RDAS: Reputation-based Resilient Data Aggregation in Sensor Networks

Carlos R. Perez

US Patent and Trademarks Office

Rajesh K. Panta, Saurabh Bagchi

Dependable Computing Systems Lab (DCSL)
School of Electrical & Computer Engineering
Purdue University

Reliability of Sensor Networks

- Sensor networks have many applications for interfacing the physical environment with the cyber world
- The goal is to design miniature devices to cooperatively monitor physical or environmental conditions in possibly remote and inaccessible locations
- Design constraints present challenges in maintaining the reliability of these networks
  - Unreliable radio communication
  - Limited computational power
  - Remote and unsupervised environment requires fault and intrusion tolerance mechanisms

Data Aggregation

- Primary task of sensor networks is to gather large amounts of data
- Independent communication from nodes to a collecting station becomes infeasible as network size grows
- Data aggregation is a prevailing approach to reduce the amount of redundant communication during the collection of sensor readings

Event Localization

- A highly desirable task of sensor networks is the detection and localization of events
- Event localization is an example of data aggregation
  - Nodes forward sensor readings to aggregator node
  - Aggregation consists of computing the location of the event
  - Event location is sent to a base station

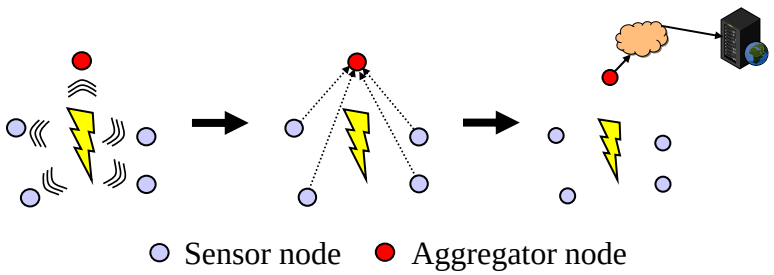

Problems with Data Aggregation

- Vulnerabilities of data aggregation
  - Faulty sensing nodes
  - Faulty aggregator nodes
  - Authentication is not enough
    - Legitimate nodes can be compromised in a WSN
    - Sensors in legitimate nodes can become damaged
- Problem Specification
  - Create a system that improves the accuracy of data aggregation in the presence of faulty and adversarial nodes

Solution Approach: Reputation System

- A reputation system collects, distributes and aggregates feedback about participants' past behavior
- They are commonly used for rating the quality of users and content in Internet communities and websites
  - eBay, Amazon, Digg, YouTube
- Advantages in sensor network context
  - Provides the ability to detect and isolate nodes behaving inappropriately
  - Allows nodes that have not previously interacted to know about each other's past behavior
- Challenges
  - Preventing false accusations or false praise by some nodes to alter the reputation of others
  - Distinguishing faulty activity from natural error

Contributions

- Designed a generic method of effectively generating, propagating and using reputation to secure data aggregation
- Applied this reputation system to improve the accuracy of event localization in the presence of adversaries
- Used a generalized fault model
  - Any node can be compromised
  - Nodes can collude in their adversarial behavior

Outline

- Introduction and Motivation
- **Reputation system**
- Data aggregation
- Experiments and Results
- Conclusions

Representing Reputation

- The *reputation rating* $R_{i,j}$ is the opinion a node ***i*** has about a node ***j*** regarding its performance in the system
- Reputation can be treated as the probability of future good behavior by quantifying past behavior as the parameters of a beta distribution.
  - Parameter α – count of good behavior
  - Parameter β – count of bad behavior
- The expected value of the beta distribution over past observations gives an estimate of the probability of good behavior into the future

$$R_{i,j} = E(Beta(\alpha + 1, \beta + 1)) = \frac{\alpha + 1}{\alpha + \beta + 2}$$

PURDUE UNIVERSITY

Generating Reputation

- Reputation is unique among any pair of nodes (*i,j*)
- A node generates reputation ratings from two sources
  - First-hand information: observations made by node
  - Second-hand information: information reported by other nodes
- Second-hand information is vulnerable to false accusation or false praise attacks
  - *Bad-mouthing*: maliciously trying to lower the reputation of well-behaving nodes
  - *Ballot-stuffing*: maliciously trying to increase the reputation of bad-behaving nodes
- Preventing these attacks is essential to enable use of second-hand information

