AI for EV Charging: Quantifying and exploiting flexibility in EV charging with data-driven modeling and reinforcement learning

Chris Develder et al.


Nasrin


Manu


Johannes


Matthias

GHENT UNIVERSITY

IDLab INTERNET & DATA LAB

imec

AI for EV Charging: C... g and
expl... arging
with ... d
reinf...

AI for EV charging:
KNOWING and
CONTROLLING electric
vehicle charging

Chris Develder et al.

Nasrin Manu Johannes Matthias

GHENT UNIVERSITY

IDLab INTERNET & DATA LAB

imec

SMART GRID

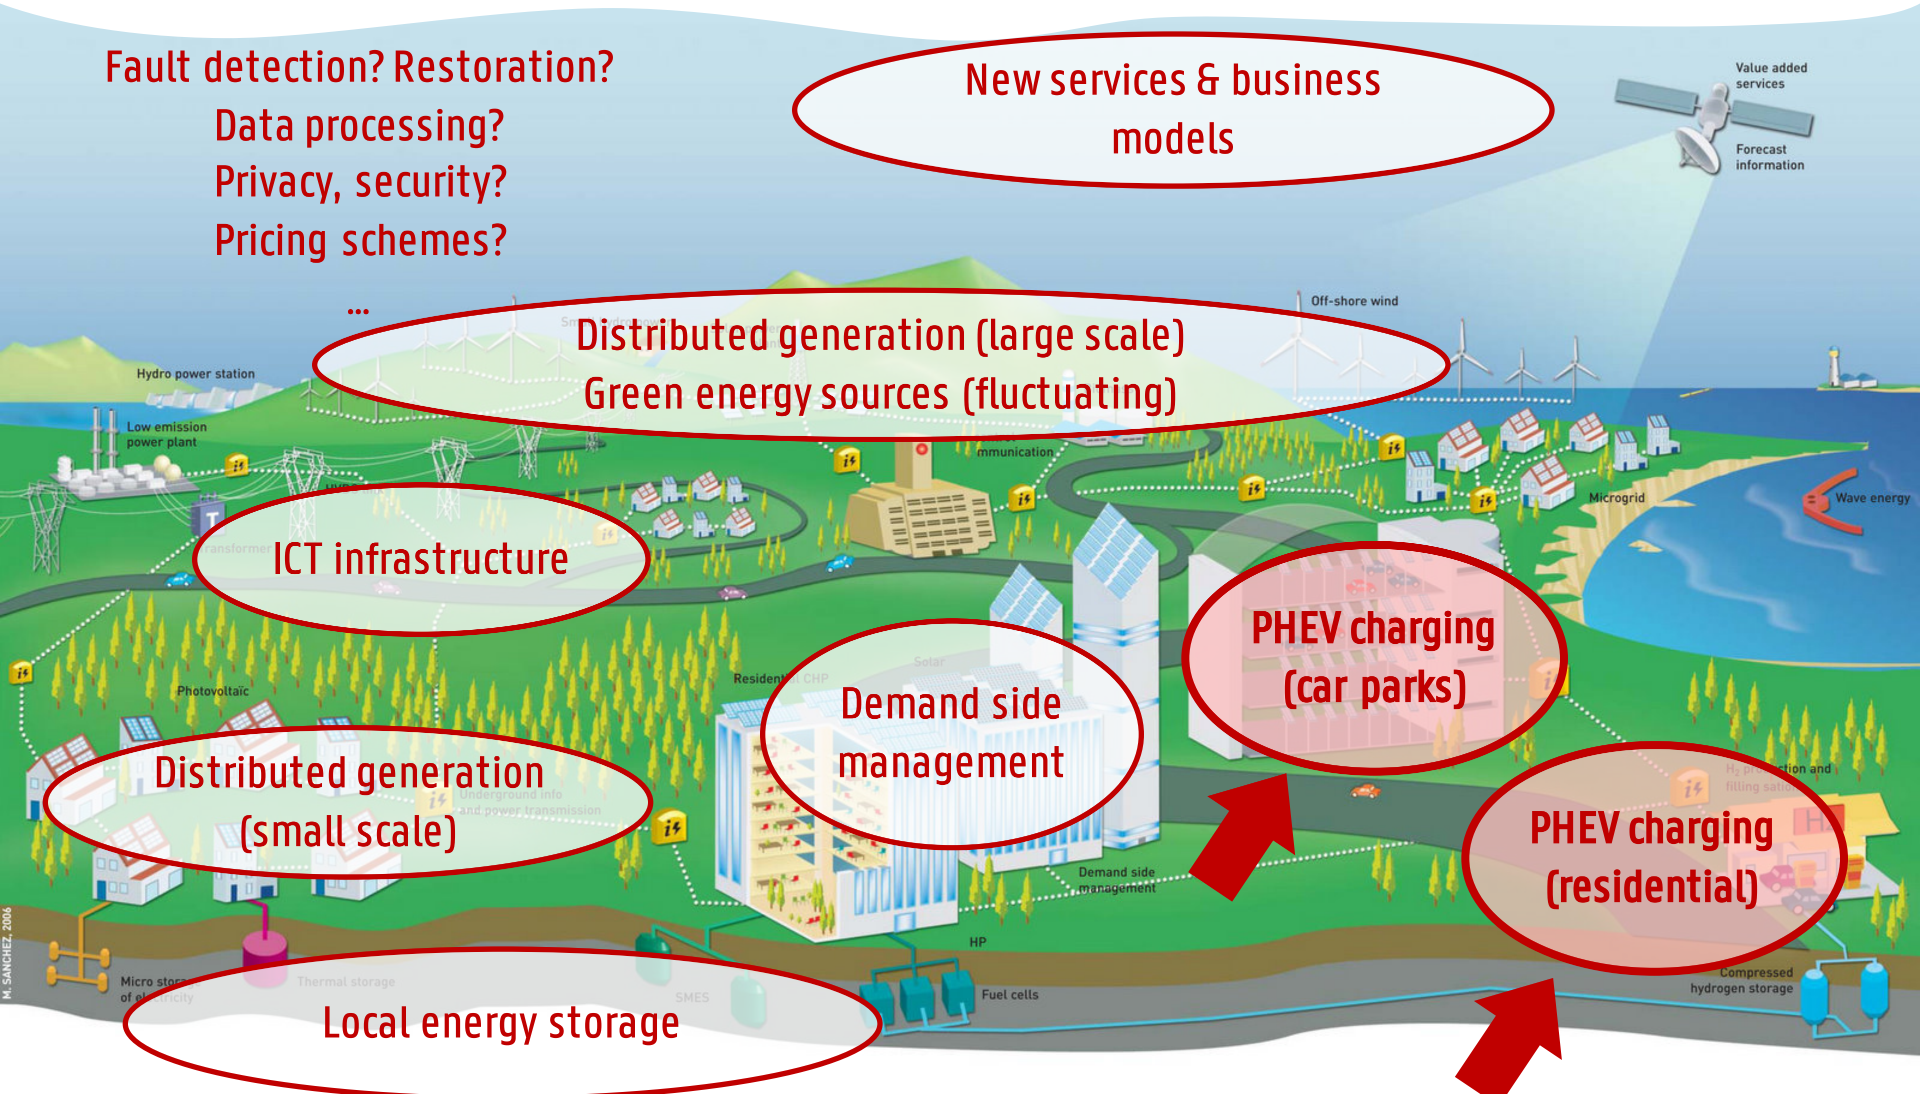

GHENT UNIVERSITY

OUTLINE

- PART I: “Knowing” EV charging session behavior
  - Quantitative analysis of real-world EV charging data
  - Flexibility <u>potential</u> quantification
  - Flexibility <u>exploitation</u> quantification

- PART II: “Controlling” EV charging sessions
  - Reinforcement learning model
  - Experimental evaluation through simulations

PART I: DATA ANALYTICS
"Knowing" EV charging session behavior

N. Sadeghianpourhamami, N. Refa, M. Strobbe and C. Develder, **"Quantitive analysis of electric vehicle flexibility: A data-driven approach"**, Int. J. Electr. Power Energy Syst., Vol. 95, Feb. 2018, pp. 451-462.

N. Sadeghianpourhamami, D.F. Benoit, D. Deschrijver and C. Develder, **"Bayesian cylindrical data modeling using Abe-Ley mixtures"**, Appl. Math. Model., Vol. 68, Apr. 2019, pp. 629-642.

C. Develder, N. Sadeghianpourhamami, M. Strobbe and N. Refa, **"Quantifying flexibility in EV charging as DR potential: Analysis of two real-world data sets"**, in Proc. 7th IEEE Int. Conf. Smart Grid Communications (SmartGridComm 2016), Sydney, Australia, 6-9 Nov. 2016, pp. 600-605.

Thanks Nasrin!

