From Conception to Retirement: a Lifetime Story of a 3-Year-Old Wireless Beacon System in the Wild

NSDI 21

Problem Statement

- Design and deploy a system for a Instant delivery platform that is capable to infer its couriers' indoor status.
- Examples of Instant delivery platform : Doordash, uberEats, Alibaba Local Services.

Most annoying behavior of delivery apps?

Basic Flow

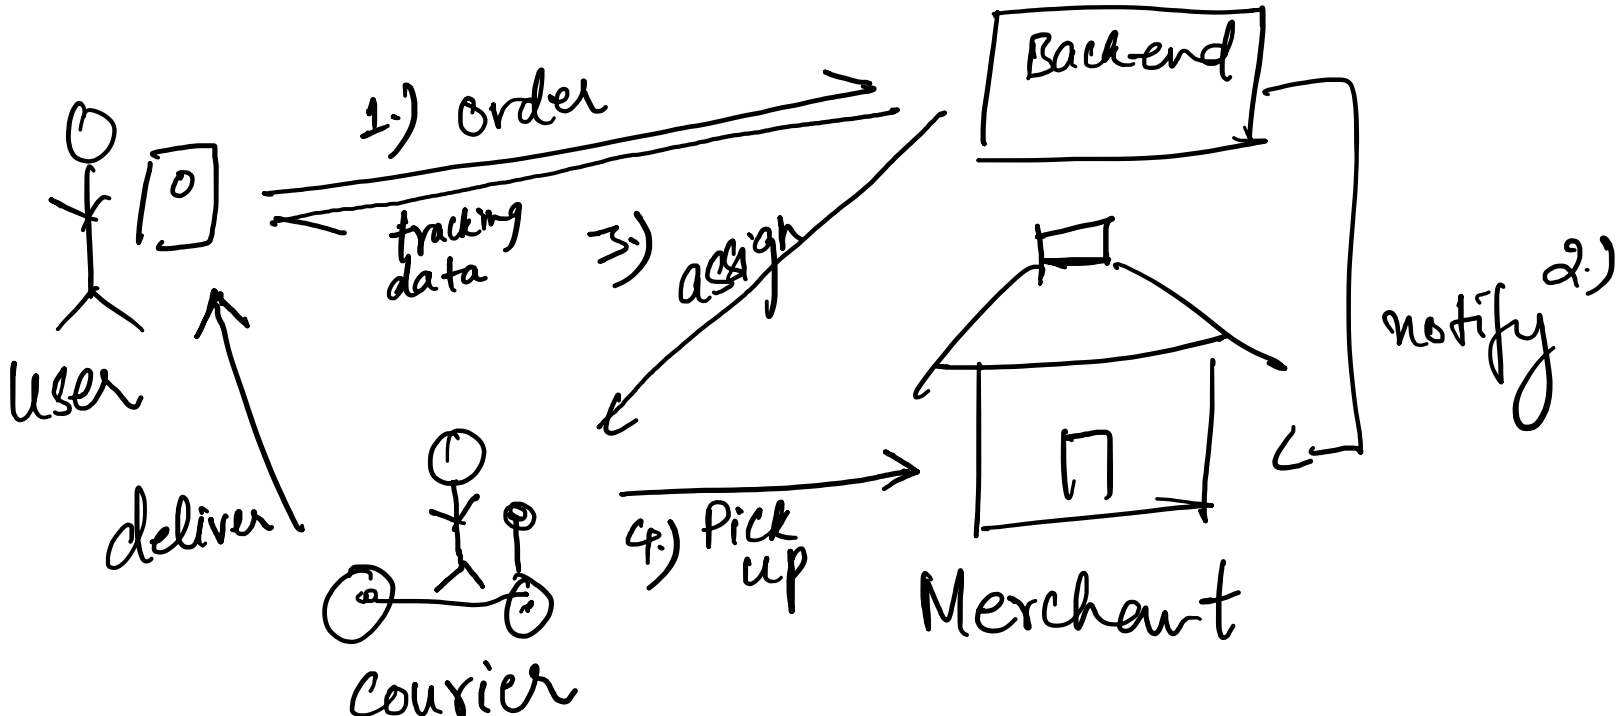

How to prevent apps from breaking promised ETA?

- Better ETA and tracking. Of course!
- Cost associated with overdue – Typically $1. If customer has insurance -> 200% x (order value) ☹
- Efficient scheduling algorithms that can assign the right courier to the task.
- In places like Shanghai, merchants are located in multistoried malls. C can be 33% of the total time. BIG CHUNK.
- How to find C?

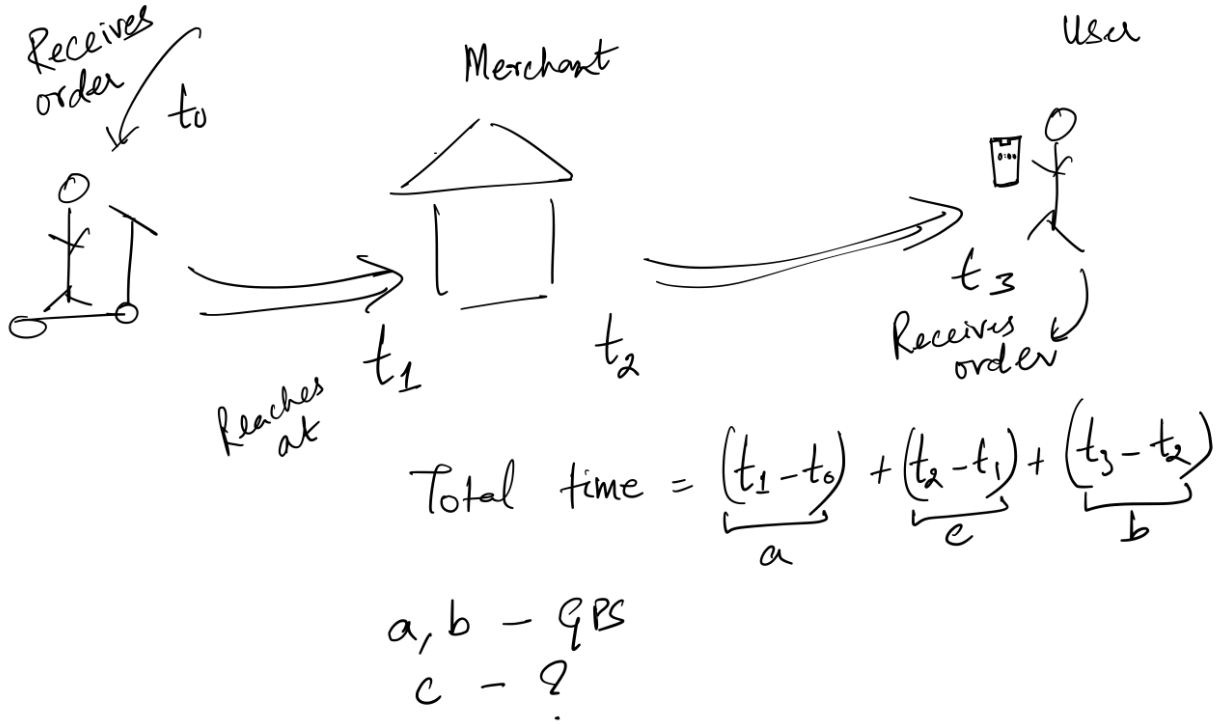

Arrival and Departure detection

- Most basic problem ever. Extensive research exists in this domain
- BUT, its usually done in a controlled environment.
- Small scale or private environment. Not “in the wild”.
- Existing solutions: Manual, WiFi, RFID, LED fixtures.

