The RopeComb: A Configurable Rope-and-Pulley Transmission for Impedance-Matched Mechanical Launch

Rami N. Mahdi

Interesting Machines LLC, Sacramento, California, USA

Abstract

Transferring kinetic energy from a slow, high-force source, such as a descending mass, a ram or a spring, to a fast, light payload requires a transmission whose mechanical advantage rises continuously through the stroke. We introduce the RopeComb: a passive, configurable rope-and-pulley transmission in which a translating carriage progressively deflects a tension member into successive fixed spans. Two independent geometric freedoms per member (span width and engagement offset) make essentially any monotonically rising profile reachable, and because the resulting displacement ratio has a closed form, matching a target is a smooth least-squares fit rather than a search.

By specifying a source deceleration rather than a payload force, and solving for the force that produces it, the geometry can be optimised for uniform payload force. Across mass ratios of 100:1, 1,000:1 and 10,000:1, arrays of eight to thirteen members track that target to within 1.5% and reach its exit velocity in arrays 1.9–2.8 m wide. Measured against the work needed to raise the source again, through the braking distance as well as the height that sets its entry speed, the transmission delivers 86.3% of the input to the payload; that figure is fixed by the entry and final speeds and the braking distance alone, and does not depend on mass ratio. The three simulated designs realise it to within 0.4%. It precedes a correction for the rotational inertia of sheaves in the fast path, which is first-order, scale-invariant, and costs a further 9 to 12% of exit velocity, leaving the payload 66 to 72% of cycle input rather than 86.3%.

The stroke is terminated by a traction-loss instability, and the specified deceleration is the hardest that leaves the tension member continuously loaded throughout, with release occurring before the terminal transient. Pairing the RopeComb with a compliant, pre-tensioned output member absorbs the transients of discrete engagement, holding peak-to-mean payload force to 1.09–1.19×, where the rigid mechanism gives 2.0–2.1×. That places its force uniformity at or beyond the nominal 1.25× of a steam catapult and approaching the 1.05× of an electromagnetic launcher, achieved with passive mechanics and no electrical storage, power conditioning or steam plant. The RopeComb figures reported here are simulated; the steam and electromagnetic figures they are compared against are measured on fielded hardware.

1. Introduction

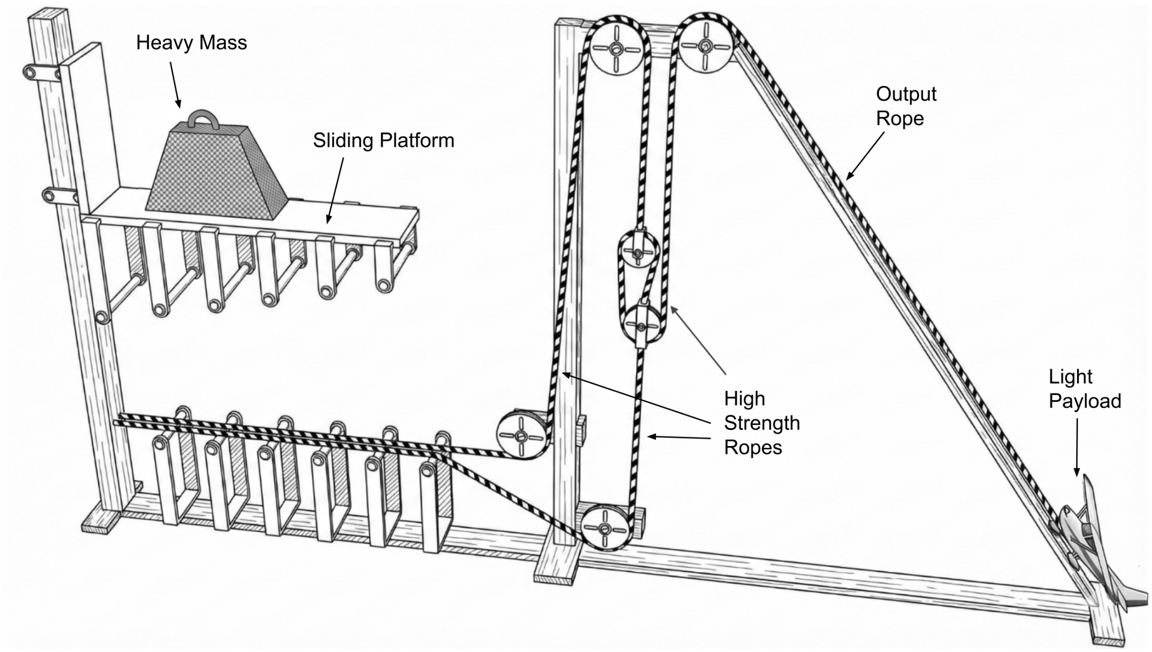

Figure 1. The RopeComb launcher. A weight rides a sliding platform down a vertical guide; pins beneath the platform drive a rope into a row of fixed brackets, each fold drawing rope in. The rope passes through a block and tackle at the mast, here a 5:1 second stage, and out along the inclined rail to the payload. Only that final run, labelled the output rope, and the sheaves carrying it move at payload speed; every other element travels at a fifth of it or slower. Member count and span spacing are illustrative, not one of the designs of §5.3.

Assisted launch transfers the energy of initial acceleration from a vehicle to ground infrastructure, permitting reduced propellant, greater payload fraction, and smaller airframes [1]. Existing systems fall into a few families, each constrained in a characteristic way.

Gravity and elastic systems (trebuchets, ballistae, bungee launchers) store potential energy and release it through a linkage whose mechanical advantage is fixed by its geometry. The force delivered to the payload is strongly non-uniform and cannot be specified independently of the linkage [2].

Pneumatic and hydraulic systems drive a piston with a working fluid. They impart a high initial shock load, requiring structural reinforcement that adds parasitic mass [3], and their attainable exit velocity is limited by the fluid and by the inertia of the piston assembly.

Electromagnetic systems can produce a substantially uniform acceleration but require large electrical infrastructure, energy storage and power conditioning, which places them out of reach at small scale or at sites with limited supply [4, 5].

A difficulty common to several of these is one of impedance. An energy source efficient at high force and low velocity is poorly matched to a light payload that must reach high velocity, and direct coupling transfers energy inefficiently. Efficient transfer requires a transmission whose ratio **rises** as the source decelerates, because a transmission of fixed ratio ceases to exert useful force once the payload velocity, referred back through the ratio, exceeds that of the source. This is as true of a trebuchet as of any other such device; the trebuchet's rising ratio is a consequence of its arm geometry rather than a designed quantity.

Known variable-ratio mechanisms, principally cammed linkages, provide a profile determined by a manufactured contour, difficult to modify and difficult to design toward an arbitrary target. What is wanted is a transmission whose ratio profile is specifiable substantially at will, adjustable without remanufacture, and realisable in components of low inertia.

This paper introduces such a transmission and, equally importantly, a method for designing it. Section 2 derives the ratio profile that produces uniform payload force. Section 3 describes the mechanism and its closed-form ratio. Section 4 presents the design method and shows why direct search fails. Sections 5 and 6 give simulation results and analyse the instability that terminates the stroke. Section 7 bounds the design space with three physical limits.

---

2. The target ratio profile

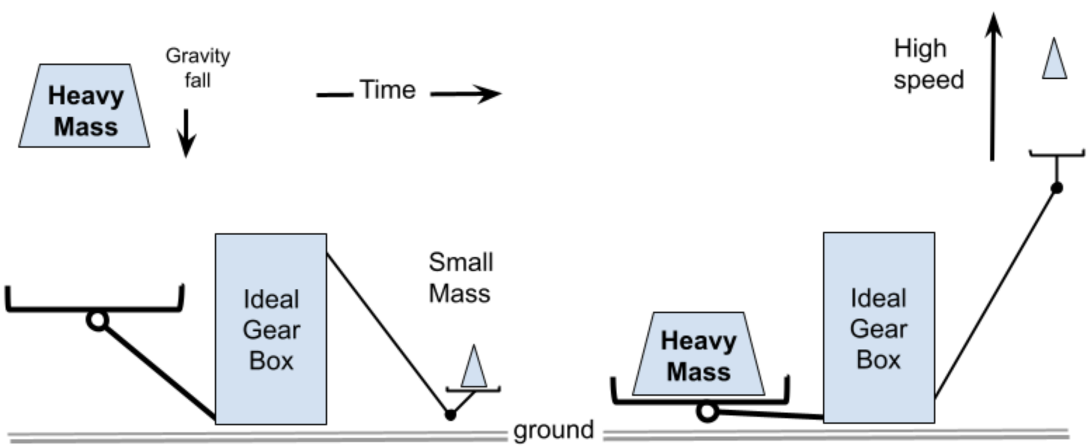

Figure 2. The transfer in the abstract: a decelerating heavy mass on one side, an accelerating light mass on the other, and a transmission between them whose ratio must rise throughout.

Consider a source mass M entering a stroke at velocity v_0 and decelerating while a payload m accelerates. Let d denote source displacement and y payload displacement, so the instantaneous ratio is $G = dy/dd$.

Uniform payload force F is desirable for several reasons. It minimises peak force for a given transferred energy, which reduces the strength, and therefore the mass, required of the payload structure, the tension member and the frame. For a given permissible peak force, it minimises the stroke length and duration. And it minimises the peak acceleration experienced by the payload, which for acceleration-sensitive payloads is the binding constraint.

Imposing uniform F and conserving power gives, for a lossless transmission,

$$\frac{dy}{dd} = \frac{\sqrt{2Fy/m}}{\sqrt{v_0^2 + 2gd - 2Fy/M}} \tag{1}$$

with $G^*(d) = dy/dd$. Equation (1) is a first-order ordinary differential equation in the **displacement** domain, containing no velocities, no discrete events and no contact conditions. The $2gd$ term presumes a gravity-driven source; for a ram or a spring it is replaced by the corresponding work term and nothing else in the method changes. It integrates in milliseconds and its solution is smooth. That the target can be expressed as a function of displacement rather than time is what makes the design problem tractable in Section 4.

The initial condition is degenerate, since $y = 0$ is a fixed point, so the integration is seeded from the small-d expansion $y \simeq (2F/m)d^2/(4v_0^2)$.

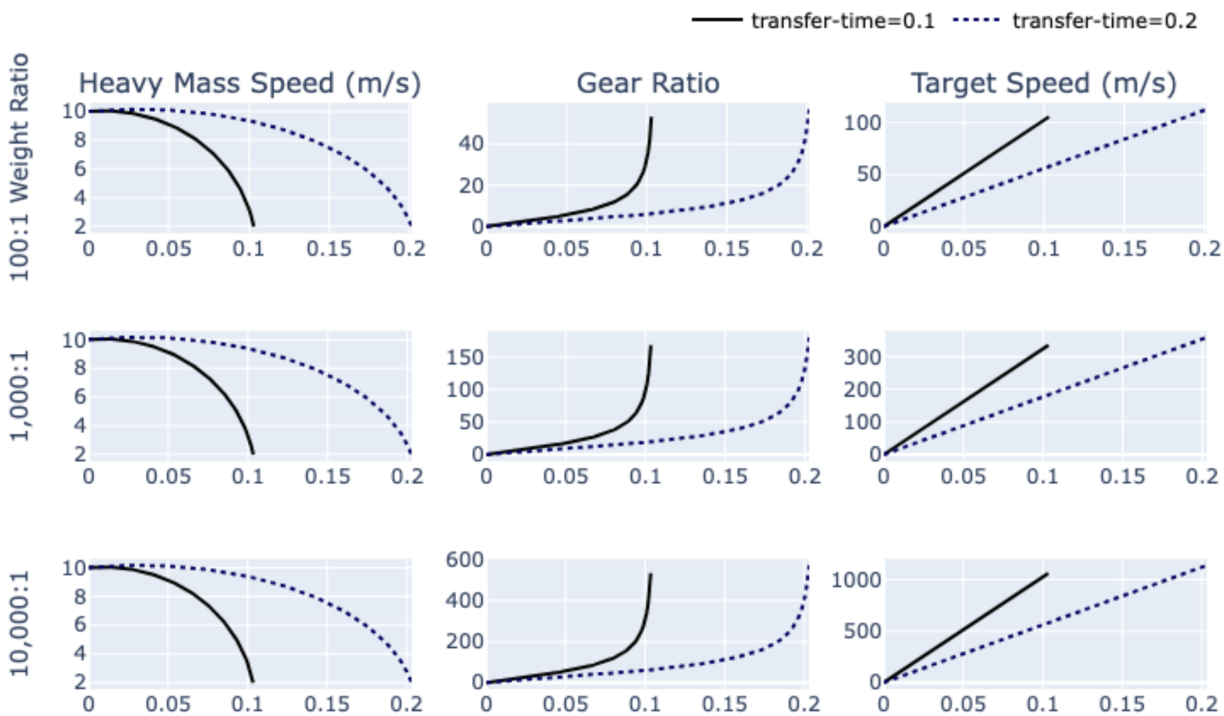

Figure 3. Ratio profiles producing uniform payload force, for mass ratios of 100:1, 1,000:1 and 10,000:1 and two transfer durations. The required ratio rises at an accelerating rate throughout.

The resulting profiles rise slowly at first and then increasingly steeply, the terminal steepness growing with mass ratio. Reproducing that shape, and in particular its accelerating rise, is the design problem.