Characterization of real-world EV charging sessions

DATASETS: IMOVE (BE) AND ELAADNL

	IMOVE	ELAADNL
PERIOD	03/2012 – 03/2013	01/2012 – present
# SESSIONS	8.5 k	> 2 M
# USERS	134	about 53 000
CAR TYPE	Full EV	Unknown mix
CHARGE POINT	At home	Public
TRIP DETAILS	Yes	No

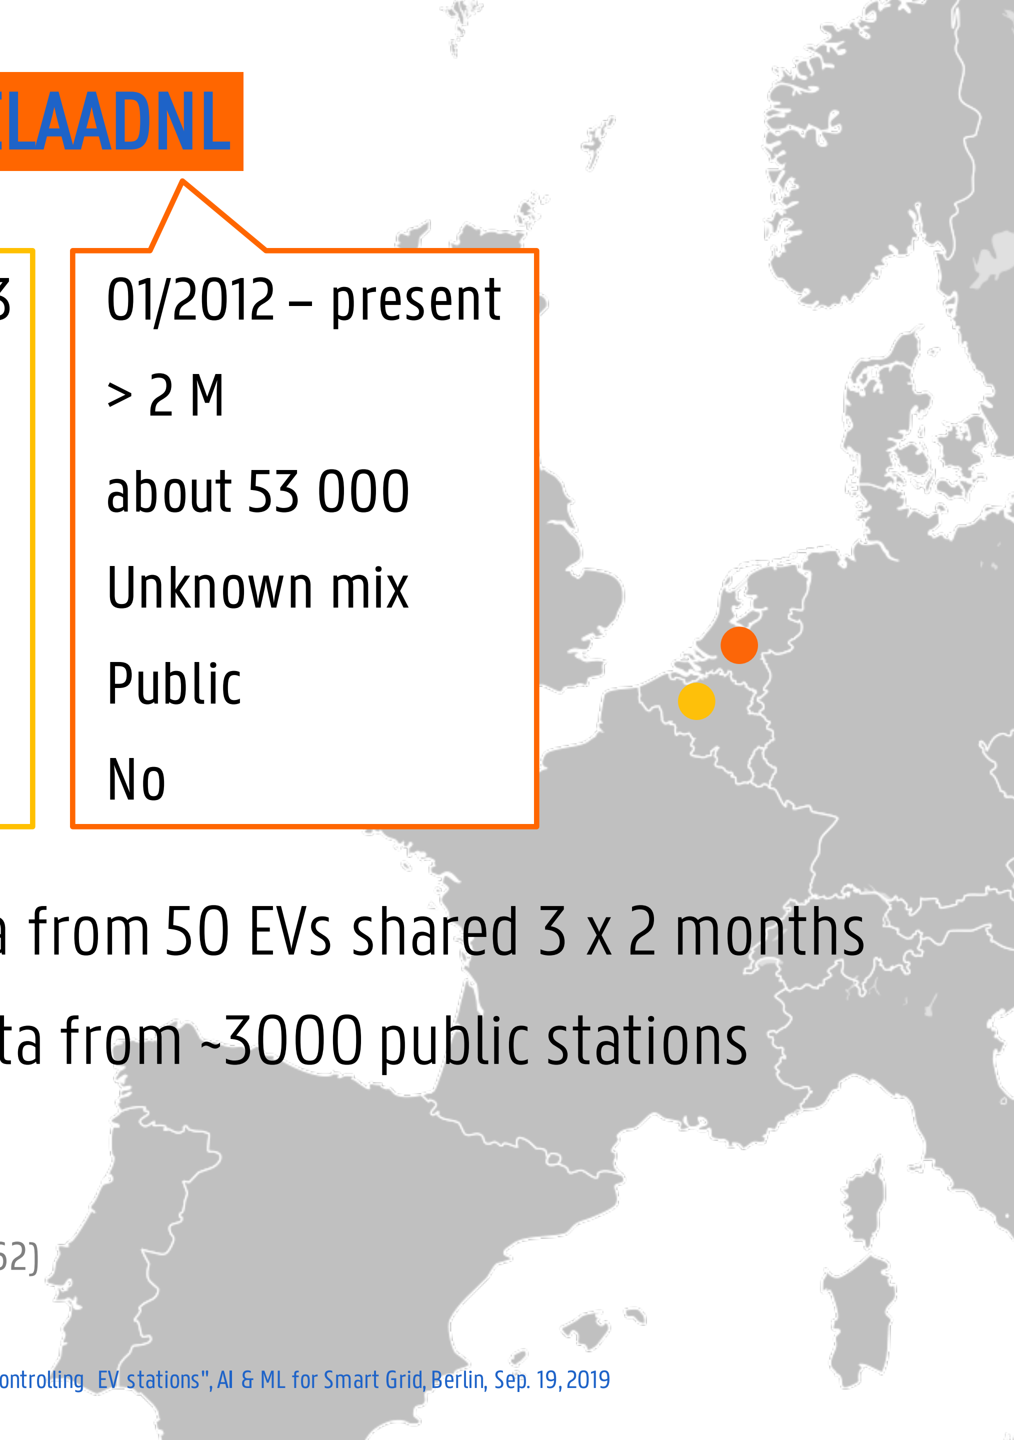

iMove: Flemish EV field trial; data from 50 EVs shared 3 x 2 months

ELaadNL: EV innovation in NL; data from ~3000 public stations

* : Analysis on data from 1 Jan.–31 Mar. 2015 (N = 90 562)

TYPICAL ARRIVAL AND DEPARTURE TIMES (1/2)

