

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

Víctor Mayoral-Vilches^{1,2}, Jason Jabbour³, Yu-Shun Hsiao³, Zishen Wan⁴, Alejandra Martínez-Fariña¹, Martiño Crespo-Álvarez¹, Matthew Stewart³, Juan Manuel Reina-Muñoz¹, Prateek Nagras¹, Gaurav Vikhe¹, Mohammad Bakhshalipour⁵, Martin Pinzger², Stefan Rass^{6,2}, Smruti Panigrahi⁷, Giulio Corradi⁸, Niladri Roy⁹, Phillip B. Gibbons⁵, Sabrina M. Neuman¹⁰, Brian Plancher¹¹, Vijay Janapa Reddi³

Abstract— 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. The suite encompasses ROS 2 packages covering the full robotics pipeline and integrates two distinct benchmarking approaches: black-box testing, which measures performance by eliminating upper layers and replacing them with a test application, and grey-box testing, an application-specific measure that observes internal system states with minimal interference. Our benchmarking framework provides ready-to-use tools and is easily adaptable for the assessment of custom ROS 2 computational graphs. Drawing from the knowledge of leading robot architects and system architecture experts, **RobotPerf** establishes a standardized approach to robotics benchmarking. As an open-source initiative, **RobotPerf** remains committed to evolving with community input to advance the future of hardware-accelerated robotics.

I. INTRODUCTION

In order for robotic systems to operate safely and effectively in dynamic real-world environments, their computations must run at real-time rates while meeting power constraints. Towards this end, accelerating robotic kernels on heterogeneous hardware, such as GPUs and FPGAs, is emerging as a crucial tool for enabling such performance [1], [2], [3], [4], [5], [6], [7]. This is particularly important given the impending end of Moore’s Law and the end of Dennard Scaling, which limits single CPU performance [8], [9].

While hardware-accelerated kernels offer immense potential, they necessitate a reliable and standardized infrastructure to be effectively integrated into robotic systems. As the industry leans more into adopting such standard software infrastructure, the Robot Operating System (ROS) [10] has emerged as a favored choice. Serving as an industry-grade middleware, it aids in building robust computational robotics graphs, reinforcing the idea that robotics is more than just individual algorithms. The growing dependency on ROS 2 [11], combined with the computational improvements offered by hardware acceleration, accentuates the community’s demand for a standardized, industry-grade benchmark to evaluate varied hardware solutions. Recently, there has been a plethora of workshops and tutorials focusing on benchmarking robotics applications [12], [13], [14], [15], [16], [17], [18], [19],




Fig. 1: A high level overview of RobotPerf. It targets industry-grade real-time systems with complex and extensible computation graphs using the Robot Operating System (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.

[20], [21], [22], and while benchmarks for specific robotics algorithms [23], [24] and certain end-to-end robotic applications, such as drones [25], [26], [27], [28], do exist, the nuances of analyzing general ROS 2 computational graphs on heterogeneous hardware is yet to be fully understood.

In this paper, we introduce *RobotPerf*, an open-source and community-driven benchmarking tool designed to assess the performance of robotic computing systems in a standardized, architecture-neutral, and reproducible way, accommodating the various combinations of hardware and software in different robotic platforms (see Figure 1). RobotPerf focuses on evaluating robotic workloads in the form of ROS 2 computational graphs on a wide array of hardware setups, encompassing a complete robotics pipeline and emphasizing

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

real-time critical metrics. The framework incorporates two distinct benchmarking methodologies that utilize various forms of instrumentation and ROS *nodes* to capture critical metrics in robotic systems. These approaches are: black-box testing, which measures performance by eliminating upper layers and replacing them with a test application, and grey-box testing, an application-specific measure that observes internal system states with minimal interference. The framework is user-friendly, easily extendable for evaluating custom ROS 2 computational graphs, and collaborates with major hardware acceleration vendors for a standardized benchmarking approach. It aims to foster research and innovation as an open-source project. We validate the framework’s capabilities by conducting benchmarks on diverse hardware platforms, including CPUs, GPUs, and FPGAs, thereby showcasing RobotPerf’s utility in drawing valuable performance insights.

RobotPerf’s source code and documentation are available at <https://github.com/robotperf/benchmarks> and its methodologies are currently being used in industry to benchmark industry-strength, production-grade systems.

II. BACKGROUND & RELATED WORK

A. The Robot Operating System (ROS and ROS 2)

ROS [10] is a widely-used middleware for robot development that serves as a *structured communications layer* and offers a comprehensive suite of additional functionalities including: open-source packages and drivers for various tasks, sensors, and actuators, as well as a collection of tools that simplify development, deployment, and debugging processes. ROS enables the creation of computational graphs (see Figure 1) that connect software processes, known as nodes, through topics, facilitating the development of end-to-end robotic systems. Within this framework, nodes can publish to or subscribe from topics, enhancing the modularity of robotic systems.

ROS 2 builds upon ROS and addresses many of its key limitations. Constructed to be industry-grade, ROS 2 adheres to industry Data Distribution Service (DDS) and Real-Time Publish Subscribe (RTPS) standards [29]. Based on the Data Distribution Service (DDS) standard, it enables fine-grained, direct, inter- and intra-node communication, enhancing performance, reducing latency, and improving scalability. Importantly, these improvements are also designed to support hardware acceleration [30], [5]. Over 600 companies have adopted ROS 2 and its predecessor ROS in their production environments, underscoring its significance and widespread adoption in the industry [11].

ROS 2 also provides standardized APIs to connect user code through language-specific client libraries, *rclcpp* and *rclpy*, which handle the scheduling and invocation of callbacks such as timers, subscriptions, and services. Without a ROS Master, ROS 2 creates a decentralized framework where nodes discover each other and manage their own parameters.

