2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 14-18, 2024. Abu Dhabi, UAE

A Voxel-Enabled Robotic Assistant for Omnidirectional Conveyance

Michael Angelo Carvajal, Katiso Mabulu, Muneer Lalji, James Flanagan, Rui Luo, Samuel Hibbard, Tanav Chinthapatla, Rohan Bettadpur, Salah Bazzi, Mark Zolotas, Kristian Kloeckl, and Taşkın Padır

***Abstract*— Conventional bidirectional conveyance platforms use a flat translating belt or a series of spinning wheels or rollers to apply a shear force to payloads to move them. Wheel/roller-based conveyors in particular cannot double as a worktop when idle, do not support collision-free multi-object manipulation by default, and are not optimized to move objects that are *either* slippery *or* pliable—let alone both. This paper introduces a Voxel-Enabled Robotic Assistant (VERA), a network of intelligent table "partitions" whose topologically dynamic worktops enable omnidirectional conveyance; each partition is composed of a 2D array of "quadrants," axisymmetric modules that can be hot-swapped for maintenance or repairs; each quadrant contains a 2D array of "cells," unitary robotic submodules; each cell houses an independently controllable "voxel," the motorized rotary element that conveys an overhead object. The efficacy of a VERA prototype was determined by evaluating waypoint error as a range of payloads were maneuvered between trajectory waypoints. By conveying *both* pliable *and* rigid payloads having slippery textures, the faceted voxels outperformed those augmented to mimic the circular-profiled wheels/rollers of competitor systems. VERA also successfully performed collision-free multi-object planar manipulations planned by its pathfinding algorithm. In light of these results, VERA emerges as a promising material handling platform for use in "Future of Work" settings as the need for multi-purpose collaborative industrial robots continues to grow.**

I. Introduction

Industry 3.0 has long advocated for the introduction of fully automated systems into labor-intensive workflows that involve repetitive yet precise actions [1], as such physical activity hastens muscle fatigue in and slows down human workers [2]. Having robots tend to such mundane actions *both* promises consistent task completion times *and* affords their human coworkers the opportunity to focus on new roles with more complex duties [3]. For those reasons, industries see an increase in annual labor productivity after adding human–robot teams to their workforce [4], negating the presumed need to upgrade to a fully autonomous one. Therefore, Industry 4.0 instead encourages the development of adaptable robotic systems that can seamlessly collaborate *with*—rather than *replace*—their human coworkers [5].

This material is based upon work supported by the National Science Foundation under Award No. 1928654. *(Michael Carvajal and Katiso Mabulu are co-first authors.) (Corresponding author: Katiso Mabulu.)*

All authors are with the Institute for Experiential Robotics, Northeastern University, Boston, MA 02120 USA (e-mail: mabulu.k@northeastern.edu).

Kristian Kloeckl is an Associate Professor in the School of Architecture and in the Department of Art + Design at Northeastern University.

Taşkın Padır holds concurrent appointments as a Professor of Electrical and Computer Engineering at Northeastern University and as an Amazon Scholar. This paper describes work performed at Northeastern University and is not associated with Amazon.com, Inc. or its affiliates.

Fig. 1. (a) Voxelizing a flat workspace for omnidirectional conveyance. (b) Trackable payload placed atop the 1×2 array of VERA partitions.

A conventional robotic conveyor's design becomes more attractive to an industrial warehouse having extensive sorting and packing processes if it adapts to deliver "*enhanced* directional control" [6] capabilities like synchronous multi-object manipulation free of collisions. Yet, most are often limited to "open-loop" control, *rigid*-body payloads, and nonplanar topographies that cannot serve as worktops. This paper introduces the Voxel-Enabled Robotic Assistant (VERA), a modular omnidirectional conveyance platform that contributes the following advancements to the arts of *both* material handling *and* human–robot interaction (HRI):

1) A flat visual-servo-controlled human–robot interface that can articulate *discrete* subsections of its worktop;
2) Manipulation of *both* rigid *and* soft bodies having slippery surfaces as effected by faceted rotary voxels;
3) Concurrent three-degree-of-freedom (3-DOF) planar maneuvering of multiple payloads *sans* collisions.

II. Related Work

This section describes the architectures and limitations of robotic platforms that feature omnidirectional conveyance.

