



November, 2022

Data processing acceleration in Traffic modelling

EBDVF 2022

RADIM CMAR (SYG)

EVEREST project manager

rcmar@sygic.com

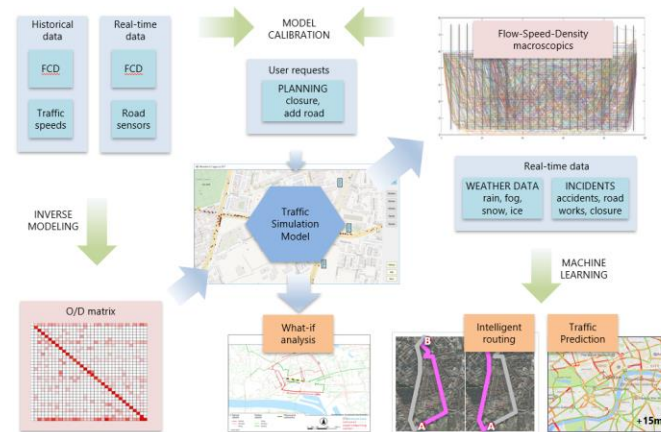
Traffic Modeling Usecase

Product: Traffic eco-system for cities

- Computation of road speed profiles over weeks
- Prediction on each road segment
- Traffic Analysis tool when closures
- Routing with collaborative intelligence

Technology challenges/optimizations

- Big data
 - FCD data: gigabytes daily
- AI computation
 - Neural Networks (NN)
 - Hidden Markov Model (HMM)




Why Traffic modelling

- Urban traffic still on rise – 75% of population in cities
- City congestion costs 1% of GDP, accounts for 40% of CO2
- If city wants to improve traffic – needs analytics & prediction

Traffic modeling – Datasets & Processing

Floating car data (FCD)



Sygić



Origin-Destination matrix (ODM)



