Dual-Array Variable Mechanical Advantage Transmissions with Guide-Load Balancing for Mechanical Launchers

Rami N. Mahdi
Interesting Machines LLC, Sacramento, California, USA

Preprint, version 1. 5 September 2026

Abstract

Variable mechanical advantage transmissions can drive a downstream fixed-ratio stage from dual arrays of engagement members on opposite sides of a central guide. Anchoring each array's tension member independently partitions the stage's k falls between two separate lines. This removes $\lfloor k/2 \rfloor$ elements from the busiest rope path, reducing friction and bending losses while enabling larger stage ratios that slow the stage blocks and reduce required member counts.

However, mirror-symmetric dual arrays cannot balance guide reactions at odd stage ratios k: because the partitioned falls differ by one ($|n_1 - n_2| = 1$), symmetry imposes a permanent, uncorrectable lateral load of $|n_1 - n_2|/k$ (e.g., 20% at $k = 5$). Dynamic guide balance is restored by configuring deliberately unequal, semi-symmetric arrays according to the generalized condition $G_1 : G_2 = n_2 T_2 : n_1 T_1$, wherein the array actuating more falls or suffering lower path losses is assigned a smaller displacement ratio. The independent geometry doubles design degrees of freedom, improving ratio profile tracking threefold. Furthermore, initiating engagement with the array serving fewer falls ($n_1 = \lfloor k/2 \rfloor$) suppresses inter-engagement force drift, cutting residual lateral loads by 39% to 57% at $k = 5$ and dominating the alternative lead assignment across almost the whole accuracy-balance trade at $k = 7$ and $k = 9$.

In simulated 1,000:1 mass-ratio launches, semi-symmetric designs at $k = 7$ (7 members/array) and $k = 9$ (5 members/array) reduce peak-to-mean compliant payload forces from 1.40 to 1.12 and 2.19 to 1.48, respectively, compared to matched single or mirror-symmetric baselines, while cutting line contacts. Against the mirror-symmetric machine the peak lateral guide load falls from 14.3% to 8.1% of the carriage reaction at $k = 7$ and from 11.1% to 7.0% at $k = 9$, reductions of 43% and 37%, or 14.6 kN and 9.3 kN on a peak reaction of about 230 kN.

1. Introduction

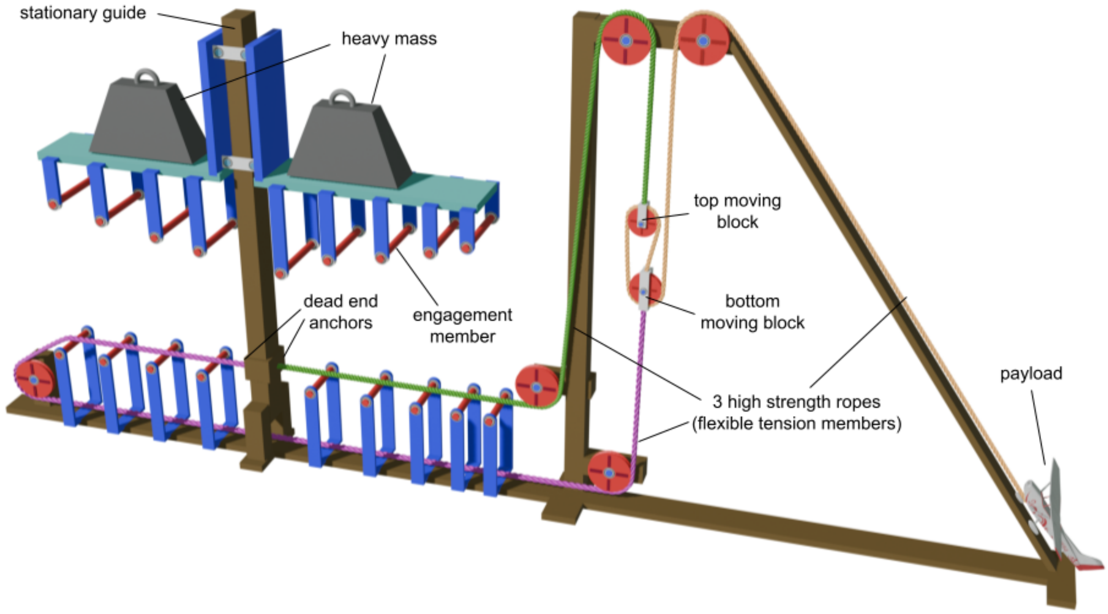

Figure 1: Schematic of the dual-array RopeComb launcher. As the carriage descends the central guide, engagement members deflect independent tension members into fixed spans on either side. Each rope actuates a separate movable block at the fixed-ratio stage, combining at the output member to launch the payload. Only the final output line and payload move at launch speed. Geometry is illustrative and not to scale.

Assisted launch systems transfer initial acceleration work from flight vehicles to stationary ground infrastructure, enabling reduced airframe structural mass and higher payload mass fractions [1]. Because coupling a slow, high-force energy source, such as a descending mass, pneumatic piston or hydraulic ram, directly to a high-speed vehicle produces extreme peak accelerations and low transfer efficiency, impedance-matching transmissions are required [2, 3]. The descending-mass form of the problem is an old one. The counterweight trebuchet drives a light projectile from a slow, heavy source through a beam and sling whose effective ratio rises through the throw [4], and is in that sense an early variable mechanical advantage launcher; what it lacks is any means of choosing that ratio profile, which is the freedom the present architecture supplies.

The recently introduced RopeComb transmission [5] accomplishes this impedance match passively through a configurable rope-and-pulley geometry. A carriage carrying an array of engagement members translates past a row of fixed supports, deflecting a flexible tension member into successive spans. Each fold draws in rope at a rate that increases with deflection depth, producing a displacement ratio that rises monotonically through the stroke. Two independent geometric degrees of freedom per member, the span half-width R_i and the engagement offset s_i, allow the mechanical advantage profile to be synthesized via least-squares fitting against an arbitrary target, approaching the force uniformity of electromagnetic catapults without electrical storage or power conditioning [5, 3].

In high-impulse linear launchers the forces opposing carriage motion are severe, roughly 23 times the weight of the driving mass in all three designs of Mahdi [5], which is 23.4 kN at a mass ratio of 100:1, 228 kN at 1,000:1 and 2.09 MN at 10,000:1. Managing how these reaction forces interface with the carriage guide and the downstream stage leads to three distinct architectures.

A single array on one side of the guide is the baseline arrangement of Mahdi [5]. It is mechanically straightforward, but incurs two severe penalties. The entire carriage reaction force is reacted at a lateral offset from the guide axis, imposing a large overturning moment on the guide bearings while they slide at velocities up to 10 m/s; bearing friction under this side load constitutes unmodelled dissipation that directly reduces exit velocity. Furthermore, a single continuous tension member must traverse every span of the array and every fall of the fixed-ratio stage, resulting in long rope paths and cumulative bending losses.

Two mirror-image arrays positioned symmetrically on opposite sides of the guide, proposed in Mahdi [5, Appendix E.2], eliminate the overturning moment by construction and halve the load carried by each array. Its principal practical merit is structural economy: the two arrays are identical parts, so an apparatus of N members per side requires only N distinct span widths and N distinct engagement heights, halving tooling, spares and machining setups.

Two arrays that are deliberately unequal, the semi-symmetric architecture investigated here. Nothing in the kinematics of the launcher requires the two opposing arrays to be identical. Because each array operates its own tension member, dead-ended on the guide and driving an independent movable block at the fixed-ratio stage, the two sides are kinematically independent.

Crucially, this paper corrects an incomplete assumption in Mahdi [5, Appendix E.2], which asserted that equal tension between opposing mirror-image arrays "is the natural state rather than something to be engineered". While mirror symmetry guarantees equal displacement rates, it does not guarantee equal reaction forces at the carriage. Opposing forces depend on the number of falls each array actuates at the stage and on the tension delivered along each independent path. Whenever the fixed-stage ratio is odd the falls cannot be divided equally, leaving a mirror-symmetric machine with a permanent, uncorrectable lateral guide load reaching 11% to 33% of total reaction force. Even at even stage ratios, unavoidable differences in path length, sheave friction and compliance unbalance the guide.

Allowing the two arrays to differ resolves this limitation. By setting array displacement ratios in inverse proportion to their fall counts and path tensions, $G_1 : G_2 = n_2T_2 : n_1T_1$, lateral guide forces cancel across the stroke. Furthermore, partitioning the stage falls between two separately anchored lines shortens the busiest rope path by $\lfloor k/2 \rfloor$ elements, enabling higher stage ratios that slow the stage sheaves and lower parasitic inertia [5].

Section 2 establishes notation and summarizes the single-array kinematics. Section 3 presents the dual-array architecture and derives the fall-count weighting. Section 4 itemizes the four mechanical benefits of the dual architecture, the first two of which follow from having two separately anchored ropes at all and the second two from letting the arrays differ. Sections 5 and 6 formulate the generalized balance condition and analyze odd-ratio guide imbalance. Section 7 gives the engagement-ordering rule that suppresses lateral excursions. Section 8 details the two-stage fitting algorithm. Section 9 reports dynamic simulation results at a 1,000:1 mass ratio, and compares the architecture against single and mirror-symmetric arrays at matched configuration. Section 10 sets out what is general and what is worked, and Section 11 concludes.

The architecture is the subject of two provisional patent applications; see *Competing interests*.

2. Notation and the array ratio

Consider one array of N engagement members carried by the carriage, interleaved with fixed supports that define spans between them. Writing d for carriage displacement, R_i for the half-width of the ith span and s_i for the displacement at which the ith member first contacts the tension member, the depth of the ith member below the baseline is $D_i = \max(0,\, d - s_i)$ and the array's displacement ratio is

$$G_{\text{array}}(d) \;=\; \sum_i \frac{2\,D_i}{\sqrt{R_i^2 + D_i^2}} \tag{1}$$

as derived in [5, Section 3]. Each engaged member contributes at most 2, approached as $D_i/R_i \to \infty$, so an array of N members has a ceiling of $2N$.

Subscripts 1 and 2 denote the two arrays. A *lateral load* stated as a percentage means the greatest value the imbalance reaches at any point of the stroke, divided by the greatest value reached at any point of the stroke by the sum of the two arrays' weighted contributions: a peak compared against a peak. It is not an instantaneous ratio, which is of no use near the beginning of the stroke where both contributions approach zero.

