2012 Ninth European Dependable Computing Conference

Security of Electrostatic Field Persistent Routing: Attacks and Defense Mechanisms

Oliviu C. Ghica
Dept. of EECS
Northwestern University
Evanston, USA
oliver@northwestern.edu

Cristina Nita-Rotaru
Dept. of CS
Purdue University
West Lafayette, USA
crisn@cs.purdue.edu

Goce Trajcevski
Dept. of EECS
Northwestern University
Evanston, USA
g-trajcevski@northwestern.edu

Peter Scheuermann
Dept. of EECS
Northwestern University
Evanston, USA
pscheuer2008@gmail.com

Abstract—**Electrostatic Field-based Routing (EFR) is a form of geographical multi-path routing where packets are routed along a collection of electrostatic field lines, defined by electrostatic charges associated with source and sink nodes. EFR provides an efficient and scalable solution to the workload balancing problem. However, it assumes that nodes behave in a cooperative manner making EFR-based routing protocols vulnerable to various attacks.**

We investigate the security aspects of EFR-based routing protocols. We focus on an instance of EFR, called Multi-Pole Field Persistent Routing (MP-FPR), for which we identify attacks that can target different components of the protocol, and propose a set of corresponding defense mechanisms. We present extensive experimental evaluations of the impact of the different attacks and the effectiveness of the proposed defense mechanisms.

I. Introduction

Wireless Sensor Networks (WSNs) [1] have emerged as a promising paradigm for many application domains that require combined capabilities of sensing, processing, and communication in different physical environments. In a typical WSN application, a user-initiated query is disseminated to the appropriate *source* nodes where the data of interest is locally collected. The resulting point-to-point data-stream is relayed back to a remote *sink* node which, in turn, interfaces with the user. Many routing protocols for WSN are designed under the location-aware assumption and rely on the *geography*-based (greedy) routing principle, according to which packets are forwarded to nodes that are physically closer to a given destination [2]. A type of geographic routing is trajectory based forwarding (TBF) [3], in which packets are routed towards the destinations along pre-defined ”virtual” trajectories.

Electrostatic Field-based Routing (EFR) [4] is a multi-path routing protocol that reduces the complexity of determining and managing the collection of underlying trajectories by representing them as electrostatic field lines, rather than relying on geometric models. The field lines originate at source nodes, where the data is produced, and lead towards a designated sink node, where the data is being consumed. The main advantage of EFR is that it creates *implicitly* spatially *disjoint trajectories* which in turn allows to achieve workload balancing in dense and uniformly distributed networks. In networks where this assumption does not hold, path-merging can occur reducing the workload balancing capabilities. Multi-Pole Field Persistent Routing (MP-FPR) protocol [5] extends EFR's applicability to less-dense and often non-uniform network distributions by actively seeking to separate any merged paths whenever network conditions allow.

MP-FPR assumes that nodes in the network always operate correctly, assumption no longer valid when MP-FPR is deployed in an adversarial environment. As many applications for WSNs require deployment in an adversarial environment, it is critical to provide mechanisms to ensure that MP-FPR operates correctly and securely.

In this article we analyze the security of the MP-FPR protocol and identify the main *data*- and *control*-level components that can be exploited by an attacker. We study disruptions to users' data streams and the system-wide performance and resource-utilization, such as the disruption of *workload-balancing*. Our main contributions are:

- We identify a set of attacks in MP-FPR and assess their impact on the entire system. We first identify a set of *control-level* attacks: *path deflection*, *path diversity deflation*, *family path intersection*, *wild-path* and *field-line hopping*, all of which are specific to electrostatic-field based routing. These attacks are carried through control messages, and can lead to quality of service degradation by disrupting the workload-balancing operation. We next identify a set of *data-level* attacks: *data denial of service* (DoS), *data pollution*, and *data stream invalidation*, which directly target users' data streams.
- We evaluate the resilience of MP-FPR to adversarial scenarios and observe the epidemic character of several attacks that can yield significant performance degradation with minimal staging efforts. For example, a *single* attack consisting of inserting eight forged charges in the system via a sink node can nearly double the standard deviation of the residual energy levels – a representative metric for describing the workload balancing performance.
- We propose two classes of defense mechanisms that target basic attacker operations. We identify TESLA [6] as a viable solution for integrity and authentication of the MP-FPR messages, and propose a set of mechanisms to provide resilience against selective forwarding of various protocol messages, denoted as k-EF, k-RPEF and PDMS. The first two mechanisms, k-EF and k-RPEF, rely on multi-path in the electrostatic context, while the third one, PDMS, is a complementary monitoring scheme to provide closed-loop control over path diversity. We report

978-0-7695-4671-1/12 $26.00 © 2012 IEEE
DOI 10.1109/EDCC.2012.23

IEEE computer society

the quantitative observations regarding the effectiveness of the proposed approaches based on an extensive set of experimental evaluations.

Outline. The rest of the article is organized as follows. In Section II we overview the main aspects of the MP-FPR multipath routing protocols. The details of the adversarial model are presented in Section III. Resilience mechanisms against attacks carried through selective message forwarding are presented in Section IV. The results of our experimental investigation are shown in Section V. We overview the related work in Section VI and conclude the article in Section VII.

II. Multi-Pole Field Persistent Routing

Below we present a detailed description of the MP-FPR protocol. We assume that a network consists of a set **SN** = $\{sn_1, sn_2, \ldots, sn_n\}$ of n wireless sensor nodes, each capable of acting both as a *relay* and a *source* of sensed data.

A. Overview and Forwarding Mechanisms

MP-FPR is based of on the EFR routing protocol. EFR routes along trajectories represented via *electrostatic field lines* originating at source nodes, which are assigned a ”positive charge”, and ending at sink nodes which are assigned a ”negative charge”. In order to route a packet to the sink, a relay-node needs to know its own location, as well as the location and the electrostatic charge information of the source and sink nodes. Permanent path deviations may occur when a given relay node cannot find subsequent relay node(s) that are along or in the vicinity of a particular electrostatic field line. Thus, adjacent paths may intersect and/or merge, resulting in overloading a subset of the downstream relay nodes.

MP-FPR recreates spatially disjoint routes via splitting previously merged routes (see Figures 1(b) and 1(c)). Unlike EFR where packets are routed only along field lines that the current relay node resides on, in MP-FPR packets are routed mainly along the original field line from which a packet may have been diverted. The *original* field line is piggy-backed on data-packets. MP-FPR forwards messages using two mechanisms: Electrostatic Field (EF) forwarding which relies on electrostatic fields and Shortest Geographical Path (SGP) which is a greedy based geographical routing.

EF forwarding: This routing is based on a discrete subset of field lines between a given *(source, sink)* pair. We refer to this set S_f as a family of paths and show an example in Figure 1(a). Each field line in S_f is uniquely identified by the value of the angle φ_j, determined by the *tangent* to a given/chosen field line at the source, and the line segment between the source and the sink . Assuming a uniform selection of the tangential-angle from $[0, 2\pi]$, a particular field line φ_j can be chosen from a family $S_f = \{k\frac{2\pi}{N_r} \mid k = \overline{1, N_r}\}$, where N_r is the desired cardinality of the family of routes S_f. Each route built along a field line φ_j is uniquely identified by a route index, denoted r_j. For simplicity, we assume $r_j = \varphi_j$.

Every node sn_i in the network can determine the tangent angle $\varphi_j \in S_f$ of the field line that it *actually* belongs to based on the (1) location and charge information of the source(s), (2) location and charge information of the sink, and (3) its own location. Once sn_i receives a packet, it piggy-backs the information about the field line that the packet is *supposed* to be forwarded along, i.e. φ_j. A particular relay node will select, as its subsequent relay node, one of its 1-hop neighbors which exhibits the smallest field line deviation $|\varphi_j - \varphi_i|$, where φ_i represents the actual field line a downstream relay sn_i actually resides on, and it is furthest away towards the sink (cf. [5]).

SGP forwarding: This is a greedy geographic routing similar to BVR [7], where packets are sent via a geographically shortest path towards a known physical destination. In MP-FPR nodes determine their own position via a lightweight localization service external to the routing protocol (see [8] for a survey), as well as the position of their 1-hop neighbors through a periodic location information exchange.

MP-FPR protocol consists of the following components: *query dissemination and charge allocation*, *route establishment*, and *data forwarding*. We overview each component and summarize the messages used by the protocol in Table I.

