Activation–Inhibition–Based Data Highways for Wireless Sensor Networks*

Daniele Miorandi[1], David Lowe[1,2], and Karina Mabell Gomez[1]

[1] CREATE-NET, v. alla Cascata 56/D, 38123 – Povo, Trento, IT
{daniele.miorandi,david.lowe,karina.gomez}@create-net.org
[2] Centre for Real-Time Information Networks, UTS, Ultimo, NSW 2007, Australia
david.lowe@uts.edu.au

Abstract. In this work we report on a method, based on the use of activation–inhibition mechanisms, for building and maintaining high–rate routing paths (data highways) in dense wireless sensor networks. We describe the algorithms devised and report on the outcomes of a simulation study, aimed at assessing the scalability and the performance of the proposed approach.

Keywords: activation–inhibition mechanisms, reaction–diffusion patterns, data highways, routing, wireless sensor networks.

1 Introduction

The design of efficient, robust and scalable methods for routing data from sources to sink(s) is a major issue in wireless sensor networks (WSNs) [1]. From a communication perspective, traffic patterns in WSNs are of the many–to–one or many–to–some type, depending on the presence of one or multiple data sinks. Data sinks can be either gateway nodes, through which the sensed information, appropriately processed, can be accessed by remote machines, or actuators (e.g., PLCs), where control decisions are taken based on the physical phenomena monitored by the WSN.

Inspired by the work of Franceschetti et al. [5] on the possibility of achieving the capacity bound in randomly deployed wireless ad hoc networks by using a set of high–rate wireless backbones or 'highways' spanning a given network strip, we introduced in a companion work [6] a self–organizing scheme for the distributed construction of data highways in dense WSNs. The scheme is based on the use of activation–inhibition mechanisms [3] for driving the emergence of suitable *spatial patterns*. Similar approaches have recently found application in the wireless networking setting for dealing with activation problems [4,8]. The scheme presented in [6] assumes that all nodes have perfect information on the relative location of data sinks. Such information, which is used to "orient"

---

* This work has been partially supported by the European Commission within the framework of the BIONETS project EU-IST-FET-SAC-FP6-027748, www.bionets.eu

E. Altman et al. (Eds.): Bionetics 2009, LNICST 39, pp. 95–102, 2010.

the activation–inhibition filter along the node–to–sink direction, may not be available in realistic deployments. We therefore investigated a refined version of such a model, where a beaconing mechanism is used to allow each node to estimate the direction towards the sink(s) (i.e. which neighbouring nodes lie on the minimum hop path towards the sink). In this work, we describe the refined model and present some early–stage results derived from an implementation in an event–driven simulator of the aforementioned techniques.

The remainder of the paper is organized as follows. In Sec. 2 we introduce the architecture and the methods for the construction and maintenance of data highways. In Sec. 3 we present the outcome of our simulative study. Sec. 4 concludes the paper presenting a brief overview of the issues left open.

2 Architecture and Methods

2.1 Highway Generation

We wish to develop self-organizing processes that lead to the emergence of data highways in a dense wireless network. These highways should be optimally spaced such that all nodes are within range of a highway (using long–range single–hop communications), but the highways themselves should utilize short–range hops to transport messages to data sinks (to optimize power consumption while limiting interference). The goal is to minimise the number of highway nodes while respecting such constraints. Our problem is therefore to approximate such a desired pattern by using decentralised mechanisms which exploit only local interactions among neighbouring nodes.

We took inspiration, for addressing such a problem, from mechanisms based on activation–inhibition reaction–diffusion techniques [2,8]. These mechanisms describe how field strengths or substance concentrations vary over space and time under two competing influences, a short range positive activation effect, and a longer range negative inhibition effect. The resultant models have been widely used to describe emergent behaviours in biological and physical processes [3]. The simplest formulation of this approach, yet enabling the emergence of a variety of patterns of interest, makes use of a single field variable and can be modelled in the discrete time domain as follows:

$$u(\mathbf{k}, t+1) = g\left[\varphi_s u(\mathbf{k}, t) + \sum_{\mathbf{j} \in R_i} \varphi_i(\mathbf{j}) u(\mathbf{k}+\mathbf{j}, t) + \sum_{\mathbf{j} \in R_a} \varphi_a(\mathbf{j}) u(\mathbf{k}+\mathbf{j}, t)\right], \quad (1)$$

where $\mathbf{k}$ is a physical location, R_i is the inhibition region, R_a the activation region, the activation coefficient φ_a and the self–activation coefficient φ_s are strictly positive, the inhibition coefficient φ_i is strictly negative and $g()$ is a normalizing function.

