SIEVE: a distributed, accurate, and robust technique to identify malicious nodes in data dissemination on MANET

Rossano Gaeta, Marco Grangetto, Riccardo Loti
Università di Torino, Dipartimento di Informatica
Corso Svizzera 185, 10129 Torino, ITALY
first.last@di.unito.it

Abstract—

In this paper we consider the following problem: nodes in a MANET must disseminate data chunks using rateless codes but some nodes are assumed to be malicious, i.e., before transmitting a coded packet they may modify its payload. Nodes receiving corrupted coded packets are prevented from correctly decoding the original chunk. We propose *SIEVE*, a fully distributed technique to identify malicious nodes. *SIEVE* is based on special messages called *checks* that nodes periodically transmit. A check contains the list of nodes identifiers that provided coded packets of a chunk as well as a flag to signal if the chunk has been corrupted. *SIEVE* operates on top of an otherwise reliable architecture and it is based on the construction of a *factor graph* obtained from the collected checks on which an incremental belief propagation algorithm is run to compute the probability of a node being malicious. Analysis is carried out by detailed simulations using ns-3. We show that *SIEVE* is very accurate and discuss how nodes speed impacts on its accuracy. We also show *SIEVE* robustness under several attack scenarios and deceiving actions.

Keywords—MANET, data dissemination, malicious node identification, statistical inference.

I. INTRODUCTION

Due to their open, fully distributed and dynamic architecture, shared wireless channel, and resource limited participants, MANET are vulnerable to attacks at any layer of the Internet model [23].

We consider a particular type of active, non-cryptography related attack: data corruption at the application level launched by insider nodes. In this paper we deem as a use case a data dissemination application over a MANET. Nodes generate data chunks to be disseminated to all participants using rateless codes; some malicious nodes deliberately modify coded packets of a chunk before relaying them to prevent honest nodes from obtaining the original information. In this paper we propose *SIEVE* a distributed, accurate and robust technique to identify malicious nodes on top of an otherwise reliable and attacker-free architecture. Each node in *SIEVE* dynamically creates a bipartite graph (*factor graph*) whose vertices are *checks*. A check is a report created by a node upon decoding a data chunk; a check contains a variable length list of nodes identifiers that provided parts of the data as well as a flag to signal if the data chunk has been corrupted. Detection of the compromised chunks is achieved

exploiting the constraints imposed by linear channel coding. The factor graph is periodically and independently analyzed by each node running an incremental version of the Belief Propagation (BP) algorithm [25]. The proposed algorithm allows each node to compute the probability of any other node being malicious; these latter probabilities are used to derive a *suspect ranking* of nodes in the MANET. The major contributions of the paper are the recasting of the problem of malicious nodes identification in terms of the estimation of the marginal probabilities on a bipartite graph and the proposal of a distributed and accurate solution based on the BP algorithm. Paper contributions are completed by a comprehensive experimental investigation of *SIEVE* capabilities. It is worth pointing out that the selected data dissemination application is just a quite popular use case [18], [17], whereas the proposed approach is by no means constrained to such particular scenario. In particular, *SIEVE* can be used in any application that uses multi-party download or collaboration, provided that is possible to detect that a given set of collaborating entities is compromised by at least one malicious node.

SIEVE fits well two key MANETS features that must be accounted for when devising any security solution: it is fully distributed and does not rely on any infrastructure (as opposed to some solutions in the area of peer-to-peer streaming where special well known nodes are necessary, e.g., [21]) and does not have a high computational cost (as opposed to on-the-fly verification techniques in the area of network coding, e.g., [9], [2], [27], [7], [26]).

The paper is organized as follows: Section 2 describes the system model we consider, Section 3 presents the *SIEVE* technique, Section 4 discusses the simulation methodology and the accuracy, reactivity, and robustness results we obtained, Section 5 summarizes other works related to *SIEVE*, finally Section 6 draws conclusions and outlines directions for future developments.

II. A USE CASE FOR SIEVE

In this paper we consider a MANET composed of N ($N = N_{fast} + N_{slow} + N_{still}$) wireless nodes on a square area whose size is l meters. Nodes move using different average speeds: N_{fast} fast moving nodes at V_{fast} m/s (e.g.,

cars, buses, motorcycles), N_{slow} slow nodes at V_{slow} m/s (e.g., pedestrian, bikes), and N_{still} still nodes do not move (e.g., sensors, relay stations, shops).