---

3. The RopeComb mechanism

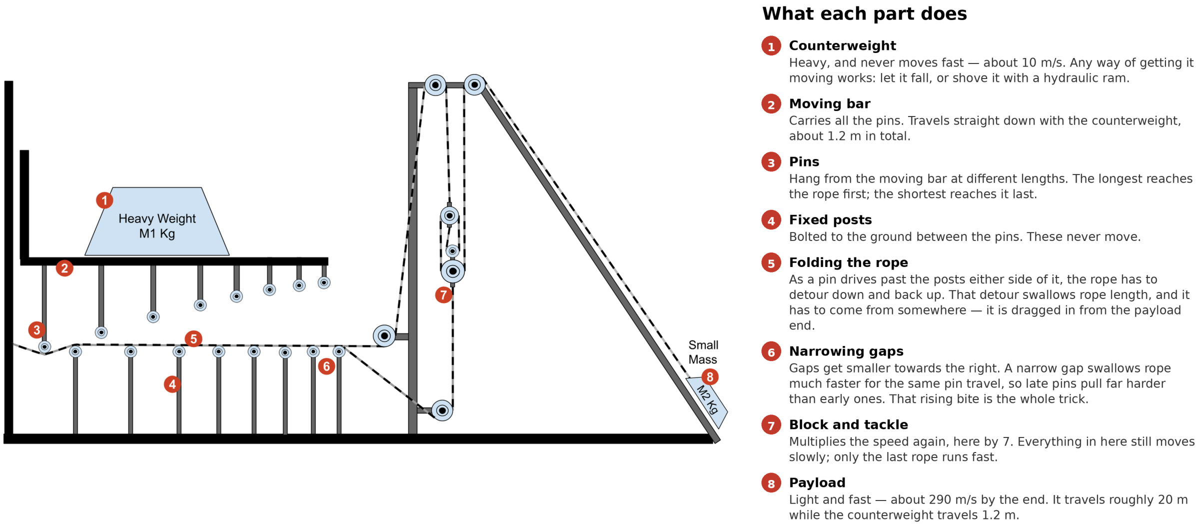

Figure 4. The same apparatus with its parts named and its action described in plain terms. The rope is folded into a row of gaps that narrow along the bar, so later pins draw rope in far faster than earlier ones.

3.1 A single engagement member

A flexible tension member spans two fixed supports separated by $2R$. An engagement member carried on a translating carriage descends into the span, deflecting the member to depth D. The member then follows two segments each of length $\sqrt{R^2 + D^2}$, so the length drawn in is

$$Y(D) = 2\left(\sqrt{R^2 + D^2} - R\right) \tag{2}$$

and the instantaneous displacement ratio is

$$\frac{dY}{dD} = \frac{2D}{\sqrt{R^2 + D^2}} = \frac{2}{\sqrt{(R/D)^2 + 1}} \tag{3}$$

This rises continuously from zero at $D = 0$ toward an asymptote of 2. The rate of rise is governed by R: a narrow span rises rapidly, a wide span slowly. A single member therefore contributes at most 2 to the ratio, and the *rate* at which it does so is a free parameter.

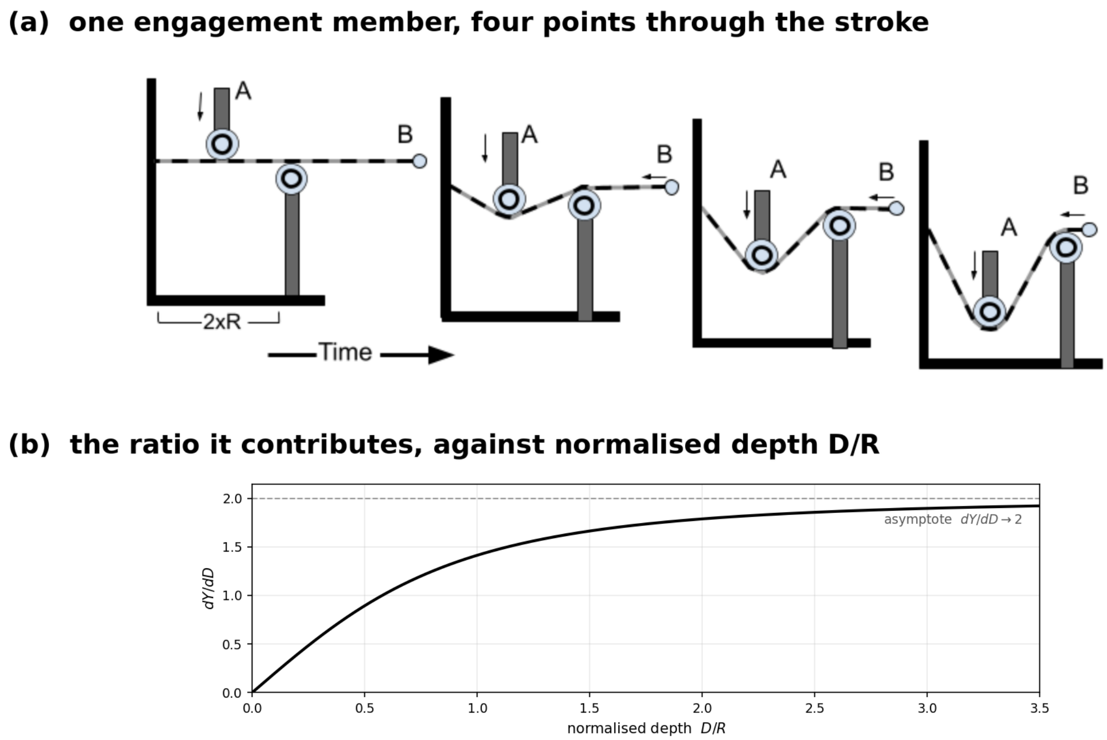

Figure 5. A single engagement member deflecting the tension member into a span of half-width R, and the resulting displacement ratio dY/dD against normalised depth D/R. The ratio rises from zero toward an asymptote of 2; a wider span rises more slowly.

3.2 Chaining

Placing N members on a common carriage driven by the source mass, so that carriage displacement and source displacement are the same quantity d, each descending into its own span of half-width R_i, and offsetting them along the direction of travel so that member i first contacts the member after the carriage has descended s_i, gives

$$G_{\text{array}}(d) = \sum_{i=1}^{N} \frac{2D_i}{\sqrt{R_i^2 + D_i^2}}, \qquad D_i = \max(0,\ d - s_i) \tag{4}$$

Equation (4) is exact and closed-form. It was verified independently by laying the member out as an explicit polyline through support and engagement-member coordinates, measuring its length, and finite-differencing against carriage displacement; agreement is to $1.5 \times 10^{-5}\%$, and cumulative take-up agrees to $7 \times 10^{-4}\%$ (Appendix C).

The array therefore provides **two independent geometric degrees of freedom per member**: the span half-width R_i, which sets how fast that member contributes, and the engagement offset s_i, which sets when. It is this pair, rather than the chaining alone, that permits an arbitrary monotonically increasing profile to be synthesised.

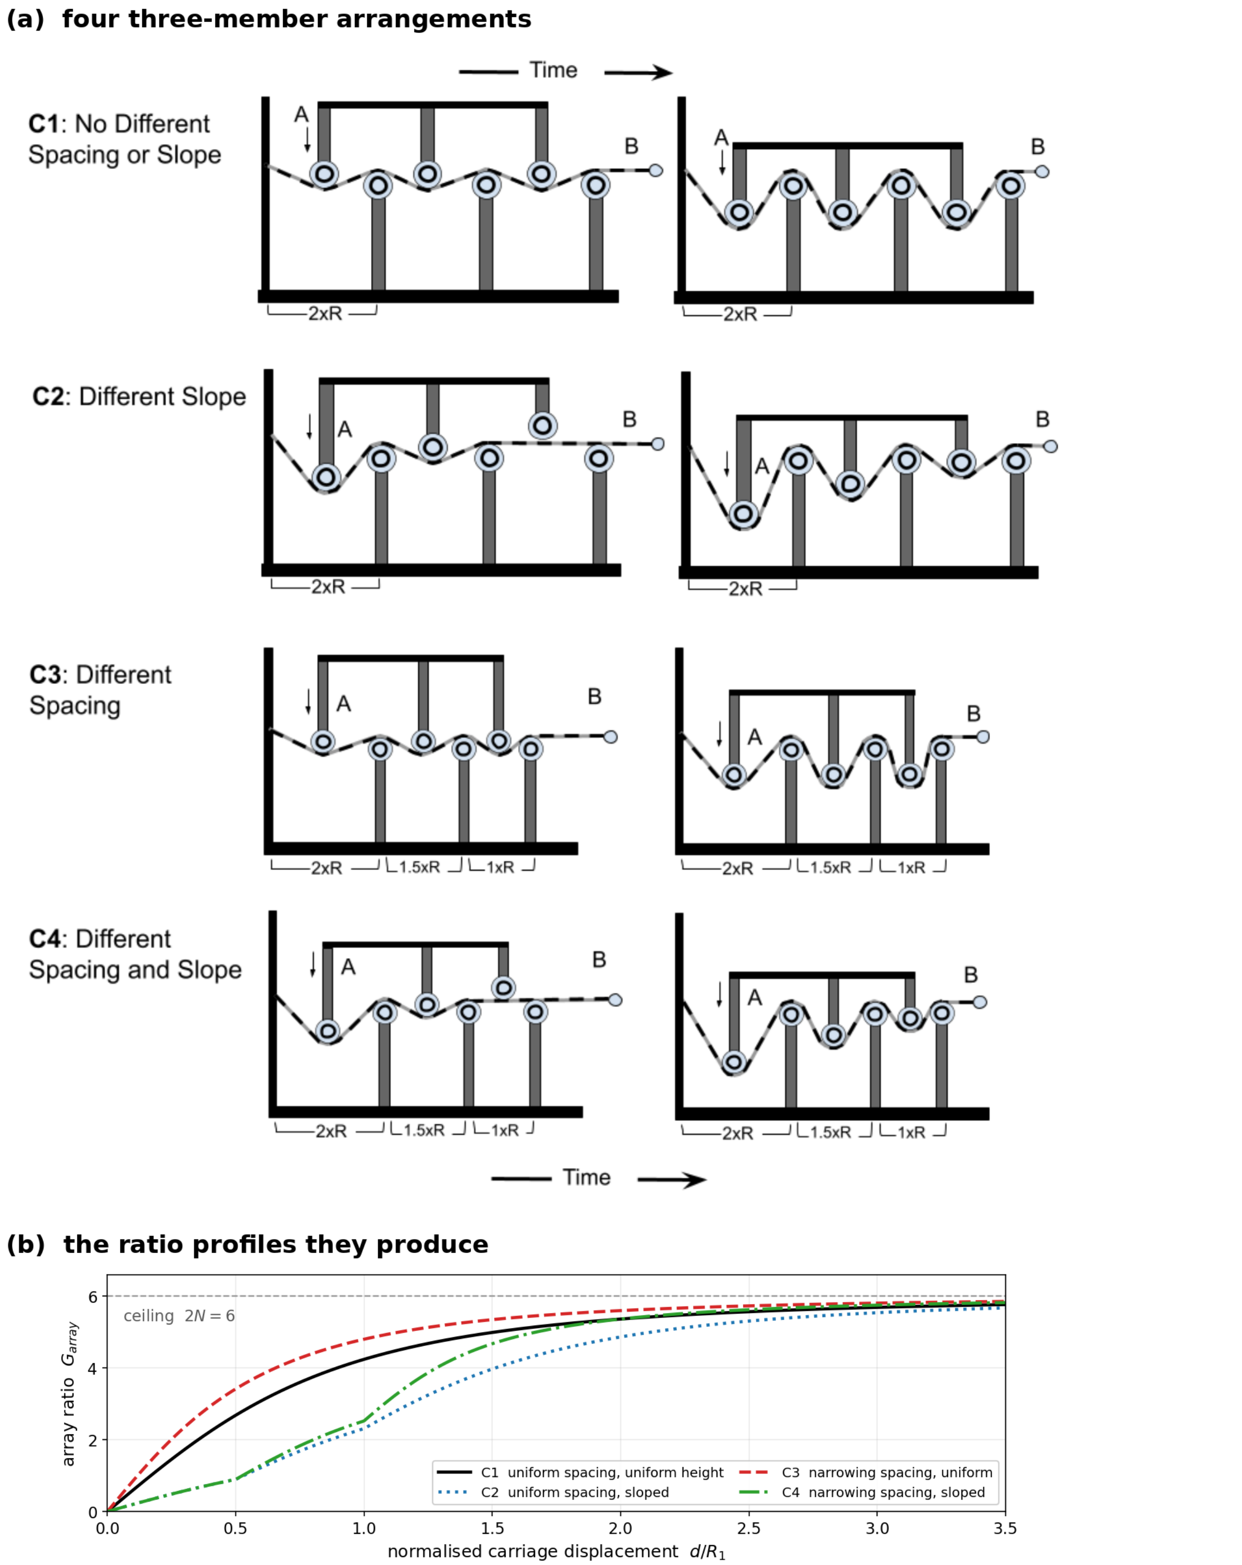

Figure 6. Four three-member configurations differing in span width and member height, and the ratio profiles they produce. Varying these two quantities alters the shape of the rise, not merely its magnitude.

3.3 Fixed-ratio stage

Reaching a high ratio by chaining alone requires $N \gtrsim G/2$ members, with attendant inertia, friction and cost. A fixed-ratio stage of ratio k interposed between the array and the payload multiplies the ratio while adding few components, giving a net ratio $G = k\,G_{\text{array}}$ and a lower bound on member count

$$N \geq \frac{G}{2k} \tag{5}$$

The stage has a second and more important function. The movable pulley assemblies translate at $1/k$ of the payload velocity, so the massive elements of the transmission are confined to low speed and only a single output member moves at payload velocity. Since kinetic energy scales as the square of velocity, this reduces parasitic energy by approximately k^2 in those elements.

