Systemy Logistyczne Wojsk
Zeszyt 65(2026)
s. 25-48

Instytut Logistyki
Wydział Bezpieczeństwa, Logistyki i Zarządzania
Wojskowa Akademia Techniczna w Warszawie

Military Logistics Systems
Volume 65(2026)
pp. 25-48

Institute of Logistics
Faculty of Security, Logistics and Management
Military University of Technology in Warsaw

DOI: 10.37055/slw/242078
Article type: Original research article

ISSN 1508-5430
eISSN 2719-7689

EMPIRICAL QUANTIFICATION OF THE MOBILITY COST OF MASS GROWTH IN MILITARY VEHICLES: IMPLICATIONS FOR FLEET READINESS AND CONFIGURATION MANAGEMENT

Przemysław Simiński
przemyslaw.siminski@witpis.pl; ORCID: 0000-0002-2323-3152
Military Institute of Armoured and Automotive Technology, Sulejówek, Poland

Abstract.
Mass growth during the service life of military vehicles, caused by add-on armour, mission equipment and protection upgrades, is usually treated as a structural and payload problem, while its effect on mobility and fleet readiness is rarely quantified from measured data. This paper provides an empirical, measurement-based assessment of the mobility cost of mass growth for three Polish Armed Forces platforms: STAR 266, STAR 944 and JELCZ 442.32. Using traction-dynamics characteristics obtained at WITPiS, the inverse scaling of the dynamic factor with mass is treated as an analytical premise, against which three hypotheses are tested: that platform-independent relative loss does not translate into operationally meaningful quantities; that the mass–gradeability relation is piecewise because of the switch between adhesion- and traction-limited regimes; and that deterministic compliance does not ensure mobility readiness. A 30% mass increase reduces the maximum dynamic factor by 23.1% for every platform, but its operational effects differ sharply: one vehicle falls below the 25% march requirement, while others retain substantial margins. A Monte Carlo simulation (N = 20,000), using terrain and mass-growth distributions, estimates the probability of retaining required mobility after up-fitting and shows that a platform with a comfortable deterministic margin retains the mobility function in only about nine cases in ten. The results are verified against source characteristics rather than independent field trials. The study reframes the mass budget as a configuration-management and fleet-readiness instrument linking vehicle dynamics to operational availability and acquisition decision support.

Keywords:
military vehicle mobility; mobility readiness; configuration management; mass growth; Monte Carlo simulation

Introduction

Military wheeled vehicles rarely remain in the configuration in which they entered service. Across their life cycle they accumulate mass: appliqué and spaced armour, mine-protection kits, remote weapon stations, electronic counter-measures, command-and-control suites and additional stowage are added in successive modernisation and force-protection programmes. Each increment improves survivability or capability, yet every kilogram is also carried by the powertrain, the running gear and, ultimately, by the logistic system that must keep the platform moving. The tension between protection and mobility is one of the defining trade-offs of modern land systems, and it is increasingly a sustainment problem rather than a purely engineering one.

From a logistics and readiness perspective, mobility is not an attribute of an individual vehicle but a function that the transport and manoeuvre fleet must deliver on demand. When mass growth erodes that function, what degrades is the operational availability of mobility across the fleet: marches take longer, gradients that were previously trafficable become marginal, and the proportion of the fleet able to meet a given tactical requirement falls. In the language of integrated logistic support, a configuration change has altered a key performance parameter, and the consequence propagates into readiness reporting, march planning and the resilience of the distribution and manoeuvre network. The mass budget, which is often handled as an engineering accounting device, is therefore better understood as a configuration-management control with direct readiness implications.

Despite the operational importance of this link, the relationship between mass growth and mobility is frequently treated qualitatively. It is common to assert that a heavier vehicle is slower and climbs less steeply, or to postulate a scaling law without anchoring it in measured data. Two gaps follow. First, the quantitative cost of a given mass increment is rarely derived from the actual traction-dynamics characteristics of fielded platforms. Second, even where a deterministic figure is offered, it is seldom propagated through the variability of real terrain and real up-fitting decisions to a readiness probability that a sustainment or acquisition authority could use.

This paper addresses both gaps for three platforms of the Polish Armed Forces. It is built on measured traction-dynamics characteristics. A distinction is drawn at the outset between the analytical premise on which the analysis rests and the hypotheses that are actually put at risk by it. The premise is stated as such and is not claimed as a result; the three hypotheses that follow are formulated so that each could have been refuted by the measured baselines or by the simulation, and none is a restatement of the definitions adopted in Section 3:

P1. The relative loss of dynamic performance with mass growth follows a single, platform-independent relation, because the dynamic factor is a surplus force normalised by weight. This is an elementary consequence of the definition adopted

in Section 3.2 and is treated throughout as a premise, not as a finding; its role is to establish the reference against which H1 to H3 are non-trivial.

H1. The platform-independence asserted in P1 does not propagate to any operationally meaningful quantity. Neither the loss of attainable gradeability nor the increase in acceleration time is platform-independent, because both depend on mass through further channels — the baseline margin against the requirement, the binding traction or adhesion regime, and the rotating-mass factor — that are platform-specific. Relative degradation of the dynamic factor is therefore not an admissible criterion for judging the operational acceptability of an up-fit.

H2. The mapping from mass growth to attainable gradeability is not smooth but piecewise, because gradeability is bounded simultaneously by traction and by adhesion. A platform whose baseline performance is limited by grip rather than by tractive capability will therefore appear robust to mass growth until the binding constraint changes, so that apparent baseline robustness may be an artefact of adhesion saturation and part of the true erosion of its reserve is concealed.

H3. Deterministic compliance at reference conditions is not a sufficient condition for mobility readiness. Because attainable gradeability is a saturating, non-linear function of the terrain parameters, the readiness probability is not a monotone image of the deterministic margin: a platform that clears the march requirement deterministically by an apparently comfortable margin may lose the required mobility function in a materially non-negligible fraction of representative operating conditions.

The contribution is fourfold. First, the mobility cost of mass growth is quantified directly from measured characteristics of fielded platforms rather than postulated, using data that already exist for every accepted platform and therefore without new instrumentation. Second, it is shown that the universal relative law of P1, although correct, is operationally uninformative on its own, because identical relative degradation produces opposite operational verdicts depending on the baseline margin and on the binding regime. Third, the deterministic result is extended into a stochastic readiness indicator through a Monte Carlo simulation over representative terrain and up-fitting distributions, and that indicator is shown to diverge materially from the deterministic verdict. Fourth, the mass budget is reinterpreted as a configuration-management and fleet-readiness instrument, providing a quantitative bridge between an elementary engineering relation and decisions about operational availability, march planning and acquisition. The remainder of the paper reviews the relevant traction-dynamics and logistics literature, sets out the model and the stochastic method, presents the deterministic and probabilistic results for the three platforms, and discusses their implications for readiness and configuration management.

