The Design and Implementation of a Constrained WSN for Permaculture Farming in Egypt

Ahmed Ismail
Electronics Engineering Department
American University in Cairo
Cairo, Egypt
a.m.ismail@aucegypt.edu

Karim Seddik
Electronics Engineering Department
American University in Cairo
Cairo, Egypt
kseddik@aucegypt.edu

***Abstract*—Since the inception of the concept of Wireless Sensor Networks (WSNs), their applicability within the context of environmental monitoring systems (EMS) has constantly been explored. Egypt stands to gain much from WSN-based EMS systems if they are properly applied within its considerably large agricultural industry. A system is developed and tested using locally available hardware within the technical, economic and social constraints of modern day Egypt. We show that the constraints in question mostly impact the routing layer of the WSN necessitating the development of a modified implementation of the LEACH routing protocol. The resulting system is thoroughly evaluated and our recommendations for its deployment are presented.**

I. Introduction

As of recent years, WSNs have oft been used to assist with the many applications of EMSs. To name but a few examples, they have been used to monitor water levels and quality, air quality, animal habitats and migrations, and agricultural parameters [1], [2], [3], [4], [5]. Out of these applications, Egypt, as a nation, stands to gain the most from a WSN-based agricultural EMS. This is mainly due to two reasons. First of all, Egypt has a sizeable agricultural industry which constitutes 17% of its GDP, employs 30% of its labour force, and takes up 81% of its water supply [6]. Second of all, WSNs are capable of fulfilling agriculturists' need to constantly and accurately observe climatic and physical environmental parameters in a convenient manner. Therefore, utilizing the technology of WSNs in the field of agriculture could potentially allow for qualitative and quantitative improvements with far-reaching effects.

Furthermore, permaculture appears to be the most suitable type of agriculture for WSN deployments. This is since, in contrast to its contemporaries, permaculture is only successful if productivity is achieved through sustainable practices by way of mindful observation. This is clear from permaculture's definition as "the conscious design of [...] infrastructure to emulate the diversity, stability, and resilience of natural ecosystems, while providing [...] food, energy, shelter, and other needs in a sustainable fashion. True permaculture stems from protracted and thoughtful observation. Its goals are water conservation, local food production, and regional self-reliance, constituting a sustainable earth-care system" [7].

Hence, a capable WSN for this type of farming may assist in the achievement of better permaculture designs whose successes may allow for the widespread adoption of permaculture practices. This may, in turn, allow for a more sustainable Egyptian agricultural industry and all around better living standards for Egyptians.

In order to develop such a robust and effective WSN, the lessons learnt from previous WSN deployments that had related goals or used similar equipment should not be ignored. In the case of COMMONSense Net, a system was designed and deployed with direct feedback from the Indian farmers of Karnataka to provide them with agriculture-relevant environmental data. Even though the system provided data pertaining to parameters that the farmers themselves requested, the system observed low levels of usage. This failure was attributed to the moderate computer illiteracy of farmers [1].

In contrast to the COMMONSense Net project, the PODs project, a large-scale EMS system developed at the University of Hawai'i at Mnoa, was designed for use by researchers and scientists. This system did not record low usage. It did however note that utilizing a two-level hierarchy, scalable routing protocols and smaller processors were required of EMS WSNs in order to avoid the costly death of WSN nodes [2].

On the other hand, in [3] the College of the Atlantic's Great Duck Island deployment of 150 Mica2Dot motes had nodes failing mainly due to harsh weather conditions and drained power supplies. The issue of exhausted supplies was found to be caused by overhearing and low power reception. During this deployment the cost of risking having a single sink node and, hence, introducing a single point of failure to the system, was shown to be disastrous. This is as when the server did fail, it caused the loss of several weeks' worth of data.

Other problems faced in [4] and [5] showed that high message loss rates could result from a number of other reasons as well. Packet losses due to node failures were due to inadequate node housing, bad contacts and software bugs. Network congestion and the variable quality of links were also found to cause packet losses. This was noted to be especially true since transmission paths inherit and compound the problems of its links. This in turn allowed for path oscillations and asymmetric paths, incurring even higher message loss rates and large message delivery delays. Further factors that increased loss rates were clock drifts, relaxed scheduling and clock resynchronization failures.

To summarize, the common themes in the aforementioned deployments may be condensed into three main points. Firstly, all the systems above developed their own implementations of the routing layer tailored to fit the hardware deployed, the task to be performed and the targeted environment. Second of all, aside from issues such as node failures due to harsh weather conditions and battery depletion, most of the lessons presented above appear to be implementation or design issues pertaining to the strength of said implemented routing protocols. Lastly, even though the systems did meet said issues, all the systems described did however achieve measurable successes in terms of performing the assigned tasks, attesting to their approach of utilizing tailored hardware and protocols.

