THE CHIPS TO SYSTEMS CONFERENCE
SHAPING THE NEXT GENERATION OF ELECTRONICS

The Magnificent Seven Challenges & Opportunities in Domain-Specific Accelerator Design for Autonomous Systems

Sabrina M. Neuman*
Boston University
sneuman@bu.edu

Brian Plancher*
Barnard College, Columbia University
bplancher@barnard.edu

Vijay Janapa Reddi*
Harvard University
vj@eecs.harvard.edu

Robotics is a **Big** and **Embodied** space

Sensors → Compute → Actuators

Autonomous Systems

~~Robotics~~ is a Big and Embodied space

Sensors → Compute → Actuators

Autonomous Systems can do amazing things …

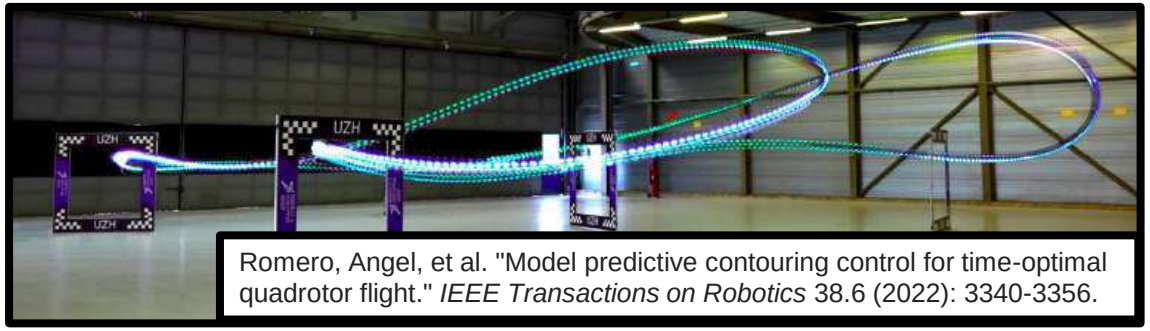

Romero, Angel, et al. "Model predictive contouring control for time-optimal quadrotor flight." *IEEE Transactions on Robotics* 38.6 (2022): 3340-3356.

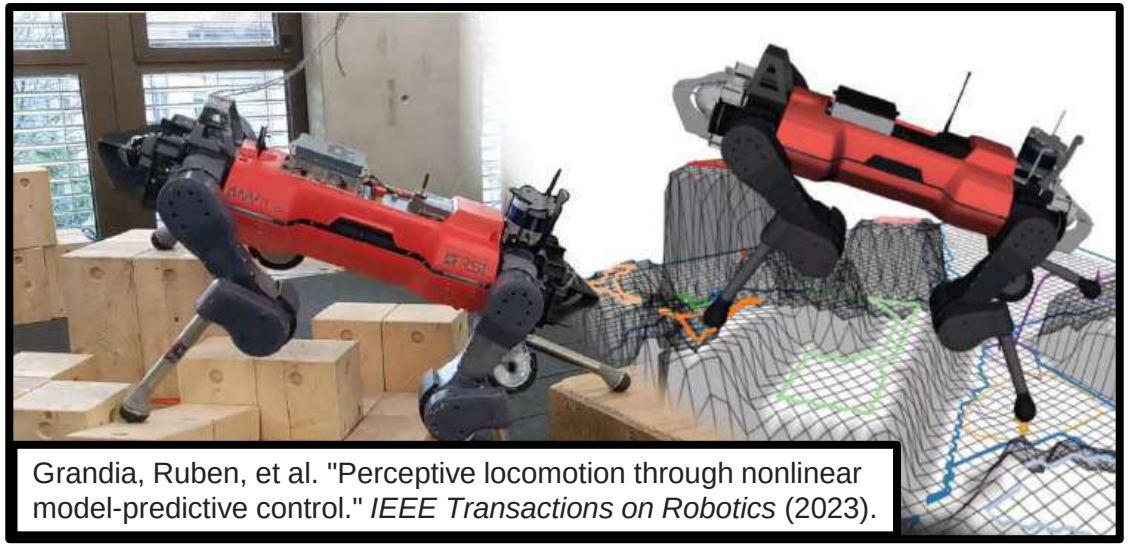

Grandia, Ruben, et al. "Perceptive locomotion through nonlinear model-predictive control." *IEEE Transactions on Robotics* (2023).

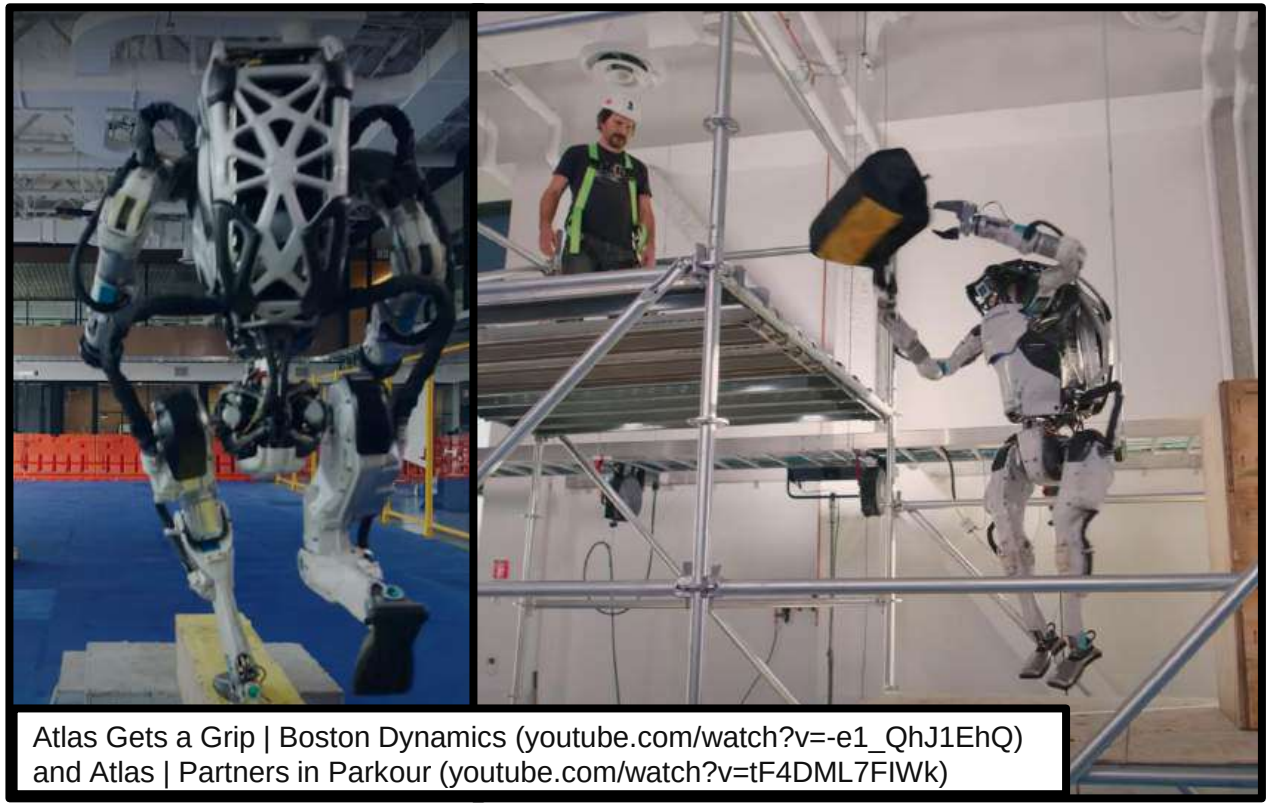

Atlas Gets a Grip | Boston Dynamics (youtube.com/watch?v=-e1_QhJ1EhQ) and Atlas | Partners in Parkour (youtube.com/watch?v=tF4DML7FIWk)

Autonomous Systems can do amazing things …

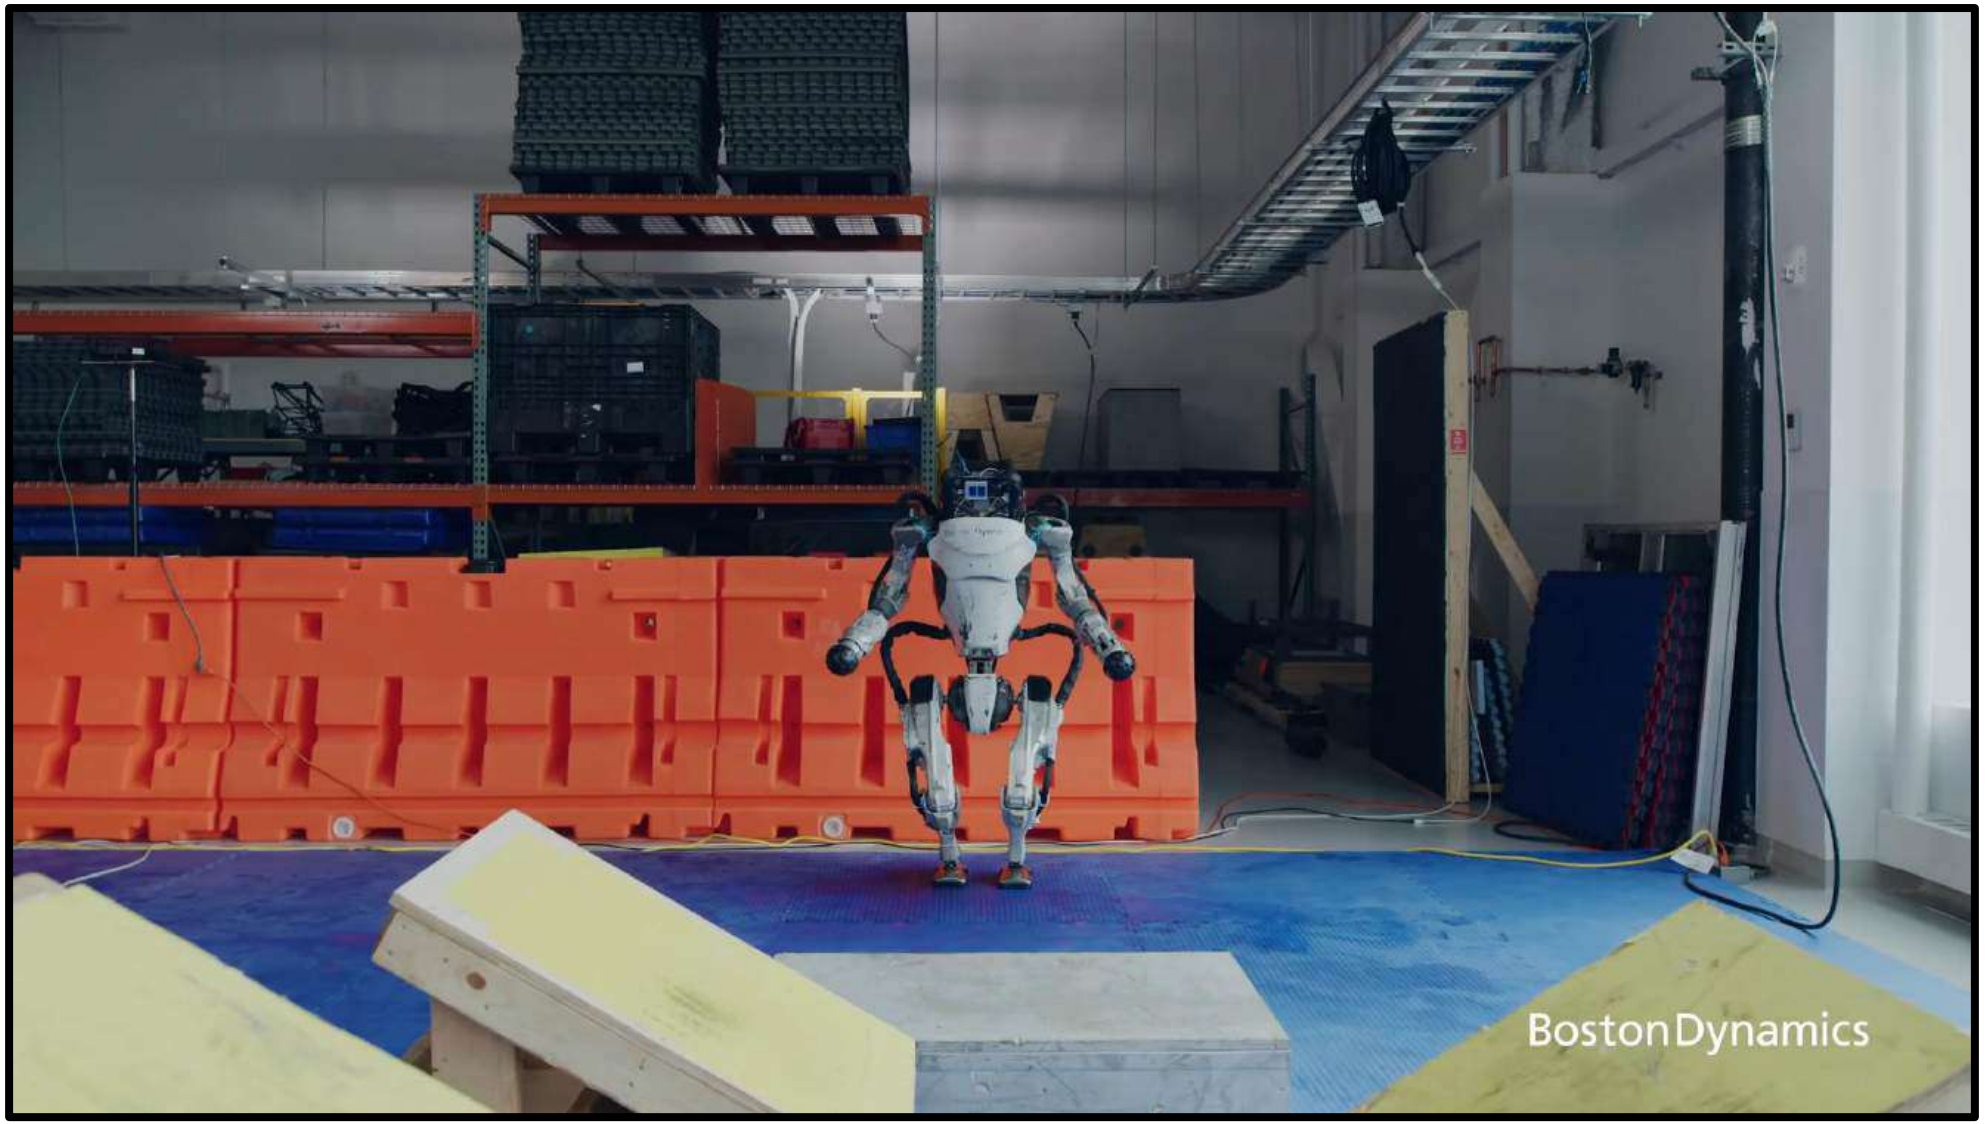

… but they still have a long way to go!

… but they still have a long way to go!

Romero, Angel, et al. "Model predic
quadrotor flight." *IEEE Transactions*

Grandia, Ruben, et al. "Perceptive locomoti
model-predictive control." *IEEE Transactio*

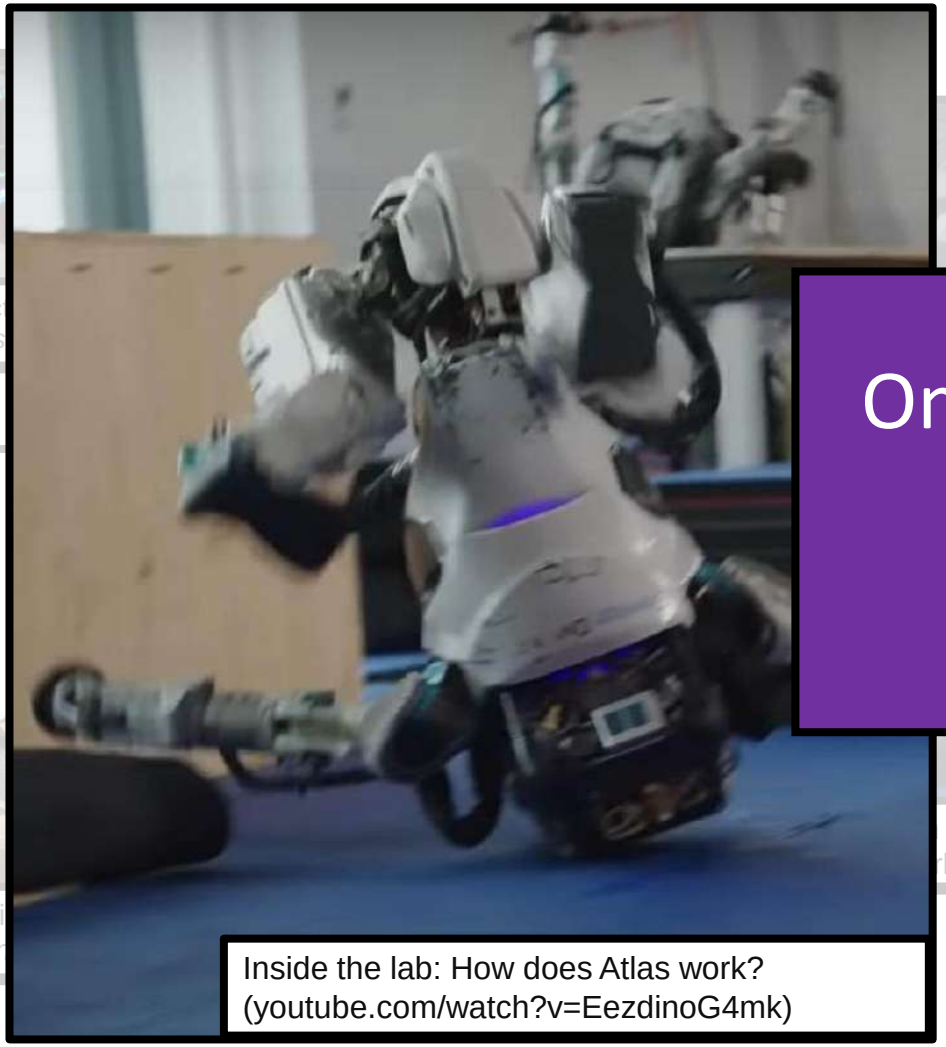

Inside the lab: How does Atlas work?
(youtube.com/watch?v=EezdinoG4mk)

Dynamics (youtube.com/watch?v=-e1_QhJ1EhQ)
kour (youtube.com/watch?v=tF4DML7FIWk)

One Major Challenge: Computational Performance!

Unfortunately, we have reached the limits of technology scaling and on-chip power density…

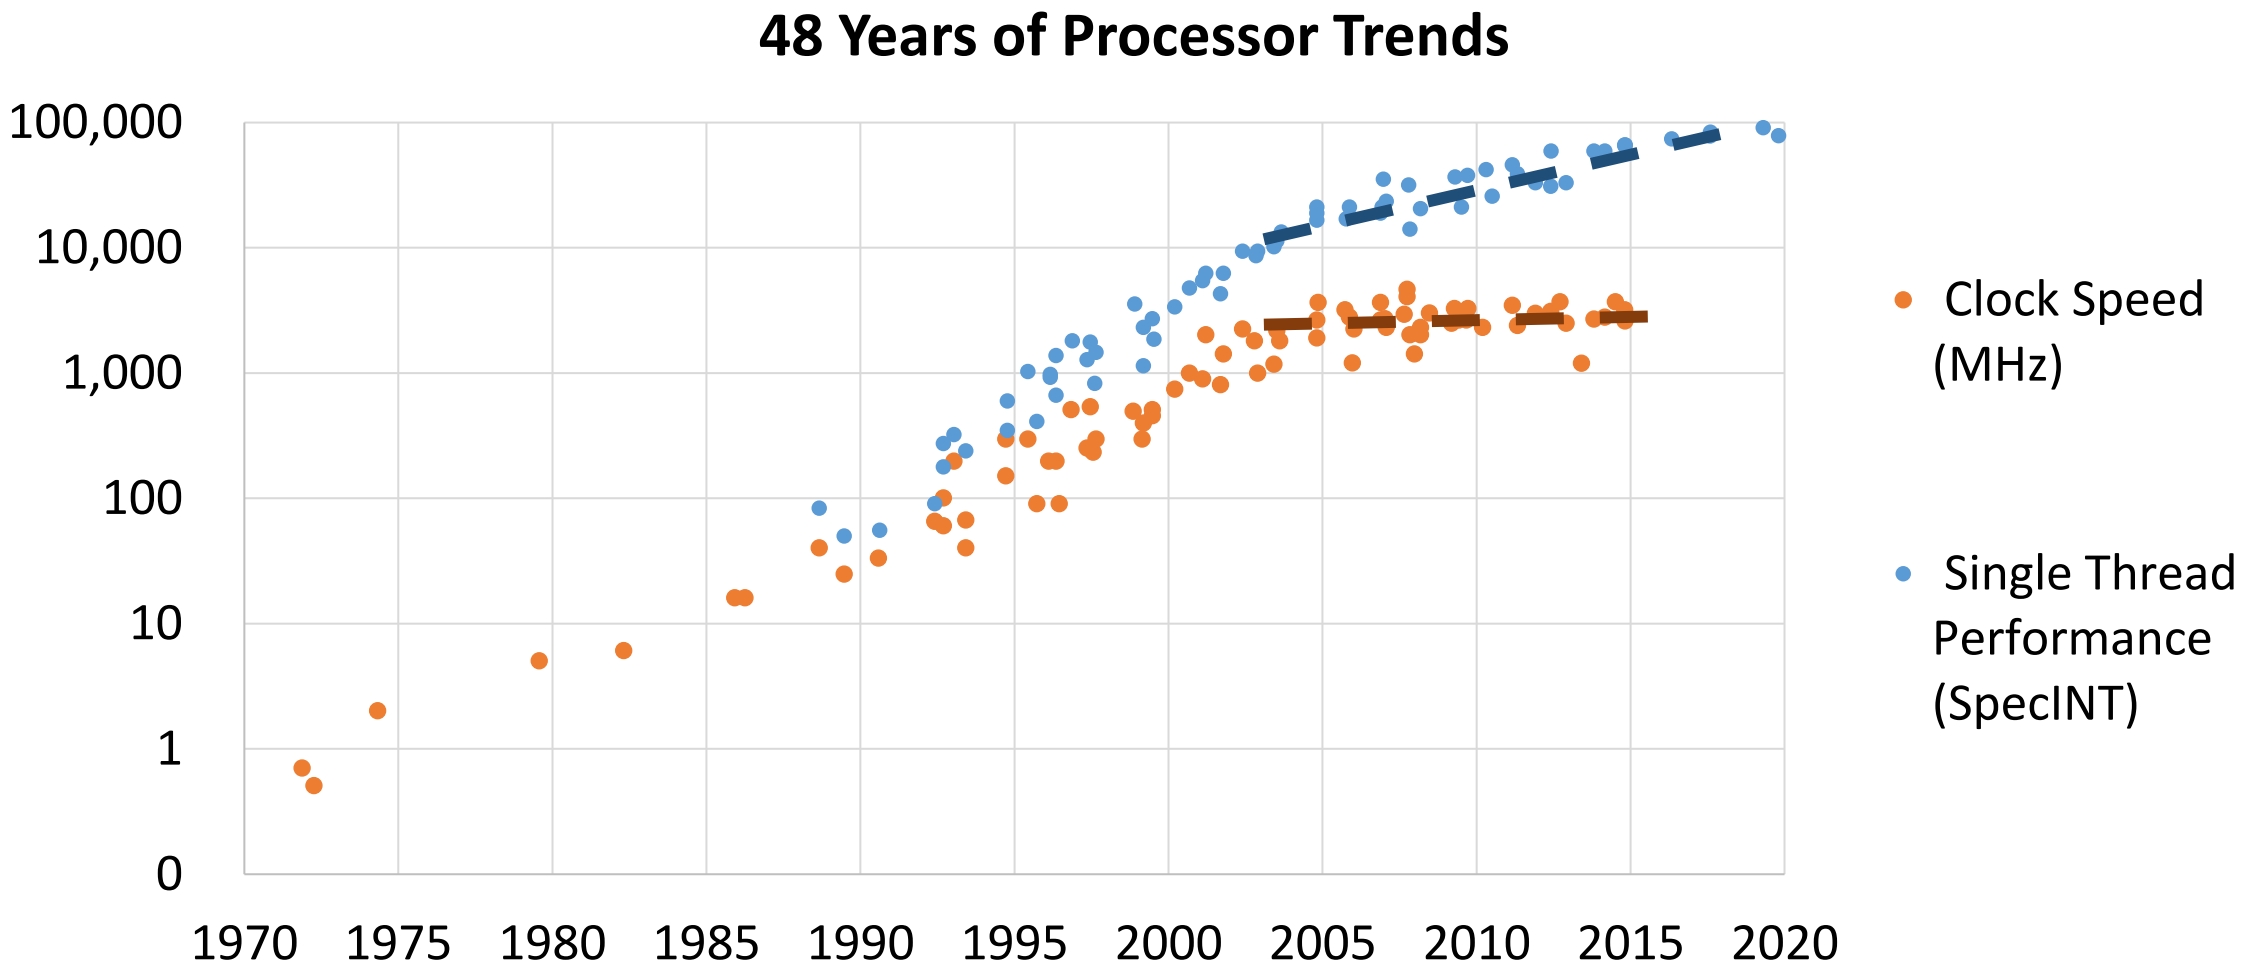

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2019 by K. Rupp https://github.com/karlrupp/microprocessor-trend-data

