

# Convex Hull Watchdog: Mitigation of Malicious Nodes in Tree-based P2P Monitoring Systems

Andreas Disterhöft

Technology of Social Networks, University of Düsseldorf  
Universitätsstr. 1, 40225 Düsseldorf, Germany  
Email: disterhoeft@cs.uni-duesseldorf.de

Kalman Graffi

Technology of Social Networks, University of Düsseldorf  
Universitätsstr. 1, 40225 Düsseldorf, Germany  
Email: graffi@hhu.de

**Abstract**—Monitoring the global state in peer-to-peer networks through decentralized mechanisms allows targeted optimization and improvement of the peer-to-peer network. However, malicious nodes could aim to distort the process of gathering the global state through monitoring. In this paper, we propose DOMiNo, a security solution for tree-based peer-to-peer monitoring mechanisms. It passively listens to incoming events, e.g. data, and rates its suspiciousness based on outlier detection, structural verification and sanity check mechanisms. For our main objective, which is to limit the monitoring error of the desired global view, we performed an extensive evaluation. Evaluation shows tolerance with normal fluctuations but effective filtering of outliers, that severely influence the global view. As our watchdog solution operates passively, we do not add any costs nor create new surface for attacks to the monitoring system.

## I. INTRODUCTION

Peer-to-peer (P2P) networks offer many advantages compared to client-server network models. They are able to build networks using only the capacities of every single peer participating in the network and build up applications with a minimum of operational costs. However, as no central instance is present it is hard to manage these networks to offer a certain Quality of Service. Monitoring systems present a solution to gain information about the status of the network so peers may adapt several parameters to improve the overall performance (e.g. reducing the output during high package loss or increasing their routing tables' size in unstable networks). In particular, structured monitoring systems provide highly precise monitoring results by building a topology to effectively aggregate and distribute data. Since every peer gathers its local view of the network (e.g. local hop count, response time etc.) and has to contribute its known information to the monitoring system to gain a profound view on the underlying P2P overlay, new attack possibilities arise. Even more as no trustworthy instance, i.e. a server, is present.

P2P overlays with controlled quality of service might play a role in future, such as for P2P-based online social networks (LibreSocial [9]), WebP2P-based communication platforms [6] or future communication platforms for decentralized energy sources. As an example the Department of Computer Science from the University of Oldenburg [26] uses a constructed communication overlay to exchange data between several power plants which monitor each other. If one of those power plants

gets compromised the results could be fatal. Not just in this cases, but in general, more information on the P2P network, in specific global information as obtained from monitoring mechanisms, are beneficial for the operation of the network. Bottlenecks can be identified, parameters configured ideally and an autonomous control loop could be implemented on the long run, as suggested in [12]. As information is crucial in the correct operation of a P2P overlay, malicious nodes could aim to distort the monitoring mechanism, thus aiming at indirect damage through the processing of false observations.

In our paper, we propose a security solution, DOMiNo, for existing structured tree-based monitoring systems. Our approach is applicable for tree-based monitoring systems which aggregate gathered system-specific information in a push-based manner. Examples for suitable systems are CONE [3], SDIMS [28], SOMO [30] and SkyEye.KOM [10]. Appropriate systems require a dissemination process, where each participating node is supplied with the monitoring mechanism's global view. The construction of the monitoring tree must be deterministic. Thus a parent node in the tree must be able to identify a certain ID space where its child nodes must be located. This requirement is interchangeable with any other criteria where a node's parent is able to verify its child nodes. Please note, if this requirement is not met by the tree-based monitoring system, malicious nodes can inject their (manipulated) data at arbitrary and even multiple points in the monitoring tree, which may diminish our system's efficiency. Also, a node can only be considered for the monitoring process if it is actively participating in the underneath p2p overlay, thus nodes participating in the monitoring process can check each node's activity. With these assumptions, our solution provides following properties:

- **Passive Rating:** No overhead is added to the existing monitoring approach. The detection does not add any overhead in terms of exchanging information, i.e. sending messages, nor does it perform any additional non-local operations like lookups. Thus, DOMiNo is a passive rating mechanism which judges by listening to inbound data messages only. By doing so, we do not introduce new possibilities for attackers to disturb or influence the monitoring system in a bad way.
- **Precision:** The main objective is to minimize the relative

This work was supported by the DFG Grant GR4498 1-1 - OverlayMeter

error from the measured global view in a malicious environment to the measured global view in a non-malicious environment respectively desired global view. Thus, we prioritize a high precision over a minimization of the false positive rate.

- Performance: The detection mechanism should need as little effort in terms of computation time as possible.

First, in Section II we introduce briefly in the literature in the field of monitoring systems in general and security of them in particular. In order to be able to evaluate our solution, we focus on the motivation of attackers and concretize a fitting attacker model in Section III. Our main contribution, called DOMiNo, which makes use of outlier detection mechanisms based on Z-Scores, structural and other plausibility checks, is presented in Section IV. The quality of our proposed system DOMiNo is evaluated in Section V through simulations. We show that the influence of the malicious node's attempt on the global view are effectively limited. This paper closes with a conclusion and comments on future work in Section VI.