Mayer [7] developed the Flexconveyor, a 2-DOF module with a 50×50 cm footprint that can connect to other copies of itself to create a "continuous material flow" system. The largest "small load bearer" (SLB), *i.e.*, payload, that a module can transport has a 40×40 cm footprint and weighs 20 kg, meaning the Flexconveyor system sports a $125\,\mathrm{kg/m^2}$ payload-carrying capacity. In its default configuration, the square-shaped module pushes SLBs along its primary Cardinal axis of motion via a series of powered parallel

cylindrical rollers positioned at a fixed elevation. To divert an SLB sideways, *i.e.*, along the module's secondary Cardinal translation axis, a series of pulley-supported elastic bands interlaced between and passively driven by said rollers rise up to engage the SLB's underside, suspending it at a higher elevation. The Flexconveyor is programmed so that its conveyance configuration can only be toggled if the SLB it is moving is completely within the module boundaries; this is effected by photoelectric barriers that can detect ingress or egress over any of the four edges of the module to which they are mounted. Thus, a single Flexconveyor cannot move an SLB along a diagonal or curved trajectory nor can it induce a yaw rotation of the SLB on its own.

Keek *et al.* [8] proposed the E-Pattern Omniwheeled Cellular Conveyor (EOCC), a conceptual platform composed of an array of cellular modules that each house one of two "E"-shaped arrangements of "small" (∅6.18 cm) and "large" (∅17.2 cm) omniwheels within a 25×21.7 cm footprint. The EOCC is essentially the discretized form of one of Omnia's 2-DOF sortation tables [6] that use orthogonally oriented omniwheels to transport a payload in the X and/or Y direction and offer carrying capacities of up to $333\,\mathrm{kg/m^2}$. Yet, unlike [6], the EOCC *also* enables *both* yaw rotation of a payload *and* concurrent handling of *multiple* payloads because its parallel-axis omniwheels are *independently* controlled instead of being driven in unison. The concept was evaluated by simulating a 150×152 cm EOCC's conveyance of a 37.5×37.5 cm carton in a virtual environment. The simulation employed camera and image processing techniques to determine the real-time poses of the carton so that the appropriate omniwheels were activated as it moved along diagonal, 8-shaped, and ∞-shaped trajectories.

Krühn *et al.* [9] interconnected an array of addressable omnidirectional "small-scale conveyor modules" (SSCMs) to form a "module matrix" that exhibited a $125\,\mathrm{kg/m^2}$ carrying capacity. An SSCM's embedded light sensor and its lone 360°-swivel roller respectively detect the presence of and convey a payload resting atop said roller. Prototype module matrices proved to be suitable for conveying, sorting, and re-directing rigid payloads whose flat bases cover more surface area than the 7×7 cm footprint of a single SSCM.

Omnidirectional conveyance features aside, neither of these platforms demonstrate the *holistic* functionality suitable for warehouse-based HRIs, *e.g.*, a mixed-use worktop that couples pathfinding algorithms with vision-based feedback to harmoniously manipulate several in-transit bodies, rigid *or* soft, across or within its top surface without collisions.

III. System Architecture

This section describes the VERA system architecture in terms of its mechanical, electrical, and software constituents.

A. Design Requirements

An assistive robotic platform is proposed to aid laborers who handle loose items atop a work surface of typical elevation. Example scenarios include: 1) dismantling sub-assemblies into its separate components; 2) grouping scattered items; or 3) cutting large objects into smaller pieces. During such workflows, a laborer may need to, either permanently or temporarily, move different items toward and/or away from their personal station, *e.g.*, to pass components down an assembly line or to set them aside for later use. Therefore, being able to consecutively move items in various directions—sans collision with other objects on its work surface—is a key design requirement of this platform.

The proposed platform must also provide users with a flat unencumbered workspace whenever payloads are motionless. Ergo, any moving part of its topside must be able to return to a pose that aligns flush with the platform's worktop. Neither platform from [7], [8], nor [9] suffice as the circular profiles of their exposed wheels/rollers preclude flat topographies.

B. General System Specifications

To achieve omnidirectional payload motion, we introduce VERA. This system abstractly discretizes a flat work surface into an imaginary grid of "**zones**," square subsections that each promote payload conveyance along a Cartesian axis coplanar with the tabletop, *i.e.*, either X_T or Y_T, as illustrated by the bidirectional arrows in Fig. 1(a). Axis assignments alternate such that zones located diagonally from each other promote parallel motions, while adjacent zones promote motion along axes perpendicular to each other.

