Scaling Energy Adaptive Applications for Sustainable Profitability

Fabien Hermenier, Giovanni Giuliani,
Andre Milani, Sophie Demassey

DC4Cities

Let existing and new data centres become energy adaptive

Adapting the power consumption to the availability of renewable energy

Being adapted to the requests of a Smart City Energy Management Authority

How to reconcile competing objectives ?

An economical approach to avoid overcommitment

Energy authority

regulate energy usage

Client

best performance

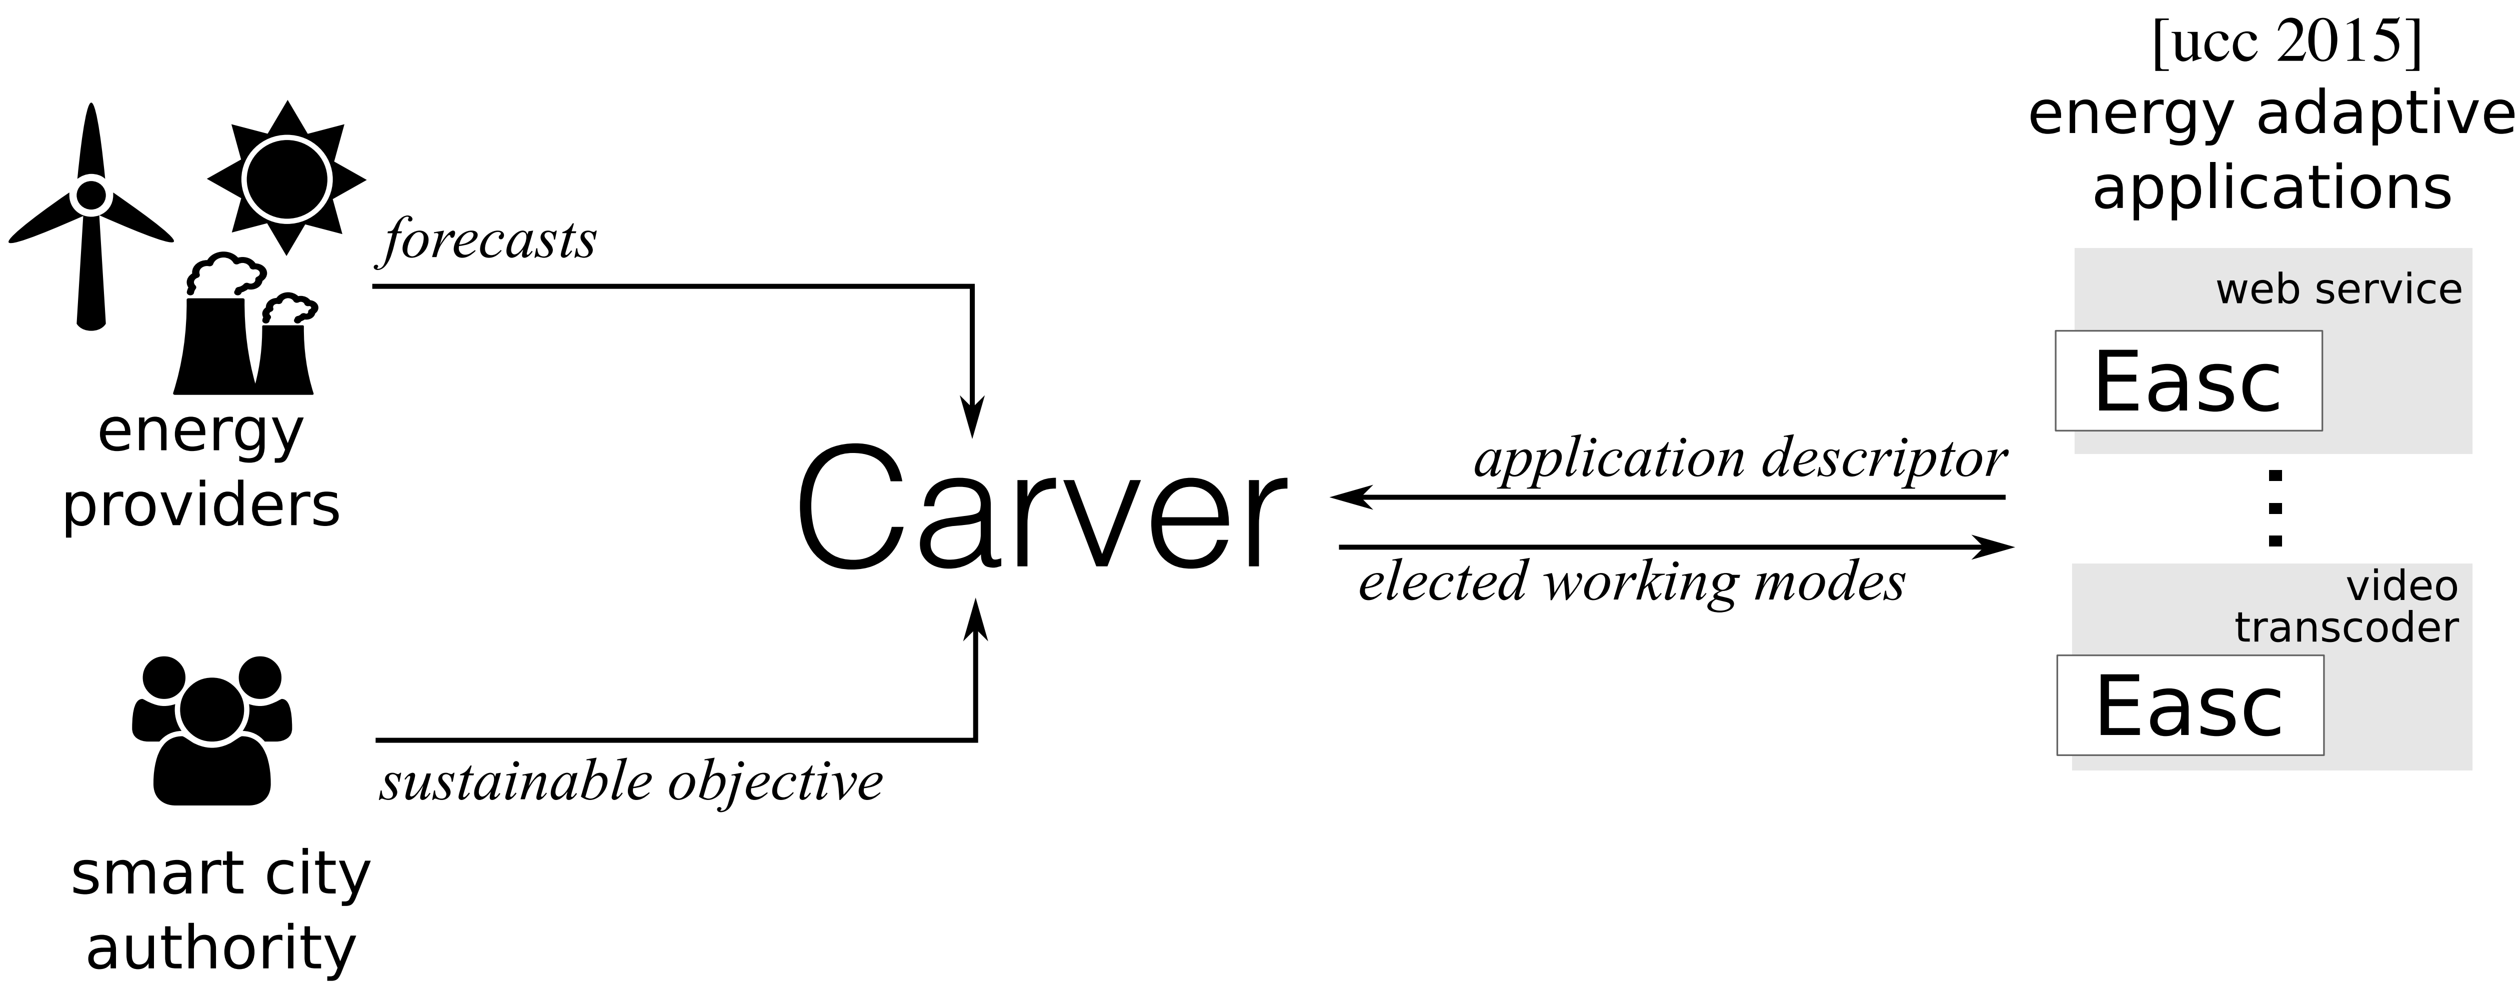

Orchestrates EASC for profitable sustainability

Optimises over 24-hrs time window
(96 time-slots of 15 minutes)

Called every 15 minutes
to accommodate uncertainty

The energy providers

Per time-slot data or forecasts
for the next 24 hours

% issues from renewable sources

price

capacity

Regulate energy usage through contracts


```
e",
"name":                gy Objective",
"descriptio            % green energy",
"type": "ENER          RTY",
"implementation        : "MUST",
"enabled": true,
"dataCenterId":          ,
"priority": 0,
"target": {
  "metric": "renewablePercentage",
  "operator": "gt",
  "value": 80
},
"timeFrame": {
  "recurrentExpression": "0 0 0 ? * SUN",
  "duration": "P1D"
},
"priceModifiers": [
  {
    "threshold": "80 %",
    "modifier": "0 EUR/pp"
  },
  {
    "threshold": "75 %",
    "modifier": "1 EUR/pp"
  }
```


day-to-day periods

Objective	At least X% renewable, Power budget, …
Incentive	Penalty function (flat, linear, composite)

```yaml
- name: pageIndexing
  businessUnit: kPage
  SLO:
    - timeFrom: 00:00:00
      timeTo: 24:00:00
      cumulativeObjective: !amount '200 kPage'
      basePrice: !amount '100 EUR'
      priceModifiers:
      - threshold: !amount '200 kPage'
        penalty: !amount '0 EUR/kPage'
      - threshold: !amount '100 kPage'
        penalty: !amount '-1 EUR/kPage'
      - threshold: !amount '0 kPage'
        penalty: !amount '-100 EUR'
  workingModes:
    - name: WM0
      actuator: bin/run.sh WM0
      performance: !amount '0 kPage/h'
      power: !amount '6 W'
      transitions:
        - target: WM1
          performanceCost: !amount '1 kPage'
        - target: WM2
          performanceCost: !amount '2 kPage'
    - name: WM1
      actuator: bin/run.sh WM1
      performance: !amount '32 kPage/h'
      power: !amount '27 W'
      transitions:
        - target: WM2
          performanceCost: !amount '2 kPage'
    - name: WM2
      actuator: bin/run.sh WM2
      performance: !amount '60 kPage/h'
      power: !amount '33 W'
```


EASC characterisation

```yaml
- name: pageIndexing
  businessUnit: kPage
  SLO:
    - timeFrom: 00:00:00
      timeTo: 24:00:00
      cumulativeObjective: !amount '200 kPage'
      basePrice: !amount '100 EUR'
      priceModifiers:
      - threshold: !amount '200 kPage'
        penalty: !amount '0 EUR/kPage'
      - threshold: !amount '100 kPage'
        penalty: !amount '-1 EUR/kPage'
      - threshold: !amount '0 kPage'
        penalty: !amount '-100 EUR'
  workingModes:
    - name: WM0
      actuator: bin/run.sh WM0