Nodes periodically produce a new data *chunk* to be disseminated to all others once every h seconds; to this end, all nodes cooperate by running a common application that assigns each node an unforgeable unique identifier, e.g., an integer value. When nodes produce or successfully obtain a chunk they turn into data sources for that chunk and help in spreading it. Data is transmitted by source nodes using rateless codes [?]: a chunk (whose size in bytes is fixed and is denoted as S) is divided in K equally sized packets (size of each coded back is $S_{cb} = \frac{S}{K}$); the source node then creates and forward coded packets using rateless LT codes [?]. According to [?] the original chunk can be obtained by any node able to collect any set of $K \cdot (1 + \epsilon)$ coded packets (ϵ is known as the code overhead). Coded packets are produced only by source nodes for a chunk and consist in simple binary XOR operations among a random set of d original data packets, provided that d is selected according to the Robust Soliton Distribution (RSD) with parameters c and δ [?]. A coded packet conveys the XORed payloads of the corresponding original packets as well as a header signalling the indexes of the combined packets. Decoding is carried out by solving a system of linear equations. In [?] it is shown that the most simple decoder based the simplification of the equations of degree 1 exhibits asymptotically optimal overhead, provided that the RSD distribution is used to encode packets.

Each node can simultaneously collect coded packets for different chunks; to this end a window based mechanism is used where the w_r most recent chunks can be concurrently downloaded. Similarly, each node can serve as data source for the w_t most recent received chunks. A simple round robin scheduling policy is used when transmitting coded packets of chunks in the upload window; coded packets are transmitted every T_{tx} msec using UDP over an 802.11g-like wireless communication interface yielding an average transmission range of r meters.

A subset of $P < N$ nodes ($P = P_{fast} + P_{slow} + P_{still}$) is composed of *malicious nodes*, i.e., they deliberately modify the payload of coded packets they produce before transmission to prevent honest nodes from correctly decoding the original chunk. In presence of coding even a single corrupted coded packet can prevent an honest node from decoding the original chunk.

III. THE SIEVE PROTOCOL

SIEVE is based on the concept of *checks* that are reports created by nodes upon decoding a chunk. A check contains a variable length list of nodes identifiers that provided coded blocks of a chunk and a flag to label the chunk as corrupted or not.

A node is able to detect such condition as soon as an inconsistency is found in the solution of the underlying system of linear equations. In particular, according to the procedure [?] the decoder keeps cancelling out all the already known data packets. This is achieved by observing that a coded packet with a degree 1 equation represents a data packet in the clear. Such data packet can be simplified from all the incoming equations. Since LT codes have a certain overhead some coded packets that are linearly dependent on the ones received previously are always collected before successful decoding; this amount to the reception of some equations whose terms are all already known. As soon as this condition is met the LT decoder can check the consistency of the payload carried by the coded packet; in other words, the same linear combination must be obtained combining a set of already known packets. If this constraint is violated the whole chunk is recognized as corrupted. Please note that the receiver node is not able to identify the corrupted block(s) but only that at least one of them has been maliciously manipulated. A check describing a corrupted chunk is called a *positive check* while it is termed a *negative check* otherwise. Each nodes n maintains a list of all checks created that is denoted as $\mathcal{L}_n$.

Each node, besides accumulating the checks from its local decoding operations, gossips them in the neighborhood. Each node n in the MANET transmits its checks in two cases:

- as soon as n decodes a chunk it inserts it in $\mathcal{L}_n$ and broadcasts it;
- once every T_s seconds n randomly selects Q checks in $\mathcal{L}_n$ and transmits them.

The checks in $\mathcal{L}_n$ and all checks received by n are used to build a *factor graph* $\mathcal{G}_n = (\mathcal{U}, \mathcal{C}, \mathcal{E})$. $\mathcal{G}_n$ is a bipartite graph where the vertex set $\mathcal{U}$ is the set of uploader nodes, the vertex set $\mathcal{C}$ is the set of checks, and an undirected edge $\{i, I\} \in \mathcal{E}$ exists if and only if check $I \in \mathcal{C}$ depends on uploader $i \in \mathcal{U}$.

In the following we will refer to the set of uploaders involved in check I as $\mathcal{U}_I$ and the set of checks that node i contributes as an uploader as $\mathcal{C}_i$. An example of factor graph with four uploaders (circles) and two checks (squares) is show in Fig. 1. Each uploader i can be in one of two states $x_i = 1$ or $x_i = 0$, depending on whether uploader i is or is not a malicious node. Each check can report one of two observations $c_I = 0$ or $c_I = 1$ in case of negative or positive pollution detection, respectively. Coming back to the example of Fig. 1 the filled circle is used to represent a polluter that in turn causes a positive check (filled square).

