ICRA2026 Late Breaking Results Poster presented at
2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Wrench-Feasible Whole-Body Planning via Time-Layered DAG Optimization for Omnidirectional Aerial Manipulation

Daum Park[1], Bohyeong Pak[1], and Sanghyun Kim[1,2,†]

***Abstract*— Omnidirectional aerial manipulators (OAMs) must coordinate a floating base and onboard arm to track end-effector trajectories under coupled geometric and dynamic constraints. In cluttered long-horizon tasks, collision-free motions may still be dynamically inadmissible and dense time-layered planning graphs can become disconnected. We present a GPU-accelerated whole-body planning framework that combines reverse-chain node sampling, collision and wrench-aware feasibility filtering, guide-based connectivity refinement, and dynamic programming on a time-layered directed acyclic graph. At each timestep, the target end-effector pose is converted to a wrist-anchored representation, from which large batches of whole-body candidates are generated in parallel on the GPU. Sampled nodes are pruned using self/environment collision checks and rotor-allocation-based wrench-feasibility tests. When local disconnections remain, sparse guides trigger targeted dense resampling to recover connectivity. On drawing and peg-in-hole tasks, the proposed method achieves 91.7% and 100.0% dynamic-safe runs, versus 0.0% and 8.3% for a wrench-unconstrained baseline. Guide refinement reduces mean broken layers by 200.2 and 49.6, yielding 100 percentage-point success gains over no refinement. GPU batching keeps planning practical for long-horizon cluttered OAM problems.**

I. INTRODUCTION

Omnidirectional aerial manipulation requires coordinated motion of a fully actuated aerial base and an onboard manipulator under collision, velocity, and actuation constraints [1]. For prescribed end-effector trajectories, this becomes a high-dimensional whole-body planning problem with two practical failure modes: dynamically unrealizable candidates and disconnected long-horizon time-layered graphs. Optimization-based whole-body planning has been explored for omnidirectional aerial manipulators [3], but nonlinear formulations may still be sensitive to initialization in cluttered scenes.

We instead use a discrete whole-body planner that builds a time-layered DAG of sampled states, applies node-wise collision and wrench-feasibility filtering, and restores missing connectivity through guide-based refinement before final dynamic programming. To our knowledge, this is the first discrete OAM planner to combine wrench-aware node filtering with guide refinement in a time-layered DAG for prescribed end-effector tracking; Fig. 1 summarizes the pipeline.

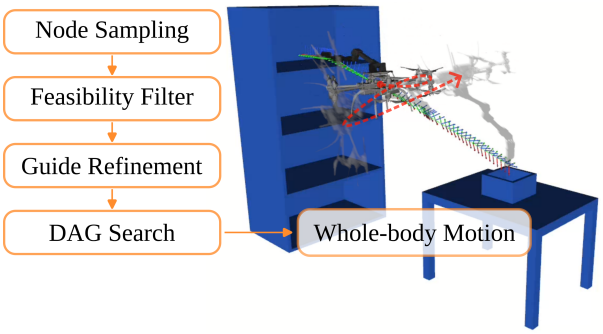

Fig. 1. Proposed GPU-accelerated DAG planning pipeline for OAM motion generation.

II. METHOD

A. Node Sampling

The planner uses a reverse-chain representation rather than directly sampling full base-and-arm poses. For each timestep, the target end-effector pose is converted to an equivalent wrist-anchored target through the fixed wrist-to-tool transform,

$$^{W}T_{w,k} = {}^{W}T_{ee,k}\,{}^{ee}T_{w}, \tag{1}$$

where ${}^{ee}T_{w}$ is constant for the robot/tool pair. This lets the planner fix the wrist target, sample candidate base-reaching configurations around it, and recover a full whole-body state. Direct sampling of base position, base orientation, and arm joints would be prohibitively inefficient for long-horizon OAM planning.