B. Query Dissemination and Charge Allocation

This component has several goals. First is to forward the user query towards the source node. This is achieved through a QUERY message, which is sent by the sink via SGP forwarding towards L_{src} – the location within the area where data relevant to the query should be collected from. A sensor node which is closest to L_{src} will assume the role of the source for the given QUERY message and initiate its processing. Second goal is to disseminate electrostatic charges information, which consists of a set of (location, magnitude and expiration) information associated with each routing end-point, i.e. source or sink node, in the network. For example, if there are m source nodes relaying data-streams to a common sink, the QUERY message contains a set $C_e = \{e_{snk}\} \cup \{e_i | i \in \overline{1, m}\}$ of electrostatic charges. Third goal is to limit the number of alternative paths to be built in order to correspondingly bound the duration of the route establishment via a numerical parameter $N_r = |S_f|$ embedded in the body of the QUERY message. We refer to this limit as the *path diversity quota*.

Whenever a new data source is added to the existing set of source-nodes, a new corresponding charge is added to the virtual electrostatic field. The charge information is updated at each of the source nodes via an UPDATE message. Upon receiving an UPDATE, the route establishing process is re-initiated by the source nodes in order to establish new families of routes that are consistent with the new charge distribution.

C. Route Establishment and Data Forwarding

Initiated upon receiving a QUERY or an UPDATE at a source, *route establishment* is a two-phase request-acknowledgment process. During the *request phase*, the source transmits a set of RREQ messages along *distinct* electrostatic field lines towards the sink. A RREQ message carries a list of network's current charges C_e as well as the field line index (equivalently route index) $r_i \in S_f$ of the field line a specific

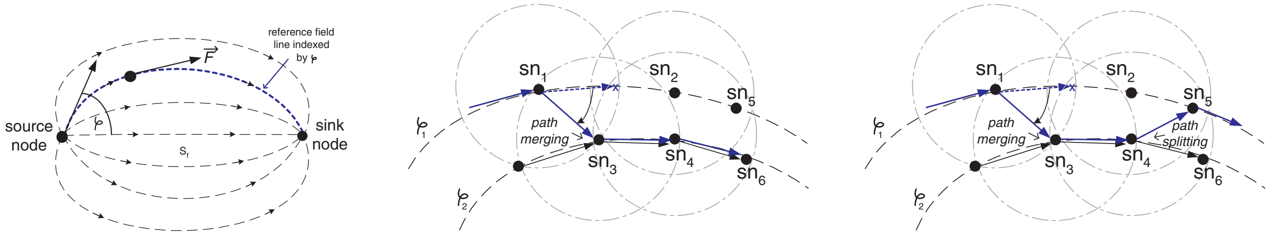

(a) Field line selection

(b) Permanent merging in EFR

(c) Field-persistency in MP-FPR

Fig. 1. MP-FPR mechanism. (a) Family of multiple field lines between a source and a sink: selection of an arbitrary angle φ and associated incidental reference field line that is followed by the corresponding indexed route; (b) Path merging: node sn_1 is unable to reach node sn_2 and redirects the route to node sn_3, which is already servicing another route r_2 associated to field line φ_2 (c) Un-merging previously merged paths: node sn_4 redirects the route r_1 (that went through sn_1) to sn_5 to resume routing along φ_1

TABLE I
MP-FPR MESSAGES

Type	Flow	Functionality	Protocol Phase	Forwarding	Fields of Interest
QUERY	Sink → Sources	Query Specification Wrapper	Query Dissemination and Charge Allocation	SGP	L_{src}, C_e, N_r,
UPDATE	Sink → Sources	Charge Information Update	Query Dissemination and Charge Allocation	SGP	L_{src}, C_e
RREQ	Sources → Sink	Route Request (Probe)	Route Establishment	EF	L_{src}, C_e, r_i, t_{sent}
ACK	Sink → Sources	Route Acknowledgment	Route Establishment	SGP	L_{src}, r_i
DATA	Sources → Sink	User Data-Payload Wrapper	Data Forwarding	EF	r_i, $Data$

RREQ is to be sent along. To amortize the transmission cost of the charges, this information is sent only once along RREQ messages, and cached locally by the relay nodes along a route; subsequent DATA messages will not carry them. The source will also incorporate its actual location information L_{src} in the RREQ message such that sinks maintain a more accurate representation of the actual sources. A timestamp t_{sent} is included in the RREQ message to assist in determining the quality (e.g. latency) of a specific route.

If, upon receiving a RREQ message, the route exhibited an admissible latency, the route is acknowledged via an ACK message to the specific source. The route index r_i of that route is included in the ACK message. Note that ACK messages are sent back via the SGP mechanism towards the actual location of the source L_{src}, and not via EF mechanism the RREQ message was sent. Every acknowledged route is added to a source-maintained set of acknowledged routes $S_f^{ack} \subseteq S_f$, i.e., a pool of routes that are available for data forwarding.

Data forwarding: The DATA messages pertaining to a data-stream as a result of query processing are forwarded back to the sink node via the EF mechanism, by using individual routes r_i from the set of acknowledged routes S_f^{ack}.

III. Taxonomy of Attacks

In this section we identify a representative set of attacks that can be carried against the MP-FPR protocol and exploit vulnerabilities introduced by the use of electrostatic field lines and by the field persistency mechanism. MP-FPR has two main system goals that can be compromised by attacks: (1) increase network lifetime by promoting delivery of the data stream in a workload balanced manner and (2) ensure soft QoS guarantees, such as bounded end-to-end data stream delivery latencies.

A. Adversarial Model

We assume that the only trusted nodes in the network are the sink and the source. Honest nodes participate correctly in the routing protocol, whereas malicious nodes acting alone or in collusion can drop, delay, modify or replay packets.

We assume the forwarding mechanisms used by MP-FPR, EF and SGP, are not secure. However, since previous work examined the security of SGP [9], we focus mainly on the security of EF. Both EF and SGP rely on a localization service. We assume security mechanisms [10], [11] are in place to protect the localization service. Similarly, we assume that the time synchronization mechanism is also secure [12], [13].

B. Attack Classification

We classify the attacks as *data-level* and *control-level* based on their target, the user-data or the network operation, respectively. An attacker can drop, delay, or modify any of the five type of messages the MP-FPR protocol relies on: QUERY, UPDATE, RREQ, ACK, and DATA. We do not consider replay-attacks as they can be easily addressed by using packet sequencing or timestamps. Table II summarizes the main attacks that MP-FPR protocol is susceptible to.

C. Query Dissemination and Charge Allocation Attacks

Attacks during the query dissemination and charge allocation protocol phase can be mounted by targeting the QUERY and UPDATE messages (see Table I). We identify the following attacks: *data DoS*, *data stream invalidation*, *path diversity deflation*, *path deflection*, and *family path intersection*.

Data DoS. This attack can be easily mounted by maliciously dropping QUERY messages and disrupting the delivery of users' data-flow. The absence of the entire data-stream can be easily detected and thus the underlying attack unveiled.

TABLE II
ATTACK TAXONOMY

Category	Attack	Drop	Delay	Modify	Primary Impact
Control Level	Path Deflection	-	-	QUERY(C_e), UPDATE(C_e)	Workload balancing
	Path Diversity Deflation	RREQ, ACK	RREQ	QUERY (N_r), ACK(r_i, L_{src}), RREQ(r_i, L_{src}, t_{sent})	– ” –
	Family Path Intersection	UPDATE	UPDATE	QUERY(C_e), UPDATE(L_{src}, C_e)	– ” –
	Wild Path	-	-	RREQ(C_e)	– ” –
	Field-Line Hopping	-	-	RREQ(r_i), DATA(r_i)	– ” –
Data Level	Data DoS	QUERY, DATA, ACK	DATA	DATA(r_i), ACK(r_i)	Data delivery reliability
	Data Pollution	-	-	DATA(payload)	Data integrity
	Data Stream Invalidation	-	-	QUERY(L_{src})	Data validity

Data Stream Invalidation. An attacker can alter the parameters of a user-submitted query, such as the L_{src} parameter in the body of the QUERY message. As a result, the user will receive an invalid data-stream. This is a stealthy attack, as opposed to the *data DoS* attack, the user does receive an uninterrupted data-stream; however, the user may not be aware that it is not the data that he requested.

Path Diversity Deflation. This attack targets the load-balancing by reducing the number of alternate paths that the protocol can use. This number is bounded by a parameter, N_r, included in the QUERY message. Decreasing N_r reduces path diversity. If N_r is maliciously set to 1, MP-FPR will effectively degrade to *single*-path routing. *Path diversity deflation* may not have an immediate, noticeable impact to the user, however, its damaging effect is visible through a significant reduction of network's lifetime.

