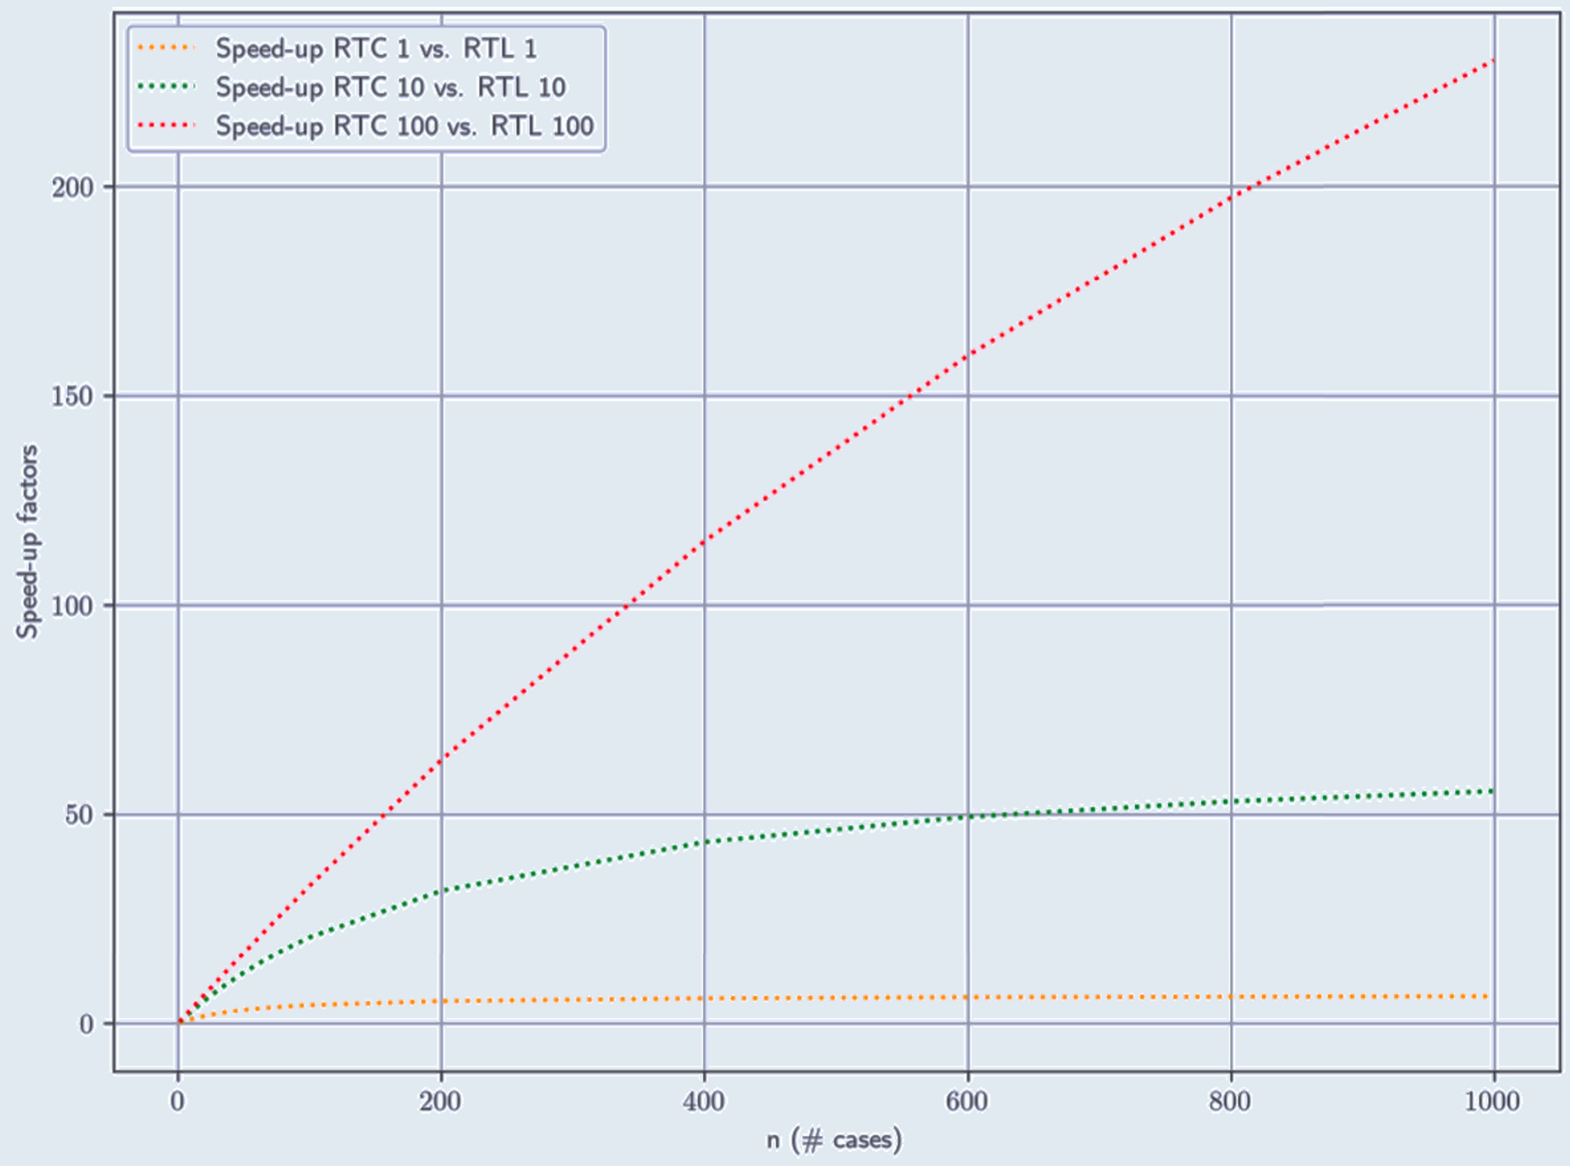

Kaleidosim *cloudCompanion* for COMSOL Multiphysics®

Introducing a novel technology that enables seamless access and orchestration of cloud computing directly out of the Graphical User Interface of COMSOL Multiphysics®.

V. Muser[2], V. Lienhard[2], D. Sharman[1]

[1]Kaleidosim Technologies AG, Zurich, Switzerland
[2]Zurich University of Applied Sciences, Zurich, Switzerland

Introduction

To this day, seamless accessibility of local or cloud hardware for high performance computing as well as time consuming workflows involving sequencing and orchestration of available hardware for multiple simulation runs, remain limiting factors within the field of Multiphysics simulation-based engineering.

The cloudCompanion tackles these issues as an add-on app – directly out of well-established simulation workflows. A particular specialty of the novel solution is its parameter study capability – the ability to coordinate cluster sweep operations on hundreds of simultaneously running cloud-based machines, such that related workflows can be accelerated by factors potentially beyond 100. Advantages of significantly more CPUs available in the cloud not yet considered.

A validated case is made for the considerable potential of the novel cloud computing technology, enhancing productivity, output, data-density as well as insights gained by applying COMSOL Multiphysics® in both industry as well as in academia.