      performance: !amount '0 kPage/h'
      power: !amount '6 W'
      transitions:
        - target: WM1
          performanceCost: !amount '1 kPage'
        - target: WM2
          performanceCost: !amount '2 kPage'
    - name: WM1
      actuator: bin/run.sh WM1
      performance: !amount '32 kPage/h'
      power: !amount '27 W'
      transitions:
        - target: WM2
          performanceCost: !amount '2 kPage'
    - name: WM2
      actuator: bin/run.sh WM2
      performance: !amount '60 kPage/h'
      power: !amount '33 W'
```

Service Level Objectives

daily-based
deferrable workloads

slot-based
non-deferrable workloads

flat or linear pricing policies

```yaml
- name: pageIndexing
  businessUnit: kPage
  SLO:
    - timeFrom: 00:00:00
      timeTo: 24:00:00
      cumulativeObjective: !amount '200 kPage'
      basePrice: !amount '100 EUR'
      priceModifiers:
        - threshold: !amount '200 kPage'
          penalty: !amount '0 EUR/kPage'
        - threshold: !amount '100 kPage'
          penalty: !amount '-1 EUR/kPage'
        - threshold: !amount '0 kPage'
          penalty: !amount '-100 EUR'
  workingModes:
    - name: WM0
      actuator: bin/run.sh WM0
      performance: !amount '0 kPage/h'
      power: !amount '6 W'
      transitions:
        - target: WM1
          performanceCost: !amount '1 kPage'
        - target: WM2
          performanceCost: !amount '2 kPage'
    - name: WM1
      actuator: bin/run.sh WM1
      performance: !amount '32 kPage/h'
      power: !amount '27 W'
      transitions:
        - target: WM2
          performanceCost: !amount '2 kPage'
    - name: WM2
      actuator: bin/run.sh WM2
      performance: !amount '60 kPage/h'
      power: !amount '33 W'
```

The working modes

discretise elastic applications states

performance level
power consumption
transition cost
actuator

manual or automatic calibration

```yaml
- name: pageIndexing
  businessUnit: kPage
  SLO:
    - timeFrom: 00:00:00
      timeTo: 24:00:00
      cumulativeObjective: !amount '200 kPage'
      basePrice: !amount '100 EUR'
      priceModifiers:
      - threshold: !amount '200 kPage'
        penalty: !amount '0 EUR/kPage'
      - threshold: !amount '100 kPage'
        penalty: !amount '-1 EUR/kPage'
      - threshold: !amount '0 kPage'
        penalty: !amount '-100 EUR'
  workingModes:
    - name: WM0
      actuator: bin/run.sh WM0
      performance: !amount '0 kPage/h'
      power: !amount '6 W'
      transitions:
        - target: WM1
          performanceCost: !amount '1 kPage'
        - target: WM2
          performanceCost: !amount '2 kPage'
    - name: WM1
      actuator: bin/run.sh WM1
      performance: !amount '32 kPage/h'
      power: !amount '27 W'
      transitions:
        - target: WM2
          performanceCost: !amount '2 kPage'
    - name: WM2
      actuator: bin/run.sh WM2
      performance: !amount '60 kPage/h'
      power: !amount '33 W'
```

3 use cases inside DC4Cities

Webservice,
video transcoding,
e-health

The underlying problem

For every time-slot

find 1 working mode per EASC

using available energy (viable dispatch)