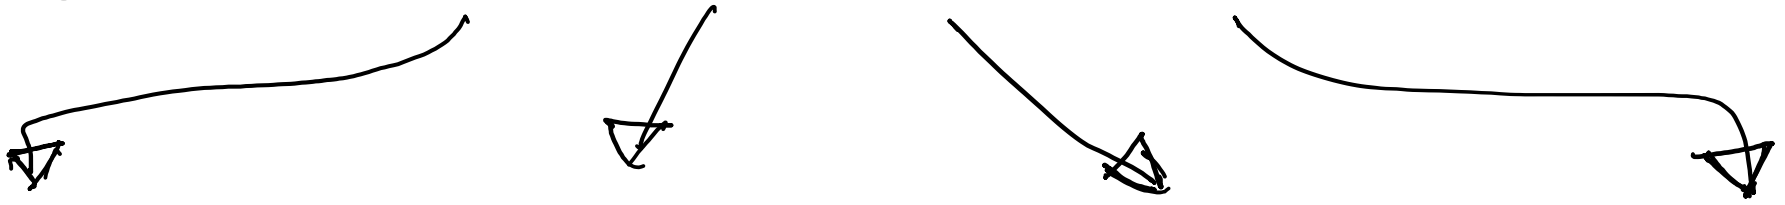

Unreliable – Intentional, unintentional QR code scanning.

1. Continuous scanning required – Power consumption
2. Access points are expensive for merchants still in the stone age

Additional equipment at both ends – Price issues.

Doesn’t scale as hardware modification is required.

BLE Beacon (Preferred Solution)

- BLE – Bluetooth Low Energy
- Easy to **deploy** (still an issue though)
- Acceptable cost ($10)
- Transparent to couriers

- Continuous scanning – 2% extra power consumptions – way less than wifi
- Only at merchant end unlike RFID
- Battery powered small devices that are portable unlike LED fixtures

aBeacon System

- Commissioned by Alibaba. Experiment took place in Shanghai.
- Uses customized BLE devices.
- Some numbers:
  - Experiment length – 3 years.
  - Total number of merchants - 12109
  - Total number of couriers - 109K
  - 64 million delivery orders for 7.3 million customers!!
  - $600 million in order values.

aBeacon Architecture

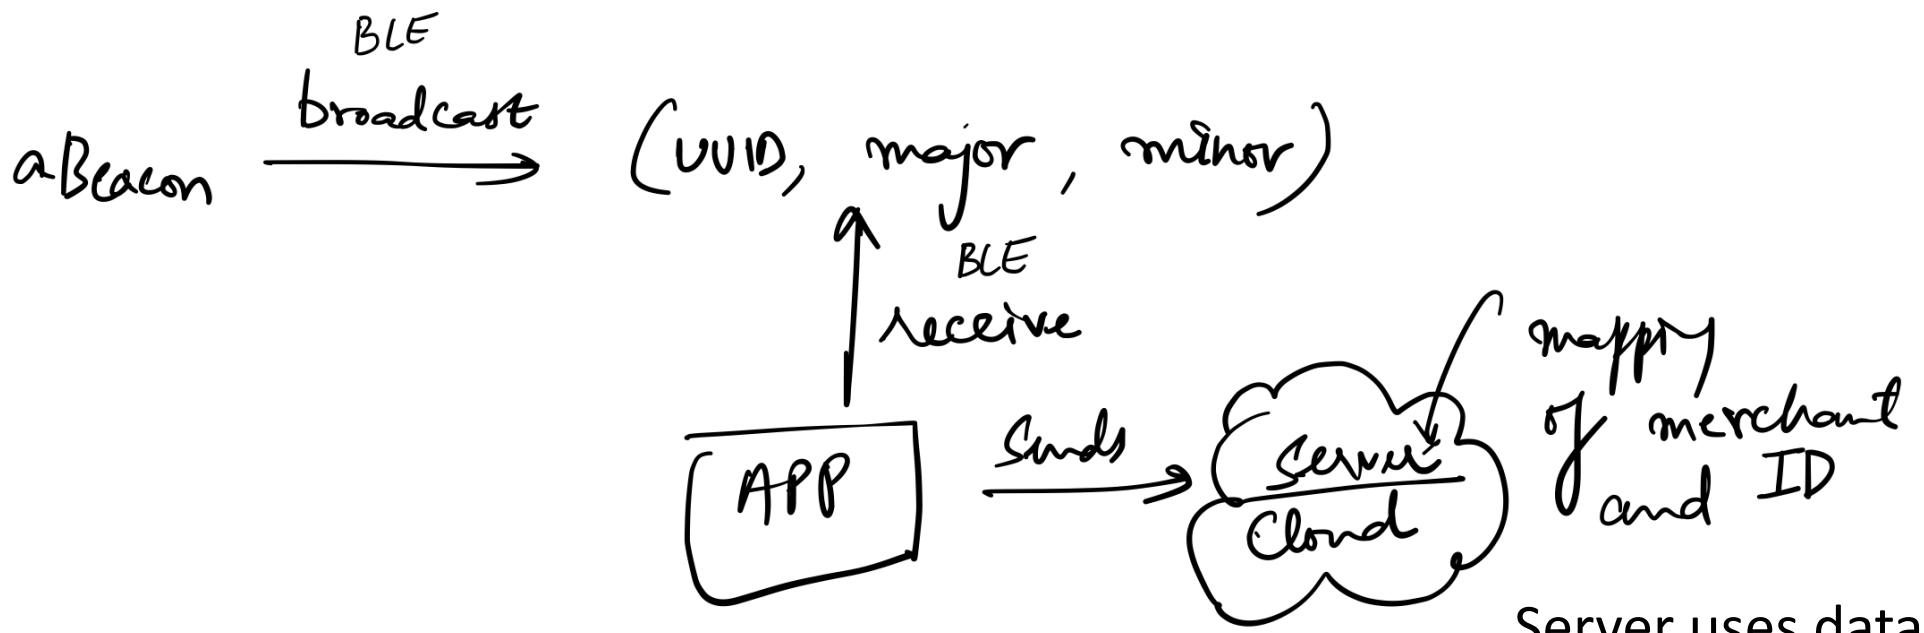

Server uses data to crunch numbers – ex : run scheduling algorithms

Justifying the cost of the experiment

- A metric based approach is used for this; to understand the cost-performance tradeoff
- Metrics –
  - Cost: Device cost and deployment cost
  - Lifetime: Battery – Estimate (2yrs)
  - Reliability: Failure rates
  - Utility: Reduction in overdue delivery rate and hence overdue cost

Justifying the cost of the experiment

- Gain Equation:

Table 1: Metric Summary

C_{Dev}:	cost of a device, i.e., hardware & deployment
C_{Over}:	cost of overdue penalty per order, e.g., \$1.
P^i_{Life}:	lifetime of a device i
P^i_{Reli}:	reliability of i
P^i_{Util}:	utility of i
t^i_0:	day of i was deployed
T:	# of days since `aBeacon` deployed
N_t:	# of deployed devices until the tth day
O^i_t:	# of orders at tth day in the merchant with i

Gain → $G_T = \sum_{t=1}^{T} \sum_{i=1}^{N_t} B_t^i - C_T$

$C_T = N_T \cdot C_{Dev}$

Total no. of devices

device cost

$B_t^i = F_1(P^i_{Life}, t, t_0^i) \cdot F_2(O_t^i, P^i_{Reli}, P^i_{Util}, C_{over})$

F_1 and F_2 are functions

$A(1,0)$ function which is 1 when device is active

$y = P^i_{Life} - (t - t_0) > 0 \; ? \; 1 : 0;$

The real benefit: How much cost saved.