Writing T for the tension in each of N_r parallel array members and θ_i for the angle of the segments at member i from the direction of travel, the force opposing the carriage is

$$F_{\text{carriage}} = N_r T \sum_{i=1}^{N} 2\cos\theta_i = N_r T\, G_{\text{array}} \tag{6}$$

since $2\cos\theta_i = 2D_i/\sqrt{R_i^2 + D_i^2}$ is exactly the term member i contributes to G_{array}. The angles differ from member to member and evolve through the stroke, so no single θ can be factored out.

The payload force is $F = N_r T/k$, so $T = kF/N_r$. Substituting recovers $F_{\text{carriage}} = F\,G$, i.e. conservation of power. Note that the tension T and the vertical reaction contributed by a segment differ by $\cos\theta$, which varies continuously through the stroke; the two are frequently conflated and are not interchangeable.

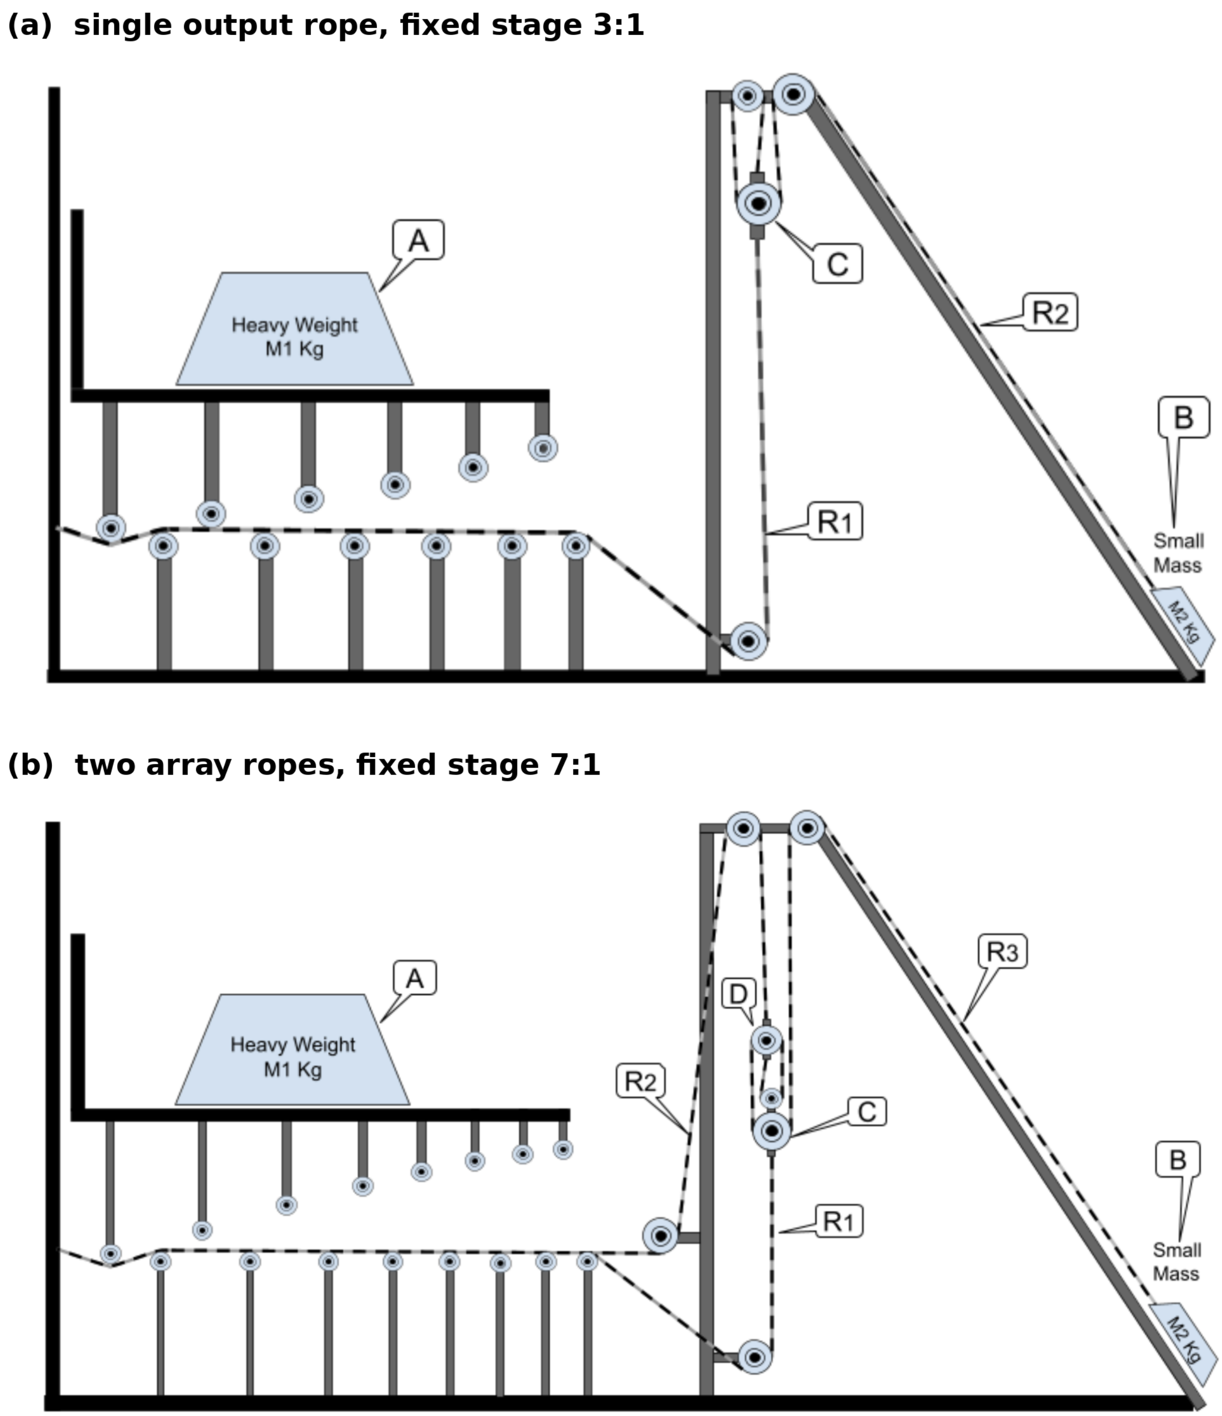

Figure 7. Apparatus. (a) A single array rope with a 3:1 second stage. (b) Two array ropes with a 7:1 second stage. In both, the movable pulley assemblies travel at $1/k$ of payload velocity, so only the output rope runs at payload speed.

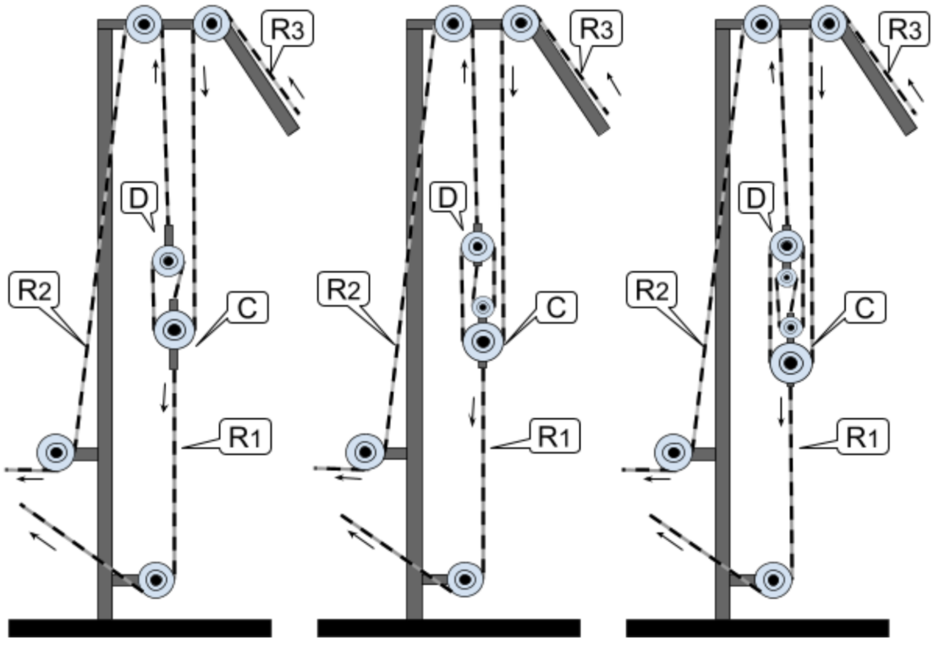

Figure 8. Second-stage arrangements giving fixed ratios of 5:1, 7:1 and 9:1.

3.4 Source of momentum

The transmission is agnostic to how the source mass acquires its velocity. Gravity is one option: a mass released from 5.1 m attains 10 m/s. But the same 50 kJ may be delivered to a 1,000 kg mass by a hydraulic or pneumatic ram over a 0.5 m stroke at 100 kN, by a spring, by combustion acting on a piston in the manner of a diesel impact hammer, or by an electric linear drive. This eliminates the requirement for elevated structure and permits mounting on a vehicle, vessel or low platform, and permits operation independent of gravitational orientation.

A characteristic of embodiments with a slowly recharged source is **power amplification**. Where the source is charged over seconds to minutes and discharged over milliseconds, peak output power exceeds mean input power by a large factor: in the worked 1,000:1 case, 50.6 kJ is delivered in 100 ms (506 kW peak) while raising the source over 60 s requires under 1 kW. High peak power is thus available at sites with only low continuous supply, without capacitor banks, pulsed alternators or comparable storage.

---

4. Design by fitting, not by search

4.1 The problem is a curve fit

Both the array ratio (4) and the target profile (1) are functions of source displacement. Determining the geometry is therefore the minimisation

$$\min_{\{R_i,\, s_i\}} \int_0^{d_{\max}} \left| k\, G_{\text{array}}(d) - G^*(d) \right|^2 dd \tag{7}$$

over $2N$ parameters. The objective is smooth and differentiable except at the engagement points, and standard least squares solves it in seconds.

The objective (7) **carries no additional terms.** It is the plain squared residual against the target profile, with no weighting, no one-sided terms and no curvature or smoothness penalties. §4.7 gives the complete procedure; §5.5 explains why nothing more is required.

Crucially, no dynamic simulation is required during fitting. The dynamics are used afterwards, to select among fitted candidates and to verify what the chosen geometry actually does.

4.2 Why direct search fails

Table 1 compares methods on the 1,000:1 case, $N = 8$, against a full-scale ratio range of 49.1.

Table 1: Least-squares fitting against direct search, on the 1,000:1 case with eight members and a full-scale ratio range of 49.1.

method	evaluations	best RMS	% of full scale
random search	300	1.474	3.00%
random search	30,000	1.040	2.12%
random search	300,000	0.683	1.39%
structured grid, 4 parameters	9,600	1.937	3.94%
least squares, 10 starts	~2,000	**0.089**	**0.18%**

Least squares improves on 300,000 random samples by a factor of 7.7 using 150× fewer evaluations. At a budget of a few hundred trials, a realistic manual effort, random search reaches RMS 1.47, some sixteen times worse.

The more consequential row is the structured grid. Parameterising the array the way a designer naturally would (a geometric progression of span widths and a power-law stagger, four parameters, 9,600 combinations) performs **worse than uniform random sampling**. Low-dimensional families of this kind, including uniform spacing, uniform slope and geometric progression, do not contain the optimum. No amount of search within them converges, which is a sufficient explanation for the difficulty of designing these arrays by hand.

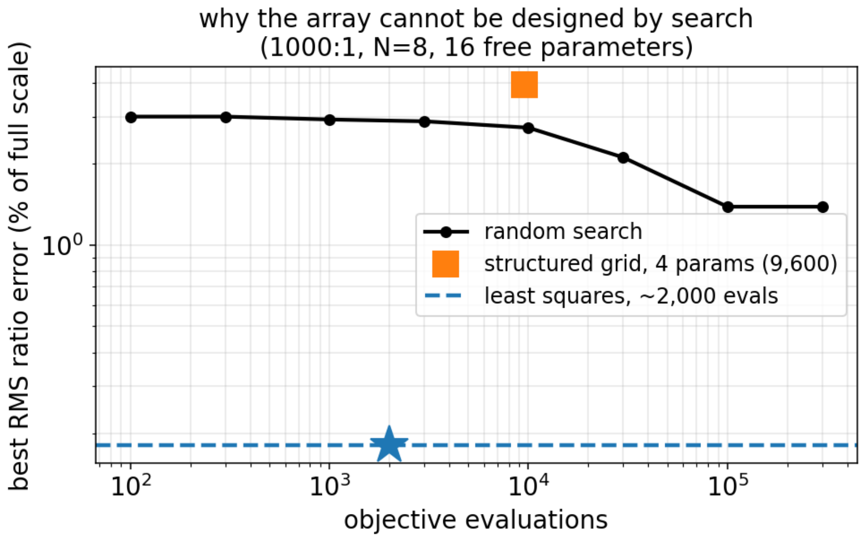

Figure 9. Best residual against objective evaluations for the 1,000:1 case with eight members (sixteen free parameters), computed against the earlier target profile of §5.5. Least squares reaches 0.18% of full scale in roughly two thousand evaluations; random search remains above 1.3% after three hundred thousand; and a structured grid over four intuitive parameters is worse than random sampling, the optimum lying outside that family.

4.3 The optimum is not intuitive

The geometry of Table 2, and the search comparison of §4.2, were computed against an earlier and more demanding target profile. They are retained because the point they make (that the optimum is not reachable by intuition, and that least squares finds it where search does not) is a statement about the objective's shape rather than about any particular target. The canonical designs of §5.3 are fitted to the target of §5.2.

