

# Defining Green Profit in Distributed Datacenters

Fereydoun Farrahi Moghaddam, Reza Farrahi Moghaddam,  
Mohamed Cheriet  
École de Technologie Supérieure  
Montreal Canada

Yves Lemieux  
Ericsson Canada  
Montreal, Canada

**Abstract**—In this paper, a profit model is introduced for distributed datacenters which work in a job-scheduling business model, in order to maximize their profit. A “virtual carbon tax” is also introduced to artificially add a virtual cost to the profit model, resulting in a virtual profit (green profit). Simulation results show that job scheduling based on green profit balances the tradeoff between real profit and environmental footprints of the datacenter (a small loss in profit results in a large environmental impact reduction).

**Keywords**—virtual carbon tax; green profit; profit maximization; sustainability; environmental footprint

## I. INTRODUCTION

By nature, datacenters (excluding their emergency power system) are carbon neutral, but the electricity consumed in them is not. The electricity mix of the power grid is determined by distribution of its power plants with diverse technologies ranging from coal to hydro. Knowing the exact mixture of electricity mix, one can calculate its environmental footprint. It is worth noting that electricity mix has a dynamic but semi-periodic nature relative to time zone, time of day, and season, among other parameters.

In reality, the electricity price only includes electricity generation cost and does not represent its real cost which includes the cost of offsetting its environmental footprint. Not many jurisdictions are imposing strict policies towards real electricity cost, and therefore for-profit organizations may not care about the greenness of the electricity mix. Having said that, some corporations such as Google, Apple, and Facebook are taking voluntary actions towards this direction [5].

This paper provides geographically-distributed datacenters (which are connected to different electricity mixes) with a model to maximize their non-green profit even with a simple job scheduler [2]. Then, it imposes a voluntary *virtual carbon tax* to the profit model to balance the tradeoff between maximizing profit and minimizing environmental impacts. When *virtual carbon tax* is added to the model, it is more profitable (green profit) for the datacenter to schedule new jobs in locations with lower environmental footprint.

## II. PROFIT MODEL

Equation (1) calculates the profit generated by a single CPU core (working at frequency  $f$ ) during an hour, namely ‘Profit per Core Hour GHz’ (PpCHG). To calculate the total profit, one must sum up all the applicable PpCHGs over all the cores and operating hours.

$$\left\{ \begin{array}{l} PpCHG(f) = \text{revenue}_{core, hour, f} - \text{cost}_{core, hour, f} \\ \text{cost}_{core, hour, f} = OPEX_{core, hour, f} + \\ \text{electricity}_{core, hour, f} * \text{electricityRate} + \\ \text{coolingElec}_{core, hour, f} * \text{electricityRate} \end{array} \right. \quad (1)$$

Having PpCHG( $f$ ), it is possible to calculate the optimum frequency which maximizes the PpCHG (PpCHG( $f_{opt}$ )). Figure 1 illustrates PpCHG( $f_{opt}$ ) where DVFS model is valid:  $\text{electricity}_{core, hour, f} = \beta + \alpha f^3$  [1].

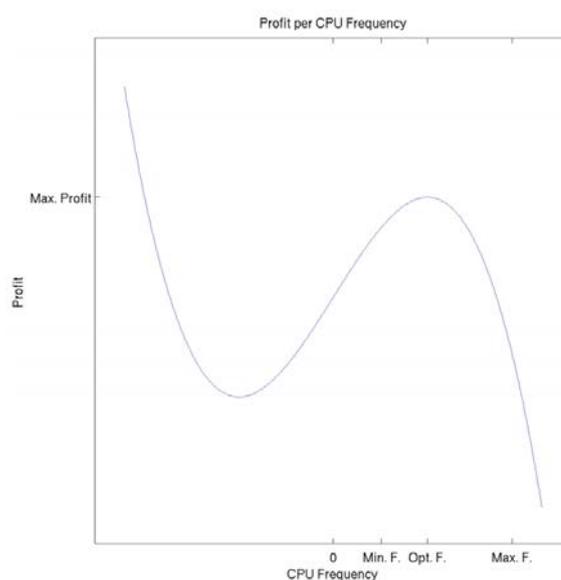


Figure 1 - PpCHG( $f_{opt}$ )

## III. SCHEDULING BASED ON PpCHG

Figure 2 shows PpCHG( $f_{opt}$ ) for several location during a week (Each 160 cores are located in a different time zone). Darker areas are associated with lower profit. A scheduler can use this map to avoid low profit areas to maximize the total profit.

