International Journal of Scientific Engineering and Technology (ISSN : 2277-1581)
www.ijset.com, Volume No.1, Issue No.4, pg : 117- 126 01 Oct. 2012

A Passive Diagnosis for Self organizing and Distributed Wireless Sensor Networks

A.Rajasekaran[1], **K.Saraswathi**[2]

[1]Assistant Professor, ECE Department, SCSVMV University, Enathur, Kanchipuram.
[2]Assistant Professor, E&I department, SCSVMV University, Enathur, Kanchipuram.
Email: shrisairaja@gmail.com,saraswathi_krani@yahoo.co.in

Abstract

Network diagnosis, an essential research topic for traditional networking systems, has not received much attention for wireless sensor networks (WSNs). Existing sensor debugging tools like sympathy or EmStar rely heavily on an add-in protocol that generates and reports a large amount of status information from individual sensor nodes, introducing network overhead to the resource constrained and usually traffic-sensitive sensor network. We report our initial attempt at providing a lightweight network diagnosis mechanism for sensor networks. We further propose PAD, a probabilistic diagnosis approach for inferring the root causes of abnormal phenomena. PAD employs a packet marking scheme for efficiently constructing and dynamically maintaining the inference model. Our approach does not incur additional traffic overhead for collecting desired information. Instead, we introduce a probabilistic inference model that encodes internal dependencies among different network elements for online diagnosis of an operational sensor network system. Such a model is capable of additively reasoning root causes based on passively observed symptoms. We can implement the PAD prototype in our sea monitoring sensor network test-bed.

Index Terms—Diagnosis, passive, sensor networks.

I. INTRODUCTION

Wireless Sensor Networks (WSNs) have been widely studied for enabling various applications such as environment surveillance, scientific observation, traffic monitoring, etc. [14], [28]. A sensor network typically consists of a large number of resource-limited sensor nodes working in a self-organizing and distributed manner. Having made increasing efforts [6], [7], [10]–[12], [16], [18]–[20], [27], [31] on the robustness and reliability of WSNs under crucial and critical conditions, researchers, however, have done little work targeting the in-situ network diagnosis for testing operational sensor networks. It is of great importance to provide system developers useful information on a system's working status and guide further improvement to or maintenance on the sensor network. Due to the *ad hoc* working style, once deployed, the inner structures and interactions within a WSN are difficult to observe from the outside. Existing works for diagnosing WSNs mainly rely on proactive approaches, which implant debugging agents into sensor nodes, periodically reporting the internal status information of each node to the sink, such as component failures,link status, neighbor list, and the like. For example, Zhao *et al.*[33] propose to scan the residual energy and monitor parameter aggregates including link loss rate and packet count. Such information is collected locally at each node and transmitted back to the sink for analysis. Sympathy [23] actively collects run-time status from sensor nodes like routing table and flow information and detects possible faults by analyzing node status together with observed network exceptions. The proactive information generation and retrieval exerts extra computational operations on sensors and imposes a large communication burden on a WSN, which is usually fragile at high-traffic loads. Those approaches work more like debugging or evaluation [26] tools before the system is released for use outside laboratory settings. While such tools are effective for offline debugging when sensor behavior and network scale can be strictly controlled, they may not be suitable for in-situ network diagnosis of an operational WSN since they continuously generate a large amount of traffic and aggressively consume computation, communication, and energy resources. Also, integrating those complex debugging agents with application programs at each sensor node introduces difficulties for system development. This work is motivated from our ongoing sea monitoring project [4], [30]. As shown in Fig. 1, for this project, we launched a working prototype WSN consisting of tens of nodes that float on the sea surface and collect scientific data such as sea depth, ambient illumination, pollution, and so on. Recently, in the field deployment tests, we often observed abnormal energy depletion that never occurred in the controlled laboratory experiments. We suspect that such a phenomenon is due to the usage of the Multi Hop Router (integrated in SURGE) component that frequently switches the optimized routing tree of the network owing to the highly instable environment of the sea. We also observed other problems on the sink side such as high delay of data sampling and unbalanced packet loss. Fast and accurate identification of the root causes is necessary before taking any further action such as issuing reboot messages to certain nodes or physically examining the suspicious links. With current debugging tools, it is indeed difficult to integrate

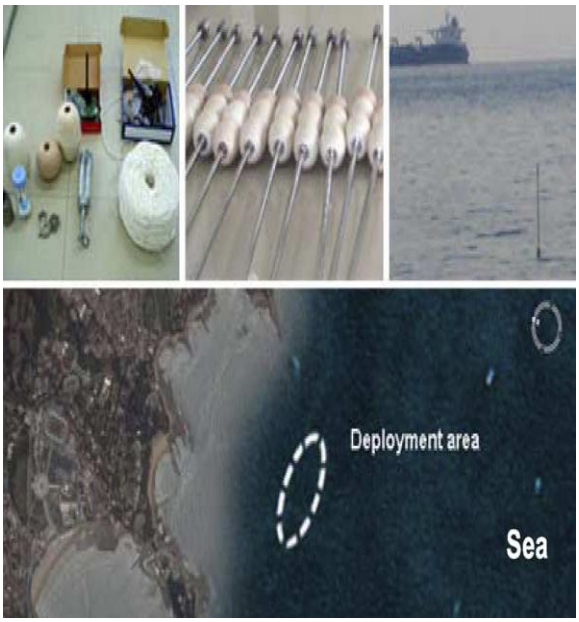

Fig. 1. OceanSense project.

their agents with our application programs. It is even worse if we implant proactive information collectors in the network, which would inevitably speed up the depletion of energy and rapidly reduce the expected lifetime of the sensor network. In this study, we propose an online diagnosis approach that passively observes the network symptoms from the sink. Using probabilistic inference models, this approach effectively deduces the root causes of abnormal symptoms in the network. Compared to proactive debugging tools, the passive diagnosis approach observes data from routine application packets for back-end analysis. It can also be maintained in a running system at lightweight cost, thus it is expected to accommodate the application system in a timely manner without degrading performance. Inference-based network diagnosis methods have been widely investigated and applied in enterprise networks [5].

