Assembling Sensor Networks

Jeffrey V. Nickerson
Stevens Institute of Technology
jnickerson@stevens.edu

Abstract

Sometimes sensors need to be assembled in response to an emergency. Such assemblage might take a variety of forms. At one extreme, sensors might converge on a single point. At the opposite extreme, sensors might fill in the position of a lattice. This paper looks at solutions in between these two extremes, in which sensors converge on several different locations. Through analysis and simulations, we evaluate the effectiveness of alternative strategies. We find there is a tradeoff: multiple rally point heuristics converge faster, at the expense of a more fragmented network. We discuss ways of bridging clustered networks. This work has pragmatic implications for emergency responders: in situations of high density, single rally point solutions are superior, and in situations of low density time-to-converge trades off with network integration.

1. Introduction

In emergencies, we may want to assemble a sensor network to understand what has happened. For example, if a bomb goes off in a city, we might want to know what the bomb is made of.

But if something big has happened, coordinated movement will be difficult, as normal communication structures will have failed, and normal radio channels that survive will be flooded.

Thus, one possible way of quickly assembling a sensor system is to create a network from a set of units that might be randomly deployed around the city. With current technology, such units can be built into police cars and fire trucks. With future technology, autonomous robotic vehicles could serve this assembly function.

Most sensor literature sensibly assumes that sensors don't move. In contrast, mobile ad hoc network literature assumes that nodes are always moving. But in many situations, a hybrid is called for. Before an emergency, we don't know where to deploy a sensor network. Once an emergency strikes, we need to deploy, and to do so the sensor units will need to be transported. Once the sensor network is deployed, the sensors can remain in place: the network becomes the kind that is normally analyzed in the sensor literature.

The broad question we ask is: how can such sensor networks be deployed? In previous research we have looked at some aspects of this problem. Here we specifically look at the problem of deploying sensors around multiple rally points. The contribution of this paper lies in the identification of tradeoffs between time to deploy and fragmentation of the network. In addition, we discuss how networks which are fragmented might be bridged. Our work has pragmatic implications for those involved in emergency response planning, and has theoretical implications for the deployment and integration of sensor networks.

We proceed in the following way. First we describe related work. Then we describe the results of a series of simulations, presenting diagrams of configurations as well as graphs relating the time-to-convene with the density of the sensors. We then discuss how separated networks might be bridged

2. Related Work

Since sensor networks use ad hoc networks to communicate information, research in ad hoc networks is relevant [1, 2]. However, ad hoc network research usually assumes the nodes move independently of each other to random destinations. In contrast, sensors in a network cooperate toward meeting a goal. If sensors are given mobility, their movement will not be to random destinations, but instead will be coordinated toward agreed-upon locations. Thus, the research that looks at intentional movement in ad hoc networks is of interest [3-5] .

Some have suggested that sensors might move once into position and then stop [6-8]. This strikes us as a good idea. Even if one is doubtful about the economics of building moving sensor chips, sensors always need to be transported before deployment. Therefore

0-7695-2755-8/07 $20.00
1530-1605/07 $20.00

consideration of initial movement is an important part of any sensor system: systems broadly construed include the processes that establish them physically.

This paper follows closely on previous work in which we looked at the use of rally points to reconvene after an emergency [9-12]. In that work, we showed that the concept of a rally point [13] – an agreed place to congregate – changes when ad hoc radios are present. Nodes possessing such radios can move toward a rally point, but can stop when they come in radio contact of an already connected node. This forms a branching structure emanating from the rally point. We showed that the time for all units to converge is

$$t = Max\left(\frac{w}{a} - \frac{c\sqrt{n}}{b}, 0\right) \quad (1)$$

where w is the width of the field, a is a constant related to the grid type, c is the radio coverage area, n is the number of sensors, and b is a scaling constant [11]. The equation fits well when the density of coverage is low, but does not work so well when the density is high, because the configurations formed are then most dependent on the initial configuration of nodes – the structure doesn't get a chance to assemble itself. The first term reflects the size of the field the sensors are dispersed in, and the second term reflects the size (as a radius of gyration) of the connected configuration. Thus, the time for the team to converge is bound by the time it takes for a unit to intersect with the growing configuration.

Equation (1) is related to phenomena which have been investigated in the physics literature: diffusion limited aggregates [14, 15]. These structures have characteristic fractal dimensions, which in turn make predictions about convergence time possible. Some have observed the similarities of these patterns to urban pedestrian movement and economic growth [16, 17]. From a broader perspective, our work is related to robotic assembly [18, 19]: the sensor network is assembled .

In order to make these ideas clearer, we provide two illustrations.

