01.04.2021

Edge Intelligence

Surprises, Risks, and Lessons

Aaron Ding

CPI Lab, TU Delft

TU Delft

CPI Lab @ TU Delft

- TU Delft
  - Staff: 5899
  - Students: 23461
  - **#1** Engineering in NL

- CPI Lab
  - Director: Aaron Ding
  - Focus: Edge AI, IoT
  - Founder of ACM EdgeSys
  - 20+ international projects

All about Edge Intelligence

What’s the challenge ?

TUDelft

Bridge the Gap

Consolidate Cloud & IoT

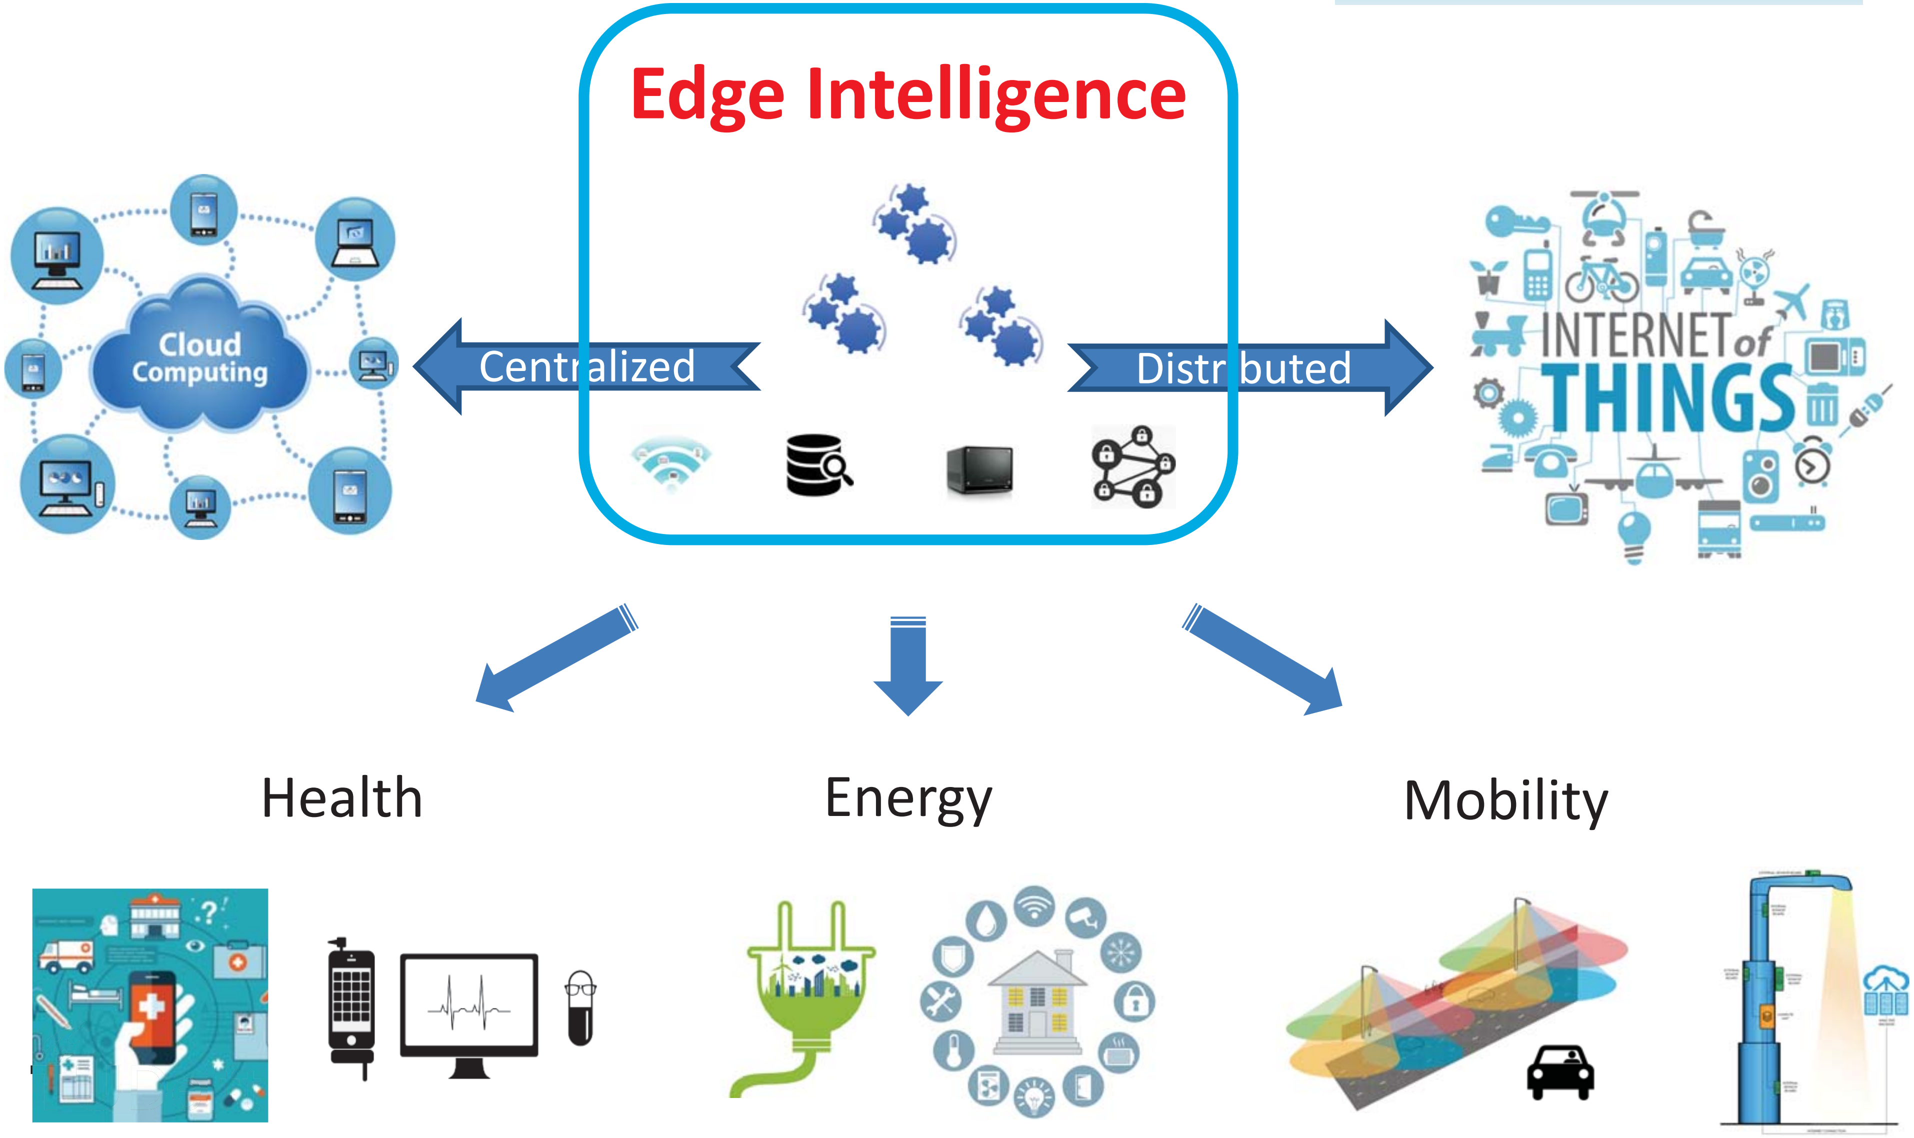

Is this Edge a real thing ?

System Package

Provider	Product
Microsoft	Azure Data Box Edge
Intel	Movidius Neural Compute Stick
NVIDIA	Jetson Nano, TX, Xavier NX
Huawei	Atlas AI Computing Platform

Computing Framework

Provider	Product
Microsoft	Azure IoT Edge
Google	Google Cloud IoT
NVIDIA	NVIDIA EGX
Amazon	AWS IoT Greengrass
Alibaba	Link IoT Edge
Linux Foundat.	EdgeX & Akraino Edge Stack
Huawei	KubeEdge

Edge is **Real**

Edge AI Hardware

Provider	Product
Google	Tensor Processing Units (TPU
Intel	Movidius Vision Processing U
Qualcomm	Qualcomm Snapdragon 8 Ser
Huawei	Ascend Series & Kirin 600/90
Samsung	Exynos 9820 Neural Processir
NVIDIA	TURING GPU

