Discovering Adaptive Heuristics for Ad-Hoc Sensor Networks by Mining Evolved Optimal Configurations

Prasanna Ranganathan, Aravind Ranganathan, Kenneth Berman and Ali Minai

Abstract*— Ad-hoc sensor networks comprising large numbers of randomly deployed wireless sensors have recently been an active focus of investigation. These networks require self-organized configuration after deployment, and ad-hoc heuristic methods for such configuration have been proposed with regard to many aspects of the networks' performance. However, systematic approaches for such configuration remain elusive. In this paper, we present a preliminary attempt towards such a systematic approach using evolutionary algorithms and reverse engineering. In particular, we focus on the problem of obtaining heterogeneous networks that optimize global functional properties through local adaptive rules. Almost all work on ad-hoc sensor network has so far involved homogeneous networks where all nodes transmit with the same power level, creating a symmetric connectivity. It is possible to construct heterogeneous networks by allowing nodes to transmit at different power levels, and such networks are known to provide improvements in network lifetime, power efficiency, routing, etc. However, such networks are difficult to build mainly because the optimal power level for each node depends on the node location and spatial context, which are not known before deployment. A few heuristic schemes focused on improving power consumption have been proposed in the literature, but the issue has not been investigated sufficiently at a general level. In this paper, we present a new and improved heuristic developed using a reverse engineered approach. A genetic algorithm is used to generate a set of heterogeneous sensor networks that are characterized by low short paths and minimal congestion. Analysis of this optimal network set yields rules that form the basis for a local heuristic. We show that networks adapted using this heuristic produce significant improvement over the homogeneous case. More importantly, the results validate the utility of the proposed approach that can be used in other self-organizing systems.

I. INTRODUCTION

The growing capabilities and miniaturization of wireless devices over the past decade has fuelled increased research interest in wireless networks. One of the most active research areas in this regard has been that of ad-hoc sensor networks. These networks comprise large numbers of wireless sensor nodes deployed randomly over an extended region, and then required to configure themselves into an ad-hoc network for applications such as event-monitoring, guidance and intrusion detection [1]. Because of their random structure, large extent and low node power, in-field configuration of these networks typically involves some variant of self-organization — obtaining global structure through local adaptive rules — making them an attractive model system for research on self-organization.

Prasanna Ranganathan, Aravind Ranganathan, Kenneth Berman and Ali Minai are with the Electrical & Computer Engineering and Computer Science Department, University of Cincinnati, Cincinnati, OH 45221. Email:{ranganp, rangana, berman, aminai}@ececs.uc.edu

A. Background and Motivation

Most research on ad-hoc sensor networks has focused on homogeneous networks where all nodes transmit signals with the same power. Since the nodes are assumed to be identical, transmission power determines the range over which each node's signal can be detected by others, giving an identical *transmission radius* for each node and creating a symmetric network (except for inhomogeneities caused by obstructions, etc., which are usually ignored). This simplifies the analysis of network performance, e.g., using the tools of percolation theory [2] and random graphs [3]. It also makes the design of configuration protocols easier by directly relating connectivity to node density, eliminating the need for individual nodes to choose their transmission power. However, as in all distributed systems, optimal configurations are likely to be highly heterogeneous [4]–[7], and the homogeneity of current models is driven more by computational expediency than by performance goals [8], [9]. This has been verified by research showing that heterogeneous networks where nodes use different transmission power levels have lower power consumption [10], [11]. Since sensor nodes are limited power devices, conservation of power typically prolongs the network lifetime [12]. Heterogeneity in sensor networks also provides better routing schemes [13], [14], and is likely to reduce channel congestion by allowing many — even most — nodes to have smaller communication radii.

In spite of the fact that heterogeneous networks are more efficient than their homogeneous counterparts, their utility is limited by the difficulty of configuration. In most applications of large scale sensor networks, the actual or relative location of the sensors cannot be determined before their deployment. Thus, each node must determine its transmission radius after deployment using information obtained in the field. The transmission radii of nodes in a heterogeneous network affect network connectivity, lifetime and robustness, which are all global properties that can be optimized. However, each node must base its decisions only on *local* information in order to keep the configuration process scalable. Thus, the central question for this system is: *How can each node determine its transmission power using local information such that certain global properties of the network are optimized?* This is an instance of a fundamental issue faced by all self-organizing systems such as swarms, multi-agent systems, etc., where individual components must adapt parameters of choice based on limited local information while seeking to optimize global performance. Discovering general rules for such choices is extremely difficult, often leading to the use of ad-hoc heuristic rules based on intuitive analysis and subsequent fine tuning. The goal in the work presented here is to propose a more systematic, knowledge-discovery oriented approach that begins to address this issue. From the application viewpoint, the work extends our earlier results [15] on a heuristic that allows each node to choose its transmission radius to optimize for mean shortest path length (MSPL) and congestion in the network as a whole. Minimizing MSPL reduces the overhead of multi-hop communication. MSPL alone is clearly minimized by allowing maximal transmission power, but this also increases congestion. Optimizing both MSPL and congestion together leads to a network where communication overhead is minimized and is evenly distributed throughout the network.

