

Cutting the Cord: System Architecture for Low-Cost, GPU-Accelerated Bimanual Mobile Manipulation

Artemis Shaw, Chen Liu, Justin Costa, Rane Gray, Alina Skowronek, Kevin Diaz, Nam Bui, Nikolaus Correll

Abstract—We present a bimanual mobile manipulator built on the open-source XLeRobot with integrated onboard compute for less than \$1300. Key contributions include: (1) optimized mechanical design maximizing stiffness-to-weight ratio, (2) a Tri-Bus power topology isolating compute from motor-induced voltage transients, and (3) embedded autonomy using NVIDIA Jetson Orin Nano for untethered operation. The platform enables teleoperation, autonomous SLAM navigation, and vision-based manipulation without external dependencies, providing a low-cost alternative for research and education in robotics and robot learning. Design files and documentation are available at doubleblind.

Index Terms—Bimanual manipulation, Mobile manipulation, Education Robotics

I. INTRODUCTION

There exists a wide range of low-cost, open-source platforms for robotics research and education [1]–[5]. Recent advances in vision-language-action (VLA) models and imitation learning have renewed interest in affordable hardware platforms capable of supporting data collection and policy deployment. Low-cost manipulators such as the ViperX/WidowX arms have been used in bimanual teleoperation setups [6], [7], including mobile variants [8] at significantly higher system cost. At the same time, increased interest in humanoid robotics is driving interest in an affordable and safe dual-arm mobile manipulation platform that can serve as a low-barrier baseline to compare capabilities and cost.

The Hugging Face “LeRobot” ecosystem [9], [10] has lowered the entry barrier by providing open-source designs for both manipulation and mobility, recently combined into the XLeRobot platform [11]. In its reference configuration, XLeRobot offers a tethered dual-arm mobile manipulator for less than \$700 in parts, making it one of the most accessible bimanual mobile systems currently available.

Yet, practical untethered deployment of low-cost mobile manipulators remains challenging. In particular, building systems based on online instructions introduces non-trivial systems engineering problems that are often under-documented. These include structural compliance in 3D-printed links, voltage instability under high transient motor loads, power distribution constraints, and thermal limits of embedded compute. While individually well understood, their interaction in tightly integrated, resource-constrained mobile manipulators can lead

We thank the original developers for creating XleRobot, which enabled this study.




Fig. 1. The evolved XLeRobot platform a low-cost bimanual platform for education and research assisting an industrial robot by removing screws from an electric vehicle battery.

to difficult-to-debug, hard-to-explain problems such as unpredictable resets of servos, control boards, or on-board compute, reduced payload capacity, and premature actuator wear. For example, insufficient conductor sizing and daisy-chained power rails can produce voltage collapse during concurrent multi-axis actuation and specific load patterns. Similarly, compliant 3D-printed links can act as unintended series-elastic elements [12], reducing effective torque transmission and increasing drivetrain stress. Off-the-shelf Lithium-Polymer battery packs provide high energy density, but require careful handling and protection circuitry. Commercial power stations offer a safer alternative, but impose strict current limits that must be respected at both the hardware and firmware levels. Finally, faulty start-up scripts might lead to brown-outs for some initial poses when all servos start up at once.

This paper addresses these practical but critical integration challenges through a rigorously characterized, untethered evolution of the XLeRobot platform (Fig. 1).

Our contributions are primarily at the systems level, resulting in an entry-level platform for research and education:

- *Optimized Structural Topology*: A graded 3D-printing strategy that maximizes stiffness-to-weight ratio, increasing functional load capacity while reducing mass.
- *Tri-Bus Power Topology*: A power distribution architecture that electrically isolates high-transient motor loads from embedded compute, combined with firmware-level current envelope to prevent brownouts under untethered operation.
- *Embedded Autonomy Stack*: A reproducible ROS2-based