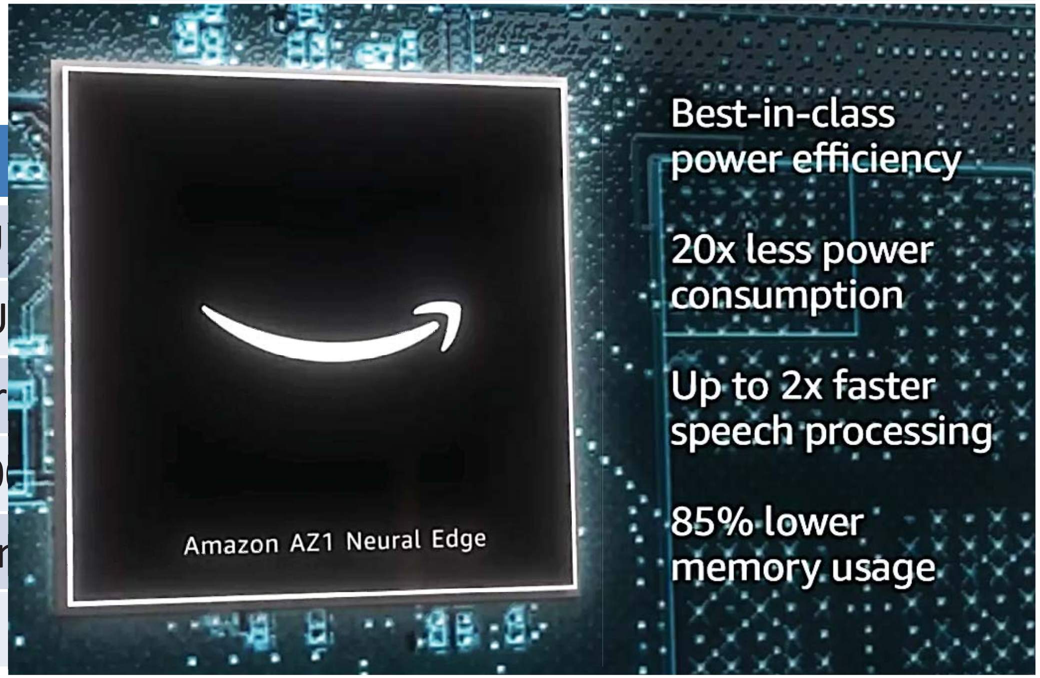

Cyber-Physical Intelligence (CPI) Lab

- PhD Researchers






TU Delft

Edge Research @ CPI Lab

Edge Analytics and Services

Edge Computing	**Edge Communication**	**Edge Security**

System Platform

Cyber-Physical Infrastructure

TUDelft

Light ?

Shed Real Light on Communication

Edge Research @ CPI Lab

Edge Analytics and Services

Edge Computing

Edge Communication

Edge Security

Architecture and Platform

Cyber-Physical Infrastructure

TU Delft

The Dark Side of Wireless

- WiFi to see through walls
  - Track position, actions, and movement of individuals
  - Even behind closed doors

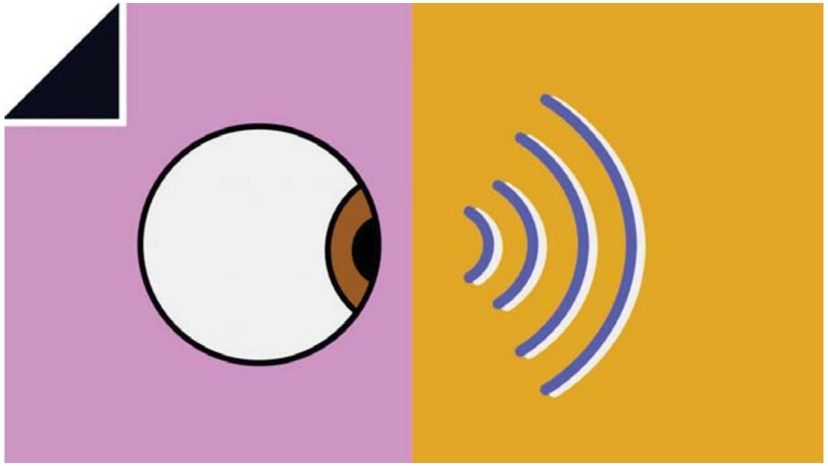

"See through walls with WiFi!", F. Adib, D. Katabi, ***ACM SIGCOMM'13***
"Through-Wall Human Pose Estimation Using Radio Signals", ***IEEE CVPR'18***

Business Impact

Using Wi-Fi to "see" behind closed doors is easier than anyone thought

With nothing but a smartphone and some clever computation, researchers can exploit ambient signals to track individuals in their own homes.

by Emerging Technology from the arXiv   November 2, 2018

Secure & Private by Design or No Design ?

- Wireless is (too) pervasive
  - Hard to control information boundary
  - Bluetooth, WiFi
  - Hard to achieve localized sharing at network edge

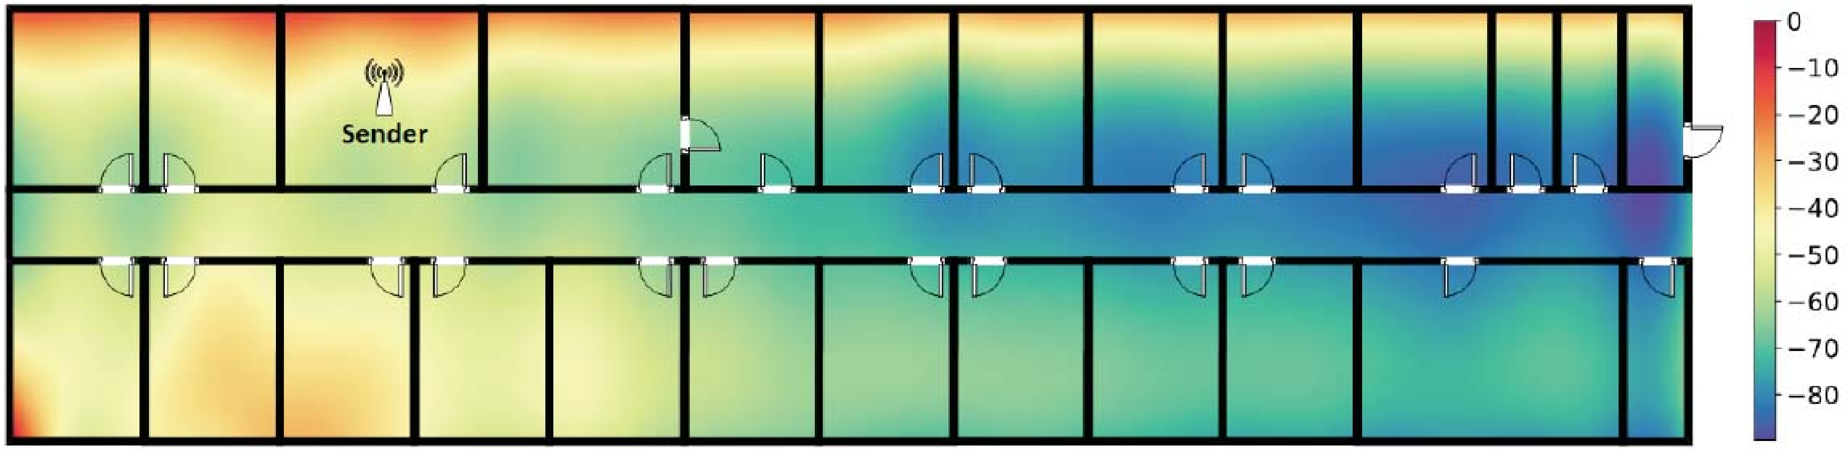

Visible Light Communication (VLC)

- Visible light is nature in distance boundary
  - Fine-grained information boundary control
  - Does not interfere with existing wireless channels
  - Many devices already have light sensors !

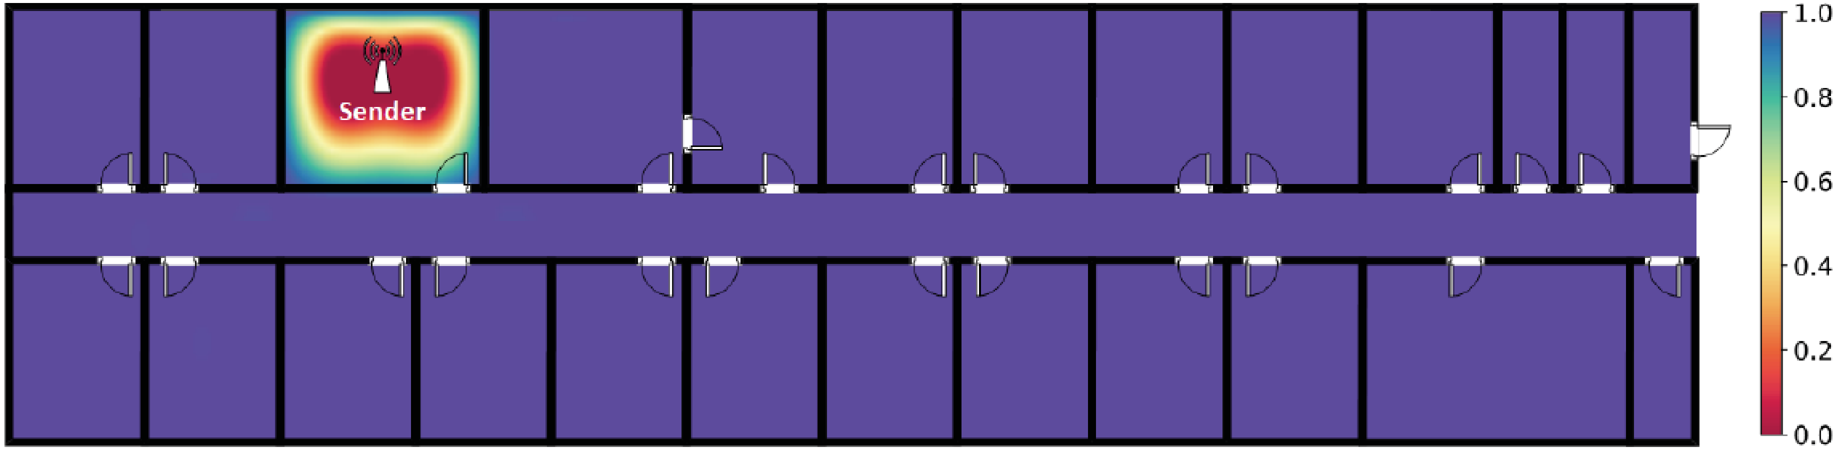

TUDelft

Attractive Features

- Distance-bounding
  - Localized service discovery and advertisement
  - Smart home services e.g., authentication

