Title

Inspect: a general framework for on-line detection and diagnosis of sensor faults

Authors

Bose, Subhonmesh
Tulone, Daniela
Zahedi, Sadaf
et al.

Publication Date

2007-10-10

Peer reviewed

Center for Embedded Networked Sensing

Inspect: A General Framework for On-Line Detection and Diagnosis of Sensor Faults

Subhonmesh Bose, Daniela Tulone, Sadaf Zahedi, Laura Balzano, Mani Srivastava

NESL-http://nesl.ee.ucla.edu/

Introduction: Detection of the Sensor Faults

Sensor Faults and misbehaviors

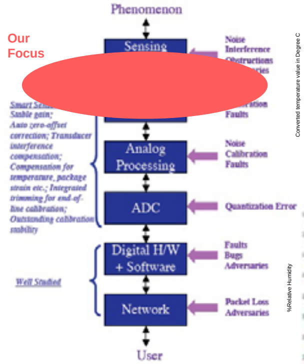

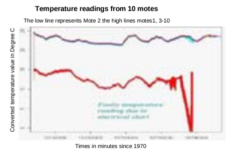

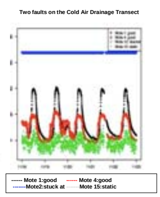

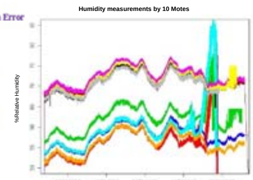

Traditional Fault Detection Architecture

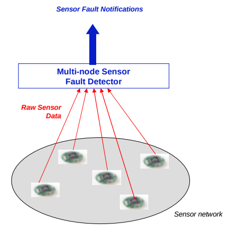

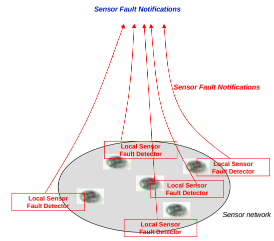

Centralized Fault Detection

- Energy and communication overhead as all sensor data extracted to a central node
- Higher fault detection latency

Distributed Fault Detection

- Sensor integrity checks based solely on local measurements is suboptimal
- Effective fault detection usually requires comparisons against readings at other sensors

Framework Overview: Inspect : Tiered Model-based fault detection Architecture

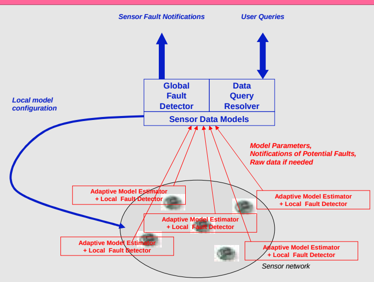

- Fault detection and diagnosis distributed across a ***local tier*** and a ***global tier***
- Each sensor node learns a model for its measured time series of sensor values, and sends model parameters to sink
- Sensor compares new measurements against the model to decide when to adapt the model and send updated parameters
- Sensor uses local integrity checkers to detect potential faults, and sends notifications of potential faults to the sink
- Sink uses models to form a global view of the phenomenon for a
  - (i) answering sensing queries
  - (ii) verifying potential fault reports and their geographical scope
- Key attributes: tiered architecture, and joint querying and integrity checks

The Inspect System Model: Local Fault Detection at Sensors and Global Fault Detection at sink, based on the Physical Model of the Phenomena

Modeling the Physical Phenomena

Current Approach: Time-series Forecasting

- Phenomenon F(t): trend function plus a residual process
- Trend m_t modeled as a low order polynomial (linear)
- Residual *X(t)* modeled as a *weakly stationary* AR(q) autoregressive time series
  - mean and variance are time invariant with zero mean Gaussian noise
  - small q ∈ [1,7] ensures cheap learning / re-learning & compactness

$$F(t) = m_t + X(t)$$
$$X(t) = \alpha_1 X(t-1) + \ldots + \alpha_q X(t-q) + b(\omega)N(0,1)$$

Temperature data with diurnal cycles

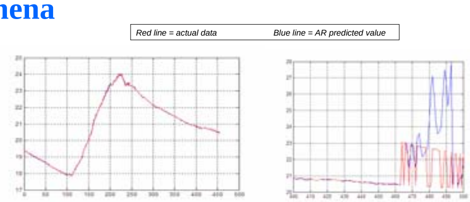

- Training phase for learning model coefficients that are sent to sink
- Prediction quality monitored, and model adapted in case of persistent deviations between the actual and predicted values
- Problems: erratic behavior near outliers and noise
  - Better models?

Local Fault Detection at Sensor Nodes

- Detectors for various faults (stuck-at, calibration etc.)
- Key: detect deviation from normal behavior
  - Local detection prevents inter-nodal analysis
  - Limited resources prevent detailed histories and complex detectors
- Using time series model for fault detection as well?
  - Works well for compression but too erratic for discriminating faults from outliers and noise
- Approach: sensors with cyclical variations
  - Divide cycle (e.g. day) in to time slots
  - Learn mean, variance, and trend (kendall-τ correlation) statistics for each slot
  - Consistent deviation from the statistical model results in notification of potential faults
  - Feedback from sink to refine the model
- Approach: sensors without definite cyclical behavior
  - Smooth using {Spline fitting, Median smoothing, Moving average} and detect change in distribution of residuals (assumed Gaussian)

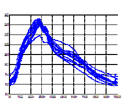

Temperature data with diurnal cycles

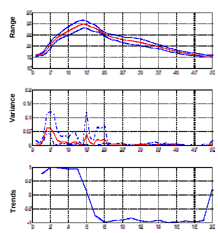

Length of stuck-at fault	10	25	50
Faults injected	45	45	15
Faults detected correctly	4	23	15
False positive	3	8	10
False negative	41	22	0

Length of high-f noise	10	30	50
Variance	0.05	0.10	0.10
Faults injected	15	15	15
Faults detected correctly	8	10	14
False positive	7	4	7
False negative	7	5	1

Global Fault Detection at Sink

- Some faults are impossible to detect without inter-node analysis, e.g. calibration faults
- Most other faults require inter-node analysis or global information to resolve ambiguity
  - e.g. stuck-at-0 light sensor vs. snow cover
- Main idea: sink combines models from individual sensors to create a model of the ground truth
- Approach #1: cluster nodes with similarly valued measurements into groups
  - calculate average divergence for neighbors
  - cluster s.t. maximum divergence < threshold
  - a "***Virtual Reference Source***" represents each cluster, and used to detect faults and verify fault reports from individual sensors
- Approach #2: model correlation between neighbors
  - trajectory of a node relative to neighbors' modeled using local linear regression

    $(x(t-k), y(t-k))$ for $k = 1, 2, \ldots, regWindow$

  - detect deviation from value predicted by neighbors' values

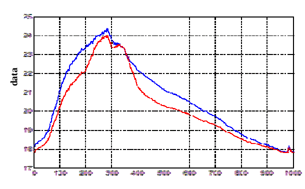

Data with calibration fault injected into one node
Red and Blue lines represent neighboring nodes

Y=aX + b from t=701.
'a' was varied from 1 to 2 in steps of 0.01
'b' was varied from -0.2 to +0.2 in steps of 0.05

Results

Cases tested = 909
False negatives = 8
False positives = 0 for this particular dataset