Figures 1 and 2 show initial randomly distributed units in red, with lines to the locations of the points after their movement is completed. Figure 2 corresponds to a discovery heuristic discussed in [11] which takes advantage of software radio to sense further than the normal radius of radio communication; one can see that the resulting configuration reaches out further, which means that units converge faster using this heuristic.

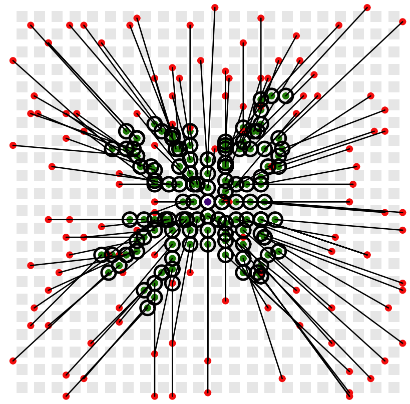

Figure 1. Smaller circles indicate initial positions, connected by lines to final positions.

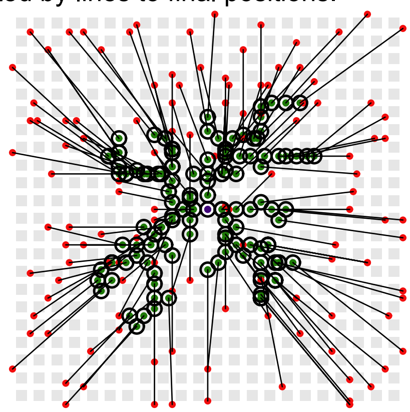

Figure 2. As with Figure 1, except using an extended discovery range.

3. The problem

Imagine a set of responders spread over an area, handling routine business. Once in a great while an event will occur which will disrupt infrastructure. For example, a terrorist bombing, a hurricane, an earthquake, or a blackout will disrupt existing ways of communicating, and, at the same time, create a need for responders to get back into communication in order to plan. Recent analyses of the response to the attack of Sept 11, 2001 and the response to Hurricane Katrina document the confusion and inability to communicate which engulfed search-and-rescue teams [20, 21].

This problem is related to sensing in the following way. In the aftermath of a disaster, communication is important for planning a response, and planning is in turn contingent on understanding the situation. For example, in the event of an explosion, responders want to know if the attack has chemical, biological, or radioactive components to it. In the event of a natural disaster, responders will want to understand water levels, structural integrity of levees, and other conditions which might hinder or help response procedures.

Thus one wants to form a communicating network to understand the situation, as well as to control the response.

Our assumption is that disasters will strike in a way that produces unexpected results, including the destruction of infrastructure: therefore, response procedures need to be established in anticipation of emergencies.

In previous work, we pointed out that responders often use the concept of a rally point [9, 13]: they decide where they will converge if normal forms of communications break.

Usually this is thought of as a singular point. What if we assigned many? There are a few ways this might work. At the extreme, let us assume that responders have radios that can consistently talk only 100 meters in any direction in a city. This is a small radius, chosen to make an argument, but there is a reason to be pessimistic about radio range in city environments during emergencies. Reflections from buildings dissipate signal strength: the strength is about $\frac{1}{d^4}$, d representing distance [22]. Communication around corners is severely restricted, and propagation through modern concrete and steel buildings is very low [22]. In addition, in emergencies, dust and weather will reduce signal strength. Furthermore, power will need to be conserved.

Then, let us assume we picked a point on the grid ahead of time, one point for each unit. Everyone knows where to go, and the end result forms a lattice.

However, this may not make much sense. First of all, since the responders' positions vary, some responders might have to move a long way to get to their assigned point. If everyone tries to go to the closest point, it is likely that several may converge on one point, leaving other parts of the lattice vacant. In addition, some responders may be trapped, and unable to fill in their part of the grid. Therefore it is likely the grid would be sparsely and unevenly populated.

Then, at the other extreme is the model of a single rally point.

We are interested in the in-between: what if we establish several different rally points?

There are two scenarios. In the first scenario, we establish ahead of time a hardened infrastructure of rally points that will be connected and withstand most disasters: for example, underground Ethernet, or protected microwave antennae.

The second scenario says we have no such infrastructure, or the infrastructure fails. It is possible that even with multiple rally points, a single cluster will form. But more likely, we will have created several isolated groups through the multiple rally point approach. Those groups might be integrated through a secondary bridging mechanism. We will discuss that possibility after examining the results of simulations with multiple rally points.

The reason for using multiple rally points may be situational. For example, it may be that for some tasks several small teams are better than one large team. Also, in search and rescue situations, a hot zone may be set up at the center of a disaster [23]. Responders might want to congregate in the warm zone, in order to decide what to do before going into the center of a situation. Therefore it would make sense to create rally points on corners of a surrounding square, a periphery, rather than in the center.