3. The dual architecture and the fall-count weighting

Each array has its own flexible tension member, dead-ended at a stationary anchor carried by the guide, running outboard through its own array and returning to the fixed-ratio stage, where it drives its own movable element. The output member is reeved with n_1 falls about the first movable element and n_2 about the second, and

$$n_1 + n_2 = k \tag{2}$$

Figure 2 shows the arrangement schematically and Figure 1 a whole machine embodying it.

Kinematic independence. Because each tension member is separately anchored and drives its own movable element, nothing requires the two arrays to draw in rope at the same rate. This is a structural precondition for everything that follows. Were the two arrays to act on a common movable element, or were their tension members portions of one continuous rope, the arrangement would be over-constrained and any difference in geometry would put one member slack. The separate anchors are what make unequal arrays possible at all.

The weighting. Taking up a length Y_1 of the first tension member advances the output member by $n_1 Y_1$, and likewise for the second. The net displacement ratio is therefore

$$G_{\text{net}}(d) \;=\; n_1\, G_1(d) \;+\; n_2\, G_2(d) \tag{3}$$

Expression (3) is exact and is a plain sum, so the net ratio may be evaluated by counting each array's members according to its fall count. It is worth seeing why it is exact, because the same counting gives the stage ratio itself. Each segment of the output member changes in length by the travel of whichever movable elements it touches, so summing over segments returns (3) with n_i the number of segments touching element i. With p sheaves distributed between the two elements the output member has $p + 1$ segments, and every one but the last touches both elements, so $n_1 + n_2 = 2p + 1$: the stage ratio of (9). No stationary sheave is needed anywhere in the stage, which is what keeps its transmission loss low.

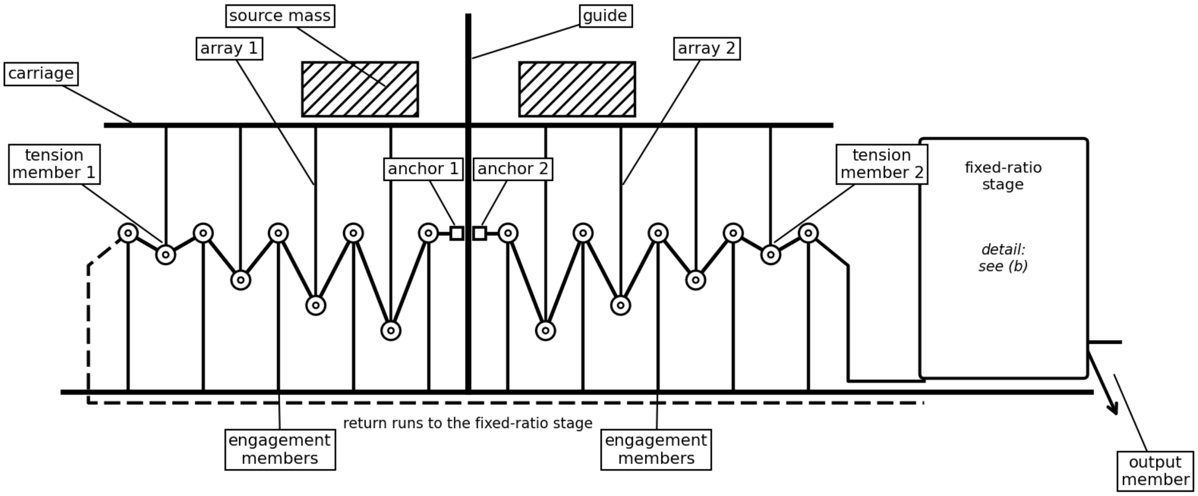

(a) the dual array, two arrays either side of the stationary guide

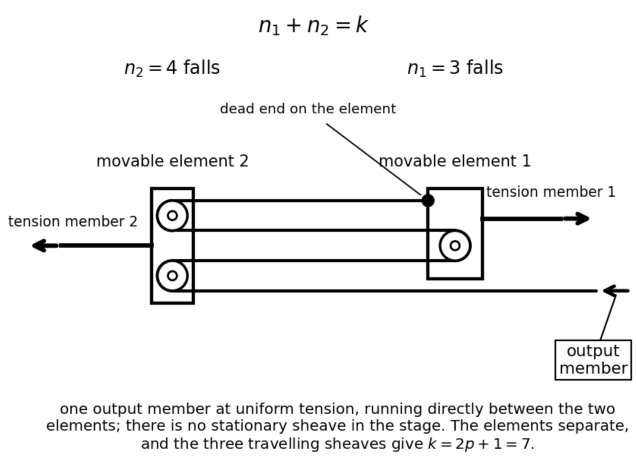

(b) the fixed-ratio stage: the k falls are partitioned between the two arrays

Figure 2: The architecture. (a) The dual array with its stationary guide, carriage, source-mass halves and both tension members dead-ended on the guide. (b) The fixed-ratio stage at $k = 7$, showing the two movable elements and the n_1 and n_2 falls of the output member about them. The output member is one continuous line at uniform tension, running directly between the two elements: it dead-ends on element 1, passes about the upper sheave of element 2, returns about element 1's single sheave, passes about the lower sheave of element 2 and leaves for the payload. Counting the segments that touch each element gives $n_1 = 3$ and $n_2 = 4$. The two tension members draw the elements apart, as in Figure 5. There is no stationary sheave in the stage, so every sheave counts toward p and $k = 2p + 1 = 7$; a stage reeved instead against fixed sheaves would reach the same ratio at the cost of their friction.

Where k is odd the fall counts cannot be equal. Throughout, subscript 1 denotes the array that engages first, which Section 7 shows must be the one serving the fewer falls, so $n_1 = \lfloor k/2 \rfloor$ and $n_2 = \lceil k/2 \rceil$. That is 2 and 3 at $k = 5$, 3 and 4 at $k = 7$, 4 and 5 at $k = 9$.

With mirror-image arrays $G_1 = G_2 = G$ and (3) returns kG for *any* split summing to k. The weighting is therefore invisible in a symmetric machine, and becomes visible only when the arrays differ. This is worth noting because it explains why the question does not arise until one asks it: within the symmetric arrangement, the quantity that distinguishes the two sides cannot be observed.

4. What the second array buys

Four benefits follow, and they are largely independent of one another. The first two follow from having two separately anchored ropes at all, whether or not the arrays differ. The second two follow specifically from letting them differ.

4.1 Shorter rope paths

Every element a tension member touches costs friction and bending loss. Each rope's path costs two contacts per member of its own array, plus one per fall it serves at the stage, plus the entry: about $2N + f + 1$ elements for an array of N members serving f falls.

A single array serves all k falls. Each of two arrays serves about half. At N members per array the busiest path therefore carries $2N + \lceil k/2 \rceil + 1$ against $2N + k + 1$, a saving of $\lfloor k/2 \rfloor$ elements that does not depend on N and grows with the stage ratio (Table 1).

Table 1: Elements removed from the busiest rope path by dividing the k falls between two separately anchored tension members. The saving is $\lfloor k/2 \rfloor$ and is independent of member count.

k	5	7	9	11	13
elements saved	2	3	4	5	6
efficiency gain at 2% per element	4.1%	6.2%	8.4%	10.6%	12.9%

4.2 A higher usable stage ratio, and a simpler first stage

A larger k is desirable on three counts. It lowers the number of members each array needs, by (7) below. It slows the movable elements of the stage, which translate at $1/k$ of payload velocity and so store energy as $1/k^2$. And fewer members means a shorter, less complex first stage.

On a single array these gains are opposed, because the one rope path that matters lengthens with k contact for contact. Divided between two arrays it does not, and the opposition is roughly halved. The dual arrangement therefore moves the useful range of k upward, and the benefit compounds: a higher k needs fewer members, which shortens the rope path again.

4.3 Twice the geometric freedom

Two arrays of N members constrained to be mirror images have N independent members between them. Two arrays free to differ have $2N$. Where that freedom is not needed for load cancellation it is available for profile accuracy. Table 2 shows the trade at $k = 8$ with seven members per array on the 1,000:1 target of Section 9.

Table 2: The accuracy–balance trade at $k = 8$, seven members per array, on the 1,000:1 target. Imposing balance drives the two arrays back onto the mirror solution and returns the profile accuracy with it.

balance weight λ	rms against target	lateral load	mean difference in array spans
0	**0.197**	6.9%	50 mm
0.3	0.603	0.2%	0.9 mm
1.0	0.603	0.0%	0.0 mm

Unequal arrays track the target about three times more closely than mirror-image arrays of the same member count and width, if a few per cent of lateral load is accepted. This is a genuine trade and both ends of it are useful: λ is a dial the designer sets, not a value to be maximised.

4.4 Balance that symmetry cannot reach

Mirror symmetry equalises the two arrays' *geometry*. It does not equalise the forces they apply unless two further conditions hold: that the two arrays serve equal numbers of falls, and that their rope paths deliver equal tension. Sections 5 and 6 take these in turn. Where either fails, only unequal arrays can cancel the lateral load, and the required inequality is small and calculable.

5. The balance condition

The force each array applies to the carriage is the product of its fall count, the tension delivered to its movable element, and its own displacement ratio. Writing T_1 and T_2 for those tensions, the net lateral load on the guide is

$$\text{imbalance}(d) \;=\; \big|\, n_1 T_1 G_1(d) - n_2 T_2 G_2(d) \,\big| \tag{4}$$

and it vanishes at every point of the stroke where

$$n_1 T_1 G_1(d) \;=\; n_2 T_2 G_2(d) \qquad \text{hence} \qquad G_1 : G_2 \;=\; n_2 T_2 : n_1 T_1 \tag{5}$$

The array serving more falls, or suffering less loss along its path, must be given the smaller displacement ratio. The direction is counter-intuitive: the array doing more work at the stage is the one to be made weaker.

Two special cases matter.

Equal tensions. If the output member is treated as carrying one uniform tension, $T_1 = T_2$, and (5) reduces to

$$G_1 : G_2 \;=\; n_2 : n_1 \tag{6}$$

Combining with (3), each array carries half the weighted target,

$$G_1(d) \;=\; \frac{G^*(d)}{2\,n_1}\,, \qquad G_2(d) \;=\; \frac{G^*(d)}{2\,n_2} \tag{7}$$

where G^* is the target net ratio. This is the form used for the designs of Section 9.

Equal tensions and equal fall counts. If in addition $n_1 = n_2$, (6) gives $G_1 = G_2$ and mirror symmetry is correct. This is the case Appendix E.2 of [5] had in view, and within it the statement made there is right.

Member counts. Each engaged member contributes at most 2 to its own array's ratio, so (7) fixes a minimum number of members per array. Table 3 gives them for the 1,000:1 target of Section 9, whose terminal net ratio is 79.3.

