high probability at least one will succeed, reaching $\Theta(n)$ nodes in all adjacent boxes, and so on. Thus, with high probability, every node u in the network is surrounded by $\Theta(n)$ nodes within transmission radius that have received the message and will try to gossip. But then, u too will be reached by the message with high probability. Thus, both gossip and irrigating are reliable in the limit, i.e. for large n. It becomes then interesting to see if an experimental evaluation can sharpen our analysis.

The main conclusion of our experimental evaluation is that indeed the two processes differ significantly:

- It is more likely that G^n_{rc} be connected than gossip reaches all nodes in the network. Thus irrigating is a more reliable protocol than gossiping.
- The number of edges along which a message is sent by gossiping is much greater than the number of edges of G^n_{rc}. Thus, in all situations where the cost of a message can be charged to the edge connecting the two nodes, irrigating is a much more efficient protocol. In particular this applies to sensor networks, for which the cost of sending a message is roughly equal to that of receiving it.
- Graphs generated using "the power of choice" retain the same good connectivity properties of G^n_{rc}, while their maximum degree drops significantly. The graphs also tend to be more regular.

We have performed our experiments by distributing points at random over a 1000×1000 points grid, varying the various parameters such as number of nodes and transmission radius between intervals that are realistic for wireless network applications. The charts we show exemplify the general trends. Essentially we are considering three protocols: flooding, corresponding to G^n_r, irrigating, corresponding to G^n_{rc} for various values of c, and gossiping.

Perhaps the most telling figure is Figure 5.1 (a). On the y-axis we report P, the relative size of the largest connected component of the graph generated by the various processes G^n_r, G^n_{rc} and gossip (with $p = 0.5$). When $P = 1$ the graph is connected and thus the corresponding protocol reaches all nodes. The striking feature is that the curves of G^n_r and G^n_{r4} overlap, meaning that G^n_{r4} is connected as soon as the full visibility graph is and thus that irrigating is as reliable as it can be. In particular, irrigating with G^n_{r4} is as reliable as flooding, but much cheaper. In contrast gossiping reaches a very small fraction of the nodes when the underlying graph becomes connected ($n = 1000$). Note also that G^n_{r2} and G^n_{r3} work quite well.

Figure 5.1 (b) compares the probability that irrigating (with G^n_{r4}) reaches all nodes to the probability of gossiping doing the same (with $p = \frac{1}{2}$). As predicted, gossiping does reach all nodes for large n, but it does so much later than the instant when the visibility graph becomes connected. Irrigating, as observed, on the other hand reaches all nodes as soon as the visibility graph becomes connected (the curves of G^n_{r4} and G^n_r are superimposed).

Figure 5.2 exemplifies the second main conclusion, namely that gossiping is much more expensive than irrigating in terms of edges used. On the y-axis we report the relative number of edges used by the two protocols (number of edges used divided by the total number of edges of the visibility graph). Clearly, as the number of nodes grows G^n_r becomes denser and therefore the proportion of edges used by the G^n_{rc} must go to zero. This prediction is fully confirmed by the data. In contrast, a quick back of the envelope calculation shows that gossiping with $p = \frac{1}{2}$ is bound to use 75% of the edges. As seen, for large n gossiping reaches all nodes. If we consider an edge uv of the visibility graph this edge *will not* be used by gossiping if both u and v fail to gossip, an event that happens with probability p^2. Thus we expect that a p^2 fraction of the edges will not be used and concentration of measure says that the actual number will be very close to this. Figure 5.2 shows this clearly: the number of edge used by gossiping is the predicted value– 75% of the total.

The effect of the "power-of-choice" is demonstrated in Figures 5.3 and 5.4. While connectivity remains unchanged, there is marked improvement in the degree profile of the networks: the maximum degree is reduced by a constant factor even at moderate sized networks, and the full degree distribution is less extreme and it is more sharply concentrated around the average value implying a more regular graph.

Finally, in Figure 5.5 we test experimentally to what extent the connectivity behaviour of G^n_{r4} and G^n_r coincide. On the z-axis we report the relative size of the largest connected component of the two graphs as a function of the number of points (x-axis) and of the transmission radius r (y-axis). Surprisingly the two surfaces coincide perfectly. It would be quite interesting to establish an analytical result in this direction.

6. REFERENCES

[1] S. Asmussen and H. Hering, *Branching Processes*, Birkhäuser, Boston, 1983.