4. Simulation results

4.1 Overview

We want to understand when multiple rally points will be better than single rally points, and what the tradeoffs are.

First, we take a particular random starting situation, and compare its convergence using first one and then four rally points. After this, we look at average performance over 30 different starting situations. We also consider the placement of rally points.

4.2 An individual case

Starting with a random configuration, we build baseline single rally point configurations, shown in Figure 3, for nine different density levels. Higher density translates into more units, which drive faster convergence. Figure 4 shows convergence when four rally points are placed in the center of the four quadrants of the square.

We see that at higher densities the number of rally points does not matter very much: many units already have overlapping radio ranges and they converge rapidly.

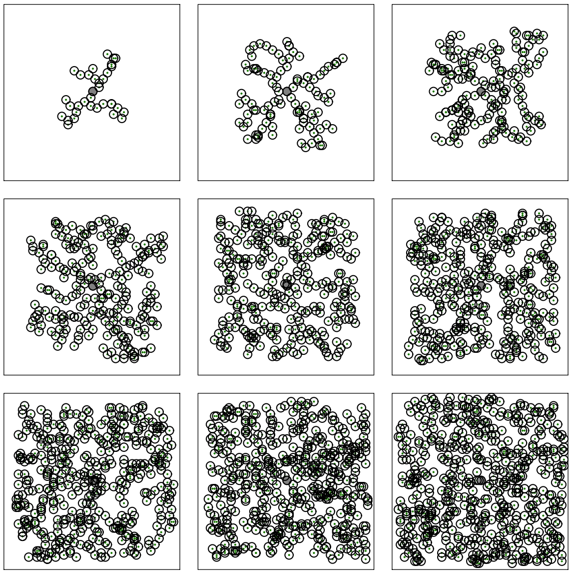

Figure 3. Single rally point convergences for varying densities (Rally points are in gray)

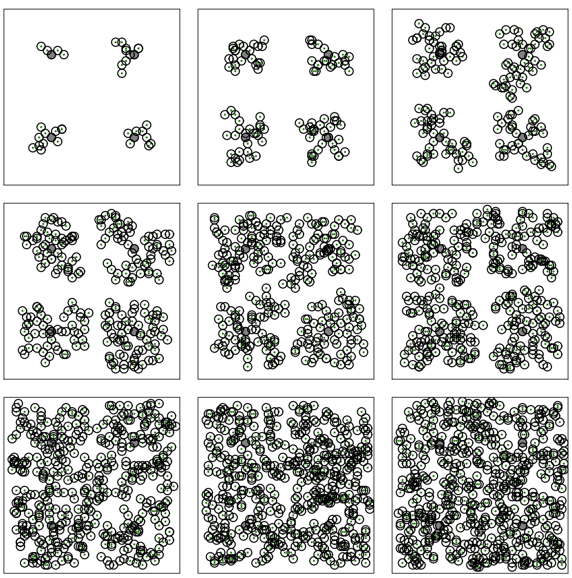

Figure 4. Convergences on four-rally-point configurations for varying densities.

Figure 5 shows the time to converge: the top line represents the single rally point scenario, and the bottom line represents the multiple rally point scenario.

In situations of low density, four rally points halve the time reconvene. This makes sense, as the distance any unit might have to travel has been split. At higher densities, however, there is not much difference.

So at low densities there is a clear tradeoff: multiple rally points produce faster convergence, but separate clusters are created. What happens to clusters with multiple rally points at higher densities?

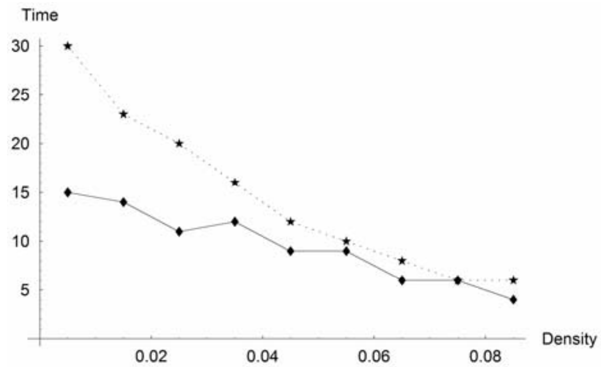

Figure 5. Times to converge for the runs shown in Figures 3 and 4. The star symbols represent single rally point scenarios, the diamonds multiple rally point scenarios.