B. Robotics Benchmarks

There has been much recent development of open-source robotics libraries and associated benchmarks demonstrat-

	Characteristics						
	Real-time Performance Metrics	Spans Multiple Pipeline Categories	Evaluation on Heterogeneous Hardware	Integration with ROS/ROS 2 Framework	Functional Performance Testing	Non-functional Performance Testing	Community Led
OMPL Benchmark [31]	✓	✗	✗	✗	✗	✓	✗
MotionBenchMaker [32]	✓	✗	✗	✗	✓	✓	✗
OpenCollBench [33]	✗	✗	✓	✗	✓	✗	✗
BARN [34]	✗	✗	✗	✓	✓	✗	✗
DynaBARN [35]	✓	✗	✗	✓	✓	✗	✗
MAVBench [25]	✓	✓	✓	✓	✓	✓	✗
Bench-MR [36]	✓	✗	✗	✗	✓	✗	✗
RTRBench [23]	✓	✓	✗	✗	✗	✓	✗
RobotPerf (ours)	✓	✓	✓	✓	✗	✓	✓

TABLE I: Comparative evaluation of representative existing robotics benchmarks with RobotPerf across essential characteristics for robotic systems.

ing their performance as well as a plethora of workshops and tutorials focusing on benchmarking robotics applications [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22]. However, most of these robotics benchmarks focus on algorithm correctness (*functional* testing) in the context of domain specific problems, as well as end-to-end latency on CPUs [31], [32], [33], [34], [37], [35], [38], [36], [39], [40], [41], [42], [43], [44], [45], [46], [47], [48]. A few works also analyze some *non-functional* metrics, such as CPU performance benchmarks, to explore bottleneck behaviors in selected workloads [23], [24], [49].

Recent work has also explored the implications of operating systems and task schedulers on ROS 2 computational graph performance through benchmarking [50], [51], [52], [53], [54] as well as by optimizing the scheduling and communication layers of ROS and ROS 2 themselves [55], [56], [57], [58], [59], [60], [61], [62]. These works often focused on a specific context or (set of) performance counter(s).

Finally, previous work has leveraged hardware acceleration for select ROS Nodes and adaptive computing to optimize the ROS computational graphs [63], [64], [65], [66], [67], [68], [69], [70], [71], [72], [73], [74], [75], [76], [77], [78], [79]. However, these works do not provide comprehensive frameworks to quickly analyze and evaluate new heterogeneous computational graphs except for two works that are limited to the context of UAVs [25], [28].

Research efforts most closely related to our work include *ros2_tracing* [80] and RobotCore [5]. *ros2_tracing* provided instrumentation that demonstrated integration with the low-overhead LTTng tracer into ROS 2, while RobotCore illuminates the advantages of using vendor-specific tracing to complement *ros2_tracing* to assess the performance of hardware-accelerated ROS 2 Nodes. Building on these

two specific foundational contributions, RobotPerf offers a comprehensive set of ROS 2 kernels spanning the robotics pipeline and evaluates them on diverse hardware.

Table I summarizes our unique contributions. It includes a selection of representative benchmarks from above and provides an evaluation of these benchmarks against RobotPerf, focusing on essential characteristics vital for robotic systems. We note that while our current approach focuses only on non-functional performance benchmarking tests, RobotPerf’s architecture and methodology can be extended to also measure functional metrics.

III. ROBOTPERF: PRINCIPLES & METHODOLOGY

RobotPerf is an open-source, industry-strength robotics benchmark for portability across heterogeneous hardware platforms. This section outlines the important design principles and describes the implementation methodology.

A. Non-Functional Performance Testing

Currently, RobotPerf specializes in non-functional performance testing, evaluating the efficiency and operational characteristics of robotic systems. Non-functional performance testing measures those aspects not belonging to the system’s functions, such as computational latency, memory consumption, and CPU usage. In contrast, traditional functional performance testing looks into the system’s specific tasks and function, verifying its effectiveness in its primary goals, like the accuracy of the control algorithm in following a planned robot’s path. While functional testing confirms a system performs its designated tasks correctly, non-functional testing ensures it operates efficiently and reliably.

B. ROS 2 Integration & Adaptability

RobotPerf is designed specifically to evaluate ROS 2 computational graphs, rather than focusing on independent robotic algorithms. We emphasize benchmarking *ROS 2 workloads* because the use of ROS 2 as middleware allows for the easy composition of complex robotic systems. This makes the benchmark versatile and well-suited for a wide range of robotic applications and enables industry, which is widely using ROS, to rapidly adopt RobotPerf.

C. Platform Independence & Portability

RobotPerf allows for the evaluation of benchmarks on a variety of hardware platforms, including general-purpose CPUs and GPUs, reconfigurable FPGAs, and specialized accelerators (e.g., ray tracing accelerators [81]). Benchmarking robotic workloads on heterogeneous platforms is vital to evaluate their respective capabilities and limitations. This facilitates optimizations for efficiency, speed, and adaptability, as well as fine-tuning of resource allocations, ensuring robust and responsive operation across diverse contexts.

D. Flexible Methodology

We offer grey-box and black-box testing methods to suit different needs. Black-box testing provides a quick-to-enable external perspective and measures performance by eliminating the layers above the layer-of-interest and replacing those

CRITERIA	Grey-Box	Black-Box
PRECISION	Utilizes tracers from in-code instrumentation.	Limited to ROS 2 message subscriptions.
PERFORMANCE	Low overhead. Driven by kernelspace.	Restricted to ROS 2 message callbacks. Recorded by userspace processes.
FLEXIBILITY	Multiple event types.	Limited to message subscriptions in current implementation.
PORTABILITY	Requires a valid tracer. Standard format (CTF).	Standard ROS 2 APIs. Custom JSON format.
EASE OF USE	Requires code modifications and data postprocessing.	Tests unmodified software with minor node additions.
REAL-ROBOTS	Does not modify the computational graph.	Modifies the computational graph adding extra dataflow.