## II. RELATED WORK

In the research field of monitoring in decentralized networks there exist several approaches but none of them handle malicious behavior. On one hand, the unstructured gossip-based approaches like [16], [18], [23] periodically exchange messages with their neighbors and average the obtained values. Over time, the system wide average is obtained, which can be used also for sum or count calculations. Non of these solutions propose a mechanism to protect the system against misbehaving nodes. On the other hand, there are the structured monitoring approaches, which are mostly tree-based. GAP [5] and its follow ups Adaptive GAP [22] and TCA-GAP [27] build up a tree to aggregate monitoring information and provide a threshold crossing alert but they do not support malicious behavior analysis which is also mentioned in TCA-GAP. Other approaches like CONE [3], SDIMS [28], SOMO [30] and SkyEye.KOM [7] also do not provide such mechanisms. In this paper, we propose an add-on which can be used by the before mentioned approaches.

In the field of sensor networks there exist some approaches like SDAP [29] or from Chan et al. [4]. In the first approach, SDAP, the authors propose a grouping technique for aggregates to determine the suspiciousness for those groups. In their approach they build up new communication paths for their group-ups, in contrast, we passively listen to incoming data and rate it. Also, we reserve the right to drop suspicious data. The second approach also builds up its own structure which violates our requirement of being passive. Our solution is a passive module which can be integrated to the already existing tree-based aggregation structure.

Another field of research are reputation systems in P2P networks, wherefore Marti and Garcia-Molina provide a survey [19]. Here, the most prominent solution is EigenTrust [17], where nodes actively distribute trust vectors whereas the trust system as a whole is transitive (if node A trusts B then A trusts all nodes B trusts). However, we focus on passive

mechanisms, which do not add any message exchange to the current monitoring service and do work only locally.

The field of (distributed) intrusion detection systems, which work passively, are well suited for our purpose. But in contrast to traditional (distributed) intrusion detection systems like mentioned in the survey from Axelsson [2], our desire is to provide a mechanism with a decision-making component which i) operates only on the local view ii) works decent on normally distributed data iii) does not need a preceding training and iv) operates reliably on small data sets. In the field of Mobile Ad Hoc Networks, Watchdog mechanisms [20], [21] have been proposed that observed the forwarding behavior of neighboring nodes and mark them as potentially malicious. While in Mobile Ad Hoc Networks, communication is assumed to be observable, in P2P networks, communication paths are hidden. Nevertheless, the idea is appealing.

Therefore we have chosen to apply methodologies from the field of *outlier detection* [13]. Due to our restrictions we can not make use of Supervised and Semi-supervised outlier detections, thus Unsupervised methods, which are proximity-based, cluster-based, density-based, distance-based or using probabilistic/statistical models, are left. Most of the mentioned methods operate reliably on bigger data sets and/or build clusters, which do not fit our needs. Except statistical models, like the Extreme Value Analysis with Z-Scores, which is used by Iglewicz and Hoaglin [15], fit our needs. Thus using the divergence to the standard deviation of the data set we are able to locally rate the badness of incoming data. Additionally, the proposed method from Iglewicz and Hoaglin works best on a normally distributed data base.

Hence, in the theory, using an appropriate divergence factor, each node should be able to distinguish between data from honest nodes including some noise and bad data, if the number of honest data exceeds the number of malicious data. In Section IV we describe how we utilized this in our solution.

## III. ATTACKER MODEL

Attackers in monitoring systems can have different intentions to behave not in accordance with the monitoring algorithm. We state briefly the motivation of attackers, classify attack patterns and extract an attacker model.

### A. Motivation of Attackers

The literature offers a wide range of an attacker's motivation to act malicious as stated in the survey from Rounds and Pendraft [24]. However, irrespective of the reason for being malicious, virulent nodes participating in a monitoring system can be categorized into the following classes.

In the first class, participants may manipulate monitoring data to hinder/influence the QoS (= Quality of Service) mechanism of the system by changing the current system view or certain node's capacities. Applying nodes may aim to attack honest nodes' global views to, for example, let QoS mechanisms regulate accordingly which may lead to an overall worse P2P experience.

The second class deals with free riders which want to maximize their benefit with as little own effort as possible. So, these malicious nodes are only interested in the global view but ideally do not want to forward any messages.

The third and last class of malicious nodes want to purposefully disable monitoring by not following rules or even shutting down specific mechanisms or sealing off single nodes. In contrast to the first group, this group tries to break links in the monitoring system and to exclude respectively eclipse specific nodes or whole sub trees from participation. This implies an incomplete global view for nodes not eclipsed and no global view for eclipsed nodes which leads to unawareness of the P2P network's current state.

### B. Attack Patterns

Next, we describe attacks that achieve the aims mentioned in the last section. For this, we want to stick to an example monitoring system that fulfills the requirements mentioned in Section I. SkyEye.KOM [10] is a tree-based monitoring system which aggregates statistics in a bottom-top manner and disseminates the global view top-bottom. For that, all nodes calculate initially their position in the tree as well as their parent's DHT ID based on the node's ID responsibility in the DHT. Periodically, each nodes captures its local statistics, aggregates it with those received from its child nodes, and sends the aggregated snapshot of its subtree to its parent node. The parent node replies with an ACK containing its currently known global view. Over time, all information is aggregated towards the root, where the global view is created and from there trickled down to all nodes using the ACKs.

Figure 1 shows an example topology in SkyEye.KOM with branching factor 2. Black lines denote relevant respectively *active* connections of an attack, where data is sent in the direction the arrow points. Dotted lines show connections which should be established following the monitoring algorithm, but will not be maintained as a result of an occurring attack. If marked with an arrow, a dotted line shows that one side tries to establish the connection unsuccessfully. Gray lines show existing connections according to the SkyEye.KOM algorithm which are not affected by either attack depicted in the Figure.

