Using Hidden Semi-Markov Models for Effective Online Failure Prediction

Felix Salfner and Miroslaw Malek

SRDS 2007

Presented by: Fahad Arshad

School of Electrical and Computer Engineering
Purdue University

Introduction

- Proactive fault handling requires prediction of *failures*
- Online failure prediction is short-term failure predictions based on the current runtime system's state

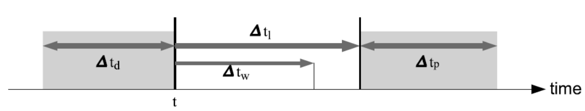

Figure 1. Time relations in online failure prediction: t – present time; Δt_l – lead time; Δt_w – warning time; Δt_p – prediction period; Δt_d – data window size

Classification of Failure prediction Techniques

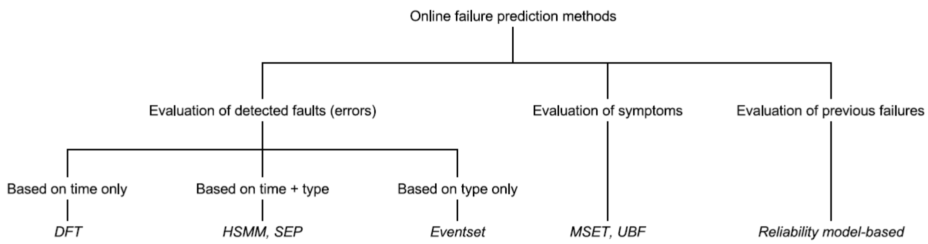

- Failure prediction based from previous failures is closely related to reliability prediction
- Majority of existing techniques are symptom based
  - Symptoms are side-effects of faults
  - Memory consumption, no of running processes etc.

PURDUE UNIVERSITY

HSMM-based Failure Prediction

- Basic assumption:
  - Failure-prone system behavior can be identified by characteristics patterns of errors.
- Approach
  - Error event timestamps and message IDs form an *error sequence*
  - After some data preprocessing, failure and non-failure error sequences are extracted
  - Two HSMMs are trained: Failure and Non-Failure sequences with parameters λ_F and $\lambda_{F'}$.

PURDUE UNIVERSITY

HSMM-based Failure Prediction

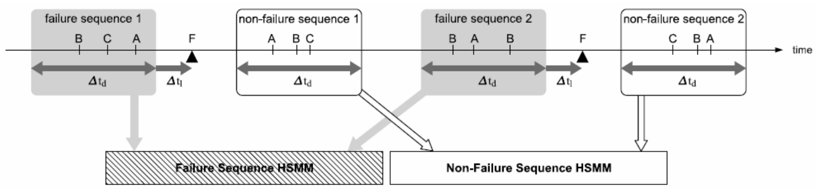

Figure 3. Two HSMMs are trained: One for failure sequences and one for non-failure sequences. Sequences consist of errors A, B, or C that have occurred in previously recorded training data. Failure sequences consist of errors that occurred within a time window of length Δt_d preceding a failure (▲) by lead time Δt_l. Non-failure sequences consist of errors that occurred at times when no failure was imminent.

PURDUE UNIVERSITY

HSMM-based Failure Prediction

- The goal is to assess, whether a given an error sequence $\boldsymbol{o} = [o0 \ldots oL]$ is failure prone or not
- Sequence likelihood $P(o|\lambda)$
  - Probability that a given model λ can generate observation sequence $\boldsymbol{o}$
  - Sequence likelihood is computed for both models

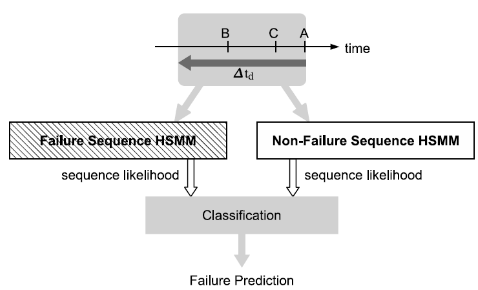

Figure 4. Online failure prediction. At the occurrence of error A (present time), the sequence under investigation consists of all errors that have occurred within the preceding time window of length Δt_d. Failure prediction is performed by computing sequence likelihood of the sequence using both the failure and non-failure model. Both likelihoods are evaluated using Bayes decision theory for classification as failure-prone or failure-free

PURDUE UNIVERSITY

HSMM-based Failure Prediction

- Criteria for classification:
  - Classify sequence ***o*** as failure-prone, iff

$$\log\left[P(\boldsymbol{o}\,|\,\lambda_F)\right] - \log\left[P(\boldsymbol{o}\,|\,\lambda_{\bar{F}})\right] \;>\; \underbrace{\log\left[\frac{c_{\bar{F}F} - c_{\bar{F}\bar{F}}}{c_{F\bar{F}} - c_{FF}}\right]}_{\in(-\infty;\infty)} + \underbrace{\log\left[\frac{P(\bar{F})}{P(F)}\right]}_{const.}$$

  - c_{ta} denotes the associated cost for assigning a sequence of type *t* to class *a*,
    - e.g., $c_{FF`}$ denotes the cost for falsely classifying a failure-prone sequence as failure free.
  - P(F) and P(F`) denote class probabilities of failure and non-failure sequences, respectively

PURDUE UNIVERSITY

Reliability-based prediction

- A simple reliability model is represented by an exponential distribution and is used for comparison

$$F(t) = 1 - e^{-\lambda t}$$

  - *F(t)* denotes the probability of a failure until time *t* and *R(t)* denotes reliability
- Failure rate is set to the inverse of mean-time-to-failure (MTTF)
- A failure is predicted according to the median of the distribution.
- After each failure that occurs in the test data set, the timer is reset and prediction starts again.