…resulting in a need to leverage parallelism…

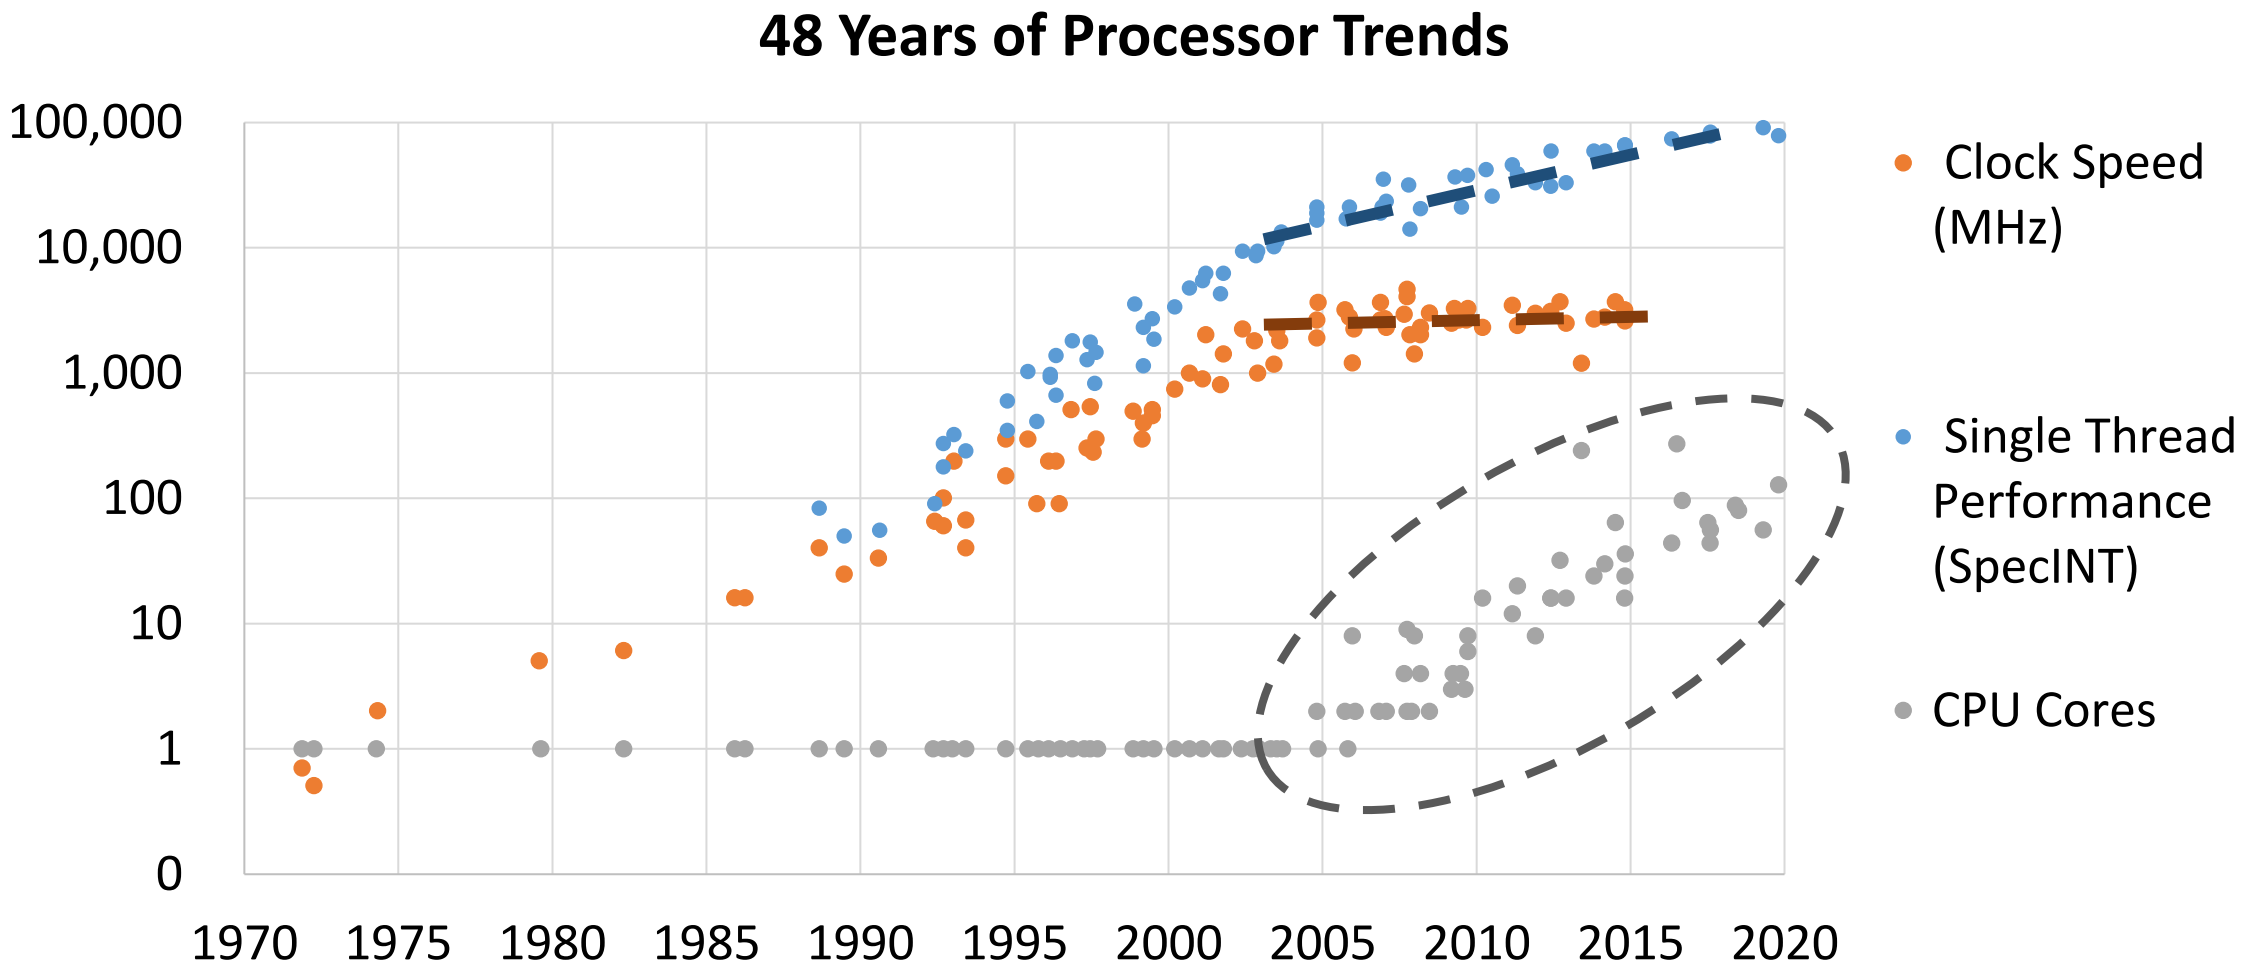

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2019 by K. Rupp https://github.com/karlrupp/microprocessor-trend-data

...and move toward specialized accelerators!

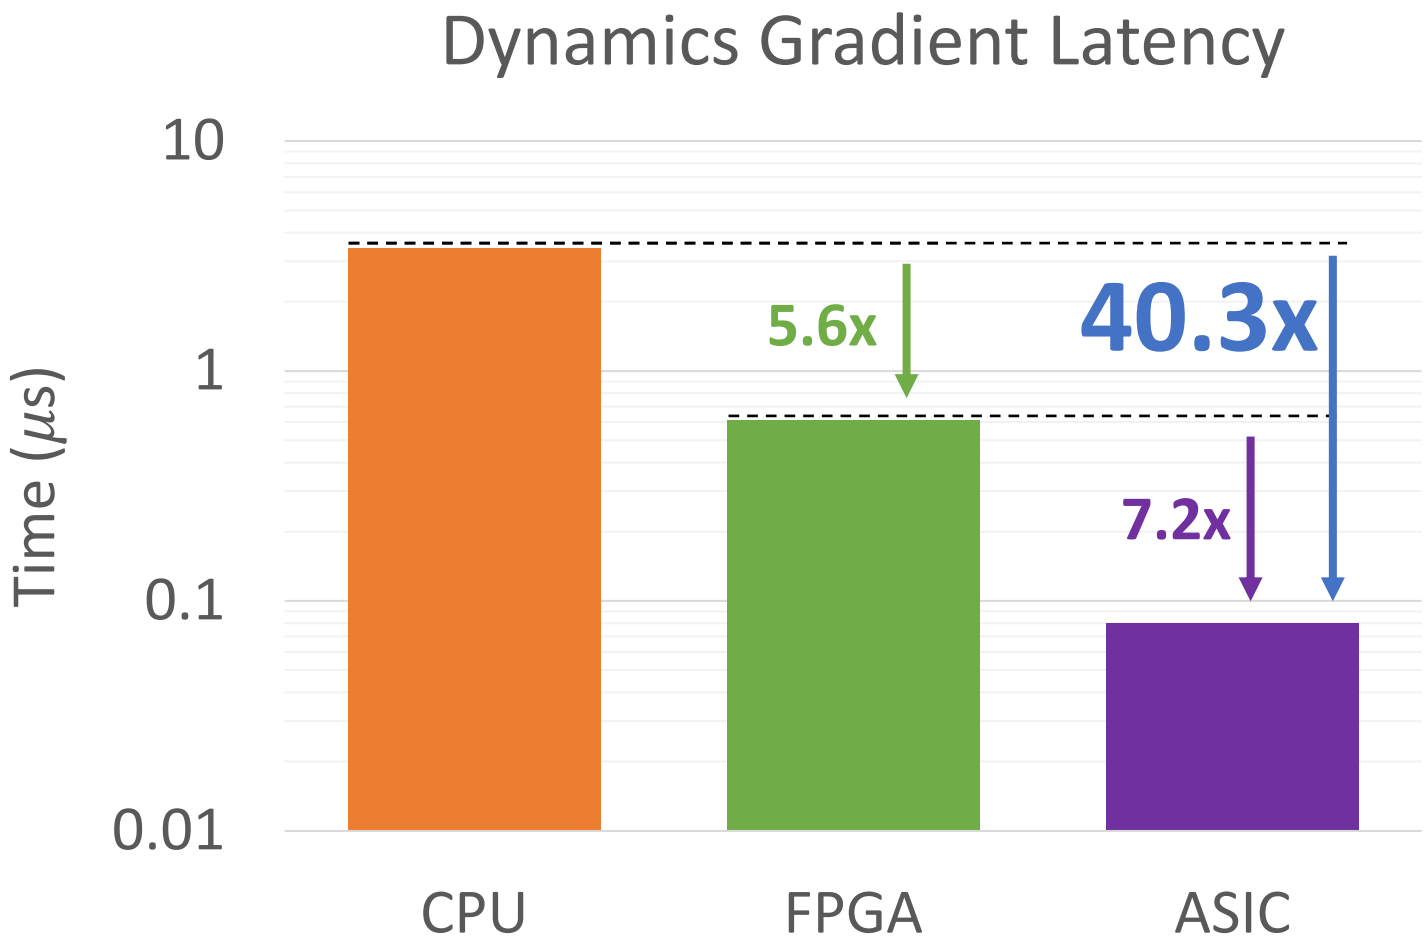

S. M. Neuman, B. Plancher, T. Bourgeat, T. Tambe, S. Devadas, V. Janapa Reddi. Robomorphic Computing: A Design Methodology for Domain-Specific Accelerators Parameterized by Robot Morphology. ASPLOS 2021.

...and move toward specialized accelerators!

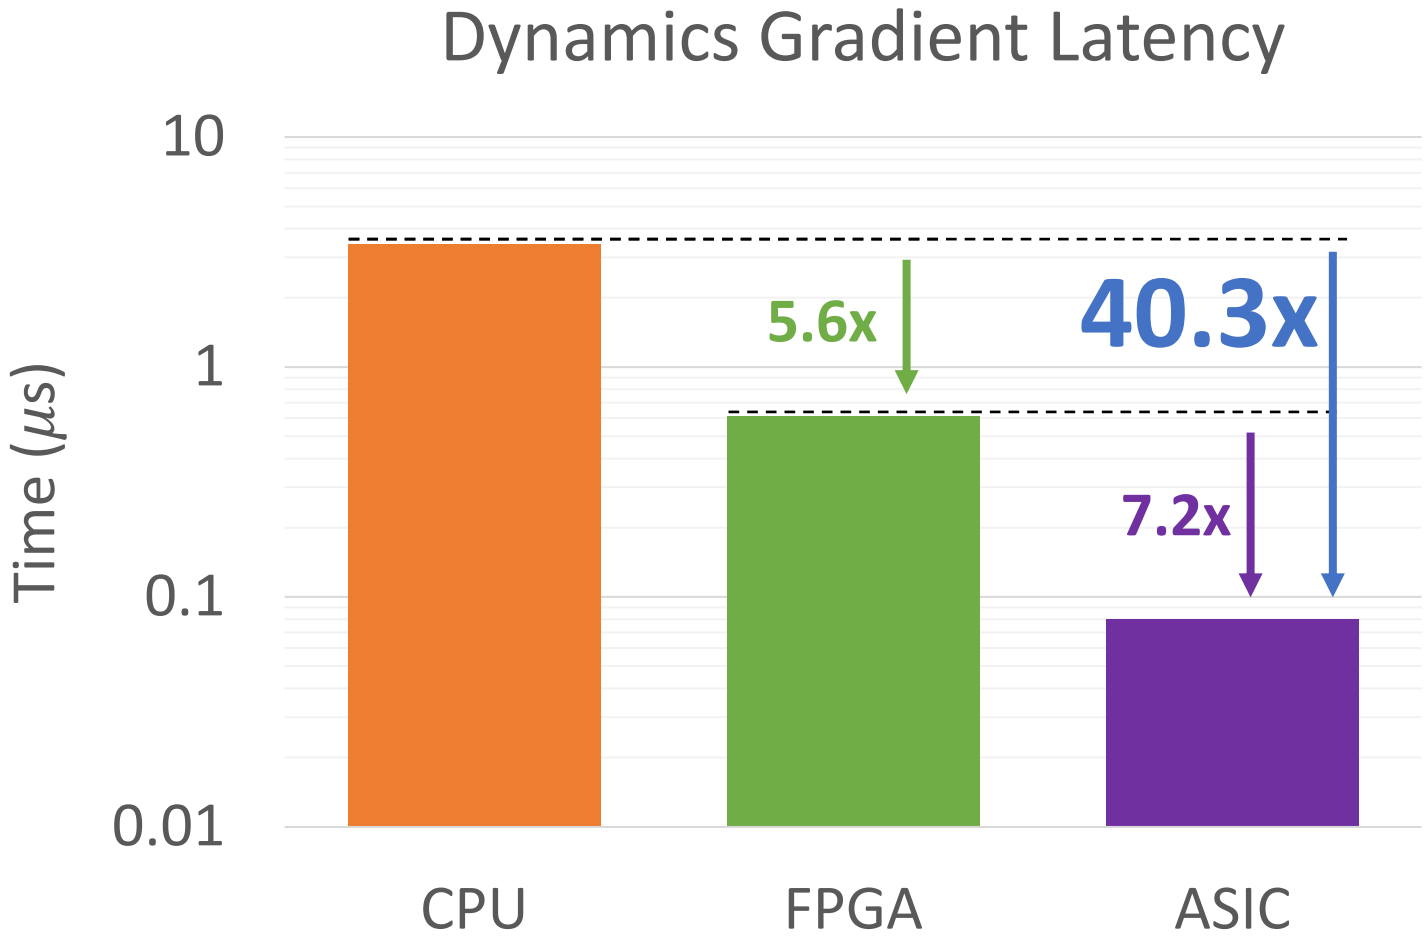

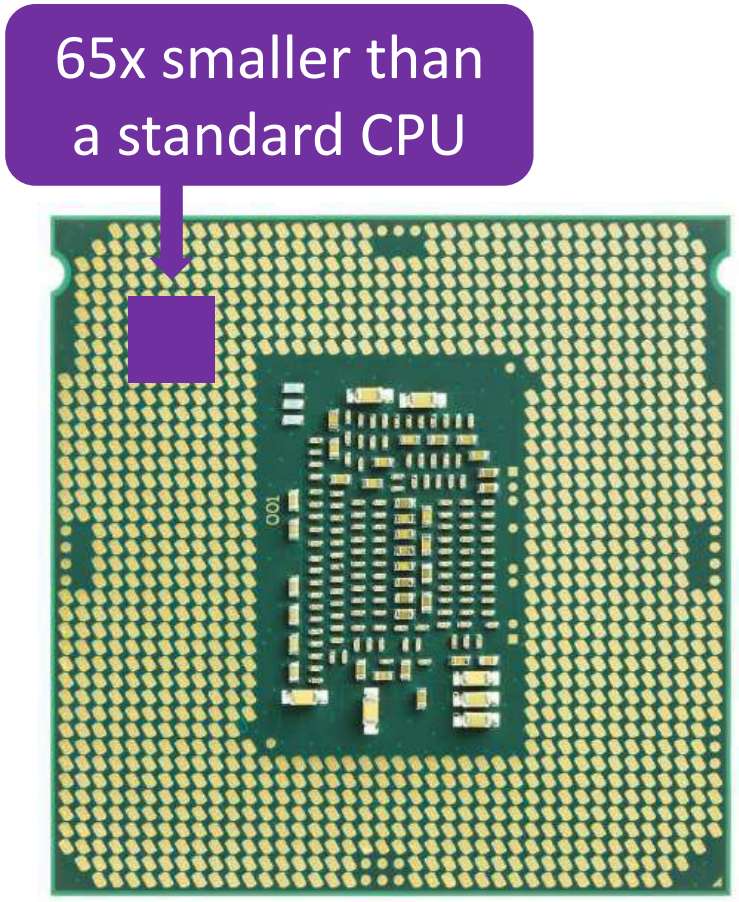

S. M. Neuman, B. Plancher, T. Bourgeat, T. Tambe, S. Devadas, V. Janapa Reddi. Robomorphic Computing: A Design Methodology for Domain-Specific Accelerators Parameterized by Robot Morphology. ASPLOS 2021.

Interest in Accelerators for Autonomous Systems is Growing!

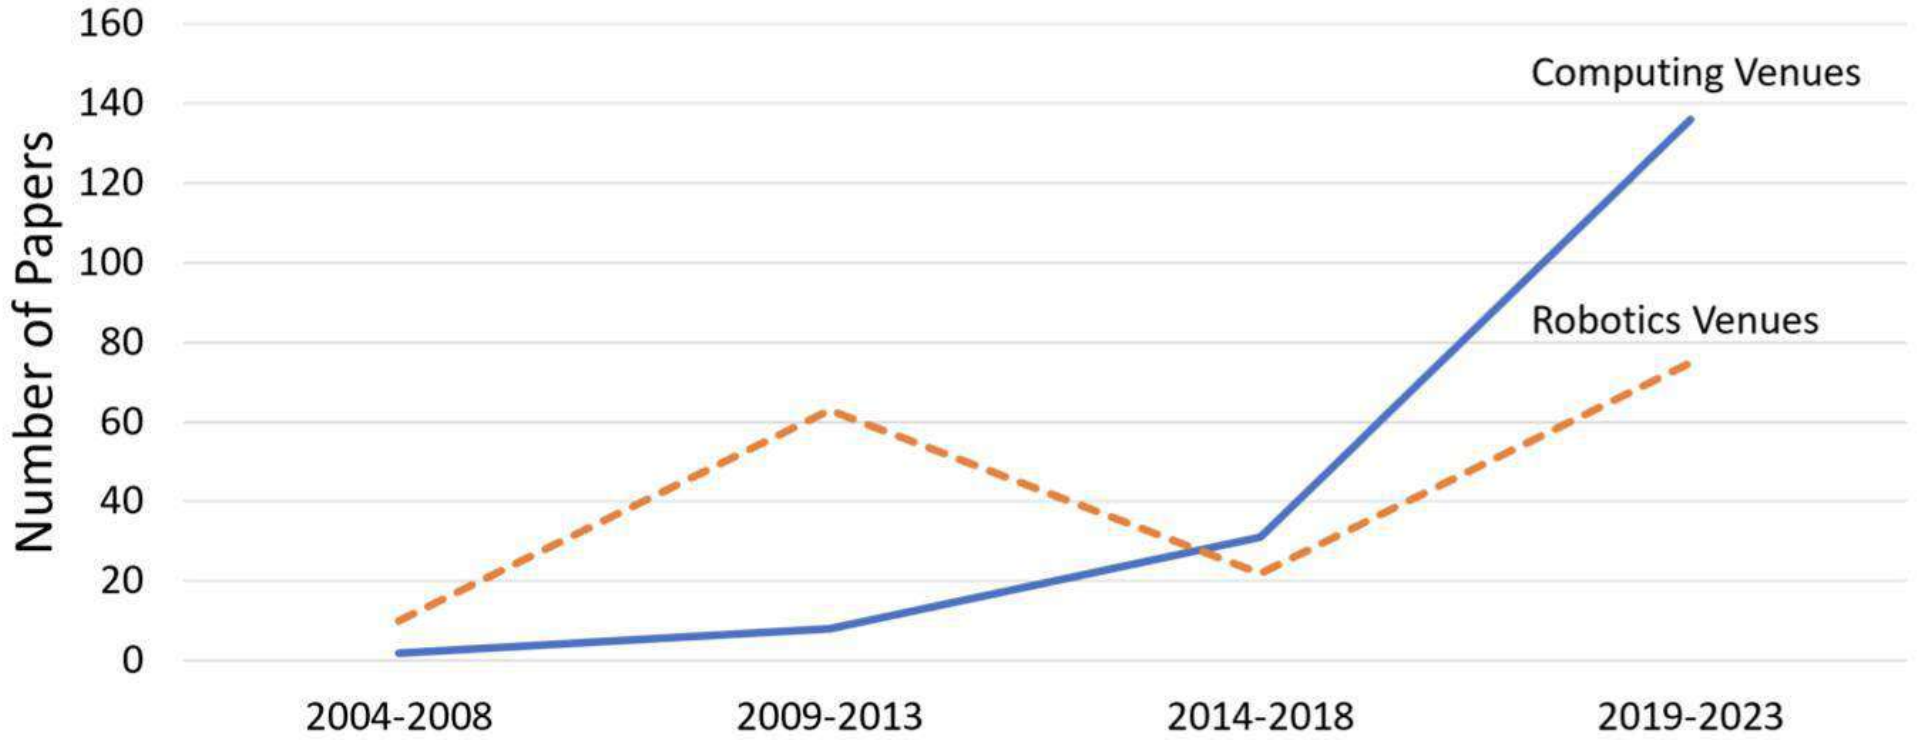

Computing Keywords: Accelerator & Robot | Autonomous System | UAV | Drone | Autonomous Vehicle | Self-Driving, in DAC, ISCA, MICRO, HPCA, and ASPLOS. Robotics Keywords: ASIC | FPGA, in ICRA, IROS, RSS, and RA-L.

Interest in Accelerators for Autonomous Systems is Growing!

How should we build domain-specific accelerators for autonomous systems?

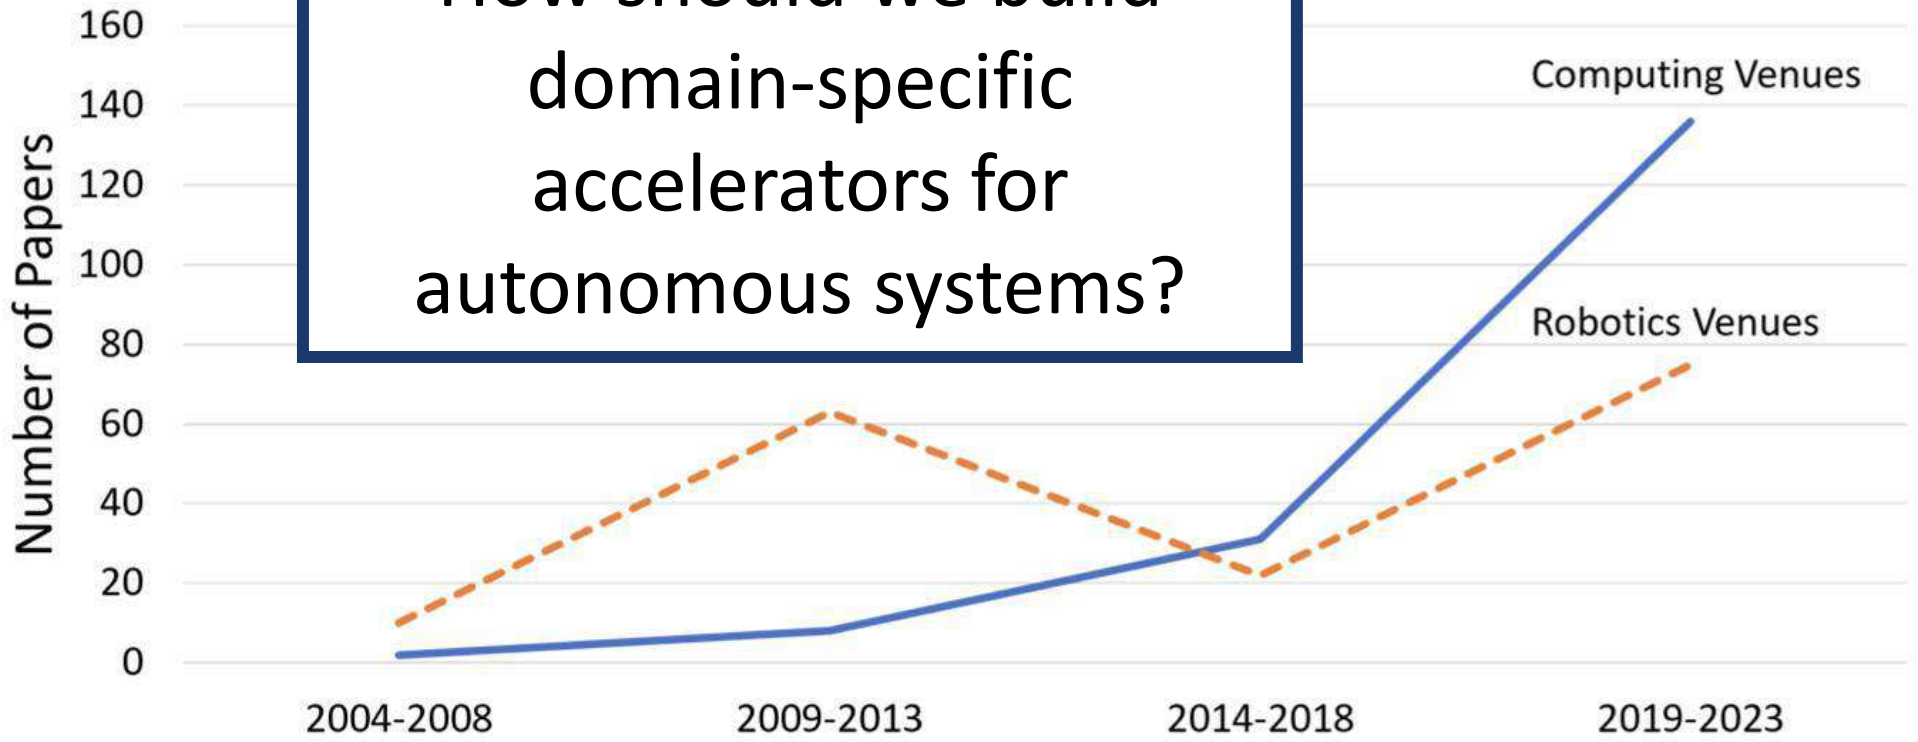

Computing Keywords: Accelerator & Robot | Autonomous System | UAV | Drone | Autonomous Vehicle | Self-Driving, in DAC, ISCA, MICRO, HPCA, and ASPLOS. Robotics Keywords: ASIC | FPGA, in ICRA, IROS, RSS, and RA-L.

