

DESIGN AND COMPARATIVE ANALYSIS OF SINGLE-PATH AND EPIDEMIC APPROACHES TO INFORMATION AND ENERGY MANAGEMENT IN WIRELESS SENSOR NETWORKS

D. Zilli, G.V. Merrett

Electronics and Computer Science, University of Southampton, Hampshire, SO17 1BJ, UK
Email: {dz2v07;gvm}@ecs.soton.ac.uk – Tel: +44 (0)23 8059 4996

Keywords: wireless sensor networks, flooding, energy management

Abstract

Intelligent energy management is a key challenge in Wireless Sensor Networks. The choice of an appropriate routing algorithm constitutes a critical factor, especially in unstructured networks where, due to their dynamic nature, a reactive routing protocol is necessary. Such networks often favour packet flooding to fulfil this need. One such algorithm is IDEALS, a technique proposed in the literature, which balances energy consumed with information delivered. This paper evaluates the use of a single-path solution with IDEALS to increase efficiency. Simulation results comparing the two approaches show that the single-path algorithm outperforms flooding in terms of energy consumption for any network size. Furthermore the benefit of IDEALS is preserved as its combination with the single-path algorithm maximises information throughput.

1 Introduction

Energy efficient operation is a key challenge in Wireless Sensor Networks as nodes are usually locally powered and hence highly energy constrained [1]. Energy can be saved through efficient hardware design, communication algorithms or adaptive sampling techniques [2]. Alternatively, energy can be obtained from the node's environment by harvesting from sources such as sun light, mechanical vibration or temperature gradients [3].

Networks featuring energy harvesting enable nodes to constantly leave and rejoin the network as they deplete and refill their energy reserves. This almost sporadic operation can be further aggravated by the algorithms operating on the node. As nodes continually drop in and out, maintaining paths from source to destination becomes challenging. Unstructured networks require reactive routing algorithms that are able to identify routes on-demand [4]. Therefore, in these networks, an epidemic “flooding” approach is often used to reactively route packets without performing continuous route discovery and repair. However, it is widely accepted that, while offering good end-to-end reliability, epidemic approaches to routing are

highly energy (and bandwidth) inefficient [2].

An example of an algorithm causing an unstructured ‘sporadic’ network, which uses flooding, is IDEALS [5], which controls packet transmission in order to be energy-efficient while maintaining sufficient resources to guarantee delivery of important packets. Other techniques have since been proposed with a similar aim, and have already adopted a single-path algorithm—an example of this is Randomised-Re-Routing [6], which attempts to improve the quality of service by transmitting routine packets on secondary paths when critical information needs to be delivered.

In this paper we investigate whether, to address the lack of energy-efficiency inherent in packet flooding, it is possible to use a single-path reactive routing algorithm with IDEALS, while still abiding by the constraints of such an unstructured network. The routing algorithm proposed is not considered a contribution of this research, rather an analysis of the impact it has on IDEALS. The two approaches are evaluated against each other through simulation, and analysed for energy consumption, packet success, and scalability. The results show that a) it is possible to use a single-path routing algorithm with IDEALS and b) considerable benefits can be obtained in terms of energy consumption and information throughput.

2 IDEALS

For a full description of IDEALS, the reader is referred to the paper by Merrett et al. [5]. However, this section provides a brief summary for conciseness. IDEALS is a technique proposed to extend the lifetime of a wireless sensor network by selectively discriminating messages based on their contents. It attempts to guarantee the delivery of important packets even when energy reserves are scarce, at the cost of sacrificing messages with lower importance. Combined with a rule-based scheme to determine the information content of a packet, IDEALS instructs a node to relay or discard the packet depending on energy reserves. It is particularly suited to networks that use energy harvesting, as the network can find an energy equilibrium and operate indefinitely.

At any one time, each node has an Energy Priority (*EP*), which reflects the energy level of the node. In the literature, this parameter is (but does not have to be) in the range 0 to 5 [5]. An




Figure 1: Mean residual energy in a network featuring combinations of energy harvesting and IDEALS.

EP of 5 indicates that the battery is full, and 0 that it is empty. When a node generates a packet to report an event, it similarly assigns the packet a Packet Priority (*PP*) according to its information content. On receipt of a packet to forward, a node decides whether to relay or discard it according to the importance of the packet (i.e. its *PP*), and the nodes' energy resources, (i.e. their *EP*). Only nodes with abundant energy will participate in routing unimportant packets, while all nodes will route important ones. This means that packets of high importance observe a dense and well-connected sub-network, while low-priority packets see a sparser and less-connected sub-network. Through this, independent and selfish decisions made at individual nodes help manage and balance the energy state of the network. This is observed in Figure 1, where the use of IDEALS ensures an equal or greater residual energy in the network at any given point in time. As nodes can regularly fall in and out of sub-networks (a problem exacerbated by the usage of energy harvesting, which permits nodes to increase as well as decrease their energy resources), the network can be considered highly unreliable, mobile and unstructured.