TABLE II: Grey-box vs. black-box benchmarking trade-offs.

with a specific test application. Grey-box testing provides more granularity and dives deeper into the internal workings of ROS 2, allowing users to generate more accurate measurements at the cost of increased engineering effort. As such, each method has its trade-offs, and providing both options enables users flexibility. We describe each method in more detail below and highlight takeaways in Table II.

1) *Grey-Box Testing*: Grey-box testing enables precise probe placement within a robot’s computational graph, generating a chronologically ordered log of critical events using a tracer that could be proprietary or open source, such as LTTng [82]. As this approach is fully integrated with standard ROS 2 layers and tools through `ros2.tracing`, it incurs a minimal average latency of only 3.3 μ s [80], making it well-suited for real-time systems. With this approach, optionally, RobotPerf offers specialized input and output nodes that are positioned outside the nodes of interest to avoid the need to instrument them. These nodes generate the message tracepoints upon publish and subscribe events which are processed to calculate end-to-end latency.

2) *Black-Box Testing*: The black-box methodology utilizes a user-level node called the `MonitorNode` to evaluate the performance of a ROS 2 node. The `MonitorNode` subscribes to the target node, recording the timestamp when each message is received. By accessing the propagated ID, the `MonitorNode` determines the end-to-end latency by comparing its timestamp against the `PlaybackNode`’s recorded timestamp for each message. While this approach does not need extra instrumentation, and is easier to implement, it offers a less detailed analysis and alters the computational graph by introducing new nodes and dataflow.

E. Opaque Performance Tests

The requirement for packages to be instrumented directly within the source code poses a challenge to many benchmarking efforts. To overcome this hurdle, for most benchmarks, we refrain from altering the workloads of interest and, instead, utilize specialized input and output nodes positioned outside the primary nodes of concern. This setup allows for benchmarking without the need for direct instrumentation of

Category	Benchmark Name	Description
Perception	a1_perception_2nodes	Graph with 2 components: <code>rectify</code> and <code>resize</code> [83], [84].
	a2_rectify	<code>rectify</code> component [83], [84].
	a3_stereo_image_proc	Computes disparity map from left and right images [85].
	a4_depth_image_proc	Computes point cloud from rectified depth and color images [86].
	a5_resize	<code>resize</code> component [83], [84].
Localization	b1_visual_slam	Visual SLAM component [87].
	b2_map_localization	Map localization component [88].
	b3_apriltag_detection	Apriltag detection component [89].
Control	c1_rrbot_joint_trajectory_controller	Joint trajectory controller [90].
	c2_diffbot_diff_driver_controller	Differential driver controller [91].
	c3_rrbot_forward_command_controller_position	Position-based forward command controller [92].
	c4_rrbot_forward_command_controller_velocity	Velocity-based forward command controller [92].
	c5_rrbot_forward_command_controller_acceleration	Acceleration-based forward command controller [92].
Manipulation	d1_xarm6_planning_and_traj_execution	Manipulator planning and trajectory execution [93].
	d2_collision_checking_fcl	Collision check: manipulator and box (FCL [94]).
	d3_collision_checking_bullet	Collision check: manipulator and box (Bullet [95]).
	d4_inverse_kinematics_kdl	Inverse kinematics (KDL plugin [96]).
	d5_inverse_kinematics_lma	Inverse kinematics (LMA plugin [97]).
	d6_direct_kinematics	Direct kinematics for manipulator [93].

TABLE III: RobotPerf beta Benchmarks (see [98]).

the target layer. We term this methodology “opaque tests,” a concept that RobotPerf adheres to when possible.

F. Reproducibility & Consistency

To ensure consistent and reproducible evaluations, RobotPerf adheres to specific common robotic dataformats. In particular, it uses ROS 2 `rosbags`, including our own available at <https://github.com/robotperf/rosbags>, as well third-party bags (e.g., the `r2b` dataset [99]).

To ensure consistent data loading and finer control over message delivery rates, we drew inspiration from [100]. Our computational graphs incorporate *modified and improved* `DataLoaderNode` and `PlaybackNode` implementations, which can be accessed at https://github.com/robotperf/ros2_benchmark. These enhanced nodes offer improvements that report worst-case latency and enable the reporting of maximum latency, introduce the ability to profile power consumption and so forth.

G. Metrics

We focus on three key metrics: latency, throughput and power consumption including energy efficiency. Latency measures the time between the start and the completion of a task. Throughput measures the total amount of work done in a given time for a task. Power measures the electrical energy per unit of time consumed while executing a given task. Measuring energy efficiency (or performance-per-Watt) captures the total amount of work (relative to either throughput or latency) that can be delivered for every watt of power consumed and is directly related to the runtime of battery powered robots [25].

H. Current Benchmarks and Categories

RobotPerf beta [98] introduces benchmarks that cover the robotics pipeline from perception, to localization, to

control, as well as dedicated benchmarks for manipulation. The full list of benchmarks in the beta release can be found in Table III. Aligned with our principles defined above, each benchmark is a self-contained ROS 2 package which describes all dependencies (generally other ROS packages). To facilitate reproducibility, all benchmarks are designed to be built and run using the common ROS 2 development flows (`ament` build tools, `colcon` meta-build tools, etc.). Finally, so that the benchmarks can be easily consumed by other tools, a description of each benchmark, as well as its results, is defined in a machine-readable format. As such, accompanying the `package.xml` and `CMakeLists.txt` files required for all ROS packages, a YAML file named `benchmark.yaml` is in the root of each benchmark which describes the benchmark and includes accepted results.