A representative fitted geometry (1,000:1, $N = 8$):

Table 2: A representative fitted geometry, 1,000:1 with eight members, computed against the earlier target of §5.5. Span width is non-monotonic and engagement intervals contract steadily; neither pattern is expressible in two or three parameters.

member	s_i (m)	R_i (m)
1	0.003	0.436
2	0.197	1.155
3	0.353	1.015
4	0.497	0.854
5	0.616	0.627
6	0.706	0.422
7	0.773	0.261
8	0.824	0.150

Span width is non-monotonic: modest at the first member, maximal at the second, thereafter decreasing eight-fold. Engagement intervals contract steadily (0.194, 0.157, 0.144, 0.119, 0.090, 0.068, 0.050 m). Neither pattern is expressible in two or three parameters.

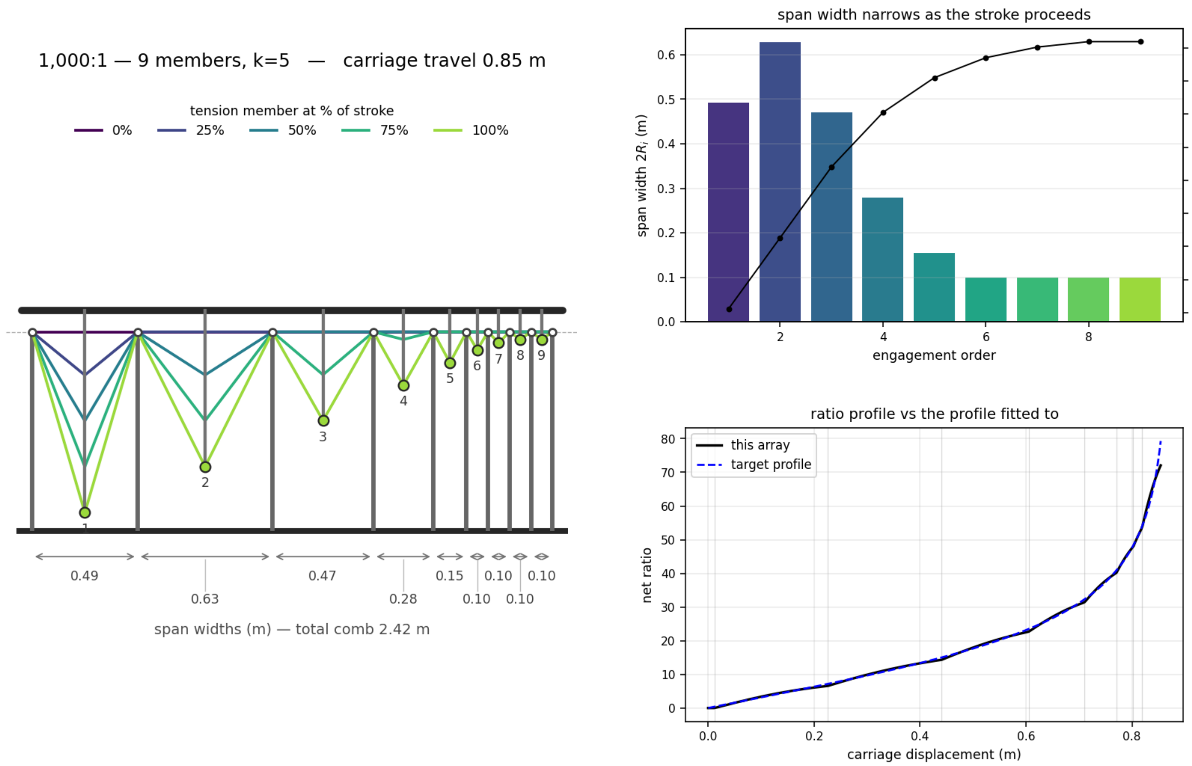

Figure 10. Fitted array geometry for the canonical 1,000:1 design of §5.3, not the eight-member fit tabulated above. The member is drawn at five points through the stroke, span widths are dimensioned beneath, and the design variables are plotted in engagement order. The dashed curve in the lower-right panel is the target *ratio profile* $G^*(d)$ that the array was fitted to, which is a function of displacement; the dashed curves in Figures 12 and 13 are the ideal *run* against time. Span width peaks at the second member and narrows to the 10 cm floor by the last.

4.4 Tolerance

Fits initiated from differing random starting configurations converge to geometries differing by 6–13% in their parameters while attaining comparable residuals. The solution manifold is therefore broad, and the apparatus is correspondingly tolerant of departures from any nominal geometry: it does not require exact positioning of its members. This is a favourable property for construction, and it is testable independently of any performance claim.

4.5 The array cannot be convex

Each engaged member contributes $f(D) = 2D/\sqrt{R^2 + D^2}$, the term summed in (4), giving

$$f''(D) = \frac{-6R^2 D}{(R^2 + D^2)^{5/2}} < 0 \quad \text{for all } D > 0 \tag{8}$$

Every engaged member is individually concave: its rate of contribution decelerates from the moment it engages. The array can therefore only approximate a convex target through the staggering of onsets: each engagement injects a fresh upward step in slope, and between engagements the curve is necessarily concave.

This is the origin of the force ripple. Between engagements dG/dd decays, the payload gains on the rope, and tension sags; a member engages, the slope steps up, and tension spikes. It also explains why the concave *fraction* is the wrong quantity to minimise: the best configurations found here are the most concave by that measure. What correlates with force quality is the **arc length** between engagements, and equivalently the slope deficit accumulated over each arc, which falls from 67% to 46% as members are added.

4.6 Member count and pruning

Members whose engagement offset lies beyond the terminal carriage displacement never engage and can be deleted without changing the simulated result. Fitting with more members than are intended to be built, then pruning, is therefore free, and it is better than fitting the smaller number directly: the surplus members act as slack variables that let the optimiser place the real ones more advantageously.

The pruning criterion must be evaluated on the **compliant** stroke, not the rigid one. A compliant output member lets the carriage travel further, so members that look redundant in the rigid model do engage in the real machine. Under a target the array cannot reach this is worth several percent of exit velocity: at 1,000:1 under the superseded target of an earlier draft, the rigid stroke ended at 0.764 m while the compliant stroke reached 0.836 m, and the two members engaging in that interval were the difference between 309.5 and 321.1 m/s at identical peak force.

Under the target of §5.2 the question does not arise. **Every fitted member engages**, in all three cases, at every member count tried, across all 1,350 simulated candidates. The fitter no longer parks surplus members beyond the end of the stroke, because it is no longer straining to reach a ratio it cannot achieve. The fitted count is then simply the member count, and the sweep at 1,000:1 with $k = 5$ reads as in Table 3:

Table 3: Member count against array width, peak force and exit velocity at 1,000:1 with $k = 5$. The trade is width against force smoothness, not member count against exit velocity.

members	built	array width	peak force, compliant	exit velocity, compliant
4	4	0.60 m	1.82×	269.7 m/s
5	5	0.84 m	1.57×	297.8 m/s
6	6	1.13 m	1.39×	310.1 m/s
7	7	1.48 m	1.26×	315.7 m/s
8	8	1.89 m	1.16×	318.1 m/s

The trade is width against force smoothness, not member count against exit velocity. Exit velocity saturates at six members and is flat thereafter; what further members buy is a lower peak force, at roughly 0.2 m of additional array each.

This sweep fixes $k = 5$ and varies only the member count. The procedure of §4.7 instead selects N and k jointly, and returns nine members at $k = 5$ for this case.

4.7 The design procedure

The pieces above combine into the single procedure of Algorithm 1. Nothing in it is tuned: the only quantities a designer supplies are the masses, the entry velocity and the deceleration to be achieved.

Algorithm 1. The RopeComb design procedure

```
INPUT   M          source mass                    OUTPUT  k          fixed-stage ratio
        m          payload mass                           {R_i}      span half-widths
        v0         entry velocity                         {s_i}      engagement offsets
        v_end, T   deceleration to achieve

1  SOLVE THE TARGET                                                       (§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 G*(d)

2  GRID SEARCH                                                            (§5.3)
   for N = 2 .. 16:
     for k in {1, 2, 3, 5, 7, 9}:
       for restart = 1 .. 50:
         w      ~ U(0.2, 0.8)                    mean span for this restart
         R_i(0) = 0.5 * w * U(0.6, 1.4)          half-widths, i = 1..N
         s_i(0) = sort( U(0, 0.85 * d_max) )     engagement offsets
         {R,s}  = argmin  SUM_d ( k*G_array(d) - G*(d) )^2
                  bounded least squares in log R; no other terms in the objective
       keep the 5 restarts of lowest residual

3  EVALUATE each kept candidate                                     (§4.6, §5.1)
   simulate rigid      -> carriage travel d_r
   simulate compliant  -> carriage travel d_c
   prune members with s_i >= max(d_r, d_c)
   re-simulate the pruned array, rigid and compliant
   p2m = max(F_payload) / mean(F_payload) over the first 99.5% of the trace
         (the release precedes the terminal traction-loss pulse, §5.6)

4  SELECT
   admissible: p2m_compliant <= 1.3  AND  v_compliant >= 0.99 * v_ideal
   score = (55 - N_eff - k - W)/50 * v_rigid * v_compliant / sqrt(p2m_r + p2m_c)
           W = total array width in metres
   return the highest-scoring admissible candidate
```

The score rewards exit velocity in both models, penalises the peak-to-mean force in both, and discounts cost through $N_{\text{eff}} + k + W$: member count, fixed-stage ratio and array width, the three quantities that set what the machine costs to build. Width enters in metres and with unit weight, which at these scales makes a metre of array about as expensive as one engagement member. The admissibility filter is applied first because the score alone will trade away exit velocity: at 10,000:1 the highest-scoring candidate overall reaches only 97% of its ideal, and is excluded. Its exact form matters less than it appears: across the 1,000:1 grid, scoring on the rigid model alone and scoring on both agree on the winner, and the top ten candidates by either measure all lie within 0.3% in exit velocity. What the score mainly does is break ties among designs that have already saturated the ideal.

Two properties of the procedure are worth stating explicitly. **No dynamic simulation enters the fit**: step 2 is pure curve fitting against the closed-form ratio (4), which is why 4,500 fits per case, 13,500 in all, are affordable; only the 450 survivors per case are ever simulated. And **the**

grid is not a refinement of a default: member count and fixed-stage ratio are not separable, so searching them jointly is necessary rather than merely thorough. Fixing k and sweeping N, as earlier versions of this work did, returns a different and worse answer at every mass ratio.

---

5. Simulation results

5.1 Model

The rigid-body model treats the tension member as inextensible and enforces the unilateral contact condition (the member may pull but not push) by advancing the payload by the greater of its free-flight displacement and the displacement required to remain in contact. It contains no friction, no rotational inertia, no member mass and no aerodynamic drag. Pseudocode is given in Appendix A. This model is used as the baseline throughout because it has no free parameters; a compliant variant, which requires stiffness and pre-tension to be specified, is deferred to Appendix B.

5.2 Target selection: how hard to decelerate the source

Every result below depends on a choice that earlier drafts of this work made badly. The ideal profile is not defined until one fixes *how much* the source mass is asked to give up, and that choice was previously made by picking a payload force and seeing what came out. It should be made the other way round: state the deceleration, then solve for the force that produces it.

We specify that **the source mass decelerates from 10 m/s to 4 m/s over a 0.1 s stroke**, and solve for the uniform payload force F that achieves exactly that. This is a one-dimensional root find on F; at the working time step it reaches the specified final speed to within 0.002 m/s. The solved values are given in Table 4.

Table 4: The design force solved from the specified deceleration, and the profile it implies. The first two rows are the specification and are identical in all three cases; everything below them follows from it. The braking distance is common to the three because the deceleration is.

	100:1	1,000:1	10,000:1
specified deceleration	10.0 -> 4.0 m/s	10.0 -> 4.0 m/s	10.0 -> 4.0 m/s
specified stroke	0.100 s	0.100 s	0.100 s
solved payload force F	999.0 N	3169.6 N	10033.8 N
braking distance	0.854 m	0.854 m	0.854 m
ideal exit velocity	99.9 m/s	317.0 m/s	1003.5 m/s
terminal ratio	25.0:1	79.3:1	251.0:1

Why 4 m/s. Stopping the source at 4 m/s leaves it with 16% of its entry kinetic energy.

The accounting must be taken against the energy put *into* the machine, which for a gravity-driven source is not $\frac{1}{2}Mv_0^2$. Before the next launch the source has to be raised again, through the braking distance as well as the height that gave it its entry speed, so the input per cycle is

$$E_{\text{in}} = Mg\,(h_0 + d_{\text{max}}), \qquad h_0 = \frac{v_0^2}{2g} \tag{9}$$

a total fall of 5.96 m in all three cases. Against that denominator (Table 5):

Table 5: Energy per launch cycle, measured against the work needed to raise the source again. The efficiency is the same in all three cases because it cannot depend on mass ratio. The residual kinetic energy the source carries away at 4 m/s is a loss unless it is recovered.

	100:1	1,000:1	10,000:1
source fall per cycle, $h_0 + d_{\max}$	5.96 m	5.96 m	5.96 m
input energy $Mg(h_0 + d_{\max})$	5,837 J	58,371 J	583,707 J
delivered to the payload	5,037 J	50,371 J	503,707 J
retained by the source at 4 m/s	800 J	8,000 J	80,000 J
transfer efficiency	**86.3%**	**86.3%**	**86.3%**

The three columns are identical, and necessarily so. Substituting $gh_0 = \frac{1}{2}v_0^2$ into the ratio gives

$$\eta = \frac{\frac{1}{2}(v_0^2 - v_f^2) + g\,d_{\max}}{\frac{1}{2}v_0^2 + g\,d_{\max}} \tag{10}$$

in which the source mass has cancelled. Transfer efficiency is fixed by the entry speed, the final speed and the braking distance alone; mass ratio sets the required force, the exit velocity and the transmission ratio, but cannot appear here. The simulated ideal runs reproduce the analytic figure to within 0.9%, 0.3% and 0.04% at 100:1, 1,000:1 and 10,000:1, the discrepancy being integration drift and largest for the smallest source mass.