Table 3: Minimum members per array set by the balance condition (7), for the 1,000:1 target of Section 9 whose terminal net ratio is 79.3. Array 1 leads and serves the fewer falls, so it carries the larger displacement ratio and the higher floor. Raising the stage ratio lowers the floor on both arrays.

k	falls n_1/n_2	G_1 required	G_2 required	min. members, array 1	min. members, array 2
5	2 / 3	19.82	13.22	10	7
7	3 / 4	13.22	9.91	7	5
9	4 / 5	9.91	7.93	5	4
11	5 / 6	7.93	6.61	4	4

Raising the stage ratio lowers the floor on both arrays, which is the mechanism behind Section 4.2. Figure 3(b) plots it.

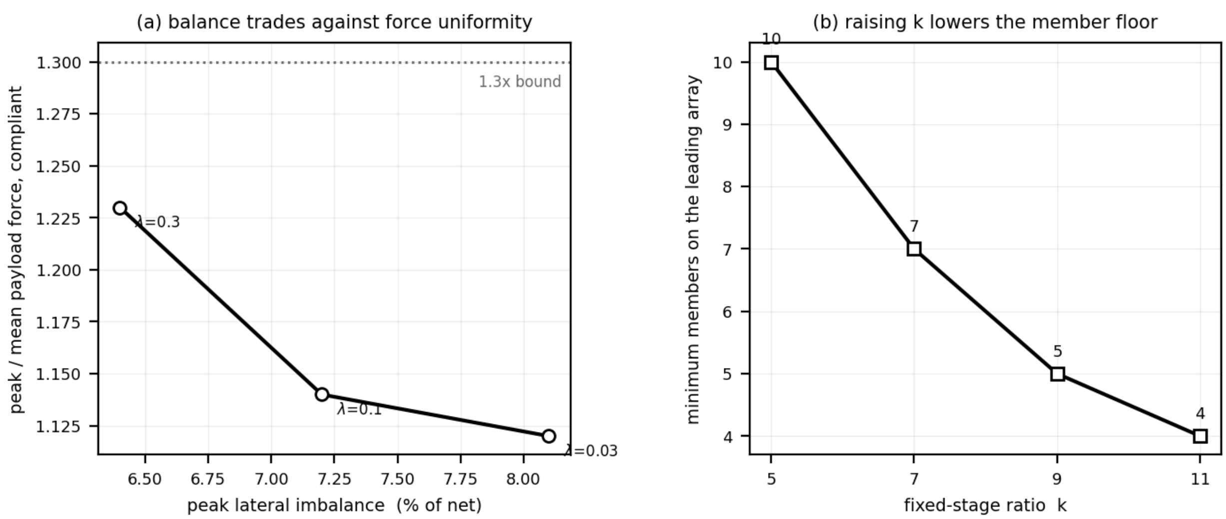

Figure 3: (a) The trade between load cancellation and profile accuracy traced by λ, for the $k = 7$ design of Section 9 over the sweep of Table 8. (b) Minimum members on the leading array against stage ratio, the array-1 column of Table 3.

6. When the two conditions fail

6.1 Unequal fall counts: odd stage ratios

Where k is odd the fall counts differ by one and (6) cannot be satisfied by equal arrays. A mirror-symmetric machine is then left with a residual lateral load of

$$\frac{|\,n_1 - n_2\,|}{k} \tag{8}$$

of the total: **33% at $k = 3$, 20% at $k = 5$, 14% at $k = 7$, 11% at $k = 9$**, zero only where k is even. Figure 4(a) plots it. This is a property of the reeving, not a tolerance: it cannot be reduced by adjustment, by matching of components, or by any care in construction. The 1,000:1 design of [5] uses $k = 5$ and built symmetrically would sit at exactly the 20% that the same paper describes as "far larger than anything a built machine should exhibit".

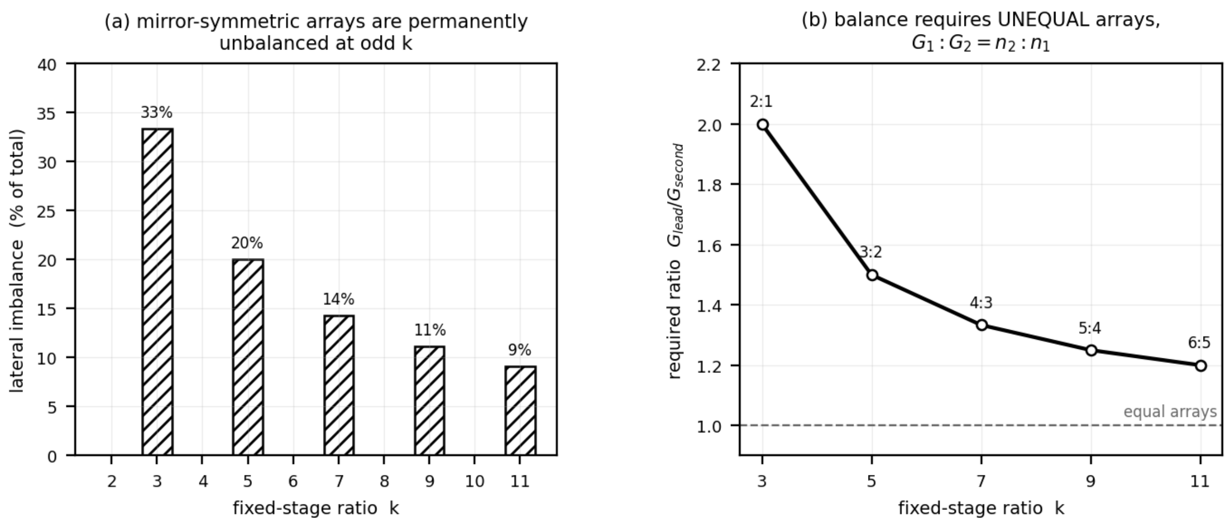

Figure 4: Lateral load and the balance condition. (a) The load a mirror-symmetric machine carries as a function of stage ratio, zero only at even k. (b) The ratio the two arrays must hold to cancel it.

Odd stage ratios are not a corner case; they are what the stage naturally produces. The fixed-ratio stage carries two travelling blocks that move in *opposite* directions, and with p sheaves distributed between them the reeving yields

$$k = 2p + 1 \tag{9}$$

Figure 5 shows the construction for the three ratios used in [5]. Blocks C and D are the two travelling assemblies: D is drawn upward by one tension member and C downward by the other, so the pair separates as the arrays take in rope. Counting the sheaves they carry gives two in panel (a), three in (b) and four in (c), and the stage ratios are 5, 7 and 9 respectively. The rule is read directly off the drawing.

Expression (9) describes this stage. A stage reeved otherwise, with a single travelling block or with the free end anchored to the frame, will not in general obey it; the even case noted below is one such arrangement.

The counter-moving arrangement is chosen for a kinematic reason. Two blocks separating at speed u each open their gap at $2u$, so the output member is drawn in at twice the rate a single block moving at u would achieve. Put the other way, a single travelling block must move at *twice* the speed to consume rope at the same rate. Since kinetic energy goes as the square of velocity, each of the two blocks then stores a quarter of what the single fast block would store, and the stage holds

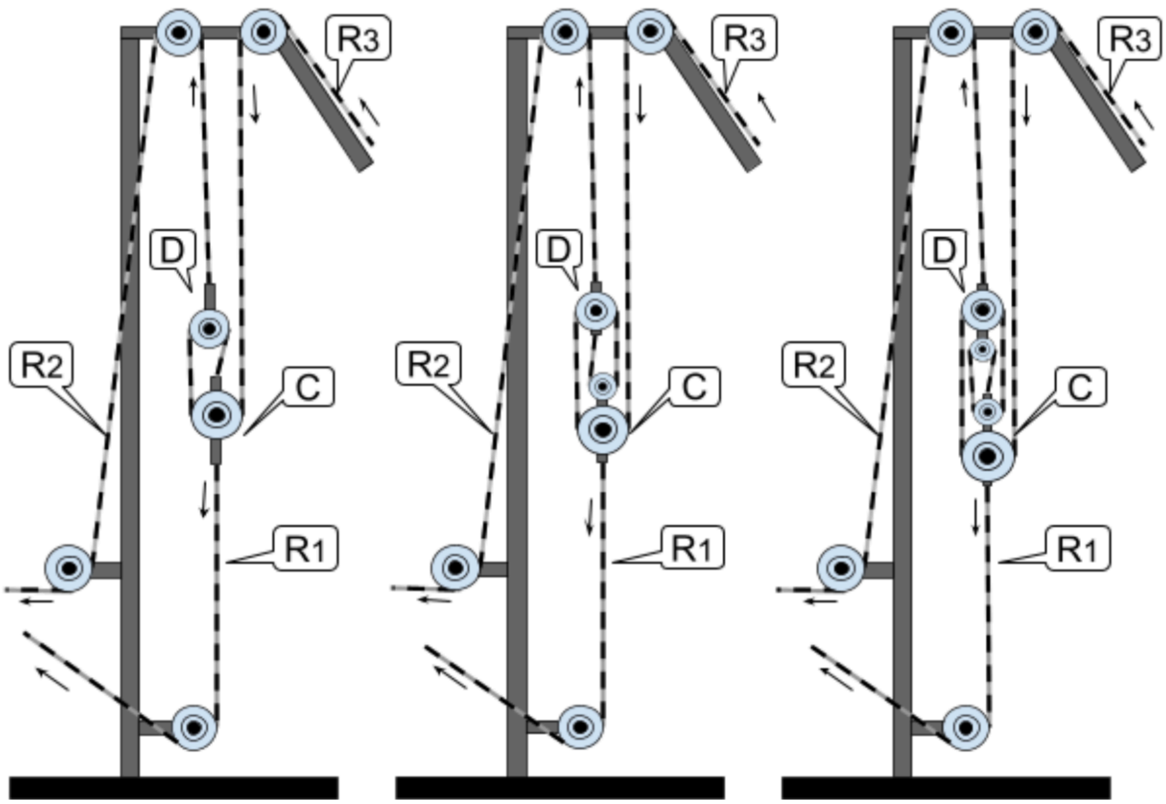

Figure 5: Fixed-ratio stages at $k = 5$, 7 and 9, after [5]. Blocks C and D travel in opposite directions, D drawn up by one tension member and C down by the other, so the pair separates as the arrays take in rope. R_1 and R_2 are the two tension members, R_3 the output member running to the payload. The travelling sheaves number two, three and four from left to right, giving $k = 2p + 1 = 5$, 7 and 9. Arrows show the direction of travel.

about half as much energy in total even though it now carries two blocks rather than one. The parasitic inertia of the stage is the quantity being minimised, so the trade is strongly favourable.