Centered within each zone is the coplanar facet of a cam that spins about an axis parallel to and below the work surface. The non-circular geometry of this "**voxel**" either protrudes above or aligns with the tabletop depending on its orientation. Non-slip contact between a voxel and the object above it results in relative motion between said object and the stationary table frame, $\{T\}$. While a single voxel can only push objects along the movement axis assigned to its zone, *multiple* voxels can work in harmony to perform in-plane yaw rotations, omnidirectional displacements, or a combination of these planar Euclidean motions on objects.

C. Concept Simulation

To validate the concept of VERA, a 6×6 array of cube-shaped voxels, and a rectangular prism body to be conveyed, were modeled in the MuJoCo virtual environment [10]. A simulation involves a user setting a target (X_T, Y_T)

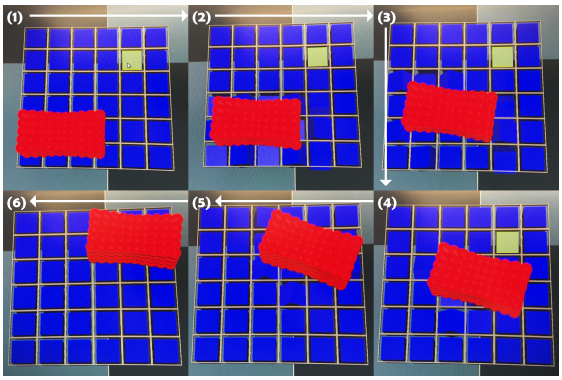

Fig. 2. Clockwise (1–6): screenshots from a MuJoCo simulation of voxels manipulating a 3D body from (1) a starting pose toward (6) a goal pose.

coordinate for the prism's center of mass by clicking on the top face of a voxel. An A* graph search algorithm [11] determines if and when a voxel rotates to contact and push the supported payload toward the selected goal position. Promising simulation results, like the sequence of manipulations depicted in Fig. 2, validated the decision to develop detailed specifications for a physical system.

D. Detailed Mechanical Specifications

In order to reduce machine downtime in the event of a hardware failure, VERA follows a nested module-based architecture. Faulty modules—from *any* nest level—can be easily removed for individual maintenance and replaced with functioning ones. The hierarchy of modules is as follows:

1) A "**cell**" subassembly houses a powered rotary voxel and is VERA's building block embodiment of a zone.
2) A "**quadrant**" is a three-layer assembly with an upper and lower lid that sandwich a 2×2 array of cells.
3) A "**partition**" is a set of four quadrants powered and controlled by the same printed circuit board (PCB).

1) Cells: An increase in VERA's manipulation resolution causes the number of cells, *i.e.*, the costs of goods, to follow suit. To strike a balance, a cell has a 6.35×6.35 cm footprint and its voxel has a 3.81×3.81 cm cam profile. While this design choice does not create a minimum footprint needed by a payload to be translated across the top of *one* cell, it *does* set 9×9 cm as the minimum footprint needed by a payload to *continuously* undergo sequential X_T *and* Y_T translations across *several* cells, as smaller objects either get trapped in the gaps between or fumble uncontrollably across voxels.

Fig. 3 depicts opposing isometric views of a cell's 3D model. The integral, cube-shaped voxel rotates concentrically about a dowel pin whose ends are supported by the cell walls. Voxel subassemblies were fabricated by press-fitting 3D-printed flanged endcaps into the open ends of square tube slices. The corners of this profile fit inside a ∅5.1 cm circumscribed circle, giving a spinning voxel 0.6 cm of radial clearance to neighboring cells. Under a no-slip condition, a driven conveyor wheel's angular displacement is linearly proportional to the tangential travel of the object it moves. Therefore, a square-profiled voxel's corner should translate an overhead object by 3.81 cm per 90° of rolling contact.

Powering a cell's integrated bidirectional gearmotor spins the rotary voxel. Each 7.5-cm-long 24VDC-rated planetary gear brushed motor is oriented vertically to fit inside its cell's footprint. The magnetic contact pads of the polarized connector protruding below the cell floor are at a higher elevation than the bottom edges of the cell walls, allowing the submodule to sit upright on a flat surface when unplugged.

Four neodymium magnets are embedded into a voxel's non-geared lateral side and are spaced axisymmetrically about its axle. Their south poles face the nearby cell wall which supports a Hall Effect sensor to detect their presence.