Quoting the payload's energy against $\frac{1}{2}Mv_0^2$ instead would give 100%, 100% and 101%, because gravity keeps doing work through the stroke. That ratio exceeds unity and is therefore not an efficiency: its denominator is not the input.

The gravitational term is close to efficiency-neutral. A horizontal ram or spring does no work during the stroke, so both the gravitational input and the gravitational contribution to the payload disappear together, and the same deceleration returns $1 - (v_f/v_0)^2 = 84\%$. Gravity raises that to 86.3% rather than transforming it, because the extra work is delivered at full transmission efficiency while the retained 16% is a fixed penalty. The choice of source therefore sets the required force and the recharge machinery, not the efficiency.

The model is lossless, so these are upper bounds; §5.9 and §7.2 set out what is omitted.

Asking for more than this is where the machine breaks. Figure 11 sweeps the specified final speed at each case's locked geometry.

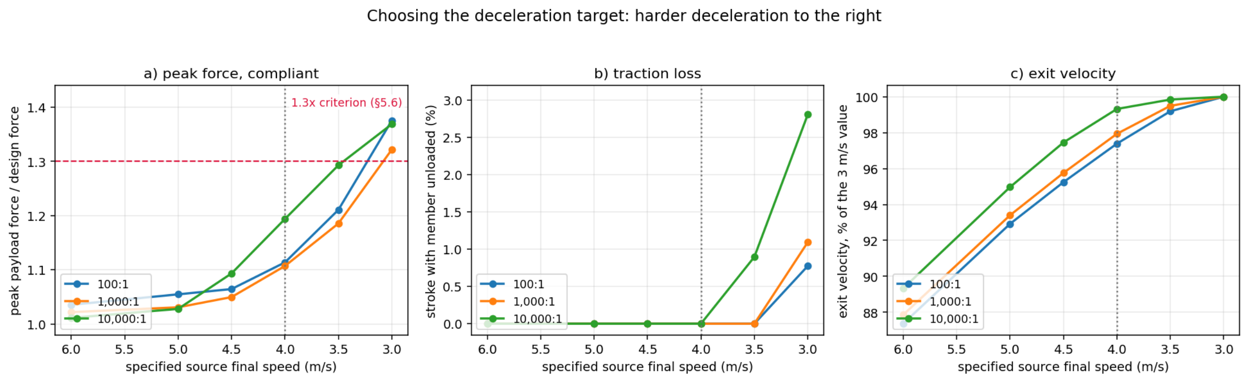

Figure 11. Why the source is stopped at 4 m/s. Left: compliant peak payload force. Centre: fraction of the stroke during which the tension member carries no load. Right: exit velocity relative to the 3 m/s value. Harder deceleration is to the right in each panel.

Three things happen together as the deceleration is made more severe. Table 6 holds N and k at the values selected in §5.3 and refits the geometry at each target with a reduced restart budget, so its 4.0 m/s column approaches but does not exactly reproduce the canonical designs:

Table 6: Peak force, traction loss and exit velocity against the specified source final speed, with N and k held at the values selected in §5.3 and the geometry refitted at each target.

final speed	3.0	3.5	**4.0**	4.5	5.0	6.0
peak force, 100:1	1.38x	1.21x	**1.11x**	1.06x	1.05x	1.04x
peak force, 1,000:1	1.32x	1.19x	**1.11x**	1.05x	1.03x	1.02x
peak force, 10,000:1	1.37x	1.29x	**1.19x**	1.09x	1.03x	1.01x
slack fraction, 100:1	0.8%	0.0%	**0.0%**	0.0%	0.0%	0.0%
slack fraction, 1,000:1	1.1%	0.0%	**0.0%**	0.0%	0.0%	0.0%
slack fraction, 10,000:1	2.8%	0.9%	**0.0%**	0.0%	0.0%	0.0%
exit, % of that case's ideal	98–100	99–100	**100**	100	100	100

Traction loss is absent at 4.0 m/s and above in all three cases, and appears below it in all three. At 3.0 m/s the member is unloaded for 0.8%, 1.1% and 2.8% of the stroke at 100:1, 1,000:1 and 10,000:1; only 10,000:1 also shows it at 3.5 m/s, at 0.9%. Peak force rises monotonically as the deceleration is made more severe, and by 3.0 m/s all three cases have crossed the 1.3x bound. Exit velocity, meanwhile, stops reaching its own ideal: at 3.0 m/s the 10,000:1 case attains only 98% of a target it can no longer follow.

Easing the deceleration in the other direction costs velocity and array width without buying much: going from 4 to 5 m/s at 1,000:1 lowers the ideal exit from 317 to 303 m/s and more than doubles the array, from 2.4 m to 5.5 m, for 0.08x of peak force.

A note on how the onset depends on member count. An earlier version of this sweep, run at seven members rather than nine for the 1,000:1 case, showed slack at 3.5 m/s as well as at 3.0. Adding members postpones the onset, because it is the discreteness of engagement that drives it (§6.2). The location of the knee is therefore not a fixed property of the mechanism but a joint property of the target and the member count; what is robust across every configuration tried is that 4 m/s is clear of it while 3.0 m/s is not.

4 m/s is therefore taken as the knee: the hardest deceleration that leaves the tension member continuously loaded in all three cases, with margin. It is a design input rather than a derived optimum, and a machine with more members could be pushed further; 4 m/s is the value at which no case in this study shows slack and every case reaches its ideal.

5.3 Three canonical cases

Each case is fitted to the profile of §5.2 by plain least squares. All enter the stroke at $v_0 = 10$ m/s with a 1 kg payload and a minimum span width of 10 cm. Force scales linearly with payload mass. Table 7 gives the geometry and performance of the three canonical designs.

Table 7: The three canonical designs, each produced by the procedure of §4.7. Peak forces are multiples of the design force; compliant figures are for an output member pre-tensioned to the working tension.

	100:1	1,000:1	10,000:1
design payload force	999 N	3,170 N	10,034 N
engagement members	**8**	**9**	**13**
fixed-stage ratio k	**2**	**5**	**9**
array width	2.82 m	2.42 m	1.92 m
narrowest span	10 cm	10 cm	10 cm
ideal exit velocity	99.9 m/s	317.0 m/s	1003.5 m/s
exit velocity, rigid	89.8 m/s	284.8 m/s	885.9 m/s
exit velocity, compliant	100.2 m/s	318.2 m/s	1004.8 m/s
the same, corrected for fast-path sheave inertia	*87.9–91.5 m/s*	*279.1–290.5 m/s*	*881.3–917.2 m/s*
fraction of ideal	**100%**	**100%**	**100%**
peak-to-mean force, rigid	2.05 x	2.07 x	2.10 x
peak-to-mean force, compliant	**1.12 x**	**1.09 x**	**1.19 x**
objective	plain least squares	plain least squares	plain least squares
selection	grid search over N and k	grid search over N and k	grid search over N and k

Peak forces are multiples of the design force. Compliant figures are for an output member pre-tensioned to the working tension (§5.7).

Three features of this table matter more than the headline velocities.

Every case reaches essentially the whole of its ideal (100.3%, 100.4%, 100.1%) and every case meets the peak-force criterion of §5.6. Neither was true of any earlier target.

Member count and fixed-stage ratio both rise with mass ratio, but far more slowly than it does. Across a hundredfold change in mass ratio, member count goes from eight to nine to thirteen, a factor of 1.6, and the stage ratio from $k = 2$ to 5 to 9, a factor of 4.5, against terminal ratios of 25:1, 79:1 and 251:1, a factor of ten. Neither is proportional to mass ratio, and the low-ratio case has no use for a large stage and pays for one in width. Array width still favours the high-ratio case, 2.82 m against 2.42 m and 1.92 m, because a modest terminal ratio spread over a full stroke demands a slow initial rise and a slow rise demands wide spans, but the gap is a factor of 1.5 rather than five once width is priced into the selection.

The rigid column no longer separates the cases: 2.05x, 2.07x and 2.10x peak-to-mean. All three sit within a factor of 1.03 of one another and within about a factor of two of uniform force even without compliance, with the *lowest* mass ratio now the best behaved. Earlier versions of this work reported 20x, 6x and 3x and concluded that low mass ratio was the difficult regime and that only the highest mass ratio would be viable on an inextensible member. Neither conclusion survives joint selection of N and k: both were artefacts of hand-chosen configurations. What compliance now supplies is a final correction from about 2x to within 9–19% of uniform, rather than the rescue of an otherwise unusable machine.

The exit velocities above are model outputs and are optimistic. The model of §5.1 omits the rotational inertia of sheaves, and for the output member that omission is first-order rather than small. Because a sheave traversed at payload speed stores $\frac{1}{2}m_s v^2$ (the translational kinetic energy of its own mass at payload velocity), fast-path inertia acts exactly as added payload mass, and

$$v = v_{\text{model}}\sqrt{\frac{m}{m + n_{\text{eff}}\, m_s}}, \qquad n_{\text{eff}} = \sum_i \left(\frac{v_i}{v}\right)^2 \tag{11}$$

for sheaves of mass m_s each, where v_i is the rim speed of sheave i. A sheave's rim speed is the speed of the member *relative to its own axle*, so only fixed-axle sheaves on the output member turn at full payload speed; the movable blocks of the fixed-ratio stage translate with the member and turn at $1/k$ of payload speed, contributing $1/k^2$ each. In the arrangement of Figure 7 two sheaves run at full speed, the orientation sheave and the first sheave taking the output member, giving $n_{\text{eff}} \approx 2$ to 3 once the slower sheaves are counted. Against a 1 kg payload that is a 20 to 30% mass increase and a **9 to 12% reduction in exit velocity** by (11), which is the italicised row above. The upper end applies at $k = 2$, where the slower sheaves contribute a quarter each rather than a fiftieth.

Two things about that correction should be stated plainly. It does **not** wash out at larger payload: a sheave must react the rope tension, so its mass scales with the force it carries and hence with payload mass, leaving $n_{\text{eff}} m_s/m$ and the correction factor roughly invariant with scale. And it is **larger than any difference between the candidate geometries of §5.3**, all of which reach their ideal to within 1%. The peak-to-mean force figures are ratios and survive a uniform added mass largely unchanged; only the velocity column needs reading as an upper bound.

Appendix E.1 removes the orientation sheave from the fast path, worth about 4% of exit velocity, and §7.2 treats the term as the largest single omission from the models used here.

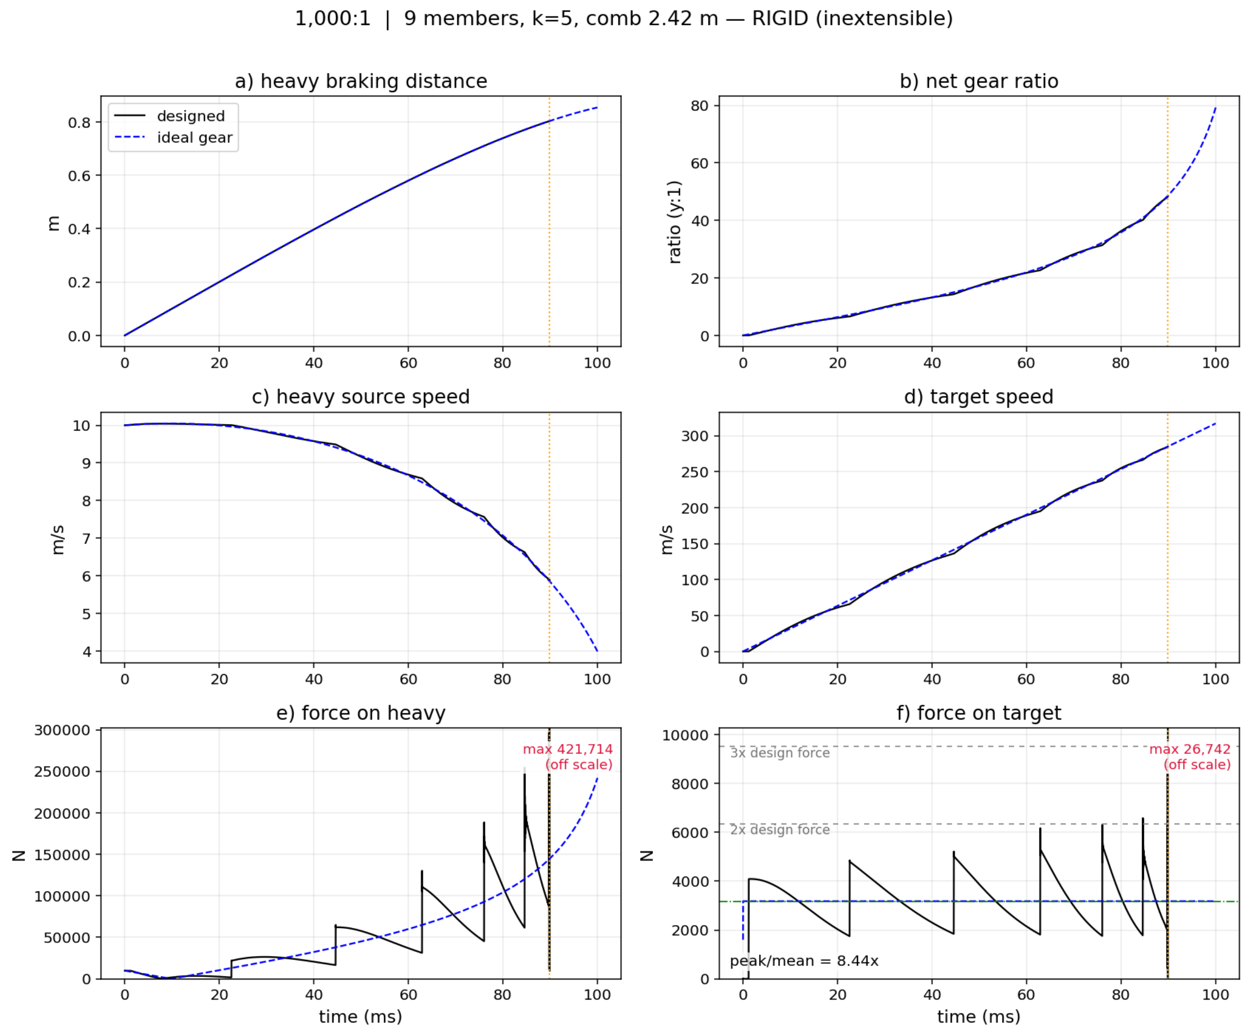

Figure 12. Rigid-body dynamics for the 1,000:1 case: nine members in a 2.42 m array. Dashed curves are the ideal profile (a uniform 3,170 N on the payload) computed by the same integration loop, differing only in that the ratio is solved for at each step rather than read from a geometry. Panel (f) shows one bounded sawtooth in payload force per engagement.