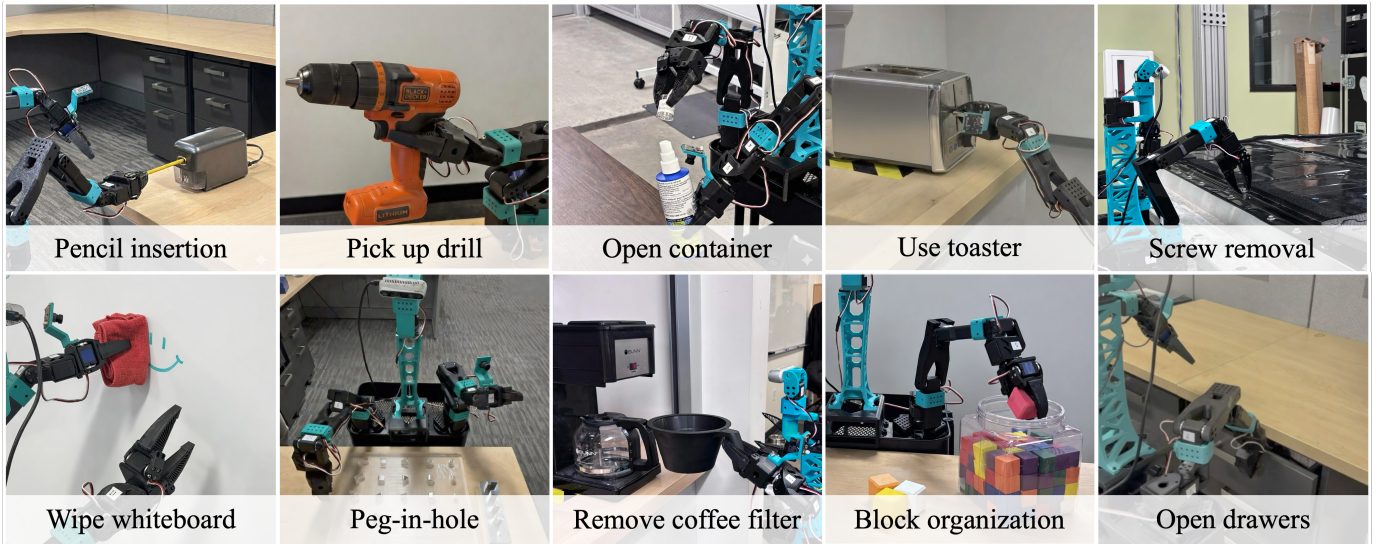


Fig. 2. Diverse Bimanual Manipulation Tasks. Demonstrations include pencil insertion into an electric pencil sharpener, picking up a YCB drill [13], bimanual container opening, toaster use, screw removal from car-battery disassembly, whiteboard wiping, peg-in-hole insertion, coffee filter removal, block organization, and drawer opening.

architecture integrating onboard perception, inverse kinematics, SLAM, navigation, and teleoperation, with benchmarking of edge compute performance using a Jetson Orin.

- *Open and Reproducible Design*: Complete CAD, electrical schematics, characterization data, and documentation to facilitate replication in research and education.

II. RELATED WORK

Table I summarizes existing dual-arm mobile manipulators and their cost.

TABLE I
COMPARISON OF DUAL-ARM MOBILE MANIPULATORS

Platform	Compute	GPU	Untethered	DoF	Payload	Est. Cost
AhaRobot [14]	Mini-ITX (i5)	RTX 4060	Yes	16	≈ 1.5 kg/arm	≈ \$2k
Cone-E [5], [15]	Intel NUC	—	Yes	16	≈ 1.5 kg/arm	< \$12k
Mobile ALOHA [8]	Laptop (onboard)	RTX 3070 Ti	Yes	16	≈ 0.75 kg/arm	≈ \$32k
TB3 Waffle Pi + arms [16], [17]	Raspberry Pi 4	—	Yes	12	≈ 0.5 kg/arm	≈ \$3k
XLeRobot [11]	External PC	—	No	15	≈ 0.8 kg	≈ \$700–1k
Ours	Jetson Orin	Ampere	Yes	17	1.0 kg/arm	≈ \$1.2k

Note: DoF includes all independently actuated joints. Costs and specifications based on 2025–2026 public BOMs and publications, although we could not corroborate the cost of the AhaRobot.

While the xLeRobot has inspired this project, achieving true autonomy requires integrating sensing and computation into the chassis. The AhaRobot [14] utilize Raspberry Pi units for onboard compute, but these setups face a significant throughput ceiling when executing high-frequency inference for Vision-Language Models (VLMs), as they lack a Graphics Processing Unit (GPU). The use of CUDA-accelerated edge modules, such as the NVIDIA Jetson Orin Nano [18], has emerged as a middle-ground solution. However, this approach introduces significant engineering trade-offs regarding energy density and power-bus stability [19].

The evolved xLeRobot aims at providing minimal requirements for untethered operation that is capable of research and education in robotics and robotic learning, i.e. providing the power infrastructure to support a GPU-enabled on-board computer, as well as providing the workspace and height that allows to perform meaningful tasks.

III. MATERIALS AND METHODS

The xLeRobot combines the LeKiwi omni-directional drive kinematics [10] with two 3 degrees-of-freedom (DoF), dual SO-101 manipulators [20] with 5+1 DoF (one DoF for the gripper), a 2-DoF neck, and a hybrid structural architecture integrating a mass-manufactured steel utility substrate (IKEA RÅSKOG) with custom FDM interfaces, reaching a height of 1.20m. To support sustained untethered autonomy, the NVIDIA Jetson Orin Nano is embedded in the chassis mid-plane using a custom enclosure with passive thermal ducting. An Intel RealSense D435 on a 2-DOF neck provides RGB-D data for SLAM navigation and manipulation policies. Table II details the system components.

A. Structural Design and Fabrications

Standard FDM profiles (2 perimeters, 15% infill, Polyactic Acid, PLA, Ouverture, 1.75mm) exhibit excessive compliance under torsional loads. We employ functionally graded topology: a High-Shell topology (4 perimeters, 15% infill) for primary links to maximize the Second Moment of Area (I) [23]. This approach provides a rigid seat for actuators, preventing motor misalignment. Conversely, compliant components, such as the joint belts stabilizing the motors, use the Baseline topology (2 perimeters, 15% infill) for assembly flexibility. Fig. 3 shows an arm along with the dimensions of individual links, with an overall reach over 40cm.

