Abstract

Nowadays data centers are attracting tremendous attention of researchers because they can accommodate a variety of applications. Data center network is an essential component to allow distributed applications to run efficiently and predictably. Performance of data centers require promising conduct round the clock. Incompetency in providing high performance may lead to dissatisfaction in customers. Customers are very sensitive to delay, and long delay implies throw away customers. Latency in data centers affects adversely on business revenue. However, reducing latency for services, which require gigantic database consultation and heavy computation, is a difficult job. Data center topologies are well structured, and this topology information of data center can be utilized for diminishing latency. This thesis is about TAP, a topology aware version of an existing flow scheduler (pFabric) and JFEPM, which takes advantage of higher capacity links of topology structures and use jumbo frames on them.

Recent Data center network topology puts forward multiple paths between the end hosts to provide high bisection bandwidth. Load-balancing proposals mainly deal with how to distribute traffic among multiple paths. In this thesis, FlowFurl, a distributed, end host driven, flow-level load-balancing technique will be discussed. This technique re-routes a flow not only based on the congestion at switch, but also takes care of the status of the links along the path to maximize the benefit of flow-level re-routing.

...