Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

NEC Laboratories Europe

Delft University of Technology

Bayesian Online Learning for Energy-Aware Resource Orchestration in Virtualized RANs

Jose A. Ayala-Romero[1], Andres Garcia-Saavedra[2], Xavier Costa-Perez[2 3], George Iosifidis[4]

[1]Trinity College Dublin, [2]NEC Laboratories Europe, [3]i2CAT Foundation and ICREA, [4]Delft University of Technology

Presenter: Jose A. Ayala-Romero

IEEE INFOCOM 2021 - 10-13 May 2021

Introduction

Why vRAN?

- Virtualization is considered today the most promising approach for bringing cellular networks up to speed with the demanding services they aspire to support.
- The softwarization of the base stations (vBS) allows their deployment in diverse platforms, but render less predictable their performance and power consumption
- In order to unleash the full potential of vRANs we need to answer two key questions:
  - *what is the performance and energy consumption profile of vBSs?*
  - *how can we optimize their operation using an adaptive and platform-oblivious approach?*

Experimental evaluation

We evaluate the power consumption of a vBS in our testbed

- The testbed comprises a virtual Base Station (**vBS**), the user equipement (**UE**), and a digital **power meter**.
- We use **2** Ettus Research **USRP** B210 (radio part) and **2** Intel **NUCs** with CPU i7-8559U (BBU).
- We use **srsLTE suite** to implement the BBU for both the eNB and UE
- We select the **10 MHz** bandwidth.
- Digital power meter **GW-Instek GPM-8213** along with the adapter GPM-001.
- We have integrated **O-RAN E2 interface** and the ability to change vBS configurations *on-the-fly*.
- We generate the traffic load for both DL and UL using *mgen*.

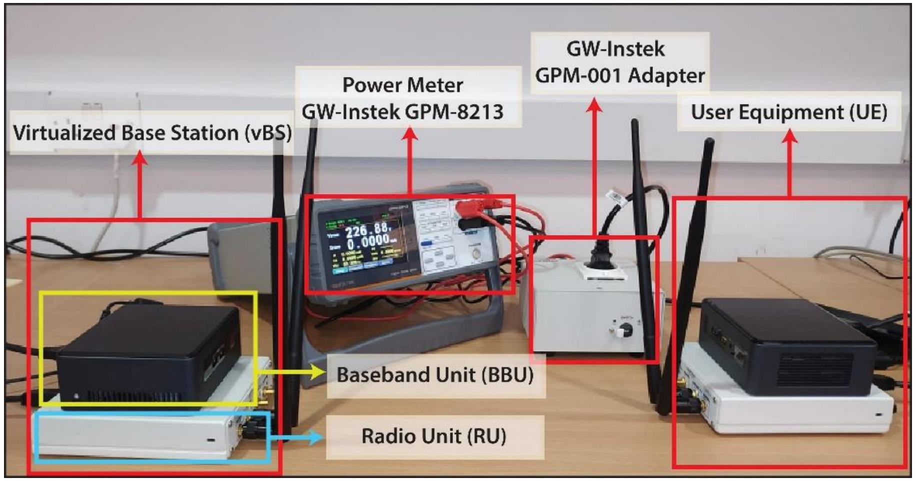

Experimental evaluation

BBU/CPU cost & impact of computing platform.