Telekom



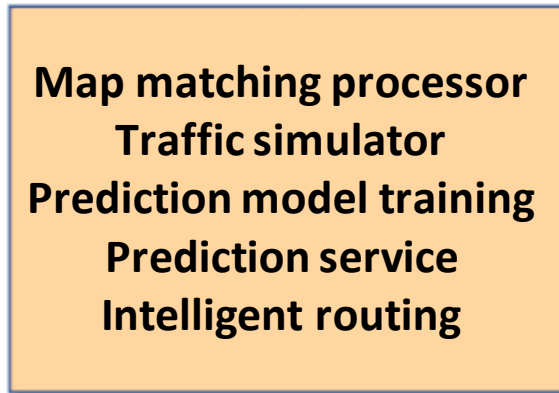
Road network




Sygić




Traffic computation eco-system



Speed profiles



3D traffic model




Analyst
Planner



Traffic services

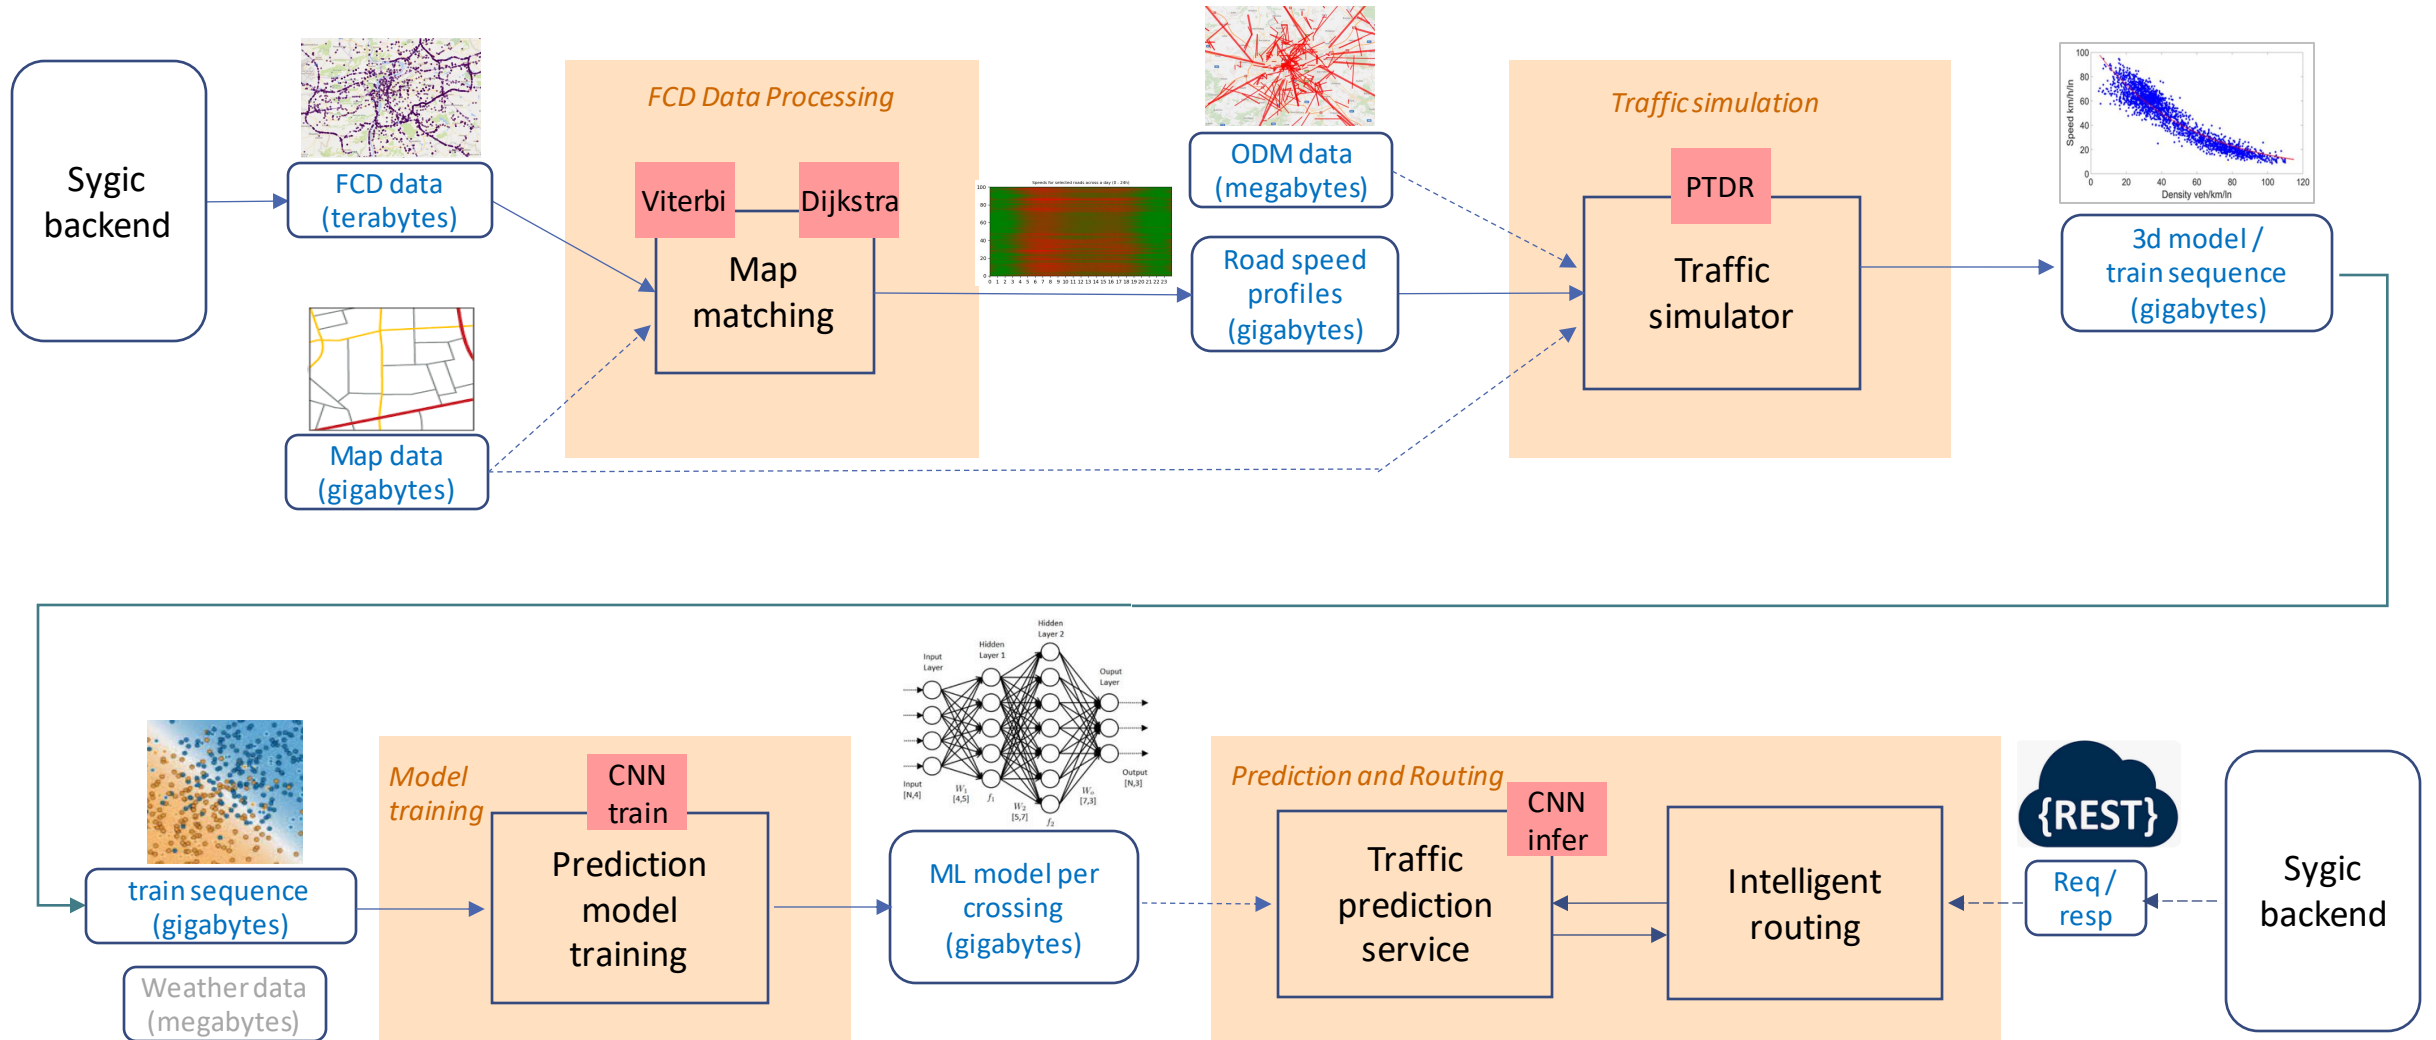


Traffic prediction
Intelligent routing




Training data
Prediction model

Architecture of the System



Designing with EVEREST SDK (1)