In the same manner, in order to properly grasp the opportunity present for WSNs in the field of Egyptian agriculture, a system is developed within Egypt's technological, sociopolitical and economic constraints using locally available hardware. The system is then fit with a specially tailored implementation of the LEACH routing protocol to enable the system to deliver within said constraints. A thorough assessment of the WSN is then performed in order to fully quantify its potential at fulfilling the needs of the task at hand.

Consequently, this paper presents said system with Section II describing the complete design of the system, including the hardware, the implementation of the routing layer, and the testing bed used to evaluate the performance of said protocol. Section III then proceeds to describe and discuss the different tests performed on the protocol generated in Section II to ensure safe and optimal deployment. Section III also presents the results attained from said tests alongside a discussion of these results. Section IV will end the paper by presenting the conclusions drawn from the study performed as well as providing recommendations for any future work and proposing future work based on the findings of this study.

II. System Design

This section is mainly concerned with describing the system design. The first subsection covers the constraints within which the design had to be developed. The second subsection described the hardware utilized to form the WSN nodes. The third portion gives an overview of the method by which LEACH was adapted to form the routing layer, and, lastly, the fourth subsection illustrates the software and hardware aspects of the testing bed developed to examine and compare the several facets of the installed routing protocol.

A. System Constraints

The constraints faced during the design of this system were namely sociopolitical, legal and economic issues in addition to those imposed by the task of designing for the application of permaculture farming.

Literacy-wise, Egypt is listed by UNESCO as one out of the 10 nations responsible for three fourths of the World's 774 million illiterate adults [8]. Therefore, the system should not expect user intervention. It should support simple maintenance procedures and should be inherently self-healing.

Further system constraints are due to the limited selection of locally available communications equipment. At the moment, three RF modules are offered for civilian use. These include Nordic Semiconductor's nRF24L01+, a longer range variant of the nRF24L01+ that includes PA and LNA circuits on boards, and, lastly, a generic RF module based on the TI CC1101 chip.

In an effort to avoid any legal issues, any system developed must comply with two points for the foreseeable future. First off, the system must be based off of any one or more of the three legally permitted chips listed above. Second of all, the system should not be deployed until approval from the authorities is attained. This means that, up until the point in time at which this paper was written, all testing has taken place in a controlled, small-range, non-deployable state.

Also, this system is being developed under strict economic constraints. Egypt's gross national income (GNI) per capita is 2980 USD, and for the lower middle income strata it is 1893 USD [6]. For widespread adoption the system should therefore have cost-effectiveness and robustness as top priorities.

Furthermore, due to several agrarian reforms that may be traced back to the 1950s, farms in Egypt have frequently been resized. Hence, according to [8] 95% of landowners have less than two hectares of farmland each. The remaining percentage of landowners have larger farms that account for almost half of Egypt's agricultural land. This implies that to support both farm sizes, the WSN should be scalable and capable of supporting long-range communications.

Lastly, away from the constraints of developing for Egypt, there are those constraints which are imposed by designing for permaculture.

Permaculture operates in a zonal manner based off of the partitioning of the design elements. This segregation is mainly dependent on the human environment and the projected frequency of human visits to each zone.

Also, functionally speaking, certain tasks to be handled by an agricultural WSN system are considered more important than others. For example, monitoring for fire outbreaks is a generally more important task for a WSN than monitoring the air's temperature or soil moisture levels.

Consequently, WSN systems are to mirror these concepts in its own design. This means that nodes should allow for their permanent assignment to clusters based on their location and function. Each cluster should not route for other clusters whom may be responsible for tasks of lesser importance, thereby conserving the channel's availability and their power supplies for their own tasks.

With these ideas in place, the hardware design to meet and operate within these constraints were collected and interfaced, as is described in Section II-B.

B. Hardware Design

As previously mentioned, local suppliers legally had three RF modules to offer. Based on the comparison shown in Table I, the nRF24l01+ was chosen as the better contender. This

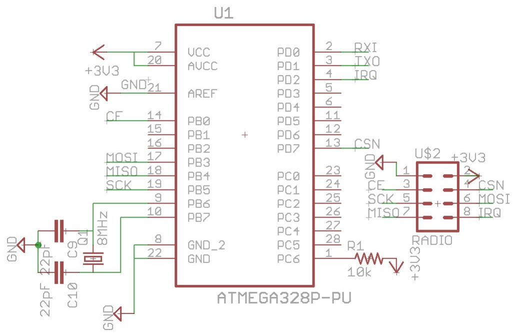

Fig. 1: The Final Hardware Design of the WSN Node

TABLE I: Transceiver Comparison

	Wireless Module	
	CC1101-Based	nRF24l01+
Frequency	433 MHz	2.4 GHz
Bandwidth	-	1000/2000 kHz
Modulation	FSK	GFSK
Sensitivity	-116 dBm	-85 dBm
RSSI	-105 - -50 dBm	1 bit RSSI
Voltage	1.8 - 3.6 V	1.9 - 3.6 V
Data Rate	600 Kbps	1 and 2 Mbps
Price	60 LE	40 LE
Range	200m	85m

is since it provides better data rates at a lower price and has a range of 85m, which is sufficient for a feddan's deployment. If longer ranges are needed, it is directly interchangeable with the nRF24l01+ with LNA and 2dB antenna, which is also locally available, has a range of 1000m, and does not need any extra hardware or configuration changes beyond that which is already in place for the shorter-range nRF24L01+.