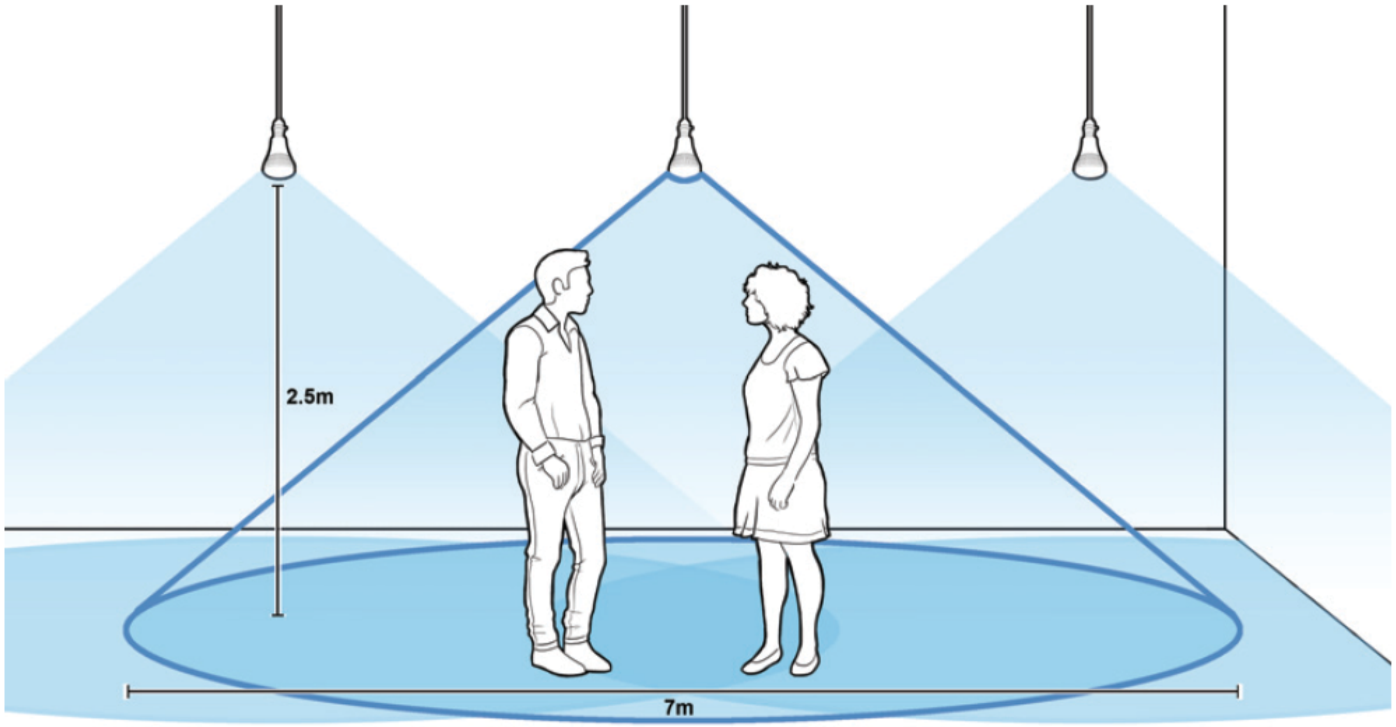

Core Design

- VLC System
  - Morse coding for light
  - Linux kernel modules
  - Service protocols