We will denote the different attack types via the schema  $AT_{Abbreviation-for-Attack}$  for future references.

1) *Manipulation of Monitoring Results*: This group contains malicious nodes which send data to other nodes to change the current view of them. With this they still keep the monitoring procedure running, but can influence the behavior of the whole network, e.g. by sending parent and child nodes the information that only little bandwidth is available in the corresponding sub tree. Another possibility would be to declare high available bandwidth which may lead to the failure of single nodes due to congestion or similar. Next, we want to differentiate the following attacks in this group.

a)  $AT_{Parent}$  – *Send modified Data Sets to the Parent Node*: If a peer sends wrong data to its parent node it is able to modify the local view in the parent node and therefore such a peer has indirect influence on adjacent sub trees, as the

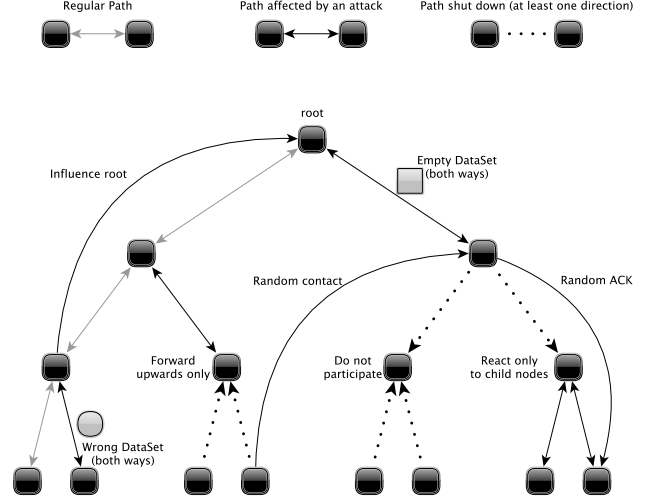


Fig. 1: Overview of attack types against SkyEye.KOM, a tree-based monitoring system fulfilling our requirements

parent node will send the influenced local view up towards the root. At some point the root will get the modified data and uses it to build the global estimation which will be sent down the tree. Furthermore peers may present themselves as weak nodes to maximize their benefit from the running P2P overlay without providing too much resources themselves as many overlays have mechanisms to balance load among the peers according to their capabilities.

This attack type is best suited for having influence on the monitoring results while being as unobtrusively as possible since the node behaves *correctly* regarding the message forwarding algorithm, i.e. it addresses the correct parent node.

b)  $AT_{Child}$  – *Send modified Data Sets to the Child Node(s)*: Peers can influence all their child nodes or only a specific child node by forwarding incorrect global views downwards the tree with the ACKs. With this they are able to reduce the profit child nodes get from the P2P overlay, as they may not be able to react on potential network changes.

Just like  $AT_{Parent}$ , malicious nodes using this approach behave correctly according to the monitoring procedure.

c)  $AT_{Random}$  – *Send Data to a Random Contact*: Sending data to random nodes corresponds to sending messages upwards the tree and lets the sender appear as a child node of the receiver. In addition, nodes applying this attack also manipulate the data to increase their impact. As depicted in Figure 1 this attack can affect arbitrary nodes at any position and may cause loops in the tree if forwarded to a node lower in the tree, which may lead to exorbitant global views. Furthermore, the number of child nodes may exceed a possible maximum, i.e. the branching factor, if all possible child nodes are present, which may lead to further actions depending on the monitoring system, e.g. reorganization respectively applying load balancing mechanisms.

d)  $AT_{RandomACK}$  – *Send ACK to a Random Contact*: In SkyEye.KOM, application level ACK messages are sent

downwards the tree to disseminate the global view. In general, this kind of message is interchangeable, but in the following, ACKs are linked to the dissemination of the global view. Consequently, sending arbitrary ACK messages will influence or even override the global view in the receiving node.

Using this attack results in a more suspicious behavior as parent nodes are usually well known since child nodes initiate the communication between themselves and the parent.

2) *Free Riders: max benefit, as little own effort as possible:* Free riders want to maximize their benefit from a mechanism with little own effort. This can be achieved requesting the global view which is calculated, e.g. at the root node.

a) *AT<sub>Root</sub> – Send Data to the Root directly:* Data sent to the root node directly by a node that is not a legitimate child is a special case of the attack *AT<sub>Random</sub>*. As demanded in Section I, a certain ID space can be identified where child nodes live in but they cannot be fully verified. In contrast to *AT<sub>Random</sub>*, this type of attack has a much higher influence (on the global view), as the root node aggregates the global view and starts the dissemination down the tree.

3) *Shutting down of Monitoring Mechanisms:* The following attack types aim to shut down specific mechanisms of the monitoring system or to seal off single nodes from it. These attacks are more aggressive towards the concerned parts of the system as they do not provide any information at all.

a) *AT<sub>Empty</sub> – Forward Empty Data Sets:* Forwarding empty data sets results in the erasure of parts of the gathered statistics if the forwarding is upwards to the root or even in the erasure of the whole statistics if empty sets get forwarded downwards the tree. Attackers are able to shut down significant parts of the monitoring tree completely as the receivers have to work with sparse data or even no data if the attacker is the only connection to the other tree parts. In Figure 1, this would mean that if the root node uses this attack, the right sub tree would not be able to gain any global data according to the message forwarding algorithm.

b) *AT<sub>Upwards</sub> / AT<sub>Downwards</sub> – Forward Data only upwards or downwards:* On the one hand, if nodes forward data only upwards, i.e. without participating in the dissemination process, they take advantage of the monitoring system without contributing much to it except the message they send towards the root and the contribution to the global view. On the other hand, if nodes forward data only downwards, they can eclipse a whole sub tree without instantly being suspicious. Nodes in the sub tree will not be able to commit to the global view and may only receive information about the sub tree itself if the attacker decides to do so.

c) *AT<sub>DoNothing</sub> – Do not participate:* Missing participation in the monitoring system will result in broken links between multiple tree parts since the denying nodes are the proper parent respectively child nodes and thus need to be contacted as part of the monitoring algorithm. This behavior will both influence the monitoring system's results and shut off specific nodes from the system. It combines both attacks *AT<sub>Upwards</sub>* and *AT<sub>Downwards</sub>*.