2) Quadrants: Since wires do not permanently connect to its cells, a quadrant is a self-contained module. As shown in Fig. 4(b), the centroids of the four cells within a quadrant are spaced 6.35 cm apart in the X_T and Y_T directions; therefore, the quadrants have a 12.7×12.7 cm footprint since quadrants are meant to sit flush and adjacent to each other while maintaining a constant spacing between neighboring voxels. The rotational orientation of each cell is also axisymmetric about the quadrant's central Z_T-parallel axis. This design choice gives quadrants an *additional* level of modularity: their appearance and functionality are agnostic to their planar orientation relative to the table; in other words, a quadrant can be unplugged, lifted off the table, and spun about its central axis by 90°, 180°, or 270° (see Fig. 5) before being lowered back into place to be reconnected without issue.

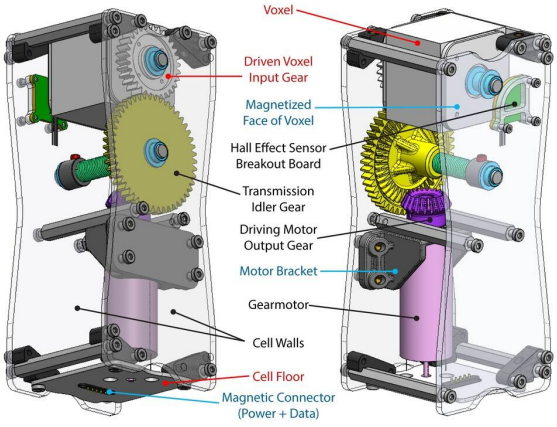

Fig. 3. 3D CAD model of the prototype VERA cell design.

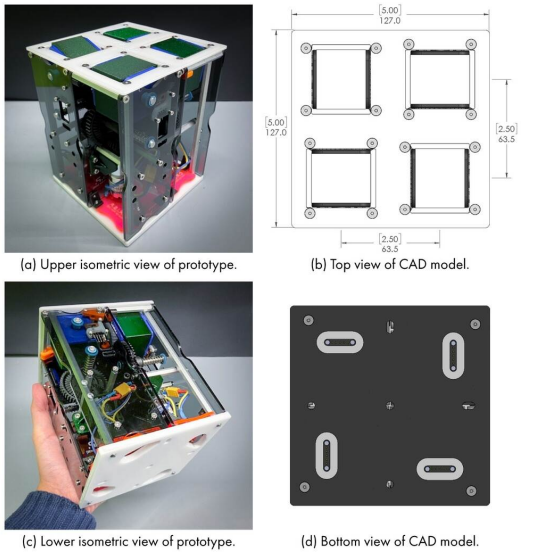

Fig. 4. Isometric and orthogonal views of a VERA quadrant.

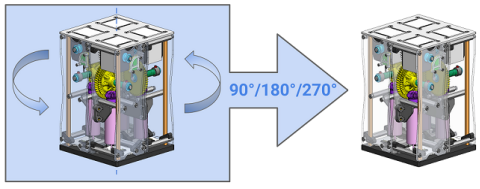

Fig. 5. Quadrants are axisymmetric hot-swappable modules.

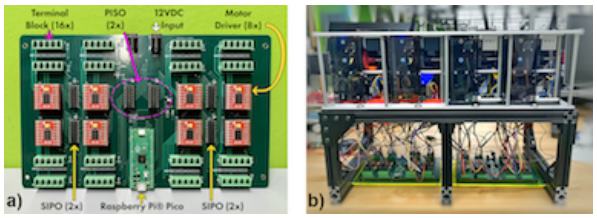

Fig. 6. (a) Partition-controlling PCB. (b) Functional 1×2 prototype.

3) Partitions: For this build, four quadrants were placed atop a square-shaped support plate that maintains them in a 2×2 array. Beneath this plate is a single PCB sporting 16 cable terminal blocks arranged in a 4×4 array [see Fig. 6(a)]. The power and data cables stemming from these blocks pass through openings in the plate to access their respectively assigned cells. The magnetic connector at the end of each cable automatically orients itself to align its contact pins with the mating pins of its assigned cell's connector. Excluding the cost of labor and secondary machining operations to modify raw materials, a single partition cost less than 1500 USD to fabricate. The VERA prototype built and tested for this research consists of a 1×2 partition array [see Fig. 6(b)].

