Strategyproof Multi-Resource Allocation for Cloud Computing under Divisible and Indivisible Task Models

The rapid growth of cloud computing platforms has made the fair allocation of multiple heterogeneous resources, such as CPU, memory, and bandwidth, a fundamental algorithmic problem. A canonical mechanism in this setting is Dominant Resource Fairness (DRF) [Ghodsi et al., 2011]. DRF models agents' preferences via Leontief utilities, where each user requires resources in fixed proportions, and extends maximin fairness to the multi-resource setting by equalizing users' dominant resource shares, the resource type they demand most intensively. DRF satisfies four foundational properties that capture the key requirements in cloud systems: sharing incentive (SI), envy-freeness (EF), strategyproofness (SP), and Pareto optimality (PO).

Despite its strong fairness and incentive guarantees, DRF performs poorly in terms of efficiency. Under the classical utilitarian approximation-ratio benchmark, Parkes et al. [2015] showed that the social welfare achieved by DRF can be a factor of m worse than optimal, where m is the number of resource types. Moreover, this limitation is inherent: no mechanism satisfying any one of SI, SP, and PO can outperform DRF under this benchmark. Consequently, meaningful efficiency guarantees appear unattainable when evaluated against the unconstrained welfare optimum.

To overcome this barrier, Bei et al. [2026] introduced a new evaluation benchmark, the fair-ratio, which compares a mechanism only against allocations that themselves satisfy basic fairness requirements such as SI and EF. Under this framework, DRF is far from optimal in the basic two-resource setting. In particular, they designed new mechanisms that retain retain SI, EF, SP, and PO, while achieving a worst-case fair-ratio of 4/3, strictly improving upon the value of 2 attained by DRF. However, no nontrivial lower bound is currently known for the achievable fair-ratio in the two-resource setting. This leads to our first question.

Question 1. *How small can the fair-ratio be in the two-resource setting?*

We address this question by introducing a unified parametric mechanism design framework, called Adaptive Speed Fairness (ASF), which captures previous mechanisms as special cases. ASF mechanisms operate in two phases: first guaranteeing SI via a baseline allocation, and then allocating remaining resources by regulating the relative growth of dominant shares according to instance-dependent imbalance parameters. We show that SI, EF, and PO hold unconditionally within ASF, and we derive sufficient conditions ensuring SP. By carefully tuning its parameters, we construct a mechanism achieving asymptotic worst-case fair-ratio guarantee of 9/8, strictly improving the previously known 4/3 bound.

The above discussion rests on a critical modeling assumption: that tasks are divisible. In practical systems, however, tasks are indivisible and users cannot derive partial utility from resources unless they receive enough to run an entire task. This motivates the study of the indivisible-task model.

Question 2. *Can meaningful efficiency guarantees be achieved when tasks are indivisible?*

The indivisible setting was studied by Parkes et al. [2015], who established strong impossibility results showing that the four classical properties cannot be achieved simultaneously, even when envy-freeness is relaxed to envy-freeness up to one good (EF1). They therefore focused on mechanisms satisfying SI, EF1, and PO. We revisit the indivisible setting from a different perspective. Rather than insisting on exact Pareto optimality, we evaluate mechanisms using the quantitative fair-ratio benchmark and focus on SI+EF1+SP. We first show that simple equal-division rules can have unbounded fair-ratio. We then design a DRF-style mechanism that satisfies SI+EF1+SP and guarantees a constant fair-ratio of 6. Finally, we demonstrate that natural extensions of efficient divisible-task mechanisms fail to remain strategyproof once tasks become indivisible, revealing a structural separation between divisible and indivisible task models.

References

Xiaohui Bei, Zihao Li, and Junjie Luo. Fair and efficient multi-resource allocation for cloud computing: Beyond dominant resource fairness. *Mathematics of Operations Research*, 2026. Forthcoming.

Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX conference on Networked Systems Design and Implementation (NSDI)*, pages 323–336, 2011.

David C Parkes, Ariel D Procaccia, and Nisarg Shah. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation (TEAC)*, 3 (1):1–22, 2015.