[2] Y.Azar, A.Z.Broder,A.R.Karlin and E.Upfal, *Balanced allocations*. SIAM Journal on Computing, Volume 29, Number 1, pp. 180-200, 1999

[3] G. Badishi, I. Keidar and R. Melamed, *Towards Survivability of Application-Level Multicast*, Proceedings of S.O.S. Survivability: Obstacles and

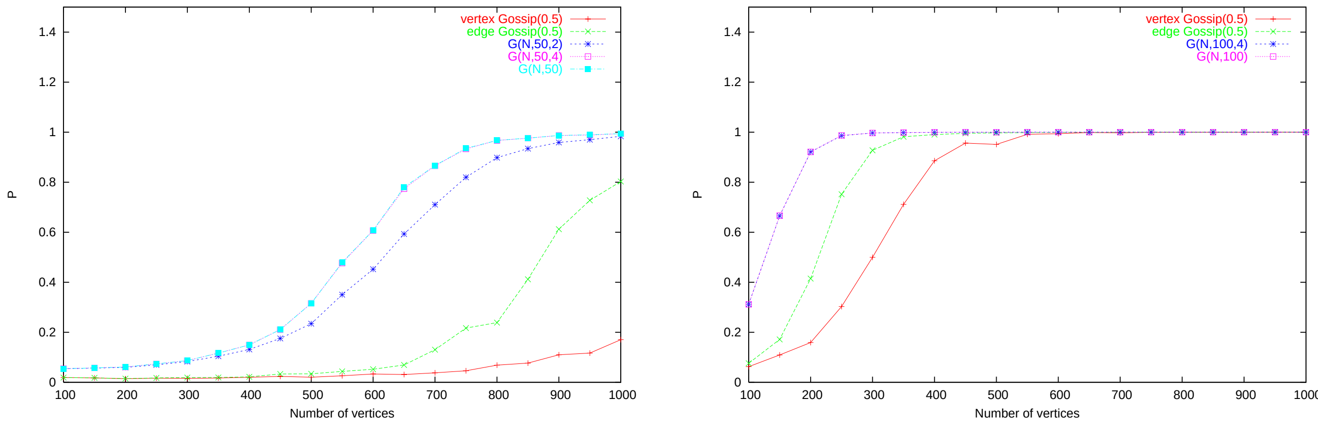

(a) (b)

Figure 5.1: Grid dimension is 1000×1000. Each point is the average over 100 runs. On the y-axis P is the ratio between the number of vertices in the largest connected component and the total number of edges. Thus, when $P = 1$ the graph is connected. On the left we used a transmission radius $r = 50$. Note that while all G^n_{rc}s perform well, G^n_{r4} behaves exactly like G^n_r. On the contrary, the performance of gossip is quite poor. On the right we used a transmission radius $r = 100$. The plot shows that for large n gossip too reaches all nodes, but it does so much later than irrigating. Here we used irrigating with G^n_{r4}.

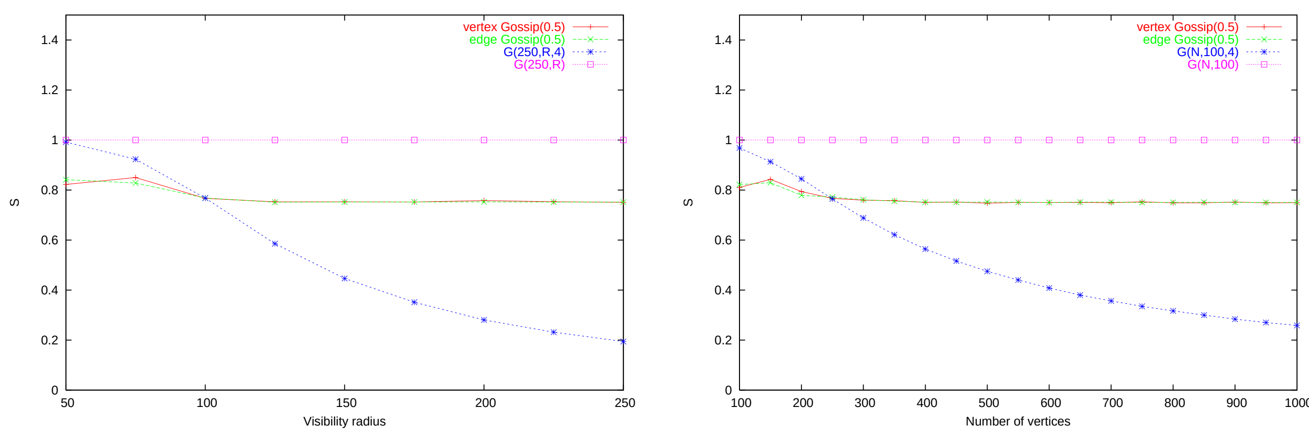

(a) (b)

Figure 5.2: Grid dimension is 1000×1000. Each point is the average over 100 runs. On the y-axis S is the ratio between the number of edges in the largest connected component and the total number of edges. Thus, the higher the S the more expensive the protocol in terms of communication. While the number of edges used by irrigating goes to zero, gossiping uses a full 75% of the edges (here $p = \frac{1}{2}$).