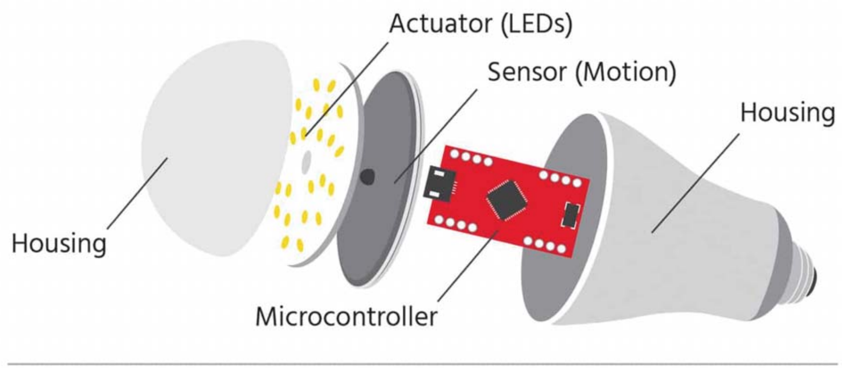

IoT-enabled LED

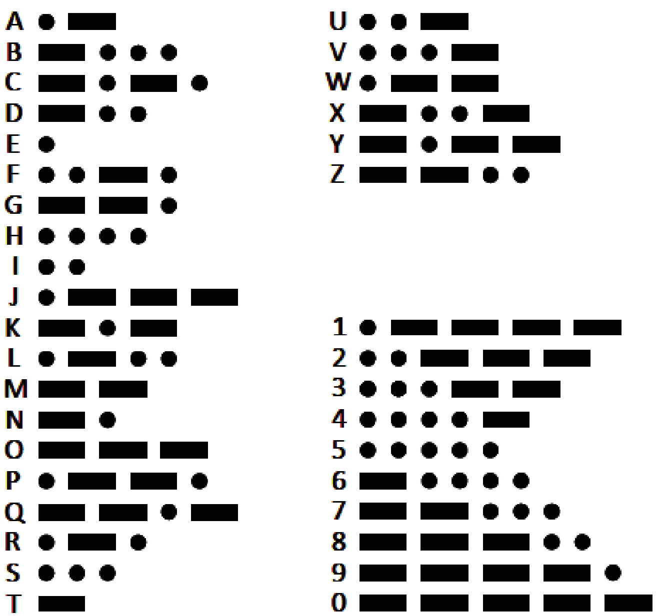

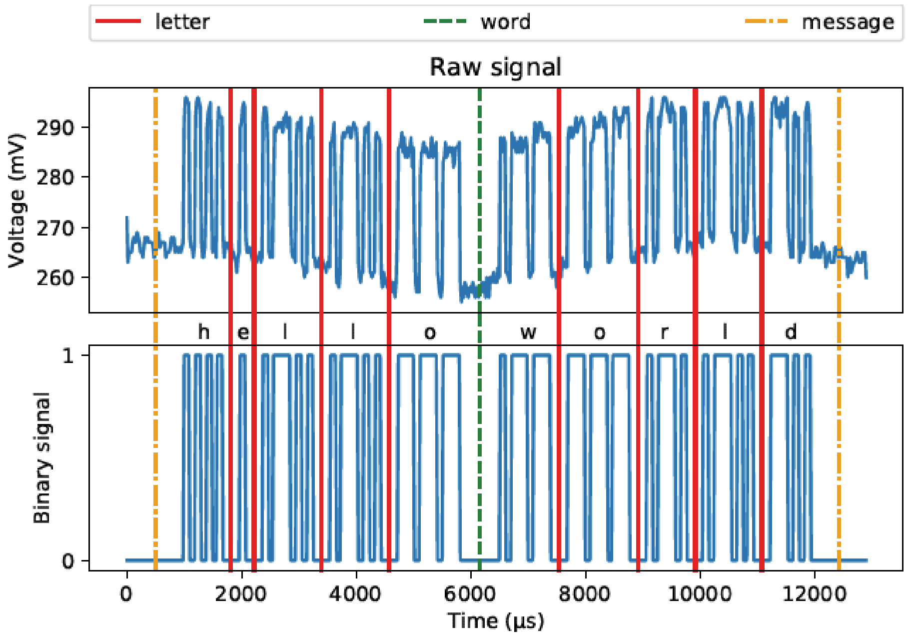

Morse code

LocalVLC

- Visible light based communication system

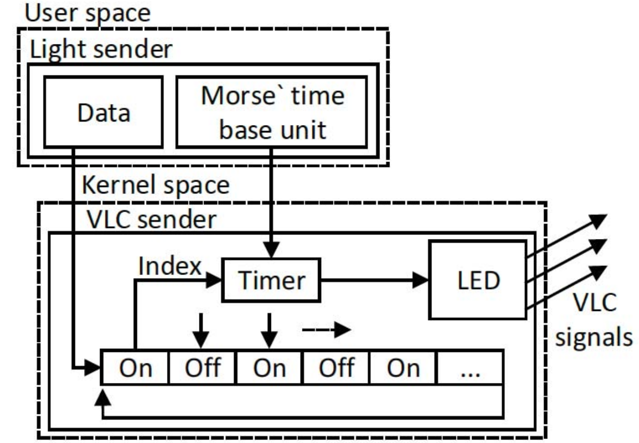

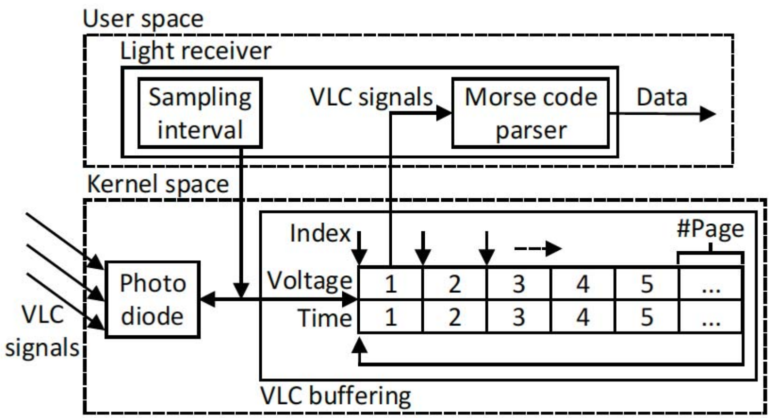

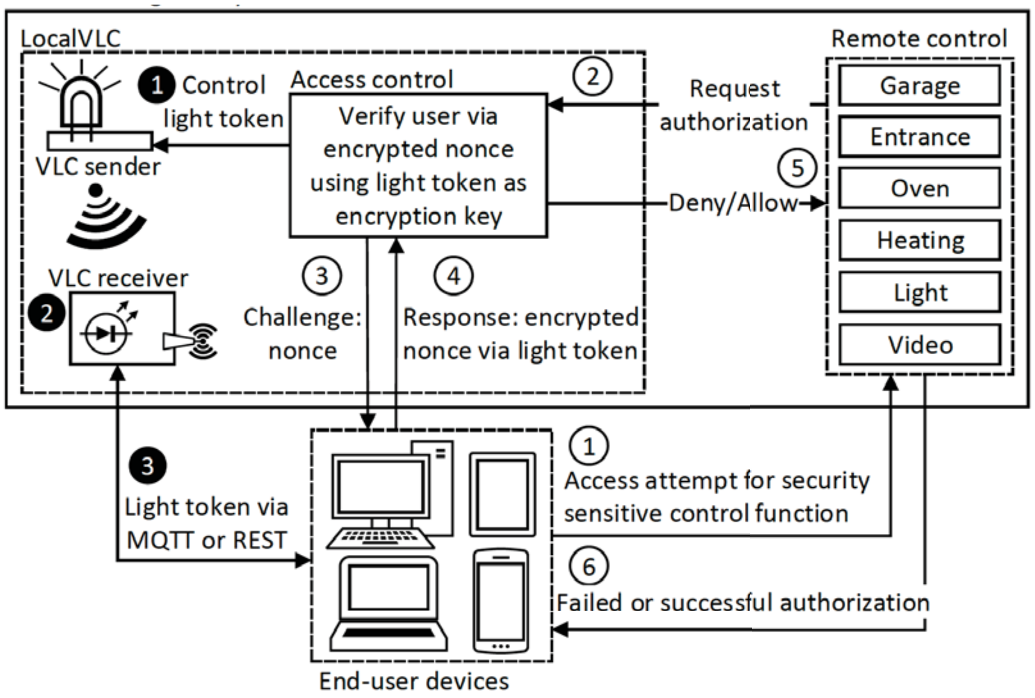

System Implementation

- Light based system
  - 3D printed bulb
  - Off-the-Shelf components

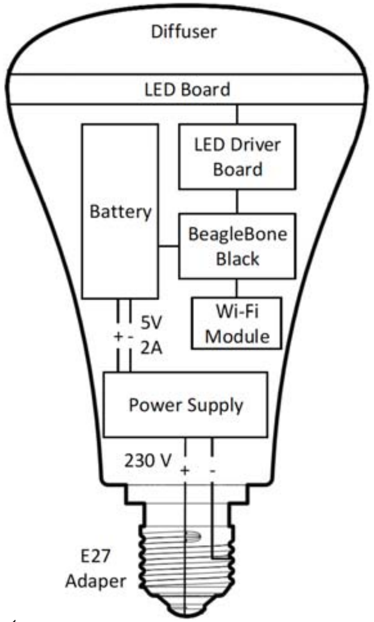

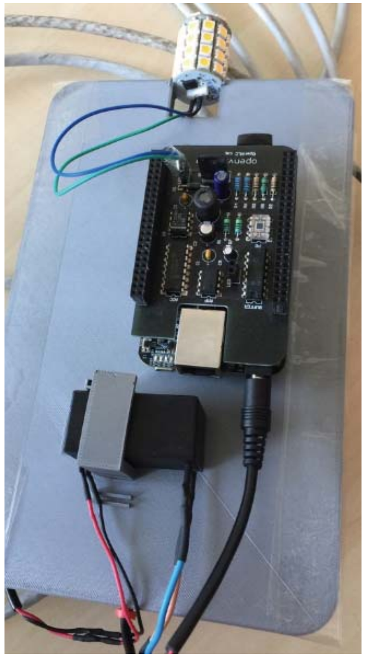

TUDelft

Evaluation

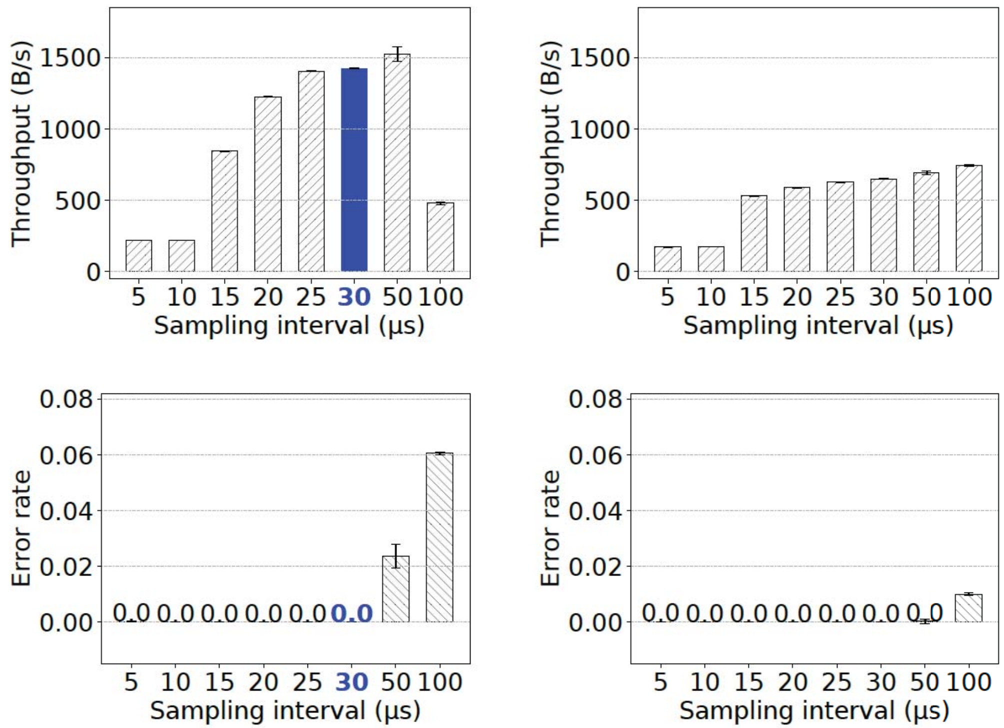

$\lambda = 50\,\mu s$

$\lambda = 100\,\mu s$

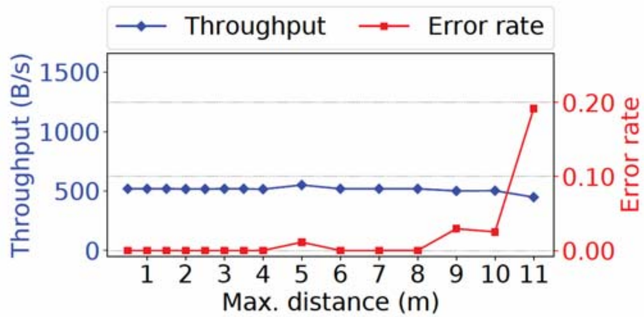

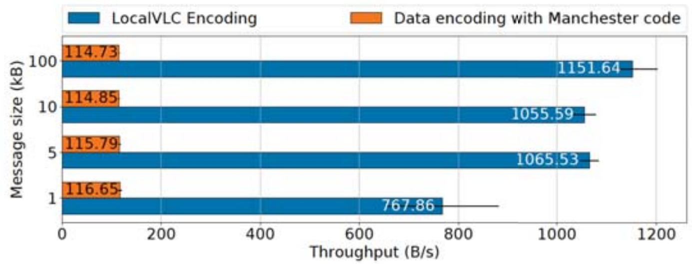

(a) Throughput (B/s)

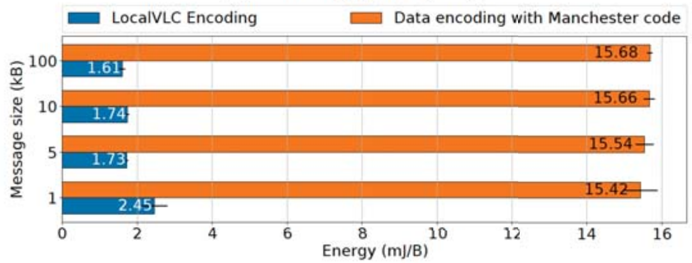

(b) Energy consumption (mJ/B)

Indoor ambient light		Directional LED		Omnidirectional LED	
Level	Intensity (lx)	Throughput (B/s)	Error rate	Throughput (B/s)	Error rate
Low	6.34 ± 0.2	517.08 ± 1.84	0.0 ± 0.0	1423.25 ± 1.84	0.0 ± 0.0
Mid	18.73 ± 0.22	575.22 ± 8.44	0.17 ± 0.01	1.79 ± 0.52	0.71 ± 0.08
High	39.53 ± 0.28	654.09 ± 20.27	0.38 ± 0.04	22.66 ± 1.52	0.35 ± 0.01

Ready to Publish !