The equation in (1) corresponds to a two–dimensional time–invariant filter applied to the random field representing the activation level $u(\cdot, \cdot)$. The filter is characterized by the shapes of activation and inhibition regions and by the value of the coefficients $\varphi_s, \varphi_a, \varphi_i$. In [6] it has been shown that, by using asymmetric (or polarised) filters (i.e., filters in which the inhibition and activation regions are not symmetric) it is possible to achieve various spatial patterns. The repeated filter convolution causes the emergence of the ridge peaks in the activation field by activating localized regions that align with the filter axis, whilst inhibiting the off–axis areas between these regions. The width of the filter's inhibition zone controls the separation of the resultant ridge peaks. It is therefore possible to select filter parameters to achieve desired ridge separations. By appropriately adapting the direction of the axis throughout the network, it is possible to change the ridge orientation in different locations within the network. If this is done appropriately then the ridges can be controlled to converge on specific locations – i.e. the data sinks within the network. If we have multiple data sinks then the filter orientation, and hence the ridge orientation, can be derived based on a gravitational attraction model [6]:

$$\mathbf{d_i} = \frac{\sum_{\mathbf{s_j} \in \mathbf{S}} (\mathbf{n_i} - \mathbf{s_j}) |\mathbf{n_i} - \mathbf{s_j}|^{-2}}{\sum_{\mathbf{s_j} \in \mathbf{S}} |\mathbf{n_i} - \mathbf{s_j}|^{-2}}, \tag{2}$$

where $\mathbf{S}$ is the set of sink nodes.

Such a method is based on the assumption that nodes are able to estimate the direction to all data sinks. In the absence of such information, nodes have to estimate which nodes, among the neighbouring ones, are along the shortest path to the sink. This can be achieved by using a suitable beaconing mechanism to acquire knowledge of the distance from the sink(s).

As we are assuming that the nodes *do not* know their own spatial location nor that of the sinks, the convolution filter cannot be oriented based on equation (2). We can however obtain an estimate of the direction to the sink from a given source node based on the sink hop count of the neighbouring nodes, and in particular whether or not they are on the shortest path to the sink. If the neighbourhood has size R (expressed in hops) and the selected node is at distance δ from a given sink, the activation region is constituted by all nodes being at distance n from the node ($n \leq R$) and at distance $\delta \pm n$ from the sink. All other nodes in the neighbourhood are considered to be in the inhibition region. One additional implication is that —rather than determining the weighted direction to all sinks as in (2)— given that each node knows the distance to all sinks, we apply the filter convolution separately for each sink, and hence determine a set of independent ridges for each sink. The highways are then only generated at each node using the ridge data for the closest sink.

2.2 Protocols Description

Formally, we assume the following:

- Nodes are assigned a unique identifier;
- Nodes can tune dynamically their transmission power level P_{tx} in the range $[P_{min}, P_{max}]$;

- The network is connected when all nodes use $P_{tx} = P_{min}$;
- Nodes transmit at P_{min} unless otherwise specified.

The algorithm works according to the following steps:

(i) **Sink announcement**: each sink broadcasts a beacon with its ID. Nodes receiving such a beacon update the `distance` field and re-broadcast it. In such a way each node will eventually have knowledge of all sinks and its distance from them.

(ii) **Neighbourhood discovery**: each node broadcasts a message to its one–hop neighbours, asking for information about them (including activation state) and about nodes which are at most $R - 1$ hops from them. In such a way, each node may acquire (by means of a gossiping mechanism) information on nodes that are within its 'neighbourhood' (at most R hops away).