TABLE II
System Bill of Materials (BOM)

Subsystem	Component	Qty	Cost (USD)
Compute	NVIDIA Jetson Orin Nano Super [18]	1	\$249.00
Compute	microSD Card	1	\$11.23
Peripherals	Anker 4-in-1 USB-C Hub	1	\$14.99
Actuation	Feetech STS-3215 Servos (12V) [21]	17	\$271.83
Actuation	Servo Wire: 50' 3-Color	1	\$15.97
Actuation	Wonrabai Serial Bus Servo Driver Board	2	\$21.10
Power	Anker SOLIX C300 Power Station [22]	1	\$159.99
Structure	PLA Filament (3kg)	1	\$45.00
Structure	M3 Screws and Nuts Set	1	\$14.99
Frame	IKEA RÅSKOG Utility Cart	1	\$39.99
Mobility	4" Omni Wheels	3	\$29.97
Perception	Intel RealSense D435*	1	\$333.75
Power Cable	USB-C to DC5521 Cable	1	\$8.99
Power Cable	USB-C to DC5525 PD 140W Cable	1	\$11.99
Power Cable	DC5521 Car Cigarette Lighter Cable	1	\$9.49
Data Cable	USB-C to USB-C Cable (2 pcs)	1	\$8.99
TOTAL	Untethered xLeRobot	-	\$1202.28




Fig. 3. Structural composition and kinematics: (left) Robotic arm with 4-wall (black)/2-wall topology (cyan), and (right) cantilever experimental setup.

A primary goal of our topology is to maximize transmission efficiency by increasing the stiffness k of each link. A compliant link (low k) acts as a series elastic actuator [12] and absorbs motor work as elastic potential energy, defined by:


$$U = \frac{1}{2}k\delta^2 \quad (1)$$

where δ represents structural deflection. This parasitic compliance wastes torque, requiring the motor to reach its stall threshold (τ_{max}) at a lower external force. By implementing rigid 4-wall links, we minimize δ , converting torque directly into end-effector force. This allows the SO-101 to utilize the full torque profile of its actuators—enabling a 1 kg payload capacity—while protecting internal gear-trains from fatigue.


Additionally, following the original design [11] the platform utilizes a compliant gripper geometry based on the Fin Ray effect—a bio-inspired principle leveraging passive morphological computation (Fig. 3) printed in Thermoplastic Polyurethane (TPU, Ouverture, 1.75mm). The triangular rib-and-strut geometry allows the fingers to wrap around objects

when force is applied, relying on controlled material deformation rather than complex sensor feedback.

B. Power Distribution and Tri-Bus Topology



(a) Logical power schematic



(b) Hardware integration

Fig. 4. Power System Architecture. (a) The logical power schematic illustrates the distribution from the Anker SOLIX C300 to the distinct motor buses and the NVIDIA Jetson controller via high-output USB-C and car outlet interfaces. (b) The hardware integration diagram shows the physical routing of power lines within the mobile manipulator chassis, supporting the Intel RealSense D435 and dual arm configurations.

The system uses an Anker SOLIX C300 (288Wh/300W) [22] Power Distribution Unit (PDU) to provide robust over-current protection while avoiding custom LiPo assembly risks [11]. Proper power distribution is crucial for mobile manipulators; simultaneous high-torque movements generate significant voltage transients that can induce compute brownouts or system resets if the Jetson Orin Nano is not properly isolated from the actuator power rails. Quantitative testing on the original daisy-chained reference design showed that the 3A and 5A current ceilings per port were insufficient to handle peak dynamic loads. Qualitatively, this manifested as frequent compute shutdowns during bimanual manipulation.

To decouple high-torque loads from the compute domain, we implemented a Tri-Bus Topology (Fig. 4). The SOLIX C300 features three USB-C ports (two 140W, one 100W) and a 12V DC car outlet capable of 10A. We consolidate the high-draw dual manipulators onto the 10A DC car outlet. The wheels and neck actuator bus (Bus A) are powered by a high-power USB-C rail, while the Jetson Orin Nano is isolated on a separate high-power USB-C rail. The system employs $N = 17$ Feetech STS3215 servos with a stall current (I_{stall}) of 2.7 A [21]. We model the load $I(\tau, \alpha)$ as a function of

Torque (τ) and Acceleration (α) to enforce a software-defined Safe Operating Envelope (SOE):

$$I_{load}(\tau, \alpha) = N_a \cdot \left[\underbrace{\left(2.52 \cdot \frac{\tau}{1000} \right) + 0.18}_{\text{Static Load Component}} + \underbrace{D(\alpha)}_{\text{Dynamic Inrush}} \right] \quad (2)$$

Here, 0.18A represents the baseline no-load running current, and 2.52A is the net current slope required to reach the rated stall torque at 2.7A [21]. The dynamic inrush $D(\alpha)$ is mitigated using the STS3215’s internal acceleration registers.

To prevent PDU shutdown, we derived firmware-level saturation limits based on the specific constraints of the SOLIX C300. Maximum torque τ is calculated by isolating it from the current model I_{load} based on the per-port limit I_L :

$$\tau \leq \left(\frac{\frac{I_L}{N_a} - 0.18 - D(\alpha)}{2.52} \right) \times 1000 \quad (3)$$

In these models, I_L represents the maximum current limit of the specific PDU port, and N_a represents the number of active actuators on that bus. Table III summarizes these hard-coded limits, acting as a “virtual fuse”. With actuators capped at $\approx 240W$ (20A at 12V), the system maintains 60W headroom for the Jetson Orin Nano, which requires between 15W and 25W depending on inference load, ensuring zero-latency computation without risk of brownouts [18].

TABLE III
FIRMWARE SATURATION SETTINGS

Bus Group	Torque (τ) [*]	Accel (α)	Load (I_{total})	Limit (Fuse)	PDU Constraint
Bus A (Wheels/Neck)	650	20 (Med)	9.84 A	10 A	USB-C2 Output
Bus B (Arms)	450	40 (Soft)	4.90 A	5 A	DC Car Outlet
Jetson Orin	N/A	N/A	<2.1 A	5 A	USB-C3 Output

^{*} Values for τ represent the integer settings within the motor firmware registers (0–1000). Torque Conversion:
 $\tau = 450 \approx 1.32 \text{ N}\cdot\text{m}$; $\tau = 650 \approx 1.91 \text{ N}\cdot\text{m}$.