PURDUE UNIVERSITY

Trust

- The *trust rating* $T_{i,j}$, is the opinion i has about j regarding its performance in reporting about other nodes
- Behavior in reporting is measured by comparing reputation reports with current reputation rating

$$a = \begin{cases} 1 & if \mid R_{i,j} - E((Beta(r_{k,j}, s_{k,j})) \mid < d \\ 0 & if \mid R_{i,j} - E((Beta(r_{k,j}, s_{k,j})) \mid \geq d \end{cases}$$

$$\gamma := v\gamma + a$$
$$\delta := v\delta + (1-a) \qquad T_{i,j} = E(Beta(\gamma+1, \delta+1)) = \frac{\gamma+1}{\gamma+\delta+2}$$

- High-level idea
  - Reputation represents confidence in entity's performance
  - Trust represents confidence in entity's reporting of performance
  - Performance reports contrary to current belief initially affect trust of reporter
  - As reports are confirmed its trust is reestablished

Fusing reputation reports

- Given first-hand observations $(r_{i,j}, s_{i,j})$ and second-hand reports $(r_{k,j}, s_{k,j})$ from N nodes about node j we fuse them as follows

$$\alpha_{i,j}^{new} = u\alpha_{i,j} + r_{i,j} + \sum_{k \in N} D(r_{k,j});$$

$$D(r_{k,j}) = \frac{\{2 * \gamma_{i,k} * r_{k,j}\}}{\{(\delta_{i,k} + 2) * (r_{k,j} + s_{k,j} + 2)\} + \{2 * \gamma_{i,k}\}}$$

- Old reputation is aged
- FH report is counted completely
- SH report is weighted using Dempster-Shafer belief discounting
  - Trust rating is used to weight down reputation report

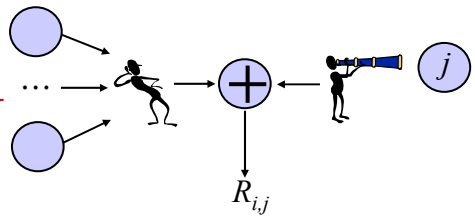

Outline

- Introduction and Motivation
- Reputation system
- **Data aggregation**
- Experiments and Results
- Conclusions

Data Aggregation: LEACH

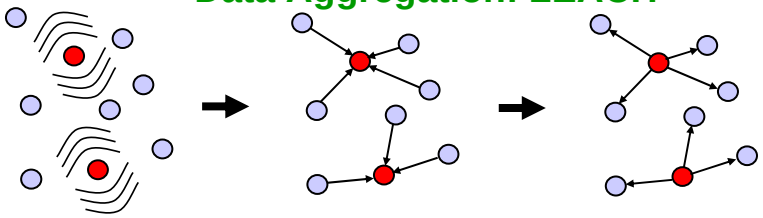

1. Cluster heads broadcast an advertisement

2. CN's compare RSSI's, send a join request to best CH and adjust tx strength of radio

3. CH's reply with a TDMA schedule for CH-CN communication

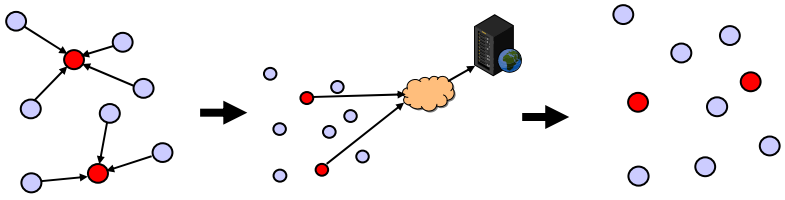

4. CN's send data during their allocated TDMA slot

5. CH's aggregate data and transmit to base station

6. New CHs are elected and a new round starts

Fault Model

- No trusted infrastructure – any node can be compromised
- Basic cryptography is in place to authenticate node-to-node communication
- Nodes know estimated location of all nodes within communication range
- Compromised nodes can exhibit the following behavior
  - *Faulty*: Non-colluding or Colluding
  - *Liar*: reports false reputation reports
  - *Malicious*: both faulty and liar behavior
  - *Malicious Cluster Head*: Alters or drops data after aggregation
- We assume in the initial state no nodes are compromised
  - Nodes become compromised at a finite rate