In terms of the microcontrollers, we chose to operate with Atmel's AVR microcontrollers given their wide-spread availability at wholesale prices.

Out of the AVR family, the only DIP package microcontroller locally available is the ATMega328-PU. In contrast to its surface-mount alternatives, with the addition of a DIP socket to the final circuit, no soldering would be required for the replacement of faulty microcontrollers. This greatly simplifies the projected maintenance process, partially fulfilling the requirement for maintenance simplicity as imposed by the aforementioned constraints.

The final node therefore consists of an ATMega328-PU, a single 10K pull-up resistor, two 22pF capacitors and a 16MHz crystal for clocking, and the nRF24L01+ module for communications, all operated using a single 3.3V supply, as shown in Fig. 1. In total, the cost of a single node, has come to be 69LE, or 10 USD.

C. Routing Layer Design

The nRF24L01+ does not allow for a true-mesh design. The nRF24L01+'s firmware is programmed to operate in a star

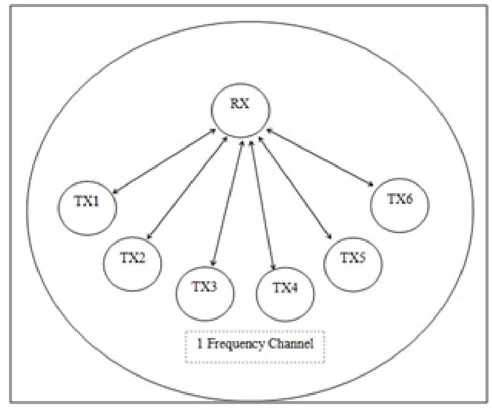

Fig. 2: The Native Topology of the NRF24L01+ Module

topology with a maximum of 7 nodes per network, as shown in Fig. 2. However, the transceiver may be programmed to switch between both receiver and transmitter roles, thereby allowing it to operate in two star topologies at once. This is at the cost of having the receiver occupy 2 of the available 7 node positions in each star, but allows for the development of a hierarchical network composed of consecutive star topologies.

Reviewing the available hierarchical protocols listed in [9], the Low-Energy Adaptive Clustering Hierarchy (LEACH) routing protocol is the most suitable for the aforementioned setup. This is since LEACH is an energy-efficient protocol that creates a dynamic single-hop routing backbone. This backbone is composed of routers known as cluster heads (CH). Given the hierarchical nature of nRF24L01+ modules' topology, every star topology's receiver may be considered a cluster head. Therefore, with the network structure already in place, by adopting the methods of LEACH, the network's lifetime may be further lengthened. This will however require a modification of the methods of LEACH.

LEACH operates based on a probabilistic function which is periodically invoked by each node to calculate if it will be a CH. This calculation is based on the node's current energy level and the number of times it has fulfilled the CH role. Once a CH role is self-assigned, the node advertises this information and awaits non-CH nodes to join its cluster. Non-CH nodes make their selection by choosing the cluster whose advertisement has the strongest RSSI. The CH then creates a TDMA schedule and broadcasts it through its cluster. Once the TDMA frame is completed, the clusters dissolve, new CHs rise out of the nodes that have not yet performed the CH role, and the process repeats [9].

Modifications to LEACH, however, are necessary. This is since, due to the constraints of Section II-A, nodes should not be permitted to leave their clusters, implying that clusters are not to dissolve and reform. Also, RSSI may not be used to determine which cluster is best for a node to join as the nRF24L01+ only has a single bit RSSI register. Furthermore, in contrast to vanilla LEACH, since a network of nRF24L01+ modules may effectively only have 5 members per star configuration, the LEACH protocol may not assign the CH role to more than 5 nodes.

With this in mind the routing layer designed adapts LEACH by utilizing various network roles and system level packets, and by altering its native methods of operation.

The network designed has three roles that may be assigned to any node. These are the sink node, parent and child roles. The sink node is where all data is directed. The parent node generates data and routes packets for other nodes. The child node generates data and transmits it to the sink node via a parent node. A parent may also be a child of another parent node, and vice versa.

These three node types utilize four system-level packet types. These are the network topology packet, the energy level packet, the CH selection packet and the duplicate mitigation packet.