IV. SOFTWARE & CONTROL ARCHITECTURE

The platform uses a hierarchical control stack with ROS 2, Pinocchio/Pink [24], [25], and Open3D.

The NVIDIA Jetson Orin Nano Super provides 67 INT8 TOPS and 102 GB/s memory bandwidth within 7–25W [18], sufficient for vision backbones (YOLO) [26], quantized LLMs (Llama 3.1) [27], transformer policies (ACT, Diffusion Policy) [6], [28] and TensorRT accelerates inference [29].

A. Task-based inverse kinematics and manipulation

In order to support classical robot control policies, we implemented a deterministic, perception-driven manipulation pipeline that converts color and depth (RGB-D) into synchronized motor commands. The pipeline is structured as a sequence of modular, independently verifiable stages: perception, coordinate transformation, task-based inverse kinematics (IK), and motor mapping (Fig. 5).




Fig. 5. Perception-driven manipulation pipeline. (a) Color-based segmentation and centroid detection in the camera optical frame. (b) 3D mapping and representation of the environment. (c) Final grasp execution of the target object following inverse kinematics planning.

The process begins by capturing an RGB-D frame from the Intel RealSense D435. An object is identified via a color-based detector—implemented through Hue, Saturation, Value (HSV) thresholding and connected-component analysis—which computes the 3D centroid of the target label (e.g., “red”) in the camera’s optical frame (Fig. 5a).

This target point is then transformed into the SO-101 arm base frame using the robot’s MuJoCo’s MJCF kinematic model. The transformation accounts for the current head joint angles to obtain the instantaneous camera pose, applies the required coordinate alignment, and adds a small, hand-tuned offset in the base frame to compensate for systematic calibration errors.

Task-based IK is available via the pink [25] library. It formulates the task-based IK problem as a quadratic program (QP) with constraints. Given the target pose for the left or right arm and posture constraints for the additional degrees of freedom, pink will generate a joint-space trajectory $\{q_0, \dots, q_T\}$ at 100 Hz ($dt = 0.01$ s). The solver minimizes end-effector position error while applying posture regularization to avoid extreme configurations. By choosing different gains for frame tasks and posture constraints, the system will automatically trade between the mobile base and using the arm kinematics. Similarly, by restricting the movement of the opposite arm during a grab, pink is able to avoid self-collisions with the other arm. Yet, formulating IK as a constrained QP is not guaranteed to find a solution, which will require motion planning.

The trajectory is streamed as synchronized position commands. Upon reaching the target pose, a stepped gripper-closure routine (incremental position commands with short delays) is executed to achieve reliable, compliant grasping that tolerates small placement inaccuracies (Fig. 5c).

B. Mobile Navigation (SLAM & Nav2)

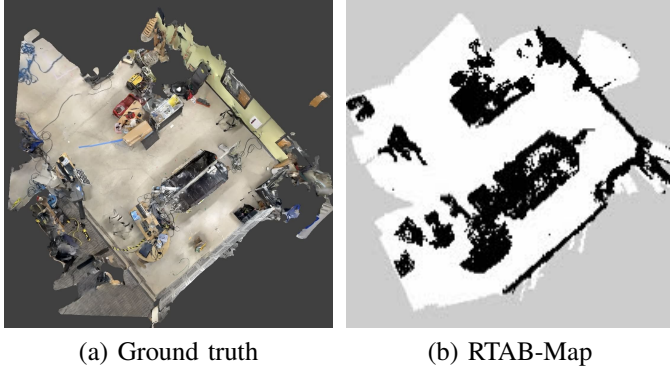


Fig. 6. SLAM Navigation Representation. (a) “Ground truth” top down view of the laboratory (Scaniverse) and (b) corresponding 2D occupancy grid for autonomous navigation.

For autonomous mobility, we implement a ROS 2 navigation pipeline that couples RGB-D perception with 2D costmap planning and control. Our current stack is designed around (1) a RealSense D435 sensor, (2) RTAB-Map [30] for visual odometry and SLAM, and (3) ROS 2 Nav2 for global planning and closed-loop execution. Fig. 6 shows a 3D map of a test environment recorded using the *Scaniverse* iPhone as a ground truth as well as a 2D occupancy map constructed using onboard sensors and RTAB-Map to the right.

We localize against a previously recorded map by loading an RTAB-Map database and running in localization-only mode (incremental mapping disabled). RTAB-Map consumes synchronized RGB, aligned depth, and camera intrinsics, and fuses visual odometry with wheel odometry (`/odom`) to produce a stable pose estimate in the `map` frame. This separation (map building vs. deployment-time localization) improves repeatability and reduces drift during long autonomous runs.

Navigation is executed with standard Nav2 servers. A `map_server` loads the static occupancy grid (Fig. 6b), the `planner_server` generates global paths, and the `controller_server` tracks paths to produce velocity commands on `/cmd_vel`. For debugging and phased deployment, the controller can be toggled on/off, allowing map/localization validation before enabling closed-loop motion.