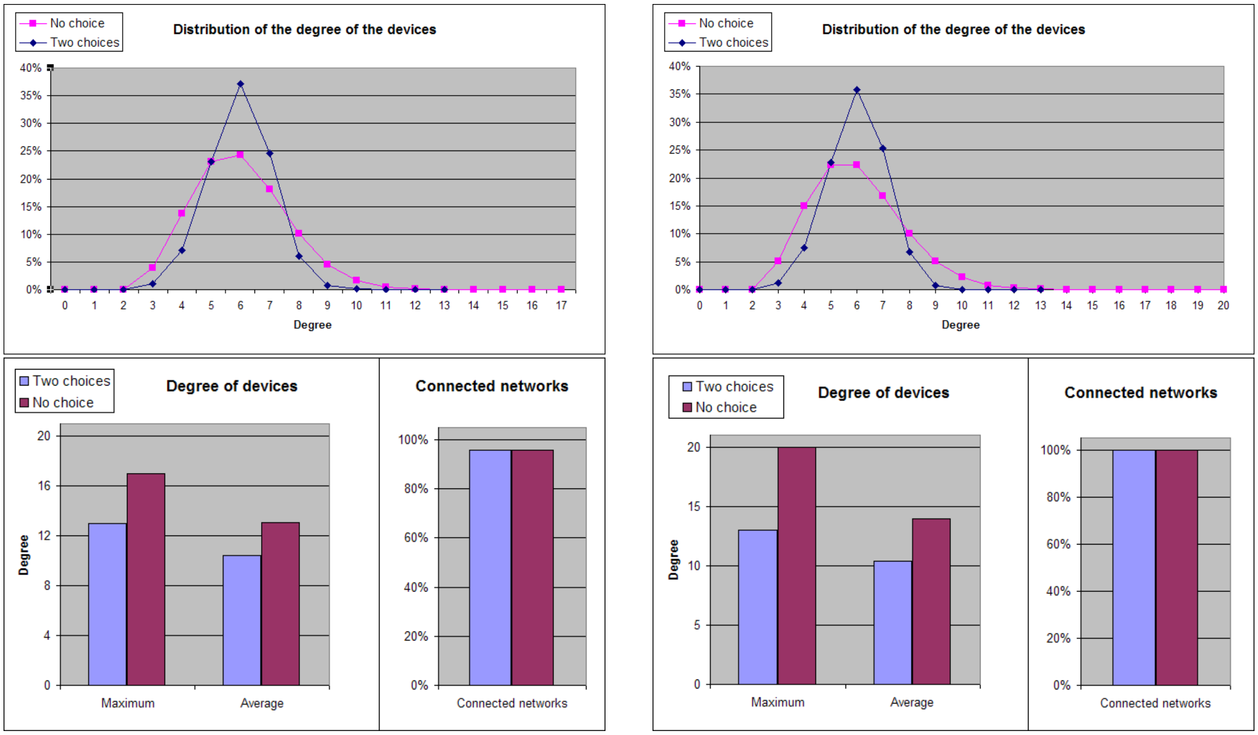

Figure 5.3: The effect of the power of choice. The figure compares the degree profile, the maximum and average degrees, and the probability of connectivity of $G^{2000}_{.07,2,2}$ and $G^{2000}_{.15,2,2}$ with, respectively, $G^{2000}_{.07,2}$ and $G^{2000}_{.15,2}$. While connectivity is preserved, the maximum degree drops and the graph becomes more regular. Each dot is the average over 10,000 runs.