O_t : number of orders

P_{reli} : % of failures

P_{util} : % reduction in overdue rate

C_{over} : Cost of overdue.

Deployment Timeline

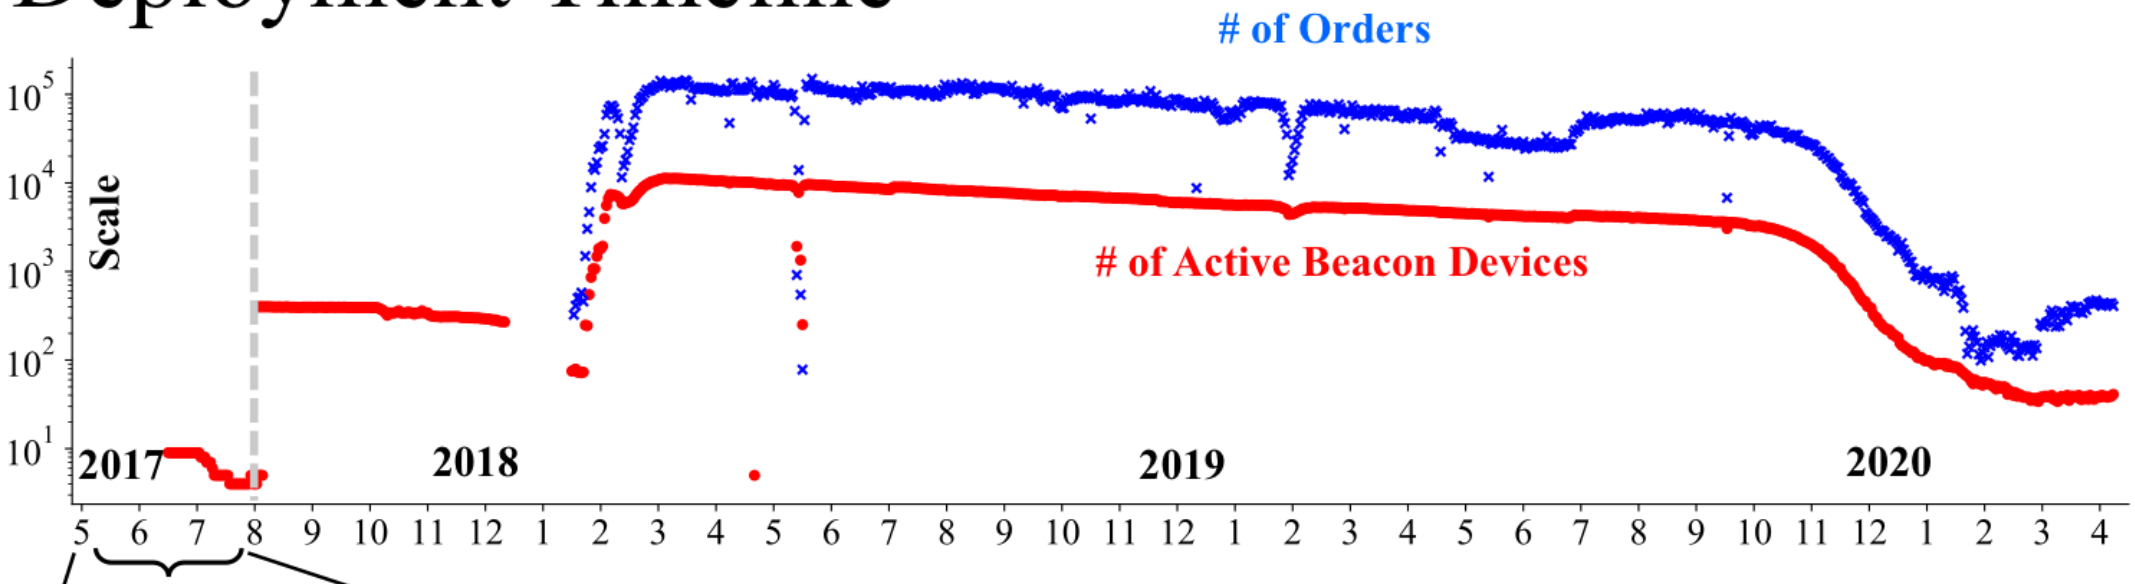

Stage 1: Conception

Size Comparison	Device 1 (T15)	Device 2 (T4)	Device 3 (T11)
Tx Power	-59 dB	-65 dB	-65 dB
Advertised Lifetime	≤ 3 yr	2 ~ 3 yr	≤ 3 yr
Cost	$11 each	$10 each	$10 each
Encapsulation	—	Water, Dust, Shock Proof	Dust Proof Only

Phase 1 – Pilot
Conception stage:

- 3 months
- 18 merchants, 3 bands = 54 devices
- Average reliability - 98%

Cost looks high. Should try to reduce it. Each dollar counts in a huge experiment like this.