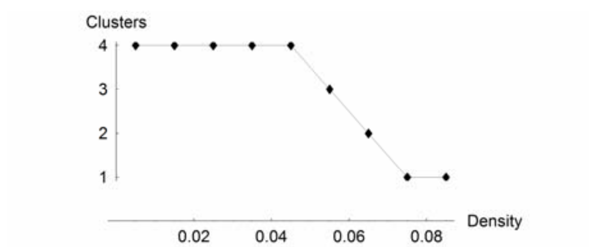

Figure 6. The number of clusters left after convergence on 4 rally points.

Figure 6 shows the number of clusters in relationship to density: the results can be confirmed by a careful inspection of Figure 4. As density increases, we get fewer clusters, but it is not until the density is quite high that a single component is built.

Therefore it would seem that we have a clear tradeoff between speed and consolidation when density is low, but when density is high a single rally point holds the advantage: it guarantees convergence to one cluster.

4.3 Thirty random scenarios

We performed simulations of time to converge at different densities on 30 starting scenarios in order to understand the variance of the results. Figure 7 shows the result for a single rally point; Figure 8 for multiple rally points. The error bars show variance. Figure 9 shows the two sets of data together. Figure 9 looks like a smoothed version figure 5, and indicates that the

0-7695-2755-8/07 $20.00 © 2007 IEEE

individual case we analyzed in section 4.2 was not an anomaly.

Figure 10 shows the ratio between the multiple and single rally point scenarios: at low density the ratio is about .5, and at high densities it nears 1.

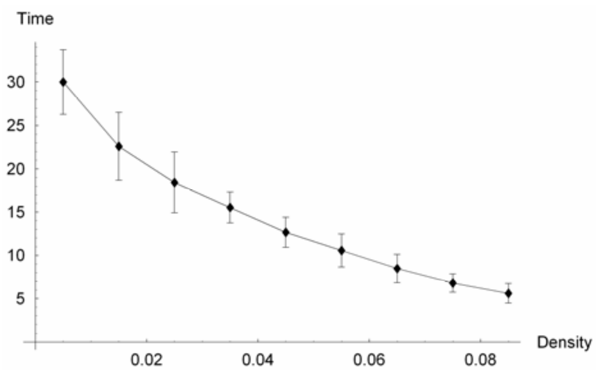

Figure 7. Time to converge by density, over 30 random starting situations with a single rally point.

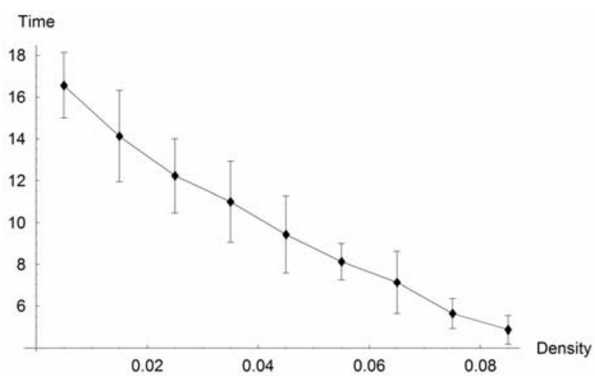

Figure 8. Time to converge by density, over 30 random starting situations with four rally points.

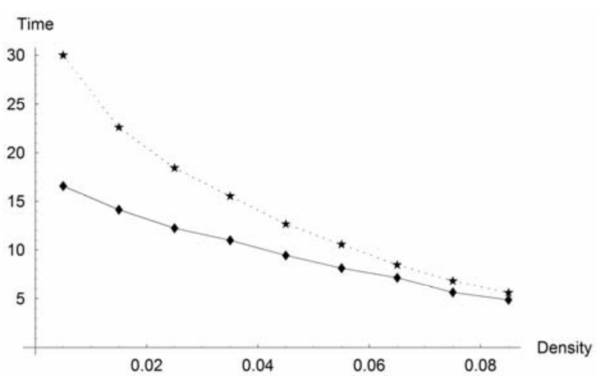

Figure 9. Multiple rally points (diamonds) and single rally points (stars) compared.

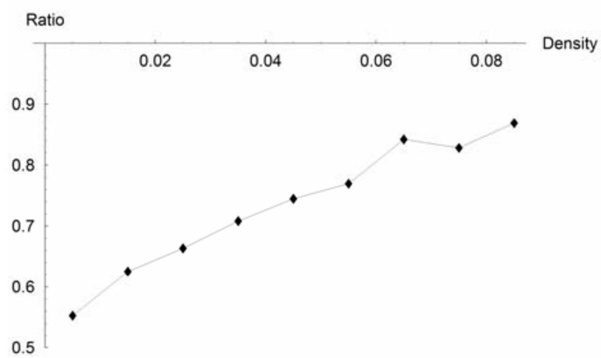

Figure 10. The ratio of multiple rally point solutions to single rally point solutions, in terms of time to convene for given densities.