I. Run Rules

To ensure the reliability and reproducibility of the performance data, we adhere to a stringent set of run rules. First, tests are performed in a controlled environment to ensure that performance data is not compromised by fluctuating external parameters. As per best practices recommended by `ros2_tracing` [80], we record and report settings like clock frequency and core count. Second, we look forward to the possibility of RobotPerf being embraced by the community and have results undergo peer review, which can contribute to enhancing reproducibility and accuracy. Finally, we aim to avoid overfitting to specific hardware setups or software configurations by encompassing a broad spectrum of test scenarios.

IV. EVALUATION

We conduct comprehensive benchmarking using RobotPerf to evaluate its capabilities on three key aspects vital for a robotics-focused computing benchmark. First, we validate

the framework’s capacity to provide comparative insights across divergent heterogeneous platforms from edge devices to server-class hardware. Second, we analyze the results to understand RobotPerf’s ability to guide selection of the optimal hardware solution tailored to particular robotic workloads. Finally, we assess how effectively RobotPerf reveals the advantages conferred by hardware and software acceleration techniques relative to general-purpose alternatives. All of our results and source code can be found open-source at: <https://github.com/robotperf/benchmarks>.

A. Fair and Representative Assessment of Heterogeneity

Assessing hardware heterogeneity in robotic applications is imperative in the ever-evolving field of robotics. Different robotic workloads demand varying computational resources and efficiency levels. Therefore, comprehensively evaluating performance across diverse hardware platforms is crucial.

We evaluated the RobotPerf benchmarks over a wide list of hardware platforms, including general-purpose CPUs on edge devices (e.g., Qualcomm RB5), server-class CPUs (e.g., Intel i7-8700), and specialized hardware accelerators (e.g., AMD Kria KR260). Figure 3 illustrates benchmark performance in robotics per category of workload (perception, localization, control, and manipulation) using radar plots, wherein the different hardware solutions are depicted together alongside different robotic workloads per category. Each hardware solution is presented with a different color, with smaller values and areas representing better performance in the respective category. Given our ability to benchmark 18 platforms (bottom of Figure 3), RobotPerf is capable of benchmarking heterogeneous hardware platforms and workloads, paving the way for community-driven co-design and optimization of hardware and software.

B. Quantitative Approach to Hardware Selection

The rapid evolution and diversity of tasks in robotics means we need to have a meticulous and context-specific approach to computing hardware selection and optimization. A “one-size-fits-all” hardware strategy would be an easy default selection, but it fails to capitalize on the nuanced differences in workload demands across diverse facets like perception, localization, control, and manipulation, each exhibiting distinctive sensitivities to hardware capabilities. Therefore, a rigorous analysis, guided by tools like RobotPerf, becomes essential to pinpoint the most effective hardware configurations that align well with individual workload requirements.

The results in Figure 3 demonstrate the fallacy of a “one-size-fits-all” solution. For example, focusing in on the latency radar plot for control from Figure 3 (col 3, row 1), we see that the i7-12700H (I7H) outperforms the NVIDIA AGX Orin Dev. Kit (NO) on benchmarks C1, C3, C4, and C5, but is 6.5× slower on benchmark C2. As such, by analyzing data from the RobotPerf benchmarks, roboticists can better determine which hardware option best suits their needs given their specific workloads and performance requirements.

One general lesson learned while evaluating the data is that each workload is unique, making it hard to general-




Fig. 2: Benchmark comparison of perception latency (ms) on AMD’s Kria KR260 with and without the ROBOTCORE Perception accelerator. The benchmarks used are a1, a2, and a5 as defined in Table III. We find that hardware acceleration can enable performance gains of as much as 11.5×.

ize across both benchmarks and categories. To that end, RobotPerf results help us understand how the use of various hardware solutions and dedicated domain-specific hardware accelerators significantly improves the performance.

C. Rigorous Assessment of Acceleration Benefits

In the rapidly advancing field of computing hardware, the optimization of algorithm implementations is a crucial factor in determining the success and efficiency of robotic applications. The need for an analytical tool, like RobotPerf, that facilitates the comparison of various algorithmic implementations on uniform hardware setups becomes important.

Figure 2 is a simplified version of Figure 3, depicting AMD’s Kria KR260 hardware solution in two forms: the usual hardware and a variant that leverages a domain-specific hardware accelerator (ROBOTCORE Perception, a soft-core running in the FPGA for accelerating perception robotic computations). The figure demonstrates that hardware acceleration can enable performance gains of as much as 11.5× (from 173 ms down to 15 ms for benchmark a5). We stress that the results obtained here should be interpreted according to each end application and do not represent a generic recommendation on which hardware should be used. Other factors, including availability, the form factor, and community support, are relevant aspects to consider when selecting a hardware solution.

V. CONCLUSION AND FUTURE WORK

RobotPerf represents an important step towards standardized benchmarking in robotics. With its comprehensive evaluation across the hardware/software stack and focus on industry-grade ROS 2 deployments, RobotPerf can pave the way for rigorous co-design of robotic hardware and algorithms. As RobotPerf matures with community involvement, we expect it to compare CPU, GPU and FPGA, exploring their power consumption and flexibility in augmenting real-world robotic computations. With a standardized robotics benchmark as a focal point, the field can make rapid progress in delivering real-time capable systems that will unlock the true potential of robotics in real-world applications.