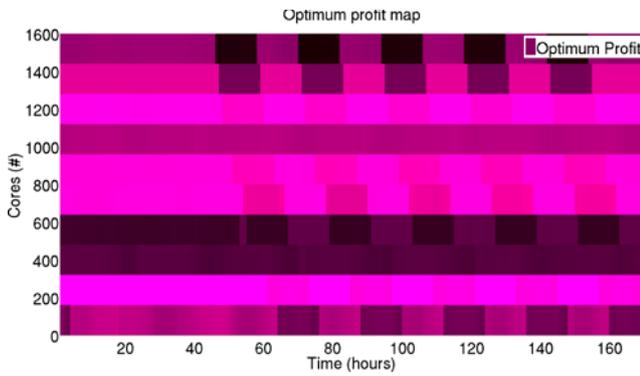


Figure 2 – Distributed datacenter maximum-profit map

IV. IMPOSING VIRTUAL CARBON TAX

Equation (2) adds an imaginary carbon tax<sup>1</sup> to Equation (1):

$$\begin{cases}
 PpCHG(f) = revenue_{core, hour, f} - cost_{core, hour, f} \\
 cost_{core, hour, f} = OPEX_{core, hour, f} + \\
 electricity_{core, hour, f} * electricityRate + \\
 coolingElec_{core, hour, f} * electricityRate + \\
 carbon_{core, hour, f} * virtualCarbonTaxRate + \\
 coolingCarb_{core, hour, f} * virtualCarbonTaxRate
 \end{cases} \quad (2)$$

V. SIMULATION RESULTS

A distributed datacenter located in ten time zones is tested under three schedulers: a profit-aware [3,4], a carbon-aware [2], and a green PpCHG-aware. Figure 3, Figure 4, and Figure 5 illustrate the actual schedule of HPC jobs by these schedulers. Each job is colored with an RGB color (red: frequency, green: greenness, and blue: electricity price).

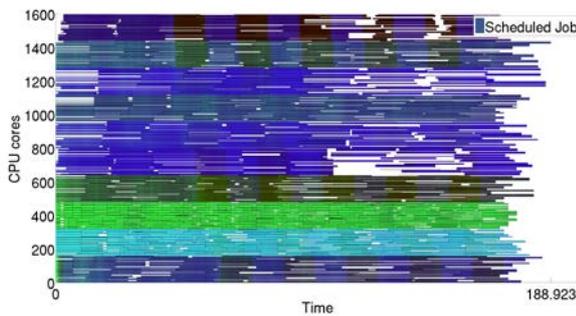


Figure 3 – Carbon-aware job scheduler

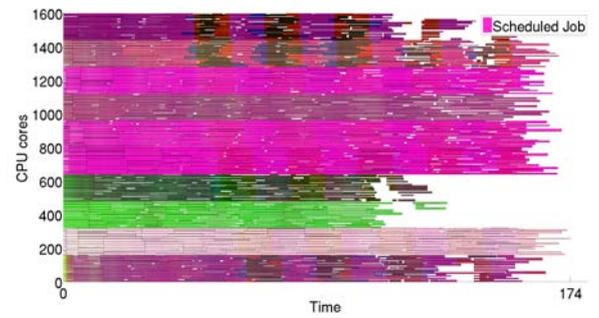


Figure 4 – PpCHG-aware job scheduler

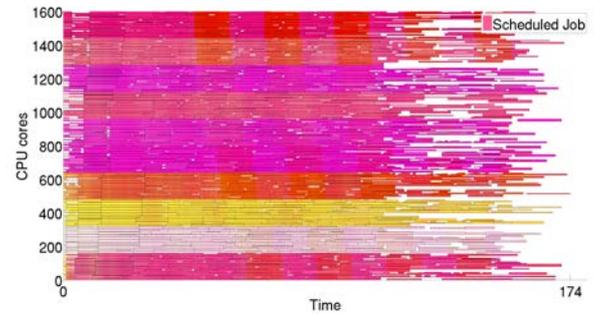


Figure 5 – Profit-aware job scheduler

Carbon-aware scheduler tries to avoid non-green areas and reduce the frequency of CPUs. Profit-aware scheduler increases the frequency of CPUs, and PpCHG-aware scheduler avoids non-profitable areas (for example, by avoiding electricity peak hours or reducing the frequency during the peak hours).

