CLASS

COORDINATING EDGE AND CLOUD
FOR BIG DATA ANALYTICS

When big data meets smart cities...

...how to master them?

Chongquing, August 28th, 2019

Paolo Burgio, UNIMORE

The CLASS project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 780622

Modena Automotive Smart Area

1 Km^2 urban area with connectivity enables IoT devices to exchange information

- Hundreds of smart cameras
- Traffic scanners
- Three highly-connected cars
  - UNIMORE's EU projects *Class & Prystine*
  - *Vehicle-to-infrastructure* (V2I), *vehicle-to-cloud* (V2C), *vehicle-to-vehicle* (V2V)
  - Cameras @4K, long-range and middle range radars and ultrasound sensors

V2V, V2I, V2C connectivity

Real value for the city: (Big) data analytics

Build a real-time dataset for *Smart City Distributed Awareness*

- Traffic/mobility analysis and prediction
- Critical scenarios identification
- Road user behaviour understanding and prediction
  - Improved algorithms for autonomous driving
- Data anonymization for GDPR compliance
  - Implemented at source/edge level

Distributed real-time urban awareness

Infrastructure sensors

In-vehicle sensors

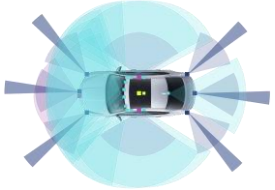

Real-time detection

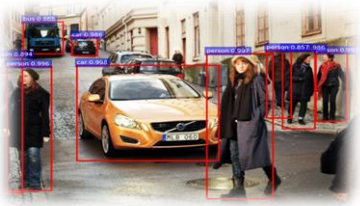

Low-latency V2X communication

Data analytics

Public authorities

Traffic enforcement

Autonomous vehicles

(Prototype) control center tools

- The presentation includes a video here
- For more information, please contact: class-project@bsc.es

HUGE credits to...

- Dott. Roberto Cavicchioli
- Prof. Nicola Capodieci
- …and all of their team


CLASS
(one possible) use-case...

Vulnerable Road User detection & avoidance

(VRUs are: other vehicles, buses, trucks, pedestrians, bicycles, etc.)

- Infrastructure cameras detect road users and obstacles
- Also, vehicle cameras detect objects!
- Sent to L3/L4 vehicles (V2I) <u>in real-time for collision avoidance</u>

The goal:
camera-to-car comm
to deliver high-criticality services
< 100ms !