The problem of identifying the malicious nodes from a given number of checks can be recast as an inference problem. The goal of the inference is the estimation of the hidden state of the nodes, i.e. being malicious or not, given a set of observations corresponding to the checks. Each check can be interpreted as an accusation raised against a set of

Figure 1: Example of factor graph.

uploader nodes by a witness node. In this paper we adopt the *Belief Propagation* (BP) algorithm [25].

The BP algorithm can be used to estimate from the factor graph the so called variable marginals $(P(x_i))_{i \in \mathcal{U}}$, i.e. the probability of node i being malicious. BP is an iterative algorithm based on the exchange of probability messages, the belief, along the edges of the bipartite graph $\mathcal{G}_n$. In case of a Bayesian network BP represents a close form solution for the marginals. The same algorithm can be applied iteratively to a general factor graph that, as opposed to a Bayesian network, includes loops; in this case BP has proven to be a robust estimator for the variable marginals [25].

In our setting it is convenient to distinguish between two class of messages: message from uploader i to check I , m_{iI}^x , that is meant to be the probability that uploader i is in state x , given the information collected via checks other than check I ($\mathcal{C}_i \setminus I$); message from check I to uploader i , m_{iI}^x is defined as the probability of check I having value c_I if uploader i is considered in state x and all the other uploaders states have a separable distribution given by the probabilities $\{m_{i'I}^x : i' \in \mathcal{U}_I \setminus i\}$.

The messages are initialized to the values $m_{iI}^0 = m_{iI}^1 = 0.5$, i.e. all nodes are equally likely to be malicious in the first run. Then, the probabilities m_{iI}^x are estimated using:

$$m_{iI}^x = \sum_{\{x_{i'} : i' \in \mathcal{U}_I \setminus i\}} P(c_I | x_i = x, \{x_{i'} : i' \in \mathcal{U}_I \setminus i\}) \prod_{i' \in \mathcal{U}_I \setminus i} m_{i'I}^{x_{i'}} \quad (1)$$

Equation (1) depends on the probability of observing a certain check value c_I , given the states of the uploaders of such check. Given that a check turns out to be positive as soon as at least one of the uploaders is a malicious node we can write:

$$P(c_I = 1 | \{x_i : i = 1, \dots, k\}) = \begin{cases} 0, & \text{if } x_i = 0, \forall i \\ 1, & \text{otherwise} \end{cases} \quad (2)$$

Analogously, observing that a check can be negative if and only if all the uploaders are not malicious we get

$$P(c_I = 0 | \{x_i : i = 1, \dots, k\}) = \begin{cases} 1, & \text{if } x_i = 0, \forall i \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

Plugging the last two expressions into Equation (1) we can simplify it as shown in Table 1 for the four possible combinations of c_I and x . The next step is constituted by

Table I: Definition of m_{iI}^x as a function of x and c_I .

$c_I \setminus x$	0	1
0	$\prod_{i' \in \mathcal{U}_I \setminus i} m_{i'I}^0$	0
1	$1 - \prod_{i' \in \mathcal{U}_I \setminus i} m_{i'I}^0$	1

the updating of the probabilities m_{iI}^x , using information from previous computation in the checks: $m_{iI}^x = \prod_{I' \in \mathcal{C}_i \setminus I} m_{iI'}^x$. Finally, the marginal $P(x_i = x)$ can be computed as $P(x_i = x) = \prod_{I' \in \mathcal{C}_i} m_{iI'}^x$. To conclude the BP algorithm initializes the values of m_{iI} , then keeps updates messages from checks to nodes and from nodes to checks. A reliable estimate of $P(x_i = x)$ is computed after a certain number of iterations. In all the experiments reported in this paper 3 iterations have been used.

A. Incremental BP estimation

In the previous description of the BP we have assumed that the factor graph $\mathcal{G}_n$ is known in advance and kept fixed for all the iterations. In practice this assumption is not met in the proposed scenario. Nonetheless, the proposed algorithm can be implemented using an incremental (or sliding window) approach as follows.

Each node n keeps receiving checks from the other ones and it is allowed to run the BP on $\mathcal{G}_n$ every T seconds considering only the checks received and created in a time window of the past w seconds. At time t , depending on the checks stored during a time window w , an updated factor graph $\mathcal{G}_{n,t,w}$ is obtained by removing the old checks and adding the new ones; then, the corresponding estimates $P_{t,w}(x_i = x)$ are computed through BP. It is worth pointing out that the belief values are initialized only once, as soon as a node shows up as an uploader for the first time in a check; then, the partial estimates of m_{iI}^x and m_{iI}^x are propagated to all the computation windows where such values are needed.