Data aspect

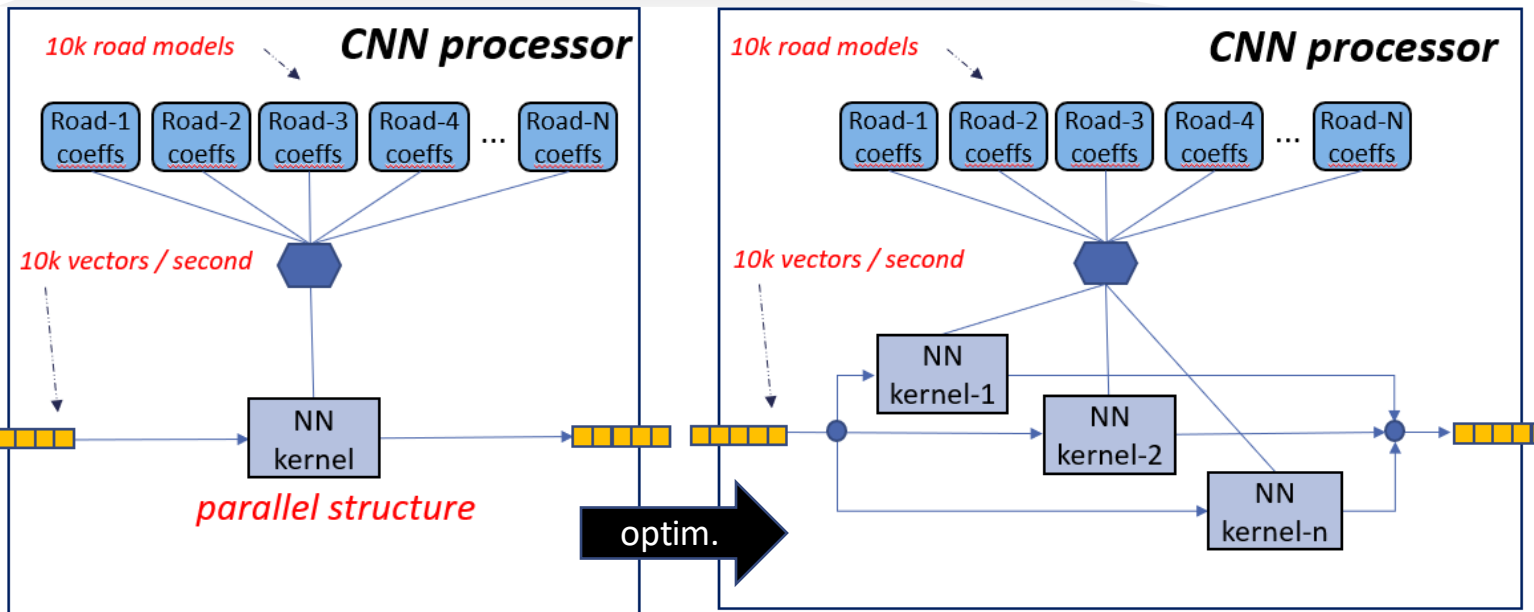
- Vector in / vector out
- Each vector - neural network processing
- Real-time client requests