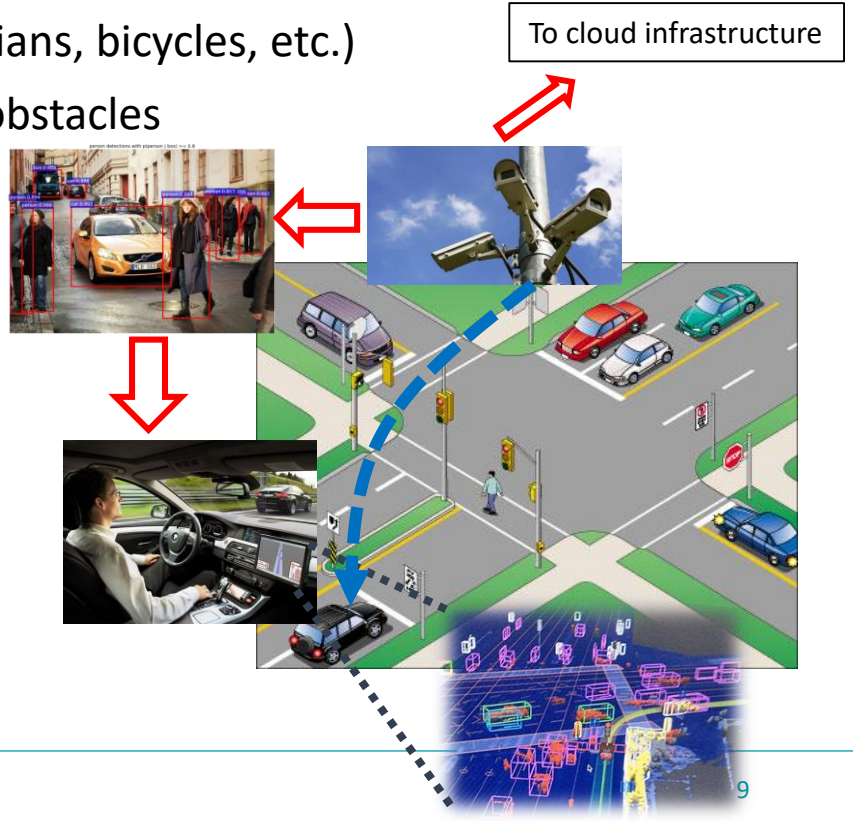

Traditional Smart City flow: high latency

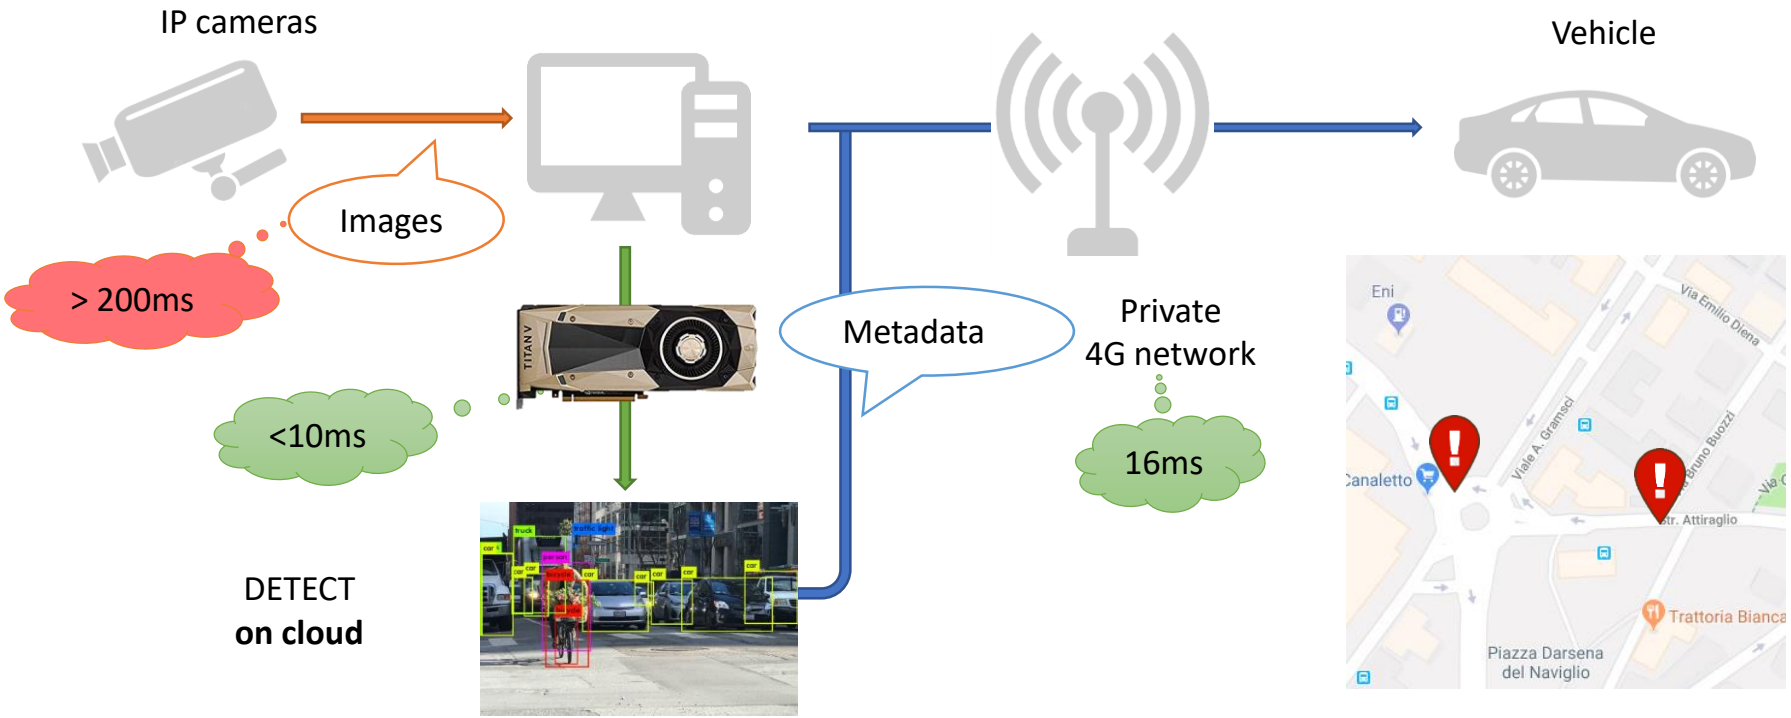

Class & MASA: fog computing in Real-Time

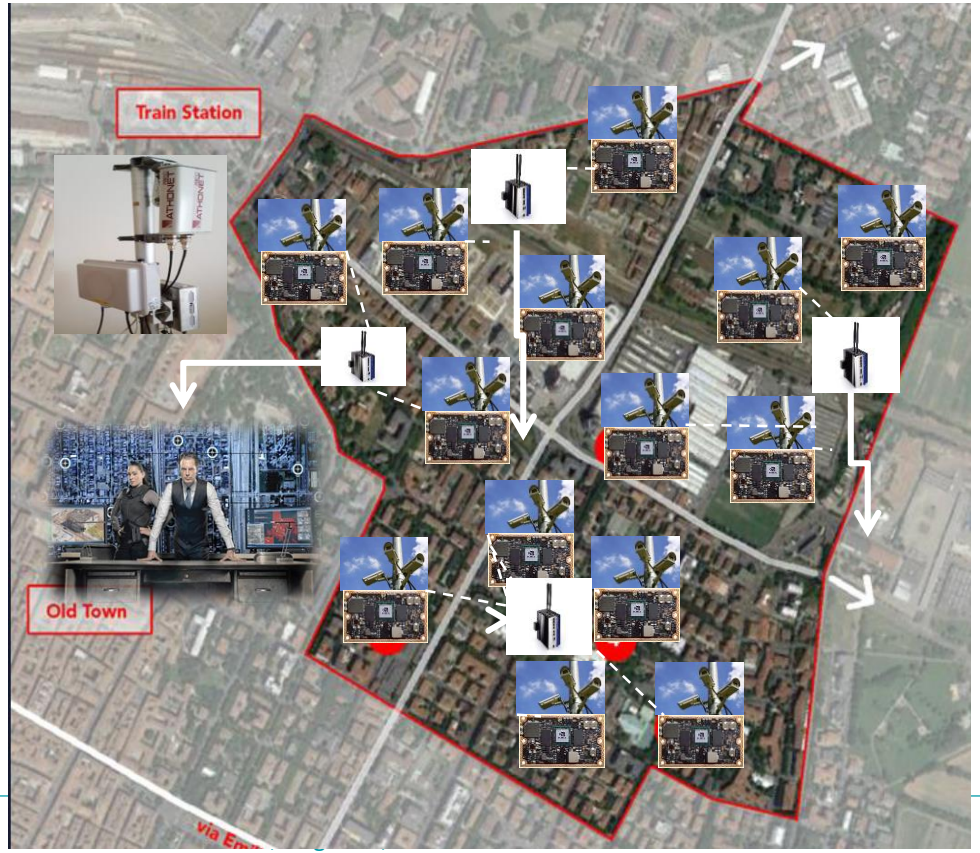

Edge nodes

- Hundreds of smart cameras detecting VRUs in real-time
- Local high-performance embedded board
- Information sent in V2I to vehicles, and in ETH to Fog nodes

Fog nodes

- Four local servers
- High-end computing system
- Additional (complex) tasks not feasible on edge

Cloud

- Fog nodes are fiber connected to **main control center**
- Data storage, learning, low-criticality