E. Detailed Electrical Specifications

When two bits of information, *IN1* and *IN2*, are sent to the motor driver regulating it, a bidirectional motor enters one of four possible modes: 1) shorted brake; 2) counterclockwise (CCW) spin; 3) clockwise (CW) spin; or 4) zero power. A partition in turn requires 32 bits to control its 16 cells in concert, hence its PCB being designed around the Raspberry Pi Pico, a low-cost high-performance 32-bit microcontroller. A daisy chain of four 8-bit Serial-In Parallel-Out (SIPO) shift registers digests the Pico's outgoing 32-bit command. Two dual-channel motor drivers divvy up a SIPO's 8-bit instruction, ergo each channel reads two bits. Pulse width modulation (PWM) is then used to augment channel outputs to ultimately control the rotational *velocity* of each voxel.

The rotational state of each voxel is constantly fed back to the PCB through the planar Hall sensor mounted inside each cell. When a voxel's facet is flush with the table surface, its corresponding embedded magnet is concentric with the Hall sensor, which detects the presence of that magnetic field, saturates, and outputs a "LOW" voltage; else, when a voxel is atilt relative to the worktop, the unsaturated sensor outputs its default "HIGH" voltage signal. These 1-bit outputs are transmitted in parallel to a daisy chain of two 8-bit Parallel-In Serial-Out (PISO) shift registers, which then concatenate and relay the 16-bit signal to the Pico's serial input pin.

F. Detailed Software Specifications

VERA's software suite is built atop ROS 2 Humble [12], establishing instantaneous communication between its main components. The system's software architecture consists of: 1) the RViz GUI; 2) a local path planner per payload; 3) a global controller for the table; 4) the micro-ROS node that operates on the two Picos; and 5) the Virtual Reality Peripheral Network (VRPN) [13] streaming engine

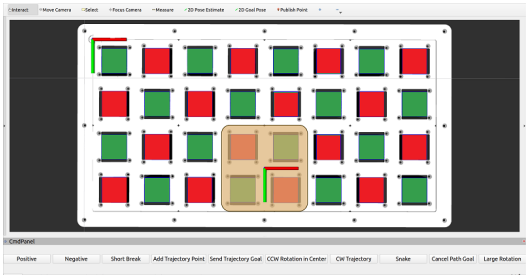

Fig. 7. Interactive 3D model of the 1×2 VERA represented in RViz.

implemented in Motive for the overhead multi-camera OptiTrack [14] V120:Trio looking down at the tabletop.

$$ {}^{T}_{P_n}\mathbf{A} = {}^{T}_{C}\mathbf{A} \cdot {}^{C}_{P_n}\mathbf{A} = ({}^{C}_{T}\mathbf{A})^{-1} \cdot {}^{C}_{P_n}\mathbf{A}. \tag{1} $$

The coordinate frames attached to the OptiTrack, the table, and its n^{th} payload are $\{C\}$, $\{T\}$, and $\{P_n\}$, respectively. The optical tracking system estimates ${}^{C}_{T}\mathbf{A}$ and ${}^{C}_{P_n}\mathbf{A}$, the corresponding camera-based poses of the table and the n^{th} payload, respectively. Said poses are relayed to that payload's local planner, which then calculates ${}^{T}_{P_n}\mathbf{A}$, the n^{th} payload's pose relative to $\{T\}$, using (1). The system then references the stored geometries of the table and said payload to identity the voxels currently covered by said payload's footprint.

The RViz GUI displayed in Fig. 7 allows a user to communicate a target composite trajectory for a selected payload to follow. By clicking on each voxel to add its centerpoint as a stored trajectory waypoint, the list of saved waypoints can then be sent to the payload's local planner as a trajectory goal. The planner iterates over the received list and uses the respective payload's current $\{T\}$-based pose to determine a path through each waypoint in the list.

Each local planner runs on a separate thread to allow concurrent operation such that the global occupancy grid can be reasoned over to ensure no collisions occur between payloads. The A* graph search algorithm is employed on every payload's planner, which runs at 1 Hz to adjust plans as objects simultaneously traverse the table. Each local planner's generated commands are sent to a global controller that determines which voxels to activate and in what direction to spin them. The global controller then combines said local commands into a single global command that is sent to the chain of microcontrollers to be distributed to the motor drivers that control the relevant voxels in the system.

IV. Experiments

This section describes the field testing conducted on the 1×2 VERA prototype to assess the platform's capabilities.