The network topology packet contains the full map of active nodes. This packet is only generated and transmitted by the sink node. It is built based on the nodes that communicate with the sink node. Once the sink node becomes aware of an active node, it adds it to its network map and disseminates the new map to all active nodes on the network.

The energy level packet contains a representation of a node's energy level, instead of the actual measured voltage of the node's power supply. This representation is in the form of a counter that tallies the number of packets transmitted and relayed by the node in question. This is since, as was shown in [10], the voltage of a discharging battery is fairly uniform throughout its lifetime. Given that the cost of a single packet transmission using the nRF24l01+ is constant, a packet counter has been used to represent the amount of energy consumed by any given node.

The last two packet types are the CH selection and duplicate mitigation packets. The former is a unidirectional packet that is transmitted when a parent is requesting that a child take on the role of CH in its stead. The latter is a packet sent by the sink node to nodes that have been found to occupy the same position in the network.

Now that the underlying framework of the protocol has been explained, the following is an account of how the protocol operates.

Once a non-sink node is powered on, it transmits its energy level to the sink node. If the node is a direct child of the sink node, then this is a single-hop transmission. If the node is a descendant, then the packet will only transmit if the intermediate nodes, its parent and every consecutive parent there on, are also powered on and able to receive. With this information, the sink node activates the node on its network graph and disseminates this information throughout the network. If a node is to transmit to any other node, it follows the same routing path with the message having to traverse to the sink node first. The only exception is if the destination is part of the same cluster. In this case, the message only continues up the tree up till the destination's parent before the parent, as its CH, forwards the message to the appropriate node.

Unlike LEACH, the process of CH selection is initiated by each respective CH once it has transmitted a certain number of packets. Once that point is reached, the parent reviews the energy levels of its children. If the average energy of its children is higher than its own energy level, then the child that has transmitted the least number of packets is chosen for the CH role. This, however, requires that each node periodically transmit the contents of its packet counter to the node's parents.

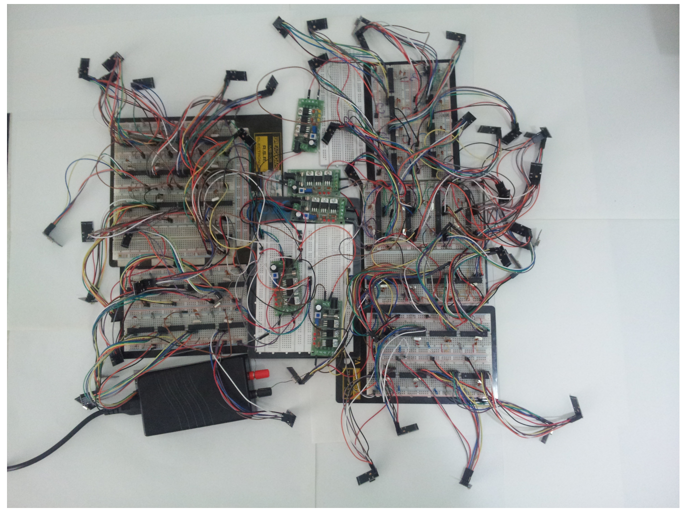

Fig. 3: The 49 Nodes Test Bed

The following checks are performed to introduce a degree of reliability to the CH selection and assignment process:

1. If a CH selects a child to take its role, the CH transmits a CH message with its new role to said child. This is repeated until an acknowledgement is received or the retransmissions limit is met. The CH then relegates itself to that child's position in the network. Once there, the node then attempts to transmit a message to the new CH using the CH's address. If no response is received, it is assumed that the child did not successfully receive the CH message and the process is reversed and repeated.

2. Any message routed to a child that goes unacknowledged by the child is taken as an indication that the child node has failed. This causes the parent to discount this child from the CH selection process. If the child acknowledges any ensuing messages addressed for it, the parent no longer discounts this child.

3. If several consecutive messages being transmitted by a child to or via its parent fail, the child assumes that the parent node has failed. To ensure that the parent is dead, further test messages are sent to the parent node in close succession to each other. If these test messages remain unacknowledged, the child node choose one of its children to replace its role before rising up to take up the role of the failed parent node. This process then repeats down the tree until the entire network self-heals.

4. If the sink detects more than a single node occupying the same position in the network, the sink transmits a duplicate mitigation packet to each of them, save one. Any node that

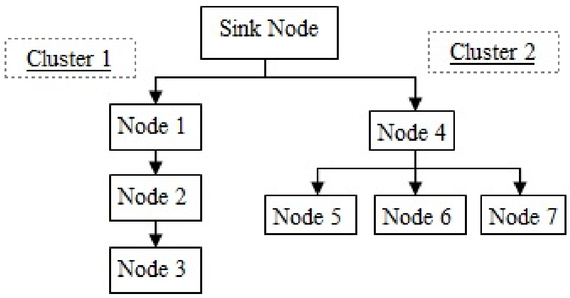

Fig. 4: The Network Topology of Tests 1 and 2

receives this packet finds an unoccupied position in the lower layers of its cluster and relegates itself to that position.