96 **NP-hard** bin-packing problems to solve

+ transition costs
+ various Energy Authority policies
+ Instant and cumulative objectives
+ generic pricing policies

= I dont want a heuristic full of corner cases

CHOCO

An Open-Source java library
for constraint programming

$$
\begin{aligned}
\mathcal{X} &= \{x_1, x_2, x_3\} \\
\mathcal{D}(x_i) &= [0, 2]\,, \forall x_i \in \mathcal{X} \\
\mathcal{C} &= \left\{ \begin{array}{r} c_1 : x_1 < x_2 \\ c_2 : x_1 + x_2 \geq 2 \\ c_3 : x_1 < x_3 \end{array} \right.
\end{aligned}
$$

deterministic composition
high-level constraints

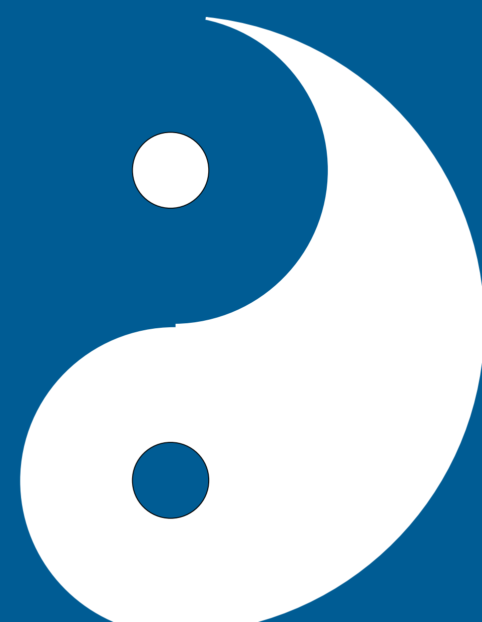

the right model
for the right problem

Modelling - TLDR;

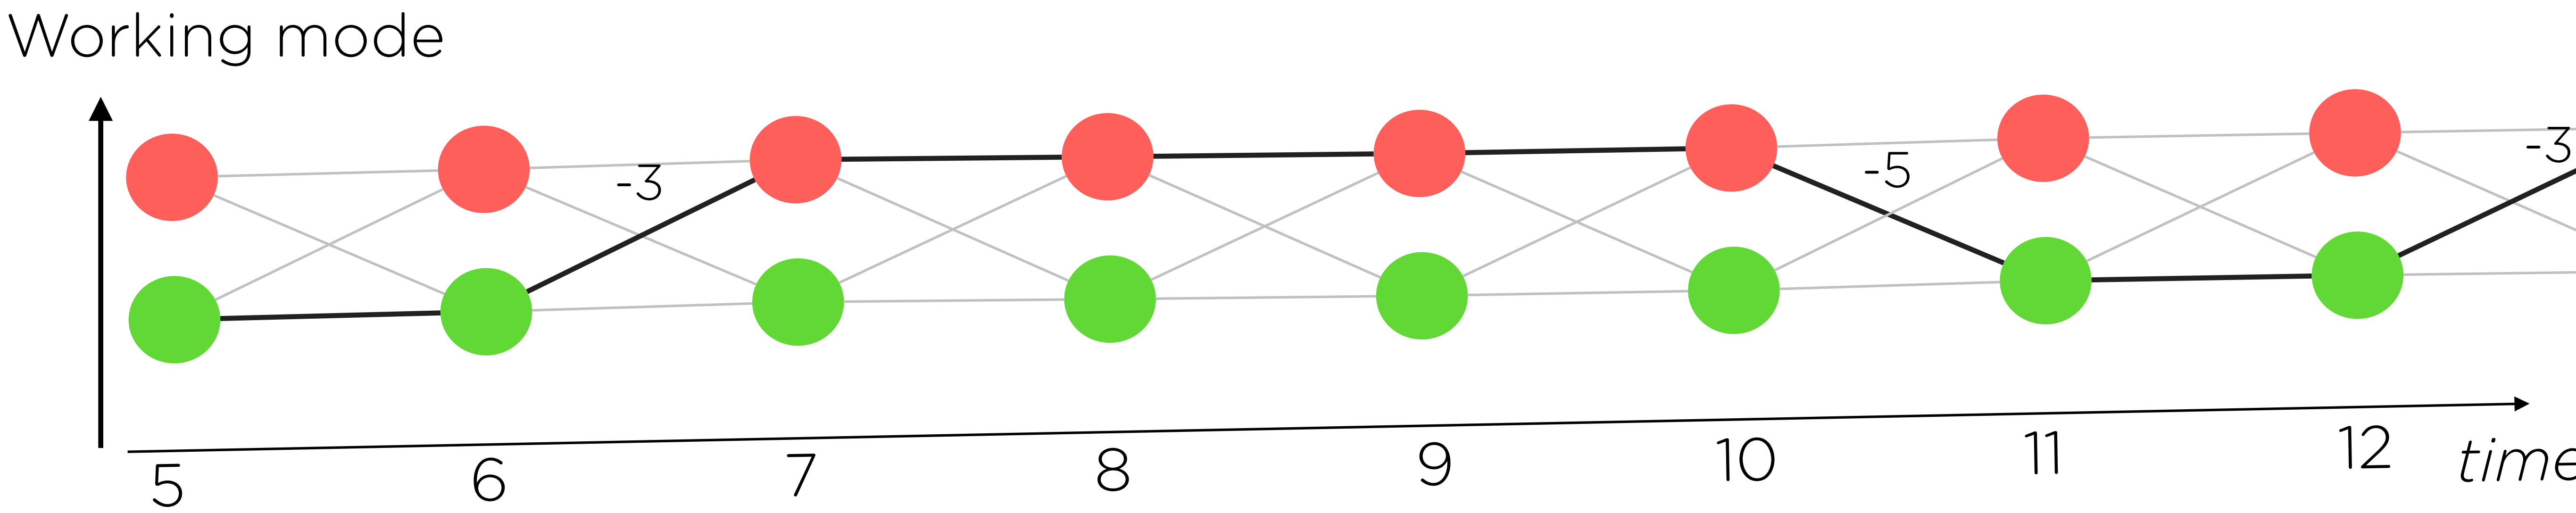

An automaton to model each EASC life-cycle

Counters to accumulate daily incomes

Energy to use is dispatched among
the sources

Energy authority pricing policy
over the cumulative daily usage

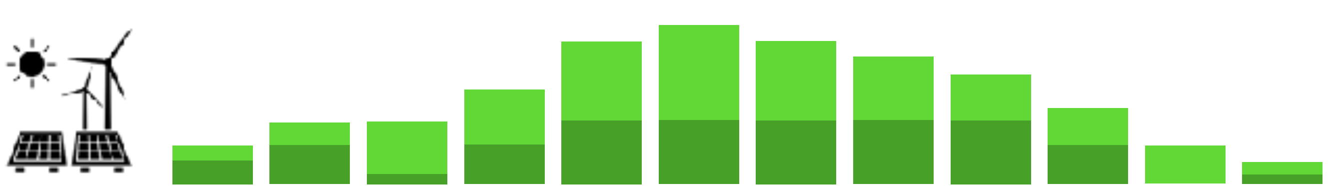

Reconciliations

“Trading energy, performance and
energy authority conformance
to maximise the daily running costs

