Architecture-Aware Online Failure Prediction for Distributed Software Systems

Teerat Pitakrat[1], Dušan Okanović[2], André van Hoorn[3], Lars Grunske[4]

Abstract: This extended abstract summarizes our article on architecture-aware online failure prediction, which has been published recently in the Journal on Software and Systems [Pi17].

Keywords: Online failure prediction; software architecture; distributed software systems

Today's software systems are complex. They comprise an immense number of distributed hardware and software components to deliver desired functionalities. Failures during production are inevitable despite successful approaches for quality assurance during software development. A failure in one component, e.g., a memory leak or slow response times, can create a chain of failures propagating to other components and the users [Av04]. Online failure prediction [SLM10] aims to foresee imminent failures by making predictions based on system parameters from monitoring data. Existing approaches employ prediction models that predict failures either for the whole system or for individual components without considering the software architecture.

We propose Hora, an architecture-aware online failure prediction approach, that combines failure prediction with architectural knowledge. The Hora approach is divided into three concurrent activities, as depicted in Fig. 1, which are component failure prediction, architectural dependency modeling, and failure propagation modeling and prediction.

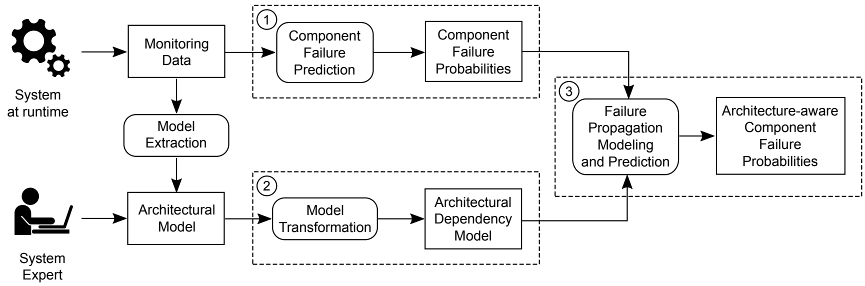

Fig. 1: Overview of the Hora approach

[1] University of Stuttgart, Institute of Software Technology, pitakrat@informatik.uni-stuttgart.de
[2] University of Stuttgart, Institute of Software Technology, okanovic@informatik.uni-stuttgart.de
[3] University of Stuttgart, Institute of Software Technology, van.hoorn@informatik.uni-stuttgart.de
[4] Humboldt University Berlin, Department of Computer Science, grunske@informatik.hu-berlin.de

The first activity of Hora is component failure prediction, denoted by (1). At runtime, the predictors receive monitoring data collected from application performance management (APM) tools [He17]. Each predictor is responsible for continuously predicting the failure of one component in the system. In the current implementation, we employ Kieker [vHWH12] to monitor method response times and resource utilization. The prediction technique used for predicting the failures is time series forecasting. Other prediction techniques, e.g., machine learning, can be employed if they can make predictions based on the monitoring data.

The second activity is architectural dependency modeling, denoted by (2). An Architectural Dependency Model (ADM) is a model that represents dependencies between components in the system. It indicates how components are connected and how likely a failure of one component will affect another. The ADM can be extracted from existing architectural models, e.g., SLAstic, or from monitoring data provided by APM tools, e.g., Kieker [vHWH12].

The third activity is failure propagation modeling and prediction, denoted by (3). The results of the first activity (component failure probabilities) and the ADM in the second activity are provided to a Failure Propagation Model (FPM). An FPM is a model that employs Bayesian network theory to model how a failure propagates through components until it reaches the system boundary and the users. The FPM continuously receives information from these two sources and periodically computes a system failure probability.

The evaluation of the Hora approach is performed on a microservice-based RSS feed reader application originally developed by Netflix. Three types of faults are injected into the system and Hora is applied to predict system failures. The results show that the prediction quality is improved when software architectural knowledge is considered explicitly.

References

[Av04] Avizienis, Algirdas; Laprie, Jean-Claude; Randell, Brian; Landwehr, Carl: Basic Concepts and Taxonomy of Dependable and Secure Computing. IEEE Transactions on Dependable and Secure Computing, 1(1):11–33, 2004.

[He17] Heger, Christoph; van Hoorn, André; Mann, Mario; Okanović, Dušan: Application Performance Management: State of the Art and Challenges for the Future. In: Proceedings of the 8th ACM/SPEC International Conference on Performance Engineering (ICPE). ACM, pp. 429–432, 2017.

[Pi17] Pitakrat, Teerat; Okanović, Dušan; van Hoorn, André; Grunske, Lars: Hora: Architecture-aware Online Failure Prediction. Journal of Systems and Software, 2017. In press. Online first: https://doi.org/10.1016/j.jss.2017.02.041.

[SLM10] Salfner, Felix; Lenk, Maren; Malek, Miroslaw: A Survey of Online Failure Prediction Methods. ACM Computing Surveys, 42(3):10:1–10:42, 2010.

[vHWH12] van Hoorn, André; Waller, Jan; Hasselbring, Wilhelm: Kieker: A Framework for Application Performance Monitoring and Dynamic Software Analysis. In: Proceedings of the 3rd ACM/SPEC International Conference on Performance Engineering (ICPE). ACM, pp. 247–248, 2012.