Deployment Timeline

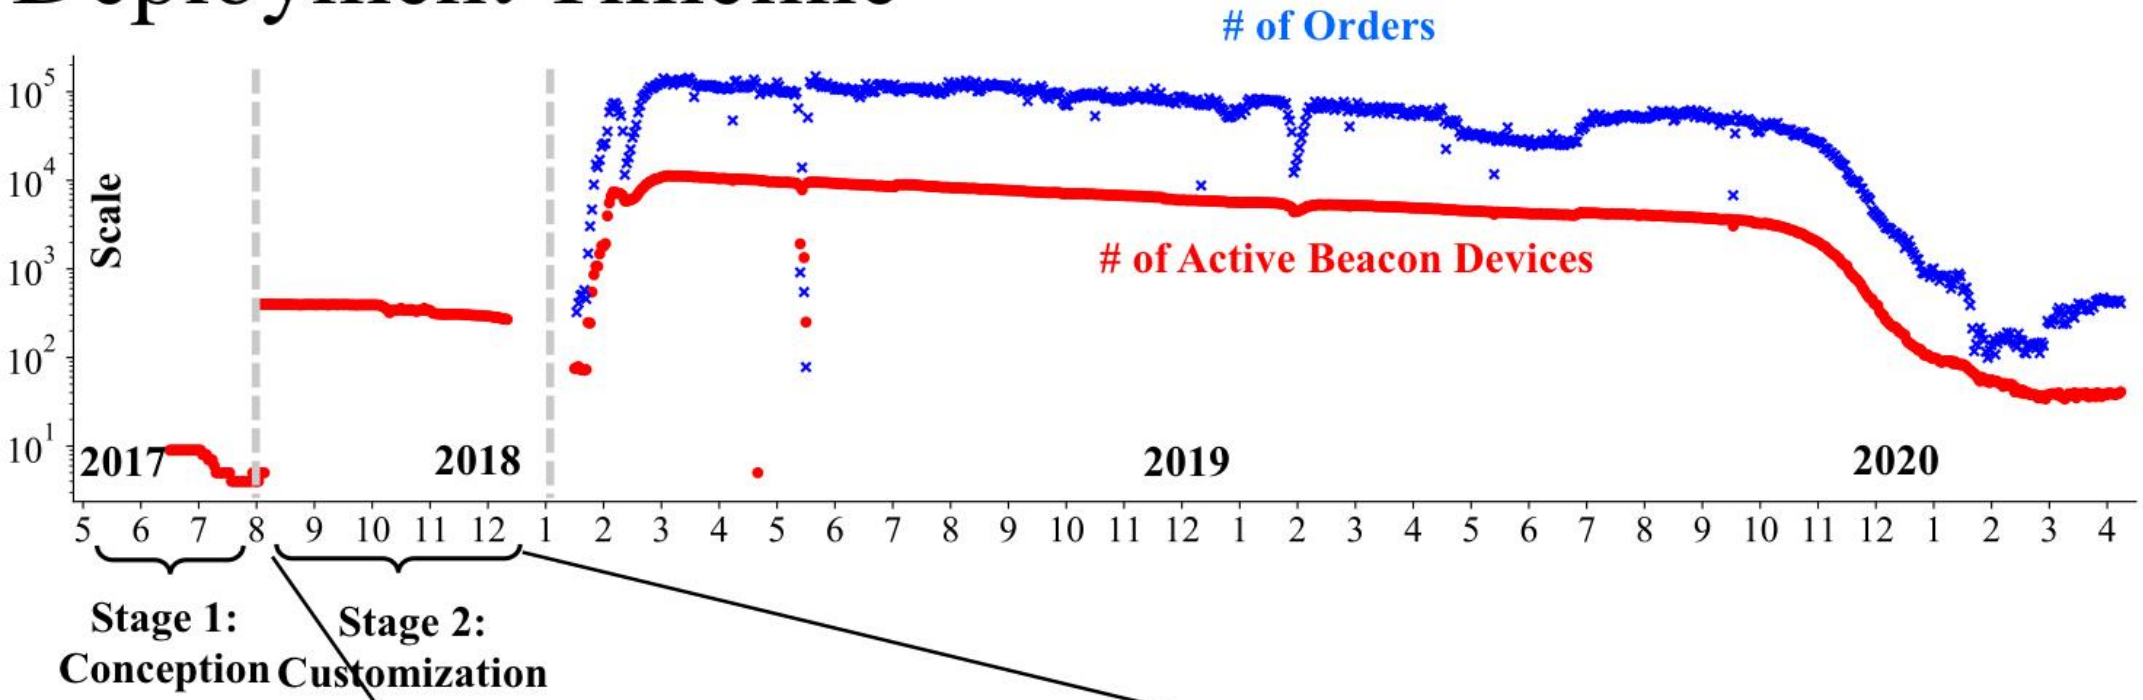

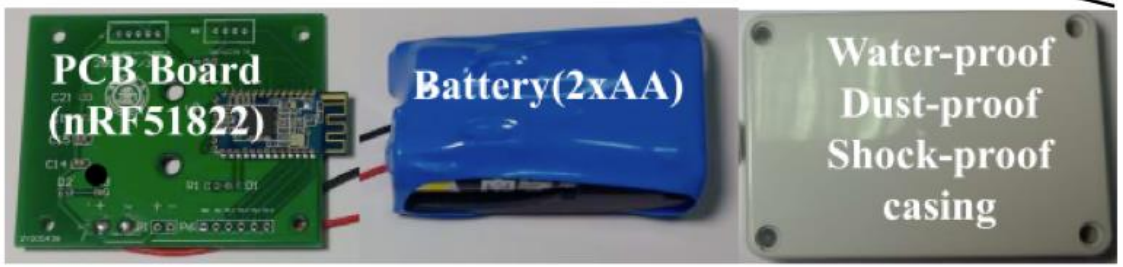

Customized beacon devices:

- Less cost ($8 each)
- Longer lifetime (>=2 years)

Phase 1 – Last 5 months

- BLE module with low power consumption
- Alkaline battery proffered over lithium since less mAh/$
- Casing for future proof – Outdoors
- A/B testing with custom and commodity sensor - 200 units in 200 stores. Similar reliability and lower cost 👍

