# Routing In Cellular Sensornets With Uniquely Identified Destination Nodes

Jonathan Tate and Iain Bate

June 2010

### Abstract

Sensornet nodes observe physical phenomena, yielding data labelled with their geographic position. Geographic context can be exploited in packet routing to minimise energy consumption, provided that the application does not need to route packets to uniquely identified nodes. The latter may be necessary for a minority of packets concerning system management and sensor tasking. We present a hybrid approach to packet routing for cellular sensornets. A geography-aware mechanism routes packets between cells, which is sufficient for the majority of packets. For the minority of packets which must be delivered to uniquely identified nodes, a geography-ignorant mechanism handles the final stages of packet delivery within the cell containing the destination node.

## 1 Introduction

Wireless sensor networks, or sensornets, compose many autonomous *motes* into ad hoc networks for distributed sensing and processing applications [1]. Motes are small, cheap computers equipped with independent power supplies, wireless communication capability, and sensors with which to passively monitor with their environment. Sensor-actuator networks also interface motes with actuators to actively influence the environment, completing the control feedback loop. Typical applications include industrial process control, habitat monitoring, precision agriculture, and military surveillance [2].

Geographical location, rather than node identity, is generally used to label sensor data and data flow endpoints [3]. This follows from the assumption that geographic locations of data production and consumption, rather than globally unique identifiers for individual sensornet nodes, are more appropriate for data labelling in geography-aware sensornet applications [4].

*Clustering* methods are commonly employed in sensornets to enable scalable distributed sensing and processing applications [5]. Each cluster composes a number of independent nodes, located within a small geographic region, into a higher level structure which can be effectively addressed as a single entity. One node from each cluster is selected as the *clusterhead*, and is responsible for managing activity within its cluster, and managing interaction with other clusters [6]. Clustering may allow sensornets to achieve greater scalability and reduce energy consumption [7]. However, it is not guaranteed that clustered sensornets outperform non-clustered sensornets [8], and few comparative studies exist.

In a geography-aware sensornet, *cellular clustering* [9] is an obvious and reasonable approach to the definition and formation of clusters. The physical region occupied by the sensornet is divided into *cells*. Any node located within the geographical boundaries of a given cell is considered to be a member of that cell. This avoids the requirement for complex and resource-consuming cluster membership assignment protocols, and avoids non-deterministic cluster allocation [10, 11]. However, this requires that nodes be aware of their location and the geographic cellular structure. Sensornet motes may be equipped with geolocation hardware, such as GPS receivers [2], or may employ other location inference methods to establish the location of each node relative to known fixed points or other network nodes [12, 13].

We assume that all nodes within a cell are equal peers, such that all cell members can support an equal share of the processing and storage burden for the distributed application. The failure of a single node may temporarily reduce the capability of the cell, but should not cause the entire cell to fail. This implies either a homogeneous cell population, or a heterogeneous population in which all members have at least the resources required to support an equal share of the application workload.

Given that nodes sharing a cell are physically close, we can route packets between *cells* rather than *nodes* [14]. As network management instructions, raw sensor data, and processed information, flow through the network, they can be directed between cells by any nodes which happen to be assigned to these cells. The decision as to which of the cell members is responsible for handling a given packet at a given time can be managed by a duty allocation protocol such as CDAP [15].

The primary contribution of this paper is a packet routing protocol which yields shorter routes, and hence reduced delivery latencies and energy costs. The *Implicit Token Ring* (ITR) approach, as defined in this paper, can be used to deliver packets to nodes with globally unique identifiers. ITR is a hybrid protocol, combining the best aspects of geography-aware and geography-ignorant protocols, and the energy efficiency of cellular network designs.

Packets are first delivered across the network from a source node to some intermediary node within the same cell as the uniquely identified destination node, employing a geography-aware and identity-ignorant approach. This is equivalent to an unmodified geographical routing approach, and makes no assumptions as to the current state of any nodes in the cell containing the destination geographical location. If the intermediary node is not the destination node the packet is then delivered within the cell, employing a geography-ignorant and identity-aware approach.

We assume that the majority of packet delivery attempts are not to uniquely-identified nodes. Although ITR will function correctly if this assumption does not hold true, the design of ITR has been optimised for this use case. We also assume packet source nodes are aware of the general geographic location of the destination node, and encode this into the packet header. This assumption is trivially true if the globally unique identifier for a node is the geographical position of that node, but position data need only be sufficiently accurate to identify the general direction of the cell containing the destination.

Low-level network issues, such as unreliable media, packet loss, congestion, and hidden terminals, are generally addressed by protocols at the *Data Link Layer* [16], and are therefore invisible to protocols such as ITR which reside in the *Network Layer*. We therefore assume that these issues are addressed by an underlying MAC protocol [17], the details of which are not considered here.

## 2 Related work

The *Token Ring* Local Area Network architecture [17] resides at the *Data Link Layer* [16]. Logically, nodes are composed into a connected ring structure, with each node connected to exactly two neighbouring nodes. A special frame called the *token* is passed from node to node around this ring, usually in a single direction, in an endless loop. When the token reaches a node waiting to transmit data, the token frame is converted into a data frame. This data frame is passed from node to node around the ring, until it reaches the node at which the data frame originated, in a similar fashion to the token frame. No other nodes are permitted to transmit data frames during this process. This continues until it reaches the node at which the data frame originated, and is replaced with another token frame.

Physically, wired connections may exist between nodes which match the logical ring topology. Alternatively, the nodes may be connected in a star topology, where each logical connection between ring neighbours is implemented by a wired connection to a central hub, and back out to the next node in the logical ring. In a wireless network either of these approaches can be emulated, with all exchange of token and data packets implemented through a shared wireless medium.

Token Ring networks have some theoretical advantages over stochastic CSMA networks [17] such as Ethernet [18] or ALOHA [19]. The most significant advantage for embedded systems is that Token Ring network behaviour is deterministic and predictable. Token Ring is egalitarian; every node is fairly allocated equal access to the network, and delivery latencies have predictable upper bounds. Stochastic protocols are unpredictable by design, and are focused on optimising the *average case* performance rather than the *worst case* performance. Stochastic protocols tend to perform very poorly when network contention is high [20], and are more difficult to analyse for safety critical applications.

There are also disadvantages with the Token Ring approach [17]. Although packet latencies in Token Ring networks can be predicted and bounded accurately in the worst case, they may be greater than in stochastic networks in the average case. Larger rings imply a greater probability of failure, if we assume each node-to-node hop may fail with non-zero probability. In unreliable networks the token may be lost, and a single failed physical connection may disable the entire ring by partitioning.

Missing tokens can be addressed by assigning one node the role of *Active Monitor*, with responsibility for detecting and replacing missing tokens. In a sensornet there is no guarantee that this single Active Monitor node will reliably detect this condition, and it is prone to failure as with any other node [20]. Failed connections can be addressed by having the token change direction when the holding node detects a missing connection to its ring successor. This assumes bidirectional communications links, which are not always available in sensornets [21].

The *Cambridge Ring* system [22] was the first practical Token Ring implementation. The topology places the entire network into a single ring. The simplicity of this structure is advantageous when reasoning about system behaviour and performance, and effective implementations can be realised with simple hardware and software. However, scalability is poor. End-to-end latency increases linearly in the number of nodes, as the token must visit all nodes while completing a single circuit of the network, with a commensurate decline in throughput.

The *IBM Token Ring Network* (TRN) [20] built upon the fundamental concepts developed in the Cambridge Ring. Point-to-point wiring between nodes was replaced with a star topology, rendering TRN much easier to deploy in realistic deployment environments. Scalability was improved by allowing topologies composed of multiple rings connected through bridges. The requirement for dedicated hardware, and particular wiring topologies, renders TRN unsuitable for the domain of sensornets.

The international standard IEEE 802.4 [23] defines a *Token Bus* architecture for Local Area Networks. Token Bus networks are similar to Token Ring networks, but all nodes are connected at all times through a shared bus. The current holder of the token is permitted to broadcast into the shared medium. The broadcast is received simultaneously by all nodes, including the intended destination node. This is in contrast to a Token Ring, in which each node is connected (either physically or logically) to only its immediate predecessor and successor in the ring.

