Digital sufficiency behaviors to deal with intermittent energy sources in data center

ICT4S'24 @ Stockholm, Sweden

Jolyne Gatt, **Maël Madon**, Georges Da Costa

IRIT, Université de Toulouse, France

June 25, 2024

Introduction

- Energy efficiency:
  Koomey's law = doubling the number of computations per kWh every 1.57 years

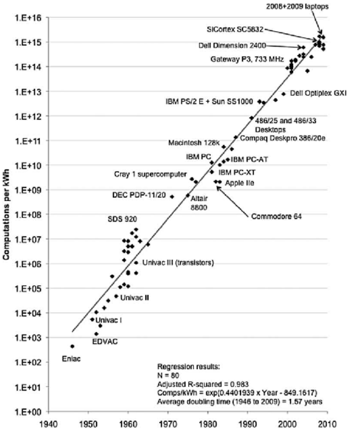

Source: Koomey et al. 2011 [1]

Introduction

- Energy efficiency:
  Koomey's law = doubling the number of computations per kWh every 1.57 years

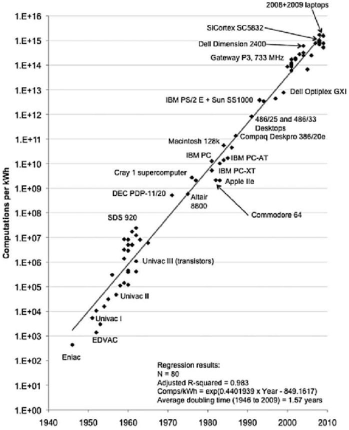

Source: *Koomey et al. 2011 [1]*

- Rebound effect:

Global trends in digital and energy indicators, 2015-2022

	2015	2022	Change
Internet users	3 billion	5.3 billion	+78%
Internet traffic	0.6 ZB	4.4 ZB	+600%
Data centre workloads	180 million	800 million	+340%
Data centre energy use (excluding crypto)	200 TWh	240-340 TWh	+20-70%
Crypto mining energy use	4 TWh	100-150 TWh	+2300-3500%
Data transmission network energy use	220 TWh	260-360 TWh	+18-64%

Source: *International Energy Agency*

Sufficiency

- Efficiency is not enough: **sufficiency**

Digital sufficiency (Santarius et al., 2022 [2])

Any strategy aimed at directly or indirectly **decreasing the absolute level of resource and energy demand from the production or application of IT**.

- **What would "sufficiency" mean for data centers?**
  $\rightarrow$ voluntary limitation, empower and involve the user
- this study: estimate the potential of "sufficiency behaviors" for data center users in a context of intermittent energy production

Model

Data center model

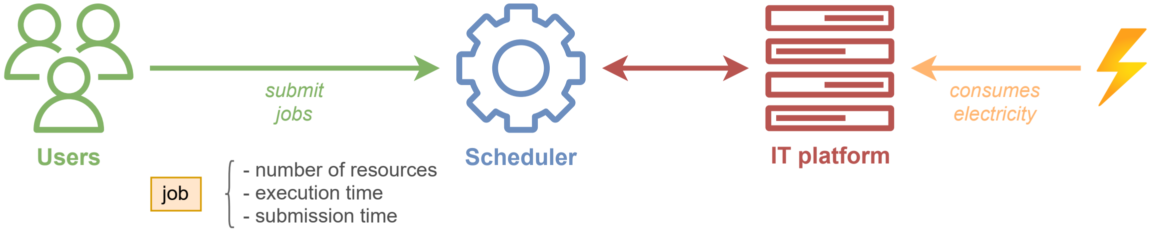

Renewable energy production

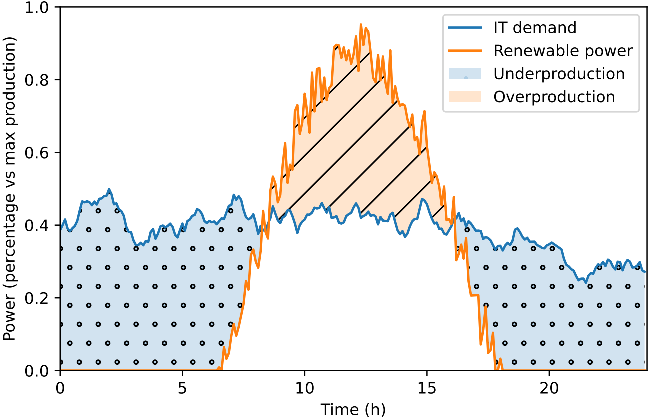

- **Objective**: minimize underproduction (a.k.a. “brown energy”)