IDEALS uses a simple reactive flooding algorithm, so that nodes deciding not to relay a packet do not impede its successful delivery and no measure is needed to find an alternative route. As the overheads required in maintaining routing tables make proactive algorithms unsuitable (as nodes are continually dropping in and out of sub-networks), packet flooding minimises network complexity though being energy inefficient.

3 A Single Path Approach to Routing

To address the shortcoming of high energy waste due to flooding, in this paper IDEALS has been extended to allow a single-path, reactive, approach to packet routing. It is assumed that nodes are initially aware of their location with respect to the base station and to their neighbours. Geographical algorithms are attractive in these scenarios, as paths can then be established when required using only local knowledge of where a node is with respect to the sink [7]. As the network is not physically mobile, this information could be programmed onto nodes at deployment and never changed, or identified at run-

time [8]. Nodes route using a reactive protocol, only searching for paths on-demand. When a node has a packet to transmit, it selects its neighbour who is closest to the sink. If the selected node decides that its energy reserves are insufficient to forward the packet, it refrains from responding to the originator (using the principles of IDEALS). In turn, having waited a set timeout, this node elects a different neighbour to route the packet, namely the second closest to the base station, and so forth until one with sufficient energy is found. If a node fails to find a willing neighbour, it sends a negative acknowledgement one step backwards along the propagated route so that the predecessor can itself find a different branch of the tree of nodes with sufficient energy to reach the destination.

The algorithm, exemplified in Figure 2, is guaranteed to succeed should there be even just one path available. This simple method, however, cannot promise that the path found to destination is necessarily the 'best' available one. Moreover, acknowledgements are costly as, effectively, two packets are sent for each one that must be delivered. In some circumstances, this could result in worse performance than flooding. This method is therefore not intended to demonstrate the best single-path routing algorithm, but aims at providing an 'upper bound', i.e. the simplest single-path solution that can compete with packet flooding in terms of packets delivered.

4 Simulation Environment

To explore the comparison between flooding and single-path approaches, a simulation model has been designed and implemented. Nodes are scattered arbitrarily across a 2-dimensional plane, but are modelled to be sensibly spaced and never initially isolated. Randomness makes the system generic, and is also realistic in several different scenarios. The base station is also placed in a random location and guaranteed to be able to communicate directly with at least one node. The model is run in a discrete time simulation where, at each time step, a node is assigned a probability of sensing a noteworthy event. In addition




Figure 2: The operation of the proposed reactive single-path routing approach. Thin red arrows show how sometimes many attempts are required before finding a suitable path.

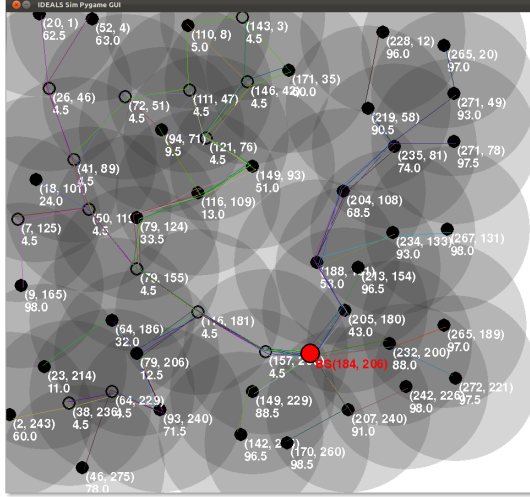


Figure 3: The simulator's User Interface. Black filled circles are active nodes, outlined ones have depleted their energy. In red is the base station. Each node shows coordinates and energy left.

to monitoring packet generation and routing, the model also considers the energy of each node. The energy E consumed by a node i at any given time step t is defined as:

$$E[i, t] = P[t] \cdot E_s[i] - E_h[i] \quad (1)$$

where E_h is the energy gained from harvesting, E_s the energy consumed sending a packet and $P[t]$ the number of packets sent, determined by the probability $p(s)$ of sending a packet in a given time step (see Table 1). It follows that the average energy in the network at a certain time step is given by:

$$\bar{E}[\tau] = E_{init} - \frac{1}{n} \sum_{t=1}^{\tau} \sum_{i=1}^n E[i, t] \quad (2)$$

where E_{init} is the initial energy of nodes and n the number of nodes.

