Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

COMPASS: centre-of-mass preview and actuator-synchronised speed for an electrified mobile manipulator on rough terrain

Mehdi Heydari Shahna

mehdi.heydarishahna@tuni.fi

Tampere University

Jouni Mattila

Tampere University

Research Article

Keywords: mobile manipulator, tip-over stability, construction robotics, electro-mechanical linear actuator, actuator rate limits, control barrier function

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11800831/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

COMPASS: centre-of-mass preview and actuator-synchronised speed for an electrified mobile manipulator on rough terrain

Mehdi Heydari Shahna*, Jouni Mattila

Faculty of Engineering and Natural Sciences, Tampere University, 33720 Tampere, Finland.

Abstract

Electrified crane manipulators can counterbalance construction carriers on rough terrain only if slow actuators reach balancing postures in time; stability-constrained planners usually assume perfect tracking. COMPASS combines terrain-preview counterweight planning, capacity-limited speed and route escalation, and execution that slows the vehicle when the manipulator lags, as established for hydraulic mobile manipulators. In simulations of a loaded 6.1 t skid-steer machine, an independent contact-wrench audit showed the acceleration-based reserve proxy deviating 8.8° RMS from a moment-inclusive diagnostic. Relative to the transport pose, the counterweight raised the audited restoring moment in all eight crest seeds across contact, tyre and time-step settings, also at equal speed; speed-compatible execution prevented a crest failure. Under drive shortfalls, the coupling reduced the tracking error by up to 0.22 s and outperformed uniform slowing at equal driving time; with documented drive ratings it cost under 1 % driving time at 0.8 and 0.35 $\mathrm{m\,s^{-1}}$.

Keywords: mobile manipulator, tip-over stability, construction robotics, electro-mechanical linear actuator, actuator rate limits, control barrier function

1. Introduction

Construction machines carry material across embankments, spoil heaps, ditches and loose fill, and compact electric carriers are increasingly combined with heavy manipulators [1, 2, 3]. A crane and its payload can shift the whole-machine centre of mass substantially relative to a short wheelbase, so on slopes and crests their posture and motion affect overturning as well as manipulation [4, 5].

Using the manipulator as a counterweight has a long history [6, 7, 8, 9]. Song and Sharf [10] combine hierarchical rough-terrain planning, stability-constrained trajectory optimisation and heavy-machine reconfiguration, and evaluate them with a known terrain map and perfect trajectory tracking. On electrified construction machines, however, the reconfiguring drives are electro-mechanical linear actuators (EMLAs) and servo slew drives that move at centimetres per second, and their achievable rates depend on load, temperature and supply voltage. A posture planned for a given vehicle speed is of little use if the vehicle reaches the crest before the drives have moved. This paper therefore asks how a counterweight plan should be executed while the carrier moves, when the drives may not keep up.

*Corresponding author: Mehdi Heydari Shahna
E-mail address: mehdi.heydarishahna@tuni.fi

Three coupled requirements follow. First, the available reserve depends on support attitude, payload, clearance, joint ranges and actuator effort limits, so the carrier must slow, change route or stop when reconfiguration alone is insufficient. Second, slow drives require terrain preview: the counterweight must start moving before the vehicle reaches the terrain that needs it. Third, execution must keep vehicle progress and counterweight motion synchronised. For the third requirement, coupling platform progress to manipulator tracking error is established for flow-limited hydraulic mobile manipulators [11]. This paper examines whether that coupling still helps when the coupled motion is a stability counterweight, and whether it is a better use of driving time than slowing the vehicle uniformly or stopping to reconfigure.

The case study is a battery-electric skid-steer carrier with interconnected oscillation beams and an electrified loader crane, empty and with a 400 kg payload. The proposed framework, COMPASS (centre-of-mass preview and actuator-synchronised speed), uses the crane and beam levelling as an active counterweight and adapts speed or route only when the estimated balancing capacity is insufficient. A path-indexed preview plan respects the actuator rates and includes a vehicle-speed profile. Execution reads the posture reference at measured vehicle progress, reduces the commanded vehicle speed when the manipulator lags (the tracking-error time scaling of [11]), and filters the commands with a control barrier function (CBF).

The contributions are:

1. An integrated preview-counterweight architecture for an electro-mechanically actuated construction carrier, in which the crane and beam levelling act as an active counterweight and speed and route are escalated only when the estimated balancing capacity is insufficient (Sections 3 and 4). Relative to the transport pose, it raises the audited restoring moment on crests in every tested seed across contact conditioning, tyre stiffness and integration step, also when a controller with the same planner drives at the nominal speed of the transport-pose baseline, and its speed-compatible execution prevents the crest failure of that equal-speed execution (Sections 6.3, 6.4 and 6.5).
2. An evaluation of tracking-error time scaling [11] when the coupled motion is a stability counterweight rather than a task trajectory. The coupling is compared at equal driving time with uniform slowing and with a stop-and-go limit, its mechanical effect is assessed with an independent audit, and its operating regime is located with drive envelopes derived from the documented motor ratings, with the travel speeds of field experiments on the same machine and with a blocked slew drive (Sections 6.6, 6.7 and 6.9). The coupling lowers the tracking error by up to 0.22 s when the drives deliver below the demanded rate, tracks better than uniform slowing at equal driving time, and costs less than 1 % of driving time with the documented drives of the machine.
3. An independent contact-wrench and centroidal-momentum audit of the simulation. It quantifies how far the acceleration-based reserve, a reduced form of the force-angle measure [12, 13], departs from the moment-inclusive wrench angle at the actual contacts, and it provides proxy-independent evidence for the counterweight effect (Sections 3.4 and 6.3).

The evaluation uses whole-body simulations of the case-study machine, with geometry, masses and drive ratings from its documentation, and relates the model and the tested travel speeds to field experiments with the same machine. The code implementing the proposed framework and a supplementary video accompany the paper.

2. Related work

2.1. Construction machines and stability assessment

Construction automation requires machines that operate in changing, unstructured environments [1, 2]. Dadhich et al. [3] identify the complexity of machine–material interaction in automated earth moving, and Eraliev et al. [14] organise autonomous-excavator research into sensing, perception, decision, planning and action; both motivate treating mobility, terrain perception and manipulation as one problem. For crane operations, Fang et al. [15] reconstruct crane poses and site geometry to warn operators of collision hazards, Kwiatoń et al. [16] assess rotary-crane stability numerically, and the HEAP walking excavator uses separate controllers for driving, chassis balancing and arm motion [17].

The force-angle measure of Papadopoulos and Rey evaluates stability about candidate tip axes from the resultant force and the equivalent contribution of moments [12, 13]. A force computed from gravity and whole-body centre-of-mass acceleration alone is therefore a reduced form of the dynamic measure, and the difference matters when articulation and terrain excitation change the angular momentum. Wheel-load models give a complementary view: Petrović and Mattila [18] derive the supporting forces of a wheeled mobile manipulator on uneven terrain, and Paz et al. [19] model the normal forces of skid-steered heavy manipulators with active suspension. Hydraulic-machine and crane studies show the influence of base compliance, actuator dynamics and payload motion on stability [4, 5]. For non-coplanar contacts, the gravito-inertial wrench cone provides a general contact-stability criterion [20]. The reserve used for planning in this paper is therefore checked against an independent contact-wrench audit (Section 3.4).

2.2. Counterweight reconfiguration and stability-constrained planning

Huang et al. [6, 7] introduced stability compensation by manipulator motion and coordinated mobile-manipulator planning, Morales et al. [8] used centre-of-gravity estimation and arm motion on a field robot with a heavy manipulator, and Ding et al. [9] combined an improved tip-over moment criterion with online avoidance actions. For the machine class considered here, Paz and Mattila [21] modelled and optimised an actively articulated suspension.

Song and Sharf [10] provide the closest planning precedent: a rapidly-exploring random tree (RRT) with zero-moment-point (ZMP) stability checks and manipulator reconfiguration initialises nonlinear time-optimal planning, which receding-horizon planning then refines. Their formulation assumes a known topographic map (their Section 4) and, in simulation, perfect trajectory tracking (their Section 5). In the illustrated relocation the joint variables change only while the base is at rest, and the receding-horizon refinement is introduced to reduce such stop-and-go motion (their Section 5.4); on-board sensing, replanning and a stability-guaranteed control law are named as future work (their Section 6). COMPASS builds on this hierarchy and addresses what the perfect-tracking assumption leaves open for electrified machines: executing the counterweight reconfiguration while the carrier moves when the drives are slower than the planner assumes. Terrain preview builds on elevation mapping, whose probabilistic formulations represent the uncertainty caused by localisation [22].

2.3. Time scaling and tracking-error coordination

Trajectory scaling relates geometric motion to actuator limits [23], and reachability-based path parameterisation provides a systematic treatment of feasible timing [24]. Online time scaling moves this into execution. Dahl and Nielsen [25] modify the path accel-



Figure 1: The case-study machine with its 3D LiDAR, camera housings, satellite-navigation (GNSS) antenna masts, slew drive, lift and tilt EMLAs, telescopic-extension EMLA, wrist/gripper and oscillation-beam bogie. The simulation model of Section 3 represents this machine; field experiments with it are reported in [31, 32, 33, 34].

eration of a nominal trajectory during motion so that a torque-limited manipulator keeps following its path, and Lampinen et al. [26] scale the velocity along the trajectory of a hydraulic manipulator so that it remains achievable under a global flow limit. Hulttinen and Mattila [11] carry the idea to a wheeled hydraulic mobile manipulator: the platform path-following controller uses velocity bounds computed from the pump flow limits, the manipulator trajectory is time-scaled from the monitored position tracking error, and the scaling factor is propagated to the platform controller so that both advance together; their results are from simulation.

The execution layer of COMPASS adopts this principle. Its lock, Eq. (16), scales the commanded vehicle speed by a function of the manipulator tracking error, and because the posture reference is read at measured vehicle progress, the manipulator reference slows by the same factor; the actuator-rate constraints of the planner play the role of the limit-based velocity bounds. What differs is the role of the coupled motion. The crane and beam levelling do not execute a task trajectory: their reference is a counterweight posture that terrain preview and a balancing-capacity estimate assign to positions on the route ahead, and the speed that the coupling scales is itself limited by that capacity. A tracking error therefore means that the counterweight is not yet in place when the vehicle reaches the terrain that requires it. This raises three questions beyond task tracking: how the tracking gained by the coupling relates to stability margin, how the coupling compares with uniform slowing at equal driving time, and how often the drives of a real machine fall short of the demanded rate. Sections 6.6 and 6.7 answer them.

Control-barrier-function quadratic programmes impose forward-invariance conditions under the assumed model and feasibility conditions [27], and robustness results require bounds on the perturbations [28]; the filter properties of Section 4.7 hold under such conditions. Split-conformal regression provides marginal predictive coverage under exchangeability [29], and distribution changes require additional treatment [30]; the conformal margin is therefore used as a calibrated margin on the capacity estimate.

3. Whole-body model and stability criteria

3.1. Case-study machine

The case-study machine (Fig. 1) is a battery-electric skid-steer carrier with a loader crane. The platform has mass 4,836 kg, dimensions $3.71 \times 2.26 \times 1.39$ m, wheel radius 0.411 m, track 1.944 m and wheelbase 1.45 m. Each side has an oscillation beam carrying its front and rear wheels. Interconnected oscillation cylinders (bore/rod/stroke 80/50/140 mm) allow differential beam rotation; pumping oil changes their common rotation and levels the hull in pitch. Per side, an electric motor drives a hydrostatic pump feeding two series-connected wheel motors through 53:1 gearboxes. The crane has 716 kg moving mass and a 175 kg fixed base. Its actuators are a servo-driven slew column, lift and tilt EMLAs closing a triangle and a dog-bone four-bar, a telescopic EMLA and a three-axis wrist/gripper. The loaded case carries a rigidly gripped 400 kg block ($0.60 \times 0.50 \times 0.35$ m). A simulated 3D LiDAR supplies the terrain preview.

The same machine has been operated in field experiments on asphalt and soft soil [32, 33] and on gravel with scattered rocks [34]. The model matches it in whole-body mass (5,727 kg empty, within 5 % of the approximately 6,000 kg reported for those experiments [31, 32, 33, 34]), wheelbase and wheel radius (1.45 and 0.411 m), track (1.944 m against 1.952 m) [34] and travel-drive architecture [31].

3.2. Multibody and actuator dynamics

The model is built from CAD, datasheets and project documents; its parameter file records the source and a confidence label of every value (Table 1). Its 30 bodies include the hull, beams, compliant tyre carriers, wheels, oscillation cylinders and 13 manipulator bodies. The closed-chain compiler of RoboCompiler [35] converts the mechanism graph to a tree with five loop cuts and mobility 17, and MuJoCo [36] simulates the floating-base dynamics [37]

$$\boldsymbol{M}(\boldsymbol{q})\dot{\boldsymbol{v}} + \boldsymbol{c}_{\mathrm{b}}(\boldsymbol{q},\boldsymbol{v}) = \boldsymbol{\tau}_{\mathrm{act}} + \boldsymbol{J}_c^{\top}\boldsymbol{\lambda}_c + \boldsymbol{J}_e^{\top}\boldsymbol{\lambda}_e + \sum_w \boldsymbol{J}_w^{\top}\boldsymbol{w}_w^{\mathrm{soil}}, \tag{1}$$

with $\boldsymbol{q} \in \mathbb{R}^{34}$, $\boldsymbol{v} \in \mathbb{R}^{33}$, bias forces $\boldsymbol{c}_{\mathrm{b}}$, rigid-contact forces $\boldsymbol{\lambda}_c$, loop and equaliser constraint forces $\boldsymbol{\lambda}_e$ and soil wrenches $\boldsymbol{w}_w^{\mathrm{soil}}$. The implicit-fast integrator uses a 1 ms step, the Newton solver and elliptic friction cones.

An EMLA motor drives a screw of lead p through ratio i and efficiency η [38]:

$$k = 2\pi i/p, \qquad F = \eta k T_m, \qquad \dot{s} = \omega_m/k, \qquad m_{\mathrm{eq}} = J_m k^2. \tag{2}$$

Its saturated position/speed servo is

$$F = \mathrm{sat}_{F_{\max}}[k_p(s_{\mathrm{ref}} - s) - k_v\dot{s}], \qquad k_v = \eta m_{\mathrm{eq}} 2\pi f_s, \quad k_p = K_v k_v, \tag{3}$$

with speed-loop bandwidth f_s, position gain K_v (Table 1), a rate-limited set-point and $F_{\max} = \eta k T_{m,\max}$. The lift and tilt EMLAs use $i = 15$, $p = 12$ mm, $\eta = 0.85$ and peak motor torque 32 N m, which gives $F_{\max} = 213.6$ kN and a reflected inertia of 52.4 t. Their configured stroke-speed cap, 82 mm s^{-1}, corresponds to the 6,150 rpm maximum of the motor table; the torque–speed curves of the sizing documentation end at 3,580–3,750 rpm (47.7–50.0 mm s^{-1}), and the installed motors are rated at 10 N m and 2,000 rpm, i.e. 66.8 kN at 26.7 mm s^{-1} of stroke (Section 6.7 tests these documented limits). The largest forces and stroke speeds of the sizing simulations are 129.9/102.1 kN and 45.27/44.37 mm s^{-1}

(lift/tilt); the planning rates are 90 % of these speeds (40.7/39.9 mm s^{-1}), of which the planner uses a share $\rho = 0.9$ (Section 4.4). The extension ($i = 3$, $p = 10$ mm) has a force limit of 20.8 kN.

The slew motor of the machine's motor list (Siemens 1FK7061) has peak torque 17.3 N m, rated torque 5.4 N m and rated speed 3,000 rpm. The model uses a slew reduction ratio of 600 with efficiency 0.85 (Table 1) and caps the motor at its rated speed, which gives 8.82 kN m and a rate cap of 0.52 rad s^{-1}; planning uses 0.25 rad s^{-1}. The load inertia ranges from 1,600 to 12,400 kg m^2, and the rotor inertia adds 180 kg m^2 at the column. The servo uses a 4 Hz speed loop tuned at 5,000 kg m^2, position gain 5 s^{-1} and set-point acceleration 0.25 rad s^{-2}. A holding brake with twice the peak drive torque engages after 0.2 s at rest. The wrist holds $q_{w2} = -(q_l + q_t)$ to keep the payload upright.

The travel-drive states are pump speed and three line pressures per side. Pump torque balances electric-motor torque and hydraulic load, and the pressure rates follow compressibility, pump and motor flow, leakage, relief and make-up flows [39, 40]. Cross-port relief and charge pressures are 310 and 30 bar, and the commanded side speed sets the pump speed; the complete equations and coefficients are documented with the code. The oscillation constraint is $s_L + s_R = c_o$; levelling changes c_o at up to 0.04 m s^{-1}, approximately 0.05 rad s^{-1} of hull pitch within $|o| \le 0.12$ rad.

3.3. Ground contact and terramechanics

Rigid fill and convex-mesh rocks use compliant MuJoCo contact with friction coefficients 0.60/0.70, rolling resistance 0.02 and radial tyre stiffness 1.2×10^6 N m^{-1} with damping ratio 0.1; the assigned ground coefficient combines with the robot coefficient 0.50 in the contact model [41]. Loose fill uses a 5 cm soil grid [42], the pressure–sinkage law $p = (k_c/b + k_\phi) z^n$ and the shear law $\tau_s = (c_s + p \tan \phi_s)(1 - e^{-j/K})$, with elastic unloading, passive earth pressure, ruts and berms [43, 44]. The fill uses the pressure–sinkage parameters of LETE sand [44], $n = 0.79$, $k_c = 102$ kN/m^{n+1} and $k_\phi = 5301$ kN/m^{n+2}, with cohesionless-gravel shear parameters $\phi_s = 38°$ and $c_s = 0$. Every 100 physics steps, the deformed soil surface, after local slope relaxation, is copied to a height field that supports contacts of the hull, beams and boom with the loose fill.

3.4. Reserve proxy and contact-wrench audit

For the centre of mass $\boldsymbol{c}$, the controller uses ideal wheel-support points constructed from wheelbase, track and beam angles rather than the actual rock-contact locations. An edge e passes through $\boldsymbol{a}_e$ along the unit vector $\boldsymbol{u}_e$, counter-clockwise from above. Let $\hat{\boldsymbol{r}}_e$ point perpendicularly from $\boldsymbol{c}$ to the edge and let $\hat{\boldsymbol{f}}_\perp$ be the unit projection of $\boldsymbol{f} = \boldsymbol{g} - \ddot{\boldsymbol{c}}$ normal to that edge. The reduced reserve is

$$\theta_e = -\operatorname{atan2}\left(\boldsymbol{u}_e \cdot (\hat{\boldsymbol{f}}_\perp \times \hat{\boldsymbol{r}}_e),\ \hat{\boldsymbol{f}}_\perp \cdot \hat{\boldsymbol{r}}_e\right), \qquad \theta = \min_e \theta_e. \tag{4}$$

With $\boldsymbol{f} = \boldsymbol{g}$ this is a model-based *static reserve*. With measured or simulated acceleration it is an *acceleration-based reserve proxy*, denoted θ_{acc}. It omits the moments included in the full force-angle construction [12, 13].

The missing mechanics follow from the whole-system force and moment balance about the centre of mass [20],

$$\sum_j \boldsymbol{f}_j = m(\ddot{\boldsymbol{c}} - \boldsymbol{g}) - \boldsymbol{F}_x, \qquad \sum_j [(\boldsymbol{p}_j - \boldsymbol{c}) \times \boldsymbol{f}_j + \boldsymbol{\mu}_j] = \dot{\boldsymbol{H}}_C - \boldsymbol{\tau}_x, \tag{5}$$

Table 1: Main model parameters (sources: D documentation and datasheets, CAD, M earlier simulation models, C computed, A assumption, L literature).

Quantity	Value	Source
Platform mass / centre-of-mass height (vehicle frame)	4,836 kg / 0.805 m	D
Wheel radius / track / wheelbase	0.411 / 1.944 / 1.45 m	M, CAD, D
Oscillation beam mass (each)	300.3 kg	C
Beam rotation relative to hull	$-0.159 \ldots 0.201$ rad	C
Levelling rate (equaliser offset)	$0.04\,\mathrm{m\,s^{-1}}$	A
Tyre radial stiffness / damping ratio	$1.2 \times 10^6\,\mathrm{N\,m^{-1}}$ / 0.1	C / A
Pump / wheel motor displacement	17.83 / 44.5 $\mathrm{cm^3/rev}$	D
Gear ratio / relief / charge pressure	53 / 310 / 30 bar	D
E-motor peak torque / pump speed limit	173 N m / 3,600 rpm	D
Speed limit	$0.83\,\mathrm{m\,s^{-1}}$	D
Crane moving mass / fixed base	716 / 175 kg	M, C
Lift, tilt EMLA: mass with motor	64.5 kg	D, L
Lift, tilt EMLA: i, p, η, $T_{\max}$	15, 12 mm, 0.85, 32 N m	D
Lift, tilt motor: rated torque / speed	10 N m / 2,000 rpm	D
Lift, tilt motor: table / curve max. speed	6,150 / 3,580–3,750 rpm	D
Lift, tilt EMLA: closed length + stroke	0.919 + 0.552 m	D
Extension EMLA: i, p, $T_{\max}$, range	3, 10 mm, 13 N m, 0–1 m	D, M
Extension motor: rated torque / speed	3.7 N m / 3,000 rpm	D
EMLA drive speed bandwidth / position gain	20 Hz / $16.7\,\mathrm{s^{-1}}$	A
Slew motor peak torque / ratio	17.3 N m / 600	D / A
Slew motor: rated torque / speed	5.4 N m / 3,000 rpm	D
Slew torque limit / holding brake	8.82 / 17.6 kN m	C / A
Slew rotor inertia / speed loop / position gain	$5\,\mathrm{kg\,cm^2}$ / 4 Hz / $5\,\mathrm{s^{-1}}$	A
Slew set-point acceleration	$0.25\,\mathrm{rad\,s^{-2}}$	A
Payload (loaded case)	400 kg, $0.6 \times 0.5 \times 0.35$ m	A
Tyre–gravel / tyre–rock friction	0.60 / 0.70	L / A
Whole-body mass (empty / loaded)	5,727 / 6,127 kg	C

where $\boldsymbol{H}_C$ is the total centroidal angular momentum, $(\boldsymbol{F}_x, \boldsymbol{\tau}_x)$ are other external loads and $(\boldsymbol{f}_j, \boldsymbol{\mu}_j)$ are ground forces and couples at the actual contacts $\boldsymbol{p}_j$. Internal actuator and loop forces cancel, although articulation changes $\dot{\boldsymbol{H}}_C$. The restoring moment about edge e is

$$S_e = -\boldsymbol{u}_e^{\top}\Big[(\boldsymbol{c} - \boldsymbol{a}_e) \times \{m(\boldsymbol{g} - \ddot{\boldsymbol{c}}) + \boldsymbol{F}_x\} + \boldsymbol{\tau}_x - \dot{\boldsymbol{H}}_C\Big]; \tag{6}$$

with gravity alone it is positive for a stationary centre of mass inside a level footprint. When $\boldsymbol{F}_x = 0$, dropping the moment term introduces the edge-moment error $|\boldsymbol{u}_e^{\top}(\boldsymbol{\tau}_x - \dot{\boldsymbol{H}}_C)|$, so the reduced reserve is reliable only where this error is small compared with the reduced restoring moment. On non-coplanar terrain the general test is whether the six-dimensional wrench required by (5) belongs to the feasible set of the actual contacts, with unilateral normal forces, friction and contact couples.