State of Research

The analysis draws on two bodies of knowledge that are seldom connected: the engineering of vehicle traction dynamics and terramechanics, and the discipline of military logistics, reliability and configuration management.

Vehicle traction dynamics and terramechanics

The dynamic-factor method used here is a standard tool of automotive performance analysis. It expresses the surplus tractive force available for acceleration and grade climbing, normalised by vehicle weight, as a function of speed; the classic treatments are given by Gillespie (1992), Wong (2008), Jazar (2017) and Genta and Morello (2009), with vehicle-dynamics control aspects in Rajamani (2012). The dependence of grade-climbing and acceleration capability on vehicle mass is elementary in this framework, since the dynamic factor is the surplus force divided by weight; this point is made implicitly throughout the standard texts and is not, in itself, novel. For off-road performance the governing discipline is terramechanics, founded by Bekker (1969) and developed extensively by Wong (2010), Wong and Huang (2006) and Maclaurin (2007), which links soil strength and running-gear geometry to drawbar pull and trafficability. Recent work has advanced high-fidelity and physics-based simulation of off-road mobility, including the studies of Wasfy and Jayakumar (2021a, 2021b), McCullough et al. (2017), He et al. (2019) and Gonzalez and Iagnemma (2018), while NATO has codified standard mobility-assessment methods through the work reported in NATO STO (2020) and the modelling and simulation standards of AMSP-06 (2021). These sources establish how mobility is computed; they do not, however, translate a mobility loss into a fleet-level readiness statement.

Logistics, reliability and configuration management

The second strand concerns the management of materiel readiness across the life cycle. Integrated logistic support and life-cycle systems engineering are treated comprehensively by Blanchard (2004) and in the integrated logistic support handbook of Jones (2006), which frame an item's performance as one input to its supportability and availability. The reliability-engineering literature, in particular O'Connor and Kleyner (2012), formalises the link between component-level behaviour and the operational availability of a function, a concept that is directly applicable when a configuration change degrades a performance parameter rather than causing an outright failure. At the doctrinal level, NATO logistics doctrine and the sustainment

doctrine of the United States Army, ADP 4-0 (US Army 2019), define readiness and operational availability as the central objects of the sustainment warfighting function. Configuration management, including the control of mass and the mass budget, is the mechanism through which such changes are governed across a fleet.

The gap addressed in this paper lies at the intersection of these two strands. Mobility engineering produces deterministic performance figures; logistics and readiness management require statements about the probability that a fleet retains a required function under representative conditions. Bridging the two means expressing the mobility cost of a configuration change as a readiness indicator that is usable in march planning, configuration control and acquisition decision support. That bridge, anchored in measured data and made stochastic, is the contribution of the present work.

Materials and Methods
Platforms and measured data

Three representative platforms of the Polish Armed Forces transport and support fleet were analysed: the STAR 266 (6×6), the STAR 944 (4×4) and the JELCZ 442.32 (4×4). For each platform the traction-dynamics characteristics, that is the dynamic factor as a function of speed in every gear, were obtained from the computational and measurement characteristics held at WITPiS (WITPiS 2024). The principal vehicle parameters used in the analysis are listed in Table 1. The remaining inputs, in particular the full engine torque maps and the per-gear dynamic-factor curves, are taken directly from the source characteristics; the dynamic-factor curves are therefore treated as the empirical input to the analysis rather than re-derived from first principles.

Table 1. Principal parameters of the three analysed platforms

Parameter	STAR 266	STAR 944	JELCZ 442.32
Engine designation	S-359	MAN D0824	MTU 6R106TD21
m_0 — baseline mass (kg)	12,500	9,700	16,000
Number of driven wheels	6	4	4
r_d — dynamic radius (m)	0.527	0.527	0.578
Frontal area A (m^2)	6.21	6.27	6.05
c_x — drag coefficient (–)	0.78	0.78	0.78
i_p — final-drive ratio (–)	6.33	5.937	6.000
i_r — transfer ratio (–)	1.000	1.061	1.000
Driveline efficiency η (–)	0.85	0.87	0.85

Table 1 (continued)

Parameter	STAR 266	STAR 944	JELCZ 442.32
Number of forward gears	5	6	9
D_{max} — baseline maximum dynamic factor (–)	0.305	0.414	0.693

Source: own elaboration based on WITPiS (2024)

Dynamic factor and its dependence on mass

The tractive force at the driven wheels follows from the engine torque Me transmitted through the gearbox ratio ig, the final-drive ratio ip, the transfer ratio ir and the driveline efficiency ηm to the dynamic radius rd:

$$F_t = (M_e \cdot i_g \cdot i_p \cdot i_r \cdot \eta_m)/r_d \quad (1)$$

The aerodynamic resistance is Fair, a function of the drag coefficient, frontal area, air density ρ and the square of speed:

$$F_air = 1/2\ \rho\ c_x\ A\ v\text{^}2 \quad (2)$$

The dynamic factor D is the surplus of tractive over aerodynamic force, normalised by vehicle weight G = m g:

$$D = (F_t - F_air)/(m\ g) \quad (3)$$

Because the numerator of Equation (3) is, to first order, independent of mass while the denominator is proportional to it, the dynamic factor at any speed scales inversely with mass. Denoting the baseline mass and dynamic factor by m_0 and D_0, and the up-fitted values by a prime, the relation is:

$$D'(v) = D_0\ (v) \cdot\ m_0/m' \quad (4)$$

3.3 The universal relative law

Let the mass-growth fraction be defined as

$$\varphi = (m' - m_0)/m_0 \quad (5)$$

so that $m' = m_0(1 + \varphi)$. Substituting into Equation (4) and evaluating at the speed of maximum dynamic factor gives the relative law adopted as premise P1:

$$[\![D']\!]_max/D_(0,max) = m_0/m' = 1/(1 + \varphi) \quad (6)$$

Equation (6) is platform-independent: it contains no vehicle-specific parameter. It is, deliberately, an elementary consequence of the definition of the dynamic factor, and it is presented as such: P1 is a premise of the analysis and is not offered

as a result. Its analytical status is precisely what gives the hypotheses their content, since H1 to H3 all assert that this exact and universal relation fails to determine the operational outcome. The novelty of the paper lies not in the relation but in its empirical anchoring in fielded platforms and in the operational and stochastic consequences developed below.

Gradeability and the two-regime limit

On a steady climb at low speed, neglecting acceleration and aerodynamic terms, the dynamic factor is balanced by the gradient and the rolling resistance with coefficient f, so that for a slope angle α:

$$D = \sin\alpha + f\cos\alpha \tag{7}$$

The gradeability is $i = \tan\alpha$. The force that the tyres can transmit is, however, bounded by adhesion: the attainable longitudinal force cannot exceed μ G, where μ is the adhesion coefficient. The attainable gradeability is therefore the lesser of the traction-limited and the adhesion-limited values:

$$i = \tan[\min(\alpha_D, \arctan\mu)] \times 100\% \tag{8}$$

where αD solves Equation (7). The reference conditions used for the deterministic comparison are a firm, well-maintained road surface with f = 0.012 and an adhesion ceiling μ = 0.75, representative of a dry, hard road. Under these conditions a platform is traction-limited when its dynamic factor is the binding constraint and adhesion-limited when grip is the binding constraint; the regime is reported explicitly for every case, because it is decisive for the interpretation of the results.