An even k is constructible but wasteful of sheaves. The one simple even case is $k = 2$: a single block carrying one sheave, with the end of the output member anchored to a stationary point. Beyond that, forcing even parity costs sheaves without buying ratio.

This is what the parent work uses: every stage ratio of three or more in [5] is odd, 3 and 7 in the apparatus figures, 5, 7 and 9 in the alternative stages, 5 in the 1,000:1 simulation and 9 in the 10,000:1. The only even value is $k = 2$ in the 100:1 case, exactly the single-sheave arrangement just described.

Adding sheaves is also the right direction. The stage energy goes as p/k^2, and since k grows linearly in p while energy falls as $1/k^2$, $p/(2p+1)^2 \approx 1/(4p)$: going from $k = 3$ to $k = 5$ doubles the sheave count and still cuts stage parasitic energy by 28%, $k = 7$ by 45%, $k = 9$ by 56%. At equal sheave count the odd reeving beats the even one by 21 to 36%, and an even k one sheave larger is no better while costing the extra sheave.

6.2 Unequal tensions: dissimilar rope paths

The assumption $T_1 = T_2$ is an idealisation, and in a built machine it is close to always false. The output member loses tension at every element it passes, so what arrives at one movable element is not what arrives at the other. The two tension members traverse different numbers of elements between their anchors and the stage, and may further differ in free length, in redirecting sheaves along their routes, in elastic extension under load, and in the mass of rope accelerated.

It follows that **mirror symmetry does not fully balance a machine even at even k.** Equal fall counts cancel the fall-count term of (5) and leave the tension term. Symmetry balances only where the two paths are additionally alike in element count, length and compliance; where they are not, a residual remains that no adjustment of the arrays' *common* geometry can remove, precisely because mirror-image arrays are constrained to share that geometry.

The correction is within reach of the same freedoms. Two parameters per member per array are ample to satisfy (5) for any realistic path asymmetry, and the objective of Section 8 extends by one term in the estimated or measured tension difference, the tensions being obtained from element counts and a per-element loss, from computed rope extension, or from measurement on a built machine.

We have not quantified this residual. The designs of Section 9 were configured on (6) with no such term, and the interleaving of Section 7 bounds the between-engagement excursion regardless of tension asymmetry, which sufficed for these cases. Quantifying it needs a per-element friction model and a fixed routing for both members, neither of which this work provides. What is recorded here is that the correction is analytically available and needs no new degrees of freedom.

7. Engagement ordering

The array ratio is continuous at the moment a member engages; it is the slope dG/dd that steps. The lateral load therefore drifts toward whichever array engaged most recently, at a rate proportional to that array's fall count. Two things follow, and they are not the same thing.

The rule: the array that engages first must be the one driving the movable element with the fewer falls, that is $\lfloor k/2 \rfloor$. Leading with the low-fall array makes the first excursion grow at the lesser rate and lets the other array close it. Leading with the high-fall array does the opposite.

The two assignments are compared in Table 4. Because each sits on its own accuracy-against-balance curve, comparing them at a single value of the weight λ is misleading; the table traces both curves instead, taking the best of fourteen restarts at each point.

Table 4: The two lead assignments as trade curves, at a 1,000:1 target. Each row is one value of the balance weight λ, and each entry the best of fourteen restarts. Bold marks the better value of each pair. Comparison is between curves, not between rows: at $k = 7$ every high-fall-leading point is beaten on both axes by the low-fall point at the same λ, and at $k = 9$, where the same-row comparison favours the high-fall column on balance, five of its six points are still beaten on both axes by some point of the low-fall curve.

	low-fall array leads		high-fall array leads	
λ	rms	lateral load	rms	lateral load
$k = 7$, seven members per array				
0.1	**0.400**	**7.2%**	0.776	9.7%
0.3	**0.436**	**6.0%**	0.886	7.5%
1	**0.558**	**4.7%**	0.998	5.9%
3	**0.807**	**3.5%**	1.106	5.0%
10	**0.995**	**1.4%**	1.523	5.4%
30	**1.110**	**0.9%**	2.193	4.4%
$k = 9$, five members per array				
0.1	**0.662**	6.6%	0.933	**6.5%**
0.3	**0.704**	6.1%	1.057	**5.0%**
1	**0.916**	4.4%	1.101	**3.7%**
3	**1.057**	3.3%	1.167	**2.3%**
10	**1.159**	1.8%	1.212	**1.4%**
30	**1.235**	**1.1%**	1.256	1.1%

At $k = 9$ the high-fall-leading column shows a lower lateral load at five of the six values of λ, which is why the comparison has to be made as a curve. Those points are bought with a worse residual against the target, and a low-fall-leading fit at a larger λ reaches the same balance at better accuracy: 3.3% at rms 1.057, against 3.7% at rms 1.101. Only at $\lambda = 10$ does the high-fall curve hold a point the low-fall curve does not dominate. The effect of the assignment is strongest at small k, where the two fall counts differ most in proportion. At $k = 5$ they are 2 and 3, and leading with the wrong array roughly doubles the lateral load, from 8.9% to 14.5% at $\lambda = 0.3$ and from 6.7% to 15.4% at $\lambda = 1$.

Alternation is a consequence, not a cause. Once the lead is assigned correctly, the fitted offsets alternate between the arrays whether or not alternation is imposed: an unconstrained fit at $k = 5$ produced the order LSLSLLSLSLLSLLLSS, with a longest run of three members from either array, and imposing strict alternation on the same problem did not lower the lateral load but raised it, from 1.8% to 3.1%. What imposing it buys is in the fitting rather than in the machine. Parametrising the merged offset sequence by strictly positive increments and dealing them alternately makes every candidate feasible by construction, which is what allows the seeded method of Section 8 to converge in a single solve where unconstrained restarts needed dozens. Coincident and near-coincident engagements remain permissible and do occur in configured designs.

8. Design method

Choose $2(N_1 + N_2)$ parameters so that $n_1 G_1 + n_2 G_2$ approximates the target while $n_1 G_1$ and $n_2 G_2$ stay close to one another:

$$\min \left\| n_1 G_1 + n_2 G_2 - G^* \right\|^2 \; + \; \lambda \left\| n_1 G_1 - n_2 G_2 \right\|^2 \tag{10}$$

subject to span-width bounds and a total-width budget, with the interleaving imposed by construction. The weight λ traverses a continuous family from best-fit to exact cancellation. Table 2 shows both ends of that family at an illustrative $k = 8$, where exact cancellation is reachable; Figure 3(a) traces the shorter interval the $k = 7$ design of Section 9 is drawn from. The two are different configurations, and the trade is a property of the family rather than of either one.

The objective is not convex. Two devices make it tractable.

Two-stage seeding. Solve first for a *single* array coupled to the same stage by the method of [5, Section 4], initialise both arrays with that solution, and displace the second array's offsets to positions intermediate those of the first. Because $n_1 + n_2 = k$ exactly, a pair so constructed reproduces the single array's net ratio exactly, so the seed is already feasible. At the $k = 7$ design of Section 9 one seeded fit reaches rms 0.380, better than the best of forty random restarts of the same objective, which plateau at 0.383; it costs about what one restart costs. `seed_benefit.py` reproduces this.

Staged solution. Within a fit, solve for span widths with offsets held fixed, then for both jointly. Fixing the offsets makes the first phase far better conditioned.

The whole procedure is set out in Algorithm 1. Nothing in it is tuned: the designer supplies the two masses, the entry and terminal velocities, the stage ratio, the members per array, a width budget and the balance weight λ that fixes where on the accuracy-balance trade the result sits, as swept for both configured designs in Table 8.

Algorithm 1. Configuring an unequal dual array

```
NOTATION  arrays are indexed 1 (leading) and 2 (trailing); [] is a subscript,
          ^ is exponentiation only, * is multiplication, Gstar(d) the target.

INPUT   M, m       source and payload mass  OUTPUT  n1, n2       falls per element
        v0         entry velocity                   R1[], s1[]   leading array
        v_end, T   deceleration to achieve          R2[], s2[]   trailing array
        k          fixed-stage ratio
        N          members per array
        L          width budget per array
        lam        balance weight

1  SOLVE THE TARGET                                            (as [4], Sec. 5.2)
   bisect on F until the ideal run, terminated at v_h = v_end, takes time T
   -> design force F, braking distance d_max, target profile Gstar(d)

2  ASSIGN THE FALLS                                               (Sec. 3, Sec. 7)
   n1 = floor(k/2)        the LEADING array, engages first
   n2 = ceil(k/2)         the trailing array
   the lead must take the SMALLER fall count; reversing it is dominated
   member floor, each array:   N >= Gstar(d_max) / (4 * n)

3  STAGE-1 SEED: SOLVE ONE ARRAY                                 (as [4], Sec. 4)
   Rs[], ss[] = argmin  SUM_d [ k * Garray(d; R, s) - Gstar(d) ]^2
                bounded least squares, total width <= L

4  BUILD THE TWO-ARRAY SEED                                             (Sec. 8)
   array 1:  R1[i] = Rs[i]        s1[i] = ss[i]
   array 2:  R2[i] = Rs[i]        s2[i] = ( ss[i] + ss[i+1] ) / 2     i < N
                                  s2[N] = ss[N] + ( d_max - ss[N] ) / 2
   the seed already reproduces the net ratio exactly, since n1 + n2 = k

5  PARAMETRISE SO THE ORDER CANNOT BE VIOLATED                          (Sec. 7)
   merged offsets from strictly positive increments:
     g[]   = softplus(p_gap) + eps                        every g[i] > 0
     s_all = d_max * sigmoid(p_span) * cumsum(g) / sum(g)
   deal s_all alternately:  slot 1 -> array 1, slot 2 -> array 2, ...
   half-widths:  R = R_floor + softplus(p_R)

6  PHASE 1: WIDTHS ONLY, OFFSETS FROZEN                                 (Sec. 8)
   minimise J over p_R alone, holding p_gap and p_span fixed
   stops two offsets migrating onto each other before the widths adapt

7  PHASE 2: JOINT SOLVE from the phase-1 result
   J =         SUM_d [ n1*G1(d) + n2*G2(d) - Gstar(d) ]^2
     + lam   * SUM_d [ n1*G1(d) - n2*G2(d) ]^2
     + w_L   * [ max(0, 2*sum(R1) - L)^2 + max(0, 2*sum(R2) - L)^2 ]

8  EVALUATE                                                             (Sec. 9)
   simulate rigid and compliant; prune members with s[i] >= carriage travel
   payload force over the first 99.5% of the trace (release precedes the
   traction-loss pulse):
     p2m rigid     = max(F_payload) / mean(F_payload)
     p2m compliant = max(F_payload) / F
   lateral load from the geometry, over the whole design stroke 0..d_max:
     lateral load  = max |n1*G1 - n2*G2| / max ( n1*G1 + n2*G2 )
```

9. Configured designs

Both are for a source mass of 1,000 kg and a payload of 1 kg, entering at 10 m/s and leaving at 4 m/s, target uniform payload force, source travel 854 mm. Simulation follows [5, Appendices A and B] unchanged: a rigid-body integrator, and a compliant integrator in which the output member is a pre-tensioned spring. Its stiffness follows [5, Appendix B]: $K_{\mathrm{rope}} = EA/L$ for 2 mm Dyneema R3 [6], with $E \approx 100$ GPa and $A \approx 2.8 \times 10^{-6}\,\mathrm{m}^2$. The output member spans the 15.8 m of payload travel, about 1 m around the sheaves of the movable elements and about 1 m spare, so $L \approx 17.8$ m and $K_{\mathrm{rope}} \approx 1.6 \times 10^4$ N/m. We take 16,471 N/m, the value used in [5], whose worked case has essentially the same payload travel. Appendix A shows the result is insensitive across the range that material and length choices actually span. Both designs are fitted at $\lambda = 0.03$.