The Magnificent Seven
Challenges & Opportunities in Domain-Specific Accelerator Design for Autonomous Systems

1. Build Bridges: Engage with Domain Experts
2. Measure Twice, Cut Once: Metrics Matter
3. “Widgetism”: Avoid Over-Specialization
4. Pump the Brakes: Do Not Always Accelerate
5. Chips and Salsa: Acceleration Beyond ASICs
6. Forest vs. Trees: Take an End-to-End View
7. Design Global: Sustainability and Impact

1 Build Bridges: Engage with Domain Experts

Pitfall: *Interact with domains exclusively through benchmarks published in computer systems, without input from domain experts.*

1 Robotics is a **Big** and **Embodied** space

Sensors → Compute → Actuators

1 Robotics is a **Big** and **Embodied** space

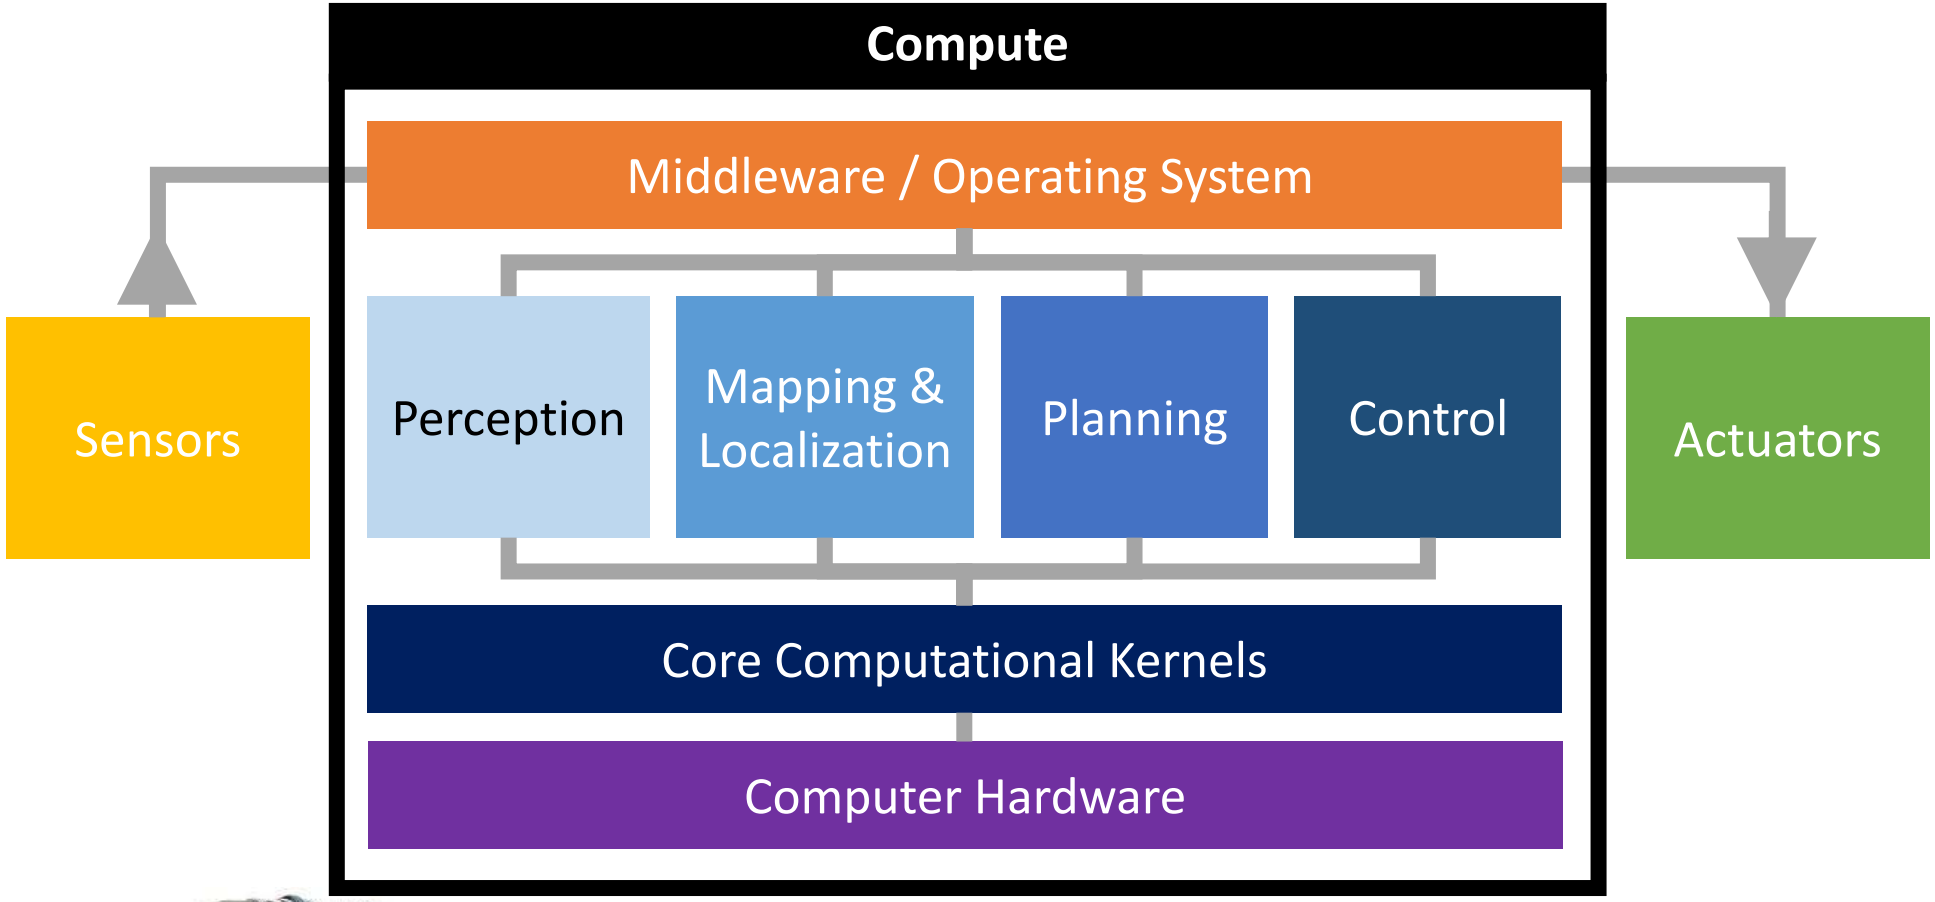

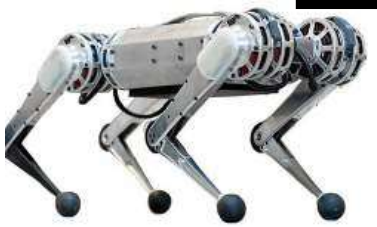

1 Robotics is a **Big** and **Embodied** space

1 Robotics is a **Big** and **Embodied** space

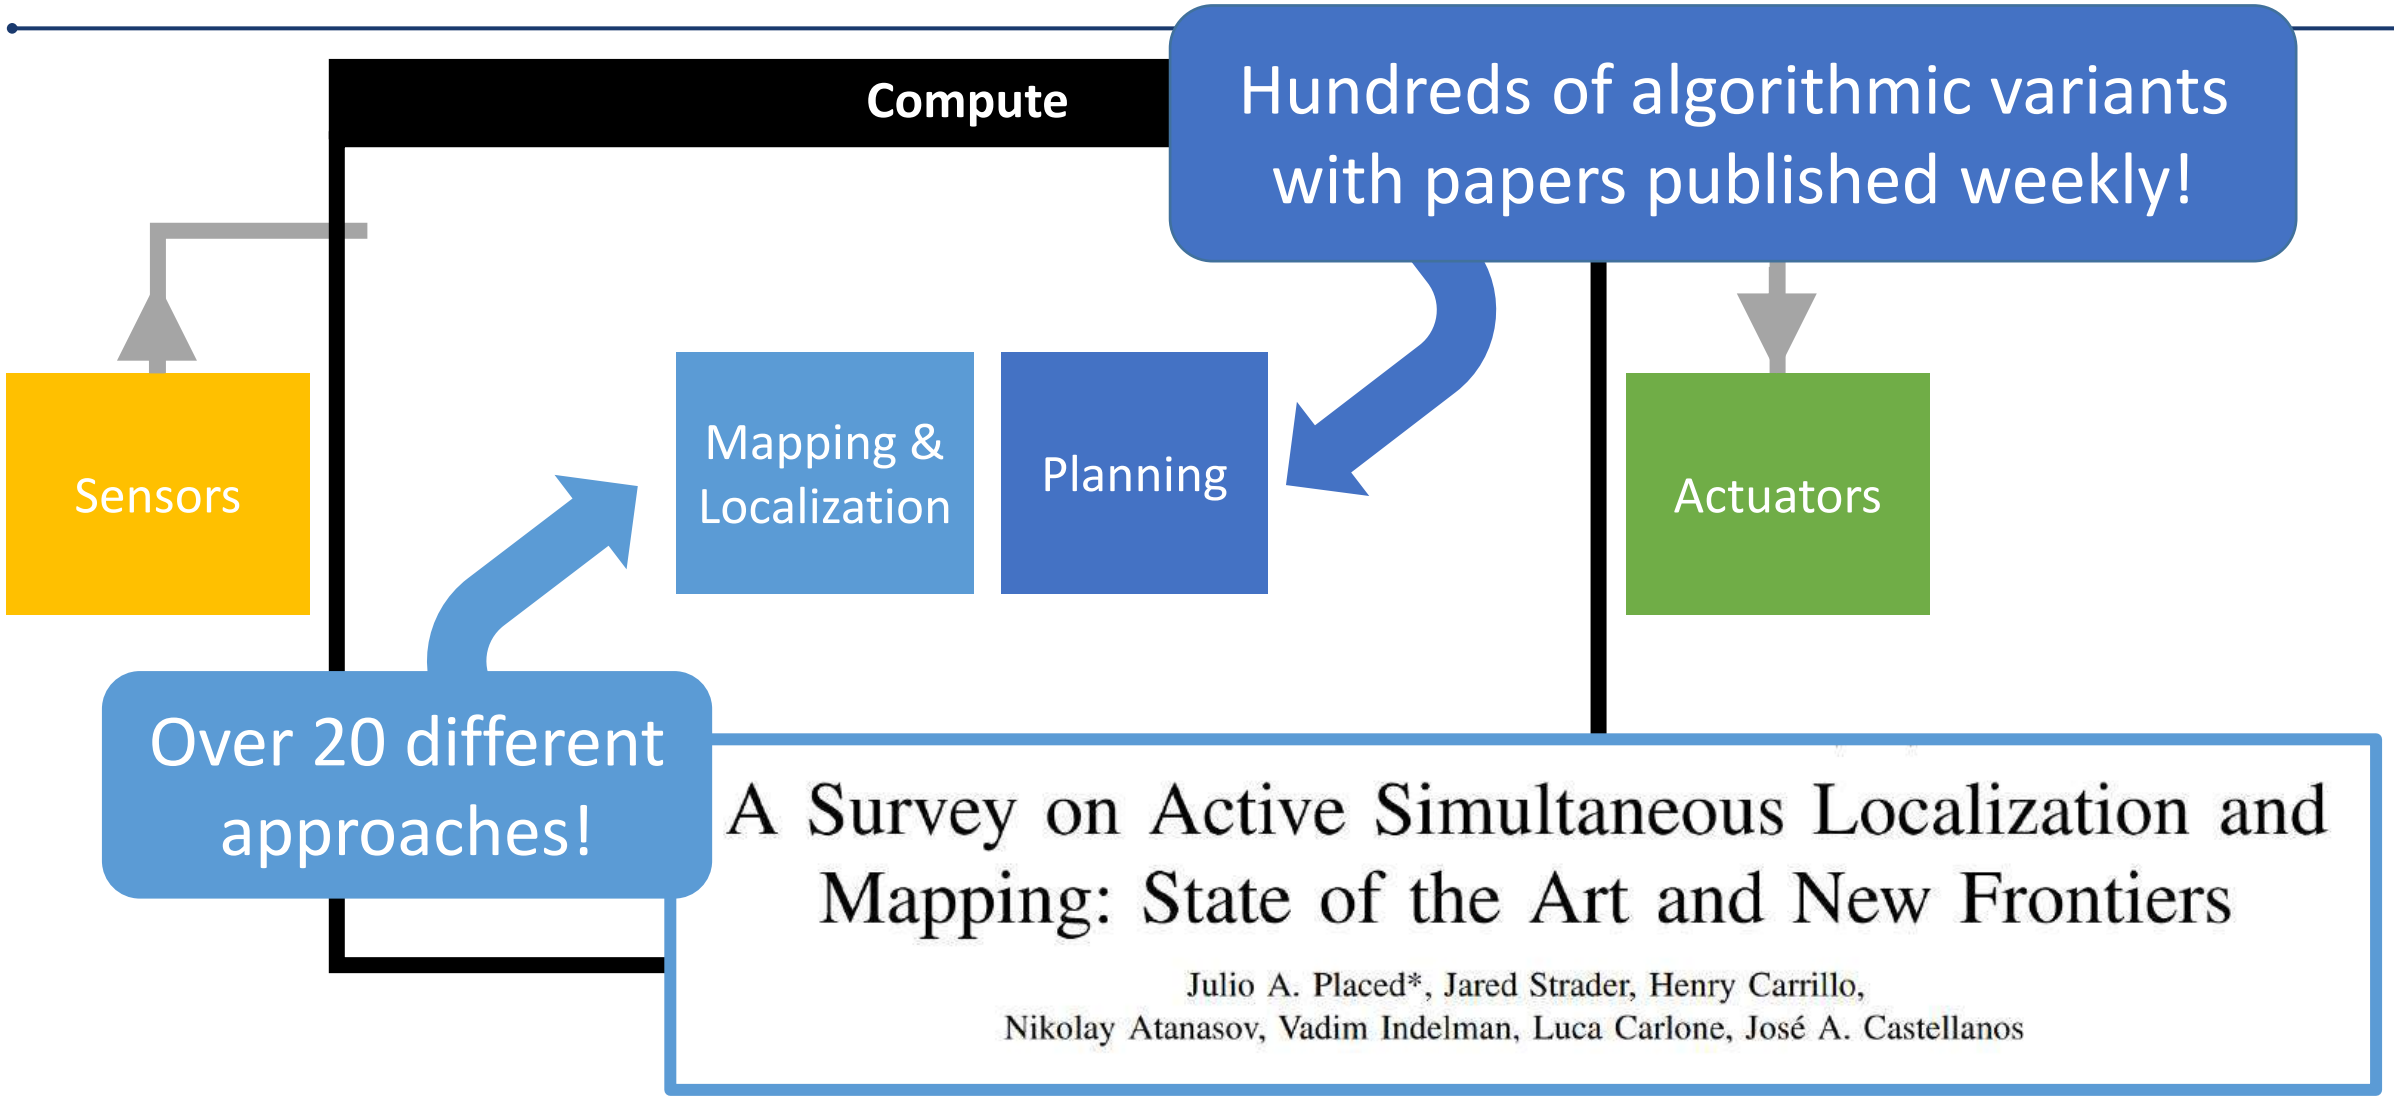

1 Build Bridges: Engage with Domain Experts

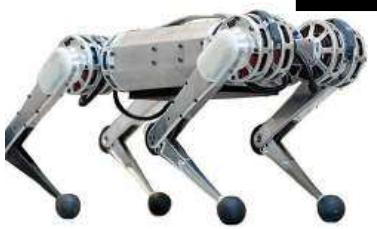

3 “Widgetism”: Avoid Over-Specialization

Pitfall: *A cycle of pick one slow algorithm, lower it to an ASIC, repeat.*

5 Chips and Salsa: Acceleration Beyond ASICs

Pitfall: *Focus on ASICs, leaving software, GPUs, and FPGAs behind.*

3,5 “Widgetism”: Avoid Over-Specialization
Chips and Salsa: Acceleration Beyond ASICs

Software Performance Engineering can get you pretty far!

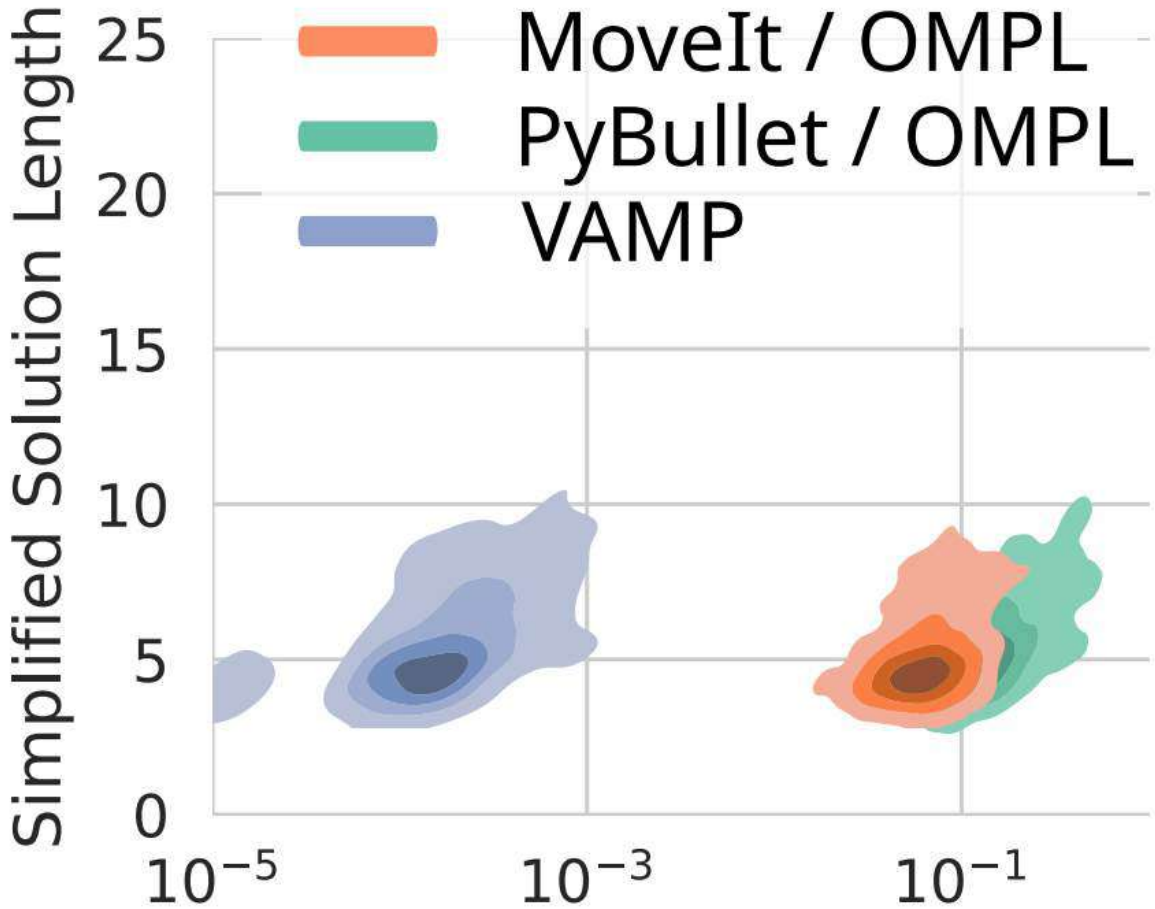

W. Thomason, Z. Kingston, L. E. Kavraki. Motions in Microseconds via Vectorized Sampling-Based Planning. ICRA. 2024.

3,5

“Widgetism”: Avoid Over-Specialization
Chips and Salsa: Acceleration Beyond ASICs

GPUs give you the flexibility to quickly iterate plus improved performance!

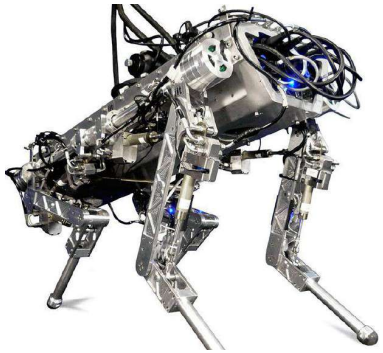

Kuka IIWA Manipulator

HyQ Quadruped

x_G

x_o

What happens when we scale robot complexity and consider multiple parallel computations of robot physics?

3,5

“Widgetism”: Avoid Over-Specialization
Chips and Salsa: Acceleration Beyond ASICs

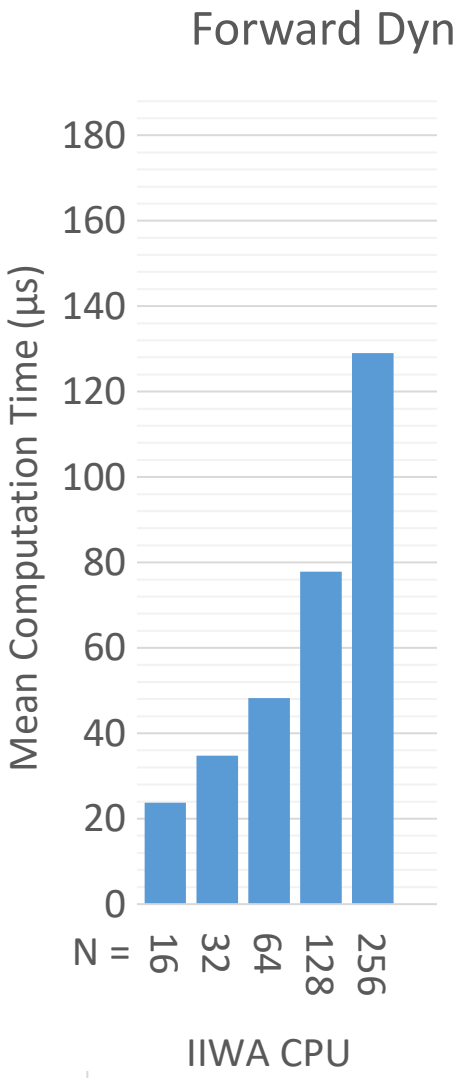

B. Plancher, S. M. Neuman, R. Ghosal, S. Kuindersma, V. Janapa Reddi.
GRiD: GPU-Accelerated Rigid Body Dynamics with Analytical Gradients. ICRA, 2022.

3,5 “Widgetism”: Avoid Over-Specialization Chips and Salsa: Acceleration Beyond ASICs

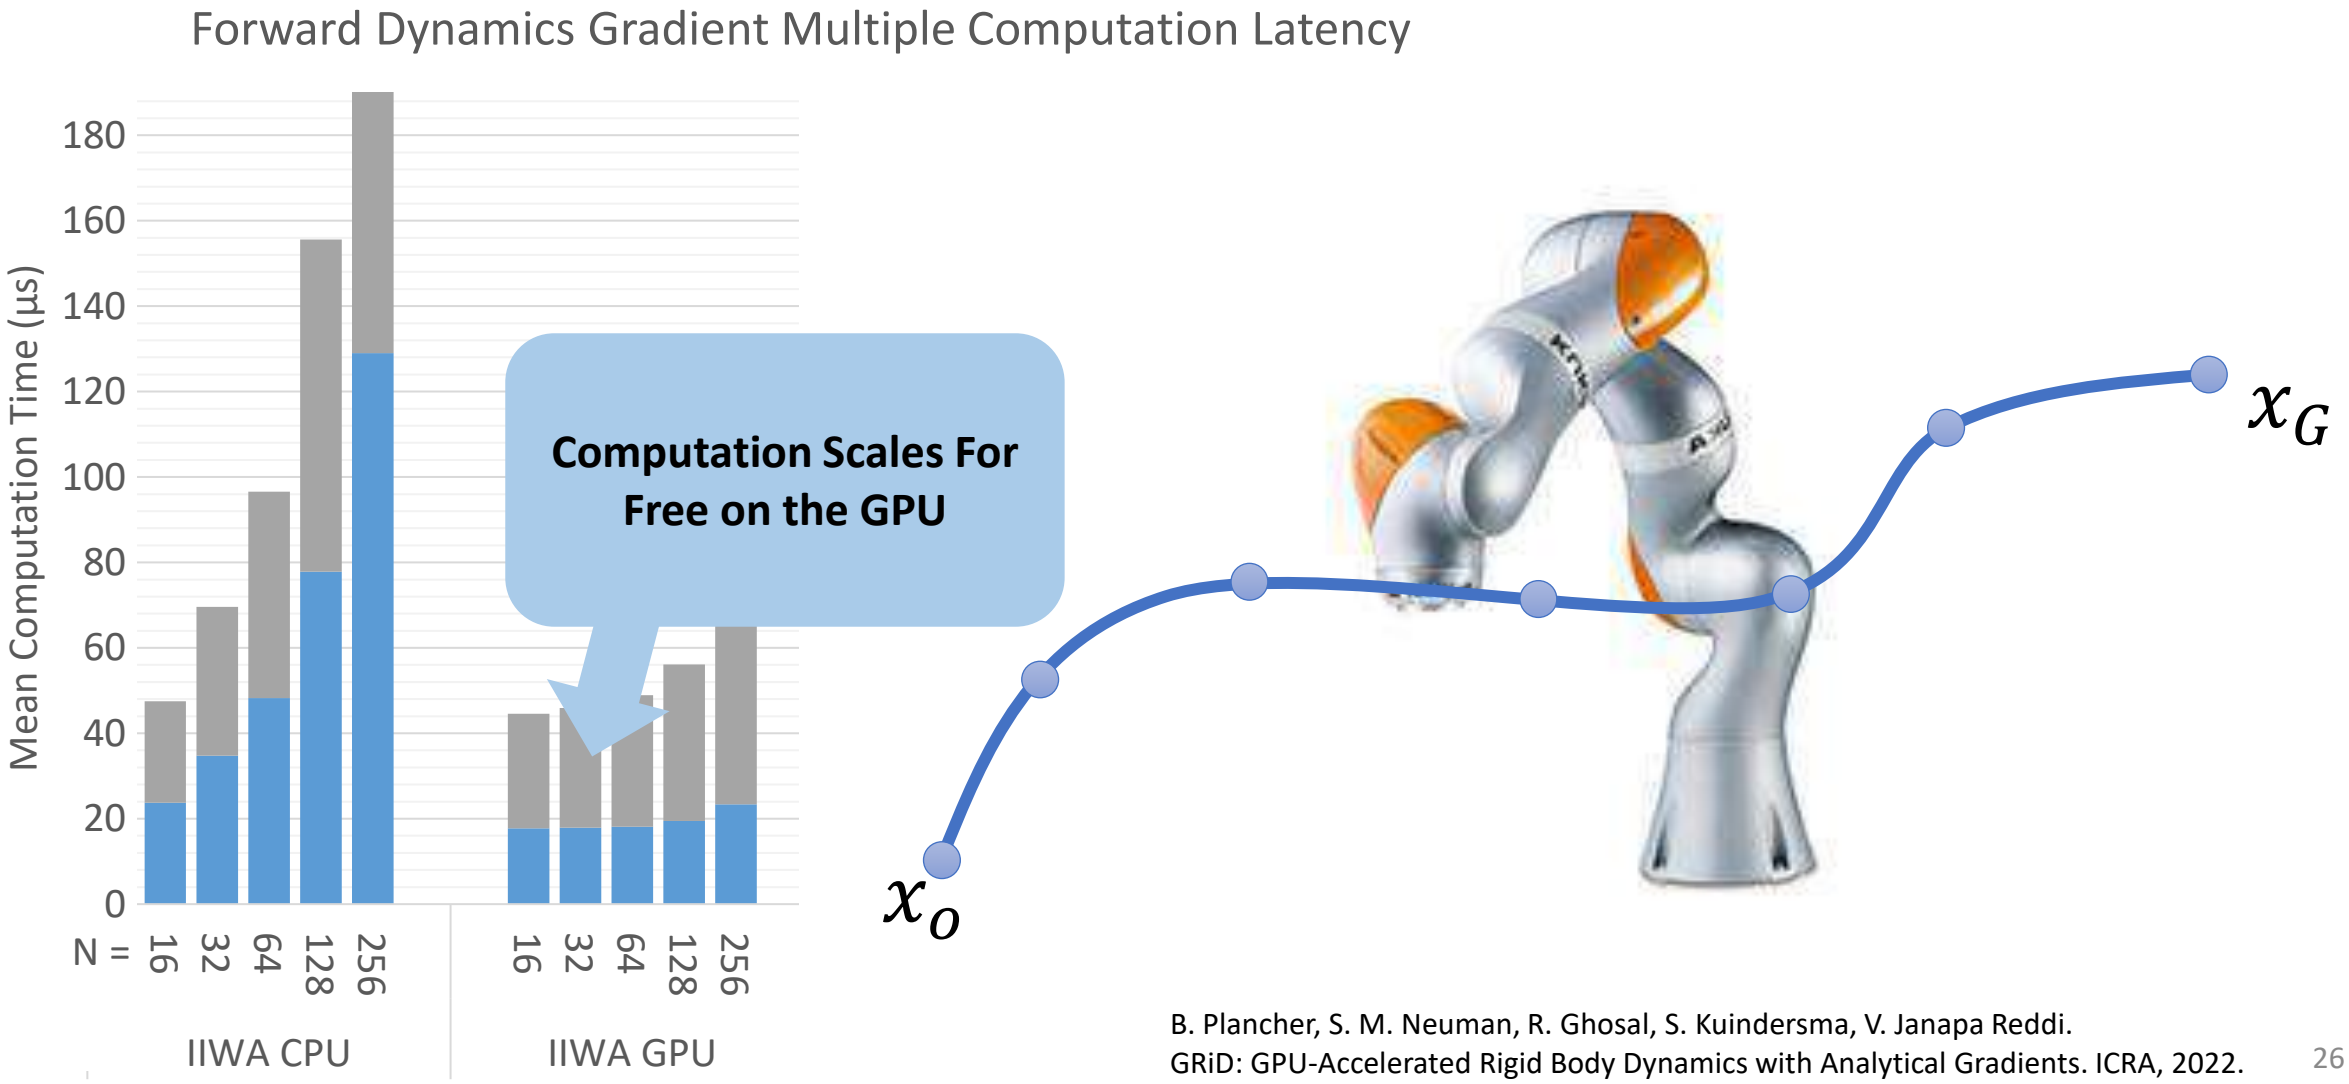

B. Plancher, S. M. Neuman, R. Ghosal, S. Kuindersma, V. Janapa Reddi.
GRiD: GPU-Accelerated Rigid Body Dynamics with Analytical Gradients. ICRA, 2022.