Path Deflection. The outcome of this attack consists of a geographical shift of the existing families of routes, or constraining the field-region in which routes can be built. This attack can be conducted by modifying charge information in either the QUERY or UPDATE messages. For example, an attacker can modify the magnitude of a particular charge, or introduce new ”fake” charges in the system. Altering the magnitude of a charge will affect the load-balancing among distinct families of routes. In extreme cases, it is possible to narrow the admissible relay field so much that most of the paths will merge, leading to a single-path routing behavior, which is the equivalent of a *path diversity deflation* attack. Adding one fake charge may result in a geographical shift of the existing families of routes, possible leading to increased routes' lengths, with a consequent increase of the end-to-end delivery latencies. Figure 2(b) presents a family-path geographical shift as a result of one fake charge.

Family Path Intersection. This attack targets the disjointness of the routing paths from the same family, as well as paths from distinct families. The attack can be mounted by either dropping UPDATE messages or by modifying the L_{src} parameter in the UPDATE message. Some of the conditions that lead to a path deflection may also create intersection between routes pertaining to different families if charge information becomes inconsistent among families. Paths pertaining to the same family of routes will continue to maintain the non-intersection property among themselves, however, distinct families of routes will cross each-others geographical bounds. Such intersections create resource utilization hot-spots with direct consequences on the overall network's lifetime. Figure 2(c) illustrates a family path intersection attack.

D. Attacks during Route Establishment

The attacks that can be carried during this phase are qualified as *control-level*, as they target RREQ and ACK control messages respectively. We analyze attacker strategies and identify, in addition to the previously described *data DoS* and *path diversity deflation* attacks, new attacks: *wild path* and *field line hopping*.

Path Diversity Deflation. Dropping either RREQ or ACK messages may result in an overall reduction of the route content within a family of routes. Since paths are designed to spread through a larger network-area for workload balancing purposes, an attacker can target an arbitrary node, without a priori insider information. Additionally, delaying RREQ messages or altering the embedded source-transmission timestamp t_{sent} may increase the latency beyond a user-defined tolerance. Changing the source location information L_{src} in the RREQ or ACK messages will cause ACK messages to be delivered to a node different than the source. Lastly, altering route index information r_i in the RREQ or ACK messages can also lead to the same outcome. For example, in either case, the (corresponding) acknowledgment will acknowledge an arbitrary route, which may have been already acknowledged, while the intended route will be dropped from usage. Figure 2(e) presents an example of a *path diversity deflation* attack where route r_0's acknowledgment is never received by the originating source node. All of these conditions can ultimately lead to diminished energy consumption balancing performance.

Data DoS. This attack can be mounted by targeting ACK messages. During the *route acknowledgment* phase, compromising ACK messages vs. RREQ messages can lead to different effects, because distinct forwarding mechanisms handle the two types of messages: ACKs are sent via a single-path (SGP), whereas RREQ via EF. Therefore, if a single node along the SGP path is compromised, *all* ACK messages can be compromised or dropped. Since path diversity can be effectively reduced to zero, the user's data-stream will be completely blocked. Alternatively, a malicious node may alter the route index r_i in the ACK message. In this case, an arbitrary route will receive an acknowledgment, possibly one that was not probed or one that may not satisfy user deliverability requirements, such as end-to-end delivery latency.

Wild Path. This attack makes a route from a family of routes to break the disjointness property of electrostatic field lines and start intersecting other routes. There are two

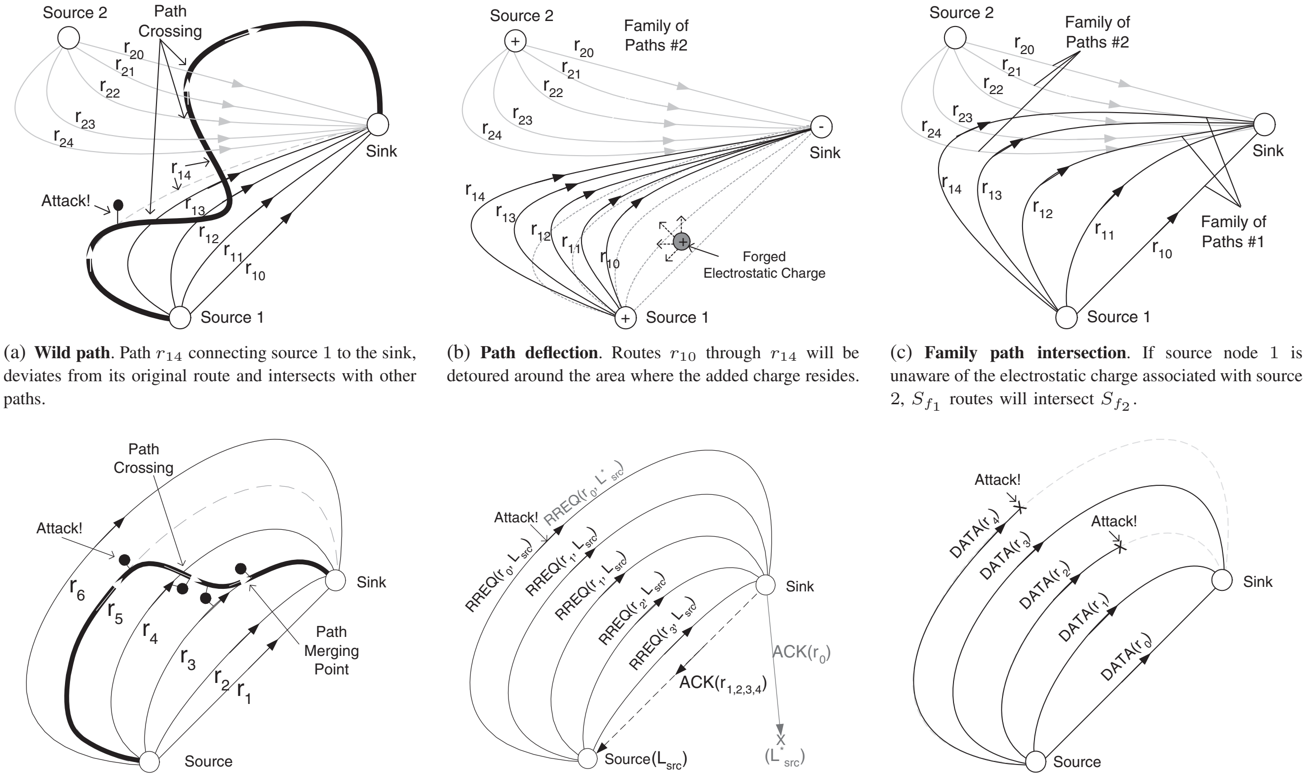

(a) **Wild path**. Path r_{14} connecting source 1 to the sink, deviates from its original route and intersects with other paths.

(b) **Path deflection**. Routes r_{10} through r_{14} will be detoured around the area where the added charge resides.

(c) **Family path intersection**. If source node 1 is unaware of the electrostatic charge associated with source 2, S_{f_1} routes will intersect S_{f_2}.

(d) **Field line hopping**. If the identification information r_5 is maliciously changed to r_3, the original path r_5 deviates, intersecting r_4 before it merges with r_3, doubling the load downstream from the merging point on r_3.

(e) **Path diversity deflation**. Path r_0 will never be used for data-forwarding since its never acknowledged: its corresponding ACK message never reaches the source located at L_{src}.

(f) **Data DoS**. Selective forwarding of DATA messages sent along routes r_2 and r_4 nullifies those routes; user will receive an incomplete data-stream at the sink node.

Fig. 2. Examples of attacks against the MP-FPR protocol

important differences from the *family path intersection* attack: (1) a wild path attack targets a single route, rather than an entire family of routes, and (2) the compromised route intersects not only other routes within the same family, but also routes pertaining to other families. This attack is carried via altering charge information within a relay node along a particular route. Recall that charge information transmitted via RREQ messages are cached by the relay nodes for subsequent use. Consequently, the attack can be carried by altering the RREQ messages before their contents are cached. The entire path downstream of the compromised node will exhibit an abrupt deviation from the designated field line. Figure 2(a) illustrates an instance of a *wild path* attack.