A. Single Soft-Body Payload

The platforms described in [7], [8], and [9] require the no-slip contact of rolling friction to apply a *shear* force to the undersides of payloads for conveyance. VERA's faceted voxels, in contrast, are meant to apply a *normal* contact force to both the bottom *and lateral* surfaces of a payload, theoretically enabling the platform to maneuver both rigid *and* pliable payloads. If all four platforms were tasked with

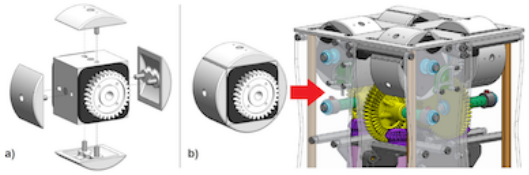

Fig. 8. (a) Attaching profile modifiers; (b) Wheel-shaped voxels.

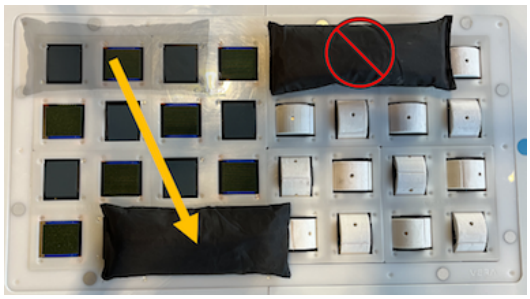

Fig. 9. Sandbag conveyance by *(left)* faceted versus *(right)* round voxels.

conveying identical pliable slippery payloads, it would be difficult to compare their performances given the unique spacing and effective rolling diameters of their respective rotary drivers. Thus, to create a *reasonable* proxy for wheel/roller-based systems, 3D-printed covers were attached to the facets of one VERA partition's cube-shaped voxels, as depicted in Fig. 8, to convert them into wheels *without* affecting their spacing or effective rolling diameter.

To experimentally assess and compare them, the faceted and round voxels were tasked with transporting a pliable payload—a sandbag with a 6.4×19.1 cm footprint equal to a 1×3 cell array—diagonally across their partitions. Fig. 9 shows how the partition with faceted voxels translated the smooth 3.81-cm-thick sandbag as desired while the wheel-shaped voxels, despite being made from the same plastic filament, slipped past and were *unable to move* the sandbag.

B. Single Rigid-Body Payload

The protocol for testing if VERA can plan a path for and perform 2-DOF translations on a rigid body begins by centering said payload frame's origin, O_P, on the quadrant closest to the table frame's origin, O_T, as Fig. 10 depicts. The system is then allowed five attempts to convey $\{P\}$ toward the voxel nearest the diagonal corner of the 1×2 partition array. Trial runs end when the OptiTrack detects O_P to be within a 1.25-cm-wide radius of the target waypoint, *i.e.*, when it is overlapping the target voxel. The perceived *waypoint error,* ε, between the expected and actual resting positions of $\{P\}$ is averaged and serves as the dependent variable to characterize the prototype's performance.

Various weighted containers, capped with lids laden with retroreflective markers for optical tracking, served as rigid body specimens. As outlined in Table I, the independent variables here were a container's: 1) *footprint*, $(L_X \times L_Y)$; and 2) *total load*, M. Prior empirical testing revealed that a single voxel can convey a 3.6 kg object, implying that the VERA platform has a carrying capacity of $450\,\mathrm{kg/m^2}$ *per* translation

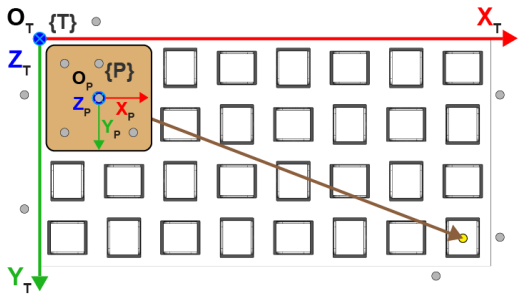

Fig. 10. Start and target poses for conveying a single rigid body diagonally.

TABLE I

RIGID-BODY CONVEYANCE: SINGLE-PAYLOAD ERROR

Body	L_X [cm]	L_Y [cm]	M [kg]	ε [cm]
$\{P_{SMALL}\}$	10.5	10.5	0.68	0.97 ± 0.76
			1.13	1.32 ± 0.69
			2.27	0.81 ± 0.23
$\{P_{MEDIUM}\}$	25.6	16.5	2.27	0.69 ± 0.03
			4.54	0.97 ± 0.53
			9.07	2.84 ± 1.57
$\{P_{LARGE}\}$	35.6	25.1	3.40	0.46 ± 0.08
			6.80	0.76 ± 0.13
			13.6	1.24 ± 0.22