For a sensornet in which all nodes are always active, and continuously interacting through a single shared wireless medium, a Token Bus regime may be viable. In a typical sensornet, however, most nodes will be dormant for a significant proportion of time in order to save energy and extend network lifetime. Without a mechanism to ensure that both the source and destination nodes are simultaneously active, it is not possible to guarantee that a Token Bus network will successfully deliver any packets at all.

The international standard IEEE 802.5 [24] defines a Token Ring architecture for Local Area Networks. It is based on the the IBM Token Ring Network. There are minor differences, but these do not result in significantly different performance. *Fiber Distributed Data Interface* (FDDI) [25] also extends the IBM Token Ring Network mechanism. FDDI adds a secondary ring, of equal performance to the primary ring. This can either be used to double network throughput, or as backup for the primary ring. In wireless networks it is difficult to implement a secondary ring without multiple transceivers.

The *Wireless Token Ring Protocol* (WTRP) [26] is a Token Ring MAC protocol for wireless Local Area Networks, based on the IEEE 802.4 Token Bus standard [23]. It supports many topologies; it is not necessary for all stations to be fully connected, or for a central base station to exist. It is intended as an alternative for the stochastic IEEE 802.11 protocol [27] in applications where guaranteed QoS is important, defined in terms of bounded latency and reserved bandwidth, and to improve power usage and reliability by reducing collisions and the retransmissions they induce. However, a special token packet must be passed from node to node, and a Management Information Base must be maintained, with concomitant energy and storage costs.

## 3 Cellular network structure

We assume that physical cell boundaries do not overlap [28]. Wireless packet broadcasts may cross boundaries into neighbouring cells, however, unless physical barriers exist around the cell perimeter. This enables packets to be exchanged between neighbouring cells, which is of critical importance in maintaining a network in which any node is reachable from any other node. Wireless broadcast power modulation schemes [29] can be employed to ensure that broadcasts are sufficiently powerful to be receivable within the same cell, and neighbouring cells, but not so powerful as to interfere with distant cells and induce long-range disruption [30]. We are not concerned with the physical distribution of nodes within cells, which may be uniform or non-uniform, depending on the design of the network, and the method of node deployment,

We are not concerned with the physical cellular structure; we need only know which cells are *adjacent*; by this we mean the set of potentially interacting pairs of cells. In most cases *logical adjacency* is implied by *physical adjacency* owing to the characteristics of wireless broadcast communications. It is possible for more complex sensornets to include one or more direct wired links between distance nodes, effectively giving *logical adjacency* to a pair of cells which do not share *physical adjacency.* We do not consider this situation explicitly as it has no bearing on our analysis.

At any given time any non-degenerate cell contains one or more nodes. Each node resides in exactly one cell. Each node knows to which cell it belongs at this time, but does not need to know the other members, if any, of this cell.

Any pair of nodes within a cell can exchange packets if simultaneously active and listening to the wireless medium [29]. However, mechanisms intended to minimise energy conservation, packet interference, and network congestion, often require nodes to spend significant proportions of total runtime in dormant states or to refrain from interaction with the shared wireless medium [15, 31, 32]. At any given time, a subset of nodes located within a cell are simultaneously active; only node pairs drawn from this active subset are able to exchange packets.

If full connectivity between all nodes in a cell is not possible for a given cellular configuration, perhaps as a consequence of physical obstacles or voids, one possible solution is to reconfigure the network with smaller cells until full connectivity is re-established. We also assume that for each pair of adjacent cells there exists at least one pair of nodes, with one node per cell, which is able to communicate. This is essential for adjacent cells to exchange application and network management data.

Unless each cell is a *singleton cell* containing exactly one node, the logical network is smaller than the physical network. It follows that protocols which become less reliable as the network size increases will benefit from addressing the network at the level of the logical cell rather than the physical node. As fewer physical nodes are involved, it is possible to achieve corresponding decreases in energy consumption and network congestion among other factors.

It is possible for the communication range of a node to cover a physical region extending beyond its own cell. This is essential for intercellular communication; it is possible only if there exists at least one pair of nodes, split between the cells, within mutual communication range. Nodes can modulate the broadcast power of their wireless communications modules to exert influence on the physical region within which pairwise exchange is feasible [29], but this is beyond the scope of this paper.

Consider a large sensornet consisting of many nodes, divided into cells containing smaller numbers of nodes in close geographic proximity [28]. Within a cell each node has a similar view of the physical environment, and similar connectivity to nearby base stations or surrounding cells [33]. All nodes within a cell are approximately equivalent with respect to extracellular entities and environmental context.

Suppose that an external entity broadcasts a message received by all members of a cell. Unless the message is intended for a specific member of that cell, it is unclear which cell member or set of cell members should respond. Data packets to be forwarded to remote destinations need only be rebroadcast once; if all cell members rebroadcast this wastes energy, increases contention for the wireless medium, and risks collisions [34]. If a tasking message requests that a sample value be read from the physical environment then all cell members will produce equivalent readings [35]. Consequently, energy and network capacity may be wasted in delivering multiple redundant messages. In-network aggregation and processing could become slower and more costly as the volume of data increases. However, a certain level of redundancy is useful, as this provides robustness against the failure of individual nodes, and multiple sensors can cooperate to yield readings of higher accuracy and stability.

## 4 Routing packets in cellular sensornets

Conventional networks generally consider each individual entity to be significant, with network routing decisions implemented on the logical structure of the network, which may or may not correlate with the physical layout of the nodes or interconnections. Section 3 discusses the concept of routing between network cells using geographical context, rather than routing between individual network nodes using globally unique identifiers.

This is usually sufficient in a sensornet in which individual nodes are unimportant, and in which application data and routing decisions are defined in terms of geographical location. If any node within a certain small range of a physical location is equally capable of taking a sensor reading, forwarding a data packet, or performing any other duty of the distributed application, the identity of the node selected from the set of all suitable candidates is unimportant.

However, under certain circumstances the network operator may wish to deliver a packet between a specific arbitrary source and destination node pair. Clearly this is impossible in the general case if the network does not have globally unique node identifiers. Assuming that globally unique node identifiers exist, there remains the problem of determining what to do with the packet at each stage of delivery along a multi-hop path. It is possible to retain most of the efficiency advantages of a cellular network if we divide the task into two subproblems:

Subproblem 1: *Deliver the packet to the cell containing the destination node.*

This subproblem is addressed by any geographical routing protocol, such as the *Implicit Geographic Forwarding* (IGF) protocol [36]. The routing and delivery process begins at the source node $s$, and terminates when the packet reaches an arbitrary node $a$ within the boundaries of the cell containing the destination node, $d$. This implies that either the packet must include the geographical location of the destination node as well as its globally unique logical identifier, or for all cells it must be true that every node within the cell knows the logical identifier of all other nodes in the cell. The internal details of geographical routing protocols are not considered further here, but are widely discussed in the literature, e.g. [37].

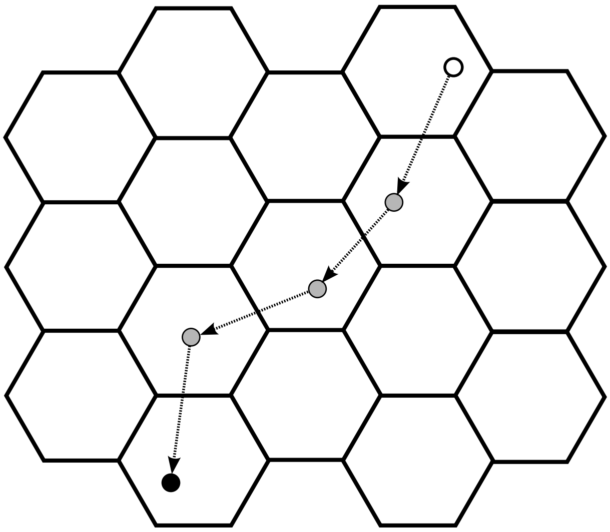

Figure 1: Delivering a packet between cells

Figure 1 illustrates the route traversed by a packet from a source node, depicted by a white circle, to some node in the cell containing the destination node, depicted by a black circle. This latter node is not necessarily the destination node. Relay nodes are indicated by grey circles. Nodes which do not participate directly are not shown.