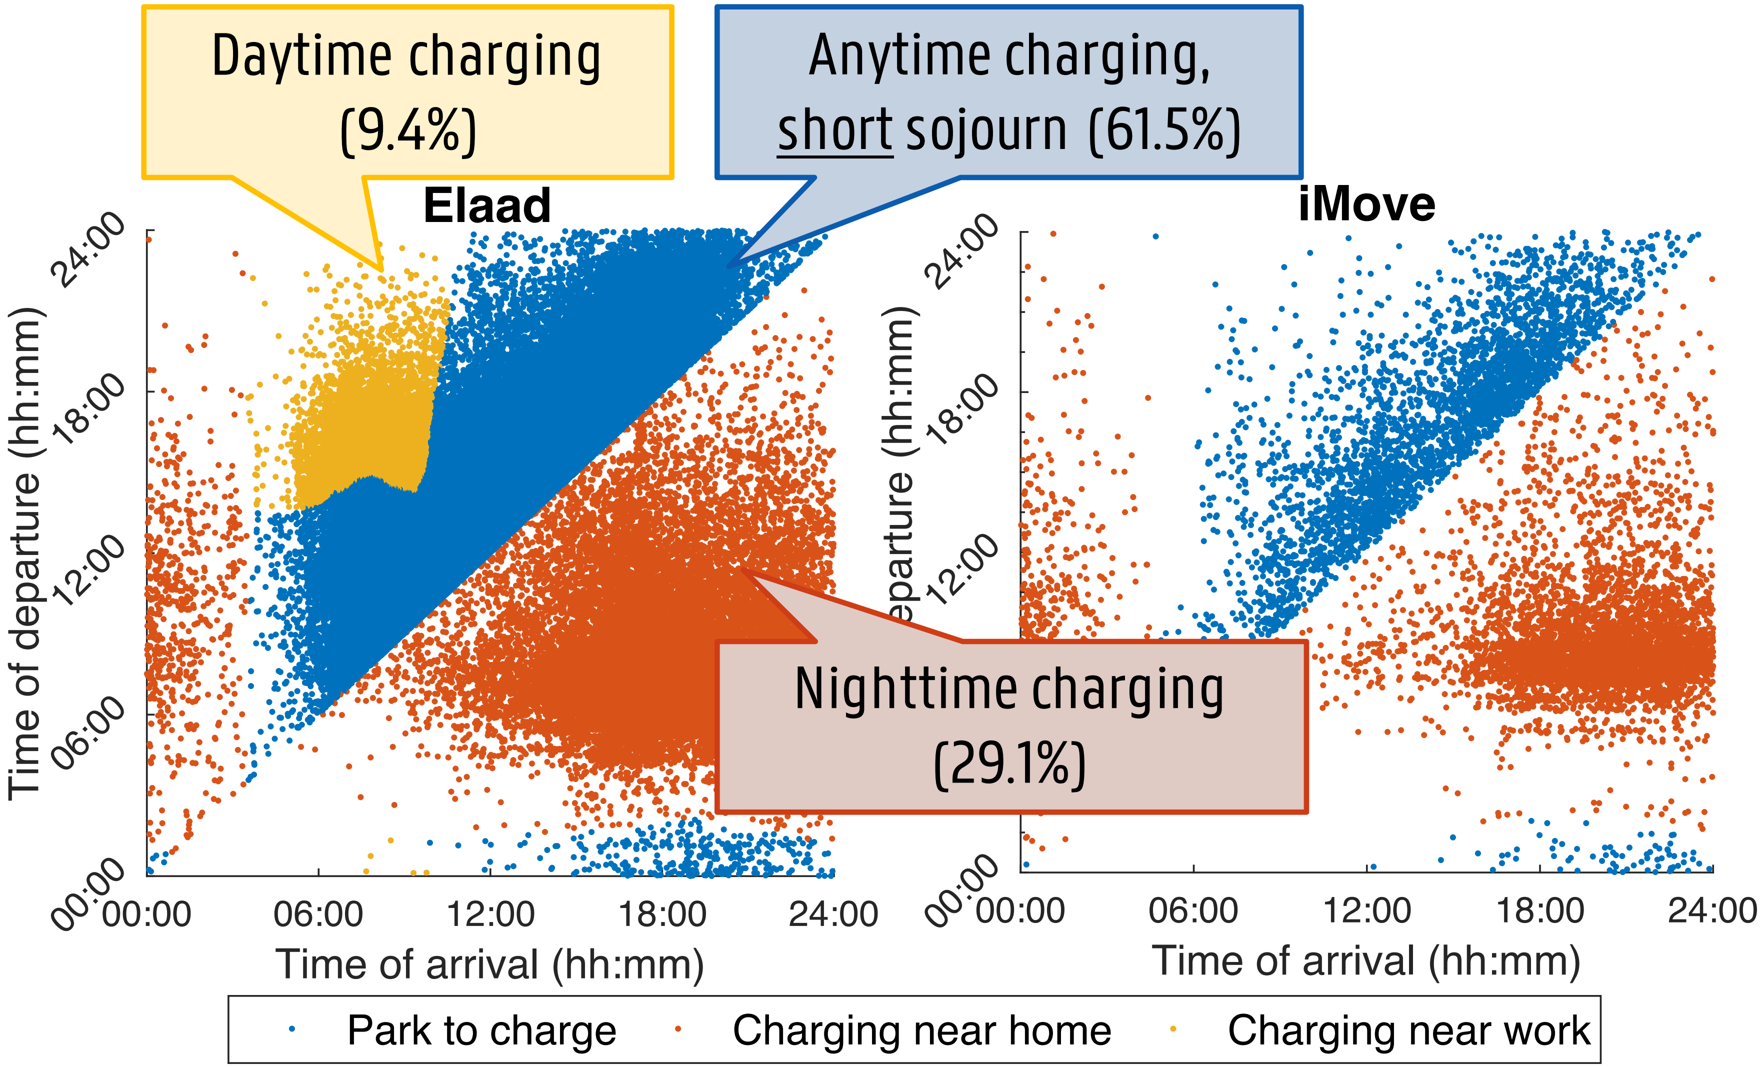

TYPICAL ARRIVAL AND DEPARTURE TIMES (2/2)

Anytime charging, short sojourn (61.5%)

Nighttime charging (29.1%)

Daytime charging (9.4%)

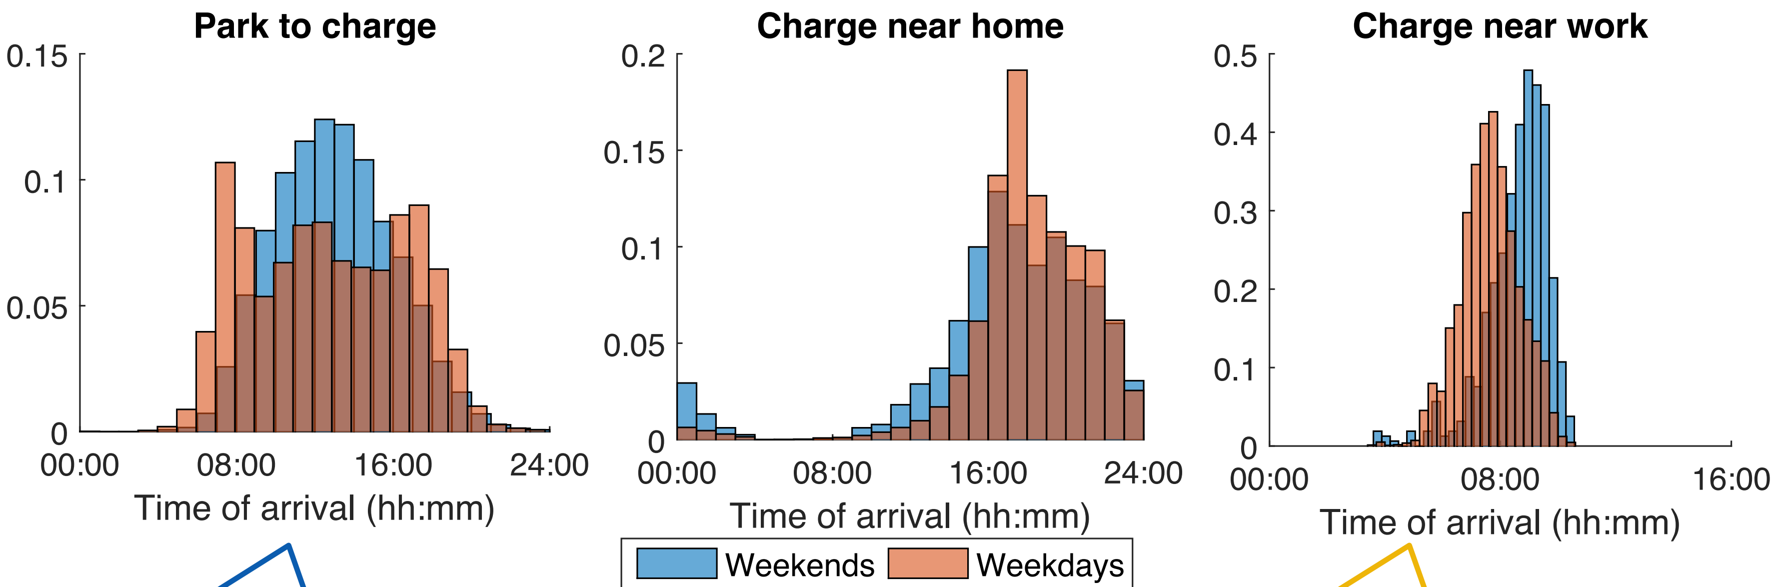

AM/Noon/PM peaks on weekdays

Shift to later times on weekends

SOJOURN AND IDLE TIMES (2/2)

Anytime charging, short sojourn (61.5%)

Nighttime charging (29.1%)

Daytime charging (9.4%)

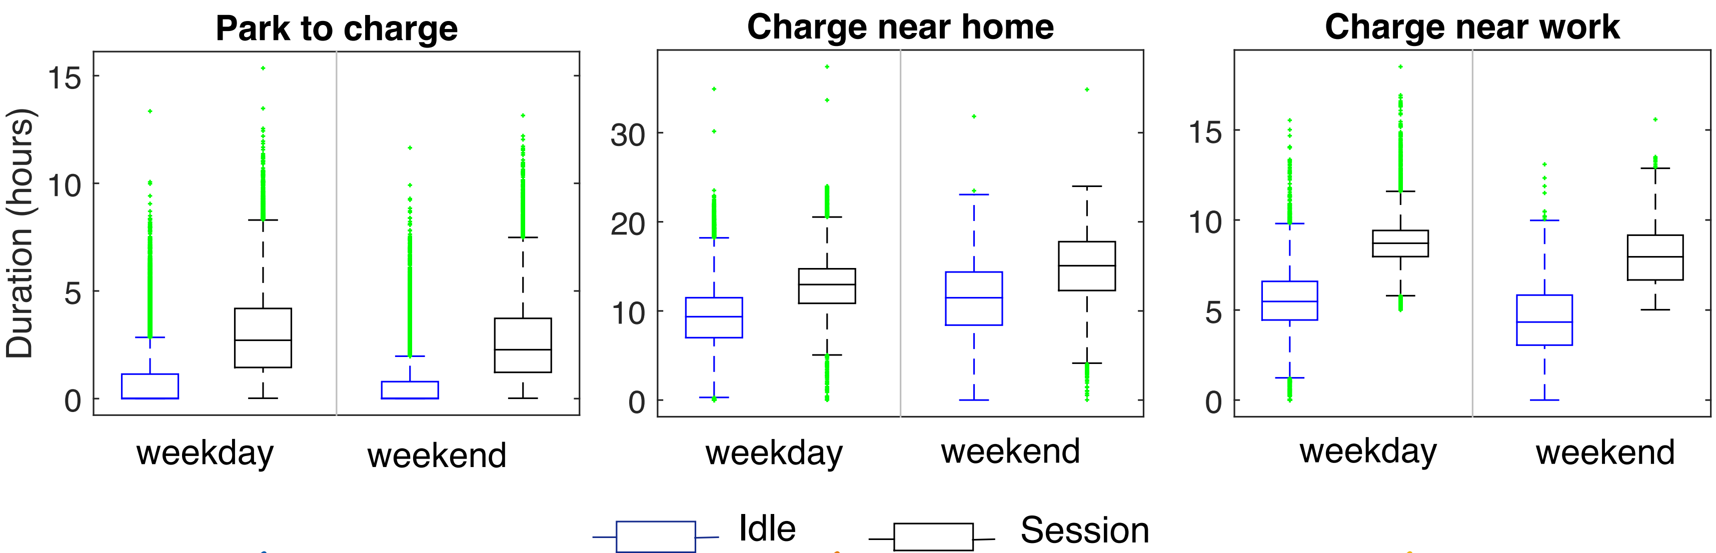

Idle: 48min
Sojourn: 2h 28min

Idle: 10h
Sojourn: 13h 24min

Idle: 5h 30min
Sojourn: 8h 42min

GHENT UNIVERSITY

SOJOURN AND IDLE TIMES (2/2)