Deployment Timeline

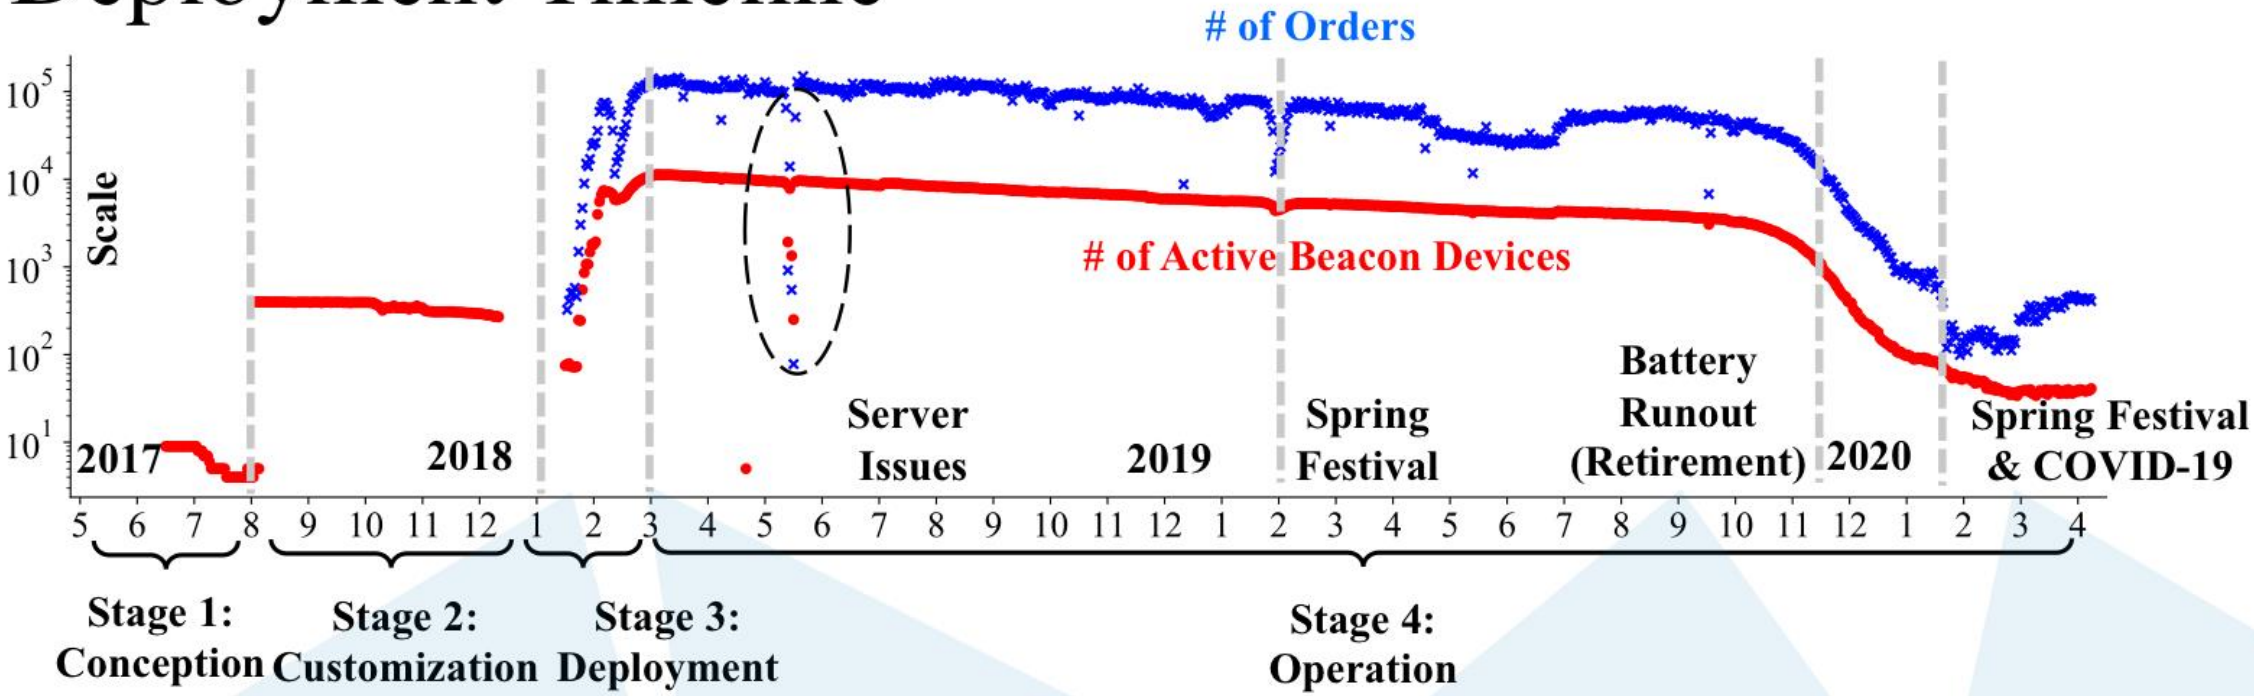

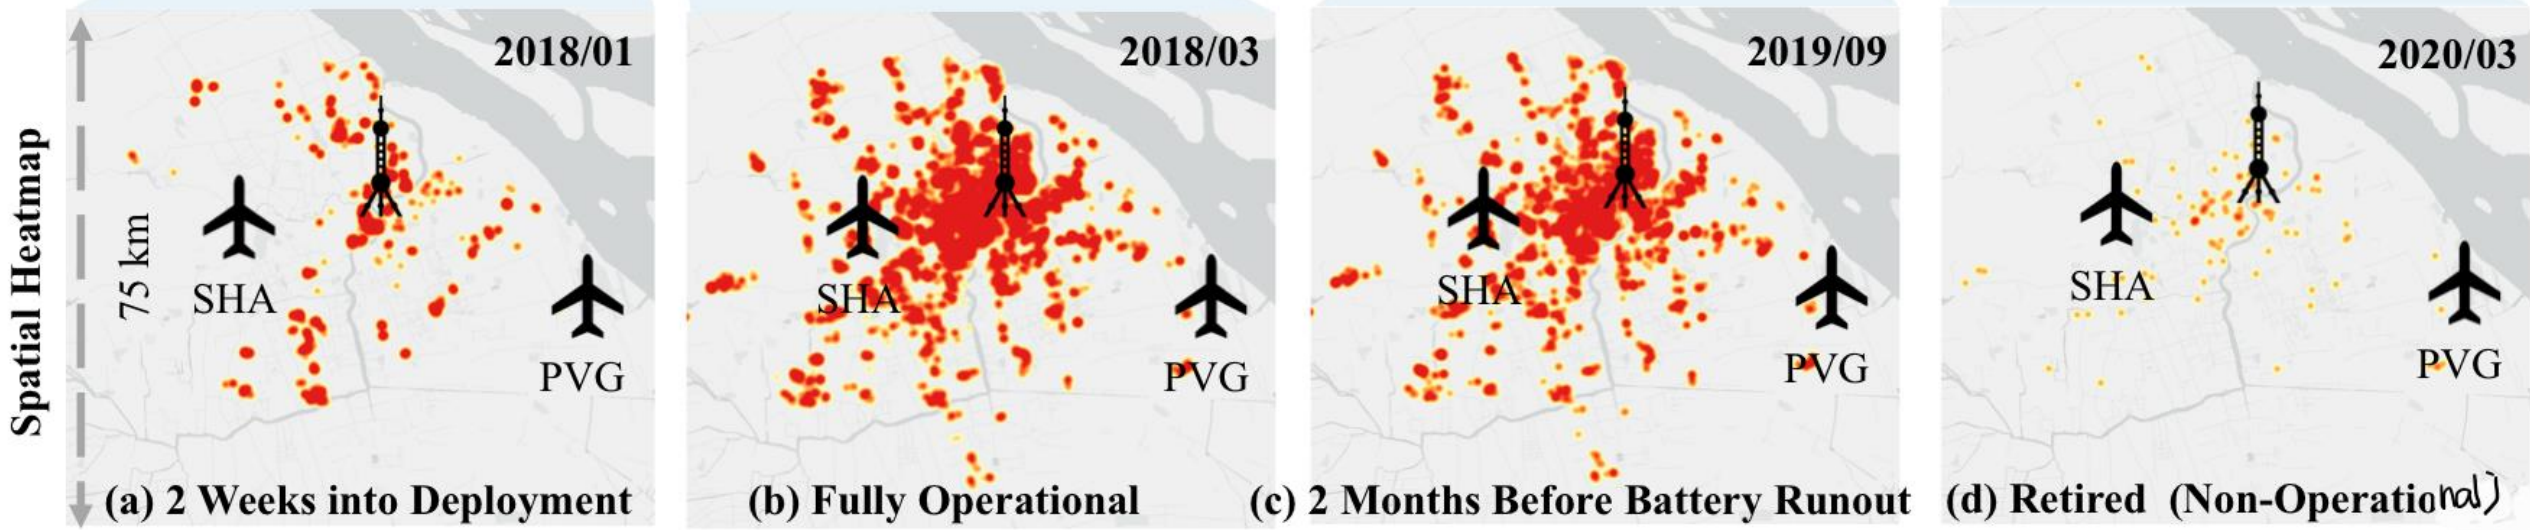

Performance: Utility (*overdue ratio reduction*)