It can be seen that for cells which contribute to the packet delivery attempt there is exactly 1 participant node per cell, and for cells which do not contribute there are exactly 0 participant nodes per cell. Note that the packet gets closer to the destination cell on each hop, as a geographical routing protocol is used for intercellular packet transport.

Subproblem 2: *Deliver the packet within the destination cell to the destination node.*

When the packet has reached some node in the destination cell there remains the problem of delivering it to the final destination. If $a = d$ then no further action is necessary. If we assume that any given pair of nodes within a cell can directly exchange packets, this is equivalent to a single hop delivery path from $a$ to $d$. If all nodes within the destination cell were active simultaneously at some time, then node $a$ could broadcast the packet directly to node $d$. When implementing a duty management protocol such as CDAP [15] to reduce energy consumption, however, some additional mechanism is required.

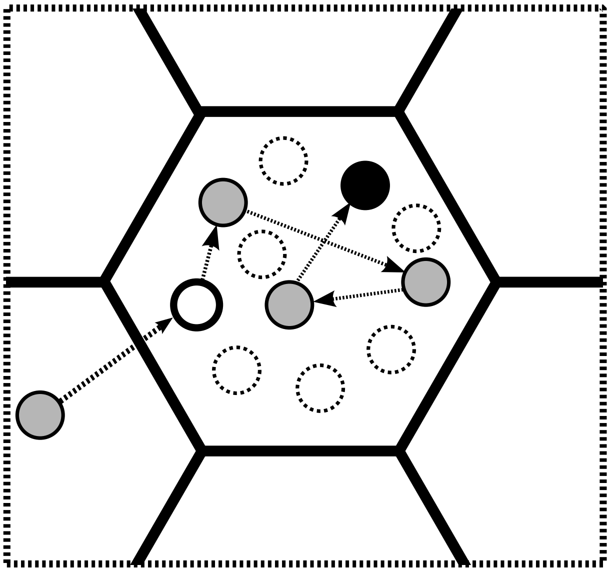

Figure 2: Delivering a packet within a cell

Figure 2 illustrates the route traversed by a packet within the destination cell, after entering this cell as shown in figure 1. The packet enters the cell when it is received by the node depicted by the white circle. The packet is then received, and retransmitted, by the relay nodes depicted by the grey circles, in the order indicated by the arrows. Finally, the packet is received by the destination node, depicted by the black circle, which consumes the packet. No further retransmissions occur. Empty circles with dotted edges represent nodes not selected for relay duty before the packet arrived at the destination node.

Note that the route traversed by the packet within the destination cell is defined only by the order in which nodes are selected for relay duty, and is independent of the relative geographic position of nodes within this cell. This is acceptable because all nodes within the cell are sufficiently close that any node pair can successfully exchange packets.

At first glance the multiple retransmissions may appear inefficient. However, the destination node is not necessarily awake when the packet reaches the destination cell, so the packet must be exchanged between active nodes. Furthermore, this applies only within the destination cell; long delivery routes span multiple cells in which only one node participates.

## 5 ITR in sensornet cells

In this section we consider delivery of packets, where the source and destination nodes reside with the same cell, using a Token Ring approach.

In a Token Ring network, the nodes are organised logically in a ring topology. A control token circulates around the ring, controlling access to the medium. Data packets are transmitted sequentially from node to node around the ring. This is a good match for nodes within a cellular sensornet in which a cyclical duty cycle is enforced, for example under the control of a protocol such as Implicit-EDF [28] or CDAP [15], with LISP synchronisation packets [38] taking the place of the dedicated token.

In the simplest form of Token Ring, each host-to-host link is unidirectional such that packets traverse the network in one direction only. This restriction may be imposed by software policy, or a consequence of hardware limitations. More complex variants of Token Ring may allow bidirectional packet exchange and traffic flows, or may have multiple tokens traversing the network at any given time, but these are beyond the scope of this paper.

### 5.1 Host interactions in a Token Ring

Consider a network composed of a set $H$ of hosts $h_1, \cdots, h_n$. Each host $h_i$ is connected to hosts $h_{i-1}$ and $h_{i+1}$, with a final connection from $h_n$ to $h_1$. A continuous closed loop is formed in which each host is connected to exactly two hosts.

A single *token* is passed from host to host in a single direction around the loop, with each host-to-host exchange following its predecessor by some delay $d$. It follows that the token traverses the network in $n$ hops in $nd$ time units, at which point the token has returned to its starting point and the process repeats.

When a given host $h_i$ possesses the token it is entitled to transmit data packets. If all hosts in $H$ are connected by a shared medium then all hosts will simultaneously receive transmissions from the current token holder $h_i$ so traffic can reach its destination in exactly one hop. If each host $h_j$ is connected only to its ring predecessor $h_{j-1}$ and successor $h_{j+1}$ then the token holder $h_i$ can transmit only to $h_{i+1}$, and this latter host must then forward the packet onward if it is not the ultimate destination.

### 5.2 Token Rings in sensornet cells

Successful and timely interaction between independent nodes requires some form of clock synchronisation protocol to operate within the network. Numerous suitable clock synchronisation protocols have been proposed in the literature [39]. The work discussed in this paper does not require any specific protocol to be selected. However, a combination of the LISP [38] and DCAP [40] protocols is well-suited to the synchronisation of periodic timing signals between the nodes within network cells, and between the cells of a larger network.

Both the packet sender $S_i$ and receiver $S_{i+1}$ nodes interact with the shared medium simultaneously during packet exchange. In a sensornet, the energy-hungry wireless communications subsystems of individual nodes are generally switched off when not strictly required. Node duty schedules must therefore be organised so as to ensure at least nodes $S_i$ and $S_{i+1}$ both have wireless communications subsystems switched on when $S_i$ may possess the token.

Any mechanism which achieves this condition is acceptable. One such mechanism is CDAP [15], if we set the CDAP parameter $m \geq 2$. Figure 3 illustrates the schedule constructed for a cell of $n = 5$ nodes, labelled *a*-*e*, where $m = 2$. A system period of length $e$ is divided into 5 equal timeslots, labelled 1-5, each of length $e/n$, centred on underlying LISP synchronisation pulses [38].

Within each system period, each node $S_i$ is assigned to the CDAP *ACTIVE* state for exactly two timeslots. During the first assigned timeslot, $S_i$ listens for packets broadcast by its ring predecessor $S_{i-1}$. If a packet not destined for $S_i$ is received from $S_{i-1}$, then $S_i$ rebroadcasts this to its ring successor $S_{i+1}$ in its second assigned timeslot. Node $S_i$ can implement any other communications duties required by the application during either assigned timeslot.

We define that node $S_i$ holds the *implicit token*, which is “transferred” from host to host with each new CDAP timeslot. Of course, no token need actually be exchanged explicitly, but if the CDAP timeslots [15] are managed by LISP synchronisation transmissions [38] then the minimal LISP synchronisation pulse packets take the role of the token.