(iii) **Filter construction and activation level update**: based on the information gathered, each node constructs its local filter and updates its activation level. Let us now consider the specific algorithm for performing the activation–inhibition convolution within each node. If a neighbouring node is n hops from the selected node ($n \leq R$), and either n hops closer to, or n hops further away from, the sink node, then the neighbour will be on the shortest path to/from the sink - and hence can be viewed as being along the filter alignment axis, and should therefore be an activator for the selected node. Conversely, a node that is n hops from the selected node, and not n hops closer to, or n hops further away from sink, will not be on the shortest path, and can therefore be considered to be orthogonal to the filter axis, and will therefore be an inhibitor for the selected node. *Steps (ii) and (iii) are repeated for a number of times in order to achieve a stable spatial pattern.*

(iv) **Ridge detection and tracing**: each node runs a procedure for deciding whether it should act as highway node, by checking whether it is in range of an already existing highway (in which case it just connects to it) or whether it is a local peak, and should therefore activate (i.e., turn into a highway node). If a node is activated, it starts a tracing process by identifying which node shall represent the next hop on the way to the closest sink.

The detailed description of the algorithms, together with their pseudo-code, are reported in [7].

3 Numerical Results

All the algorithms and protocols devised have been implemented in an event–driven network simulator, Omnet++ [9]. The model used was based on the IEEE 802.15.4 standard for PHY and MAC protocols. Our primary goal was to understand how the methods introduced scale with respect to the number of nodes in the network. For scaling the system in a consistent way, we used a fixed node density (set to 0.1429 nodes/m^2, corresponding to an average neighbourhood size equal to 7 when transmitting at the minimum power). The larger the number of nodes, the larger the playground size on which they were placed. Nodes were placed according to a uniform distribution. We used one single sink, located at $(1, 1)$. The maximum and minimum communication distance were setup considering, in the absence of interference, a power transmission range of $(-25, 0)$ dBm, a signal attenuation threshold of -120 dBm and a path loss coefficient of 2. The following set of parameters were used:

Parameter	**Value**
Number of runs	10
Duration of each run (s)	3600
Playground area (m^2)	70/350/700/1400/3500/5250/7000/10500
Number Hosts(n)	10/50/100/200/500/750/1000/1500
Maximum communication distance (m)	12.0
Minimum communication distance (m)	4.0
Neighbourhood size ($hops$)	4
Self-activation coefficient(φ_s)	2.0
Activation coefficient(φ_a)	1.0
Inhibition coefficient(φ_i)	-0.1

In all simulation runs our protocol showed to be able to build valid routes, i.e., all nodes had a valid `nextHop` field, highways were connected to the sink and nodes not on highways were within the maximum communication distance of 12 m from a highway node.

In order to provide insight into the patterns arising in the network as a consequence of the activator–inhibitor mechanism, we report in Fig. 1 the following graphs, related to the case with 750 nodes:

(a) Connectivity graph among nodes when transmitting at minimum transmission power;
(b) Contour plot of resulting distance from the sink;
(c) Resulting activation field after 25 iterations;
(d) Data highways resulting from tracing activation field ridges to the sink.

As it can be seen, highways tend to arise where dense clusters of nodes are present; the ridge tracing process ensures then ability to reach back to the sink. The control of the distance among highways is achieved by setting appropriately selecting the neighbourhood size for the activation filter.

In order to evaluate the performance of our protocol, we preliminarily focused on two metrics. The first one is the time needed by a node in the network to achieve a valid path to the sink node. This corresponds to the time needed to bootstrap a WSN. We considered the minimum, average and maximum value attained for any node over 10 runs. The results are reported in Fig. 2. Both the minimum and average number turn out to be only slightly sensitive to the number of nodes in the system. This is due to the fact that the time needed to construct routing paths from any node to the sink turns out to be dominated by the values of some timers which are part of the routing framework (see Fig. 4 for a workflow–like representation of the overall operations of the protocol). The case of 10 nodes show significantly better performance, due to the simple topology achieved (in most runs all nodes were directly connected to the sink).