Various types of inference models, both deterministic and nondeterministic, have been proposed for inferring the root causes of service failures. Most models are built on expert knowledge or trained from historical data from the networks. The construction of such models can be very complicated, and once constructed, the models are often viewed as maining unchanged for a relatively long period [5], as enterprise networks are usually stable with few dynamics in their structures. Compared to enterprise or static networks, however, sensor networks have the following unique features: 1) sensor nodes have extremely limited computational and energy resources; 2) the network topology is highly dynamic due to the instable environment and acquiring prior knowledge of the network is difficult; 3) the individual sensor nodes are error-prone. Such conditions make existing active approaches for static network diagnosis infeasible. Thus, WSNs cannot easily adapt to such slow start approaches as sensors are self-organized without any prior information on the dependencies among network elements. The high dynamics of the WSN structure also leads to the infeasibility of those inference models built from static data. We address the above challenges as follows. First, we introduce a packet marking scheme, which marks the regular routine communicating packets to continuously reveal their communication dependencies within the network. Using the output of the scheme, the sink constructs and dynamically maintains a probabilistic inference model. This scheme works in a lightweight manner without any extra transmission in the network and can adapt to frequent network changes. Second, we employ a hierarchical inference model that captures multilevel dependencies in the network. The hierarchical model can be constructed based on incomplete information, and it is able to efficiently handle the network dynamics by updating only the changed parts. This model takes both positive and negative symptoms as input and reports the inferred posterior probability of possible root causes. Third, we design an online inference engine capable of additively reasoning the root causes such that it works even with incomplete or suspicious inputs in a nondeterministic manner. The major contributions of this study are as follows. To the best of our knowledge, we are the first to investigate a passive method of diagnosing the wireless sensor networks. 1) According to the unique features of sensor networks, we design an efficient packet marking scheme that dynamically reveals the inner dependencies of sensor networks without injecting extra transmissions.

2) We propose hierarchical inference models that capture the multilevel dependencies among the network elements and achieve high accuracy. We further introduce a fast inference scheme that reduces the computational complexity and is thus scalable for online diagnosis in large-scale WSNs.

3) We implement our diagnosis approach, PAD, and test its effectiveness in our sea monitoring project with 24 sensors. The results of our field test show that PAD indeed helps in exploring the root causes of observed symptoms. Relying on the output of PAD, we have successfully improved our application programs. 4) We further analyze and evaluate the scalability and effectiveness of PAD design through extensive simulations under varied conditions using the trace we collect from the prototype implementation. The rest of this paper is organized as follows. Section II introduces related work. Section III describes the framework of our system. We introduce the packet marking scheme in Section IV and discuss the two inference models based on Belief Network and Causality Diagram in Section V. In Section VI, we present our implementation and simulation results. We conclude this work in Section

II. Related Work

Most existing approaches for sensor network diagnosis are proactive, in which each sensor employs a debugging agent to collect its status information and reports to the sink by periodically transmitting specific control messages. Some researchers propose to monitor sensor networks by scanning the residual energy [33] of each sensor and collecting the aggregates of parameters of sensors where in-network processing is leveraged. By collecting such information, the sink is aware of the network conditions. Some debugging systems [23], [29] aim to detect and debug software failures in sensor nodes. For example, Clairvoyant [29] focuses on debugging sensor nodes at source level and enables developers to wirelessly connect to a remote sensor in the network and execute standard debugging commands on that node including break, step, and the like. Sympathy [23] is an advanced debugging tool that detects and debugs the failures in a sensor network. It actively collects in-network information periodically from each sensor node such as neighbor list, traffic flow, and the like and analyzes the network status at the sink. By carefully selecting an optimal set of information metrics, Sympathy aims at minimizing the diagnosis cost so as to be applicable to resource-limited sensor networks. It also applies an empirical decision tree to determine the most likely root causes for an observed exception. Much effort has been expended on network diagnosis for for enterprise networks. Commercial tools [1]–[3] independently monitor servers and routers with various control messages, and alerts are automatically generated from the implanted agents in different network equipment. Those tools, being effective for diagnosing large-scale networks, are too complicated and energy-consuming for resource-constrained sensor networks. There have been some passive diagnosis approaches proposed for enterprise networks that collect a network's operational enterprise networks. Commercial tools [1]–[3] independently monitor servers and routers with various control messages, and alerts are automatically generated from the implanted agents in different network equipment. Those tools, being effective for diagnosing large-scale networks, are too complicated and energy-consuming for resource-constrained sensor networks. There have been some passive diagnosis approaches proposed for enterprise networks that collect a network's operationalstatus from routine data packets so as to deduce the possible root causes of exceptions by an inference model. For example, Score [17] troubleshoots via shared risk modeling. It adopts a simplified two-level graph as the inference model and formulates the problem of locating fault roots as a minimal set cover problem. Kandula *et al.* explore the bipartite graph inference model and propose Shrink, introducing a probabilistic inference scheme [15]. The bipartite graph model approximates the dependencies in enterprise networks and greatly simplifies the complexity of the inference process. Steinder and Sethi [24], [25] also assume a bipartite graph model and apply Belief Networks [21] with the bipartite graph to represent relations among links and end-to-end communications. Shi *et al.* [22] present a fault diagnosis approach for general static complex systems based on Causality Diagram. The above schemes either require pre knowledge of the network dependencies, which are obtained through Shared Risk Link Groups or SNMP in a relatively stable enterprise network, or adopt simplified models to approximate the network dependencies. A WSN, however, is featured by its hierarchical multilevel structures, which can hardly be approximated by the bipartite graph model. It is also unpractical to maintain the network dependencies as stable inputs in highly dynamic and self-organized sensor networks.