After every BP algorithm run a list of suspect nodes is obtained by setting a threshold on the probability $P_{t,w}(x_i = 1) \geq \eta$. Each node n keeps a counter for each of its uploader nodes i : the nodes in the list of suspects after the BP run have their counter increased by 1. Finally, a *suspects ranking* $\mathcal{R}_n$ over uploader nodes is defined by sorting their counters in decreasing order. As an example, the first node in the suspects ranking at time t is the uploader that more often has been included in the list of suspects after all the BP runs performed up to time t .

IV. RESULTS

In this section we present results on the accuracy and robustness of *SIEVE*. The performance of *SIEVE* have been investigated by simulations using ns-3 version 3.12 on a

Table II: Parameter values for the reference scenario.

N_{fast}, N_{slow}	100	K	100
N_{still}	50	S_{cb}	500 Bytes
P_{fast}, P_{slow}	5	T_{tx}	100 ms
P_{still}	0	c, δ	0.01
r	175 m	l	2000 m
V_{fast}	uniform [10-40] m/s	h	600s
V_{slow}	uniform [1-5] m/s	w_r	50
Mobility model	2D random walk	w_t	6

Red Hat 4.4.6-3 machine using standard variable settings. In particular, we developed a node object implementing the sender and receiver functionalities according to the protocol described in Sections 2 and 3. All nodes movements are described by the same mobility model and moving nodes are randomly distributed across the simulation area at the beginning of the experiment. Still nodes are distributed on a grid to avoid clustering in a particular region. We used the ns-3 default transmission model composed of a propagation delay model (a constant value) and a propagation loss model (based on a log distance model with a reference loss of 46.677 dB at a distance of 1 meter). Under these settings we derived the maximal transmission range r as 175 meters.

We conducted terminating simulations (not steady-state) runs that lasted for 1 hour of simulated time. We used the independent replication method to obtain $N_{EXP} = 30$ realizations and computed 95% confidence intervals. Each realization is obtained with different streams of the random number generator provided by ns-3. The simulation output is made of a log file containing all checks received by each node that is post-processed by our analysis software to run the *SIEVE* protocol to detect malicious nodes.

All results have been obtained by setting as data sources only the still nodes; all N nodes contribute to the spreading of the information according to the policy described in Section 2. All the system parameter values of the reference scenario considered in the following are summarized in Table 2.

For each node n in the system we define the following performance indexes:

- $p(n)$ an indicator function whose value is equal to 1 if node n is not malicious and 0 otherwise;
- $c_n(t)$ an indicator function whose value is equal to 1 if $\mathcal{R}_n$ is non empty at time t and 0 otherwise;
- $a_n(t)$ the number of actual malicious nodes correctly identified by node n at time t . More precisely, $a_n(t) = x$ if x nodes in the top P positions of $\mathcal{R}_n$ are truly malicious;
- $r_n(t)$ the number of actual malicious nodes correctly identified in the top positions of $\mathcal{R}_n$ at time t . More precisely, $r_n(t) = x$ if *top* x nodes in $\mathcal{R}_n$ are all actually malicious. If *SIEVE* at time t has identified $P - 1$ out of P malicious nodes ($a_n(t) = P - 1$) but the first node in $\mathcal{R}_n$ is not malicious then $r_n(t) = 0$.

Figure 2: Completeness for $w = 60$ s with 95% confidence intervals.

- $t_{n,trigger}$ the time node n received the very first positive check that triggered the *SIEVE* activation;
- $t_{n,hit}(x) = \min_t \{t : r_n(t) \geq x\}$, that is, the minimum time to identify at least x malicious nodes in the top x positions of $\mathcal{R}_n$.

In each independent replication trial we compute the following averages:

- $c^{(i)}(t) = \frac{\sum_{n=1}^N p(n)c_n(t)}{\sum_{n=1}^N p(n)}$, i.e., the fraction of honest nodes that detected the attack and have a non empty $\mathcal{R}$;
- $a^{(i)}(t) = \frac{1}{P} \frac{\sum_{n=1}^N p(n)c_n(t)a_n(t)}{\sum_{n=1}^N p(n)c_n(t)}$, i.e., the average accuracy of honest nodes with non empty $\mathcal{R}$. When $a^{(i)}(t) = 1$ it means that *all* honest nodes with non empty $\mathcal{R}$ have correctly identified *all* malicious nodes.
- $tsi^{(i)}(x) = \frac{\sum_{n=1}^N p(n)(t_{n,hit}(x) - t_{n,trigger})}{\sum_{n=1}^N p(n)}$, i.e., the average time honest nodes require to unambiguously identify x malicious nodes.