RDAS: Data Aggregation by Cluster Head

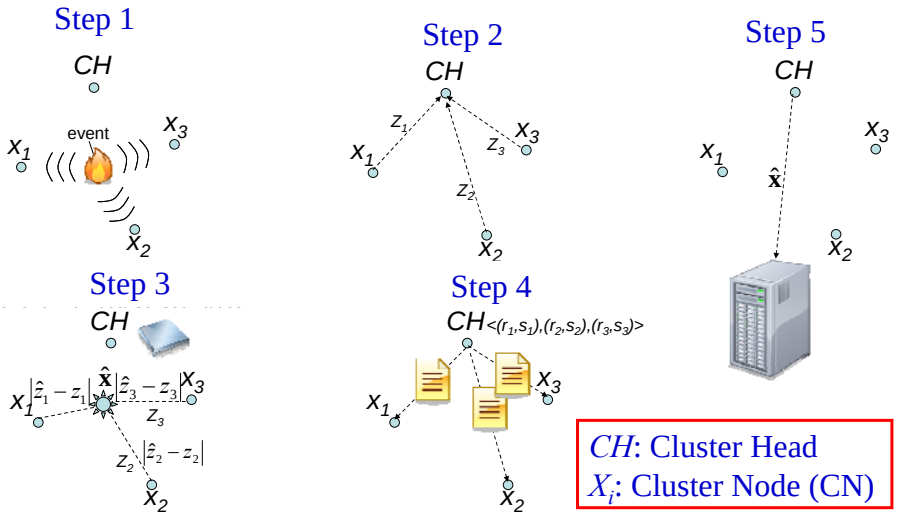

Cluster Head Attack

- Protocol is still vulnerable to CH modifying or dropping data
- Cluster Monitors (CMs)
  - Elected at time of cluster head election
  - Overhear CN-to-CH communication, perform event localization, and generate data reputation reports
  - Overhear CH-to-Base communication and generate cluster head reputation reports
- Separate reputation rating for cluster heads
  - Generated and propagated by CM
  - Used during CH election to filter out malicious nodes

Outline

- Introduction and Motivation
- Reputation system
- Data aggregation
- **Experiments and Results**
- Conclusions

Simulation

- Protocol was implemented in TinyOS 2.0
- Simulated using TOSSIM 2.0, which runs TinyOS code on a PC
- Simulation details
  - Nodes placed in uniformly distributed random locations
  - Every nodes knows locations of neighbor nodes with normal error

Network parameters

Parameter	Value
Sensor field size(m)	300x300
Tx range (m)	100
# Nodes	50
# Clusters	~3

Reputation parameters

Parameter	Value
k constant	1.414
Filtering threshold (*FT*)	0.5
Reputation Aging	0.99
Trust deviation (*d*)	~3

Experiments

- Network becomes compromised in 6% increments at a fixed rate
- Experiments conducted
  1. Faulty nodes, compromised up to 72%
  2. Malicious nodes, compromised up to 72%
- Experiments were repeated for non-colluding and colluding nodes

Metrics

- Localization Error – distance between event location estimate and real event location
- Accuracy – ratio of reduction in event localization error
  - Shorthand for relative improvement in localization accuracy
  - Range: $(-\infty,1]$
  - Accuracy above 0 indicates improvement over baseline system

$$Accuracy = 1 - \left(\frac{\text{Localization error using reputation}}{\text{Localization error using baseline}} \right)$$

- Average reputation and trust rating
  - Separately computed for legitimate and compromised nodes