4.4 Placement of rally points

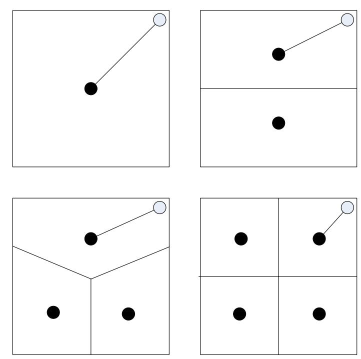

Figure 11. The distance to the furthest unit for 1, 2, 3, 4 rally points.

Our earlier work showed that units reconvene in roughly linear time with respect to the width of the field. Therefore, as we add rally points, we wish to know to what extent this will reduce the maximum distance to a rally point. For example, two rally points on a square grid may not reduce the distance much; see Figure 6. We can calculate that the distance to the split rally point will be reduced by about 20% $(\sqrt{2}-\sqrt{5/4})/\sqrt{2}$). However, four rally points will halve the distance to be covered, and therefore will halve the amount of time to reconvene. That is what we saw in Figure 10 when density was low.

What happens if we move the rally points to different locations in the field? Figure 12 shows the results of moving the rally points in increments of one eighth of the field at a give density: .035. (we don't reshow the single rally point example, which can be seen in Figure 3, frame four) .

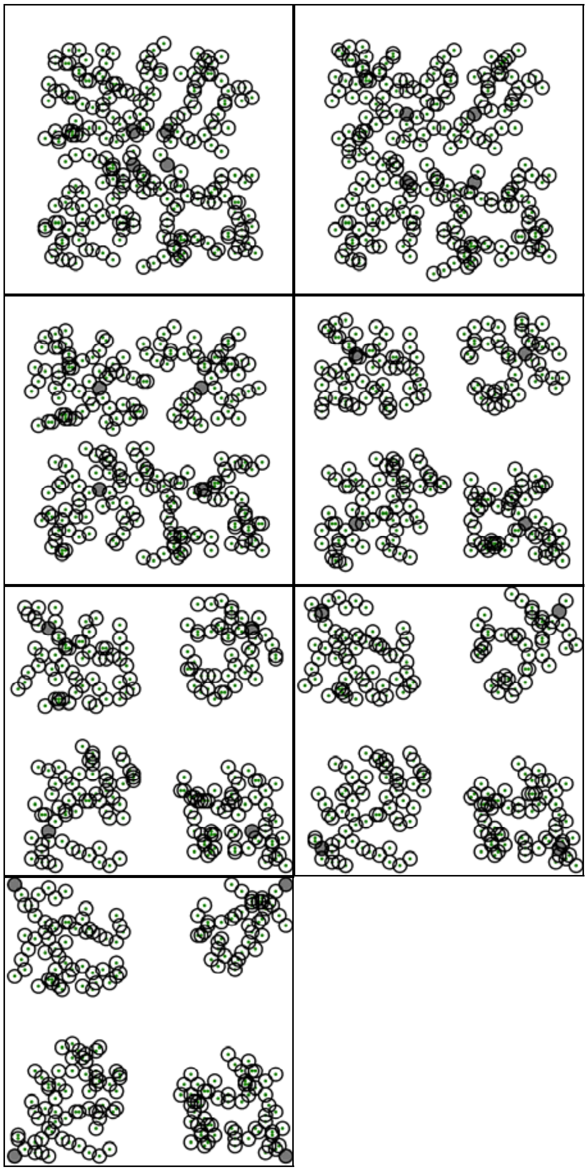

Figure 12. The effects of rally point placement.

What happens if we move the rally points to different locations in the field? Figure 12 shows the results of moving the rally points in increments of one eighth of the field at a give density: .035. (we don't reshow the single rally point example, which can be seen in Figure 3, frame four) .

Figure 13 shows the results in relationship to the time to converge. The first point shown is for zero spacing: in other words, for a single rally point. The results make sense: the times get better, but then get worse as the distance to the rally point for the average point begins to increase. The spacing of 20 is the geometric center of the quadrants that we used in all the previous tests.