- Measurements
  - Wireless vs VLC
- System design
  - “New” encoding
- Implementation
  - Running system !

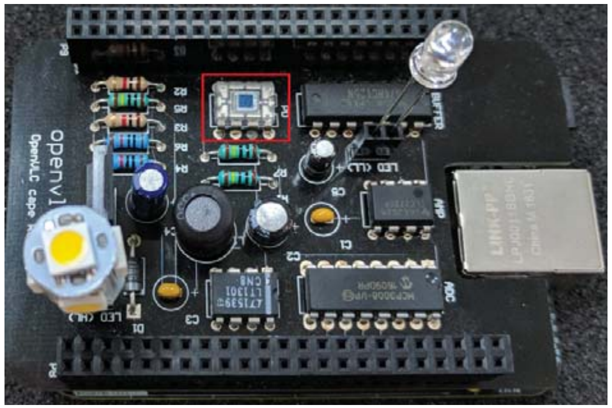

Algorithm 1: Morse Code Processing of LocalVLC

input : voltage $\leftarrow (v_1, v_2, \ldots, v_n)$, t $\leftarrow (t_1, t_2, \ldots, t_n)$,
morse-code-dict
output: message

Step 1: preprocessing

$\bar{v} \leftarrow \frac{1}{n} (\sum_{i=0}^{n} \text{voltage}_i)$
for $i \leftarrow 0$ **to** n **do**
    voltage-on-off [i] $\leftarrow \begin{cases} 1, & \text{if } v_i > \bar{v} \\ 0, & \text{otherwise} \end{cases}$
end
for $i \leftarrow 0$ **to** n **do**
    changepoint [i] $\leftarrow$ voltage-on-off [i+1] - voltage-on-off [i]
end
changepoint-pos $\leftarrow$ `seek` (changepoint = 1 or = -1)
for $i \leftarrow 0$ **to** n **do**
    duration [i] $\leftarrow$ t[changepoint-pos [i+1]] - t[changepoint-pos [i]]
end

Step 2: parsing

voltage-on-off $\leftarrow$ voltage-on-off [changepoint-pos]
voltage-off-pos $\leftarrow$ `seek` (voltage-on-off == 0)
$\theta_{\text{letter}} \leftarrow \frac{1}{n} (\sum_{i=0}^{n} \text{duration[voltage-off-pos}_i])$
$\theta_{\text{word}} \leftarrow 2.5 \cdot \theta_{\text{letter}}$, $\theta_{\text{msg}} \leftarrow 4.5 \cdot \theta_{\text{letter}}$
voltage-on-pos $\leftarrow$ `seek` (voltage-on-off == 1)
duration-on $\leftarrow$ duration [voltage-on-pos]
$\theta_{\text{dash}} \leftarrow \frac{1}{n} (\sum_{i=0}^{n} \text{duration-on}_i)$
dash-pos $\leftarrow$ `seek` (voltage-on-off == 1 and duration $> \theta_{\text{dash}}$)
voltage-on-off [dash-pos] = dash
letter-pos $\leftarrow$ `seek` (voltage-on-off == 0 and duration $> \theta_{\text{letter}}$)
letters $\leftarrow$ `split` (voltage-on-off, letter-pos)
duration-off $\leftarrow$ duration [letter-pos]

Step 3: translation

message $\leftarrow \emptyset$
for $i \leftarrow 0$ **to** n **do**
    letter-on-pos $\leftarrow$ `seek` (letters$_i$ == 1 or == 3)
    letter-pattern $\leftarrow$ letters$_i$ [letter-on-pos]
    message $\leftarrow$ `append` (morse-code-dict [letter-pattern])
    if duration-off$_i$ $> \theta_{msg}$ **then**
        message $\leftarrow$ `append` ("\n")
    end
    else if duration-off$_i$ $> \theta_{word}$ **then**
        message $\leftarrow$ `append` (" ")
    end
end

TU Delft

Securing Edge Communication for IoT

- Distance-bounding secure services
  - Security

Reject x 2

All security properties are inherited from the out-of-band VLC, which by default is assumed not to be accessible to any adversary.

There is no adversary model or security analysis to adequately determine how secure the system actually is.

It's obvious to me that VLC is distance-bounded by nature and any service using this technology is distance bounding. Thus, it appears to me as if the paper promoted VLC technology rather than introducing any novelty.

Classic Mistake…

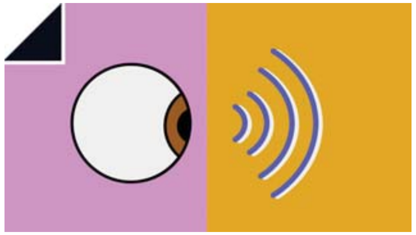

Business Impact

Using Wi-Fi to “see” behind closed doors is easier than anyone thought

With nothing but a smartphone and some clever computation, researchers can exploit ambient signals to track individuals in their own homes.

by Emerging Technology from the arXiv   November 2, 2018

- Property is not contribution
  - Over emphasis on the “attribute”
  - Attractive feature without “novelty”
  - System evaluation mismatch

- Place on the wrong plate
  - Hard to justify
  - Hard to compare

What to Do ?

- Out-of-Box reflection
  - Annoying VLC blinking
  - Darklight hints

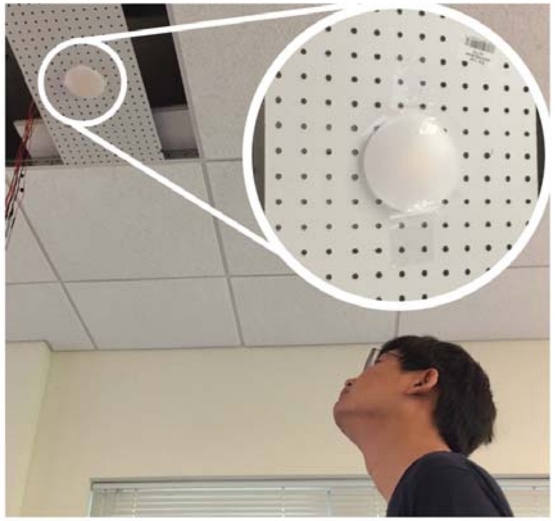

- Usability ?
  - Eliminate light flickering effect
  - Morse code efficiency
  - Energy efficiency
  - Programmable

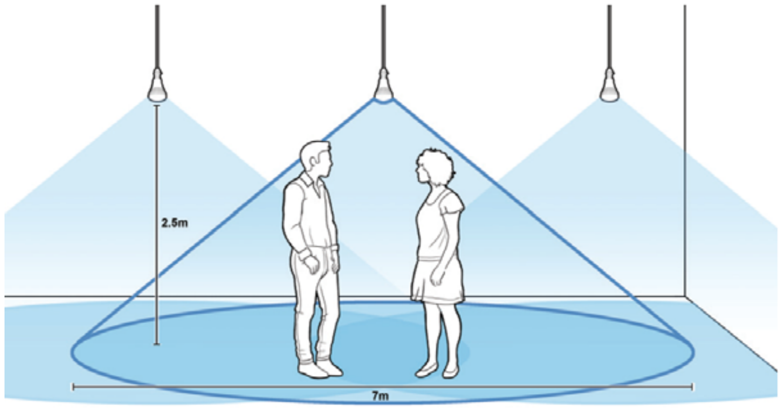

Things start Rolling

LocalVLC in Practice

- Hands-free wireless authentication
  - Interaction and authentication with IoT
- Live demo in Munich

[1] "Demo Touchless Wireless Authentication via LocalVLC" ***ACM MobiSys 2018***

[2] "Enhancing Indoor IoT Communication with Visible Light and Ultrasound" ***IEEE ICC 2019***

[3] "LocalVLC: Augment Smart IoT Services with Practical Visible Light Communication" ***IEEE WoWMoM 2019***

[4] "DevLoc: Light Bulb Networks for IoT" ***ACM/IEEE IoTDI 2020***

Lessons

- Risk 1: where to position core competence
  - Be careful with first impression
  - Do not mix ‘properties’ with scientific contributions
- Risk 2: match or mismatch
  - What’s done shall match with what we claim
  - What’s evaluated shall match the focus

Peter Hofstee also agrees ☺

What about Intelligence ?