Candidate generation is implemented with a FABRIK-based geometric IK solver. FABRIK iteratively updates joint positions through alternating forward and backward passes while preserving link lengths, as illustrated in Fig. 2. Starting from the initial chain and target, the forward stage moves the end effector to the target and updates intermediate joints along link-length-preserving lines, after which the backward stage re-anchors the root and refines the final configuration. FABRIK is adopted because it provides lightweight batched inverse kinematics for large candidate sets without repeated Jacobian solves. In the planner, base targets are sampled in the wrist frame and FABRIK recovers corresponding manipulator configurations before the resulting candidates are inserted into the current DAG layer.

*This research was supported by the Ministry of Trade, Industry & Energy (MOTIE, Korea) through the Industrial Strategic Technology Development Program (No. RS-2024-00416441) and the KEIT/MOTIE Program (No. RS-2025-11082970), and by the Korea Basic Science Institute (National Research Facilities and Equipment Center) grant funded by the Ministry of Science and ICT (No. RS-2025-00564593).

[1]Daum Park, Bohyeong Pak, and Sanghyun Kim are with the School of Mechanical Engineering, Kyung Hee University, Yongin, Republic of Korea doumpork@khu.ac.kr, bhe1004@khu.ac.kr, kim87@khu.ac.kr

[2]Sanghyun Kim is also with the Digital Human Research Center, Advanced Institute of Convergence Technology (AICT), Suwon, Republic of Korea

[†]Corresponding Author: Sanghyun Kim.

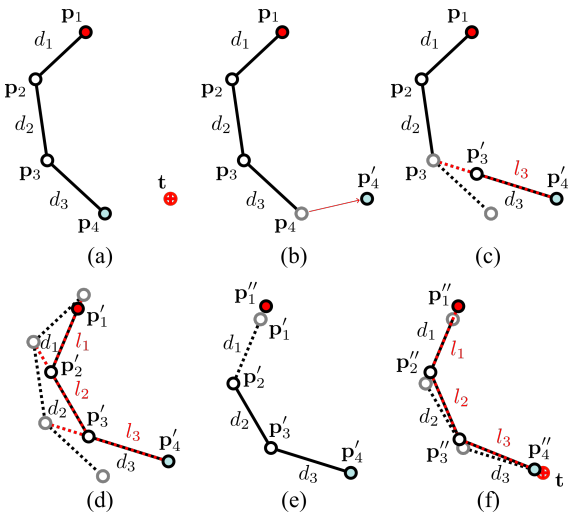

Fig. 2. Geometric workflow of FABRIK for a 4-DoF serial chain. (a) Initial joint positions, link lengths, and target. (b)–(d) Forward reaching updates the chain toward the target while preserving link lengths. (e)–(f) Backward reaching re-anchors the root and refines the final configuration [4].

B. Feasibility Filter and DAG Optimization

Collision clearance is evaluated from environment SDF values and self-collision sphere-pair distances:

$$d_{\text{env}}(\mathbf{s}) = \min_{m} \left(\phi(\mathbf{p}_m) - r_m\right), \\ d_{\text{self}}(\mathbf{s}) = \min_{(m,n)\in\mathcal{P}} \left(\|\mathbf{p}_m - \mathbf{p}_n\| - r_m - r_n\right). \tag{2}$$

Here, $\phi(\mathbf{p}_m)$ is the environment signed distance at sphere center $\mathbf{p}_m$, r_m is the collision-sphere radius, and $\mathcal{P}$ is the set of self-collision sphere pairs. Dynamic admissibility is enforced through the rotor allocation constraints

$$A\mathbf{u} = \mathbf{w}_{\text{req}}, \qquad \mathbf{u}_{\min} + \delta \le \mathbf{u} \le \mathbf{u}_{\max} - \delta, \tag{3}$$

where A is the allocation matrix and $\mathbf{u}$ is the rotor-thrust vector. The margin term δ enforces thrust reserve rather than checking feasibility exactly at the actuator limits. Here $\mathbf{w}_{\text{req}}$ stacks the required net force and torque induced by the planned base motion, gravity, and task wrench. Surviving nodes are connected by feasible one-step transitions and optimized by dynamic programming on the layered DAG:

$$J_t(j) = \min_{i\in\mathcal{N}_{t-1}} \left[J_{t-1}(i) + c_{\text{edge}}(i,j) + c_{\text{node}}(j)\right]. \tag{4}$$

In (4), $\mathcal{N}_{t-1}$ denotes feasible predecessor nodes, while c_{edge} and c_{node} penalize transition and configuration cost.