Release convention. The stroke is terminated by the traction-loss instability analysed in [5, Section 6], and the payload is released immediately before the terminal transient that instability produces. Payload forces are accordingly measured over the first 99.5% of the simulated time. The integrators step at a fixed interval, so this is a window in time and not in carriage displacement; the two are not interchangeable near the end of the stroke, where the source mass has slowed and each unit of time buys less travel than it did at entry. The convention matters for the rigid figures and not for the compliant ones: on the $k = 7$ design the rigid peak-to-mean is 1.93 measured this way and 6.14 if the final sample is included, that sample being the traction-loss spike itself. The figure is insensitive to where the cut is placed, giving 1.93 at the first 99.9%, 99.5%, 99%, 98% and 95% of the simulated time; only the last sample changes it. `check_release_window.py` reproduces this.

The lateral loads below follow a different convention, and it is worth stating which. They are geometric quantities, taken from the fitted ratios against carriage displacement over the whole design stroke 0 to $d_{\max}$, not from a time trace, so no traction-loss transient enters them and there is nothing to truncate. The choice is not quite immaterial, because the imbalance peaks late in the stroke. Translating the release window into displacement: under the compliant integrator the first 99.5% of the simulated time carries the carriage to 99.8% of $d_{\max}$, and restricting the geometry to that span lowers the reported figures by 0.4 and 0.2 percentage points, to 7.7% at $k = 7$ and 6.8% at $k = 9$, with the load removed relative to mirror symmetry unchanged to within 0.7 kN. Under the rigid integrator the same window reaches only 93% of $d_{\max}$, because that stroke ends early at the traction-loss instability rather than at $d_{\max}$, so the late-stroke imbalance never develops at all. The figures quoted are the full-stroke ones, which are the conservative choice and the ones the fit itself minimises.

9.1 Seven members per array, $k = 7$

Three falls about the movable element driven by the leading array, four about the other. The design is given in Table 5 and plotted in Figures 6 and 7.

Table 5: Configured design at $k = 7$, seven members per array, 1,000:1 mass ratio, fitted at $\lambda = 0.03$. The greater array width is 1.98 m; the terminal ratio reaches 94% of target.

	array serving 3 falls	array serving 4 falls
span half-widths (mm)	367, 279, 127, 55, 50, 50, 50	367, 267, 134, 73, 50, 50, 50
engagement offsets (mm)	0, 247, 552, 713, 778, 807, 825	22, 407, 643, 757, 807, 825, 831

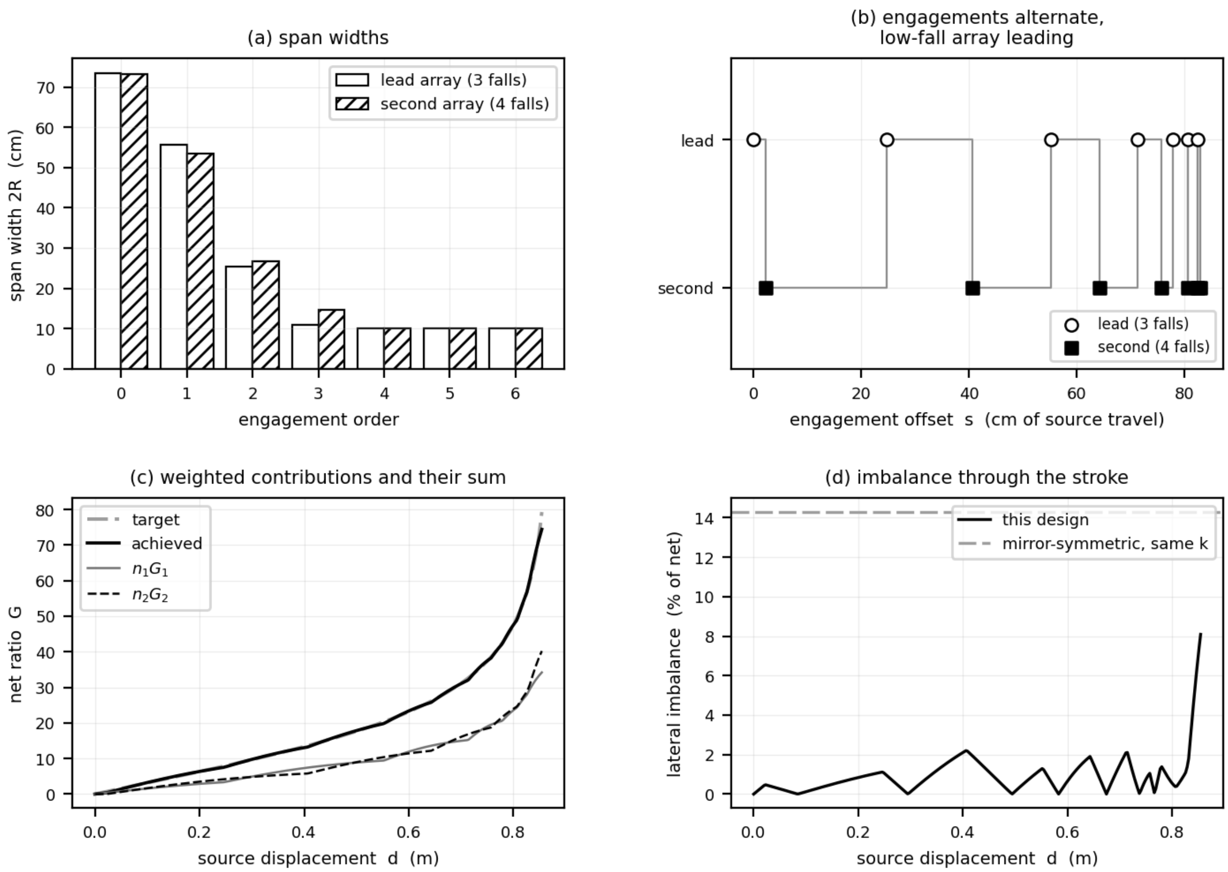

Figure 6: The seven-per-array $k = 7$ design. (a) span widths per array; (b) engagement offsets, showing alternation with the low-fall array leading; (c) the weighted contribution of each array and their sum against the target; (d) lateral load through the stroke against the mirror-symmetric value.

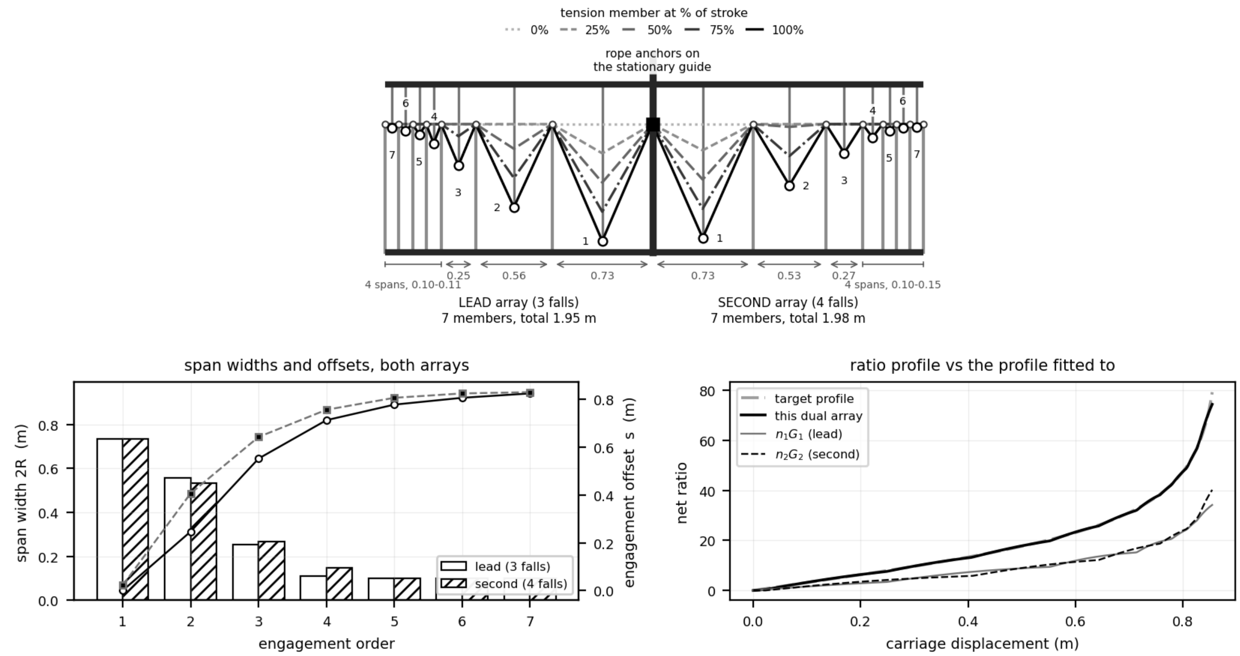

Figure 7: The seven-per-array $k = 7$ design drawn to scale: both arrays either side of the central guide, anchors on the guide, the tension member at 0, 25, 50, 75 and 100% of the stroke, members numbered in engagement order.