An independent audit evaluates these quantities without using the controller's computation. At 50 Hz, a separate MuJoCo data instance is forwarded at the simulated configuration, velocity, inputs and applied soil loads; it does not feed the controller. Rigid-

contact forces and couples are summed at their actual positions and frames, soil forces and couples at the wheel centres, and the body inertias, velocities and accelerations give $\dot{\boldsymbol{H}}_C$ and the required force in (5) independently; normalised balance residuals check this accounting. For each nondegenerate horizontal convex envelope of the recorded contact positions, with the full 3D edge heights and $\boldsymbol{r}_e = (\boldsymbol{I} - \boldsymbol{u}_e\boldsymbol{u}_e^\top)(\boldsymbol{c} - \boldsymbol{a}_e)$, the equivalent force including the edge moment is

$$\boldsymbol{f}^{\mathrm{w}}_{e,\perp} = (\boldsymbol{I} - \boldsymbol{u}_e\boldsymbol{u}_e^\top)\boldsymbol{F}^* + \frac{\boldsymbol{u}_e^\top \boldsymbol{N}^*}{\|\boldsymbol{r}_e\|^2}(\boldsymbol{u}_e \times \boldsymbol{r}_e), \quad \boldsymbol{F}^* = -\boldsymbol{F}_{\mathrm{ground}}, \quad \boldsymbol{N}^* = -\boldsymbol{M}_{\mathrm{ground},C}. \tag{7}$$

Substitution into (4) gives the moment-inclusive *wrench angle* θ_{w}, and repeating it with $\boldsymbol{N}^* = 0$ gives θ_{F} on identical geometry and force. Thus $\theta_{\mathrm{F}} - \theta_{\mathrm{w}}$ isolates moment inclusion, whereas $\theta_{\mathrm{acc}} - \theta_{\mathrm{w}}$ also includes geometry, filtering and force-estimation differences. The equivalent force can point away from the edge, which produces an angle near $\pm 180°$ even when the edge moment is small, so angular comparisons use the compressive branch $\boldsymbol{f}^{\mathrm{w}}_{e,\perp} \cdot (-\boldsymbol{r}_e) > 0$ on every edge, whereas moment and event summaries keep all samples. The principal audit outcome is the signed restoring moment $\min_e S_e$ of (6), normalised by body weight times one metre (BW m), and its first percentile over an episode, $S_{1\%}$, evaluated on samples with nondegenerate wheel support. For the loaded machine, 10^{-3} BW m is about 60 N m, whereas the static restoring moment about a side edge on level ground is about 1 BW m; $S_{1\%}$ thus measures how closely the lower tail of the restoring moment approaches zero. Line or point contact envelopes are counted as degenerate support, and samples with additional body or payload support are reported separately. Before the envelope is evaluated, contact points closer than 10 µm are merged, because nearly coincident contacts create numerically spurious tipping axes; results with this *contact conditioning* at 1 and 100 µm are also reported. Because $S_{1\%}$ is evaluated only on nondegenerate wheel support, it is read together with the exposure to degenerate support and with the edge-rotation events defined below.

Two audit outcomes do not use the envelope. On a regular ten-second grid, a six-dimensional wrench-feasibility calculation maximises the additional moment about each coordinate axis at fixed force and other moment components, subject to unilateral contacts and inner and outer polyhedral bounds on the recorded friction cones; where both problems are feasible, they bracket the additional moment capacity of this instantaneous friction-cone relaxation. An *edge-rotation event* requires a single supported wheel edge to carry more than 50 % of the total upward wheel reaction, less than 2 % on its opposite wheels, and an outward hull angular velocity above $0.02\,\mathrm{rad\,s^{-1}}$ for at least 0.1 s.

The model package includes checks of kinematics, static wheel loads, drive speed and stall torque, the equaliser and closed-form soil cases. The compiled soil and wheel forces used in the simulations differ from the reference implementation by at most 2.4×10^{-16} relative, and the drive torques agree exactly. On a reference mission across a cross slope with the payload, steps of 1 and 0.5 ms give a roll RMS difference of 0.37° and the same maximum roll, 31.6°.

4. COMPASS framework

4.1. Objective and priority structure

The objective is to traverse a nominal site route Γ_0, empty or with a payload, while maintaining a reserve and respecting the actuator limits. The balancing coordinates are

$$\boldsymbol{x} = (\psi, q_l, q_t, q_e, o) \in \mathcal{X}, \tag{8}$$

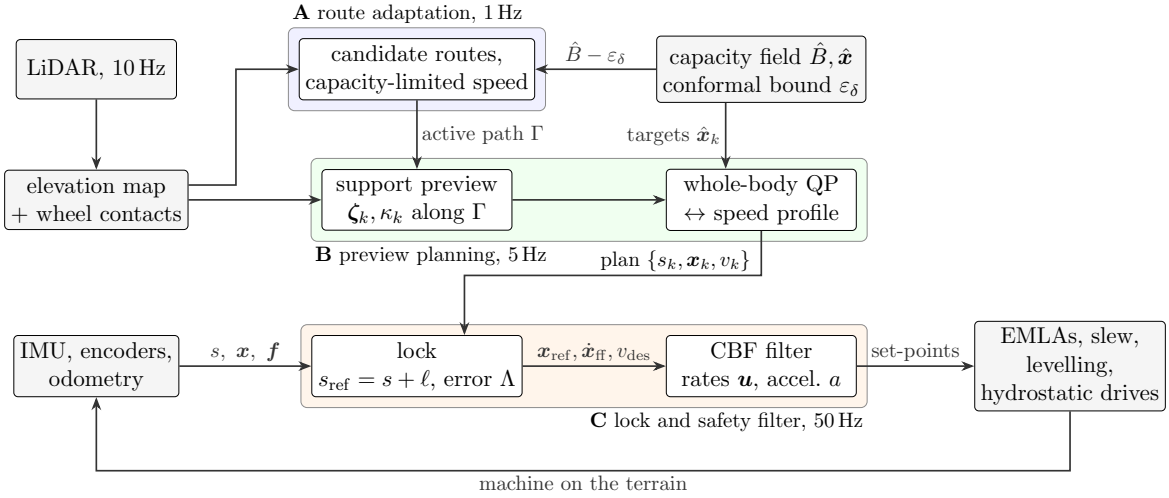


Figure 2: COMPASS. The LiDAR map feeds the support preview; the calibrated capacity estimate supports route selection and posture targets. Layer B couples a path-indexed posture plan with a speed profile. Layer C reads the plan at measured progress, reduces the commanded speed when the manipulator lags, and filters the commands using the acceleration-based reserve proxy.

namely slew, lift, tilt, extension and hull levelling relative to the oscillation beams. The actuator coordinates are $\boldsymbol{a} = (\psi, s_l(q_l), s_t(q_t), q_e, c_o(o))$; the closed-loop maps give the EMLA strokes and the equaliser offset. With $\boldsymbol{G} = \partial \boldsymbol{a}/\partial \boldsymbol{x}$, the actuator-speed limits give the componentwise joint-rate limits $\bar{\boldsymbol{r}} = |\boldsymbol{G}|^{-1}\bar{\boldsymbol{a}}$. The support attitude $\boldsymbol{\zeta} = (\phi, \vartheta, \chi)$ comprises beam-plane roll, pitch and beam twist.

COMPASS gives priority to counterweight motion, then to changes of the traverse, then to progress. The manipulator and the beam levelling remain active to seek $\theta \geq \theta_{\min}$; if the predicted counterweight capacity is insufficient, the platform slows, shifts its route or receives a stop profile. Three layers implement this structure (Fig. 2): Layer A adapts the route at 1 Hz, Layer B plans path-indexed postures and speeds at 5 Hz, and Layer C applies the lock and a CBF-based command filter at 50 Hz. Planning uses the quasi-static reserve plus an empirical dynamic allowance; filtering uses the measured acceleration-based reserve proxy. Appendix A lists the parameters of the framework.

4.2. Balancing capacity and learned surrogate

With the wrist levelled, the first mass moment of the manipulator about a point $\boldsymbol{b}$ on the slew axis is

$$\boldsymbol{S}(\psi, q_l, q_t, q_e) = \boldsymbol{R}_z(\psi)\big(\boldsymbol{A}(q_l, q_t) + q_e \boldsymbol{B}(q_l, q_t)\big). \tag{9}$$

The fields $\boldsymbol{A}, \boldsymbol{B}$ are sampled from the closed-chain model and interpolated with bicubic splines. The whole-body centre of mass, its Jacobian and the static actuator efforts follow from the mass moments and virtual work; levelling rotates the hull and crane about the beam pivots. Agreement with the full model is 1×10^{-9} m for the centre of mass and 2×10^{-6} relative for the efforts, at 0.16 ms per evaluation.

For each support attitude, the posture map solves

$$\boldsymbol{x}_{\mathrm{map}}(\boldsymbol{\zeta}) = \arg \max_{\boldsymbol{x} \in \mathcal{X}_{\mathrm{adm}}(\boldsymbol{\zeta})} \left\{ \min_e \theta_e(\boldsymbol{x}; \boldsymbol{\zeta}) - \lambda \|\boldsymbol{x} - \boldsymbol{x}_{\mathrm{tr}}\|^2_{\boldsymbol{W}_t} \right\}, \tag{10}$$

and the quasi-static capacity is $B(\boldsymbol{\zeta}) = \min_e \theta_e(\boldsymbol{x}_{\mathrm{map}}; \boldsymbol{\zeta})$. The admissible set includes joint and stroke ranges, beam end stops, EMLA design forces, 80 % of the extension-force

Table 2: Learned capacity field: training samples (teacher grid), calibration/test samples (independent random attitudes), test errors of the predicted capacity [deg], and conformal margins ε_δ [deg] with the empirical test coverage [%] in parentheses (target $1-\delta$).

Payload	Train	Cal./test	RMSE	MAE	$\delta = 0.10$	$\delta = 0.05$	$\delta = 0.01$
empty	2871	400/400	0.61	0.46	0.27 (91.8)	0.49 (95.8)	0.95 (99.0)
400 kg	2871	400/400	1.23	0.78	0.55 (87.2)	0.98 (95.5)	5.72 (100.0)

capacity, slew holdability, and a tool clearance of 0.75 m above the support plane and 0.10 m above the deck. The quasi-static slew torque is $T_s = -\boldsymbol{g} \cdot \partial_\psi \boldsymbol{S}$, with amplitude

$$\hat{T}_s = \|\boldsymbol{g}_{xy}\| \|\boldsymbol{S}_{xy}\|. \tag{11}$$

Requiring $\hat{T}_s \le 0.8 T_{s,\max}$ avoids relying on a low-torque but unstable uphill equilibrium whose disturbance could exceed the holding torque of the drive. The proximity penalty ($\lambda = 0.002$, weights based on actuator travel time) favours postures close to the transport pose $\boldsymbol{x}_{\mathrm{tr}}$, the driving posture with the payload ahead of the hull (slew 90°, lift 0.85 rad, tilt -1.8 rad, extension retracted, no levelling). Sequential least-squares quadratic programming (SLSQP) starts from the transport pose; if a grid of 23,040 postures improves the reserve by more than 1°, its three best candidates are also refined. A multistart search without the proximity penalty gives the comparison capacity B^*, the best reserve found over all starts.

A multilayer perceptron (three inputs, three hidden layers of 64 tanh units, eight outputs) learns B, the posture map (slew represented by sine and cosine) and B^* from 2,871 grid attitudes per payload (roll ±40°, pitch ±35° and twist). Where no admissible posture exists, the target is the brake-held transport pose and its static reserve (17 % of the payload grid and none of the empty grid). Split conformal calibration [29] uses the $\lceil (n+1)(1-\delta) \rceil$-th ordered score $\hat{B} - B$, denoted ε_δ, so that for a new attitude exchangeable with the n calibration attitudes

$$\Pr\{B(\boldsymbol{\zeta}) \ge \hat{B}(\boldsymbol{\zeta}) - \varepsilon_\delta\} \ge 1 - \delta. \tag{12}$$

Layer A uses $\hat{B} - \varepsilon_{0.05}$; the margins are 0.49° empty and 0.98° loaded (Table 2). The coverage is marginal over calibration and a new exchangeable attitude. Errors are largest near the payload feasibility boundary: RMS 3.2° within 5° of an infeasible attitude versus 0.9° elsewhere. Layer C acts on the measured reserve proxy rather than on the learned capacity.

4.3. Terrain preview

LiDAR returns form a 0.1 m elevation grid [22]: the highest return per scan is averaged over scans, small gaps are filled from neighbours and larger gaps remain unknown. Wheel heights supplement the map. At each path pose the highest terrain point in each wheel footprint predicts the contact height z, and

$$\vartheta = \tfrac{1}{2}(\vartheta_L + \vartheta_R), \quad \phi = \operatorname{atan2}\left(\tfrac{1}{2}(z_{FL} + z_{RL} - z_{FR} - z_{RR}), w\right), \quad \chi = \tfrac{1}{2}(\vartheta_L - \vartheta_R), \tag{13}$$

where $\vartheta_s = \operatorname{atan2}(z_{Rs} - z_{Fs}, \ell_w)$, w is the track and ℓ_w the wheelbase. The terrain excitation is $\kappa(s) = \|\partial_s^2(\phi, \vartheta)\|$ after smoothing. The dynamic allowance

$$\Delta\theta(v, \kappa) = k_\kappa \kappa v^2 + k_v v + k_0 \tag{14}$$

Algorithm 1 Posture–speed synchronisation (Layer B, 5 Hz)

Require: path Γ, progress s_0, posture $\boldsymbol{x}_0$, speed v_0, elevation map
predict $\boldsymbol{\zeta}_k, \kappa_k$; query $\hat{\boldsymbol{x}}_k$ at $s_k = s_0 + k\Delta s$
initialise $v_k = v_{\text{nom}}$; cap to v_{unk} from three samples before unknown terrain
for $j = 1, \dots, 4$ **do**
    $\Delta t_k \leftarrow 2\Delta s/(v_{k-1} + v_k)$; solve (15) twice
    $\theta_k \leftarrow \min_e \theta_e(\boldsymbol{x}_k; \boldsymbol{\zeta}_k)$ (nonlinear model evaluation)
    if $\theta_k \geq \theta_{\min} + \Delta\theta(v_k, \kappa_k)$ for all k **then break**
    end if
    $k^* \leftarrow$ first violation; $v_i \leftarrow 0.6 v_i$ for $i \in [k^* - 5, k^*]$
end for
for each sample still violating its requirement **do**
    if $\theta_k \geq \theta_{\min} + \Delta\theta(v_{\min}, \kappa_k)$ **then**
        $v_i \leftarrow \min(v_i, v_{\min})$ for $i \in [k-5, k]$
    else
        $v_i \leftarrow 0$ for $i \geq k - 5$; mark infeasible; **break**
    end if
end for
backward pass $v_k \leq \sqrt{v_{k+1}^2 + 2a_{\text{dec}}\Delta s}$; forward pass $v_{k+1} \leq \sqrt{v_k^2 + 2a_{\text{acc}}\Delta s}$
return $\{s_k, \boldsymbol{x}_k, v_k, \theta_k\}$

is fitted to the 95th conditional quantile of the observed difference between the static reserve and the acceleration-based proxy (Section 5.5), giving the requirement $\theta_{\text{req}} = \theta_{\min} + \Delta\theta$.

4.4. Layer B: posture and speed planning

At $s_k = s_0 + k\Delta s$, $k = 1, \dots, N$, Layer B predicts $\boldsymbol{\zeta}_k, \kappa_k$, queries the targets $\hat{\boldsymbol{x}}_k$ and assigns speeds v_k, hence intersample times Δt_k. Two sequential-quadratic-programming (SQP) iterations about the previous plan solve

$$
\begin{aligned}
\max_{\boldsymbol{x}_{1..N}, t, \sigma \geq 0} \quad & w_t t - \sum_k \|\boldsymbol{x}_k - \hat{\boldsymbol{x}}_k\|_{\boldsymbol{W}}^2 - \sum_k \|\boldsymbol{x}_k - \boldsymbol{x}_{k-1}\|_{\boldsymbol{R}}^2 - w_\sigma(\sigma^2 + \sigma) & \\
\text{s.t.} \quad & \theta_e(\bar{\boldsymbol{x}}_k; \boldsymbol{\zeta}_k) + \nabla\theta_e(\boldsymbol{x}_k - \bar{\boldsymbol{x}}_k) \geq \theta_{\text{req},k} + t & \forall k, e \\
& |\boldsymbol{G}(\bar{\boldsymbol{x}}_k)(\boldsymbol{x}_k - \boldsymbol{x}_{k-1})| \leq \rho\bar{\boldsymbol{a}}\Delta t_k & \forall k \\
& |\boldsymbol{E}(\bar{\boldsymbol{x}}_k) + \nabla\boldsymbol{E}(\boldsymbol{x}_k - \bar{\boldsymbol{x}}_k)| \leq (1+\sigma)\widetilde{\boldsymbol{E}}_{\max,k} & \forall k \\
& h(\bar{\boldsymbol{x}}_k) + \nabla h(\boldsymbol{x}_k - \bar{\boldsymbol{x}}_k) + \ell_\sigma\sigma \geq \widetilde{h}_{\min,k} & \forall k \\
& \boldsymbol{x}_{\min,k} \leq \boldsymbol{x}_k \leq \boldsymbol{x}_{\max,k}, \qquad t \leq t_{\max}. &
\end{aligned}
\tag{15}
$$

Here $\boldsymbol{E}$ contains the slew-torque amplitude and the lift and tilt forces, and h is the clearance of the lowest tool or payload point above the mapped terrain. The soft rows use $\widetilde{E}_{\max,k,j} = \max(E_{\max,j}, |E_j(\bar{\boldsymbol{x}}_k)| + \epsilon_j)$ and $\widetilde{h}_{\min,k} = \min(0.55\,\text{m}, h(\bar{\boldsymbol{x}}_k) - 0.05\,\text{m})$, with $\epsilon_j = 1$ in the SI unit of each effort and $\ell_\sigma = 1\,\text{m}$; they keep each quadratic programme (QP) feasible when the current posture already exceeds the nominal design effort or clearance. The joint and trust-region box includes the starting posture and a predicted beam-end-stop margin of 0.05 rad, relaxed where needed to include that posture. There are $5N + 2$ variables and approximately $21N$ rows; PIQP solves a QP in about 2 ms [45]. A negligible

quadratic regularisation of t is omitted from the display. The surplus $t \in \mathbb{R}$ has no lower bound, and its cap is $t_{\max} = 6°$.

Algorithm 1 couples posture and speed: a missed reserve requirement reduces the approach speeds, which gives the actuators more time, and a sample that fails even at the creep speed $v_{\min}$ receives a stop profile; v_{unk} caps the speed ahead of unmapped terrain, and a_{acc} and a_{dec} are the acceleration limits (Table A.1). Backward and forward passes impose the commanded acceleration limits. This is path time scaling [23, 24] with quasi-static balance and actuator rates. A plan starts from its previous posture at the current progress if the tracking error Λ of (16) is below its lower threshold Λ_0, and otherwise from the measurement.

4.5. Layer C: lock and command filter

The reference is read at measured progress plus a lead, $s_{\text{ref}} = s + \ell_0 + \tau_\ell v$, with feed-forward $\dot{\boldsymbol{x}}_{\text{ff}} = (\mathrm{d}\boldsymbol{x}/\mathrm{d}s)v$, so a delayed platform delays the posture reference. Conversely, a posture tracking error slows the platform command. The tracking error is normalised by the planning rate of each coordinate and expressed in seconds of actuator travel,

$$\Lambda = \max_j \frac{|x_{\text{ref},j} - x_j|}{\bar{r}_j(\boldsymbol{x})}, \qquad v_{\text{des}} = v_{\text{plan}}(s)\,\mathrm{sat}_{[0,1]}\left(\frac{\Lambda_{\max} - \Lambda}{\Lambda_{\max} - \Lambda_0}\right). \tag{16}$$

We call Λ the *rate-normalised tracking error* and the speed reduction in (16) the *lock*; it implements the tracking-error time scaling of [11], and the paper calls its use *locking*. Λ is a posture error in time units, not a delay; a directly measured time-shift delay τ^* is defined in Section 5.4. The CBF-based filter [27] optimises the manipulator and levelling rates $\boldsymbol{u}$ and the commanded platform acceleration a:

$$\begin{aligned}
\min_{\boldsymbol{u},a,\varsigma\geq 0} \quad & \|\boldsymbol{u} - \boldsymbol{u}_{\text{nom}}\|^2_{\boldsymbol{W}_u} + w_a(a - a_{\text{nom}})^2 + w_\varsigma \varsigma^2 \\
\text{s.t.} \quad & \theta_e + \tau_c \nabla_{\boldsymbol{x}} \theta_e \boldsymbol{u} + \partial_a \theta_e (a - a^-) + \varsigma \geq \theta_{\min} + (1-\gamma)(\theta_e - \theta_{\min}) \quad \forall e \\
& h + \tau_c \nabla h \boldsymbol{u} + \varsigma \geq h_{\min} + (1-\gamma)(h - h_{\min}) \\
& \hat{T}_s + \tau_c \nabla \hat{T}_s \boldsymbol{u} \leq \max(0.8 T_{s,\max}, \hat{T}_s), \quad |\boldsymbol{u}| \leq \bar{\boldsymbol{r}}_{\text{hw}} \\
& \boldsymbol{x}_{\min} \leq \boldsymbol{x} + \tau_c \boldsymbol{u} \leq \boldsymbol{x}_{\max}, \quad a_{\min} \leq a \leq a_{\text{nom}}.
\end{aligned} \tag{17}$$

Here θ_e is the controller's acceleration-based proxy, evaluated with the measured hull specific force rather than the centre-of-mass acceleration used for evaluation (Section 5.4), $\partial_a \theta_e$ approximates its sensitivity to the commanded acceleration, a^- is the previous command, $\boldsymbol{u}_{\text{nom}} = \dot{\boldsymbol{x}}_{\text{ff}} + \boldsymbol{K}(\boldsymbol{x}_{\text{ref}} - \boldsymbol{x})$ and $a_{\text{nom}} = \mathrm{sat}_{[-2a_{\text{dec}}, a_{\text{acc}}]}((v_{\text{des}} - v_{\text{cmd}})/\tau_v)$, with $\tau_v = 0.4\,\text{s}$. The lower bound is $a_{\min} = \min(a_{\text{nom}}, \max(-2a_{\text{dec}}, -v/t_b - a_b))$, with $t_b = 0.5\,\text{s}$ and $a_b = 0.5\,\text{m}\,\text{s}^{-2}$, so the acceleration command cannot exceed the nominal command. The slew row is relaxed if the slew torque is already beyond its nominal limit, and the slack, scaled in SI units, keeps the filter feasible. On solver failure, clipped nominal commands are returned. The filtered rates become drive set-points, and pure pursuit gives the skid-steer side speeds [46].