axis, *i.e.*, $900\,\mathrm{kg/m^2}$ if powering both axes. The footprints of $\{P_{SMALL}\}$, $\{P_{MEDIUM}\}$, and $\{P_{LARGE}\}$ approximate that of one, two, and six quadrants, respectively; thus, the voxels covered by each container should support loads up to 14.4 kg, 28.8 kg, and 86.4 kg, respectively. Typically, the carried test load continually doubled for each container, with one's maximum load being reused as the lightest load for the next-size-up container. However, the interior volume of $\{P_{LARGE}\}$ could not fit more than 13.6 kg of the dead weights used in these tests, thereby forcing its lighter loads to be derived by continually halving said maximum. As a result, $\{P_{LARGE}\}$'s lightest load *is less than* $\{P_{MEDIUM}\}$'s heaviest.

Smaller footprints overlap fewer voxels and are thereby intermittently supported, *i.e.*, *stabilized*, by fewer contact points during each conveyance step. Thus, it was intuited that, for a given load, an object's footprint is inversely proportional to ε. Extrapolating $\{P_{LARGE}\}$'s data to infer its ε from carrying the same 2.27 kg as $\{P_{MEDIUM}\}$ and $\{P_{SMALL}\}$ yields an error of 0.39 cm, appearing to validate said theory.

C. Multiple Rigid-Body Payloads

The protocol for testing if VERA can simultaneously plan paths for and maneuver multiple payloads—*sans* collision—begins by placing two identical quadrant-sized square plates, $\{P_1\}$ and $\{P_2\}$, on the table, against adjacent corners, as depicted in Fig. 11. The system is allowed five attempts to convey the pair in concert to waypoints on opposite corners, a task that normally induces intersecting motion paths.

Fig. 12 plots the OptiTrack's record of $\{P_1\}$'s and $\{P_2\}$'s *non-intersecting* trajectories as faded green and blue traces, respectively, while averaged trajectories are colored solid. A payload's path variations are the result of the underlying reasoning of its local planner's A* search algorithm.

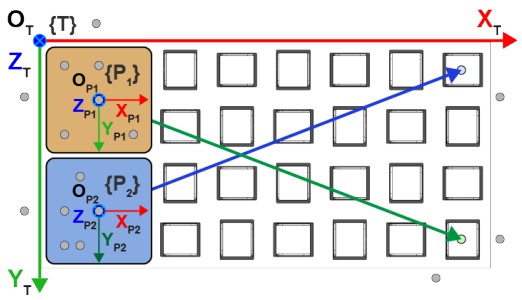

Fig. 11. Start and target poses for conveying two rigid bodies diagonally.

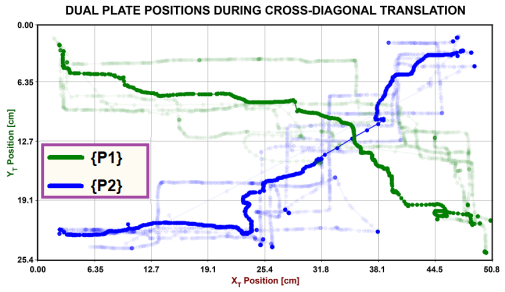

Fig. 12. Trajectories of two payloads moving diagonally simultaneously.

TABLE II

RIGID-BODY CONVEYANCE: MULTI-PAYLOAD ERROR

Body	L_X [cm]	L_Y [cm]	M [kg]	ε [cm]
$\{P_1\}$	12.4	12.4	0.1	2.83 ± 0.51
$\{P_2\}$	12.4	12.4	0.1	1.31 ± 0.46

Table II summarizes each payload's parameters and averaged waypoint error. Despite its similar footprint, the relatively higher loads carried by $\{P_{SMALL}\}$ could be why the plates saw more error. $\{P_1\}$ also saw double the error of $\{P_2\}$, suggesting that VERA's path-planning algorithm's choices can affect its multi-object conveyance accuracy.

V. CONCLUSION

This paper presented a functional prototype of VERA, a novel robotic omnidirectional conveyance platform, and demonstrated its ability to maneuver multiple payloads across its partitioned worktop using a 2D array of articulating rotary voxels under visual servo control. A collision-free route was generated for each of these payloads by its respective path planner, which reasoned over a global occupancy grid. Initial testing also suggests that VERA has a $450\,\mathrm{kg/m^2}$ payload capacity *per* translation axis, which *bests* the ratings of the benchmarked platforms described in [7], [8], *and* [9]. Surrogate testing of these and other wheel/roller-based conveyance systems was facilitated by converting the voxels in one VERA partition to have circular profiles. While that *retrofitted* form of the modular robotic assistant failed to convey a soft slippery sandbag, its *unaltered* form successfully controlled *all three* planar motion DOFs of *both* rigid *and* soft-body slippery payloads.