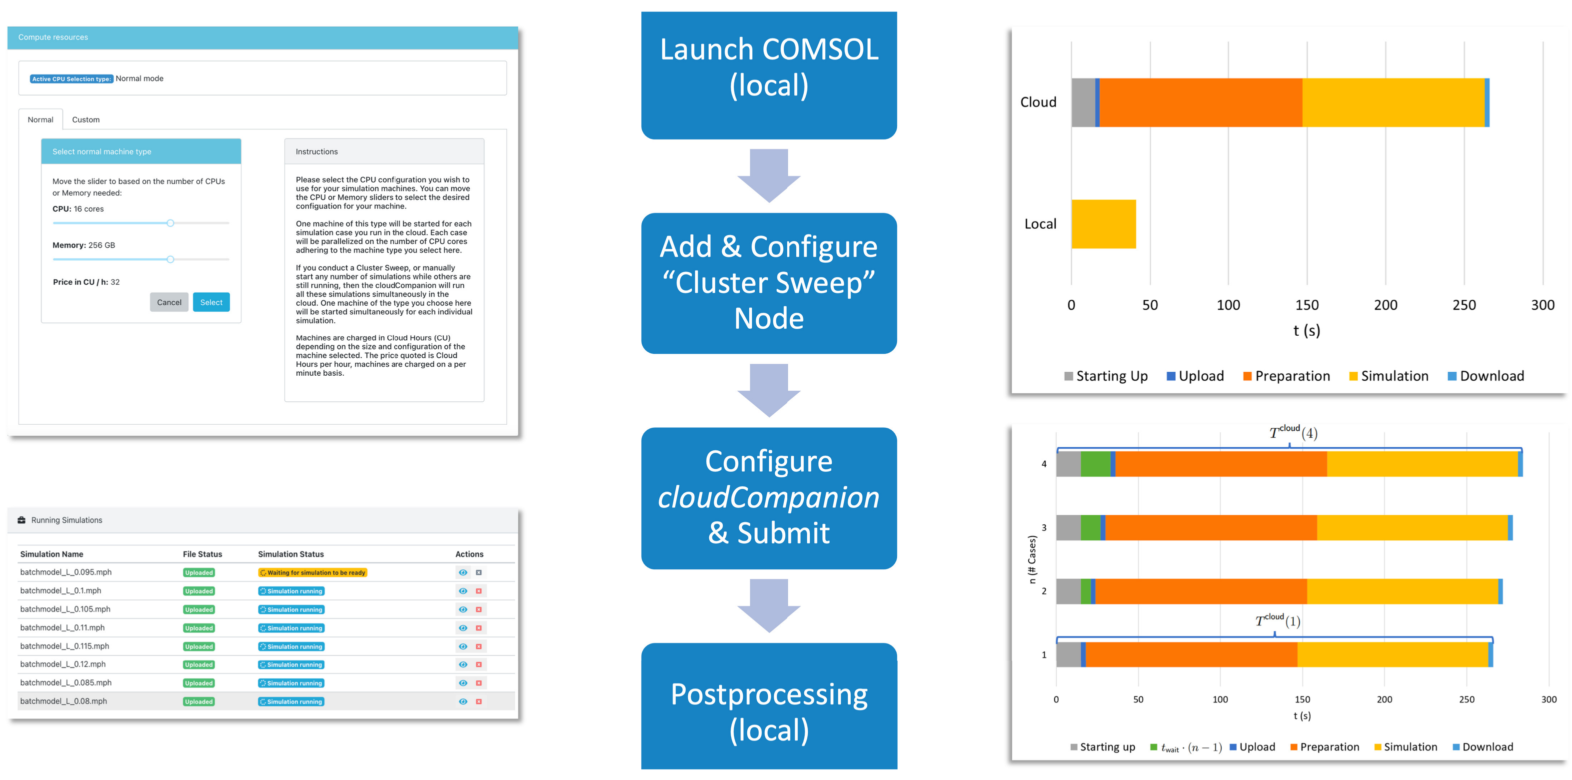

Figure 1. Left: selected app screenshots, Middle: applied workflow, Right: computing times for example & parameter run

Methodology & Workflow

In this showcase the scaling effect by using the *cloudCompanion* was analyzed, based on the intentionally unfavorable "Electrical Heating in a Busbar Assembly" example by COMSOL Multiphysics®[1].

Based on robustness and single case analysis, a theoretical prediction model for total relative time consumption of cloud computation as a function of the number of parallel cases n and total local/cloud single case computation time $T^{local}(1)/\,T^{cloud}(1)$ could be validated:

$$\tilde{t}^{cloud}\left(n, T^{local}(1)\right) = \frac{T^{cloud}(1) - t_{wait}}{T^{local}(1)} + \frac{t_{wait}}{T^{local}(1)}n$$

Hardware 2 CPU: of Intel i9-12900H (local), AMD EPYC Milan 3G 7B13 (cloud)[2]

Benchmarking Results

In the current setup, the waiting time t_{wait} = 6s was measured and the mean of ten 1x1 runs 1 for the single case duration was taken – resulting in $T^{cloud}(1)$ = 345.35s and $T^{local}(1)$ = 40.05s.

This means that even for this highly unfavorable showcase it is advised to use the cloud for an $n_{critical}$ ≈ 10 already.

It also becomes apparent that the potential time savings increase significantly when choosing cases with longer local reference computation time. This is sketched in Fig. 2 by RTC/RTL 10 & 100 where assumed reference times $T^{cloud}(1)$ & $T^{local}(1)$ are increased by factors of 10x/100x respectively.