Edge Research @ CPI Lab

Edge Analytics and Services

Edge Computing	Edge Communication	Edge Security

System Platform

Cyber-Physical Infrastructure

Wireless Edge Analytics

- Demand for wireless edge analytics
- Mobility + Traffic
  - Interplay
  - Across device types
  - Modeling insights

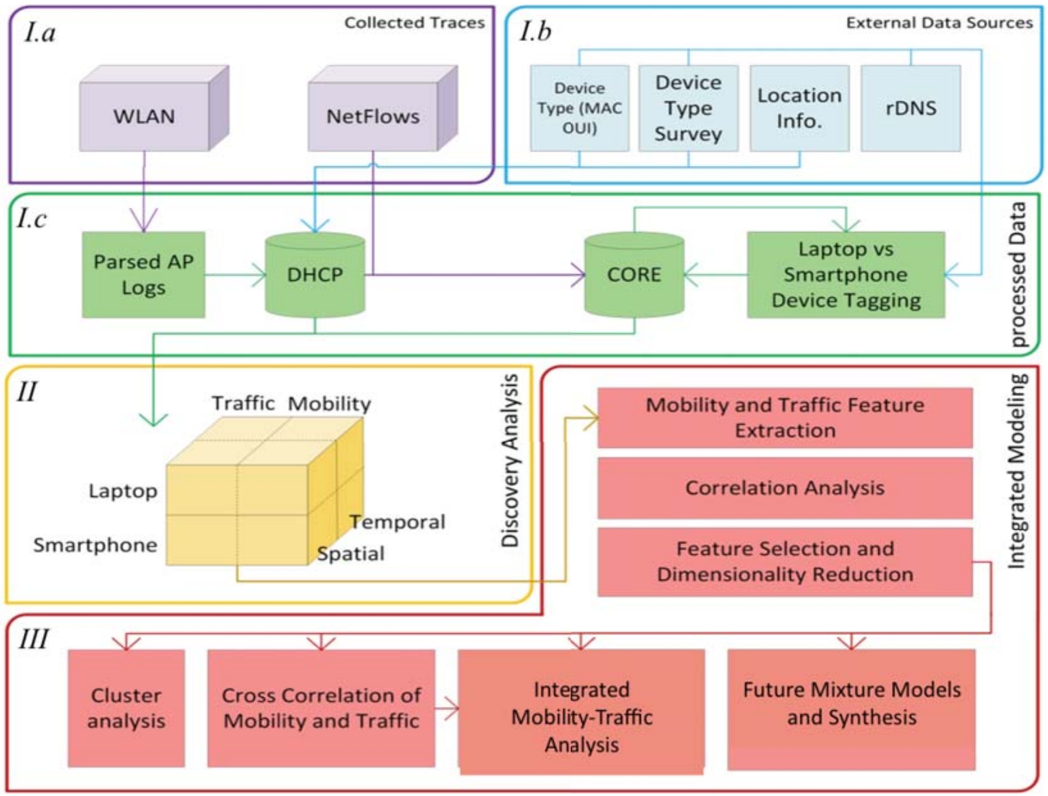

Motivation

- **Mobility-Traffic Interdependence** is not well-studied

- Major factors affecting mobile network performance are **mobility** and **traffic** patterns

We Need Dataset

  - Simulations, analytical based performance evaluations, and future predictive caching schemes rely on **models** to approximate factors affecting the network

- Many earlier mobility modeling studies use pre-smartphone WLAN traces

FLAMeS Dataset

- Size of raw dataset
  - 30+ TB, 1760 APs, 138 buildings, over 479 days
  - 76 billion NetFlow records, 555 million AP traces, 316k devices
- Device categorization
  - MAC address survey
  - OUI matching
  - Web domain analysis

	# Records		Traffic Vol. (TB)		# MAC	
	DHCP	CORE	TCP	UDP	WLAN	CORE
Flutes	412.0 M	2.13 B	56.18	4.50	186.0 K	50.3 K
Cellos	101.0 M	4.20 B	73.85	12.90	93.2 K	27.1 K
Total	557.5 M	6.53 B	134.39	17.61	316.0 K	80.0 K

TU Delft

Good Dataset is indeed Nice

- Down to details

TCP accounts for 78.5% of cello flows (84.6% of bytes) and 98.2% of flute flows (91.6% of bytes). The higher presence of UDP in cellos is reasonable, considering that UDP applications (e.g., multi-player games, video conferencing and file sharing) are more likely to be used with cellos. Comparing the number of packets in flows, in case of TCP, the average number of packets in cello flows is almost half that of a flute flow (4.6 vs 8.8), and the average packet size of flutes is 22% higher than that of cellos. This supports our earlier observation regarding the bigger flow sizes of flutes. However, for UDP, the two device types are similar in terms of average packet count per flow (2.5 for cellos & 2.87 for flutes) and average packet size (119 for both). This conforms to low latency requirements of many UDP applications.

Time

Discovery and Insights

- Mobility analysis
  - Session start probability, radius of gyration, visit preference, sessions per building, etc.
- Traffic analysis
  - Flow level, spatial, temporal behavior
- Integrated analysis
  - Feature engineering, modeling insights

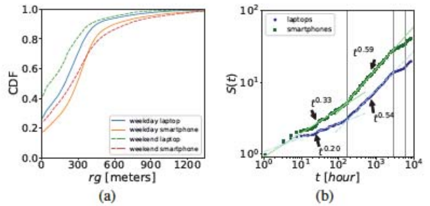

Fig. 4: (a) Radius of gyration (rg for the device types). (b) Visited locations $S(t)$. Vertical lines at 7, 120 and 240 days.

session at a building b, here referred as DLT. Interestingly, cellos have slightly longer stays but both have medians around 2:40 hours. The similarity of the distributions, combined with a lower number of visited locations indicate that cellos are used mostly when users remain longer periods at places.

Fig. 4b highlights the differences between *flutes* and *cellos* on the required time t to visit $S(t)$ locations. *After an initial exploration period of one week the rates of new visits change similarly for both device types, and new exploration rates show up at 120 and 240 days.* These could be explained by the weekly schedules of the university as well as the usual length of a lecture term (≈ 4 months).

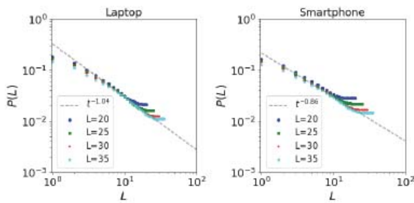

Fig. 5: Zipf's plot on L visited access points.

We also consider the number of unique APs a device associates with, APC, which provides a finer spatial resolution than the building level. Furthermore, the probability of finding a device at its *L-th* most visited access point is shown in Fig. 5. When taking buildings as aggregating points for location, the values become $L^{-1.36}$ for *cellos* and $L^{-1.16}$ for *flutes*. These approximations validate previous work on human mobility [8], yet highlight differences between device types.

D. Sessions per building

To study AP utilization over time, we look at the session duration distribution, or session duration dispersal kernel P(t), depicted in Fig. 6. The smaller inner plots represent the same metric, limited to four types of buildings.

We noted that the five-minute spikes correspond to default idle-timeout for the used WiFi routers. On the other hand, the *knees* at 1 and 2 hours could be explained by the typical duration of classes. They are only noticeable at Academic buildings (shown inside inner plots) and during weekdays (not shown). This leads us to conclude that despite the differences in distributions of device types, *flutes* and *cellos* present *certain similarities in their usage, such as during classes.* To differentiate *pass-by* access points, we examine all sequences of three unique APs where all session durations are lower than 5 minutes (typical idle-timeout). We observed these APs clustered at buildings that also had major bus stops nearby.

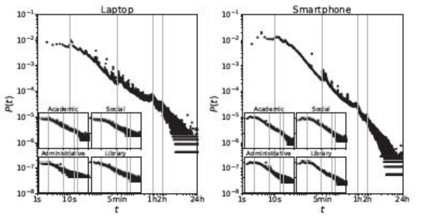

Fig. 6: Probability $P(t)$ of session duration t.

VI. Traffic Analysis

In this section, we compare different *traffic* characteristics, across *device types*, *time* and *space*. For this purpose, we start with statistical characterization of *individual* flute and cello flows. Next, we measure how these flows, *put together*, affect the network patterns across APs and buildings. Finally, *user behavior* is analyzed by monitoring weekly cycles, data rates, and active durations. By quantifying *temporal* and *spatial* variations of traffic across device types, we make a case for new models to capture such variations based on the most relevant attributes. Table IV summarizes the results.

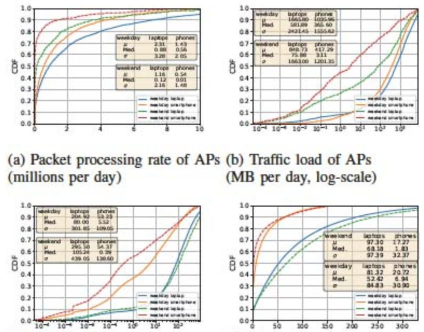

(a) Packet processing rate of APs (millions per day)

(b) Traffic load of APs (MB per day, log-scale)

(c) User data consumption (MB per day, log-scale)

(d) User active time (minutes per day)

Fig. 7: Distribution plots

Big Edge Data for the Win ?