In accordance with the findings of [2] and [3], the protocol is designed to support three hops. This limits the number of nodes per channel to 157 nodes. Although any number of hops beyond then are possible, this is neither supported nor recommended due to the incurred path-long transmission costs, possible message delivery failures, and possible uncapped growth of a network using a single sink node. Instead, it is recommended to have another sink node operating on another channel with its own 157 children and descendants. With the nRF24l01+ being able to operate in 125 channels, the maximum number of nodes within a single locale is 19,625 nodes.

D. Testing Bed Design

In order to test the WSN, the nodes are loaded on five breadboards in a lab environment as shown in Fig. 3. Two further breadboards are used as power rails.

In order to power the entire test bed, a single voltage 220V-to-12V step-down voltage converter with a 2.5A output was used. The output was then divided between 5 separate variable voltage regulators. A pot resistor is connected to each regulator to fine tune the voltage supplied to each breadboard to match the 3.3V supply needed by each node. In the case of non-sink nodes, in order to limit the simulation time and gather comparable results, each node is given the same number of packets that it may transmit or relay. Once the node reaches that number, the node halts its transceiver, records its results on the internal EEPROM and awaits either the manual collection of data or a signal to initiate a new simulation.

The sink node, on the other hand, does not receive a pre-set cap, and continues to operate as long as there is a functional node. This provides real time, time-stamped data for analysis throughout the simulation. This includes every packet sent and received by the sink node with source and destination addresses, node energy levels, message success rates and any soft resets or reboots self-initiated by the sink node.

The data recorded by end nodes at the end of every simulation consists of four pieces of information. These are the total number of messages sent by the node, the total number of successful messages, the node's personal history, which is every position in the network that the node has occupied, and the amount of packets left in its packets counter, which is relevant only for soft reboots.

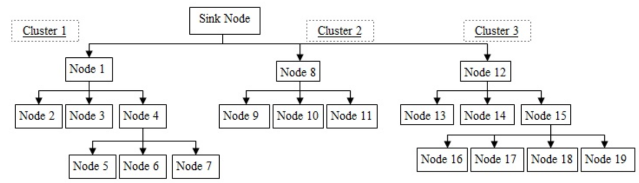

Fig. 5: The Network Topology of Test 3

A faulty node was also programmed for testing purposes. This node, modelled on the behaviour of the several partially failed nRF24L01+ transceivers found while prototyping, was set to transmit and receive, but to refuse to relay messages on behalf of other nodes.

III. Results

The first three tests performed explored the optimal topology characteristics, and the first two in specific also tested the effect of including a faulty transceiver. The fourth test sought to observe the effect of having a four layer cluster on the network. The fifth test observed the effect of incorporating parent failure detection and mitigation functions. The sixth test tested said functions on a network deployed with the optimal topological characteristics determined from the previous tests. Finally, the seventh test added node duplicate detection and mitigation functions and a faulty transceiver to the same network of test six.

Each respective test is run five times with packet counters set to a limit of 2000 packets. Each node is set to trigger its CH assignment functions every 50 packet transmissions. During the tests all nodes are programmed to transmit a packet to every other node in the network with a spacing of (2xN) seconds between each transmission, where N is the total number of nodes. These settings have resulted in over 840 hours of simulated run time.

Due to the large amount of data collected during the simulations, it was chosen to best represent this information in the summarized forms shown in Tables II and III. Distributing nodes based on their success rates into success rate intervals for each test, Table II portrays this distribution in terms of percentages. Table III shows the percentage of the total number of packet failures in each test attributable to each type of failure.

Observing the results of the first two tests, Table II shows that they both recorded high failure rates. The majority of these failures were due to cluster isolations and the faulty transceiver.

Cluster isolations resulted when a cluster head, which in this case would be either node 1 or 4 shown in Fig. 4, failed

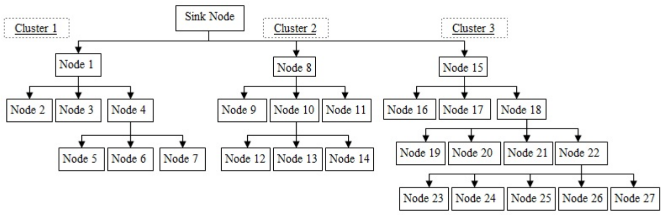

(a) Test 4

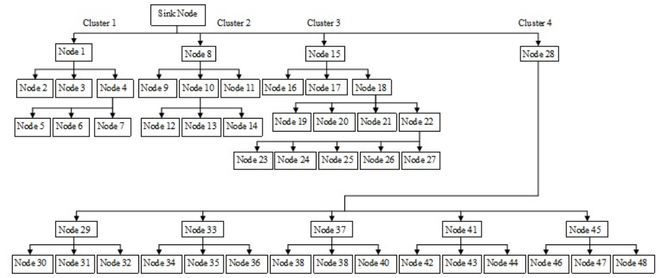

(b) Test 5

Fig. 6: The Network Topology of Tests 4 and 5

to select a replacement before it depleted its packet tokens. With no child node rising up to take its place, the cluster in question would effectively be isolated from the rest of the network leading to high message failures.

