Timo Warns: Structural Failure Models for Fault-Tolerant Distributed Computing

1. Gutachter: Prof. Dr. Wilhelm Hasselbring (Universität Kiel)
2. Gutachter: Prof. Dr. Oliver Theel (Universität Oldenburg)
Datum der Prüfung: 28.9.2009

Zusammenfassung:
The dependability of a distributed system strongly depends on the occurrence of faults and on the ability of the system to cope with them. A fault-tolerant system is capable of providing service as expected even if some components have failed. Unfortunately, no system can tolerate arbitrary severe and arbitrary many faults. Engineering faulttolerant systems, therefore, require a fault model that describes the faults to tolerate. A good fault model must be accurate for the relevant aspects of faults, but abstract away irrelevant details. There is empirical evidence that, in particular, dependences and propagation of faults are relevant in real-world systems. In this thesis, we address the questions of how to model such faults and how to tolerate them.

For a fault model, we distinguish functional from structural failure models. A functional failure model describes how a component that is failed may behave. A structural failure model describes the extent of component failures. We investigate different classes of nonprobabilistic structural failure models and, in particular, introduce two new ones: set-based models for dependent faults and sequence-based models for dependent and propagating faults. Both classes close a gap between probabilistic models that cover dependent and propagating faults and previous nonprobabilistic models that do not. The new classes and several previous ones are compared with respect to their expressiveness resulting in a comprehensive hierarchy of nonprobabilistic structural failure models. All of the considered previous classes are strictly less expressive than the new set-based class, which is strictly less expressive than the new sequence-based class.

For many problems of distributed computing, there exist solutions that rely on quorums and, in particular, on highly available coteries to achieve fault tolerance. We illustrate how to solve distributed computing problems under the new model classes using highly available coteries and probing quorums. More precisely, we give characterisations of highly available coteries that show how to construct such a coterie from a set-based model if a highly-available coterie exists. Considering sequence-based models, we introduce the quality measure refined probe complexity that gives a tight bound on the number of required probes to find a quorum of noncrashed processes or to reveal that no such quorum exists. Additionally, we present a new probe strategy that is defined for all quorum sets and is more efficient in the number of required probes than previous strategies.

The considerations of quorums are independent of a particular fault tolerance problem. As a concrete problem, we show how to reach consensus in the presence of faults. In particular, we demonstrate that the new model classes do not require solutions developed from scratch: Adapting and transforming previous solutions for previous model classes suffice to reach consensus. Using the new model classes turns out to be beneficial as it allows more resilient and/or more efficient solutions.

Veröffentlicht als:
Timo Warns: Structural Failure Models for Fault-Tolerant Distributed Computing, Vieweg+Teubner Verlag, Wiesbaden, 2010 (ISBN: 978-3-8348-1287-2).