Ritwik Takkar
09/05/22

Paper Report:
Use of Formal Methods at Amazon Web Services by Chris Newcombe et al.

ritwiktakkar.com

1 Summary

Newcombe et al. cite myriad benefits related to TLA+ to argue in favor of implementing formal methods to verify the design of complex distributed systems. Specifically, they highlight how the formal methods have: (1) aided the teams in catching subtle, serious bugs prior to releasing to production, and (2) allowed them to implement "aggressive optimizations" thanks to testing changes on a model that would otherwise not exist. The authors, however, acknowledge the primary limitation of implementing formal methods: i.e., they deal with models rather than real-world systems themselves. But they also counter this point by arguing that increasing an engineer's familiarity with the design (a necessity for them to deal with formal methods) "can only increase the chances that the engineers will get the code right."

2 Strengths of the paper

This paper (purposely?) lacks the formal, academic tone that those firmly stationed within industry may find off-putting. Although this results in a less technical paper, it also increases readability and, therefore, probably reaches a larger audience. Moreover, given that this paper directly comes from AWS engineers, there is really no argument to be made against the practicality of implementing whatever it is they're suggesting in large-scale, industry-type environments. The more detailed personal experiences shared with respect to the implementation of formal methods enforces that idea (e.g., one of the authors, T.R., applied TLA+ to uncover subtle bugs/problems that not only did their extensive testing miss, but so too did conventional informal proofs).

3 Major weakness of the paper

While the authors claim that including snippets of the formal specifications may be off-putting to potential new users, and the means through which they convey their appreciation for the method is by discussing the value added, I think diving into at least one tutorial/example specific to AWS may have reinforced the point they were trying to make while also appealing to a more technical audience. Given that this paper comes from AWS and is authored by several team members, I don't think anyone thinks they lack credibility; I was convinced that formal methods offered great value a few pages in, but I would've liked to see a detailed example proving that instead of positive reviews by several individuals.

4 Future work opportunities

Since the purpose of AWS engineers was to find a tool and get on with their practical engineering, and they happened to find TLA+ which satisfied their use case, it seems they didn't spend significant effort evaluating other methods or perhaps building one from scratch. Not that I am familiar with TLA+, but I wonder whether there is enough demand, or, lack of features offered by today's formal methods to verify the design of complex distributed systems, to actually build one from scratch. On an unrelated note, the authors also mention a research opportunity to apply the TLC model checker to discover 'edge cases' in design on which to test the code on software.