Cumulative Gain = Cumulative Utility - Cost

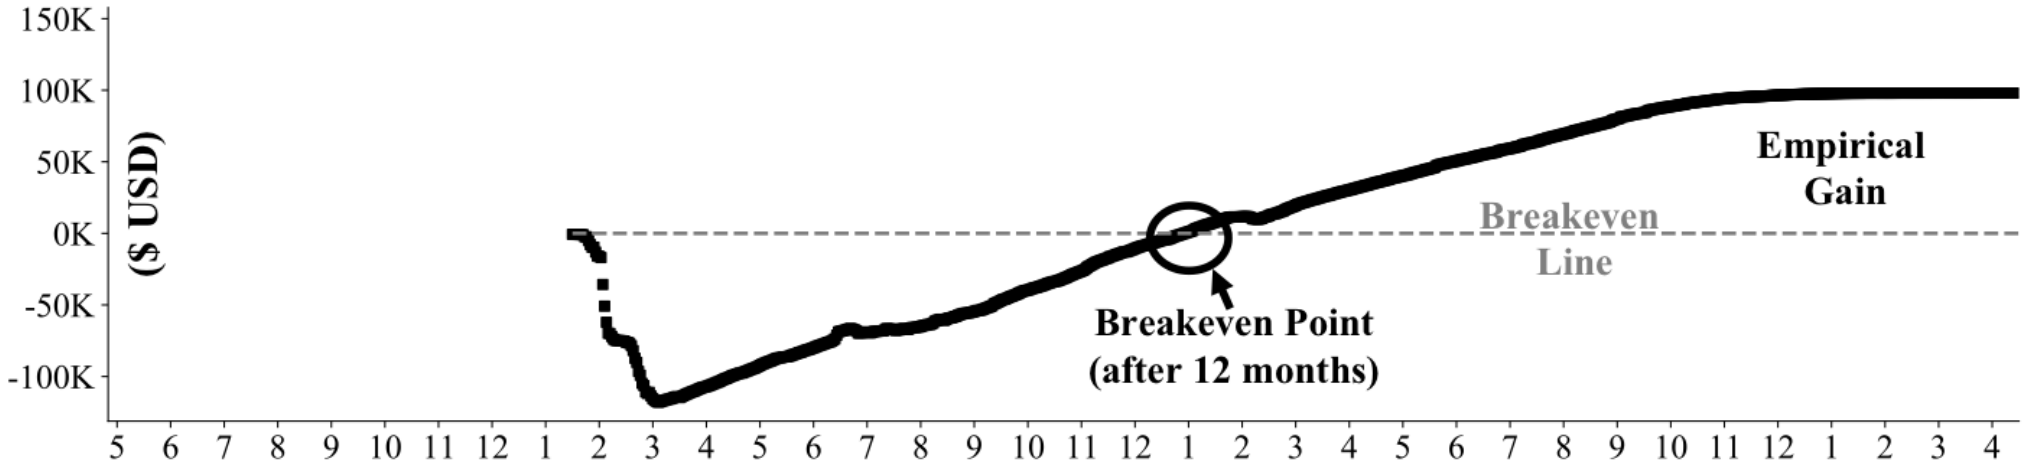

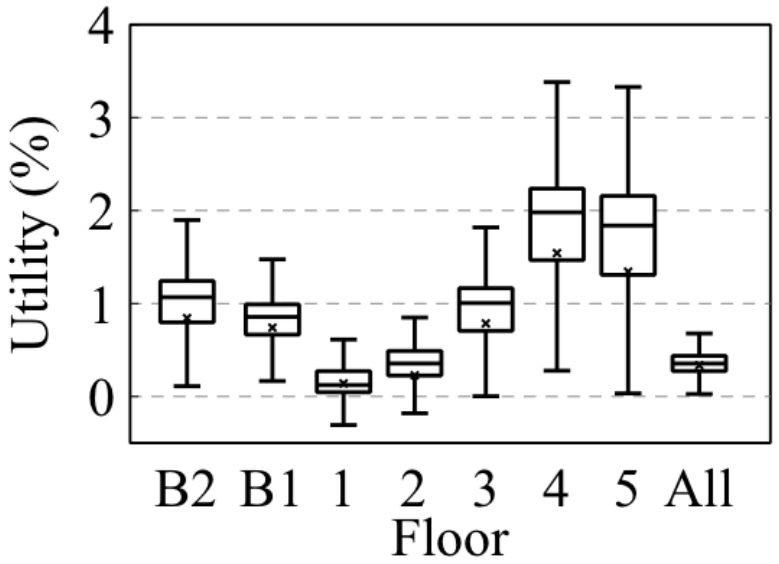

Impact of Floor

Observation: aBeacon is more beneficial in ***higher floors*** *and* ***basements***

Impact of area: Densely populated areas have more utility.

Lesson: More orders might not imply more utility, it's the uncertainty of courier behaviors that lets aBeacon shine => more utility.

Performance: Reliability

(how many arrival events can be detected among all events?)