Edge + Fog + Cloud

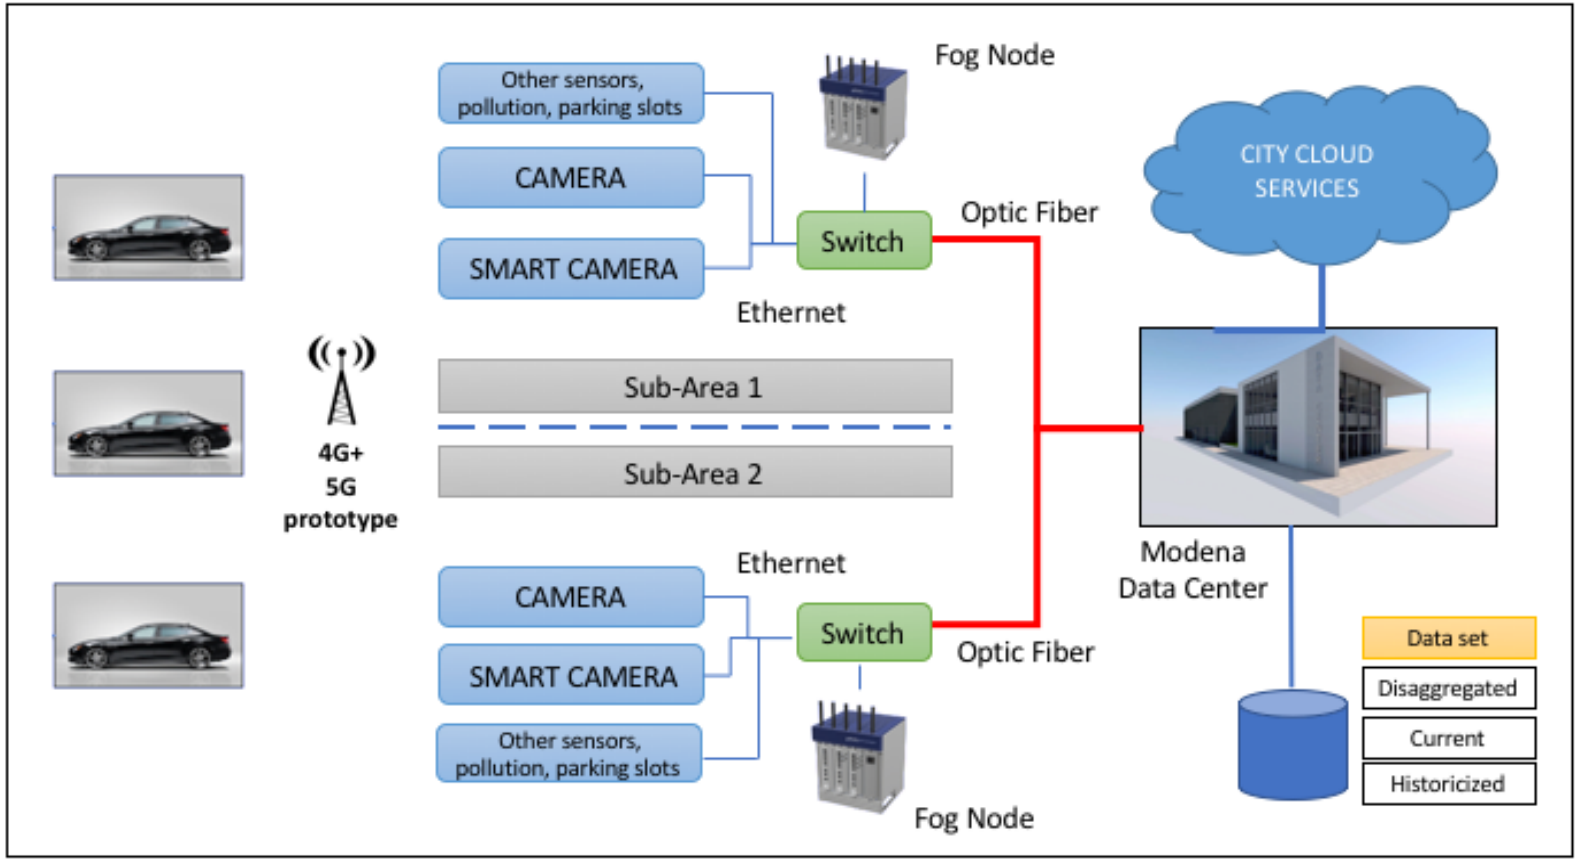

CLASS Smart City flow: **under 100 *ms***

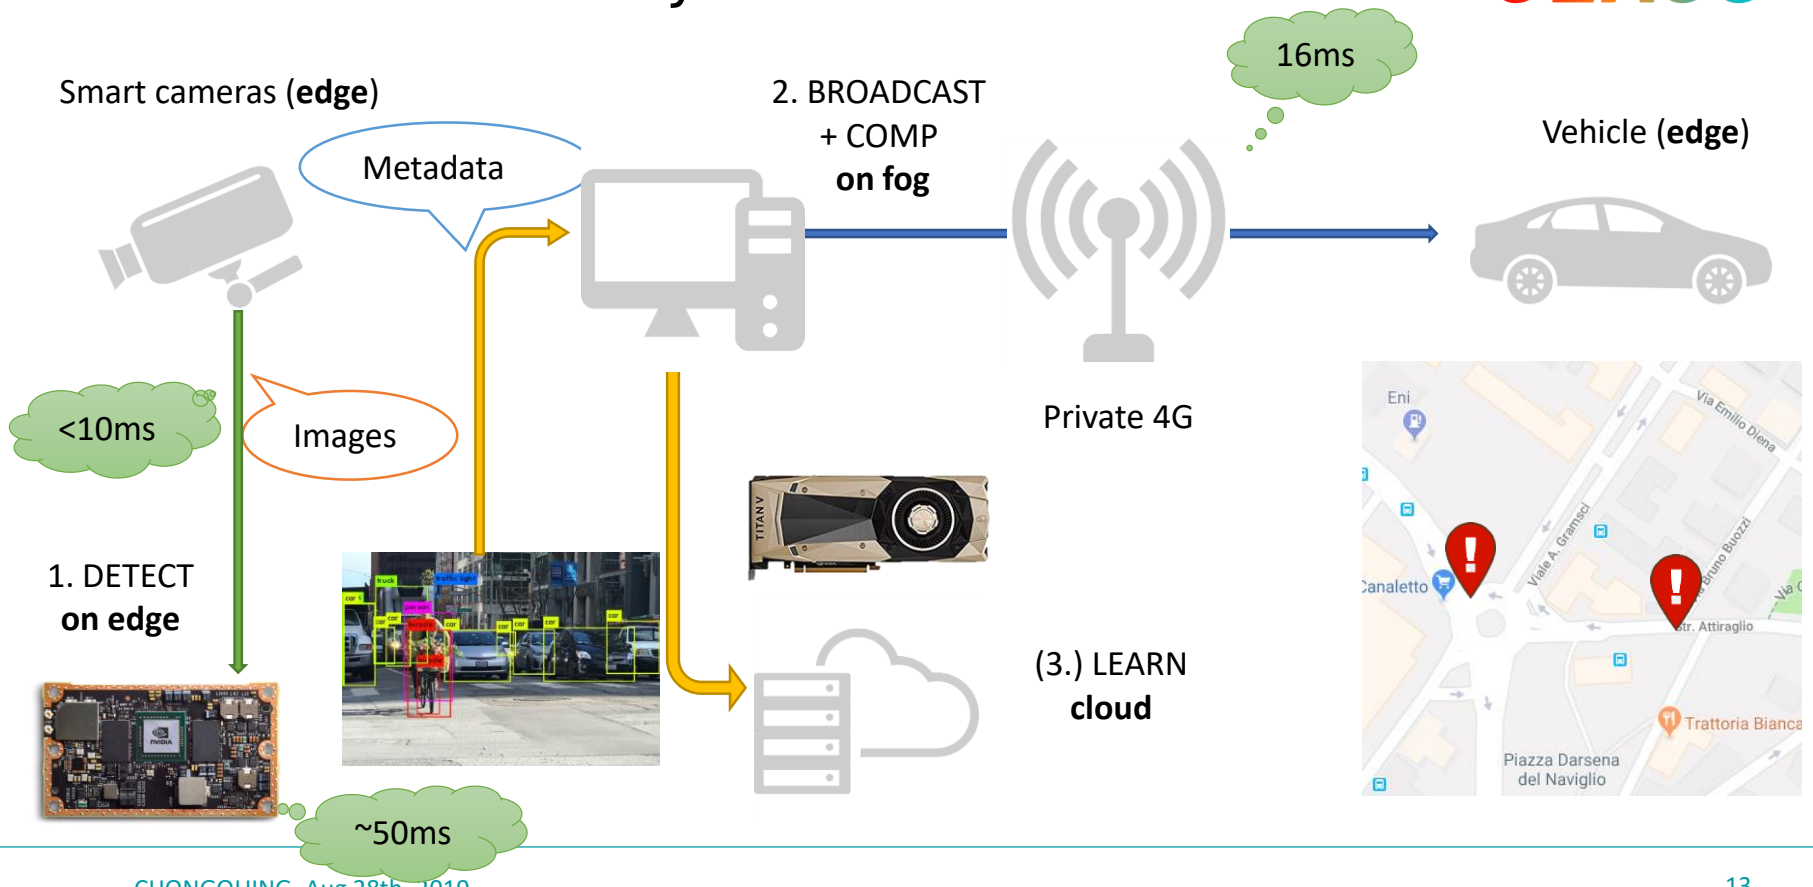

How much data?