Nav2 outputs standard `/cmd_vel`. We forward these commands to the mobile base through a bridge process that supports linear/angular scaling and optional direction inversion. This isolates drivetrain-specific details from Nav2 and enables rapid adaptation across base variants without modifying planners or controllers.

C. VR Teleoperation Architecture

To facilitate bimanual control, we implemented a Virtual Reality (VR) teleoperation system. Primarily, it provides an intuitive, human-in-the-loop interface that allows empirical validation of the kinematic workspace and dexterity of the bimanual mobile manipulator (Fig. 2 and 9). Secondly, and

crucially for modern embodied AI research and education, the system functions as a high-fidelity data collection pipeline. By recording synchronized streams of the operator’s control inputs and the robot’s sensor data (e.g. ego-centric camera feeds and joint states), the framework generates rich multimodal demonstration datasets required for large-scale imitation learning.

Building up on the Open-TeleVision architecture [31], our teleoperation framework processes both bare-hand tracking (wrist poses and pinch gestures) and standard VR controller actions. These inputs are dynamically mapped to the dual SO-101 manipulators using the real-time Inverse Kinematics (IK) solver that reliably operates at control frequencies up to 60 Hz, guaranteeing low-latency actuation. To provide an immersive telepresence experience, the robot’s onboard camera feeds are projected directly into the operator’s Head-Mounted Display (HMD). Furthermore, our framework is hardware-agnostic, supporting cross-platform deployment on headsets such as the Meta Quest, Apple Vision Pro, and Pico 4 Ultra.

V. EXPERIMENTAL EVALUATION

We evaluate: (1) structural properties and resulting manipulation capabilities, (2) computational benchmarking, (3) system stress testing, and (4) full-stack functional validation.

A. Structural Rigidity and Load Capacity Analysis

To evaluate our hybrid topology, we compared three printing strategies: Baseline (2-wall, 15% infill), High-Infill (2-wall, 50%), and High-Shell (4-wall, 15%).

We first measured structural rigidity by applying a 100g static load to the cantilevered arm at 18.6cm extension and measuring the downward deflection. As shown in the table in Fig. 7, the High-Shell configuration matched the deflection of the High-Infill profile (2.0mm) while reducing total mass by 15.8% (838g), confirming that perimeter reinforcement is the superior strategy for maximizing stiffness-to-weight ratio.

Additionally, we characterized the maximum load capacity ($n = 5$) of different topologies using a digital force gauge (SHIMPO FGV-10XY) to measure the force required to stall the actuator gearbox (rated at 1Nm, 3Nm peak). The arm was fixed in a cantilevered position while a downward vertical force was incrementally applied until reaching the actuator’s holding torque limit, manifested by either audible motor stalling or internal gearbox failure (Fig. 3). As illustrated in the box-plots in Fig. 7, High-Shell achieved a mean System Yield Force of 18.49 ± 2.15 N, a 67.1% increase over the Baseline (11.06 ± 0.63 N). These results validate that increasing wall density significantly improves functional load capacity (≈ 1 kg) for grasping tasks while optimizing material usage.

B. Grasp Reliability Validation

We evaluated gripper reliability through pick-and-place trials ($N = 75$, 15 per object) across five objects (Fig. 8) spanning various sizes and weights (17–858g), utilizing the validated 4-wall design. The system achieved a 98.7% success rate. Compliant Fin Ray mechanics provided 100% success on rigid objects (cube, bottle) by leveraging passive

Profile	Walls	Infill	Mass	Deflection (Δz)
Baseline	2	15%	758g	2 mm
High Infill	2	50%	996g	2 mm
High Shell (Ours)	4	15%	838g	2 mm

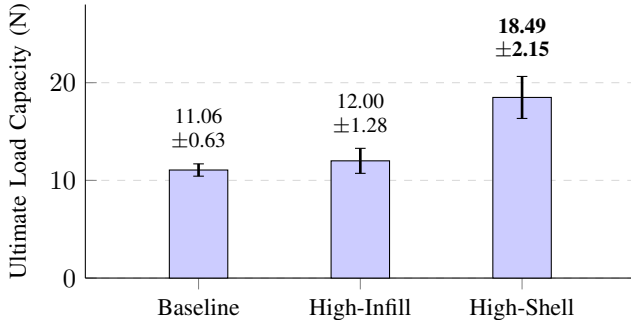


Fig. 7. Structural Performance Characterization. Comparison of link topologies across rigidity (top) and ultimate load capacity (bottom, $n = 5$).

morphological adaptation to maximize contact area. The 858g drill was successfully lifted but showed marginal stability due to high inertial loads and low-friction surfaces, leading to slip during rapid accelerations. A single failure occurred with the plastic apple due to an improper initial grasp.

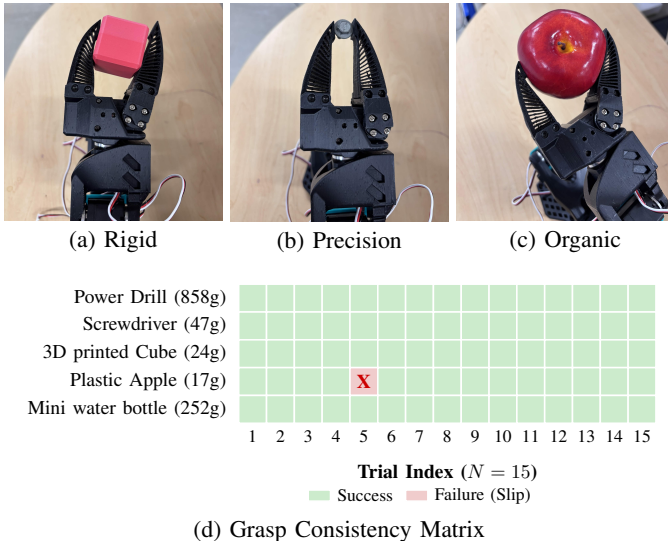


Fig. 8. Grasp Performance and Reliability. (a-c) Compliant mechanical adaptation. (d) Experimental results demonstrating a 98.7% success rate.

