

Budget-aware scheduling algorithms for scientific workflows on IaaS Cloud platforms

Yves Caniou
LIP, ENS Lyon, France
yves.caniou@ens-lyon.fr

Aurélié Kong Win Chang
LIP, ENS Lyon, France
aurelie.kong-win-chang@ens-lyon.fr

Eddy Caron
LIP, ENS Lyon, France
eddy.caron@ens-lyon.fr

Yves Robert
LIP, ENS Lyon, France
Univ. Tenn. Knoxville, USA
yves.robert@inria.fr

This work introduces several budget-aware algorithms to deploy scientific workflows on IaaS Cloud platforms, where users can request Virtual Machines (VMs) of different types, each with specific cost and speed parameters. We use a realistic application/platform model with stochastic task weights, and VMs communicating through a datacenter. We extend two well-known algorithms, HEFT and MIN-MIN, and make scheduling decisions based upon machine availability *and* available budget. The new budget-aware algorithms are called HEFTBUDG and MIN-MINBUDG. During the mapping process, these budget-aware algorithms make conservative assumptions to avoid exceeding the initial budget; we further improve our results with refined versions that aim at re-scheduling some tasks onto faster VMs, thereby spending any budget fraction leftover by the first allocation. These refined variants are called HEFTBUDG+ and HEFTBUDG+INV and differ by the order in which they process tasks. The refined variants are much more time-consuming than the former algorithms, so there is a trade-off to find in terms of scalability. We report an extensive set of simulations with workflows from the Pegasus benchmark suite. We have designed a simulator based on SimDag, an extension of the discrete event simulator SimGrid, to evaluate all algorithms. Budget-aware algorithms generally succeed in achieving efficient makespans while enforcing the given budget, and despite the uncertainty in task weights.

The schedules obtained for both refined algorithms HEFTBUDG+ and HEFTBUDG+INV have a shorter makespan than HEFTBUDG (see Figure 1 a, d, g). Their makespan can be up to one third shorter than for HEFTBUDG, *e.g.*, for MONTAGE. Surprisingly the refined algorithms manage to achieve a smaller makespan using *fewer* VMs than HEFTBUDG. This is mostly because they succeed in assigning interdependent tasks onto the same VM. The initial budget needed to obtain the same makespan as HEFT is the same for HEFTBUDG, HEFTBUDG+ and HEFTBUDG+INV. Moreover, the budget is respected overall, as was already the case with HEFTBUDG.

To compare HEFTBUDG+ and HEFTBUDG+INV, we observe that their makespans are very similar, apart from the case where a budget near the minimal one is needed to complete a schedule: in that case, HEFTBUDG+ obtains an average makespan twice shorter than HEFTBUDG+INV. This difference seen in this exact configuration may be due to the particular structure of MONTAGE workflows: they have a lot of initial tasks, and the amount of work for every

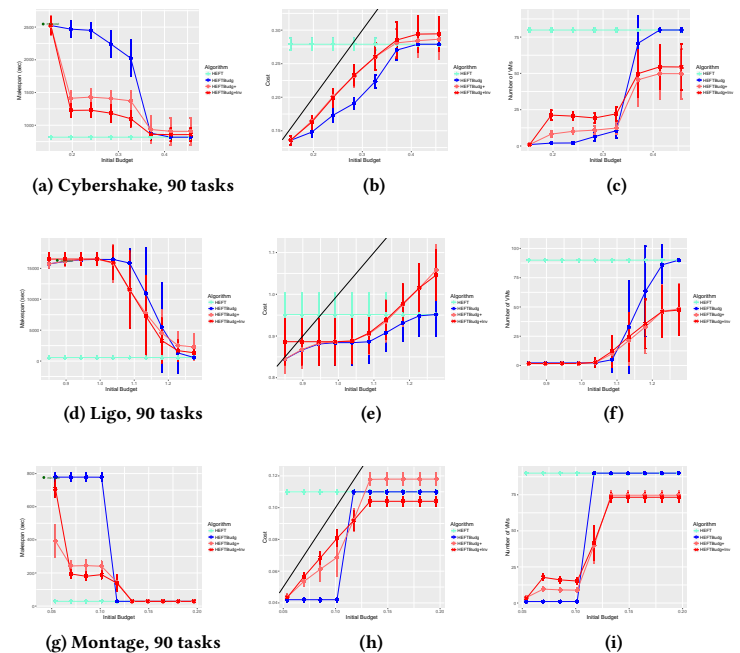


Figure 1: HEFTBUDG+, HEFTBUDG+INV compared to HEFT and HEFTBUDG for the three workflow types with 90 tasks.

kind of tasks is of the same magnitude. However, HEFTBUDG+ and HEFTBUDG+INV achieve far better makespans than HEFTBUDG but for a higher computational cost. For instance, for MONTAGE with 90 tasks and a high budget, HEFTBUDG finds a solution in 2.60 ± 0.28 seconds while HEFTBUDG+ needs 379.45 ± 44.20 seconds, and HEFTBUDG+INV needs 382.29 ± 43.25 seconds. We refer the reader to the extended version [1] for all details and bibliographical references.

REFERENCES

- [1] Y. Caniou, E. Caron, A. K. W. Chang, and Y. Robert. Budget-aware scheduling algorithms for scientific workflows on IaaS cloud platform. Research Report 9088, INRIA, Aug. 2017.