Field Line Hopping. Consider a route indexed by r_j, which is built along a reference field line φ_j. If the route index from in the RREQ message is altered, the original route will suddenly change its reference field line and ”hop” to a different one within the same family. The immediate consequence is path intersection or merging. This situation is different from a *wild path* situation, because field lines do not change; rather the actual route changes field lines. Figure 2(d) shows an example of *field line hopping* attack. Field line hopping creates relay node overload, resulting in degraded energy consumption balancing and reduction of lifetime expectancy.

E. Data Forwarding

DATA messages carry the information-load resulting from processing a user-submitted query. Since DATA messages follow probed and acknowledged paths, they are susceptible to the same attacks as those against RREQ messages.

Data DoS. This attack blocks a user data-stream. It can be mounted by selectively dropping DATA messages along a path, i.e. if one of the relay nodes along the path is compromised. Figure 2(f) illustrates this scenario, in which two different compromised nodes along different routes drop all incoming DATA messages, effectively nullifying those paths. In some instances, altering the route index information r_i in the DATA messages, which can redirect the message along non-probed and possibly long paths, or simply delaying these messages, may similarly lead to a *data DoS* attack. In both cases, it is likely that the message will be discarded at the sink node if not received within certain admissible delay tolerances.

Field Line Hopping. Analogous to attacks carried through RREQ messages, DATA messages can be maliciously ”re-routed” along different routes than the originally prescribed ones, resulting in path merging and overloading of some of the downstream relay nodes. The net effect consists of energy consumption balancing disruption and a reduction of network's lifetime. This attack can be achieved by modifying the route index r_i embedded in the DATA message.

Data Pollution. The attacker may directly alter the user-payload within the DATA message itself. This attack can be severe, since the user may not be able to distinguish valid data from faux, and it may require advanced data analysis to detect anomalies in the data-stream.

IV. Defense Against Attacks

There are two main fundamental causes of the identified attacks: (1) the lack of message *authentication* and *integrity* mechanisms, and (2) the lack of a robust delivery mechanism resilient to malicious message dropping. Authentication and integrity can be provided with existent cryptographical approaches such as symmetric key-based *HMAC*[14], public-key based *digital signatures* [15], [16], or a hybrid solution like *TESLA* [6]. We conducted a performance and overhead analysis and concluded that TESLA is feasible across broader real sensor platforms, efficient and computationally tractable. For lack of space we could not include the results, a detailed description is presented in [17].

Below we present three new defenses against selective data dropping: k-EF, k-RPEF, and Path Diversity Monitoring Scheme (PDMS). These solutions are also applicable to selective *delaying*: k-EF provides resilience against delaying DATA messages, k-RPEF addresses the delays of QUERY, ACK and UPDATE messages, while PDMS addresses the delaying of RREQ messages.

A. Our Approach

MP-FPR uses five type of messages sent via two forwarding mechanisms, EF and SGP. Consequences of attacks carried through selective forwarding of the MP-FPR protocol messages are presented in Table II.

We propose a proactive defense mechanism that uses replication of outgoing messages in order to improve resilience to adversarial activities. The solution aims at providing redundancy in the forwarding mechanism. Instead of one message, a number of k-copies of a certain message may be sent along k-distinct routes, significantly reducing the probability that an attacker will successfully manage to drop all k such copies. We refer to the parameter k as the *degree of replication*. This approach is appealing because the required underlying support, i.e. multi-path routing, is readily available in MP-FPR and thus requires minimal changes.

Both source-to-sink and sink-to-source traffic must be augmented with resilient forwarding mechanisms. The source-to-sink traffic consists of DATA messages, for which resilient forwarding can be easily provided: these messages can be sent along subsets of already constructed routes. We refer to this mechanism as *k-EF*. Note that these subsets of routes continue to be used in alternation for workload balancing purposes.

Sink-to-source, reverse-traffic, comprises QUERY, UPDATE and ACK messages. The challenge here is that these messages rely on SGP forwarding mechanism and no routes are readily available as in the EF mechanism. There are two possible solutions that can be considered to provide k-resilience to reverse-path selective forwarding in MP-FPR: (1) replacement of the standard SGP mechanism with a k-shortest path routing [18] (which we refer to as k-SGP), and (2) adapt MP-FPR protocol to rely directly on the field-based forwarding provided by EF to forward copies along multiple *on-the-fly* built routes, which we will refer to as k-RPEF (Reverse Path Electrostatic Forwarding). We adopt the secondary approach, i.e. k-RPEF, for the following three reasons: (1) it is relatively easy to implement since it relies on the same forwarding mechanism as in EF, (2) it simplifies the network-protocol stack by removing the SGP component altogether, and (3) its redundant paths inherit the non-braiding property of field-based routing, which cannot be guaranteed with k-SGP.

RREQ messages, however, cannot benefit from a redundancy mechanism, since RREQ messages are bound to the route they probe and implicitly construct, i.e. copies of RREQ messages cannot be sent on different routes. For RREQ messages, we propose a reactive mechanism, namely the Path Diversity Monitoring Scheme (PDMS). This monitoring scheme reactively attempts to compensate for any deficiencies in path diversity by persisting in building more routes until the user defined path diversity quota is met.

B. k-EF Defense Mechanism

The k-EF mechanism provides replication of DATA messages using the set of active routes resulting from the route establishment phase. The degree of replication is given by the value of $k \leq N_r$, where N_r represents the maximum number of routes that can be established. We use a random selection scheme to select k paths from the total of N_r possible. We remind that the N_r routes are uniquely identified via a route index $r_i \in \varphi_{N_r} = \{1\frac{2\pi}{N_r}, 2\frac{2\pi}{N_r} \ldots N_r\frac{2\pi}{N_r}\}$, i.e. equally distributed across the $\varphi \in (0 \ldots 2\pi]$ domain, hence in a k-redundant scheme, the indexes of the k routes should be randomly picked from the φ_{N_r} set without replacement.

C. k-RPEF Defense Mechanism

k-RPEF provides redundant forwarding of QUERY, UPDATE and ACK messages towards the source nodes. Forwarding is still based on electrostatic field lines, but traverse in opposite direction of the field vectors, towards the source. In order to enable reverse electrostatic field lines traversal we reverse the algebraic sign of the charge's magnitudes corresponding to the sink and specific source charges for reverse path forwarding only. For example, if a sink and a source have charges of $Q_{src} = -1 \cdot 10^{-19}$ coulombs and $Q_{snk} = +1 \cdot 10^{-19}$ coulombs respectively, k-RPEF's field lines will be built on the set of charges $Q_{src} = +1 \cdot 10^{-19}$ coulombs and $Q_{snk} = -1 \cdot 10^{-19}$ coulombs instead. We note here that only the source's charge towards which we intend of forwarding the message gets the magnitude reversed, whereas other source nodes remain unchanged – this is required in order to prevent messages from reaching other source nodes by hopping on their field lines. Also, the algebraic magnitude's sign reversal is performed in isolation from other sources, i.e. such information is not broadcasted and it is only used locally. Charge magnitude reversal forces the field line vectors to point towards the source node rather than the sink, guiding the associated routes accordingly, without further modification of the forwarding algorithm.

D. Path Diversity Monitoring Scheme (PDMS)

Dropping of RREQ messages critically affects path diversity and, consequently, the energy balancing. Although the k-RPEF mechanism addresses the path diversity deflation problem from the perspective of attacks against ACK messages, it cannot be used for attacks against RREQ messages, because RREQ messages are uniquely associated to the routes they are forwarded through, hence replicas of a RREQ message cannot follow a different route. The idea in PDMS is to enable the source node to persist in probing for new routes until the user-specified *path diversity quota*, i.e. number of distinct routes N_r the user demands, is being met. PDMS relies on the observation that distinct routes will map to distinct sets of nodes, hence bypassing of compromised nodes can be achieved in subsequent attempts.

PDMS cannot be used as a standalone solution for path diversity deflation attacks carried out via ACK messages, for the following reason. Recall that, in the absence of k-RPEF mechanism, ACK messages are sent via SGP forwarding, therefore compromising the single reverse path will block the acknowledgment phase completely. In this case, regardless of the number of attempted routes to be built, routes will never get acknowledged. PDMS, however, can provide *compensatory* benefits if the k-RPEF resilient mechanism is already employed for ACK messages, and our experimental results will demonstrate this benefit.

MP-FPR tries to evenly distribute the workload by building evenly distributed routes in the physical field. We seek to either maintain this property or gracefully degrade it under adversarial conditions. Thus, the *sequence* of routes that will be probed must take into consideration the existing distribution of routes and attempt to fill any existing "gaps". Recall that MP-FPR adopts an *angular model* for route-indexing cf. Section II-A. Consequently, we rely on the assumption that the distribution of the routes indexes is representative for the distribution of the actual routes.

