

GEOMBINATORICS OF “SMART DUST”

by Edward Vidal¹, Luc Longpré²,
Vladik Kreinovich², and Huang Haitao²

¹Army Research Laboratories
White Sands, NM, USA, evidal@arl.mil

²Department of Computer Science
University of Texas at El Paso
El Paso, TX 79968, USA
emails {longpre, vladik}@cs.utep.edu

Abstract. *Smart Dust is a collection of small sensor-equipped leaves which send their information to two or more receivers. When a receiver gets a signal from a sensor, it can determine the direction from which this signal came. By combining the directions from two different receivers, we can determine the 3-D locations of all the leaves, and thus, transform their sensor readings into a 3-D picture of the corresponding parameters (temperature, moisture, etc.). The more leaves we send, the more information we gather. However, since the direction can only be measured with a certain accuracy, when we send too many leaves, we lose the ability to match their directions and thus, we can no longer reconstruct the leaves’ 3-D locations. Thus, there is the optimal number of leaves for which we can get the largest amount of information. Determining this optimal number is an open problem.*

What is “Smart Dust”. Smart Dust is a project developed by the University of California at Berkeley under the US DARPA funding [Kahn et al. 1999], [Pister et al. 1999].

In this project, small surfaces shaped like maple leaves are equipped with temperature and moisture sensors; each particle costs about \$30. A small automatic 8-in simple plane lifts a bunch of these leaves up and throws them down. The leaves slowly descend and as they descend they send signals to Earth-based receivers.

One of the potential future applications of this system is to trace wind profiles in the Bay area; it is important for the US Environmental Protection Agency.

Weather application of Smart Dust: problems. Since the leaves which form the Smart Dust are spread around the 3-D zone, they provide us with a unique opportunity to measure weather parameters (temperature, moisture, wind, etc.) in different points within this zone and thus, to create a 3-D weather map.

With respect to measuring temperature and moisture, the main difficulty of creating such a 3-D weather map is that when we receive a signal from a leaf, we know the *direction* from which we received this signal, but we do not know the *distance* to the location of this leaf, and therefore, we do not know the exact 3-D position of a point where the measurements were made.

Similarly, since the leaves do not measure the wind velocity or direction directly, a natural indirect way to determine these parameters is to trace how the location of the leaves change in time. For that, we also need to know their exact 3-D locations (we can also use Doppler measurements of leaves' velocity).

Natural solution to the problem: use two or more receivers. Since by using a single receiver, we can only determine the direction from which the leaf is sending this information but not the exact 3-D location of the leaf, a natural idea is to use two or more receivers.

If we know the exact direction from two different receivers, then we know, for each receiver, the straight line on which this leaf is located, and thus, we can, in almost all cases, uniquely determine the 3-D location of the leaf as the unique point which is the intersection of the corresponding two straight lines. The only case when we cannot uniquely determine this location is when these two lines coincide, i.e., when the leaf is located exactly on the line which connects the two receivers. (This is a rare possibility, but if we want to have a unique reconstruction for all the leaves, then we need to add the third receiver; this receiver will lead to a guaranteed uniqueness, and it will also increase the accuracy with which we measure the leaves' locations.)

An alternative solution would be, for a leaf, to pick the signals from several different beams, use this comparison to determine its exact coordinates (like in GPS), and transmit these coordinates to-

gether with its readings. Unfortunately, this would require adding a lot of sophisticated equipment to the leaf, and it is still not even clear how to place the existing equipment within the required size parameters.

Related problem: matching signals coming from the same leaf to different receivers. When we use two receivers, then for each receiver, we get a lot of signals from different leaves. To process this information, we must match the signals coming from the same leaf to two different receivers.

We can try to match leaves which send the exact same sensor information, but it is possible that two nearby locations have the same temperature and moisture.

For this match, it is thus beneficial to assign a unique ID to each leaf, an ID which is transmitted together with the sensor information, so that we will be able to trace individual leaves. However, as we have mentioned, at present, no additions to the leaves are possible. Therefore, we must match the leaves without such ID's.

In principle, the matching problem is solvable. Let R_1 and R_2 be 3-D locations of receivers, and L_1, \dots, L_n be 3-D locations of leaves. For each leave L_i , the first receiver detects the direction to L_i ; based on this information, we can conclude that this sensor is located on the straight-line ray $R_1 \rightarrow L_i$ which goes from this receiver R_1 in the observed direction.

The second receiver R_2 also detects the directions to different leaves. We can describe these directions by placing a plane ("screen") S near the receiver and describing each direction by the unique point of intersection S_j between the ray $R_2 \rightarrow L_j$ and this screen S .