PURDUE UNIVERSITY

Experiment 1: Faulty Nodes, non-colluding

Localization Error with and without reputation at 6% per 800 events

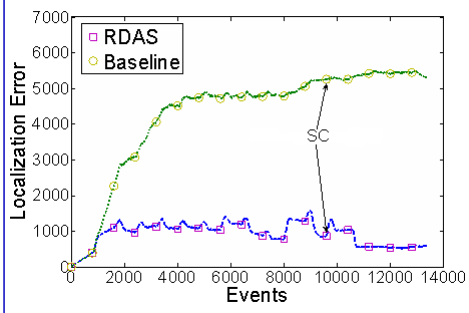

Accuracy with reputation

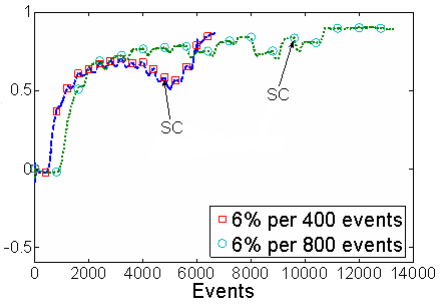

- Accuracy stays above 0 as network is becoming compromised
- In steady-state with 72% compromise, accuracy stays at 90%
- Addition of compromised nodes causes transient increases in error

PURDUE UNIVERSITY

Experiment 1: Faulty Nodes, colluding

Localization Error with and without reputation at 6% per 800 events

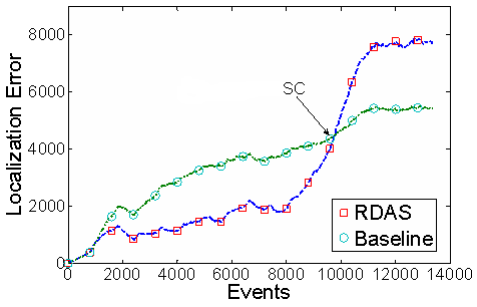

Accuracy with reputation

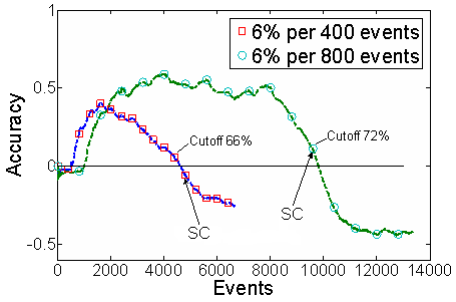

- Collusion affects accuracy but system still improved localization with > 50% compromised
- Accuracy stays above 0 up to 72% of network compromised
- Accuracy drops below 0 only after very high localization error

Experiment 1: Faulty Nodes

Avg Reputation with non-colluding Nodes at 6% per 800 events

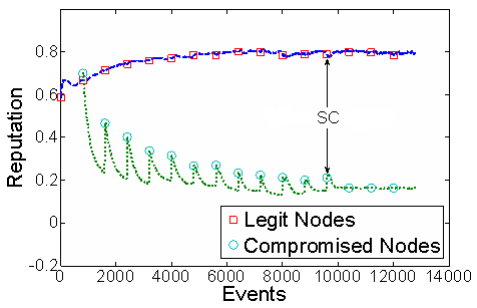

Avg Reputation with colluding Nodes at 6% per 800 events

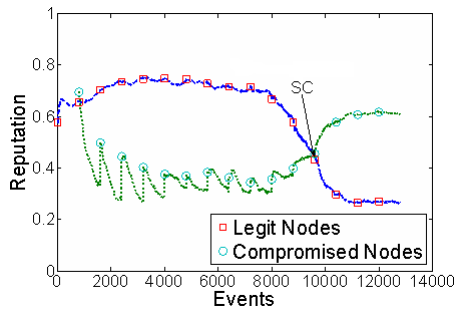

- Collusion directly affects reputation
- System performs better than baseline as long as reputation of legitimate nodes stays above compromised nodes