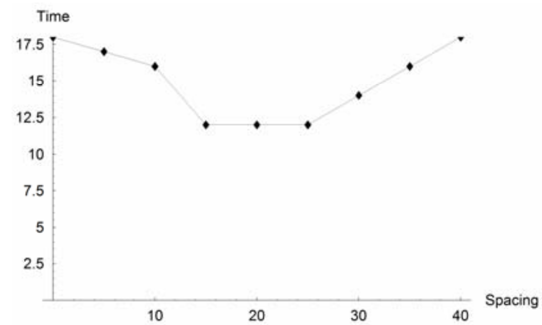

Figure 13. Time to converge for the graphs above

The graph accentuates the tradeoff between time and integration: the point at which the times get a lot better is also the point at which the graph breaks into separate clusters (see frame 3 of Figure 12, which links to the 4th point of Figure 13, at a spacing of 15). Now we turn to the problem of bridging such clusters.

5. Bridging

5.1 Overview

In studying multiple rally points, we originally assumed that the rally points would be joined by LAN or microwave. For example, a city, anticipating potential disasters, might create rally points with hardened communication infrastructure designed to survive emergencies. Then convergence on several different rally points would be the same as convergence on one point, in the sense that all nodes would be connected to each other – through their neighbors to their rally point, and through the infrastructure to other clusters. Alternatively, responders might carry portable microwave equipment that would be used to construct temporary links between rally points.

While these ideas are feasible, they are not without problems. Emergencies by definition cannot be fully anticipated. Underground LANs would be disrupted by earthquakes, floods, and bombs. Microwaves would be disrupted by weather or dust in the air triggered by any number of unanticipated events.

Thus, there maybe a need to establish other kinds of ad hoc bridges between separated clusters.

5.2 Chain bridges

In the absence of infrastructure, we might construct ad hoc network bridges through the proper positioning of units, as in Figure 14. We would favor the use of robots for this task; in many scenarios, the area to be bridged may be dangerous, and the job is repetitive.

Essentially, a bridge might be constructed by sending out members of one cluster toward the known location of another cluster. For example, several of the configurations in Figure 4 could be bridged if separated clusters slightly reconfigured their arms to reach out toward each other.

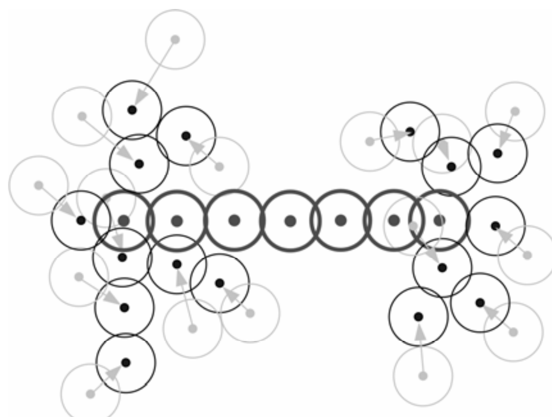

Figure 14. A robotic network bridge

In some cases, it may not be possible to form such a bridge. For example, there may not be enough idle units available, or the area to be bridged may be turbulent.

5.3 Couriers

Then we might consider using a courier technique, as in Figure 15, where a unit in one rally point cluster carries messages back and forth between another rally point cluster. In many search and rescue operations, this is communication method used – human couriers often bridge the gap between workers in the void of a building collapse and those outside the void. We can reason about these exchanges – the communication distance is the amount of time it will take the courier to traverse the gap [9].

In related work, we have looked at how to model the couriers [24-26]. By using more couriers, the average latency of messages can be reduced, but only to a certain point. No matter how many couriers are used, messages will be delayed by the time it takes to traverse the gap. Yet this technique may still be the most pragmatic one in an emergency.

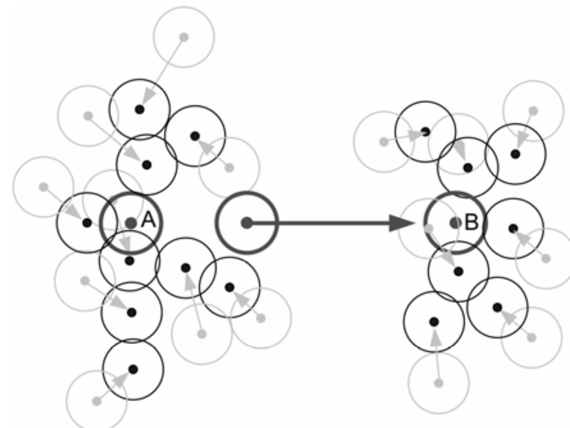

Figure 15. A robotic courier

5.4 Assemblies as bridges

Up until this time, we have been modeling homogeneous units. Let us assume now that we have two kinds of responders: humans and robots. Then the nature of bridging changes.

Robots are expendable in an emergency, and are better able to handle harsh environmental conditions. Thus, we can think of robots as forming a bridge for humans to communicate over.

More radically, we can think of a single rally point convergence of robots as forming a tree-like bridge between all humans.