### 5.3 ITR packet relay actions

Consider a data packet $P$ which is received by the node $S_i$ in cell $C_j$, which holds the *implicit token* as a consequence of having just entered the *ACTIVE* state in the current CDAP timeslot

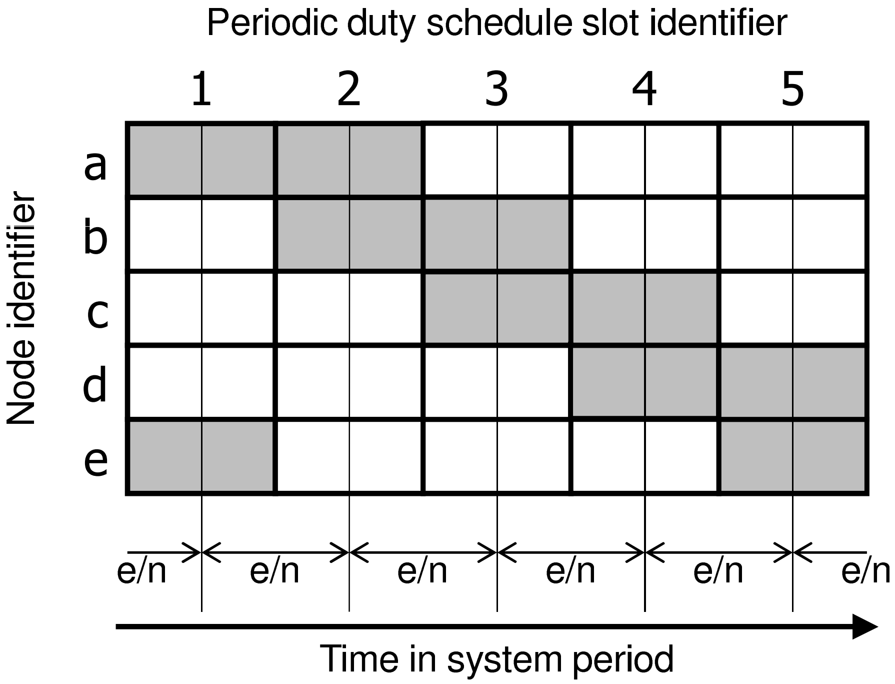

Figure 3: CDAP periodic duty schedule

as described above. Packet $P$ specifies a destination node $S_x$ which resides in cell $C_y$. Node $S_i$ can determine whether node $S_x$ also resides in cell $C_j$ such that $C_j = C_y$, or whether $S_x$ resides in some other cell $C_y$ such that $C_j \neq C_y$. This decision might be made by considering the geographical boundaries of cells, by a node identification scheme which encodes cell ID into node ID, by a locally held table of other cell members, or by some other mechanism.

Node $S_i$ must determine into which of the following three categories the ultimate specified destination $S_x$ of the packet $P$ falls:

Category 1: $S_i = S_x$

Packet $P$ has arrived at its destination $S_x = S_i$ at which it can be consumed. No forwarding is required for $P$. This does not involve ITR routing within cell $C_j$.

Category 2: $S_i \neq S_x \wedge C_j \neq C_y$

Packet $P$ has not arrived at its destination $S_x$, and $S_x$ resides in a different cell $C_j \neq C_y$. Packet $P$ should be passed to the routing mechanism used for normal intercellular packet forwarding. This will schedule $P$ for retransmission for the attention of neighbouring cells, which may or may not be the destination cell $C_y$. This does not involve ITR routing within cell $C_j$.

Cells located within adjacent cells which observe packet transmission emanating from cell $C_j$ should take no further action if the packet header specifies that the ultimate destination is also located in cell $C_j$, or the observing node $S_i$ has already seen at least one copy of packet $P$ previously. Although no incorrect behaviour would be induced by any such redundant retransmission, this would waste energy and occupy the shared wireless medium unnecessarily.

Category 3: $S_i \neq S_x \wedge C_j = C_y$

Packet $P$ has not arrived at its destination $S_x$, but $S_x$ resides in the current cell $C_j = C_y$. Packet $P$ should be forwarded within this cell until it reaches its destination $S_x$. This involves ITR routing within cell $C_j$ as described below.

If $P$ specifies more than one destination, for example as a consequence of a fuzzy routing algorithm or a multicast policy, then the decision process associated with more than one of the categories specified above may be implemented as appropriate to the scope of the destination definition.

## 5.4 Implementing ITR in sensornets

Correct functioning of the packet forwarding mechanism described herein is dependent on the validity of the following preconditions. Both of these hold under the circumstances outlined in section 5.1:

1. $S_i \neq S_x \wedge C_j = C_y$ - packet $P$ has arrived at some node in the cell containing the destination, but has not yet reached the destination itself.
2. $m \geq 2$ under CDAP - at least two nodes are simultaneously in the $ACTIVE$ state at any given time, and this overlapping shared state passes between cells over time, such that data packets can be exchanged between nodes in the cell. It is always eventually true that the destination node will be $ACTIVE$ at some time after some other node in the cell has received the packet while in the $ACTIVE$ state, with an unbroken sequence of mutually active node pairs leading to the destination node's $ACTIVE$ period.

The ITR routing mechanism is very simple. Packet $P$ is held at some node $S_i$ in the destination cell $C_j$, having been received in the first CDAP duty period after entering the $ACTIVE$ state. If $S_i = S_x$, the packet has reached the destination and the algorithm terminates. Otherwise, node $S_i$ waits until the current CDAP duty period completes. At this point, another node $S_{i+1}$ enters the $ACTIVE$ state; node $S_i$ rebroadcasts packet $P$, which is heard by node $S_{i+1}$ which is now listening to the shared wireless medium.

At this point, the involvement of node $S_i$ in the delivery of packet $P$ ends. Node $S_{i+1}$ now takes the role outlined for node $S_i$ above, and so on with nodes $S_{i+2}$, $S_{i+3}$, and so on until packet $P$ reaches node $S_{n-1}$. At this point, the packet has been processed by all nodes in the destination cell $C_j$, and must therefore have reached its destination, so the algorithm necessarily terminates at this point. It is possible that the packet need not touch all $n$ nodes in cell $C_j$ prior to reaching the destination node $S_x$, in which case the algorithm terminates earlier.

# 6 Algorithm

**Algorithm 1** : ITR packet relay

```
Require: Packet $P$ resides at node $S_i$ inside cell $C_j$
Require: Geographic location of $S_i$ is $G_i$
Require: Packet $P$ was transmitted by node $S_t$
Require: Packet $P$ specifies location $G_t$ of node $S_t$
Require: Packet $P$ specifies destination node $S_x$
Require: Packet $P$ specifies location $G_x$ of node $S_x$
Require: Packet $P$ specifies TTL in destination cell, $c$
Require: Maximum TTL in destination cell is $r_{max}$
if $S_i = S_x$ then
  {ITR relay destination category 1}
  Consume $P$ at node $S_i = S_x$
else
  if $G_x$ is outside $C_j$ boundary then
    {ITR relay destination category 2}
    Enqueue $P$ for intercellular retransmission to $G_x$
  else
    {ITR relay destination category 3}
    if node $S_t$ resides outside cell $C_j$ then
      Initialise $P$ intracellular TTL, $c' \Leftarrow r_{max}$
      Enqueue $P$ for retransmission after delay $e$
    else
      Decrement $P$ intracellular TTL, $c' \Leftarrow c - 1$
      if $P$ intracellular TTL $c > 0$ then
        Enqueue $P$ for retransmission after delay $e$
      else
        Drop $P$ at non-destination node $S_i$
      end if
    end if
  end if
end if
```

Algorithm 1 defines the ITR mechanism employed by a node $S_i$ upon receiving a packet, $P$, eventually destined for node $S_x$. The packet originates at source node $S_s$, and was most recently transmitted by node $S_t$; it may or may not be the case that $S_s = S_t$. Upon receipt at $S_i$, packet $P$ may be forwarded to another cell, forwarded within the current cell, consumed if $S_i = S_x$, or dropped if no further intracellular forwarding is permitted. This is determined by the packet TTL, which is specified in discrete node-to-node hops.

We assume each node has a globally unique identifier, known to the packet source and destination, but may not be known by any potential intermediary nodes. We also assume that node $S_i$ at geographic location $G_i$ can identify whether nodes $S_t$ and $S_x$ reside in different cells. This is trivially addressed by labelling each packet $P$ with the geographic location $G_x$ of the destination node $S_x$, and the geographic location $G_t$ of the most recent relay node $S_t$, as packet recipient node $S_i$ can easily determine whether these locations fall within its own cell, $C_j$.

# 7 Costs and overheads

We now consider the costs and overheads associated with operating the packet forwarding algorithm discussed in this paper.