Experiment 2: Malicious Nodes

Accuracy, non-colluding

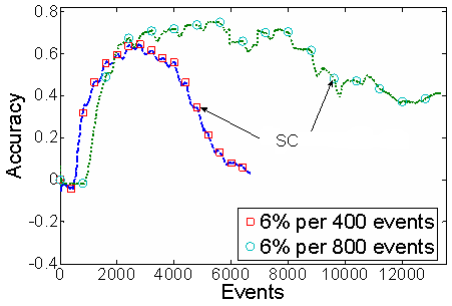

Accuracy, colluding

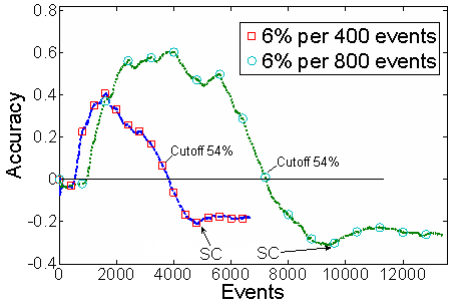

- Liar nodes affect accuracy
- System always outperforms baseline for non-colluding nodes
- System outperforms baseline up to 54% for colluding nodes

Insights

- Reputation systems are a feasible way of improving the reliability of data aggregation in wireless sensor networks
- Generating reputation by analysis of reported data allows the accuracy of event localization to be improved even for substantial percentages of the network compromised
- Influence of liar nodes is mitigated by keeping track of behavior in the reporting of reputation ratings
- Results provide insight into the effect of compromise rates in data aggregation accuracy

Backup Slides

The Beta Reputation System

- The behavior of a node can be described as a binary event, since the possible outcomes are cooperative or non-cooperative behavior
- The Beta distribution is used to represent posteriori probability distributions of binary events
- It is parameterized by two positive integers numbers, which represent the prior outcomes of the observed event
- For example:
  - A process with 2 possible outcomes $\{\alpha,\beta\}$ has produced outcome α seven times and β once
  - The pdf of observing outcome α based on past experience is expressed as $f(p \mid 8, 2)$

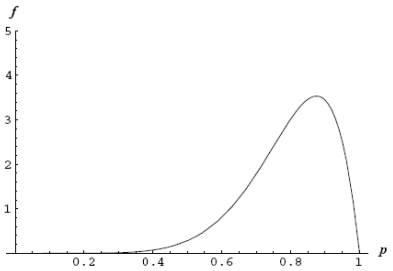

[1]

LEACH Overview

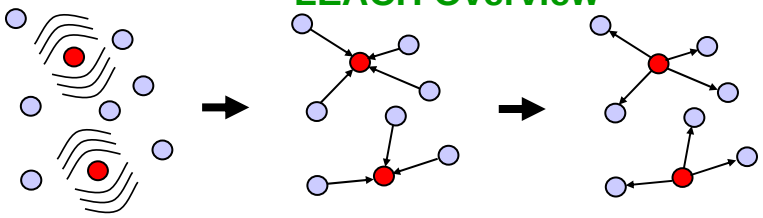

1. Cluster heads broadcast an advertisement

2. CN's compare RSSI's, send a join request to best CH and adjust tx strength of radio

3. CH's reply with a TDMA schedule for CH-CN communication

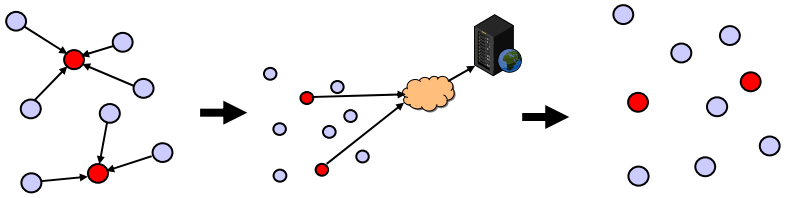

4. CN's send data during their allocated TDMA slot

5. CH's aggregate data and transmit to base station

6. New CHs are elected and a new round starts

Event Localization

- We assume nodes follow the sensor model given by:

$$z_i = \frac{a}{\|\mathbf{x} - \xi_i\|^2},$$

  - Where z_i is the sensed value, a is fixed amplitude, ξ_i is node location and $\mathbf{x}$ is the event location