While there is a large body of work on "evolving networks", most of it refers simply to networks with dynamic structure rather than actual evolutionary optimization. However, such optimization has been applied extensively in the areas of neural networks [16]–[18], gene regulation networks [19]–[24] and metabolic networks [5], [6]. Evolutionary algorithms are usually considered too slow for problems requiring real-time self-organization without multiple trials, such as configuration of sensor networks, but these algorithms can play an extremely important role in developing the self-organization algorithms themselves. This, after all, is exactly the function evolution serves in the biological context (e.g., optimizing developmental programs). In this work, we use evolutionary algorithms in a more limited design context: Providing optimal solutions from which desirable system features can be extracted by knowledge discovery. The larger project of optimizing self-organization rules through evolution is also underway.

The organization of the paper is as follows. In the following section the network model, optimization strategy followed and other details from our earlier work are explained. Section III describes the radius heuristic in detail. The simulation results and analysis are provided in Section IV. The conclusions and possible future work are discussed in Section V.

II. Description of Approach

The primary difficulty in systematically obtaining local rules for global self-organization is that the relationship between local features that can be directly addressed by some rules and the corresponding global properties being optimized are not available. The system is too complex to be analyzed tractably except under simplifying assumptions such as uniform transmission radius, which renders the exercise moot. We attempt to circumvent this problem by adopting a reverse engineering approach that works by first trying to obtain optimal configurations in a broad population of systems using a global search method and then *mining* these configurations to determine their useful features. This is described below.

A. Network Model

The network is modeled as a unit square where n sensor nodes are distributed in a uniform random fashion with density λ. Each node is able to choose its operating transmission radius from the range $[r_{min},\ r_{max}]$. For simplicity, we only allow nodes to choose transmission radii from a set of discrete values, $R = \{r_{min}, r_2, r_3, \ldots, r_{max}\}$. The values for r_{min} and r_{max} for the heterogeneous sensor network are chosen such that, if r_{perc} is the percolation radius [2], [25] of the corresponding homogeneous network, then r_{min} is less than r_{perc} and r_{perc} is one of the possible choices for node radius from the set R. The homogeneous network used to compare performance has all its radii set to r_{perc} to ensure connectivity. The expectation is that, in optimal heterogeneous networks, many nodes will be able to use radii lower than r_{perc} because a few nodes with large radii can compensate to maintain connectivity.

It is assumed that each node knows its global location, and that protocols exist to allow reliable communication. These protocols are not the object of our study and are not simulated explicitly.

B. Optimization Strategy

The general strategy used to discover local features of optimal networks and to obtain local adaptive rules involves the following steps:

1) A configuration space and a fitness (objective) function are defined for the class of networks considered.
2) A large number of random node distributions are generated and an evolutionary algorithm is used to obtain (possibly multiple) optimal (or near-optimal) node radius configurations for each node distribution. The resulting set of optimal network configurations is termed the *optimal network set (ONS)* and is used for feature mining. Multiple runs, random mutational shocks and local search are used to ensure that the configurations obtained can reasonably be considered near-optimal.
3) The optimal configurations are analyzed to identify properties that differentiate them from surrogate non-optimal configurations with the same radius statistics.
4) Based on the results, local rules are developed that allow a node to make decisions leading to networks with features similar to the evolved optimal configurations.
5) Finally, the rules are implemented and the performance of the resulting network is compared with that of networks obtained using the evolutionary algorithm.

C. Generation and Analysis of the Optimal Network Set

For our specific problem, we used the approach described above to optimize MSPL and congestion in networks. The MSPL, M_j, of a given network, j, was computed as the average of the lengths of the shortest paths between all pairs of nodes in the network. Congestion, C_j, was modeled as the mean in-degree of all nodes in the network. The evolutionary algorithm (EA) coded the network as an ordered vector of the transmission radii of its nodes. This vector formed the

