

# Seer: A Lightweight Online Failure Prediction Approach

Cemal Yilmaz

## Address:

Computer Science and Engineering Program  
Faculty of Engineering and Natural Sciences  
Sabanci University  
Orhanli, Tuzla, Istanbul, 34956, TURKEY  
Office: FENS G019

Voice: +90 (216) 483 9532

Fax: +90 (216) 483 9550

**E-mail:** [cyilmaz@sabanciuniv.edu](mailto:cyilmaz@sabanciuniv.edu)

## Abstract

Online failure prediction approaches aim to predict the manifestation of failures at runtime before they occur so that preventive measures, such as system reboots, or protective measures, such as checkpointing, can be proactively taken to improve software dependability. Existing approaches generally refrain themselves from collecting internal execution data, which can further improve the prediction quality. One reason behind this general trend is the runtime overhead incurred by the measurement instruments that collect the data. Since these approaches are targeted at deployed software systems, excessive runtime overhead is generally undesirable. I conjecture that large cost reductions in collecting internal execution data for online failure prediction may derive from pushing the substantial parts of the data collection work onto the hardware. In this talk, I will present a lightweight online failure prediction approach, called Seer, in which most of the data collection work is performed by fast hardware performance counters -- CPU resident counters that record various low level events occurring on a CPU. The hardware-collected data is augmented with further data collected by a minimal amount of software instrumentation that is added to the system's software. In the empirical evaluations conducted on three open source projects, Seer performed significantly better than other related approaches in predicting the manifestation of failures.