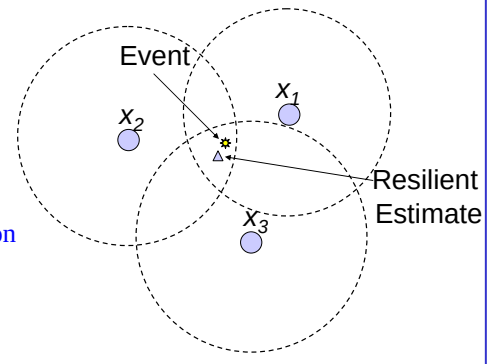

- Location is computed by generating a system of redundant equations computing the least-squares fit

$$\hat{\mathbf{x}} = \left(\mathbf{A}^T\mathbf{A}\right)^{-1}\mathbf{A}^T\mathbf{b}$$

Assumptions

- All nodes are of identical capacity (no super-nodes or trusted infrastructure)
- Any node can become compromised
- Basic cryptography is in place to authenticate node-to-node communication
- Nodes know estimated location of all nodes within communication range
- Events of known amplitude can occur at any location within sensing range
- Network density is large enough that events will be detected by most nodes in the same cluster

PURDUE UNIVERSITY

Overview of Securing Data Aggregation

- Cluster heads generate reputation for nodes in the network
  - Collect data from CNs and estimate the event location
  - Use estimated event location to measure accuracy of individual nodes' sensor data values
  - Generate a reputation report about each node
  - Broadcast report to the cluster
- Cluster heads use reputation to improve accuracy
  - Eliminate nodes with low reputation from the computation of event localization

PURDUE UNIVERSITY

Generating Reputation

- A reputation rating is kept for performance in reporting of data
  - Good performance: $a = 1$
  - Bad performance: $a = 0$
- Data aggregation
  - Using the sensor model and event location estimate, the CH determines what sensor data value ($\hat{z}_i$) should have been reported by every node
  - With this estimate it calculates the error percentage of each node's sensor measurement
- By treating E_i as a random variable with mean μ, it tests Chebyshev's inequality to determine if the node's error falls within the distribution of error for the cluster

$$r_i := r_i + a$$
$$s_i := s_i + (1 - a)$$

$$E_i = \frac{\left|\hat{z}_i - z_i\right|}{\hat{z}_i}$$

$$a = \begin{cases} 1 & \text{if } \left|E_i - \mu\right| \le k\sigma \\ 0 & \text{otherwise} \end{cases}$$

$$\mu = mean(E_i) \,\forall\, i$$
$$\sigma = \text{StdDev}(E_i) \,\forall\, i$$

Resilient Event Localization

- Event Detection
  - Weight vote of each node by reputation
- Event Localization
  - Eliminate nodes with reputation below filtering threshold from redundant system of equations