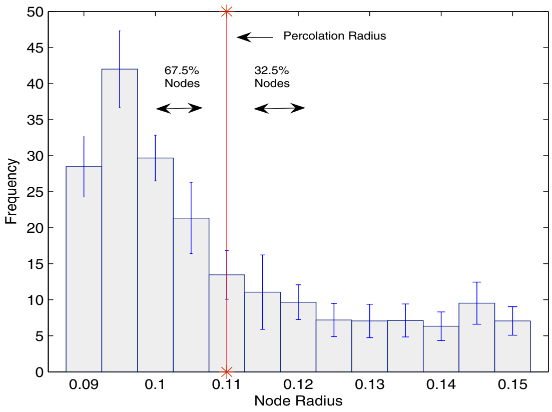

Fig. 1. Mean Radius Distribution of GA networks with error bars

network's chromosome. The fitness function of the network was taken to be $\phi_j = (M_j C_j)^{-1}$. Initial candidate networks were built by assigning each node a uniformly randomly chosen radius from the range $[r_{min},\ r_{max}]$.The candidate networks for the next generation were obtained by crossing over chromosomes of parent networks from the current generation. Selection for both mating and survival was stochastically elitist, with a fraction of the fittest networks guaranteed selection (the parameters varied between runs because the goal was to obtain the best networks rather than systematically studying the the evolutionary algorithm.) Each trial was run until a stable, reliable maximum was obtained, but at least for 500 generations. The resulting optimal networks formed the ONS.

Analysis of the ONS to obtain useful features is largely an inductive process which, for this study, was still done mostly by informed trial-and-error. A more automated approach is being developed and will be reported in the future. The analysis comprised two steps:

1) The statistical distribution of radii for ONS members was studied and was found to show a characteristic pattern. Most nodes had radii between r_{min} and r_{perc}, while a few key nodes had large radii (see Figure 1). This architecture is somewhat reminiscent of the hub-based configurations studied in the scale-free networks literature [26], [27]. These results provided the principle that small radii could be considered the default in optimal networks, with a few well-chosen nodes given large radii, i.e., what was needed was a bimodal heterogeneity. Interestingly, several researchers have recently obtained similar results in random networks with non-geometric (non-local) connectivity [28], [29]. The issue, of course, is to identify the large-radius nodes based on local features.
2) Next, the radius values for nodes in each optimal configuration were correlated with several local properties of the node's neighborhood configuration. It was found that the radius value chosen by a node j was linked to two local properties: A) The density, N^j, of neighbors within close proximity of j, and B) The the maximum separation, d^j, among nodes within j's transmission radius. These two properties A and B were then used to obtain the local rules for adaptation as described below.

III. The Radius Heuristic

Let $R = \{r_{min} = r_1, r_2, \ldots, r_n = r_{max}\}$ denote the set of n different radius values that can be chosen by a node. The maximum neighborhood of a node j consists of all nodes reachable from j along with their relative positions when j has the maximum transmission radius r_{max}. Since the strategy is to optimize global fitness, ϕ_j, by having each node choose its radius using properties A and B, define a local fitness measure for each node in terms of these properties. Each node j computes this fitness measure, f_i^j, for each value $r_i \in R, i = 1, 2, ..., n$, and chooses its transmission radius $r^j = r_{i^*}$ such that $f_{i^*}^j \geq f_i^j, i = 1, 2, ..., n$.

The fitness measure f_i^j for a given radius value r_i is the sum of two components:

1) Base fitness measure, b_i^j.
2) Gain fitness measure,g_i^j.

Both fitness components are a function of the transmission radius being considered by the node. Intuitively, the base fitness captures the quality of the node at that radius while the gain fitness indicates whether this quality is disproportionately better than that at the best available lower radius.

The base fitness, b_i^j, of node j for a radius r_i, is the sum of two components $b_i^{N^j}$ and $b_i^{d^j}$. The first component, $b_i^{N^j}$ is the ratio of number of neighbors of the node j at this radius, N_i^j, to the expected number of neighbors of j. The expected number of neighbors of a node j depends on the radius r_i and the node density, λ. The second component $b_i^{d^j}$ is the maximum neighbor separation d_i^j at the radius r_i normalized by the diameter. Thus,

$$b_i^j = b_i^{N^j} + b_i^{d^j}$$

$$b_i^{N^j} = \frac{N_i^j}{(\pi(r_i)^2\lambda) - 1}, \quad b_i^{d^j} = \frac{d_i^j}{2r_i}$$

The gain fitness, g_i^j, of node j is computed only for radius values $r_i > r_1$ and again is the sum of two individual components: $g_i^{N^j}$ and $g_i^{d^j}$. The first component, $g_i^{N^j}$, is the increase in the number of neighbors expressed as a percentage of the expected increase in the number of neighbors between radii r_{i-1} and r_i. The expected increase in the number of neighbors is given by $\pi[(r_i)^2 - (r_{i-1})^2]\lambda$. The second component, $g_i^{d^j}$ is computed only if there is an increase in the maximum neighbor separation after increasing the radius from r_{i-1} to r_i. It is computed as the percentage increase in the maximum neighbor separation at radius r_i compared to the previous radius r_{i-1} with respect to the increase in diameter which is $2(r_i - r_{i-1})$. Thus,

$$g_i^j = g_i^{N^j} + g_i^{d^j}$$

$$g_i^{N^j} = \frac{(N_i^j - N_{i-1}^j) - \pi[(r_i)^2 - (r_{i-1})^2]\lambda}{\pi[(r_i)^2 - (r_{i-1})^2]\lambda}$$

$$g_i^{d^j} = \begin{cases} \frac{(d_i^j - d_{i-1}^j) - 2(r_i - r_{i-1})}{2(r_i - r_{i-1})} & \text{, if } d_i^j > d_{i-1}^j \\ 0 & \text{, else.} \end{cases}$$

A. Heuristic Implementation

The heuristic rule used by each node is described below. It uses only local information obtained by each node through local messages. All nodes make their initial choices autonomously, followed by a coordination step to improve performance as described in the next section.