The order alternates throughout, the three-fall array leading. Peak lateral load 8.1%, against 14.3% for a mirror-symmetric machine at the same stage ratio. The output member carries the design tension of 3,169.6 N, so on a peak carriage reaction of 236 kN those are 19.1 kN and 33.7 kN: the unequal arrays remove 14.6 kN of transverse load from the guide bearings.

9.2 Five members per array, $k = 9$

Four falls about the movable element driven by the leading array, five about the other. The design is given in Table 6 and plotted in Figures 8 and 9.

Table 6: Configured design at $k = 9$, five members per array, 1,000:1 mass ratio, fitted at $\lambda = 0.03$. The greater array width is 1.89 m; the terminal ratio reaches 91% of target.

	array serving 4 falls	array serving 5 falls
span half-widths (mm)	488, 272, 83, 50, 50	418, 190, 65, 50, 50
engagement offsets (mm)	0, 332, 662, 779, 810	38, 523, 737, 810, 822

Peak lateral load is 7.0%, against 11.1% for a mirror-symmetric machine: 16.0 kN against 25.3 kN on a peak carriage reaction of 228 kN, a reduction of 9.3 kN.

9.3 Comparison at matched configuration

The relevant comparison is against a single array of the *same member count and stage ratio*, since the second array is added hardware. Table 7 and Figure 10 give the comparison.

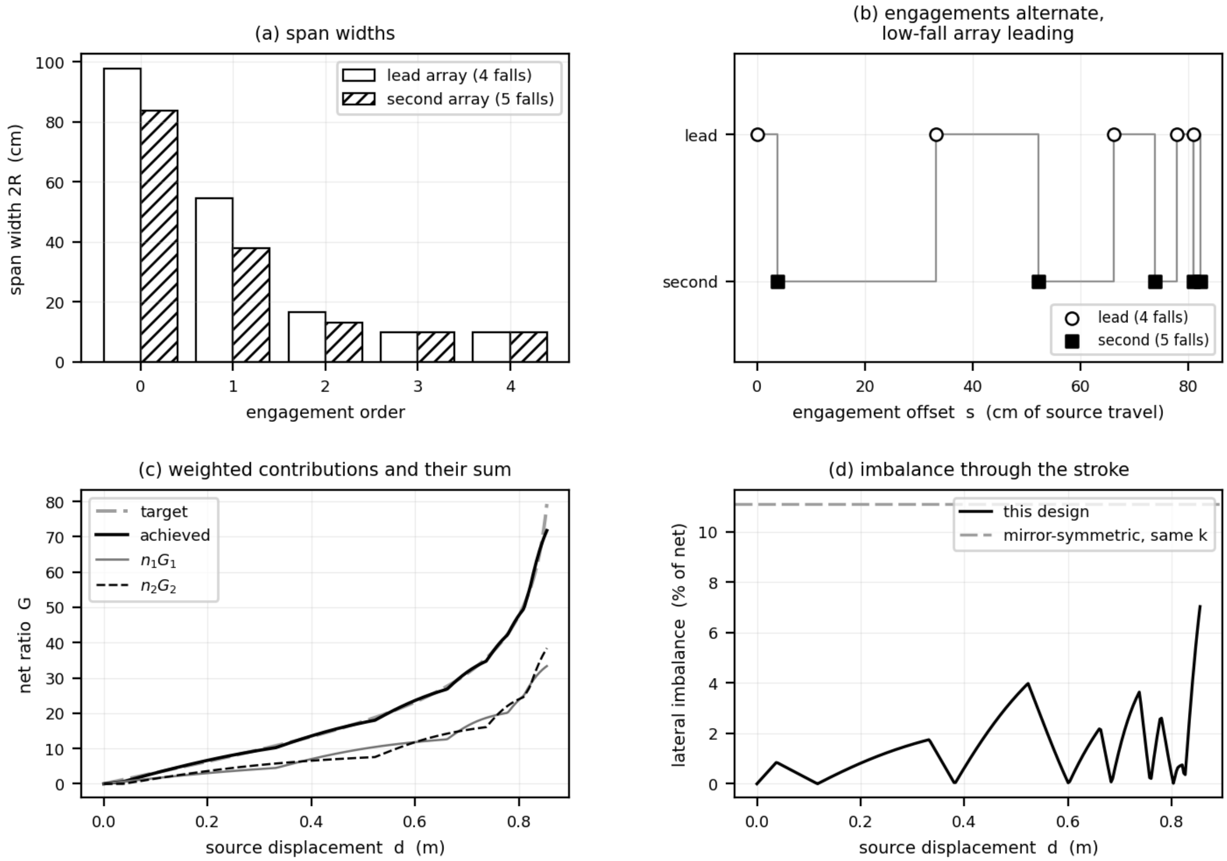

Figure 8: The five-per-array $k = 9$ design, panels as Figure 6.

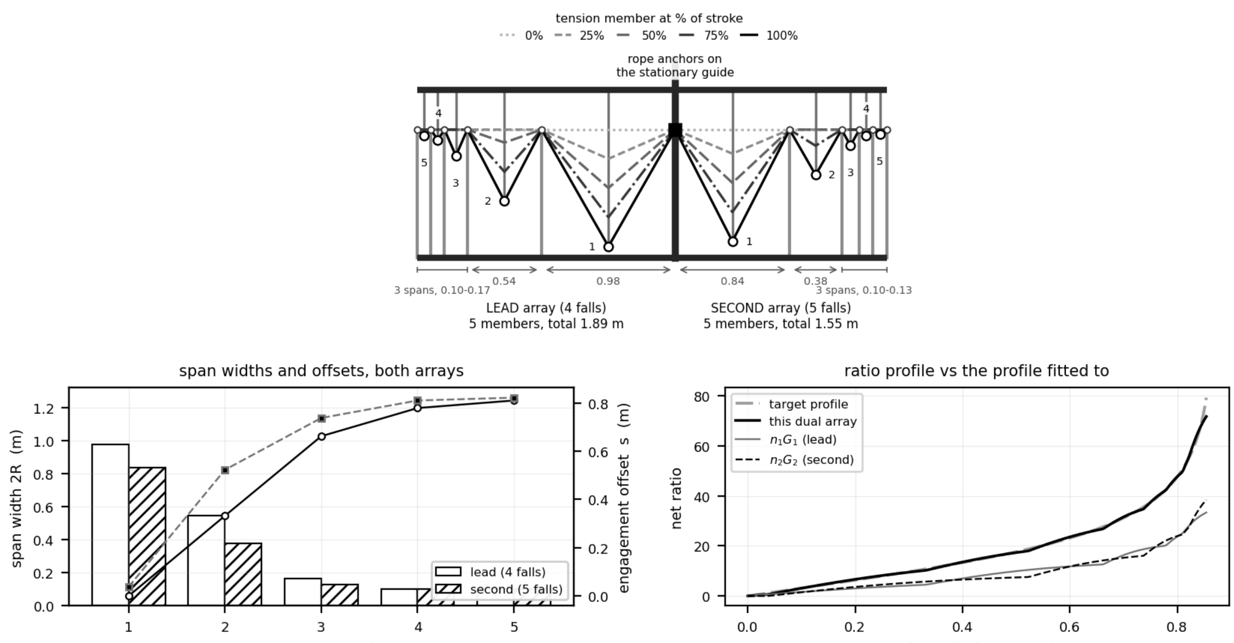

Figure 9: The five-per-array $k = 9$ design drawn to scale, as Figure 7.

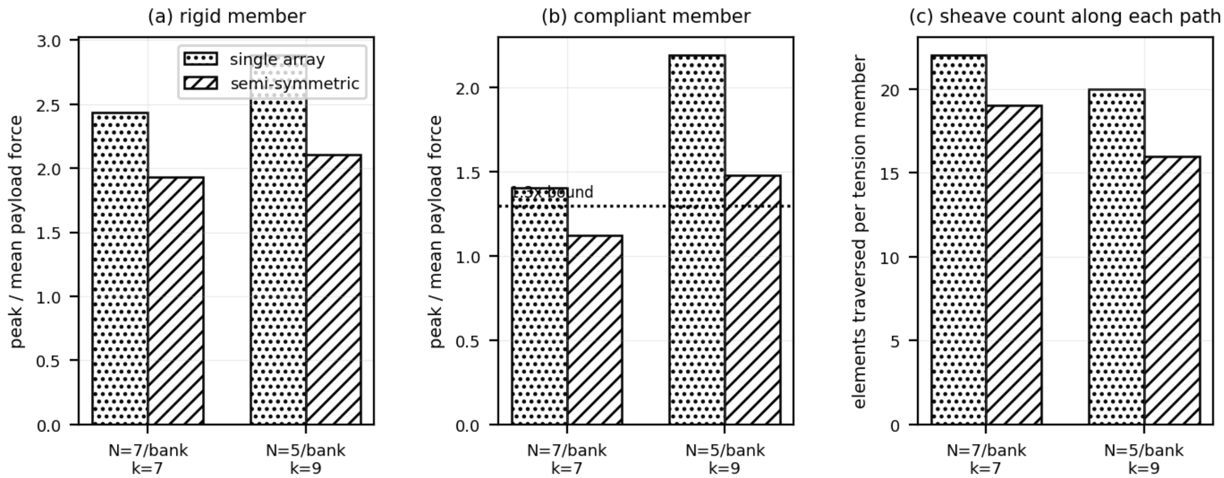

Figure 10: Comparison against a single array at the same member count and stage ratio. (a) peak-to-mean payload force, rigid; (b) the same, compliant, against the 1.3 bound; (c) elements traversed per tension member.

Table 7: Dual arrays against a single array at matched member count and stage ratio. Peak-to-mean forces are measured over the release window of Section 9. Both designs are fitted at $\lambda = 0.03$ and simulated at 16,471 N/m.

	$\langle 7,7\rangle$		$\langle 5,9\rangle$	
	single	dual	single	dual
peak-to-mean, rigid	2.44	**1.93**	2.88	**2.11**
peak-to-mean, compliant	1.40	**1.12**	2.19	**1.48**
exit velocity (m/s)	321.0	318.7	328.1	321.2
elements per tension member	22	**18 / 19**	20	**15 / 16**
contact efficiency, 2% per element	0.641	**0.688**	0.668	**0.731**

At $\langle 5,9\rangle$ neither arrangement reaches the compliant peak-to-mean of 1.3 that [5] adopts as an admissibility threshold, but the dual array cuts the excess over it by about four fifths, from 0.89 to 0.18. The comparison also covers mirror symmetry, which is worth making explicit. Two identical arrays serving n_1 and n_2 falls give a net ratio $n_1G + n_2G = kG$, exactly that of a single array of the same N members driving the same stage. A mirror-symmetric dual machine therefore delivers the *same* payload force profile as the single array, and the "single" column of Table 7 is equally its column: mirror symmetry removes the guide moment and shortens the rope paths, but it cannot alter the force profile, because the two arrays contribute indistinguishably. Only unequal arrays change it. The improvement from 1.40 to 1.12 at $\langle 7,7\rangle$, and from 2.19 to 1.48 at $\langle 5,9\rangle$, is therefore an improvement over mirror symmetry as much as over the single array.