Anytime charging, short sojourn (61.5%)

Nighttime charging (29.1%)

Daytime charging (9.4%)

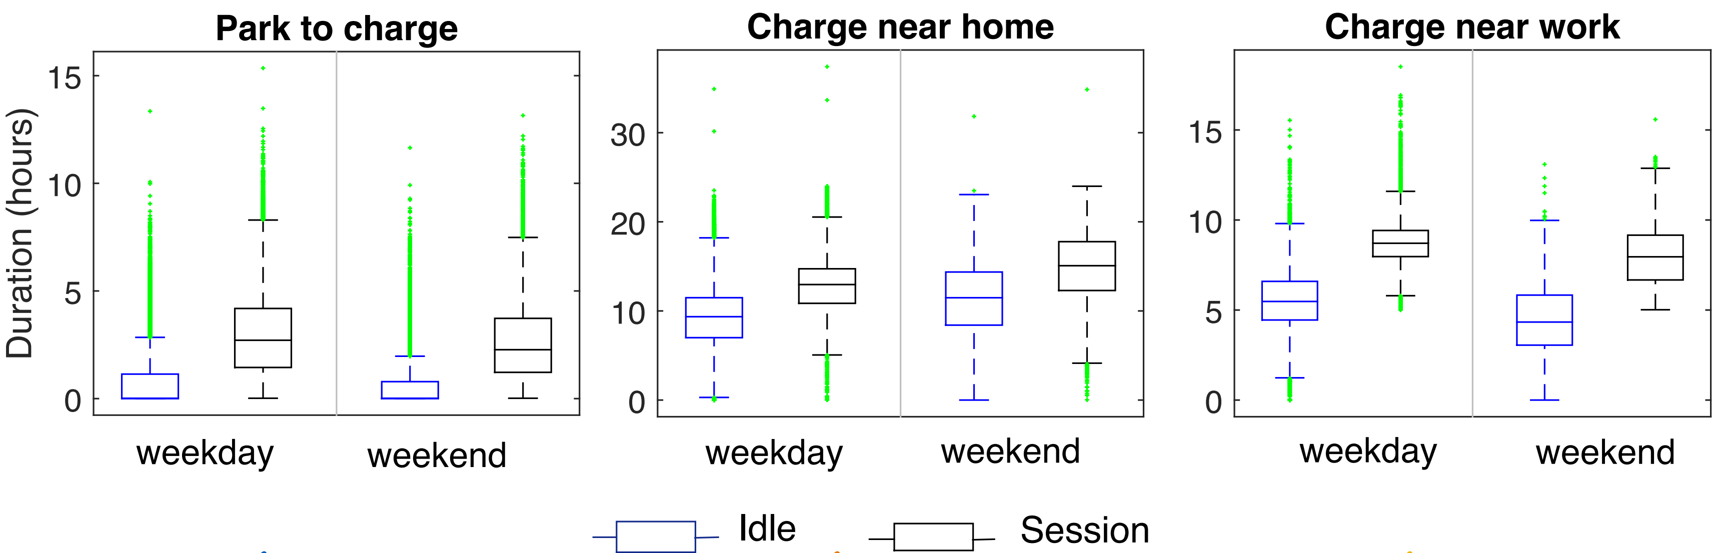

Week ≈ Weekend

Longer in weekend

Shorter in weekend
Lower var. in week

GHENT UNIVERSITY

Quantifying flexibility potential

GHENT
UNIVERSITY

QUANTIFICATION OF FLEXIBILITY: CALCULATION

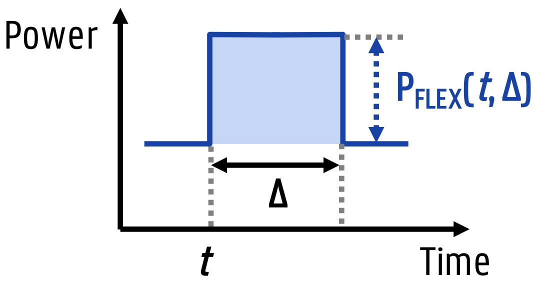

Upper bound: we disregard impact of using/suppressing power in [t, t+Δ] on flexibility at other times t'

$P_{FLEX}(t, \Delta)$ = Maximal power that DR could either consume constantly, or not at all, in interval [t, t+Δ]

- Charging session has to include [t, t+Δ]
  [else demand is not available for DR for the whole interval]
- Charging duration ≥ Δ
  [else we could not consume power for the full interval]
- Flexibility = session duration − Δ ≥ charging duration
  [we can also move the charging power entirely outside the interval]

QUANTIFICATION OF FLEXIBILITY: RESULT

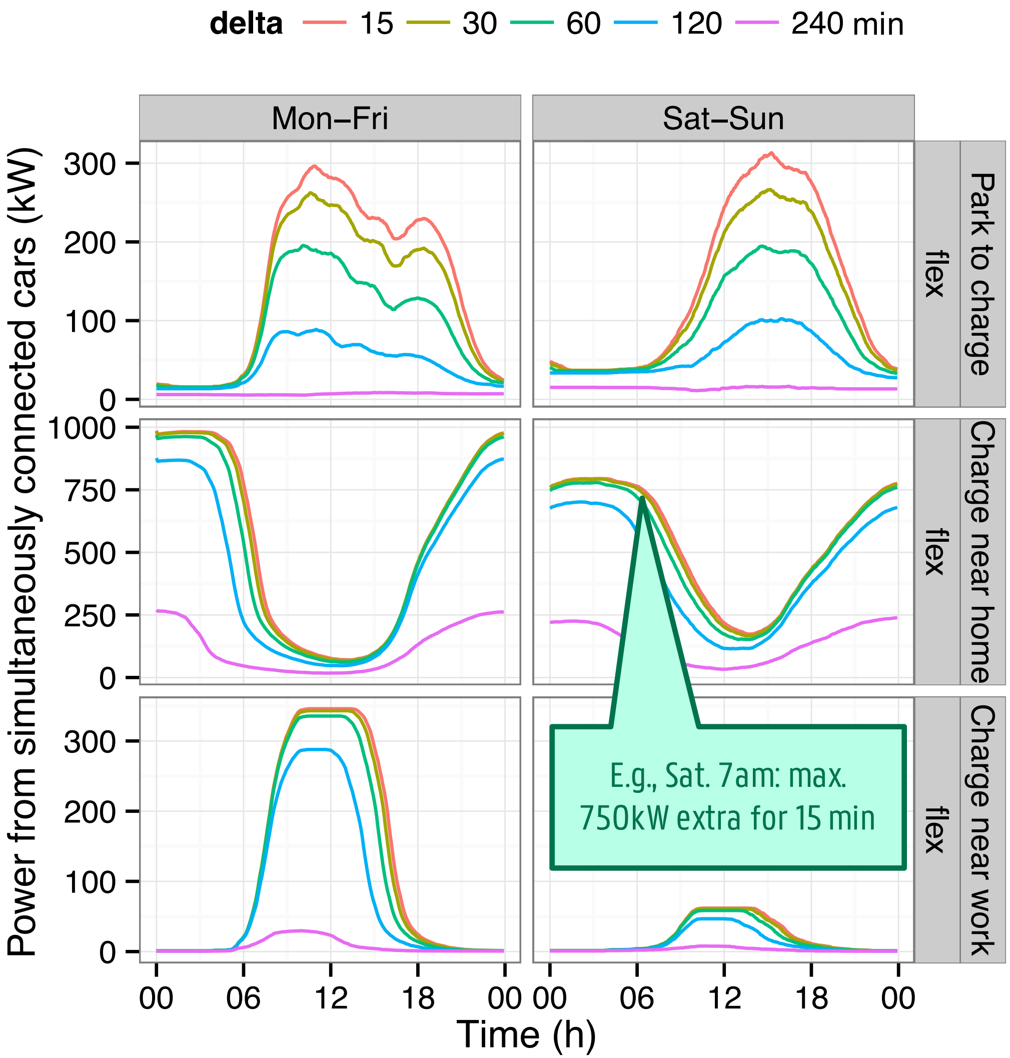

- Park to charge:
  - Daytime flexibility
  - Weekend: ≈ volume, but ≠ timing
- Near home:
  - Nighttime flex
  - Weekend: lower & more spread
- Near work:
  - Daytime flex
  - Low in weekend