Rotating masses and acceleration

For acceleration the inertia of the rotating masses must be included through the rotating-mass factor δ. Because added armour and equipment are non-rotating, δ decreases as mass grows; it is therefore evaluated per gear and per mass level:

$$\delta_g(m) = 1 + (I_e\ (i_g\, i_p\, i_r)\text{^}2\ \eta + n_w\ I_w)/(m\ [\![r_d]\!]\text{^}2) \tag{9}$$

where Ie is the engine and transmission moment of inertia, Iw the moment of inertia of one wheel and nw the number of wheels. The time to accelerate from rest to a target speed V follows from integrating the reciprocal of the instantaneous acceleration:

$$t_(0\rightarrow V) = \int_0\text{^}V' dv/(a(v)), a(v) = g[D(v) - f] / \delta_g \tag{10}$$

The per-gear accelerations reproduce the source characteristics exactly. For the absolute acceleration times the baseline values are anchored to the source characteristics and the effect of mass growth is applied through the mass-dependent factor of Equation (9), so that the reported time increases reflect both the reduced surplus force and the changing rotating-mass factor.

3.6 Stochastic readiness model

The deterministic analysis fixes a single mass increment and a single surface. Real up-fitting decisions and real terrain are variable, and readiness is intrinsically a probabilistic property of a fleet operating over a range of conditions. To capture this, a Monte Carlo simulation was performed for a representative up-fitting programme in which each platform receives a mass increase centred on 30%.

In each of N = 20,000 trials the mass-growth fraction was drawn from a truncated normal distribution with mean 0.30 and standard deviation 0.05, bounded to the interval [0.15, 0.45]. The operating surface was drawn from a three-component terrain mixture: firm road (probability 0.60, rolling resistance f ∈ [0.012, 0.025], adhesion μ ∈ [0.65, 0.80]); hard unpaved ground (probability 0.30, f ∈ [0.030, 0.060], μ ∈ [0.55, 0.70]); and soft ground (probability 0.10, f ∈ [0.080, 0.150], μ ∈ [0.35, 0.50]). For each trial the attainable gradeability was computed from the realised mass, rolling resistance and adhesion using Equations (4), (7) and (8). A trial is a success when the attainable gradeability meets or exceeds the march requirement ireq = 25%. The mobility-readiness indicator is the success probability:

$$\Phi = P(i \geq i_req) = 1/N \sum_(k = 1)\text{^}N' [i_k \geq i_req] \quad (11)$$

The square-bracket term is the Iverson indicator, equal to one when the condition holds and zero otherwise. The simulation used a fixed random seed and is fully reproducible. Confidence intervals for Φ are reported using the Wilson score interval, which is appropriate for proportions near the boundaries of the unit interval; where no failures were observed, a one-sided interval is reported rather than a point value of 100%.

Methodological note

The procedure is a verification of internal consistency against the source traction-dynamics characteristics, not an independent experimental validation. All deterministic outputs were checked for exact agreement with the per-gear characteristics from which they derive, and the stochastic layer is a transparent propagation of documented distributions. Independent field measurement of the up-fitted platforms remains necessary to confirm the absolute figures; the validation programme required for that purpose is specified in Section 5.4 and the limitations of the present evidence are set out there.

Results and analysis
Baseline dynamic characteristics

The baseline dynamic-factor envelopes of the three platforms are shown in Figure 1. The maximum dynamic factor, attained in first gear, is 0.305 for the STAR 266, 0.414 for the STAR 944 and 0.693 for the JELCZ 442.32. The ordering reflects the installed surplus power and the gearing of each platform: the JELCZ, with the most powerful engine and the deepest first gear, has by far the largest reserve of tractive capability, whereas the STAR 266 has the smallest. As will be seen, these baseline values, and the margins they imply against the march requirement, are decisive for the operational interpretation of mass growth.

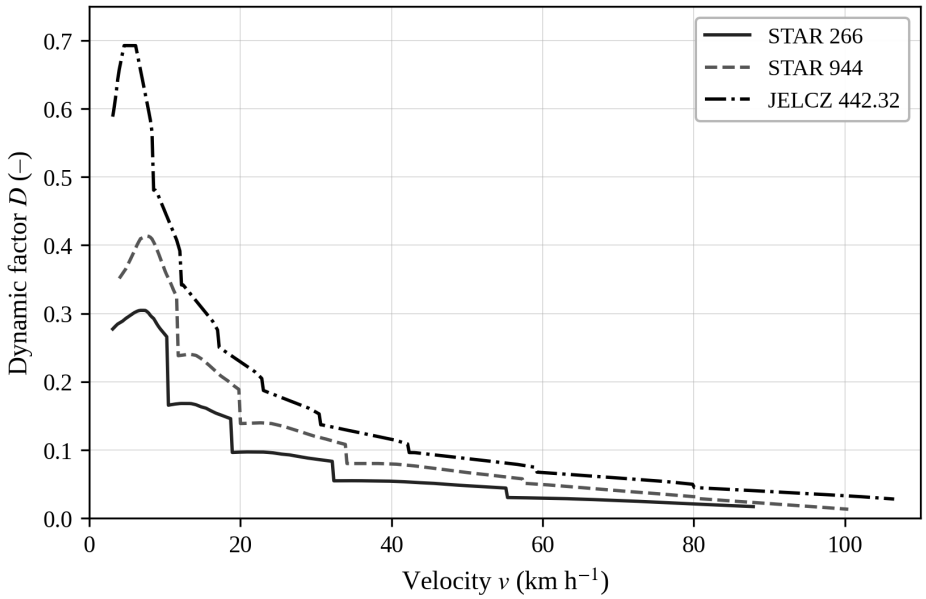


Figure 1. Dynamic-factor envelopes D(v) of the three baseline platforms
Source: own elaboration based on WITPiS (2024)

The universal relative relation as an analytical premise (P1)

The relative effect of mass growth on the maximum dynamic factor is identical for all three platforms, as premise P1 requires. A 15% mass increase reduces the maximum dynamic factor to 87.0% of its baseline value, a 30% increase to 76.9%, and a 45% increase to 69.0%. Equivalently, a 30% mass increase reduces the maximum dynamic factor by 23.1%, and a 45% increase by 31.0%, on every platform without

exception. The left-hand panel of Figure 2 shows the three platforms collapsing exactly onto the single curve $1/(1 + \varphi)$, which is the graphical statement of P1. The relevant values are listed in the relative-factor column of Table 2.