Finally, we computed averages and 95% confidence intervals over N_{EXP} realizations and considered $c(t)$, $a(t)$, and $tsi(x)$. In all computations we set $\eta = 0.99$ (η is the threshold on the probability $P_{t,w}(x_i = 1)$ for suspect identification) and iterated each run of the BP algorithm on a specific factor graph three times.

A. Sensitivity results

The first analysis we conduct is on the sensitivity of $c(t)$, $a(t)$, and $tsi(x)$ to parameters w , T , T_s , and Q used by *SIEVE* as defined in Section 3. Due to space limitations, in the following we omit graphs for $tsi(x)$. To this end we considered integer multiples of 10s values for w ranging from 10s to 90s, $T = \{5, 10\}$ s, $T_s = \{5, 10, 15\}$ s, and $Q = \{5, 10\}$. We analyzed all 108 combinations and observed that *SIEVE* performance is most sensitive to the value w .

Figure 3: Accuracy as a function of w .

Figure ?? show the completeness $c(t)$ for $w = 60s$ when $T = T_s = 10s$, and $Q = 10$ (all values of w yield very similar results so we chose a representative case to avoid cluttering the figures). It can be noted that $c(t)$ approaches 1 after about 4 minutes; the behavior of $c(t)$ clearly depends on the delay after which a node has collected a sufficient amount of checks to start its own suspect ranking. All experiments worked out in the following exhibit a very similar behavior for $c(t)$ so we avoid repeating similar figures. Furthermore, all following graphs compare different scenarios so we plot average results without confidence intervals for better readability.

Figure 2 shows the performance index $a(t)$ with the same settings (all 108 configurations yielded similar results so we selected a representative scenario). It can be noted that too small or too large values for w yield the worst performance. Indeed, small window sizes do not allow the factor graph to include enough checks to accurately infer the node status; on the other hand, large values of w provide more checks in $\mathcal{G}_n$ but increase the number of short loops in the factor graph which in turn impact on the accuracy of the probability estimates yielded by the BP algorithm [15]. As for *SIEVE* reactivity, we observe that $w = 60s$ yields the lowest values for $tsi(x)$ (we obtain $tsi(1) = 33s$).

We selected $w = 60s$ as the value for all successive evaluation. Performance of *SIEVE* are less dependent on the values of other parameters. Actually, small values for T improve reactivity by decreasing $tsi(x)$ but increase the frequency of BP computation; this may cause concerns for battery operated nodes whose lifetime could be shortened. Battery lifetime can be prolonged by using large values for T_s and small values of Q . Nevertheless, performance of *SIEVE* deteriorates for the same reasons observed for small values of w . We selected a compromise between performance and energy consumption awareness using the following setting for the rest of the paper: $w = 60s$,

Figure 4: Accuracy for moving nodes ranging from all fast to all slow.

$T = T_s = 10s$, $Q = 10$.

B. Mobility and *SIEVE* performance

The first interesting observations we made is that nodes mobility affects *SIEVE* performance. Figure 5 shows the overall performance of *SIEVE* in the reference scenario and in scenarios with different mixes of fast and slow nodes. The system where all moving nodes are fast yields much higher accuracy and much lower reaction times with respect to other extreme case where all nodes move slowly. Nevertheless, *SIEVE* accuracy is 0.83 at the end of the one hour long experiments and approaches 1 for longer runs.

Figure 6 shows results in the reference scenario and in scenarios where all malicious nodes are either fast or slow. The best way for malicious nodes to delay identification by honest nodes is to move as slow as possible; indeed, if all malicious nodes move fast all of them are quickly identified.

Although *SIEVE* is able to identify all malicious nodes in the long run in any setting, in all three cases we can conclude that *high speed is key to obtain both high accuracy and low reaction delays by honest nodes*. Why is high speed beneficial for *SIEVE* performance of honest nodes but detrimental for malicious nodes? Consider the case where different mixes of fast and slow nodes (Figure 5) are compared and assume an extreme scenario composed of all still nodes. Since nodes do not move it is possible to define a static geometric graph $\mathcal{O}$ describing connections among nodes where vertices are nodes and an undirected edge between two vertices exist if the corresponding nodes are within the transmission radio range r . For any node $n \in \mathcal{O}$ let $\mathcal{N}(n)$ denote the set of nodes connected to n (the neighborhood of n).