The recently proposed Sherlock is the only work that adopts a multistate and multilevel inference graph for the network diagnosis [5]. They use a scoring function to derive the best explanations(root causes) for observed service exceptions. In order to avoid NP-hard computation complexity, they assume that there are at most a small constant number of failures in the enterprise network. This assumption is not valid for the unreliable and lossy WSNs. Guo *et al.* 13] tackle the problem of detecting nodes with faulty readings.

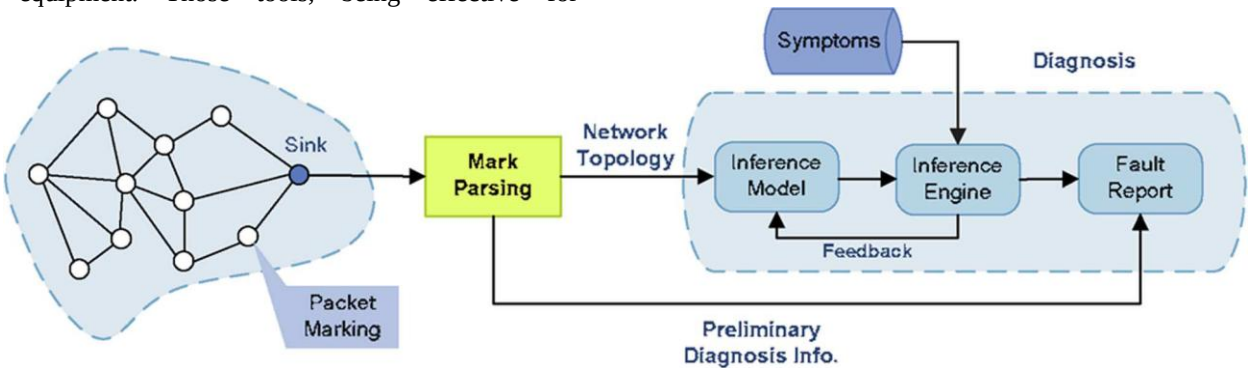

III. System Framework

We view the sensor network as a method for data acquisition in which source nodes periodically sample data and deliver

them back to the sink through multi hop communication. We do not assume any specific routing strategy, that is, our approach deals with networks of various communication topologies such as spanning tree or directed acyclic graph (DAG). We design a passive diagnosis approach, PAD, for such sensor networks. PAD aims to help network managers explore the root causes of exceptions in a running sensor system. PAD implants
a tiny lightweight probe into each sensor node that sporadically marks routine application packets passing by so that the sink can reassemble a big picture of the network conditions from those small clues. Nevertheless, information from marking probes is quite limited and not sufficiently accurate. PAD employs a probabilistic model to infer the statuses of unobservable network elements and reveal the root faults in the network. PAD denotes the observed abnormal situations as negative symptoms such as a long time delay of data arrival or frequent packet loss. It denotes any successful packet reception as positive symptoms. The inference model inputs both negative and positive symptoms to derive network statuses. As illustrated in Fig. 2, PAD is mainly composed of four components: a packet marking module, a mark parsing module, a probabilistic inference model, and an inference engine. The packet marking module resides in each sensor node and sporadically marks routine application packets passing by. At the sink side, the mark parsing module extracts and analyzes the marks
carried by the received data packets. The network topology can thus be reconstructed and dynamically updated according to the analysis results. The mark parsing module also generates preliminary diagnosis information such as packets loss on certain links, route dynamics, and so on. The inference model builds a graph of dependencies among network elements based on the outputs from the parsing module. Using the inference model and observed negative and positive symptoms as inputs, the inference engine is able to yield a fault report, which reveals the root causes of exceptions by setting the posterior probabilities of each network component being problematic. The inference results are also taken as feedback to help improve and update the inference model.

IV. PACKET MARKING

Since a sensor network has a self-organized time-varying network structure, unlike the case in an enterprise network, no prior knowledge can be obtained for constructing the inference model. Also, as a WSN topology is highly dynamic, we need to acquire the network statuses continually to maintain the topology in real time. To address the above requirements, we design a packet marking algorithm in PAD, which dynamically captures the network topology and extracts the inner dependencies among network components. Before the analysis results are directed to the inference engine for further reasoning, we can generate a preliminary diagnosis report on some basic network exceptions. The main operation of this marking algorithm is to let sensor nodes stamp their IDs on passing data packets. Due to the size limitations of the data packets used in sensor networks, however,the marking scheme only adds 2 bytes to each data packet that records one node ID. During the packet delivery, each packet is marked by only one selected sensor node based on a set of rules.
At the sink side, the mark parsing module traces back the paths from each source node through analyzing sporadically marked packets. Through assembling the paths from different source nodes, the network topology can be reconstructed along with the regular data delivery of the system. If the network remains
static, the packet marking process automatically converges and stops after the entire network topology is constructed. When network conditions vary, such as when packet loss or route changes occur, the packet marking process restarts somewhere close to the exceptional event. A strength of this design is that it does not inject any extra message into the network and strictly limits the overhead of marks attached to each data packet.