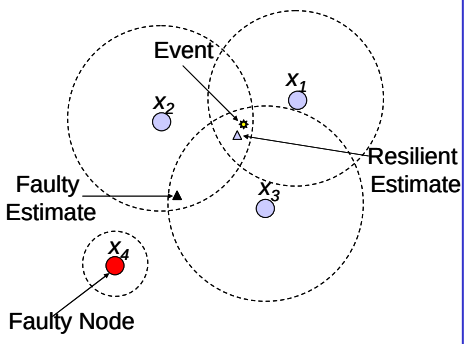

Experiment 2: Malicious nodes, non-colluding

Avg Reputation with non-colluding nodes at 6% per 800 events

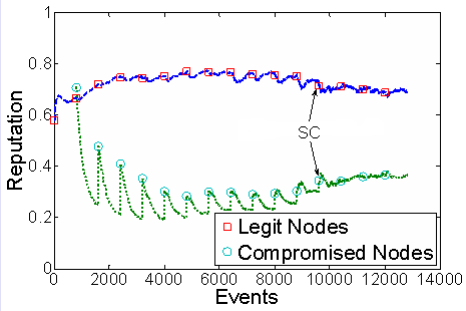

Avg Trust with non-colluding nodes at 6% per 800 events

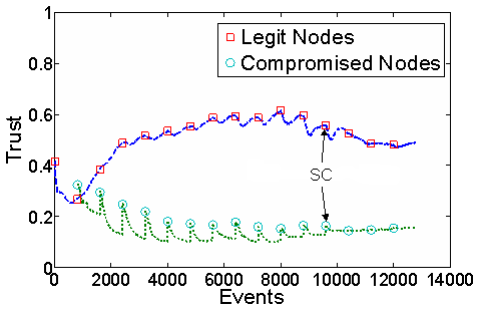

- Presence of liar nodes slightly affects reputation and trust

Experiment 2: Malicious nodes, colluding

Avg Reputation with colluding nodes at 6% per 800 events

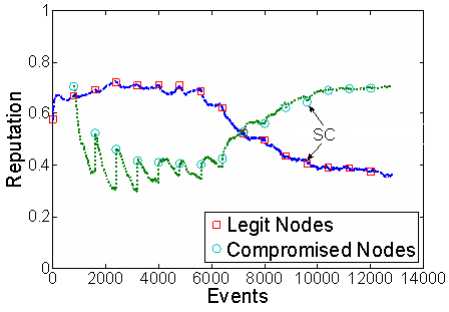

Avg Trust with colluding nodes at 6% per 800 events

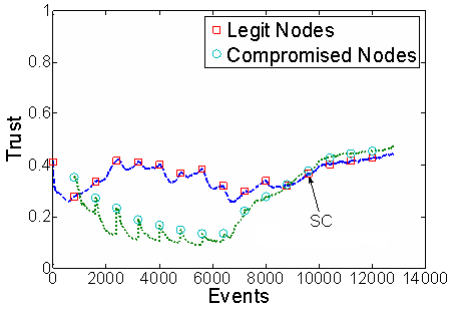

- Increase in malicious nodes affects cutoff point of reputations, 54% from 72% without liars
- Trust ratings cross at 66%, but become close at 54%

Experiment 3: Malicious Nodes, cap at 30%

Localization Error, colluding

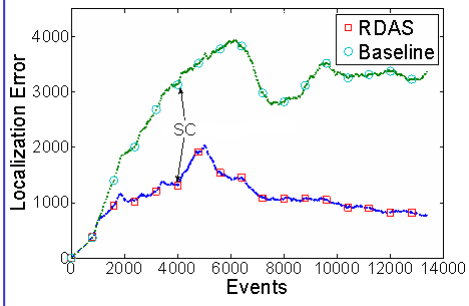

Accuracy, colluding

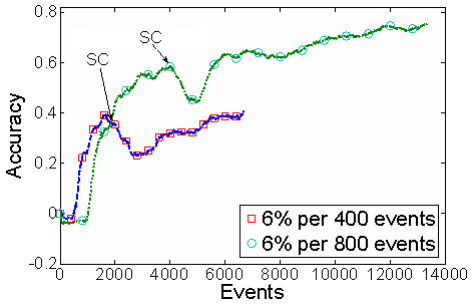

- Localization error decreases while it stabilizes for baseline system
- Securing portion of network guarantees low error

Experiment 3: Malicious nodes, cap at 30%

Avg Reputation with colluding nodes at 6% per 800 events

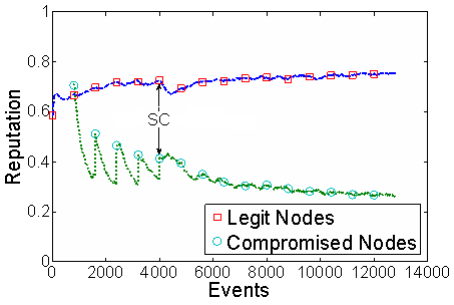

Avg Trust with colluding nodes at 6% per 800 events

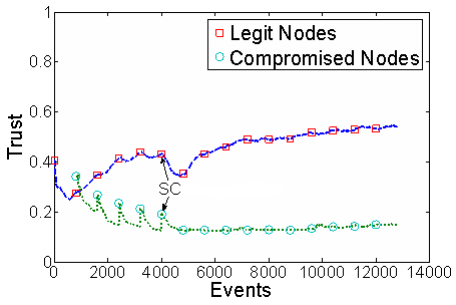

- Reputations and trusts diverge with time as no new malicious nodes enter