CHOCO

Evaluating Carver

real deployments
3 testbeds running production softwares

HP Lab @ Milan

Web service
Mixed EASCs
Bare metal

CSUC @ Barcelona

Video transcoding
openNebula stack

Trento @ Italy

e-health
Openstack

Testbed

20 HP moonshot cartridges in 2 chassis
no virtualisation layer
20 Watts peak

The grid

un-pure energy
market price

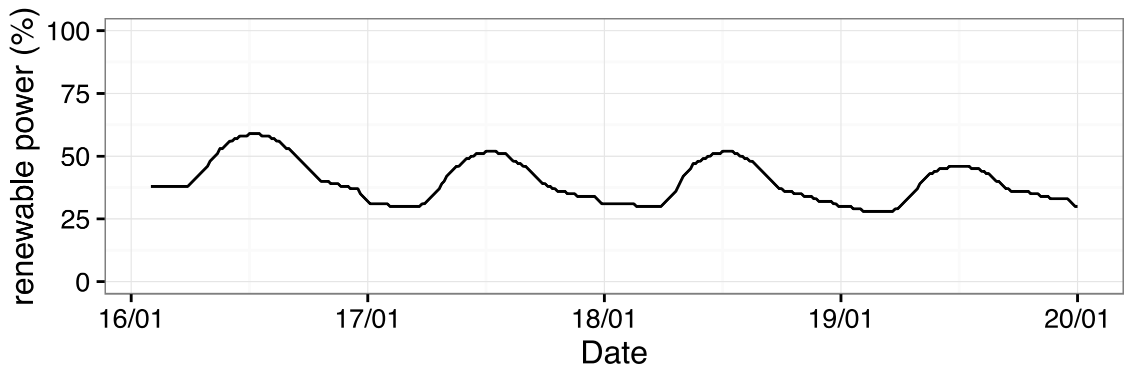

The green

8m2 of solar panels

4 typical days
from a 1 y. data collect

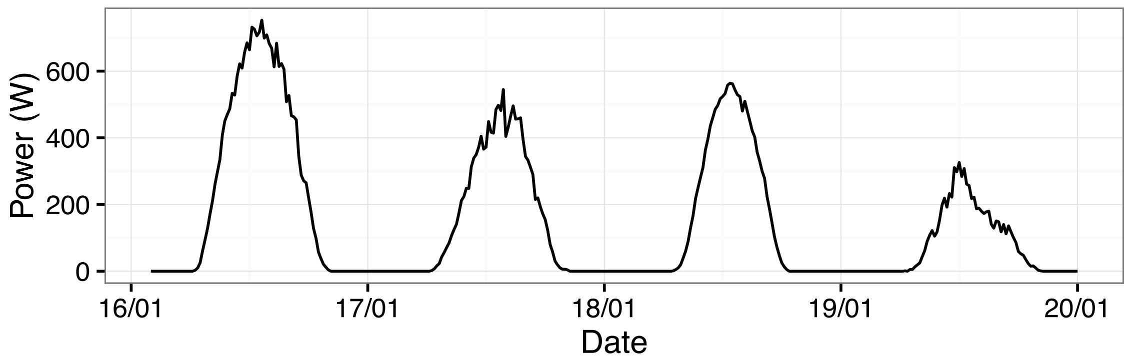

Energy authority expectations

65% renewable target
100€ / pp.