Benchmarking

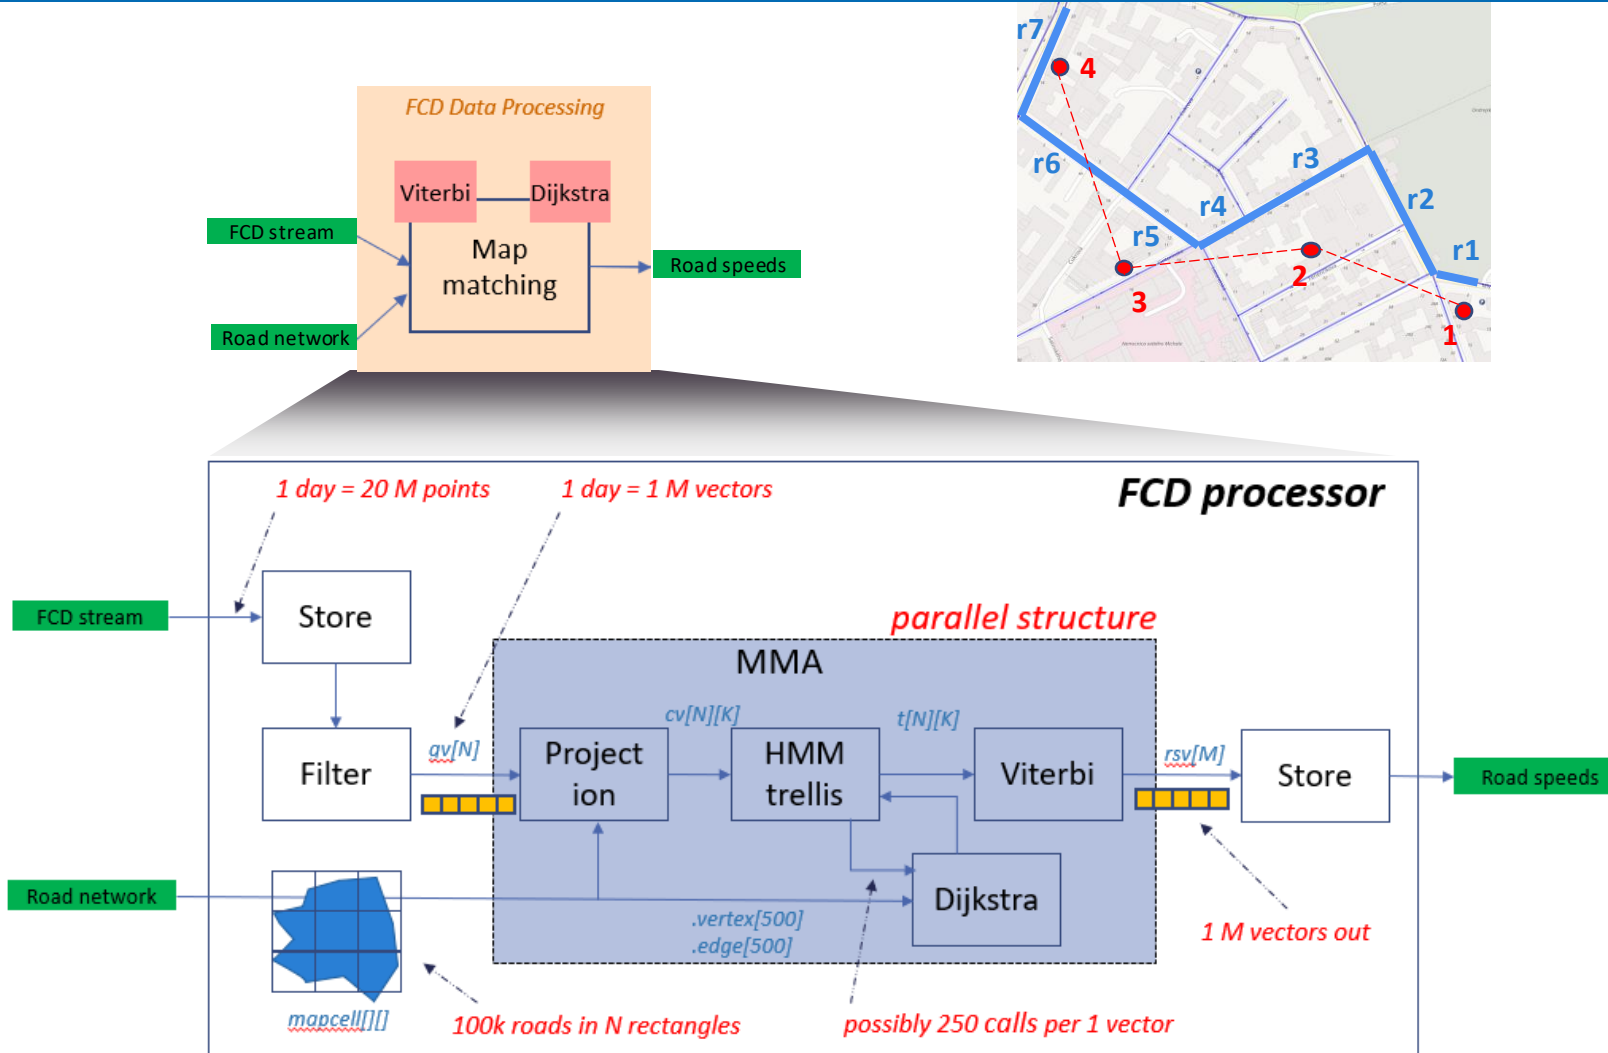
- 100k roads/vectors ~ 1 city
- Intel i5 2.3GHz 4 cores ~ 9s
- Goal: > 10x speed up

Optimization with EVEREST

- Python transformation and DF optim.
- Kernel parallelization
- Kernel mapping to FPGA(s)
- Memory access optimization (coeffs)



Designing with EVEREST SDK (2)



Data aspect

- Vector in / vector out
- with each vector - complex processing
- Batch processing