Fig. 3: Benchmarking results on diverse hardware platforms across perception, localization, and manipulation workloads defined in RobotPerf β Benchmarks. Radar plots illustrate the latency, throughput, and power consumption for each hardware solution and workload, with reported values representing the maximum across a series of runs. The labels of vertices represent the workloads defined in Table III. Each hardware platform and performance testing procedure is delineated by a separate color, with darker colors representing Black-box testing and lighter colors Grey-box testing. In the figure’s key, the hardware platforms are categorized into four specific types: general-purpose hardware, heterogeneous hardware, reconfigurable hardware, and accelerator hardware. Within each category, the platforms are ranked based on their Thermal Design Power (TDP), which indicates the maximum power they can draw under load. The throughput values for manipulation tasks and power values for localization tasks have not been incorporated into the β version of RobotPerf. As RobotPerf continues to evolve, more results will be added in subsequent iterations.

REFERENCES

- [1] S. M. Neuman, B. Plancher, T. Bourgeat, T. Tambe, S. Devadas, and V. J. Reddi, "Robomorphic computing: a design methodology for domain-specific accelerators parameterized by robot morphology," in *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2021, pp. 674–686.
- [2] W. Liu, B. Yu, Y. Gan, Q. Liu, J. Tang, S. Liu, and Y. Zhu, "Archytas: A framework for synthesizing and dynamically optimizing accelerators for robotic localization," in *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, 2021, pp. 479–493.
- [3] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, et al., "Isaac gym: High performance gpu-based physics simulation for robot learning," *arXiv preprint arXiv:2108.10470*, 2021.
- [4] B. Plancher, S. M. Neuman, R. Ghosal, S. Kuindersma, and V. J. Reddi, "Grid: Gpu-accelerated rigid body dynamics with analytical gradients," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6253–6260.
- [5] V. Mayoral-Vilches, S. M. Neuman, B. Plancher, and V. J. Reddi, "Robotcore: An open architecture for hardware acceleration in ros 2," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 9692–9699.
- [6] Z. Wan, A. Lele, B. Yu, S. Liu, Y. Wang, V. J. Reddi, C. Hao, and A. Raychowdhury, "Robotic computing on fpgas: Current progress, research challenges, and opportunities," in *2022 IEEE 4th International Conference on Artificial Intelligence Circuits and Systems (AICAS)*. IEEE, 2022, pp. 291–295.
- [7] S. Liu, Z. Wan, B. Yu, and Y. Wang, *Robotic computing on fpgas*. Springer, 2021.
- [8] H. Esmailzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger, "Dark Silicon and the End of Multicore Scaling," in *Proceedings of the 38th Annual International Symposium on Computer Architecture*, ser. ISCA '11. ACM, pp. 365–376.
- [9] G. Venkatesh, J. Sampson, N. Goulding, S. Garcia, V. Bryksin, J. Lugo-Martinez, S. Swanson, and M. B. Taylor, "Conservation Cores: Reducing the Energy of Mature Computations," in *Proceedings of the Fifteenth Edition of ASPLOS on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS XV. ACM, pp. 205–218.
- [10] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.
- [11] V. Mayoral-Vilches, "ros-robotics-companies," <https://github.com/vmayoral/ros-robotics-companies>, [Accessed: July 9, 2023].
- [12] "Icra2021 workshop cloud-based competitions and benchmarks for robotic manipulation and grasping," June 2021. [Online]. Available: <https://sites.google.com/view/icra2021-workshop/home>
- [13] "Icra 2022 workshop determining appropriate metrics and test methods for soft actuators in robotic systems," May 2022. [Online]. Available: <https://sites.google.com/andrew.cmu.edu/softactuatorsmetrics/>
- [14] "Icra 2022 workshop on releasing robots into the wild: Simulations, benchmarks, and deployment," May 2022. [Online]. Available: <https://www.dynsyslab.org/releasing-robots-into-the-wild-workshop/>
- [15] "Iros 2020 workshop on benchmarking progress in autonomous driving," Oct. 2020. [Online]. Available: <https://www.robotics.qmul.ac.uk/events/iros-2021-workshop/>
- [16] "Iros 2021 workshop - benchmarking of robotic grasping and manipulation: protocols, metrics and data analysis," Sept. 2021. [Online]. Available: <https://www.robotics.qmul.ac.uk/events/iros-2021-workshop/>
- [17] "Evaluating motion planning performance," Oct. 2022. [Online]. Available: <https://motion-planning-workshop.kavrakilab.org/>
- [18] "Methods for objective comparison of results in intelligent robotics research," Oct. 2023. [Online]. Available: http://www.robot.t.u-tokyo.ac.jp/TCPEBRAS_IROS2023/index.html
- [19] "Benchmarking tools for evaluating robotic assembly of small parts," July 2020. [Online]. Available: <https://www.uml.edu/research/nerve/assembly-workshop-rss-2020.aspx>
- [20] "2021 rss workshop on advancing artificial intelligence and manipulation for robotics: Understanding gaps, industry and academic perspectives, and community building," July 2021. [Online]. Available: <https://sites.google.com/view/rss-ai-manipulationperspective/home>
- [21] "Robot learning in the cloud: Remote operations and benchmarking," July 2022. [Online]. Available: <https://sites.google.com/andrew.cmu.edu/cloud-robotics-benchmarking/>
- [22] "Datasets and benchmarking tools for advancing and evaluating robotic manufacturing," July 2023. [Online]. Available: <https://sites.google.com/view/rss-2023-nist-moad>
- [23] M. Bakhshalipour, M. Likhachev, and P. B. Gibbons, "Rtrbench: A benchmark suite for real-time robotics," in *2022 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 2022, pp. 175–186.
- [24] S. M. Neuman, T. Koolen, J. Drean, J. E. Miller, and S. Devadas, "Benchmarking and workload analysis of robot dynamics algorithms," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 5235–5242.
- [25] B. Boroujerdian, H. Genc, S. Krishnan, W. Cui, A. Faust, and V. Reddi, "Mavbench: Micro aerial vehicle benchmarking," in *2018 51st annual IEEE/ACM international symposium on microarchitecture (MICRO)*. IEEE, 2018, pp. 894–907.
- [26] S. Krishnan, Z. Wan, K. Bhardwaj, P. Whatmough, A. Faust, S. M. Neuman, G.-Y. Wei, D. Brooks, and V. J. Reddi, "Automatic domain-specific soc design for autonomous unmanned aerial vehicles," in *2022 55th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2022, pp. 300–317.
- [27] S. Krishnan, Z. Wan, K. Bhardwaj, N. Jadhav, A. Faust, and V. J. Reddi, "Roofline model for uavs: A bottleneck analysis tool for onboard compute characterization of autonomous unmanned aerial vehicles," in *2022 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 2022, pp. 162–174.
- [28] D. Nikiforov, S. C. Dong, C. L. Zhang, S. Kim, B. Nikolic, and Y. S. Shao, "Rosé: A hardware-software co-simulation infrastructure enabling pre-silicon full-stack robotics soc evaluation," in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, 2023, pp. 1–15.
- [29] Object Management Group (OMG), "Data Distribution Service (DDS) Interoperability Wire Protocol RTPS (Real-Time Publish-Subscribe) Protocol Specification Version 2.5," <https://www.omg.org/spec/DDS-I-RTPS/2.5/>, [Accessed: July 9, 2023].
- [30] V. Mayoral-Vilches and G. Corradi, "Adaptive computing in robotics, towards ros 2 software-defined hardware," *Xilinx, WP537*, 2021.
- [31] I. A. Sucan, M. Moll, and L. E. Kavraki, "The open motion planning library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [32] C. Chamzas, C. Quintero-Pena, Z. Kingston, A. Orthey, D. Rakita, M. Gleicher, M. Toussaint, and L. E. Kavraki, "Motionbenchmarker: A tool to generate and benchmark motion planning datasets," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 882–889, 2021.
- [33] T. Tan, R. Weller, and G. Zachmann, "Opencollbench-benchmarking of collision detection & proximity queries as a web-service," in *The 25th International Conference on 3D Web Technology*, 2020, pp. 1–9.
- [34] D. Perille, A. Truong, X. Xiao, and P. Stone, "Benchmarking metric ground navigation," in *2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, 2020, pp. 116–121.
- [35] A. Nair, F. Jiang, K. Hou, Z. Xu, S. Li, X. Xiao, and P. Stone, "Dynabarn: Benchmarking metric ground navigation in dynamic environments," *navigation*, vol. 7, p. 9.
- [36] E. Heiden, L. Palmieri, L. Bruns, K. O. Arras, G. S. Sukhatme, and S. Koenig, "Bench-mr: A motion planning benchmark for wheeled mobile robots," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4536–4543, 2021.
- [37] M. Moll, I. A. Sucan, and L. E. Kavraki, "Benchmarking motion planning algorithms."
- [38] Z. Kingston and L. E. Kavraki, "Rowflex: Robot motion planning with moveit made easy," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3108–3114.
- [39] M. Ahn, H. Zhu, K. Hartikainen, H. Ponte, A. Gupta, S. Levine, and V. Kumar, "Robel: Robotics benchmarks for learning with low-cost robots," in *Conference on robot learning*. PMLR, 2020, pp. 1300–1313.
- [40] J. Weisz, Y. Huang, F. Lier, S. Sethumadhavan, and P. Allen, "Robobench: Towards sustainable robotics system benchmarking,"