C. Guide Refinement

The time-layered DAG gives the method a useful failure mode: if no feasible transition exists between adjacent layers, the graph becomes explicitly disconnected instead of silently failing inside a nonlinear optimizer. When dense-layer connectivity remains broken, a sparse guide identifies the difficult interval and targeted dense samples are added only around that guide before the final DP pass. This concentrates computation on structurally difficult segments rather than uniformly increasing sampling over the full horizon, and is also consistent with recent refinement-based end-effector tracking methods [2].

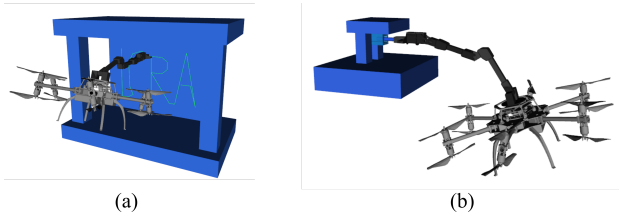

Fig. 3. Representative whole-body motion planning results. (a) Cluttered long-horizon drawing. (b) Peg-in-hole with added obstacles.

TABLE I

DYNAMIC SAFETY AND GUIDE-REFINEMENT EFFECTS ON REPRESENTATIVE OAM TASKS.

Task	Dynamic-safe runs (%)		Guide refinement effect	
	Proposed	Wrench-unconstrained	Success gain (pp)	Broken-layer reduction (layers)
Drawing	91.7	0.0	+100.0	200.2
Peg-in-hole	100.0	8.3	+100.0	49.6

III. RESULTS AND DISCUSSION

Fig. 3 shows drawing and peg-in-hole results, and Table I summarizes the quantitative outcomes. The proposed wrench-aware planner achieves 91.7% dynamic-safe runs on drawing and 100.0% on peg-in-hole, while a wrench-unconstrained baseline achieves 0.0% and 8.3%, respectively. In drawing, the wrench-unconstrained baseline often selects kinematically valid states that exceed rotor lateral-force limits, explaining its 0.0% dynamic-safe rate. Relative to no refinement, guide refinement yields 100 percentage-point success gains on both tasks and reduces the mean number of broken layers by 200.2 and 49.6. In a drawing-task node-budget ablation, success reaches 100% from 512 nodes/layer, with best runtime near 4–5 s at 512–1024 nodes/layer. These results show that wrench-aware filtering improves dynamic admissibility and guide refinement resolves dense-graph disconnections.

IV. CONCLUSION

We presented a GPU-accelerated whole-body planner for omnidirectional aerial manipulators that combines reverse-chain node sampling, collision and wrench-aware feasibility filtering, guide-based connectivity refinement, and DAG optimization. The method improves dynamic safety and restores graph connectivity in long-horizon cluttered tasks. A natural next step is extension to richer tasks and higher-DOF onboard arms.

REFERENCES

[1] M. Allenspach *et al.*, "Design and optimal control of a tiltrotor micro-aerial vehicle for efficient omnidirectional flight," *Int. J. Robot. Res.*, vol. 39, no. 10–11, pp. 1305–1325, 2020.

[2] Y. Wang and M. Gleicher, "Anytime planning for end-effector trajectory tracking," *IEEE Robot. Autom. Lett.*, vol. 10, no. 4, pp. 3246–3253, 2025.

[3] D. Lee, B. Kim, and H. J. Kim, "Autonomous aerial manipulation at arbitrary pose in $SE(3)$ with robust control and whole-body planning," *Int. J. Robot. Res.*, early access, 2025, doi: 10.1177/02783649251378195.

[4] A. Aristidou and J. Lasenby, "FABRIK: A fast, iterative solver for the inverse kinematics problem," *Graph. Models*, vol. 73, no. 5, pp. 243–260, 2011.