- What boasted, all **fired back**

“ Your data is not new enough ”

“ Your findings may not reflect the latest situation ”

“ Your analysis coverage is limited ”

“ Your insights for modeling are incomplete ”

“ Your work impact is not … ”

…

TU Delft

What Went Wrong ?

- Method or dataset ? What is the core value ?

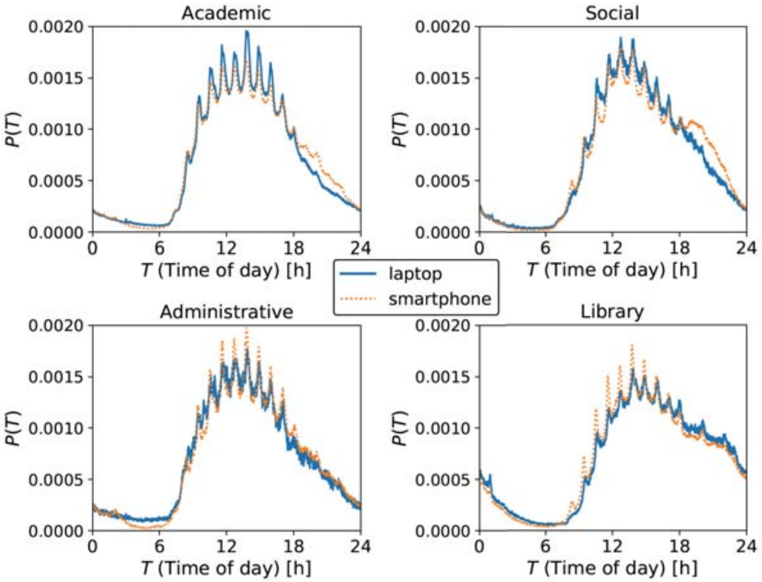

	Flutes (F)			Cellos (C)			Ratio (C/F)	
	μ	*mdn*	σ	μ	*mdn*	σ	μ	*mdn*
LJM	435 350	296 168	813 683	178 97	1 1	624 312	0.409 0.277	**0.003** 0.006
DIA	549 425	411 179	874 739	195 107	1 1	642 338	0.355 0.252	**0.002** 0.006
TJM	1582 1036	707 279	2336 1793	378 252	1 1	1444 1766	0.239 0.243	**0.001** 0.004
GYR	396 330	290 248	2725 1368	321 178	191 65.1	3265 1800	1.102 1.247	1.019 1.4
BLD	5.4 2.8	3 2	5.6 4.1	1.8 1.5	1 1	2.1 1.8	0.811 0.539	0.659 0.262
APC	11.8 7.2	6 4	13.3 8.8	3.7 3	2 2	4.8 3.8	0.333 0.536	0.333 0.5
PDT	225 223	161 135	219 272	248 278	164 189	254 292	0.314 0.417	0.333 0.5
DTL	316 326	235 247	302 308	316 316	217 221	305 309	1 0.97	0.92 0.89

Start time	Finish time	Duration	Source IP	Destination IP	Protocol	Source port	Destination port	Packet count	Flow size
1334332274.912	1334332276.576	1.664	173.194.37.7	10.15.225.126	TCP	80	60482	157	217708

User IP	User MAC	AP name	AP MAC	Lease begin time	Lease end time
10.130.90.3	00:11:22:33:44:55	b422r143-win-1	00:1d:e5:8f:1b:30	1333238737	1333238741

TU Delft

Back to the Basics

- Wireless edge analytics

Data	Analysis	Modeling
• NetFlow • WLAN Traces • Device Classes • DHCP • Core (merged)	• Cellos vs Flutes • Spatio-Temporal Characteristics • Traffic and Mobility Features	• Mobility and Traffic feature extraction • Correlation/Cross-Correlations • Mixture Models and Synthesis

TU Delft

Framework for Wireless Edge Analytics

- Feature extraction
  - WLAN logs and NetFlows

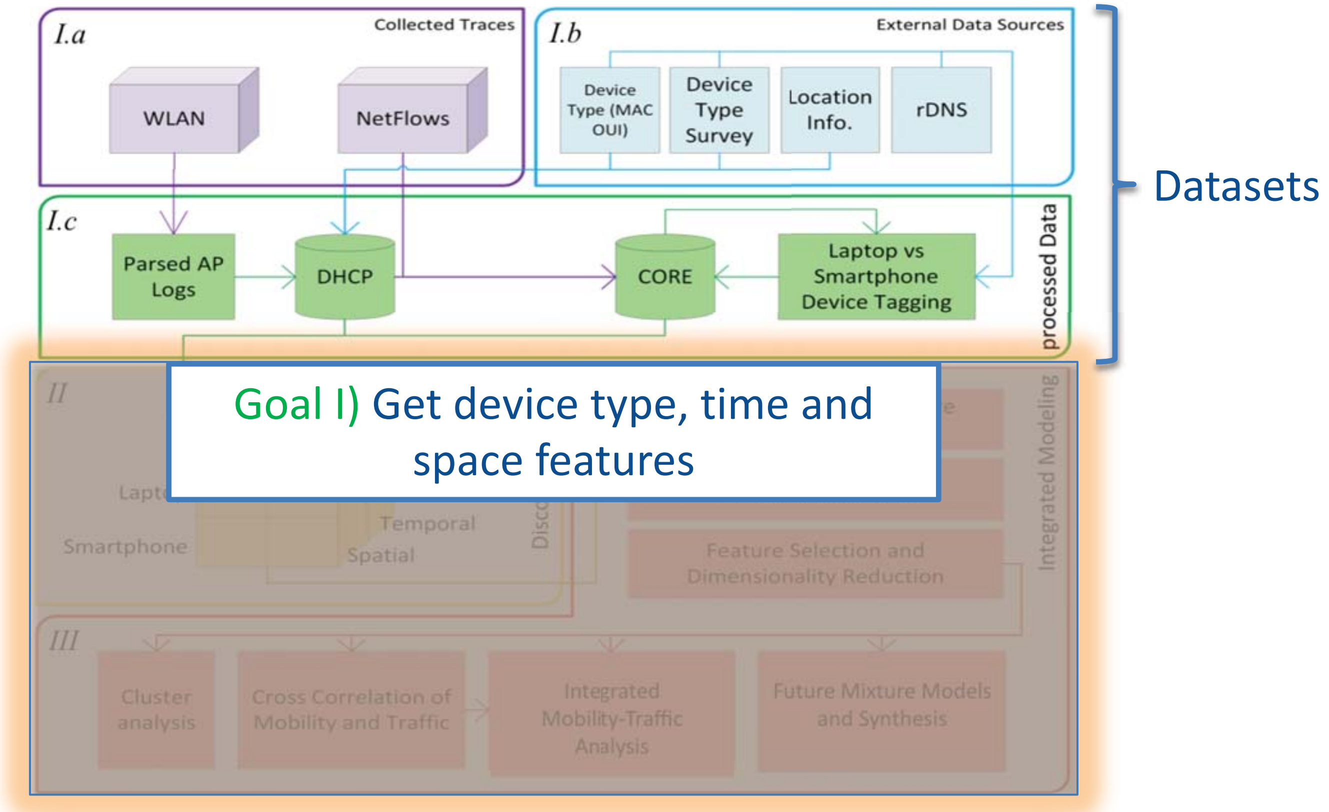

FLAMeS

- Data traffic and mobility interdependency

Data cube, **traffic/mobility** analyzed **temporally, spatially**, and **per device type**