## 7.1 Algorithmic costs

The intracellular packet forwarding mechanism outlined above is simple, and therefore exhibits low computational complexity. Each of the $n$ nodes in the destination cell need only decide if the packet $P$ has reached its destination or must be forwarded to the next node in the ring. This is independent of cell and network population, and is performed at most once at each of the $n$ nodes for any given packet $P$.

The algorithm is O(1) in cell population as executed at each node, and is O($n$) in node count for the entire cell. No state information is held at any node, other than the O(1) cost of holding a packet which awaits forwarding to the next node in the ring.

The packet forwarding mechanism is optimised for distributed applications in which most packets are routed between cells rather than between specific uniquely identified nodes. This is a reasonable assumption for most sensornet applications, which are inherently data-centric and geography-centric designs. Where many packets are sent to permanently active hosts, for example base stations acting as data sinks for the entire sensornet or gateways to conventional wired networks, ITR implies no additional intracellular rebroadcast penalty as such hosts are always directly reachable within their cell.

In common with other Token Ring mechanisms [17], costs grow in proportion to the number of nodes in the ring; large rings may yield undesirable costs and latencies. Under ITR, the size of the ring within a given cell is equal to the number of nodes actively participating in that cell. Cellular sensornets are generally designed so as to restrict active cell membership [9], perhaps by implementing mechanisms such as ADCP [31]. Cell populations of around $10-20$ are common in sensornets [36] as this is generally an energy-efficient cluster size for sensornets containing hundreds to thousands of nodes [41], and is within the *ad hoc horizon* limit of $10-20$ nodes collaborating independently without hierarchical or external control [42].

Note that the number of uniquely identified sources is irrelevant; it is uniquely identified destinations for which ITR may imply multiple rebroadcasts within the destination cell. If a packet must be rebroadcast several times within a destination cell, this implies an energy cost which may be significant but is nevertheless bounded by the intracellular cost of ITR and the cost of the selected network-wide intercellular routing protocol. If the number of packet deliveries between uniquely identified source-destination pairs is relatively small, then the cost impact of permitting these is commensurately small.

## 7.2 Energy costs

We assume each of the $n$ nodes consume energy at the same rate, where the actual value of that rate depends on the current state of a given node. We also assume the additional energy cost of transmitting some packet $P$ is equal for any node in the destination cell, though of course this may vary between different packets of different length.

As nodes will incur a given energy cost for each period regardless of network traffic, and each rebroadcast of a given packet $P$ implies a fixed energy cost, we need consider only the number of rebroadcasts required for a given packet when analysing the energy efficiency of the ITR packet forwarding policy within the packet destination cell.

Assume that some packet $P$ enters the destination cell $C_y$ at some node $S_i$. In the best case $S_i = S_x$ such that 0 rebroadcasts are required. In the worst case the packet must be rebroadcast $n-1$ times before it has touched all nodes in the cell. It follows that the number of rebroadcasts, $r$, which is proportional to the total energy consumed by this intracellular packet forwarding, is given by the interval $[0, n-1]$. The delivery latency experienced within the cell is given by $^{re}/_{n}$, as a delay equal to the length of one duty period $e$ is observed between each rebroadcast (see section 5.2).

If we assume that the destination node $S_x$ is equally likely to be any of the $n$ nodes located in the destination cell $C_y$, the probability that a given packet $P$ has reached the destination node is proportional to the number of rebroadcasts since entering the cell. Packet $P$ will be successfully forwarded to its destination node $S_x$ within the destination cell $C_y$ with probability $p$ if $r = \lceil p\ (n-1) \rceil$.

We can take advantage of this to artificially bound $r$ to some $r_{max} \in [0, n-1]$, perhaps to prevent a single packet delivery attempt consuming unacceptable energy or incurring unacceptable delay. This exchanges the deterministic delivery guarantees of unbounded ITR with a probabilistic delivery guarantee, but reduces the maximum energy cost and delivery latency associated with a given packet. The network designer must decide whether the trade-off implied by selecting $r_{max} < (n-1)$ is acceptable in a given application.

## 7.3 Worst-case latency

Now consider a real-time sensornet application in which packets must be delivered from source to destination within a defined deadline, $\tau$. Consider a network composed of $c$ cells, each of which contains $n$ nodes. In the worst case, a packet originating in cell $C_j \neq C_y$ will be forwarded through every other cell in the network before reaching the destination cell $C_y$, traversing $c-1$ intercellular hops. Upon arriving at node $S_i \neq S_x$ in cell $C_y$, in the worst case the packet will be forwarded through every other node in the cell before reaching the destination node $S_x$, traversing $n-1$ intracellular hops. Assuming each hop requires time $d$, the worst-case delivery time $t_w$ is given by equation 1. Provided that $t_w \leq \tau$, the packet is guaranteed to be delivered within the deadline.

$$\begin{aligned} t_w &= ((c-1)+(n-1))e \\ &= (c+n-2)e \end{aligned} \qquad (1)$$

However, the actual packet delivery latency is likely to be considerably less. The number of cell-to-cell hops, $c-1$, is unlikely to exceed the cell network diameter, $d$ [17]. This is the number of cells crossed by the longest straight line within the deployment region, which is dependent on the spatial configuration of the network nodes. IGF generates approximately straight delivery routes [36], so the number of cells involved in any delivery attempt is unlikely to be significantly greater than the minimum implied by a perfectly straight route. The number of intracellular hops traversed by packets is evenly distributed in $[0, n-1]$, so the expected value is $\frac{n-1}{2}$.

Now consider the worst case packet delivery latency for plain IGF, without the ITR element. Here we find that the worst case delivery latency $t_x = (c-1)e$, as there is no intracellular component. Clearly $t_x \leq t_w$, but it is no longer guaranteed that the final intercellular hop will deliver the packet to the uniquely identified destination; for a cell of $n$ nodes, the probability of successful delivery is exactly $1/n$, assuming no packet loss from other causes.

The probability of the network delivering the packet to the destination cell, but then failing to deliver it to the destination node, is given by $p_l = 1 - 1/n$. All energy and other resources invested by the network in bringing the packet to the destination cell will be wasted. For cells composed of numerous nodes, $p_l \to 1$ as $n \to \infty$, which is obviously undesirable. It follows that a mechanism which prevents this wasteful behaviour, such as the ITR mechanism discussed in this paper, becomes increasingly beneficial as the network size increases.

Now consider the general worst case packet delivery time, $t_y$, given by equation 2. Here we assume that the packet traverses $n-1$ intracellular hops within each cell, including the cells containing the source and destination nodes. It is always true that $t_w \leq t_y$, and often $t_w \ll t_y$.

$$\begin{aligned} t_y &= (c(n-1))e \\ &= (cn-c)e \end{aligned} \tag{2}$$

### 7.4 Network loading

In our analysis we assume that the network loading is such that within each system period of length $e$ it is always possible for a pair of cells to exchange a given packet. Notably, the packet forwarding algorithm discussed in this paper does not require the transmission or exchange of network management or coordination packets, and hence does not itself contribute toward network loading beyond that implied by the forwarding of application data packets.

If the network loading is particularly high, such that the total transmission time required for all packets to be exchanged within cells exceeds $e$, this assumption may not hold. A node which is unable to forward a given packet within some system period may either drop the packet, or retain the packet for further transmission attempts when the node is next assigned to active duty.

We do not consider this issue further in this paper, as the selection of an appropriate strategy is highly dependent on the specific characteristics of the sensornet application and its sensitivity to deadlines misses. However, sensornets typically do not exchange large volumes of data for extended periods, as this quickly exhausts scarce energy resources [3]; in-network data processing and volume reduction is generally favoured [43]. It follows that the influence of this potential issue is less acute in sensornets than in general wireless networking applications.

## 8 Evaluation

We now evaluate the ITR mechanism by simulation experiment, to confirm that protocol behaviour matches the theoretical expectations.

### 8.1 Experimental configuration

The modelled network is a sensornet in which 1000 nodes are distributed with uniform spatial density within a planar deployment region. This region is divided into 100 hexagonal cells of equal size, arranged into a 10×10 hexagonal grid. A distributed sensing and processing application is simulated, in which the start and end points of each traffic flow are selected with uniform random probability within the deployment region.