In addition, the network was unable to cope with the inclusion of a faulty transceiver. Over 50% of the packet failures observed in test 1 were due to the faulty transceiver. Once removed, the cumulative percentage of nodes in test 2 achieving less than 50% success rates halved while those achieving between 50 and 80% and over 90% almost doubled in number. This difference in success rates is due to the faulty transceiver always having the largest store of packet tokens left as it did not relay packets for other nodes. It was therefore almost always selected as the new CH for cluster 1, which consistently resulted in the isolation of cluster 1.

These tests also observed a small portion of failures resulting from the relaxed scheduling used. This property of the routing layer made it possible for several nodes to attempt to transmit at the same time, resulting in them interfering with each others' transmissions.

It is worth noting that cluster 2 was found to be more successful in both tests. By having a larger pool of nodes at layer 2 of the cluster for the CH to select from, cluster 2 had a longer lifetime than cluster 1, which delayed the onset of cluster isolation.

Capitulating on the topological findings of test 1 and 2, test 3 adopted the more horizontal distribution shown in Fig. 5. Multi-hop transmissions were increased to a maximum of 3 hops to and from the sink node. The increase in number of nodes increased the overall number of packet tokens available and, consequently, lengthened the testing scenario. Due to the relaxed scheduling used, this larger number of nodes also increased the chances of CH assignment failures and of message delivery failures and collisions.

Even with these risks, overall, most of the distribution lay between 50 and 90% success rates. However, with the elimination of the predominantly vertical nature of cluster 1 from this network, test three had 55.8% of nodes achieving success rates between 0 and 80%. This is only a slight improvement on test two's 65%. This was mainly due to the bottom heavy characteristic of this network with the number of nodes increasing with every layer. This effectively meant that, more often than not, CHs would quickly be drained of packet tokens before being able to select a replacement, causing early network isolations, which meant that, like test 2, cluster isolations were still the largest cause of packet delivery failures.

Test 4 increased the maximum number of hops from a node to the sink to 4 hops by adding a fourth layer, as shown in Fig. 6b. This was expected to further exacerbate the problem of cluster isolations. Rather, compared to test 3, the distribution of nodes achieving success rates above 90% and between 80 and 90% were about the same, but less achieved 50-80% and more achieved less than 50%.

In this test, only 54% of packet failures were due to CH isolations. This is since 43% of the packet failures were directly caused by the 4th layer nodes. With the large number of hops required, their messages were rarely delivered successfully. Alternatively, when nodes 23 to 27 of layer 4 did rise up to layer 3 or above, their success rates improved, but at the cost of reducing other node's success rates. This is since they would be trading positions with other more successful nodes which reduced the overall effectiveness of the network. Another reason for the marked increase in the number of node achieving less than 50% success rates is due to the increase in the number of possible points in the network where a CH selection process may fail, and the larger

In order to test the remedies to the aforementioned problems, test 5 examined the effect of adding the CH failure and mitigation functions. The tests were run with the worst case scenario of having the fourth layer nodes of cluster 3 still a part of the topology and with an added 21 nodes as part of cluster 4. Cluster 4, however, was distributed in the manner seen as the most ideal based on the previous tests, as shown in Fig. 6(b).

The results of test 5 show huge improvements in the achieved success rates. Including the parent death detection and replacement functions allowed for 50% of the distribution to achieve success rates above 90%, and 75% over 80%, respectively. This, however, was at the cost of 5% of the packet failures being directly attributable to these functions. This is since, to operate these functions, many packets are