C. Computational Benchmarking

To evaluate system performance on edge hardware, we benchmarked three state-of-the-art manipulation models within an end-to-end control loop (encompassing camera capture, image preprocessing, batch construction, and policy inference). The evaluated models include Action Chunking with

Transformers (ACT) [6]¹, Diffusion Policy [28]², and the 450M-parameter SmolVLA [32]³.

All evaluations were conducted directly on the Jetson Orin Nano, configured in the maximum performance MAXN SUPER Mode. To optimize latency, models were run in half-precision (FP16) utilizing a fused Scaled Dot-Product Attention (SDPA) kernel [33]. Models were configured for chunked action execution [6], with varying action horizons (H) and execution prefixes (K). For both SmolVLA and Diffusion Policy, we fixed the number of sampling steps to $T=10$.

Based on our empirical measurements, we report compute-limited upper bounds for the system’s operational frequencies. Assuming policy inference represents the primary bottleneck and streaming previously planned actions incurs negligible overhead, the maximum theoretical replanning frequency is defined as:

$$f_{\text{replan}}^{\text{max}} = \frac{1}{\bar{t}_{e2e}}$$

where \bar{t}_{e2e} is the measured mean end-to-end latency.

A detailed breakdown of the latency and theoretical frequency limits for each architecture is presented in Table IV. Note that these calculations exclude downstream physical limitations, such as servo constraints and serial communication overhead, and thus represent theoretical upper bounds rather than guaranteed closed-loop performance.

TABLE IV
INFERENCE BENCHMARKS ON JETSON ORIN NANO (MAXN SUPER MODE)

Model	H	K	T	t_{e2e} (ms)	$f_{\text{replan}}^{\text{max}}$
ACT	100	50	-	36.0 ± 0.9	27.8
Diffusion	20	10	10	539.6 ± 0.1	1.8
SmolVLA	20	10	10	713.8 ± 9.3	1.4

Note: E2E latency t_{e2e} covers the loop from camera acquisition to action output. H=Horizon, K=Prefix, T=Sampling steps.

Crucially, our architectural comparison reveals that the primary computational bottleneck on edge devices stems from the iterative action experts and the requisite denoising steps they demand. The data demonstrates that visual-semantic processing in SmolVLA introduces only minor overhead compared to a standard Diffusion Policy, proving that high-parameter semantic heads are not the only prohibitive factor for high-frequency robotic control.

D. System Stress Tests

We subjected the system to high-duty cycle stress tests, simulating concurrent multi-axis movements (e.g., bimanual arm actuation and neck rotation) to evaluate power stability and thermal management, monitored via oscilloscope and onboard telemetry.

In the original shared-bus design, these loads induced a current surge and voltage collapse from 12.2V to 306mV at

¹https://huggingface.co/lerobot/act_aloha_sim_transfer_cube_human

²https://huggingface.co/lerobot/diffusion_push_t

³https://huggingface.co/lerobot/smolvla_base

$t = 19$ s, necessitating a full power-cycle for recovery. This failure mode repeated across all 5 trials. The revised Tri-Bus architecture maintained stable voltage (mean 12.0V, variance $\leq 0.1V$) under identical conditions for 1-minute trials ($n = 5$), eliminating transients and brownouts.

Thermal tests confirmed no Jetson throttling after 30 minutes of continuous operation at peak inference load while running the 450M-parameter SmolVLA model [32] (maximum GPU temperature 54.62°C), validating the passive ducting design for untethered deployment. Furthermore, the robot maintained operational stability with a battery discharge of only 5% over 30 minutes of full functional load. These results address common voltage instability in low-cost systems [19].

E. Teleoperation Dexterity




Fig. 9. VR Teleoperation Setup: (Left) The operator interacting with the system; (Right) a detailed view of the ManipulationNet peg-in-hole board.

We evaluated the effectiveness of the teleoperation interfaces for precise bimanual manipulation using a peg-in-hole insertion task, requiring the insertion of a 12mm peg into a 15mm hole (Fig. 9). The study measured success rate (defined as full insertion without dropping the workpiece), task completion time, and the number of required correction attempts. To account for user variability, we conducted the study with three participants of varying robotics experience, totaling 25 trials across three modalities. Aggregated results for intuitiveness, ergonomics, and temporal performance are reported in Table V.

TABLE V
TELEOPERATION PERFORMANCE ON HIGH-PRECISION PEG-IN-HOLE TASK ($N = 25$ TOTAL TRIALS)

Modality	Intuitiveness	Ergonomics	Time (s)
VR Controller*	2.67	2.33	48.22
Hand Tracking*	1.33	1.67	68.30
Xbox Controller Baseline	1.50	3.67	248.73

Note: *Utilized the Meta Quest 3 AR system. Intuitiveness and Ergonomics represent mean scores on a 5-point Likert scale (1: Poor, 5: Excellent). Time represents the mean duration of trials.