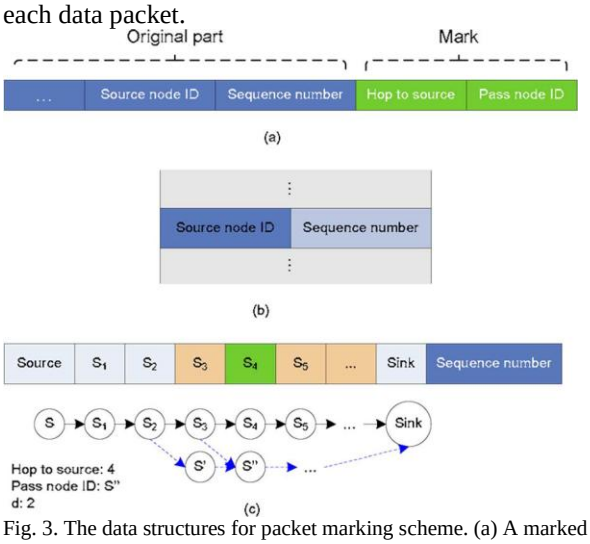

Fig. 3. The data structures for packet marking scheme. (a) A marked data packet. (b) Cache in sensor node. (c) *Path* updating

A. Marking Scheme on Sensor Nodes

Fig. 3(a) depicts an example of marked data packet. We assume that each original data packet contains: 1) a source node ID denoting the source node of this packet; and 2) a sequence number identifying the packet. If there is no such information recorded in the application, the marking scheme adds them to the packets. The mark added to the original packet consists of a pass node ID field that records the ID of a sensor that participates in delivering this packet and a hop to source field recording the number of hops from the source node to the marking node. When the source node issues a new data packet, it leaves the pass node ID field empty and sets the hop to source field to 0. Every intermediate node maintains a cache for its downstream source nodes. As illustrated in Fig.

3(b), each cache entry consists of a source node ID and the sequence number of the recently received packet from the source. We call two sequence numbers of a source *continuous* if the first sequence
number is larger than the latter one by 1. As shown in Algorithm 1, upon receiving a packet, an intermediate node first checks whether the packet has been marked. If yes (the pass node ID is not empty), it forwards the packet with no further operations. Otherwise, the node checks its own cache. If there is no entry for the source node ID of this packet, it marks the packet by filling the pass node ID field with its own ID. It also creates a new entry for this source node in its cache and records the sequence number for the packet. If there exists an entry in the cache for the source node and the sequence number in the packet is *continuous* with the cache entry, the intermediate node updates the cache entry with the new sequence number. To prevent duplicate marking, the intermediate node does not fill the pass node ID field, instead it increments the hop to source field in the packet by 1 and forwards the packet. If the sequence number of the packet is not *continuous* with that recorded in the cache entry, it might be due to the packet loss or routing dynamics. The intermediate node marks the packet by filling the packet pass node ID field with its own ID. The node then updates its cache entry with the new sequence number of this packet and forwards it. The sink also participates in the marking process and creates a table recording source nodes and their packet sequence numbers. Using this marking scheme, the received packet in the sink records the ID of one intermediate node in the routing path together with its hop distance to the source node.We avoid duplicate marks of the same node on the same path to save communication costs. We can further reduce the memory usage in each sensor node by organizing its cache table into bloom filters. Each intermediate node inserts and extracts the source node information on the bloom filter. The error rate introduced by the bloom filter introduces negligible adverse impact in the lossy by-nature sensor network.

Algorithm 1 Packet Marking (packet)

1: **if** has been marked
2: **return**;
3: **else**
4: check cache;
5: **if** no entry for source node of
6: mark ;
7: create entry with *source node ID* and *sequence number* in ;
8: **else if** entry exists **and** *sequence numbers* are *continuous*
9: update entry with new *sequence number*;
10: increase *hop to source* in by 1;
11: **else if** entry exists and *sequence numbers* are not *continuous*
12: mark ;
13: update entry with new *sequence number*;
14: **end if**
15: **end if**
16: **return**;

B. Parsing the Marks

At the sink, the mark parsing module extracts and parses the marks piggybacked from the received packets. For each source node, we keep a data structure denoted as *path* to record node IDs along the path from the source node to the sink. As shown in Fig. 3(c), a *path* contains an array of slots and each slot records a node ID along the routing path hop by hop. The *path* also has a field that records the sequence number of the latest arrived packet from each source. On receiving a new packet, the mark parsing module checks the existence of a *path* structure associated with its source node. If there is no such *path*, it means it is the first time the sinkas received packets from that source. The sink creates a new *path* for the source node and records the source node ID at the first slot. The mark parsing module then examines whether the packet has been marked (the pass node ID field has been filled). If it has been marked, the sink updates the associated slot in the *path* to be the recorded node ID according to the hop to source field in the packet. For the packets from the recorded *path*, the parsing module operates according to the recorded sequence number.We denote as the difference between the sequence number of the received packet and the sequence number recorded in the *path*. If the sequence number of the new packet is equal to or less than that recorded in the *path* , it means that this is a duplicate or delayed packet. As information in the duplicate and delayed packets is usually outdated and may lead to errors in the mark
parsing process, we ignore marks in such packets and do not update sequence number or other slots for the *path*. As a matter of fact, according to our deployment experiences in an operational sensor network, with a relatively long sampling interval, this kind of situation is rare. If , the sequence number recorded in the *path* is updated by the newly received
packet, and then other slots of the *path* are accordingly updated by parsing the mark as Algorithm 2. Normally, without packet loss, , and we directly add the marked node ID into the *path*. Discontinuousness of the sequence numbers indicates
that the packet loss occurs, which triggers a preliminary diagnosis report on packet loss. Besides, the number of packet losses is quantified as . A mismatch of the recorded pass node ID in the packet and the recorded node ID in corresponding slot in the *path* indicates a route alternation happening at the position between the hop to source recorded in the packet and its hops upward. If not so, the marking should have been taken earlier. The parsing algorithm then generates a preliminary report of a route switch. In such a case, the slots in recorded