GHENT UNIVERSITY

Quantifying flexibility <u>exploitation</u>

SAMPLE CASE STUDIES

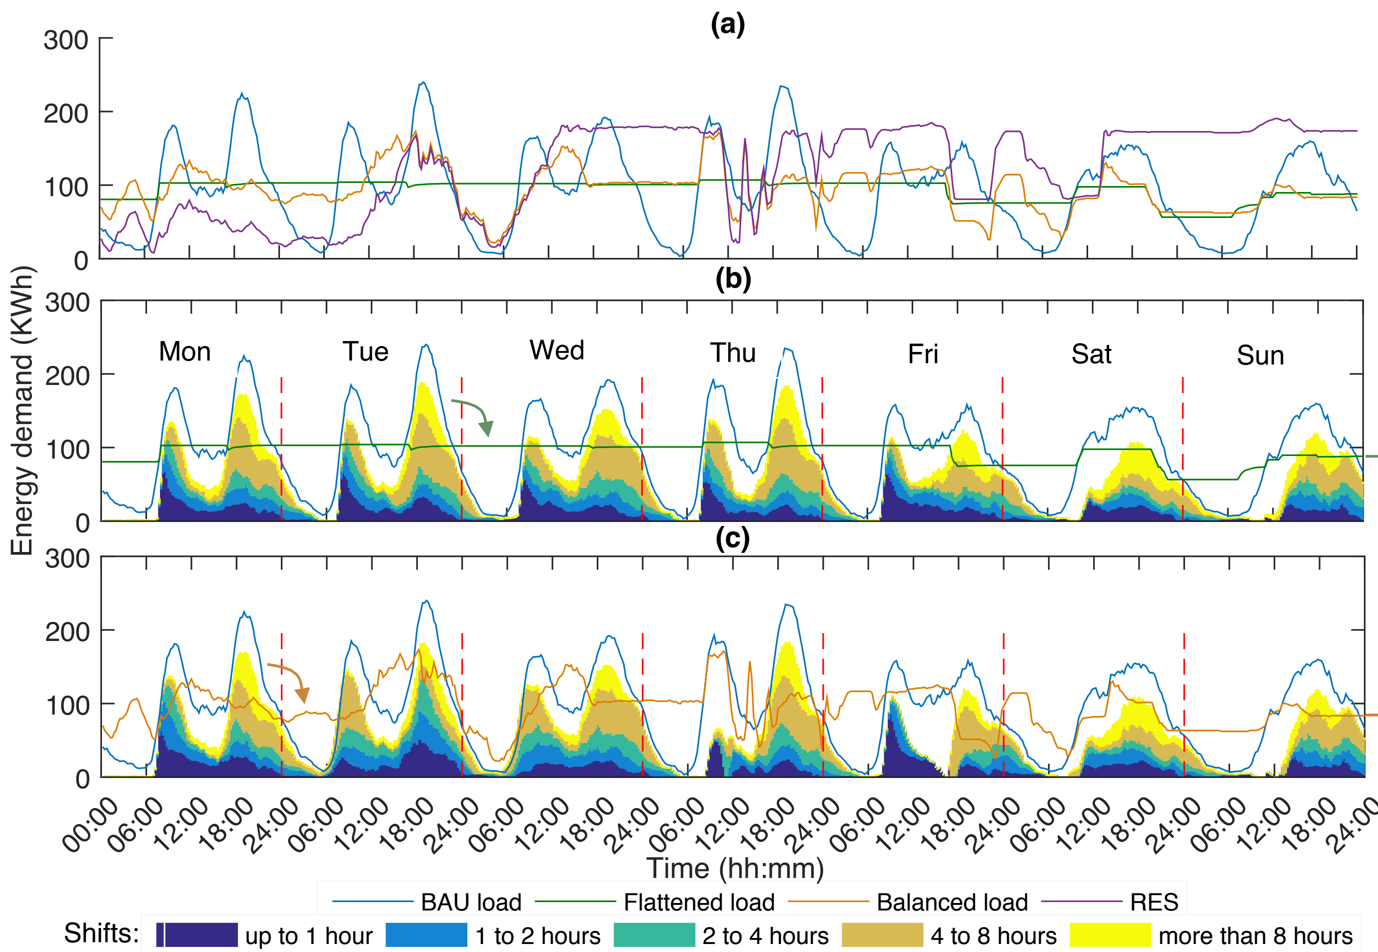

Case (1): Load flattening

Case (2): RES balancing

Fill = over what time span load is deferred

GHENT UNIVERSITY

SO … WHAT FLEXIBILITY IS ACTUALLY USED?

Quantification of use of flexibility in relevant use cases:

$$E_{flex} = \frac{\text{Energy beyond } t_{BAU}}{\text{Maximal energy beyond } t_{BAU}} \Rightarrow 1 - E_{flex} = \text{fraction charged at } t_{BAU}$$

$$T_{flex} = \frac{t_{coordinated} - t_{BAU}}{t_{depart} - t_{BAU}} = \text{fraction of idle time exploited to delay}$$

E.g., $E_{flex} = 0.2 \Rightarrow$ only 20% of charge volume is delayed

E.g., $T_{flex} = 0.8 \Rightarrow$ end-of-charge at 80% of flexibility time window

CASE STUDIES: (1) Load flattening, (2) RES balancing

SAMPLE CASE STUDIES

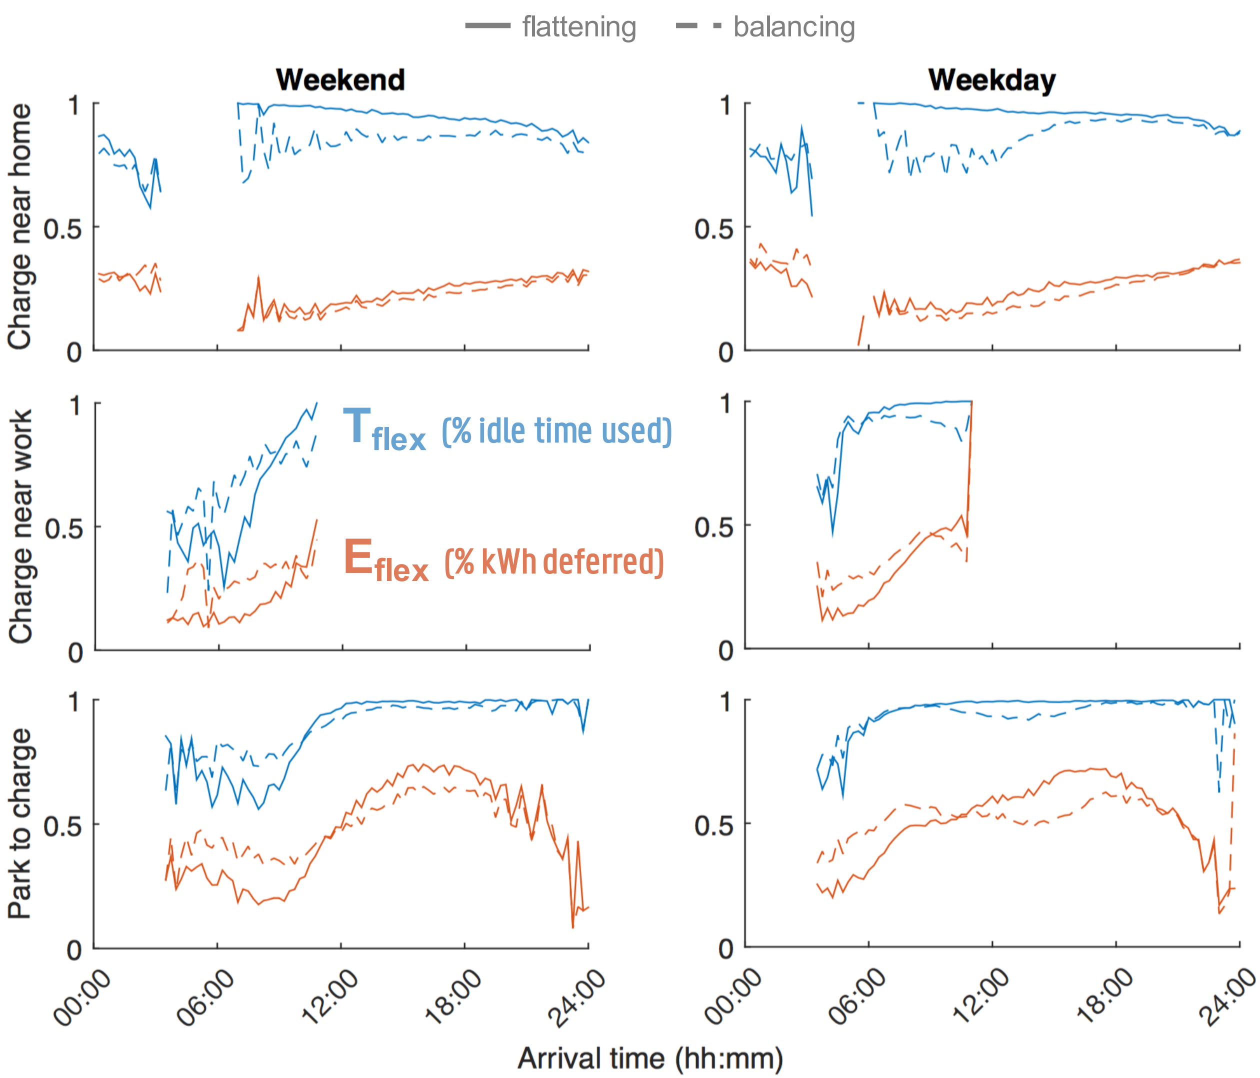

- Near home:
  - T_{flex} close to 1: charging till last moment, but…
  - E_{flex} low: reasonable SoC at t_{BAU}
- Near work:
  - Higher T_{flex} in weekend
  - Reasonable SoC at t_{BAU}
- Park-to-charge:
  - T_{flex} close to 1
  - Peaked E_{flex} during daytime