4.6. Layer A: route adaptation

Candidate paths are lateral offsets d_j at 0.4 m spacing in the permitted corridor, joined by a 6 m quintic blend. Over an 18 m mapped horizon, the capacity-limited speed is

$$\bar{v}_j(s) = \max\{v \in [0, v_{\text{nom}}] : \hat{B}(\boldsymbol{\zeta}_j(s)) - \varepsilon_\delta \geq \theta_{\min} + \Delta\theta(v, \kappa_j(s))\}, \tag{18}$$

with zero assigned if no speed satisfies the inequality. A route is model-feasible when all observed samples allow $\bar{v}_j \geq v_{\min}$, and capacity-limited when any allows less than the nominal speed. The active route is kept while it is not capacity-limited, except that the route layer may return towards the nominal route along candidates that are not capacity-limited either. Otherwise, model-feasible candidates are compared by $T_j = \sum_s \Delta s/\bar{v}_j(s)$ and by $T_j + w_d|d_j| + w_\zeta\overline{\phi^2 + \vartheta^2} + w_\kappa\overline{\kappa}$. If the current route is feasible, a candidate must first improve the traversal time by at least 10 % (and 0.5 s). Unknown samples are priced at v_{unk} and do not determine feasibility, and eligible horizons must be 60 % observed. Switches wait one blend length unless the active route becomes infeasible. If all routes are infeasible, a switch requires an improvement of the smallest capacity margin by at least twice the switching margin $m_{\text{sw}} = 1°$, and Layer B issues a stop profile.

4.7. Conditional planning and filtering properties

The following statements concern the stated models and command laws; the proofs are in Appendix B.

Proposition 1 (Feasible QP and checked planning samples)**.** *If the joint box contains* $\boldsymbol{x}_0$, (15) *has a feasible hold trajectory for every linearisation and speed profile. Every future planning sample* $k = 1, \ldots, N$ *with* $v_k > 0$ *satisfies* $\min_e \theta_e(\boldsymbol{x}_k; \boldsymbol{\zeta}_k) \geq \theta_{\min} + \Delta\theta(v_k, \kappa_k) - \epsilon_\theta$ *in the predicted quasi-static model, with implementation tolerance* $\epsilon_\theta = 10^{-4}$ *rad. If the creep speed fails at* k^**, the profile is zero from* $k^* - 5$ *(*2.5 m *earlier), and its backward pass satisfies* $v_k \leq \sqrt{2a_{\text{dec}}(s_{\text{stop}} - s_k)}$ *upstream.*

Proposition 2 (Lock for the command law)**.** *For the continuous command law* $\dot{v}_{\text{cmd}} = a$*, with nonnegative speed saturation and* $a \leq a_{\text{nom}}$ *as in* (17)*, an interval with* $\Lambda \geq \Lambda_{\max}$ *has commanded progress at most* $d_b(v_1) = \tau_v v_1$ *if* $v_1 \leq 2a_{\text{dec}}\tau_v$*, and otherwise* $d_b(v_1) = (v_1^2 - 4a_{\text{dec}}^2\tau_v^2)/(4a_{\text{dec}}) + 2a_{\text{dec}}\tau_v^2$*. When* $\Lambda < \Lambda_{\max}$*,* $|x_j - x_{\text{plan},j}(s + \ell)| < \bar{r}_j\Lambda_{\max}$*. At the same support attitude, Lipschitz constants* L_θ, L_x *bound the static-reserve difference from the plan at* s *by* $L_\theta\|\bar{\boldsymbol{r}}\|\Lambda_{\max} + L_\theta L_x \ell$*.*

Proposition 3 (Sampled proxy bound with slack)**.** *Let* $\beta_{e,i} = \theta_e(t_i) - \theta_{\min}$*,* $\alpha = \gamma/\tau_c > 0$*, and* $q = 1 - \alpha T_{\text{C}} \in [0, 1)$*. Assume* $\beta_{e,i+1} = \beta_{e,i} + (T_{\text{C}}/\tau_c)\hat{\Delta}_{e,i} + T_{\text{C}} w_{e,i}$*, where* $\hat{\Delta}_e = \tau_c \nabla_{\boldsymbol{x}}\theta_e \boldsymbol{u} + \partial_a\theta_e(a - a^-)$ *and* $|w_{e,i}| \leq \bar{w}$*. A solved filter with slack* ς_i *gives*

$$\beta_{e,i} \geq q^i \beta_{e,0} - T_{\text{C}} \sum_{j=0}^{i-1} q^{i-1-j} \left(\bar{w} + \frac{\varsigma_j}{\tau_c} \right). \tag{19}$$

For zero slack this is $\beta_{e,i} \geq q^i\beta_{e,0} - (\bar{w}/\alpha)(1 - q^i)$*.*

Proposition 1 establishes sampled reserve checks and a commanded stopping profile for the planning model; the algorithm displays the sample tests without the numerical tolerance. Proposition 2 bounds commanded progress, not actual progress. In Proposition 3, T_{C} is the filter period and the residual w collects terrain excitation, differences between specific force and centre-of-mass acceleration, actuator tracking and linearisation error, so the result is a conditional sampled bound [28] in terms of the residual bound $\bar{w}$.

5. Simulation studies

5.1. Simulator and scenarios

All runs use MuJoCo 3.15 [36] with the whole-body model of Section 3. The physics step is 1 ms, the drive and actuator models are updated every 10 ms, and the controllers

run at 50 Hz on noisy measurements: hull attitude (0.2° standard deviation), specific force of an inertial measurement unit (IMU) at the hull reference point (noise 0.05 m s^{-2}, low-pass filter with a 0.05 s time constant), joint encoders and a planar pose (2 cm). The LiDAR is modelled after the sensor in the machine's CAD (32 channels over ±22.5°, 180 columns over the front 180°, 30 m range, 2 cm range noise, 10 Hz), mounted at the front of the hull and pitched 10° down, by ray casting against the visible terrain surface (deformed loose layer and rock tops). Actuator speed caps are enforced as rate limits on the drive set-points, together with the slew set-point acceleration limit of Section 3.2.

Four construction-site scenarios (Fig. 3) are generated from a seed, which sets the feature severity, the roughness of the compacted fill, the rocks (8–18 cm high, partly in the wheel tracks) and the loose-soil patches. The nominal route is the straight connection a site-logistics planner would give the machine; the route layer may move it within a scenario-specific corridor.

S1 embankment: a 63 m traverse along a 6 m wide embankment face whose cross slope rises to $\alpha_{\max} \sim \mathcal{U}(24°, 29°)$ over 20 m, with rocks in the wheel tracks and a loose strip on the downhill side; corridor −2.6 m to 2.0 m.

S2 crest: climb of an 18°–23° spoil-heap ramp to a 1.8–2.3 m high top 2–3.5 m long, and descent (pitch reversal); rocks and loose material on the top; a gentler access ramp of 11° lies 7 m to the side; corridor −3 m to 8 m.

S3 ditch: diagonal crossing (30°–45°) of a 0.25–0.35 m deep drainage ditch and a 0.25–0.35 m high berm (sides up to 23°) on a fill with a cross slope of 9°–13°, with rocks and a loose patch; the crossing twists the beams and excites roll and pitch together; corridor ±3 m.

S4 stress: a 48 m embankment face whose cross slope increases from 16° to $\alpha_{\max} \in \{26°, 30°, 34°, 38°\}$.

Every scenario is run with the empty manipulator and with the 400 kg payload held in the gripper. The nominal speed is 0.8 m s^{-1}, just below the 3 km h^{-1} of the operating instructions; one study repeats the execution comparison at 0.35 m s^{-1}, the speed limit used in the field data collection with the machine [33].

5.2. Controllers and execution variants

Table 3 lists the compared controllers. FIX holds the manipulator in its transport pose and follows the route at the nominal speed; SLOW adds an operator-style speed rule on the measured static reserve. REACT is a reactive active counterweight: the manipulator and the levelling are commanded to the posture map of the measured (filtered) support attitude within the configured drive limits, and REACT-S adds the speed rule. PREV plans with the same preview, posture targets and quadratic programme as COMPASS, but without the iterative speed adjustment; it ignores the planned speed profile, indexes its references by time at the nominal speed and has no lock, safety filter or route layer. Each ablation removes one component of COMPASS; COMPASS−S keeps the path-indexed reference and the acceleration limits and unmapped-terrain cap of the speed profile, but neither lowers the speed for the actuators or the reserve nor waits for a lagging manipulator. All balancing controllers use the same capacity network.

PREV changes several execution mechanisms together; a set of execution variants isolates them. All variants keep the COMPASS planner with its iterative speed adjustment,

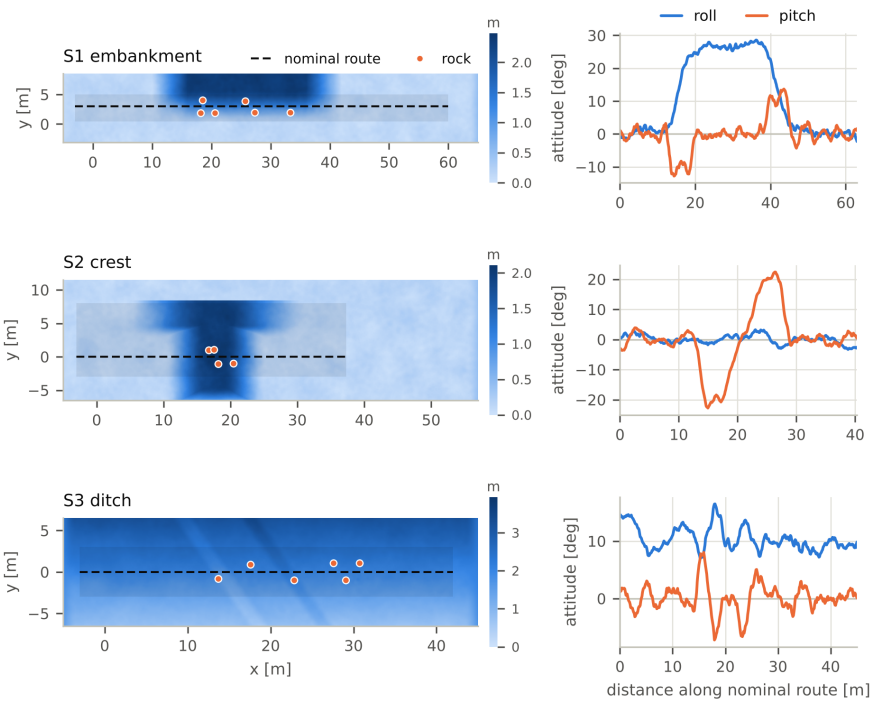


Figure 3: Scenarios S1–S3 (seed 0) with the nominal route (dashed) and the corridor of the route layer (shaded); colour: terrain height. Right: terrain roll and pitch along the nominal route.

Table 3: Compared controllers. P: preview from the LiDAR map; S: speed synchronisation (Algorithm 1) and lock (16); C: safety filter; L: active levelling; R: route layer; A: actuator-aware planning.

Controller	Manipulator	P	S	C	L	R	A
FIX	transport pose	–	–	–	–	–	–
SLOW	transport pose, speed rule	–	–	–	–	–	–
REACT	posture map of the measured attitude	–	–	–	✓	–	–
REACT-S	as REACT, speed rule	–	–	–	✓	–	–
PREV	preview plan, time-indexed	✓	–	–	✓	–	✓
COMPASS	full framework	✓	✓	✓	✓	✓	✓
COMPASS−P	no preview (present attitude extrapolated)	–	✓	✓	✓	–	✓
COMPASS−S	no synchronisation loop, sample check or lock	✓	–	✓	✓	✓	✓
COMPASS−C	no safety filter	✓	✓	–	✓	✓	✓
COMPASS−L	no levelling	✓	✓	✓	–	✓	✓
COMPASS−R	no route layer	✓	✓	✓	✓	–	✓
COMPASS−conf.	capacity without conformal margin	✓	✓	✓	✓	✓	✓
COMPASS−A	planner without actuator speed limits	✓	✓	✓	✓	✓	–

actuator constraints, CBF, routing, posture gains and spatial lead, and all execute the planned speed profile. They differ in two factors. The posture reference is indexed either by measured path progress (P) or by a compatible clock (T) that integrates the planned speed and resets to the measured progress at each 5 Hz replan. The commanded speed is either reduced by the lock of (16) (L) or not (N). P-L reproduces COMPASS: both give identical logged signals in a noisy verification run. P-N is COMPASS without the lock, and COMPASS−S in addition removes the iterative speed adjustment of Algorithm 1.

P-U is P-N with the planned speed multiplied by a constant factor f before the CBF, i.e. uniform slowing. The letters of the variants are unrelated to the component letters of Table 3. We call execution *speed-compatible* when the vehicle follows the planned speed profile, including its cap ahead of unmapped terrain, and the posture reference is read at the progress or the clock of that profile. All variants and COMPASS−S are speed-compatible, whereas PREV drives at the nominal speed.

5.3. Studies

Table 4 lists the studies; about 1,700 runs in total. The main study, the ablations, the stress envelope and the derated-drive study compare the controllers. The derated-drive study caps all manipulator and levelling drives at 30 % or 15 % of their configured speed, below the planning rates and without informing the planner. These caps are stress tests of a drive shortfall; Section 6.7 relates them to the documented drives. The model mismatch study runs COMPASS with the 400 kg payload undeclared to the controller, with a declared payload that is absent, with the assigned ground coefficient reduced from 0.60 to 0.45, with the lift, tilt and slew drives at half their maximum speed ($41\,\mathrm{mm\,s^{-1}}$, $0.26\,\mathrm{rad\,s^{-1}}$, still just above the planning rates), and with a slew drive that stops 12 s into the drive and holds its position, unknown to the controller. The *drive-fault* study repeats this slew fault with P-L and P-N (Section 5.2).

The execution studies use the variants of Section 5.2 on S1 and S2 with the 400 kg payload and, except for the plant-variation anchors and the drive-fault study, on seeds 8–23, which are separate from the seeds of the main study. The design and primary outcomes of every study condition were fixed before its outcomes were examined.

- The *matched execution* study crosses indexing (P/T) and locking (L/N) with configured drives and with 15 % rate caps.

- The *equal-time execution* study tests whether locking reduces the tracking error more than uniform slowing at the same achieved driving time; the tracking outcome is Λ_{95}, the 95th percentile of Λ over an episode (Section 5.4). With 15 % rate caps, it runs P-N, P-U with $f \in \{0.95, 0.891, 0.85, 0.80\}$ and P-L with $\Lambda_{\max} \in \{4, 2, 1, 0.5\}$ s, with $\Lambda_0 = 0.4$ s unchanged; $\Lambda_{\max}$ enters only the lock statement of (16). The factor $f = 1/1.1221 = 0.891$ matches the mean time cost of locking in the matched study. The threshold $\Lambda_{\max} = 0.5$ s commands the vehicle to stop whenever the tracking error exceeds about 0.5 s and to resume when the drives catch up; it is the stop-and-go limit of locking, which reconfigures at standstill. For each scenario and seed, the uniform-slowing frontier $\Lambda_{95}^{U}(T)$ interpolates piecewise-linearly through the completed P-N and P-U runs ordered by achieved driving time T. For a locking run with driving time T^L and tracking outcome Λ_{95}^{L}, the primary outcome is the equal-time contrast $\Lambda_{95}^{L} - \Lambda_{95}^{U}(T^L)$, evaluated without extrapolation.

- The *survey-map* study repeats P-L, P-N and P-U ($f = 0.891$) with the elevation map pre-filled at $t = 0$ with the true visible surface; scans and wheel contacts then update it as before. Without a survey map, the planner caps the speed at $0.3\,\mathrm{m\,s^{-1}}$ ahead of unobserved terrain, which slows every variant before the hidden crest descent and may mask the consequence of a rate shortfall.

- The *drive-envelope* study replaces the configured caps by envelopes derived from the documented motor ratings (Section 3.2). The short-time envelope allows the peak lift/tilt motor torque of 32 N m up to a corner speed and decreases linearly

to zero torque at 3,580 rpm; the corner, 3,276 rpm, places the largest force and the largest speed of the documented sizing simulations (129.9 kN and 45.3 mm s^{-1}) jointly on the curve, which is the most generous reading of the short-time capability. The sustained-duty envelope limits the lift/tilt motor power to the rated power $P_{\mathrm{N}} = 10\,\mathrm{N\,m} \times 2{,}000\,\mathrm{rpm} = 2.09\,\mathrm{kW}$, so that the stroke speed is at most $\min(47.7\,\mathrm{mm\,s^{-1}}, \eta P_{\mathrm{N}}/|F|)$; it is a constant-power model of continuous operation. At every 10 ms servo update, the set-point rate limit of each lift/tilt EMLA is set to the envelope speed at its actuator force, low-pass filtered with a 50 ms time constant; the controller keeps a static limit of 47.7 mm s^{-1} and the planner its planning rates. Because the slew can dominate the tracking error, a further rated-power condition also holds the slew to its documented rated power (5.4 N m at 3,000 rpm, 1.70 kW): the column rate is limited to $\eta P_{\mathrm{N}}/|T_{\mathrm{col}}|$ at the filtered column torque T_{col} while the drive is not braked. This bound is independent of the slew reduction ratio. The extension and the levelling keep their configured caps. A 30 % cap on all manipulator and levelling drives completes the conditions, and the 15 % caps of the equal-time study complete the comparison.

- The *travel-speed* study repeats P-L and P-N at a nominal speed of 0.35 m s^{-1} with the configured caps, the rated-power drives and the 30 % and 15 % caps, with the time limit raised from 200 to 460 s.

- The *drive-fault* study runs P-L and P-N on S1 and S2 with seeds 0–3 under the slew fault of the model-mismatch study: the slew stops 12 s into the drive and holds its position, while all other drives keep their configured caps. The P-L runs repeat the COMPASS episodes of the model-mismatch study.

The independent audit of Section 3.4 records every run of the studies marked in Table 4. The *mechanics set* re-runs 32 main-study episodes with the recorder: FIX, PREV and COMPASS on S2 with seeds 0–7, and FIX and COMPASS on S1 and S3 with seeds 0 and 6, all with the payload. Recording is observational: every re-run reproduces the native 20 Hz signals of its main-study episode exactly. When the 45° threshold ends an episode, the closed-loop controller continues for five seconds to observe the subsequent contact and motion; these continuations are analysed separately. The *crest robustness* study re-simulates the 16 FIX and COMPASS crest episodes of the mechanics set four times: unchanged; with the tyre radial stiffness reduced to $0.33 \times 10^6\,\mathrm{N\,m^{-1}}$, the alternative inference recorded in the machine documentation, or raised to $1.8 \times 10^6\,\mathrm{N\,m^{-1}}$, the upper bound of the plant-variation range (both with damping ratio 0.1); and with a 0.5 ms physics step, for which the drive and servo updates stay at 10 ms, the controller at 50 Hz, and the soil–height-field synchronisation runs every 50 instead of 100 ms. The counterweight result is called robust under a treatment if at least seven of eight seed pairs favour COMPASS at every conditioning and COMPASS has less negative-moment exposure and fewer edge-rotation events; this criterion was fixed before the re-simulations. Because COMPASS crosses the crests more slowly than FIX, the same treatments are also applied to PREV, which uses the same planner but drives at the nominal speed of FIX; PREV is paired with a second set of re-runs of the FIX episodes and judged by the same criterion.

The *plant-variation* study uses 16 Latin-hypercube draws over 13 parameters of the slew drive and brake, the tyres, the soil and the ground friction, a nominal row and two combined corners; Appendix C gives the ranges, which include slew reduction ratios of

Table 4: Simulation studies. Drives: configured caps unless stated; caps are percentages of the configured manipulator and levelling rate caps. Audit: runs recorded with the independent 50 Hz contact-wrench audit.

Study	Controllers / variants	Scenarios, payload, seeds	Drives	Runs	Audit	Sect.
Allowance identification	FIX at 0.3, 0.55, 0.8 $\mathrm{m\,s^{-1}}$	S1–S3; 0, 400 kg; 100–101		36		5.5
Main	six controllers (Table 3)	S1–S3; 0, 400 kg; 0–7		288		6.2
Ablations	six ablations	S1–S3; 400 kg; 0–7		144		6.8
Derated drives	PREV, COMPASS−S, COMPASS	S1–S3; 400 kg; 0–7	30, 15 %	144		6.5
Stress envelope	FIX, REACT-S, COMPASS, −R, −conf.	S4 26–38°; 0, 400 kg; 0–1		80		6.8
Components	COMPASS, −R, −conf.	S4 34, 38°; 0, 400 kg; 8–15		96	✓	6.8
Model mismatch	COMPASS	S1, S2; 0, 400 kg; 0–3	5 settings	40		6.9
Mechanics set	FIX, PREV, COMPASS	S2 (0–7), S1, S3 (0, 6); 400 kg		32	✓	6.3
Crest robustness	FIX (two sets), COMPASS, PREV	S2; 400 kg; 0–7	4 treatments	128	✓	6.4
Matched execution	P-L, P-N, T-L, T-N	S1, S2; 400 kg; 8–15	100, 15 %	128	✓	6.6.1
Equal-time execution	P-N; P-U, 4 factors; P-L, 4 thresholds	S1, S2; 400 kg; 16–23	15 %	144	✓	6.6.2
Survey map	P-L, P-N, P-U	S1, S2; 400 kg; 16–23	15 %	48		6.6.2
Drive envelopes	P-L, P-N	S1, S2; 400 kg; 16–23	5 conditions	160		6.7
Travel speed	P-L, P-N at 0.35 $\mathrm{m\,s^{-1}}$	S1, S2; 400 kg; 16–23	4 conditions	128		6.7
Plant variation	P-L, P-N	S1, S2; 400 kg; 8–15, 6	100, 15 %	100	✓	6.9
Drive fault	P-L, P-N	S1, S2; 400 kg; 0–3	slew blocked	16		6.9

450–750 and tyre stiffnesses of 0.6–1.8 $\mathrm{MN\,m^{-1}}$, and the compositions of the corners. Two *anchors*, settings taken from sources rather than sampled, use seed 6. The first uses the manufacturer's data for the slew motor with brake, a rotor inertia of 0.00051 $\mathrm{kg\,m^2}$ and a holding torque of approximately 13 N m [47], i.e. 6.63 kN m at the column instead of the nominal 17.6 kN m, with the brake delays of a variant datasheet of the motor family; the second anchor adds the 0.33 $\mathrm{MN\,m^{-1}}$ tyre. Every setting is paired between P-L and P-N on S1 and S2, and both corners and both anchors are repeated with 15 % rate caps. The controller keeps its nominal capacity and effort models throughout.