path ranging from hops before the hop to source position to the sink become inaccurate, so we clear all those slots. Let us look at the example in Fig. 3(c), where a new mark is received. The pass node ID is , four hops away from the source. The mismatch between and current node in the same position of this path indicates a route variation. Now, the issue is how to determine where the route variation occurs. If there is no packet loss, it must be node that changes its route from to. In this case equals 2, indicating that one packe this just been lost. The situation can be more complicated indeed. As illustrated in Fig. 3(c), the route variation can happen at ; for example, changes its parent node from to , and then marks the consequent packet. The packet, however, gets lost on its way to the sink, so before the next packet marked
by arrives, the sink cannot be aware of the route ariation. Another possible case is that the route switch happens at , but fails to send the consequent packet to , and then has to mark the second packet. As the route variation happens, slots ranging from to sink are suspicious. We update to and clear other slots, expecting further information. The reception of the packet without any marks triggers a preliminary report of a successful delivery. The mark parsing function is presented in Algorithm 2. The mark parsing module constructs and updates the network topology with the recorded *path*s. Once a newpacket is received, the *path* associated to its source node is updated. This indicates that all links along the current path have just participated in the transmission of a packet. For each link in the network topology, we keep a counter to count the number of transmissions experienced by this link. Such information facilitates the construction of the inference model as it tells the strength of the dependency between the parent and its successive nodes. Since links in sensor networks are usually shared by multiple paths, we do not need to collect complete path information for all paths before revealing the entire network topology. Indeed, this scheme captures the network topology with a small number of packet receptions, as demonstrated in our field experiment. One potential issue is that when the sink fails to learn the information of some path segments and the network topology is stable, few marks are received. As a result, it will take really a long time for the sink to learn the missing path segments. Such a drawback, however, is alleviated due to the sharing feature of network links, i.e., the missing links can be recovered from other paths that share them. Such a feature definitely alleviates, but does not completely avoid, this problem. To actively
eliminate such a problem, in our implementation we let the intermediate nodes periodically clear their local caches.With this operation, new marks are inserted to packets, and the path information at sink can be periodically refreshed even when the network topology is static.

Algorithm 2 Mark Parsing(packet)

1: **if** *p.sourceNodeID* has no associated *path*
2: create new *path* for *p.sourceNodeID*;
3: **end if**
4: ;
5: **if** //duplicate packet
6: **return**;
7: **else**
8: ;
9: **end if**
10: **if** //no packet loss
11: **if** //route
switch
12: ;
13: clear all slots in *path* after ;
14: generates route switch report;
15: **end if**
16: **else if** //packet loss detected
17: generate packet loss report;
18: **if** //route
switch
19: clear all slots in *path* after ;
20: ;
21: **end if**
22: **end if**

Clearly, in this design we propose to mark simple messages only, but if we insert more marks into the data packets, we obtain richer information on the network statuses and make the diagnosis process more straightforward. Nevertheless, in resource constrained sensor networks, we have to minimize the communication overhead introduced by our diagnosis model. Therefore, we choose to only use simplified marks to additively reconstruct the network. We give details about this issue in later discussions. Compared to existing approaches, our approach with quick reactivity and fast convergence is thus more suitable for highly dynamic environments.

C. Preliminary Diagnosis Reports

Before the final diagnosis results are obtained from the inference engine, some preliminary diagnosis reports can be yielded from the mark parsing module, which help to analyze the network statuses. The preliminary diagnosis briefly infers the following reports. 1) Success delivery report. When the sink receives a packet without any mark, it indicates a successful delivery along the current path. This report tells us that the route from the source sensor node to the sink is still the same and all links along this path have just conducted a successful transmission that confirms the active state of those links.

2) Packet loss report. As described above, if the difference between the sequence number recorded in the *path* and the sequence number of the packet is more than one, it can be inferred that the packet loss occurs. The number of packet loss is quantified as . In this case, according to our marking scheme, the packet must have been marked by some

intermediate node. This report can further locate the packet loss location if there is no route switch accompanying the packet loss.
3) Route Switch Report. The mismatch of the pass node ID in the packet and the recorded ID in the corresponding slot in the *path* indicates that the previous routing path has been altered. The position of the switch is between the hop to source recorded in the packet and its hop upward. V. PROBABILISTIC INFERENCE The packet mark parsing module provides a coarse abstraction and incomplete report.

At the sink, the successive probabilistic inference helps to reveal the inner dependencies among different network elements in the sensor network and expose the hidden root causes of the exterior symptoms. Network elements are inner correlated, for example, the crash of an upstream node causes all its children to disconnect from the sink. In contrast, simultaneous congestion of multiple paths may indicate a high probability of a malfunction at a common link. Based on such observations, we explore the dependencies among network elements (link status, sensing function, path status, etc.) on the constructed communication topology and encode them with a probabilistic model.

Exterior symptoms like delay or loss of data samples are considered as inputs. When specific symptoms are observed by our inference algorithm, we can deduce the probability of the failures of each network element and find the most probable root causes in real time.Most existing inference schemes for static enterprise networks use the simplified bipartite graph or tree-based inference model.
As the network topologies in sensor networks are highly dynamic and no prior knowledge can be acquired in advance, it is difficult to apply the models for static networks in sensor networks. Instead, we apply a hierarchical inference model to capture the inner dependencies in sensor networks. The hierarchical model is good for encoding indirect dependencies with its hierarchical structure and can be constructed without complete information. Also, being assembled by many subparts, it can easily handle the network dynamics efficiently by updating the changed parts only.
We first apply the Belief Network [21] as our inference model. Belief Network is a well-known probabilistic model that has been widely used in research domains like artificial intelligence and system engineering. In Belief Network, each possible root cause or symptom is represented by a variable. Each variable might have multiple values (e.g., 1 for a link in active state and 0 for in trouble). Causal relationships between different variables are denoted as directional arcs. Inferences can be conducted on this model to deduce the probability of particular values to our interested variables once the values of some other variables have been observed (e.g., symptoms like the high delay of data samplings). To further speedup the process, we propose a simplified inference model, Causality Diagram. According to the characteristics of sensor networks, we can design a simplified Causality Diagram that accurately approximates the inference results and reduces the overhead.