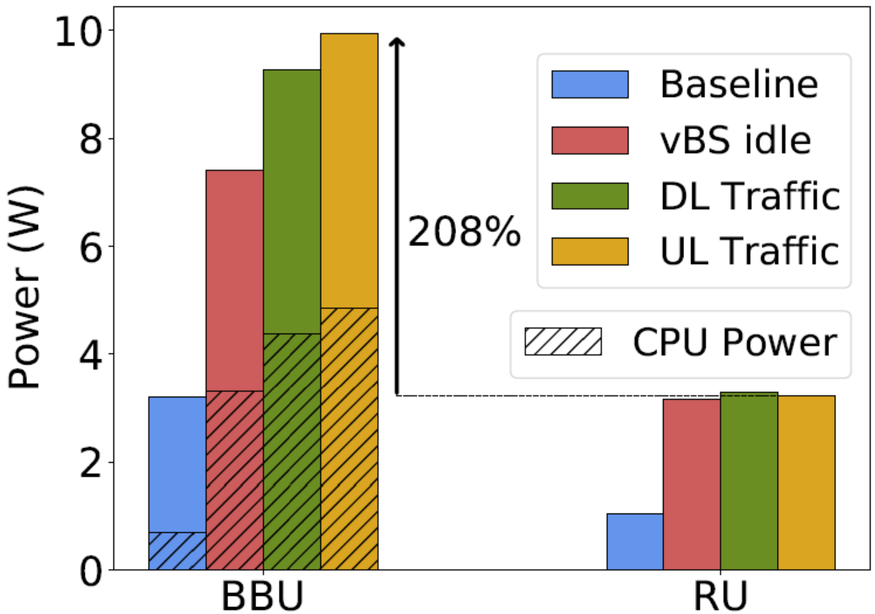

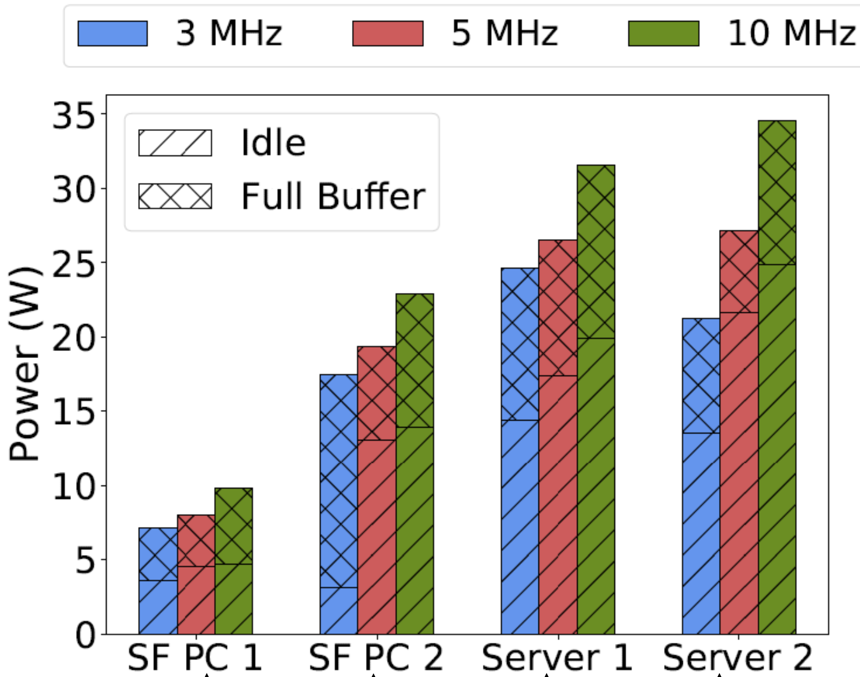

Computational contribution

Radio contribution

Different computing platforms

Intel NUCs for the SF PCs, and two high-performance computation servers.