## IV. OUR SOLUTION: DOMiNo

*Detection Of Malicious Nodes* (DOMiNo) is our approach to identify inappropriate behavior in a tree-based structured monitoring system like SkyEye.KOM, which meets the requirements stated in Section I. Our approach is based on a rating system containing multiple rating mechanisms, each returning a numerical rating for an examined data set or peer, respectively. Ideally, once a peer behaves maliciously, it would trigger one or more active rating systems and be detected.

In the following, we discuss the detection criteria and identify which attacks these criteria are able to detect. We divide our rating mechanisms into *instant rating mechanisms* which instantly determine malicious behavior and *collecting rating mechanisms* which operate on a collected data base.

### A. Instant Rating Mechanisms

Instant rating mechanisms check different properties of a received message, e.g. its origin or the sender's properties. Furthermore they review the containing data sets regarding several points, e.g. the completeness or the plausibility of data, using simple but effective rules.

1) *Check Data's Origin:* This mechanism checks each received data's origin whether its ID lies in the acceptable ID range where data is expected to come from, as demanded in Section I. By doing so, false negatives cannot arise while false positives are still possible. For this, nodes should not be able to forge their IDs. To achieve this, several approaches can be applied like using a PKI, where the public key equals the node ID, thus encrypting data using the public key can only be decrypted by the node owning the corresponding private key (See Graffi et al. [11]). By doing this, IDs are coupled to an identity and its asymmetric key pair, i.e. to a user. IP/MAC spoofing for example is useless as incoming data can not be decrypted due to the missing private key. This mechanism can be exploited to detect *AT<sub>Random</sub>* attacks.

2) *Identifying correct Parent Nodes via Nonces:* By introducing numbers-used-once (nonces) set by the requester and used by the responding node to identify a pair of request and corresponding response message, the monitoring mechanism can cope with the attack type *AT<sub>RandomACK</sub>*. So, the requesting node is able to identify the response to a sent request and can ignore or filter out ACK messages with an unknown nonce.

3) *Implausible Values:* Based on the knowledge we have about the data we receive respectively expect we can perform several checks for implausible values. Depending on each data set's values, numerous sanity checks like  $\min \leq \max$ ,  $\min \leq \text{mean} \leq \max$ ,  $\text{sumOfSquares} \leq \max^2 * \text{count}$ , etc. should hold. These criteria support the detection of attacks which send insufficient data, i.e. *AT<sub>Empty</sub>*, but also every attack which carelessly manipulates data.

4) *Awareness of Missing ACKs:* As part of the monitoring mechanism, once contacted, each node can expect a reaction from the contacted node if it is active. A node ignoring this policy is acting highly suspicious. In order to not produce false positives, as a message can get lost or just delayed, an

active node is identified to be malicious if it repeatedly does not answer, e.g. after three unsuccessful tries.

By doing so, we are able to detect those attacks where nodes do not or partly participate in the monitoring system, i.e.  $AT_{Upwards}$  and  $AT_{DoNothing}$  attacks. Please note, that leaf nodes running  $AT_{DoNothing}$  cannot be detected with a passive approach such as DOMiNo. Depending on the used monitoring system, parents may be not aware of their children. In order to rectify this, new lookups would be needed producing a significant amount of new overhead to the overlay.

### B. Collecting Rating Mechanisms

In contrast to instant rating mechanisms, collecting rating mechanisms need a certain data base in order to work. With the help of this data base, our proposed mechanisms determine the likeliness that the data was produced by the same system and try to differentiate manipulated data from correct data.

In the following, we propose our collecting rating component which uses Z-Scores.

**Outlier detection: Z-Scores:** Our outlier detection, which has its root in the Modified Z-Scores introduced by Iglewicz and Hoaglin [15], rates each incoming data set by using its local data base. This data base contains all received data sets which are not older than a certain maximum age, thus, the branching factor of the monitoring tree affects the data base used by our algorithm.