5.4. Outcomes and statistics

The evaluation uses simulator states that the controllers cannot access. The static reserve uses (4) with gravity. The acceleration-based proxy θ_{acc} uses the whole-body centre-of-mass acceleration, differentiated at 200 Hz, transformed to hull axes and low-pass filtered with a 0.05 s time constant; it is logged at 20 Hz. Unlike the controller's proxy (Section 4.5), it is free of sensing-point and estimation errors. Three outcomes are distinguished:

- a *proxy event* is a negative θ_{acc} for at least 0.1 s after $t \geq 3$ s;
- a *45° termination* ends the episode when |roll| or |pitch| exceeds 45°;

- a *wheel-unloading event* is a minimum off-edge load fraction $\min_e \eta_e < 0.02$ for at least 0.1 s, where η_e sums the reactions of the two wheels opposite edge e along hull z and divides by the total wheel load, averaged over each 50 ms logging interval.

Wheel unloading also occurs during bridging, impacts or beam end-stop contact; the audit evaluates the full wrench of (5), including payload contact. Completion means reaching the route end within the time limit (200 s; 460 s at $0.35\,\mathrm{m\,s^{-1}}$); a nonprogress stop means less than 0.3 m of additional best progress for more than 25 s. The audit outcomes are $S_{1\%}$, the exposure to a negative minimum restoring moment and to degenerate support (shares of the evaluated time), and edge-rotation events (Section 3.4).

The tracking outcome is Λ_{95}, the 95th percentile of the rate-normalised tracking error over an episode. Because Λ is not a delay, a time-shift delay is also computed from the 20 Hz logs: for each balancing coordinate j with a moving reference, τ_j^* minimises $\mathrm{RMS}_t\,[x_j(t) - x_{\mathrm{ref},j}(t-\tau)]$ over $\tau \in [0, 8]\,\mathrm{s}$, using samples with $t \geq 3\,\mathrm{s}$ and $t - \tau \geq 3\,\mathrm{s}$, and τ^* is the maximum over the coordinates. A windowed variant, the 95th percentile of τ^* over 10 s windows with 50 % overlap, is reported as well; single windows can return optima of several seconds, so the whole-episode τ^* is the primary delay measure.

All primary metrics exclude the first 3 s and any continuation after a 45° termination. In the main study, bootstrap 95 % intervals describe episode variation, paired Wilcoxon tests compare continuous metrics and exact McNemar tests compare the composite endpoint of a proxy event or a 45° termination. Simulation frames are not independent units: execution effects are paired by scenario and seed, and their percentile intervals (10,000 bootstrap draws) resample seed blocks with both scenarios kept together; plant-variation summaries resample parameter rows. Exact sign tests count seed pairs.

5.5. Identification of the dynamic allowance

The coefficients of (14) were identified once, before all other studies, from FIX runs at 0.3, 0.55 and $0.8\,\mathrm{m\,s^{-1}}$ on S1–S3 with both payloads, on seeds 100 and 101, which are not used in the evaluation. For each logged sample, the loss is the difference between the static reserve and the acceleration-based proxy, and the excitation is computed from the simulated terrain with the planner's support predictor along the driven path. A linear 95 % quantile regression with non-negative coefficients gives $k_\kappa = 1.88\,^\circ\,\mathrm{s^2\,rad^{-1}}$, $k_v = 12.0\,^\circ\,\mathrm{s\,m^{-1}}$ and $k_0 = 0$, with an empirical coverage of 95.6 % on all 75,267 samples (93.2–97.3 % per speed; 95th percentile of the loss 6.4°, largest loss 18.9°). The excitation term is small: for the excitations met in these runs (99th percentile $0.17\,\mathrm{rad\,m^{-2}}$) it adds at most 0.2° at the nominal speed. At the nominal speed the allowance is 9.6°, so the planner requires a static reserve of about 17.6°; at the creep speed it requires 9.8°.

6. Results

6.1. Balancing capacity

Figure 4 compares the static reserve of the transport pose with the calibrated lower estimate of the counterweight capacity over roll ±40° and pitch ±35°. The transport pose exceeds $\theta_{\min} = 8°$ on 55 % of this attitude box empty and on 52 % with the 400 kg payload; $\hat{B} - \varepsilon_{0.05}$ exceeds it on 96 % and 81 %, respectively (geometric fractions of the sampled box). The posture map raises the static reserve by 11.7° on average (maximum 19.5°) empty and by 10.8° (maximum 27.6°) loaded, principally by moving mass fore or aft on steep pitch. With the payload, slew holdability restricts the postures on steep cross slopes: no admissible counterweight posture exists in 13 % of this attitude box (17 % of

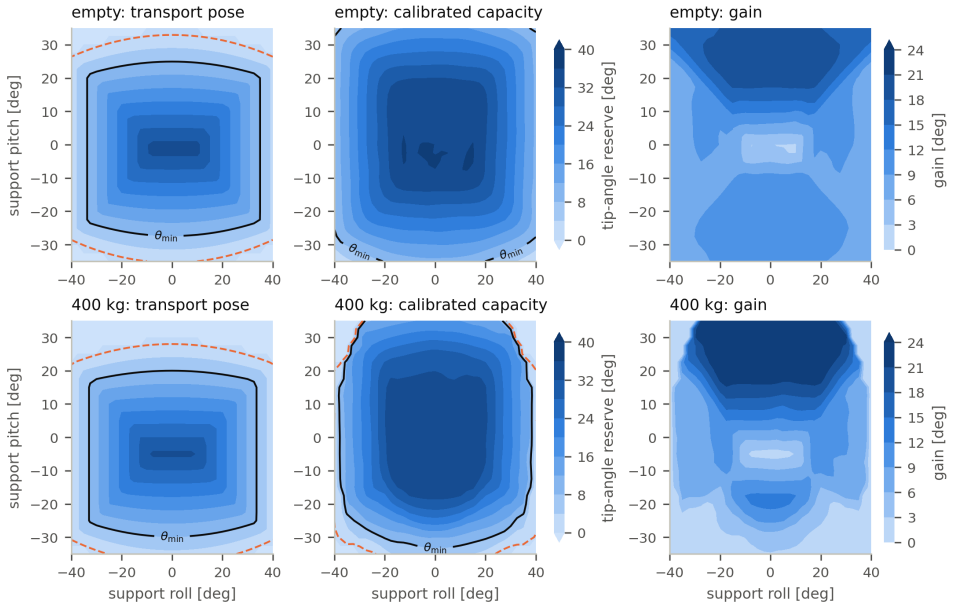


Figure 4: Static reserve of the transport pose (left), calibrated lower capacity estimate $\hat{B} - \varepsilon_{0.05}$ (centre) and their difference (right), empty (top) and with the 400 kg payload (bottom), at zero beam twist. Positive pitch is nose down. Black: $\theta = \theta_{\min}$; dashed: zero static reserve.

the training grid, which also covers twist), and the controller then falls back to the braked transport pose and requests platform intervention.

6.2. Main comparison

Table 5 and Fig. 5 summarise the 288 episodes of the main study; Table D.1 gives the breakdown by scenario. COMPASS completed all 48 episodes without a proxy event, a 45° termination, a payload contact or a stop. Its smallest static reserve was 14.2°, and the episode-wise 1st percentile of the acceleration-based reserve averaged 15.9°. Its mean driving time was 67.6 s, 8 % longer than that of FIX (62.8 s); the relative increase was largest on the crests (13–15 %), which it crossed more slowly.

FIX crossed the ditch without a proxy event but had one on the embankment. On the crest it reached the 45° termination twice and had 18 proxy events in six of 16 episodes, five of them with the payload. SLOW reached the 45° termination in seven crest episodes and took 18 % longer than FIX: its response to the measured static reserve began near the descending edge, where braking further unloaded the rear wheels. REACT and REACT-S raised the mean reserves but had proxy events in two and three episodes, respectively; their terrain-unaware postures also caused 46 and 48 payload contacts and one stop each. PREV differed from COMPASS mainly on the crest (Section 6.5). All controllers had wheel-unloading events (3–13 per episode on average), including on rocks and at beam end stops.

Over matched episodes, the composite endpoint of a proxy event or a 45° termination occurred in none of the COMPASS episodes, in seven FIX episodes (exact McNemar $p = 0.016$) and in nine SLOW episodes ($p = 0.004$). REACT and REACT-S had two and three affected episodes ($p \geq 0.25$), and COMPASS had higher reserves than both (Wilcoxon $p = 0.001$ for the smallest static reserve and $p < 10^{-6}$ for the smallest acceleration-based reserve). PREV and COMPASS had similar pooled reserves ($p = 0.08$ and $p = 0.42$; mean smallest static reserves 21.4° and 21.7°).

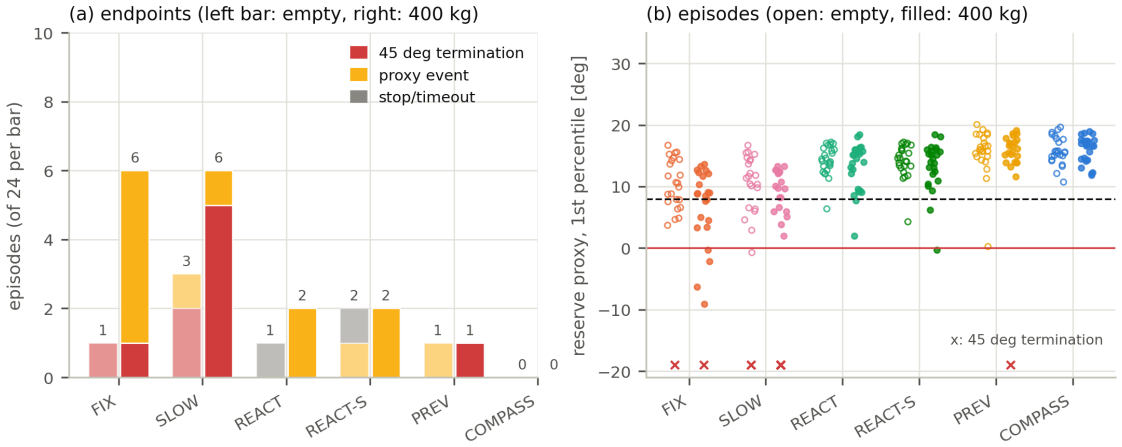


Figure 5: Main study, S1–S3 pooled. (a) Outcome counts by controller and payload (24 episodes per bar): 45° termination, proxy events without 45° termination, and stops. (b) Episode-wise 1st percentile of the acceleration-based reserve proxy; crosses: 45° termination; dashed: $\theta_{\min}$.

Table 5: Main comparison over S1–S3 (eight seeds each, 24 episodes per row). Compl.: reached the goal; 45°: termination when roll or pitch exceeds 45°; Stop: nonprogress stop or time limit; Proxy events: mean number per episode (episodes with at least one); Unloading: mean number of wheel-unloading events per episode; θ_{acc}: mean over episodes of the 1st percentile of the acceleration-based reserve proxy [deg]; θ_{st}: mean over episodes of the smallest static reserve [deg]; Time: mean driving time of completed episodes [s]; Slew sat.: share of time with the slew drive at its torque limit [%]; Λ_{95}: 95th percentile of the rate-normalised tracking error [s].

Controller	Compl.	45°	Stop	Proxy events	Unloading	θ_{acc}	θ_{st}	Time	Slew sat.	Λ_{95}
empty manipulator										
FIX	23/24	1	0	0.08 (1)	9.2	10.0	14.8	62.8	0.6	–
SLOW	22/24	2	0	0.33 (3)	8.6	9.3	13.5	71.5	0.6	–
REACT	23/24	0	1	0.00 (0)	12.5	14.2	20.9	61.7	6.8	–
REACT-S	23/24	0	1	0.04 (1)	13.2	14.1	20.9	61.8	6.8	–
PREV	24/24	0	0	0.17 (1)	8.0	15.6	21.9	65.1	4.3	0.24
COMPASS	24/24	0	0	0.00 (0)	7.6	15.8	21.9	67.9	0.9	0.44
400 kg *payload*										
FIX	23/24	1	0	0.71 (6)	7.5	6.3	9.4	62.8	0.9	–
SLOW	19/24	5	0	0.38 (6)	6.9	5.3	6.8	77.7	0.9	–
REACT	24/24	0	0	0.17 (2)	5.4	13.1	18.6	62.2	6.0	–
REACT-S	24/24	0	0	0.21 (2)	5.3	13.2	18.6	62.8	5.9	–
PREV	23/24	1	0	0.12 (1)	2.8	15.0	20.8	64.4	2.5	0.21
COMPASS	24/24	0	0	0.00 (0)	3.1	16.1	21.6	67.4	2.0	0.32

6.3. Independent mechanical audit

In the 32 runs of the mechanics set, the independent force and moment accounting closes with RMS residuals of 5.82×10^{-6} body weights (BW) and 6.19×10^{-6} BW m (maxima 5.35×10^{-5} BW and 7.51×10^{-5} BW m). In the 404 audited runs of the mechanics, matched-execution, plant-variation and component studies and of the P-L, P-N and P-U ($f = 0.891$) runs of the equal-time study, the residuals stay at most 1.32×10^{-4} BW and 2.02×10^{-4} BW m, and no run with recorded solver diagnostics reported a numerical warning. Of the 1,852.82 s evaluated in the mechanics set, 1,765.86 s have resolved, nondegenerate wheel support, 86.64 s degenerate support and 0.32 s additional body or payload

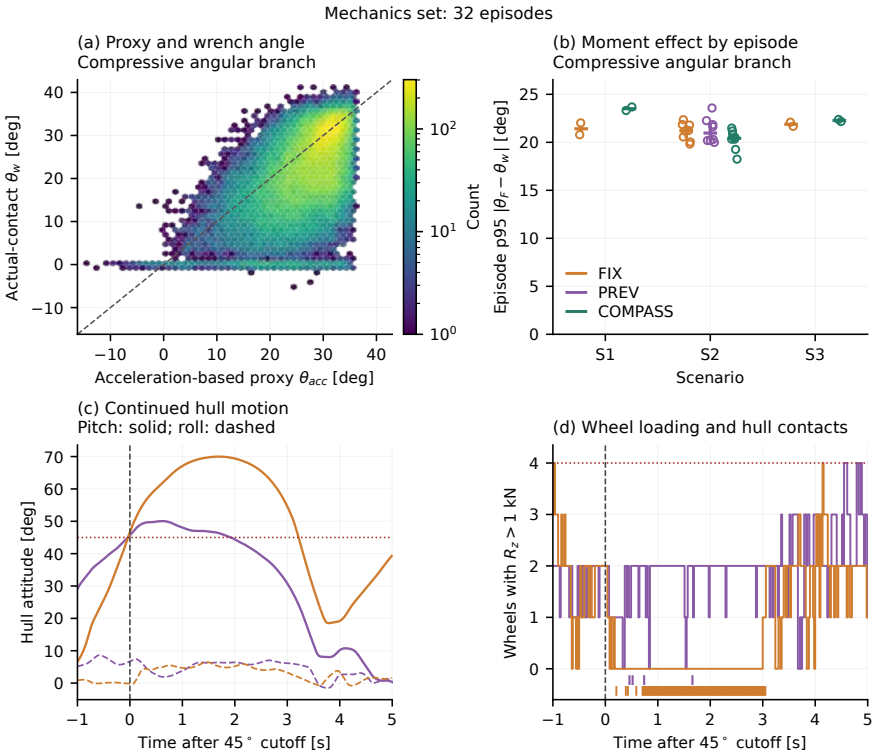


Figure 6: Independent audit of the 32-run mechanics set. (a) Acceleration-based proxy against the actual-contact moment-inclusive angle on compressive, nondegenerate wheel support; logarithmic sample counts; native 20 Hz proxy samples are paired with the nearest 50 Hz audit states. (b) Episode-wise 95th-percentile moment contribution on identical force and contact geometry; bars show the median and interquartile range. Branch conditioning excludes 3 of 35,343 matched samples in (a) and 7 of 88,293 audit samples in (b); all adverse moments and support strata remain in the full analysis. (c, d) Five-second continuations after the two 45° terminations: pitch solid, roll dashed; the loaded-wheel count uses an upward reaction of 1 kN, and rugs mark hull contact.

support.

On the compressive, nondegenerate wheel-support samples, the acceleration-based proxy differs from the wrench angle θ_w by 8.84° RMS, with an absolute 95th-percentile difference of 18.15° (Fig. 6a). Holding the actual geometry and the instantaneous force fixed, omitting the moment changes the angle by 9.18° RMS and 21.42° at the 95th percentile (Fig. 6b), so angular-momentum-rate effects are not negligible in these traversals. Proxy and wrench-angle signs agree during 97.75 % of the nondegenerate wheel-supported time, but this agreement is dominated by jointly nonnegative samples: the proxy is negative during only 3.38 s of the 42.92 s with a negative wrench angle (7.88 %). A negative minimum restoring moment occupies 42.92 s with the full moment and 169.24 s when the moment is omitted. The reserve that the controller actually uses, computed from the hull specific force and encoder estimates, differs from the wrench angle by 9.25° RMS and 19.31° at the 95th percentile on the wheel-only, nondegenerate samples of the 356 audited runs of the mechanics, matched-execution, plant-variation and component studies; unlike the compressive-branch figure above, this comparison includes all angular branches.

The edge-rotation definition identifies 72 events in 22 of the 32 episodes (54–77 events in 19–22 episodes when the opposite-wheel threshold is varied over 1, 2 and 5 % and the angular-velocity threshold over 0.01, 0.02 and 0.04 rad s^{-1}). On the crest, FIX has 47 events in all eight episodes, PREV five in two and COMPASS five in five; the proxy flags

Table 6: Independent audit of the counterweight (mechanics set, 400 kg, configured drive rates). $S_{1\%}$: first percentile of the minimum signed wheel-edge restoring moment on resolved, nondegenerate, wheel-supported samples at 10 µm conditioning [10^{-3} BW m] (range over episodes). Exposures are percentages of the evaluated time. ER: edge-rotation events (episodes with at least one). PE: proxy events (episodes).

Scen.	Controller	n	$S_{1\%}$ range	ER	Neg. moment [%]	Degenerate [%]	PE
S2	FIX	8	−7.22 to −5.39	47 (8)	3.98	6.74	16 (5)
S2	PREV	8	−4.31 to −0.01	5 (2)	1.42	3.09	3 (1)
S2	COMPASS	8	−2.53 to 83.85	5 (5)	1.20	2.49	0 (0)
S3	FIX	2	−6.55 to −6.24	8 (2)	3.33	7.95	0 (0)
S3	COMPASS	2	−3.76 to −3.60	2 (2)	2.18	4.15	0 (0)
S1	FIX	2	−4.43 to −1.50	4 (2)	2.75	7.76	0 (0)
S1	COMPASS	2	−5.53 to −2.69	1 (1)	2.78	5.27	0 (0)

none of the COMPASS events. Total upward wheel load stays at least 0.384 BW during the 12.24 s in which the event conditions hold, so these events occur under substantial wheel support. The contact conditioning changes absolute exposures considerably: from 1 to 10 to 100 µm, degenerate exposure grows from 72.36 to 86.64 to 97.86 s, and negative-moment exposure falls from 57.16 to 42.92 to 31.98 s. Absolute edge diagnostics are thus reported with their conditioning, and Section 6.4 tests the counterweight contrast across conditionings.

The audit gives a proxy-independent view of the counterweight (Table 6). On the crest, every COMPASS episode has a higher first-percentile restoring moment $S_{1\%}$ than every FIX episode, so all eight seed pairs favour COMPASS (exact sign test $p = 0.0078$); paired by seed, COMPASS exceeds FIX by a mean of 16.0×10^{-3} BW m (seed-block interval 4.7–37.4; median 5.7) at 10 µm. COMPASS also has fewer edge-rotation events (5 versus 47), lower negative-moment exposure (1.20 versus 3.98 % of the evaluated time) and less degenerate wheel support (2.49 versus 6.74 %). PREV, which uses the same planner but drives at the nominal speed, also separates from FIX ($S_{1\%}$ from −4.31 to −0.01, 5 events) and has exposures similar to those of COMPASS apart from its seed-6 termination; Section 6.4 tests this speed-matched contrast as well. On the ditch, COMPASS has the higher $S_{1\%}$ in both seeds. On the embankment, it has fewer edge-rotation events (1 against 4) and less degenerate support (5.27 against 7.76 %), with $S_{1\%}$ and negative-moment exposure close to those of FIX (Table 6); the crest is where the counterweight separates the controllers. In the six-dimensional friction-cone calculation, the outer approximation is feasible at all 197 grid snapshots and the inner one at 178; where both are feasible, the median minimum roll/pitch torque capacity is 0.523/0.532 BW m (inner/outer), a narrow bracket. Five-second continuations follow the two terminated episodes (Fig. 6c, d): PREV reaches 50.07° and returns to a nearly level attitude, whereas FIX reaches 69.99° with 2.44 s of hull contact.

6.4. Robustness of the crest counterweight result

Table 7 tests whether the crest contrast between COMPASS and the transport pose depends on the contact conditioning, the tyre stiffness or the integration step. In the reference records, every COMPASS episode has a higher $S_{1\%}$ than every FIX episode at 1, 10 and 100 µm, with 3.1–3.7 times less negative-moment exposure; the conditioning shifts both controllers in the same direction. Two audit outcomes do not use the contact envelope: the edge-rotation events (5 versus 47) and the friction-cone capacity, for which

COMPASS has the larger median minimum roll/pitch torque capacity in all eight seed pairs under both approximations (0.60–0.61 versus 0.43 BW m).

The unchanged re-simulation reproduces the reference ranges of $S_{1\%}$, the exposures, the events and the terminations. COMPASS exceeds FIX in all eight seed pairs at every conditioning (exact sign test $p = 0.0078$), by a mean of 13.5, 16.0 and 25.8×10^{-3} BW m at 1, 10 and 100 µm. With the tyre stiffness reduced to 0.33 or raised to 1.8 MN m^{-1}, and with the 0.5 ms step, the outcome is the same in all nine treatment–conditioning combinations: every COMPASS episode has a higher $S_{1\%}$ than every FIX episode, all intervals of the paired difference exclude zero, COMPASS has the lower negative-moment exposure in every seed pair (3.3–4.7 times lower pooled), FIX has 36–43 edge-rotation events against one or two for COMPASS, and only FIX episodes have proxy events or end with the 45° termination (one to three per treatment). The sign and the separation of the contrast hold under every treatment, and its size varies with the treatment: at 10 µm, the mean difference ranges from 6.63 (stiff tyre) to 56.27×10^{-3} BW m (soft tyre) and the median from 5.7 (reference) to 76.4×10^{-3} BW m (soft tyre). The force and moment residuals stay below 6.8×10^{-6} BW and BW m RMS in every run. The crest counterweight result meets the robustness criterion under every treatment. In the reference records, COMPASS is also higher on the ditch at all three conditionings (+2.36, +2.71 and $+2.84 \times 10^{-3}$ BW m); on the embankment, the two-seed mean difference is small (+0.29, −1.14 and -0.84×10^{-3} BW m), and COMPASS has fewer edge-rotation events and less degenerate support there.