production software stack
with 3 EASCs

Injectors mimic the
production workload

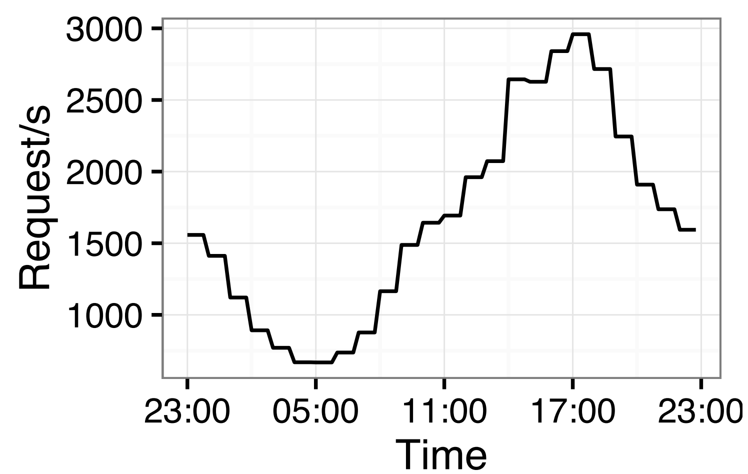

Web service

instant SLO
6 working modes
penalty from a step function

E-learning, G-learning

cumulative SLOs
3 working modes
penalty from a linear function

Carver

Sustainable profitability

Perf

Production mode

Cache built at midnight

Lowest suitable working mode
for the WebSite

VS