Sufficiency behaviors

Rigid (baseline)

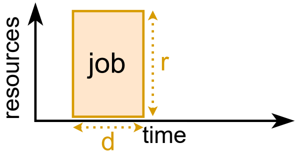

Sufficiency behaviors

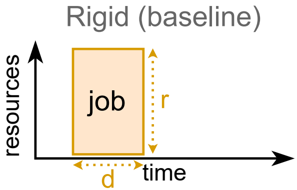

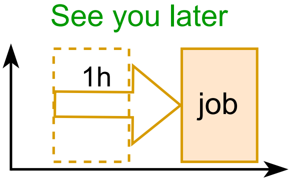

Sufficiency behaviors

Rigid (baseline)

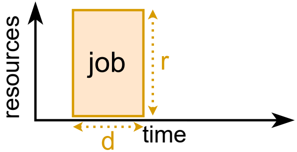

See you later

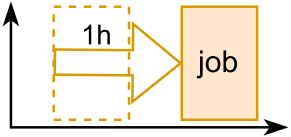

Reconfig

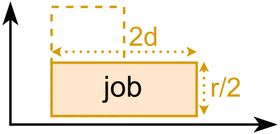

ex: *fewer nodes for image processing*

Sufficiency behaviors

Rigid (baseline)

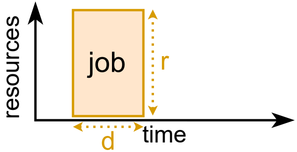

See you later

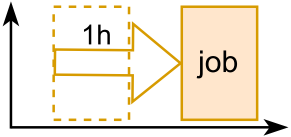

Reconfig

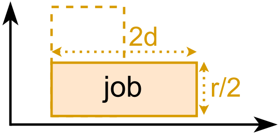

ex: *fewer nodes for image processing*

Space Degrad

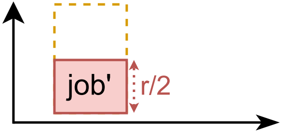

ex: *only 5 outputs instead of 10*

Sufficiency behaviors

Rigid (baseline)

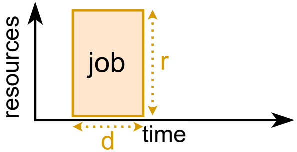

See you later

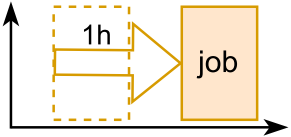

Reconfig

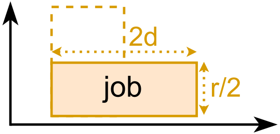

ex: *fewer nodes for image processing*

Space Degrad

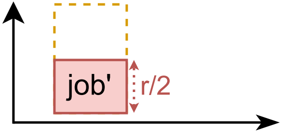

ex: *only 5 outputs instead of 10*

Time Degrad

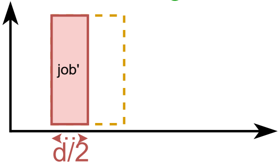

ex: *lower accuracy in a linear solver*

Sufficiency behaviors

Rigid (baseline)

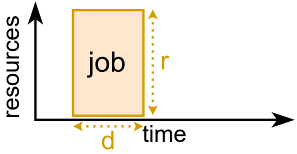

See you later

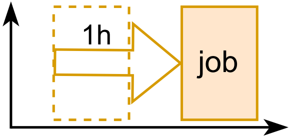

Reconfig

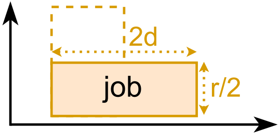

ex: *fewer nodes for image processing*

Space Degrad

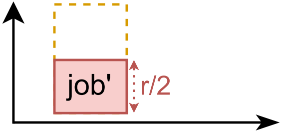

Time Degrad

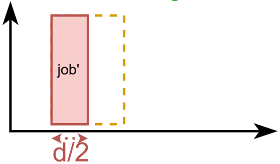

Renounce

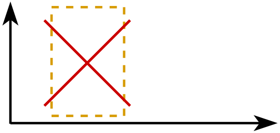

ex: *only 5 outputs instead of 10*

ex: *lower accuracy in a linear solver*

Sufficiency behaviors

Rigid (baseline)