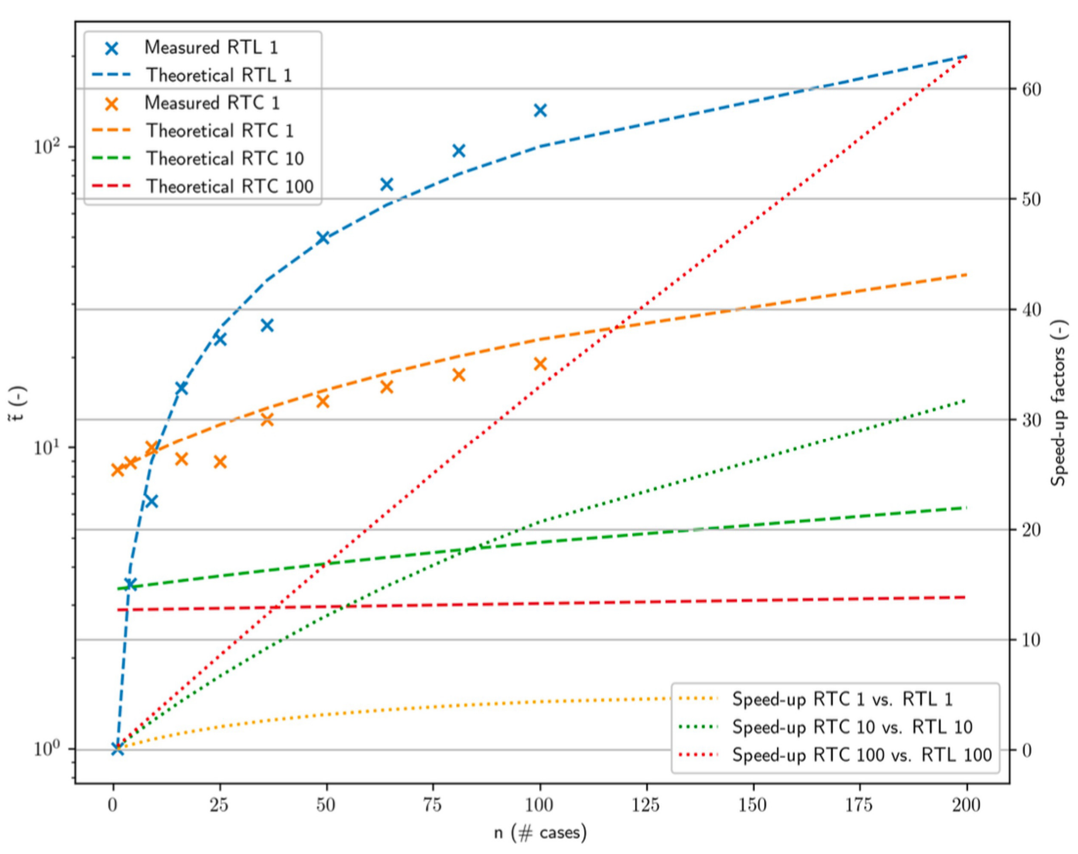

Cases	1x1	5x5	15x15	1x1 (local)
Run 1	383	472	2258	49
Run 2	374	545	2137	35
Run 3	253	499	2273	34
Run 4	378	445	2312	44
Run 5	282	537	2231	45
Run 6	372	472	2176	44
Run 7	402	538	2120	45
Run 8	277	433	2322	39
Run 9	399	431	2192	34
Run 10	378	364	2277	35
Mean	345.35	470.28	2228.78	40.05
SD	53.09	54.96	67.19	5.33
%RSD	0.15	0.12	0.03	0.13

Figure 2. Left: relative computation times for local/cloud (RTL/RTC) & speedups. Right: n=1 prediction model validation

REFERENCES

[1]COMSOL. Electrical heating in a busbar assembly, 2023. https://www.comsol.com/model/electrical-heating-in-a-busbar-assembly-56231 [Accessed: 23.06.2023)].
[2]Google. Compute-optimized machine family for compute engine, 2023. https://cloud.google.com/compute/docs/compute-optimized-machines [Accessed: 23.06.2023)].