PURDUE UNIVERSITY

Dispersion Frame Technique (DFT)

- A Dispersion Frame (DF) is the interval time between successive error events
- The Error Dispersion Index (EDI) is defined to be the number of error occurrences in the later half of a DF
- A failure is predicted if one of five heuristic rules fires
  - one rule puts a threshold on error-occurrence frequencies
  - another on window-averaged occurrence frequency

Dispersion Frame Technique (DFT)

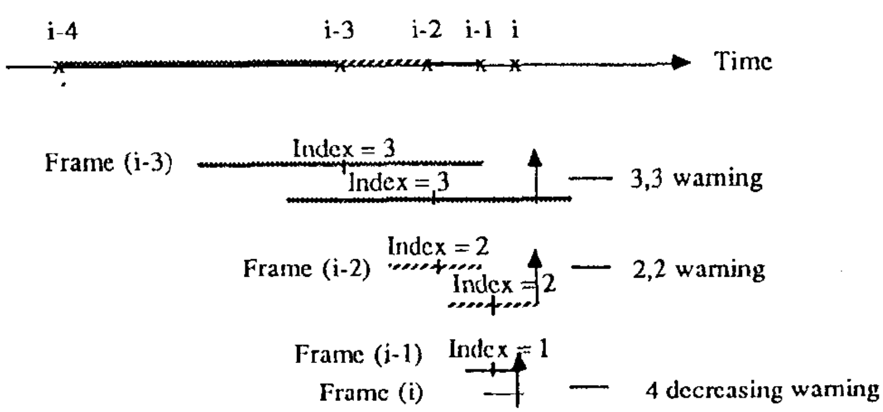

Figure 3. Dispersion Technique

Eventset-based Method

- This method is based on a set of events (errors) preceding a target event (failure).
- The goal is to set up a rule-based failure prediction system containing a database of indicative eventsets
- For each error, current set of events *Z* is formed from all errors that have occurred within Δt_d before present time.
- Database *DB* of indicative eventsets is then checked whether *Z* is a subset of any $D \varepsilon DB$. If so, a failure warning is raised at time Δt_l in the future.

PURDUE UNIVERSITY

Eventset-based Method

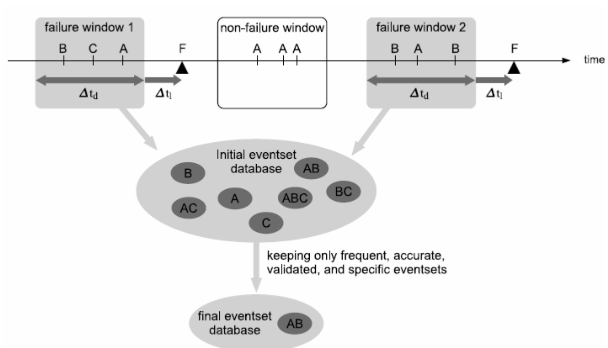

Figure 6. Eventset-based method. An eventset is the set of error types that have occurred within a time window of length Δt_d preceding a failure (▲) by lead time Δt_l

PURDUE UNIVERSITY

Experiment Description

- The prediction techniques are applied to a commercial telecommunication system
- Failure definition:
  - An event such that for a non-overlapping five minute intervals, the fraction of calls having response time longer than 250ms exceeds 0.01%
- Metrics
  - *Precision*: fraction of correctly predicted failures in comparison to all failure warnings
  - *Recall*: fraction of correctly predicted failures in comparison to the total number of failures.
  - *F-Measure*: harmonic mean of precision and recall
  - *False positive rate*: fraction of false alarms in comparison to all non-failures.

Contingency Table and Metrics formulas

		Truth	
		Failure	Non-failure
Predicted	Failure	True positive (TP)	False positive (FP)
	Non-Failure	False negative (FN)	True negative (TN)

(a)

Metric	Formula
precision	$p = \frac{TP}{TP+FP}$
recall=true positive rate	$r = tpr = \frac{TP}{TP+FN}$
false positive rate	$fpr = \frac{FP}{FP+TN}$
F-measure	$F = \frac{2*p*r}{p+r}$

(b)

Figure 8. Contingency table (a) and definition of metrics (b)

Results

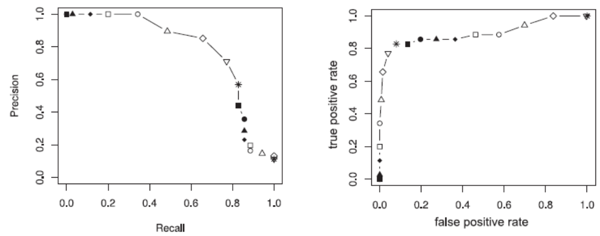

Figure 12. Prediction performance of the HSMM approach. precision-recall plot (left) and ROC curve (right). The various symbols denote different classification threshold values

Prediction technique	Precision	Recall	F-Measure	FPR
reliability-based	0.214	0.154	0.1791	n/a
DFT	0.314	0.458	0.3729	0.0027
Eventset	0.242	0.917	0.3826	0.1068
HSMM (max. F-measure)	0.852	0.657	0.7419	0.0145

Figure 13. Summary of prediction results

Conclusion

- A new approach to online failure prediction that forecasts failures by recognition of failure-prone patterns of error events
- The approach allows to employ a customizable threshold by which the tradeoff between, e.g., precision and recall can be controlled.
- HSMM gives better prediction performance in comparison to reliability based prediction, DFT and eventset based method