First, based on the Modified Z-Scores, the median absolute deviation (in short MAD) is calculated using the Equation 1, where  $x$  is the whole data base and  $x_i$  is the  $i$ 'th entry which gets rated. Please note, the use of the median is crucial in order to not get biased by extreme values. The Modified Z-Score is calculated using the Equation 2. The authors propose the threshold, for labeling a data set as a potential outlier, to be 3.5. This threshold identifies when noise turns into suspiciousness, which is defined by a convex hull around a multiple of the median absolute deviation as visualized in Figure 2. By increasing this threshold the algorithm gets increasingly tolerant to noise.

$$MAD = \text{median}(|x_i - \text{median}(x)|) \quad (1)$$

$$Score_i = \frac{|0.6745(x_i - \text{median}(x))|}{MAD} \quad (2)$$

In the rare case where at least 50 percent of the provided data sets report the same value, Equation 1 turns to 0, thus we would divide by 0 in Equation 2. In this case all points that do not equal the median could be marked as an outlier. In our context this is not a desired behavior as we require to operate on diverse data input with a possibly high deviation. Especially when using discrete data, e.g. the hop count, this case might happen. To be able to make a meaningful decision nevertheless, we use a solution used by IBM in their software SPSS Statistics [14]. They use the mean absolute deviation (MeanAD) in case the MAD equals 0. The corresponding calculations for the  $MeanAD$  respectively  $Score_i$  are denoted in Equations 3, 4. MeanAD can only be 0, if all data set values

### Algorithm 1 Modified Z-Median Score Algorithm

```

1: Input: A set of collected data  $D$ , the data point to examine  $d_i$ 
2: procedure CALCULATEMODIFIEDZMEDIAN
3:    $dataSetMedian \leftarrow \text{median}(D)$ 
4:   for all  $d_x$  in  $D$  do
5:     add  $|dataSetMedian - d_x|$  to a median deviation list  $MD$ 
6:   end for
7:    $MAD \leftarrow \text{median}(MD)$ 
8:   if  $MAD == 0$  then ▷ If more than 50% of the value were equal
9:      $MeanAD \leftarrow \text{average}(MD)$ 
10:    if  $MeanAD == 0$  then ▷ All values are equal
11:       $Score_i \leftarrow \text{NOTSUSPICIOUS}$ 
12:    else
13:       $Score_i \leftarrow |(0.7979 * (d_i - dataSetMedian))| / MeanAD$ 
14:    end if
15:  else
16:     $Score_i \leftarrow |(0.6745 * (d_i - dataSetMedian))| / MAD$ 
17:  end if
18:  return  $Score_i$ 
19: end procedure

```




Fig. 2: Visualization of a node rating its incoming data sets.

are the same and therefore no outlier is present.

$$MeanAD = \text{average}(|x_i - \text{median}(x)|) \quad (3)$$

$$Score_i = \frac{|0.7979(x_i - \text{median}(x))|}{MeanAD} \quad (4)$$

Algorithm 1 presents our method to get a rating for a single data set. Due to the possibility of false positives, especially given the reduced data base, one should keep in mind to control a node's behavior over a certain time interval and not prejudge a node because of a single statistical abnormality. There is no guarantee that a node may not look suspicious for a short amount of time due to sudden changes in the network or other reasons. We therefore propose to evaluate a node's behavior several times before making a final judgment when using statistical approaches. This countermeasure is most suited for manipulation attacks sending data upwards the tree, i.e.  $AT_{Parent}$ ,  $AT_{Random}$ , and  $AT_{Root}$ .

Additionally, all attacks mentioned in Section III may be applied over a long period of time which could increasingly influence the global view created by the monitoring mechanism. In order to minimize the impact of false values on the calculation of the convex hull, old monitoring data is cleaned periodically by using a sliding window approach, i.e. only a limited number of data sets, in particular the latest, is taken into account.

### C. Blocking Data Sets

In order to reduce or even clear the effect of abnormal behaving nodes, we propose to drop malicious data. If a participant does not behave correctly, we do not want him to influence the monitoring mechanism.

We propose to discard suspicious data according to three policies implementing an increasing cautiousness against the impact of malicious data. Discard *none* is the incautious approach which does not drop data. The *discard\_safe* policy only discards data sets which are certainly malicious. These are attacks against the monitoring structure and rated by our instant rating mechanism, thus false positives are impossible, as the participant's behavior is checked. Lastly, *discard\_all* drops each suspicious data set. Please note, due to the use of an outlier detection system, we may drop false positive data.

**Data Set Discard Theory:** If aggregated data sets have to be dropped due to a discard policy, then the data of a complete subtree will be discarded. To have a rough estimation how many data sets are expected to be dropped we introduce a formula modeling the tree with following simplifications: i) all malicious data sets have been found ii) the tree is balanced and iii) complete. Here is the formula:

$$\phi = \text{Round}(\log_{bF}(N * p * (bF - 1))) - 1 \quad (5)$$

$$\sum_{i=0}^{\phi} \left( \frac{bF - 1}{bF^{i+1}} * p * N * \sum_{j=0}^i (bF^j - p * bF^j) \right) \quad (6)$$

Here,  $N$  is the total number of nodes,  $p$  the ratio of malicious nodes and  $bF$  the branching factor of the tree. The idea behind Equation 6 is to traverse the tree bottom-top and calculate on each level how many nodes respectively data sets should be dropped if this node has been detected from his parent. The inner sum calculates how many nodes this dropped sub tree has and subtracts already dropped nodes. The outer sum stops when the probability that a malicious node is located at that level is close to zero.

### V. EVALUATION

To evaluate our developed rating system, we simulated several scenarios using our implementation of SkyEye.KOM, which meets our requirements to the monitoring system (Section I), in the PeerfactSim.KOM simulator [8]. As a preparation, we marked all (aggregated) statistics containing information from malicious nodes with an "evil bit" [1] which tells us the number of malicious nodes that influenced this piece of information. Ideally, all information marked with evil bits should be discarded by DOMiNo.

For the evaluation, we consider the following metrics: i) number of evil bits in the the global view, which equals the detection rate (in short #evilBits), ii) the number of discarded data sets, which gives together with #evilBit a hint about the false positives, iii) precision, where we compare the measured global view with the desired one and iv) at a certain point the costs respectively the fairness of the system's load distribution. Next, we describe the different used simulation parameters along with our scenarios.