Two qualifications must be attached to this result, and they are what the remainder of the analysis is about. The first is that the collapse of the three platforms onto a single curve is not an empirical discovery: the up-fitted values in Table 2 are obtained by applying Equation (4) to the measured baseline characteristics, so the relative-factor column reproduces P1 by construction and could not have done otherwise. What is empirical in Figure 2(a) is the set of baseline values on which the curve is anchored — 0.305, 0.414 and 0.693 — and it is these, not the scaling, that carry the information. The second is that the universality is operationally empty: it states the relative penalty without reference to the level from which a platform starts, and that level is exactly what determines whether the platform remains fit for its task. The next subsection makes both points explicit and is where H1 and H2 are tested.

Non-propagation of the universal relation and the role of the binding regime (H1, H2)

When the same mass growth is expressed as attainable gradeability under the reference conditions, the platforms diverge sharply. The right-hand panel of Figure 2 plots gradeability against mass increase, together with the 25% march requirement and the adhesion ceiling. The deterministic results are collected in Table 2.

Table 2. Deterministic effect of mass growth on the dynamic factor, gradeability and acceleration

Platform	φ (%)	m′ (kg)	D_max (–)	D_max/D₀ (%)	i (%)	Regime	t_0' ₅₀ (s)
STAR 266	0	12,500	0.305	100.0	30.7	traction	18.3
	15	14,375	0.265	87.0	26.2	traction	21.4
	30	16,250	0.234	76.9	22.8	traction	24.7
	45	18,125	0.210	69.0	20.2	traction	28.2
STAR 944	0	9,700	0.414	100.0	44.0	traction	11.9
	15	11,155	0.360	87.0	37.2	traction	13.7
	30	12,610	0.318	76.9	32.2	traction	15.5
	45	14,065	0.285	69.0	28.5	traction	17.4

Table 2 (continued)

Platform	φ (%)	m′ (kg)	D_max (–)	D_max/D_0 (%)	i (%)	Regime	$t_0'_{50}$ (s)
JELCZ 442.32	0	16,000	0.693	100.0	75.0	adhesion	8.4
	15	18,400	0.602	87.0	73.6	traction	9.4
	30	20,800	0.533	76.9	61.3	traction	10.4
	45	23,200	0.478	69.0	52.8	traction	11.4

Source: own elaboration; reference conditions f = 0.012, μ = 0.75

Three features of Table 2 are decisive. First, the relative-factor column is identical across platforms, as required by H1. Second, the gradeability column is not. At a 30% mass increase the STAR 266 falls from 30.7% to 22.8%, a loss of 25.5% of its baseline gradeability that takes it below the 25% march requirement; the STAR 944 falls from 44.0% to 32.2%, a loss of 26.7% that nonetheless leaves a comfortable margin; and the JELCZ falls from 75.0% to 61.3%, an apparent loss of only 18.3%. Third, the regime column explains the JELCZ figure: at baseline the JELCZ is adhesion-limited, its enormous traction reserve capped at the 75% grip ceiling, so part of the underlying loss is masked until the platform drops into the traction-limited regime above a modest mass increase.

These results confirm H2 and refute any uniform characterisation of the penalty. It is not correct to say that mass growth reduces gradeability by a fixed fraction such as one third; the operational loss is platform-specific and is governed by the baseline margin and the binding regime, not by the universal relative law. Identical relative degradation produces opposite verdicts: the same 30% mass increase removes the required mobility function from the STAR 266 while leaving the STAR 944 and the JELCZ fit for task. The operationally meaningful question is therefore not how much the dynamic factor falls in relative terms, which is always 23.1% at 30% growth, but whether the platform retains enough absolute capability to meet the requirement, which depends entirely on where it started.

The acceleration results in the final column of Table 2 reinforce the picture. A 30% mass increase lengthens the time to reach 50 km h'^{1} by 34.8% for the STAR 266, 30.1% for the STAR 944 and 23.4% for the JELCZ, a range of approximately 23% to 35%. The spread again follows the surplus-power reserve, and the rotating-mass factor of Equation (9), which falls as non-rotating mass is added, modestly tempers the increase for the heavier configurations.

Taken together, these two columns test H1 directly. If the platform-independence of P1 propagated to the operational quantities, the loss of gradeability and the increase in acceleration time would be the same for all three platforms, because their relative loss of dynamic factor is exactly the same. It is not. At a 30% mass increase the loss of attainable gradeability spans 18.3% to 26.7% across the three platforms, and the increase in the time to 50 km h'^{1} spans 23.4% to 34.8%, against a dynamic-factor

loss that is 23.1% for every platform without exception. The dispersion is not noise: it is produced by three identifiable and platform-specific channels — the baseline margin, the binding regime and the rotating-mass factor of Equation (9) — none of which appears in Equation (6). H1 is therefore confirmed, and the practical consequence is that a relative degradation figure quoted for the dynamic factor cannot be used as a criterion for accepting or rejecting an up-fit.

The JELCZ figures test H2. This platform is the only one that is adhesion-limited at baseline: its dynamic factor of 0.693 would support a far steeper climb than the grip ceiling permits, so its baseline gradeability of 75.0% is set by μ rather than by the powertrain. The consequence is that its first increment of mass costs it almost nothing in gradeability — 75.0% to 73.6% between φ = 0 and φ = 0.15 — even though its dynamic factor has already fallen by the full 13.0% that P1 prescribes. The loss is not absent; it is concealed, absorbed by the saturated branch of Equation (8) until the platform crosses into the traction-limited regime. Beyond that crossing the same platform loses 12.3 percentage points of gradeability over the next 15% of mass growth. H2 is confirmed: the mapping from mass to gradeability is piecewise, and a platform that appears insensitive to mass growth may simply be spending a reserve that is not visible in the quantity being observed. For configuration management this is a practical warning, because it is precisely the platform with the largest apparent margin whose early up-fits look free of charge.

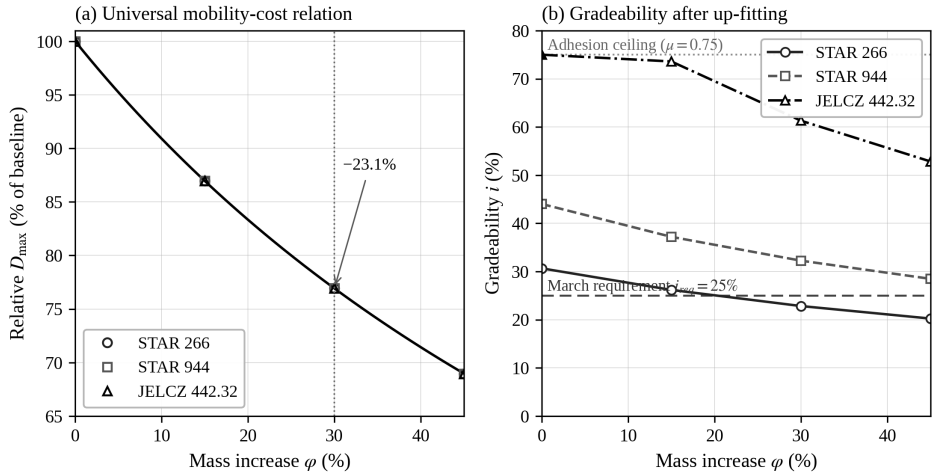


Figure 2. Universal relative degradation of the dynamic factor (a) and absolute gradeability against mass growth with the march requirement and adhesion ceiling (b)

Source: own elaboration based on WITPiS (2024)

Stochastic mobility readiness and the insufficiency of deterministic compliance (H3)

The deterministic analysis fixes a single surface and a single mass increment. The Monte Carlo simulation of Section 3.6 replaces both with representative distributions and yields, for each platform, the probability of retaining the required mobility after a nominal 30% up-fit. The resulting distributions of attainable gradeability are shown in Figure 3, and the summary statistics are given in Table 3.