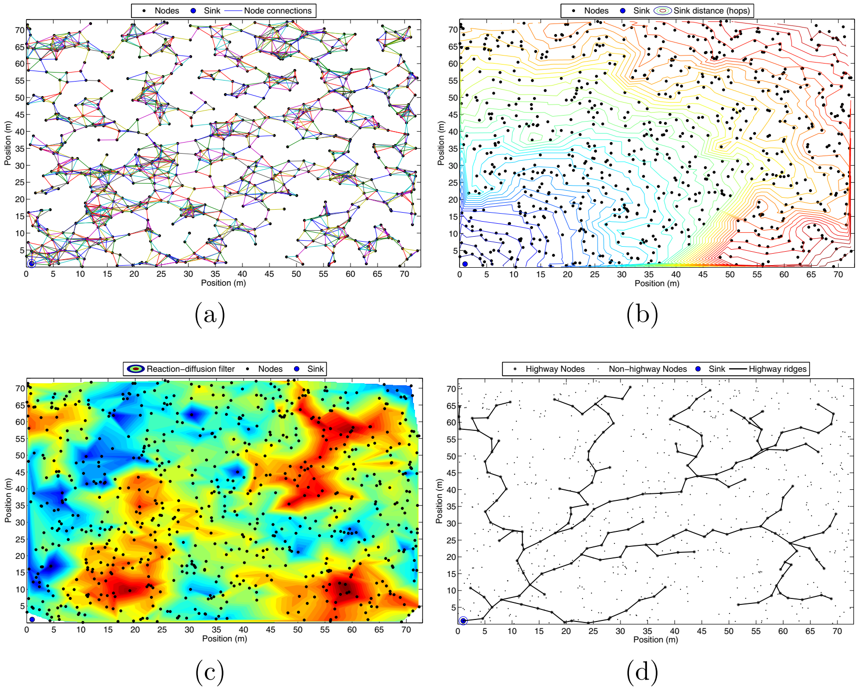

(a) (b)

(c) (d)

Fig. 1. Pattern arising for $N = 750$: (a) Connectivity graph when transmitting at the minimum power level; (b) Contour plot of distance from the sink; (c) Activation field after 25 iterations of the reaction–diffusion filter; (e) Data highways resulting from tracing activation field ridges

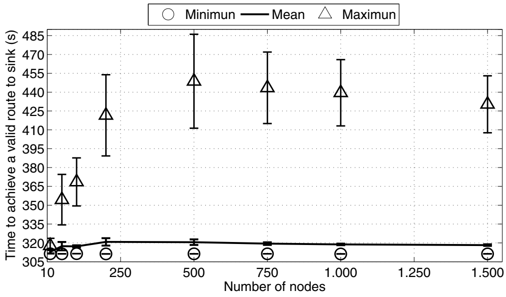

Fig. 2. Bootstrapping time (s) as a function of the network size

The maximum value increased as a function of the number of nodes. A more detailed analysis revealed that this was due to problems related to interference, which prevented some nodes to correctly decode messages destined to them, causing therefore a delay in the setup time.

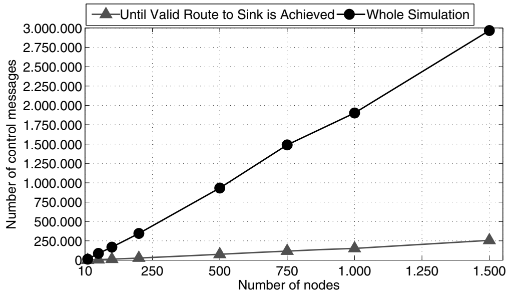

Fig. 3. Number of control messages (exchanged throughout the whole simulation and until a valid route to sink is achieved) as a function of the network size.