- in 2016 *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 3383–3389.
- [41] A. P. del Pobil, R. Madhavan, and E. Messina, “Benchmarks in robotics research,” in *Workshop IROS*. Citeseer, 2006.
- [42] O. Michel, F. Rohrer, and Y. Bourquin, “Rat’s life: A cognitive robotics benchmark,” in *European Robotics Symposium 2008*. Springer, 2008, pp. 223–232.
- [43] A. Murali, T. Chen, K. V. Alwala, D. Gandhi, L. Pinto, S. Gupta, and A. Gupta, “Pyrobot: An open-source robotics framework for research and benchmarking,” *arXiv preprint arXiv:1906.08236*, 2019.
- [44] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison, “Rlbench: The robot learning benchmark & learning environment,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3019–3026, 2020.
- [45] J. Leitner, A. W. Tow, N. Sünderhauf, J. E. Dean, J. W. Durham, M. Cooper, M. Eich, C. Lehnert, R. Mangels, C. McCool, et al., “The acrv picking benchmark: A robotic shelf picking benchmark to foster reproducible research,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4705–4712.
- [46] Y. Zhu, J. Wong, A. Mandelkar, R. Martín-Martín, A. Joshi, S. Nasiriany, and Y. Zhu, “robosuite: A modular simulation framework and benchmark for robot learning,” *arXiv preprint arXiv:2009.12293*, 2020.
- [47] L. Fan, Y. Zhu, J. Zhu, Z. Liu, O. Zeng, A. Gupta, J. Creus-Costa, S. Savarese, and L. Fei-Fei, “Surreal: Open-source reinforcement learning framework and robot manipulation benchmark,” in *Conference on Robot Learning*. PMLR, 2018, pp. 767–782.
- [48] M. Althoff, M. Koschi, and S. Manzi, “Commonroad: Composable benchmarks for motion planning on roads,” in *2017 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2017, pp. 719–726.
- [49] J. Delmerico and D. Scaramuzza, “A benchmark comparison of monocular visual-inertial odometry algorithms for flying robots,” in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 2502–2509.
- [50] M. Reke, D. Peter, J. Schulte-Tiggies, S. Schiffer, A. Ferrein, T. Walter, and D. Matheis, “A self-driving car architecture in ros2,” in *2020 International SAUPEC/RobMech/PRASA Conference*. IEEE, 2020, pp. 1–6.
- [51] S. Barut, M. Boneberger, P. Mohammadi, and J. J. Steil, “Benchmarking real-time capabilities of ros 2 and orocos for robotics applications,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 708–714.
- [52] L. Puck, P. Keller, T. Schnell, C. Plasberg, A. Tanev, G. Heppner, A. Roennau, and R. Dillmann, “Distributed and synchronized setup towards real-time robotic control using ros2 on linux,” in *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2020, pp. 1287–1293.
- [53] Y. Yang and T. Azumi, “Exploring real-time executor on ros 2,” in *IEEE International Conference on Embedded Software and Systems (ICSS)*, 2020, pp. 1–8.
- [54] A. A. Arafat, S. Vaidhun, K. M. Wilson, J. Sun, and Z. Guo, “Response time analysis for dynamic priority scheduling in ros2,” in *Proceedings of the 59th ACM/IEEE Design Automation Conference*, 2022, pp. 301–306.
- [55] Y. Sugata, T. Ohkawa, K. Ootsu, and T. Yokota, “Acceleration of publish/subscribe messaging in ros-compliant fpga component,” in *International Symposium on Highly Efficient Accelerators and Reconfigurable Technologies*, 2017, pp. 1–6.
- [56] T. Ohkawa, Y. Sugata, H. Watanabe, N. Ogura, K. Ootsu, and T. Yokota, “High level synthesis of ros protocol interpretation and communication circuit for fpga,” in *IEEE/ACM International Workshop on Robotics Software Engineering (RoSE)*, 2019, pp. 33–36.
- [57] H. Choi, Y. Xiang, and H. Kim, “Picas: New design of priority-driven chain-aware scheduling for ros2,” in *IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2021, pp. 251–263.
- [58] Y. Suzuki, T. Azumi, S. Kato, and N. Nishio, “Real-time ros extension on transparent cpu/gpu coordination mechanism,” in *IEEE International Symposium on Real-Time Distributed Computing (ISORC)*, 2018, pp. 184–192.
- [59] C. S. V. Gutiérrez, L. U. S. Juan, I. Z. Ugarte, and V. Mayoral-Vilches, “Time-sensitive networking for robotics,” *arXiv preprint arXiv:1804.07643*, 2018.
- [60] —, “Real-time linux communications: an evaluation of the linux communication stack for real-time robotic applications,” *arXiv preprint arXiv:1808.10821*, 2018.
- [61] —, “Towards a distributed and real-time framework for robots: Evaluation of ros 2.0 communications for real-time robotic applications,” *arXiv preprint arXiv:1809.02595*, 2018.