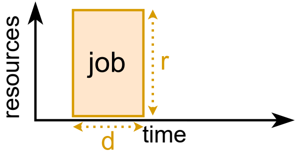

See you later

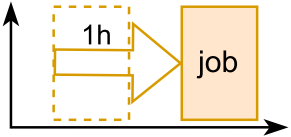

Reconfig

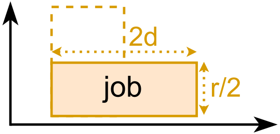

ex: *fewer nodes for image processing*

Space Degrad

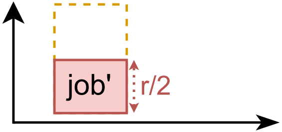

ex: *only 5 outputs instead of 10*

Time Degrad

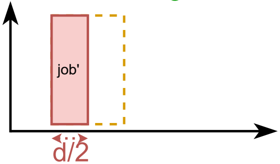

ex: *lower accuracy in a linear solver*

Renounce

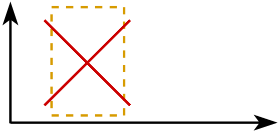

- **job final state** $= n \times \textit{see\_you\_later} + b,\ b \in \{$ Rigid, Reconfig, Space Degrad, Time Degrad $\}$

3-state energy model

- 3-color state for energy production:
  - green state: everything is fine (production $\geq$ 100% max conso)
  - yellow state: some disturbance (production $\geq$ 50% max conso)
  - red state: system critical (production $<$ 50% max conso).

3-state energy model

- 3-color state for energy production:
  - green state: everything is fine (production $\geq$ 100% max conso)
  - yellow state: some disturbance (production $\geq$ 50% max conso)
  - red state: system critical (production $<$ 50% max conso).