5 “Widgetism”: Avoid Over-Specialization
Chips and Salsa: Acceleration Beyond ASICs

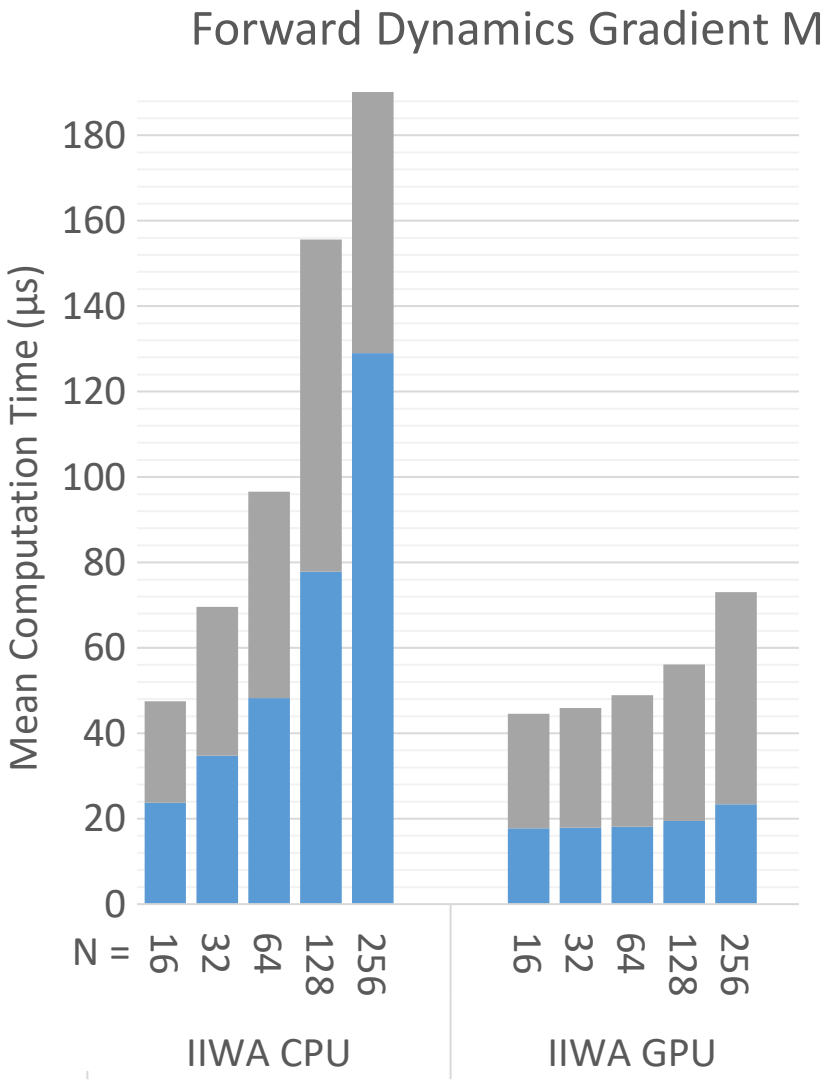

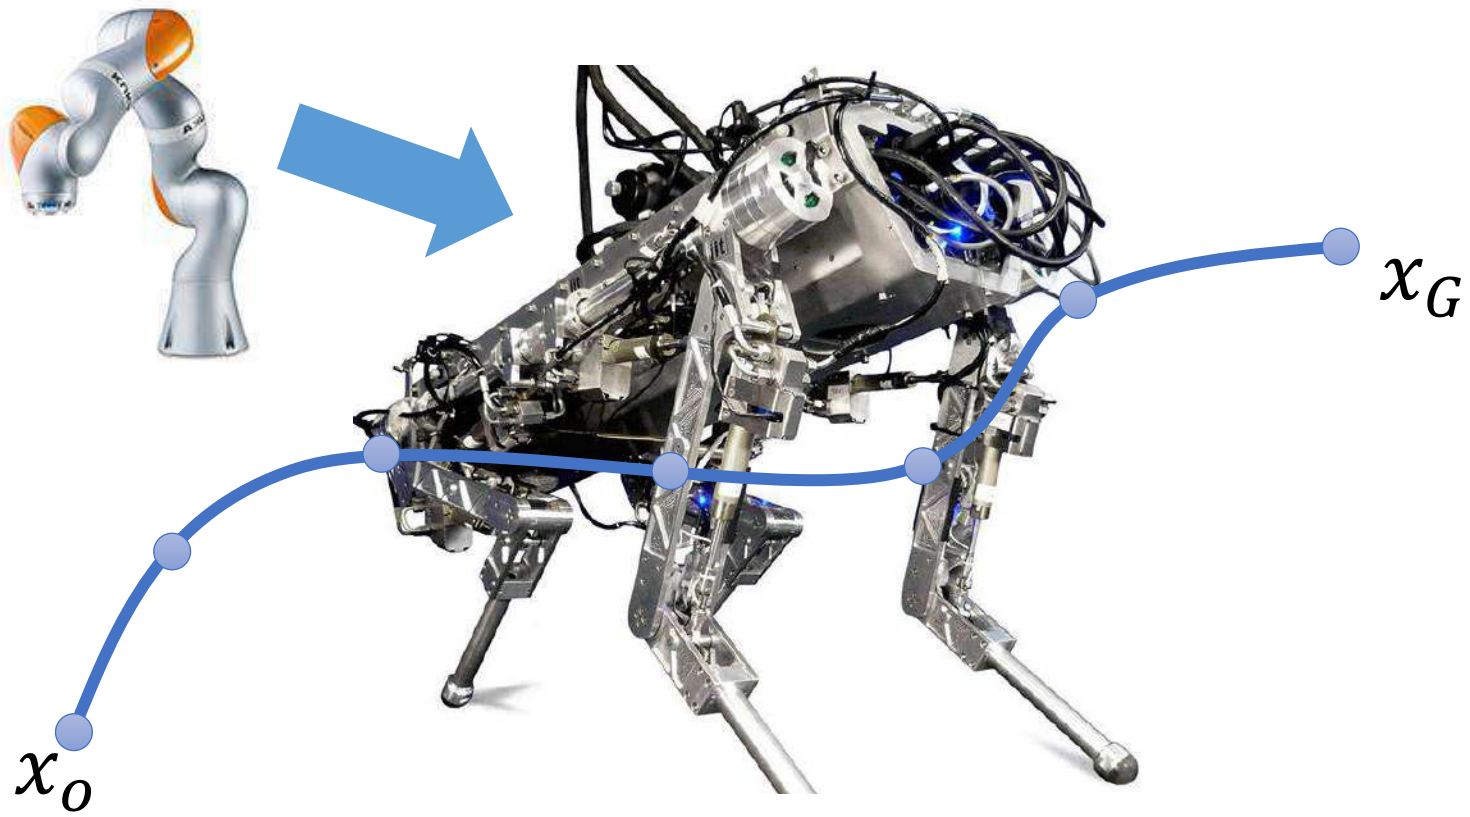

B. Plancher, S. M. Neuman, R. Ghosal, S. Kuindersma, V. Janapa Reddi.
GRiD: GPU-Accelerated Rigid Body Dynamics with Analytical Gradients. ICRA, 2022.

3,5

“Widgetism”: Avoid Over-Specialization
Chips and Salsa: Acceleration Beyond ASICs

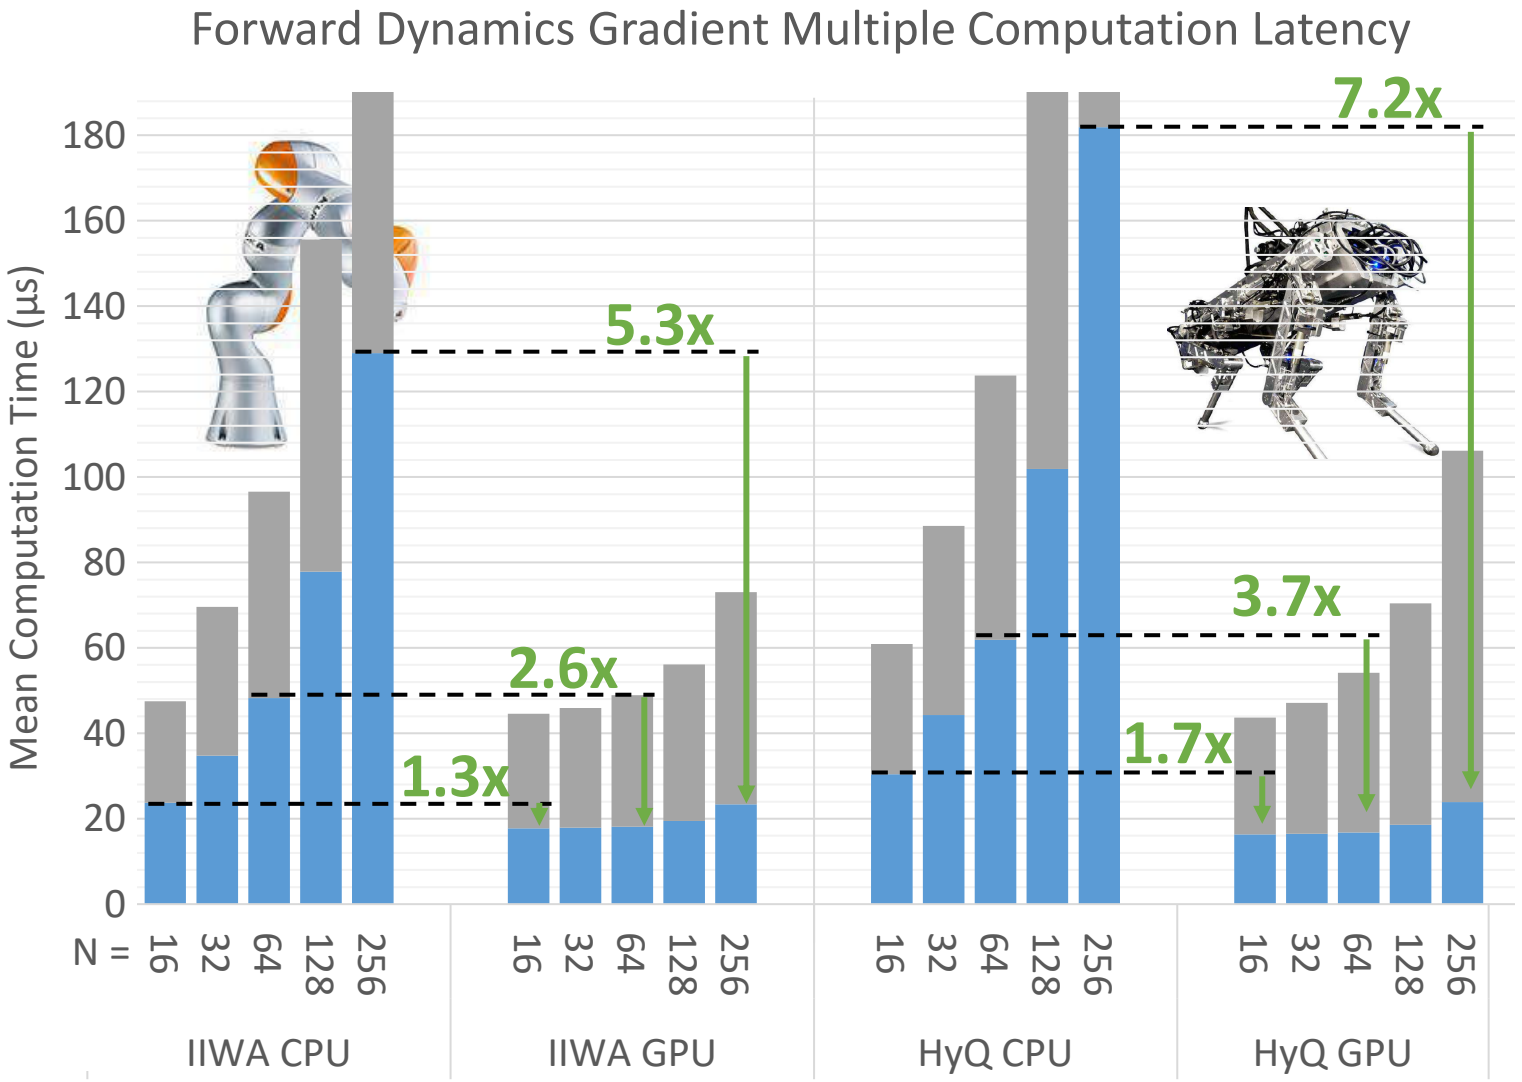

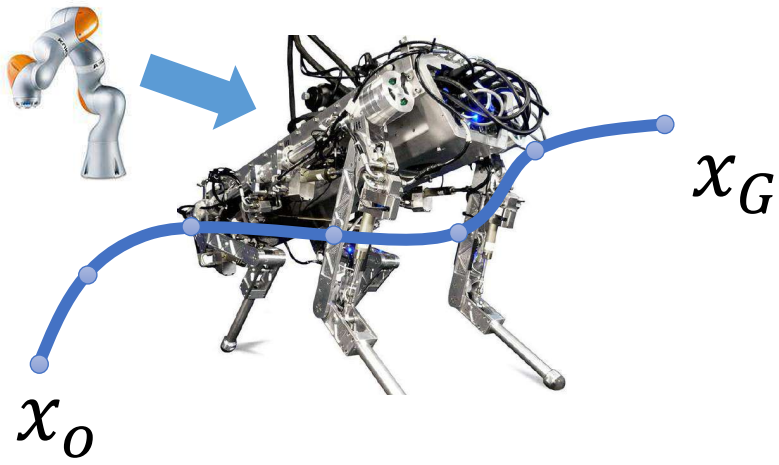

3,5 “Widgetism”: Avoid Over-Specialization
Chips and Salsa: Acceleration Beyond ASICs

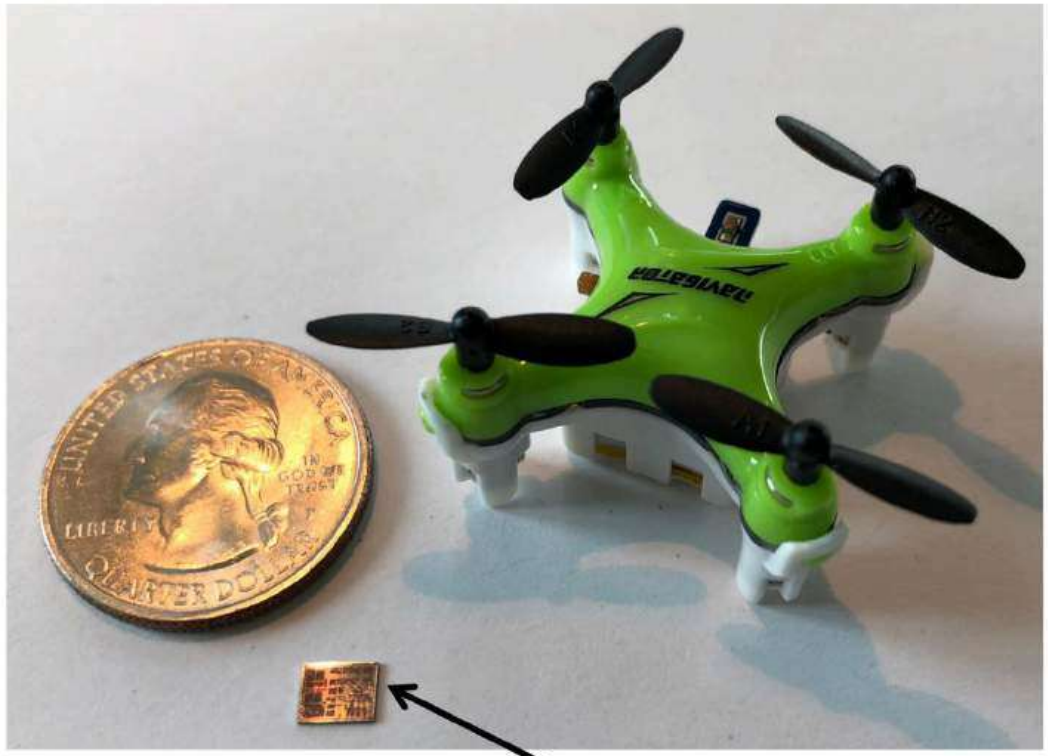

Visual-Inertial Odometry on Chip: An Algorithm-and-Hardware Co-design Approach

Zhengdong Zhang*, Amr Suleiman*, Luca Carlone, Vivienne Sze, Sertac Karaman
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
Emails: {zhangzd,suleiman,lcarlone,sze,sertac}@mit.edu, **Website: http://navion.mit.edu**
**These authors contributed equally to this work*

Navion: A 2mW Fully Integrated Real-Time Visual-Inertial Odometry Accelerator for Autonomous Navigation of Nano Drones

Amr Suleiman, *Member, IEEE,* Zhengdong Zhang, *Student Member, IEEE,* Luca Carlone, *Member, IEEE*
Sertac Karaman, *Member, IEEE* and Vivienne Sze, *Senior Member, IEEE*

3,5 “Widgetism”: Avoid Over-Specialization
Chips and Salsa: Acceleration Beyond ASICs

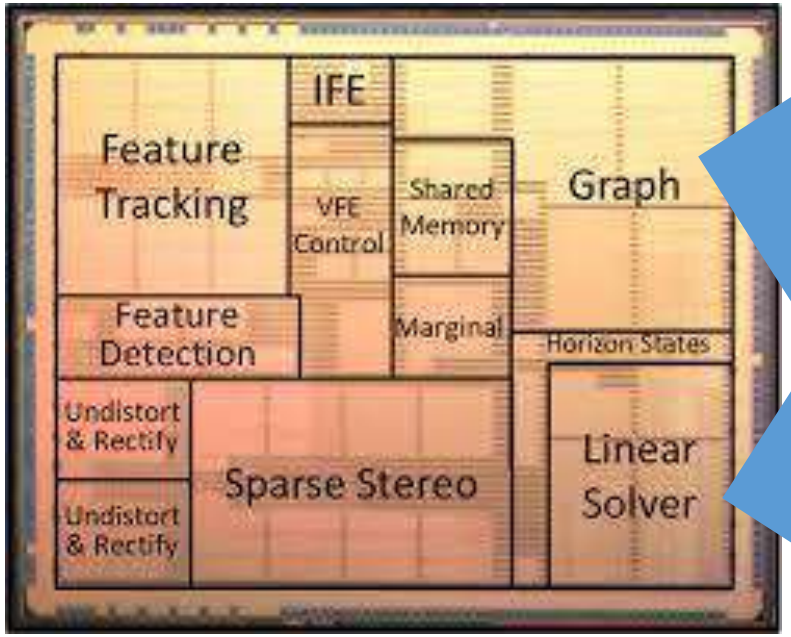

Fundamental Graph Operations and Linear Algebra **WILL** be highly portable!

A. Suleiman, Z. Zhang, L. Carlone, S. Karaman, V. Sze. Navion: A 2mW Fully Integrated Real-Time Visual-Inertial Odeometry Accelerator for Autonomous Navigation of Nano Drones. JSSC. 2019.

3,5

“Widgetism”: Avoid Over-Specialization
Chips and Salsa: Acceleration Beyond ASICs

Transition to industry and real impact is a challenge!

GPUs in robots is a **VERY NEW** thing!

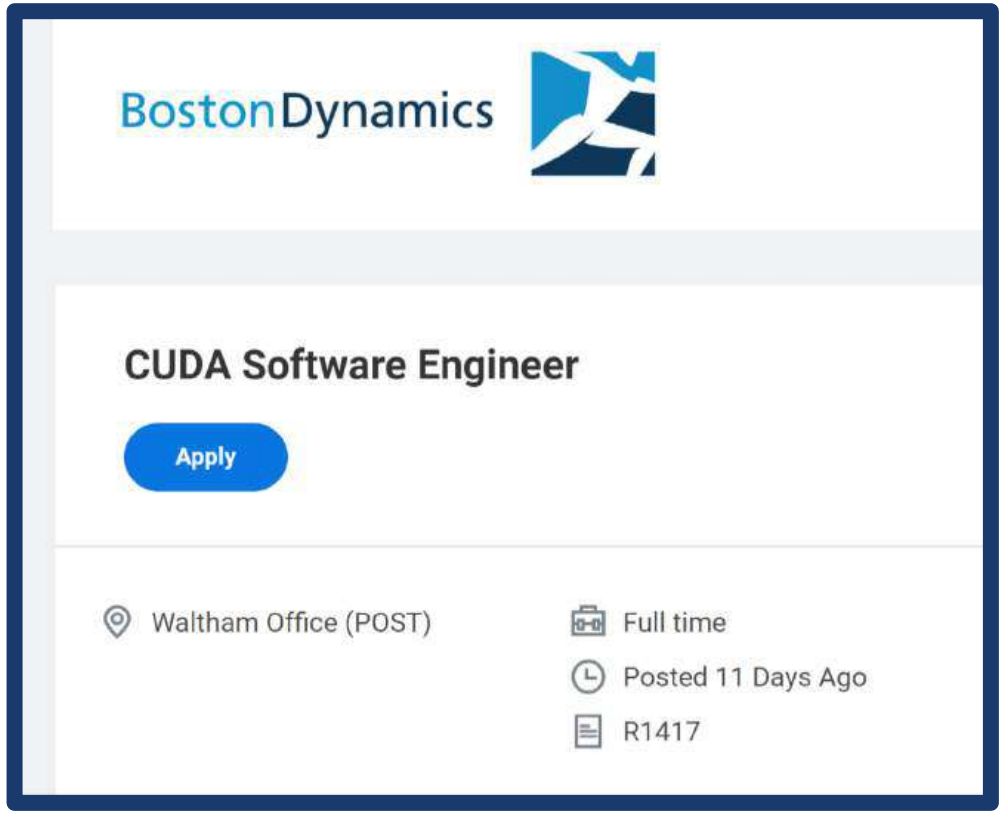

The Magnificent Seven
Challenges & Opportunities in Domain-Specific Accelerator Design for Autonomous Systems

1. **Build Bridges: Engage with Domain Experts**
2. Measure Twice, Cut Once: Metrics Matter
3. **“Widgetism”: Avoid Over-Specialization**
4. Pump the Brakes: Do Not Always Accelerate
5. **Chips and Salsa: Acceleration Beyond ASICs**
6. Forest vs. Trees: Take an End-to-End View
7. Design Global: Sustainability and Impact

4 Pump the Brakes: Do Not Always Accelerate

Pitfall: *Assume accelerators always improve total system performance.*

6 Forest vs. Trees: Take an End-to-End View

Pitfall: *A narrow scope: acceleration begins and ends with compute.*

4,6 Pump the Brakes: Do Not Always Accelerate
Forest vs. Trees: Take an End-to-End View

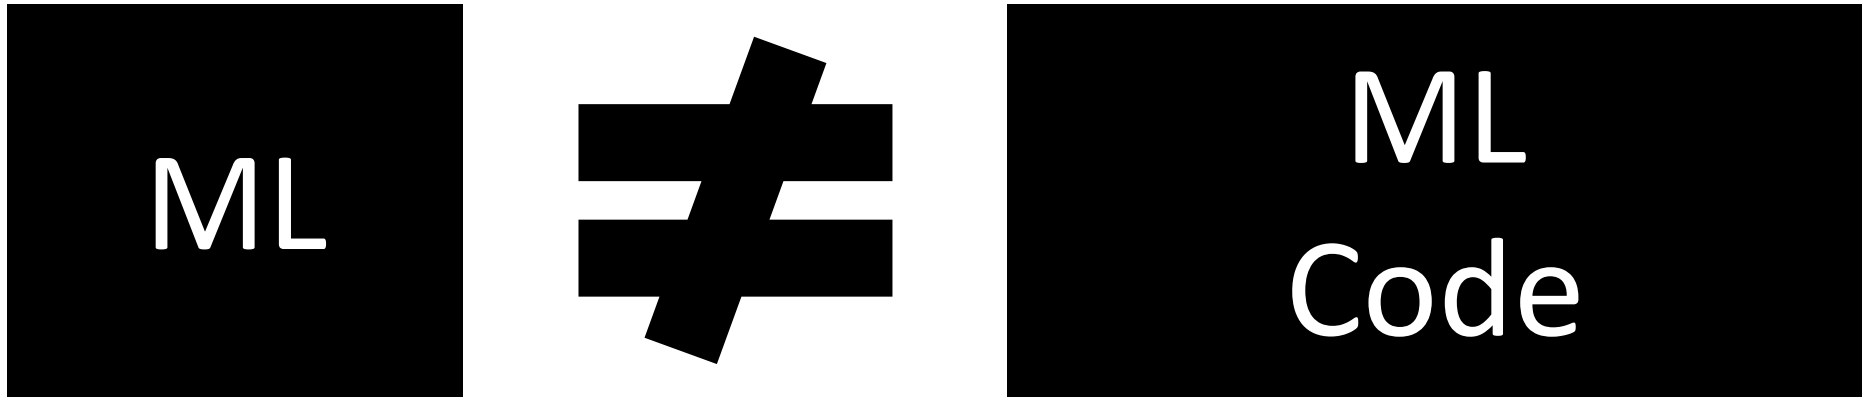

D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, D. Dennison. Hidden Technical Debt in Machine Learning Systems. NeurIPS, 2015.

4,6 Pump the Brakes: Do Not Always Accelerate Forest vs. Trees: Take an End-to-End View

D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, D. Dennison. Hidden Technical Debt in Machine Learning Systems. NeurIPS, 2015.