We model *data packets*, which are routed between geographic endpoints independent of node identity, and *non-data packets*, which are routed between uniquely identified nodes. It follows that *data packets* can be delivered to any node within the cell which encloses the geographic destination, whereas *non-data packets* must be delivered to the destination cell and thence onward to the specific node.

A set $P$ of $x$ packets, labelled $p_1 \cdots p_x$, is routed through the cellular sensornet. To avoid confounding effects arising from network congestion, each delivery attempt was allowed to complete, either successfully or unsuccessfully, before the next began. The plain IGF protocol [36] is used to relay packets between cells, and ITR is used to relay packets within the destination cell. We label the combination of protocols as ITR-IGF.

## 8.2 Traffic distribution analysis

For each of the $x$ packets $p_1 \cdots p_x$ which are successfully delivered to an appropriate destination, there is an associated cost. In sensornets the most significant cost is generally the energy cost of packet broadcast into a shared wireless medium. If a given packet is rebroadcast with similar power at each relay node, then the energy cost is approximately proportional to the number of broadcasts. It follows that we would like to minimise the number of broadcasts, so as to minimise the cost to the sensornet.

In this section we measure the number of *intercellular hops*, $q$, and the number of *intracellular hops*, $r$, traversed by each packet $p_1 \cdots p_x$ where $x = 1 \times 10^6$. We evaluate the distribution of $q$, $r$, and $s = q + r$, for each packet. $r$ gives the total delivery cost for the packet.

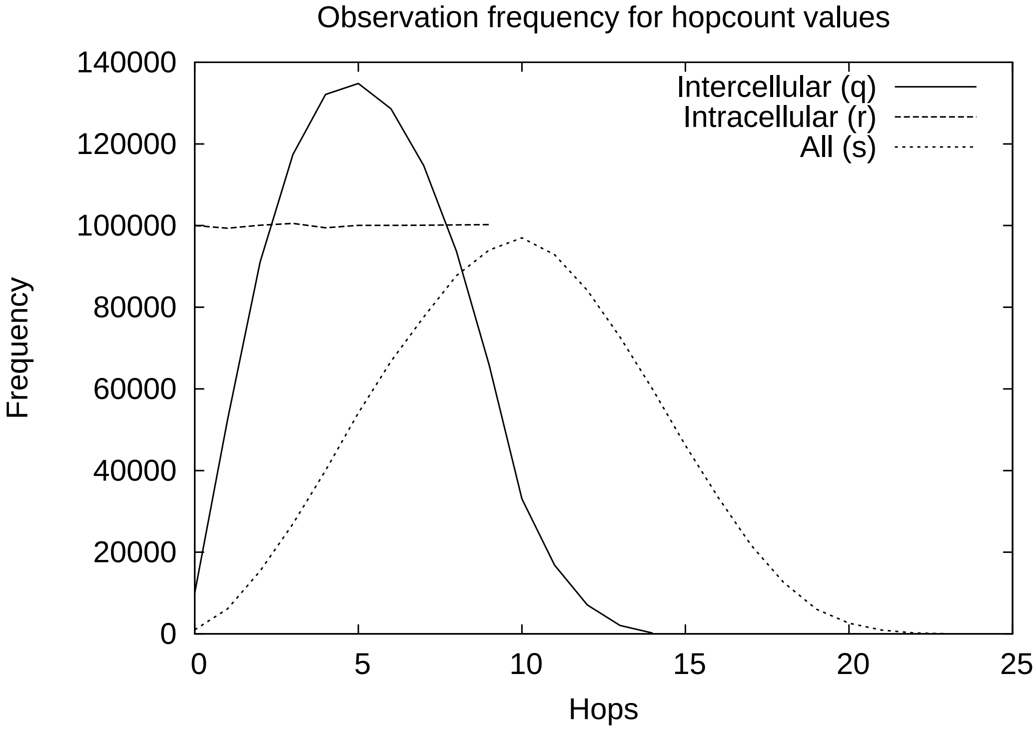

Figure 4: Hopcount distributions

Figure 4 shows the distribution of $q$, $r$, and $s$ for *non-data* packets with uniquely identified source and destination node. Similar distributions of $q$, $r$, and $s$, can be found for *data* packets with non-uniquely identified destination nodes by setting $r = 0$ such that $s = q$, as no intracellular packet relay is required. It follows that the $q$ trace given in figure 4 for *non-data* traffic is equivalent to the $s$ trace for *data* traffic, so the characteristics of *data* traffic can be understood through this interpretation.

In figure 4 we see that the distribution of $q$ observes a reasonable approximation of a Gaussian distribution. The mean is 5.34 and the standard deviation is 2.61, to 2 decimal places. This is as expected as the distribution of distances between pairs of randomly selected points within a planar region observes an approximately Gaussian distribution, and the length of each node-to-node hop is expected to be approximately equal for a uniform spatial distribution of nodes and cells.

The distribution of $r$ is approximately uniform, with mean of 4.50 and standard deviation of 2.87 to 2 decimal places. This is as expected, as the number of node-to-node hops required within the destination cell is independent of geographic position, and is distributed randomly in the interval $[0, n-1]$. This is dependent on the size of cell populations, and is independent of total network size.

We see that $s$ also observes a reasonable approximation of a Gaussian distribution. The mean is 9.84 and the standard deviation is 3.88, to 2 decimal places. The overall characteristics of the $s$ distribution are very similar to those of $q$, as $r$ is bounded by $n-1$ and hence exerts limited influence. Because end-to-end delivery routes incorporating both inter- and intra-cellular sections tend to be longer than either section considered in isolation, and because the distribution of $r$ is approximately uniform, the distribution of $s$ has both greater mean and standard deviation than that of $q$.

Figure 5 shows the distribution of end-to-end route lengths to illustrate the benefit offered by ITR in a cellular network, as compared to an alternative geography-aware protocol which is unaware of any cellular logical structure. Delivery under ITR-IGF in a cellular network, as described in this paper, is compared to delivery under plain IGF [36] through the same underlying set of network nodes, and for the same set of *non-data* packets with uniquely-specified destinations, $P$. For the plain IGF instance we assume that all nodes are awake as candidates for packet relay duty at all times, to ensure that the destination node is always reachable.

The distributions of route lengths for ITR-IGF and plain IGF are similar, in that they both observe an approximately Gaussian distribution, which is as expected from the spatial distribution of randomly selected pairs of source and destination nodes. The distribution for ITR-IGF, with mean of 17.20 and standard deviation of 8.31 to 2 decimal places, however, has a narrower distribution which is centred on a lower mean. This is desirable, and indicates that ITR-IGF tends to select shorter routes in this sensornet configuration.

