

Title:  
Selective Monitoring

Abstract:

Runtime monitoring achieves a good trade-off in terms of developer effort versus bugs found. Most monitors, however, are used only during development but are turned off in shipped products, because the overhead of monitoring is significant. We formalize this problem as finding an observation policy for a product between a Markov chain and a finite automaton. Our theoretical results suggest several lessons for a potential implementation of a selective monitor; for example, it is very important to record not only which method was called but also what the call site was. This work has some limitations, which would need to be addressed before we can close the loop back to practice. In particular, the formal model we use has no notion of data.

This is ongoing joint work with Stefan Kiefer.