

DeepDish on a diet: low-latency, energy-efficient object-detection and tracking at the edge


Matthew Danish Rohit Verma Justas Brazauskas
Ian Lewis Richard Mortier

Cambridge University

EdgeSys 2022: Rennes, France


Motivation

- ▶ Ubiquity, e.g.
 - ▶ Traffic sensors
 - ▶ Movement studies
 - ▶ Public safety
- ▶ Maintaining privacy
- ▶ Power draw



Motivation

- ▶ Ubiquity, e.g.
 - ▶ Traffic sensors
 - ▶ Movement studies
 - ▶ Public safety
- ▶ Maintaining privacy
- ▶ Power draw





© Statista 2022

Price of electricity 2015–2021




DeepDish

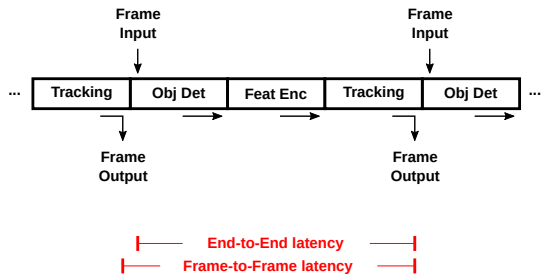
- ▶ Raspberry Pi 4B
- ▶ Coral EdgeTPU
- ▶ TensorFlow-Lite
- ▶ (also full GPU support)




Tracking-by-detection




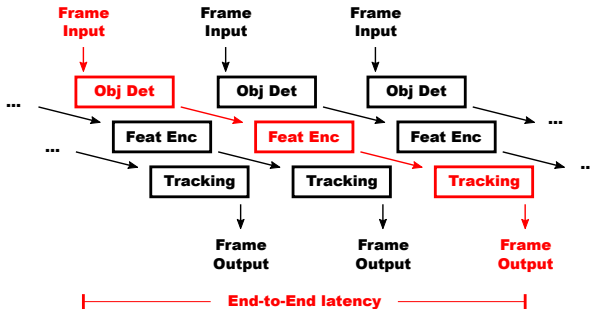
No pipelining



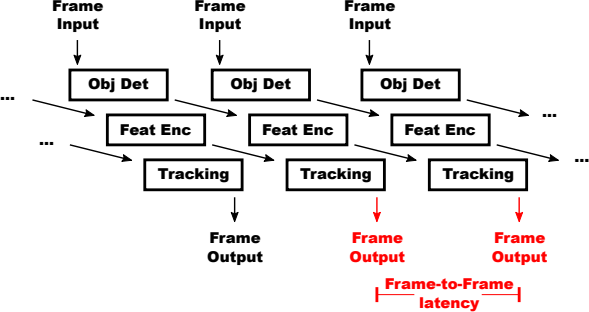
Pipelining



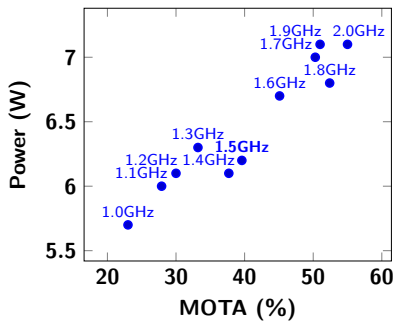
Pipelining



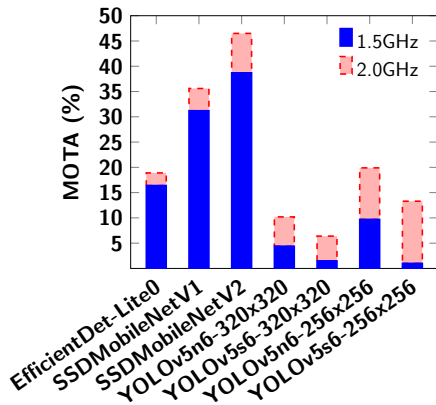
Pipelining



CPU clockspeed and power




All tested videos: power/accuracy trade-off




Video 'TUD-Campus': model, clockspeed and accuracy

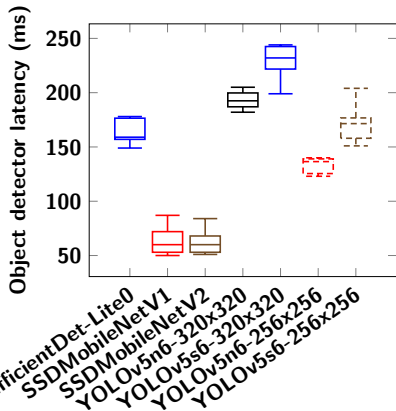

EdgeTPU effect



EdgeTPU effect



Model choice



Looking forward

- ▶ 6-7W power usage: low but not low enough
- ▶ SSDMobileNetV2 still holding its own
- ▶ Object Detection / Feature Encoding: separate or combined?
- ▶ LSTM: does it offer hope?

Limitations & thoughts



- ▶ Tracking-by-detection: popular but flawed
- ▶ Context, motion & comprehension
- ▶ Ubiquity & privacy

Thank you!

Questions?

Acknowledgements

This research forms part of Centre for Digital Built Britain's work within the Construction Innovation Hub. The funding was provided through the Government's modern industrial strategy by Innovate UK, part of UK Research and Innovation. This work is funded in part by EPSRC grant EP/T022493/1.