Table 3. Stochastic mobility readiness after a 30% mass increase (Monte Carlo, N = 20,000)

Platform	Median i (%)	P5 i (%)	P95 i (%)	$\Phi = P(i \geq 25\%)$ (%)	95% CI for Φ (%)	Failures (of 20,000)
STAR 266	21.3	12.2	23.7	0.24	[0.2, 0.3]	19,951
STAR 944	30.5	21.1	33.7	90.24	[89.8, 90.6]	1,952
JELCZ 442.32	58.6	41.8	65.4	> 99.98	≥ 99.98 (one-sided)	0

Source: own elaboration; seed fixed for reproducibility, Wilson score confidence intervals

The three platforms occupy three distinct readiness regimes. The STAR 266 fails the march requirement in all but a small fraction of conditions: its mobility-readiness indicator is only 0.24% (49 successes in 20,000 trials, Wilson interval [0.2, 0.3]), and even its 95th-percentile gradeability, attained only on the best surfaces, barely reaches the requirement. The STAR 944 retains the function in the great majority of conditions, with a readiness indicator of 90.2% (interval [89.8, 90.6]); its failures are confined to the soft-ground tail of the terrain mixture. The JELCZ retains the function in every one of the 20,000 trials; because no failure was observed, its readiness is reported as greater than 99.98% with a one-sided 95% confidence bound, rather than as an unqualified 100%, which would overstate the statistical evidence.

The displacement of the JELCZ median to 58.6%, below its deterministic firm-road value of 61.3%, is not an inconsistency but a direct consequence of averaging over the terrain mixture: the simulation samples surfaces less favourable than the reference road, which lowers the median while leaving the platform comfortably above the requirement. The value of the stochastic layer is precisely that it exposes the shape of the distribution, and in particular the left tail, which the deterministic figure conceals. For the STAR 266 the left tail is the operational reality; for the JELCZ it is irrelevant. The readiness indicator Φ compresses this distinction into a single number that a sustainment or acquisition authority can act upon.

The STAR 944 is the decisive case for H3. Evaluated deterministically at the reference conditions, this platform clears the march requirement after a 30% up-fit with an attainable gradeability of 32.2% against a requirement of 25%, a margin of

7.2 percentage points or 28.8% of the requirement itself. On a deterministic march table it is unambiguously fit for task. Evaluated over the terrain mixture, it retains the mobility function in 90.24% of trials [89.8, 90.6]: roughly one movement in ten, on a representative mixture of surfaces, would not meet the requirement. The deterministic margin and the readiness probability are therefore not monotonically related in any usable sense, and the discrepancy is not a rounding effect but an order-of-magnitude difference in the residual risk that a planner would infer. H3 is confirmed. The mechanism is the min operator of Equation (8) acting on a distribution rather than on a point: because the poor-surface component of the mixture depresses both the traction-limited and the adhesion-limited branches, the left tail of the gradeability distribution is far heavier than a symmetric error band around the deterministic value would suggest. The practical implication is that a mobility requirement expressed as a single deterministic gradeability figure, as such requirements normally are in technical specifications, does not by itself assure the readiness that the requirement is intended to secure; the surface distribution against which it is to be met has to be specified with it.

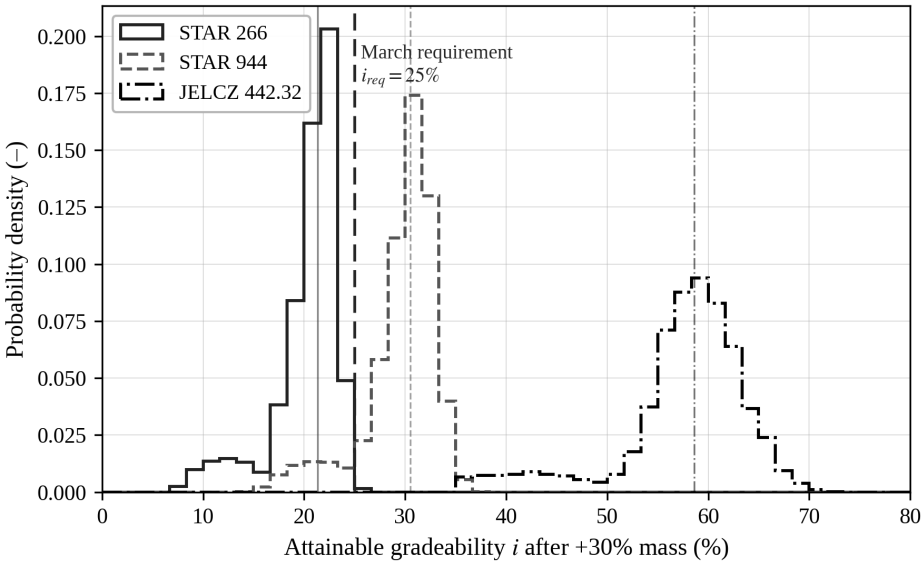


Figure 3. Monte Carlo distributions of attainable gradeability after a 30% mass increase, with the 25% march requirement

Source: own elaboration based on WITPiS (2024)

Discussion

The results invite a reinterpretation of mass growth as a readiness and configuration-management problem rather than a purely engineering one. The mobility-readiness indicator Φ introduced in Equation (11) is, in the vocabulary of reliability and logistics engineering, the operational availability of the mobility function: the probability that a platform in a given configuration can perform the required movement under representative conditions. Expressing the consequence of a configuration change in these terms connects the traction-dynamics analysis directly to the sustainment warfighting function, in which readiness and operational availability are the central objects of management (O'Connor and Kleyner 2012; US Army 2019).