### Simulation Parameters and Scenarios

For our evaluation, we simulated beforehand a lot of simulation setups and extracted the most important simulation parameters which significantly influenced the outcome. As the configuration file of the simulator is very extensive we only describe the relevant parameters for our simulation results, which are summarized by the Table I. We simulated a Pastry [25] overlay network consisting of 5000 nodes over a total of 200 minutes without churn. Exemplarily, we let SkyEye.KOM create statistics about the average number of bytes sent per node on the network layer.

The actions performed during a simulation are the same in each scenario. The first 60 minutes are the join phase where all nodes join the overlay and start distributing data in the monitoring system from minute 90 on. In parallel, we had a file sharing application running, publishing data sets between minute 60 to 80 and starting lookups after them every 10 minutes to feed the monitoring system. From minute 100 on, each node starts the rating phase. Note that malicious nodes also use DOMiNo. After feeding DOMiNo with monitoring data for 10 minutes, we get first results at minute 110.

In the scope of this paper we performed an extensive evaluation with 228 parameter combinations, drafted by Table I, with at least 10 seeds each. Out of these, we present six scenarios. In Scenario A (Section V-A) we show the impact of a data manipulating attack with DOMiNo deactivated. The impact of DOMiNo on a non-malicious environment is shown in Scenario B (Section V-B). In Scenario C.1 to C.3 (Section V-C) we expose a parameter study when performing data manipulation attacks with DOMiNo activated, which gives a first impression on how well our proposed system works. Lastly, Scenario D (Section V-D) reveals the outcome when releasing all attacks with DOMiNo activated while discarding all malicious rated data and discarding only data which is marked as assured malicious, respectively. At the end, in Section V-E, we present a discussion on our proposed system's characteristics.

Table II presents the simulation outcome ordered by our scenarios. For visualization purposes for each metric we present the absolute and relative error of the simulations average from minute 110 till 200 ( $\Delta_{\phi}$  and  $\delta_{\phi}$ ). Additionally, for the data set discard count and global view metrics we exhibit the standard deviation, denoted as  $\sigma_X$  with  $X$  being the value we create the standard deviation for. To enhance the overview, for each scenario and metric we color the best error value with light gray respectively the worst with darker gray.

#### A. Scenario A: Manipulation attacks / DOMiNo deactivated

Scenario A demonstrates the impact of manipulation attacks on an unprotected monitoring system. The simulations' outcome can be taken from Table II. Here, the branching factor (=  $bF$ ) is set to 96 for future comparisons and we vary the malicious ratio and malicious data manipulation factor. We encounter no thrown away data sets and  $5000 * m\_ratio$  evil bits in the global view. The global view's mean attribute shows we can arbitrarily influence the global view, which scales with the number of malicious nodes and their manipulation factor.






Fig. 3: Per peer sent bytes on net layer

referring to Scenario C.3 we conclude that the parameter setting  $bF=96$  and Z-score threshold=3.5 works best in terms of precision and false positive discards.

Lastly, in Scenario C.3, we apply optimal parameters for  $bF$  and Z-score threshold (96/3.5) and perform a parameter variation on the malicious ratio and manipulation factor. Throughout all malicious ratios, we can observe similar effects like in Scenario C.2. For the small multiplication factor only 30% of all evil data sets have been discarded but the precision is quite decent compared to Scenario A. For malicious ratios of 5%, 10% respectively 20% we encounter the following relative errors: 1.6% / 6.2% / 17.5% with DOMiNo enabled and 11.4% / 23.8% / 52.5% with DOMiNo disabled. With a higher manipulation factor the full potential is unfolded and we hit a relative error of -0.9%, 0.8% and 9.4% in Scenario C.3 with DOMiNo enabled compared to 118.6% / 334% / 997.5% in Scenario A without protection. As the number of evil bits is close to zero, the discard Equation 6 can be applied, which supplies 473 for 5%, 895 for 10% and 2392 for 20%.

We conclude that the influence of malicious manipulation attacks is limited to the bounds of the convex hull around the median. Thus, small changes can influence the precision, but the degree of influence is limited. Furthermore greater manipulation factors are filtered out efficiently by DOMiNo.

#### D. Scenario D: All attacks / DOMiNo activated

In our last scenario, we apply the parameters found till Scenario C.3 ( $bF=96$  / Z-score threshold=3.5) and see how the system is performing with all attacks enabled. Here, we want to figure out how the malicious node ratio, the discard policy and the manipulation factor influence DOMiNo's efficiency. Preliminary we have to state, if  $p\%$  of all nodes behave maliciously then these nodes are equally distributed on the active attacks, which are nine (see Section III-B).