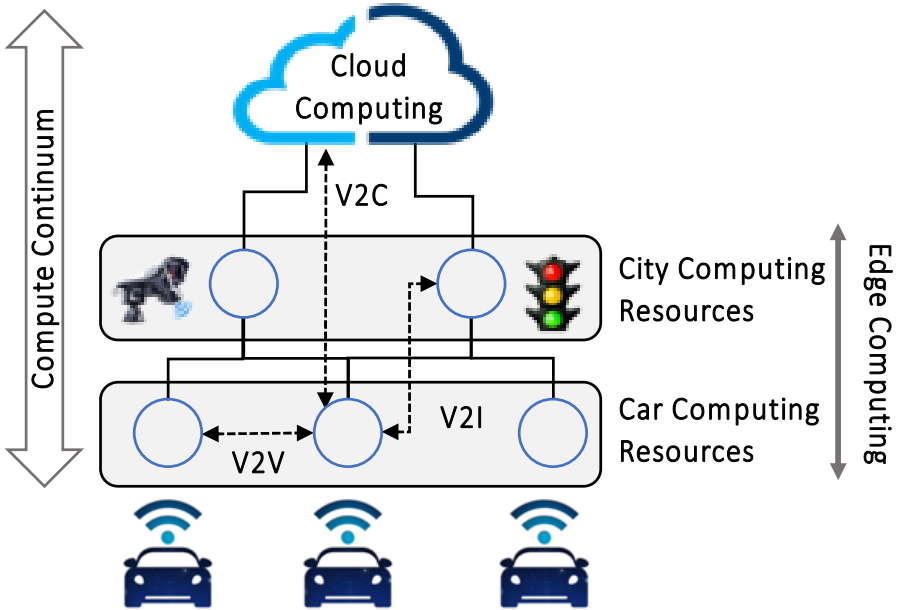

11.4 GB/s data-sets

(3 cars + 1 km2 sensing area)

Goal: services in 100ms

- Acquisition
- Processing
- Transmission


CLASS
How to master this data?