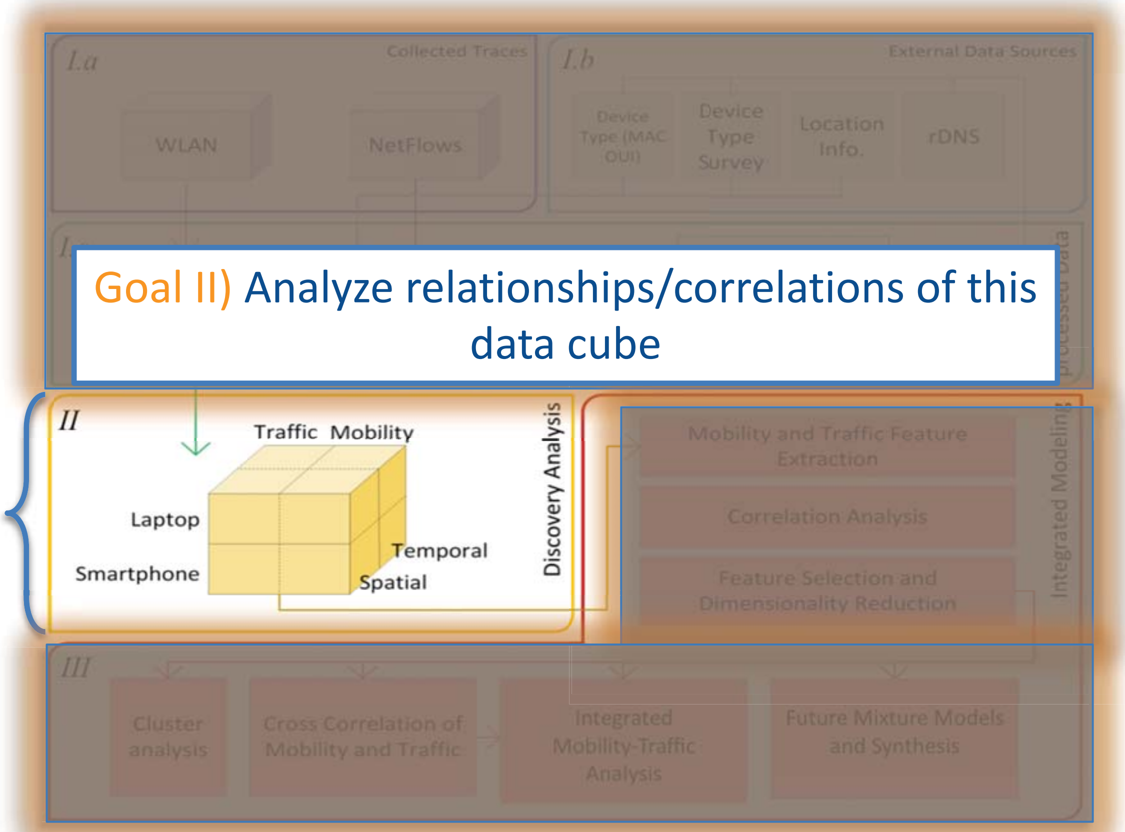

FLAMeS

- Towards integrated modeling

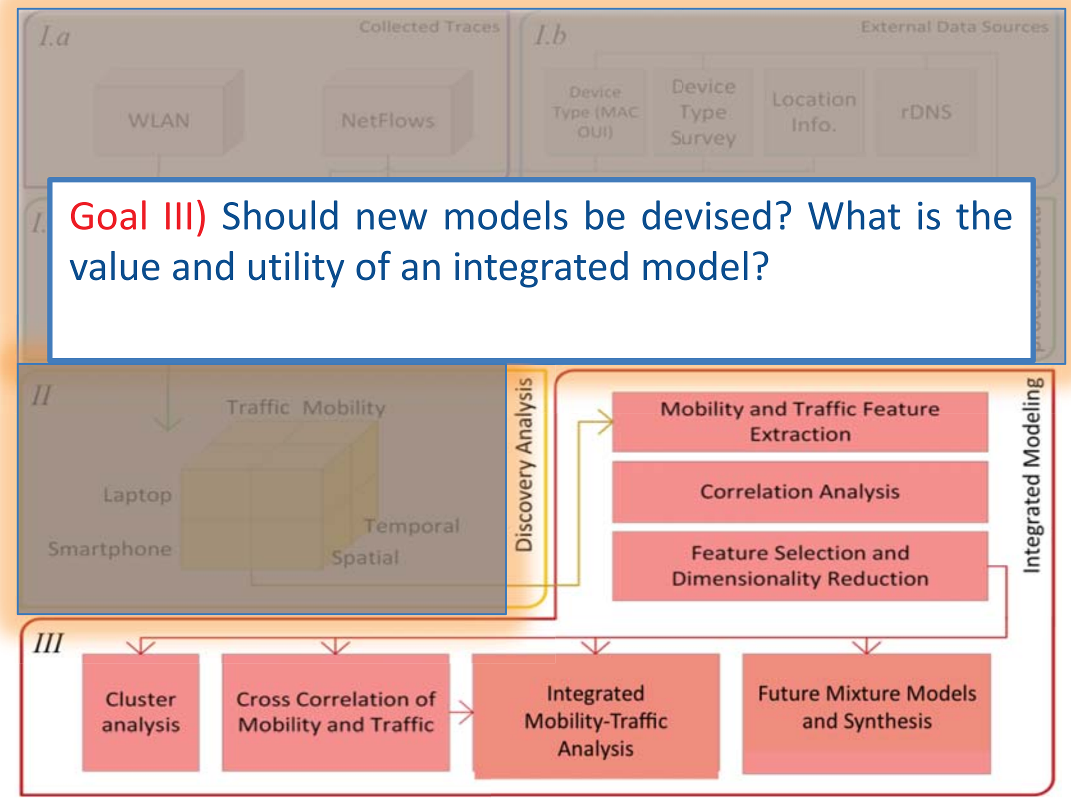

TU Delft

Adjust the Focus

- Methodology and framework
  - dataset mainly as a tool to verify our assumption and investigations

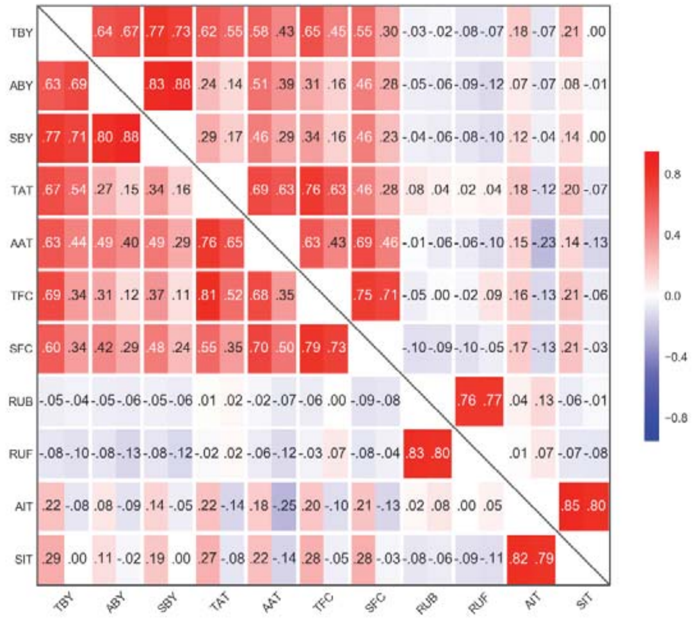

Abbr.	Description
TBY	Total flow bytes
ABY	Avg. flow bytes
SBY	Std. flow bytes
TAT	Total active time
AAT	Avg. active time
TFC	Total flow count
SFC	Std. flow counts
RUB	UDP bytes / total bytes
RUF	UDP flows / total flows
AIT	Avg. IAT
SIT	Std. IAT

We made it !

IEEE INFOCOM 2018

ACM MSWiM 2019

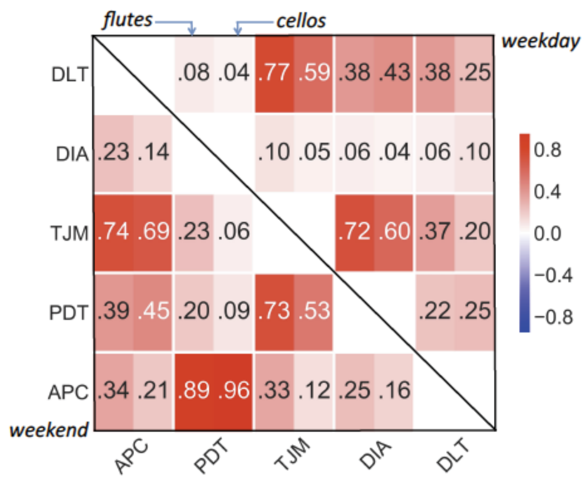

Abbr.	Description
APC	AP Count (unique)
PDT	Preferred building Δt
TJM	Total (sum) jumps
DIA	Diameter of mobility
DLT	Delta time (time of network association)

Publications on Edge Analytics

[1] "Transfer Learning-Based Outdoor Position Recovery with Telco Data" ***IEEE Transactions on Mobile Computing (TMC) 2020***

[2] "Where Are You Going Next? A Practical Multi-dimensional Look at Mobility Prediction" ***ACM MSWiM 2019***

[3] "Flutes vs. Cellos: Analyzing Mobility-Traffic Correlations in Large WLAN Traces" ***IEEE INFOCOM 2018***

Lessons

- Risk 1: Boasting dataset value
  - Avoid over claiming
  - Proper focus/position is crucial
- Risk 2: Good material needs less polishing ?
  - Can block the work from top venues
  - Balance and structure

Toolkit and in-depth study are appreciated

🏆 Best Presentation Award at IEEE INFOCOM 2018

Outlook

- Community & Collaboration
  - EdgeSys 2021
  - Dagstuhl 2021
  - Lorentz 2021

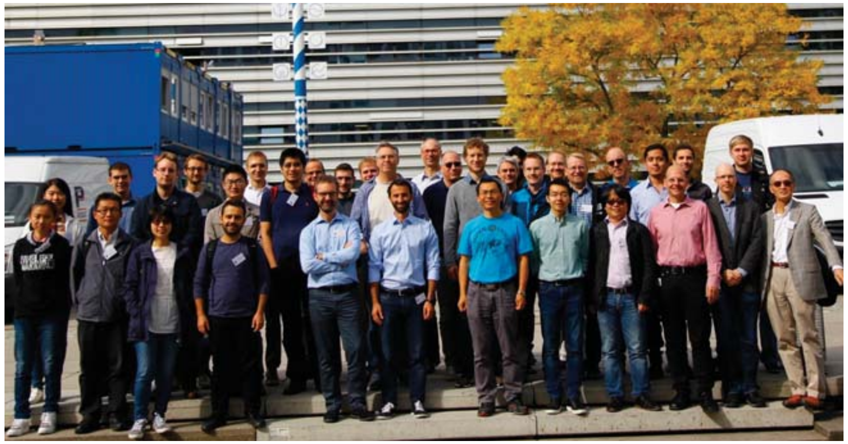

TU Delft