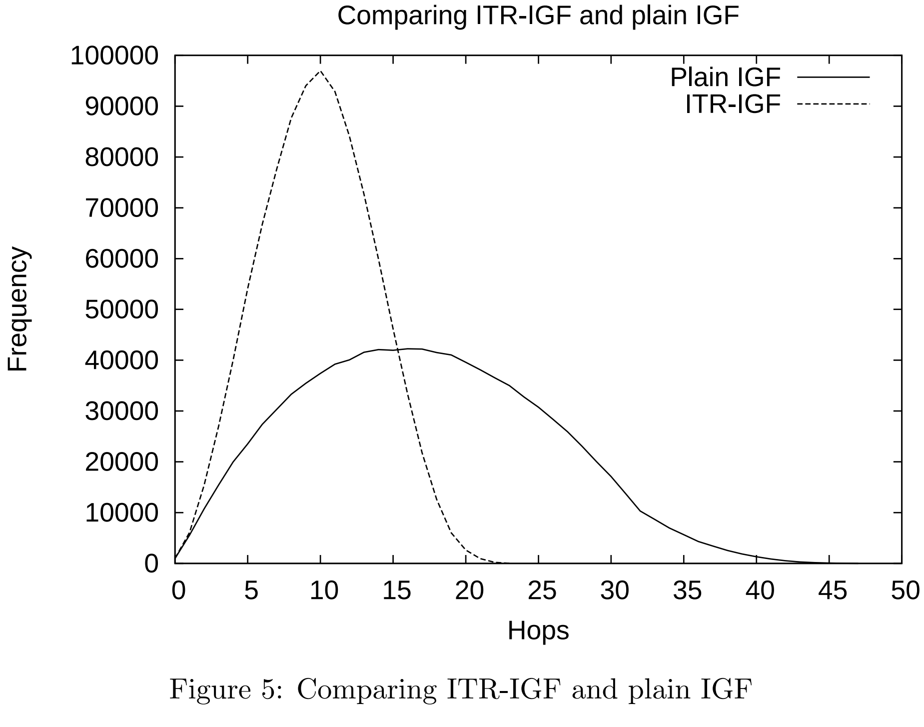

Figure 5: Comparing ITR-IGF and plain IGF

This is primarily a consequence of the cellular structure, and allocating the packet relay role to exactly one node per cell at all times. Under plain IGF any active node is a relay candidate. The selection decision is based on angle rather than distance [36], it is common for a packet to be relayed by multiple nodes within a single cell, which is not globally optimal. The cellular approach employed by ITR-IGF is also not guaranteed to be globally optimal, but allows large spatial distances to be traversed with relatively few logical node-to-node hops.

An alternative approach would be a hierarchical variant of IGF. Packets would be delivered to the destination cell as under ITR-IGF, and then a local instance of IGF would deliver the packet within the destination cell. This would generally require fewer intracellular node-to-node hops than ITR. However, this would impose the additional constraint that the destination node be active at all times. Whereas this may be acceptable for highly-resourced base station destinations, it is not acceptable for general sensor mote destinations which must often enter low-power inactive states to provide acceptable lifetime [32].

### 8.3 Impact of intracellular TTL bounds

As described in sections 6 and 7, it is possible to configure the ITR mechanism to provide upper bounds on intracellular delivery costs by specifying a finite value of intracellular packet TTL, $r_{max}$, within the destination cell. This allows the network designer to manage the trade-off between the probability that a given packet is successfully delivered, and the worst case delivery cost and worst case latency for that packet. In energy-constrained networks with cells containing many nodes, or networks in which the freshness of delivered data is of prime importance, this reduces the likelihood of stale packets consuming disproportionate resources.

In this section we consider the delivery of *non-data packets*, for which the start and end points are uniquely specified nodes, each of which is selected at random from the set of all nodes. All packets reach the destination cell, but successful delivery to the specific destination node is not guaranteed. Of the $x$ packets $p_1 \cdots p_x$, the total number which arrive at the intended destination is given by $y$, and the total number dropped within the destination cell owing to bounded TTL is given by $z$.

Figure 6 illustrates the relationship between the proportion of packets successfully delivered to the specified destination, $d = {}^{y}\!/_{y+z}$, and the maximum permitted packet TTL within the destination cell, $r_{max}$. The experiment is repeated for different values of $r_{max}$ in the interval $[0, n-1]$, where the number of nodes within each cell is $n = 10$. In each experiment the set of packets, $P$, is identical, where $x = 1 \times 10^6$.

It can be seen in figure 6 that there exists a near-linear relationship between $r_{max}$ and $d$. This is as expected, as the uniquely specified packet destination nodes are selected at random, so the required number of node-to-node hops within the destination cell, $h$, is evenly distributed in the interval $[0, n-1]$. Given $r_{max} \in [0, n-1]$, the probability that $h > r_{max}$ is equal to ${}^{1}\!/_{n-r_{max}}$. This is the probability that the packet is dropped within the destination cell prior to reaching the destination node owing to TTL.

## 9 Conclusions

Network behaviour can be defined in terms of cells, rather than individual nodes, where multiple nodes within a cell are equally able to respond to a request such as obtaining a sensor reading. This reduces logical network size, making it easier to employ protocols for which performance degrades with increasing network size. However, it may sometimes be necessary to route packets between a specific pair of nodes with globally unique identifiers, for example when issuing commands to manage the network or to activate uniquely-located sensors.

This problem can be decomposed into two subproblems. Firstly, the packet is delivered to the cell containing the destination node. Secondly, the packet is delivered within the cell to the destination node itself. The *Implicit Token Ring* protocol solves this second subproblem

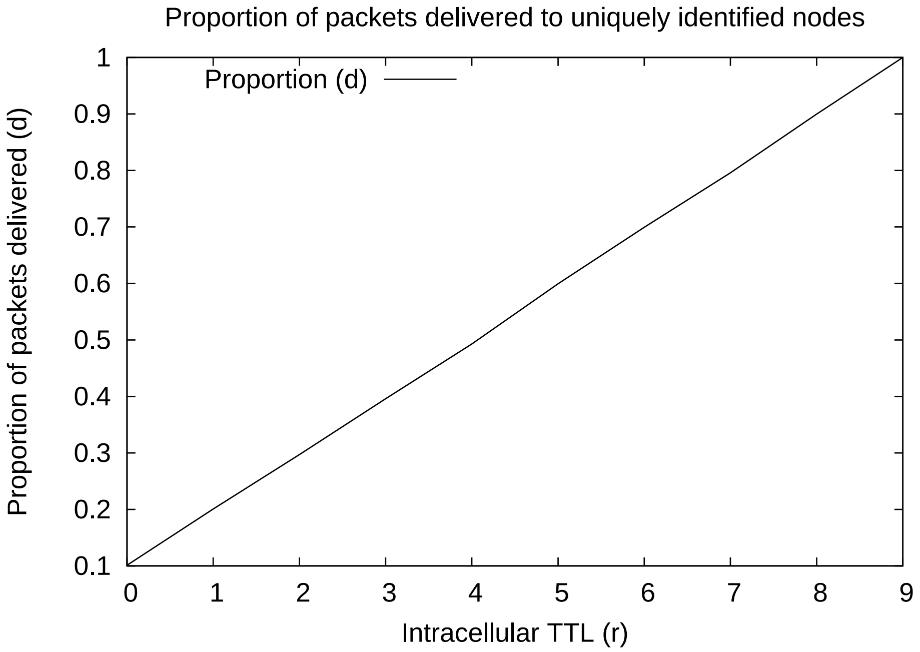

Figure 6: $d$ versus $r_{max}$

with bounded costs and delays, without compromising the efficiency or correctness of protocols which address the first subproblem.

It was shown that packets which specify a uniquely identified destination can be delivered successfully, without increasing the delivery cost of the majority of typical sensornet traffic packets which do not specify a uniquely identified destination. This is a significant advantage over other protocols which do not differentiate between these traffic categories, and hence either do not support delivery to uniquely identified destination nodes, or incur an unnecessarily high delivery cost where this is not required.

## Acknowledgements

This work is partially supported by an EPSRC grant, and BAE Systems plc by the grant "Hierarchical System Management for Integrated Modular Systems".

## References

[1] D. Estrin, L. Girod, G. Pottie, and M. Srivastava, "Instrumenting the world with wireless sensor networks," in *Proc. 11th IEEE International Conference on Acoustics, Speech, and Signal Processing*, 7-11 May 2001, pp. 2033–2036.

[2] I. Akyildiz, W. Su, and Y. Sankarasubramaniam, "Wireless sensor networks: a survey," *Computer Networks*, vol. 38, no. 4, pp. 393–422, 2002.

[3] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar, "Next century challenges: Scalable coordination in sensor networks," in *Proc. 5th International Conference on Mobile Computing and Networking*, 17-19 August 1999, pp. 263–270.

[4] L. Ni, Y. Zhu, J. Ma, Q. Luo, Y. Liu, S. Cheung, Q. Yang, M. Li, and M. Wu, "Semantic Sensor Net: an extensible framework," *International Journal of Ad Hoc and Ubiquitous Computing*, vol. 4, no. 3/4, pp. 157–167, 2009.

[5] J. Chen, E. Shen, and Y. Sun, "The deployment algorithms in wireless sensor net works: A survey," *Information Technology Journal*, vol. 8, no. 3, pp. 293–301, 2009.