VERA's contributions could impact a wider audience if its system architecture were further enhanced. For instance, requiring tracked objects to sport retroreflective markers is cost prohibitive of the OptiTrack system, especially when extended to large-scale use in industrial warehouse settings. Akin to the method in [8], a more practical embodiment of vision-based feedback would use modern computer vision algorithms to analyze visuals from an affordable overhead RGB camera to detect and track *both* objects *and* humans.

With the power to locally manipulate the topography of its worktop to promote omnidirectional multi-object conveyance on command, VERA opens up research avenues for future HRI applications. One prominent direction to pursue is adding camera-based gesture controls to enable VERA's tableside human coworkers to directly—and intuitively—interact with it. In industrial settings, this adaptation could increase ergonomic efficiency and safety by enhancing production pipelines and reducing strain on human posture when handling objects, potentially complementing related forms of handover between humans and robotic arms [15].

REFERENCES

[1] K. H. Tantawi, A. Sokolov, and O. Tantawi, "Advances in industrial robotics: From industry 3.0 automation to industry 4.0 collaboration," in *2019 4th Technology Innovation Management and Engineering Science International Conference (TIMES-iCON)*, 2019, pp. 1–4.

[2] L. Ma, D. Chablat, F. Bennis, and W. Zhang, "A new simple dynamic muscle fatigue model and its validation," *International Journal of Industrial Ergonomics*, vol. 39, no. 1, pp. 211–220, 2009.

[3] P. Fleming, "Robots and Organization Studies: Why Robots Might Not Want to Steal Your Job," *Organization Studies*, vol. 40, no. 1, pp. 23–38, 2019.

[4] G. Graetz and G. Michaels, "Robots at Work," *The Review of Economics and Statistics*, vol. 100, no. 5, pp. 753–768, Dec. 2018.

[5] M. Fechter, P. Foith-Förster, M. S. Pfeiffer, and T. Bauernhansl, "Axiomatic design approach for human-robot collaboration in flexibly linked assembly layouts," *Procedia CIRP*, vol. 50, pp. 629–634, 2016, 26th CIRP Design Conference.

[6] Industry Update, "Omnia brings intuitive mobility to global markets," *Industry Update Manufacturing Magazine*, pp. 48–49, Apr./May 2024.

[7] S. H. Mayer, "Development of a completely decentralized control system for modular continuous conveyors," Ph.D. dissertation, Karlsruher Institut für Technologie (KIT), 2009.

[8] J. S. Keek, S. L. Loh, and S.-H. Chong, "Design and Control System Setup of an E-Pattern Omniwheeled Cellular Conveyor," *Machines*, vol. 9, no. 43, 2021.

[9] T. Krühn, S. Falkenberg, and L. Overmeyer, "Decentralized control for small-scaled conveyor modules with cellular automata," in *2010 IEEE International Conference on Automation and Logistics*, 2010, pp. 237–242.

[10] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 5026–5033.

[11] D. Foead, A. Ghifari, M. B. Kusuma, N. Hanafiah, and E. Gunawan, "A systematic literature review of A* pathfinding," *Procedia Computer Science*, vol. 179, pp. 507–514, 2021, 5th International Conference on Computer Science and Computational Intelligence 2020.

[12] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall, "Robot Operating System 2: Design, architecture, and uses in the wild," *Science Robotics*, vol. 7, no. 66, May 2022.

[13] R. M. Taylor, T. C. Hudson, A. Seeger, H. Weber, J. Juliano, and A. T. Helser, "VRPN: A Device-Independent, Network-Transparent VR Peripheral System," in *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST)*, 2001, pp. 55–61.

[14] G. Nagymáté and R. M. Kiss, "Application of optitrack motion capture systems in human movement analysis: A systematic literature review," *Recent Innovations in Mechatronics*, vol. 5, no. 1, pp. 1–9, 2018.

[15] M. Zolotas, R. Luo, S. Bazzi, D. Saha, K. Mabulu, K. Kloeckl, and T. Padır, "Productive inconvenience: Facilitating posture variability by stimulating robot-to-human handovers," in *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE Press, 2022, pp. 122—-128.