MANGO: scalable modularity for transparently elastic cloud applications

Wei-Chiu Chuang (student,with demo), Milind Kulkarni
Purdue University
{chuangw, milind}@purdue.edu

Charles Killian
Purdue University and Google
ckillian@google.com

1 Introduction

Service providers love the cloud because its *elasticity* allows flexible resource management: an application can use cloud resources, such as virtual machines, storage or network bandwidth, on demand, which are billed by the amount of usage.

But programming elasticity proves to be a challenge for application developers: developers would have to explicitly handle nodes that join or leave, ensuring system state is consistent. While not impossible, manually adding elasticity to a system requires careful design and tremendous implementation efforts. Therefore, programmers often use programming models that exhibits *transparent elasticity*, such as EventWave [1] to simplify their daunting task: programmers write inelastic applications for a *logical node*, and elasticity is provided by the runtime system. In this manner, the runtime system distributes the application state on a varying number of *physical nodes*, which is transparent to the programmers.

Unfortunately, emerging cloud application developers face greater challenges than just elasticity. As applications become both larger and more complex, scalability becomes even more important. Meanwhile, complex application development requires modular design, but prior works for transparent elasticity did not address modularity. Existing solutions for modularity, such as loosely-coupled service oriented architecture, either does not guarantee consistency across services or does not scale. Therefore, a new solution is needed to fill the vacuum.

2 Mango

Our answer to the problem is MANGO: a framework for transparently elastic applications that enables both modular design and scalable performance. Similar to EventWave, MANGO embraces event-driven programming style. Programmers write code with atomic, sequential semantics to handle events generated from incoming requests, such as network messages, or timer notification.

MANGO sees an application as a collection of modules, or *services*, and each service is associated with state (which is a hierarchy of objects) and functions. A service may invoke functions in other services synchronously to accomplish a task, and the relation of the function invocation forms a directed acyclic graph. Although oblivious to programmers, our investigation found that implicit invocation rules exist in several distributed system implementations. MANGO uses this fact to enforce sequential semantics while enabling multiple events executing simultaneously.

When a logical node distributes its state on multiple physical nodes, each physical node creates events independently. Unlike EventWave, events in MANGO are mostly partially-ordered. Dependency tracking is necessary for events that access multiple objects to ensure consistency and sequential semantics, and a new dependency tracking protocol based on vector clock is used to reduce the space for recording dependency.

3 Results

We have built several real world applications from computational benchmarks to distributed services using MANGO, and successfully deployed prototypes on Amazon EC2 cloud infrastructure. Our preliminary result shows MANGO has better scalability than prior work, and programming is easier.

One of the distributed service we built using MANGO is ParkRanger, an elastic version of ZooKeeper-like distributed configuration service. We designed it as multiple services: a transport service, a primary-backup service and a service to create high-level constructs.

4 Conclusion

Existing programming models for developing cloud applications failed to address multiple challenges, including transparent elasticity, modularity and scalability, but all are crucial for emerging cloud applications.

We presented Mango, a framework whose programming model offers modularity for transparently elastic applications. Its runtime system enables consistent scalability. Finally, cases studies show that MANGO is useful for real world applications.

References

[1] CHUANG, W.-C., SANG, B., YOO, S., GU, R., KULKARNI, M., AND KILLIAN, C. Eventwave: programming model and runtime support for tightly-coupled elastic cloud applications. In *Proceedings of the 4th annual Symposium on Cloud Computing* (2013), ACM, p. 21.