Consider one data source s and a chunk c ; clearly, when s has transmitted $K \cdot (1 + \epsilon)$ coded packets, c is decoded by all nodes in $\mathcal{N}(s)$; in turn, these nodes start transmitting fresh coded packets for c and, if $\mathcal{O}$ is connected, all nodes

Figure 5: Accuracy for malicious nodes that are either all fast or all slow.

will be able to decode c after some time. It can be noted that all coded packets produced by s will follow the same paths in $\mathcal{O}$. This means that for any node $n \in \mathcal{O}$ chunks originating from s will be provided to n by the same set of uploaders $\mathcal{U}_s(n)$ that is a subset of $\mathcal{N}(n)$. Of course, the same reasoning is valid for all checks describing chunks produced by any data source other than s . The final effect is that the set of checks created by each node n upon decoding chunks ($\mathcal{L}_n$) is such that strong intersections exist among the uploaders of different checks (given that $\mathcal{O}$ is static and $\mathcal{N}(n)$ does not change over time); this translates into a high number of short loops in the factor graph used by *SIEVE* that is a well known cause of poor performance of the BP algorithm [15]. The shortest length of a loop in a bipartite graph with at most one edge between any two nodes is 4 and is due to the presence of at least two common uploaders in a pair of checks, i.e. the cardinality of the intersection between the uploaders of a pair of checks is at least 2. When the cardinality of the intersection is $x > 2$ then each of the $\binom{x}{2}$ pairs defines a length 4 loop.

When nodes move the geometric graph $\mathcal{O}$ becomes dynamic, i.e., for any node n its $\mathcal{N}(n)$ varies with time. Higher speed translates into higher rate of changing in $\mathcal{N}(n)$ that, in turn, lowers the number of short loops in $\mathcal{G}_n$ increasing accuracy of the BP algorithm. For the reference scenario we computed the average number of checks in $\mathcal{G}_n$ for fast, slow, and still nodes: we obtained 435, 423, and 742, respectively. We also computed the average total number of length 4 loops obtaining 29978, 41574, and 76563, respectively. It can be noted that still nodes must work on a larger $\mathcal{G}_n$ (the average neighborhood size is 7.37, 7.69, and 12.73, respectively so still nodes receive more random selections of Q checks from their neighbors) that results in a much higher number of short loops. Furthermore, the average number of checks for a node (the degree of a node) is higher for still

Figure 6: Accuracy for different pollution intensities.

nodes (24.75) with respect to fast and slow nodes (16.49 and 16.56, respectively). Of course, the higher the presence of a node in different checks, the higher the probability of short loops forming in $\mathcal{G}_n$ because of intersection of the set of uploaders.

C. Deceiving actions and *SIEVE* robustness

Besides following the data dissemination and *SIEVE* protocols, malicious nodes may also implement disturbing actions hoping to avoid or delay identification by honest nodes. Malicious nodes may deceive honest nodes by:

- reducing pollution intensity: a coded packet is modified before transmission only with probability p_{poll} by flipping a coin. Figure 7 shows how *SIEVE* performs when nodes lower their pollution intensity by 50% ($p_{poll} = 0.5$) in the attempt of making their identification harder. It can be noted that *SIEVE* is able to quickly identify all malicious nodes also for $p_{poll} = 0.5$; indeed, malicious nodes that reduce their pollution intensity are less effective in damaging honest nodes and do not succeed in substantially delay identification.
- lying on the status of the check: they flip a coin and with probability p_{lie} a positive check is sent out as negative and viceversa. Figure 8 shows that this trick is almost ineffective.
- disparaging honest nodes: malicious nodes always produce dummy positive checks involving a set of honest nodes. In this case malicious nodes flip a coin and with probability $p_{disparage}$ the actual uploader nodes of a check are replaced by honest nodes; the detection flag is marked as positive. Figure 9 shows *SIEVE* performance when $p_{disparage} = 1$ and honest nodes are either randomly chosen or chosen in the same fixed order followed by all malicious nodes (this is a colluding attack to honest nodes). It can be noted that, by colluding, malicious nodes only succeed in slightly delay their identification. As a matter of fact, *SIEVE* is

Figure 7: Accuracy for different lying intensities.

Figure 8: Accuracy for different disparaging attacks.

able to correctly identify all malicious nodes anyway at the end of the simulation experiments.

The last stress test for *SIEVE* is to consider its performance for an increasing number of malicious nodes that coordinate to launch the colluding attack. In Figure 10 we show results for up to $P = 40$ malicious nodes in the system (in all cases, $P_{fast} = P_{slow} = \frac{P}{2}$). It can be noted that *SIEVE* accuracy reaches 0.9 at the end of the simulation experiments.