The counterweight contrast is also present at equal speed. PREV uses the same planner, drives at the nominal speed and crosses the crest in about the same time as FIX (54.8–54.9 against 53.2–54.2 s across the treatments). Under all four treatments, every PREV episode has a higher $S_{1\%}$ than every FIX episode at every conditioning ($p = 0.0078$ in each of the twelve treatment–conditioning combinations), PREV has the lower negative-moment exposure in every seed pair, and it has 2–5 edge-rotation events against 36–47 for FIX; the contrast thus meets the robustness criterion with PREV in place of COMPASS (Table 8). The re-runs reproduce the $S_{1\%}$ ranges, exposures and events of the FIX records of Table 7 and of the PREV records of the mechanics set. At 10 µm, the mean PREV minus FIX difference is $4.7–19.2 \times 10^{-3}$ BW m against 6.6–56.3 for COMPASS, and the medians are 4.8–7.1 against 5.7–76.4 (5.7 for both in the reference records). The benefit is thus present without the slower crossing, and COMPASS, which adds the lock, the filter, the route layer and the speed profile, reaches the larger values, most clearly with the soft tyre and the 0.5 ms step. PREV still reaches the 45° termination on seed 6 under every treatment, which COMPASS avoids (Section 6.5); $S_{1\%}$ covers each episode up to its termination.

6.5. Speed-compatible execution

PREV uses the same preview, posture targets, planning programme and actuator limits as COMPASS, but executes a time-indexed reference at the nominal speed. Both controllers completed the 32 embankment and ditch episodes without proxy events, PREV 3–4 % faster. The crest of seed 6 (22.4° climb, 20.8° descent, 2.5 m top) exposes the speed mismatch (Figs. 7 and 8). Moving the payload rearward requires approximately 50° of slew, or 3.5 s at the planning slew rate. Because the descent is initially hidden, the speed profile drops to the unmapped-terrain cap of 0.3 m s^{-1}; COMPASS follows it and keeps at least 14.4° of static reserve. PREV instead drives onto the descent at the nominal speed with the payload still forward; the hull pitches over the edge, the speed rises to 1.7 m s^{-1}

Table 7: Robustness of the crest counterweight contrast (S2, 400 kg, configured rates, seeds 0–7; eight FIX/COMPASS seed pairs per treatment). Cond.: contact conditioning [µm]. $S_{1\%}$: range over episodes of the first-percentile minimum signed wheel-edge restoring moment [10^{-3} BW m]. C>F: seed pairs in which COMPASS has the higher $S_{1\%}$. $\Delta S_{1\%}$: mean COMPASS minus FIX with 95 % seed-block interval [10^{-3} BW m]. Neg., Deg.: negative-moment and degenerate-support exposure, FIX/COMPASS [% of evaluated time]. ER: edge-rotation events, FIX/COMPASS (independent of the conditioning).

Treatment	Cond.	$S_{1\%}$ FIX	$S_{1\%}$ COMPASS	C>F	$\Delta S_{1\%}$	Neg. [%]	Deg. [%]	ER
Reference	1	−8.02 to −5.92	−3.21 to 68.68	8/8	13.53 [4.04, 31.46]	5.02 / 1.63	5.70 / 2.05	
	10	−7.22 to −5.39	−2.53 to 83.85	8/8	15.95 [4.73, 37.36]	3.98 / 1.20	6.74 / 2.49	47 / 5
	100	−6.66 to −4.25	−1.38 to 95.87	8/8	25.81 [9.50, 49.21]	3.12 / 0.85	7.61 / 2.84	
Tyre 0.33 MN m^{-1}	1	−6.29 to −3.89	−0.83 to 83.34	8/8	18.41 [6.02, 39.68]	3.31 / 1.00	4.15 / 1.11	
	10	−5.98 to −3.56	−0.10 to 99.90	8/8	56.27 [26.62, 82.47]	2.61 / 0.70	4.85 / 1.41	36 / 1
	100	−5.46 to −3.12	6.40 to 109.49	8/8	82.23 [55.99, 104.58]	2.00 / 0.44	5.46 / 1.66	
Tyre 1.8 MN m^{-1}	1	−7.89 to −6.10	−3.18 to −0.11	8/8	5.28 [4.34, 6.21]	5.60 / 1.69	6.32 / 2.24	
	10	−7.22 to −5.72	−2.65 to 3.36	8/8	6.63 [5.26, 8.06]	4.41 / 1.23	7.51 / 2.70	42 / 1
	100	−6.81 to −5.29	−0.23 to 63.85	8/8	31.19 [16.12, 47.88]	3.47 / 0.78	8.46 / 3.15	
Step 0.5 ms	1	−7.67 to −6.08	−2.72 to 83.13	8/8	16.38 [5.31, 37.54]	4.58 / 1.38	5.72 / 2.15	
	10	−6.95 to −5.68	−1.28 to 95.83	8/8	26.51 [9.35, 50.40]	3.60 / 0.92	6.69 / 2.61	43 / 2
	100	−6.26 to −5.27	−0.51 to 98.41	8/8	48.37 [25.65, 71.73]	2.72 / 0.58	7.58 / 2.97	

Table 8: The counterweight at the nominal speed: PREV, which uses the COMPASS planner but drives time-indexed at the nominal speed of FIX, against a second set of FIX re-runs under the treatments of the crest robustness study (Table 7; S2, 400 kg, configured rates, seeds 0–7). Cond.: contact conditioning [µm]. $S_{1\%}$: range over episodes [10^{-3} BW m]. P>F: seed pairs in which PREV has the higher $S_{1\%}$. $\Delta S_{1\%}$: mean PREV minus FIX with 95 % seed-block interval. Neg., Deg.: negative-moment and degenerate-support exposure, FIX/PREV [% of evaluated time]. ER: edge-rotation events and 45°: terminations, FIX/PREV; $S_{1\%}$ is evaluated up to a termination. Mean driving times (FIX/PREV): Reference 53.3/54.9 s; Tyre 0.33 MN m^{-1} 54.2/54.8 s; Tyre 1.8 MN m^{-1} 53.8/54.9 s; Step 0.5 ms 53.2/54.9 s.

Treatment	Cond.	$S_{1\%}$ FIX	$S_{1\%}$ PREV	P>F	$\Delta S_{1\%}$	Neg. [%]	Deg. [%]	ER	45°
Reference	1	−8.02 to −5.92	−5.08 to −1.08	8/8	4.34 [3.74, 4.93]	5.02 / 1.98	5.70 / 2.52		
	10	−7.22 to −5.39	−4.31 to −0.01	8/8	5.10 [4.21, 5.87]	3.98 / 1.42	6.74 / 3.09	47 / 5	1 / 1
	100	−6.66 to −4.25	−3.45 to 37.64	8/8	12.92 [5.06, 23.45]	3.12 / 1.00	7.61 / 3.52		
Tyre 0.33 MN m^{-1}	1	−6.29 to −3.89	−2.48 to 3.09	8/8	4.49 [3.30, 5.64]	3.31 / 1.41	4.15 / 1.52		
	10	−5.98 to −3.56	−0.83 to 72.88	8/8	19.23 [6.09, 38.23]	2.61 / 0.97	4.85 / 1.96	36 / 2	3 / 1
	100	−5.46 to −3.12	−0.45 to 97.51	8/8	43.27 [17.95, 71.74]	2.00 / 0.64	5.46 / 2.29		
Tyre 1.8 MN m^{-1}	1	−7.89 to −6.10	−5.49 to −2.01	8/8	3.97 [3.20, 4.52]	5.60 / 2.12	6.32 / 2.91		
	10	−7.22 to −5.72	−5.17 to −0.09	8/8	4.68 [3.55, 5.65]	4.41 / 1.61	7.51 / 3.41	42 / 4	2 / 1
	100	−6.81 to −5.29	−3.64 to 0.14	8/8	5.19 [4.20, 6.04]	3.47 / 1.15	8.46 / 3.88		
Step 0.5 ms	1	−7.67 to −6.08	−3.50 to 0.04	8/8	5.62 [5.12, 6.02]	4.58 / 1.44	5.72 / 2.03		
	10	−6.95 to −5.68	−0.80 to 37.59	8/8	14.35 [6.90, 24.24]	3.60 / 0.98	6.69 / 2.49	43 / 3	1 / 1
	100	−6.26 to −5.27	−0.11 to 75.88	8/8	37.88 [20.64, 55.37]	2.72 / 0.65	7.58 / 2.82		

and the episode ends with the 45° termination. Its static reserve falls to −11.5°, although its plan predicted approximately 20°. Without the payload, the same seed gives PREV four proxy events. A counterweight plan executed at a speed incompatible with it can fail.

COMPASS−S is the closer comparison: it keeps path indexing, the speed profile and the filter, and removes the iterative speed adjustment and the lock. With configured drives these two mechanisms act briefly: the iterative loop reduces the speed below the unmapped-terrain cap for 1.4 s over all 48 episodes, and the lock acts during 6 % of the driving time; together they add 0.3 s to the mean loaded driving time (67.4 against 67.1 s for COMPASS−S).

In Table 9, the manipulator and levelling drives are derated without informing the planner. At 30 % of the configured caps, COMPASS and COMPASS−S complete all

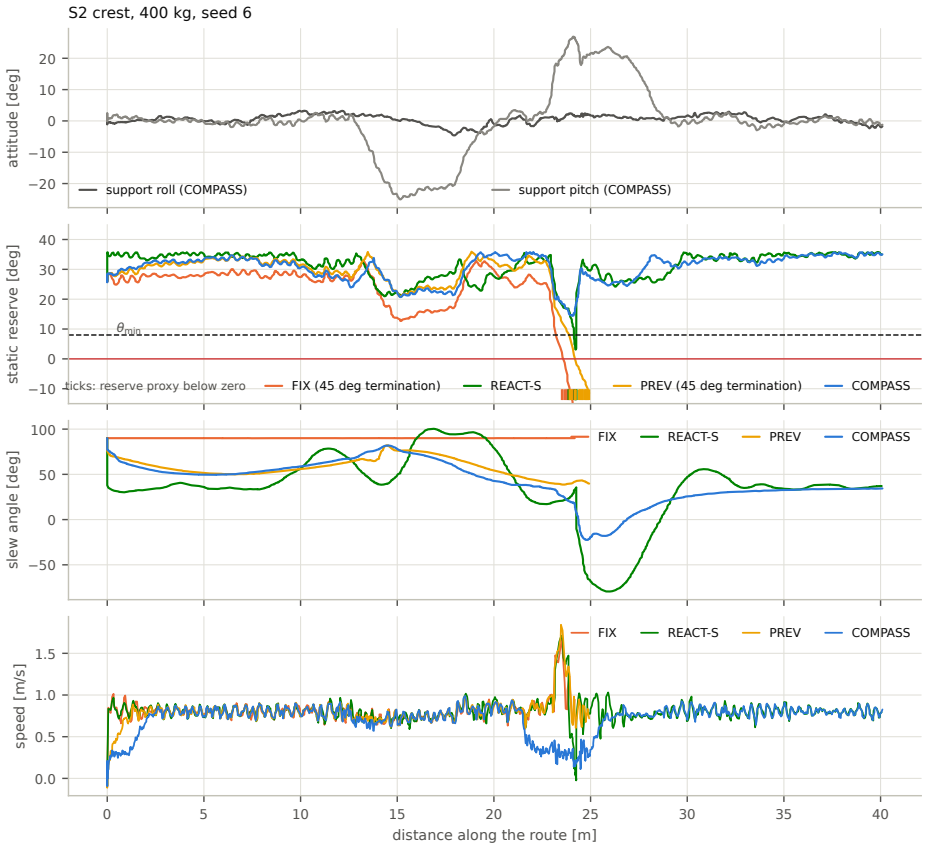


Figure 7: Crest S2, seed 6, 400 kg: support attitude in the COMPASS episode, static reserve (ticks: acceleration-based proxy below zero), slew angle (90°: payload ahead) and speed for FIX, REACT-S, PREV and COMPASS. FIX and PREV reach the 45° termination on the descent. Positive pitch is nose down.

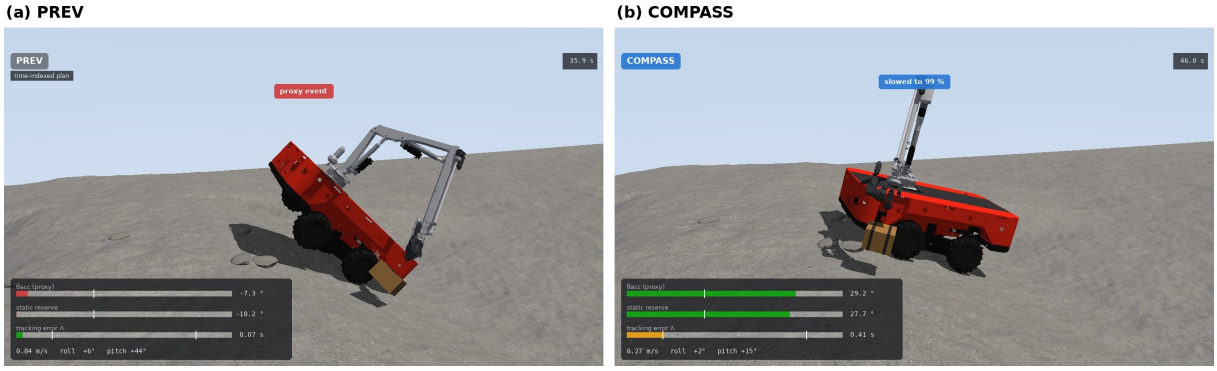


Figure 8: Crest S2, seed 6, 400 kg, in the supplementary video. (a) PREV at $t = 35.9$ s, just before the 45° threshold at $t = 36.0$ s. (b) COMPASS at $t = 46.0$ s near the same position on the route; it completes the route. Video times: 11.9 and 16.1 s. The head-up display shows the acceleration-based proxy θ_{acc}, the static reserve and the tracking error Λ.

24 episodes without proxy events, whereas PREV again reaches the 45° termination on seed 6. At 15 %, PREV reaches the same termination and makes 32 payload contacts in three episodes; its seven completed crest episodes average 13.0° for the 1st-percentile acceleration-based reserve, against 16.2° for COMPASS (paired Wilcoxon signed-rank test over all eight crest seeds, $p = 0.008$). COMPASS−S still completes every episode without a proxy event or termination. Its mean Λ_{95} is 0.87 s against 0.73 s for COMPASS ($p < 0.001$), and its crest reserve is 14.5° ($p = 0.016$ against COMPASS). Under this

Table 9: Derated drives: lift, tilt, extension, slew and levelling drives at their configured speed and at 30 % and 15 % of it (below the planning rates, unknown to the planner); 400 kg, S1–S3, seeds 0–7 (24 episodes per row; the configured rows are episodes of the main study and the ablations). Λ_{95}: mean over episodes [s], for PREV relative to its time-indexed reference. Gap: 95th percentile of the difference between planned and executed static reserve [deg]; other columns as in Table 5.

Drives	Controller	Compl.	45°	Stop	Proxy events	θ_{acc}	θ_{st}	Time	Λ_{95}	Gap
configured	PREV	23/24	1	0	0.12 (1)	15.0	20.8	64.4	0.21	3.5
	COMPASS−S	24/24	0	0	0.00 (0)	16.1	21.5	67.1	0.32	4.0
	COMPASS	24/24	0	0	0.00 (0)	16.1	21.6	67.4	0.32	4.0
30 %	PREV	23/24	1	0	0.12 (1)	14.9	20.9	64.3	0.23	3.5
	COMPASS−S	24/24	0	0	0.00 (0)	16.0	21.6	67.1	0.53	4.0
	COMPASS	24/24	0	0	0.00 (0)	15.8	21.4	68.2	0.46	4.2
15 %	PREV	23/24	1	0	0.12 (1)	14.7	19.8	64.4	0.35	3.5
	COMPASS−S	24/24	0	0	0.00 (0)	15.9	21.6	66.9	0.87	3.6
	COMPASS	24/24	0	0	0.00 (0)	16.5	21.9	72.5	0.73	3.8

derating, COMPASS needs 72.5 s on average against 66.9 s for COMPASS−S. On seed 6, the time-indexed execution thus failed at every rate setting, whereas speed-compatible execution completed the crest with or without the iterative speed adjustment and the lock; the ablations without the filter or the route layer (Section 6.8) also completed it. The two mechanisms improve tracking and crest reserve under slow drives for 8 % more driving time, and Section 6.6 isolates the lock.

6.6. Effect of the lock

6.6.1. Matched comparison

All 128 matched-execution episodes complete without a proxy event or a 45° termination (Table 10; Fig. 9 shows the paired effects). Because every variant executes the planned speed profile, this comparison separates the lock from the speed mismatch of PREV and from the iterative speed adjustment that COMPASS−S also removes.

With configured drives, P-L reduces Λ_{95} relative to P-N by 0.0162 s (0.00145–0.0322 s) at 0.464 % more driving time (0.022–0.884 %) and changes $S_{1\%}$ by -0.353×10^{-3} BW m (-0.879 to 0.057). With 15 % rate caps, P-L reduces Λ_{95} by 0.222 s (0.197–0.248 s) at 12.21 % more driving time (10.96–13.46 %). It raises $S_{1\%}$ by 3.77×10^{-3} BW m (1.03–8.79) at 10 µm, with positive point estimates at 1 and 100 µm (1.13 and 7.86×10^{-3} BW m), raises $C_{1\%}$ by 0.257° (0.069–0.607°) and lowers the degenerate-support exposure by 0.353 s (0.141–0.561 s); the edge-event difference is −0.188 per episode (−0.438 to 0.063). This moment effect depends on the tyre stiffness (Section 6.9). Clock-indexed locking gives similar tracking and time effects (edge-event difference −0.438, −0.875 to −0.063), and path and clock indexing give comparable restoring moments (P minus T with locking at 15 %: -0.809×10^{-3} BW m, −2.219 to 0.234). Because the moment gain at 15 % comes with 12 % more driving time, Section 6.6.2 asks whether uniform slowing with the same time gives the same gain.

6.6.2. Equal-time comparison with uniform slowing

In the equal-time study (Table 11), every variant completes 15 of 16 episodes; S1 seed 16 ends in a nonprogress stop for every variant, and no run reaches the 45° termination.

Table 10: Matched execution on S1 and S2, eight seeds per scenario, 400 kg. Rate: percentage of the configured drive-rate caps. P/T: measured-path/compatible-clock indexing; L/N: lock/no lock. Done: completed/available episodes; 45°, PE and ER: episodes with a 45° termination, a proxy event and an edge-rotation event. Episode means: $S_{1\%}$ [10^{-3} BW m] and $C_{1\%}$ [deg] are the first percentiles of the minimum signed restoring moment and of θ_w on resolved, nondegenerate, wheel-supported samples (the angular branch can wrap); Λ_{95} [s]; T: driving time of completions [s]; D: degenerate-support exposure [s].

Rate [%]	Variant	Done	45°	PE	ER	$S_{1\%}$	$C_{1\%}$	Λ_{95}	T	D
100	P-L	16/16	0	0	8	−2.19	−0.15	0.40	72.4	3.3
100	P-N	16/16	0	0	8	−1.84	−0.14	0.42	72.1	3.2
100	T-L	16/16	0	0	7	−1.76	−0.13	0.40	72.3	3.3
100	T-N	16/16	0	0	8	−1.75	−0.13	0.40	71.9	3.3
15	P-L	16/16	0	0	5	1.56	0.11	0.93	79.4	2.6
15	P-N	16/16	0	0	9	−2.21	−0.15	1.15	70.7	2.9
15	T-L	16/16	0	0	4	2.37	0.14	0.94	79.8	2.6
15	T-N	16/16	0	0	8	−1.88	−0.13	1.15	71.1	3.2

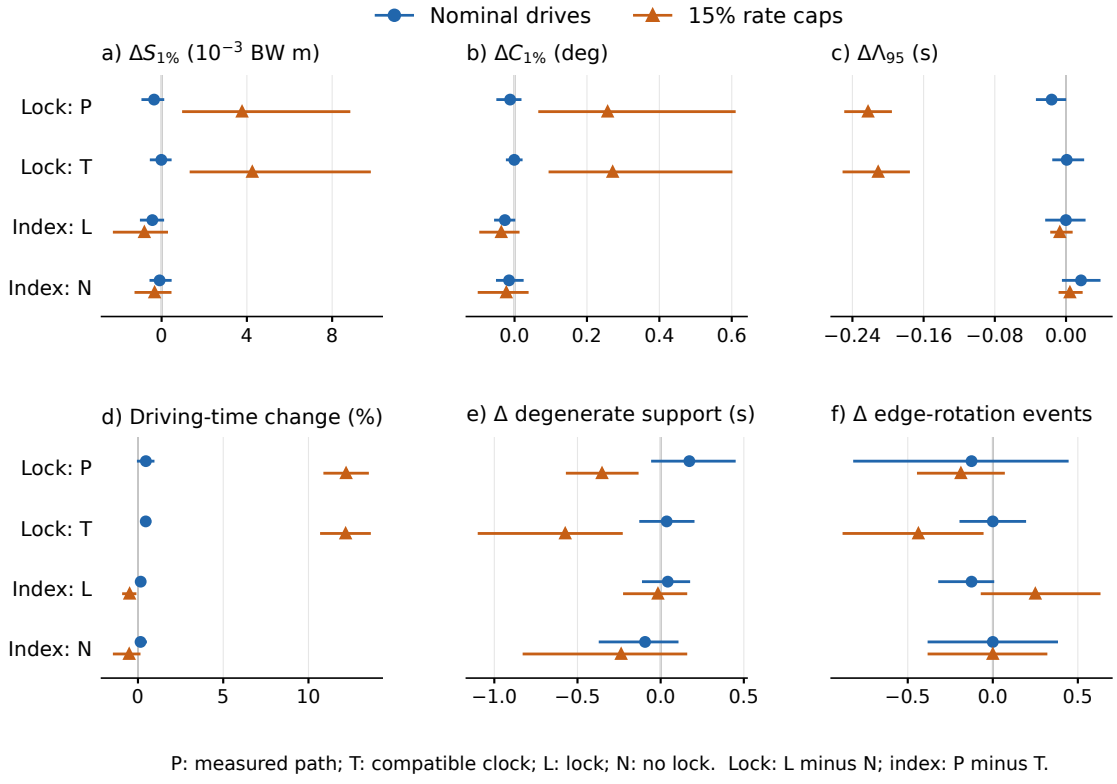


Figure 9: Matched execution study (S1 and S2, seeds 8–15, 400 kg): paired effects of the lock (L minus N) and of the indexing (P minus T) with configured drives and with 15 % rate caps, as means with 95 % seed-block bootstrap intervals. P: measured-path indexing; T: compatible-clock indexing; L: lock; N: no lock. $\Delta C_{1\%}$: change of the first percentile of θ_w (Table 10).

At equal achieved driving time, locking with $\Lambda_{max} = 2$ s has a lower Λ_{95} than the uniform-slowing frontier in all 15 seed–scenario pairs (exact sign tests: 7/7 on the embankment, $p = 0.016$; 8/8 on the crest, $p = 0.0078$). The pooled equal-time contrast is −0.136 s (−0.156 to −0.118 s), −0.142 s on the embankment and −0.131 s on the crest (Fig. 10). No speed factor gives both a shorter driving time and a lower tracking error than locking

Table 11: Equal-time execution study (seeds 16–23, 400 kg, 15 % rate caps; eight seeds per scenario). Done: completed/available. T: mean driving time of completions [s]. Λ_{95} and τ^*: mean episode 95th-percentile rate-normalised tracking error and whole-episode time-shift delay [s] over all episodes, including the S1 seed-16 nonprogress stop; changes quoted in the text are paired over completed episodes. PE/45: episodes with a proxy event / a 45° termination. Equal-time $\Delta\Lambda_{95}$: locking minus the uniform-slowing frontier interpolated at the driving time of the locking run, pooled over seed–scenario pairs whose locking time lies inside the observed uniform range (number in parentheses), with 95 % seed-block interval.