The nerd corner: CLASS SW infrastructure

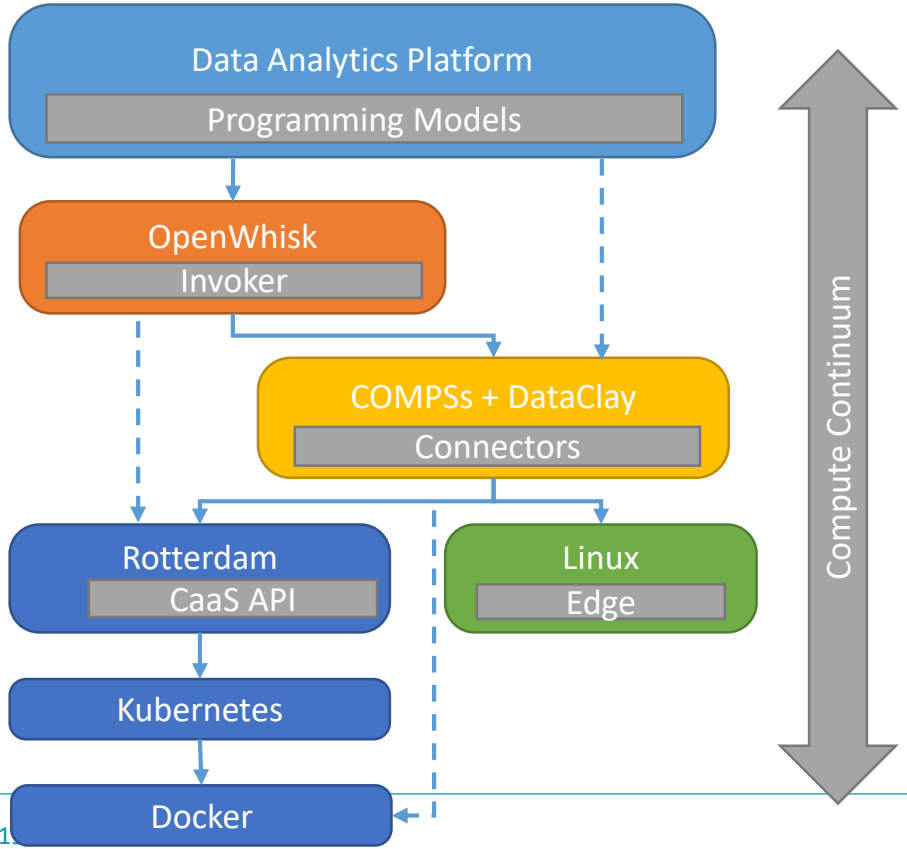

Edge computing nodes for smart cameras

NVIDIA Tegra X2 SoM

- Dev platform @~350€
- Production platform for autonomous driving @~5000€

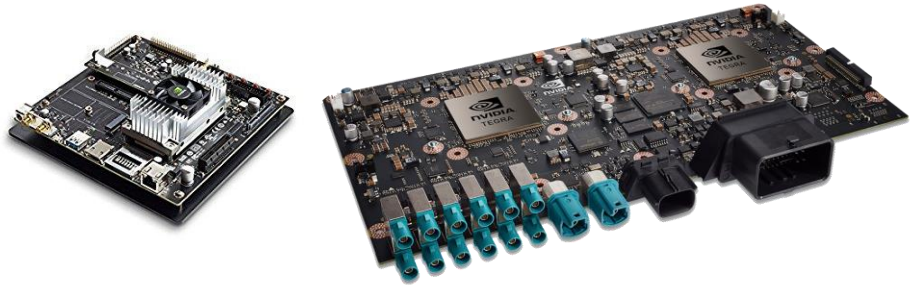

- Complete Sw dev toolchain
  - GNU/Linux
- Support by NVIDIA
- Already two generation ahead!
  - Xavier SoM, Pegasus board

Multiple trackers: view from camera

- The presentation includes a video here
- For more information, please contact: class-project@bsc.es