- [62] C. S. V. Gutiérrez, L. U. S. Juan, I. Z. Ugarte, I. M. Goenaga, L. A. Kirschgens, and V. Mayoral-Vilches, “Time synchronization in modular collaborative robots,” *arXiv preprint arXiv:1809.07295*, 2018.
- [63] K. Yamashina, T. Ohkawa, K. Ootsu, and T. Yokota, “Proposal of ros-compliant fpga component for low-power robotic systems: case study on image processing application,” *International Workshop on FPGAs for Software Programmers (FSP)*, 2015.
- [64] K. Yamashina, H. Kimura, T. Ohkawa, K. Ootsu, and T. Yokota, “cre-comp: Automated design tool for ros-compliant fpga component,” in *IEEE International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSOC)*, 2016, pp. 138–145.
- [65] A. Podlubne and D. Göhringer, “Fpga-ros: Methodology to augment the robot operating system with fpga designs,” in *IEEE International Conference on ReConfigurable Computing and FPGAs (ReConFig)*, 2019, pp. 1–5.
- [66] M. Eisoldt, S. Hinderink, M. Tassemeier, M. Flottmann, J. Vana, T. Wiemann, J. Gaal, M. Rothmann, and M. Pörmann, “Reconfros: Running ros on reconfigurable socs,” in *Drone Systems Engineering and Rapid Simulation and Performance Evaluation: Methods and Tools*, 2021, pp. 16–21.
- [67] C. Lienen, M. Platzner, and B. Rinner, “Reconros: Flexible hardware acceleration for ros2 applications,” in *International Conference on Field-Programmable Technology (ICFPT)*, 2020, pp. 268–276.
- [68] D. P. Leal, M. Sugaya, H. Amano, and T. Ohkawa, “Automated integration of high-level synthesis fpga modules with ros2 systems,” in *International Conference on Field-Programmable Technology (ICFPT)*, 2020, pp. 292–293.
- [69] T. Ohkawa, K. Yamashina, T. Matsumoto, K. Ootsu, and T. Yokota, “Architecture exploration of intelligent robot system using ros-compliant fpga component,” in *IEEE International Symposium on Rapid System Prototyping (RSP)*, 2016, pp. 1–7.
- [70] S. Panadda, J. Nattha, P. L. Daniel, and O. Takeshi, “Low-power high-performance intelligent camera framework ros-fpga node,” in *Asia Pacific Conference on Robot IoT System Development and Platform*, no. 2020, 2021, pp. 73–74.
- [71] J. P. Queralt, F. Yuhong, L. Salomaa, L. Qingqing, T. N. Gia, Z. Zou, H. Tenhunen, and T. Westerlund, “Fpga-based architecture for a low-cost 3d lidar design and implementation from multiple rotating 2d lidars with ros,” in *IEEE SENSORS*, 2019, pp. 1–4.
- [72] T. K. Maiti, “Ros on arm processor embedded with fpga for improvement of robotic computing,” in *International Symposium on Devices, Circuits and Systems (ISDCS)*, 2021, pp. 1–4.
- [73] T. Ohkawa, K. Yamashina, H. Kimura, K. Ootsu, and T. Yokota, “Fpga components for integrating fpgas into robot systems,” *IEICE Transactions on Information and Systems*, vol. 101, no. 2, pp. 363–375, 2018.
- [74] D. P. Leal, M. Sugaya, H. Amano, and T. Ohkawa, “Fpga acceleration of ros2-based reinforcement learning agents,” in *International Symposium on Computing and Networking Workshops*, 2020, pp. 106–112.
- [75] H. Amano, H. Mori, A. Mizutani, T. Ono, Y. Yoshimoto, T. Ohkawa, and H. Tamukoh, “A dataset generation for object recognition and a tool for generating ros2 fpga node,” in *IEEE International Conference on Field-Programmable Technology (ICFPT)*, 2021, pp. 1–4.
- [76] Y. Nitta, S. Tamura, and H. Takase, “A study on introducing fpga to ros based autonomous driving system,” in *IEEE International Conference on Field-Programmable Technology (FPT)*, 2018, pp. 421–424.
- [77] K. E. Chen, Y. Liang, N. Jha, J. Ichnowski, M. Danielczuk, J. Gonzalez, J. Kubiatowicz, and K. Goldberg, “Fogros: An adaptive framework for automating fog robotics deployment,” in *IEEE International Conference on Automation Science and Engineering (CASE)*, 2021, pp. 2035–2042.
- [78] NVIDIA, “NVIDIA Isaac ROS,” Accessed 2022, github.com/NVIDIA-A-ISAAC-ROS.
- [79] Z. Wan, K. Swaminathan, P.-Y. Chen, N. Chandramoorthy, and A. Raychowdhury, “Analyzing and improving resilience and robustness of autonomous systems,” in *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, 2022, pp. 1–9.
- [80] C. Bédard, I. Lütkebohle, and M. Dagenais, “ros2.tracing: Multi-purpose low-overhead framework for real-time tracing of ros 2,” Accessed 2022, gitlab.com/ros-tracing/ros2.tracing.