Relation to existing approaches to mobility assessment

The stochastic treatment of mobility is not itself new. Lessem et al. (1992) converted the NATO Reference Mobility Model from a deterministic to a stochastic code, and the same reasoning underpins the current generation of stochastic mobility mapping, in which uncertainty in terrain elevation and soil strength is propagated to a distribution of achievable speed or trafficability (Gonzalez et al. 2017; McCullough et al. 2017). The next-generation programme has since moved towards physics-based, three-dimensional vehicle–terrain simulation (NATO STO 2020; Wasfy and Jayakumar 2021a, 2021b; He et al. 2019) and towards a standardised set of mobility metrics (Wong et al. 2020), with verification and validation governed by the NATO modelling and simulation standards profile (AMSP-06 2021). The present results sit alongside that body of work rather than in competition with it, and it is worth stating precisely where the difference lies.

The first difference concerns which variable is held stochastic and which is treated as the decision. In the reference-mobility-model family the vehicle is a fixed configuration and the terrain is the source of variability; the output is a mobility map, and the question answered is where a given vehicle can go. Here the terrain remains stochastic, but the vehicle configuration — specifically its mass — is the decision variable, and the output is a readiness probability attached to a proposed configuration change. That inversion is what makes the result usable in configuration control, because the question a configuration control authority has to answer is not where a vehicle can go but whether a proposed up-fit leaves the fleet able to do what doctrine requires of it.

The second difference concerns the input requirement, and it is the reason the method is deliberately elementary. Physics-based mobility models require terrain databases, soil parameters and validated tyre–soil models that are rarely available

for a proposed up-fit at the moment the decision has to be taken; the empirical traction models of the vehicle-cone-index and mean-maximum-pressure families require measured cone-index data for the terrain in question (Maclaurin 2007). The indicator developed here requires only the traction-dynamics characteristic, which exists for every accepted platform as a by-product of type testing (Simiński 2015; WITPiS 2024), together with a terrain mixture that can be specified doctrinally. It is therefore a screening instrument, intended to be applied early and often and to be cheap enough to apply to every proposed change; where a change survives screening and its consequences are severe, high-fidelity analysis of the kind described by Wasfy and Jayakumar (2021a, 2021b) remains the appropriate next step. The two are complementary, and the screening layer is the one that is currently missing from configuration-control practice.

The third difference concerns the specific finding. That a probabilistic treatment of terrain yields a different answer from a nominal-case evaluation is established in the mobility-reliability literature (Gonzalez et al. 2017), and slip and traction losses on deformable ground have been shown to be systematically underestimated by deterministic treatments (Gonzalez and Iagnemma 2018). What the present results add is a quantification of how large the discrepancy is at the point where it matters for a specification: a platform certified against a deterministic 25% gradeability requirement, and clearing it by more than a quarter of the requirement, retained the required mobility function in roughly nine cases out of ten. A margin expressed deterministically does not translate into a readiness figure without the distribution over which it is to be evaluated, and the gap between the two is large enough to change procurement and fielding decisions.

Implications for operational decisions

Three practical uses follow. First, the mass budget becomes a quantitative configuration-management control. Instead of recording added mass as an accounting entry, a configuration-control authority can score each proposed up-fit by its effect on Φ, accepting or rejecting changes according to their readiness impact and the margin that remains. Second, the analysis provides acquisition decision support. The decisive variable is the baseline margin, and the cheapest moment to buy margin is at acquisition; a platform procured with a thin mobility margin, such as the STAR 266 in this study, will lose its required mobility function under modest and entirely predictable in-service growth, whereas a platform procured with ample margin, such as the JELCZ, can absorb the same growth and remain fit for task. The cost of inadequate baseline margin is therefore a foreseeable readiness liability that should be priced into the original selection. Third, the stochastic results inform march planning and the resilience of the manoeuvre and distribution network: because Φ is derived

from a distribution over terrain, its left tail defines worst-surface readiness, which is the relevant quantity when planning movements over mixed or degraded ground.

The distinction between deterministic compliance and readiness becomes sharper in deployed and wartime conditions, because the terrain mixture on which readiness depends is not fixed. Damage to routes, the loss of hard-surfaced roads to interdiction, seasonal thaw and the forced use of secondary or cross-country routes all shift weight from the firm-road component of the mixture towards its unpaved and soft-ground components. Because attainable gradeability enters the calculation through the minimum of a traction-limited and an adhesion-limited branch, that shift does not degrade readiness proportionally. For a platform whose distribution already straddles the requirement, a modest change in the mixture moves a large fraction of the distribution across the threshold; for a platform with ample margin the same shift is immaterial. Two fleets that are indistinguishable on a peacetime march table can therefore behave very differently once the road network deteriorates, and the indicator makes that divergence visible in advance rather than after the first failed march.

Three uses follow for deployed conditions. First, the terrain mixture should be re-specified for the theatre and the season rather than taken from a peacetime default, and Φ recomputed for the fleet in its actual configuration; the computation is inexpensive and can be repeated whenever the operating environment changes. Second, march planning should read the left tail of the distribution rather than its median, because it is the worst surface on a route, and not the average one, that determines whether a column completes the march without recovery assistance; the fifth-percentile values in Table 3 are the planning quantity, not the medians. Third, where an up-fit is fielded under operational urgency, as add-on protection characteristically is, the readiness cost should be recorded at the moment of fielding, so that the resulting change in the mobility posture of the fleet is visible to the commander as an explicit trade rather than absorbed silently into the availability return.

Mobility is only one of the channels through which mass growth propagates. The same increment raises axle loads, and with them the military load classification that governs bridge, ferry and raft crossing under NATO STANAG 2021, the required tyre load rating, the class of recovery vehicle needed to extract the platform, and the transportability of the vehicle by rail and by air. It also raises fuel consumption, and therefore the resupply tonnage that the distribution network must move to sustain a given rate of manoeuvre. None of these channels is modelled here, and the mobility figure should not be read as the whole cost of an up-fit. The wider point is that a single control — the mass budget — governs all of them simultaneously, which strengthens rather than weakens the case for managing it as a readiness instrument rather than as an engineering accounting entry.

Contribution and its boundaries

It is important to be clear about what is and is not novel here. The inverse dependence of the dynamic factor on mass, and hence the universal relative law of Equation (6), is elementary; it is adopted as premise P1 and is not claimed as a result. It is stated because it must be stated correctly and because its very exactness is what gives the hypotheses their force: a relation that holds without exception across every platform is nonetheless shown, by H1 and H2, to determine nothing about the operational outcome. The contribution of the paper is fourfold. The mobility cost of mass growth is anchored empirically in the measured characteristics of three fielded platforms, using data that already exist for every accepted vehicle. It is demonstrated that the universal relative law is operationally uninformative without the baseline margin and the binding regime, so that identical relative degradation yields opposite operational verdicts. It is shown that the binding regime can conceal a real loss of reserve, so that apparent insensitivity to mass growth is not evidence of robustness. And the deterministic result is extended into a reproducible stochastic readiness indicator which is shown to diverge materially from the deterministic verdict, so that compliance and readiness are established as distinct properties requiring distinct specification. The reframing of the mass budget as a readiness instrument is the practical synthesis of these points.