4,6 Pump the Brakes: Do Not Always Accelerate
Forest vs. Trees: Take an End-to-End View

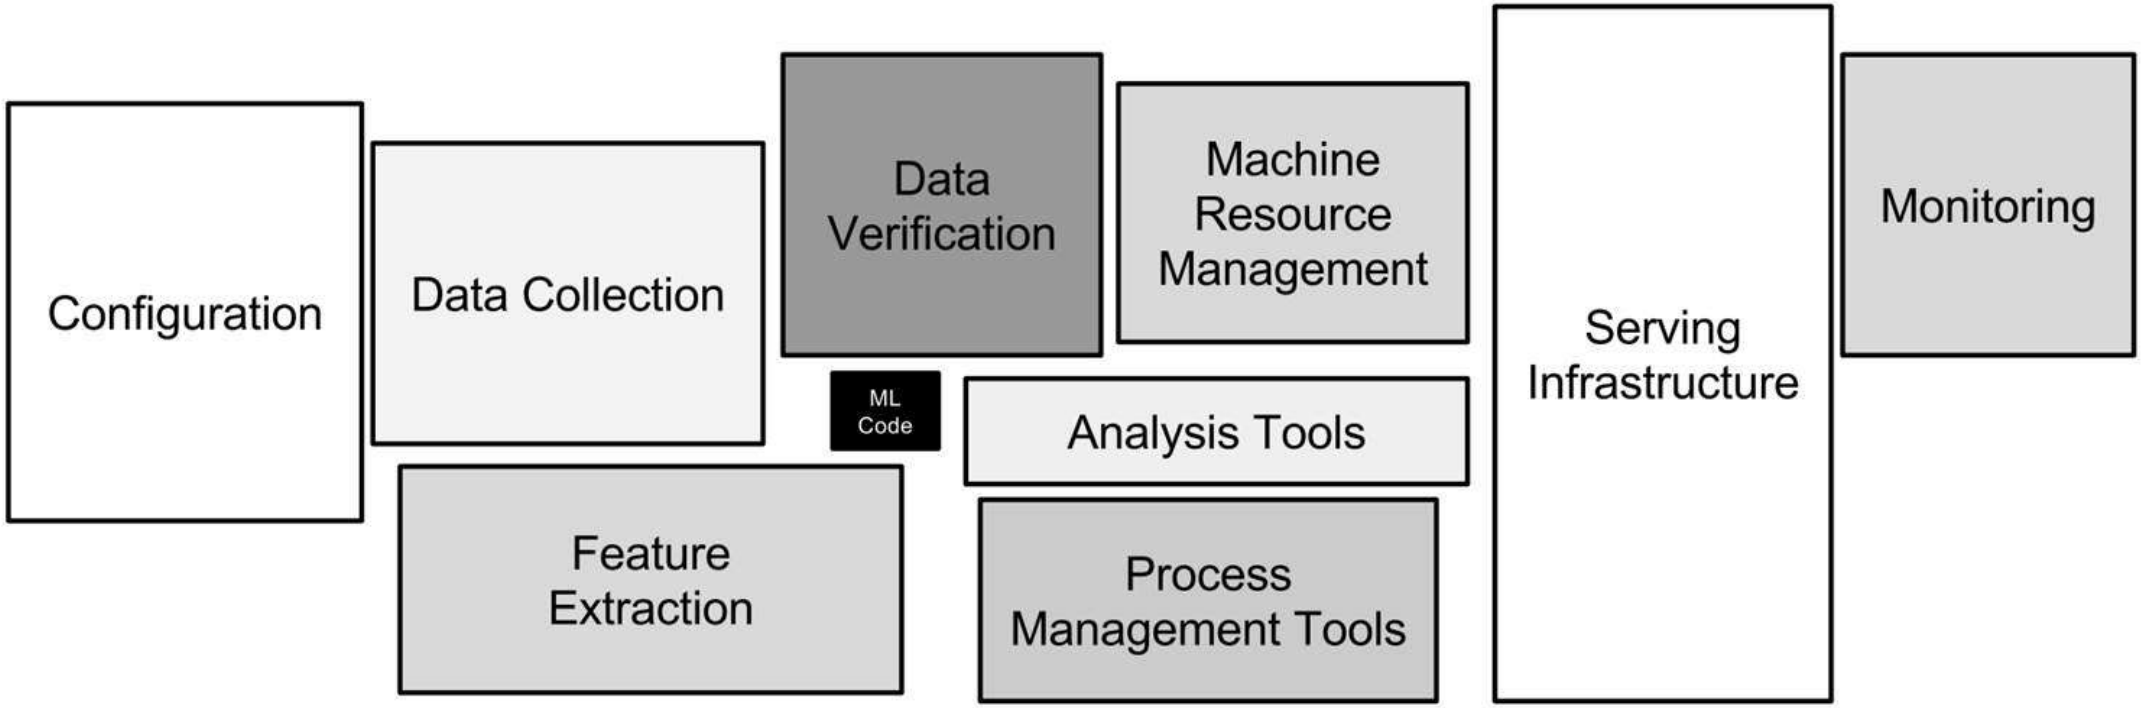

D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, D. Dennison. Hidden Technical Debt in Machine Learning Systems. NeurIPS, 2015.

Pump the Brakes: Do Not Always Accelerate
Forest vs. Trees: Take an End-to-End View

Amdahl's Law

$$\text{speedup}(f,n) = \frac{1}{(1-f)+{}^{f}/_{n}}$$

f = fraction of the program that is parallelizable
n = parallel processors

4 Pump the Brakes: Do Not Always Accelerate
Forest vs. Trees: Take an End-to-End View

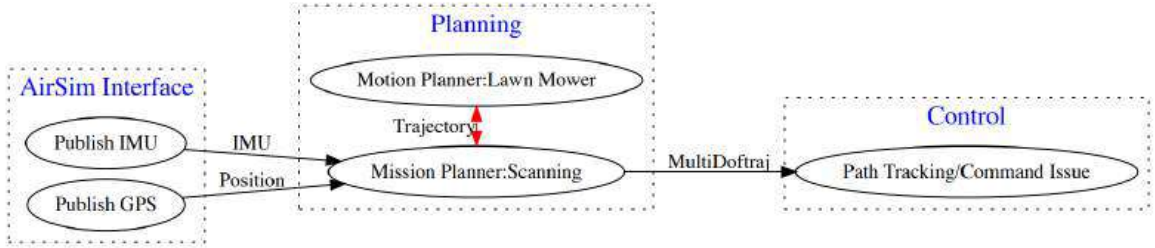

(a) Scanning.

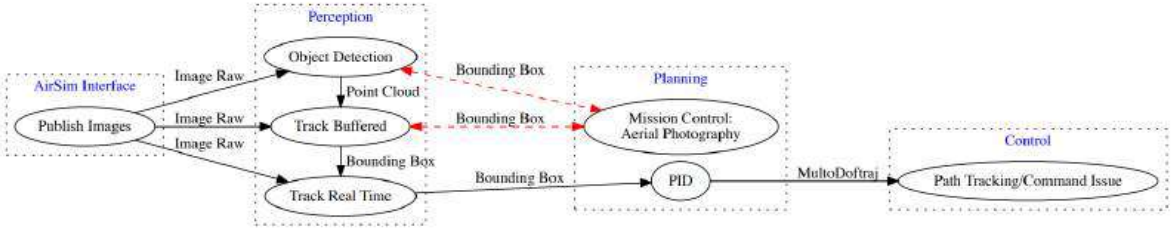

(b) Aerial Photography.

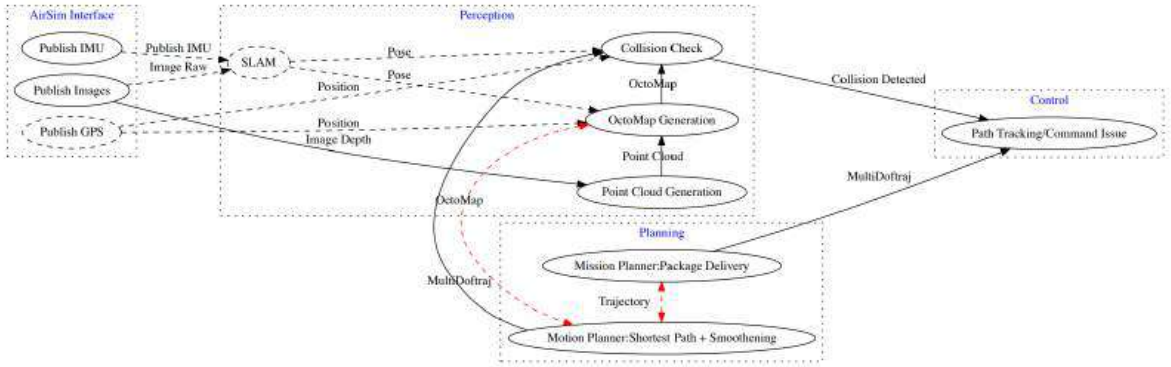

(c) Package Delivery.

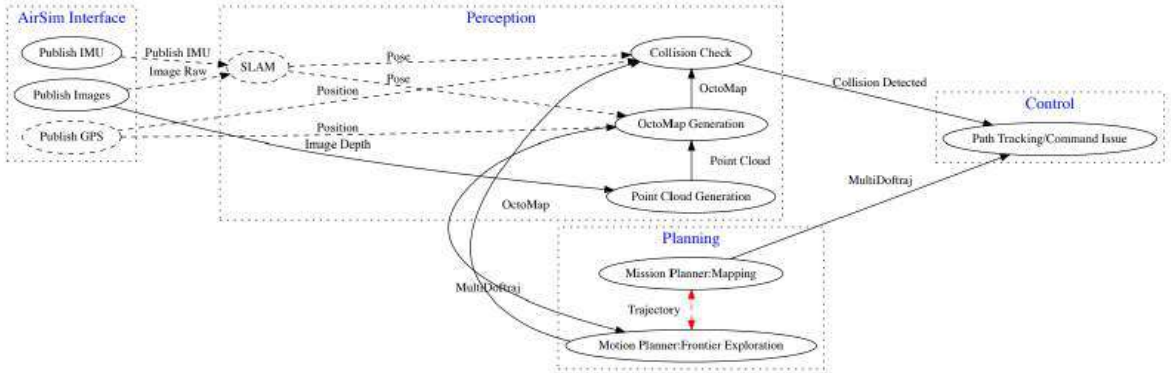

(d) 3D Mapping.

B. Boroujerdian, H. Genc, S. Krishnan, W. Cui, A. Faust, V. Janapa Reddi. MAVBench: Micro Aerial Vehicle Benchmarking. MICRO. 2018.

6 Pump the Brakes: Do Not Always Accelerate Forest vs. Trees: Take an End-to-End View

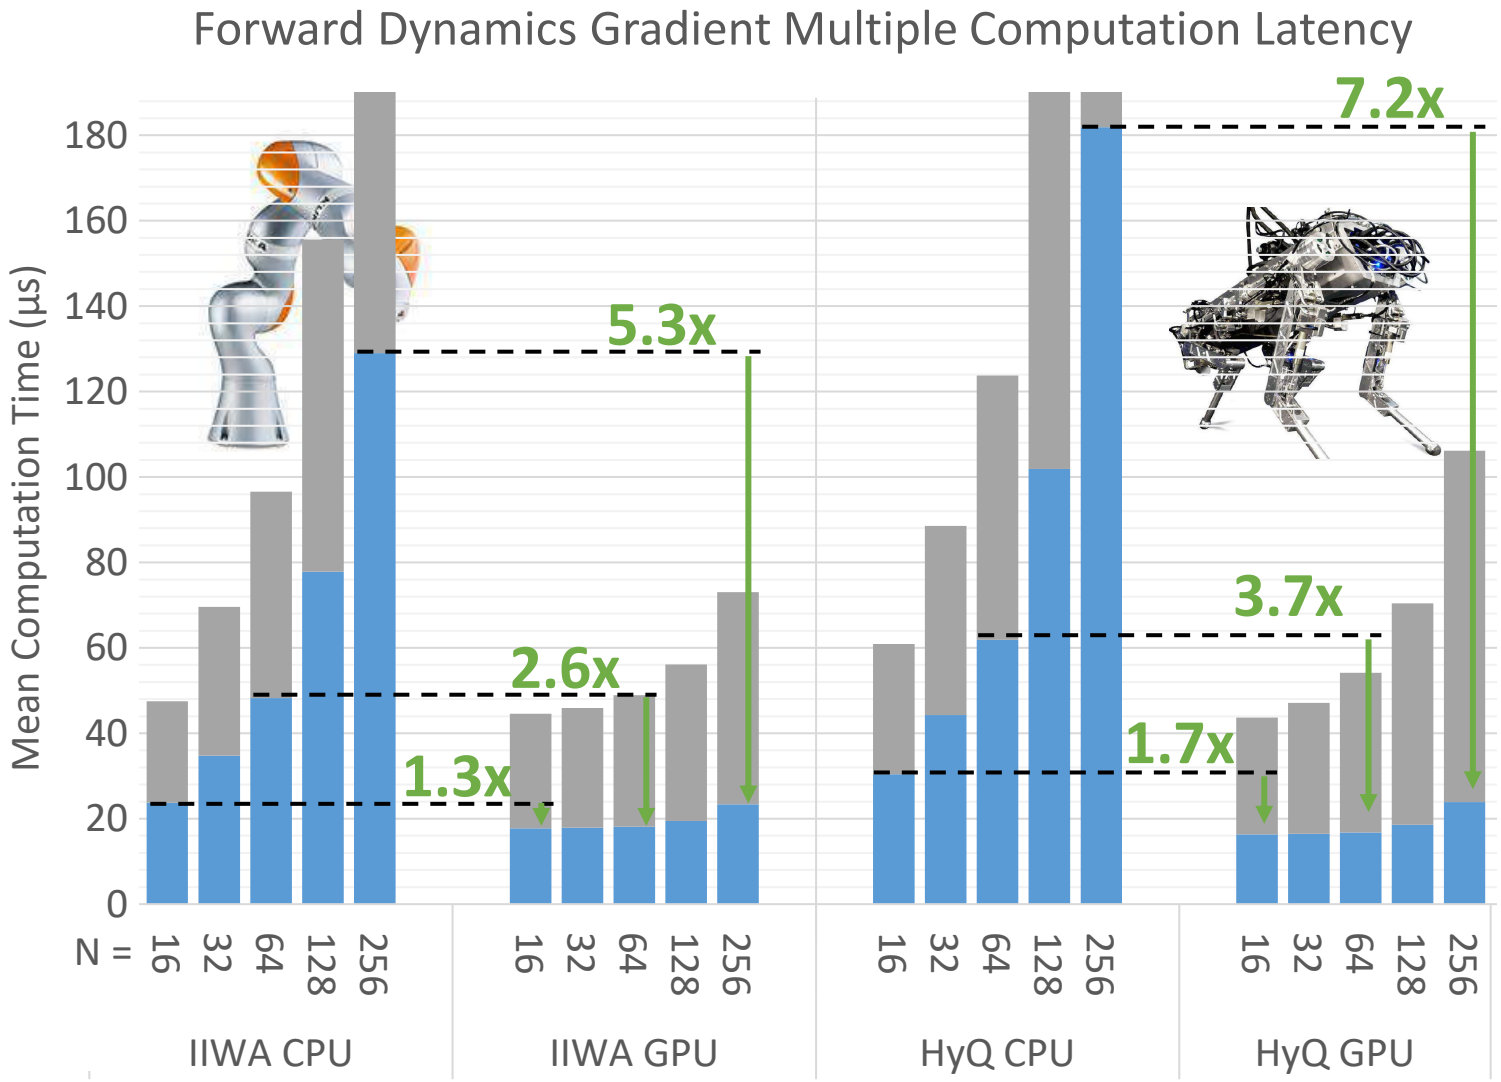

What about I/O?

6 Pump the Brakes: Do Not Always Accelerate Forest vs. Trees: Take an End-to-End View

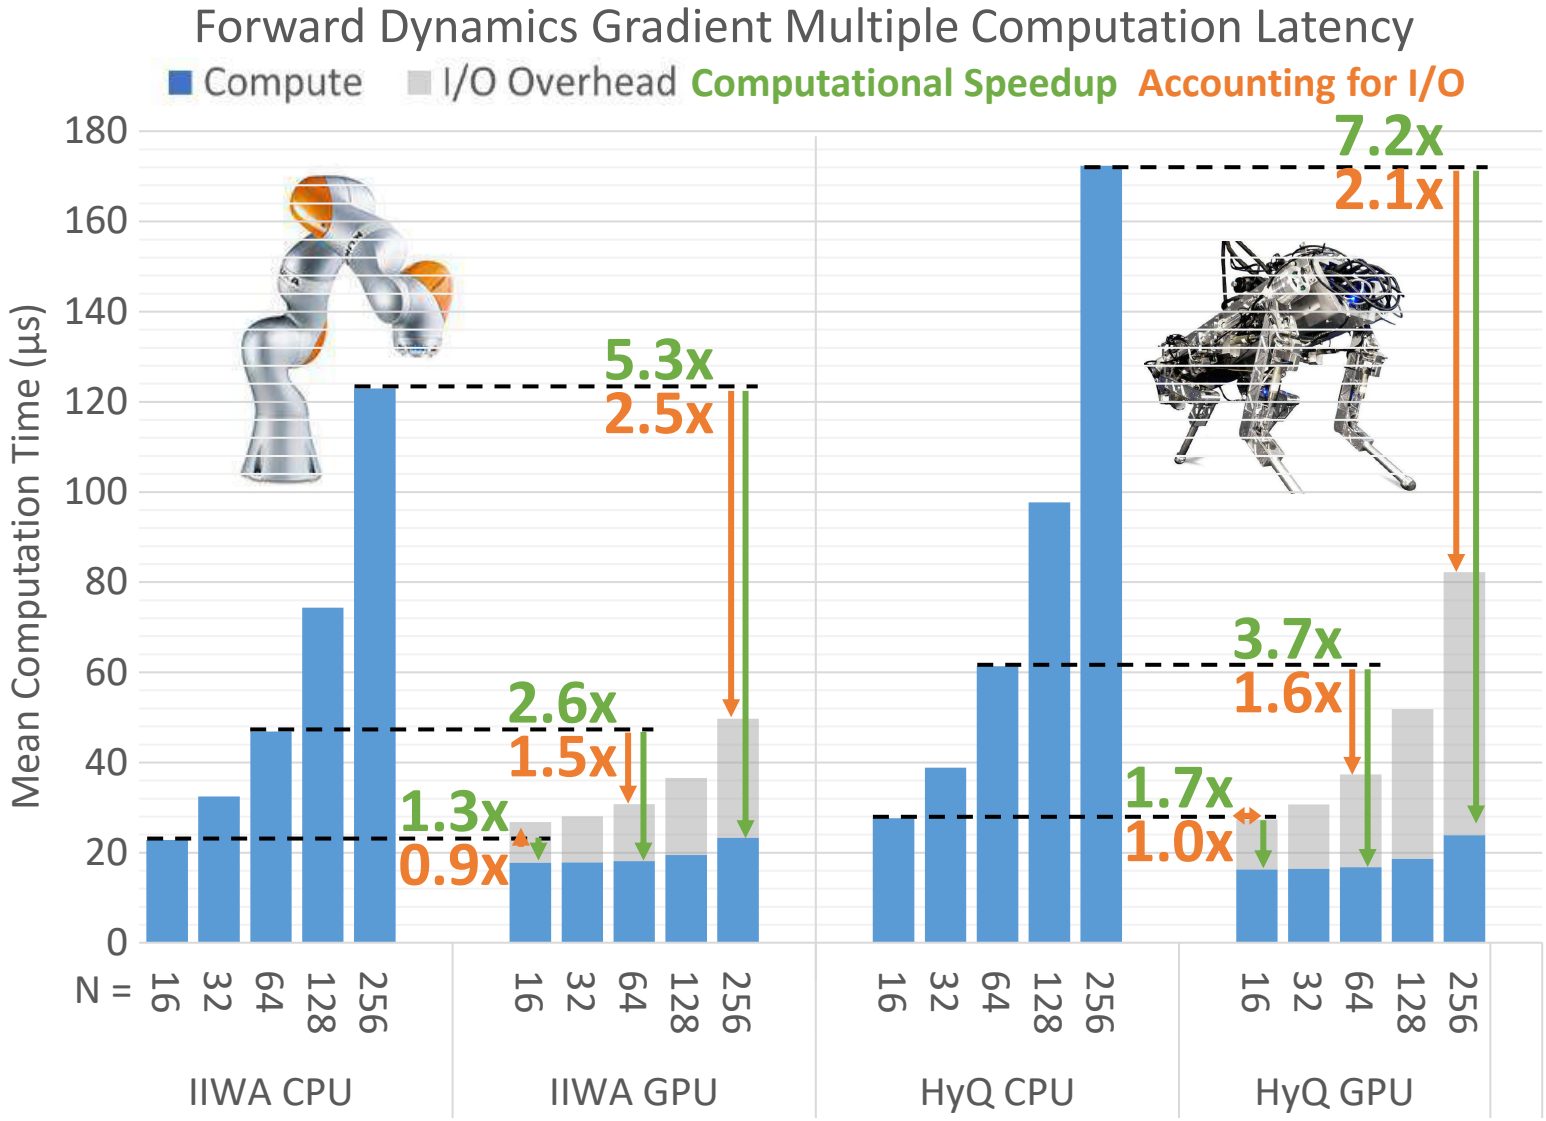

I/O Matters A LOT!

4 Pump the Brakes: Do Not Always Accelerate
Forest vs. Trees: Take an End-to-End View

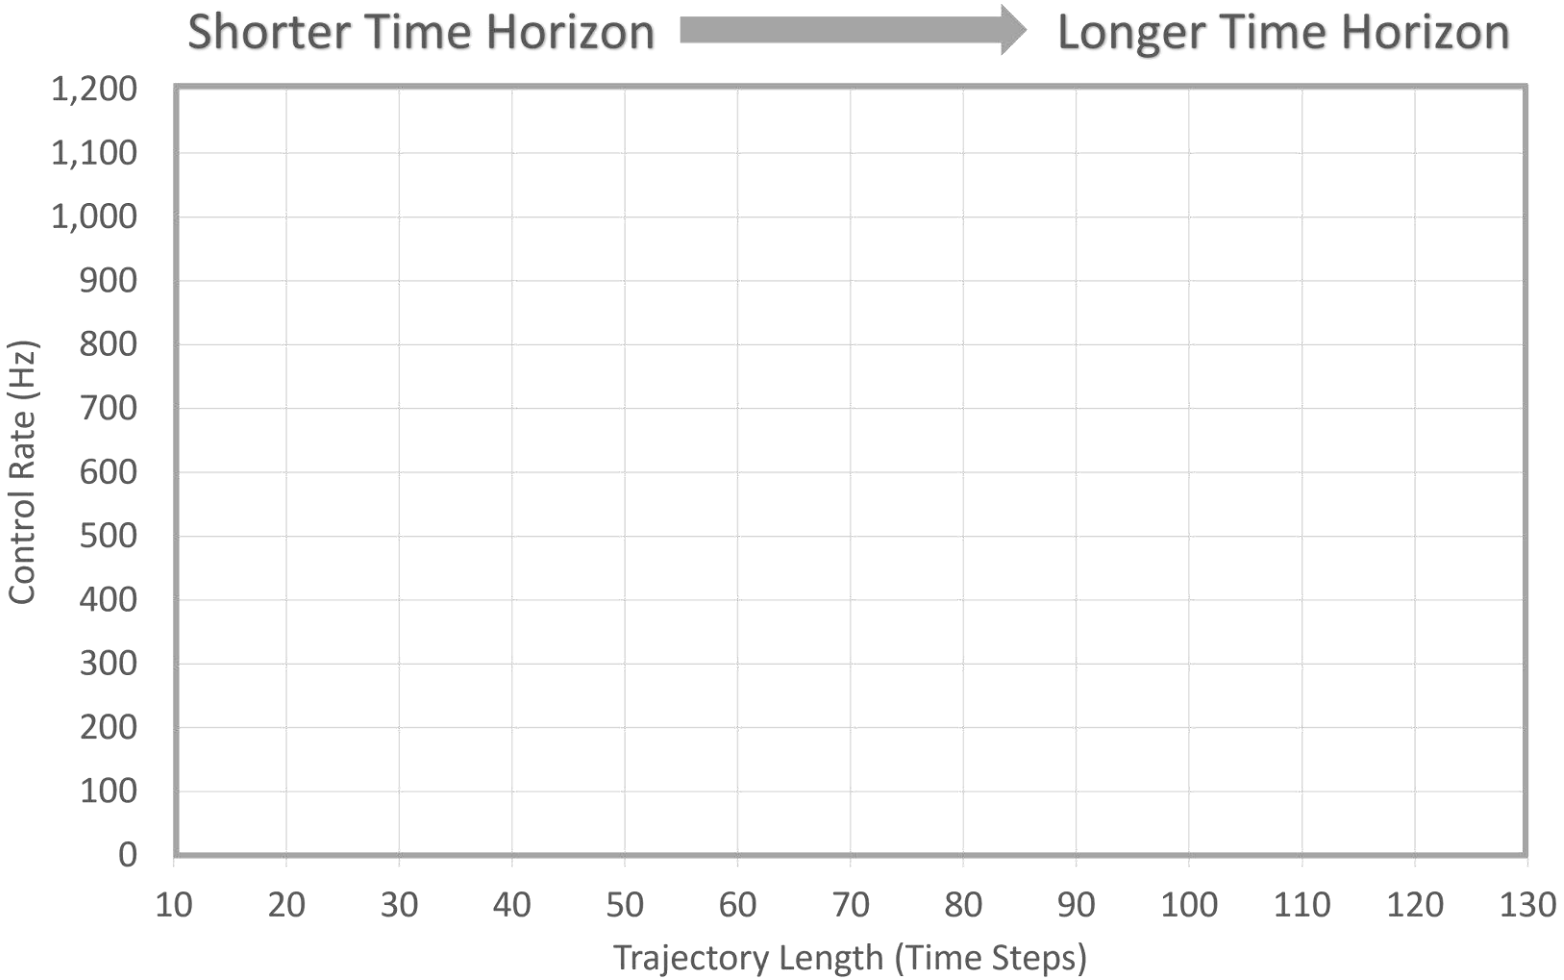

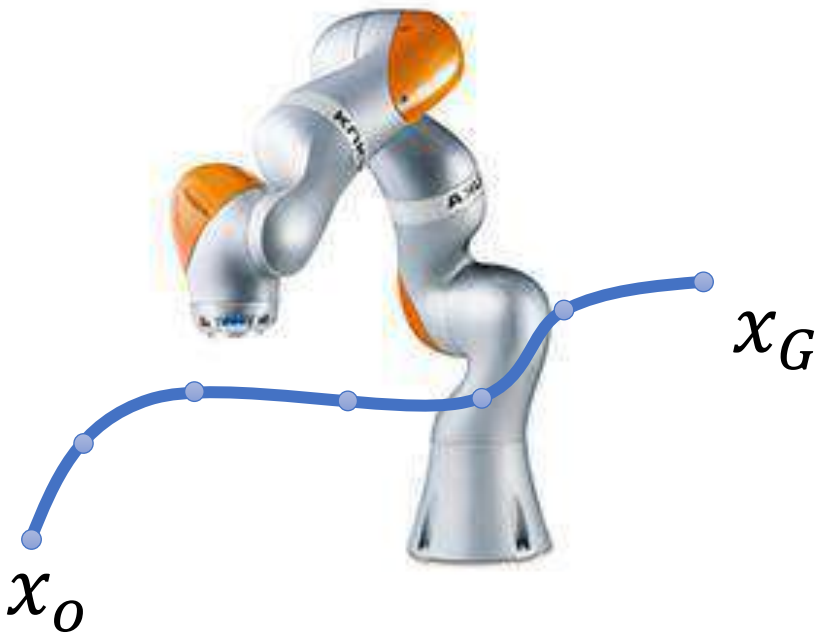

S. M. Neuman, B. Plancher, T. Bourgeat, T. Tambe, S. Devadas, V. Janapa Reddi. Robomorphic Computing:
A Design Methodology for Domain-Specific Accelerators Parameterized by Robot Morphology. ASPLOS 2021.