D. Computational cost

One of the main requirements a solution for MANET must satisfy is to have a low computational and memory cost. We measured the average CPU time experienced to run *SIEVE* on a single factor graph in our C++ implementation on an Intel(R) Core i5 2.80GHz CPU: we obtained 25ms. Of course, MANET nodes do not have the same computational power of a desktop PC but newer CPUs equipping tablets and smart phones are reducing the gap, e.g., the ARM Cortex-A9 MPCore has up to 4 cores, 2 GHz clock, and 10,000 DMIPS. Furthermore, storage requirements are very

Figure 9: Accuracy for increasing number of colluding malicious nodes.

low: the average number of checks in $\mathcal{G}_n$ for fast, slow, and still nodes is 435, 423, and 742, respectively. The average number of nodes in $\mathcal{G}_n$ is equal to 227, 224, and 228, respectively.

V. RELATED WORKS

MANET are vulnerable to attacks at any layer of the Internet model [23]. In the area of network coding several efforts have been devoted to devise on-the-fly verification techniques carried out by participants [9], [2], [27], [7], [26] to identify the sources of corrupted data. The major drawbacks of these elegant methods is the high computational costs for verification and the communication overhead due to pre-distribution of verification information. Error correction is another approach to deal with data corruption attacks in network coding [3], [4], [8]; these methods introduce coding redundancy to allow receivers to correct errors but their effectiveness depends on the amount of corrupted information. The data corruption attack we consider in this paper is a well-known plague in peer-to-peer streaming systems. Unfortunately, all solutions developed in that research are not easily adoptable in MANET. The work by Wang et al. [21] proposes a detection scheme where each peer is able to detect receipt of corrupted blocks by checking the adherence of the decoded chunk to the specific formats of the video stream. Peers detecting polluted chunks send alert messages to the video server and the tracker. Upon receipt of an alert the server computes a checksum of the original chunk and disseminates it to all peers in the overlay. The checksum is used by peers to identify which uploader actually sent a corrupted block. Peers report their suspects to the server and a true polluters cannot lie (the authors develop a non repudiation protocol to ensure that peers cannot lie when reporting suspects to the servers). Sequence numbers are used to tag alerts to deal with cycles in the overlay. The lack

of a centralized monitoring and management point make it difficult to adapt this solution to MANET.

The work by Li and Lui [13] presents a distributed detection algorithm and analyzes its performance. The technique is based on simple intersection operations performed by peers: each peer starts with a set of suspects that is equal to the entire neighborhood that is shrunk as long as chunks are downloaded from a random subset of uploaders independently chosen from the entire set of neighbors. The scheme allows malicious nodes to send corrupted blocks using a pollution probability. The technique is analyzed when the number of malicious nodes in the neighborhood is known in advance and an approximation is proposed when this quantity is unknown. The technique is attractive thanks to its simplicity and fully distributed nature. Performance deteriorates when multiple polluter exist and it is not clear if the technique can work when then neighborhood of a peer varies with time.

The work by Jin et al [5] proposed a monitoring architecture to build and maintain a reputation system that peers use to select neighbors. The focus of the paper is on reputation computation, storage and load balancing among monitoring nodes. The results show that the system is able to detect malicious nodes up to a certain degree of lies. Nevertheless, the technique relies on the assumption that each peer is able to compute the amount of corrupted blocks received by each uploader during a monitoring period. This appears to be a rather unrealistic assumption.

VI. CONCLUSIONS

In this paper we have proposed the novel *SIEVE* technique for the identification of malicious nodes performing an active, non-cryptography attack within a MANET. *SIEVE* is a completely distributed technique that, using statistical inference based on the belief propagation algorithm, allows each node to independently analyze local snapshots of the bipartite graph of the collected checks to estimate the probability of nodes being malicious and to perform a progressive ranking of the suspect nodes. Our results, worked out using detailed ns-3 simulations, show that *SIEVE* is accurate in letting each honest node identify all malicious nodes under several scenarios. We have analyzed the sensitivity of *SIEVE* performance to the nodes mobility; we have discovered that *SIEVE* is very robust to several deceiving actions, colluding attacks launched by malicious nodes, and amount of malicious nodes inside the network.