	S1 embankment				S2 crest					
Variant	Done	T	Λ_{95}	τ^*	Done	T	Λ_{95}	τ^*	PE/45	Equal-time $\Delta\Lambda_{95}$ [s]
P-N ($f = 1$)	7/8	82.1	1.305	2.31	8/8	60.0	0.926	1.38	0/0	–
P-U ($f = 0.95$)	7/8	86.0	1.261	2.23	8/8	63.4	0.916	1.33	0/0	–
P-U ($f = 0.891$)	7/8	91.6	1.210	2.09	8/8	67.6	0.816	1.24	0/0	–
P-U ($f = 0.85$)	7/8	96.2	1.180	2.03	8/8	70.6	0.708	1.10	0/0	–
P-U ($f = 0.80$)	7/8	102.6	1.121	1.91	8/8	76.9	0.736	1.12	0/0	–
P-L ($\Lambda_{\max} = 4$ s)	7/8	87.8	1.188	2.13	8/8	61.5	0.849	1.33	0/0	−0.063 [−0.098, −0.037] (15)
P-L ($\Lambda_{\max} = 2$ s)	7/8	96.8	1.040	1.95	8/8	62.4	0.772	1.25	1/0	−0.136 [−0.156, −0.118] (15)
P-L ($\Lambda_{\max} = 1$ s)	7/8	101.3	0.660	1.41	8/8	65.1	0.642	1.19	0/0	−0.339 [−0.398, −0.259] (14)
P-L ($\Lambda_{\max} = 0.5$ s)	7/8	102.0	0.458	1.04	8/8	69.4	0.504	0.94	0/0	−0.461 [−0.564, −0.323] (11)

in any pair. The directly measured delay agrees in sign: at equal time, τ^* is 0.074 s lower with locking (0.036–0.117 s; 12 of 15 pairs), and the windowed variant has the same sign (−0.29 s, −0.76 to 0.23 s). At $f = 0.891$, uniform slowing takes nearly the same time as locking when pooled over both scenarios (78.8 versus 78.4 s over completed episodes; paired time ratio −0.43 %, −1.82 to 0.54 %), although locking takes 5.7 % longer on the embankment and 7.7 % less on the crest; locking lowers Λ_{95} by 0.107 s (0.071–0.152 s). A comparison over the progress that the P-L, P-N and P-U runs of each seed and scenario have in common, which also includes the S1 seed-16 stops, gives 0.109 s (0.073–0.157 s).

Locking and uniform slowing allocate time differently. Relative to P-N, locking with $\Lambda_{\max} = 2$ s adds 18.0 % driving time on the embankment, where the lock is active during 42 % of driving, and 4.0 % on the crest, where it is active during 20 %; uniform slowing adds time everywhere. Varying the threshold traces the trade-off of locking itself, and every tested threshold lies below uniform slowing on average (Fig. 10): on the embankment, $\Lambda_{\max} = 0.5$ s lowers Λ_{95} by 0.87 s at 24.3 % more time, whereas the slowest uniform setting lowers it by 0.18 s at 25.0 %. The equal-time contrasts are −0.063, −0.136, −0.339 and −0.461 s, with 14 of 15, 15 of 15, 14 of 14 and 11 of 11 eligible pairs below the uniform frontier (0, 0, 1 and 4 pairs lie outside the observed uniform time range and are excluded). Stop-and-go execution, which reconfigures at standstill as in the relocation example of [10], is the end point of this curve: it gives the lowest tracking error at the largest time cost among the thresholds. What separates locking from uniform slowing is that the slowdown is spent where the measured tracking error requires it; the threshold selects a point on one time–tracking curve.

The equal-time advantage of locking is in tracking, and the audited restoring moment follows the driving time (Table 12): on the crest, slowing uniformly from $f = 1$ to $f = 0.80$ raises $S_{1\%}$ from −0.69 to 34.99×10^{-3} BW m and lowers the degenerate-support exposure from 2.75 to 1.85 %. At 1 and 10 µm, locking lies on the uniform-slowing frontier at equal driving time for all four thresholds: the frontier lies within the seed-block intervals of the mean equal-time differences, which are within $\pm 0.12 \times 10^{-3}$ BW m at 1 µm and between -0.81 and 1.91×10^{-3} BW m at 10 µm.

With a survey map, the planner no longer caps the speed ahead of the hidden crest

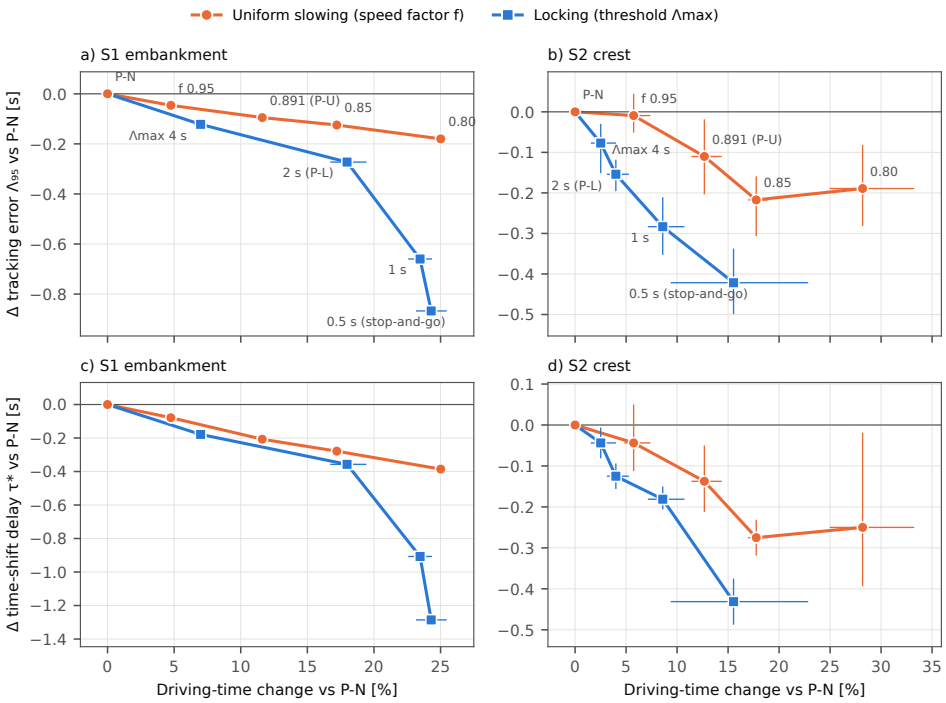


Figure 10: Tracking–time trade-off (S1 and S2, seeds 16–23, 400 kg, 15 % rate caps). Each point is the mean over seeds of the per-seed change relative to execution without the lock (P-N); bars are 95 % seed-block intervals. Orange: uniform slowing with speed factor f; blue: locking with threshold $\Lambda_{\max}$. (a, b) 95th-percentile rate-normalised tracking error Λ_{95}; (c, d) whole-episode time-shift delay τ^*. Lower left is better for tracking and time.

Table 12: Audited mechanics of the equal-time study (seeds 16–23, 400 kg, 15 % rate caps). $S_{1\%}$: mean episode first percentile of the minimum signed wheel-edge restoring moment at 10 µm conditioning [10^{-3} BW m]. ER: edge-rotation events (episodes with at least one). D: degenerate-support exposure [% of evaluated time]. Equal-time $\Delta S_{1\%}$: locking minus the uniform-slowing frontier interpolated at the driving time of the locking run, pooled, with 95 % seed-block interval (eligible pairs in parentheses).

	S1 embankment			S2 crest			
Variant	$S_{1\%}$	ER	D [%]	$S_{1\%}$	ER	D [%]	Equal-time $\Delta S_{1\%}$
P-N ($f = 1$)	−2.45	8 (6)	5.33	−0.69	3 (2)	2.75	–
P-U ($f = 0.95$)	−2.85	6 (5)	5.72	2.42	1 (1)	3.35	–
P-U ($f = 0.891$)	−2.41	6 (4)	5.36	7.42	1 (1)	3.32	–
P-U ($f = 0.85$)	−1.38	3 (3)	3.87	19.98	0 (0)	2.06	–
P-U ($f = 0.80$)	−0.50	4 (3)	3.34	34.99	0 (0)	1.85	–
P-L ($\Lambda_{\max} = 4$ s)	−1.95	6 (4)	4.69	−1.40	1 (1)	2.74	−0.81 [−2.29, 0.20] (15)
P-L ($\Lambda_{\max} = 2$ s)	−1.06	3 (3)	4.06	2.77	0 (0)	2.87	0.50 [−0.49, 1.78] (15)
P-L ($\Lambda_{\max} = 1$ s)	−0.68	6 (4)	3.94	6.04	1 (1)	2.35	1.91 [−0.22, 5.55] (14)
P-L ($\Lambda_{\max} = 0.5$ s)	−1.09	9 (5)	3.81	15.49	1 (1)	2.70	−0.76 [−2.65, 0.50] (11)

descent, and the crest crossings become shorter (P-N: 55.5 instead of 60.0 s; P-L: 57.4 instead of 62.4 s). Even so, with drives at 15 %, none of the 48 survey-map runs has a proxy event, a 45° termination or a payload contact, with or without the lock, and each variant has only the S1 seed-16 stop. Locking again lowers Λ_{95} relative to P-N, by 0.185 s (0.168–0.203 s), at 7.4 s longer mean driving time. Relative to uniform slowing, it gives the same crest tracking within its interval (+0.008 s, −0.022 to 0.034 s) with 4.2 s less driving,

and a 0.178 s lower embankment tracking error with 5.7 s more.

6.7. Documented drive capability and travel speed

The lock acts when the drives fall short of the demanded counterweight rate, so its role on the machine depends on how often the documented drives do. The demanded rate of a drive is the rate at which its counterweight coordinate moves in the configured-cap runs, which the documented envelopes do not limit. The slew dominates the tracking error: in every condition of the drive-envelope study and under the 15 % caps, the slew attains the maximum in (16) during 75–96 % of the time in which Λ exceeds Λ_0; the levelling contributes under the 30 % and 15 % caps (9–11 and 23–24 %), the extension under the drive envelopes (up to 22 %), and lift and tilt at most 6 %. The demanded rates also differ between the drives: in the 32 configured-cap runs at $0.8\,\mathrm{m\,s^{-1}}$, the slew rate has a median of $0.037\,\mathrm{rad\,s^{-1}}$ and a 99th percentile of $0.33\,\mathrm{rad\,s^{-1}}$, above its planning rate of $0.25\,\mathrm{rad\,s^{-1}}$: the planning rate bounds the planned posture, whereas the feedback term of $\boldsymbol{u}_{\mathrm{nom}}$ in (17) lets a lagging slew catch up at up to its configured cap of $0.52\,\mathrm{rad\,s^{-1}}$. The lift stroke speed has a median of $2.1\,\mathrm{mm\,s^{-1}}$ and a 99th percentile of $14.4\,\mathrm{mm\,s^{-1}}$, far below its planning rate.

Figure 11 places the caps against the documented ratings, which the loaded drives often exceed in force but rarely in demanded rate. The lift EMLA carries more than its rated force (66.8 kN) during 45 % of the driving time (RMS force 1.01 times the rated force), and the tilt EMLA during 1 %. The short-time envelope supplies at least 107 % of the lift planning rate at every logged force; a lift drive held to its rated power would supply a median 67 % of the planning rate, but the counterweight motion would exceed that envelope during only 0.21 % of the driving time. The slew column torque exceeds the rated motor torque at the model ratio during 36 % of the time in which the slew is not braked. Because a power limit bounds the product of torque and rate independently of the ratio, a slew held to its rated power would fall below the planning rate during 6.3 % of its unbraked time and below the demanded slew rate during 1.6 % of the driving time. The extension EMLA carries more than the force of its rated motor torque (5.9 kN) during 35 % of the driving time, yet its rated power covers the demanded extension rate during 99.94 % of it. In total, drives held to their rated power fall short of the demanded counterweight rate during 1.8 % of the driving time, the 30 % caps during 9.9 % and the 15 % caps during 27.9 % (Table 13); in terms of the demanded rate, the 15 % caps thus describe drives far below their documented continuous ratings.

The closed-loop runs confirm this ordering (Table 13). With the configured caps and the short-time and sustained-duty lift/tilt envelopes, locking changes Λ_{95} by +0.001 to +0.002 s at 0.38–0.65 % longer driving, and with the slew also held to its rated power by +0.010 s (−0.007 to 0.032 s) at 0.45 % longer driving. Drives that meet their documented ratings thus keep up with the counterweight plan, and the lock is active during about 4–5 % of the driving time. Under the 30 % caps it lowers Λ_{95} by 0.114 s (0.073–0.145 s; 13 of 14 completed pairs) and τ^{*} by 0.075 s (0.056–0.093 s) at 1.3 % longer driving, and under the 15 % caps by 0.209 s (0.184–0.233 s) and 0.233 s (0.196–0.256 s) in all 15 pairs at 11.6 % longer driving. No run reaches the 45° termination. With locking, the minimum static reserve changes by 0.2° or less on average in every condition, and the 1st-percentile acceleration-based reserve by at most 0.29°. Proxy events occur only on embankment seeds 16 and 17; under the 30 % caps, the seed-17 episode without locking stops after two proxy events with a minimum static reserve of −0.8°, whereas the locked episode completes without one.

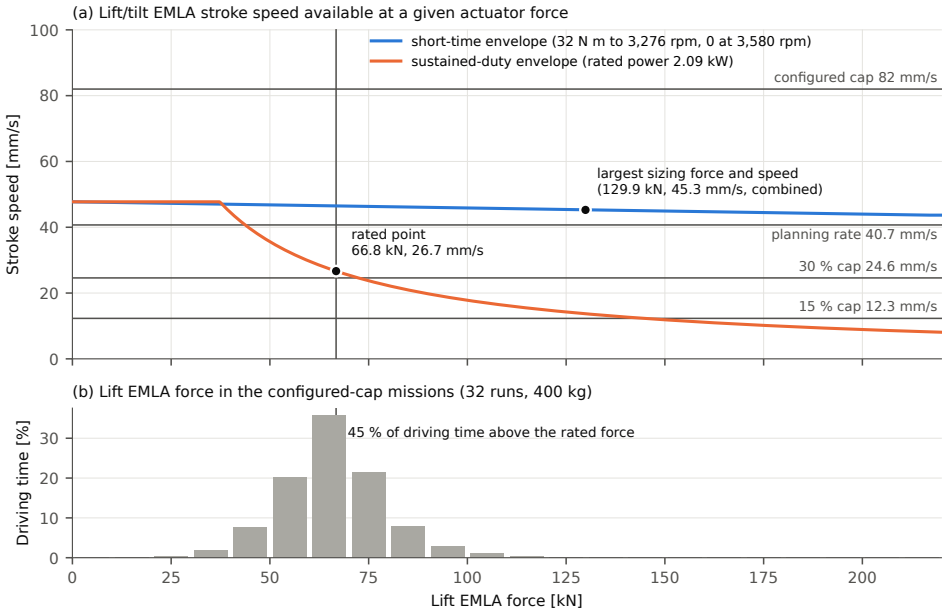


Figure 11: Documented drive capability and the tested rate caps. (a) Stroke speed of a lift/tilt EMLA available at a given actuator force under the short-time and sustained-duty envelopes of Section 5.3; grey lines: configured cap, planning rate and the 30 % and 15 % caps; vertical line: rated force 66.8 kN. (b) Lift EMLA force over the driving time of the 32 configured-cap runs at 0.8 m s^{-1} (S1 and S2, seeds 16–23, 400 kg).

At the 0.35 m s^{-1} speed limit of the field experiments (Table 13, lower part), the counterweight plan demands much lower rates: in the 32 configured-cap runs, the 99th percentile of the demanded slew rate is 0.18 rad s^{-1} instead of 0.33 rad s^{-1}, and drives held to their rated power fall short of the demanded rate during 0.12 % of the driving time instead of 1.8 %, the 30 % caps during 1.9 % instead of 9.9 % and the 15 % caps during 7.3 % instead of 27.9 %. All 128 runs complete without a proxy event, a payload contact or a 45° termination, including S1 seed 16, which stops at 0.8 m s^{-1}. With the configured caps and with rated-power drives, the lock is active during 1.3 and 2.0 % of the driving time and adds 0.14 and 0.13 % driving time, with Λ_{95} changes of −0.001 and −0.005 s. Under the 30 % caps it changes Λ_{95} by 0.000 s (−0.005 to 0.004 s) at 0.17 % longer driving, and under the 15 % caps it lowers Λ_{95} by 0.064 s (0.036–0.095 s; 14 of 16 pairs) at 0.71 % longer driving. At the field speed, the documented drives thus keep up with the counterweight plan, and the lock serves drives far below their rating at a small cost in driving time.

6.8. Ablations, route escalation and calibration

In the 24 loaded nominal episodes (Table 14), removing the preview (−P) causes a 45° termination on crest seed 6 and 25 payload contacts in five episodes; its crest response starts only after the edge is reached, with a tracking error Λ of nearly 3 s, and the mean 1st-percentile acceleration-based reserve falls from 16.1° to 14.7° ($p = 0.03$). Removing the filter (−C) produces nine payload contacts on two crests. Removing synchronisation (−S), levelling (−L), routing (−R) or the planner's actuator limits (−A) leaves the event outcomes unchanged, and the pooled reserves and times differ by at most 0.2° and 0.3 s; levelling raises the smallest static reserve on the crest by 0.7°. Preview and clearance filtering therefore make the largest contributions in these nominal tests, whereas synchronisation acts under slow drives (Sections 6.5–6.7) and routing on steep slopes.

Table 13: Lock effect under documented drive envelopes and rate caps at the nominal speeds of 0.8 and $0.35\,\mathrm{m\,s^{-1}}$ (S1 and S2, seeds 16–23, 400 kg). Short.: share of the driving time of the 32 configured-cap runs at the same nominal speed in which the slew, lift, tilt or extension rate exceeds the rate available under the condition at the logged forces [%]; the levelling is not included. Λ_{95}: mean P-N value over completed pairs [s]. Δ: P-L minus P-N, mean over completed seed–scenario pairs with 95 % seed-block interval; L<N: completed pairs in which locking has the lower Λ_{95}, of all completed pairs; ΔT in % of P-N driving time; τ^* in s. θ_{st}: mean minimum static reserve over completed pairs, P-N/P-L [deg]. PE: proxy events, P-N/P-L. Short-time and sustained duty: lift/tilt envelopes of Fig. 11; rated power, lift/tilt/slew: lift, tilt and slew drives held to their documented rated power. At $0.8\,\mathrm{m\,s^{-1}}$, the S1 seed-16 episode stops without progress in every condition, and the 15 % row uses the P-N and P-L ($\Lambda_{\max} = 2\,\mathrm{s}$) runs of the equal-time study. No run reaches the 45° termination.

Condition	Short.	Λ_{95}	$\Delta\Lambda_{95}$ [s]	L<N	ΔT [%]	$\Delta\tau^*$ [s]	θ_{st}	PE
nominal speed 0.8 m s^{-1}								
Configured caps	0	0.381	0.001 [−0.008, 0.009]	6/15	0.46 [0.29, 0.70]	−0.003 [−0.019, 0.010]	19.4/19.5	0/0
Short-time, lift/tilt	0.03	0.407	0.002 [−0.004, 0.011]	9/15	0.65 [0.36, 0.97]	−0.010 [−0.019, −0.003]	18.4/18.5	2/1
Sustained duty, lift/tilt	0.21	0.398	0.002 [−0.008, 0.014]	8/15	0.38 [0.16, 0.61]	0.013 [−0.004, 0.034]	19.0/18.9	0/1
Rated power, lift/tilt/slew	1.8	0.399	0.010 [−0.007, 0.032]	9/15	0.45 [0.30, 0.63]	0.033 [0.000, 0.083]	19.0/18.9	0/1
30 % caps	9.9	0.634	−0.114 [−0.145, −0.073]	13/14	1.26 [1.06, 1.45]	−0.075 [−0.093, −0.056]	19.6/19.5	2/0
15 % caps	27.9	1.113	−0.209 [−0.233, −0.184]	15/15	11.62 [10.12, 12.64]	−0.233 [−0.256, −0.196]	19.3/19.1	0/1
nominal speed 0.35 m s^{-1}								
Configured caps	0	0.225	−0.001 [−0.008, 0.005]	9/16	0.14 [0.09, 0.19]	−0.025 [−0.059, 0.003]	19.1/19.3	0/0
Rated power, lift/tilt/slew	0.12	0.253	−0.005 [−0.020, 0.009]	10/16	0.13 [0.03, 0.20]	−0.075 [−0.156, −0.013]	19.0/19.1	0/0
30 % caps	1.9	0.234	0.000 [−0.005, 0.004]	6/16	0.17 [0.12, 0.22]	−0.009 [−0.031, 0.013]	19.0/19.1	0/0
15 % caps	7.3	0.447	−0.064 [−0.095, −0.036]	14/16	0.71 [0.52, 0.92]	−0.075 [−0.100, −0.050]	19.1/19.3	0/0

Table 14: Ablations with the 400 kg payload over S1–S3 (seeds 0–7, 24 episodes per row; variants as in Table 3). Gap: 95th percentile of the difference between planned and executed static reserve [deg]; Strikes: payload–ground contacts; other columns as in Table 5.

Variant	Compl.	45°	Stop	Proxy events	θ_{acc}	θ_{st}	Time	Λ_{95}	Gap	Strikes
COMPASS	24/24	0	0	0.00 (0)	16.1	21.6	67.4	0.32	4.0	0
COMPASS−P	23/24	1	0	0.04 (1)	14.7	20.6	66.0	0.67	2.0	25
COMPASS−S	24/24	0	0	0.00 (0)	16.1	21.5	67.1	0.32	4.0	0
COMPASS−C	24/24	0	0	0.00 (0)	16.1	21.5	67.3	0.32	4.0	9
COMPASS−L	24/24	0	0	0.00 (0)	15.9	21.6	67.7	0.42	3.8	0
COMPASS−R	24/24	0	0	0.00 (0)	16.1	21.6	67.4	0.32	4.0	0
COMPASS−A	24/24	0	0	0.00 (0)	16.0	21.6	67.5	0.32	4.0	0

S4 increases the cross slope along the route (two seeds per condition; Table 15 and Fig. 12). Up to 30°, all controllers complete the routes without proxy events, and COMPASS changes its route in one of eight episodes. At 34° and 38°, FIX reaches the 45° termination in three of eight episodes and has proxy events in two others, and REACT-S

has one termination and one further proxy-event episode. COMPASS completes all eight episodes after shifting its route 2.6 m downhill to gentler ground, with one proxy event (Section 6.9); at 38° its mean smallest static reserves are 14.7° empty and 15.2° loaded, against 0.3° and 0.2° for FIX, and driving takes 2–3 % longer than at 26°. Without routing, six of the eight episodes at the two steepest slopes end in a stop or timeout, four have proxy events, and the mean smallest static reserve is 5.3°.