1) Each sensor node j, after deployment, obtains information on the nodes in its r_{max}-neighborhood and their relative positions by means of "hello" messages.
2) The *current radius index* variable, α, is set to 1.
3) The number of neighbors, N_α^j, and the maximum separation among the neighbors, d_α^j, for radius r_α are determined. The node then computes the fitness measure $f_{r_\alpha}^j$ for radius r_α. If $\alpha = 1$, this consists of just the base fitness measure.
4) The *next radius index* value, β, is set as $\alpha + 1$. The number of neighbors, N_β^j, and the maximum separation among the neighbors, d_β^j, for radius r_β are determined.
5) If there is no increase in the number of neighbors at radius r_β compared to radius value r_α, then r_β is dropped from the set of possible radius values for node j and the total number of radius values, n, is decremented.
6) In case there is an increase in the number of neighbors with the increase in radius, the gain fitness value g_β^j is calculated for r_β^j as described earlier, and α is incremented by 1.
7) The base fitness b_β^j and the final fitness value, $f_\beta^j = b_\beta^j + g_\beta^j$ is calculated.
8) If $\alpha < n$, the process is repeated from Step 4. Else, the radius r_{i^*} with the highest fitness value $f_{i^*}^j$ is chosen as the operating radius by node j.

The logic of the heuristic is simple — the base fitness measures how good a radius value, in itself, is for a node and mainly contributes to the selection of smaller radius values. The gain fitness measure is used to quantify the advantage/disadvantage of increased radius value. It must be noted that the gain fitness value can be negative thus decreasing the overall fitness of certain high radius choices for a node. This helps ensure that a large radius is chosen only in cases where it would be beneficial.

B. Improving the Heuristic

It was found through simulations (discussed in detail in the next section) that the networks obtained using the heuristic, while giving MSPL and mean node radius comparable with the evolved optimal networks in the ONS, were extremely congested. The source of additional congestion in heuristic networks is understood from the comparison of the radius distribution in the ONS networks and heuristic networks.

Figures 2 and 3 and show the distribution of radii in the evolved and heuristically organized networks respectively for a particular node layout. It is clearly seen that in the case

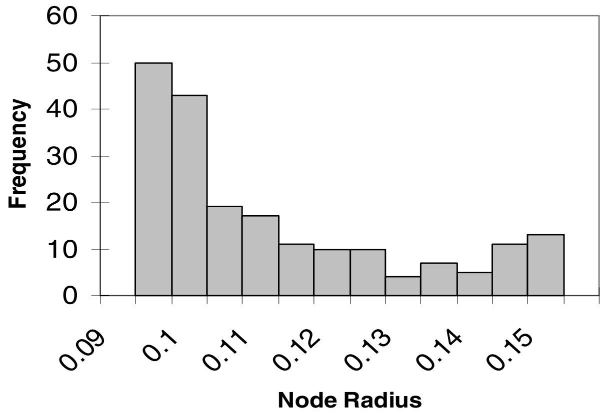

Fig. 2. Radius Distribution of GA networks

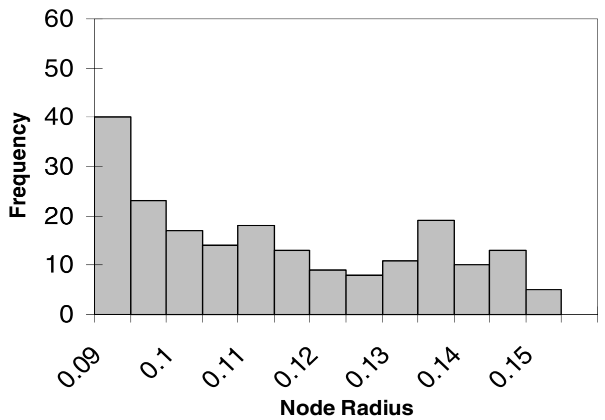

Fig. 3. Radius Distribution of Heuristic networks

of heuristic networks, more nodes have radius values that are in the middle of the possible range of values. Since the heuristic is necessarily local and node-centered, congestion minimization does not enter it explicitly as it does in the fitness function for the evolutionary algorithm. Thus, some localized coordination among nodes is necessary to reduce congestion. One possible strategy for this is to reduce the radius of those nodes in the intermediate range that will not cause a drastic change in the MSPL of the network. We have developed a collaboration mechanism based on this and show that it improves the fitness of the heuristic networks.