Experimental evaluation

Impact of SNR and MCS

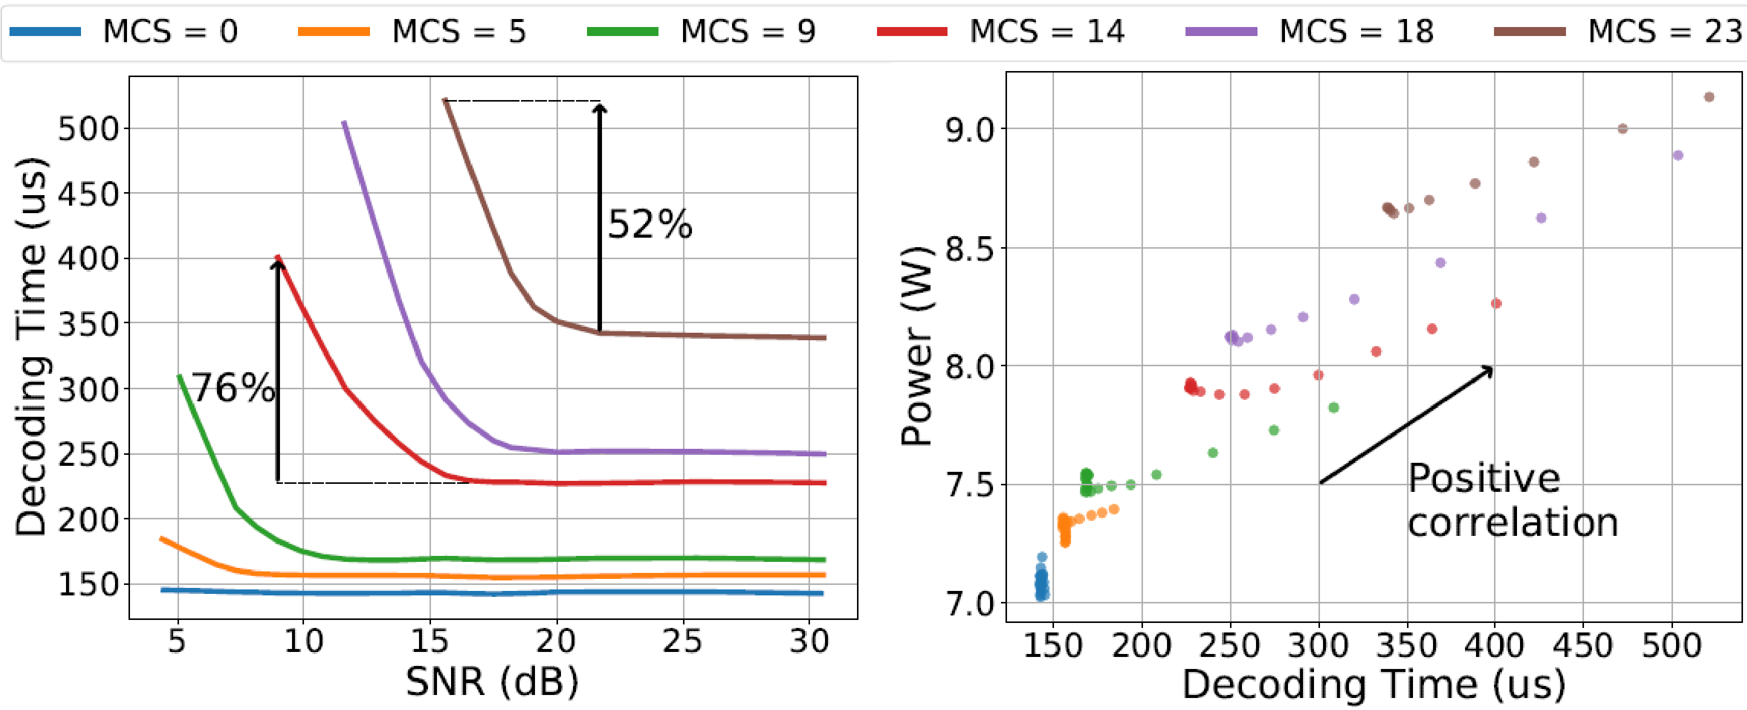

Decrease of channel quality

The higher the decoding time the higher the consumed power

Experimental evaluation

Configuration options and impact of scheduler

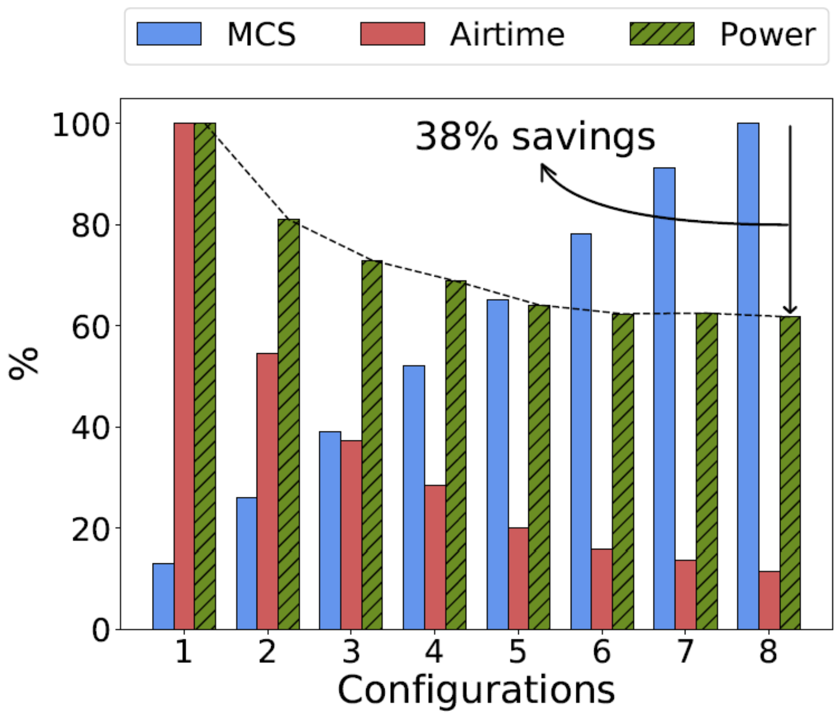

Eight different configurations with the same Throughput in the UL

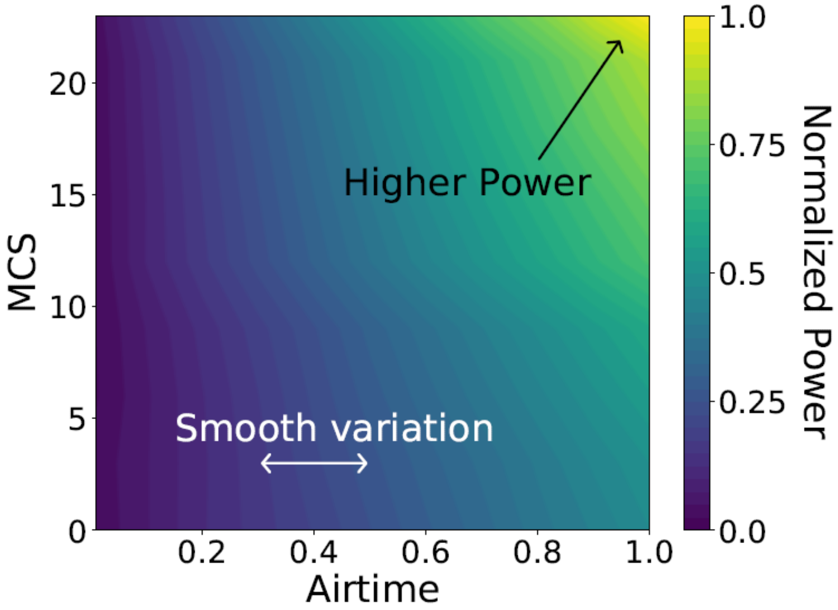

Joint effect of MCS and airtime on the consumed power

Experimental evaluation

Coupling UL and DL

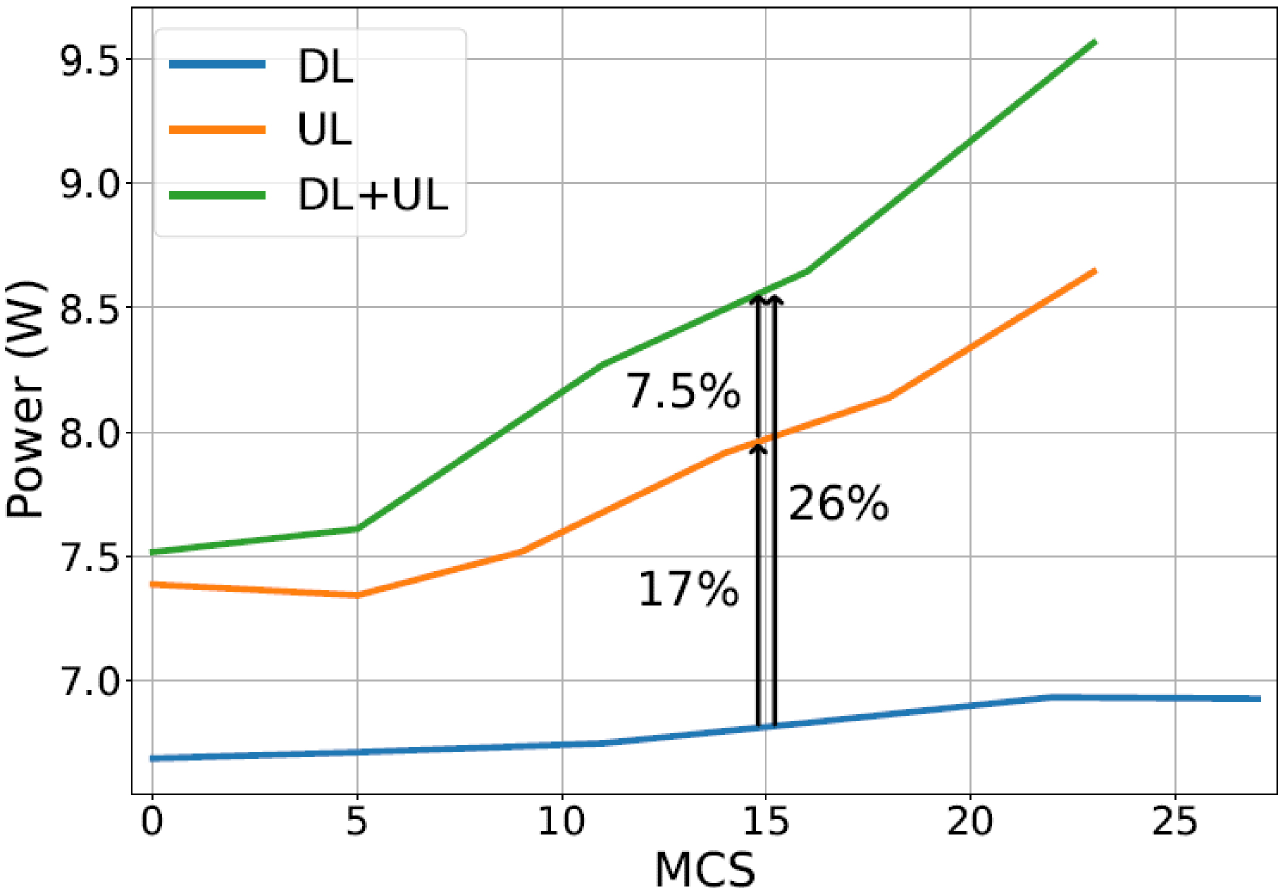

Experimental evaluation

Conclusions

- Characterizing the vBS power consumption is **intricate** as it depends on traffic, SNR, MCS and airtime.
- There are many DL and UL configurations and some of them present **non-linear and non-monotonic relations** with power and throughput.
- The **power** consumption **depends** on the BBU **platform** and radio **scheduler**.
- This hinders the derivation of general consumption models.
- We propose the use of ***online learning*** to profile each vBS power cost and performance, and devise goal-driven configuration policies.