We propose the PDMS mechanism as a multi-phase process. The first construction phase performs the same functions as in the original MP-FPR protocol, namely a *sequence* S_1 of N_r evenly distributed route indexes are generated and iteratively probed, $S_1 = \langle r_i | r_i = \frac{2\pi}{N_r} i, i \in \overline{1, N_r} \rangle$. If the path diversity quota is not met during the first phase, subsequent construction phases are invoked. The followings apply to every phase $j \geq 1$. We refer to S_j as the *base routing sequence of phase* j. Let A_j be the *set* of *active* routes that have been successfully acknowledged up to phase j. If and only if the path diversity quota is not being met at a certain phase j, i.e. $|A_j| < |N_r|$, a subsequent phase $j+1$ is initiated. In each subsequent phase $j > 1$, a new distinct sequence S_j is being generated such that $|S_j| = N_r$ (the generation method will be addressed shortly). As opposed to the very first phase however, not all routes in S_j need to be probed, and the probing process can be interrupted at any time if the path diversity quota is being met. To prevent wasteful energy resources under severe adversarial conditions, we limit the number of phases that can be executed to a predefined value $K \geq 2$.

The base routing sequence at phase $j > 1$ is generated as a counter-clockwise rotation of the base sequence of angular indexed routes from previous phase, i.e. all route indexes from current phase are obtained by incrementing the route indexes of the previous phase by a fixed amount δ. Considering the maximum number of admissible probing phases K, in the worst case scenario, the union of all base routing sequences is $\bigcup_{j=1}^{j=K} S_j = \langle r_i | r_i = \frac{2\pi}{K \cdot N_r} i, i \in \overline{1, N_r} \rangle$, hence a total of $N_p = K \cdot N_r$ distinct and evenly distributed routes may be probed by PDMS.

In order for the PDMS to ensure even distribution of the resulting routes, the base routing sequence generation mechanism is necessary, but not sufficient. Namely, since subsequent route construction phases can be terminated immediately when path diversity quota is being met, priority must be given to routes situated in the vicinity of a failed route, whose omissions has created a "gap". The intuition is as follows: if originally the base routes led to evenly distributed routes with the exception of one route, it is desirable to build a replacement route as close as possible to the original failing one, such that the deviation from the targeted distribution is minimized. This motivates the phased generation of the base sequence of routes, where δ represents the deviation added to the routes from original location.

The advantage of the proposed PDMS scheme versus a purely random one, in which route indexes are randomly, with uniform distribution, generated, is twofold: (1) PDMS maintains full control of the probed routes by primarily targeting areas with lower densities of routes (i.e. in *immediate* vicinity of failed routes) to improve route distribution, and (2) it avoids route merging effects caused by new routes that may be randomly chosen "too close" to existing ones by guaranteeing a minimum path-spacing through δ.

The prioritization mechanism that is applied to the base routing sequence of phase j, $\overline{S}_j$, is summarized as follows. The key idea is to determine the angular-gap size between any two adjacent route indexes from the ordered set of active routes A_j, and store these gaps' information in an ordered set G_j in descending order of the gap-size. Given a base routing sequence $\overline{S}_j$, we reorder the sequence such that the i^{th} element in $\overline{S}_j$ is situated within the bounds of the i^{th} gap in G_j.

V. Experimental Evaluation

In this section we evaluate the effectiveness of the proposed defense mechanisms and demonstrate their viability.

A. Simulation Settings and Metrics

We use SIDnet-SWANS [19], [20] which is an open-source large scale sensor network simulator. SIDnet-SWANS is built on the scalable architecture of JiST-SWANS [21], which in turn is based on a high-performance JiST engine. It carries adapted version of ns-2's MAC802.15.4 protocol and same signal propagation models.

Network Configuration. The simulated environment consists of a set of 750 homogeneous nodes having the following configuration: (1) 20 kbps transmission/reception rate, (2) MAC802.15.4 protocol, (3) 5 seconds idle-to-sleep interval (i.e., nodes that are not actively involved in routing enter a low energy consumption state after 5 seconds of continuous idling, in order to preserve battery power), and (4) power consumption characteristics based on Mica2 Motes specifications [22]. To reduce the simulation time while preserving the validity of the observations, nodes were configured to use a small battery with an initial capacity of 35 mAh, for a projected lifespan of several tens of hours under moderate load.

Application Settings. The tested scenario consists of four distinct, long-term, continuous, point-to-point queries rooted at a common sink node. The sink is centrally located within the network. The four corresponding source nodes are evenly distributed around the sink node. Each experiment captures 8 hours of simulated time. Data transmission interval of the point-to-point queries to the designated sink is 4 seconds. Attackers are randomly and uniformly selected, ranging from 5% to 30% of the total sensors in the network.

Metrics. We monitor the *successful query dissemination rate*, expressed as the ratio between the number of queries received at the corresponding source nodes for processing and the total number of queries submitted through the sink node. We monitor the average residual energy levels $\overline{E}$ in the entire network, normalized relative to the capacity of a fully charged battery E_{max}. The effectiveness of the workload balancing paradigm and its associated energy consumption distribution is measured by means of the standard deviation of the percentage-representation of the residual energy reserves E_{σ}. Namely, if $E_i(t) \leq E_{max}$ is the residual energy level of a sensor node sn_i at a discrete time t, then the average energy level in a network of N nodes is $\overline{E}(t) = \frac{1}{N}\sum_{i=1}^{N} E_i(t)/E_{max}$. The standard deviation of the energy level is computed as follows:

$$E_{\sigma}(t) = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(E_i(t) - \frac{1}{N}\sum_{j=1}^{N} E_j(t))^2} \qquad (1)$$

We measure the packet-delivery ratio η = n_{rcv}/n_{exp}, established between the number of packets actually received n_{rcv} by the sink node and the total number of packets sent n_{exp} by the source node and expected at the sink over an interval of time. In multipath settings, the delivery ratio accounts for the successful transmission of one (of the possible many) copies of a packet. Also, the (depreciation of the) packet delivery latency is also monitored as part of the overhead analysis.

B. Evaluation of TESLA for Integrity and Authentication

To demonstrate the effectiveness of TESLA for integrity and authentication, we mounted a path deflection attack via altering of electrostatic charge information in network via either QUERY or UPDATE messages. Path deflection is the most representative attack to be considered because (1) it is an attack that targets unique characteristics of MP-FPR, (2) it requires very little resources to mount and (3) it can yield most damaging effects over the energy consumption patterns.

Energy balancing and data delivery rate performance evaluation. The path deflection attack is constructed as follows: forged-charges are generated and randomly placed in various areas of the network through the UPDATE messages. Various levels of attack efforts are considered, by varying the number of forged charges between 4 and 24, the upper bound value being enough to create major loss of connectivity in the network, as the experiments will show.

Figure 3(a) shows the impact of inserting invalid charge information in the network: disruption of the energy balancing. MP-FPR is very sensitive to this type of attack: even few number of forged charges, for example 4 such charges, are enough to drastically affect the evenness of the energy consumption, as the standard deviation of residual energy reserves nearly doubles. The reason behind is the severe path deflection and agglomeration of routes in narrow physical areas, as a result of the repulsive effect of multiple forged charges . In these conditions, most, if not all, of the alternate paths within a family merge and converge towards a single path type of routing in the relay area. MP-FPR effectivelly degrades towards a single-path routing behavior.

When a larger body of forged charges are considered, i.e. more than 8 such charges, there exists an apparent improvement of the energy-balance, as it can be observed in Figure 3(a). This observation surfaces, in fact, an extreme side effect of charge forgery attack: user perceived data DoS. Namely, it is possible that field lines are deflected enough that *all* of the associated routes are too long to be accepted in the route construction phase. The net result is a complete isolation between affected source nodes and their targeted sink. This lack of connecting routes prevents the data-stream from being sent to the sink, resulting in energy-savings by not performing the required workload. To demonstrate that this is the case, we capture the impact over the data-delivery rate in Figure 3(b). As it can be observed, data-delivery rate drops because of this effect. Correspondingly, network wide average of residual energy levels improves by up to 12%, conform Figure 3(c), when 24 forged charges are randomly injected in the network.