4 Pump the Brakes: Do Not Always Accelerate
Forest vs. Trees: Take an End-to-End View

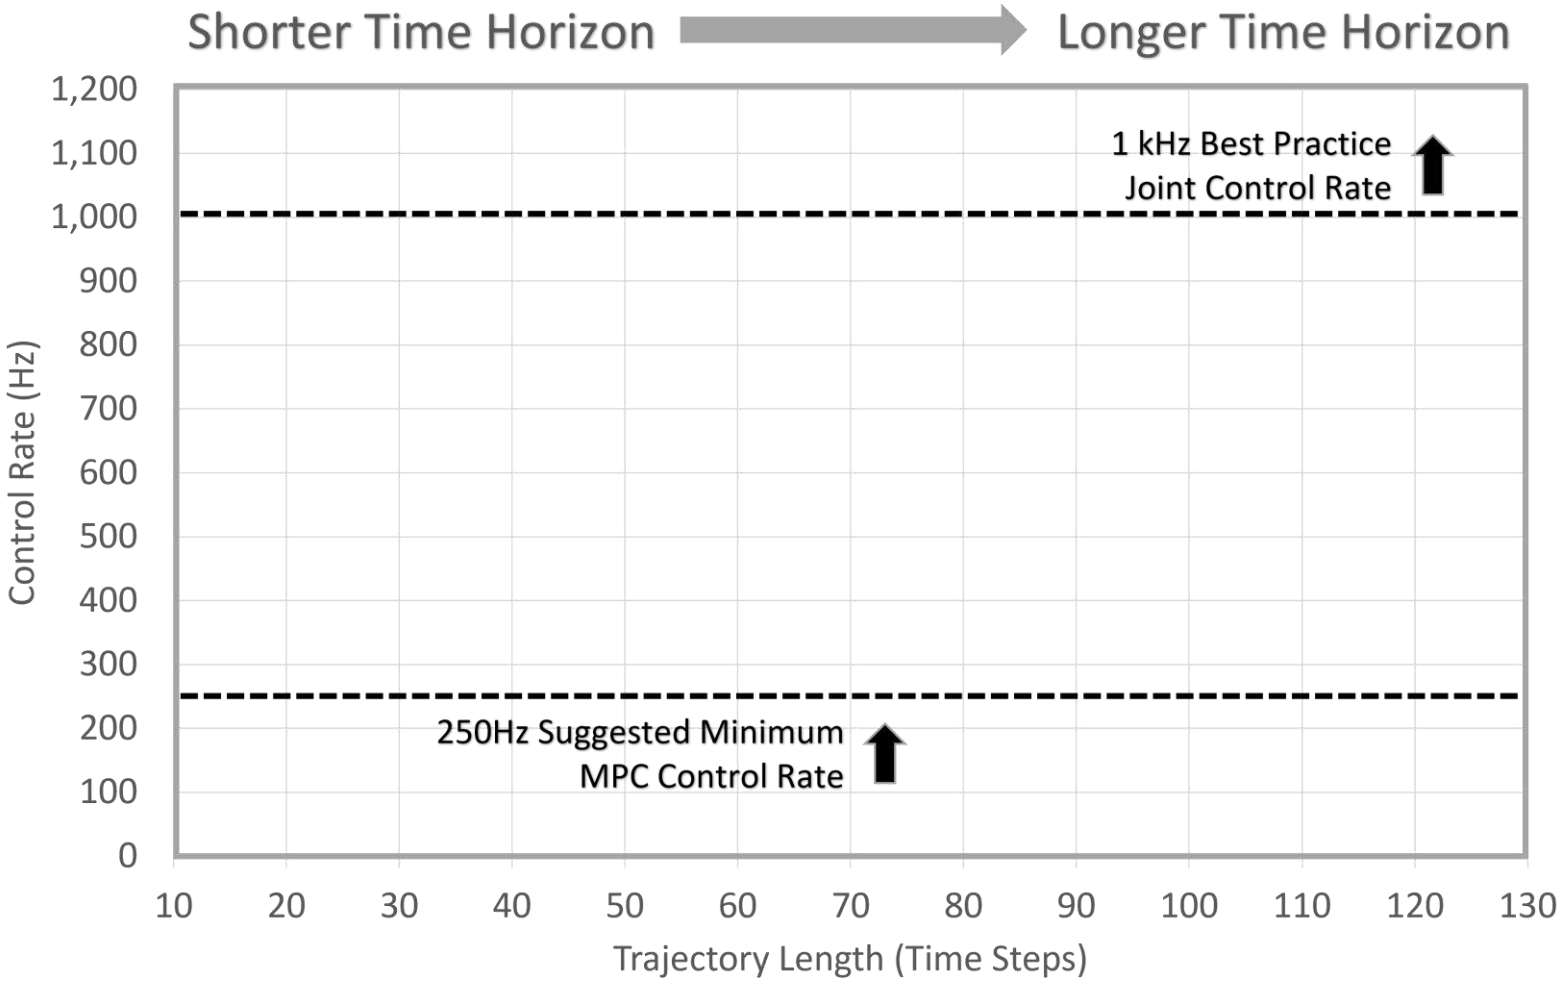

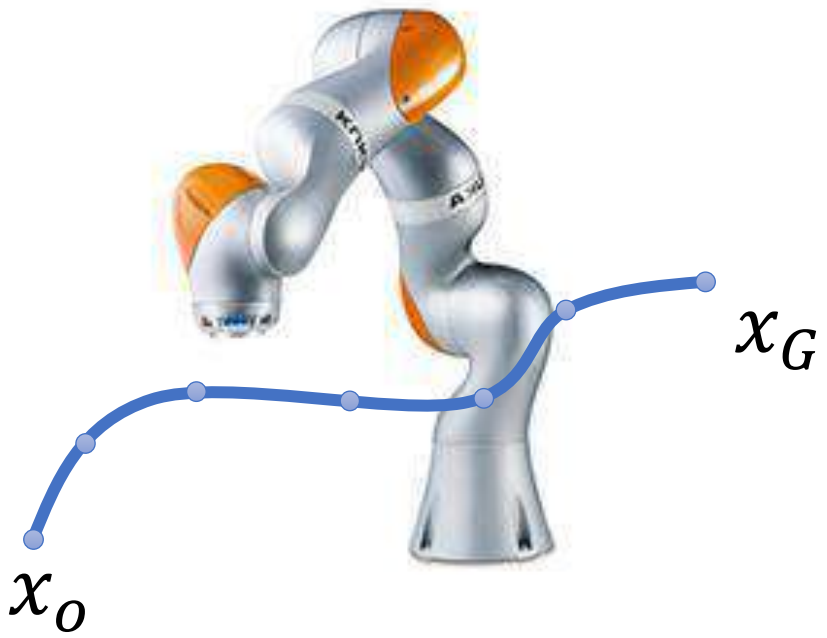

S. M. Neuman, B. Plancher, T. Bourgeat, T. Tambe, S. Devadas, V. Janapa Reddi. Robomorphic Computing:
A Design Methodology for Domain-Specific Accelerators Parameterized by Robot Morphology. ASPLOS 2021.

4 Pump the Brakes: Do Not Always Accelerate
Forest vs. Trees: Take an End-to-End View

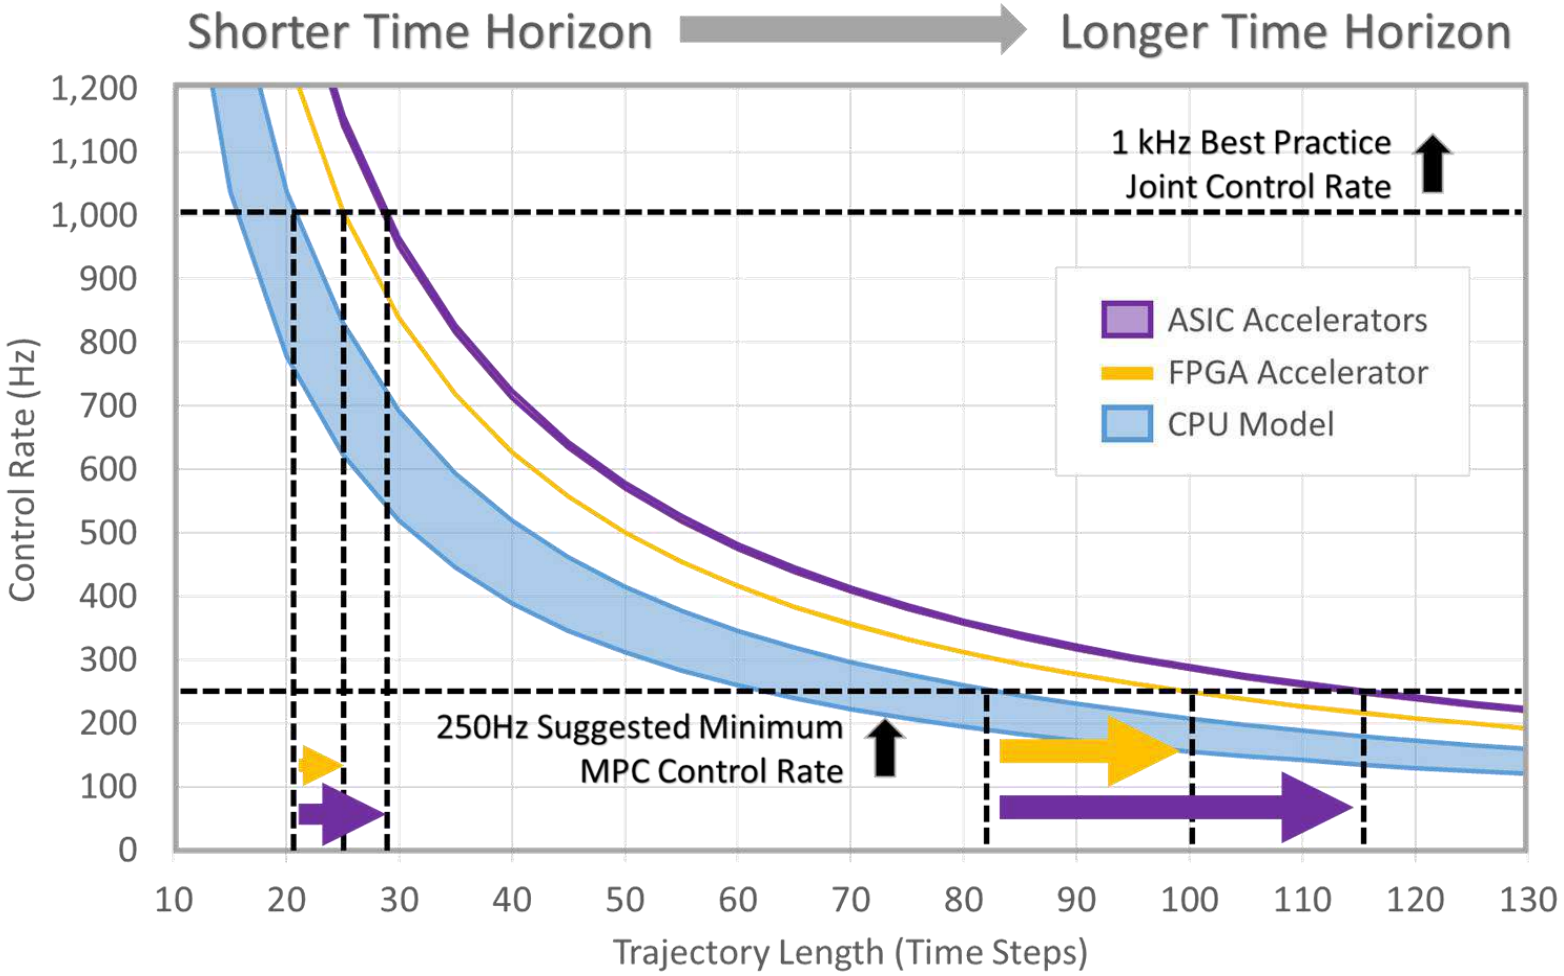

Does this matter?
How can we evaluate this impact?

S. M. Neuman, B. Plancher, T. Bourgeat, T. Tambe, S. Devadas, V. Janapa Reddi. Robomorphic Computing:
A Design Methodology for Domain-Specific Accelerators Parameterized by Robot Morphology. ASPLOS 2021.

2 Measure Twice, Cut Once: Metrics Matter

Pitfall: *Only focus on improving throughput or energy-delay product.*

2 Measure Twice, Cut Once: Metrics Matter

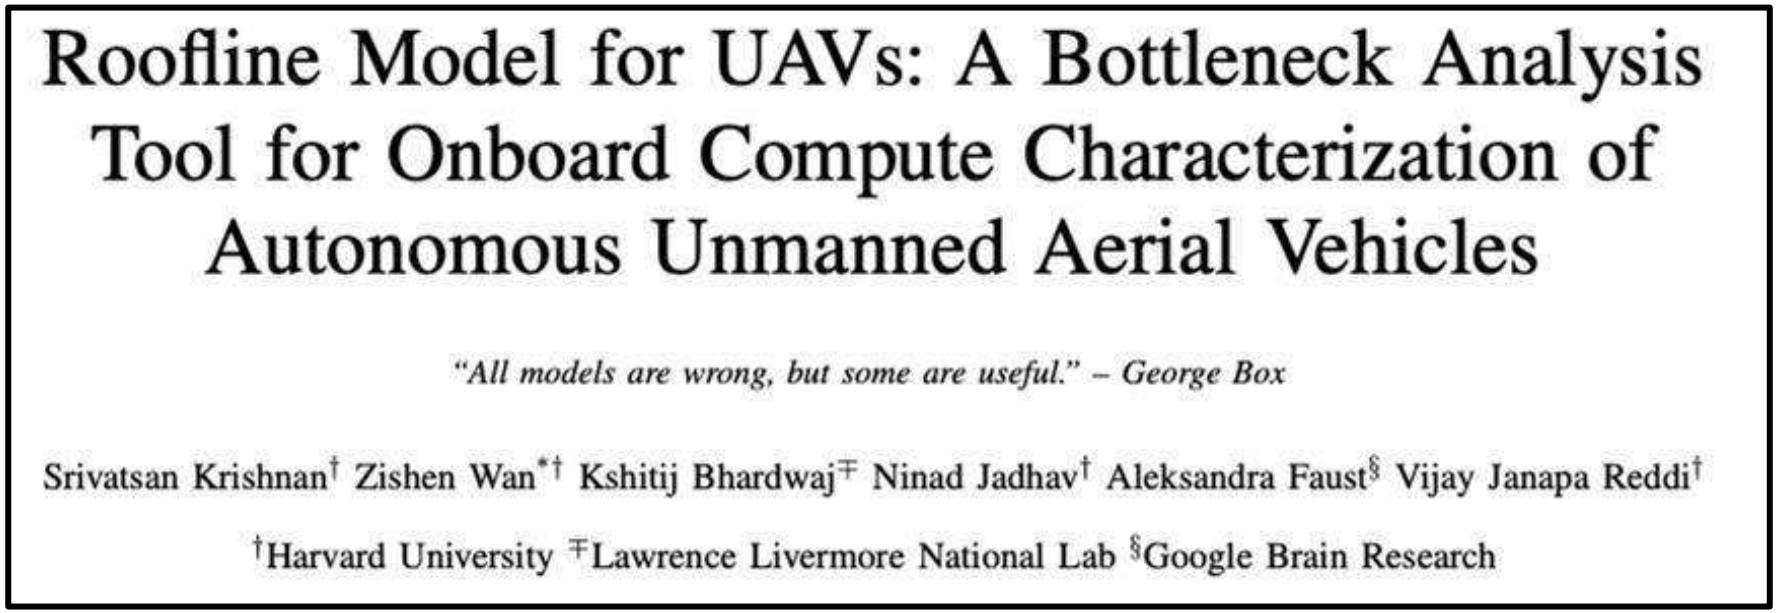

Roofline Model for UAVs: A Bottleneck Analysis Tool for Onboard Compute Characterization of Autonomous Unmanned Aerial Vehicles

"All models are wrong, but some are useful." – George Box

Srivatsan Krishnan[†] Zishen Wan[*†] Kshitij Bhardwaj[∓] Ninad Jadhav[†] Aleksandra Faust[§] Vijay Janapa Reddi[†]

[†]Harvard University [∓]Lawrence Livermore National Lab [§]Google Brain Research

S. Krishnan, Z. Wan, K. Bhardwaj, N. Jadhav, A. Faust, V. Janapa Reddi. Roofline Model for UAVs: A Bottleneck Analysis Tool for Onboard Compute Characterization of Autonomous Unmanned Aerial Vehicles. ASPLOS 2021.

2 Measure Twice, Cut Once: Metrics Matter

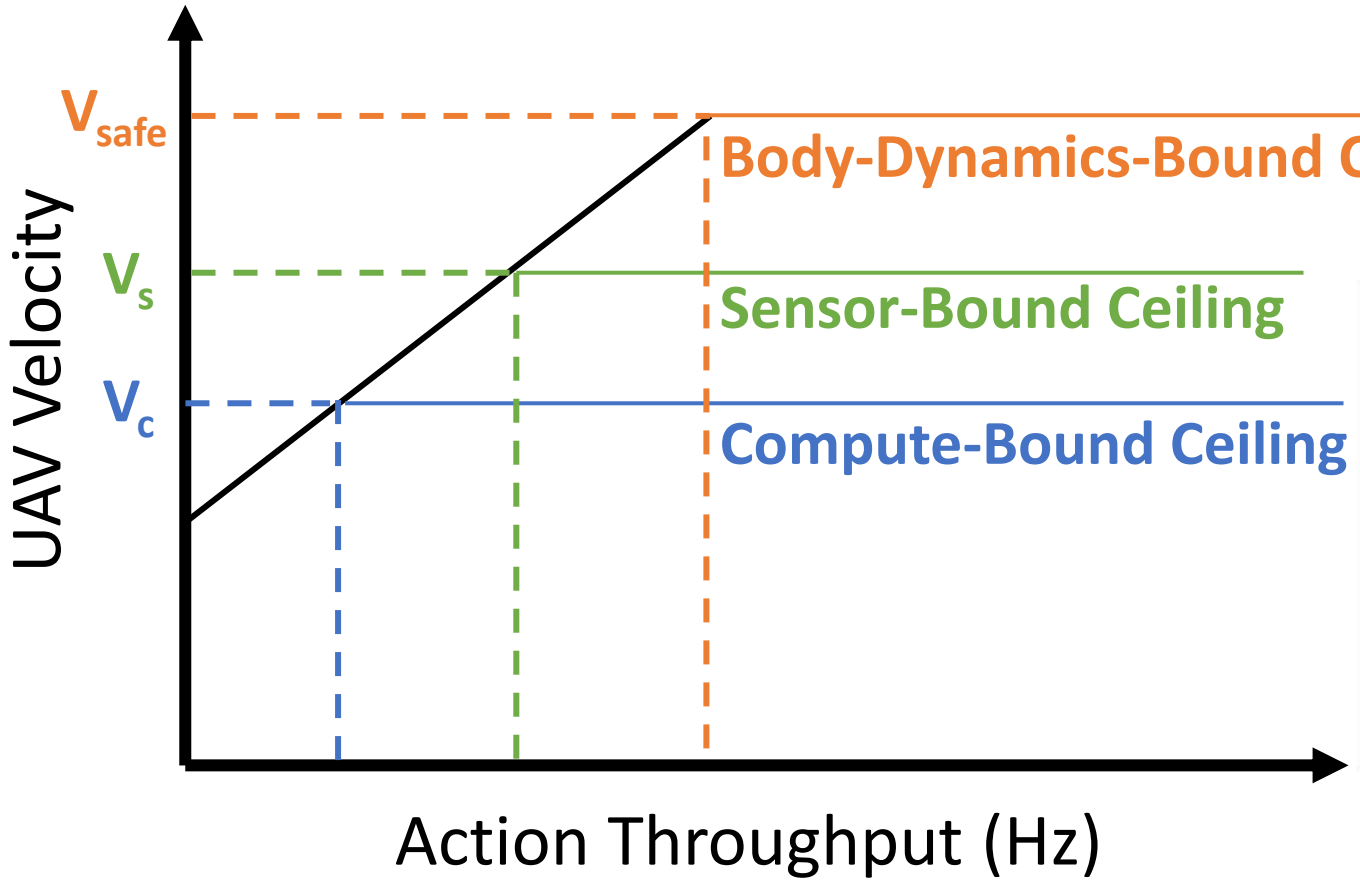

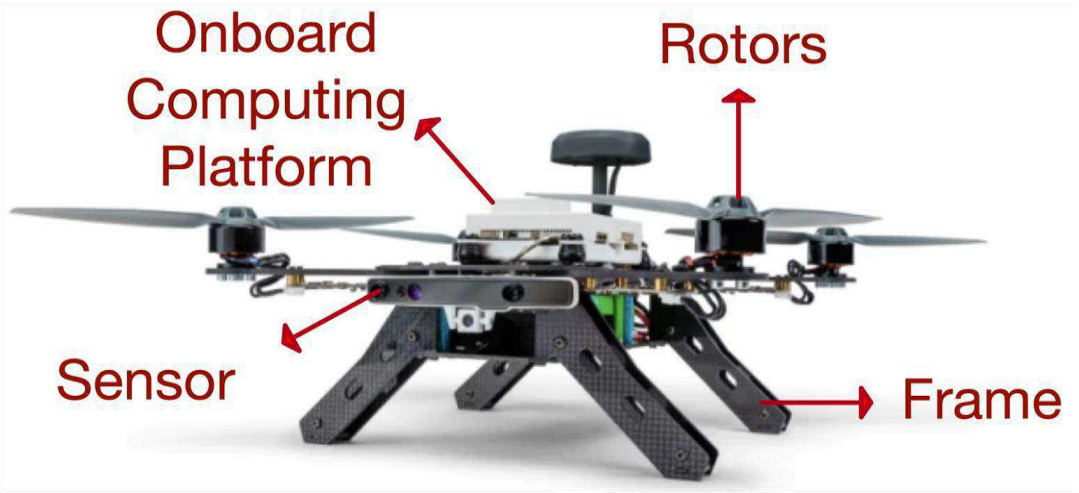

S. Krishnan, Z. Wan, K. Bhardwaj, N. Jadhav, A. Faust, V. Janapa Reddi. Roofline Model for UAVs: A Bottleneck Analysis Tool for Onboard Compute Characterization of Autonomous Unmanned Aerial Vehicles. ASPLOS 2021.

2 Measure Twice, Cut Once: Metrics Matter

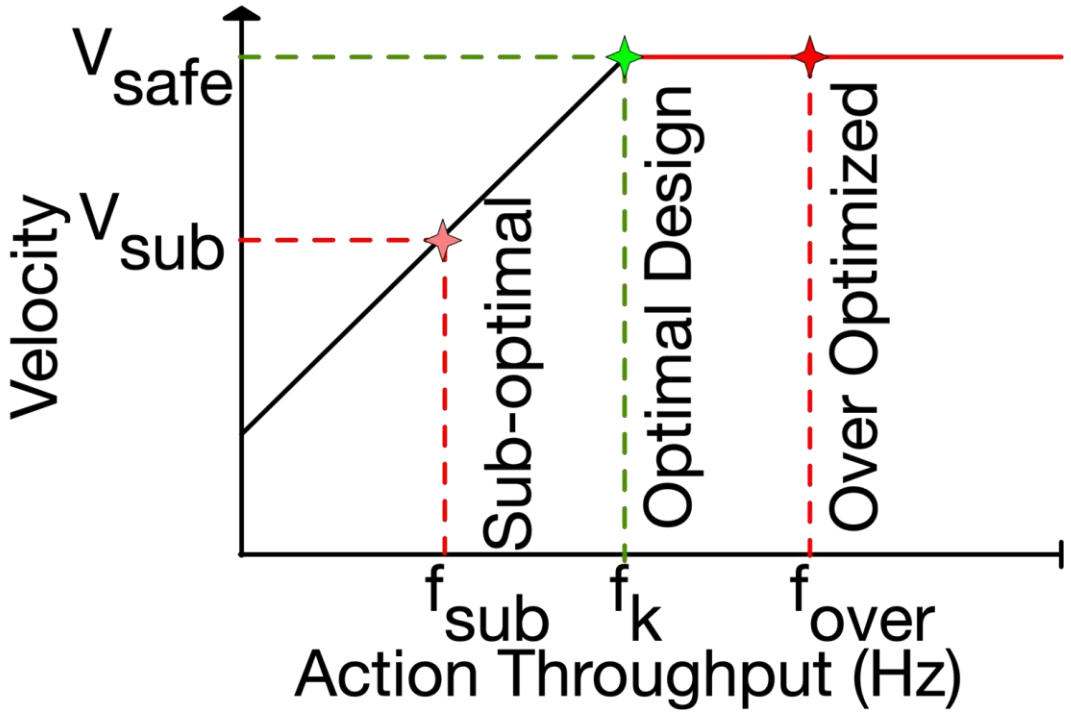

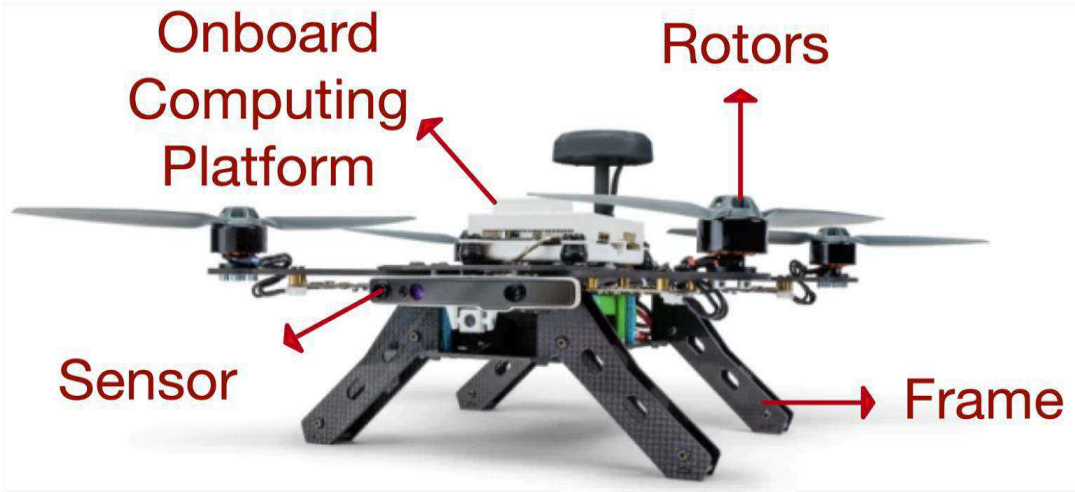

S. Krishnan, Z. Wan, K. Bhardwaj, N. Jadhav, A. Faust, V. Janapa Reddi. Roofline Model for UAVs: A Bottleneck Analysis Tool for Onboard Compute Characterization of Autonomous Unmanned Aerial Vehicles. ASPLOS 2021.