In the simplest case, the model simulates a network where flooding is in use. When a node wishes to send a packet, it relays it to all its neighbours with the request of delivering it to the sink. Each of them performs the same action. A packet ID prevents the same packet from being forwarded more than once by each node, breaking an otherwise infinite loop of packets. Acknowledgement are not sent, as this would be unnecessary and counterproductive. A node's routing behaviour is summarised by the following recursive functions:

```
function RECEIVE(PP, ID):
  if not seen ID yet then
    if node is alive and node's EP > PP then
      if node is not sink then
        send(PP, ID)
function SEND(MP, ID):
  consume_energy()
  for all node in neighbours do
    node.receive(PP, ID)
```

IDEALS is implemented on top of this basic model, with five levels of Energy Priority to reflect the literature (see section 2). Five different Packet Priorities are assigned with equal probability. Additionally, nodes harvest energy from the environment; this further extends the battery lifetime, to the extent that careful energy use can allow perpetual operation [5]. For simplicity and ease of evaluation, energy harvesting provides each node with a fixed amount of energy, E_h , each timestep.

The single-path routing approach is also implemented. Here the cost of sending a packet is doubled, as each must be acknowledged. A node generating a packet discovers a path, sends the packet and receives an acknowledgement according to the following recursive functions:

```
function SEND(PP, path):
  path ← path + self
  if node is sink then return ackn(PP, reversed(path))
  else
    while node in neighbours do
      success ← node.send(PP, path)
      node.consume_energy()
      if success then return ackn(PP, path)
      else
        neighbours ← neighbours - node
function ACKN(PP, path)
  if self is originator then return success
  else
    self.consume_energy()
    success ← next.ackn(PP, path-self)
    if success then return success
    elsereturn fail
```

The use of IDEALS and energy harvesting can be combined to allow four different simulations: basic, IDEALS only, energy harvesting only, and both. Furthermore, these are implemented for the two routing algorithms, generating a total of eight scenarios. These have been run across several different parameters and results are presented in the following section.

For visualisation, a two-dimensional Graphical User Interface (GUI) represents the network with sensing nodes, their energy, the base station and the paths between them, when available (Figure 3). The GUI can be run interactively to allow for easy inspection of the model's operation. The model, written in Python, fully leverages the multi-core environment through use of the multiprocessing library.

5 Simulation Results

The aim of IDEALS is to manage energy resources while ensuring important packets are received. This section evaluates the ability of a single-path approach to routing to fulfil this aim and compares performance with flooding. Note, results are only shown for a network with energy harvesting. In order to reduce the significance of randomness in the model, the simulation has been repeated over 10,000 iterations and results averaged and analysed in R . Table 1 summarises key parameters.

Name	Default	Description
nodes (n)	5-100	nodes in the network
range	100	range of nodes
E_{init}	100	starting energy per node
E_s	0.5	cost of sending a packet
E_h	0.004	energy harvested per timestep
$p(s)$	0.1	prob. of sending packet in timestep
EP	$0.2 \cdot E_{init}$	length of EP interval (here 5 intervals)
area_height	$\text{int}(\text{floor}(\text{sqrt}(1500 \cdot \text{nodes})))$	
area_width	$\text{int}(1500 \cdot \text{nodes} / \text{area_height})$	

Table 1: Simulation parameters used to obtain the results shown.

5.1 Energy

Figure 4 shows the average energy in the network for combinations of flooding and single-path routing and with and without IDEALS. To demonstrate how the two algorithms scale, the simulation has been run with a number n of nodes in the range from 5 to 100, of which the instances with $n = 30$ and $n = 80$ are shown here. These effectively demonstrate the scalability of the two approaches. It can be observed how in both simulation runs shown, flooding is the first to deplete the energy of the network. Where IDEALS is not present, residual energy drops almost linearly and once depleted, nodes continue to occasionally send packets due to energy harvesting. In the presence of IDEALS, however, a smoother curve can be observed. Nodes falling below a certain battery level stop forwarding packets whose Packet Priority is lower than their Energy Priority. The energy save is sufficiently efficient that it prevents the network from dying on $n = 30$ in the window of 150 time steps. However this could—and does—affect the number of packets received, as shown in subsection 5.2.

The single-path approach proves to be more efficient. Here the nodes do not all have the same amount of energy at any given time step. This is because those close to the base station will have received and forwarded packets from the entire network, while nodes that are further away may not have trans-

mitted any packets except those generated by themselves. For all network sizes, however, results show that the average energy present in the network is higher for the single-path approach than for flooding. IDEALS is beneficial in single-path too, where more energy is saved compared to when it is not present. It is important to notice that IDEALS becomes extremely useful when, as in the case of single-path $n = 80$, the network has found an *equilibrium* so that sufficient energy is harvested to be able to send all high-importance packets without exhausting the nodes.

