

data processing in space

Pinar Tözün

Associate Professor
IT University of Copenhagen

this talk is not about ...

Google Research

Who we are ▾

Research areas ▾

Our work ▾

Programs & events ▾

Careers

Blog



[Home](#) > [Blog](#) >

Exploring a space-based, scalable AI infrastructure system design



November 4, 2025 · Travis Beals, Senior Director, Paradigms of Intelligence



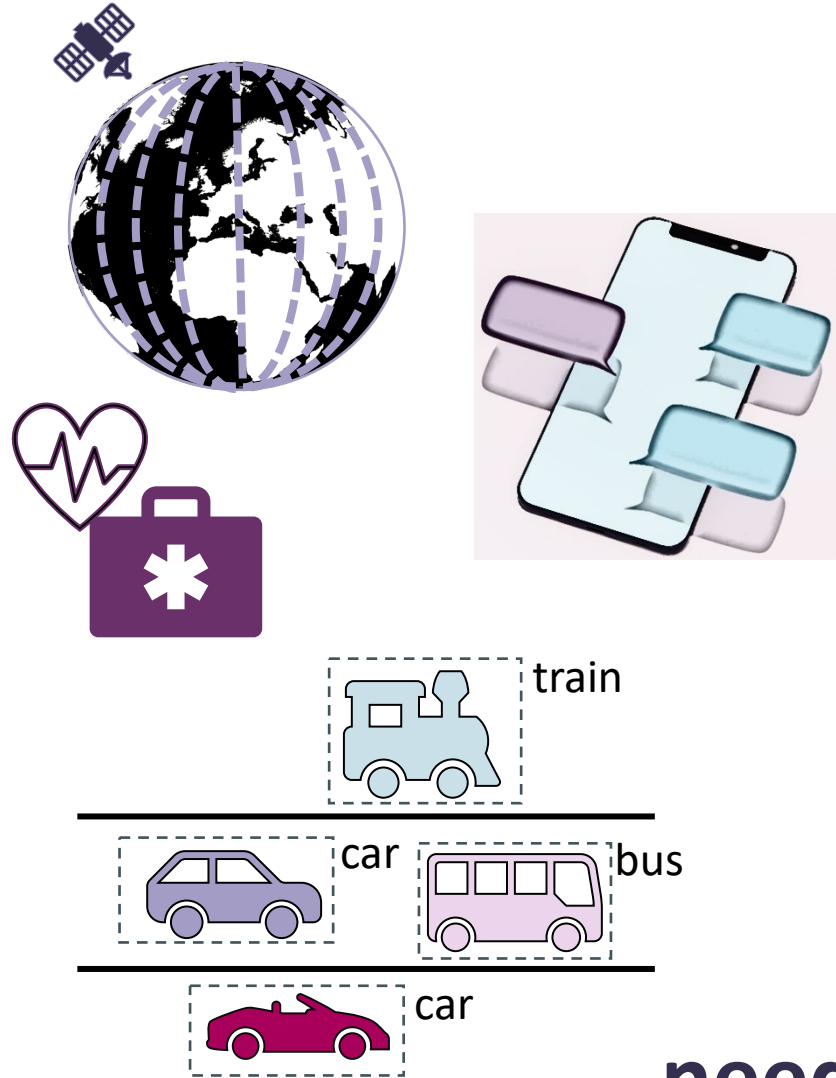
Google plans to put datacentres in space to meet demand for AI

US technology company's engineers want to exploit solar power and the falling cost of rocket launches

Robert Booth *UK technology editor*

Tue 4 Nov 2025 21.08 CET

data processing @ the edge



data source

edge



traditional-approach

- do (most) data processing in the cloud

cannot satisfy

- low-latency & real-time applications
- poor / non-existing connectivity
- legal restrictions & privacy

**need for efficient & complex data processing
closer to data sources; *at the edge!***

DISCO: Danish student CubeSat organization

collaboration across Danish universities

goal: education & research

- students learn how to build small satellites
- researchers experiment with ML models for Earth observation

use-case:


- observation of landmasses (especially snow, ice ...) in the Arctic



small satellites


benefits

\$\$



\$

reduced cost



shrinking
&
standardization


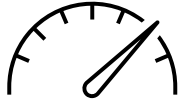



*e.g., CubeSats
10cm cube*

compromises



reduced power
generation

network on small satellites

|  |  |  |  |  |
|---|---|---|---|---|
| uhf | 9600 bit/s | low-power | low | 0.05 / day* |
| s-band | 10 Mbit/s | high-power | high | 49.1 / day* |

