Percolation for Power Control
(Abstract)

Emilio De Santis
Department of Mathematics
University La Sapienza of Rome

Fabrizio Grandoni and Alessandro Panconesi
Department of Computer Science
University La Sapienza of Rome

May 6, 2005

Loosely speaking, the problem of power control for wireless communication networks is that of adjusting the transmission power of the devices in a network in such a way that two conflicting objectives are met. On the one hand the communication radius of each node should be as small as possible, for power consumption grows (at least) quadratically with the radius. On the other, the radius should be large enough to ensure global connectivity or the emergence of a giant component. Energy efficient protocols ensuring good connectivity properties are particularly desirable in the case of ad-hoc wireless communication, for these networks have severe resource constraints. This problem can be formulated mathematically in different (not necessarily equivalent) ways. One of the most interesting models is the so-called *nearest-neighbor* model. Here we have N points uniformly distributed at random in the unit square. Each node u in the network decides a number $k(u)$ and sets its transmission radius in order to reach the closest $k(u)$ neighbors. Links are then established according to different rules. Some authors consider the resulting oriented network. If node u is within the ball of radius $r(v)$ centered at v, then there is a direct link $v \to u$, meaning that v can send a message to u but the reverse is not necessarily true [3]. In this case strong connectivity is analyzed. Others remove the orientations altogether [1]. In this paper uv is an undirected edge if u is inside v's ball, and viceversa. This model is the most realistic from the point of view of wireless applications since the communication primitives of standards such as IEEE 802.11 and Bluetooth rely on ack messages, and therefore a communication link exists only when the two nodes are both within radius of each other. In any case the three models are equivalent from the point of view of the question studied in this paper.

Assuming N points uniformly distributed within a region is the standard assumption that is made when analyzing this kind of problems. In some cases this is quite realistic. Sensor networks for instance are small, inexpensive devices that can be deployed massively within a certain region, often at will. But the point is that proofs and properties proved in this model oftentimes do not really need the uniformity assumption, but something weaker that makes the results and the properties more widely applicable than originally envisaged.

The question we consider in this paper is, how large should $k(u)$ be *on average* for connectivity and/or the emergence of a giant component?

Xue and Kumar [1] show that, if $k(u) = C$ for all u, then C must grow like $\log N$. Note that this does not imply that the average degree must be $\Omega(\log N)$, since values for the $k(u)$ need not be equal. Indeed, Kucera [3] shows that on average $k(u)$ can be constant, even though some $k(u)$ will have to grow like $\log N$. This result however is existential in flavor, since the protocol used in [3], although distributed, is not local, for a node can explore a linear-size component of the network and this requires linearly many communication rounds. In wireless networks, and distributed computing in general, it is crucial to bound the number of communication rounds of the protocol. In this paper we show that an exponential speed-up is possible by exhibiting a distributed protocol for power control that uses $O(\log^4 N)$ communication rounds. The average $k(u)$ is constant as in Kucera's protocol and the maximum $k(u)$ is $O(\log N)$.

To prove the result we establish along the way a useful fact. Namely, that there is a protocol to generate a giant component where each node has constant degree, and this protocol uses no communication at all. Moreover the giant component is uniformly spread within the area. This is a useful property because the main application envisaged so far for sensor networks is monitoring of a certain area by means of a connected network, and a giant component is enough for this purpose. The result we prove is a typical percolation result. We prove that there is a universal threshold K, that is, independent of N, such that a giant component arises if $k(u) \geq K$ for all u.

From the technical point of view we reduce our problem to bond percolation. Well-known results of percolation theory show that for independent, bond percolation a unique infinite component arises with probability one as the bond probability p exceeds a threshold $p_c \in (0, 1)$. In a finite box, we will have a unique giant component, while the remaining components will have size $O(\log^2 N)$. To use a metaphor, we will have land (sites that are on) and water (sites that are off). Land is organized in a unique super-continent of linear size and a set of very small islands, of poly-logarithmic size. These islands lie in seas (water between the border and the super-continent) or in lakes inside the super-continent. It can be proven that each sea and each lake, together with the islands they contain, are of poly-logarithmic size. All this happens with probability going to one as long as N goes to infinity, provided that the bond probability exceeds a threshold.

In the original graph this translates to a probabilistic test for membership in the giant component. If a node is in a component of size $\Omega(\log^3 N)$, then it can safely assume to be part of the giant component. Otherwise it will increase its transmission radius to reach a new node, until, within $O(\log N)$ such iterations, it will be connected with the super-continent.

References

[1] Feng Xue and P. R. Kumar, The number of neighbors needed for connectivity of wireless networks, Wireless Networks, Vol 10, No. 2, pp.169-181, 2004.

[2] Grimmett, G. Percolation. New York: Springer-Verlag, 1989.

[3] L.Kucera, Low Degree Connectivity in Ad-hoc Networks, DELIS Tech Report.