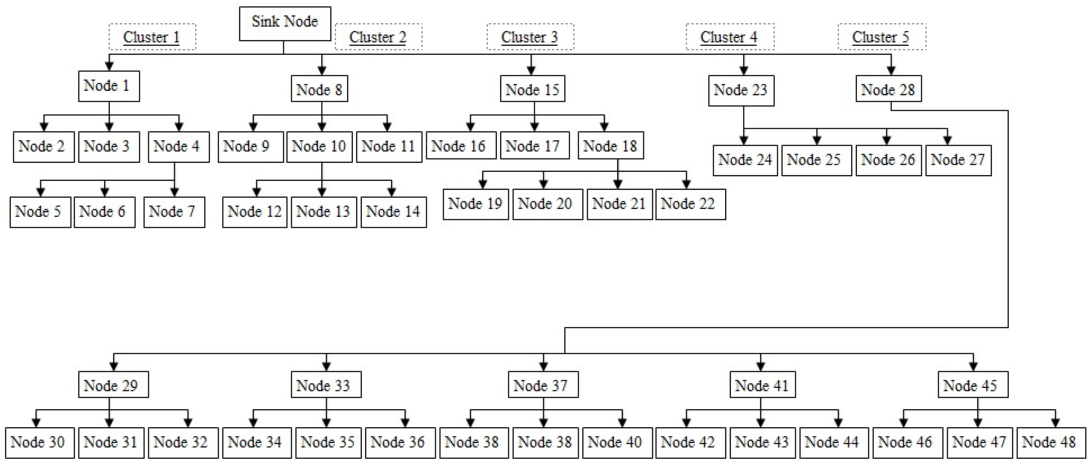

Fig. 7: The Network Topology of Tests 6 and 7

transmitted and expected to fail to confirm that nodes are down and that the node failure mitigation procedures should be initiated. Also, while these functions are executing, the cluster in question suffers a temporary period of isolation, incurring further increases in message failure counters. The inclusion of the CH failure and mitigation functions did also cause the appearance of a new sort of behaviour. This new problem is the rise of node duplicates and their effects on network efficiency.

Node duplicates, which were responsible for almost 18% of failures, occurred when several nodes simultaneously detected a CH failure. These nodes would then attempt to replace the CH resulting in the existence of several CHs per cluster. These duplicate nodes would then be found to interfere with each other's acknowledgement messages. This caused false increments in nodes' failure counters, which led to more duplicate CHs, as well as allowing for a rise in the number of message duplicates traversing the network.

Layer 4 nodes also reduced the success rates in the same manner that was observed in test 4. This layer was directly responsible for 72% of the packet failures observed, causing 10.9% of the distribution to lie within the 0-50% success rate interval. This layer was also partially responsible for the decrease in the success rates of other nodes to the 50-80% segment of the distribution due to their trading of network positions with more successful nodes.

Observing the effect of eliminating the fourth layer on the issue of node replication, test 6 commenced with the topology seen in Fig. 7. The five nodes that were previously used to make up layer four of cluster 3 were used to create an evenly spread fifth cluster.

The simulated runs of test 6 did note a marked increase in the number of nodes that achieved success rates between 50% and 90%, and a decrease in those that attained success rates above 90% and below 50%, respectively. These results signify an overall decrease in the number of message delivery failures and are due to the change in topology and the newly added functions. These changes, however, also caused an increase in the opportunity for node duplicates to exist. This raised the number of node duplicates witnessed and the percentage of failures associated with their existence.

The overall decrease in the packet failures also meant that the percentage of failures attributed to the relaxed scheduling was almost double the percentage witnessed in previous tests. This does not signify that relaxed scheduling caused more failures, but only that it is now responsible for a larger portion of the total.

Test 7 then repeated test 6 with the inclusion of the node duplicate mitigation functions. This test also reintroduced the faulty transceiver as node 3 in order to witness the behaviour of the added procedure in the presence of a faulty node.

Results showed that around 11.9% achieved less than 50%. This was equally caused by the mitigation functions, the faulty transceiver and overhead required to operate the protocol. The faulty transceiver, always having the largest number of tokens left, would consistently be chosen to fulfill the role of CH. However, as members of the cluster detect the message delivery failures associated with having a faulty node as the CH, they rise up to take on the role of CH. This meant that two CHs would exist for a short period of time until the duplicate mitigation function executed and caused the faulty node to relegate to a lower position in the network. Although this sequence of behaviour has successfully curbed the percentage of failures normally associated with transceiver faults, the aforementioned process repeats indefinitely, accumulating its own overhead of packet failures.

It was also found that, occasionally, while a node executed its duplicate node resolution functions, it would not find a free position to occupy within the top three layers of its cluster. Hence, the node would be forced to relegate itself to the fourth layer of the network. This was responsible for about 10% of the failures witnessed. The majority of failures, however, were due to the overhead required to operate the protocol. These

TABLE II: Distribution of Nodes into Success Rate Intervals

	Percentage of Nodes Per Success Rate Interval (%)			
Test #	>90%	80-90%	50-80%	<50%
1	12.5	10.0	12.5	65.0
2	22.5	12.5	32.5	32.5
3	15.8	28.3	35.0	20.8
4	17.0	24.1	25.9	33.0
5	54.4	21.1	13.6	10.9
6	40.7	31.4	25.7	2.1
7	6.5	22.9	59.6	11.0