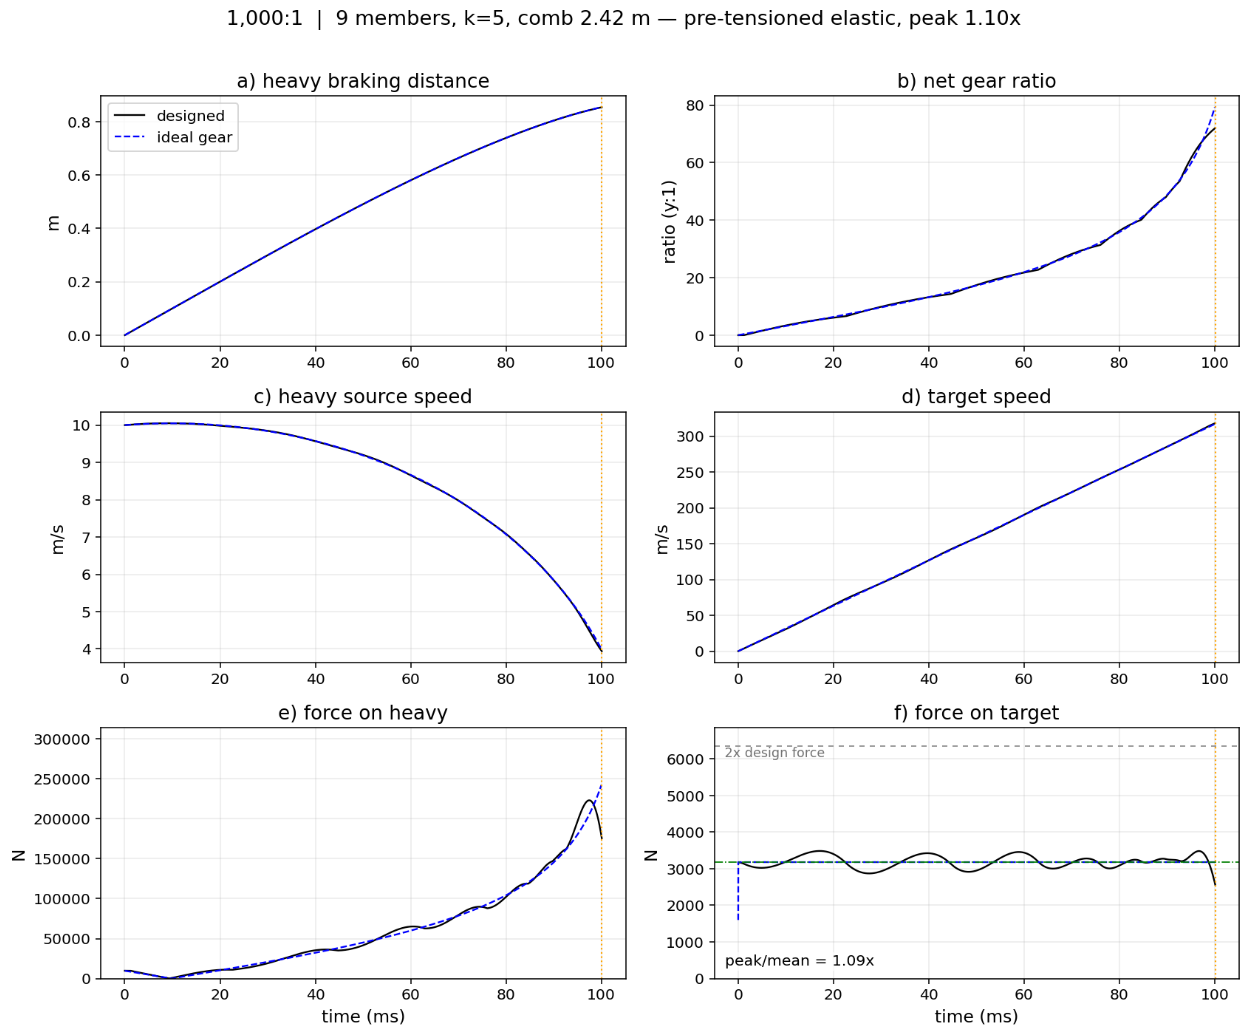

Figure 13. The same array with a compliant, pre-tensioned output member. The sawtooth is replaced by a smooth ripple between 2.8 and 3.5 kN about the 3,170 N design force, a peak of 1.10x. Panel (f) is drawn from zero so the ripple can be read against the absolute scale; the member is pre-tensioned to the working load, so the payload sees the design force from the first instant rather than ramping up from zero.

5.4 The ideal profile, and how it is computed

The profile an array is fitted to is a design choice, not a physical constant. We adopt the convention of imposing a uniform payload force F and solving, at each step, for the ratio needed to sustain it while the source mass responds to the reaction:

$$G_{new} = 2\left(\frac{\frac{1}{2}\Delta t\,\left(\frac{F}{m}\Delta t + 2v_t\right)}{\Delta d}\right) - G \tag{12}$$

which is the trapezoidal relation $\Delta y/\Delta d$ inverted for the new ratio, with v_t the payload velocity and F/m its acceleration under the uniform force. This is evaluated inside the same integration loop used for a real array (identical contact handling, identical reaction force, identical one-step lag), differing only in that the ratio is solved for rather than read from a geometry. The ideal and any simulated array are therefore produced by the same code and are directly comparable, which removes a class of error that is otherwise easy to make.

5.5 Why an unreachable target degrades the machine

§5.2 fixes the target from a physical specification. It is worth recording what happens when it is not, because the failure is counter-intuitive: asking the ideal to extract more energy than the array can deliver does not merely fall short of the request; it produces a materially **worse** machine than asking for less.

The figures in this subsection predate the target rule of §5.2 and are retained as the evidence that motivated it. At 10,000:1 with the member count capped at fourteen, an ideal demanding 96% extraction requires a net ratio of 534:1 against an array ceiling of $2Nk = 196$. The fitter responds by crowding members into narrow spans in an attempt to climb faster, engagement becomes abrupt, traction is lost at 45 ms, and the payload leaves at 394 m/s. An ideal demanding a reachable 156:1 gives 895 m/s from the same member count and array width, a factor of 2.3 in exit velocity, obtained purely by asking for less (Table 8).

Table 8: What an unreachable target costs, at 10,000:1 with the member count capped at fourteen. Asking for less returns a factor of 2.3 in exit velocity.

ideal demands	required ratio	exit velocity
96% extraction	534:1	394 m/s
intermediate	253:1	588 m/s
matched to the array	156:1	895 m/s

This effect is the origin of most of the apparatus that earlier versions of this work carried. When the target is out of reach the least-squares landscape becomes multi-modal and the fit residual stops predicting anything useful: at 100:1 under an unreachable target, twenty random restarts spanning a residual range of only 11% produced peak forces from 1.24x to 4.72x. Several additions to the objective were introduced to manage that scatter. **Under the target of §5.2 none of them is needed.** The landscape is effectively unimodal and restarts agree: of the 450 candidates simulated per case, only about 320 are distinct, and within the cell finally selected the five kept restarts agree on compliant peak-to-mean force to within 2.4%, 0.7% and 0.1% at 100:1, 1,000:1 and 10,000:1 respectively. The plain squared residual gives designs that meet every criterion in all three cases. The objective used throughout §5.3 carries no additional terms.

The practical statement of the rule is that the return from tuning the objective measures how badly the target is specified, and is therefore a diagnostic rather than a design tool.

5.6 A criterion on peak payload force

For an acceleration-sensitive payload the relevant specification is not mean force but peak.

Convention. All peak-to-mean figures in this work are taken over the first 99.5% of the force trace. The release mechanism is specified to operate before the onset of the terminal traction-loss pulse (§6.3), so that pulse is not experienced by the payload and including it would misstate the load the payload actually sees. Full-stroke figures, which do include it, are larger by a factor of 1.9

at 10,000:1, 4.1 at 1,000:1 and 21.7 at 100:1, and are reported in the data release for completeness. §6.4 gives a second reason for the truncation: the terminal pulse is also the one part of the trace that is not time-step converged.

Taking a bound of 30% above the design force, i.e. $F_{peak} \leq 1.3F$, **no rigid configuration examined meets it at any member count.** Discrete engagement produces a step in dG/dd each time a member takes up load, and the resulting sawtooth reaches 2.05, 2.07 and 2.10 times the mean in the three canonical designs. Across all 1,350 candidates simulated, the lowest rigid peak-to-mean found anywhere is 1.31.

That bound is not arbitrary. Force uniformity is the recognised figure of merit for launch systems, because for a payload able to survive a peak force F_{max} the energy delivered over a stroke of length L is $\bar{F}L$, so the stroke length required scales directly with the peak-to-mean ratio. The US Navy's replacement of steam catapults by the electromagnetic system of [6] was motivated substantially by this one property, with a reported gain of up to 31% in airframe service life. Reported peak-to-mean tow force ratios place the present designs as in Table 9:

Table 9: Reported peak-to-mean tow force for fielded launch systems, against the designs of this work.

system	peak / mean force
pneumatic and elastic catapults	3–10x
gun launch, progressive propellant	~1.7–2.5x
steam catapult	1.25x nominal, excursions to 2.0x [6]
RopeComb, compliant (this work)	**1.09–1.19x**
electromagnetic (EMALS)	1.05x [6]

The 1.3x criterion is therefore set at approximately the nominal performance of the best fielded non-electromagnetic launcher. The designs of §5.3 meet it while requiring no electrical energy storage, no power conditioning and no steam plant, the infrastructure that the electromagnetic system trades for its superior 1.05x. We note the asymmetry in evidence: the steam and electromagnetic figures are measured on operational hardware, whereas ours are simulated, and our compliant result is sensitive to output-member stiffness (§5.7).

With a compliant, pre-tensioned output member all three cases meet the bound (Table 10):

Table 10: Peak-to-mean payload force with and without a compliant, pre-tensioned output member.

	peak-to-mean, rigid	peak-to-mean, compliant
100:1	2.05x	**1.12x**
1,000:1	2.07x	**1.09x**
10,000:1	2.10x	**1.19x**

Member count is the effective lever, and the grid of §5.3 bounds it directly. The smallest member count that meets the 1.3x bound while still reaching at least 99% of the ideal exit velocity is **six at both 100:1 and 1,000:1, and twelve at 10,000:1**. Below those counts the bound is unreachable at any fixed-stage ratio in the grid: the best five-member design gives 1.43x at 100:1 and 1.61x at 1,000:1. The designs of §5.3 use eight, nine and thirteen members respectively, one to three above the floor, which buys margin in peak force at a modest cost in array width.

The force criterion is a requirement on compliance, not on geometry. This is the strongest argument in this work for treating pre-tensioned compliance as a necessary feature of the apparatus rather than a refinement of it, and the rigid column above is the measure of how necessary: once N and k are selected jointly, all three machines are within a factor of about two of uniform even without compliance, and compliance then carries them to within 9–19% of it.

5.7 Fixed-stage ratio

The fixed-stage ratio k trades member count against array width. By (5), a larger k needs fewer members, but each must then span a wider gap to cover the same rise. Array width accordingly passes through a minimum (Table 11):

Table 11: Fixed-stage ratio against member floor, smallest qualifying member count, narrowest qualifying array and best peak-to-mean force, for the 1,000:1 case over the grid of §5.3.

k	1	2	3	**5**	7	9
member floor $G/2k$	39.6	19.8	13.2	**7.9**	5.7	4.4
smallest N meeting the criterion	—	—	11	**7**	7	6
narrowest array meeting it	—	—	1.56 m	**1.44 m**	2.63 m	3.12 m
best compliant peak-to-mean	—	—	1.066x	**1.009x**	1.007x	1.011x

The table is computed for the 1,000:1 case over the grid of §5.3, counting only candidates that meet the criterion of §5.6 and reach at least 99% of the ideal exit velocity. Below $k = 3$ the member floor cannot be satisfied within the sixteen-member limit and no candidate qualifies at all. Above it, array width passes through a minimum at $k = 5$, which is the value the search selects.

The minimum moves with mass ratio. The canonical designs of §5.3 use $k = 2$, 5 and 9 at 100:1, 1,000:1 and 10,000:1. Earlier drafts of this work used $k = 7$ throughout on the strength of a single case; that generalisation was wrong, and the correct statement is that k must be chosen jointly with N rather than fixed in advance.

Exit velocity is essentially flat across k in every case, being set by the target and the instability rather than by the transmission.