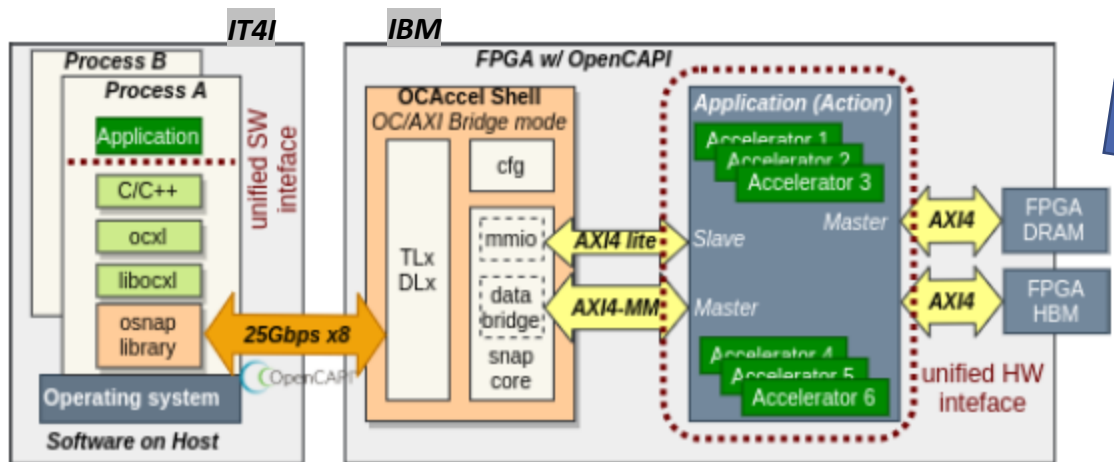
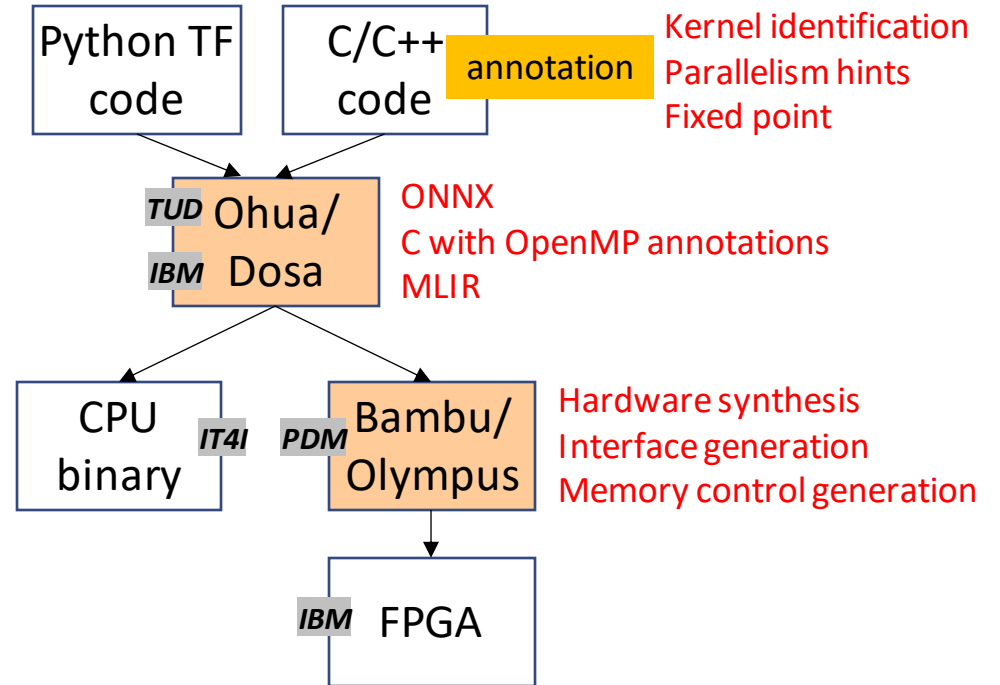
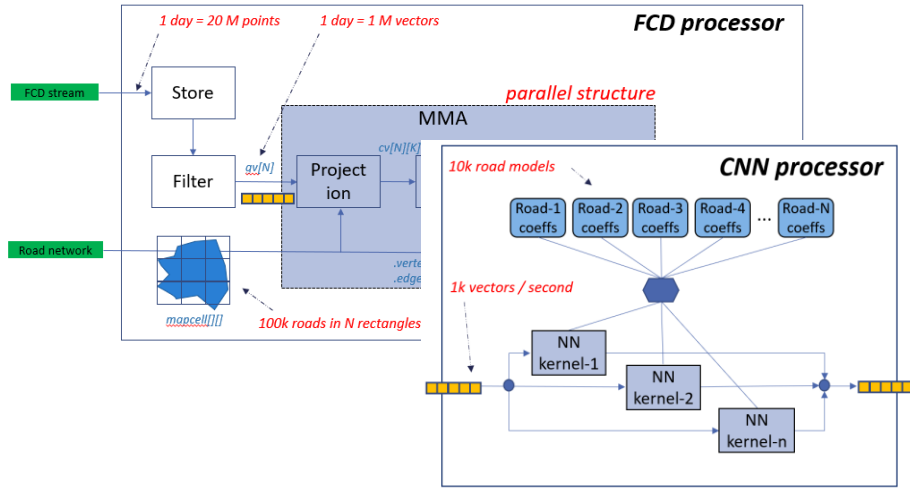
Benchmarking