TABLE III: Failure Rates vs. Failure Cause

Failure Type	**Test 1**	**Test 2**	**Test 3**	**Test 4**
Scheduling	0.86%	4.86%	2.05%	3.16%
Overhead	0%	0%	0%	0%
Isolation	43.67%	95.14%	97.95%	53.9%
Transceiver	55.47%	0%	0%	0%
Layer 4	0%	0%	0%	42.94%
Duplicates	0%	0%	0%	0%
Failure Type	**Test 5**	**Test 6**	**Test 7**	-
Scheduling	5.67%	9.82%	0.25%	-
Overhead	4.90%	37.63%	72.16%	-
Isolation	0%	0%	0%	-
Transceiver	0%	0%	17.40%	-
Layer 4	71.83%	0%	10.19%	-
Duplicates	17.60%	52.55%	0%	-

failures were due to the large number of packets that were transmitted by nodes to fail such that they may confirm that clusters were isolated or that specific network positions were unoccupied and available for relegation. This behaviour was responsible for 72.16% of packet failures and for 59.6% of the distribution to lie within the 50-80% success interval.

IV. Conclusion

A multi-hop WSN is built using nodes composed of ATMega328-PU microcontrollers and nRF24L01+ transceivers. These nodes have realized the goals of low monetary and power costs with each node having a price of 10USD and consuming 25.6mA, respectively.

The LEACH protocol was subsequently tailored to meet the system in question and tested extensively. These tests have made it apparent that the distributed approach of having end nodes detect and mitigate CH failure works better than having the CH perform said task. This, however, increases the complexity of the protocol by requiring further functions in order to ensure adequate network performance levels.

Also, given the transceiver's property of fixed costs per transmission, topologies are advised to have a maximum of three hops from the sink node. This is in concurrence with results which observed a sharp drop in success rates when nodes were deployed at a distance of four hops from the sink node.

Therefore, from these conclusions it is apparent that from a hardware standpoint the system is affordable and meets its constraints. From the viewpoint of the protocol, it is currently at a promising stage of development as further research may find more efficient methods for the routing layer of this WSN.

References

[1] J. Panchard, P. S. Rao, M. Sheshshayee, P. Papadimitratos, and J. -P. Hubaux, Wireless Sensor Networking for Rain-fed Farming Decision Support User Survey, *ACM SIGCOMM Workshop on Networked Systems for Developing Regions (NSDR)*, Seattle-USA, Aug 2008.

[2] E. Biagioni, "Pods: Issues in the Design of Practical Ad-Hoc Wireless Sensor Networks", SSGRR Summer 2003 Conference, L'Aquila-Italy, 2003.

[3] J. Polastre, R. Szewcyk, A. Mainwaring, D. Culler, and J. Anderson. Analysis of wireless sensor networks for habitat monitoring. In Raghavendra, Sivalingam, and Znati, editors, Wireless Sensor Networks, pages 399423. Kluwer Academic Pub, 2004.

[4] F. Ingelrest, G. Barrenetxea, G. Schaefer, M. Vetterli, O. Couach, and M. Parlange, "SensorScope: Application-Specific Sensor Network for Environmental Monitoring", *ACM Transactions on Sensor Networks*, Vol. 6, No. 2, Article 17, pp. 1-32, Feb. 2010.

[5] L. O. Urdiain, C. P. Romero, J. Doggen, T. Dams, and P. Van Houtven, "Wireless Sensor Network Protocol for Smart Parking Application, Experimental Study on the Arduino Platform", *Proceedings of the 2nd International Conference on Ambient Computing, Applications, Services and Technologies*, Barcelona-Spain, Sept. 2012.

[6] The World Bank Group. (2014, Mar.) *Egypt, Arab Rep. at a glance*. [Online]. Available: http://devdata.worldbank.org/aag/egy_aag.pdf.

[7] "What's "Permaculture"?", *Countryside and Small Stock Journal*, Vol. 86, No. 3, pp. 80, May 2002.

[8] M. A. El-Nahrawy. (2011, Jun).*Country Pasture/Forage Resource Profile: Egypt*. Food and Agriculture Organization of the United Nations. [Online]. Available: http://www.fao.org/ag/AGP/AGPC/doc/Counprof/Egypt/Egypt.html.

[9] A. M. Zungeru, L. Ang, and K. P. Seng, "Classical and Swarm Intelligence Based Routing Protocols for Wireless Sensor Networks: A Survey and Comparison", *J. of Network and Computer Applications*, Vol. 35, pp. 1508-1536, Mar. 2012.

[10] J. Zhang, S. Ci, H. Sharif, M. Alahmad, "Modeling Discharge Behaviour of Multicell Battery", *IEEE Transactions on Energy Conversion*, vol. 25, no.4, pp. 1133-1141, Dec. 2010.