SUMMARY OF PART I

- Real world data set
- Three major types of charging sessions
- Statistical models of user behavior
- Methodology to quantify flexibility potential
- Metrics to evaluate exploitation of flexibility

Application?

E.g., extrapolation of iMove data to 3% of Flemish fleet by 2020:

- ~100k cars out of ~3.2M
- E.g., noon in weekend $\Rightarrow$ can have ~7MW extra for 2h

PART II:
“Controlling” EV charging sessions

N. Sadeghianpourhamami, J. Deleu and C. Develder, **“Definition and evaluation of model-free coordination of electrical vehicle charging with reinforcement learning”**, IEEE Trans. Smart Grid, 2019, pp. 1-12. (In Press)

M. Lahariya, N. Sadeghianpourhamami and C. Develder, **“Reduced state space and cost function in reinforcement learning for demand response control of multiple EV charging stations”**, in Proc. 6th ACM Int. Conf. Sys. for Energy-Effic. Build., Cities and Transp. (BuildSys 2019), New York, NY, USA, 13-14 Nov. 2019.

Also: thanks Manu!

CONTROLLING EV CHARGING?

Possible objectives

Possible objectives:

- **Flatten load**: peak shaving & valley filling
- **Balance** renewable sources
- Avoid **voltage violations**

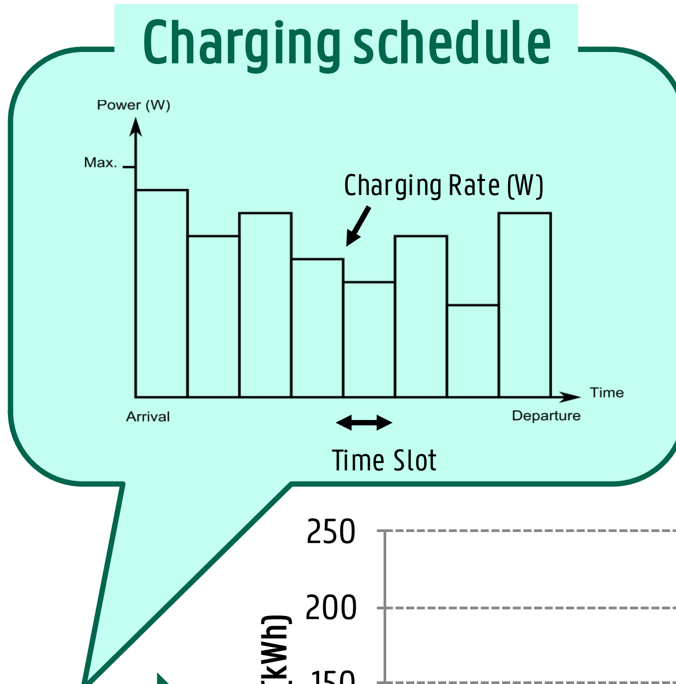

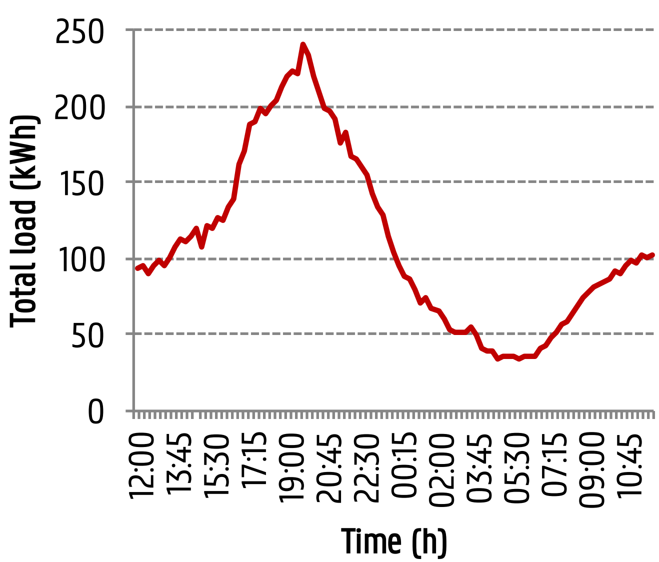

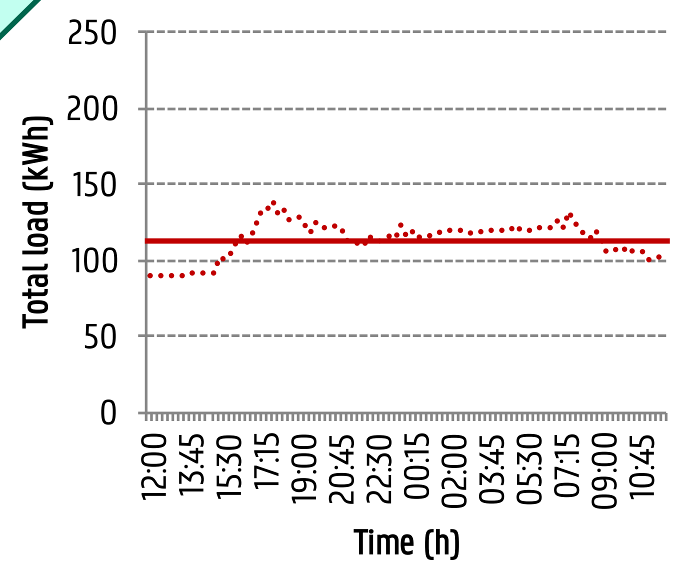

CONTROLLING EV CHARGING?

Problem statement

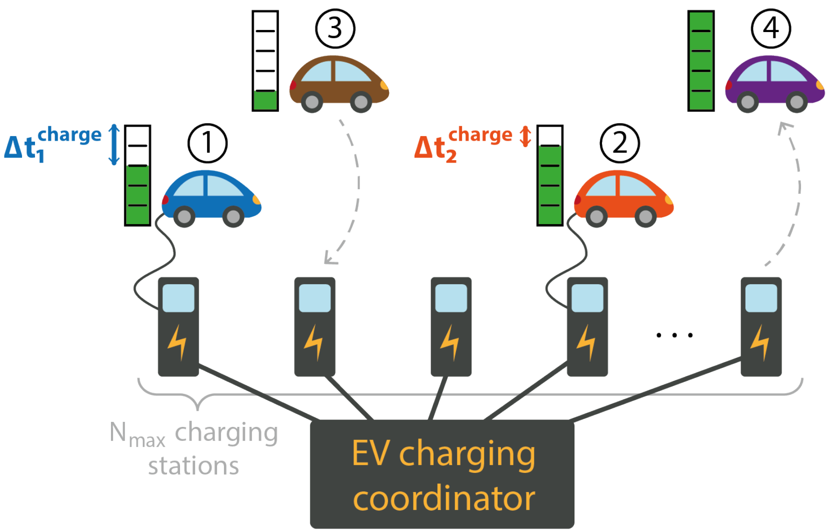

input: arrival & departure times, charging needs
output: which EVs to charge now?