Despite the positive results in terms of average energy, this comparison must consider the number of packets received to draw fair conclusions. Nodes could consume less energy just by sending fewer packets, which would clearly not provide the same service. The following section discusses results in terms of packets sent and received.

5.2 Packet success

Figure 5 shows the distribution of packets received in a network of 50 nodes where energy harvesting is in use. The total number of packets is divided into the 5 IDEALS *PP* considered, which are shown in different colour shades. Over the 10,000 simulation runs, on average one packet of each Packet Priority will be sent per time step (provided that there is sufficient energy in the network).

In the non-IDEALS flooding model (Figure 5a), nodes continue sending packets until there is no energy left, which quickly drains the energy. After, harvesting allows for the occasional packet to be received but, as there is no concept of information content, packets of any Packet Priority have equal likelihood of being delivered. In Figure 5b, IDEALS provides graceful degradation, selectively discarding low importance packets as the energy in the network drops. The changes in Power Priority are therefore clearly visible. Unimportant packets stop being received at an earlier stage, but this guarantees that an equilibrium is found (around $t = 80$), whereby critical information is always delivered. The single-path approach in Figure 5c




Figure 4: Energy residual in nodes during the simulation for the two different routing algorithms with and without the use of IDEALS, in a network featuring 30 nodes (left) and 80 nodes (right). The lines represent the mean energy across all nodes.




Figure 5: Total number of packets received per time step at the sink, with $n = 50$. It can be observed how single-path improves the number of packets delivered compared to flooding, while IDEALS improves the information contents (i.e. the amount of highly-important packets).

performs better than flooding, as the more efficient algorithm allows for longer network lifetime. While improving on the total number of packets received, single-path alone cannot guarantee that the most information is transmitted to the sink, as all packets have equal likelihood of being relayed. Its combination with IDEALS (Figure 5d) therefore brings the best results by maximising both the total number of packets received and the information content of those.

It can be concluded that the combination of the energy results (Figure 4) and packets received (Figure 5) confirm the validity of a single-path approach improving the existing solution.

6 Conclusions

This paper has compared and contrasted through simulation two routing approaches for an existing WSN energy management technique, called IDEALS. The first, packet flooding, replicates the original model on which IDEALS was built. The second, a single-path approach, aims at reducing the energy consumption without negating the value of IDEALS.

Results show that the single path approach is considerably more energy efficient than flooding. The number of packets being successfully received is not reduced and, to some extent, improved as a result of the energy savings. The combination of energy harvesting and IDEALS with the single-path algorithm proved to deliver the highest information throughput, confirming the validity of the solution. In future work, the model will be extended by considering the amount of energy consumed by nodes when receiving packets. Comparison with similar systems in the literature will also be carried out to test the performance of single-path IDEALS against the state-of-the-art.

References

- [1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks : a survey," *Computer Networks*, vol. 38, pp. 393–422, 2002.
- [2] S. Beeby and N. White, *Energy Harvesting for Autonomous Systems*. Artech House Publishers, 2010.
- [3] A. Kansal and M. Srivastava, "An environmental energy harvesting framework for sensor networks," *Int'l Symp. on Low Power Electronics & Design*, pp. 481–486, 2003.
- [4] Q. Zhao and L. Tong, "Energy efficiency of large-scale wireless networks: proactive vs. reactive networking," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 5, 2005.
- [5] G. Merrett, N. Harris, B. Al-Hashimi, and N. White, "Energy managed reporting for wireless sensor networks," *Sensors & Actuators A*, vol. 142, no. 1, pp. 379–389, 2008.
- [6] E. Gelenbe and E. C.-H. Ngai, "Adaptive QoS routing for significant events in wireless sensor networks," in *2008 5th IEEE International Conference on Mobile Ad Hoc and Sensor Systems*, pp. 410–415, IEEE, Sept. 2008.
- [7] J. Wan, D. Yuan, and X. Xu, "A review of routing protocols in wireless sensor networks," in *Int'l Conf. on Wireless Communications, Networking and Mobile Computing (WiCOM'08)*, pp. 1–4, 2008.
- [8] F. Franceschini, M. Galetto, D. Maisano, and L. Mastrogiacomio, "A review of localization algorithms for distributed wireless sensor networks in manufacturing," *Int'l J. Computer Integrated Manufacturing*, vol. 22, no. 7, pp. 698–716, 2009.