Validation status and limitations

The validation status of the results requires an explicit statement, in the sense in which verification and validation are distinguished in the NATO modelling and simulation standards profile (AMSP-06 2021). What has been performed is verification. The deterministic outputs were checked for agreement with the per-gear traction-dynamics characteristics from which they derive, and the stochastic layer is a transparent and reproducible propagation of explicitly documented distributions through Equations (4), (7) and (8). No independent experimental validation of an up-fitted configuration has been carried out, and none is claimed. The figures in Tables 2 and 3 should accordingly be read as the output of a documented model evaluated on measured baseline data, and not as certified performance of up-fitted vehicles. The comparative ordering of the three platforms, which follows from their baseline margins and from the regime in which each operates, is considerably more robust than any individual absolute value, and it is on that ordering, rather than on the individual figures, that the operational conclusions of this paper rest.

The experimental validation that would be required is well defined, and it is specified here so that it can be executed and so that the present results can be falsified. It comprises four elements. First, ballasted trials at mass-growth fractions

of 0, 0.15 and 0.30, with the ballast distributed so as to reproduce the axle-load distribution and the centre-of-gravity height of a representative protection package rather than merely its total mass. Second, measurement of attainable gradeability on a certified gradient track and of the time to 50 km h′ ¹ on an instrumented level straight, repeated at each mass level. Third, independent determination of the rolling-resistance coefficient and of the adhesion coefficient on the surfaces actually used, since the reference values f = 0.012 and μ = 0.75 adopted here are assumptions rather than measurements, and the second of them fixes the position of the regime switch on which H2 turns. Fourth, at least one trial series on soft or degraded ground, because the firm-road reference case exercises the adhesion-limited branch of Equation (8) only for the JELCZ 442.32 at baseline, whereas the Monte Carlo mixture relies on that branch throughout its lower tail. Agreement should be assessed on two quantities in particular: the predicted attainable gradeability, and the mass at which the regime switch occurs.

Several limitations of the model itself qualify the findings. The aerodynamic and maximum-speed estimates rest on standard approximations and were not the focus of this study. The rotating-mass treatment assumes that added mass is non-rotating, which is appropriate for armour and mission equipment but would require adjustment wherever an up-fit also changes wheels, rims or run-flat inserts, as protection packages frequently do. More significantly, added mass is treated as a scalar: an up-fit in practice raises the centre of gravity and redistributes load between axles, and both effects influence the utilisation of the available adhesion, the lateral stability margin and the roll-over threshold, none of which is captured by Equations (4), (7) and (8). For protection packages concentrated high on the hull this is the most consequential omission, and it will tend to make the readiness figures reported here optimistic rather than pessimistic. Finally, the gradeability model is quasi-static: it excludes transient wheel slip, the action of differential locks and traction control, and the thermal and durability limits of the driveline, which on a sustained climb may bind before the traction limit is reached.

Further limitations concern the scope of the scenario rather than the model. The terrain mixture used in the Monte Carlo simulation is representative rather than site-specific, and a single nominal up-fit centred on 30% growth is examined; a theatre-specific mixture or a multi-stage growth trajectory would shift the readiness figures, although the qualitative ordering of the platforms, set by their baseline margins, would be expected to persist. The march requirement of 25% is treated as a fixed doctrinal parameter, and every readiness figure reported here is conditional upon it. Three wheeled platforms of a single national fleet are analysed, so the generality of the conclusions across tracked platforms, across other drive configurations and across other fleets is indicated by the reasoning but is not established by the data. These limitations bound the numerical claims made in this paper. They do not affect the structural results, because H1, H2 and H3 are statements about the form

of the relationship between mass, gradeability and readiness, and would survive a change in any of the parameter values.

Conclusions

This paper has quantified the mobility cost of mass growth for three platforms of the Polish Armed Forces from measured traction-dynamics characteristics, and has reframed that cost as a fleet-readiness and configuration-management problem. The principal conclusions are as follows.

1. Premise P1 is adopted, not claimed as a result. The maximum dynamic factor scales inversely with mass through a single, platform-independent relation: a 30% mass increase reduces it by 23.1%, and a 45% increase by 31.0%, on every platform. This is an elementary consequence of the definition of the dynamic factor and is stated only because the operational argument depends on its exactness.
2. Hypothesis H1 is confirmed: the platform-independence of the relative loss does not propagate. At a 30% mass increase the loss of attainable gradeability spans 18.3% to 26.7% across the three platforms and the increase in the time to 50 km h'^{1} spans 23.4% to 34.8%, against a dynamic-factor loss that is 23.1% for every platform without exception. The same increase drives the STAR 266 below the 25% march requirement while leaving the STAR 944 and the JELCZ 442.32 fit for task. A uniform characterisation of the penalty, such as a fixed fractional loss of gradeability, is therefore incorrect.
3. Hypothesis H2 is confirmed: the mapping from mass growth to gradeability is piecewise. The JELCZ 442.32 is adhesion-limited at baseline and loses only 1.4 percentage points of gradeability over its first 15% of mass growth, although its dynamic factor has by then already fallen by 13.0%. The loss is not absent but concealed by adhesion saturation, and once the platform crosses into the traction-limited regime it loses a further 12.3 percentage points over the next 15% of growth. Apparent insensitivity to mass growth may therefore be an artefact of the binding constraint rather than evidence of a genuine reserve.
4. Hypothesis H3 is confirmed: deterministic compliance is not a sufficient condition for readiness. After a 30% up-fit the STAR 944 clears the march requirement deterministically by 7.2 percentage points, yet its mobility-readiness indicator over the representative terrain mixture is 90.24% [89.8, 90.6]. A mobility requirement expressed as a single deterministic gradeability figure does not by itself assure the readiness it is intended to secure, and the surface distribution against which it is to be met must be specified with it.

5. The mobility-readiness indicator Φ, obtained by Monte Carlo simulation over representative terrain and up-fitting distributions, provides a statistically sound measure of the operational availability of the mobility function. After a nominal 30% up-fit it is 0.24% for the STAR 266, 90.2% for the STAR 944 and greater than 99.98% for the JELCZ 442.32.
6. The mass budget should be managed as a readiness instrument. Mobility-margin requirements ought to be embedded at acquisition together with the terrain distribution against which they are to be met, proposed up-fits ought to be scored by their effect on Φ within the configuration-management process, and worst-surface readiness ought to be read from the left tail of the Φ distribution in march planning, particularly in deployed conditions, where route damage and seasonal degradation shift the terrain mixture away from its peacetime composition.
7. The results are obtained by verification against the source characteristics and not by independent field trials. The comparative ordering of the platforms is robust to the assumptions made; the absolute figures require experimental confirmation, for which a specific validation programme is set out in Section 5.4.

Future work should execute the validation programme of Section 5.4 on up-fitted platforms, extend the model to the centre-of-gravity and axle-load effects that the scalar treatment of mass omits, extend the stochastic model to multi-stage growth trajectories and site-specific terrain mixtures, and integrate the readiness indicator into fleet-level availability and sustainment models, so that the mobility consequences of configuration changes can be assessed alongside the other determinants of materiel readiness.

