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], Alejandra Martínez-Fariña[1], 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]

1: Acceleration Robotics 2: Alpen-Adria-Universitat Klagenfurt, 3: Harvard University, 4: Georgia Institute of Technology, 5: Carnegie Mellon University, 6: Johannes Kepler University, 7: Ford Motor Company, 8: AMD, 9: Intel, 10: Boston University, 11: Barnard College, Columbia University

The Big Picture:

We introduce **RobotPerf**, a vendor-agnostic benchmarking suite designed to evaluate robotics computing performance across a diverse range of hardware platforms using ROS 2 as its common baseline. Emphasizing adaptability, portability, and a community-driven approach, RobotPerf aims to provide fair comparisons of ROS 2 computational graphs across CPUs, GPUs, FPGAs and other accelerators through both grey-box and black-box approaches.

github.com/robotperf/benchmarks

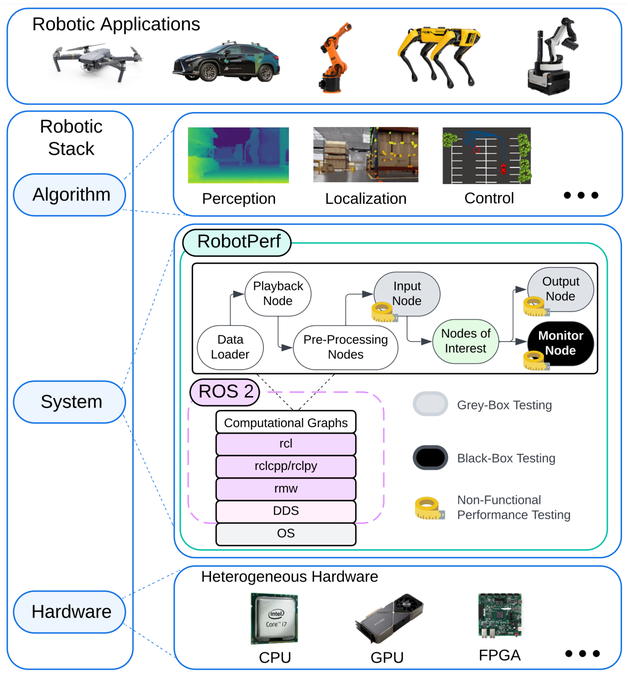

Principles and Methodology:

RobotPerf's key design principles include:

- Non-Functional Performance Testing
- ROS 2 Integration & Adaptability
- Platform Independence & Portability
- Reproducibility & Consistency
- Flexible Methodologies for Benchmarking and Opaque Performance Tests that do NOT Alter Workloads of Interest and can be accomplished through either Grey-Box or Black-Box Approaches

Evaluation Highlights :

- RobotPerf's Beta release benchmarks across a wide variety of heterogeneous hardware platforms (18) and workloads (16).
- The "one-size-fits-all" strategy is suboptimal. E.g., the control latency plot (col 3, row 1), shows that the I7H outperforms NO on benchmarks C1, C3, C4, and C5, but is 6.5× slower on benchmark C2.
- Hardware acceleration improves performance. E.g., in the perception benchmarks (col 1), the ROBOTCORE Perception accelerator (KR) provides its hardware platform with gains of as much as 11.5x over the non-accelerated variant (KK).