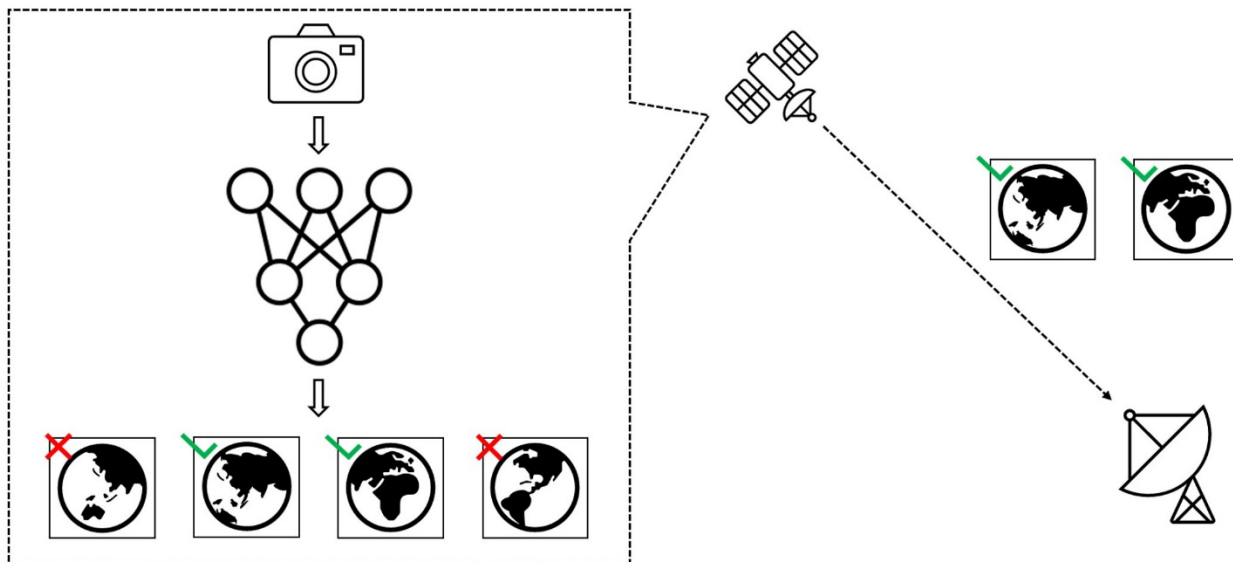
* based on 4512x4512px image
& short-lived connection

cannot send all the images to Earth!

image filtering on satellites

in DISCO: ML-based image classification

➔ **send only the relevant images to ground & minimize data movement!**



our task: build the image processing unit on the satellite
➔ **determine the edge device that can satisfy the requirements of this task!**

image processing unit requirements



real-time imaging
< 4.42 s latency



Arctic region
< 71.74 s latency



Greenland
< 270 s latency

**max 49.1 images can be
transferred per day!**

**min 320 images
captured per day.**

max 5 watts




average 2 watts

mass: 150grams

dimensions: 10x70x80mm

flexible software upload!




| ARM Cortex-M7 | Jetson Nano | Toradex Verdin | CoralAI Micro | CoralAI Mini | CoralAI USB Stick | Neural Compute Stick |
|-----------------------|----------------------|---|--|------------------------|------------------------------------|-----------------------------|
| ARM Cortex-M7 @300MHz | ARM A57 @1.43GHz | ARM Cortex-A53 @1.8GHz, ARM Cortex-M7 @800MHz | ARM Cortex-M7 @800MHz, ARM Cortex-M4 @400MHz | ARM Cortex-A35 @1.5GHz | Raspberry Pi 3 BCM2837 ARM @1.2GHz | |
| 384KB SRAM, 32KB FRAM | 4GB | 4GB | 64MB | 2GB | 1GB | |
| none | 128-core Maxwell GPU | NPU (2.25 TOPS) | CoralAI Edge TPU (4 TOPS) | | | Intel Movidius Myriad X VPU |

