Abstract

This thesis deals with the quantitative and functional analysis of hybrid sys- tems, their models, and their execution traces as time-series.

Hybrid automata representations of hybrid systems are common in the aca- demic community and are well studied. Standard analyses deal with answering set-based reachability questions. Restricting the analysis to a subset of automaton runs is not always straightforward and can be cumbersome. Furthermore, lan- guages for expressing properties with quantitative semantics exist and are power- ful. However, in practice writing properties in these languages can be complicated and non-intuitive. In this thesis, we propose a quantitative property speci_cation
language called Feature Indented Assertions (FIA) for hybrid system contexts. The semantics of FIA enable us to propose a monitor automaton structure and a product construction with a hybrid automaton having fewer locations in the product than a standard product. We show that the syntax still allows us to specify various quantitative properties that appear in practice. A mechanism for reach-set formal analysis for generating quick, yet over-approximated feature ranges, and an SMT solver approach for computing tighter bounds at the expense of longer
analysis times is proposed. The SMT solution also computes concrete traces representing behaviours of the hybrid automaton for extremal feature values.

In situations where execution traces of a hybrid system are available, we aim to mine properties that characterize these traces. Existing work has focused on classi_cation, summarization, and pattern extraction. Explanations of key events in a time-series can help in design debugging, anomaly detection, and root-cause analysis. We propose a technique for mining properties from time-series. We make use of decision trees and interval arithmetic to mine properties with the aim of explaining de_ning events in the time-series. We present challenges that exist in using existing metrics for information entropy and information gain, and introduce adapted metrics that suit our requirements of mining temporal properties from time-series data. The mined properties are temporal in nature and presented in a formal language similar to FIA. We demonstrate the proposed techniques on a number of experiments, both synthetic and from key industrial applications.

Keywords: Formal Veri_cation, Formal Languages, Hybrid Automata, Reachability Analysis, Time-Series, Interval-Arithmetic, Property Mining.