Green

Most efficient working mode
when renewable energy is here

Lowest otherwise

Perf

Ignores energy availability

Workload affinity:
human work during days

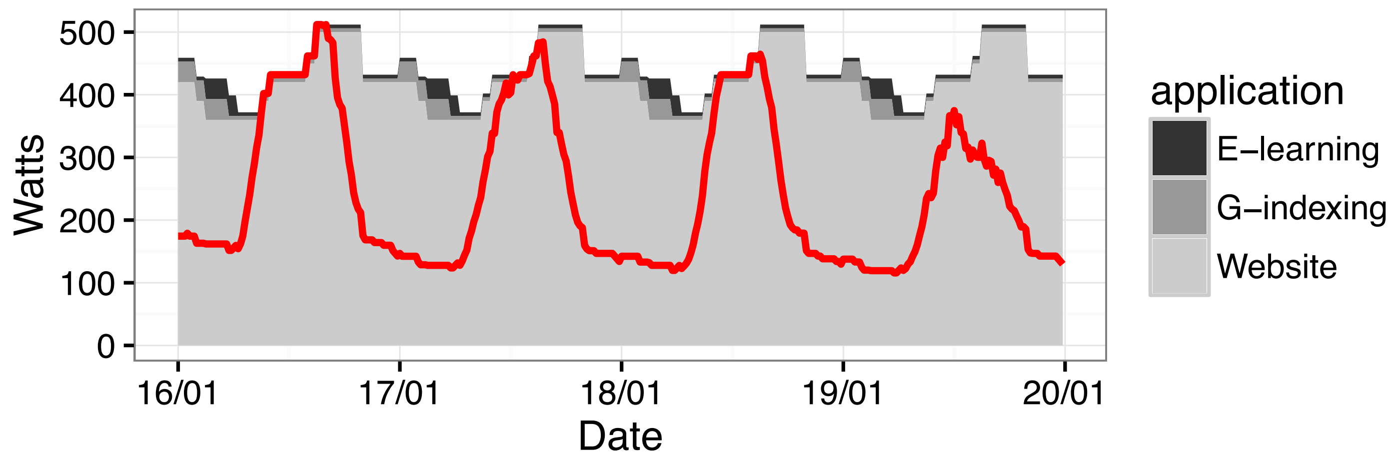

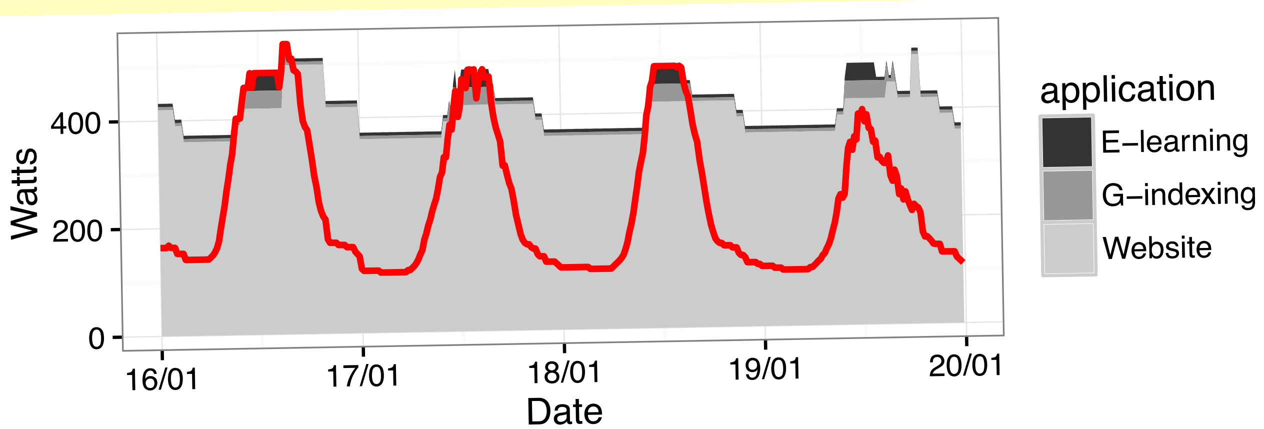

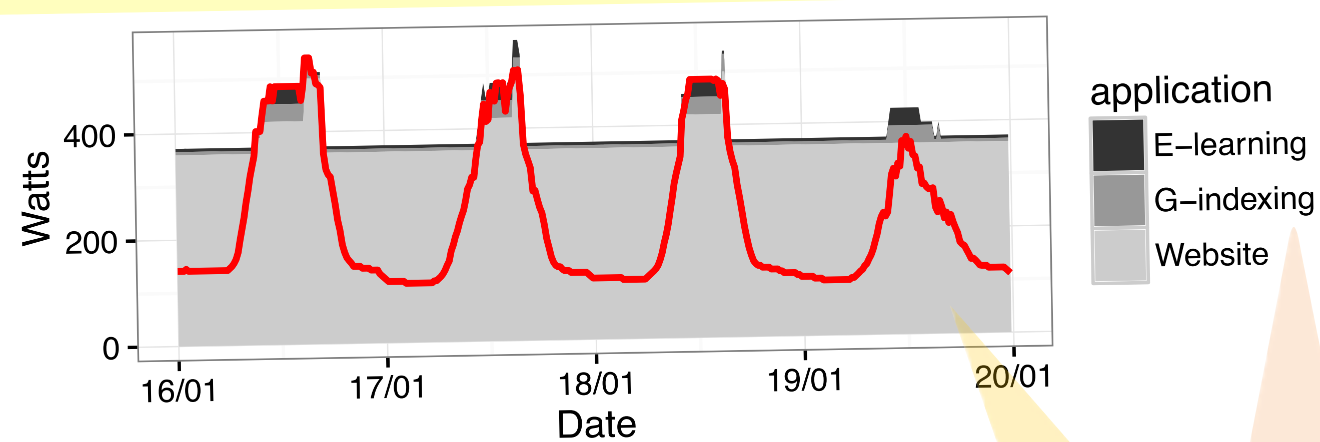

Carver

Green

Batch activities follow the sun
Green is binary

Carver do not over-commit

Carver sticked to the green threshold