When applying the discard\_all policy, DOMiNo clears more evil bits as it significantly discards more data sets. Furthermore, the relative error for the global view is smaller or equals the discard\_safe policy's one. Increasing the malicious node ratio parameter while fixing the manipulation factor to two seems to have not a great impact for the discard\_all policy. In numbers we got -2.9%, 3.5% and 1.3% for discard\_all and 2%, 4.1% and 8.4% for discard\_safe. For the manipulation factor of ten we encounter a peak when simulating with 20% of

malicious nodes with a relative error of 63.2% for discard\_all and inf (=infinity) for discard\_safe while the other ratios perform much better. Here we encounter -3.2% vs. 22.5% for 5% and -3.8% vs. 41.8% for 10% of malicious nodes for discard\_all respectively discard\_safe. The infinite value is most probably due to the  $AT_{Root}$  and  $AT_{Random}$  attacks. In the first attack,  $\frac{5000 \cdot 0.2}{9} = 111$  nodes send manipulated data to the root node which has  $bF=96$  honest child nodes, which supply data. This imbalance causes the system to crash if there also exist a loop in the tree due to the second attack. Owing to our requirement, that children must be located in a certain ID space, we are not able to detect all structural attacks. Keep in mind that loops are still possible due to the goal of having a passive rating, which does not allow to verify the potential child node which pushes monitoring data. Lastly, a comparison to an unprotected system is not useful as the relative error permanently increases, due to no mechanisms to detect loops in the monitoring tree.

#### E. Discussion

In this section we want to discuss and emphasize a few characteristics of DOMiNo.

In our proposed system attackers can only inject data which fit into their parents' continuously calculated bounds, as anything beyond these bounds is dropped by default. Please note that malicious nodes' data can be also within the convex hull and vice versa, deviated honest nodes' data may be considered as outliers and may be emitted. But as long as the majority of the nodes respectively each parents' nodes report similar data, our convex hull ensures that outlier will be filtered out.

Also, it is quite difficult to manipulate data which is still plausible and accepted by the parent, as nodes can not obtain their parent's data base to calculate the current convex hull. So single attackers must guess their parent's data base or control respectively cooperate with all other children of a particular parent. Though, assuming attackers are able to calculate their parent's convex hull, the influence on the global view greatly depends on the number of nodes performing this kind of attack. So it is possible to inject plausible but false data, and by doing this shifting the global data to the attacker's desire, if attackers 1) are able to calculate their parents' convex hull and 2) are cooperating with other attackers evenly distributed throughout the monitoring tree. Please note, that especially the first assumption is hard to achieve.

In conclusion, our proposed system DOMiNo performs efficiently, if the majority ( $> 50\%$ ) of the participants are honest, parent nodes can fully verify their children nodes and children are not able to calculate their parents' convex hull bounds. If parent nodes are not able to verify their children nodes, DOMiNo's reliability decreases when the number of cooperating malicious nodes exceeds the branching factor of the monitoring tree, i.e. the number of potentially benign children of a parent node.



## VI. CONCLUSION AND FUTURE WORK

In this paper, we presented our security system for tree-based monitoring systems with certain requirements to their architecture. The system's main objectives are to serve the service by only passively listening to incoming events, i.e. monitoring data sets, and to minimize the relative error to the global view by keeping the costs low. We developed an attacker model including manipulation attacks, structural attacks and free-riding patterns and proposed our rating system, called DOMiNo, which includes an outlier detection based on Z-Scores, structural and other plausibility checks. We implemented it in the P2P simulator PeerfactSim.KOM [8].

We evaluated our system with 228 configurations with at least 10 seeds each in a Pastry overlay with 5000 nodes. For the attackers the configurations cover variations of the malicious node ratio, their used attack types and, for manipulation attacks, the manipulation factor. On the detecting side we alter the outlier detection's Z-Score threshold, the discard policy after positively rating a node's data set and, for the monitoring structure, the branching factor. The outcome shows that we are able to limit the manipulation attackers' influence on the global view by defining a convex hull for acceptable values. Thus, normal fluctuations are tolerated but outlier, that harm the global view, can be filtered out effectively. For small manipulation factor attacks we face a relative error between 11% to 53% in an unprotected monitoring environment versus 2% to 18% with DOMiNo enabled. For higher manipulation factor attacks the unprotected system's relative error raises to 998% while we can correct it to 9%. Enabling the whole attacker model's repertoire we are able to decrease the relative error to values between 1% to 4% from potentially infinity, caused by loops through structural attacks. If we encounter more than the tree's branching factor cooperating malicious nodes, they are potentially able to break the system. To circumvent this, we propose stricter tree structural properties, where parent nodes can verify their children nodes' positions.

In the future we will work on an approach to loosen the passive rating requirement and compare it with DOMiNo. By doing this, we also plan to investigate systems with machine learning components.

## ACKNOWLEDGMENT

We would like to thank Sebastian Brink for his great work.