TESLA energy overhead evaluation. Figures 3(a), 3(b) and 3(c) demonstrate that TESLA not only provides the required protection against all path-related attacks, namely path deflection, path diversity deflation, family path intersection, wild-paths and field-line hopping, but the energy-overhead is minimal and independent of the dimension of the attack. Namely, it can be observed that TESLA's impact over the energy-balancing mechanism is below 3%, whereas, conform

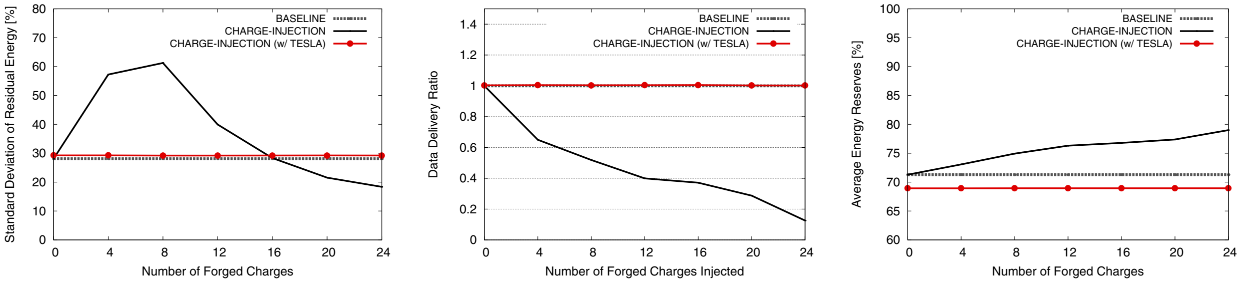

(a) Standard deviation of residual energy (b) Impact on data stream delivery (c) Average residual energy

Fig. 3. Tesla effectiveness and overhead evaluation on path deflection attacks via charge forgery

Figure 3(c), the impact over the network-wide average residual energy levels is maintained below 5%.

C. *Effectiveness of k-RPEF Against Selective Forwarding*

Selective forwarding of QUERY messages. Figure 4(a) shows that targeting the QUERY messages represents an easy and effective way to block query processing capabilities in the network. For example, by targeting 5% of the sensor nodes, an attacker can expect to impact 30% of the queries submitted. To demonstrate the effectiveness of the k-RPEF replication mechanism, we test against settings with degrees of replication of $k = 2, 4$ and 6. For example, when 6 replicas of QUERY messages are sent, MP-FPR proves to become nearly insensitive to the same small-base of attacks against QUERY messages (5%), with fewer than 1% query dissemination failures. Overall, we note an approximate reduction of successful attacks by 5% for every additional path used for replication, slightly lower under very intense attack settings of more than 25% compromised nodes. This information is relevant for deciding the number of replicas and multi-paths a query message will be sent along, when specific security needs and risk factors are known. Since query submission is an infrequent event, the number of k-RPEF multi-paths can be increased solely based on the security requirement, as the impact on the energy reserves is negligible.

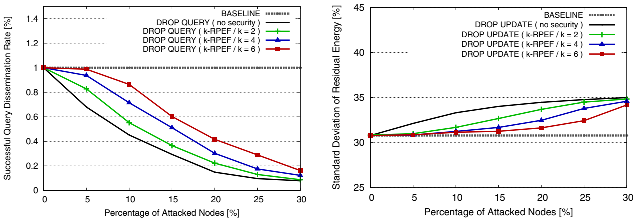

(a) QUERY messages (b) UPDATE messages

Fig. 4. Impact on energy consumption for selective forwarding of QUERY and UPDATE messages and the effectiveness of k-RPEF defense

Selective forwarding of UPDATE messages. Figure 4(b) illustrates the impact of the wild-path condition attack over the residual energy balancing property. As it can be observed, attacks carried during route establishment phase may yield up to 15% degradation of energy consumption balancing for the data-rate considered. It is important to note that the relative proximity of the source nodes determines the fraction of paths that may intersect and consequently can further impact the level energy imbalance. Figure 4(b) also demonstrates that employing the k-RPEF mechanism effectively alleviates the family path intersection attacks. Namely, when the degree of replication is set to $k = 6$, the degradation of energy balancing is maintained below 2% for bases of attacks that cover up to 15% of the nodes, and below 5% degradation when 20% of nodes are compromised.

Selective forwarding of ACK messages. Dropping ACK messages leads to a similar outcome as to the attacks carried via selective forwarding of RREQ messages, as comparing Figure 5(a) with Figure 6(a) demonstrates. Namely, with only a base of 5% of compromised nodes, the effective number of routes have been reduced by nearly 50%, slightly worse than the selective forwarding of RREQ messages.

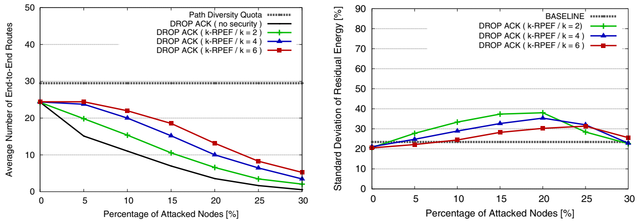

(a) Path diversity (b) Energy consumption balance

Fig. 5. Impact of selective forwarding of ACK messages to path diversity and effectiveness of k-RPEF defense mechanism

One distinction between ACK and RREQ messages in the MP-FPR protocol is that ACK message are not tightly coupled to a particular field line to be forwarded along, hence replicas can be created and forwarded along distinct paths. To this end, Figure 5(a) demonstrates a significant improvement provided by the k-RPEF mechanism, ranging from approximatively 30% improvement when the degree of replication is $k = 2$, to nearly 100% improvement as the degree of replication is increased to $k = 6$. We can also observe a linear dependency of the improvement to the number of replicas, each additional replica providing a benefit of 15%, on average, from a resilience perspective to these types of attacks.

The selection of the degree of replication k also impacts the energy balancing, as illustrated in Figure 5(b). Namely, larger number of replicas promote larger set of routes that improve energy consumption balancing at a rate of approximatively 8% for each additional replica, consistent for attacks comprised of less than 20% nodes. When the attacking base is increased beyond the 20% mark, an apparent improvement of the energy balancing situation similar with the one discussed under the RREQ message dropping manifests.

D. Effectiveness of PDMS Against Selective Forwarding

Selective forwarding of RREQ messages. We have simulated path-diversity deflation attacks via selective forwarding of RREQ messages. Figure 6(a) shows the sensitivity to path diversity deflation attacks, as even with a small base of 5% compromised nodes, the number of paths is effectively reduced by 40% as compared to the non-adversarial settings. PDMS improves the resilience to route establishment attacks, as for the same base of attacking nodes, the reduction of alternative paths is of only 6%. Thus, the attacker needs to consider tripling the attacking base, i.e. targeting approximatively 15% sensor nodes instead of 5% nodes, to achieve the same damaging effect as in the unprotected MP-FPR.

Figure 6(a) illustrates an additional benefit of PDMS: improving path diversity even under non-adversarial conditions. Namely, even when there are no compromised nodes, MP-FPR yields an average of 17% fewer routes than the user-specified quota ($N_r = 30$ in these settings). This is because MP-FPR discards routes that do not meet the end-to-end latency requirements (cf. Section II), such as overly long paths caused due to bandwidth starvation, long field lines or link quality issues and it does not compensate for. PDMS implicitly addresses this issue by persisting in probing routes until the path diversity quota is being met, as PDMS is oblivious of the underling reasons for which certain routes are not acknowledged.

Figure 6(b) illustrates the depreciation of energy-balancing as the number of compromised nodes is increased, where it can be observed that there is an 110% increase in standard deviation of the residual energy levels when only 10% of the nodes are compromised. PDMS helps maintaining even energy consumption distribution, achieving below 15% depreciation under the same scenarios – a significant improvement over the unprotected MP-FPR. The workload imbalance tops with 175% depreciation when 20% nodes maliciously drop RREQ messages, and ”recover” as the number of attacks is further increased. We recall that the apparent recovery is due to the loss of end-to-end connectivity. When absolutely no routes can be established between the source and sink nodes due to very large base of compromised nodes, the data stream becomes virtually absent and the afferent messages are dropped at the source. Energy savings are being achieved in the relay-area due to the lack of the data stream workload. To demonstrate that this is the case, we analyze in sequel the impact of attacks carried via selective forwarding of RREQ messages over the data delivery ratio.

