

1 Summary

This paper presents *OpenLambda*: an open-source platform to build web services and applications in accordance with the *serverless computation* model. First, the authors introduce the serverless computation model while describing associated research challenges needed to be addressed in the design and implementation of such systems. Then, the authors discuss a commercial offering, *Amazon Lambda*, demonstrate some of its advantages of server-based models, and refer to its style of service construction as the *Lambda* model which OpenLambda is based on.

2 Strengths of the paper

It's commendable that the authors sensed the promise of serverless computing reinforced by the performance and success of Amazon Lambda, and then actually went ahead and open-sourced OpenLambda for the benefit of the broader research community. There are times when folks look at a paper or use some commercial technology and wish they had access to the source code, so they could modify or, at least, toy around with it for their use case or learning. The authors did exactly that for serverless computing through their interaction with Amazon Lambda.

All six figures shown in this paper are easy to understand and serve a clear purpose. For instance, Figure 1 portrays why the Lambda model is an obvious evolution to next generation sharing between applications, from hardware to operating systems to the runtime environments themselves.

I consider §4 the “meat” of the paper, given that it describes the challenges the authors wish to tackle in the serverless-computing space with OpenLambda. These challenges include improving base execution time for Lambdas to compete with VM and container platforms, evaluating performance tradeoffs with expensive profiling, and calculating design choices as to which third party libraries should be cached inside their handlers.

3 Weakness of the paper

Though the paper was short and crisp, I did miss having a proper Conclusion section. It may sound monotonous to adhere to the standard flow of academic papers, but the formula, at least in terms of having some sections as default like a conclusion and intro, helps readers like myself not feel “lost” on this new journey that is each new paper. Plus, given this paper summarized, at the time especially, burgeoning ideas, it would have made life easier for a reader to have key concepts summarized at the end.

4 Future work opportunities

Of relative interest to me is the specialized Lambda support required to support complex search applications. The cooperation of different Lambda leaves to filter and process data locally sounds interesting, and I wonder how it's implemented today; my guess would be a platform like MapReduce to preprocess the data before involving Lambdas.