general-purpose

higher specialization for ML

hardware setup


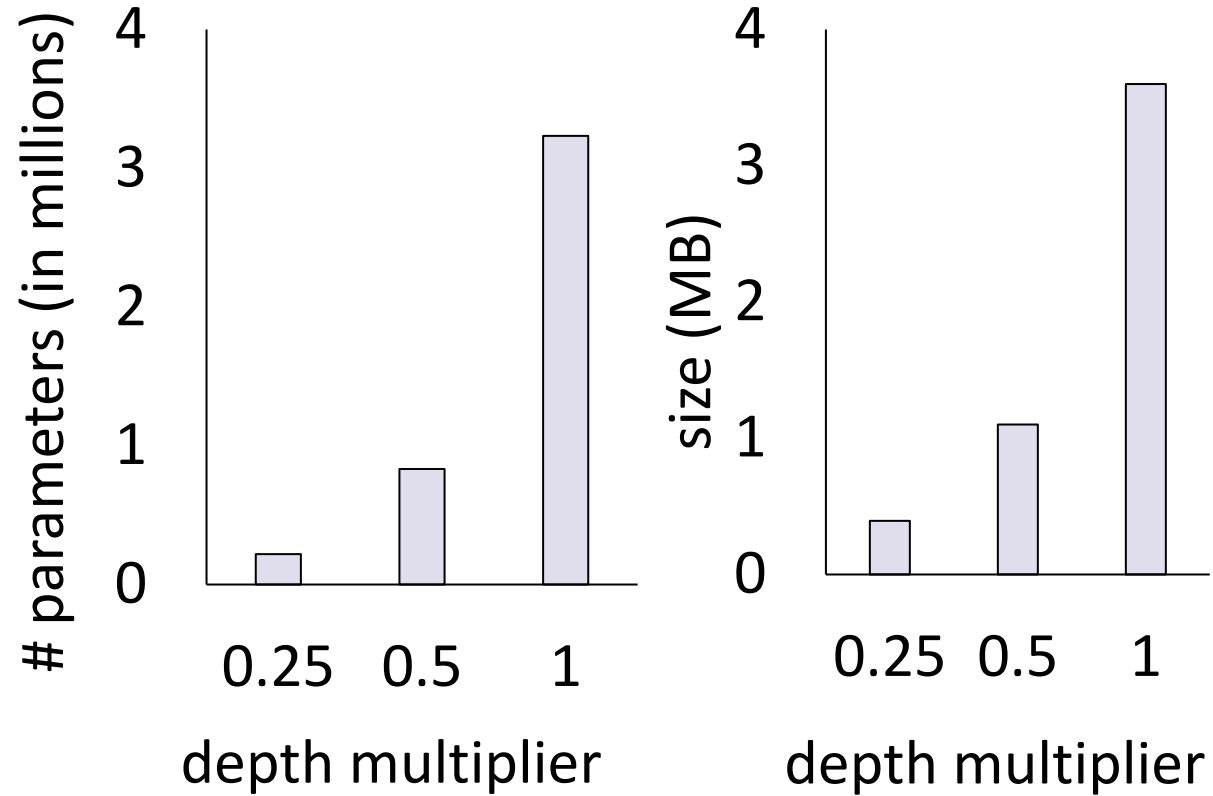


| | ARM Cortex-M7 | Jetson Nano | Toradex Verdin | CoralAI * | Neural Compute Stick 2 |
|---------------------|--|----------------------------------|---|--------------------------------------|--|
| framework | TensorFlow Lite for Microcontrollers | TensorRT | TensorFlow Lite with OpenVX backend | TensorFlow Lite | OpenVino |
| quantization | 8bit (to fit the device memory) | 16bit | 8bit (only supports 8bit ints) | 8bit (only supports 8bit ints) | 16bit (only supports 16bit floats) |
| batching | not enough memory to do batching | batch size per inference = 16 | doesn't support batching | doesn't support batching | number of concurrent inference requests = 4 |

each require separate optimizations!

model & data


pretrained MobileNetV1




***model size* impacts resource consumption & updatability**

***patching* helps with parallel inference & sending partial images**

latency & power draw

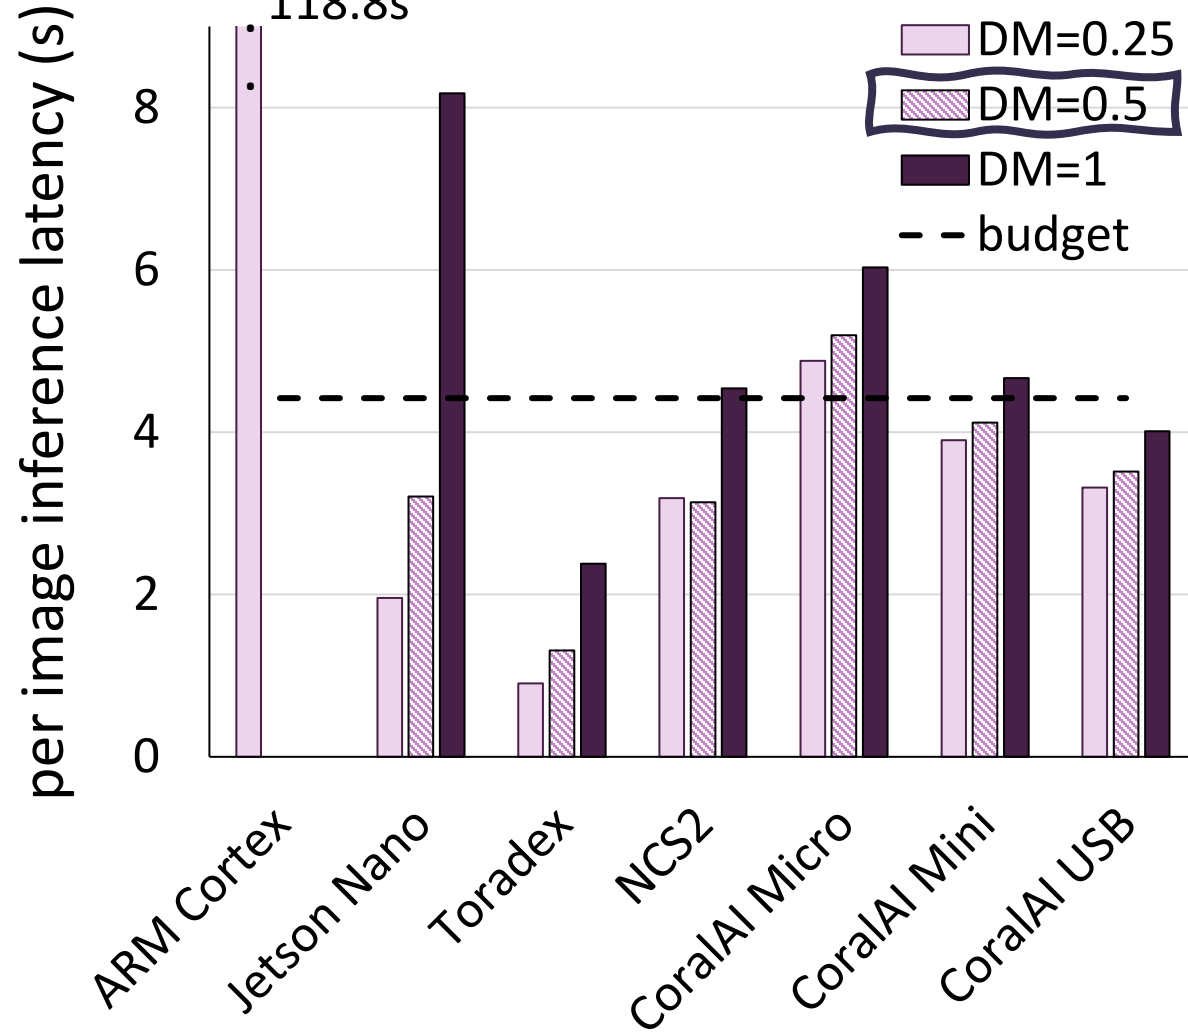


DM = depth multiplier




Jetson & NCS2 have low-latency for smaller models, Toradex is best faster devices fail the peak power budget

latency & power draw





DM = depth multiplier



CoralAI Mini & USB satisfy both latency and power budget.