C. The Collaboration Step

Once all the nodes have been deployed and have selected an operating radius based on the heuristic detailed earlier, the nodes enter into a second collaboration step. In this collaboration step, each node j communicates its chosen operating radius to other nodes within its minimum radius r_{min}. Once a node has received information regarding the chosen operating radii of all its minimum radius neighbors, it makes a decision on whether to reduce its radius or not.

The decision on reducing the operating radius is governed by the following rules.

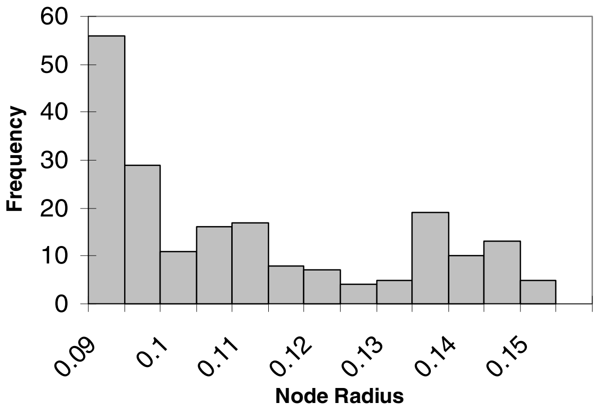

Fig. 4. Radius Distribution of C-Heuristic networks

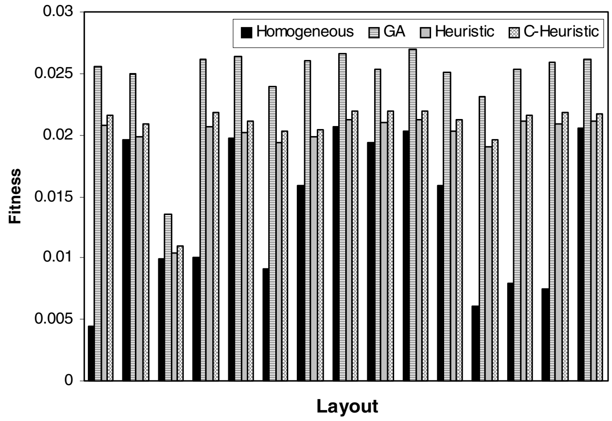

Fig. 5. Comparison of Network Fitness

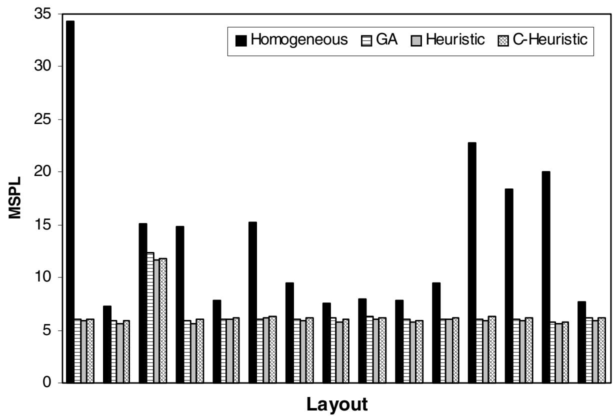

Fig. 6. MSPL Comparison

1) Only nodes with chosen operating radius greater than the percolation radius, r_{perc} and below a fixed cutoff radius value, r_{cutoff}, are allowed to reduce their radius values.
2) A node, satisfying the first rule, decides to reduce its radius if and only if it finds another node within its minimum radius neighbors with chosen operating radius greater than r_{cutoff}.
3) A node that makes the decision to reduce its radius value selects the next fittest lower radius as its final operating radius.

We denote the heuristic with the collaboration step as *C-Heuristic*. The collaboration is designed such that only those nodes that are less important reduce their radius values. The impact of the reduction in radius of these nodes on MSPL is mitigated by two things: 1) The presence of a nearby node with radius greater than the fixed high cutoff radius, and 2) The fact that a node reduces its radius only down to its next fittest lower radius. The collaboration step, after the required information has been obtained, is done independently by each node and the decision made is also independent. Thus, an a priori cutoff radius is needed to preclude a chain reaction that causes most of the nodes to reduce their radii, greatly affecting the network's global fitness. The collaboration step involves minimal communication overhead but produces a significant improvement in network fitness, congestion and mean node radius as detailed in the *Simulation and Results* section. The radius distribution of the same network as in Figure 3 after the collaboration step is shown in Figure 4. It shows that the number of nodes with radius values less than percolation radius has increased after the collaboration step. Also, this radius distribution better resembles the radius distribution of the optimal evolved network shown in Fig 2.