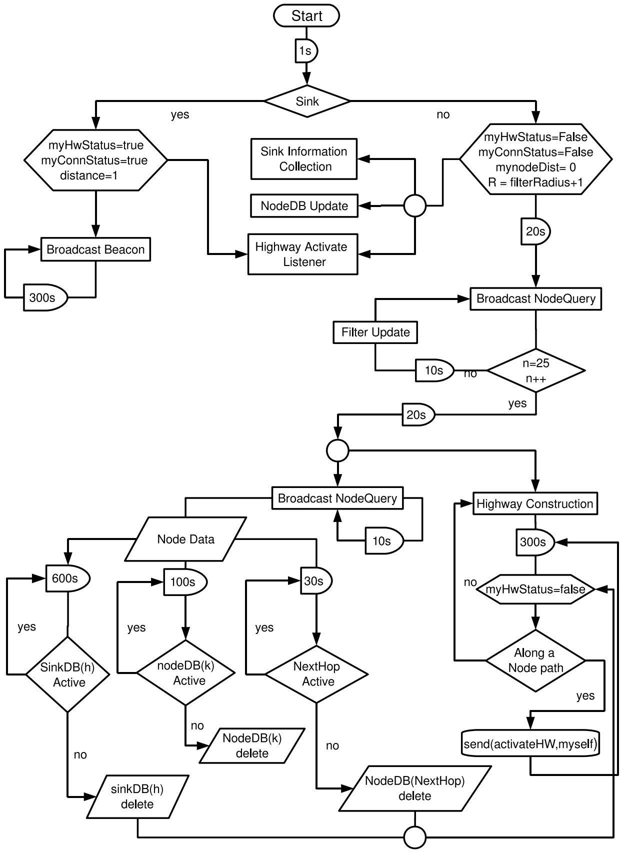

Fig. 4. Workflow–like representation of the operations of the whole routing framework. (Values for duration of timeouts are indicative and have been derived for the network sizes considered in this paper.)

The second performance metric we considered was related to the overhead induced by the protocol in terms of the number of messages exchanged until a valid route to sink is achieved and throughout the whole simulation run (of the duration of 3600 s). As it may be seen from Fig. 3, the number of such a messages scales (slightly) superlinearly in the number of nodes. A more detailed analysis, based on the different types of messages involved in the routing protocol, revealed that all messages involved but sink beacons grow linearly with the number of nodes. The number of sink beacon messages tends on the other hand to grow in a superlinear fashion (actually exponentially), which explains the behaviour observed in our experiments.

4 Conclusions and Discussion

In this paper, we have reported some early–stage numerical results on the performance attainable by a biologically–inspired scheme for the construction of high–rate data highways in dense wireless sensor networks. The results confirmed (i) the correctness of the procedures designed and their ability to meet the requirements in terms of maximal distance to highways (ii) good scalability properties with respect to the number of nodes in the system. Further simulation studies are needed to assess the performance attainable in terms of energy consumption as well as in the ability of the mechanisms designed to recover quickly and with limited overhead from nodes' failures.

References

1. Akyildiz, I.F., Su, W., Sankarasubramaniam, Y., Cayirci, E.: Wireless sensor networks: a survey. Computer Networks 38(4), 393–422 (2002)
2. Bar-Yam, Y.: Dynamics Of Complex Systems. Westview Press (2003)
3. Deutsch, A., Dormann, S.: Cellular automaton modeling of biological pattern formation: characterization, applications, and analysis. Birkhäuser, Basel (2005)
4. Durvy, M., Thiran, P.: Reaction-diffusion based transmission patterns for ad hoc networks. In: Proc. of IEEE INFOCOM, pp. 2195–2205 (2005)
5. Franceschetti, M., Dousse, O., Tse, D.N.C., Thiran, P.: Closing the gap in the capacity of wireless networks via percolation theory. IEEE Transactions on Information Theory 53(3), 1009–1018 (2007)
6. Lowe, D., Miorandi, D.: All roads lead to rome: Data highways for dense wireless sensor networks. In: Proc. of S-Cube 2009: The first international conference on Sensor Systems and Software, ICST (2009)
7. Miorandi, D., Lowe, D., Gomez, K.M.: Distributed bio-inspired methods for data highways construction and maintainance in WSNs. Technical Report 200900022, CREATE–NET (September 2009), `http://www.create-net.org/~dmiorandi/TR200900022.pdf`
8. Neglia, G., Reina, G.: Evaluating activator-inhibitor mechanisms for sensors coordination. In: Proc. of Bionetics, Budapest, Hungary, ICST (2007)
9. OMNeT++ discrete event simulation system, `http://www.omnetpp.org`