In other words, let robots converge on a point, forming a linked configuration. Then let humans move to attach to this configuration. Humans have to move less, because the robots formed an infrastructure which people can attach to at will.

We make some specific assumptions before modeling this situation. We assume a ratio of 1:4 humans to robots– a team of 50 humans becomes a team of 250 units when augmented by 200 robots.

We assume the robots move first. If the humans don't move until all the robots are connected (perhaps the robots can signal visually when their movement is complete) then the humans generally end up as leaves on a tree, as in Figure 16. The initial configuration of robots becomes a configuration of rally points which humans can move to. Such a heuristic slows up the reconvening; Figure 17 shows the impact. Yet, this is still faster than connecting with just 50 humans, and there are other benefits; the humans expend less energy, and are less likely to be trapped in the center of a disaster before gaining information on the environment through sensing.

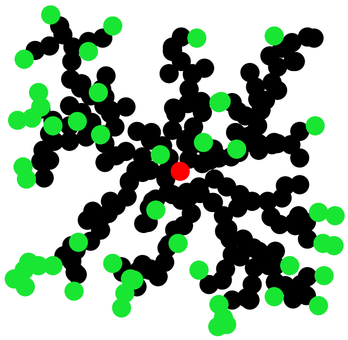

Figure 16. Humans stay still until the robots have connected

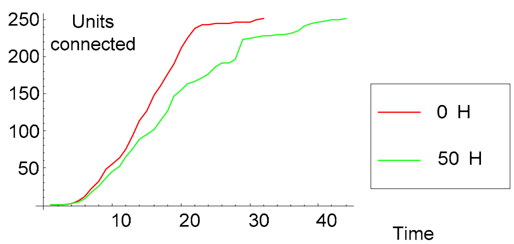

Figure 17. The red line shows 250 units converging; the green line represents a run in which 50 of the 250 units as human, and they move only after the robots are in place.

6. Future Work

The tradeoffs between the risk of remaining disconnected, and the time to reconvene can be explored in the framework of multi-attribute decision theory [27]. There are risks associated with a single rally point: if it is not reachable a network will not form. And there are risks related to multiple rally points, as the ability to bridge the rally points may be low in a true emergency environment.

There is also an aspect of sampling at work: it may be better to have disconnected entities more widely dispersed that share information over time than a highly connected set in a very small area. This is situation-dependent: sometimes sampling at the epicenter of a disaster is all that is needed, and sometimes sampling over a wider area will yield better results. For example, in identifying the outbreak of a virus, many samples dispersed over a wide area may be needed to find the extent of the problem. So sensors sent to many places, disconnected, may make more sense than a connected network in the center of a city. Then, of course, the disconnected sensors will need a way to integrate over time.

We used simple heuristics to form configurations; these heuristics have the advantage of being easy for both machines and humans to implement: all they involve is moving toward an agreed-on location. In discussing bridging, we suggested that already formed clusters might extend themselves to meet other clusters. If such algorithms could be simply implemented, then the fragmentation of multiple rally point approaches might be reduced. Other more complex assembly heuristics might yield a variety of different solutions in the time vs. integration trade-off space.

The idea of moving robots and then moving people to connect with the robotic infrastructure can be developed further. Should people go into the center of a disaster? Stay where they are? Move away from the center? Depending on the circumstance, all three approaches may make sense, and the sensors might be optimized to support all three strategies.

7. Conclusions

In emergencies, responders can converge on a single rally point. Or they can converge on several. There is a tradeoff: converging on the center will in many cases be the most robust solution, insuring everyone will connect. But converging on several rally points gets sub-teams connected faster.

Through simulation, we have illustrated the tradeoff: the time to reconvene will be faster with more rally points, but the end configuration will be many separated clusters. This tradeoff is clear when the density of sensors is low. When the density is high, there is little difference in the time to convene, and single rally point solutions have an edge in that they avoid fragmentation of the network.

Moving the rally point locations does not change the tradeoff: when the locations are close together, they function much like a single rally point.

Multiple rally points will reduce the time for a team to reconvene. There is another reason to use multiple rally points: converging on the periphery of a toxic disaster is less risky for the responders than converging on the center. Because multiple rally points may create fragmentation, we discussed several techniques for linking disconnected segments.

There are implications for sensor networks: such networks will often be isolated and clustered, and therefore alternative ways of assembling and integrating them are needed. There are implications for emergency responders: convening into several geographically dispersed groups first and deferring integration of the groups until later will create functional teams faster than attempts to connect everyone right away.

Acknowledgements

This research was supported in part by the Office of Naval Research, grant #N00014-05-1-00632.