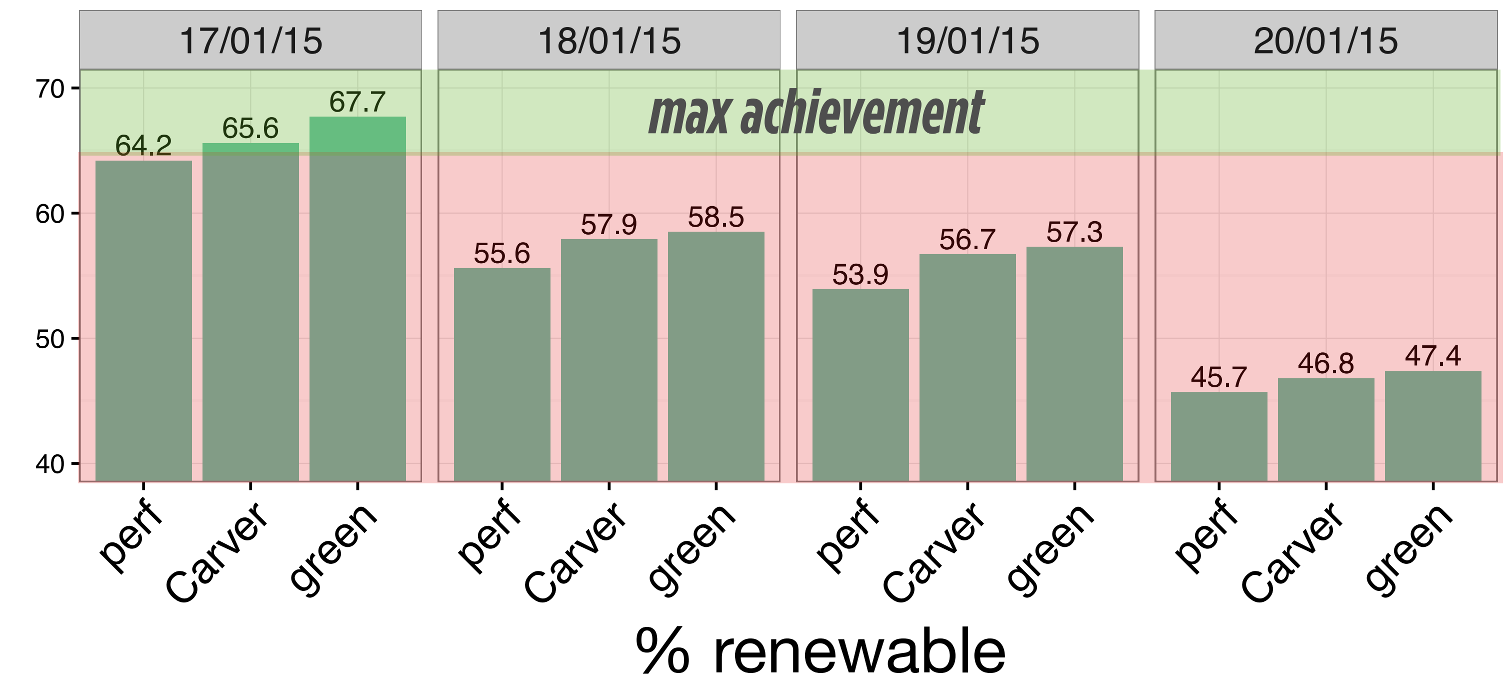

Carver sticked to the SLO

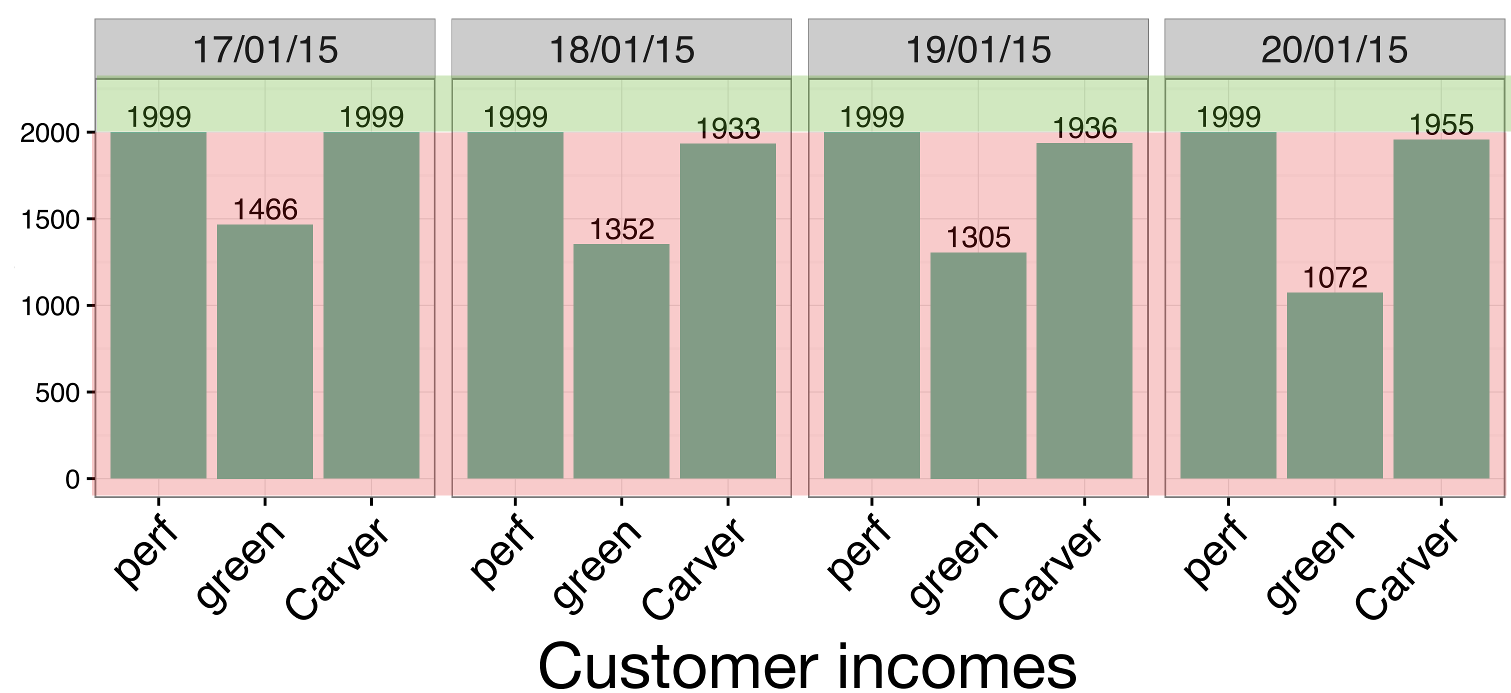

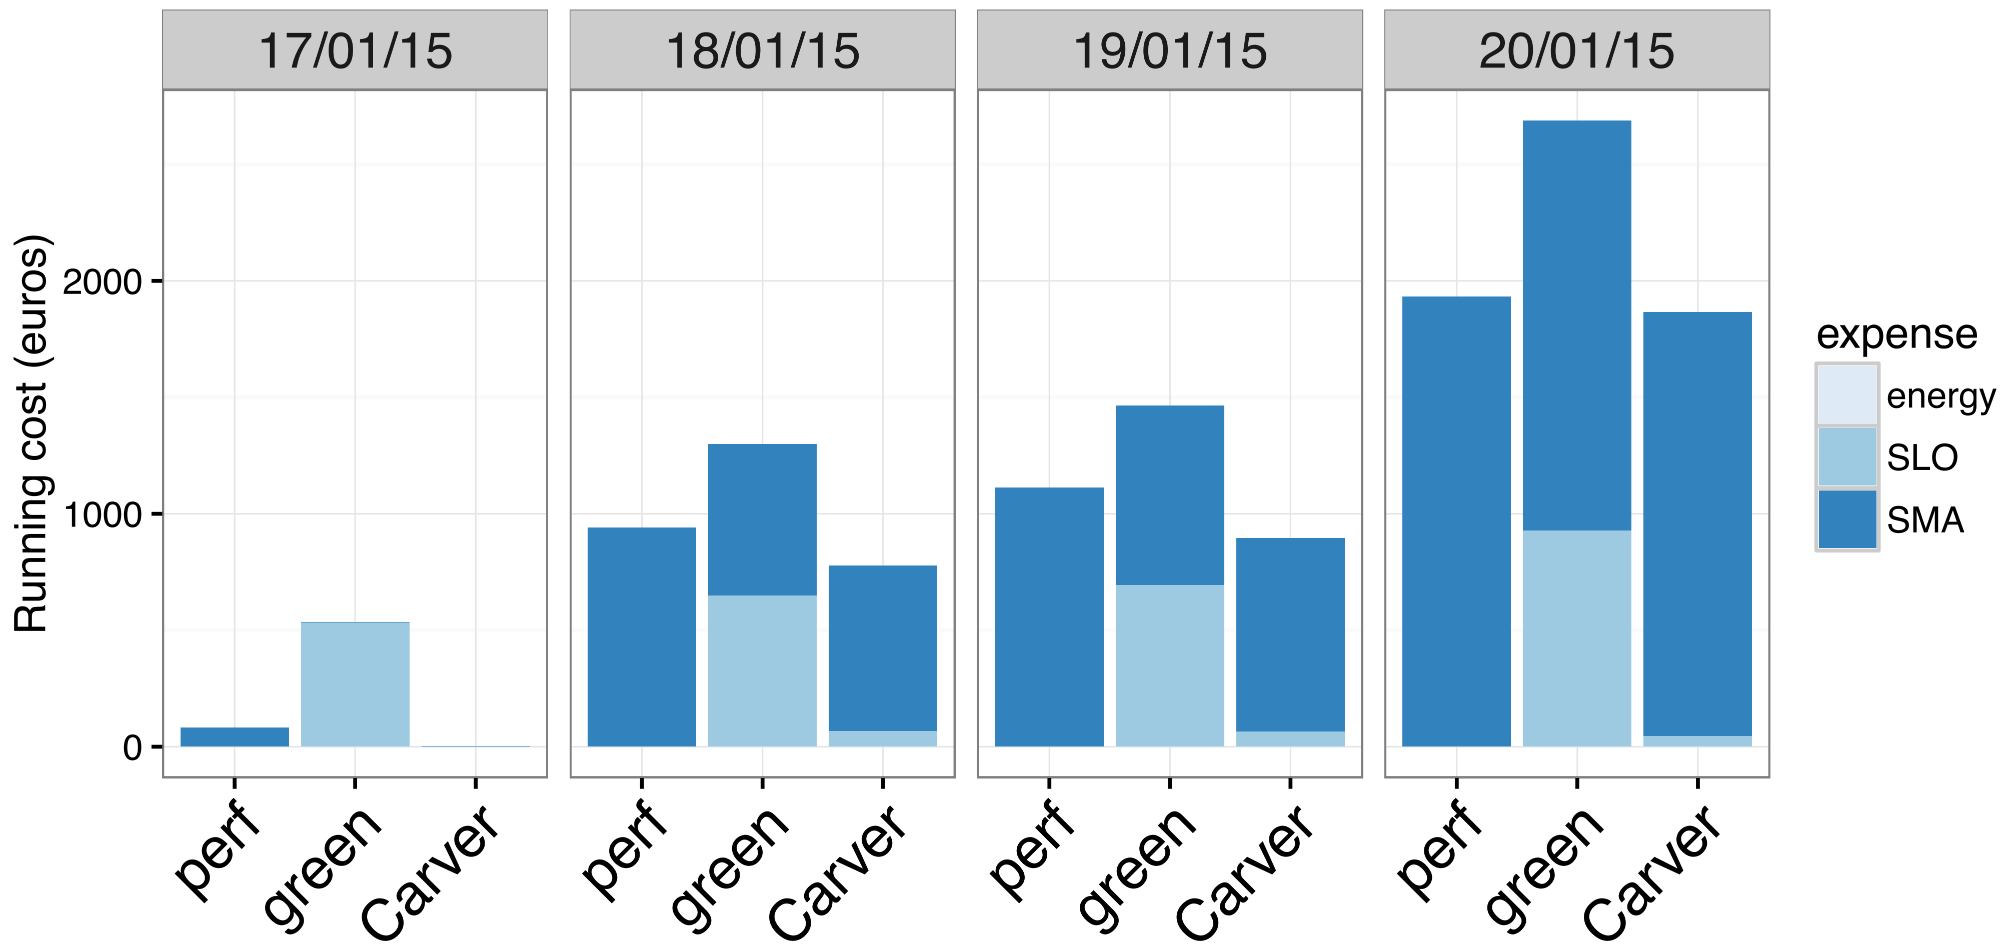

Green neglects the clients

Perf neglects the energy authority

Carver trades

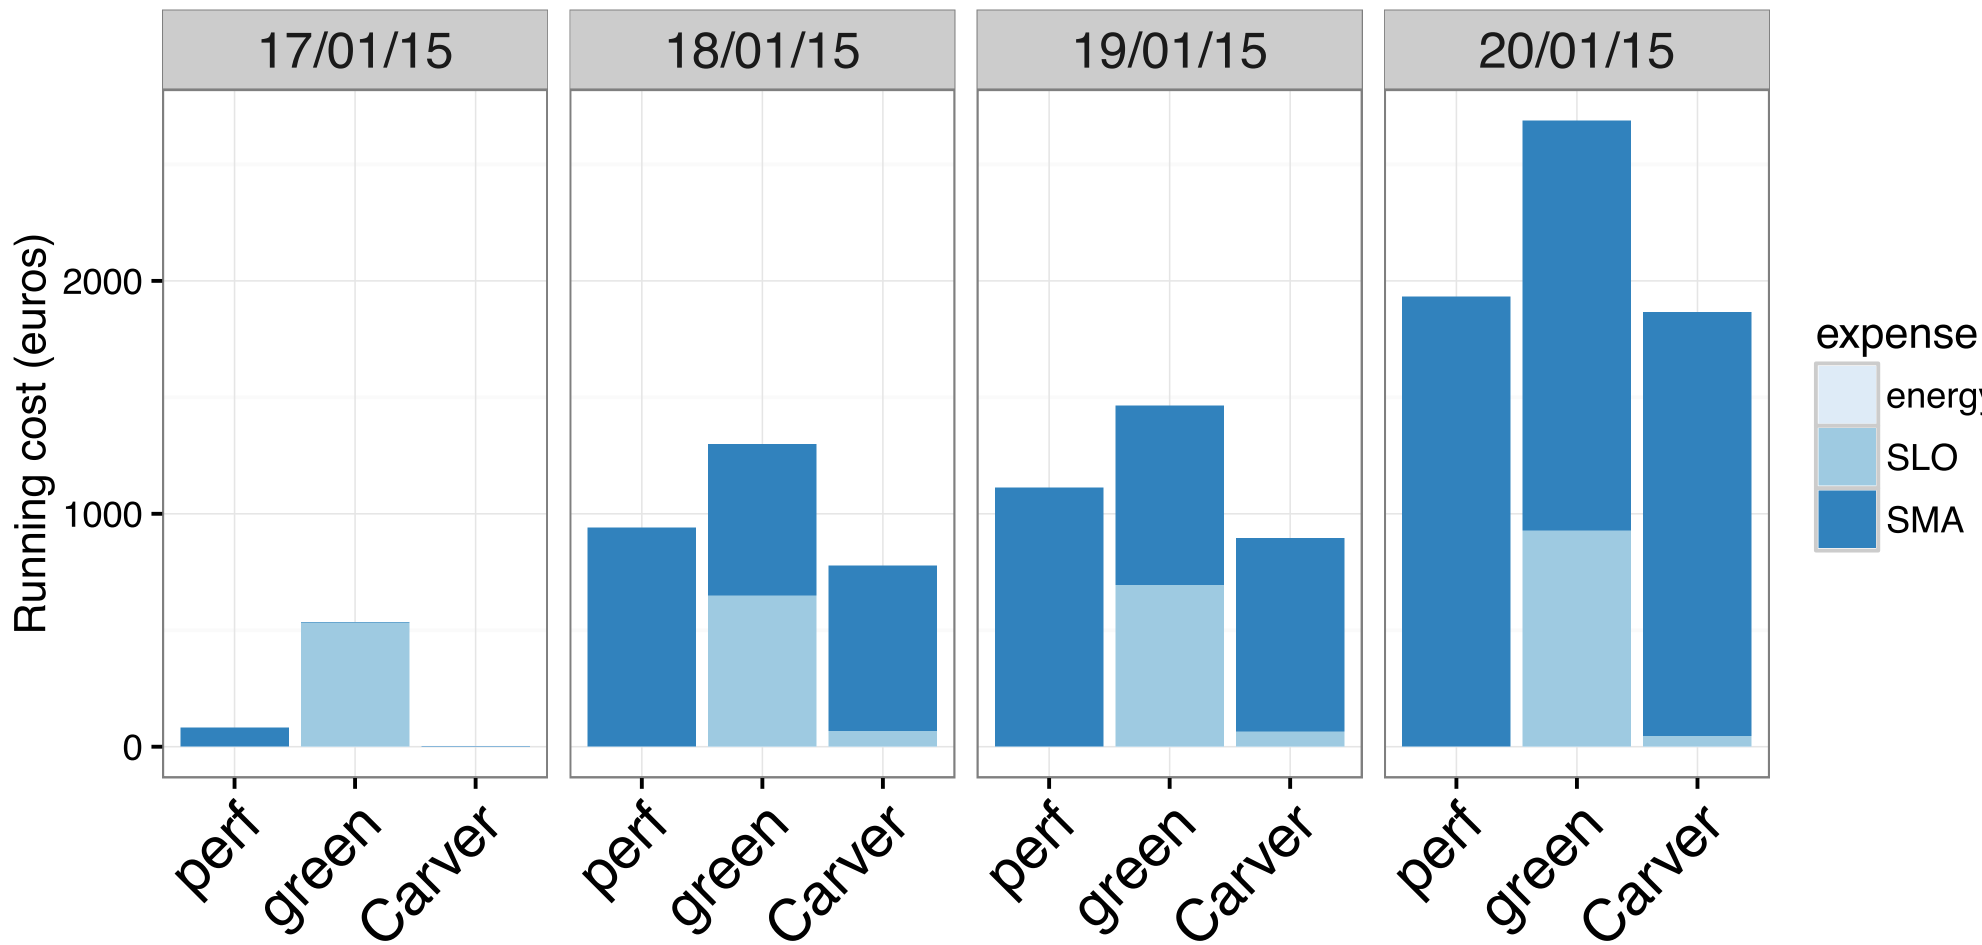

The balance is still slightly pure performance oriented

@Day 1  nothing to do to please the energy authority (natural workload affinity)