As it can be observed in Figure 6(c), the sensitivity to message-dropping of RREQ messages is significantly reduced when compared to the reduction in path diversity under the same settings. Namely, when 5% of nodes are compromised, the impact to message dropping is below 1%. This is because the diminution of path-diversity does not affect message delivery, but the total absence of connecting routes does. As it can be observed, when the base of attacks is increased to 30% nodes, the average number of disconnected source-to-sink topologies is around 50%. The PDMS enables higher data-message delivery ratios since the family of routes it yields is consistently larger and the risk of non-connectivity is consequently lowered. PDMS forces an attacker to consider a much larger base of attacking nodes, an average of 20% more, to render PDMS scheme just as ineffective in achieving end-to-end connectivity as with the unprotected MP-FPR, with respect to the data stream deliverability.

Compensatory effect of PDMS to k-RPEF during attacks via selective forwarding of ACK messages. Both k-RPEF and PDMS mechanisms provide protection against path diversity deflation under adversarial conditions. However, these two mechanisms are fundamentally different: k-RPEF is a *proactive* mechanism, whereas PDMS is *reactive*. Namely, k-RPEF attempts to reduce the risk of failing to acknowledge a route, while PDMS attempts to build a new route if one has already failed. Since dropping either of ACK and RREQ messages leads to a route construction failure, PDMS will compensate for both in an attempt to meet the path diversity quota. That is, PDMS, when employed, will react to dropping of ACK messages as well. While we have analyzed k-RPEF and PDMS solutions in isolation, we do make note of this compensatory effect of the PDMS mechanism to the k-RPEF. Therefore, we are compelled to present an experimental analysis where both of these methods are concomitantly employed.

Figure 7(a) illustrates the improvement in path diversity when PDMS mechanism is enabled to provide compensation to the standalone k-RPEF mechanism. As it can be seen, this combination provides a virtually perfect defense against selective forwarding of ACK messages when the base of compromised nodes is below 10% as path diversity remains unaffected. Moreover, the PDMS component enables MP-FPR to reach the path diversity quota even under this adversarial scenario. It takes a large base of compromised nodes, i.e. at least 30% of the total number of sensor nodes, to achieve comparative protection of k-RPEF running in isolation against 20% of compromised nodes. From the perspective of sheer resilience to adversarial activity, PDMS improves the performance of k-RPEF, on average, by 90%.

PDMS, in isolation, cannot provide any benefit against selective forwarding of ACK messages. This is due to the SGP mechanism employed for relaying ACK messages in the original MP-FPR, as it was previously discussed. That is, if the SGP established sink-to-source path is compromised, *all* ACK messages will be dropped, including those acknowledging routes that PDMS attempts to build as replacement. In other words, compromising the unique route in SGP mechanism effectively nullifies the PDMS's benefits with respect to selective forwarding of ACK messages.

Energy balancing also benefits by enabling the PDMS to operate in conjunction with the k-RPEF solution. As Figure 7(b) demonstrates, considering a degree of replication of $k = 6$ and an attacking base of 20%, the disruption of energy balancing is of only 16%, i.e. a nearly 50% improvement when

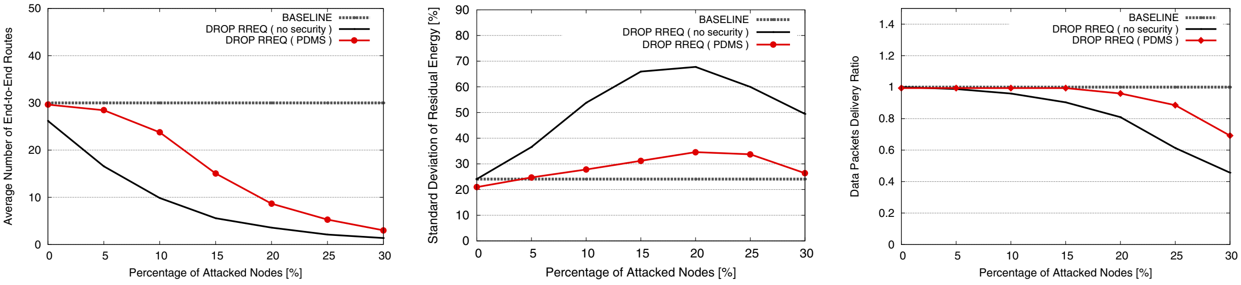

(a) Path diversity (b) Energy consumption balance (c) End-to-end data-stream delivery ratio

Fig. 6. Impact of selective forwarding of RREQ messages and the effectiveness of PDMS mechanism

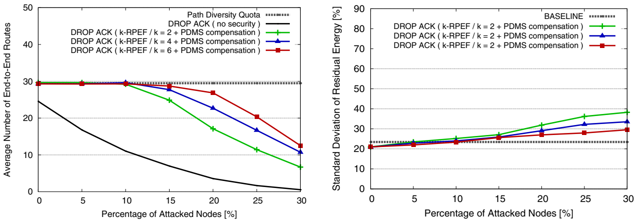

(a) Path diversity (b) Energy consumption balance

Fig. 7. Impact of selective forwarding of ACK messages and the effectiveness of k-RPEF and PDMS solution mix

compared to the equivalent performance of running k-RPEF in isolation (cf. previous results in Figure 5(b)).

E. Effectiveness of k-EF Against Selective Forwarding

Selective forwarding of DATA messages. Figure 8(a) illustrates the consequence of increasing the number of attacking nodes that target DATA messages: a 45% degradation in DATA packet delivery with a only a small base of 5% nodes, and nearly 90% degradation when the number of compromised nodes is increased to 15%. Adopting a multipath approach proves to be beneficial: at the minimum, the effect is reduced by a factor of two, i.e. from 45% to 23% message drops when only 2 replication paths are used, and less than 2% when 4 replication paths are used, considering 5% compromised nodes.

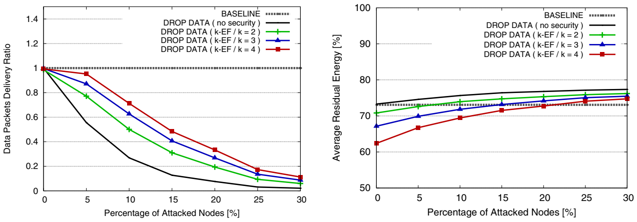

(a) Data delivery (b) Energy consumption balance

Fig. 8. Impact of selective forwarding of DATA messages and the effectiveness of the k-EF defense

Figure 8(b) shows an overhead varying between 3% and 15% as the number of multipaths is increased from $k = 2$ to $k = 4$. While the overhead is small, it can be much higher if the rate of transmission of data messages increases, currently set at .25 messages per seconds. The number of compromised nodes does not have a direct negative impact over the energy consumption. Energy savings are achieved when DATA messages are being dropped along a path due to an undesirable reduction of the workload. As it can be observed in Figure 8(b), the residual energy reserves increases monotonically corresponding to the reduction of the successful delivery of data messages from Figure 8(a).

VI. Related Work

Recent work on the security of sensor networks [23] has focused on proposing key management schemes that can be used to bootstrap other services [24], [25], [26], [27], [28], addressing general attacks such as Sybil [29] and replication [30] attacks.

The security of geographical routing protocols using physical nodes' locations was studied in [31] for sensor networks and in [32], [11] for ad-hoc networks. Most of the works focus on preventing malicious modifications of the destination location in packets, verifying neighbor location information, and preventing message dropping. Another main area of work in securing geographic routing is the protection of the location service, which includes [10], [9].

Security of the potential-field routing for WSNs has been investigated in [33]. This work distinguishes from our approach in the following aspects: (1) the work surveys a generic list of attacks and countermeasures that do not focus on the specifics of the potential-field routing, while we address specific risks introduced by the MP-FPR protocol in all phases of the protocol operation, from query dissemination and charge allocation to route establishment and data forwarding, and analyze these risk factors through extensive experimental analysis; (2) although potential-field routing and electrostatic field-based routing are both instances of the gradient based routing, their implementation is fundamentally different: the former is a *stateful* protocol, where routes are established based on distance metrics obtained by means of hop-counting, while MP-FPR does not maintain routing information and relies only on the distribution of discrete charge information for forwarding purposes; (3) field-based routing has been proposed initially in the context of large scale, dense mesh networks and there is no focus on energy consumption and workload distribution, whereas MP-FPR generalizes the usability of gradient based routing to arbitrary distributions with possible low densities of nodes and focuses on the energy aspect.