DISCO satellite



Students launch a satellite to test artificial intelligence in space

On April 14, students from ITU will contribute to writing space history. The satellite, DISCO-1, is launched into space and it carries a microcomputer to test artificial intelligence outside the atmosphere. The satellite is developed by the space program, DISCO, which is a collaboration between students from four Danish universities.

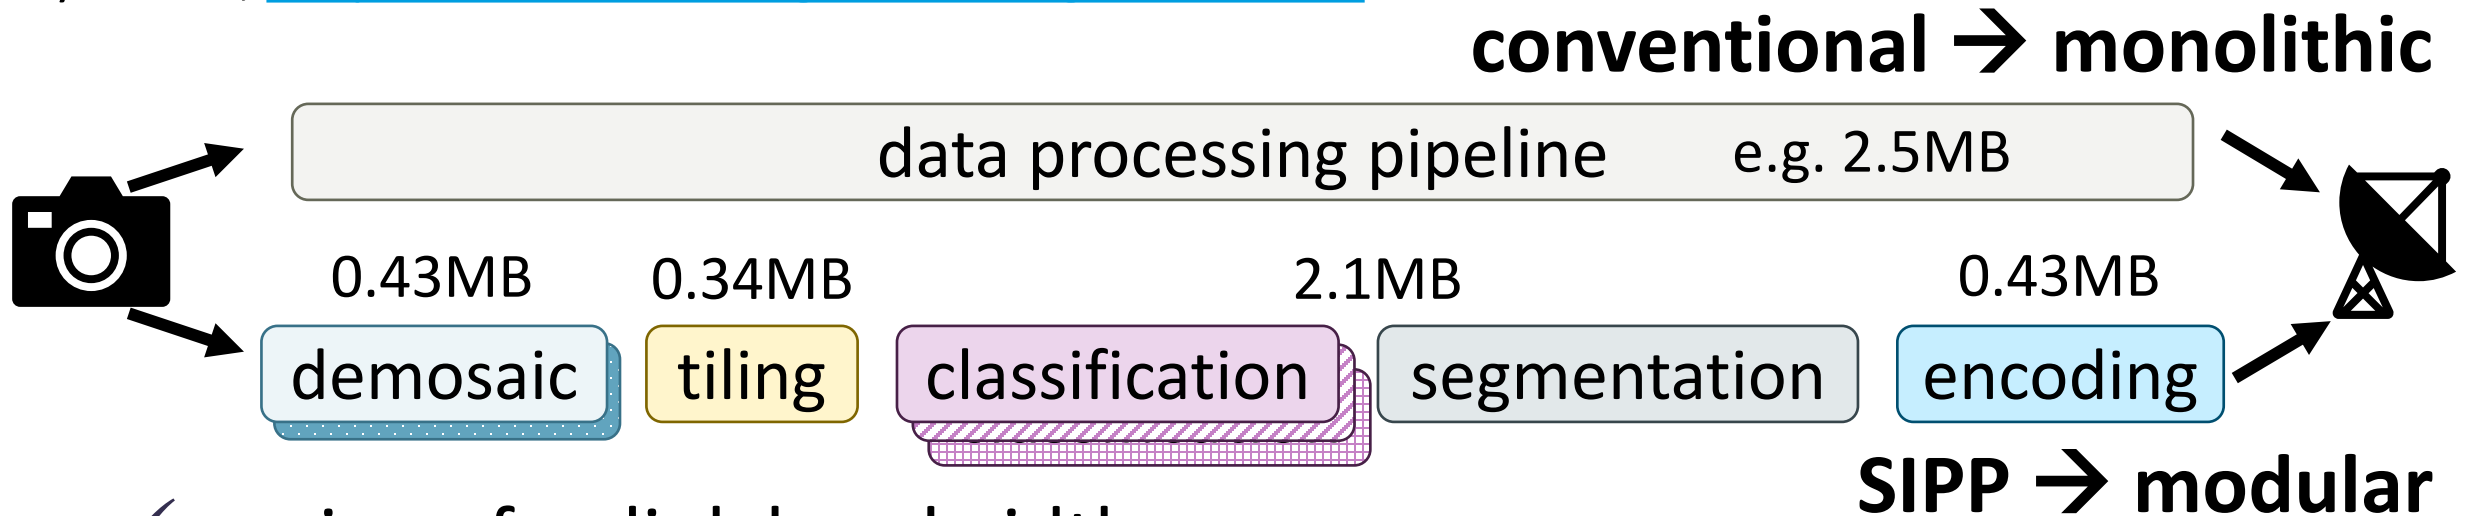
Reaching the Edge of the Edge: Image Analysis in Space