A. Belief Network

A Belief Network (or Bayesian Network) is a directed acyclic graph (DAG) that represents a set of variables and their probabilistic relationships. Each vertex in the graph denotes a random variable. In the rest of this paper, we use "vertex" and "variable" interchangeably. A directional arc from vertex to indicates a causal relation between the two variables in which the variable associated with the starting vertex acts as the cause and the variable of is the effect. The cause is called a parent of the outcome .

The strength of the relation between a parent and its child is defined by the conditional probabilities. We then formulate a Belief Network as a binary , where is a DAG and specifies a conditional probability distribution (CPD) in . Here, represents the set of vertices in , and denotes all arcs (or edges). specifies the conditional probability distribution of each variable given its parents. When the value domain of variable is discrete, the CPD can be represented as a conditional probability table (CPT). Given certain evidence (values of some variables), the Belief Network can answer three major types of queries [21]: 1) posterior probability assessment; 2) maximum posterior hypothesis; and 3) most probable explanation. The first type of query, which estimates posterior probabilities of certain variables given some evidence variables, best fits our requirements in this work.

B. Inferring Through Belief Network

Our inference model automatically constructs and maintains a Belief Network from the output of the mark parsing module. The inference engine accordingly infers from this model hidden statuses of the network. In our PAD approach, the Belief Network structure is assembled from the current network topology obtained from the mark parsing module.

1) Constructing a Belief Network:

Fig. 4(a) depicts a simple example topology composed of a sink and three sensor nodes. The directional edge between two nodes denotes a wireless link and the direction of data transmitting along the link. There are five types of variables in our Belief Network, each of which has the value domain of that denotes a normal or abnormal working status, respectively. For each source node, we add a variable to the BeliefNetwork, which denotes the status of the data reception of the

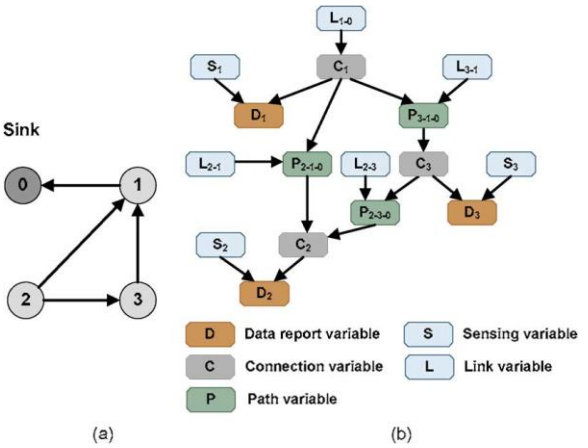

Fig. 4. Belief Network constructed from the communication topology. (a) Network
topology. (b) Belief Network.

source node. For example, if the sink observes a long time delay in the data reception from a source, the corresponding variable of this node will be set to *Down*. Note that, in many applications, some sensor nodes do not sample data, but only relay messages for other nodes. Some of the nodes simply relay packets for other sensors, so there are no data reception variables for those nodes. The status of the data report variable depends on two parent variables, the sensing variable and the connection variable . The sensing variable *Si* indicates the sensing function of the corresponding source node, and the connection variable describes the condition of the network connectivity from the source node to the sink. We add two arcs from and to to represent the dependencies between them. And are thus called the parent variables of in the Belief Network. Both the sensing functionality and the network connectivity condition will affect the success of the data reported from the source node. The connectivity from a source node to the sink relies on one or more paths connecting them. For example, node 2 in Fig. 4(a) can choose to deliver packets through two parents, node 1 and node 3, so in the corresponding Belief Network, the connection variable has two parent variables and . They are called path variables. The subscript of each path variable sequentially denotes the ID of the start node on the path, the ID of the next hop node from the start node, and the ID of the end node on the path. As illustrated in Fig. 4(b), the status of each path variable depends on two parent variables. One is the link variable on the first hop from the start node, and the other is the connection variable of its parent node. The link variable represents the communication conditions of a wireless link between two nodes and . We connect each pair of variables that has a dependency with a directional arc from the cause variable to the outcome variable. Eventually, we obtain a hierarchical network composed of these five types of variables in which dependencies among network elements are encoded. Among the five types of variables, the statuses of the link and sensing variables are hidden from the exterior observations that most need to be inferred. The path and connection variables are intermediate variables that are usually combinational results of other parent variables. The data report variables are outputs of the mark parsing module that we directly observe at the sink. The Belief Network structure consisting of Fig. 5. CPTs of (a) *noisy-OR* and (b) *Select* gates. the variables is automatically maintained and updated when network topology and communication conditions vary over time.

2) Inference on Belief Network:

Once the Belief Network structure is constructed, a critical issue is how to assign CPTs for variables that specify the conditional probabilities between parents and their children. Different logistic relations between parents and their children lead to different methods for calculating the CPT. For example, the sensing variable and connection variable affect their children variable of data report in a logical *OR* manner, i.e., if one of them is in the *Down* state, the data report variable should be switched onto the *Down* state. Due to the diverse routing schemes and high dynamics in sensor networks, a sensor may maintain multiple parents for relaying its data. Consequently, in our inference model, multiple path variables affect the same connection variable in *Select* mode where the status of selected paths will determine the status of the connection variable.