## REFERENCES

- [1] "The Security Flag in the IPv4 Header," RFC 3514, Jun. 2015. [Online]. Available: <https://rfc-editor.org/rfc/rfc3514.txt>
- [2] S. Axelsson, "Intrusion Detection Systems : A Survey and Taxonomy," *Computer Engineering*, 2000.
- [3] R. Bhagwan, G. Varghese, and G. Voelker, "CONE: Augmenting DHTs to Support Distributed Resource Discovery," Tech. Report CS2003-0755, University of California, San Diego, 2003.
- [4] H. Chan, A. Perrig, and D. X. Song, "Secure hierarchical in-network aggregation in sensor networks," in *Proc. of the Conference on Computer and Communications Security, Alexandria, VA, USA*, 2006.
- [5] M. Dam and R. Stadler, "A Generic Protocol for Network State Aggregation," in *Proc. of the Radiotekenskap och Kommunikation*, 2005.
- [6] A. Disterhoft and K. Graffi, "Protected chords in the web: Secure p2p framework for decentralized online social networks," in *Proc. of the IEEE Int. Conference on Peer-to-Peer Computing*, 2015.
- [7] K. Graffi, "Monitoring and Management of Peer-to-Peer Systems," Ph.D. dissertation, Technische Universität, Darmstadt, August 2010. [Online]. Available: <http://tuprints.ulb.tu-darmstadt.de/2248/>
- [8] —, "PeerfactSim.KOM: A P2P System Simulator - Experiences and Lessons Learned," in *IEEE P2P '11: Proceedings of the International Conference on Peer-to-Peer Computing*, 2011.
- [9] K. Graffi, C. Groß, D. Stingl, D. Hartung, A. Kovacevic, and R. Steinmetz, "LifeSocial.KOM: A Secure and P2P-based Solution for Online Social Networks," in *Proc. of the IEEE Consumer Communications and Networking Conferences*, 2011.
- [10] K. Graffi, A. Kovacevic, S. Xiao, and R. Steinmetz, "SkyEye.KOM: An Information Management Over-Overlay for Getting the Oracle View on Structured P2P Systems," in *Proc. of the Int. Conf. on Parallel and Distributed Systems*. IEEE, 2008.
- [11] K. Graffi, P. Mukherjee, B. Menges, D. Hartung, A. Kovacevic, and R. Steinmetz, "Practical security in p2p-based social networks," in *Proc. of the IEEE Conference on Local Computer Networks, LCN*, 2009.
- [12] K. Graffi, D. Stingl, J. Rückert, and A. Kovacevic, "Monitoring and Management of Structured Peer-to-Peer Systems," in *IEEE P2P '09: Proc. of the Int. Conference on Peer-to-Peer Computing*, 2009.
- [13] C. Hayward and A. Madill, "A Survey of Outlier Detection Methodologies," *Social Science and Medicine*, 2003.
- [14] IBM, "IBM Knowledge Center," [http://www.ibm.com/support/knowledgecenter/SSWLVS\\_1.0.0/com.ibm.spss.analyticcatalyst.help/analytic\\_catalyst/modified\\_z.html](http://www.ibm.com/support/knowledgecenter/SSWLVS_1.0.0/com.ibm.spss.analyticcatalyst.help/analytic_catalyst/modified_z.html), accessed: 2016-06-10.
- [15] B. Iglewicz and D. Hoaglin, *How to Detect and Handle Outliers*, ser. ASQC basic references in quality control. ASQC Quality Press, 1993.
- [16] M. Jelasity, A. Montresor, and Ö. Babaoglu, "Gossip-based aggregation in large dynamic networks," *ACM Trans. Comput. Syst.*, 2005.
- [17] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "The Eigentrust algorithm for reputation management in P2P networks," in *Proceedings of the Twelfth International World Wide Web Conference, WWW 2003, Budapest, Hungary, May 20-24, 2003*, 2003, pp. 640–651.
- [18] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-Based Computation of Aggregate Information," in *IEEE FOCS '03: Proc. of the Symposium on Foundations of Computer Science*, 2003.
- [19] S. Marti and H. Garcia-Molina, "Taxonomy of trust: Categorizing P2P reputation systems," *Computer Networks*, vol. 50, no. 4, pp. 472–484, 2006.
- [20] S. Marti, T. J. Giuli, K. Lai, and M. Baker, "Mitigating routing misbehavior in mobile ad hoc networks," in *Proc. of the Annual Int. Conference on Mobile Computing and Networking*, 2000.
- [21] P. S. Mogre, K. Graffi, M. Hollick, and R. Steinmetz, "AntSec, WatchAnt, and AntRep: Innovative Security Mechanisms for Wireless Mesh Networks," in *Proc. of Int. Conf. on Local Comp. Networks*, 2007.
- [22] A. G. Prieto and R. Stadler, "Distributed Real-Time Monitoring with Accuracy Objectives," in *Proc. of the Int. Net.g Conf. 2006 - New. Techn., Services, and Protocols; Performance of Comp. and Comm. Net.; Mobile and Wireless Communications Systems*, 2006.
- [23] V. Rapp and K. Graffi, "Continuous Gossip-based Aggregation through Dynamic Information Aging," in *Proc. of the Int. Conference on Computer Communications and Networks*, 2013.
- [24] M. Rounds and N. Pendgraff, "Diversity in network attacker motivation: A literature review," *Proc. of the IEEE Int. Conference on Computational Science and Engineering*, 2009.
- [25] A. I. T. Rowstron and P. Druschel, "Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems," in *IFIP/ACM Middleware '01: Proc. of the Int. Conference on Distributed Systems Platforms*, 2001.
- [26] University of Oldenburg, "Cohda project," <https://www.uni-oldenburg.de/en/computingscience/ui/research/topics/cohda>, accessed: 2016-06-10.
- [27] F. Wuhib, M. Dam, and R. Stadler, "Decentralized Detection of Global Threshold Crossings using Aggregation Trees," *Comp. Networks*, 2008.
- [28] P. Yalagandula and M. Dahlin, "Research Challenges for a Scalable Distributed Information Management System," The Univ. of Texas at Austin, Dept. of Computer Sciences, Tech. Rep. CS-TR-04-48, 2004.
- [29] Y. Yang, X. Wang, S. Zhu, and G. Cao, "SDAP: A secure hop-by-hop data aggregation protocol for sensor networks," *ACM Trans. Inf. Syst. Secur.*, 2008.
- [30] Z. Zhang, S. Shi, and J. Zhu, "SOMO: Self-Organized Metadata Overlay for Resource Management in P2P DHT," in *IPTPS '03: Proc. of the Int. Workshop on Peer-To-Peer Systems*, 2003.