Problem formulation

Contexts, actions and rewards

Context downlink:

New bit arrivals

$$\omega_t^{dl} := [\bar{c}_t^{dl}, \tilde{c}_t^{dl}, d_t^{dl}]$$

CQI (mean and var)

Context uplink:

$$\omega_t^{ul} := [\bar{c}_t^{ul}, \tilde{c}_t^{ul}, d_t^{ul}]$$

Context:

$$\omega_t := [\omega_t^{dl}, \omega_t^{ul}] \in \Omega$$

Action downlink:

MCS

$$x_t^{dl} := [p_t^{dl}, m_t^{dl}, a_t^{dl}]$$

Transmission power

Action uplink:

$$x_t^{ul} := [m_t^{ul}, a_t^{ul}]$$

airtime

Action:

$$x_t := [x_t^{dl}, x_t^{ul}] \in \mathcal{X}$$

Reward:

Throughput

$$r(\omega_t, x_t) := \log\left(1 + \frac{R^{dl}(\omega_t^{dl}, x_t^{dl})}{d_t^{dl}}\right) + \log\left(1 + \frac{R^{ul}(\omega_t^{ul}, x_t^{ul})}{d_t^{ul}}\right)$$

Problem formulation

Case 1: Balancing performance and cost

- The power supply is scarce, or the operator simply wants to reduce the power costs.
- Trade-off *throughput vs. power consumption*
- Objective function:

Weighting factor

Consumed power

$$u(\omega_t, x_t) := r(\omega_t, x_t) - \delta B\big(P(\omega_t, x_t)\big)$$

monetary cost associated with power consumption

- Contextual regret:

$$R_T := \sum_{t=1}^{T} \left(\max_{x' \in \mathcal{X}} u(\omega_t, x') - u(\omega_t, x_t) \right)$$

- Objective: *Find a seq. of decisions* $\langle x_t \rangle_{t=1}^{T}$ *that* $\lim_{T \to \infty} R_T / T = 0$

Problem formulation

Case 2: Hard power budget

- The vBS operates under a hard power budget P_{max}, e.g., when powered over Ethernet (PoE).
- Find for maximum throughput configuration meeting the budget
- Contextual regret:

$$R_T^s := \sum_{t=1}^{T} \left(\max_{x' \in S_t(\omega_t)} r(\omega_t, x') - r(\omega_t, x_t) \right)$$

- Available actions:

$$S_t(\omega_t) = \{x \in \mathcal{X} \mid P(\omega_t, x) \leq P_{\max}\}$$

- Objective: *Find a seq. of decisions* $\langle x_t \rangle_{t=1}^{T}$ *that* $\lim_{T\to\infty} R_T^s / T = 0.$

Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Problem formulation

How do we solve the contextual bandit problem?

Challenges:

- Most contextual bandit algorithms assume that there is a feature vector associated with each action, and the relation between the feature vector and the objective is linear or known.
- In our case, non linear the *contexts-actions present **non-linear and non-monotonic** relations with power and throughput*.
- Moreover, the objective function values of two different nearby actions are **correlated**.

Proposal:

- Bayesian online learning.
- We propose two algorithms: **BP-vRAN** and **SBP-vRAN**

Bayesian Online Learning Solutions

BP-vRAN: Balancing performance and cost

Function estimator: Gaussian Process (GP)

Context-action pair: $z \in \mathcal{Z} = \Omega \times \mathcal{X}$

Noisy observations $y_T = [u_1, \ldots, u_T]$

Corresponding to the points $Z_T = [z_1, \ldots, z_T]$

The posterior distribution of the objective function follows a GP distribution with mean and covariance:

$$\mu_T(z) = k_T(z^\top)(K_T + \zeta^2 I_T)^{-1} y_T$$

$$k_T(z, z') = k(z, z') - k_T(z^\top)(K_T + \zeta^2 I_T)^{-1} k_T(z')$$

where $k_T(z) = [k(z_1, z), \ldots, k(z_T, z)]^\top$ $K_T(z) = [k(z, z')]_{z, z' \in Z_T}$

These equations allow us **to estimate the distribution of unobserved values** of $z \in \mathcal{Z}$ based on the prior distribution, the vector Z_T, and the function observations y_T

Bayesian Online Learning Solutions

BP-vRAN: Balancing performance and cost

Kernel function.

We need the kernel to be *stationary* and *anisotropic*.

We select the anisotropic version of Matérn kernel with $\nu = \frac{3}{2}$

$$k(z, z') = (1 + \sqrt{3}d(z, z')) \exp(-\sqrt{3}d(z, z'))$$

Distance between to point according the length scale vector

Acquisition function.

