

1139-68-653

**Li Zhang\*** (zhangli@us.ibm.com). *Multi-resource fair sharing for multiclass workflows.*

Computer jobs in Big Data systems such as the Cloud, commonly require resources of multiple dimensions, including, CPU, memory, network, etc. Multi-resource sharing for concurrent workflows necessitates a fairness criteria to allocate multiple resources to workflows with heterogeneous demands. Recently, this problem has attracted increasing attention and has been investigated by assuming that each workflow has a single class of jobs and that each class contains jobs of the same demand profile. The demand profile of a class represents the required multi-resources of a job. However, for typical applications in cloud computing and distributed data processing systems, a workflow usually needs to process multiple classes of jobs. Relying on the concept of slowdown, we characterize fairness for multi-resource sharing and address the scheduling of multi-class workflows. We optimize the mixture of different classes of jobs for a workflow as optimal operation points to achieve the least slowdown, and discuss desirable properties for these operation points. (Received February 20, 2018)