IV. Simulation and Results

We performed extensive simulations on networks with 200 nodes. The percolation radius for these network was found to be approximately 0.11 which also worked out to be the best radius for the homogeneous case. The other values used were $r_{min} = 0.09$ and $r_{max} = 0.15$ for the heuristic and the evolved networks, with $r_{cutoff} = 0.135$ for the collaboration step. The performance of the heuristic, C-heuristic, evolved and homogeneous networks was evaluated on the following parameters: Fitness (as defined above), MSPL, mean node radius, and robustness. The mean node radius is an important criterion since it corresponds to the overall energy consumption of the network. Also, the propensity of sensor nodes to fail makes the robustness of a network to node failure a good indicator of the network's utility.

Figure 5 shows the fitness of 15 different networks for the homogeneous ($r = 0.11$), GA, and the heuristic and C-heuristic cases, while Figure 6 shows the plot of the network MSPL for the same networks. Each group of bars corresponds to a different deployment, that is, different uniform random placements of the 200 nodes. The values of r_{min}, r_{max} and r_{cutoff}, listed earlier, are the same for all these 15 different deployments and the fitness in each case is evaluated using the same measure. It is seen that the heuristic produces networks with better fitness and better network MSPL than homogeneous networks. Also, the collaboration step is seen to improve the overall fitness of the network even though there is a slight increase in the MSPL of the

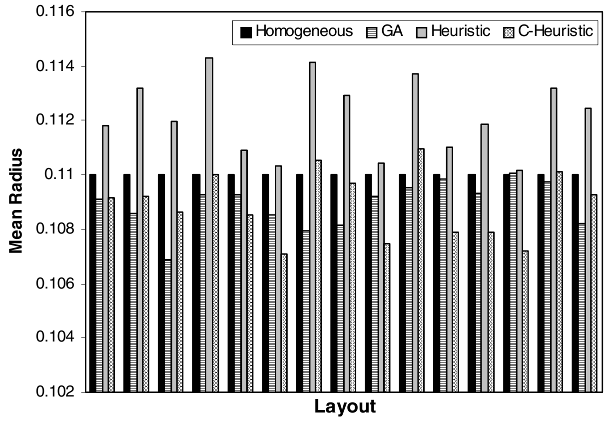

Fig. 7. Comparison of Mean Node Radius

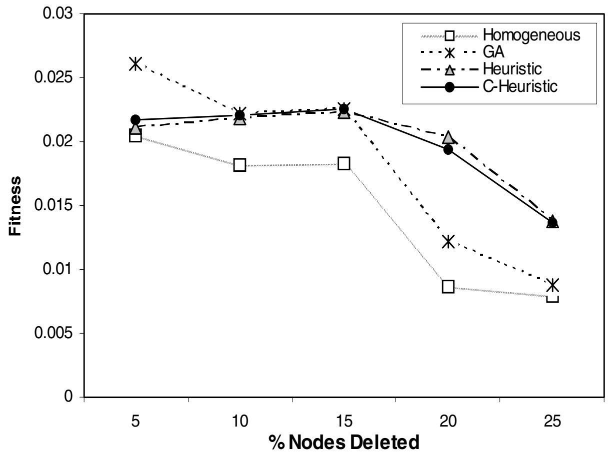

Fig. 8. Comparison of Robustness

network. Homogeneous networks at $r = 0.11$ are not always connected and hence their MSPL values are high due to the penalty incurred for disconnected nodes as seen in Figure 6.

The plot of mean node radius of the homogeneous, GA and heuristic networks for different deployments is shown in Figure 7. Here we see that the heuristic networks have a mean node radius comparable to homogeneous networks. More importantly, the collaboration step greatly reduces the mean node radius thus minimizing the overall energy consumption of the network.

It is seen that that the congestion in the heuristic networks is on the higher side even after collaboration. This additional congestion translates into an increased number of edges or connections among the nodes. Networks with comparable mean node radius are expected to have comparable congestion but the heuristic networks have unexpectedly higher congestion. Intuitively, a node or a network managing to make more connections at the same radius should be more robust. That this is indeed the case is verified through the calculation of robustness for the different networks. We measure robustness to node failure by randomly deleting 5% of the nodes and then calculating the network fitness. Figure 8 shows the plot of network fitness (averaged over multiple trials) versus percentage of nodes deleted for a particular node deployment. It is seen that the heuristic networks are very robust and outperform the GA networks considerably, while both are much better than homogeneous networks. Heuristic networks not only do better in terms of the percentage change in fitness but emerge as the fitter ones after node deletion for nearly all the deployments. Since robust networks are of great interest in many wireless network applications, we find this ancillary benefit of our heuristic very intriguing. It will be explored more fully in future research.