5.8 A lower bound on span width

A span must clear the sheave diameter, the member and working clearance; 10 cm is a reasonable floor. Imposing it **improves** the designs rather than costing performance. Unconstrained fitting produces spans as narrow as 2 cm, with slope discontinuities $2/R$ of order 200, and these are precisely what makes engagement violent enough to lose traction early. At 10,000:1, under the earlier target and a hand-chosen configuration, the floor raised exit velocity from 860 to 927 m/s while reducing array width from 2.04 m to 1.45 m. The floor is active in all three current designs: every one ends in a run of 10 cm spans.

5.9 On the reported efficiency

The quoted extraction fraction is

$$\eta = \frac{\frac{1}{2}Mv_0^2 + Mg\,d_{\max} - \frac{1}{2}Mv_f^2}{\frac{1}{2}Mv_0^2 + Mg\,d_{\max}} \tag{13}$$

that is, the fraction of the source's *available* energy (initial kinetic plus gravitational work over the stroke) removed from it during the stroke. For the designs of §5.3 this is 86.3%, identical in all three cases because the deceleration target and the braking distance are. Its denominator, $\frac{1}{2}Mv_0^2 + Mg\,d_{\max} = Mg(h_0 + d_{\max})$, is exactly the input energy (9) of §5.2, so (13) and (10) are the same expression, so in this lossless model the extraction fraction and the transfer efficiency of Table 5 are the same quantity, 86.3%, and for the same reason neither depends on mass ratio. They would separate once losses are included, extraction being what leaves the source and transfer being what arrives at the payload. The figure above is exact for the ideal profile. The simulated designs of §5.3, on a compliant pre-tensioned member, deliver 86.0%, 86.7% and 86.4% of the same input at 100:1, 1,000:1 and 10,000:1. The spread is not a difference between the designs. Each run's energy balance, payload energy plus the kinetic energy the source carries away, closes on the input to 99.3%, 100.0% and 100.4%, and the departures from 86.3% track those closure errors exactly; the 1,000:1 case, whose balance closes to within 0.04%, returns 86.7% against a ceiling of 86.7% for the 3.94 m/s it actually leaves the source at. Values slightly above 86.3% therefore reflect a design braking the source marginally below 4.00 m/s, or integration drift, not a transmission exceeding its lossless bound.

Real efficiency will be materially lower once bearing friction, rotational inertia, member mass and drag are included, none of which are modelled here. Section 7.2 argues that rotational inertia alone may be the dominant omission.

5.10 Timing and efficiency are set by braking distance alone

The following sweeps braking distance at a fixed fractional reduction in source velocity, which is a different parameterisation from the deceleration target of §5.2 and gives correspondingly different extraction fractions. For a given entry velocity and a given fractional reduction in source velocity, stroke duration and extraction fraction are independent of mass ratio. Across all three mass ratios, a 0.50 m braking distance gives 56 ms and 68.8%, and a 3.00 m distance gives 319 ms and 78.5%. Mass ratio determines exit velocity, payload force and required ratio; it does not affect timing or efficiency. Braking distance and mass ratio may therefore be selected independently.

Increasing braking distance improves several quantities at once. At 1,000:1, going from 0.86 m to 1.50 m raises exit velocity from 287 to 309 m/s, raises extraction from 70.7% to 73.6%, and reduces payload force by 37%. The cost is payload travel, which rises from 13.8 to 25.5 m.

5.11 Payload travel is the binding practical constraint

Payload travel scales as $d_{\max}\sqrt{M/m}$ and payload force as $\sqrt{M/m}/d_{\max}$. A limit on track length therefore imposes a maximum braking distance and hence a minimum force, and the penalty grows with mass ratio: relative to a 3 m braking distance, the force penalty at the track limits above is 1.45× at 100:1, 2.02× at 1,000:1 and 2.43× at 10,000:1. Extraction is barely affected. A track limit exacts its cost in force, not efficiency.

---

6. The traction-loss instability

6.1 Mechanism

Every configuration examined terminates in the same way. Through the stroke, each engagement produces a bounded sawtooth in payload force: the force steps up as a member engages, then decays as the array ratio momentarily lags. Toward the end of the stroke these excursions grow, and eventually the payload velocity exceeds the rope-tip velocity Gv_h. The member goes slack, the source mass, no longer resisted, re-accelerates, the rope tip catches up, and contact is regained as an impulse. Each cycle delivers a larger impulse than the last and the sequence diverges within a few milliseconds.

For the 1,000:1 case, first slack occurs at 89.81 ms of an 89.89 ms stroke, after which successive re-contacts deliver roughly three and then eight times the 3,175 N mean. Those multiples are indicative only: they move by about 20% between time steps of 10^{-5} and 10^{-6} s (§6.4). The whole divergence occupies the final 0.1% of the stroke, which is why the peak-to-mean figures of §5.6 are taken over the first 99.5%.

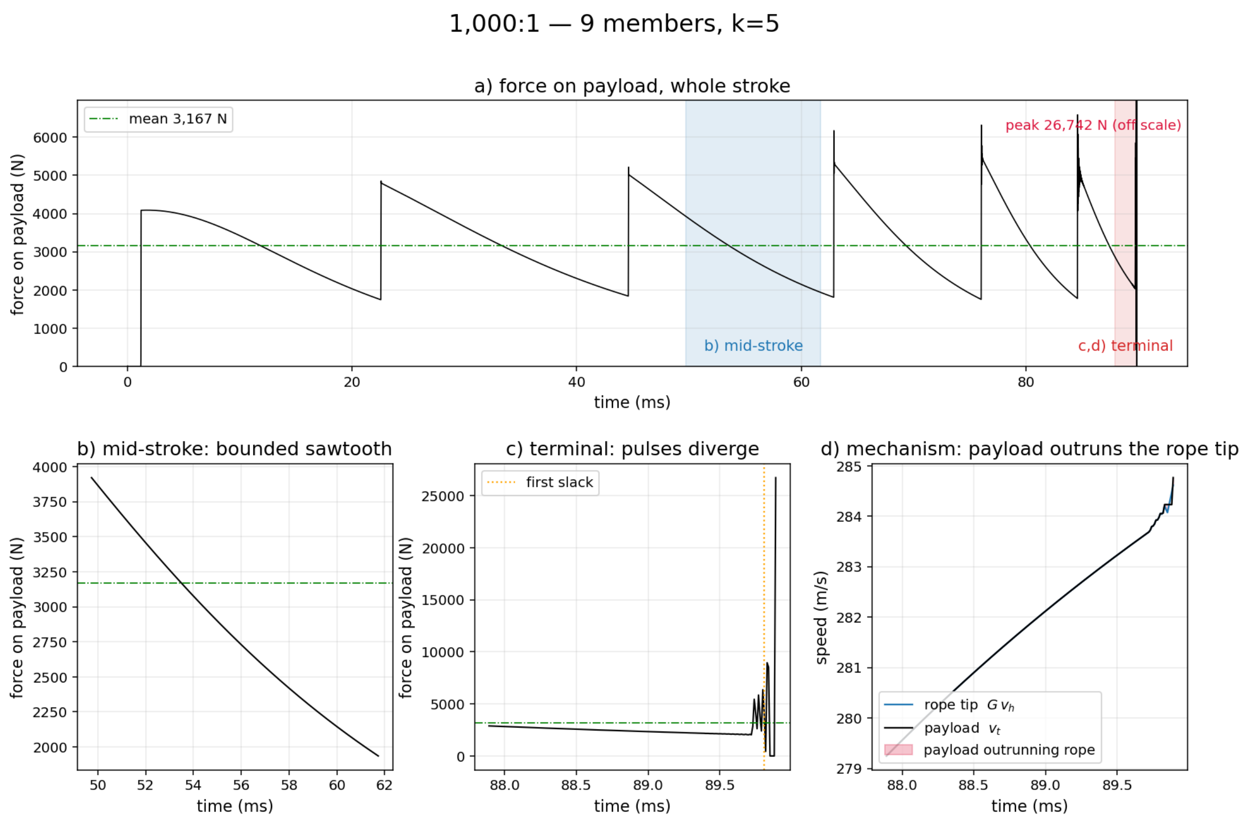

Figure 14. The traction-loss instability, 1,000:1 case, nine members. (a) Payload force over the whole stroke. (b) A mid-stroke engagement: bounded. (c) The terminal window: successive pulses diverge. (d) The mechanism: payload velocity rising above rope-tip velocity Gv_h, shaded where the member is slack.

6.2 Severity is not governed by fitting fidelity

Earlier versions of this work reported that instability severity was ordered inversely to mass ratio, mildest at 10,000:1, worst at 100:1, by a factor of seven. **That ordering does not survive joint**

selection of member count and fixed-stage ratio. With N and k chosen by the grid search of §5.3 rather than fixed by hand, rigid peak-to-mean force is as in Table 12:

Table 12: Fit residual against instability severity. Residual varies by a factor of 2.6 across the three cases while rigid peak-to-mean force varies by 2.5%.

	fit residual	peak / mean force, rigid	peak / mean force, compliant
100:1	0.57%	2.05x	1.12x
1,000:1	0.78%	2.07x	1.09x
10,000:1	1.48%	2.10x	1.19x

The spread is now a factor of 1.02 rather than seven, and the sign has reversed: the *lowest* mass ratio is now the best behaved rigidly, not the worst. The earlier ordering was an artefact of the hand-chosen configurations, not a property of the mechanism.

This also disposes of the residual hypothesis more cleanly than before. Residual varies by a factor of 2.6 across the three cases, from 0.57% to 1.48%, while rigid severity varies by 2.5%: a five-fold change in fit quality moves the rigid peak-to-mean force by almost nothing. Nor does residual order the compliant column, where the best-fitting case (100:1, 1.12x) is not the mildest; 1,000:1 is, at 1.09x, on a residual half again as large. A controlled test agrees: holding the mass ratio at 1,000:1 and improving the residual sevenfold changed the time to divergence by under 4% at a common time step, and slightly for the worse.

What remains true is the mechanism of §4.5: severity follows the discreteness of engagement (the arc length between successive members relative to the stroke) rather than how closely the ratio curve is reproduced. What is no longer true is that mass ratio orders it. We report the revised finding and note that the earlier claim was a selection artefact.

6.3 Consequences

The stroke has a useful length terminating at the onset of divergence, and a release mechanism operating at or before that point is required. Releasing early costs exit velocity, but less than the geometry suggests: because the payload force is uniform by design, $v = \sqrt{2Fy/m}$, so exit velocity follows the square root of payload travel and a fractional shortfall costs only half as much in velocity.

Releasing 5% early on the payload's own run costs about 2% of exit velocity: 2.5%, 2.4% and 1.7% at 100:1, 1,000:1 and 10,000:1, against 2.53% from the square-root law. Ten percent early costs about 5%. Five percent of the payload's run is 0.25 m, 0.79 m and 2.52 m in the three cases, and in every case that is 11 mm of the source's 0.854 m braking distance, or 2.5 ms of timing.

The tolerance is therefore centimetre-scale on the source stroke, decimetre-scale or more on the launch track, and of order a millisecond, which is undemanding for a mechanical or pyrotechnic release at the tensions involved. Conversely, exit velocity may be trimmed in service by adjusting the release position at the same rate.

6.4 Numerical status

All results in this paper use a time step of 10^{-5} s for the rigid model and 2×10^{-5} s for the compliant one. Refining by a factor of twenty, to 10^{-6} s, changes them by the amounts in Table 13:

Table 13: Time-step convergence: change in each reported quantity on refining the step by a factor of twenty, to 10^{-6} s.

	rigid	compliant
exit velocity	within 1.8%	within 0.07%
peak-to-mean payload force	within 4.3%	within 0.23%
energy drift	below 0.9%	below 0.6%

Peak-to-mean is taken over the first 99.5% of the trace throughout, per §5.6. The compliant model, from which the reported peak forces come, is converged to a fraction of a percent; the rigid model, which must resolve a hard contact switch, to a few percent. A step of 10^{-4} s is not adequate and was not used: at 100:1 it overstates exit velocity by 26% with an energy error of 35%.

The terminal traction-loss pulse itself is not converged, but it is excluded from every reported figure, by the truncation convention and by the release specification of §6.3. Establishing its magnitude and the precise time of divergence would require sub-stepping to the contact transition, which is left to future work.

7. Physical limits

7.1 Critical speed of the tension member

A member of linear density μ under tension T traversing a curved guide requires centripetal force to follow the curve. The contact force per unit length is proportional to $(T - \mu v^2)$, vanishing at

$$c = \sqrt{T/\mu} = \sqrt{\sigma/\rho} \tag{14}$$

Above this speed the guide cannot redirect the member. Transmitted power $P = (T - \mu v^2)v$ is maximised at $v = c/\sqrt{3} \approx 0.577c$ and falls to zero at c.

This depends on working *stress*, not cross-section: increasing diameter raises T and μ together and leaves c unchanged. At breaking stress, $c \approx 1{,}900$ m/s for UHMWPE [7] and comparable values obtain for PBO and carbon fibre, since these materials cluster in specific strength. **The ceiling is therefore a property of high-performance fibre in general, not of any particular choice**, and lies near 2,000 m/s at breaking stress, considerably lower at a working safety factor.

The 10,000:1 case sits at $v/c = 0.53$ at breaking stress, and at 0.92 of critical speed once a safety factor of 3 is applied, well past the $0.577c$ power optimum. It is rope-limited, and no braking distance rescues it: lengthening the stroke raises exit velocity as fast as it lowers tension.