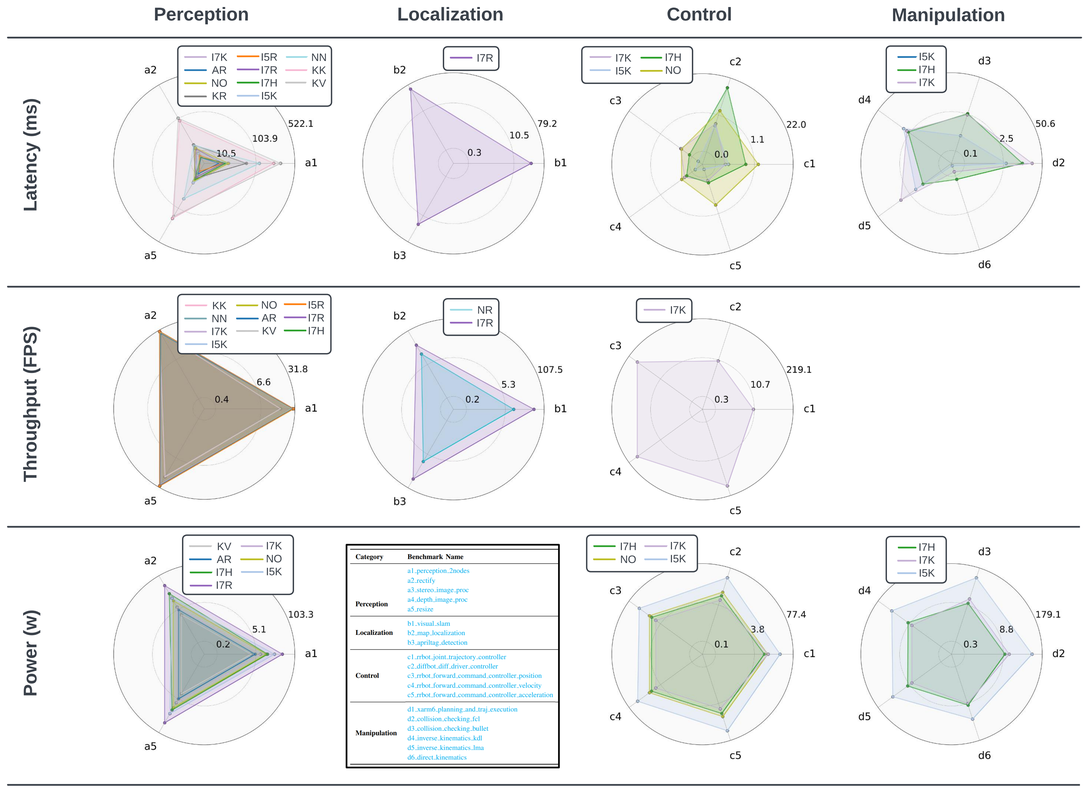

Category	Benchmark Name
Perception	a1_perception_2nodes a2_rectify a3_stereo_image_proc a4_depth_image_proc a5_resize
Localization	b1_visual_slam b2_map_localization b3_apriltag_detection
Control	c1_rrbot_joint_trajectory_controller c2_diffbot_diff_driver_controller c3_rrbot_forward_command_controller_position c4_rrbot_forward_command_controller_velocity c5_rrbot_forward_command_controller_acceleration
Manipulation	d1_xarm6_planning_and_traj_execution d2_collision_checking_fcl d3_collision_checking_bullet d4_inverse_kinematics_kdl d5_inverse_kinematics_lma d6_direct_kinematics

● [I5U] (15W)	Intel i5-8250U	○ [I5U] (15W)	Intel i5-8250U
● [AR] (65W)	AMD Ryzen 5 PRO 4650G	○ [AR] (65W)	AMD Ryzen 5 PRO 4650G
● [I7K] (95W)	Intel i7-8700K	○ [I7K] (95W)	Intel i7-8700K
● [I7H] (125W)	Intel i7-12700H	○ [I7H] (125W)	Intel i7-12700H
● [I5K] (125W)	Intel i5-13600K	○ [I5K] (125W)	Intel i5-13600K
● [I9K] (125W)	Intel i9-12900KF	○ [I9K] (125W)	Intel i9-12900KF
			General-Purpose Hardware
● [NN] (5W)	NVIDIA Jetson Nano	○ [NN] (5W)	NVIDIA Jetson Nano
● [QR] (5W)	Qualcomm RB5 Robotics Kit	○ [QR] (5W)	Qualcomm RB5 Robotics Kit
● [JX] (30W)	Jetson AGX Xavier	○ [JX] (30W)	Jetson AGX Xavier
● [NO] (60W)	NVIDIA AGX Orin Dev. Kit	○ [NO] (60W)	NVIDIA AGX Orin Dev. Kit
● [I7N] (295W)	Intel i7-12700H + NVIDIA GeForce RTX 3060	○ [I7N] (295W)	Intel i7-12700H + NVIDIA GeForce RTX 3060
			Heterogeneous Hardware
● [KK] (15W)	Kria KR260	○ [KK] (15W)	Kria KR260
● [KV] (36W)	Kria KV260	○ [KV] (36W)	Kria KV260
			Reconfigurable Hardware
● [KR] (35W)	Kria KR260 (ROBOTCORE® Perception)	○ [KR] (35W)	Kria KR260 (ROBOTCORE® Perception)
● [NR] (80W)	NVIDIA AGX Orin Dev. Kit (ROBOTCORE® Perception)	○ [NR] (80W)	NVIDIA AGX Orin Dev. Kit (ROBOTCORE® Perception)
● [I7T] (145W)	Intel i7-12700H (ROBOTCORE® Transforms)	○ [I7T] (145W)	Intel i7-12700H (ROBOTCORE® Transforms)
● [I5R] (315W)	Intel i5-13600K + NVIDIA GeForce RTX 3060 (ROBOTCORE® Perception)	○ [I5R] (315W)	Intel i5-13600K + NVIDIA GeForce RTX 3060 (ROBOTCORE® Perception)
● [I7R] (315W)	Intel i7-12700H + NVIDIA GeForce RTX 3060 (ROBOTCORE® Perception)	○ [I7R] (315W)	Intel i7-12700H + NVIDIA GeForce RTX 3060 (ROBOTCORE® Perception)
			Accelerator Hardware

Benchmarking Methodology: ● Black-box Testing ○ Grey-box Testing