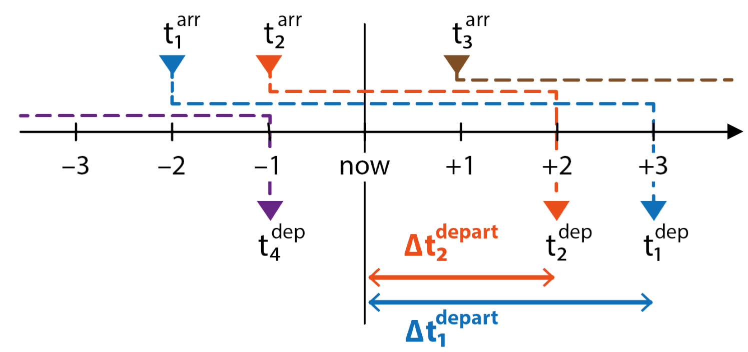

objective: meet target load schedule

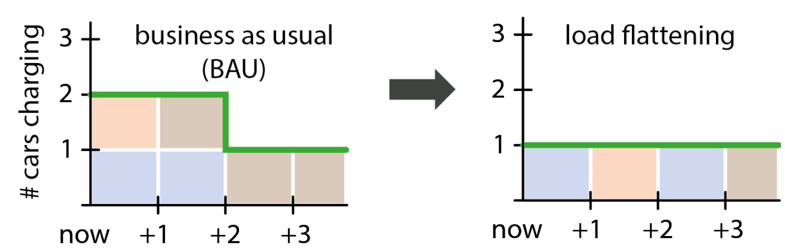

Reinforcement learning model

REINFORCEMENT LEARNING MODEL

State space

- Characterization of EVs in the system:
  - Time of arrival
  - Time till departure
  - Required charging & charging rate

$$\mathcal{V}_t = \{(\Delta t_1^{\text{depart}}, \Delta t_1^{\text{charge}}), \ldots, (\Delta t_{N_s}^{\text{depart}}, \Delta t_{N_s}^{\text{charge}})\}$$

- Aggregated state representation:
  Bin cars with similar (1) time till departure, (2) requested charging

car $c_1 : (\Delta t_1^{\text{depart}}, \Delta t_1^{\text{charge}}) = (3, 2)$

car $c_2 : (\Delta t_2^{\text{depart}}, \Delta t_2^{\text{charge}}) = (2, 1)$

Time left

Requested charging

	1	2	3
1	0	c_2	0
2	0	0	c_1
3	0	0	0

$$(t, \mathbf{x}_s) = \left(1, \begin{bmatrix} 0 & 0.5 & 0 \\ 0 & 0 & 0.5 \\ 0 & 0 & 0 \end{bmatrix}\right)$$

REINFORCEMENT LEARNING MODEL

Making charging decisions

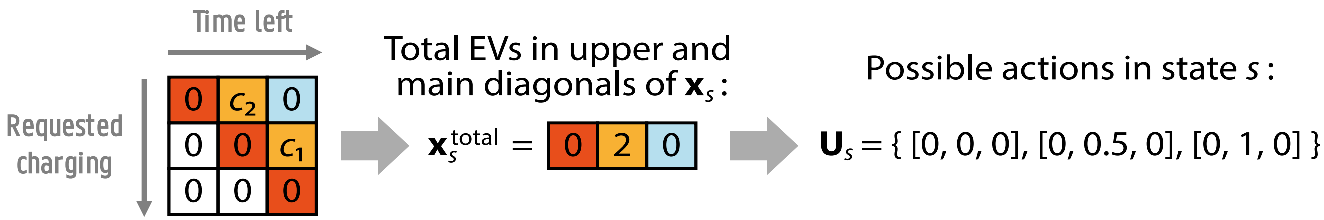

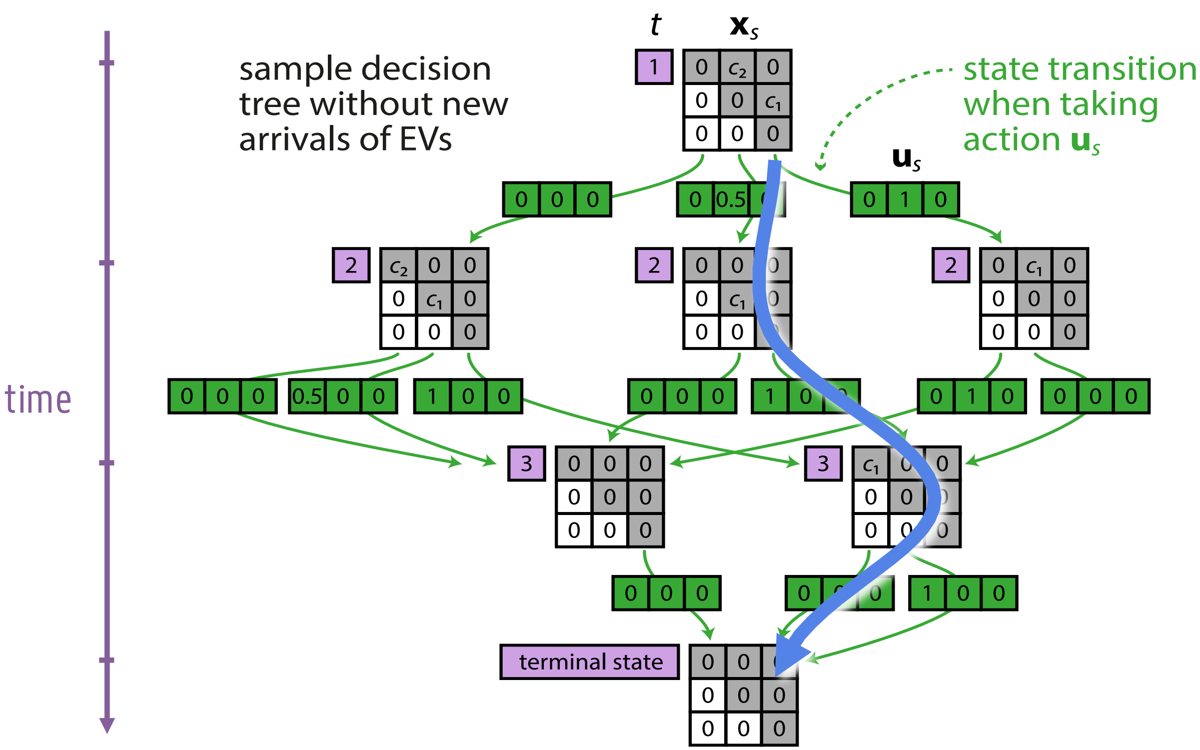

Train the system using sample trajectories from possible decision trees

Simulation results

EXPERIMENTAL EVALUATION
Experiment settings and dataset

- Arrival data from ElaadNL data, top-10 or top-50 stations from 2015
- Control/decision granularity of 2h (to compute fast)
- Set maximum connection time to 24h
- Ensure empty system in between consecutive days

- "Cost" = squared power per timeslot, summed over all timeslots
- We plot costs relative to optimal, all-knowing oracle decision

EXPERIMENTAL EVALUATION

Q1: Appropriate training data settings?

- Training period of 3 months suffices
- Sampling 5k trajectories per day suffices

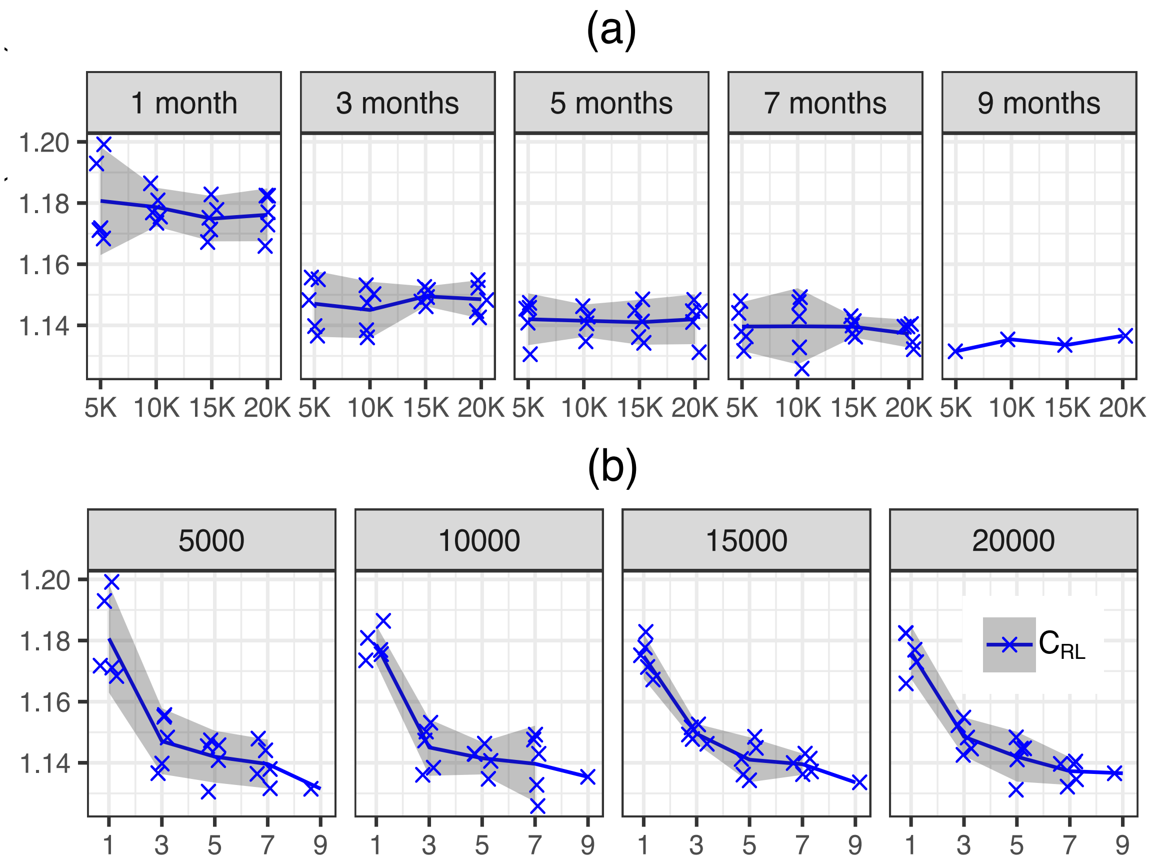

GHENT UNIVERSITY

EXPERIMENTAL EVALUATION
Q2: RL performance compared to oracle benchmark?