In PAD, we employ the *noisy-OR* gate [21] and *Select* gate [5] to encode these operations. Fig. 5(a) shows the CPT in a *noisy-OR* gate, where any one of the parent variables in *Down* status results in the *Down* status of the child variable. In Fig. 5, and represent the noisy property that means even if both parent variables are in the *Up* status, the child variable still has a chance to fail (in *Down* status). In PAD, *noisy-OR* gates exist in several cases such as when the sensing and connection variables affect the data report variables, the link and connection variables affect the path variables, and so on. The relation between multiple path variables and a connection variable is represented by the *Select* gate as illustrated in Fig. 5(b). Here, denotes the dependency strength of each parent, and in the case of Fig. 5(b), is the probability that the child connection variable selects a certain path to relay data. Thus, the probability that a connection variable is in the *Up* status is given by In the Belief Network, each *noisy-OR* gate connects two parent variables to a child variable, so the CPT calculation is quick. The *Select* gate might connect more parent variables to a child variable, but the maximum number of parent variables for one gate is bounded by the number of neighbors for a sensor node. The number of neighbors is normally treated as a constant. Hence, the CPT calculation for *Select* gate is also

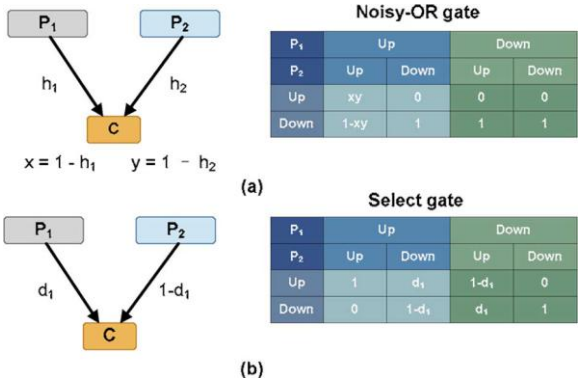

Noisy-OR gate

P_1	Up		Down	
P_2	Up	Down	Up	Down
Up	xy	0	0	0
Down	1-xy	1	1	1

Select gate

P_1	Up		Down	
P_2	Up	Down	Up	Down
Up	1	d_1	$1-d_1$	0
Down	0	$1-d_1$	d_1	1

Fig. 5. CPTs of (a) *noisy-OR* and (b) *Select* gates

the variables is automatically maintained and updated when network topology and communication conditions vary over time. *2) Inference on Belief Network:* Once the Belief Network structure is constructed, a critical issue is how to assign CPTs for variables that specify the conditional probabilities between parents and their children. Different logistic relations between parents and their children lead to different methods for calculating the CPT. For example, the sensing variable and connection variable affect their children variable of data report in a logical *OR* manner, i.e., if one of them is in the *Down* state, the data report variable should be switched onto the *Down* state. Due to the diverse routing schemes and high dynamics in sensor networks, a sensor may maintain multiple parents for relaying its data. Consequently, in our inference model, multiple path variables affect the same connection variable in *Select* mode where the status of selected paths will determine the status of the connection variable. In PAD, we employ the *noisy-OR* gate [21] and *Select* gate [5] to encode these operations. Fig. 5(a) shows the CPT in a *noisy-OR* gate, where any one of the parent variables in *Down* status results in the *Down* status of the child variable. In Fig. 5, and represent the noisy property that means even if both parent variables are in the *Up* status, the child variable still has a chance to fail (in *Down* status). In PAD, *noisy-OR* gates exist in several cases such as when the sensing and connection variables affect the data report variables, the link and connection variables affect the path variables, and so on. The relation between multiple path variables and a connection variable is represented by the *Select* gate as illustrated in Fig. 5(b). Here, denotes the dependency strength of each parent, and in the case of Fig. 5(b), is the probability that the child connection variable selects a certain path to relay data. Thus, the probability that a connection variable is in the *Up* status is given by

In the Belief Network, each *noisy-OR* gate connects two parent variables to a child variable, so the CPT calculation is quick. The *Select* gate might connect more parent variables to a child variable, but the maximum number of parent variables for one gate is bounded by the number of neighbors for a sensor node. The number of neighbors is normally treated as a constant. Hence, the CPT calculation for *Select* gate is also efficient. In the initial stages, the prior fault probability distribution of the link and sensing variables are assigned according to experience data. The value of each is assigned by estimating the percentage of transmissions delivered through each path in a connection. Such information is input from the mark parsing module.

VI. Evaluation

We conduct comprehensive simulations and implement field experiments to evaluate the performance of PAD. For the implementation, we used the BNJ implementation of the Belief Network

inference as part of our inference engine.We implemented the packet marking scheme for TelosB motes on the TinyOS platform with nesC language.We implemented the mark parsing module on the java based back end.

A. Simulations

We first examine the effectiveness and efficiency of PAD through simulations. We simulate a sensor network in which sensor nodes are deployed on a two-dimensional space, with the sink located at the center. Sensors periodically data and deliver to the sink through multihop routes. Two routing schemes are applied in the simulation. Various types of faults are inserted into links or nodes according to different test settings. We use a cutoff threshold to detect the failures. In the following tests, if the output posterior probability of a certain network element (for example a link) to be faulty is higher than 50%, we will regard this element as failure. We apply two metrics for estimating the performances of inference models, the detection ratio and false positive ratio. Detection ratio is the ratio between the number of faults founded and the number of all faults. False positive ratio is the ratio between the real failures detected and all failure reports generated by our diagnosis system.