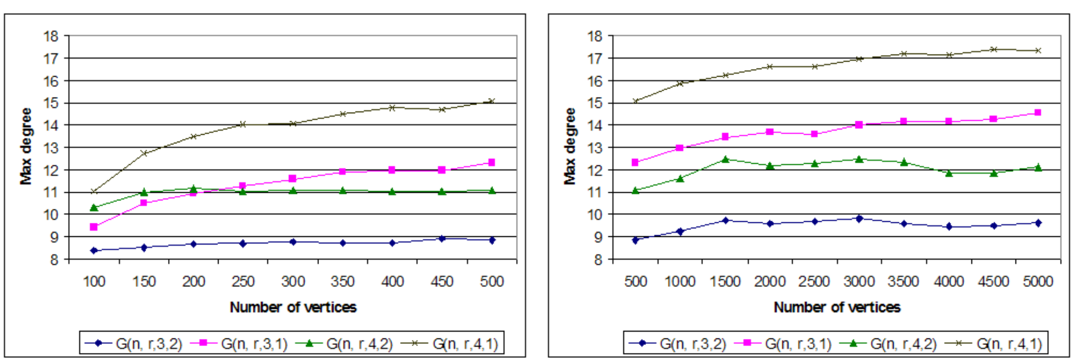

Figure 5.4: The effect of the power of choice. The figure compares the maximum degree of $G^n_{r,c,d}$ with that of $G^n_{r,c}$ (denoted here as $G^n_{r,c,1}$) for $c = 3, 4$ and $d = 2$. The maximum degree is always substantially smaller. Each dot is the average over 1,000 runs.

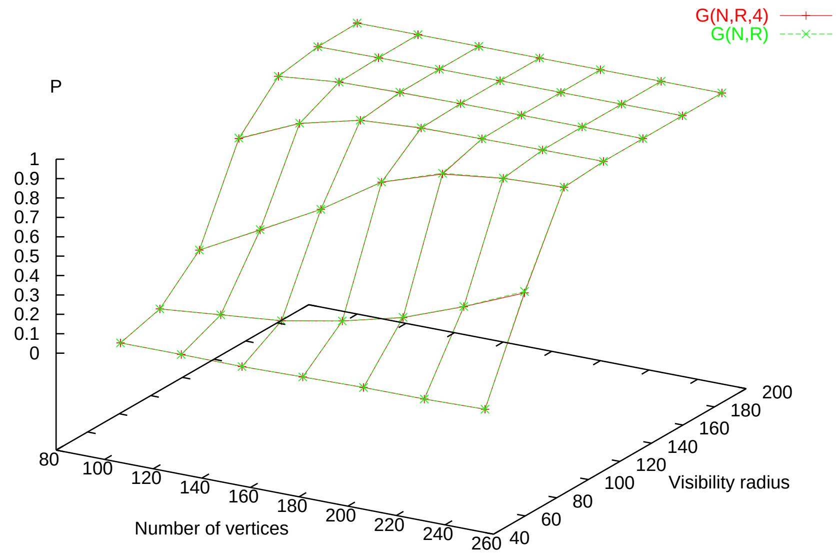

Figure 5.5: Grid dimension is 1000×1000. Each point is the average over 100 runs. On the z-axis the relative size of the largest connected component of G^n_r and G^n_{r4} is reported as a function of the number of nodes (x-axis) and of the visibility radius r (y-axis). The two surfaces coincide.

Solutions 2nd Bertinoro Workshop on Future Directions in Distributed Computing, `http://www.cs.utexas.edu/users/lorenzo/sos`

[4] G. Badishi, I. Keidar, and A. Sasson. *Exposing and elim- inating vulnerabilities to denial of service attacks in secure gossip-based multicast*. In The International Conference on Dependable Systems and Networks (DSN), 2004. To appear.

[5] S.Basagni, R.Bruno and C.Petrioli. *A Performance Comparison of Scatternet Formation Protocols for Networks of Bluetooth Devices*, Proceedings of the IEEE International Conference on Pervasive Computing and Communications, PerCom 2003, 341–350.

[6] D. Dubhashi, O.Häggström, G. Mambrini, A. Panconesi and C. Petrioli, *Blue Pleiades, a new solution for device discovery and scatternet formation in multi-hop Bluetooth networks*, accepted for publication in ACM Wireless Networks.

[7] I. Gupta, K.P. Birman, and R. van Renesse, *Fighting fire with fire: using randomized gossip to combat stochastic scalability limits*, Quality and reliability Engineering International, 18, 2002.

[8] Z. J. Hass, J.Y. halpern and L. Li, *Gossip-Based Ad Hoc Routing* IEEE Infocom, June 2002.

[9] S. Basagni, I. Chlamtac, V.R. Syrotiuk and B.A. Woodward, *A distance routing effect algorithm for mobility (DREAM)* CM/IEEE Intl. Conference on Mobile Computing and Networking (MobiCom), 1998.

[10] K.P. Birman, M. Hayden, O. Ozkasap, Z. Xiao, M. Budiu and Y. Minsky, *Bimodal multicast*, ACM Transactions on Computer systems, 17(2), 1999.