References

AMSP-06. 2021. *NATO Modelling and Simulation Standards Profile*. Brussels: NATO Standardization Office.

Bekker, M. G. 1969. *Introduction to Terrain-Vehicle Systems*. Ann Arbor, MI: University of Michigan Press.

Blanchard, B. S. 2004. *Logistics Engineering and Management*. 6th ed. Upper Saddle River, NJ: Pearson Prentice Hall.

Genta, G., and L. Morello. 2009. *The Automotive Chassis. Volume 2: System Design*. Dordrecht: Springer.

Gillespie, T. D. 1992. *Fundamentals of Vehicle Dynamics*. Warrendale, PA: Society of Automotive Engineers.

Gonzalez, R., and K. Iagnemma. 2018. "Slippage Estimation and Compensation for Planetary Exploration Rovers. State of the Art and Future Challenges." *Journal of Field Robotics* 35 (4): 564–577. DOI:10.1002/rob.21761.

Gonzalez, R., P. Jayakumar, and K. Iagnemma. 2017. "Generation of Stochastic Mobility Maps for Large-Scale Route Planning of Ground Vehicles: A Case Study." *Journal of Terramechanics* 69: 1–11. DOI:10.1016/j.jterra.2016.10.001.

He, R., C. Sandu, A. K. Khan, A. G. Guthrie, P. S. Els, and H. A. Hamersma. 2019. "Review

of Terramechanics Models and Their Applicability to Real-Time Applications." *Journal of Terramechanics* 81: 3–22. DOI:10.1016/j.jterra.2018.04.003.
Jazar, R. N. 2017. *Vehicle Dynamics: Theory and Application*. 3rd ed. Cham: Springer.
Jones, J. V. 2006. *Integrated Logistics Support Handbook*. 3rd ed. New York: McGraw-Hill.
Lessem, A., Ahlvin, R., Mason, G. and Mlakar, P., 1992. *Stochastic Vehicle Mobility Forecasts Using the NATO Reference Mobility Model. Report 1: Basic Concepts and Procedures*. Technical Report GL-92-11. Vicksburg, MS: U.S. Army Engineer Waterways Experiment Station, Geotechnical Laboratory.
Maclaurin, B. 2007. "Comparing the NRMM (VCI), MMP and VLCI Traction Models." *Journal of Terramechanics* 44 (1): 43–51. DOI:10.1016/j.jterra.2006.01.005.
McCullough, M., P. Jayakumar, J. Dasch, and D. Gorsich. 2017. "The Next Generation NATO Reference Mobility Model Development." *Journal of Terramechanics* 73: 49–60. DOI:10.1016/j.jterra.2017.06.002.
NATO STO. 2020. *Next-Generation NATO Reference Mobility Model (NG-NRMM)*. STO-TR-AVT-248. Neuilly-sur-Seine: NATO Science and Technology Organization.
O'Connor, P. D. T., and A. Kleyner. 2012. *Practical Reliability Engineering*. 5th ed. Chichester: Wiley.
Rajamani, R. 2012. *Vehicle Dynamics and Control*. 2nd ed. New York: Springer.
Simiński, P. 2015. *Wojskowe pojazdy kołowe* [Military Wheeled Vehicles]. Sulejówek: Wojskowy Instytut Techniki Pancernej i Samochodowej.
US Army. 2019. *ADP 4-0: Sustainment*. Washington, DC: Headquarters, Department of the Army.
Wasfy, T. M., and P. Jayakumar. 2021a. "Next-Generation NATO Reference Mobility Model Complex Terramechanics – Part 1: Definition and Literature Review." *Journal of Terramechanics* 96: 45–57. DOI:10.1016/j.jterra.2021.02.002.
Wasfy, T. M., and P. Jayakumar. 2021b. "Next-Generation NATO Reference Mobility Model Complex Terramechanics – Part 2: Requirements and Prototype." *Journal of Terramechanics* 96: 59–79. DOI:10.1016/j.jterra.2021.02.007.
WITPiS. 2024. *Traction-Dynamics Characteristics of the STAR 266, STAR 944 and JELCZ 442.32 Platforms*. Unpublished measurement and computation datasets. Sulejówek: Military Institute of Armoured and Automotive Technology.
Wong, J. Y. 2008. *Theory of Ground Vehicles*. 4th ed. Hoboken, NJ: Wiley.
Wong, J. Y. 2010. *Terramechanics and Off-Road Vehicle Engineering*. 2nd ed. Oxford: Butterworth-Heinemann.
Wong, J. Y., and W. Huang. 2006. "'Wheels vs. Tracks' – A Fundamental Evaluation from the Traction Perspective." *Journal of Terramechanics* 43 (1): 27–42. DOI:10.1016/j.jterra.2004.08.003.
Wong, J. Y., P. Jayakumar, and E. Toma. 2020. "A Review of Mobility Metrics for Next Generation Vehicle Mobility Models." *Journal of Terramechanics* 87: 11–20. DOI:10.1016/j.jterra.2019.10.003.

Author Contributions

The article is the work of a single author, who was responsible for the conceptualisation, methodology, formal analysis, software, data curation, visualisation, preparation of the original draft, and review and editing of the final version.

Funding

This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Data Availability Statement

The traction-dynamics characteristics on which the study is based are held by the Military Institute of Armoured and Automotive Technology (WITPiS), Sulejówek, and are not publicly available, owing to

institutional restrictions on the release of test data for platforms in service with the Polish Armed Forces. All derived data required to reproduce the deterministic results are reported in full in Tables 1 to 3. The Monte Carlo procedure is described in Section 3.6 in sufficient detail for independent reimplementation; a fixed random seed was used, and the simulation code is available from the author on reasonable request.

Conflict of Interest Statement
The author declares no conflict of interest. The author is employed by the Military Institute of Armoured and Automotive Technology, which holds the source characteristics used in the study; the institute had no role in the design of the study, in the analysis or interpretation of the data, or in the decision to submit the work for publication.

Use of AI Tools
Generative artificial-intelligence tools were used solely for language editing and for checking the internal consistency of the reference list. No part of the scientific content — the model, the analysis, the results or their interpretation — was generated by such tools. The author reviewed and verified the entire text and takes full responsibility for the content of the article.

Ethics Approval / Consent
Not applicable. The study did not involve human participants, human data or animals.

Rights and Permissions for Third-Party Material
All tables and figures are the author's own elaboration, prepared for this article on the basis of the source characteristics cited. No third-party copyrighted material is reproduced, and no permissions are therefore required.

Received: 29.06.2026
Accepted: 24.10.2026

© Przemysław Simiński, 2026. This article is published in "*Systemy Logistyczne Wojsk*" / *"Military Logistics Systems*" under the Creative Commons Attribution 4.0 International licence (CC BY 4.0).