- 20 M points ~ 1 city
- Intel i5 2.3GHz 4 cores ~ **4:35 h**
- Goal: 5x speed up

Optimization with EVEREST

- C/C++ code transformation and optim.
- Kernel parallelization
- Kernel synthesis to FPGA(s)
- Memory access optimization (Map cells)


EVEREST tooling



Consortium



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 957269



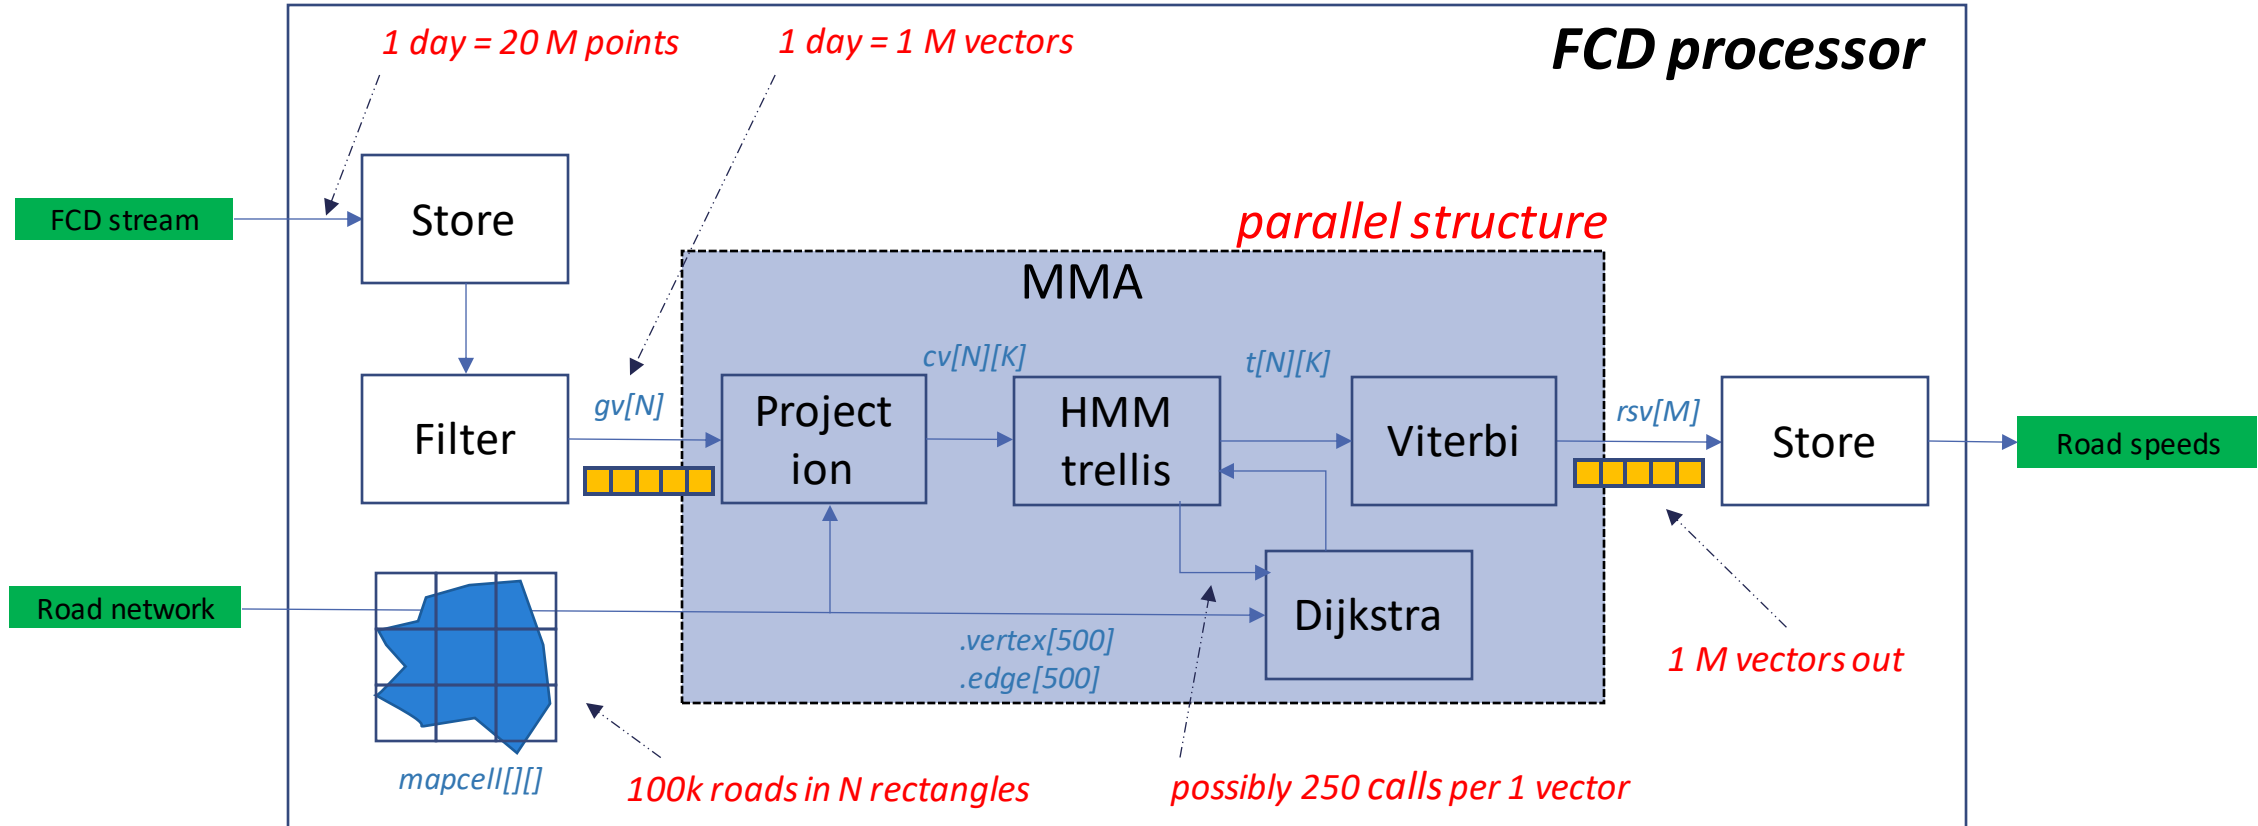
Thank you !

Backup

Final word

- Big data processing requires acceleration of kernels
 - dataflow operations, ML algorithms
- EVEREST = SDK to accelerate generic code
 - Supporting C/C++, Python, Rust
 - use of IBM FPGA and IT4I HPC
 - Accelerations > 10x

SOFTWARE ARCHITECTURE



HARDWARE ARCHITECTURE