[6] D. Lacks, M. Chatterjee, and T. Kocak, "Design and evaluation of a distributed clustering algorithm for mobile ad hoc networks," *The Computer Journal*, vol. 52, no. 6, pp. 656–670, 2009.

[7] C. Wen and W. Sethares, "Automatic decentralized clustering for wireless sensor networks," *EURASIP Journal on Wireless Communications and Networking*, vol. 2005, no. 5, pp. 686–697, 2005.

[8] N. Vlajic and D. Xia, "Wireless sensor networks: to cluster or not to cluster?" in *Proc. 7th International Symposium on a World of Wireless, Mobile and Multimedia Networks.* Buffalo-Niagara Falls, NY: IEEE Computer Society, New York, NY, 26-29 June 2006, pp. 268–277.

[9] G. Yu, "Cell-based coverage management and routing protocol for wireless sensor networks," in *Proc. 3rd International Conference on Mobile Technology, Applications, and Systems*, September 2008, pp. 1–7.

[10] M. Handy, M. Haase, and D. Timmermann, "Low energy adaptive clustering hierarchy with deterministic cluster-heads selection," in *Proc. 4th International Workshop on Mobile and Wireless Communications Network.* Stockholm: IEEE Communications Society, Piscataway, NJ, 9-11 September 2002, pp. 368–372.

[11] K. Shin, A. Abraham, and S. Han, "Self organizing sensor networks using intelligent clustering," in *Proc. 6th International Conference on Computational Science and its Applications.* Glasgow: Springer, Berlin, 8-11 May 2006, pp. 40–49.

[12] S. Al-Kuwari and S. Wolthusen, "A survey of forensic localization and tracking mechanisms in short-range and cellular networks," in *Proc. 1st International Conference on Digital Forensics and Cyber Crime*, 30 September - 2 October 2009, pp. 19–32.

[13] M. Rudafshani and S. Datta, "Localization in wireless sensor networks," in *Proc. 6th international conference on Information Processing in Sensor Networks*, 25-27 April 2007, pp. 51–60.

[14] A. Bonivento, C. Fischione, L. Necchi, F. Pianegiani, and A. Sangiovanni-Vincentelli, "System level design for clustered wireless sensor networks," *IEEE Transactions on Industrial Informatics*, vol. 3, no. 3, pp. 202–214, August 2007.

[15] J. Tate and I. Bate, "Energy efficient duty allocation protocols for wireless sensor networks," in *Proc. 14th IEEE Conference on Engineering of Complex Computer Systems*, 2-4 June 2009, pp. 58–67.

[16] H. Zimmerman, "OSI reference model - the ISO model of architecture for open systems interconnection," *IEEE Transactions on Communications*, vol. 28, no. 4, pp. 425–432, 1980.

[17] A. Tanenbaum, *Computer Networks*, 3rd ed. Harlow: Pearson, March 1996.

[18] IEEE 802.3 Working Group, *IEEE Std. 802.3-2008.* Piscataway, NJ: IEEE Press, 2008.

[19] N. Abramson, "The ALOHA system: another alternative for computer communications," in *Proc. AFIPS (Fall)*, 17-19 November 1970, pp. 281–285.

[20] M. Willett, "The IBM Token-Ring Network - a functional perspective," in *Proc. IEEE/ACM Fall Joint Computer Conference*, November 1986, pp. 489–497.

[21] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker, "Complex behavior at scale: An experimental study of low-power wireless sensor networks," UCLA, Los Angeles, CA, Tech. Rep. CSD-TR 02-0013, February 2002.

[22] R. M. Needham, "Systems aspects of The Cambridge Ring," in *Proc. 7th ACM Symposium on Operating Systems Principles*, December 1979, pp. 82–85.

[23] IEEE 802.4 Working Group, *IEEE Std. 802.4-1990.* Piscataway, NJ: IEEE Press, 1990.

[24] IEEE 802.5 Working Group, *IEEE Std. 802.5-1998.* Piscataway, NJ: IEEE Press, 1998.

[25] R. Jain, *FDDI Handbook: High-Speed Networking Using Fiber and Other Media.* Boston, MA: Addison-Wesley Longman, 1994.

[26] M. Ergen, D. Lee, R. Sengupta, and P. Varaiya, "WTRP - wireless token ring protocol," *IEEE Transactions on Vehicular Technology*, vol. 53, no. 6, pp. 1863–1881, November 2004.

[27] IEEE 802.11 Working Group, *IEEE Std. 802.11-1999.* Piscataway, NJ: IEEE Press, 1999.

[28] M. Caccamo and L. Zhang, "The capacity of implicit EDF in wireless sensor networks," in *Proc. 15th Euromicro Conference on Real-Time Systems*, 3-5 September 2003, pp. 267–265.

[29] V. Kawadia and P. Kumar, "Power control and clustering in ad hoc networks," in *Proc. 22nd IEEE Conference on Computer Communications*, 1-3 April 2003, pp. 459–469.

[30] A. Woo, T. Tong, and D. Culler, "Taming the underlying challenges of reliable multihop routing in sensor networks," in *Proc. 1st international conference on Embedded Networked Sensor Systems.* Los Angeles, CA: ACM Press, New York, NY, 5-7 November 2003, pp. 14–27.

[31] J. Tate and I. Bate, "Maintaining stable node populations in long-lifetime sensornets," in *Proc. 15th IEEE Conference on Engineering of Complex Computer Systems*, 22-26 March 2010, pp. 159–168.

[32] P. Dutta and D. Culler, "System software techniques for low-power operation in wireless sensor networks," in *Proc. 16th IEEE/ACM International Conference on Computer-Aided Design*, 6-10 November 2005, pp. 925–932.

[33] V. Paruchuri, S. Basavaraju, A. Durresi, R. Kannan, and S. Iyengar, "Random asynchronous wakeup protocol for sensor networks," in *Proc. 1st International Conference on Broadband Networks*, 25-29 October 2004, pp. 710–717.

[34] S. Ni, Y. Tseng, Y. Chen, and J. Sheu, "The broadcast storm problem in a mobile ad hoc network," in *Proc. 5th ACM/IEEE International Conference on Mobile Computing and Networking.* Seattle, WA: ACM Press, New York, NY, 15-19 August 1999, pp. 151–162.

[35] Y. Gao, K. Wu, and F. Li, "Analysis on the redundancy of wireless sensor networks," in *Proc. 2nd International Workshop on Wireless Sensor Networks and Applications*, 19 September 2003, pp. 108–114.

[36] T. He, B. Blum, Q. Cao, J. Stankovic, S. Son, and T. Abdelzaher, "Robust and timely communication over highly dynamic sensor networks," *Real-Time Systems*, vol. 37, no. 3, pp. 261–289, 2007.

[37] H. Karl and A. Willig, *Protocols And Architectures For Wireless Sensor Networks.* Wiley, 2005.

[38] J. Tate and I. Bate, "An improved lightweight synchronisation primitive for sensornets," in *Proc. 6th IEEE International Conference on Mobile Ad-hoc and Sensor Systems*, October 2009, pp. 448–457.

[39] B. Sundararaman, U. Buy, and A. Kshemkalyani, "Clock synchronization for wireless sensor networks: a survey," *Ad Hoc Networks*, vol. 3, no. 3, pp. 281–323, 2005.

[40] J. Tate and I. Bate, "Synchronising timing signals in cellular sensornets using a hybrid algorithm," University of York, York, Tech. Rep. YCS-2009-444, Nov 2009.

[41] D. Wang, "An energy-efficient clusterhead assignment scheme for hierarchical wireless sensor networks," *International Journal of Wireless Information Networks*, vol. 15, no. 2, pp. 61–71, June 2008.

[42] P. Gunningberg, P. Lundgren, H. Nordstroem, and C. Tschudin, "Lessons from experimental MANET research," *Ad Hoc Networking Journal*, vol. 3, no. 2, pp. 221–233, 2005.

[43] C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva, "Directed diffusion for wireless sensor networking," *IEEE Transactions on Networking*, vol. 11, no. 1, pp. 2–16, 2003.