The single-array column should not be read as a restatement of the results in [5]. That paper selects N and k jointly and reports compliant peak-to-mean forces of 1.09 to 1.19 and rigid forces of 2.0 to 2.1 across its three designs, returning nine members at $k = 5$ for the 1,000:1 case. The single arrays in Table 7 are instead *constrained* to the member count and stage ratio of the dual designs beside them, so that the comparison is like-for-like; $\langle 7,7\rangle$ and $\langle 5,9\rangle$ are not configurations that paper's procedure would choose, and the single array is correspondingly worse there than at its own optimum.

That is the substance of the comparison. Held at a configuration a single array handles poorly, the dual arrangement returns $\langle 7, 7\rangle$ to the class of performance [5] reports, 1.12 compliant and 1.93 rigid against its published bands of 1.09 to 1.19 and 2.0 to 2.1, and brings $\langle 5, 9\rangle$ most of the way there. The second array does not merely trim the $\langle 5, 9\rangle$ configuration; it removes most of the margin by which that configuration is otherwise inadmissible. The departure from symmetry that achieves this is modest: corresponding spans of the two arrays differ by at most 19 mm at $k = 7$ and 82 mm at $k = 9$, against array widths of about 2 m.

10. Scope and limitations

What is general and what is worked. The fall-count weighting (3), the balance condition (5) with its special cases (6) and (7), the member floors, the interleaving rule and the contact saving of Section 4.1 are algebraic. They hold at any mass ratio, any stage ratio and any member count, and nothing in their derivation depends on the case chosen to illustrate them. What is specific is the worked material: two configured designs at a single mass ratio of 1,000:1, taken at two stage ratios. That case is not a trivial one, and the paper is confined to it for brevity rather than because the method has been tested no further.

No hardware. Every number here is a simulation output. No dual-array RopeComb has been built.

The configured designs sit at a chosen point on the balance trade, not at exact cancellation. Fitted at a λ that favours profile accuracy, they reach G_1/G_2 of 1.13 at $k = 7$ where exact balance asks 1.33, and 1.09 at $k = 9$ against 1.25, leaving the 8.1% and 7.0% residual lateral loads reported in Section 9. Closer cancellation is available further along the sweep of Table 8: 6.4% at $k = 7$, at a residual against the target 14% worse. Which point to take is the designer's choice, not a property of the architecture.

The path-loss correction is derived but not quantified. Section 6.2 gives condition (5) in its general form, and the degrees of freedom needed to satisfy it are already present. Putting numbers to the residual requires a per-element friction model and a fixed routing for both tension members, neither of which this work provides.

Inherited modelling assumptions. The rope stiffness is derived in Section 9 from $K_{\text{rope}} = EA/L$ for a specified rope and length rather than assumed, but E and the exact routing are taken from [5] and not measured here; Appendix A bounds the consequence. Fast-path rotational inertia is applied as a post-hoc correction, as in [5], rather than carried in the model. The contact-efficiency figures use a nominal 2% loss per element, so they compare architectures rather than predict absolute performance.

11. Conclusion

Dividing a variable mechanical advantage transmission into dual arrays disposed about a central guide resolves the primary physical bottlenecks of the single array before any question of balance arises. Anchoring two tension members independently partitions the k falls of the fixed-ratio stage between separate lines, shortening the busiest rope path by $\lfloor k/2 \rfloor$ contact elements. Because a higher stage ratio simultaneously reduces the required member count and slows the movable stage blocks, whose parasitic kinetic energy scales as $1/k^2$, this partitioning extends the viable operating envelope of k while cutting frictional and bending losses along each rope path.

Permitting the two opposing arrays to be geometrically unequal transforms this dual architecture into a dynamically balanced system. Where mirror symmetry leaves an uncorrectable lateral guide load of $|n_1 - n_2|/k$ at every odd stage ratio (20% at $k = 5$) and cannot accommodate unequal path tensions, semi-symmetry makes complete lateral load cancellation attainable. The governing balance condition,

$$G_1 : G_2 = n_2 T_2 : n_1 T_1$$

dictates that the array actuating more falls, or experiencing lower transmission losses, must be assigned the smaller displacement ratio. The engagement ordering carries a second rule: the array that engages first must drive the movable element with the fewer falls. Leading with the wrong array raises the residual lateral load at $k = 5$ by half again to more than double, and is dominated on both accuracy and balance across almost the whole trade curve at $k = 7$ and $k = 9$.

Beyond reaction force cancellation, the independent geometry doubles the available degrees of freedom to $2N$, yielding a threefold improvement in target profile tracking compared to mirror-symmetric counterparts of equal member count and width. In dynamic simulations of 1,000:1 mass-ratio launches, semi-symmetric configurations improve compliant peak-to-mean payload forces from 1.40 to 1.12 at $k = 7$ and from 2.19 to 1.48 at $k = 9$, and rigid figures from 2.44 to 1.93 and 2.88 to 2.11, while reducing the elements each rope traverses.

Mirror symmetry remains an attractive choice in the regime where stage ratios are even ($n_1 = n_2$), line friction and compliances are closely matched ($T_1 \approx T_2$), and part-count economy outweighs profile accuracy; there it halves the number of distinct parts for no loss of balance. Because practical installations routinely require odd reeving geometries ($k = 2p+1$) and asymmetric routing, however, that regime is narrower than it appears. Semi-symmetry is the general kinematics for high-impulse, multi-line rope transmissions, and mirror symmetry the special case it contains.

Code and data availability

The method of this paper is implemented as a documented Python package, released under the MIT licence, in the same repository as the single-array implementation of [5]: https://github.com/ramimahdi/ropecomb. The dual-array extension is the module `dualarray_ropecomb`, which covers the fall-count weighting and balance condition of Sections 3 and 5, the staged solve and order-preserving parametrisation of Algorithm 1, the lead-assignment and interleaving checks of Section 7, and the stage-parity comparison of Sections 4.2 and 6.1. It calls the single-array package rather than reimplementing it, so both papers reproduce against one simulator.

The example `examples/09_reproduce_paper2_designs.py` reproduces the two configured designs of Section 9 together with their single-array baselines, and is the fastest check that an installation is behaving. Their full spans, offsets, engagement orders and simulated performance are frozen in `designs_16471.json`. The scripts that produced every remaining number and figure in this paper accompany them under `paper2/`, with a per-file description in the repository README.

The rigid and compliant integrators are those of [5], unchanged. Figure 5 is Figure 8 of that paper.

An interactive browser tool for designing and simulating RopeComb arrays is available at https://ropecomb.com. It is a convenience interface to the same method and is not a substitute for the package: the numbers reported in this paper were produced by the code above.

Licence

Competing interests and intellectual property

The author is the sole member of Interesting Machines LLC and has a financial interest in the commercialisation of the transmission described here. Two United States provisional patent applications, both filed before this article was made public, are directed to aspects of it: No. 64/124,997, filed 3 August 2026, to the single-array architecture; and No. 64/139,129, filed 23 August 2026, to the dual-array architecture, including the balance condition, the interleaving rule and the design method. Neither has been assigned at the time of writing. The Creative Commons licence above applies to this article and its figures; it does not grant rights under any patent, and neither does the MIT licence on the accompanying software.

AI usage disclosure

The author wrote the original simulation and fitting code. An AI assistant was used for transcribing and packaging that code, for exploratory analysis, and for drafting and editing the text of this paper, under the author's direction and review. All figures were produced by the author. All results were verified by the author against the accompanying code.

Correspondence

Rami N. Mahdi, Interesting Machines LLC, Sacramento, California, USA. ramimahdi@interestingmachines.com
ORCID: https://orcid.org/0009-0000-0817-3887

References

[1] U.S. Army Unmanned Aircraft Systems Center of Excellence. “Eyes of the Army”: U.S. Army Roadmap for Unmanned Aircraft Systems 2010–2035. Technical report, U.S. Army, Fort Rucker, AL, 2010.

[2] Krzysztof Sibilski and Grzegorz Jastrzębski. Analysis of pneumatic catapult launch system parameters. *Journal of KONES Powertrain and Transport*, 24, 2017.

[3] Michael R. Doyle, Douglas J. Samuel, Thomas Conway, and Robert R. Klimowski. Electromagnetic aircraft launch system—EMALS. *IEEE Transactions on Magnetics*, 31(1):528–533, 1995.

[4] Paul E. Chevedden, Les Eigenbrod, Vernard Foley, and Werner Soedel. The trebuchet. *Scientific American*, 273(1):66–71, 1995.

[5] Rami N. Mahdi. The RopeComb: A configurable rope-and-pulley transmission for impedance-matched mechanical launch. *engrXiv*, August 2026. doi: 10.31224/7913. URL https://doi.org/10.31224/7913.

[6] Marlow Ropes. D12 Max 99 technical datasheet, 2024. Dyneema® SK99 fibre is a product of Avient Protective Materials.

Appendix A. Sensitivity to rope stiffness and to the balance weight

The stiffness of the output member is set by $K_{\text{rope}} = EA/L$, written with a capital to keep it clear of the stage ratio k. Across the grades a designer would actually choose, and the routing lengths the machine actually needs, that spans a narrow band: 15,730 N/m for SK75 over 17.8 m, 16,471 N/m for the length of [5], and 19,663 N/m for a top-end SK99 or DM20 fibre at $E = 125$ GPa [6]. Over that band the compliant peak-to-mean of the configured designs moves by 0.02 at $k = 7$ and 0.04 at $k = 9$, which confirms the insensitivity asserted in [5, Appendix B]. The figure is not insensitive to stiffness in general: taken over an unphysical 8,000 to 30,000 N/m it ranges from 1.04 to 1.40 at $k = 7$ and 1.11 to 1.61 at $k = 9$. What makes it stable is that material and length choices are themselves constrained.