- Relative cost saving compared to **business-as-usual** scenario (= immediate charging) of the order of 39% or 30% for 10 or 50 EV stations
- Relative difference compared to all-knowing **optimal charging** strategy (with perfect knowledge of future arrivals) of the order of 13% to 15.6% for 10 or 50 EV stations

GHENT UNIVERSITY

EXPERIMENTAL EVALUATION

Q3: Variance of RL performance over time (a year)?

- Flexibility varies: higher cost difference between BAU and optimal
- RL policy exploits flexibility to varying degree

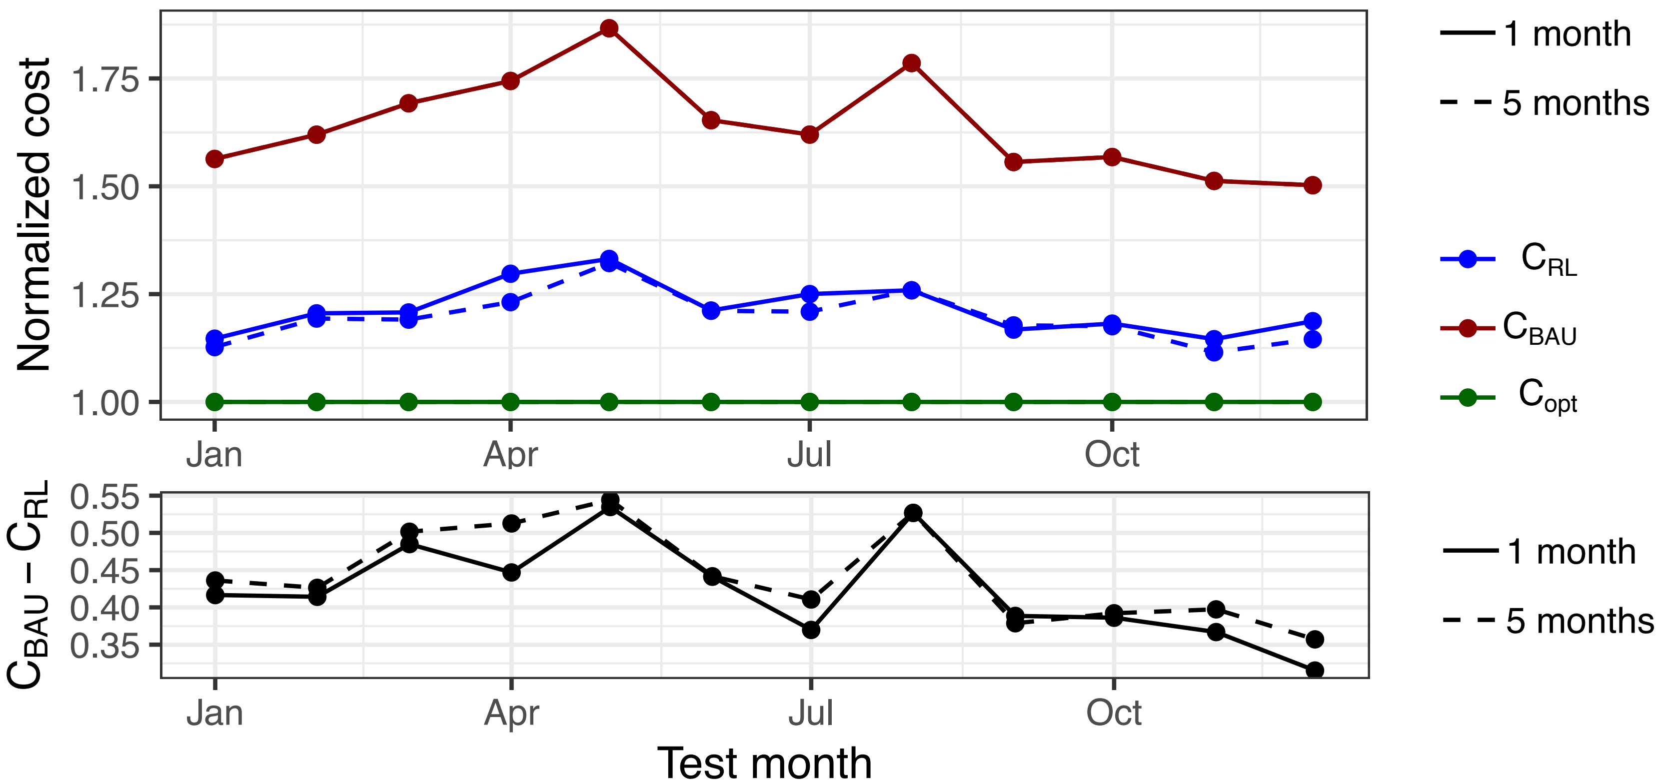

EXPERIMENTAL EVALUATION

Q4: Generalization to different number of stations?

- **Experiment**: scale up car arrivals (and # EV stations) by factor up to 10x
  … **but** follow policy from original system size
- **Result**: RL policy still achieves largely same cost reduction

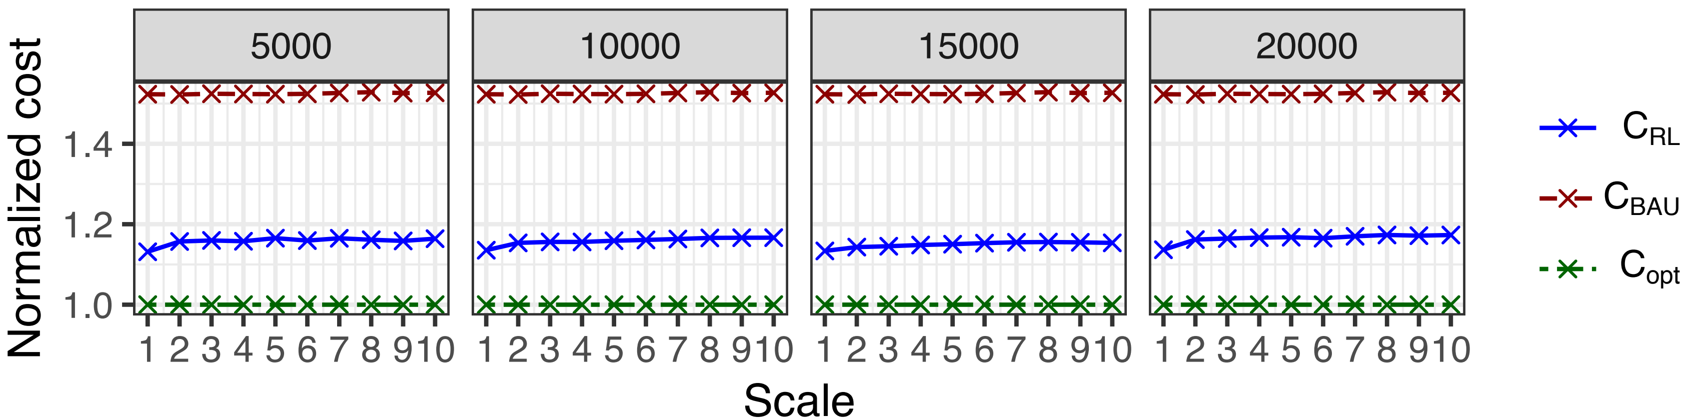

SUMMARY OF PART II

- **Challenge**: coordinate EV charging (for peak shaving & valley filling, for balancing renewable energy sources) of an EV fleet
- **Reality check**: real-world charging behavior shows flexibility potential
- **Reinforcement learning** looks like a viable approach to solve the challenge with consistent performance over time & when upscaling system

  ⟷ **MPC**: RL is purely data-driven

  ⟷ **SotA**: Our RL solution = single-step decision for aggregate of EVs
- What's next?
  - Effective exploration of vast state-action space
  - More advanced neural net architectures for Q-function approximation

GHENT UNIVERSITY

Thank you.
Any questions?

chris.develder@ugent.be

http://users.ugent.be/~cdvelder