2 Measure Twice, Cut Once: Metrics Matter

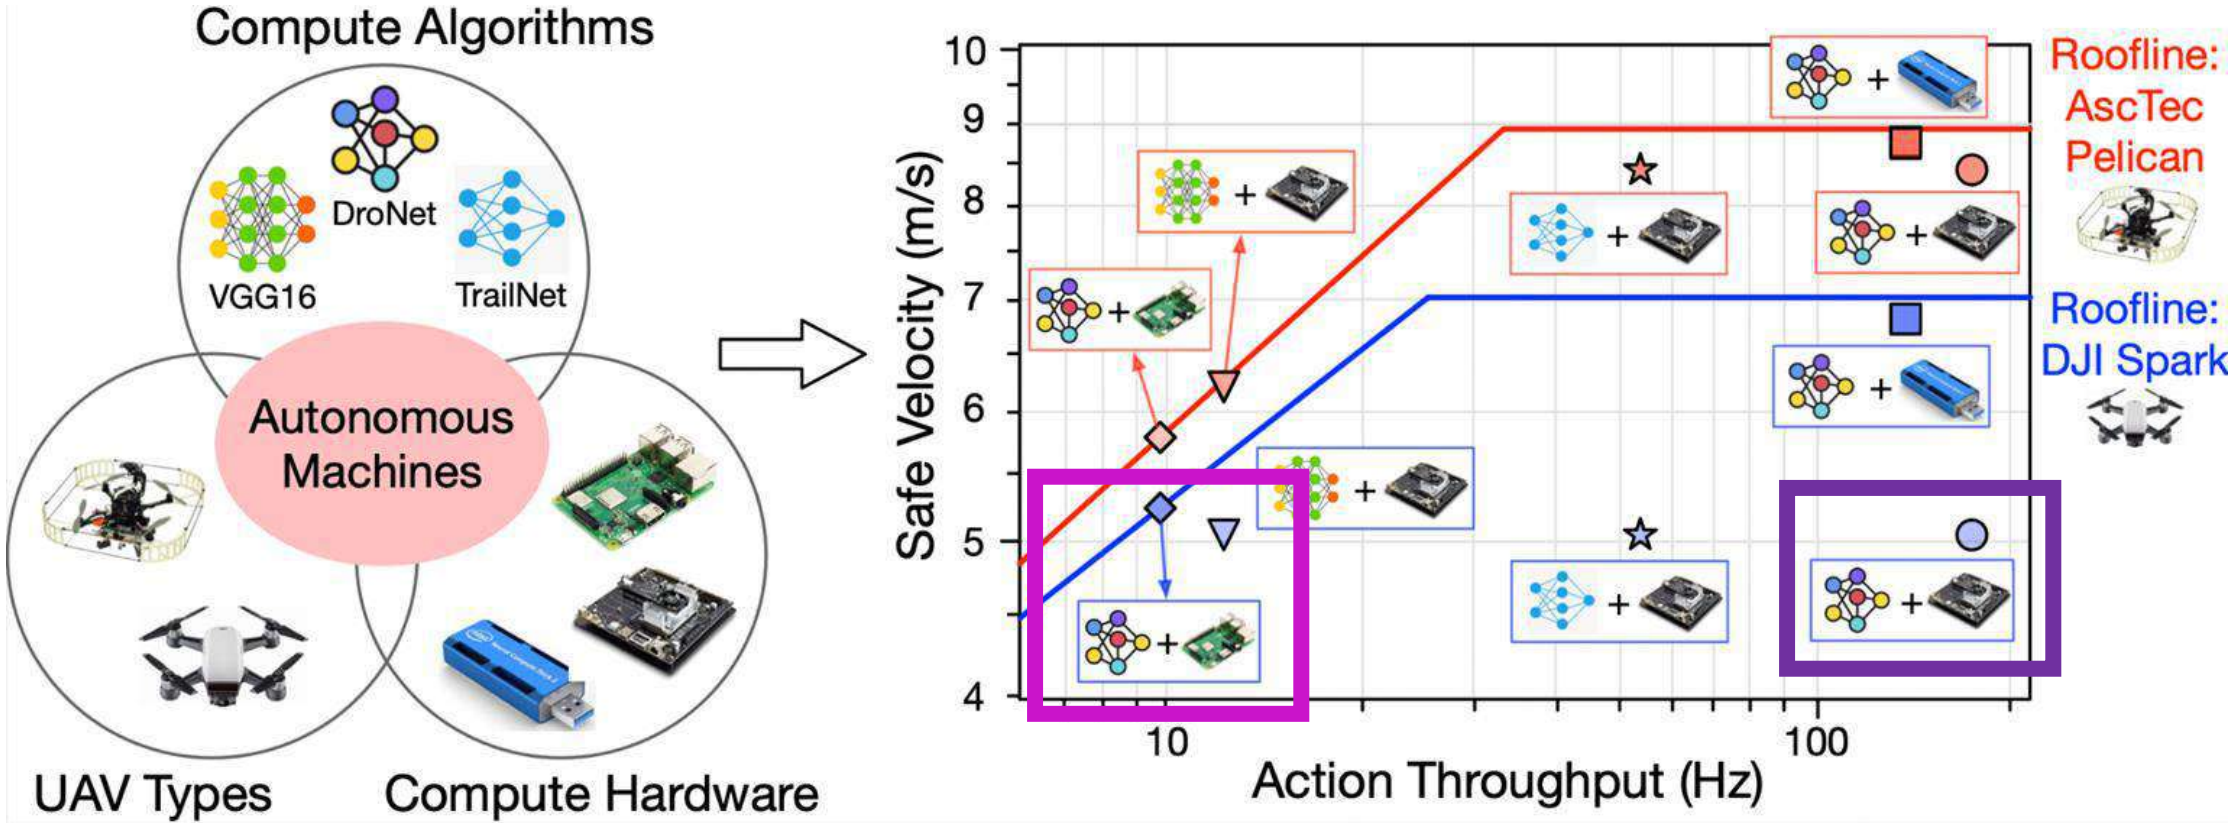

S. Krishnan, Z. Wan, K. Bhardwaj, N. Jadhav, A. Faust, V. Janapa Reddi. Roofline Model for UAVs: A Bottleneck Analysis Tool for Onboard Compute Characterization of Autonomous Unmanned Aerial Vehicles. ASPLOS 2021.

The Magnificent Seven
Challenges & Opportunities in Domain-Specific Accelerator Design for Autonomous Systems

1. Build Bridges: Engage with Domain Experts
2. **Measure Twice, Cut Once: Metrics Matter**
3. "Widgetism": Avoid Over-Specialization
4. **Pump the Brakes: Do Not Always Accelerate**
5. Chips and Salsa: Acceleration Beyond ASICs
6. **Forest vs. Trees: Take an End-to-End View**
7. Design Global: Sustainability and Impact

7 Design Global: Sustainability and Impact

Pitfall: *Design compute in isolation from its global and societal impact.*

7 Design Global: Sustainability and Impact

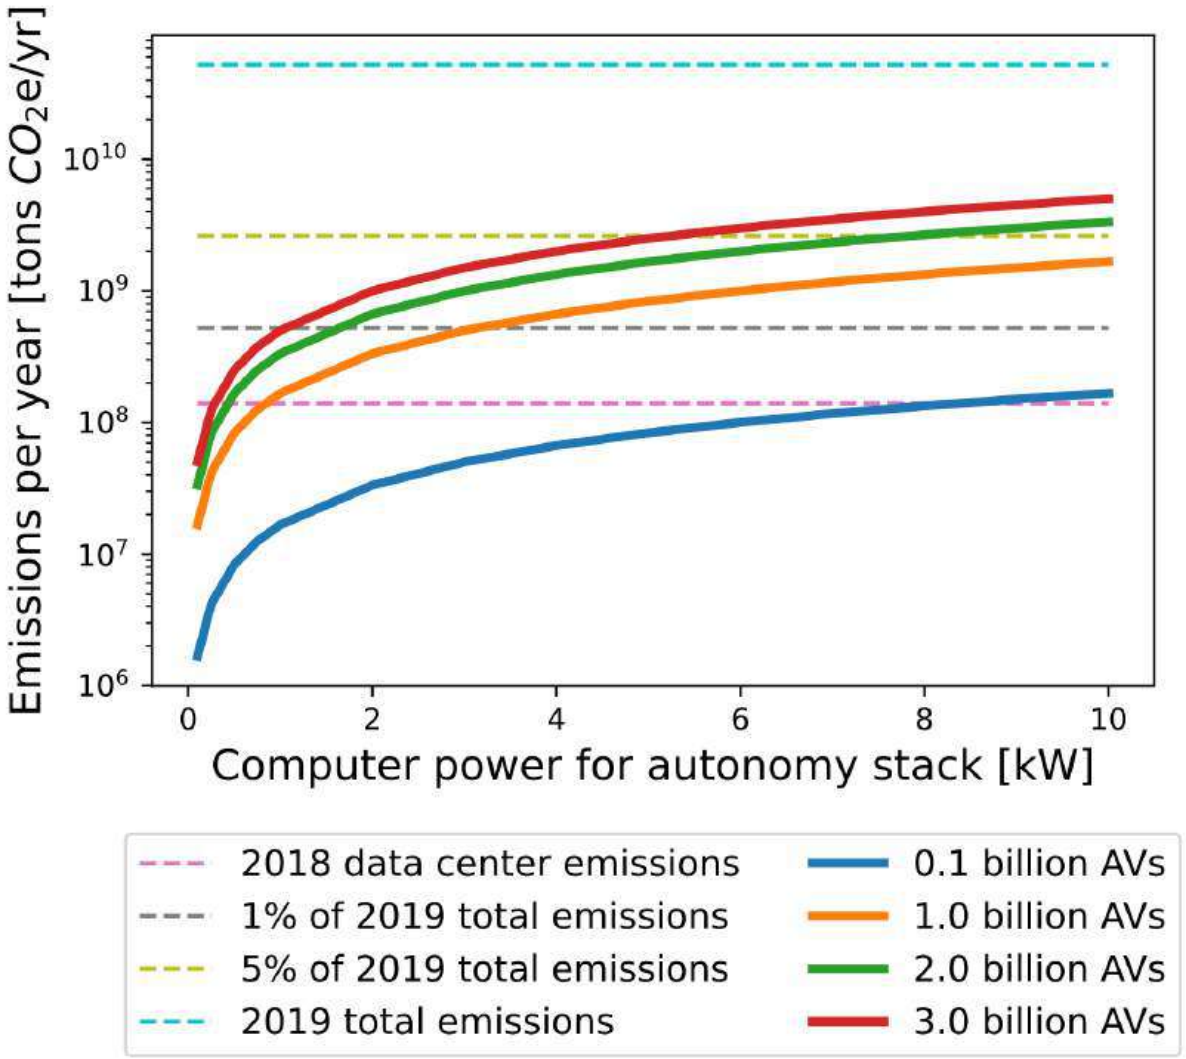

S. Sudhakar, V. Sze, S. Karaman. Data Centers on Wheels: Emissions From Computing Onboard Autonomous Vehicles. IEEE Micro 2023.

7 Design Global: Sustainability and Impact

TABLE 1. FAR for CS and engineering subfields based on prior work and including our result for robotics [1], [3], [4] (data from 2017 to 2023).

FIELD	FAR (%)
CS education	42
Human–computer interaction	26
CS overall average	**16–26**
Knowledge systems	19
Software engineering and languages	14
Artificial intelligence	12
Robotics	**11–12 (our analysis)**
Computer systems	10
Theory and algorithms	8

As has been noted in related works, this kind of methodology has many flaws and does not take into account much of the nuance in gender, including issues of bias, misperception, and nonbinary identities [7], [8]. However, we hope that this initial study will help add to the robotics community's understanding of the current state of gender diversity and, at a minimum, provide directionally correct data to help with future diversity, equity, and inclusion efforts.

E. Fields, C. Ho, M. J. Kim, Z. Wu, B. Plancher. Underrepresentation of Women in Robotics Research. IEEE RAM 2024.

The Magnificent Seven
Challenges & Opportunities in Domain-Specific Accelerator Design for Autonomous Systems

1. Build Bridges: Engage with Domain Experts
2. Measure Twice, Cut Once: Metrics Matter
3. "Widge tion
4. Pump t celerate
5. Chips a d ASICs
6. Forest vs. Trees: Take an End-to-End View
7. Design Global: Sustainability and Impact

So What Can We Do?

The Magnificent Seven
Challenges & Opportunities in Domain-Specific Accelerator Design for Autonomous Systems

Future Directions and Opportunities

- A Enabling Technologies and Methodologies
- B Fostering a Robust Research Ecosystem
- C Sustainable & Responsible Hardware Design

A Enabling Technologies and Methodologies

Opportunity: *Reduce Complexity and Time-to-Market with Agile Design Tools and End-to-End Simulation Environments*

A Enabling Technologies and Methodologies

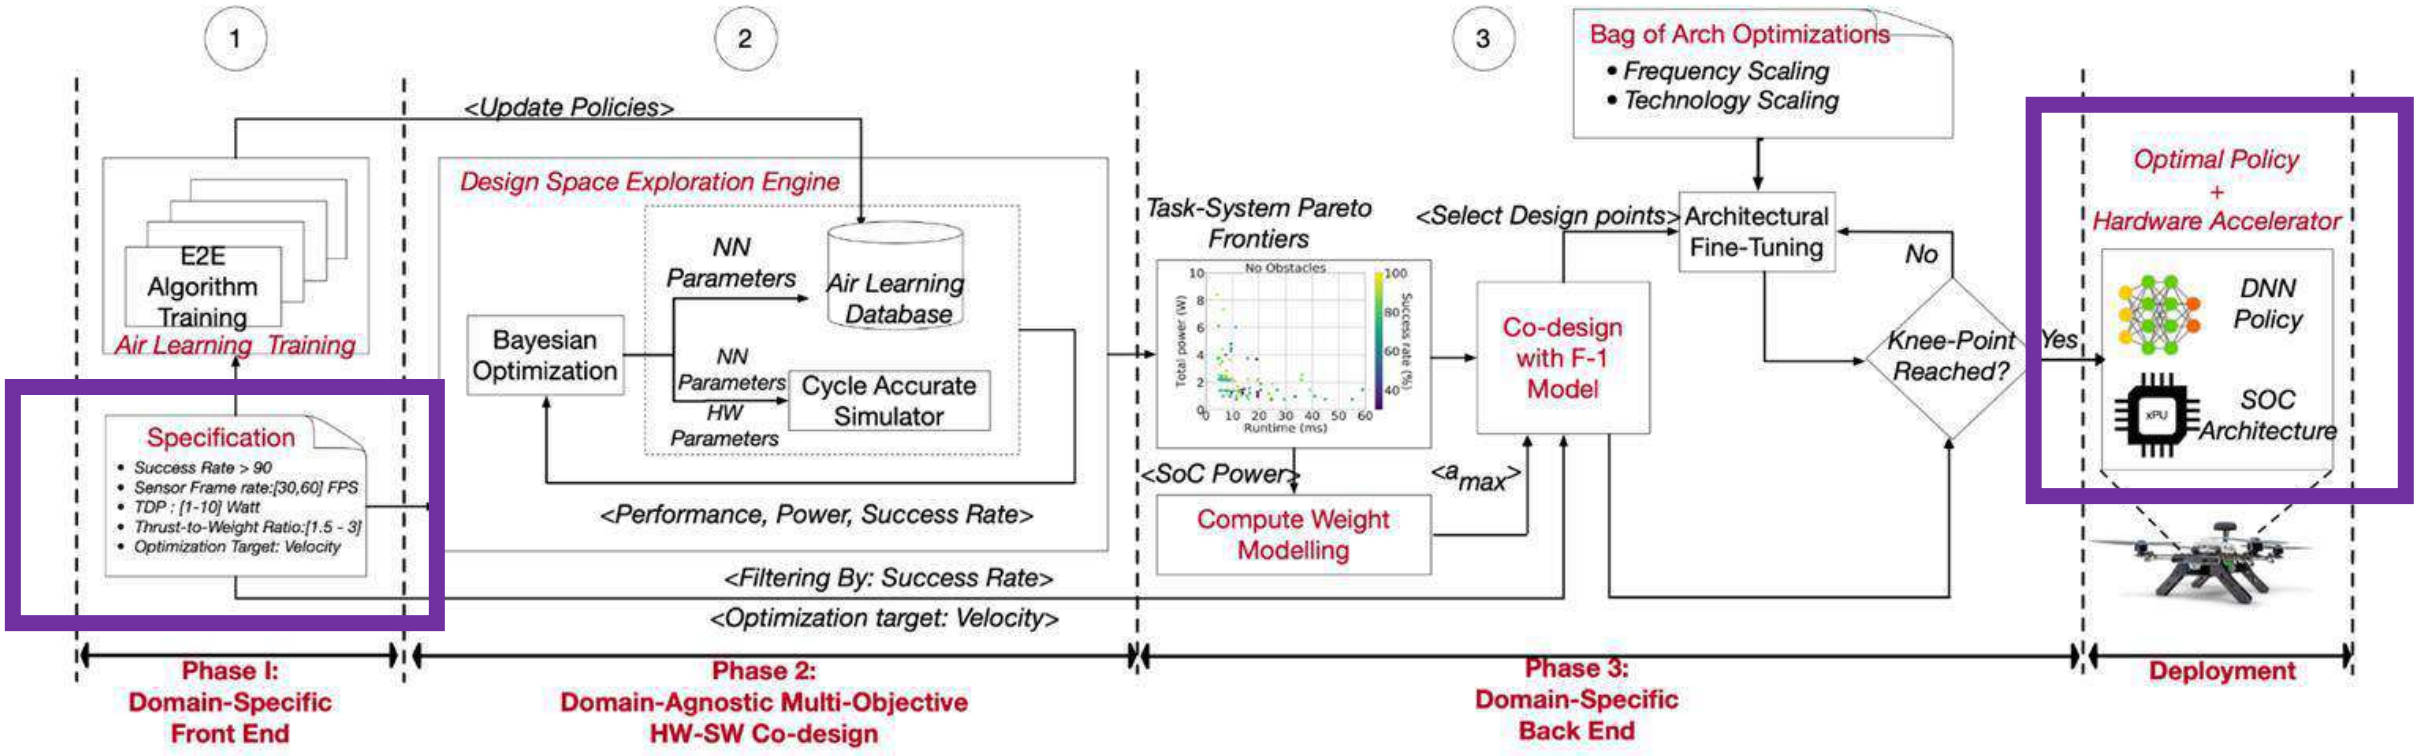

S. Krishnan, Z. Wan, K. Bhardwaj, P. Whatmough, A. Faust, S. M. Neuman, G.-Y. Wei, D. Brooks, V. Janapa Reddi.
Autopilot: Automatic domain-specific soc design for autonomous unmanned aerial vehicles. MICRO. 2022.

A Enabling Technologies and Methodologies

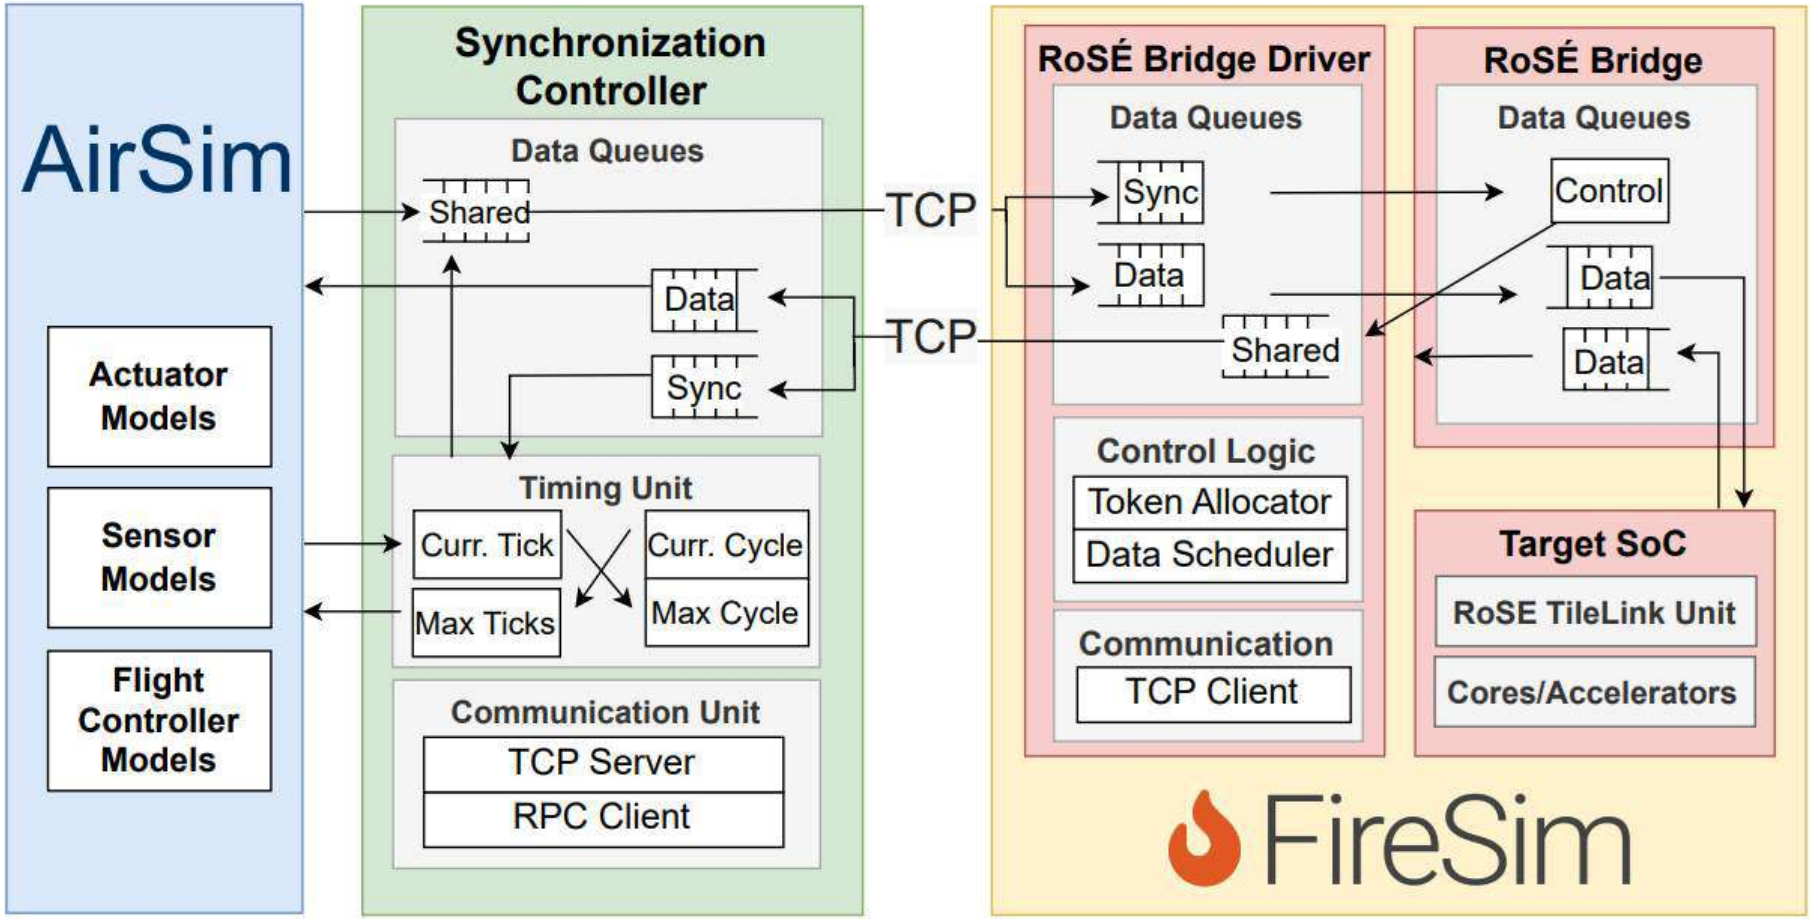

D. Nikiforov, S. C. Dong, C. L. Zhang, S. Kim, B. Nikolic, Y. S. Shao. Rosé: A hardware-software co-simulation infrastructure enabling pre-silicon full-stack robotics soc evaluation. ISCA. 2023.

B Fostering a Robust Research Ecosystem

Opportunity: *Increase Cross-Domain Collaborations and Develop Open Source Resources and Benchmarks*

B Fostering a Robust Research Ecosystem

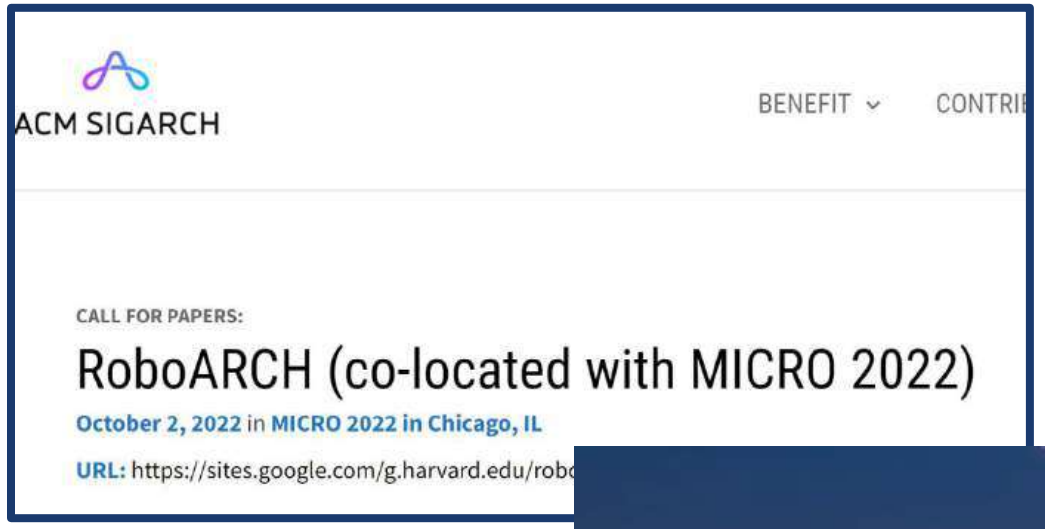

We'll be at MICRO 2024!