References

[1] Q. Li and D. Rus, "Message Relay in Disconnected Ad-hoc Networks," *IEEE MASCOTS Workshop on Mobility and Wireless Access*, pp. 14--21, 2002.

[2] M. Grossglauser and D. N. C. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Transactions on Networking*, vol. 10, pp. 477-486, 2002.

[3] I. Chatzigiannakis, S. Nikoletseas, N. Paspallis, P. Spirakis, and C. Zaroliagis, "An Experimental Study of Basic Communication Protocols in Ad-hoc Mobile Networks," *Lecture Notes in Computer Science*, vol. 2141, pp. 159-169, 2001.

[4] D. K. Goldenberg, J. Lin, A. S. Morse, B. E. Rosen, and Y. R. Yang, "Towards Mobility as a Network Control Primitive," *MobiHoc*, 2004.

[5] D. J. Goodman, J. Borras, N. B. Mandayam, and R. D. Yates, "INFOSTATIONS: a new system model for data and messaging services," *Vehicular Technology Conference*, pp. 969 -973 vol.2, 1997.

[6] G. Wang, G. Cao, and T. L. Porta, "Movement-Assisted Sensor Deployment," *Infocom*, 2004.

[7] Y. Zou and K. Chakrabarty, "Sensor Deployment and Target Localization Based on Virtual Forces," *Infocom*, 2003.

[8] Y. Zou and K. Chakrabarty, "Sensor Deployment and Target Localization in Distributed Sensor Networks," *ACM Transactions on Embedded Computing Systems*, vol. 3, pp. 61-91, 2004.

[9] J. V. Nickerson, "Robots and Humans Reconvening," *IEEE International Conference on Systems, Man and Cybernetics*, 2004.

[10] J. V. Nickerson, "A Concept of Communication Distance and its Application to Six Situations in Mobile Environments," *IEEE Transactions on Mobile Computing*, vol. 4, pp. 409-419, 2005.

[11] J. V. Nickerson, "Reconvening after a Break in Communication," *Proceedings of the Sixteenth Annual Workshop on Information Technologies and Systems*, 2006.

[12] J. V. Nickerson, "Rally Point Simulations," http://www.stevens.edu/jnickerson/rallypoints, 2004.

[13] DOD, "US Army Survival Manual: FM 21-76," US Department of Defense 1992.

[14] P. Meakin, *Fractals, scaling, and growth far from equilibrium*. Cambridge, U.K. ; New York: Cambridge University Press, 1998.

[15] T. A. Witten and L. M. Sander, "Diffusion-limited aggregation," *Phys. Rev. B*, vol. 27, pp. 5686, 1983.

[16] M. Batty and P. Longley, *Fractal Cities*. London: Academic Press, 1994.

[17] P. Longley and M. Batty, *Advanced spatial analysis: the CASA book of GIS*. Redlands, Calif.: ESRI Press, 2003.

[18] D. N. Coore, "Botanical Computing," in *Computer Science*: MIT, 1999.

[19] H. Abelson, D. Allen, D. Coore, C. Hanson, G. Homsy, T. Knight, R. Nagpal, E. Rauch, G. J. Sussman, and R. Weiss, "Amorphous computing," *Communications of the ACM*, vol. 43, pp. 74-82, 2001.

[20] *National Commission on Terrorist Attacks Upon the United States, The 9/11 Commission Report*. New York: W.W. Norton, 2004.

[21] *Select Bipartisan Committee to Investigate the Preparation for and Response to Hurricane Katrina, A Failure of Initiative*, 2006.

[22] D. M. Dobkin, *RF engineering for wireless networks: hardware, antennas, and propagation*. Boston: Elsevier, 2005.

[23] R. R. Murphy, "Human-robot interaction in rescue robotics," *IEEE Transactions on Systems, Man and Cybernetics, Part C*, vol. 34, pp. 138-153, 2004.

[24] J. V. Nickerson, "Flying Sinks: Heuristics for movement in sensor networks," *Proceedings of the 39th Annual Hawaii International Conference on System Sciences*, 2006.

[25] J. V. Nickerson and S. Olariu, "Courier Assignment in Social Networks," *Proceedings of the 40th Annual Hawaii International Conference on System Sciences*, 2007.

[26] J. V. Nickerson and S. Olariu, "A Measure for Integration and its Application to Sensor Network," *Workshop on Information Technology and Systems*, 2005.

[27] R. L. Keeney and H. Raiffa, *Decisions with multiple objectives: preferences and value tradeoffs*. Cambridge: Cambridge University Press, 1993.

0-7695-2755-8/07 $20.00 © 2007 IEEE