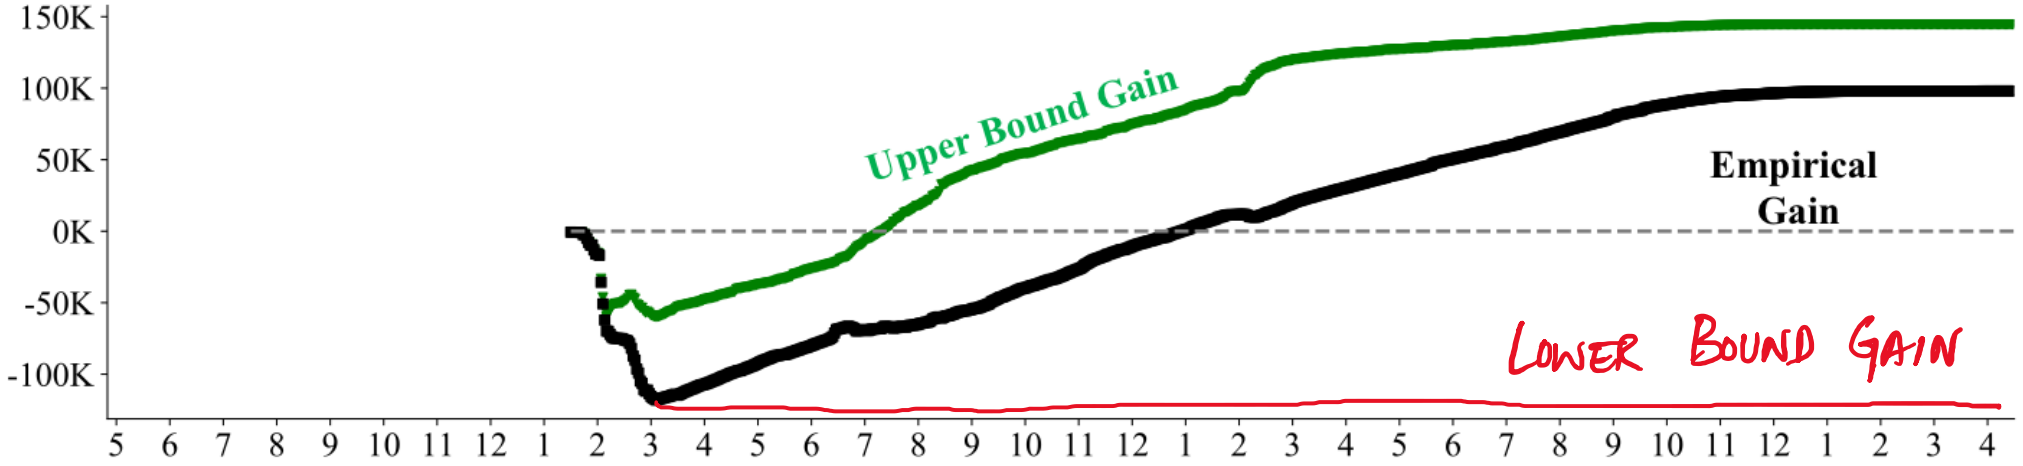

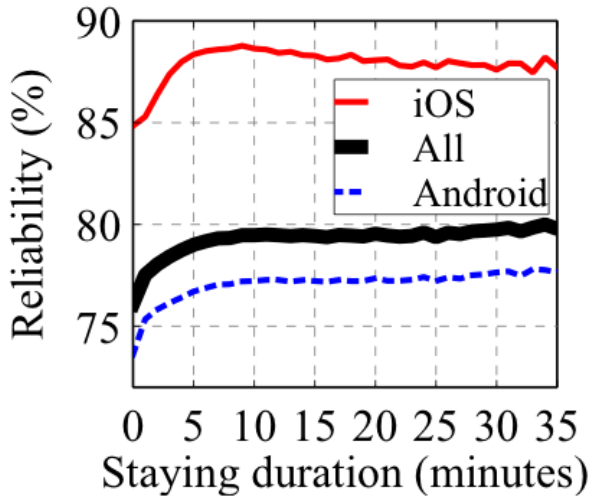

Impact of Staying Duration

Android Reliability	iOS Reliability
75%	85%

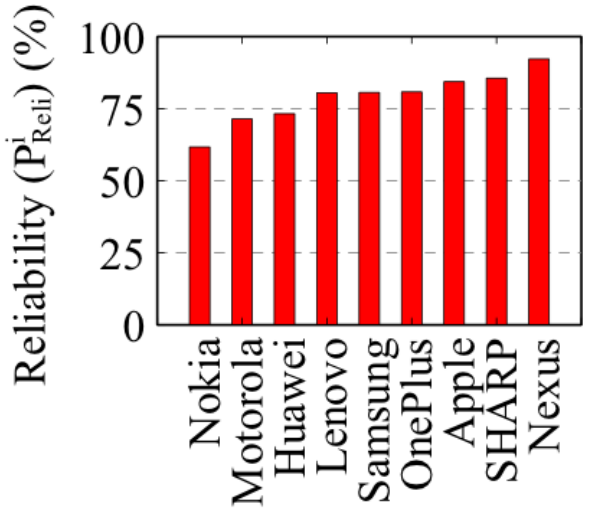

Impact of Smartphone Hardware

Performance: Reliability

(*how many arrival events can be detected among all events?*)

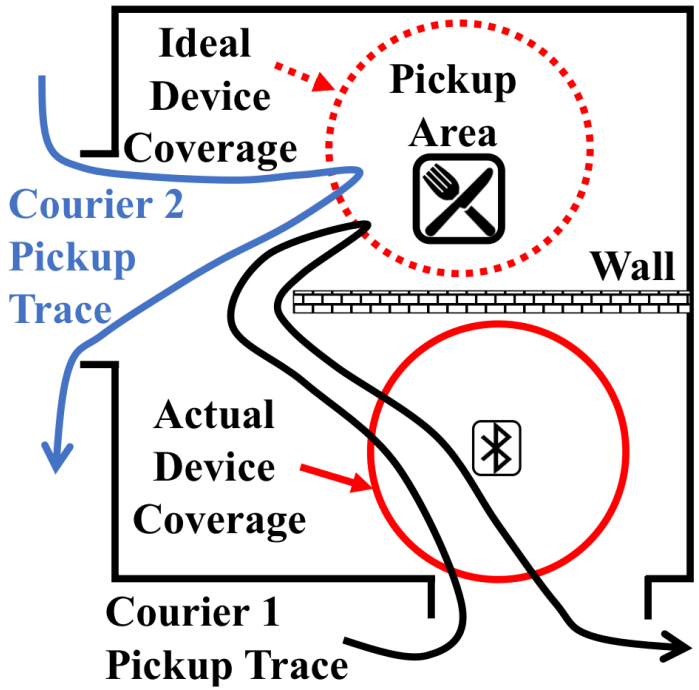

$P_{Reli} = 46\%$

Impact of device placement

Performance: Reliability

(*how many arrival events can be detected among all events?*)

Lesson Learned: Reliability in the Wild

Even *for arrival detection, the reliability is far from guaranteed in the wild due to multiple factors.*

Performance: Lifetime (*the lifetime of each device*)

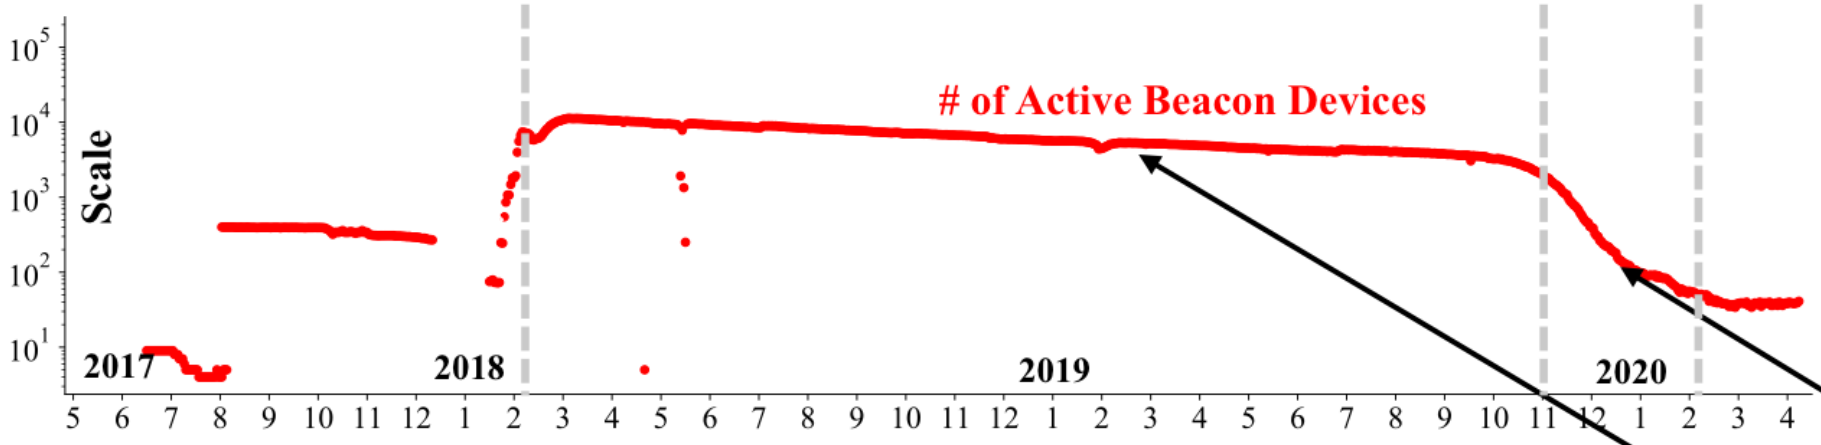

Lesson Learned: Lifetime in the Wild

- *Battery may NOT be the major constraint for mobile/wireless devices, since 40% devices survive longer than the environment (i.e., shops).*