We use the Upper Confidence Bound (UCB) method:

$$x_t = \underset{x \in \mathcal{X}}{\operatorname{argmax}} \; \mu_{t-1}(\omega_t, x) + \sqrt{\beta_t}\sigma_{t-1}(\omega_t, x)$$

where $\sigma_t(z) = k_t(z, z)$

Bayesian Online Learning Solutions

BP-vRAN: Balancing performance and cost

Algorithm 1 BP-vRAN: Performance and cost balancing

Inputs: Control Space $\mathcal{X}$, kernel k, β
Initialize: $y_0 = \emptyset$, $Z_0 = \emptyset$
for $t = 1, \ldots, T$ **do**
  Observe the context ω_t
  $x_t = \text{argmax}_{x \in \mathcal{X}} \ \mu_{t-1}(\omega_t, x) + \sqrt{\beta_t}\sigma_{t-1}(\omega_t, x)$
  Measure $R_t^{dl}(\omega_t^{dl}, x_t^{dl})$, $R_t^{ul}(\omega_t^{ul}, x_t^{ul})$ and $P_t(\omega_t, x_t)$ at the end of the decision period t
  Compute $u_t(\omega_t, x_t)$ using (1), (2) and (3)
  Update $Z_t \leftarrow Z_{t-1} \cup [\omega_t, x_t]$
  Update $y_t \leftarrow y_{t-1} \cup u_t(\omega_t, x_t)$
  Perform Bayesian update to obtain μ_t and σ_t
end for

The contextual regret is upper bounded with high probability:

$$P\left(R_T \leq \sqrt{C_1 T \beta_T \gamma_T} + 2 \ \forall T \geq 1\right) \geq 1 - \zeta$$

Bayesian Online Learning Solutions

SBP-vRAN: Safe Bayesian Optimization

- Similar formulation to BP-vRAN.
- We use **two GPs** to approximate: 1) reward function; 2) consumed power
- We use the second one to create a **set of safe actions**:

$$S_t = \{x \in \mathcal{X} \mid \mu_{t-1}^c(\omega_t, x) + \beta_t \sigma_{t-1}^c \omega_t, x) \leq P_{\max}\}$$

- And then, we use UCB over this set:

$$x_t = \underset{x \in S_t}{\operatorname{argmax}} \; \mu_{t-1}^f(\omega_t, x) + \sqrt{\beta_t} \sigma_{t-1}^f(\omega_t, x)$$

We have observed empirically that given the structure of our problem and thanks to UCB exploration, **SBP-vRAN expands the safe set efficiently**.

Bayesian Online Learning Solutions

SBP-vRAN: Safe Bayesian Optimization

Algorithm 2 SBP-vRAN: Safe online optimization

Inputs: Control Space $\mathcal{X}$, Initial safe set S_0, kernel k, β, $P_{\max}$
Initialize: $y_0^f = \emptyset$, $y_0^c = \emptyset$, $Z_0 = \emptyset$
for $t = 1, \ldots, T$ **do**
  Observe the context ω_t
  $S_t = S_0 \cup \{x \in \mathcal{X} \mid \mu_{t-1}^c(\omega_t, x) + \beta_t \sigma_{t-1}^c(\omega_t, x) \leq P_{\max}\}$
  $x_t = \operatorname{argmax}_{x \in S_t} \ \mu_{t-1}(\omega_t, x) + \sqrt{\beta_t}\sigma_{t-1}(\omega_t, x)$
  Measure $R_t^{dl}(\omega_t^{dl}, x_t^{dl})$, $R_t^{ul}(\omega_t^{ul}, x_t^{ul})$ and $P_t(\omega_t, x_t)$ at the end of the decision period t
  Compute $r_t(\omega_t, x_t)$ using (1)
  Update $Z_t \leftarrow Z_{t-1} \cup [\omega_t, x_t]$
  Update $y_t^f \leftarrow y_{t-1}^f \cup r_t(\omega_t, x_t)$
  Update $y_t^c \leftarrow y_{t-1}^c \cup P_t(\omega_t, x_t)$
  Perform Bayesian update to obtain μ_t^f, σ_t^f, μ_t^c and σ_t^c
end for