REFERENCES

- [1] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael. Learning low-level vision. *International Journal of Computer Vision*, 40:25–47, 2000.
- [2] C. Gkantsidis and P. Rodriguez. Cooperative security for network coding file distribution. In *IEEE INFOCOM*, 2006.
- [3] T. Ho, B. Leong, R. Koetter, M. Medard, M. Effros, and D. Karger. Byzantine modification detection in multicast networks with random network coding. *Information Theory, IEEE Transactions on*, 54(6):2798–2803, june 2008.
- [4] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, M. Medard, and M. Effros. Resilient network coding in the presence of byzantine adversaries. *Information Theory, IEEE Transactions on*, 54(6):2596–2603, june 2008.
- [5] X. Jin and S.-H. G. Chan. Detecting malicious nodes in peer-to-peer streaming by peer-based monitoring. *ACM Trans. Multimedia Comput. Commun. Appl.*, 6:9:1–9:18, March 2010.
- [6] D. Kamal, D. Charles, K. Jain, and K. Lauter. Signatures for network coding. In *In Proceedings of the fortieth annual Conference on Information Sciences and Systems*, 2006.
- [7] E. Kehdi and B. Li. Null keys: Limiting malicious attacks via null space properties of network coding. In *INFOCOM 2009, IEEE*.
- [8] R. Koetter and F. Kschischang. Coding for errors and erasures in random network coding. *Information Theory, IEEE Transactions on*, 54(8):3579–3591, august 2008.
- [9] M. N. Krohn, M. J. Freedman, and D. Mazieres. On-the-fly verification of rateless erasure codes for efficient content distribution. *Security and Privacy, IEEE Symposium on*, 2004.
- [10] S. Kurkowski, T. Camp, and M. Colagrosso. Manet simulation studies: the incredibles. *SIGMOBILE Mob. Comput. Commun. Rev.*, 9:50–61, October 2005.
- [11] Q. Li, D.-M. Chiu, and J. Lui. On the practical and security issues of batch content distribution via network coding. In *Network Protocols, 2006. ICNP '06. Proceedings of the 2006 14th IEEE International Conference on*, 2006.
- [12] X. Li, C.-C. Wang, and X. Lin. Throughput and delay analysis on uncoded and coded wireless broadcast with hard deadline constraints. In *IEEE INFOCOM, 2010*, pages 1–5, march 2010.
- [13] Y. Li and J. C. Lui. Stochastic analysis of a randomized detection algorithm for pollution attack in P2P live streaming systems. *Performance Evaluation*, 67(11):1273–1288, 2010.
- [14] M. Luby. Lt codes. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science, IEEE FOCS 2002*, pages 271–280.
- [15] M. Luby, M. Mitzenmacher, M. Shokrollahi, and D. Spielman. Improved low-density parity-check codes using irregular graphs. *Information Theory, IEEE Transactions on*, 47(2):585–598, feb 2001.
- [16] D. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
- [17] D. Nguyen, T. Tran, T. Nguyen, and B. Bose. Wireless broadcast using network coding. *IEEE Transactions on Vehicular Technology*, 58(2):914–925, feb. 2009.

- [18] J.-S. Park, M. Gerla, D. Lun, Y. Yi, and M. Medard. Codecast: a network-coding-based ad hoc multicast protocol. *Wireless Communications, IEEE*, 13(5):76–81, october 2006.
- [19] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.
- [20] G. F. Riley and T. R. Henderson. The ns-3 network simulator. In *Modeling and Tools for Network Simulation*, pages 15–34. Springer Berlin Heidelberg, 2010.
- [21] Q. Wang, L. Vu, K. Nahrstedt, and H. Khurana. MIS: Malicious nodes identification scheme in network-coding-based peer-to-peer streaming. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–5, march 2010.
- [22] B. Wu, J. Chen, J. Wu, and M. Cardei. A survey of attacks and countermeasures in mobile ad hoc networks. In Y. Xiao, X. S. Shen, and D.-Z. Du, editors, *Wireless Network Security, Signals and Communication Technology*, pages 103–135. Springer US, 2007.
- [23] H. Yang, H. Luo, F. Ye, S. Lu, and L. Zhang. Security in mobile ad hoc networks: challenges and solutions. *Wireless Communications, IEEE*, 11(1):38–47, 2004.
- [24] J. Yedidia, W. Freeman, and Y. Weiss. Understanding belief propagation and its generalizations. In *Exploring Artificial Intelligence in the New Millennium*, Science & Technology Books, chapter 8. Elsevier, 2003.
- [25] J. Yedidia, W. Freeman, and Y. Weiss. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Transactions on Information Theory*, 51(7):2282 – 2312, 2005.
- [26] Z. Yu, Y. Wei, B. Ramkumar, and Y. Guan. An efficient scheme for securing xor network coding against pollution attacks. In *INFOCOM 2009, IEEE*.
- [27] Z. Yu, Y. Wei, B. Ramkumar, and Y. Guan. An efficient signature-based scheme for securing network coding against pollution attacks. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, 2008.