Authors:  [Robert Bayer](#),  [Julian Priest](#),  [Pinar Tözün](#) | [Authors Info & Claims](#)

[DEEM '24: Proceedings of the Eighth Workshop on Data Management for End-to-End Machine Learning](#)
<https://doi.org/10.1145/3650203.3663330>

SIPP: Satellite Image Processing Pipeline

Bayer et al., [Adaptive and Robust Image Processing on CubeSats](#)



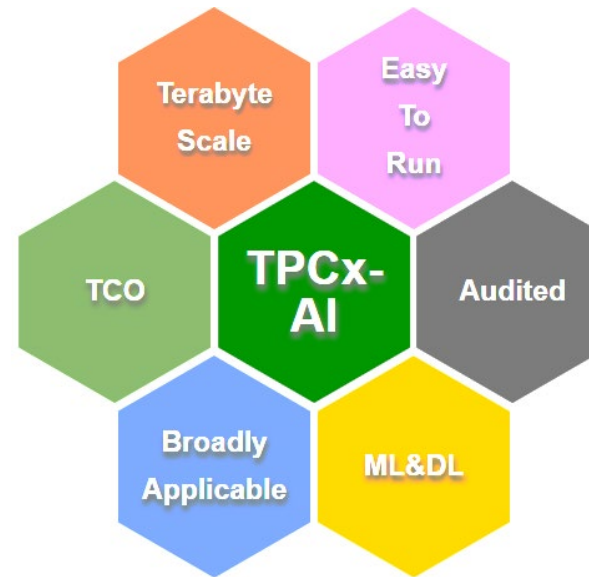
- ✓ saving of uplink bandwidth
 - upload only the updated module, not the whole pipeline
- ✓ flexible module execution based on use case / available power
 - each module can have variants power vs accuracy trade-off
- ✓ fault tolerance – restart the pipeline after a faulty module

overhead of modularity gets amortized in typical runs (1.25%)
for low-power runs it can be up to 20%

future steps – data processing @ the edge

- broader benchmarking

ML
● Commons
tiny, edge, mobile



Towards A Modular End-To-End Machine Learning Benchmarking Framework

Authors: [Robert Bayer](#), [Ties Robroek](#), [Pinar Tözün](#) | [Authors Info & Claims](#)

[TDIS '25: Proceedings of the 3rd International Workshop on Testing Distributed Internet of Things Systems](#) • Pages 23 - 26

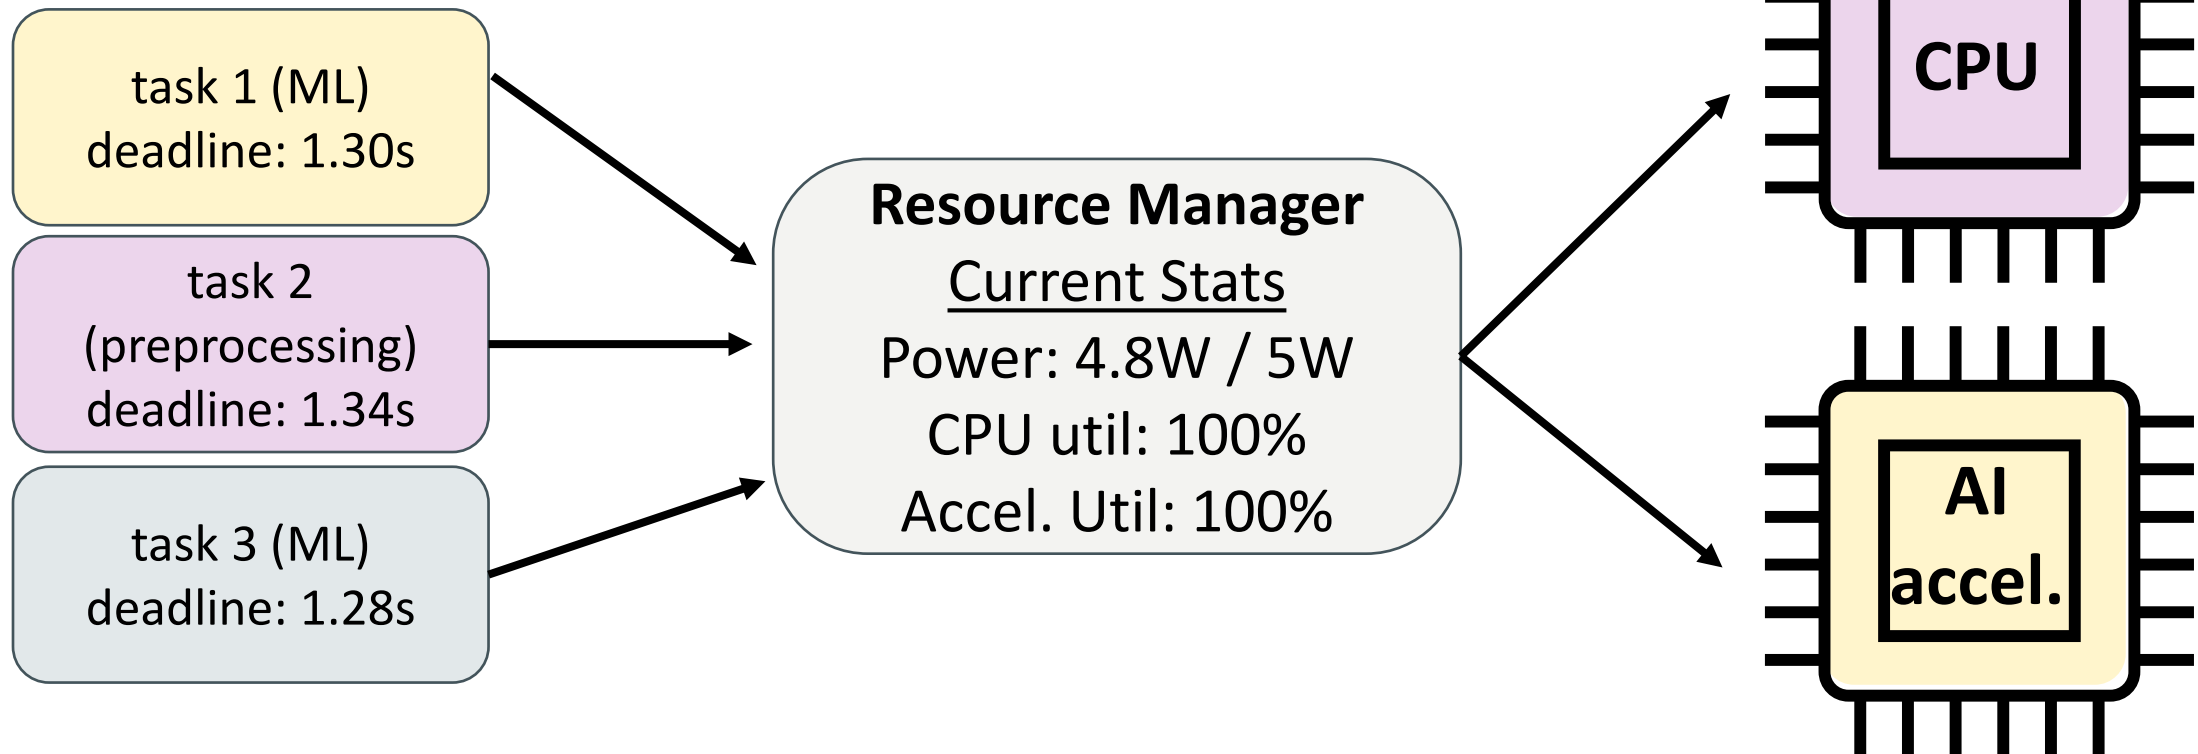
<https://doi.org/10.1145/3719159.3721223>