Position of objects

- Detection on edge
  - Bounding boxes on the camera image
- GPS conversion
  - (x, y) ➔ (GPS latitude, GPS longitude)
- Also GPS-to-meters
  - Tracking from a map point of view and not from the camera itself
  - Faster to fuse data from multiple cams

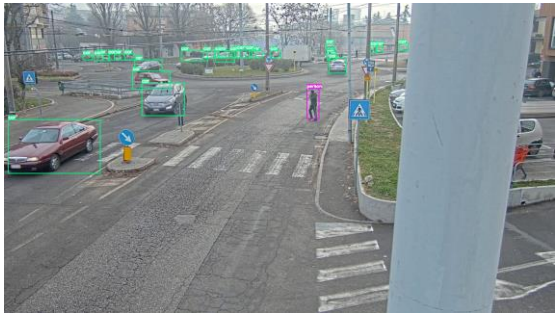

Detection performance (80 objects)

- NVIDIA's high-end Titan V (fog node): ~ 3 ms
- NVIDIAS's Tx2 (edge node): ~ 10 ms

Comparable!!

- (and, still under 100ms)


CLASS
Communication time

IP camera transmission latency

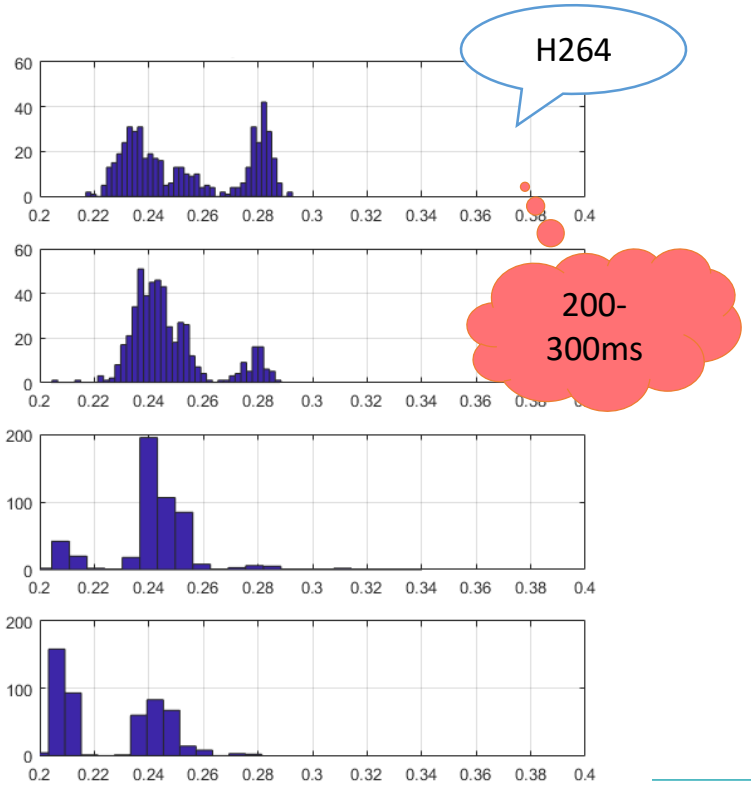

1920 x 1080

1280 x 720

768 x 432

640 x 480

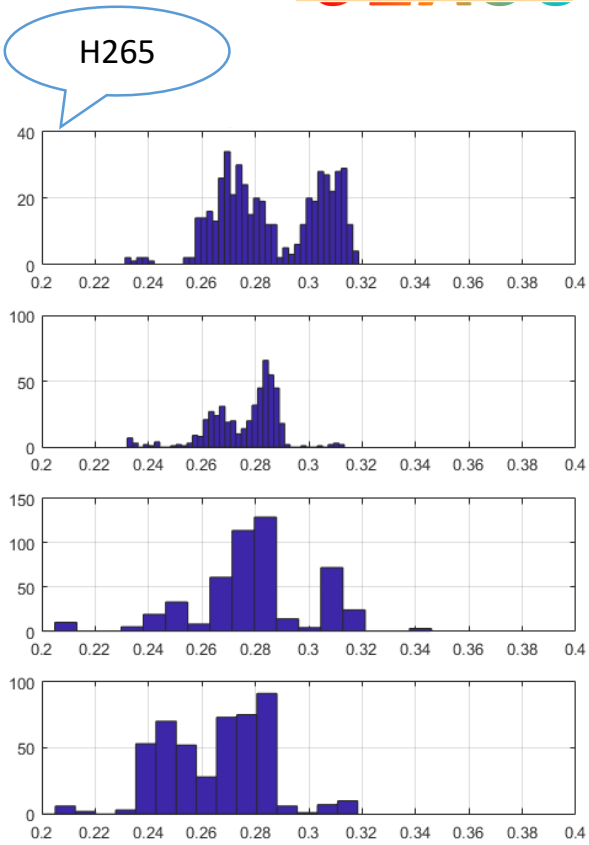

USB3 camera latency (Zed Stereo Camera)