On the screen plane, the projections of the points from each ray $R_1 \rightarrow L_i$ form a 1-D ray r_i which all start at the same point P – projection of R_1 on this plane S . So, we have:

- n points S_1, \dots, S_n which reflects the directions from this receiver, and
- n rays r_i starting at P .

Rays describe results of the first receiver, points describe the results of the second receiver. To find the 3-D location of each leaf, we

must therefore find out which measurements correspond to which, i.e., we must put points S_i in 1-to-1 correspondence with rays r_j .

If a point S_i describes the same leaf as the ray r , then the point belongs to the ray. For almost all configurations $R_1, R_2, L_1, \dots, L_n$, a point S_i does not belong to the ray r_j if $i \neq j$. Thus, we can match each point S_i with the unique ray to which this point belongs.

An asymptotically optimal matching algorithm. If we simply check, for each of n points S_i , whether it belongs to the 1st, 2nd, \dots , n^{th} ray, then in the worst case, we will need n^2 checks. For a large number n of leaves, this will mean a lot of computations.

To decrease the computation time, we can use, in S , polar coordinates with a center in P . Then, S_i belongs to the ray r_j if and only if their angles coincide. So, to match points with rays, we sort the rays by the angle (which takes $n \cdot \log(n)$ time; see, e.g., [Cormen et al. 1994]), and then, for each of n points S_i , we use the binary search ($\log(n)$ steps) to find the ray with the same angle. The total time is thus $2n \cdot \log(n) \ll n^2$.

Open problems. The above simple geometric considerations assume that we can measure the exact direction to each leaf. In this case, the more leaves we send, the more points we cover by our measurements and thus, the better the resulting 3-D weather description.

In reality, we can only measure this direction with a certain accuracy $\varepsilon > 0$. So, if the measured direction to the leaf corresponds to a point \tilde{S}_i , the (unknown) point S_i which corresponds to the actual direction to this leaf may lie anywhere within a certain distance from \tilde{S}_i ; in other words, the only information that we have about this point is that lies within a *disk* D_i with a center in \tilde{S}_i .

Similarly, on this screen S , the directions from the leaves to the first receiver are described not by rays r_i , but by *sectors* s_i bounded by two close rays which start at P . In this case, we match a disk D_i and a sector s_j when $D_i \cap s_j \neq \emptyset$.

The more leaves we take, the larger the area covered by the corresponding sectors. For a certain number m of leaves, these sectors will cover the whole screen S ($s_1 \cup \dots \cup s_m = S$). Thus,

if we add one more leaf, we will not be able to match it properly, because the corresponding disk D_{m+1} will intersect not only with its corresponding sector s_{m+1} , but also – since

$$D_{m+1} \subseteq S = s_1 \cup \dots \cup s_m$$

– with a sector s_i ($1 \leq i \leq m$) corresponding to one of the previous m leaves.

Thus, if we take measurement uncertainty into consideration, adding new leaves does not necessarily make the resulting 3-D picture better: if we add too many leaves, we lose the ability to match the signals on two receivers and thus, we do not get any 3-D picture at all. So, if we start with a single leaf and add one leaf at a time, then at first, we get better and better pictures, but after a certain number of leaves, we reach an optimum after which adding further leaves would only decrease the resulting number of matched measurements.

A natural question is: for a given measurement accuracy ε , what is the optimal number of leaves? For example, if we assume that n leaves are uniformly distributed in a given 3-D area, for what n is the expected number of matched leaves the largest possible?

Since the relative size of each sector is $\sim 2\varepsilon$, the leaves start covering the entire plane when $n \cdot 2\varepsilon \approx 2\pi$. So, intuitively, we expect that this optimal number is $n \sim 1/\varepsilon$. It is desirable to confirm (or disprove) this intuitive estimate, and to get more accurate results.

Acknowledgments. This work was supported in part by NASA under cooperative agreement NCC5-209, by NSF grant No. DUE-9750858, by United Space Alliance, grant No. NAS 9-20000 (PWO C0C67713A6), by the Future Aerospace Science and Technology Program (FAST) Center for Structural Integrity of Aerospace Systems, effort sponsored by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grant number F49620-95-1-0518, and by the National Security Agency under Grant No. MDA 904-98-1-0561.

References

Th. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to algorithms*, MIT Press, Cambridge, MA, and Mc-Graw Hill Co., N.Y., 1994.

J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Mobile networking for Smart Dust", *ACM/IEEE Int'l Conference on Mobile Computing and Networking MobiCom'99*, Seattle, WA, August 17–19, 1999.

K. S. J. Pister, J. M. Kahn, and B. E. Boser, "Smart dust: wireless networks of millimeter-scale sensor nodes", *Electronics Research Laboratory (ERL) Research Summary*, 1999.