Table 15: Stress envelope on S4 (cross slope increasing to $\alpha_{\max}$; seeds 0–1, two episodes per cell). Each cell: completed/45°-terminated/stopped episodes (nonprogress stop or 200 s time limit) and, in parentheses, the mean smallest static reserve [deg].

Payload	Controller	$\alpha_{\max} = 26°$	30°	34°	38°
empty	FIX	2/0/0 (15.5)	2/0/0 (12.7)	1/1/0 (4.3)	1/1/0 (0.3)
	REACT-S	2/0/0 (24.3)	2/0/0 (19.5)	2/0/0 (13.4)	2/0/0 (13.5)
	COMPASS−R	2/0/0 (23.9)	2/0/0 (18.5)	2/0/0 (5.7)	0/0/2 (5.5)
	COMPASS−conf.	2/0/0 (23.9)	2/0/0 (18.5)	2/0/0 (8.6)	2/0/0 (12.7)
	COMPASS	2/0/0 (23.9)	2/0/0 (18.5)	2/0/0 (8.6)	2/0/0 (14.7)
400 kg	FIX	2/0/0 (13.9)	2/0/0 (11.3)	2/0/0 (8.6)	1/1/0 (0.2)
	REACT-S	2/0/0 (23.2)	2/0/0 (17.5)	2/0/0 (12.0)	1/1/0 (−3.9)
	COMPASS−R	2/0/0 (21.7)	2/0/0 (10.8)	0/0/2 (5.7)	0/0/2 (4.5)
	COMPASS−conf.	2/0/0 (21.7)	2/0/0 (16.3)	2/0/0 (14.5)	2/0/0 (13.5)
	COMPASS	2/0/0 (21.7)	2/0/0 (13.3)	2/0/0 (15.4)	2/0/0 (15.2)

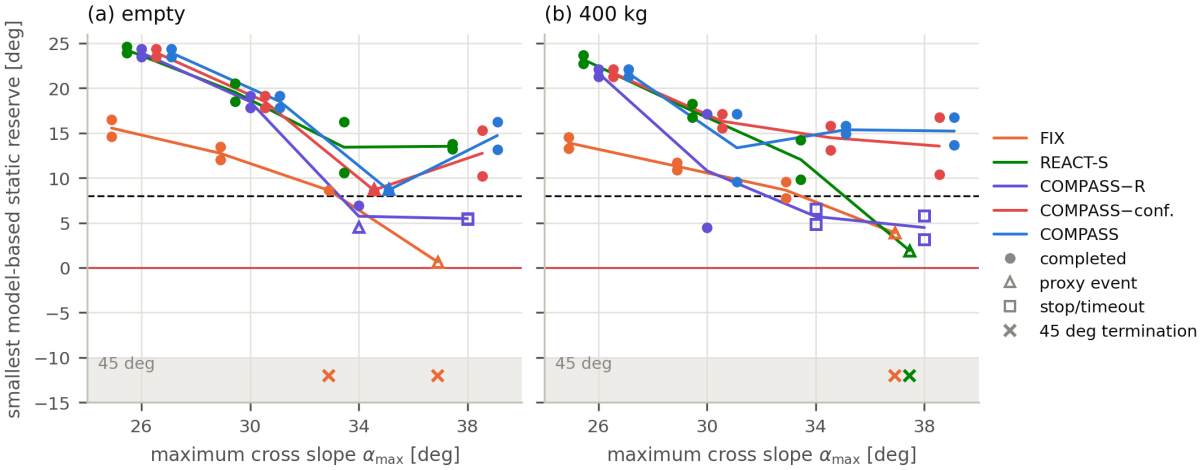


Figure 12: S4 stress envelope: smallest static reserve per episode (two seeds per controller, payload and maximum slope). Lines average the episodes without a 45° termination; markers show the outcomes. Dashed: $\theta_{\min}$; red: zero static reserve.

The component study repeats the steepest conditions with eight seeds per cell (Table 16). Full COMPASS completes 31 of 32 runs, against 6 of 32 without routing; the 26 noncompletions without routing comprise 16 nonprogress stops and ten timeouts. No variant reaches the 45° termination. Full and no-margin execution have no proxy events, whereas the variant without routing has proxy events in 18 of 32 episodes. Edge-rotation events occur in 13, 15 and 21 episodes for full, no-margin and no-routing execution (15, 20 and 29 events). Without routing, $S_{1\%}$ is highest in all four cells, together with the most stops and the longest degenerate support (13.6–20.2 s per episode, which $S_{1\%}$ excludes); $S_{1\%}$ is therefore read together with the degenerate-support exposure and the edge-rotation

events. Routing thus provides a clear mission-progress benefit on steep slopes, and the conformal margin lowers the edge-rotation count from 20 to 15 with unchanged proxy and termination outcomes.

Table 16: Component study on S4, eight seeds per slope, payload and variant. Done: completed/available episodes; Stop: nonprogress stops and timeouts; PE and ER: episodes with a proxy event and with an edge-rotation event; m_p: payload mass; $S_{1\%}$ [10^{-3} BW m] and D [s]: episode means, as in Table 10.

$\alpha_{\max}$ [°]	m_p [kg]	Variant	Done	Stop	45°	PE	ER	$S_{1\%}$	D [s]
34	0	Full	7/8	1	0	0	2	−3.96	7.6
34	0	−route	6/8	2	0	2	5	−2.64	15.1
34	0	−conf.	8/8	0	0	0	2	−3.98	8.0
34	400	Full	8/8	0	0	0	5	−3.07	6.5
34	400	−route	0/8	8	0	5	6	−1.76	14.5
34	400	−conf.	8/8	0	0	0	6	−3.12	6.8
38	0	Full	8/8	0	0	0	3	−4.50	8.0
38	0	−route	0/8	8	0	6	6	−1.06	20.2
38	0	−conf.	8/8	0	0	0	4	−3.89	8.2
38	400	Full	8/8	0	0	0	3	−3.74	6.0
38	400	−route	0/8	8	0	5	4	−1.18	13.6
38	400	−conf.	8/8	0	0	0	3	−3.91	6.4

6.9. Model mismatch and plant variation

Table 17 varies the declared payload, the ground friction and the drive operation on S1 and S2. A declared but absent payload, the assigned ground coefficient reduced from 0.60 to 0.45 (an effective rigid tyre–ground friction of at least 0.50, because the contact model combines it with the robot coefficient) and half-speed drives cause no proxy events; the lower friction reduces the mean smallest static reserve by 2.6°. Blocking the slew at 12 s increases the completed driving times by 61 % on the embankment and 15 % on the crest, and one episode reaches the 200 s time limit while creeping. An undeclared 400 kg payload exceeds the holdability limits of the declared model: on embankment seed 0, the slew saturates at 8.8 kN m, the hull reaches 29° nose-up pitch and 36° roll, and the machine slides 1.1 m downhill. Payload contact props the machine while its static reserve stays negative for 16 s and later for 6 s, and the route is completed despite 17 proxy events and a minimum static reserve of −18°; crest seed 3 has one further proxy event under this mismatch. The payload declaration is thus an operating input on which the counterweight relies.

In the drive-fault study, the lock changes the response to the blocked slew (Table 17). With the lock, the vehicle creeps while the slew cannot follow its plan: the lock is active during 67 % of the driving time, the seven pairs in which both variants complete take 36 % longer with the lock (96.8 against 71.3 s), and one embankment episode reaches the time limit. Without the lock, all eight episodes complete, but the 1st-percentile acceleration-based reserve is lower in all eight pairs (10.4° against 13.5° on average) and wheel-unloading events are more frequent (2.9 against 0.5 per episode); neither variant has a proxy event, a payload contact or a 45° termination.

Across the main, stress-envelope, derated-drive and model-mismatch studies, COMPASS has no 45° termination in 152 episodes, and proxy events occur in three of them:

Table 17: Model mismatch and drive fault, COMPASS on S1 and S2 (seeds 0–3, eight episodes per row; P-L is identical to COMPASS); the nominal rows are the matching main-study episodes. Slew blocked: the slew drive stops 12 s into the drive and holds its position, unknown to the controller; the last row repeats this fault without the lock. Strikes: payload–ground contacts; other columns as in Table 5.

Condition	Compl.	45°	Stop	Proxy events	θ_{acc}	θ_{st}	Time	Strikes
nominal, 400 kg	8/8	0	0	0.00 (0)	15.1	19.0	71.6	0
400 kg, declared empty	8/8	0	0	2.25 (2)	8.8	10.9	76.9	15
nominal, empty	8/8	0	0	0.00 (0)	14.9	19.5	72.6	0
empty, declared 400 kg	8/8	0	0	0.00 (0)	13.3	17.3	71.6	0
ground 0.45 (effective ≥ 0.50), 400 kg	8/8	0	0	0.00 (0)	14.4	16.4	72.7	0
half-speed drives, 400 kg	8/8	0	0	0.00 (0)	14.9	18.9	71.8	0
slew blocked at 12 s, 400 kg (P-L)	7/8	0	1	0.00 (0)	13.5	14.2	96.8	0
slew blocked, without the lock (P-N)	8/8	0	0	0.00 (0)	10.4	13.2	73.7	0

the two episodes with an undeclared payload and, for 0.1 s, the empty machine sliding over a rock on the 34° stress slope (seed 1). In the main study, the proxy meets the 8° target during 99.98 % of the driving time.

In the plant-variation study, both execution variants complete all 84 runs at the configured rates without proxy events or 45° terminations (Table 18). Over the 16 Latin-hypercube settings, P-L and P-N have nine and seven edge-rotation events (in eight and six of 32 episodes), and the pooled P-L minus P-N effects, with intervals over parameter rows, are $+0.73 \times 10^{-3}$ BW m (-4.97 to 6.24) on $S_{1\%}$, -0.001 s (-0.014 to 0.011) on Λ_{95}, +0.013 % (-0.60 to 0.53) on driving time and -0.173 s (-0.412 to 0.020) on degenerate-support exposure. The nominal row, the corners and the anchors change Λ_{95} by -0.012 to $+0.021$ s, and $S_{1\%}$ by less than 1×10^{-3} BW m except with the softer tyre ($+20.02 \times 10^{-3}$ BW m). At the configured rates, locking thus leaves tracking, time and mechanics unchanged within these intervals across the plant variation, as with the documented drive envelopes (Section 6.7).

Table 18: Plant variation at the configured drive rates and anchors, S1 and S2, 400 kg. Counts are P-L/P-N episode counts; n is the number of paired episodes. PE and ER: episodes with a proxy event and an edge-rotation event. LHS: the 16 Latin-hypercube settings. Effects are P-L minus P-N; the LHS intervals resample the 16 parameter rows. The named corners and the motor-family anchors are deterministic checks. $\Delta S_{1\%}$: change of the first-percentile minimum signed wheel-edge restoring moment at 10 µm conditioning [10^{-3} BW m].

Setting	n	Done	45°	PE	ER	$\Delta S_{1\%}$	$\Delta\Lambda_{95}$ [s]
LHS	32	32/32	0/0	0/0	8/6	+0.73 [−4.97, 6.24]	−0.001 [−0.014, 0.011]
Nominal	2	2/2	0/0	0/0	0/0	−0.43	−0.012
Soft corner	2	2/2	0/0	0/0	0/0	+0.06	−0.005
Stiff corner	2	2/2	0/0	0/0	1/1	−0.67	+0.021
Motor family	2	2/2	0/0	0/0	1/1	−0.11	+0.013
Family + soft tyre	2	2/2	0/0	0/0	0/1	+20.02	−0.009

With 15 % rate caps, the corners and anchors again complete without proxy events or terminations (Table 19). Locking lowers Λ_{95} in every setting, by 0.137 s (soft corner), 0.264 s (stiff corner), 0.228 s (motor-family anchor) and 0.256 s (with the softer tyre), at 5.18, 16.85, 11.70 and 12.31 % longer driving. Its effects on $S_{1\%}$ depend on the conditioning

and the tyre (with the softer tyre +0.59, −17.20 and +10.08 × 10^{-3} BW m at 1, 10 and 100 µm). The tracking benefit of locking under slow drives thus holds in all four settings, and its effect on the restoring moment varies with the plant, in agreement with the equal-time comparison (Table 12).

Table 19: Plant corners and anchors with 15 % rate caps (S1 and S2; parameters and seeds as in the configured-rate runs). Counts are P-L/P-N episodes; PE and ER as in Table 10. Effects are P-L minus P-N: $\Delta S_{1\%}$ at 1/10/100 µm conditioning [10^{-3} BW m], $\Delta\Lambda_{95}$ [s] and driving time ΔT [%].

Setting	Done	45°	PE	ER	$\Delta S_{1\%}$ (1/10/100 µm)	$\Delta\Lambda_{95}$	ΔT
Soft corner	2/2	0/0	0/0	0/0	+0.30 / −0.04 / −0.02	−0.137	+5.18
Stiff corner	2/2	0/0	0/0	0/0	+1.30 / +1.67 / −0.24	−0.264	+16.85
Motor family	2/2	0/0	0/0	1/1	+0.09 / +0.56 / +0.31	−0.228	+11.70
Family + soft tyre	2/2	0/0	0/0	0/0	+0.59 / −17.20 / +10.08	−0.256	+12.31

6.10. Computation

On one core of a 2.1 GHz server, the nominal main-study COMPASS episodes need 7.0 ms of computation per 20 ms control tick on average (episode means at most 9.0 ms). Per invocation, perception takes 8.2 ms at 10 Hz, routing 6.6 ms at 1 Hz, planning 15.5 ms at 5 Hz and Layer C 2.1 ms at 50 Hz (5.4 ms per tick amortised; the rest is untimed tick overhead). Neither QP reports a solver failure, and the filter slack is used in 0.07 % of the cycles.

7. Discussion

7.1. Implications of the audit

The acceleration-based proxy is useful as a planning and filtering signal, but on this machine it departs from the moment-inclusive wrench angle θ_{w} by about 9° RMS, and omitting the moment alone, on identical contact geometry and force, changes the angle by 9.18° RMS. The proxy is negative during only a small fraction of the time in which the wrench angle is negative. A controller that certifies tip-over margins with such a proxy needs an explicit bound on the omitted moment term of (6) or an independent check. The closure of the force and moment accounting shows that the audit is consistent within the simulator, and its restoring moments, friction-cone capacities and edge-rotation events distinguish recoverable wheel unloading, outward rotation, body support and 45° terminations. Absolute exposures and the magnitudes of paired contrasts depend on the contact conditioning; the crest contrast between the counterweight and the transport pose keeps its sign and separation under every treatment, for COMPASS and for PREV.

7.2. Contribution relative to prior coordination methods

Table 20 summarises the findings on the components of COMPASS with the comparisons that support them. Relative to hierarchical rough-terrain planning [10], COMPASS contributes the integration of preview counterweight planning with capacity-limited speed and route escalation on an electro-mechanically actuated construction carrier, and an evaluation of its execution when the drives are slower than the plan assumes. COMPASS adopts the tracking-error time scaling of [11] to execute a stability counterweight, and the results confirm the tracking–time advantage of that principle over uniform slowing in

Table 20: Summary of findings and the comparisons that support them.

Component	Comparison	Evidence	Finding
Terrain preview	COMPASS−P, 24 loaded episodes	one 45° termination and 25 payload contacts without preview; 1st-percentile proxy reserve 16.1° versus 14.7° ($p = 0.03$)	supported
Clearance filter	COMPASS−C	9 payload contacts on two crests without it	supported
Counterweight versus transport pose	FIX vs COMPASS and vs PREV (same planner at the nominal speed): main study, mechanics set, crest robustness	composite endpoint 7 vs 0 episodes ($p = 0.016$); higher audited $S_{1\%}$ in 8/8 crest seed pairs; 47 vs 5 edge-rotation events; complete separation at 1–100 µm, with 0.33 and 1.8 MN m^{-1} tyres and a 0.5 ms step, for COMPASS and for PREV	supported on crests; robust to contact, tyre and time-step treatment; present without the slower crossing
Speed-compatible execution	PREV vs COMPASS and COMPASS−S	PREV 45° termination on crest seed 6 at configured, 30 % and 15 % rates; none for speed-compatible execution	supported (crest seed 6)
Lock (tracking-error time scaling [11])	P-L vs P-N; equal-time frontier; drive envelopes; travel speed; drive fault	0.8 m s^{-1}: Λ_{95} 0.11–0.22 s lower at 1.3–12.2 % more time with 30/15 % caps, 0.136 s below uniform slowing at equal time; rated-power drives: lock active during 4–5 % of the time, at most 0.65 % more time; 0.35 m s^{-1}: 15 % caps −0.064 s at 0.7 % more time; blocked slew: automatic slowdown, proxy reserve +3.1° at 36 % more time; restoring moment at equal time on the uniform-slowing frontier (1, 10 µm)	tracking benefit under slow drives; better tracking than uniform slowing at equal time; under 1 % time cost with documented drives
Path indexing	P vs T	comparable restoring moment and tracking	both indexings suitable
Route escalation	COMPASS−R, S4 34°/38°	31/32 vs 6/32 completions	supported on steep slopes
Conformal margin	COMPASS−conf., S4	15 vs 20 edge-rotation events; unchanged proxy and termination outcomes	fewer edge-rotation events

this new role. What the results add beyond [11] are terrain preview (without it, a 45° termination and 25 payload contacts), capacity-limited route escalation (31/32 versus 6/32 completions on the steepest slopes) and a path-indexed speed profile that slows ahead of unmapped terrain, which on crest seed 6 avoided the time-indexed failure even without capacity- and rate-based slowing (COMPASS−S). The mechanical margin comes from the counterweight and, under slow drives, from the driving time: at equal driving time, the audited moment of locking lies on the uniform-slowing frontier, while the lock spends the time where the tracking error requires it. Speed-compatible execution avoided the crest failure of seed 6 with or without the lock, and the lock acts when the drives fall short of the demanded rate; the documented drives of this machine meet that rate during more than 98 % of the driving time (Section 7.3).

7.3. Relevance for the machine and its operation

The role of the lock on the machine depends on the rate that the loaded drives deliver relative to the rate that the counterweight plan demands, and the slew dominates this ratio. With the documented motor ratings (and, for the slew speed cap, the ratio of 600

of the model), drives held to their rated power meet the demanded rate during 98.2 % of the driving time at 0.8 m s^{-1}, and the lock then acts during about 4–5 % of the driving time (Section 6.7). The 30 % and 15 % caps are stress tests. The 15 % caps correspond to drives that deliver far less than their continuous rating, for example after a partial drive fault, with a weak supply, under thermal derating, or with a planner given rates that the drives cannot reach. The 30 % caps lie in between: for the lift, 24.6 mm s^{-1} is close to the stroke speed at the rated operating point (26.7 mm s^{-1}), whereas for the slew they correspond to 30 % of the rated motor speed at the ratio of the model. At the rated motor speed, the 30 % and 15 % slew caps correspond to reduction ratios of 2,000 and 4,000 instead of 600.

In the field experiments with the same machine, the linear speed was limited to 0.35 m s^{-1} during approximately 3.5 km of data collection on asphalt and soft soil, under conservative motion limits chosen because of the size and mass of the platform [33], and the gravel trials of [34] used reference speeds of 0.11–0.15 m s^{-1}. At this speed, rated-power drives meet the demanded rate during 99.88 % of the driving time, the lock costs 0.13 % of driving time, and under the 15 % caps it lowers Λ_{95} by 0.064 s at 0.7 % longer driving (Section 6.7). In the field runs of [34], the travel drives followed wheel-speed commands with fitted dead times of 0.21–0.28 s in nine of ten first-order-plus-dead-time fits. While the vehicle follows the planned speed, the posture reference moves at no more than about 90 % of the planning rate ($\rho = 0.9$), so even a stalled actuator needs about 1.8 s to raise Λ from Λ_0 to $\Lambda_{\max}$, and a speed reduction commanded by the lock takes effect well within this interval.

For construction operation, the consequences are as follows. With drives that meet their documented ratings, the lock costs under 1 % of driving time on these routes (at most 0.65 %). When the drives fall short of the demanded rate, the lock keeps the counterweight closer to its terrain-synchronous plan at a cycle-time cost that grows with the shortfall: under the 30 % caps the time-shift delay of the posture fell from 0.57 to 0.50 s at 1.3 % longer driving, and under the 15 % caps from 1.82 to 1.59 s at 11.6 % longer driving. At the mean progress speed of these runs, about 0.7 m s^{-1} (route length over driving time), the larger of these reductions (0.23 s) means that the posture lags its path-indexed plan by about 0.16 m less. Under the 30 % caps, the lock also turned a stop after two proxy events on embankment seed 17 into a completion (Section 6.7). When a drive fails outright, as with the blocked slew of Section 6.9, the lock slows the vehicle for most of the remaining drive, and the locked runs keep a 3.1° higher 1st-percentile proxy reserve at 36 % longer driving. The lock is therefore a low-cost tracking aid for drives that fall short of the demanded rate and an automatic slowdown when a drive fails. A design implication is to plan the counterweight rates within the rated-power envelope of the loaded drives, the slew in particular. Because the demanded rate grows with the travel speed and with the change of the planned posture along the path, and the rated-power envelope falls with load (Fig. 11), the lock gains importance with the travel speed, as the comparison of 0.35 and 0.8 m s^{-1} shows, and by the same mechanism with more abrupt terrain and heavier payloads.

7.4. Outlook

The whole-body model uses the documented geometry, masses and drive ratings of the machine and matches the field-tested machine in mass, wheel radius, track and travel-drive architecture. The next step is the validation of the counterweight on this machine. There, wheel-load or contact-force sensing will provide the hardware counterpart of the

independent audit, drive rates measured under load will refine the drive envelopes of Section 6.7, and the identification of the slew reduction ratio and efficiency, the levelling rate and the tyre damping will update the model parameters (Table 1). A volumetric terrain representation will extend the preview to overhangs and soil strength, an online payload estimate will complement the payload declaration (Section 6.9), and conformal methods for distribution shift [30] will carry the calibrated capacity margin to new sites.

8. Conclusion

COMPASS integrates terrain preview, crane and beam counterweight motion, capacity-limited speed and route escalation, and execution that scales vehicle progress with the manipulator tracking error, the coordination principle of [11] applied to a stability counterweight. The simulations of an electrified skid-steer construction carrier support three conclusions. First, an independent contact-wrench and centroidal-momentum audit shows that the acceleration-based reserve departs from the moment-inclusive wrench angle by 8.84° RMS, and it detects unloading with outward rotation that the proxy does not flag. Second, relative to the transport pose, COMPASS raises the audited lower-tail restoring moment on crests in all eight seeds and reduces edge-rotation events from 47 to 5; this contrast holds at contact conditionings from 1 to 100 µm, with tyre stiffnesses of 0.33 and 1.8 $\mathrm{MN\,m^{-1}}$ and with a 0.5 ms step, and also for PREV, which uses the same planner but drives at the nominal speed of FIX. At that speed, however, PREV failed on one crest seed, a failure that speed-compatible execution avoided. Third, the lock, which couples the vehicle speed to the tracking error, improves the counterweight timing when the drives fall short of the demanded rate: it lowers the 95th-percentile tracking error by 0.11–0.22 s at 1.3–12.2 % longer driving, and by 0.136 s relative to uniform slowing at equal driving time. With drives held to their documented ratings it costs under 1 % of driving time, both at 0.8 $\mathrm{m\,s^{-1}}$ and at the field speed of 0.35 $\mathrm{m\,s^{-1}}$, where under 15 % caps it still lowers the tracking error by 0.064 s at 0.7 % longer driving, and under a blocked slew it slows the vehicle automatically. Beyond the coordination principle of [11], terrain preview and route escalation also mattered: without preview, one 45° termination and 25 payload contacts occurred, and without the route layer only 6 of 32 episodes on the steepest slopes were completed, against 31. Field validation of the counterweight on the machine is the next step.