**need something that can appeal to
different sizes of devices & more complex pipelines!**

future steps – data processing @ the edge

**adjust resources assigned to tasks,
considering power & latency
requirements!**

- broader benchmarking
- smart resource manager at the edge



future steps – data processing @ the edge

- broader benchmarking
- smart resource manager at the edge
- model optimizations
 - e.g., quantization, pruning, operator fusion ...

reduce model sizes with low impact on accuracy!

future steps – data processing @ the edge

- broader benchmarking
- smart resource manager at the edge
- model optimizations
- cross coordination across devices

Lightspeed Data Compute for the Space Era

Thomas Sandholm*, Bernardo A. Huberman*, Klas Segeljakt*, and Paris Carbone*[†]

* *RISE - Research Institutes of Sweden*

thomas.sandholm@ri.se, bernardo.huberman@ri.se, klas.segeljakt@ri.se, paris.carbone@ri.se

[†] *KTH Royal Institute of Technology*

parisc@kth.se

future steps – data processing @ the edge

- broader benchmarking
- smart resource manager at the edge
- model optimizations
- cross coordination across devices
- ...

RAD - resource-aware data systems

postdocs



Ties
Robroek



Ehsan
Yousefzadeh-Asl-Miandoab

phd students



Robert
Bayer



Jens Birk
Andersen

ongoing collaborations




Pamela Delgado, Büşra Karatay
HEIG-VD



Tilmann Rabl, Marcel Weisgut
HPI



Julian Priest
ITU



Vivek Shah
Samsung

ITU students who worked on DISCO

- Adam Julian Hvitsted Rose
- Anton Skou Møller
- Bergur Davidsen
- Burak Özdemir
- Christina Gaitanou
- Daniel Gjesse Kjellberg
- David Nikolaj Vinje
- Diana Alexandra Buras
- Emil Hinge Voigt
- Emil Houlborg
- Frederik Badegaard Avlund
- Gianluca Feltrin
- Ívar Óli Sigurðsson
- Jacob Sonne
- Jacob Søgaard Bøgh
- Jens Hvass Jørgensen
- Jeppe Lindhard
- Jeppe Theiss Kristensen
- Jesper Terkildsen
- Joakim Goyle Dissing
- Jonas Fuhr Høyer
- Jonas Nielsen
- Julian Hvitsted Rose
- Julian Kristoffer Valbjerg Pedersen
- Katarina Christovova
- Kristian Ullum Vind
- Kristoffer Borgstrøm
- Mahmood Mohammed Seoud
- Maja Gréta Szász
- Martin Pinholt
- Mathias M. A. Larsen
- Nicolaj Valsted
- Nikolaj Sørensen
- Orusha Thapa Magar
- Richard Kentos
- Robert Bayer
- Rokas Kasperavicius
- Ryan Williams
- Selim Yetkin
- Su Mei Gwen Ho
- Thomas Meibom
- Victor Popp Henriksen
- Viktor Horváth
- ...

data processing @ the edge thank you!