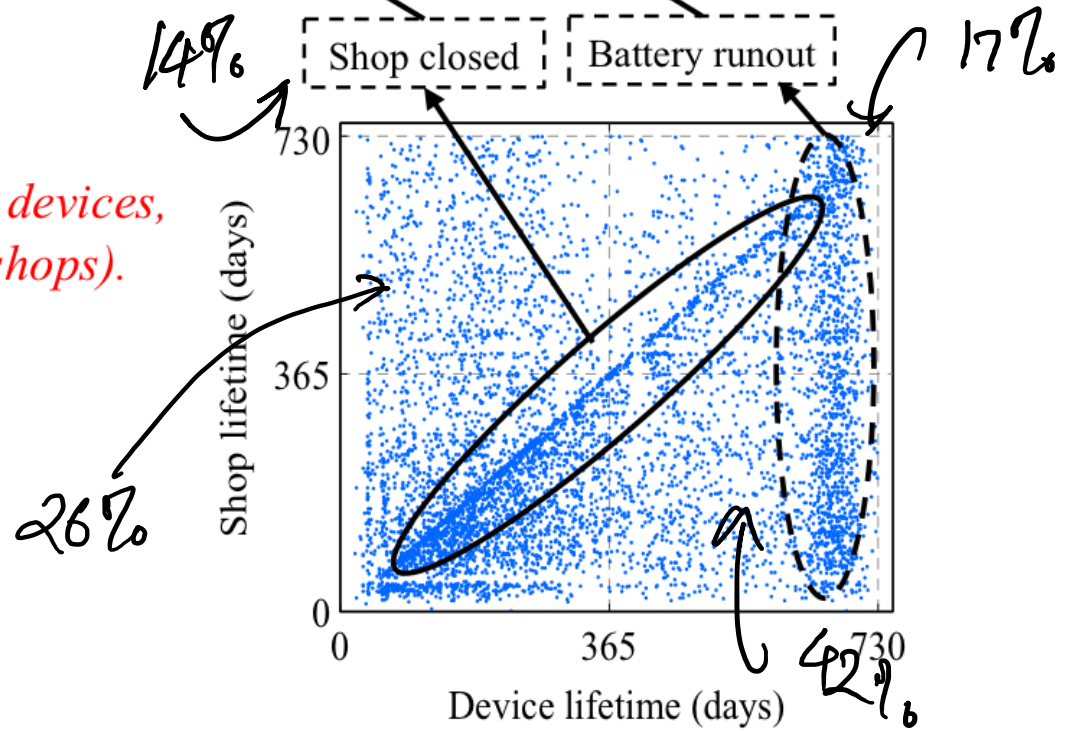

Implication for Building Industrial Systems

	System Evolution	Reliability	Lifetime
Lessons	Physical devices fail earlier than expected.	Wireless beacon devices are NOT reliable (for regulation).	Device lifetime are significantly affected by the environment.
Implication	Adopt existing devices.	Hybrid solutions (BLE+GPS+Manual Report)	Adaptive battery design.

Additional Application of aBeacons

- Order delivery time estimation
- Merchant location correction
- Anomaly detection

Next Generation of aBeacon: aBeacon+

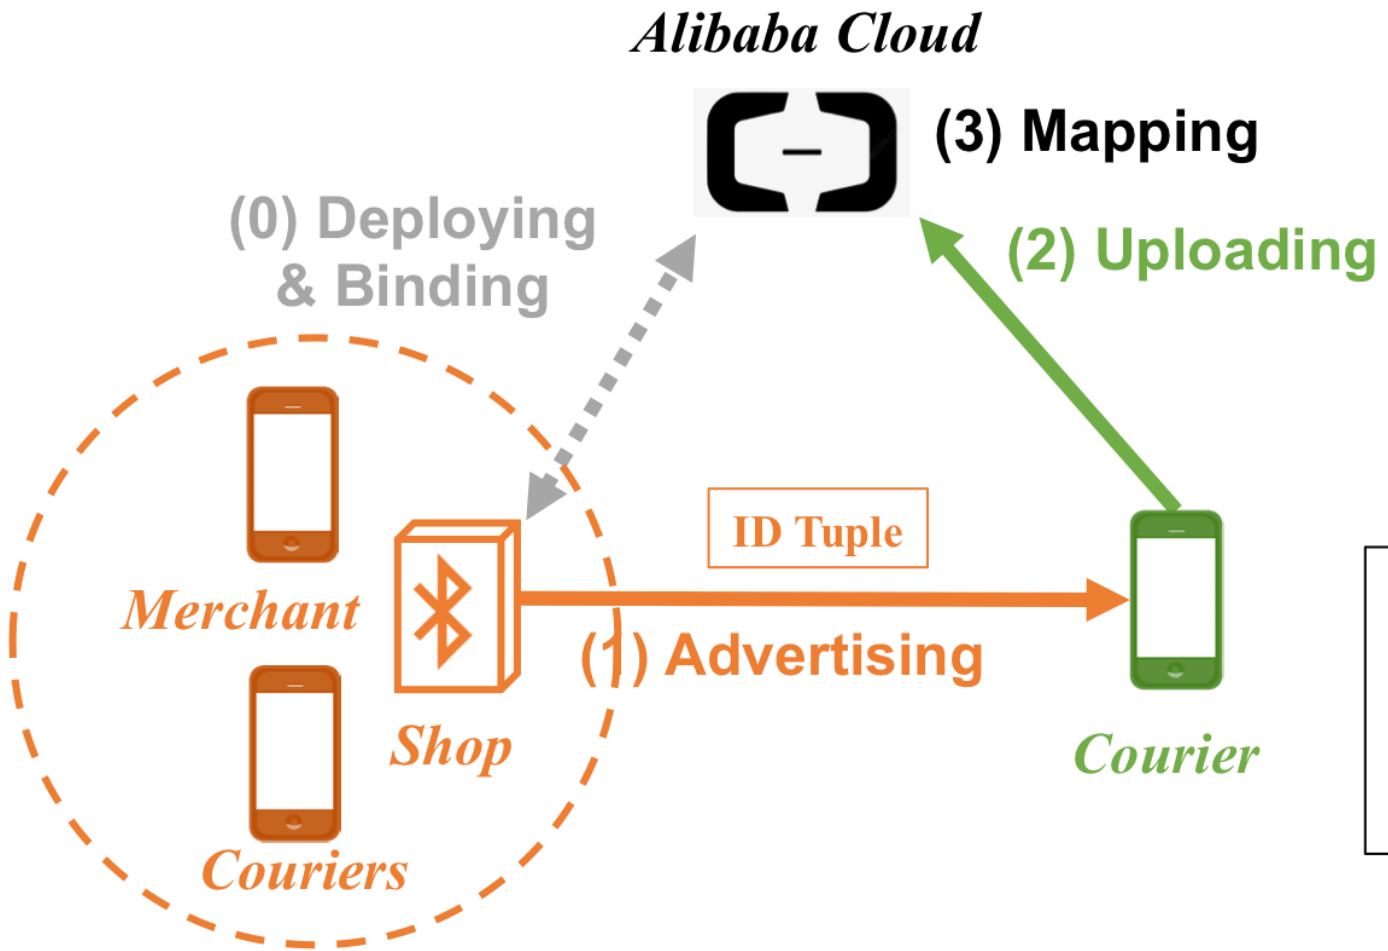

- No hardware or deployment cost.
- No lifetime worries.
- No battery worries.
- Hybrid solution.

Thank You

- Metric based evaluation of networked systems to determine gain is a takeaway

- Why does aBacon+ require batch rollouts?