[11] P T. Eugster, R. Guerraoui, A-M Kermarrec and L. Massoulié, *Epidemic Information Dissemination in Distributed Systems*, IEEE Computer, May 2004.

[12] P T. Eugster, R. Guerraoui, S.B. Handurukande and A-M Kermarrec *Lightweight Probabilistic Broadcast*, ACM Transactions on Computer Systems, 21:4, Nov. 2003.

[13] A-M Kermarrec, L. Massoulié and A.J. Ganesh, *Probabilistic Reliable Dissemination in large-Scale Systems*, IEEE Transactions on Parallel and Distributed Systems, 14:3, March 2003.

[14] T.E. Harris, *The Theory of Branching Processes*, Springer, Berlin, 1963.

[15] T. Lindvall, *Lectures on the Coupling Method*, Wiley, New York, 1992.

[16] R. Meester and R. Roy, *Continuum Percolation*, Cambridge University Press, 1996.

[17] M. Mitzenmacher, A. Richa, and R. Sitaraman. *The Power of Two Random Choices: A Survey of Techniques and Results* Book chapter, in Handbook of Randomized Computing: volume 1, edited by P. Pardalos, S. Rajasekaran, and J. Rolim, pp. 255-312.

[18] L. Orecchia, A. Panconesi, C. Petrioli and A. Vitaletti. *Localized techniques for broadcasting in wireless sensor networks*. Proceedings of ACM Dial-M POMC 2004, the Joint Workshop on Foundations of Mobile Computing.

[19] A. Panconesi and J. Radhakrishnan. *Expansion properties of (secure) wireless networks*, Proceedings of SPAA 2004, Barcelona, Spain.

[20] C.Petrioli, S.Basagni and I.Chlamtac. *BlueMesh: Degree-Constrained Multihop Scatternet Formation for*

Bluetooth Networks, ACM/Kluwer Journal on Special Topics in Mobile Networking and Applications (MONET), Special Issue on Advances in Research of Wireless Personal Area Networking and Bluetooth Enabled Networks (G. Zaruba and P. Johansson, eds.), 2003

[21] G.H. Hardy, J.E. Littlewood and G. Polya, *Inequalities* Cambridge U. Press, London, 1952.

[22] A.W. Marshall and I. Olkin, *Inequalities: Theory of Majorization and its Applications*, Academic Press, NY. 1979.

[23] S. Y. Ni, Y. C. Tseng, Y. S Chen, and J. P. Sheu. *The broadcast storm problem in a mobile ad hoc network*. In *Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking (MobiCOM '99)* (Seattle, Washington, August 15–20 1999), pp. 151–162.

[24] W. Peng, and X. Lu. *On the reduction of broadcast redundancy in mobile ad hoc networks*. In *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2000)* (Boston, Massachusetts, August 11, 2000), pp. 129–130.

[25] B. Williams, and T. Camp. *Comparison of broadcasting techniques for mobile ad hoc networks*. In *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2002)* (Lausanne, Switzerland, June 9–11, 2002), pp. 194–205.

[26] Y. Yi, M. Gerla, and T. J. Kwon. *Efficient flooding in ad hoc networks: a comparative performance study* . In *Proceedings of the IEEE International Conference on Communications (ICC 2003)* (Anchorage, Alaska, May 11–15 2003).

[27] Z. Hass and M. Pearlman, *The performance of query control schemes for the zone routing protocol*,in Proceedings ACM SIGCOMM, Aug. 1998.

[28] Y.B. Ko and N.H. Vaidya, *Location-aided Routing (LAR) in mobile ad-hoc networks*, ACM/IEEE Intl. Conference on Mobile Computing and Networking (MobiCom), 1998.

[29] A. Marcucci, M. Nati, C. Petrioli and A. Vitaletti, *Directed Diffusion Light: Low Overhead Data Dissemination, in Wireless Sensor Networks* , an IEEE Vehicular Technology Conference (IEEE VTC 2005), Stoccolma May 2005.

[30] V. Paruchuri, A. Durresi and R. jain, *Optimal Flooding Protocol for Routing in Ad-hoc Networks*, submitted to Computer Networks.

[31] C.E. Perkins, and E.M. Roper, *Ad hoc on demand distance vector routing* in IEEE Workshop on Mobile Systems and Applications, Feb. 1999.

[32] W. Vogels, R. van Renesse and K. Birman, *The power of epidemics: robust communication for large-scale distributed systems* ACM SIGCOMM Computer Communication Review archive Volume 33 , Issue 1 (January 2003)