A. Optimal Heterogeneity

In a series of seminal papers, Doyle, Carlson and colleagues have recently raised the issue of optimality in complex systems [4], [5], [7], [30], [31]. Their main point is that optimal configurations in such systems are likely to be *atypical* rather than generic (as suggested by other researchers [32]). Thus, any adaptive rule for self-organization must be able to reach the rare optimal (or near-optimal) configurations. Since there has been great interest recently in whether the distribution of connectivity in networks determine global properties such as efficiency and robustness [28], [29], [33], [34], we decided to use this distribution to examine the issue raised by Carlson et al. Taking the set of radii found by the evolutionary algorithm for an optimal network, we scrambled the radius assignments randomly to generate a set of *surrogate networks* with the same radius distribution as the optimal one. We then looked at the distribution of fitness for these surrogates and how the various optimized networks (evolved, heuristic and C-heuristic) as well as the homogeneous network fit into this distribution. Figure 9 shows the result for two typical networks.

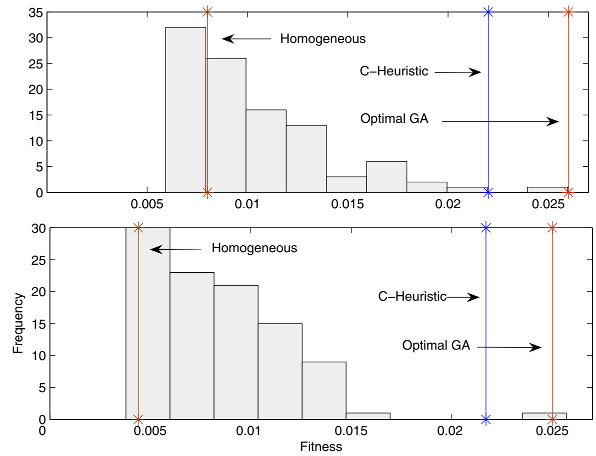

Fig. 9. Fitness distribution of networks that were assigned scrambled radius values from the optimal GA radius configuration

The fitness distribution in both cases has an exponentially decaying form, with the homogeneous network falling towards the lower end. However, the optimized networks are all well outside the typical fitness range, demonstrating their fundamentally atypical character, and the fact that it is not the distribution of radii but their assignment to specific nodes that matters. Interestingly, a very small fraction of scrambled networks also had near-optimal fitness, indicating that the optimal solution found is rare but not unique.

V. Conclusion

We have demonstrated the utility of a reverse-engineering approach using evolutionary algorithms for obtaining local adaptive rules in self-organizing systems. We have shown through simulation results that a heuristic rule obtained for setting transmission radii of nodes in a random sensor network produces heterogeneous networks that are fit, robust and have overall energy consumption comparable to homogeneous networks. We also found that the optimized networks were atypical, and that their performance was not a generic function of their radius distribution. Thus, the heuristic derived through the reverse-engineering process was successful in discovering truly unusual, empirically rejecting the null hypothesis that comparable solutions could have been discovered by random search or a less informed adaptive rule.

The work reported here has suggested two major lines of future work. First, as pointed out above, the process of analyzing the evolved networks to extract the adaptive rule is not sufficiently automated. We are considering more systematic mechanisms for this. Second, the somewhat unexpected increase in robustness found for C-heuristic networks suggests that the adaptive rule could be modified to explicitly produce extremely robust networks. We intend to explore this important issue in future work.

References

[1] T. Arampatzis, J. Lygeros, and S. Manesis, "A survey of applications of wireless sensors and wireless sensor networks," *Proceedings of the IEEE International Symposium on Intelligent Control*, 2005.

[2] D. Stauffer and A. Aharony, *Introduction to Percolation Theory*. London: Taylor & Francis, 1994.

[3] B. Bollobas, *Random Graphs*. London: Academic Press, 1985.

[4] J. Doyle, D. Alderson, L. Li, S. Low, M. Roughan, , S. Shalunov, R. Tanaka, and W. Wilinger, "The "robust yet fragile" nature of the Internet," *Proceedings of the National Academy of Sciences*, vol. 102, no. 41, pp. 14 497–14 502, 2005.

[5] R. Tanaka, M. Csete, and J. Doyle, "Highly optimised global organisation of metabolic networks," *IEE Proceedings - Systems Biology*, vol. 152, no. 4, pp. 179–184, Dec 2005.

[6] R. Tanaka, "Scale-rich metabolic networks," *Physical Review Letters*, vol. 94, no. 16, 2005.

[7] L. Li, D. Alderson, R. Tanaka, J. Doyle, and W. Willinger, "Towards a theory of scale-free graphs: Definition, properties, and implications (extended version)," *arXiv:cond-mat/0501169v1*, 2004.

[8] "Heterogeneous Sensor Networks - Sensor Nets / RFID," *http://www.intel.com/research/exploratory/heterogeneous.htm*, 2005.

[9] C. Bettstetter, "On the connectivity of wireless multihop networks with homogeneous and inhomogeneous range assignment," *Proceedings of the IEEE Vehicular Technology Conference*, 2002.