1) The Efficiency of the Packet Marking Scheme: We evaluate the convergence time of the packet marking scheme under various network conditions. In these tests, we simulate a data acquisition network using both spanning tree-based routing and DAG-based routing schemes. Each source node samples environment data and generates a new packet every 2 s. Different routing schemes lead to different types of topologies. The notation *Tree* denotes a spanning tree topology rooted at the sink, and the notation *DAG* denotes a multipath routing strategy where each sensor node has multiple parents. Besides the routing topologies, the link loss rate also impacts the topology reconstruction. Thus, we evaluate the performance of our approach with different link loss rates. Here, no link loss

indicates that all packet transmissions are guaranteed to deliver, and 10% link loss means that each link has 10% packet loss rate. Under the latter setting (10% link loss), it is difficult for a source node far away from the sink to deliver its packets to the sink since a packet has high probability to get lost on a long path.

VII. Conclusion and Future Work

Although there have been many approaches proposed for debugging the operation of sensor network systems in a controlled laboratory, few works have been done toward an in-situ diagnosis tool for monitoring the statuses of operational systems in the field. In this paper, we propose PAD, a passive diagnosis approach that can be efficiently implemented and applied to a normally working sensor network system providing in-situ network diagnosis. The proposed lightweight packet marking scheme collects necessary hints without injecting extra traffic overhead to the original system. The probabilistic inference model residing at the sink captures the unique features of the sensor networks and yields accurate results. The inference engine works well even with incomplete or suspicious inputs in a nondeterministic manner.We implement our diagnosis approach and validate its effectiveness in a field test in our sea monitoring project. The sea monitoring project is an ongoing project. We are currently utilizing PADas an important diagnosis tool to detect possible faulty components in the system and guarantee its correct operations. On the other hand, we are relying on such a platform to further test the effectiveness and efficiency of PAD and hope to improve it according to our future observations.

REFERENCES

[1] HP Openview, [Online]. Available: http://www.openview.hp.com
[2] IBM Tivoli, [Online]. Available: http://www.ibm.com/software/tivoli
[3] Microsoft Operations Manager, [Online]. Available: http://www.microsoft.com/mom
[4] OceanSense: Sensor Network for Sea Monitoring, [Online].
Available:http://www.cse.ust.hk/~liu/Ocean/index.html
[5] P. Bahl, R. Chandra, A. Greenberg, S. Kandula, D. A. Maltz, and M. Zhang, "Towards highly reliable enterprise network services via inference of multi-level dependencies," in Proc. ACM SIGCOMM, 2007, pp. 13–24.
[6] X. Bai, D. Xuan, Z. Yun, T. H. Lai, and W. Jia, "Complete optimal deployment patterns for full-coverage and k-connectivity □_ □_ __ wireless sensor networks," in Proc. ACMMobiHoc, 2008, pp. 401–410.
[7] J. Cao, L. Zhang, J. Yang, and S. K. Das, "A reliable mobile agent communication protocol," in Proc. IEEE ICDCS, 2004, pp. 468–475. [8] G. F. Cooper, "Probabilistic inference using belief networks is NP-hard," Stanford Knowledge Systems Laboratory, Tech. Rep., 1987.
[9] P. Dagum and M. Luby, "Approximately probabilistic reasoning in Bayesian belief networks is NP-hard," Artif. Intell., pp. 141–153, 1993. [10] Q. Fang, J. Gao, and L. J. Guibas, "Locating and bypassing routing holes in sensor networks," in Proc. IEEE INFOCOM, 2004, vol. 4, pp. 2458–2468.
[11] R. K. Ganti, P. Jayachandran, H. Luo, and T. F. Abdelzaher, "Datalink streaming in wireless sensor networks," in Proc. ACM SenSys, 2006, pp. 209–222.
[12] B. Gedik, L. Liu, and P. Yu, "ASAP: An adaptive sampling approachto data collection in sensor networks," IEEE Trans. Parallel Distrib.Syst., vol. 18, no. 12, pp. 1766–1783, Dec. 2007.
[13] S. Guo, Z. Zhong, and T. He, "FIND: Faulty node detection for wirelesssensor networks," in Proc. ACM SenSys, 2009, pp. 253–266.
[14] T. He, S. Krishnamurthy, J. A. Stankovic, T. Abdelzaher, L. Luo, R.Stoleru, T. Yan, L. Gu, J. Hui, and B. Krogh, "Energy-efficient surveillancesystem using wireless sensor networks," in Proc. ACM MobiSys,2004, pp. 270–283.
[15] S. Kandula, D. Katabi, and J.-P. Vasseur, "Shrink: A tool for failurediagnosis in IP networks," in Proc. MineNet, 2005, pp. 173–178.
[16] K. Klues, G. Hackmann, O. Chipara, and C. Lu, "A component-basedarchitecture for power-efficient media access control in wireless sensornetworks," in Proc. ACM SenSys, 2007, pp. 59–72.
[17] R. R. Kompella, J. Yates, A. Greenberg, and A. C. Snoeren, "IP fault localizationvia risk modeling," in Proc. USENIX NSDI, 2005, pp. 57–70.[18] S. Lim, C. Yu, and C. R. Das, "Rcast: A randomized communicationscheme for improving energy efficiency in MANETs," in Proc. IEEE

ICDCS, 2005, pp. 123–132.[19] H. Liu, P. Wan, C.-W. Yi, X. Jia, S. A. M. Makki, and N. Pissinou,"Maximal lifetime scheduling in sensor surveillance networks," inProc. IEEE INFOCOM, 2005, vol. 4, pp. 2482–2491.
[20] Y. Liu, Q. Zhang, and L. Ni, "Opportunity-based topology control inwireless sensor networks," in Proc. IEEE ICDCS, 2008, pp. 421–428.
[21] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks ofPlausible Inference. San Mateo, CA: Morgan Kaufmann, 1988.