While both VR-based modalities reduced average task completion time by over 70% compared to the joypad baseline (248.73s), the choice of input method proved critical for

reliability. Using Meta Quest 3 (MQ3) with dedicated controllers achieved the highest performance with an 80.6% time reduction (48.22s), whereas Hand Tracking—despite a 72.5% reduction (68.30s)—suffered from significantly higher failure rates. These findings highlight a distinct trade-off: spatial hand mapping is more intuitive than coordinate-based joysticks, yet optical tracking latency and occlusion issues bottleneck high-precision bimanual tasks compared to physical controllers.

VI. DISCUSSION AND LIMITATIONS

The robot’s fully untethered operation, 1kg payload and 40cm reach have the potential to demonstrate simple home automation and manufacturing use cases (Fig. 1 and 2). Here, the platform can introduce a new baseline in terms of capability and cost.

Its low cost makes it suitable as an educational platform that can be shared among student groups. Finally, its modular architecture makes it possible to teach even larger classes in which students share a large number of arms, mobile bases, and a few complete xLeRobots, allowing to implement a curriculum such as described in [2].

In addition to teaching basics ranging from vision and inverse kinematics to machine learning, students can get first-hand experience in fabricating their own hardware. Here, learning goals would not be limited to coding, but experiencing design trade offs in electrical and mechanical design such as described here, as well as experimenting with entirely novel designs including adding tactile sensors, additional cameras, or articulated hands, e.g. Several limitations present opportunities for future work. The Jetson Orin Nano’s limited CUDA kernels restricts high-frequency multi-modal transformers while a Jetson Thor (40-130W) exceeds the power budget, prompting exploration of tiered compute architectures or additional power supplies, including using an additional Anker power supply. Additionally, hobby-grade servos, while cost-effective, exhibit mechanical sensitivity; upgrading to industrial-grade servo motors would enhance durability and payload, but significantly increase cost by a factor 3-10x. Finally, the current configuration favors table-top tasks—adding a vertical linear axis like in the AhaRobot [14] will enable ground-level manipulation.

VII. CONCLUSION

This paper presents an untethered evolution of the XLeRobot platform that combines bimanual mobile manipulation with fully onboard GPU compute at a total system cost below \$1300. By jointly addressing mechanical stiffness, power integrity, and embedded compute integration, the proposed system transforms a previously tethered research platform into a self-contained mobile manipulation system capable of teleoperation, SLAM-based navigation, and vision-driven manipulation using onboard neural network inference. The tri-bus power architecture and improved structural design significantly increase system robustness while maintaining the accessibility and low cost that make the LeRobot ecosystem attractive for research and education.

Beyond the specific hardware design, this work highlights how modern embedded GPUs enable untethered robot learning systems at price points accessible to individual labs, classrooms, and distributed data-collection efforts. The resulting platform provides a practical testbed for emerging paradigms such as vision-language-action models, imitation learning, and large-scale teleoperation data collection on mobile manipulators. By lowering both the financial and technical barrier to entry, the platform can help standardize hardware for robot learning and facilitate reproducible research across institutions.

Future work will explore tighter integration with learning pipelines in the LeRobot ecosystem, including onboard policy deployment, large-scale teleoperation data collection, and benchmarking of learning-based mobile manipulation tasks. We hope that the open design and low cost of the platform will encourage community contributions and accelerate experimentation in embodied AI and robot learning.

ACKNOWLEDGMENT

The authors would like to thank Amine Boubrit for his contributions to the inverse kinematics (IK) algorithm development, and Lyle Antieau for his valuable technical support to our platform. We also extend our gratitude to the open-source communities of LeRobot and XLeRobot for providing the foundational frameworks for this research. ARPA-E contract. Generative AI (Gemini AI) was utilized for structural refinement, grammatical proofreading, and the generation of visualization code (Tikz) throughout this manuscript.