[10] E. Duarte-Melo and M. Liu, "Analysis of energy consumption and lifetime of heterogeneous wireless sensor networks," *IEEE Global Telecommunications Conference (GLOBECOM)*, vol. 1, pp. 21–25, Nov 2002.

[11] I. Chatzigiannakis, A. Kinalis, and S. Nikoletseas, "Power conservation schemes for energy efficient data propagation in heterogeneous wireless sensor networks," *Proceedings of Simulation Symposium*, pp. 60–71, April 2005.

[12] Y. Ma and J. Aylor, "System lifetime optimization for heterogeneous sensor networks with a hub-spoke technology," *IEEE Transactions on Mobile Computing*, vol. 3, no. 3, pp. 286–294, July-Aug 2004.

[13] X. Du and F. Lin, "Improving routing in sensor networks with heterogeneous sensor nodes," *IEEE Vehicular Technology Conference*, 2005.

[14] ——, "Designing efficient routing protocol for heterogeneous sensor networks," *IEEE Performance, Computing, and Communications Conference (IPCCC)*, April 2005.

[15] P. Ranganathan, A. Ranganathan, A. Minai, and K. Berman, "A self-organizing heuristic for building optimal heterogeneous ad hoc sensor networks," *Accepted for publication in the Proceedings of the IEEE International Conference on Networking, Sensing and Control (ICNSC)*, Nov 2005.

[16] K. Stanley and R. Miikkulainen, "Evolution of neural networks through augmenting topologies," *Evolutionary Computation*, vol. 10, no. 2, pp. 99–127, 2002.

[17] R. Beer, H. Chiel, and J. Gallagher, "Evolution and analysis of model cpgs for walking: Ii. general principles and individual variability," *Journal of Computational Neuroscience*, vol. 7, no. 2, pp. 119–147, 1999.

[18] S. Psujek and R. Beer, "Phenotypic variability in canalized developmental systems," *Proceedings of ALIFE 9*, vol. 9, pp. 415–420, Sep 2004.

[19] T. Rohlf and S. Bornholdt, *Gene regulatory networks: A discrete model of dynamics and topological evolution*. Birkhauser Verlag: Basel, 2004, pp. 233–239.

[20] S. Bornholdt and T. Rohlf, "Topological evolution of dynamical networks: Global criticality from local dynamics," *Physical Review Letters*, vol. 84, no. 26, 2000.

[21] S. Bornholdt and K. Sneppen, "Neutral mutations and punctuated equilibrium in evolving genetic networks," *Physical Review Letters*, vol. 81, no. 1, 1998.

[22] R. Sole, R. Ferrer-Cancho, J. Montoya, and S. Valverde, "Selection, tinkering, and emergence in complex networks," *Complexity*, vol. 8, no. 1, pp. 20–33, 2003.

[23] R. Pastor-Satorras, E. Smith, and R. Sole, "Evolving protein interaction networks through gene duplication," *Journal of Theoretical Biology*, vol. 222, pp. 199–210, 2003.

[24] A. Wagner, "Evolution of gene networks by gene duplications: A mathematical model and its implications on genome organization," *Proceedings of the National Academy of Sciences*, vol. 91, pp. 4387–4391, 1994.

[25] M. Zuniga and B. Krishnamachari, "Optimal transmission radius for flooding in large scale sensor networks," *Journal of Cluster Computing*, vol. 8, no. 2-3, pp. 167–178, July 2005.

[26] A. L. Barabasi and E. Bonabeau, "Scale-free networks," *Scientific American*, pp. 60–69, 2003.

[27] A. L. Barabasi, Z.Dezso, E. Ravasz, S.H.Yook, and Z.Oltvai, "Scale-free and hierarchical structures in complex networks," in *AIP Conference Proceedings*, 2003.

[28] G. Paul, T. Tanizawa, S. Havlin, and H. Stanley, "Optimization of robustness of complex networks," *The European Physics Journal B*, vol. 38, no. 2, pp. 187–191, March 2004.

[29] A. Valente, A. Sarkar, and H. Stone, "Two-peak and three-peak optimal complex networks," *Physical Review Letters*, vol. 92, no. 11, March 2004.

[30] J. Carlson and J. Doyle, "Highly optimized tolerance: Robustness and design in complex systems," *Physical Review Letters*, vol. 84, no. 11, pp. 2529–2532, 2000.

[31] ——, "Complexity and robustness," *Proceedings of the National Academy of Sciences*, vol. 99, no. 1, pp. 2538–2545, 2002.

[32] A.-L. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.

[33] R. Albert, H. Jeong, and A.-L. Barabasi, "Attack and error tolerance of complex networks," *Nature*, vol. 406, no. 6794, pp. 378–382, 2000.

[34] B. Shargel, H. Sayama, I. Epstein, and Y. Bar-Yam, "Optimization of robustness and connectivity in complex networks," *Physical Review Letters*, vol. 90, 2003.