CRediT authorship contribution statement

Mehdi Heydari Shahna: Writing – original draft, Validation, Methodology, Investigation, Formal analysis, Software, Data curation, Conceptualization.

Jouni Mattila: Writing – review & editing, Resources, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This work was supported by the Finnish Foundation for Technology Promotion (Tekniikan Edistämissäätiö) through the Post Docs in Companies (PoDoCo) program.

Data availability

The code implementing the proposed framework and the supplementary video are provided as supplementary material and will be made publicly available on GitHub upon publication. With the pinned environment (Python 3.12), a complete episode reproduces bit-for-bit on a second computer.

References

[1] T. Bock, The future of construction automation: Technological disruption and the upcoming ubiquity of robotics, Automation in Construction 59 (2015) 113–121. doi:10.1016/j.autcon.2015.07.022.

[2] N. Melenbrink, J. Werfel, A. Menges, On-site autonomous construction robots: Towards unsupervised building, Automation in Construction 119 (2020) 103312. doi:10.1016/j.autcon.2020.103312.

[3] S. Dadhich, U. Bodin, U. Andersson, Key challenges in automation of earth-moving machines, Automation in Construction 68 (2016) 212–222. doi:10.1016/j.autcon.2016.05.009.

[4] R. F. Abo-Shanab, N. Sepehri, Tip-over stability of manipulator-like mobile hydraulic machines, Journal of Dynamic Systems, Measurement, and Control 127 (2) (2005) 295–301. doi:10.1115/1.1898239.

[5] A. Rauch, W. Singhose, D. Fujioka, T. Jones, Tip-over stability analysis of mobile boom cranes with swinging payloads, Journal of Dynamic Systems, Measurement, and Control 135 (3) (2013) 031008. doi:10.1115/1.4023276.

[6] Q. Huang, S. Sugano, K. Tanie, Stability compensation of a mobile manipulator by manipulator motion: Feasibility and planning, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Vol. 3, 1997, pp. 1285–1292. doi:10.1109/IROS.1997.656417.

[7] Q. Huang, K. Tanie, S. Sugano, Coordinated motion planning for a mobile manipulator considering stability and manipulation, The International Journal of Robotics Research 19 (8) (2000) 732–742. doi:10.1177/02783640022067139.

[8] J. Morales, J. L. Martinez, A. Mandow, J. Seron, A. Garcia-Cerezo, A. Pequeño-Boter, Center of gravity estimation and control for a field mobile robot with a heavy manipulator, in: Proceedings of the IEEE International Conference on Mechatronics (ICM), 2009, pp. 1–6. doi:10.1109/ICMECH.2009.4957112.

[9] X. Ding, Y. Liu, J. Hou, Q. Ma, Online dynamic tip-over avoidance for a wheeled mobile manipulator with an improved tip-over moment stability criterion, IEEE Access 7 (2019) 67632–67645. doi:10.1109/ACCESS.2019.2915115.

[10] J. Song, I. Sharf, Stability-constrained mobile manipulation planning on rough terrain, Robotica 40 (11) (2022) 4090–4119. doi:10.1017/S0263574722000777.

[11] L. Hulttinen, J. Mattila, Flow-limited path-following control of a double Ackermann steered hydraulic mobile manipulator, in: Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM), 2020, pp. 625–630. doi:10.1109/AIM43001.2020.9158987.

[12] E. G. Papadopoulos, D. A. Rey, A new measure of tipover stability margin for mobile manipulators, in: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Vol. 4, 1996, pp. 3111–3116. doi:10.1109/ROBOT.1996.509185.

[13] E. Papadopoulos, D. A. Rey, The force-angle measure of tipover stability margin for mobile manipulators, Vehicle System Dynamics 33 (1) (2000) 29–48. doi:10.1076/0042-3114(200001)33:1;1-5;FT029.

[14] O. M. U. Eraliev, K.-H. Lee, D.-Y. Shin, C.-H. Lee, Sensing, perception, decision, planning and action of autonomous excavators, Automation in Construction 141 (2022) 104428. doi:10.1016/j.autcon.2022.104428.

[15] Y. Fang, Y. K. Cho, J. Chen, A framework for real-time pro-active safety assistance for mobile crane lifting operations, Automation in Construction 72 (2016) 367–379. doi:10.1016/j.autcon.2016.08.025.

[16] P. Kwiatoń, D. Cekus, P. Waryś, Computer-aided rotary crane stability assessment, Automation in Construction 162 (2024) 105370. doi:10.1016/j.autcon.2024.105370.

[17] D. Jud, S. Kerscher, M. Wermelinger, E. Jelavic, P. Egli, P. Leemann, G. Hottiger, M. Hutter, HEAP – the autonomous walking excavator, Automation in Construction 129 (2021) 103783. doi:10.1016/j.autcon.2021.103783.

[18] G. R. Petrović, J. Mattila, Analytic solutions for wheeled mobile manipulator supporting forces, IEEE Access 10 (2022) 43235–43255. doi:10.1109/ACCESS.2022.3169766.

[19] A. Paz, M. Bahari, J. Mattila, Full-dynamics analytical modeling of normal forces for skid-steering mobile heavy-duty manipulators with actively articulated suspension, in: Proceedings of the IEEE Conference on Decision and Control (CDC), 2025, pp. 5419–5424. doi:10.1109/CDC57313.2025.11312016.

[20] S. Caron, Q.-C. Pham, Y. Nakamura, Leveraging cone double description for multi-contact stability of humanoids with applications to statics and dynamics, in: Proceedings of Robotics: Science and Systems, Rome, Italy, 2015. doi:10.15607/RSS.2015.XI.028.

[21] A. Paz, J. Mattila, Mechanic modeling and nonlinear optimal control of actively articulated suspension of mobile heavy-duty manipulators, in: Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM), 2025, pp. 1–8. doi:10.1109/AIM64088.2025.11175803.

[22] P. Fankhauser, M. Bloesch, M. Hutter, Probabilistic terrain mapping for mobile robots with uncertain localization, IEEE Robotics and Automation Letters 3 (4) (2018) 3019–3026. doi:10.1109/LRA.2018.2849506.

[23] J. M. Hollerbach, Dynamic scaling of manipulator trajectories, Journal of Dynamic Systems, Measurement, and Control 106 (1) (1984) 102–106. doi:10.1115/1.3149652.

[24] H. Pham, Q.-C. Pham, A new approach to time-optimal path parameterization based on reachability analysis, IEEE Transactions on Robotics 34 (3) (2018) 645–659. doi:10.1109/TRO.2018.2819195.

[25] O. Dahl, L. Nielsen, Torque-limited path following by online trajectory time scaling, IEEE Transactions on Robotics and Automation 6 (5) (1990) 554–561. doi:10.1109/70.62044.

[26] S. Lampinen, J. Niemi, J. Mattila, Flow-bounded trajectory-scaling algorithm for hydraulic robotic manipulators, in: Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM), 2020, pp. 619–624. doi:10.1109/AIM43001.2020.9158851.

[27] A. D. Ames, X. Xu, J. W. Grizzle, P. Tabuada, Control barrier function based quadratic programs for safety critical systems, IEEE Transactions on Automatic Control 62 (8) (2017) 3861–3876. doi:10.1109/TAC.2016.2638961.

[28] X. Xu, P. Tabuada, J. W. Grizzle, A. D. Ames, Robustness of control barrier functions for safety critical control, IFAC-PapersOnLine 48 (27) (2015) 54–61. doi:10.1016/j.ifacol.2015.11.152.

[29] J. Lei, M. G'Sell, A. Rinaldo, R. J. Tibshirani, L. Wasserman, Distribution-free predictive inference for regression, Journal of the American Statistical Association 113 (523) (2018) 1094–1111. doi:10.1080/01621459.2017.1307116.

[30] R. J. Tibshirani, R. F. Barber, E. Candès, A. Ramdas, Conformal prediction under covariate shift, in: Advances in Neural Information Processing Systems 32 (NeurIPS 2019), 2019, pp. 2530–2540.

[31] M. Heydari Shahna, J. Mattila, Synthesis of deep neural networks with safe robust adaptive control for reliable operation of wheeled mobile robots, IEEE Transactions on Automation Science and Engineering 23 (2026) 7456–7473. doi:10.1109/TASE.2026.3679545.

[32] M. Heydari Shahna, P. Mustalahti, J. Mattila, NMPC-augmented visual navigation and safe learning control for large-scale mobile robots, IEEE Robotics and Automation Letters 11 (4) (2026) 5182–5189. doi:10.1109/LRA.2026.3669802.

[33] M. Heydari Shahna, P. Mustalahti, J. Mattila, Robust vision-based goal-reaching control for mobile robots using a hierarchical learning framework, Robotics and Autonomous Systems 206 (2026) 105710. doi:10.1016/j.robot.2026.105710.

[34] T. Mattila, M. Heydari Shahna, P. Mustalahti, J. Mattila, Vehicle-specific configuration space distance fields for terrain-aware predictive control of heavy robots, Research Square preprint (2026). doi:10.21203/rs.3.rs-11401816/v1.

[35] M. Heydari Shahna, J. Kim, J. Mattila, RoboCompiler: graph-native compilation of closed-chain robots for consistent modeling, control, and simulation, software

repository, https://github.com/Mehdi-Heydari-Shahna/robocompiler, commit a4e4437 (2026).

[36] E. Todorov, T. Erez, Y. Tassa, MuJoCo: A physics engine for model-based control, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2012, pp. 5026–5033. doi:10.1109/IROS.2012.6386109.

[37] R. Featherstone, Rigid Body Dynamics Algorithms, Springer, Boston, MA, 2008. doi:10.1007/978-1-4899-7560-7.

[38] M. Bahari, P. Mustalahti, Á. Paz Anaya, J. Mattila, Electrification of heavy-duty robotic manipulators via PMSM-driven electromechanical linear actuators, IEEE Access 14 (2026) 21805–21832. doi:10.1109/ACCESS.2026.3662148.

[39] N. D. Manring, G. R. Luecke, Modeling and designing a hydrostatic transmission with a fixed-displacement motor, Journal of Dynamic Systems, Measurement, and Control 120 (1) (1998) 45–49. doi:10.1115/1.2801320.

[40] M. Heydari Shahna, P. Mustalahti, J. Mattila, Robust torque-observed control with safe input–output constraints for hydraulic in-wheel drive systems in mobile robots, Control Engineering Practice 164 (2025) 106459. doi:10.1016/j.conengprac.2025.106459.

[41] MuJoCo contributors, MuJoCo 3.15.0 documentation: Modeling, https://mujoco.readthedocs.io/en/3.15.0/modeling.html, accessed 28 October 2026.

[42] R. Krenn, A. Gibbesch, Soft soil contact modeling technique for multi-body system simulation, in: G. Zavarise, P. Wriggers (Eds.), Trends in Computational Contact Mechanics, Vol. 58 of Lecture Notes in Applied and Computational Mechanics, Springer, 2011, pp. 135–155. doi:10.1007/978-3-642-22167-5_8.

[43] M. G. Bekker, Introduction to Terrain-Vehicle Systems, University of Michigan Press, Ann Arbor, MI, 1969.

[44] J. Y. Wong, Theory of Ground Vehicles, 4th Edition, John Wiley & Sons, 2008.

[45] R. Schwan, Y. Jiang, D. Kuhn, C. N. Jones, PIQP: A proximal interior-point quadratic programming solver, in: Proceedings of the 62nd IEEE Conference on Decision and Control (CDC), 2023, pp. 1088–1093. doi:10.1109/CDC49753.2023.10383915.

[46] A. Mandow, J. L. Martínez, J. Morales, J. L. Blanco, A. García-Cerezo, J. González, Experimental kinematics for wheeled skid-steer mobile robots, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2007, pp. 1222–1227. doi:10.1109/IROS.2007.4399139.

[47] Siemens AG, Encoder systems for synchronous motors 1FT7 and 1FK7 G2; Migration to 1FT7/1FK7 G2 with DQI encoder, technical documentation, 17 August 2011, https://cache.industry.siemens.com/dl/files/520/55738520/att_30080/v1/encoders_and_migration_en.pdf.

Table A.1: Nominal framework parameters; named ablations modify the applicable mechanisms.

Layer	Parameter	Value
all	target proxy reserve $\theta_{\min}$	8°
	nominal, creep, unmapped-terrain speed	0.8, 0.15, $0.3\,\mathrm{m\,s^{-1}}$
	acceleration, deceleration limits	0.25, $0.5\,\mathrm{m\,s^{-2}}$
A	period, horizon, offset step, blend length	1 s, 18 m, 0.4 m, 6 m
	conformal level δ; observed share of an eligible candidate	0.05; 60 %
	switching margin m_{sw}, time gain; dwell between switches	1°, 10 % ($\geq$0.5 s); 6 m
	tie-breaking weights w_d, w_ζ, w_κ	$0.25\,\mathrm{s\,m^{-1}}$, $40\,\mathrm{s\,rad^{-2}}$, $0.5\,\mathrm{s\,m^2\,rad^{-1}}$
B	period, horizon $N \times \Delta s$; share of actuator speed limits ρ	0.2 s, 24×0.5 m; 0.9
	planning rates: slew; lift/tilt stroke	$0.25\,\mathrm{rad\,s^{-1}}$; $40.7/39.9\,\mathrm{mm\,s^{-1}}$
	planning rates: extension; levelling	$0.15\,\mathrm{m\,s^{-1}}$; $0.035\,\mathrm{rad\,s^{-1}}$
	synchronisation iterations, slow-down factor	4, 0.6
	surplus cap $t_{\max}$	6°
	tool clearance above mapped terrain (planner)	0.55 m
C	period; commanded-speed time constant τ_v	0.02 s; 0.4 s
	reference lead $\ell_0 + \tau_\ell v$; thresholds Λ_0, $\Lambda_{\max}$	0.5 m + 0.5 s v; 0.4, 2.0 s
	CBF look-ahead τ_c, decay γ; tool clearance (filter)	0.25 s, 0.5; 0.45 m

Appendix A. Parameters of the framework

Table A.1 lists the nominal parameters of the three layers.

Appendix B. Proofs of the conditional properties

Proof of Proposition 1. The hold trajectory $\boldsymbol{x}_k = \boldsymbol{x}_0$ satisfies the rate rows and the joint box. All reserve rows hold for a sufficiently negative t because t has no lower bound; sufficiently large σ satisfies the soft effort and clearance rows. Thus the QP is feasible. Its maximisation objective is concave, penalises unbounded slack and is bounded above by $t \leq t_{\max}$ and bounded posture variables; a maximiser exists. The optional negative quadratic regularisation of t preserves this conclusion. After the synchronisation loop, nonlinear model evaluations check each future reserve to tolerance ϵ_θ. A failing sample is capped at creep speed if that speed passes to the same tolerance, otherwise the profile is zero from five samples earlier. Subsequent passes only decrease speeds. Because $\Delta\theta$ is nondecreasing in speed and available actuator time grows as speed decreases, these operations preserve passed sample checks. Induction on the backward acceleration pass from the zero-speed sample gives $v_k^2 \leq 2a_{\mathrm{dec}}(s_{\mathrm{stop}} - s_k)$. This is a property of the returned profile. □

Proof of Proposition 2. For $\Lambda \geq \Lambda_{\max}$ the desired speed is zero, hence $\dot{v}_{\mathrm{cmd}} \leq -2a_{\mathrm{dec}}$ above $v_c = 2a_{\mathrm{dec}}\tau_v$ and $\dot{v}_{\mathrm{cmd}} \leq -v_{\mathrm{cmd}}/\tau_v$ below it. Integrating the two phases gives $(v_1^2 - v_c^2)/(4a_{\mathrm{dec}})$ when $v_1 > v_c$, followed by at most $\tau_v \min(v_1, v_c)$, which yields d_b. The posture inequality follows directly from the definition of Λ and the path-indexed reference. The static-reserve inequality follows from the triangle inequality and Lipschitz continuity in posture and path coordinate, with support attitude held fixed. □

Proof of Proposition 3. The reserve row gives $\hat{\Delta}_{e,i} \geq -\gamma\beta_{e,i} - \varsigma_i$. Substitution into the assumed update yields $\beta_{e,i+1} \geq q\beta_{e,i} - T_{\mathrm{C}}(\bar{w} + \varsigma_i/\tau_c)$. Iterating proves (19); setting all slacks to zero and summing the geometric series gives the stated corollary. If $\beta_{e,0} \geq -\bar{w}/\alpha$, that corollary remains above $-\bar{w}/\alpha$ at the sample times. □

Appendix C. Plant-variation design

The 16 Latin-hypercube draws of the plant-variation study (Section 5.3; scrambled, seed 20261028) vary 13 plant parameters jointly; nominal values in parentheses. Slew drive: reduction ratio 450–750 (600), speed-loop bandwidth 2–8 Hz (4 Hz), position gain 2.5–10 s^{-1} (5 s^{-1}), holding-brake capacity 0–4 times the peak drive torque (2) and brake engagement delay 0–0.6 s (0.2 s); quantities derived from the slew drive use the varied ratio. Tyres: radial stiffness 0.6–1.8 MN m^{-1} (1.2) and damping ratio 0.03–0.30 (0.10). Soil: pressure coefficients, shear-displacement length, unloading stiffness and damping at 0.5–2 times nominal, and friction angle 32°–44° (38°). Ground: effective rigid tyre–ground friction 0.40–0.80 (0.60), set by lowering the robot-side coefficients as well, because the contact model uses the larger of the two coefficients. The soft corner takes the lower end of every range except the soil shear-displacement length (upper end); the stiff corner takes the upper ends of the slew ratio, slew position gain, brake capacity, tyre stiffness, soil coefficients, friction angle and ground friction, no brake delay, and the lower ends of the slew bandwidth and the damping ratios. The nominal row, the draws and the corners are run with seeds $8 + (i \bmod 8)$ in this order ($i = 0, 1, \dots$). The motor-family anchor uses the documented rotor inertia and holding torque of the slew motor with brake and adds a 0.05 s closing interval and a 0.10 s release delay from a variant datasheet of the motor family to the 0.20 s dwell (sources in the parameter audit of the code).

Appendix D. Results per scenario

Table D.1 gives the results of the main study by scenario and payload.

Table D.1: Main study by scenario and payload (eight seeds each); columns as in Table 5 (Completed: reached the goal; 45°: termination when roll or pitch exceeds 45°; Stop: nonprogress stop or time limit; Proxy events: mean number per episode; Unloading: mean number of wheel-unloading events per episode; θ_{acc}: mean 1st percentile of the acceleration-based reserve proxy [deg]; θ_{st}: mean smallest static reserve [deg]; Time: mean driving time of completed episodes [s]).

Controller	Completed	45°	Stop	Proxy events	Unloading	θ_{acc}	θ_{st}	Time
S1 embankment, empty								
FIX	8/8	0	0	0.00	11.1	10.0	13.8	78.2
SLOW	8/8	0	0	0.00	8.8	11.0	14.1	86.0
REACT	7/8	0	1	0.00	13.5	14.4	20.8	78.1
REACT-S	7/8	0	1	0.00	14.4	14.6	20.8	78.1
PREV	8/8	0	0	0.00	8.8	15.5	21.3	79.7
COMPASS	8/8	0	0	0.00	12.2	14.6	19.6	82.7
S1 embankment, 400 kg								
FIX	8/8	0	0	0.12	6.6	8.1	11.4	78.3
SLOW	8/8	0	0	0.00	8.4	7.9	10.5	95.7
REACT	8/8	0	0	0.38	6.2	9.5	12.3	78.2
REACT-S	8/8	0	0	0.25	6.0	11.0	13.1	79.4
PREV	8/8	0	0	0.00	4.9	14.7	19.9	79.8
COMPASS	8/8	0	0	0.00	5.8	13.9	18.8	82.5
S2 crest, empty								
FIX	7/8	1	0	0.25	9.9	4.9	6.5	53.3
SLOW	6/8	2	0	1.00	10.6	2.1	2.4	73.3
REACT	8/8	0	0	0.00	13.9	12.3	16.7	53.2
REACT-S	8/8	0	0	0.12	15.1	11.9	16.5	53.4
PREV	8/8	0	0	0.50	10.1	13.1	17.1	58.5
COMPASS	8/8	0	0	0.00	6.0	14.7	18.8	61.3
S2 crest, 400 kg								
FIX	7/8	1	0	2.00	9.2	-1.6	-2.6	53.3
SLOW	3/8	5	0	1.12	5.4	-4.5	-9.6	88.7
REACT	8/8	0	0	0.12	5.1	13.8	18.2	52.9
REACT-S	8/8	0	0	0.38	5.2	12.7	17.2	53.4
PREV	7/8	1	0	0.38	1.1	12.5	15.9	54.9
COMPASS	8/8	0	0	0.00	1.1	16.2	19.3	60.0
S3 ditch, empty								
FIX	8/8	0	0	0.00	6.5	15.0	23.9	55.6
SLOW	8/8	0	0	0.00	6.5	15.0	23.9	55.6
REACT	8/8	0	0	0.00	10.1	15.9	25.3	56.0
REACT-S	8/8	0	0	0.00	10.1	15.9	25.3	56.0
PREV	8/8	0	0	0.00	5.1	18.2	27.4	57.2
COMPASS	8/8	0	0	0.00	4.5	18.1	27.3	59.5
S3 ditch, 400 kg								
FIX	8/8	0	0	0.00	6.8	12.3	19.5	55.6
SLOW	8/8	0	0	0.00	6.9	12.4	19.4	55.7
REACT	8/8	0	0	0.00	4.8	16.0	25.3	55.6
REACT-S	8/8	0	0	0.00	4.8	16.0	25.3	55.6
PREV	8/8	0	0	0.00	2.2	17.8	26.6	57.2
COMPASS	8/8	0	0	0.00	2.4	18.0	26.5	59.6



Mehdi Heydari Shahna received the D.Sc. (Tech.) degree with distinction in Automation Science and Engineering from Tampere University, Finland, in 2025. He is currently a postdoctoral researcher at Tampere University. His research interests include physical AI, robot learning, safe control, and robust control.



Jouni Mattila received a Ph.D. in automation engineering from Tampere University of Technology, Tampere, Finland, in 2000. He is currently a professor of machine automation with the Unit of Automation Technology and Mechanical Engineering at Tampere University. His research interests include machine automation, nonlinear-model-based control of robotic manipulators, and energy-efficient control of heavy-duty mobile manipulators.