Fostering a Robust Research Ecosystem

MLPerf

RobotPerf

RobotPerf: An Open-Source, Vendor-Agnostic, Benchmarking Suite for Evaluating Robotics Computing System Performance

Víctor Mayoral-Vilches[1,2], Jason Jabbour[3], Yu-Shun Hsiao[3], Zishen Wan[4], Martiño Crespo-Álvarez[1], Matthew Stewart[3], Juan Manuel Reina-Muñoz[1], Prateek Nagras[1], Gaurav Vikhe[1], Mohammad Bakhshalipour[5], Martin Pinzger[2], Stefan Rass[6,2], Smruti Panigrahi[7], Giulio Corradi[8], Niladri Roy[9], Phillip B. Gibbons[5], Sabrina M. Neuman[10], Brian Plancher[11], Vijay Janapa Reddi[3]

ACCELERATION ROBOTICS

ALPEN-ADRIA UNIVERSITÄT

JKU JOHANNES KEPLER UNIVERSITY LINZ

AMD

CMU

intel

BOSTON UNIVERSITY

Ford

ROS.org

C Sustainable & Responsible Hardware Design

Opportunity: Strategically Develop Reusable Designs and Deployments for Full-Lifecycle Sustainability

C Sustainable & Responsible Hardware Design

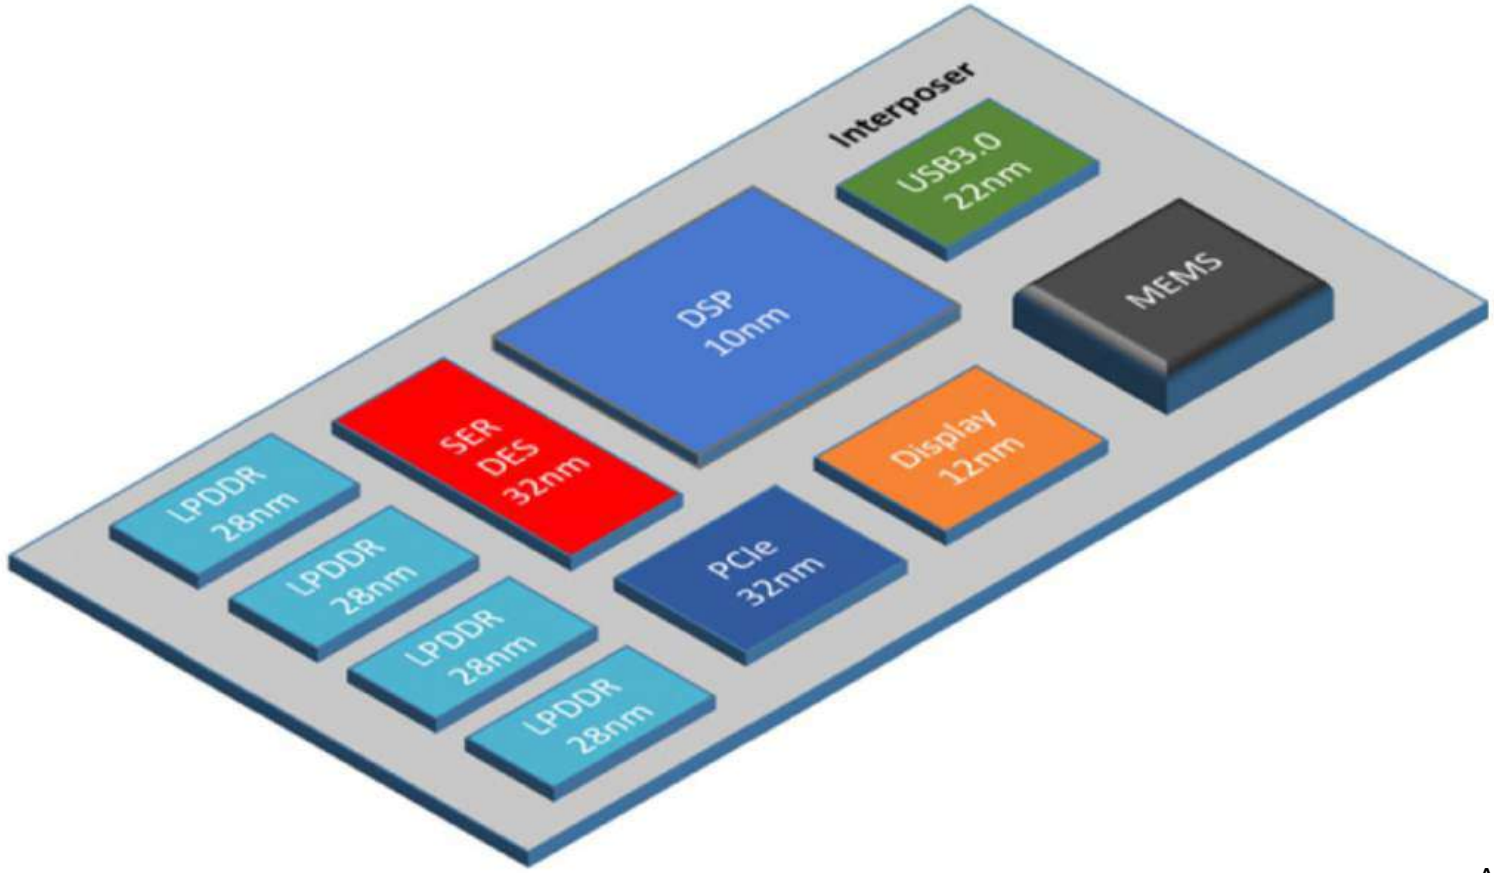

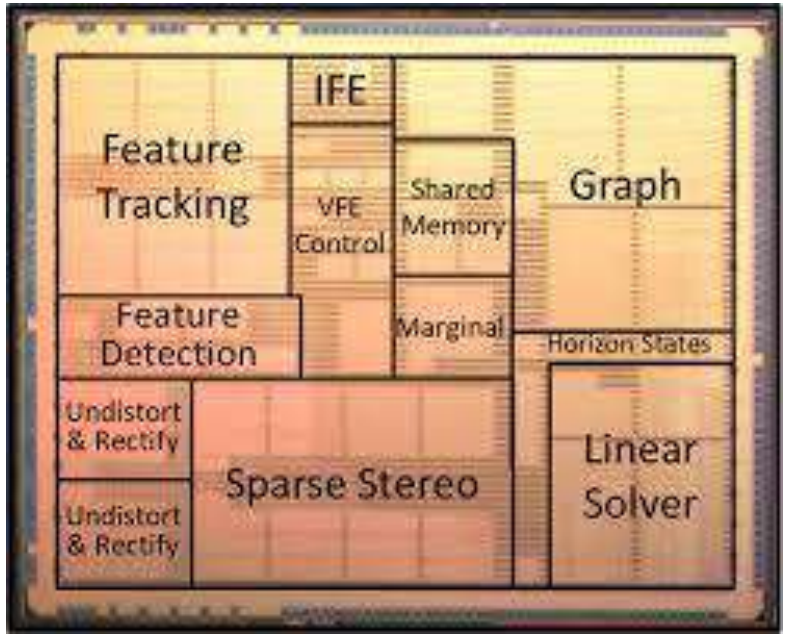

A. Suleiman, Z. Zhang, L. Carlone, S. Karaman, V. Sze. Navion: A 2mW Fully Integrated Real-Time Visual-Inertial Odeometry Accelerator for Autonomous Navigation of Nano Drones. JSSC. 2019.

https://semiengineering.com/knowledge_centers/packaging/advanced-packaging/chiplets/

C Sustainable & Responsible Hardware Design

DOI:10.1145/3608473

Assessing the environmental impacts of machine learning on microcontrollers.

BY SHVETANK PRAKASH, MATTHEW STEWART, COLBY BANBURY, MARK MAZUMDER, PETE WARDEN, BRIAN PLANCHER, AND VIJAY JANAPA REDDI

Is TinyML Sustainable?

C Sustainable & Responsible Hardware Design

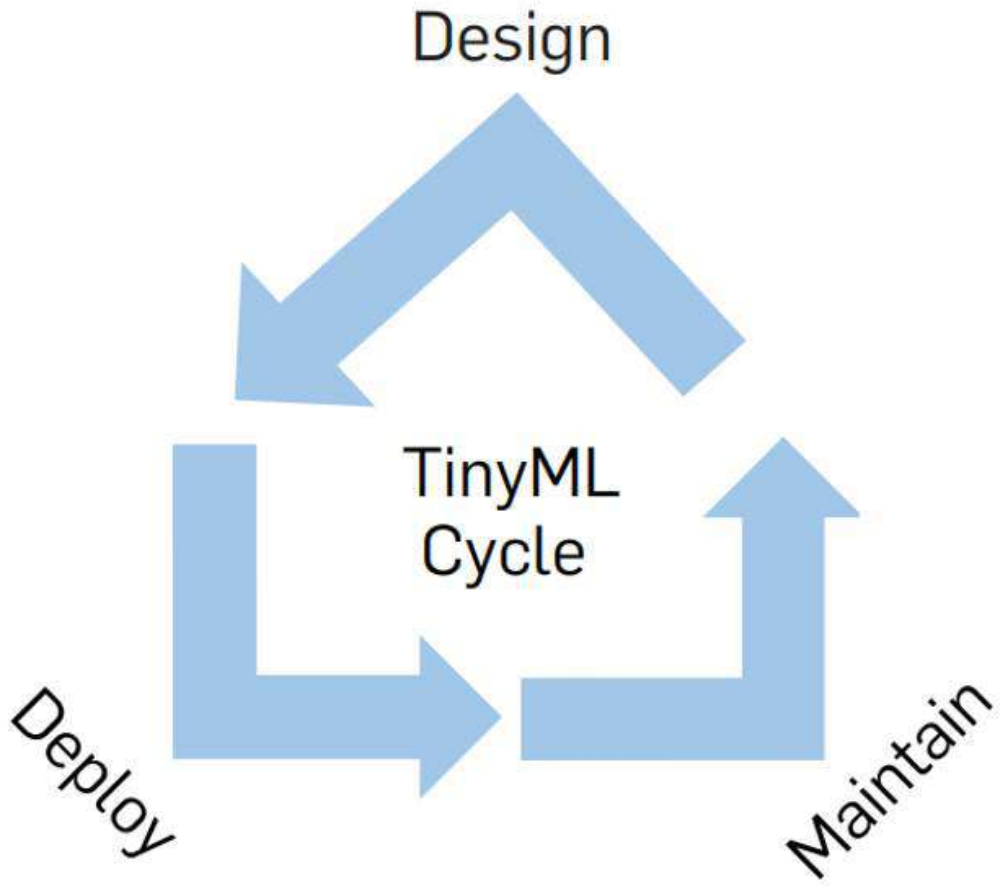

C Sustainable & Responsible Hardware Design

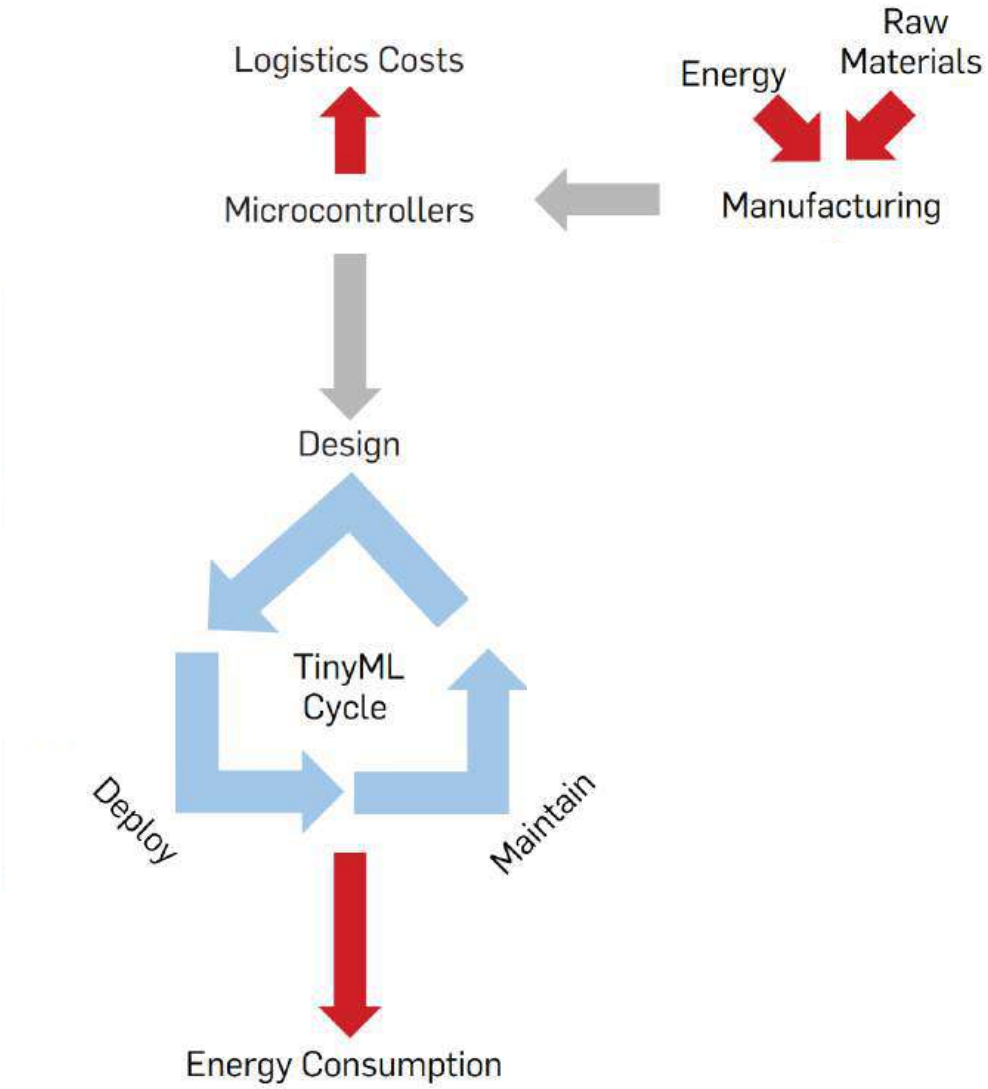

C Sustainable & Responsible Hardware Design

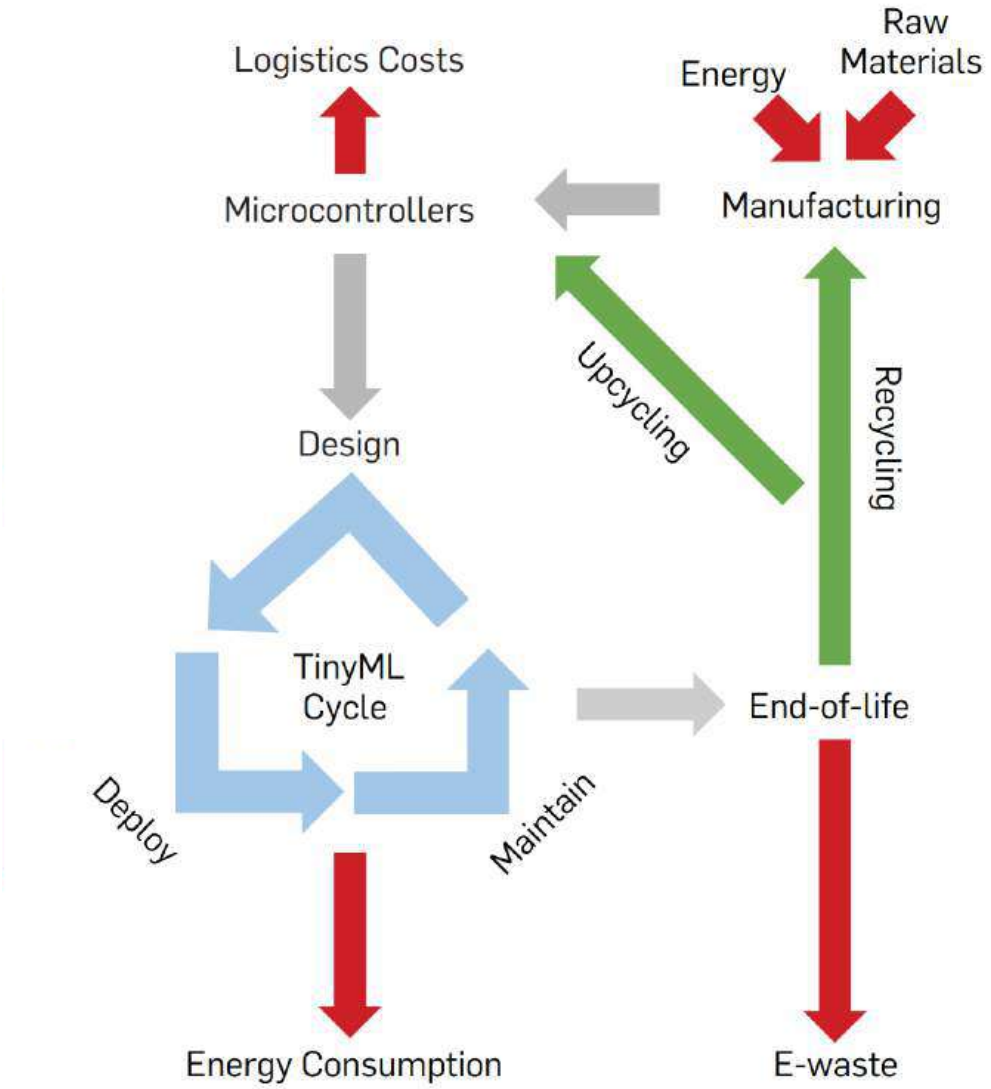

C Sustainable & Responsible Hardware Design

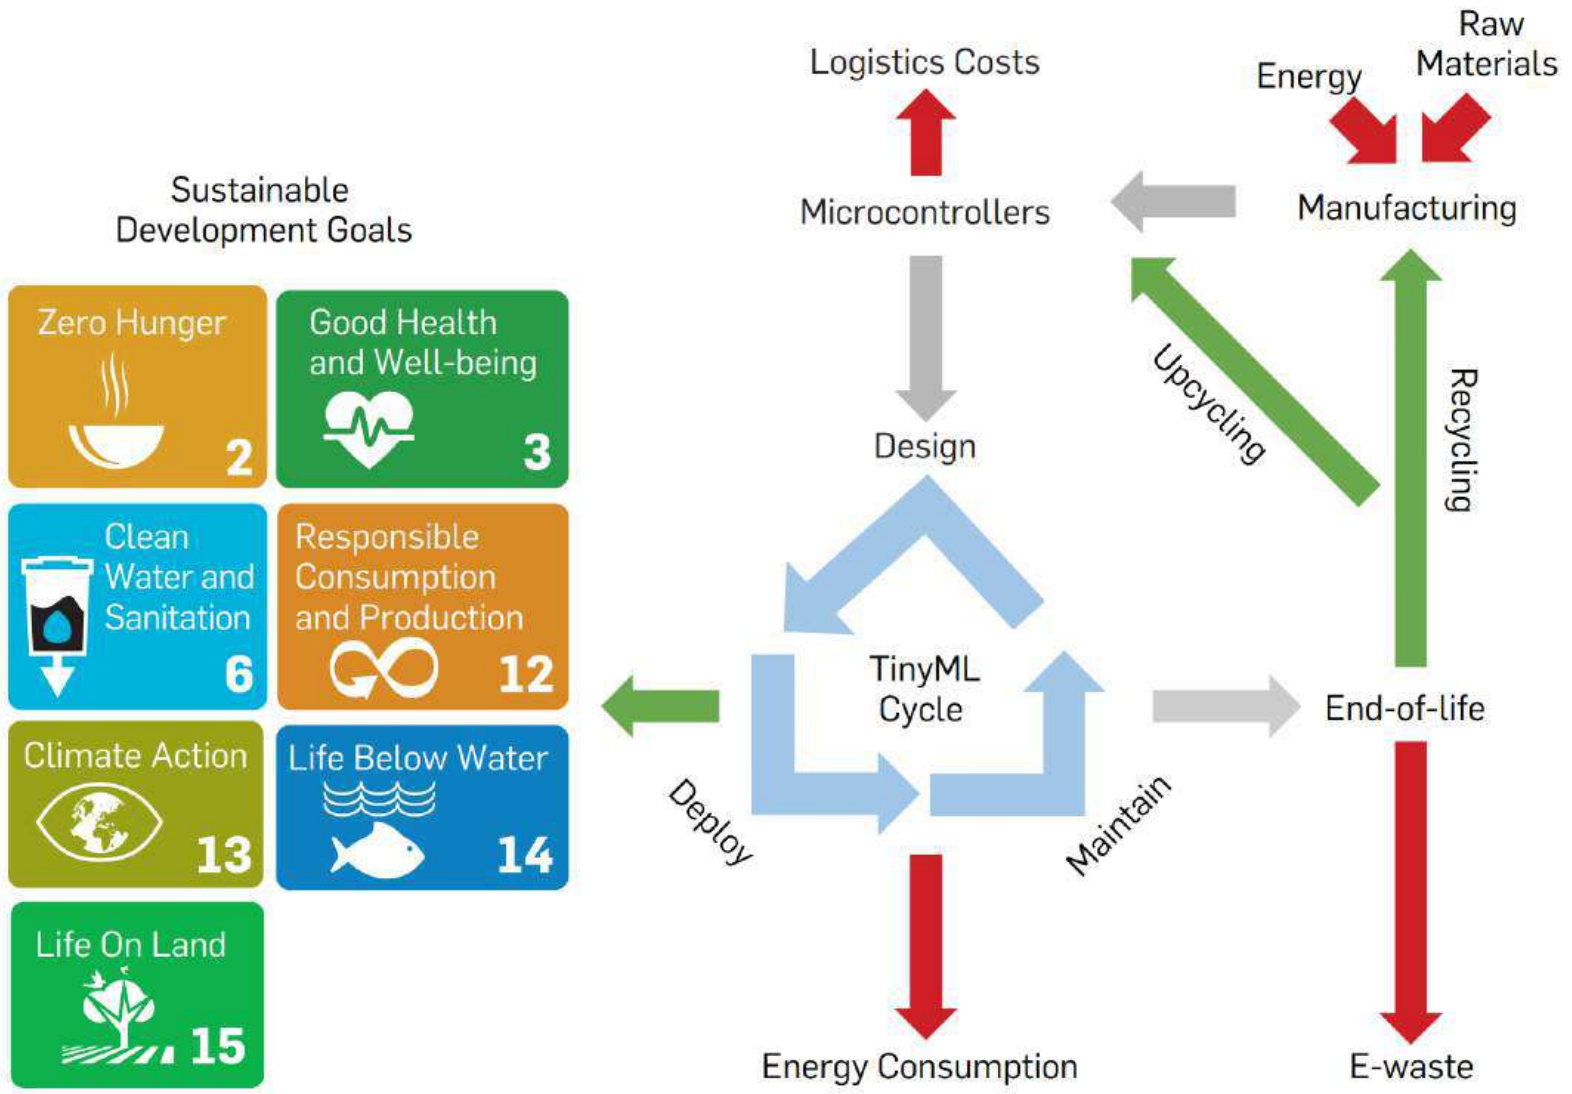

C Sustainable & Responsible Hardware Design

If deployed for good, TinyML can be net positive! Can autonomous systems be as well?

Lets Keep Building Accelerators for Autonomous Systems!

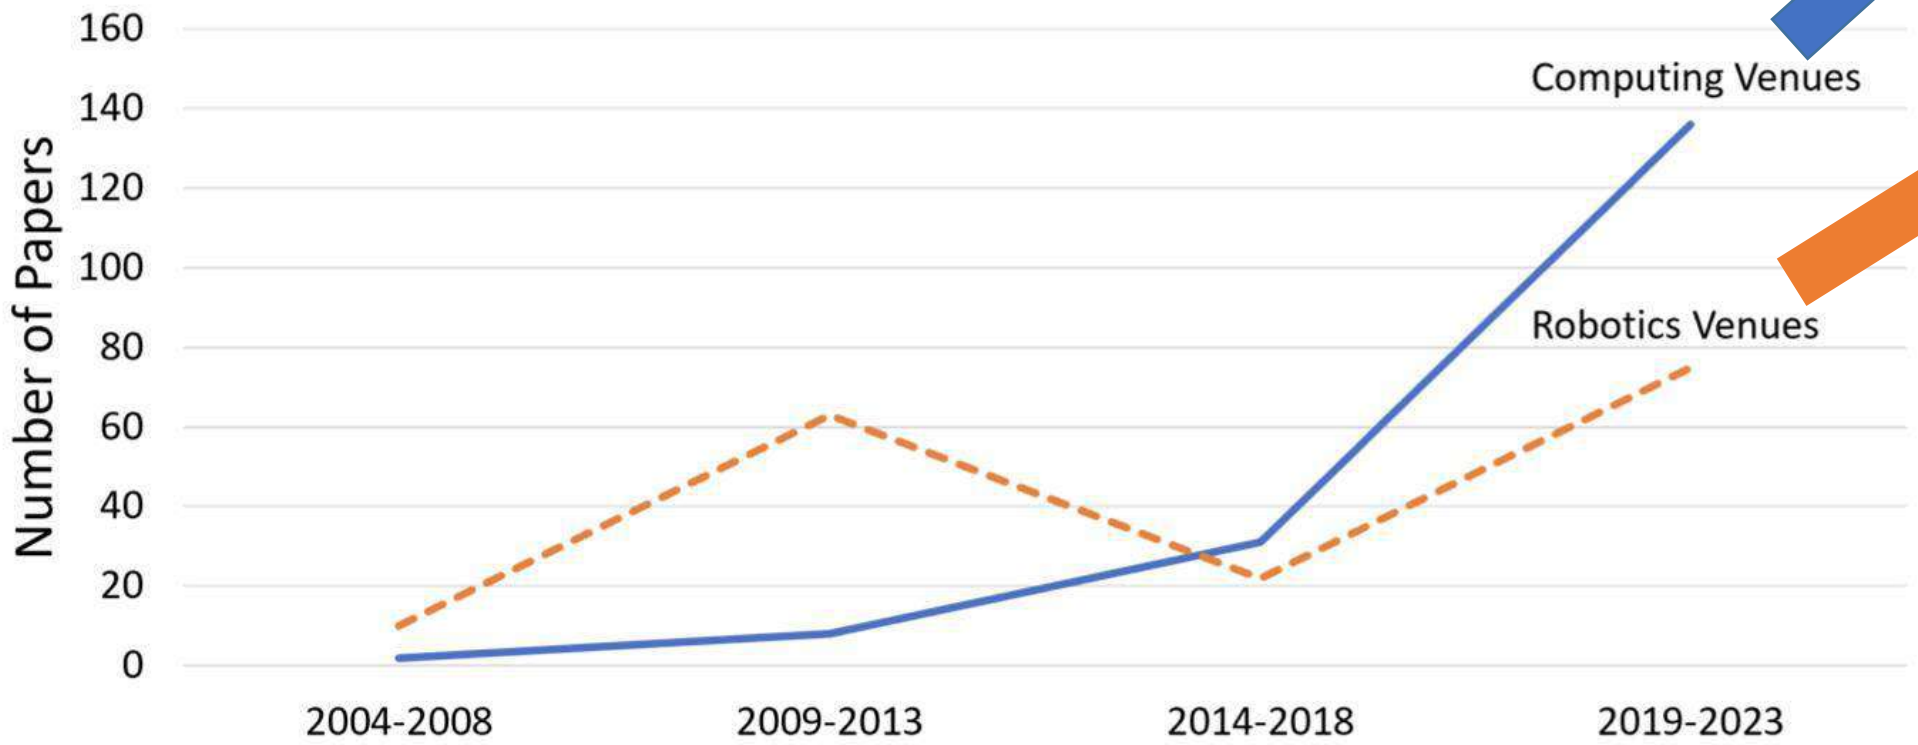

Computing Keywords: Accelerator & Robot | Autonomous System | UAV | Drone | Autonomous Vehicle | Self-Driving, in DAC, ISCA, MICRO, HPCA, and ASPLOS. Robotics Keywords: ASIC | FPGA, in ICRA, IROS, RSS, and RA-L.

The Magnificent Seven
Challenges & Opportunities in Domain-Specific Accelerator Design for Autonomous Systems

1. Build Bridges: Engage with Domain Experts
2. Measure Twice, Cut Once: Metrics Matter
3. "Widgetism": Avoid Over-Specialization
4. Pump the Brakes: Do Not Always Accelerate
5. Chips and Salsa: Acceleration Beyond ASICs
6. Forest vs. Trees: Take an End-to-End View
7. Design Global: Sustainability and Impact

Sabrina M. Neuman*
Boston University
sneuman@bu.edu

Brian Plancher*
Barnard College, Columbia University
bplancher@barnard.edu

Vijay Janapa Reddi*
Harvard University
vj@eecs.harvard.edu