Experimental evaluation

Convergence evaluation

BP-vRAN

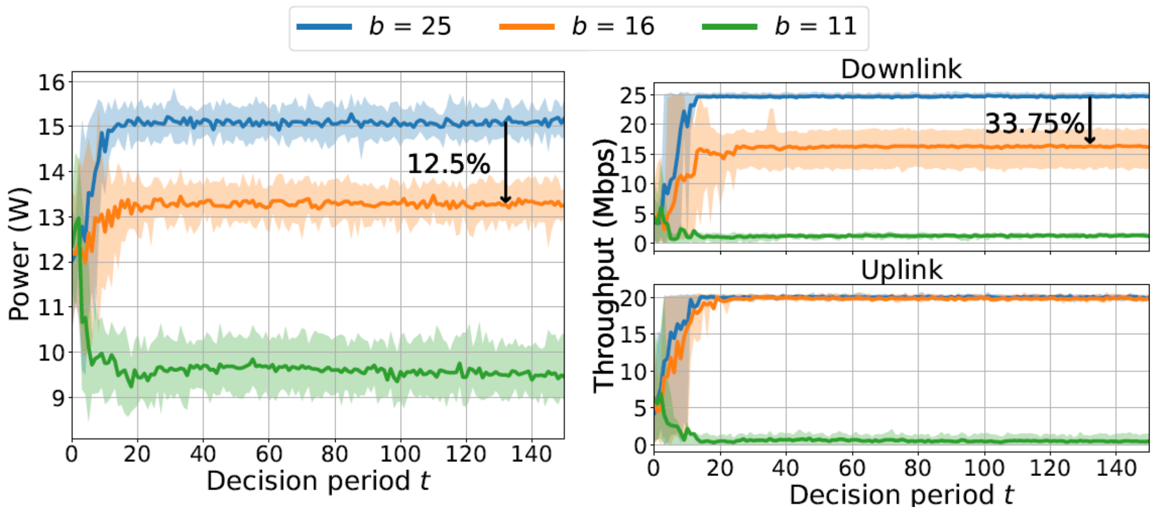

$$|\mathcal{X}| \approx 1.6 \cdot 10^6$$

SBP-vRAN

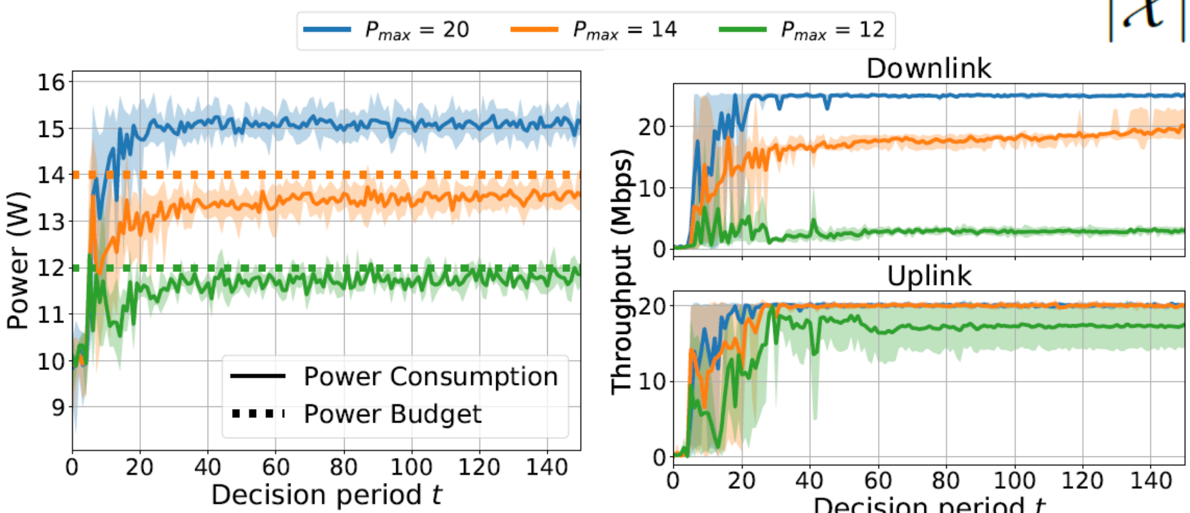

Experimental evaluation

Performance in real network contexts

Realistic context pattern:

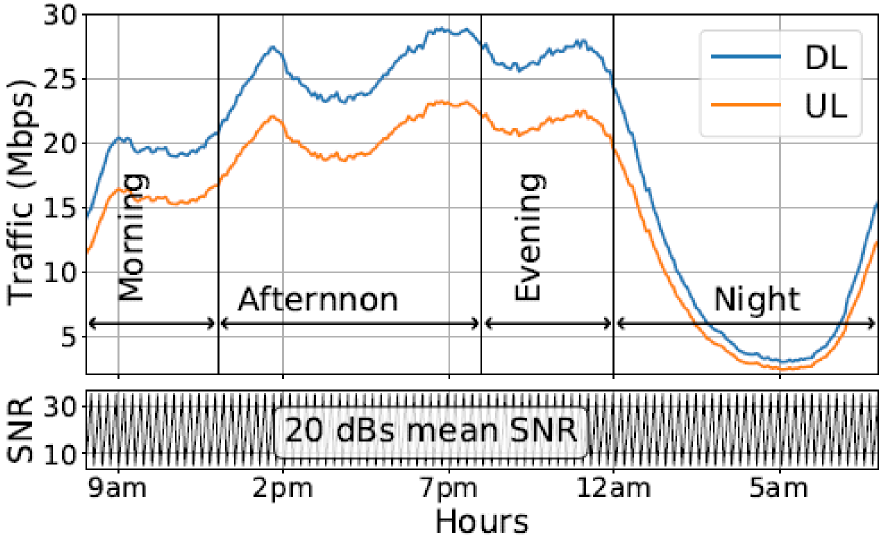

Experimental evaluation

Performance in real network contexts

BP-vRAN

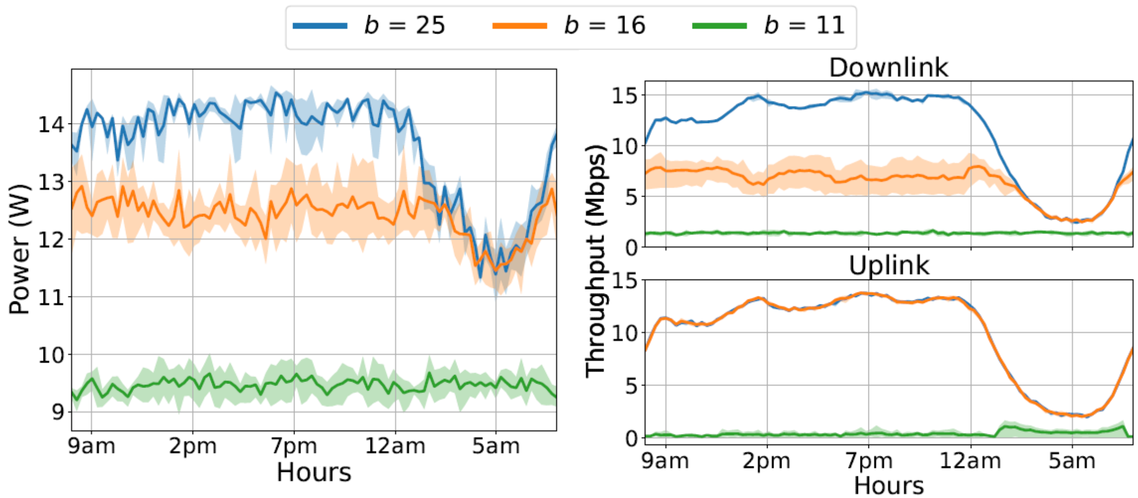

SBP-vRAN

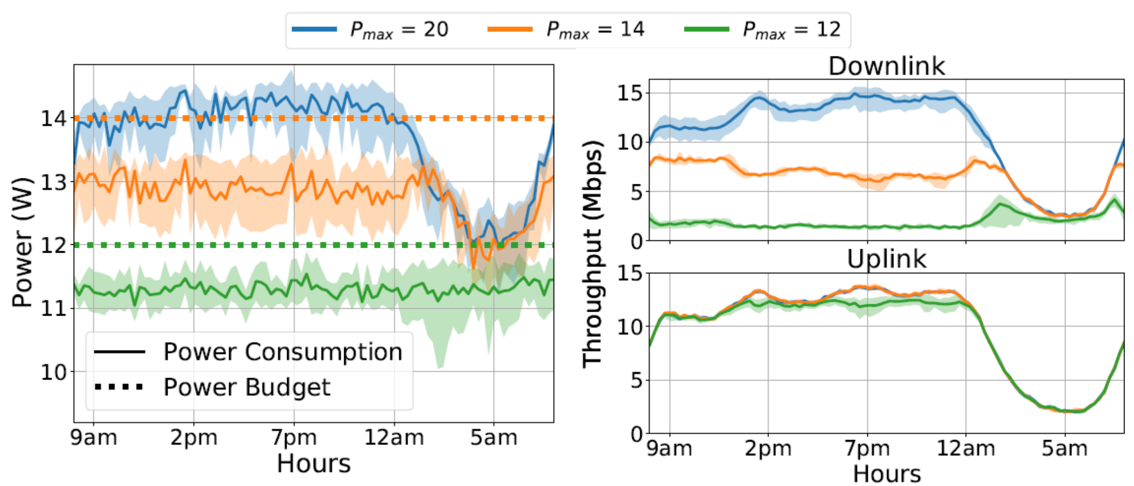

Conclusions

- We have presented an in-depth **experimental study** of the energy behavior of **vBSs**.
- Our results made evident the **complex relationship** between **performance**, **power consumption**, and different **vBS configurations**.
- Such complexity can only be tamed **with data-driven machine-learning solutions**.
- We have proposed an **online learning framework** to achieve **two goals**:
  - **Balance performance and power consumption** in unconstrained platforms such as data centers.
  - Maximize performance subject to **power constraints** vBS, e.g., PoE.
  - We proposed two algorithms based on Bayesian optimization**: BP-vRAN** and **SBP-vRAN**
- They achieve the goals with **theoretical performance** guarantees, with high **data-efficiency** and convergence speed, and **respecting power constraints** even during learning.
- We presented a thorough **experimental evaluation** of our algorithms using real-life traffic load and signal quality patterns. Our results demonstrated the ability of our approach to **converge quickly to optimal policies**.
- We have **released the source code** of BP-vRAN and SBP-vRAN along with the **dataset** used in this work to foster future research in this area.

Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

NEC Laboratories Europe

TU Delft
Delft University of Technology

i2cat

Thank You