REFERENCES

- [1] E. Vrochidou, M. Manios, G. A. Papakostas, C. N. Aitsidis, and F. Panagiotopoulos, "Open-source robotics: Investigation on existing platforms and their application in education," in *2018 26th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*. IEEE, 2018, pp. 1–6.
- [2] N. Correll, R. Wing, and D. Coleman, "A one-year introductory robotics curriculum for computer science upperclassmen," *IEEE Transactions on Education*, vol. 56, no. 1, pp. 54–60, 2013.
- [3] J. Vega and V. Pérez, "G-arm: An open-source and low-cost robotic arm integrated with ros2 for educational purposes," *Multimedia Tools and Applications*, vol. 84, no. 33, pp. 40 683–40 705, 2025.
- [4] R. A. Bindu, S. Alam, and A. A. Neloy, "A cost-efficient multipurpose service robot using raspberry pi and 6 dof robotic arm," in *Proceedings of the 2019 2nd international conference on service robotics technologies*, 2019, pp. 16–22.
- [5] M. E. Erçiyas, H. Etukuru, S. Chintala, N. M. M. Shafullah, and L. Pinto, "Cone-e: An open source bimanual mobile manipulator for generalizable robotics," in *RSS 2025 Workshop on Whole-body Control and Bimanual Manipulation: Applications in Humanoids and Beyond*.
- [6] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2023. [Online]. Available: <https://arxiv.org/abs/2304.13705>
- [7] J. Aldaco, T. Armstrong, R. Baruch, J. Bingham, S. Chan, K. Draper, D. Dwibedi, C. Finn, P. Florence, S. Goodrich *et al.*, "Aloha 2: An enhanced low-cost hardware for bimanual teleoperation," *arXiv preprint arXiv:2405.02292*, 2024.
- [8] Z. Fu, T. Z. Zhao, and C. Finn, "Mobile aloha: Learning bimanual mobile manipulation using low-cost whole-body teleoperation," in *8th Annual Conference on Robot Learning*, 2024.
- [9] R. Cadene, S. Aliberts, F. Capuano, M. Aractingi, A. Zouitine, P. Kooijmans, J. Choghari, M. Russi, C. Pascal, S. Palma *et al.*, "Lerobot: An open-source library for end-to-end robot learning," *arXiv preprint arXiv:2602.22818*, 2026.
- [10] Hugging Face, *LeKiwi Hardware Documentation*, Hugging Face, Inc., 2024. [Online]. Available: <https://tinyurl.com/2t9dad7n>
- [11] XLeRobot Contributors, "XLeRobot Documentation," <https://tinyurl.com/4ec444mu>, 2024.
- [12] G. Pratt and M. Williamson, "Series elastic actuators," in *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*, vol. 1, 1995, pp. 399–406 vol.1.
- [13] B. Calli, A. Singh, J. Bruce, A. Walsman, K. Konolige, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Yale-CMU-Berkeley dataset for robotic manipulation research," *The International Journal of Robotics Research*, vol. 36, no. 3, pp. 261–268, 2017.
- [14] H. Cui, Y. Yuan, Y. Zheng, and J. Hao, "Aharobot: A low-cost open-source bimanual mobile manipulator for embodied ai," *arXiv preprint arXiv:2503.10070*, 2025.
- [15] AgileX Robotics, "PiPER: 6-axis lightweight robotic arm," <https://tinyurl.com/48s5sy3r>, 2025.
- [16] R. Amsters and P. Slaets, "Turtlebot 3 as a robotics education platform," in *International Conference on Robotics in Education (RIE)*. Springer, 2019, pp. 170–181.
- [17] ROBOTIS, "OpenMANIPULATOR-X (RM-X52-TNM)," <https://tinyurl.com/2vfkxsmf>, 2024.
- [18] NVIDIA, "Jetson Orin Nano Super Developer Kit," <https://tinyurl.com/5d6jyae9>, 2024.
- [19] T. Mikołajczyk, D. Mikołajewski, A. Kłodowski, A. Łukaszewicz, E. Mikołajewska, T. Paczkowski, M. Macko, and M. Skornia, "Energy sources of mobile robot power systems: A systematic review and comparison of efficiency," *Applied Sciences*, vol. 13, no. 13, 2023. [Online]. Available: <https://www.mdpi.com/2076-3417/13/13/7547>
- [20] Hugging Face, "SO-101 Arm Documentation," <https://tinyurl.com/fpfd27d6>, 2024.
- [21] Shenzhen FEETECH RC Model Co., Ltd., *STS3215 12V 30kg.cm Magnetic Encoding Serial Bus Servo Specification*, 2024. [Online]. Available: <https://www.feetechr.com/525603.html>
- [22] Anker SOLIX, "Anker SOLIX C300 DC Portable Power Station," <https://tinyurl.com/2zn6py2m>, 2024.
- [23] M. A. Mazlan, M. A. Anas, N. A. N. Izmin, and A. H. Abdullah, "Effects of infill density, wall perimeter and layer height in fabricating 3D printing products," <https://doi.org/10.3390/ma16020695>, 2023, materials, vol. 16, no. 2, p. 695.
- [24] J. Carpentier, G. Saurel, G. Buondonno, J. Mirabel, F. Lamiroux, O. Stasse, and N. Mansard, "The Pinocchio C++ library – A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives," in *International Symposium on System Integration (SI)*, 2019.
- [25] S. Caron, Y. De Mont-Marin, R. Budhiraja, S. H. Bang, I. Domrachev, S. Nedelchev, P. Du, A. Escande, J. Vaillant, B. Wingo, and S. Patapati, "Pink: Python inverse kinematics based on Pinocchio," 2026. [Online]. Available: <https://tinyurl.com/mrf56sm4>
- [26] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, C.-Y. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, "Segment anything," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [Online]. Available: <https://arxiv.org/abs/2304.02643>
- [27] Meta AI, "Introducing Llama 3.1," <https://tinyurl.com/3yctubzb>, 2024.
- [28] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," in *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [29] NVIDIA, "TensorRT documentation (Overview / Quick Start)," <https://tinyurl.com/2rktc9rr>, 2025.
- [30] M. Labbé and F. Michaud, "Rtab-map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation," *Journal of field robotics*, vol. 36, no. 2, pp. 416–446, 2019.
- [31] X. Cheng, J. Li, S. Yang, G. Yang, and X. Wang, "Open-television: Teleoperation with immersive active visual feedback," in *8th Annual Conference on Robot Learning (CoRL)*, 2024.
- [32] M. Shukor, D. Aubakirova, F. Capuano, P. Kooijmans, S. Palma, A. Zouitine, M. Aractingi, C. Pascal, M. Russi, A. Marafioti, S. Alibert, M. Cord, T. Wolf, and R. Cadene, "Smolvla: A vision-language-action model for affordable and efficient robotics," 2025. [Online]. Available: <https://arxiv.org/abs/2506.01844>

- [33] T. Dao, D. Y. Fu, S. Ermon, A. Rudra, and C. Ré. “FlashAttention: Fast and memory-efficient exact attention with IO-awareness,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.