@Day 4  Only minor trading possibilities

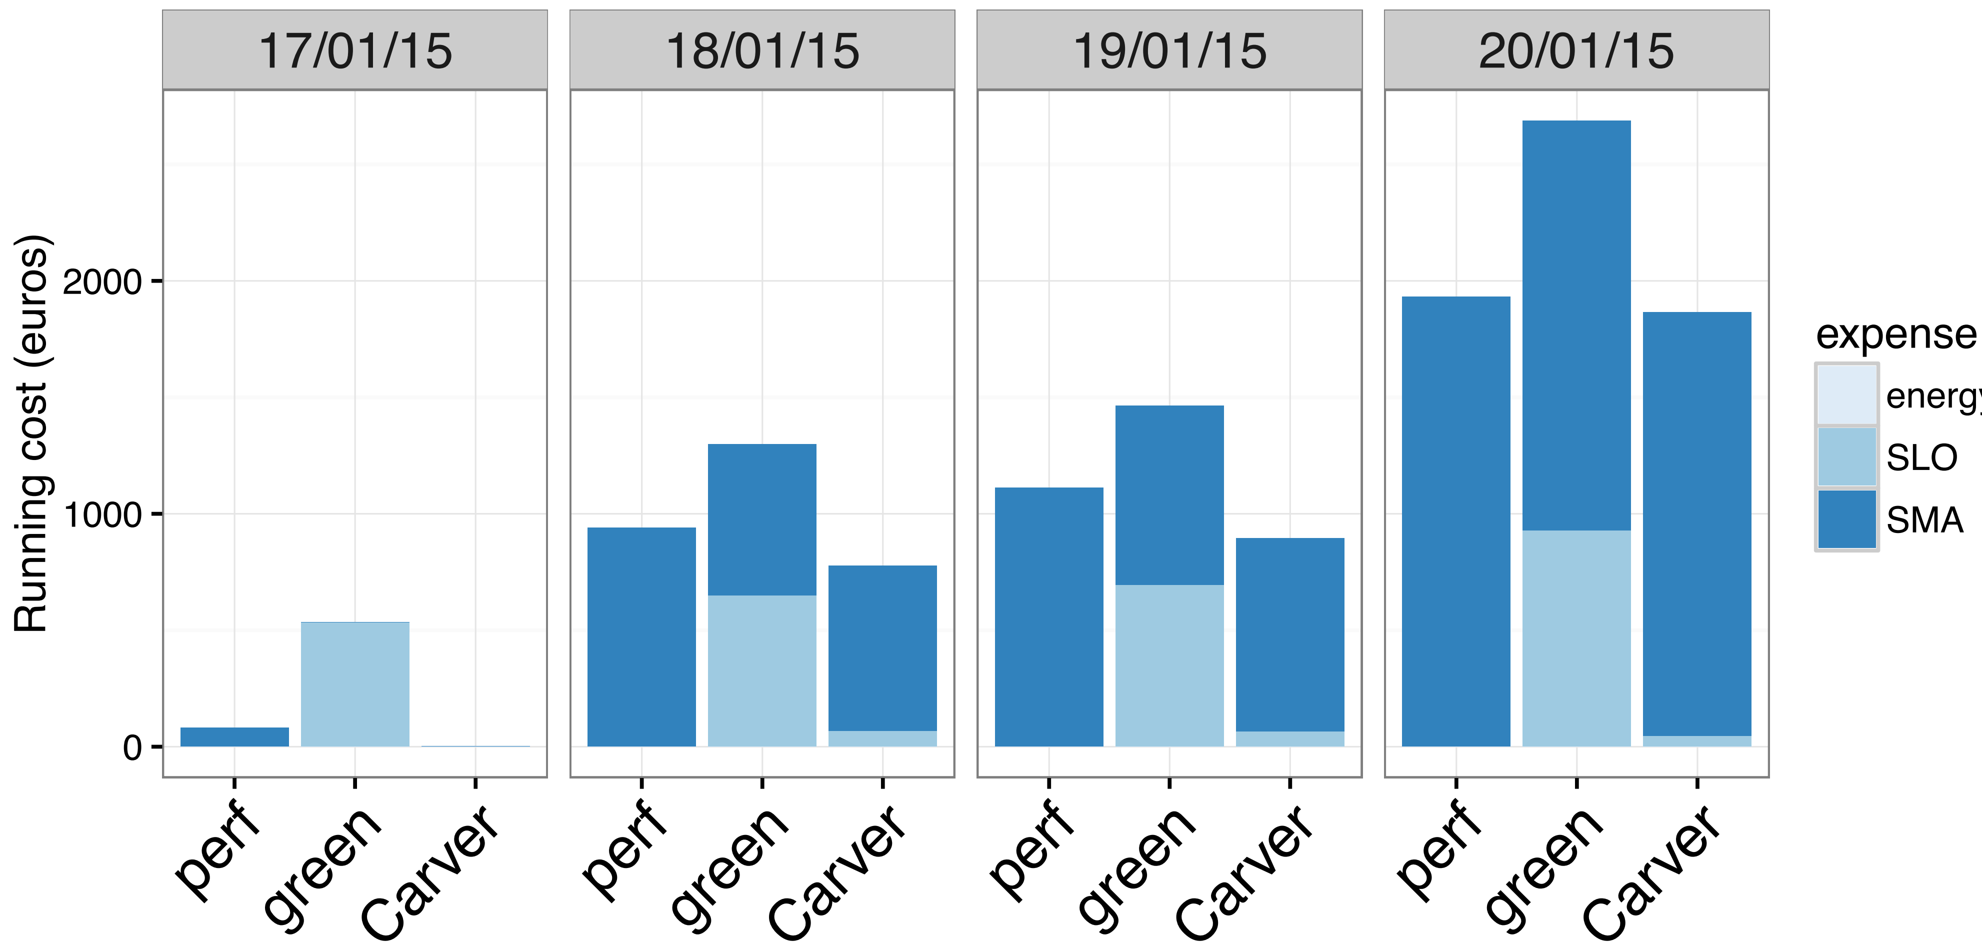

Economically speaking

Energy price does not play a role

every price <<< pricing policies
to avoid bankruptcy

small data centres
human resources / software support dominates

Lesson learned

being generic is costly	CP composability helps but data makes the problem static analysis for stronger models
require favorable hardware and software	variable working modes energy-efficient hardware
a multifacet tool	PV array sizer cost modelling prospective deployment

Sustainable profitability to motivate energy transition

Carver is looking for sun and profit

A flexible solving algorithm to cope with the problem variability

Conciliation possibilities validated on industrial testbeds

A multi-faceted tool

http://www.dc4cities.eu

deliverables, scientific publications, trial results, software repositories