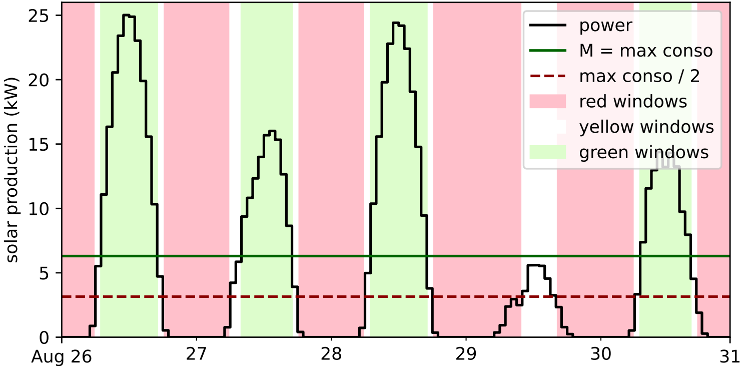

Energy-aware behaviors

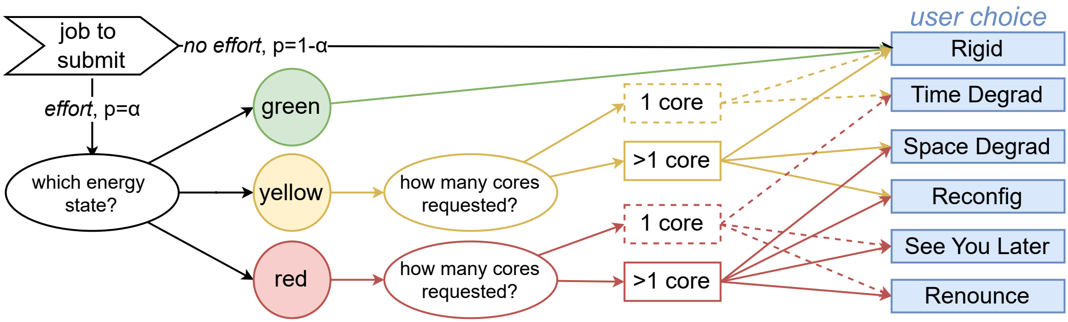

Energy-aware behaviors

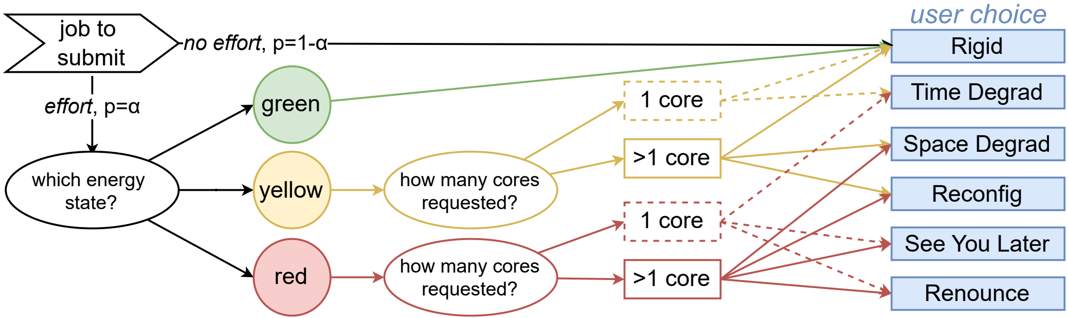

- choice of behavior at random depending on the state

Experimental campaign

Experimental setup

- **Software:** Batsim + Batmen

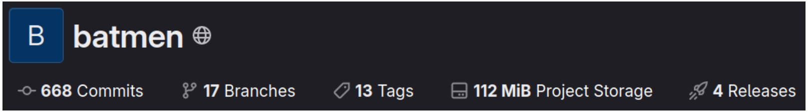

- **IT workload:** filtered version of MetaCentrum from Parallel Workload Archive
  - June 1 to November 11, 2014 (4.5 months)
  - 650000 jobs and 500 users
- **Energy production data:**
  - 145 m^2 solar panels
  - weather data Toulouse 2019 from Renewable Ninja (days aligned with IT)
- **IT platform:**
  - 42 18-core machines
- **Scheduler:** bin-packing scheduler which shutdown machine when idle.

Experimental campaign

- $\alpha =$ probability of modifying a job in red / yellow
- 6 scenari:
  - full rigid ($\alpha = 0$)
  - low effort ($\alpha = .25$)
  - medium effort ($\alpha = .5$)
  - big effort ($\alpha = .75$)
  - max effort ($\alpha = 1$)
  - full renounce/degrad/reconfig in red
- each scenario run 30 times to minimize the effect of randomness

Results

Results

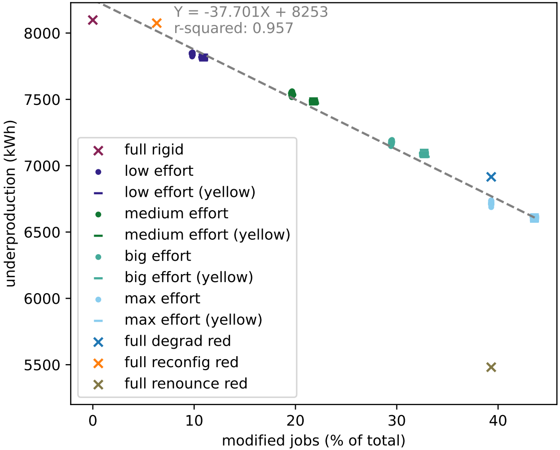

Y = -37.701X + 8253
r-squared: 0.957
8000
7500
7000
6500
6000
5500
underproduction (kWh)
0
10
20
30
40
modified jobs (% of total)
full rigid
low effort
low effort (yellow)
medium effort
medium effort (yellow)
big effort
big effort (yellow)
max effort
max effort (yellow)
full degrad red
full reconfig red
full renounce red

Results

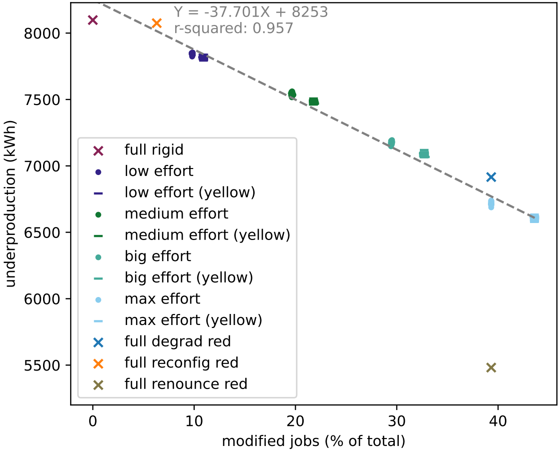

- *How much does user effort impact energy consumption?*
  - → if 50% jobs modified in red/yellow (medium effort), underproduction reduced by 8%
  - → if 100% jobs modified in red/yellow (max effort), underproduction reduced by 18%