- [81] Y. Deng, Y. Ni, Z. Li, S. Mu, and W. Zhang, "Toward real-time ray tracing: A survey on hardware acceleration and microarchitecture techniques," *ACM Computing Surveys (CSUR)*, vol. 50, no. 4, pp. 1–41, 2017.
- [82] M. Desnoyers and M. R. Dagenais, "The lttng tracer: A low impact performance and behavior monitor for gnu/linux," in *OLS (Ottawa Linux Symposium)*, vol. 2006. Citeseer, 2006, pp. 209–224.
- [83] R.-A. Community, "ros-acceleration," 2023, gitHub repository. [Online]. Available: <https://github.com/ros-acceleration>
- [84] R. P. Developers, "image_proc subdirectory on humble branch of image pipeline," GitHub Repository, 2023. [Online]. Available: https://github.com/ros-perception/image_pipeline/tree/humble/image_proc
- [85] —, "stereo_image_proc subdirectory on humble branch of image pipeline," GitHub Repository, 2023. [Online]. Available: https://github.com/ros-perception/image_pipeline/tree/humble/stereo_image_proc
- [86] —, "depth_image_proc subdirectory on humble branch of image pipeline," GitHub Repository, 2023. [Online]. Available: https://github.com/ros-perception/image_pipeline/tree/humble/depth_image_proc
- [87] N. I. R. Developers, "isaac_ros_visual_slam," GitHub Repository, 2023. [Online]. Available: https://github.com/NVIDIA-ISAAC-ROS/isaac_ros_visual_slam
- [88] —, "isaac_ros_map_localization," GitHub Repository, 2023. [Online]. Available: https://github.com/NVIDIA-ISAAC-ROS/isaac_ros_map_localization
- [89] —, "isaac_ros_apriltag," GitHub Repository, 2023. [Online]. Available: https://github.com/NVIDIA-ISAAC-ROS/isaac_ros_apriltag
- [90] R. C. Developers, "joint_trajectory_controller subdirectory in ros 2 controllers," GitHub Repository, 2023. [Online]. Available: https://github.com/ros-controls/ros2_controllers/tree/master/joint_trajectory_controller
- [91] —, "diff_drive_controller subdirectory in ros 2 controllers," GitHub Repository, 2023. [Online]. Available: https://github.com/ros-controls/ros2_controllers/tree/master/diff_drive_controller
- [92] —, "forward_command_controller subdirectory in ros 2 controllers," GitHub Repository, 2023. [Online]. Available: https://github.com/ros-controls/ros2_controllers/tree/master/forward_command_controller
- [93] M. Maintainers, "Moveit," <https://moveit.ros.org/>, 2023, official website.
- [94] F. C. L. Developers, "Flexible collision library (fcl)," <https://github.com/flexible-collision-library/fcl>, 2023, gitHub repository.
- [95] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016–2021.
- [96] M. Documentation, "The kdl kinematics plugin," MoveIt Documentation, 2023, available online: https://moveit.picknik.ai/main/doc/examples/kinematics_configuration/kinematics_configuration_tutorial.html#the-kdl-kinematics-plugin.
- [97] —, "The lma kinematics plugin," MoveIt Documentation, 2023, available online: https://moveit.picknik.ai/main/doc/examples/kinematics_configuration/kinematics_configuration_tutorial.html#the-lma-kinematics-plugin.
- [98] Robotperf. (Year of access) Robotperf Benchmarks Repository. GitHub repository directory. [Online]. Available: <https://github.com/robotperf/benchmarks/tree/main/benchmarks>
- [99] Nvidia, "R2B Dataset 2023," 4 2023. [Online]. Available: <https://catalog.ngc.nvidia.com/orgs/nvidia/teams/isaac/resources/r2bdataset2023>
- [100] N. I. ROS, "Ros2 benchmark," https://github.com/NVIDIA-ISAAC-ROS/ros2_benchmark, 2023.