Figure 6 shows the tradeoff between profit (\$) and carbon (Kg) for all three scheduler. PpCHG results are shown with triangles for a range of virtual carbon tax rate from 0 to \$0.60/kgCO<sub>2</sub>-eq.

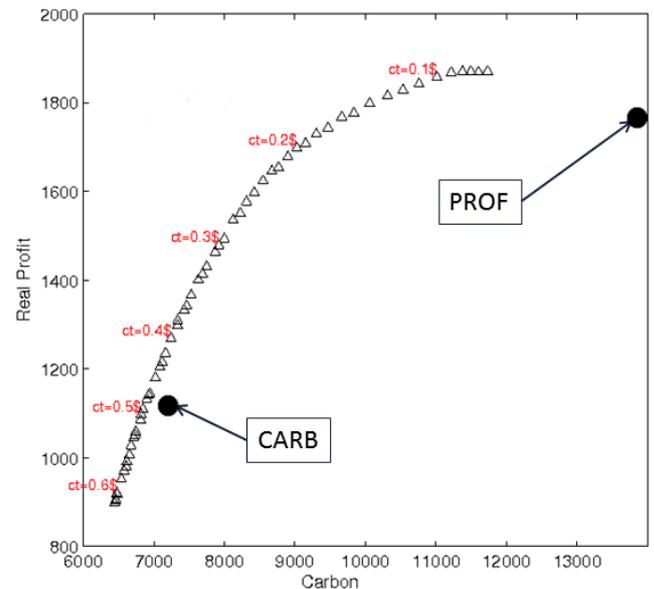


Figure 6 – Carbon-profit tradeoff

<sup>1</sup> Virtual Carbon Tax (\$) = Carbon Footprint (kgCO<sub>2</sub>-eq) \* Virtual Carbon Tax Rate (\$/kgCO<sub>2</sub>-eq)

## VI. CONCLUSION

This paper uses a profit model for a core-hour unit to calculate and maximize the total profit of a distributed datacenter. A *virtual carbon tax* was also added to the profit model in order to enable voluntary minimization of the carbon footprint of the datacenter.

A typical profit-aware scheduler will maximize the profit, but it also maximizes the carbon footprint. A typical carbon-aware scheduler will minimize the carbon footprint, but it also minimizes the profit. In a PpCHG-aware scheduler, a balanced solution can be achieved by choosing the right amount of *virtual carbon tax rate*.

Other Virtual Sustainability Taxes will be addressed in our future works such as Virtual Manufacturing Tax, Virtual Toxic-Wastes Tax, and Virtual End-of-Life Tax, among other parameters in Life Cycle Assessment of datacenters. We also plan to apply Virtual Carbon Tax to larger use cases using real-world datacenter traces.

## REFERENCES

- [1] Wang, Leping, and Ying Lu. "Efficient power management of heterogeneous soft real-time clusters." Real-Time Systems Symposium, 2008. IEEE, 2008.
- [2] Garg, Saurabh Kumar, et al. "Environment-conscious scheduling of HPC applications on distributed cloud-oriented data centers." Journal of Parallel and Distributed Computing 71.6 (2011): 732-749.
- [3] Kim, Jong-Kook, et al. "Dynamic mapping in a heterogeneous environment with tasks having priorities and multiple deadlines." Parallel and Distributed Processing Symposium, 2003. Proceedings. International. IEEE, 2003.
- [4] Qureshi, Asfandyar, et al. "Cutting the electric bill for internet-scale systems." ACM SIGCOMM Computer Communication Review. Vol. 39. No. 4. ACM, 2009.
- [5] Google, Facebook and Apple lead on green data centers, <http://www.theguardian.com/sustainable-business/greenpeace-report-google-facebook-apple-green-data-centers>, (May 2015)