Results

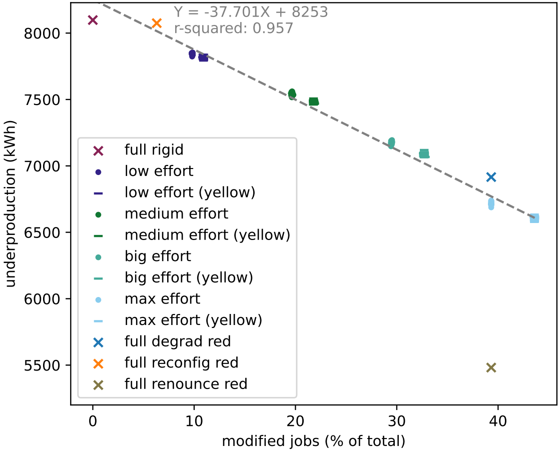

- *How much does user effort impact energy consumption?*
  - → if 50% jobs modified in red/yellow (medium effort), underproduction reduced by 8%
  - → if 100% jobs modified in red/yellow (max effort), underproduction reduced by 18%
- Energy savings linear with effort

Results: ratio energy/effort

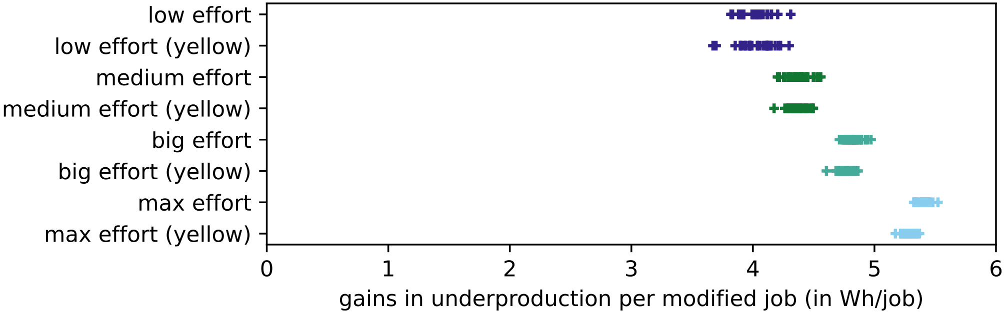

Results: ratio energy/effort

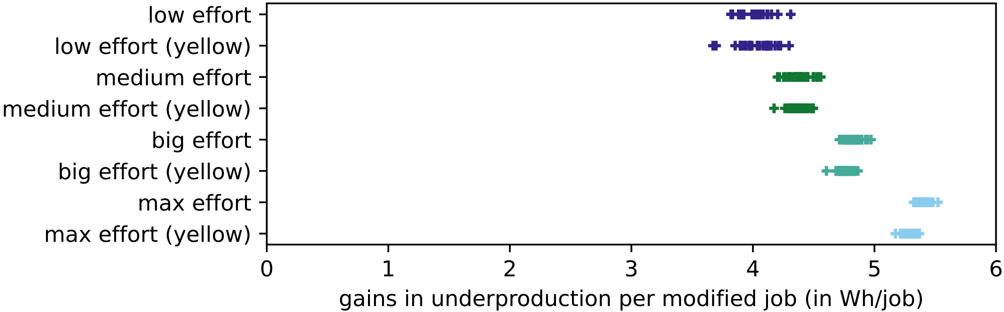

$\rightarrow$ marginal gains increase with α: "the more people who make an effort, the greater the impact of a user's additional effort"

$\rightarrow$ gains with yellow windows of the same scale than with red

Conclusion

- 3-state energy model and user behaviors to adapt job to energy consumption

Conclusion

- 3-state energy model and user behaviors to adapt job to energy consumption
- Possible improvements:
  - thresholds on instantaneous available energy
  - collaboration with the scheduler
  - more realistic replay method
  - social science studies (willingness to adopt behaviors, impact of eco-feedback)

Conclusion

- 3-state energy model and user behaviors to adapt job to energy consumption
- Possible improvements:
  - thresholds on instantaneous available energy
  - collaboration with the scheduler
  - more realistic replay method
  - social science studies (willingness to adopt behaviors, impact of eco-feedback)
- First step towards studying **sufficiency** and not **efficiency**

Conclusion

- 3-state energy model and user behaviors to adapt job to energy consumption
- Possible improvements:
  - thresholds on instantaneous available energy
  - collaboration with the scheduler
  - more realistic replay method
  - social science studies (willingness to adopt behaviors, impact of eco-feedback)
- First step towards studying **sufficiency** and not **efficiency**
- Simulation campaign **reproducible**

Merci !

- Questions?
- Do not hesitate to contact me :-)
  - www.irit.fr/~Mael.Madon
  - mael.madon@m4x.org