Geographic routing remains a promising and active area of research due to intrinsic benefits of exploiting location relationships for routing purposes. A complete survey of geography-based single-path routing approaches can be found in [34], whereas a newer approach that particularly considers the challenges of large scale sensor networks is presented in [35]. Other works have also recognized the benefits of using multipath routing in large-scale sensor networks for improving workload balancing and delivery robustness. For example, trajectory-based forwarding approaches, which rely on multiple non-braided paths via parametric curves for single source and sink scenarios, have been presented in [36], [37]. A natural extension to multiple sink, multiple-path is challenging because route disjointness cannot be easily guaranteed when adopting parametric trajectory models, therefore field, potential and gravity-based routing methodologies, which exploit physical phenomena properties to facilitate the creation of non-braiding paths, have been recently investigated [38], [39], [40]. Despite the broad interest in gradient based routing, very little work has been done to address the security aspect of such advanced protocols, which constitutes the motivational support for this body of work.

VII. Conclusions

In this article, we have presented an in-depth analysis regarding the feasibility of providing security semantics to MP-FPR – an instance of the electrostatic field based routing for location-aware sensor networks. We have identified the attacking model and the core system properties that uniquely characterize MP-FPR's settings. We discuss cryptographic mechanisms to provide integrity and authentication, and three mechanism providing resilience to selective data forwarding attacks. We have experimentally demonstrated that MP-FPR energy provisions can be significantly affected under an adversarial environment, however, effective security solutions that exploit MP-FPR's multi-path routing model can be implemented with minimal overhead.

References

[1] I. F. Akyildiz and M. C. Vuran, *Wireless Sensor Networks*. Willey, 2010.

[2] C. Lemmon, S. M. Lui, and I. Lee, "Geographic forwarding and routing for ad-hoc wireless network: A survey," in *NCM*, 2009, pp. 188–195.

[3] D. Niculescu and B. Nath, "Trajectory based forwarding and its applications," in *MOBICOM*, 2003.

[4] N. T. Nguyen, A. Wang, P. Reiher, and G. Kuenning, "Electric-field-based routing: a reliable framework for routing in MANETs," *SIGMOBILE Mobile Computing Communication Review*, vol. 8, no. 2, pp. 35–49, 2004. [Online]. Available: http://portal.acm.org/citation.cfm?id=997122.997129

[5] G. Trajcevski, O. C. Ghica, P. Scheuermann, M. Zuniga, R. Schubotz, and M. Hauswirth, "Improving the energy balance of field-based routing in wireless sensor networks," in *IEEE Globecom*, 2010.

[6] A. Perrig, R. Canetti, J. D. Tygar, and D. X. Song, "Efficient authentication and signing of multicast streams over lossy channels," in *IEEE Symposium on Security and Privacy*, 2000, pp. 56–73.

[7] R. Fonseca, S. Ratnasamy, J. Zhao, C. T. Ee, D. Culler, S. Shenker, and I. Stoica, "Beacon vector routing: Scalable point-to-point routing in wireless sensornets," in *NSDI*, 2005.

[8] T. He, C. Huang, B. M. Blum, J. A. Stankovic, and T. F. Abdelzaher, "Range-free localization schemes for large scale sensor networks," in *MOBICOM*, 2003, pp. 81–95.

[9] J. Dong, K. E. Ackermann, B. Bavar, and C. Nita-Rotaru, "Mitigating attacks against virtual coordinate based routing in wireless sensor networks," in *WiSec*. New York, NY, USA: ACM, 2008, pp. 89–99.

[10] X. Wu and C. Nita-Rotaru, "On the security of distributed position services," in *SecureComm*, 2005.

[11] J.-H. Song, V. W. S. Wong, and V. C. M. Leung, "Secure position-based routing protocol for mobile ad hoc networks." *Ad Hoc Networks*, vol. 5, no. 1, pp. 76–86, 2007.

[12] S. Ganeriwal, C. Pöpper, S. Capkun, and M. B. Srivastava, "Secure time synchronization in sensor networks," *ACM TISSEC*, vol. 11, no. 4, 2008.

[13] J. Barnickel and U. Meyer, "Secsywise: A secure time synchronization scheme in wireless sensor networks," in *ICUMT*, 2009, pp. 1–8.

[14] H. Krawczyk, M. Bellare, and R. Canetti, "Hmac: keyed-hashing for message authentication," *RFC*, vol. 2104, pp. 1–12, 1997.

[15] N. Koblitz, "Elliptic curve cryptosystems," *Mathematics of Computation*, vol. 48, pp. 203–209, 1987.

[16] D. Johnson, A. Menezes, and S. A. Vanstone, "The elliptic curve digital signature algorithm (ecdsa)," *International Journal of Information Security*, vol. 1, no. 1, pp. 36–63, 2001.

[17] O. C. Ghica, "Extending lifetime in large scale wireless sensor networks: Secure multipath routing and load-aware tracking," PhD Thesis UMI 3456555, Northwestern University, 2011.

[18] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin, "Highly-resilient, energy-efficient multipath routing in wireless sensor networks." *Mobile Computing and Communications Review*, vol. 5, no. 4, pp. 11–25, 2001.

[19] O. Ghica, G. Trajcevski, P. Scheuermann, Z. Bischoff, and N. Valtchanov, "Sidnet-swans: A simulator and integrated development platform for sensor networks applications," in *SenSys*, 2008.

[20] http://www.eecs.northwestern.edu/~ocg474/SIDnet.html.

[21] Http://jist.ece.cornell.edu/index.html.

[22] Http://www.xbow.com.

[23] C. Karlof and D. Wagner, "Secure routing in wireless sensor networks: Attacks and countermeasures," in *WSNA*, 2003. [Online]. Available: citeseer.ist.psu.edu/article/karlof02secure.html

[24] L. Eschenauer and V. Gligor, "A key management scheme for distributed sensor networks," in *CCS*, 2002.

[25] H. Chan, A. Perrig, and D. Song, "Random key predistribution schemes for sensor networks," in *SP*, 2003.

[26] H. Chan and A. Perrig, "PIKE: Peer intermediaries for key establishment in sensor networks," in *INFOCOM*, 2005.

[27] W. Du, J. Deng, Y. S. Han, P. K. Varshney, J. Katz, and A. Khalili, "A pairwise key predistribution scheme for wireless sensor networks," *TISSEC*, vol. 8, no. 2, 2005.

[28] D. Liu, P. Ning, and R. Li, "Establishing pairwise keys in distributed sensor networks," *TISSEC*, vol. 8, no. 1, 2005.

[29] J. Newsome, E. Shi, D. Song, and A. Perrig, "The sybil attack in sensor networks: analysis & defenses," in *IPSN*, 2004.

[30] B. Parno, A. Perrig, and V. Gligor, "Distributed detection of node replication attacks in sensor networks," in *SP*, 2005.

[31] N. Abu-Ghazaleh, K.-D. Kang, and K. Liu, "Towards resilient geographic routing in wsns," in *Q2SWinet*, 2005.

[32] T. Leinmüller, C. Maihöfer, E. Schoch, and F. Kargl, "Improved security in geographic ad hoc routing through autonomous position verification," in *VANET*, 2006.

[33] D. Sharma, "Security of field based routing," *Student Thesis SA-2008-08*, 2008.

[34] S. Ruehrup, "Theory and practice of geographic routing," in *Ad Hoc and Sensor Wireless Networks: Architectures, Algorithms and Protocols*. Bentham Science, 2009.

[35] A.-M. Kermarrec and G. Tan, "Greedy geographic routing in large-scale sensor networks: A minimum network decomposition approach," in *MobiHoc*, 2010.

[36] M. Desai and N. Maxemchuk, "Polar cooordinate routing for multiple paths in wireless sensor networks," in *WOWCOM*, 2010, pp. 1–9.

[37] O. Ghica, G. Trajcevski, P. Scheuermann, N. Valtchanov, and Z. Bischof, "Controlled multi-path routing in sensor networks using bezier curves," *The Computer Journal*, vol. 54, no. 2, pp. 230–254, 2011.

[38] C. Wu, R. Yuan, and H. Zhou, "A novel load balanced and lifetime maximization routing protocol in wireless sensor networks," in *VTC Spring*, 2008, pp. 113–117.

[39] J. Li, S. Ji, H. Jin, and Q. Ren, "Routing in multi-sink sensor networks based on gravitational field," in *ICESS*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 368–375.

[40] S. Rührup, H. Kalosha, A. Nayak, and I. Stojmenovic, "Message-efficient beaconless georouting with guaranteed delivery in wireless sensor, ad hoc, and actuator networks," *IEEE/ACM Transactions Networking*, vol. 18, no. 1, pp. 95–108, 2010.