- demand for more data analysis closer to the data source
 - reduces data movement & privacy concerns
 - helps with real-time decisions
- variety of edge devices to choose from offering increasingly powerful hardware but still resource-constrained
 - requires not just latency-efficient, but also energy-efficient data processing
- hardware specialization helps with latency & power budget
 - though, we need more flexibility




pito@itu.dk
pinartozun.com


need for methods that can deal with resource management & program updates at the edge!

backup

latency & power draw



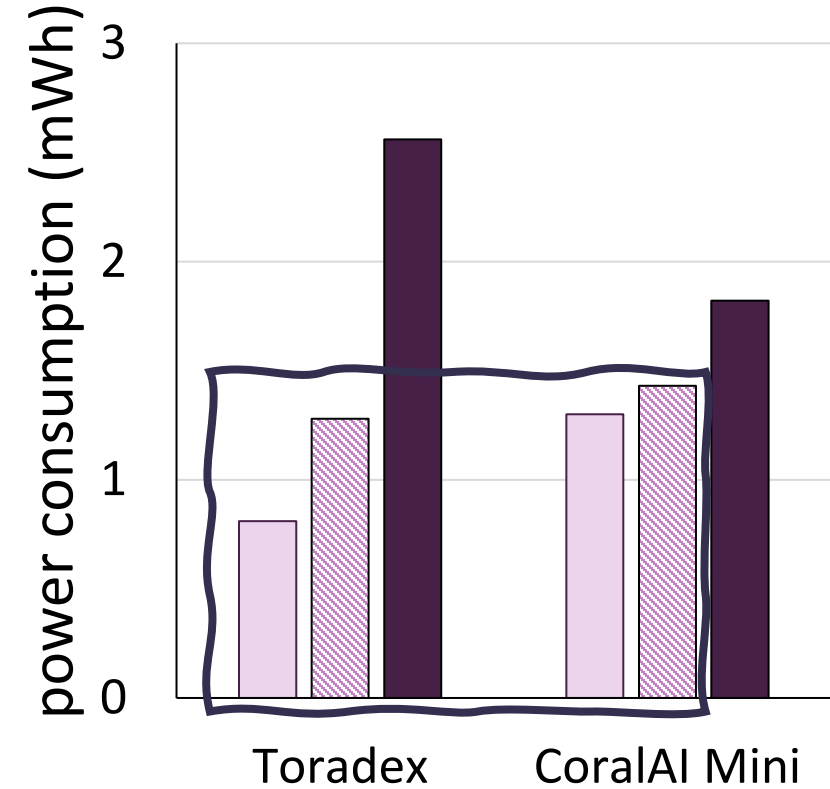


DM = depth multiplier



**ARM-based microcontroller draws little power per unit time
but per inference power need is higher than the rest!**

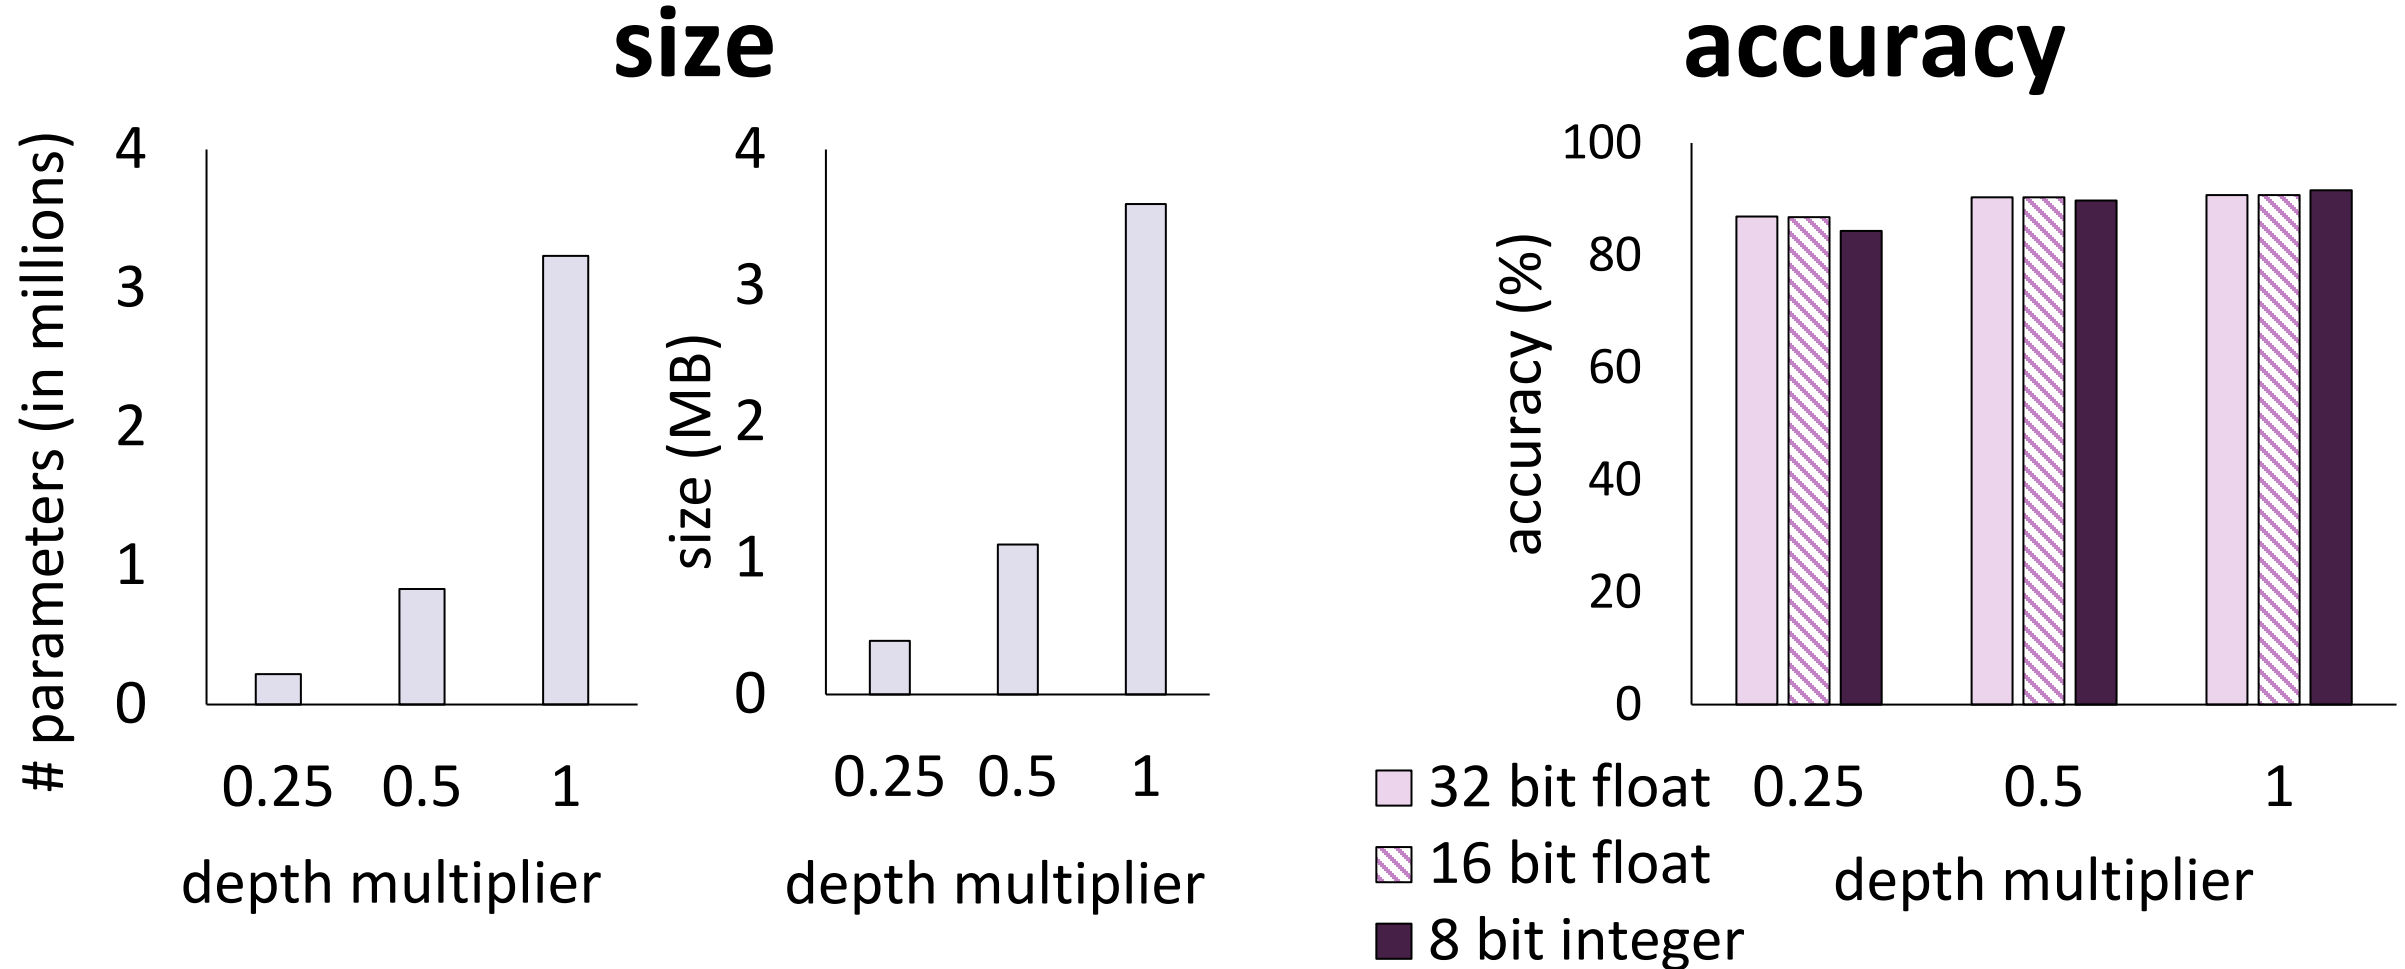
DISCO 2 – 3 CubeSats

- Toradex device – non-dev-board version



**Toradex is the most energy-efficient for the small & mid-sized models!
non-dev-board version reduces the peak power.**

model – pretrained MobileNetV1



model size impacts resource consumption & updatability.
quantization does not have big impact on accuracy.

image processing unit = IPU