2560 x 720

~40ms

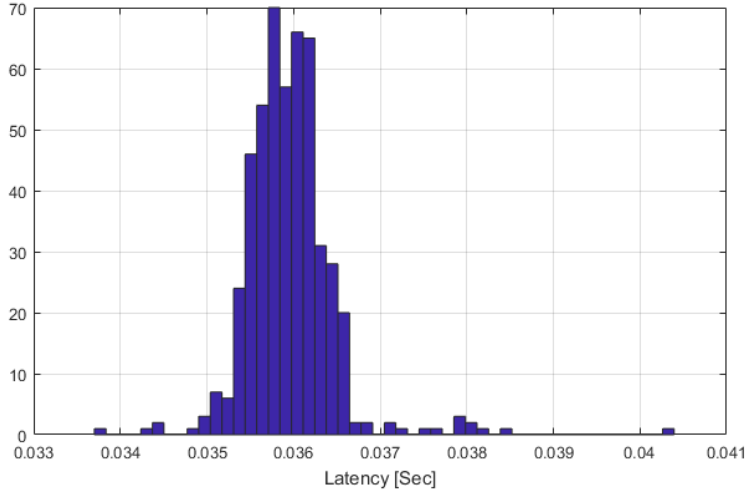

Data format and privacy-by-design

Network Connection capabilities

Edge-to-fog (cameras)

- Ethernet LAN
  - Up to **1Gbit/s**

Fog-to-cloud

- Optical Fiber
  - Up to **100Gbit/s**

Edge-to-all (V2X)

- Private **4G/LTE**
  - **16 Mbit/s**
  - <20 ms for CLASS metadata transmission
- **5G** to come soon (?)
  - Negligible overhead for CLASS metadata transmission (antenna in low-power mode!)
  - Not for single-device BW, but for multple-edge scalability!

Compact data format (aka: nerd time #2)

Standard UDP (will enhance with cryptography)

- Sender Identifier: 1B
- Timestamp: 1B
  - 1/200 sec resolution
- # objects that will be sent: 1B
- For each object:
  - **Latitude + longitude**: 8B
  - **Speed**: 4B (0 to 128 KM/h with 0,5 KM/h resolution)
  - **Orientation**: 1B (0 to 360° with ~2 degrees resol.)
  - **Category identifier**: 1B (people, dog, vehicle, bike…256 categories)

Around half KB for 50 objects sent

Data Anonymization

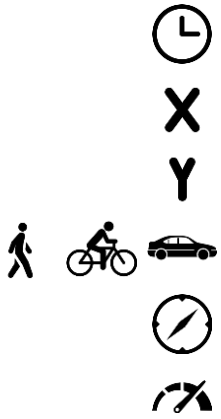

time
latitude
longitude
classification
orientation
speed

WHAT THE NUMBER OF DIGITS IN YOUR COORDINATES MEANS

LAT/LON PRECISION	MEANING
28°N, 80°W	YOU'RE PROBABLY DOING SOMETHING SPACE-RELATED
28.5°N, 80.6°W	YOU'RE POINTING OUT A SPECIFIC CITY
28.52°N, 80.68°W	YOU'RE POINTING OUT A NEIGHBORHOOD
28.523°N, 80.683°W	YOU'RE POINTING OUT A SPECIFIC SUBURBAN CUL-DE-SAC
28.5234°N, 80.6830°W	YOU'RE POINTING TO A PARTICULAR CORNER OF A HOUSE
28.52345°N, 80.68309°W	YOU'RE POINTING TO A SPECIFIC PERSON IN A ROOM. BUT SINCE YOU DIDN'T INCLUDE DATUM INFORMATION, WE CAN'T TELL WHO
28.5234571°N, 80.6830941°W	YOU'RE POINTING TO WALDO ON A PAGE
28.523457182°N, 80.683094159°W	"HEY, CHECK OUT THIS SPECIFIC SAND GRAIN!"
28.523457182818284°N, 80.683094159265358°W	EITHER YOU'RE HANDING OUT RAW FLOATING POINT VARIABLES, OR YOU'VE BUILT A DATABASE TO TRACK INDIVIDUAL ATOMS. IN EITHER CASE, PLEASE STOP.

https://git.hipert.unimore.it/rcavicchioli/masa_protocol

Thanks for your attention.

Stay tuned!

www.class-project.eu
Twitter: **@EU_CLASS**
LinkedIn: **http://bit.ly/CLASS-project**