Figures 11 and 12 show both designs run against the ideal profile, in the layout of [5, Figures 12 and 13]. Panels (a) to (d) are nearly indistinguishable between the two designs and the ideal, which is the point: the geometry is fitted to reproduce the ideal ratio profile and does so. The architectures separate in panel (f). Under the rigid integrator both show the bounded sawtooth of discrete engagement, one tooth per member, larger at $\langle 5, 9\rangle$ because five members must cover the same profile that seven cover at $\langle 7, 7\rangle$. With the compliant member the sawtooth becomes a smooth ripple about the design force, and the $\langle 5, 9\rangle$ ripple is nearly four times the amplitude of the $\langle 7, 7\rangle$ one, which is the whole of the difference between 1.48 and 1.12.

The balance weight is the more consequential choice, and the designs of Section 9 take $\lambda = 0.03$. Table 8 gives the alternatives at 16,471 N/m.

Table 8: Both configured designs against the balance weight, at 16,471 N/m. Larger λ buys lateral balance and costs force uniformity. The two peak-to-mean measures do not move together: at $\langle 5, 9\rangle$ the rigid figure is best at $\lambda = 0.1$ and the compliant figure at $\lambda = 0.03$.

λ	rms	lateral load	p2m rigid	p2m compliant	exit (m/s)
$\langle 7, 7\rangle$, *seven members per array,* $k = 7$					
0.03	0.380	8.1%	1.93	**1.12**	318.7
0.1	0.399	7.2%	1.92	1.14	319.0
0.3	0.432	6.4%	2.23	1.23	319.9
$\langle 5, 9\rangle$, *five members per array,* $k = 9$					
0.03	0.640	7.0%	2.11	**1.48**	321.2
0.1	0.659	6.6%	1.88	1.62	322.2
0.3	0.699	6.0%	2.00	1.82	324.6

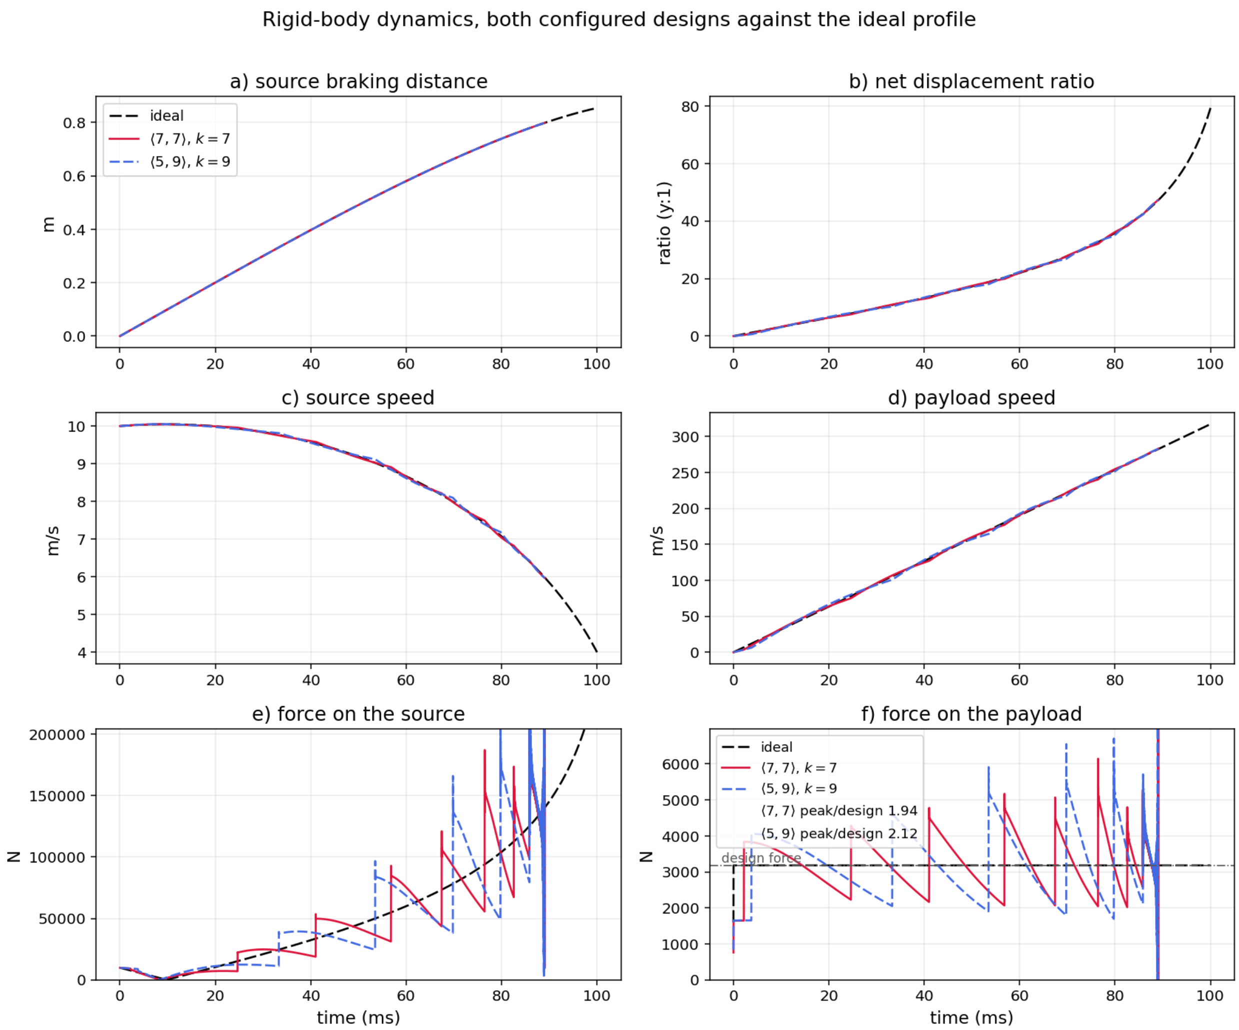

Figure 11: Rigid-body dynamics for both configured designs against the ideal profile, in the layout of [5, Figure 12]. Dashed blue is the ideal, computed by the same integration loop with the ratio solved for at each step rather than read from a geometry. Panel (f) shows one bounded sawtooth in payload force per engagement.

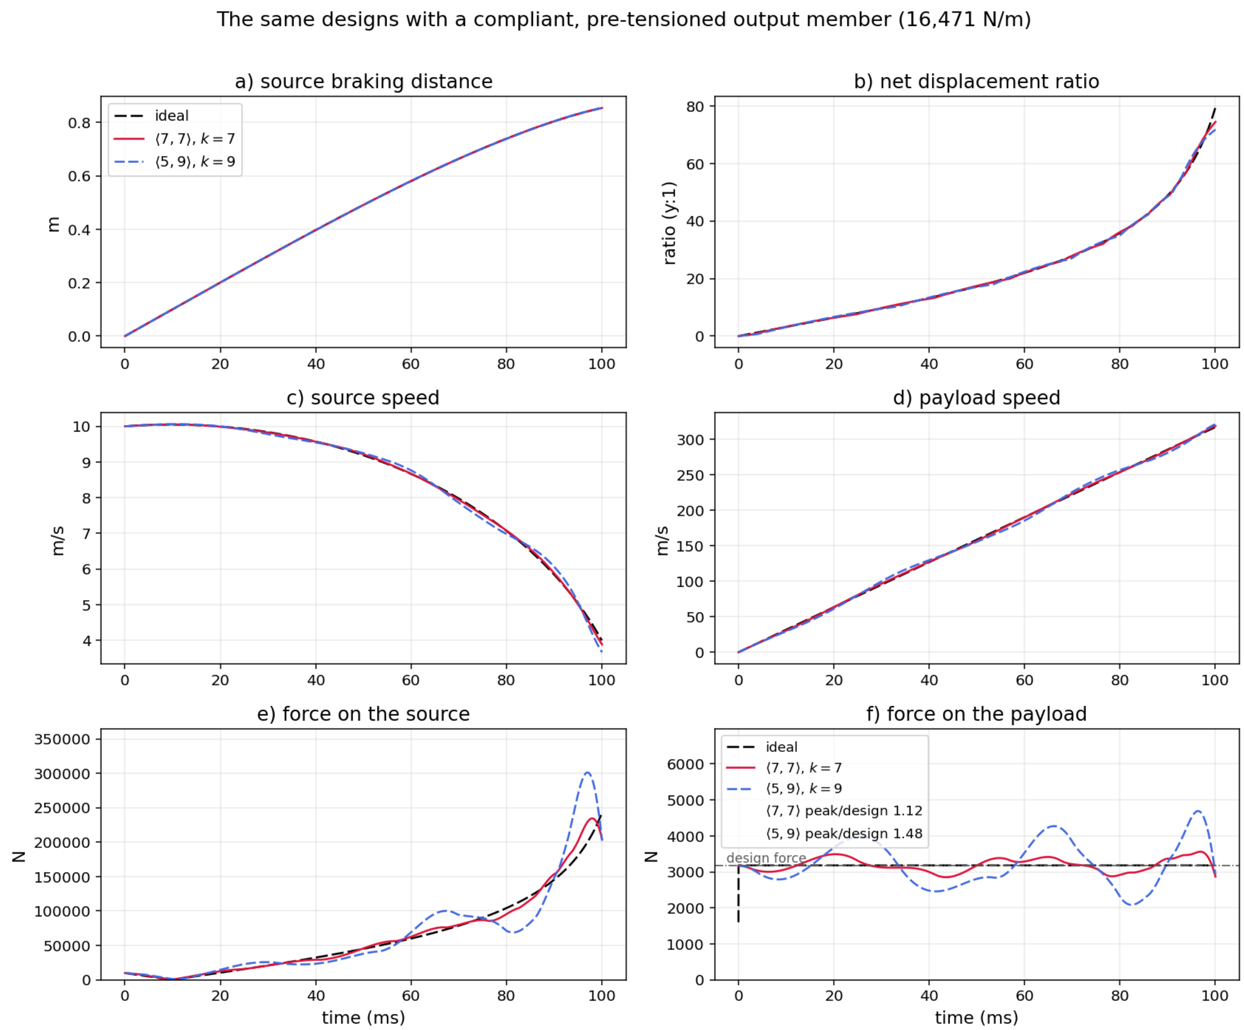

Figure 12: The same two designs with a compliant, pre-tensioned output member at 16,471 N/m. The sawtooth of Figure 11(f) is replaced by a smooth ripple about the design force. Panel (f) is drawn from zero so the ripple can be read against the absolute scale; the member is pre-tensioned to the working load, so the payload sees the design force from the first instant.