7.2 Rotational energy of high-speed sheaves

For a rotating element of mass m_s with mass concentrated near its rim, at a rim velocity equal to the member velocity v,

$$E_{\text{rot}} = \tfrac{1}{2} I\omega^2 = \tfrac{1}{2} m_s v^2 \tag{15}$$

equal to the translational kinetic energy of the same mass at member velocity. Rim speed is the speed of the member *relative to the sheave's own axle*, so **only fixed-axle sheaves on the output**

member turn at full payload speed: the movable blocks of the fixed-ratio stage translate with the member and turn at $1/k$ of it, storing $1/k^2$ as much. Two sheaves run at full speed in the arrangement of Figure 7, the orientation sheave and the first sheave taking the output member, and Appendix E.1 removes the former. Counting the slower sheaves through n_{eff} of (11) gives $n_{\text{eff}} \approx 2$ to 3. On that basis 100 g sheaves at 318 m/s store 10 to 15 kJ against a 50.6 kJ output (20 to 30%); at 10,000:1, the same fraction, 101 to 151 kJ against a 505 kJ output.

This term is absent from the simulations presented here, §5.3 and Appendix E applying it only as a post-hoc correction, and it is in the author's estimation the largest single omission. It also imposes a rotational-speed constraint: a 100 mm sheave at 318 m/s turns at 61,000 rev/min, and at 1,005 m/s at 192,000 rev/min. There is a second and harder limit: a rotating rim carries hoop stress $\sigma = \rho v^2$, so its burst speed is $\sqrt{\sigma/\rho}$, the *same expression* as the member's critical speed. Sheaves in the fast path therefore sit at the same fraction of their strength as the member does of its own. Steel gives 358 m/s and carbon composite about 1,100 m/s, so the 10,000:1 case is at the limit of available materials on this count as well as on member critical speed. Larger sheaves reduce rotational speed but increase inertia, directly opposing the design objective. Appendix E.1 gives an arrangement that removes one sheave from the fast path. Quantifying this is the most important extension to the present work.

7.3 Array width

Reproducing a slow initial rise requires wide spans, since a member's rate of contribution is governed by $1/R_i$. An embodiment requiring a modest net ratio over a long braking distance therefore needs a proportionally wider array than one requiring a high ratio over a short distance. In the three cases the ratio of array width to braking distance is 3.3, 2.8 and 2.3 respectively, the *lowest* mass ratio demanding the widest array relative to its stroke.

Subject to a given total width, fidelity improves with member count and then plateaus; additional members are neither harmful nor beneficial, and the plateau is set by the width available. Array width, not member count, is the governing practical constraint, since the array is carried by the translating carriage and determines the size and mass of the moving structure.

---

8. Discussion

The contribution here is twofold. The mechanism provides a rising ratio profile that is specifiable by geometry rather than inherited from a linkage, adjustable by repositioning rather than remanufacture, and realisable in components most of which move slowly. The design method converts what is otherwise an intractable high-dimensional search into a smooth fit against a target derived in the displacement domain.

The second may matter more than the first. Section 4.2 shows that the optimum lies outside the geometric families a designer would naturally parameterise, which explains why arrays of this kind are difficult to design by inspection and why the mechanism's capability is easy to underestimate.

Three qualifications bear emphasis. The models are lossless, and the omission of sheave rotational inertia in particular is likely to be first-order rather than a correction. The terminal transient is not time-step converged, though the truncation convention of §5.6 excludes it from every reported figure and the reported quantities are converged (§6.4). And no physical apparatus has been built or measured; every result reported here is theoretical or numerical.

Applications for which the acceleration levels are appropriate include launch as an alternative to gun or artillery delivery (where setback accelerations of order 10^4 g would otherwise require hardened components), impact and materials test facilities, target drones, and acceleration of material in vacuum or low-gravity environments where the absence of drag on the payload and the high-speed member removes a principal loss.

On the comparison with fielded systems. The peak-to-mean figures of §5.6 place these designs between the two systems that bracket the field. That is the substantive claim of this work: force uniformity comparable to a steam catapult, obtained from a passive mechanical transmission whose only stored energy is the kinetic and potential energy of a descending mass. The electromagnetic system remains better by roughly a factor of 1.2 in uniformity, and is programmable per launch, which no fixed geometry can be. What the present mechanism offers instead is the removal of the infrastructure (the energy store, the power conditioning, the cooling) that makes electromagnetic launch impractical below a certain scale and at sites without heavy electrical supply.

The comparison should be read with its asymmetry in mind. Steam and electromagnetic figures are measured on operational hardware over many thousands of launches; ours are simulated in a model that omits friction, rotational inertia, member mass and drag, and that depends on an assumed output-member compliance to which the result is sensitive. Closing that gap requires hardware, and no refinement of the model substitutes for it.

Open problems

1. **Sheave inertia.** Incorporate rotational inertia in the model itself rather than as the post-hoc correction of (11), and establish n_{eff} for a given reeving rather than bounding it at 2 to 3; determine the sheave mass and effective count at which the transmission ceases to be advantageous.
2. **Contact transition.** Sub-step the traction switch to obtain converged peak forces and divergence times.
3. **Release mechanism.** Design and analyse a release repeatable to about 10 mm of source travel, equivalently 2 ms, at the required tension. Per §6.3 that holds exit velocity to within roughly 2%.
4. **Instrumented apparatus.** Measure the ratio profile quasi-statically against prediction, which tests the central geometric claim independently of any dynamic behaviour.
5. **Distributed compliance.** Model the array-side members, which are an order of magnitude stiffer locally than the output member and are where discrete engagement occurs.

---

9. Conclusion

The RopeComb is a passive, configurable rope-and-pulley transmission that matches the impedance of a slow, high-force source to a fast, light payload. A tension member reeved alternately over a staggered array of engagement members and fixed supports is progressively deflected into successive spans, giving a mechanical advantage that rises continuously with carriage displacement. The array has a closed-form displacement ratio (4) and two independent geometric freedoms per member: span width, which sets how fast a member contributes, and engagement offset, which sets when. Those two freedoms are what permit an essentially arbitrary rising profile to be synthesised, and with it the pursuit of uniform payload force that no fixed-geometry linkage can undertake.

The profile required for uniform force satisfies a first-order equation in the displacement domain, so matching it is a smooth least-squares fit rather than a search. The fit reaches a residual nearly eight times lower than three hundred thousand random samples while evaluating the objective a hundred and fifty times less often, and it reaches configurations that low-dimensional parameterisations cannot express at all: a structured grid over four intuitive parameters performs worse than uniform random sampling, because the optimum lies outside that family. Given a target specified by the deceleration to be achieved rather than by a chosen force, arrays of eight, nine and thirteen members reproduce it to within 1.5% and reach the whole of the ideal exit velocity at mass ratios from 100:1 to 10,000:1, surrendering 86% of the energy available to the source in an idealised model. Those exit velocities are model outputs: the rotational inertia of sheaves in the fast path, omitted here, is first-order and costs a further 9 to 12% of exit velocity.

The usable stroke is terminated by a traction-loss instability whose severity follows the discreteness of engagement rather than fitting residual or mass ratio. Pairing the array with a compliant, pre-tensioned output member absorbs those transients and holds peak-to-mean payload force to 1.09–1.19×, a uniformity at or beyond the nominal figure for a steam catapult and obtained with entirely passive mechanics. That comparison should be read with its asymmetry in mind, since the present figures are simulated and the catapult's are measured on fielded hardware. The accessible design space is bounded by three limits: the tension member's critical speed at high exit velocity, fast-path sheave inertia, and array width at low mass ratio.

Appendix A — Rigid-body simulation

A.1 Conventions and assumptions

- Displacements and velocities of the source mass are positive **downward**, in the direction of its travel; gravity is therefore a positive acceleration on the source.
- `get_gear_fn(d)` returns the **net** ratio (the array ratio multiplied by any fixed stage) as a function of source displacement. It is a pure function of displacement, not of time or velocity.
- The payload is taken to travel against gravity, so the tension required is $m(g + a)$ and its free-flight deceleration is g. For a payload on a track inclined at ϕ, substitute $g \sin \phi$ in both places.
- The tension member is inextensible and massless; sheaves are frictionless and of zero rotational inertia; there is no aerodynamic drag.
- Integration is explicit and first-order in force, with trapezoidal displacement updates. The force acting on the source at the start of each step uses the ratio and reaction force carried over from the previous step. This one-step lag is intentional and is what makes the scheme explicit; it is noted because it is otherwise easily mistaken for an error.

A.2 Pseudocode

```
FUNCTION RopeCombRigidBodySimulation(
    heavy_weight,           // source mass M, kg
    target_weight,          // payload mass m, kg
    heavy_height,           // drop height setting entry speed, m
    max_time,               // stop if exceeded, s
    max_target_acc,         // stop if payload acceleration exceeds this, m/s^2
    get_gear_fn,            // net ratio as a function of source displacement
```

```
    time_unit = 1e-5        // fixed step, s
)
    g = 9.8

    // --- initial state at the start of the stroke ---------------------
    gear                 = 0.0      // no member engaged yet
    average_gear         = 0.0
    heavy_speed          = SQRT(2 * g * heavy_height)
    heavy_acc            = g
    target_speed         = 0.0
    target_acc           = 0.0
    target_reaction_force = 0.0

    heavy_dec_dist       = 0.0      // source displacement
    target_acc_dist      = 0.0      // payload displacement
    target_rope_end_dist  = 0.0      // displacement of the rope tip
    total_time           = 0.0

    // energy invariant: constant in a lossless model, used only as a check
    E0 = 0.5 * heavy_weight * heavy_speed^2

    WHILE (total_time < max_time AND target_acc < max_target_acc) DO

        total_time = total_time + time_unit

        // --- source mass ----------------------------------------------
        // Reaction is (ratio x payload tension), both carried over from the
        // previous step. Net force goes negative once reaction exceeds weight.
        net_force_on_heavy = (g * heavy_weight)
                           - (average_gear * target_reaction_force)
        heavy_acc          = net_force_on_heavy / heavy_weight

        new_heavy_speed      = heavy_speed + heavy_acc * time_unit
        heavy_dec_dist_delta = time_unit * (new_heavy_speed + heavy_speed) / 2
        heavy_dec_dist       = heavy_dec_dist + heavy_dec_dist_delta
        heavy_speed          = new_heavy_speed

        // --- ratio ----------------------------------------------------
        new_gear     = CALL get_gear_fn(heavy_dec_dist)
        average_gear = (gear + new_gear) / 2
        gear         = new_gear

        // rope tip advances by (source displacement x ratio)
        target_rope_end_dist = target_rope_end_dist
                             + heavy_dec_dist_delta * average_gear

        // --- payload: unilateral contact ------------------------------
        // The member may pull but not push, so the payload advances by
```

```
        // whichever is greater: its free flight, or the amount the rope
        // permits. When free flight wins, the member is slack.
        free_flight   = (target_speed - 0.5 * g * time_unit) * time_unit
        rope_limited  = target_rope_end_dist - target_acc_dist
        target_acc_dist_delta = MAX(free_flight, rope_limited)
        is_slack      = (free_flight > rope_limited)

        // invert the trapezoidal rule for the new velocity
        new_target_speed = 2 * (target_acc_dist_delta / time_unit) - target_speed
        new_target_acc   = (new_target_speed - target_speed) / time_unit

        target_speed    = new_target_speed
        target_acc_dist = target_acc_dist + target_acc_dist_delta

        // --- tension for the next step ------------------------------------
        // Averaged over the step; only positive accelerations correspond to
        // the member pulling.
        a_old = IF target_acc     > 0 THEN (g + target_acc)     ELSE 0
        a_new = IF new_target_acc > 0 THEN (g + new_target_acc) ELSE 0
        target_reaction_force = target_weight * (a_old + a_new) / 2

        target_acc = new_target_acc

        // --- diagnostics (no effect on the integration) -------------------
        E = 0.5 * heavy_weight * heavy_speed^2
            + 0.5 * target_weight * target_speed^2
            - heavy_weight * g * heavy_dec_dist
        RECORD total_time, heavy_dec_dist, target_acc_dist,
               heavy_speed, target_speed, gear,
               net_force_on_heavy, target_reaction_force,
               is_slack, (E - E0) / E0

    END WHILE

    RETURN recorded history
END FUNCTION
```

A.3 Known limitations of this scheme

1. **The contact switch is discontinuous.** `MAX(free_flight, rope_limited)` flips between two regimes without resolving the instant of transition. This is the principal source of the time-step sensitivity reported in §6.4. Sub-stepping to the crossing would fix it.
2. **Termination lags by one step**, since `target_acc` is tested at the head of the loop using the previous value.
3. **The energy residual `(E - E0)/E0` should remain at zero** in exact arithmetic. It is recorded rather than enforced, and provides a cheap check on any reimplementation: drift indicates an integration error, and non-monotonic drift under time-step refinement indicates the contact switch is being resolved inconsistently.

4. **`max_target_acc` conflates two stopping conditions**: a genuine physical release and a numerical blow-up. Distinguishing them requires the release logic to be modelled explicitly.

Appendix B — Compliant tension member

The output member has longitudinal stiffness $k = EA/L$; for the worked case, approximately 1.6×10^4 N/m, giving 0.24 m extension at 4 kN, or 1.4% strain. Since $A = F_{\mathrm{break}}/\sigma$, we may write $k = (E/\sigma)F_{\mathrm{break}}/L$, and E/σ varies only between about 34 and 66 across UHMWPE [7], steel wire rope, carbon fibre and PBO. Working strain is therefore of order 1% whatever the material; compliance cannot be engineered away.

Simulation with a compliant member and a slack start excites the payload mode on that compliance (approximately 20 Hz against a 96 ms stroke, hence about two oscillations), giving a peak-to-mean force ratio near 2.3 with the member unloaded for some 11% of the stroke. Pre-tensioning to the working tension before the stroke removes the step input, and the same apparatus then gives a peak-to-mean ratio near 1.03 with no unloaded interval and exit velocity within 0.1% of ideal. The result is insensitive to stiffness over the plausible range and to damping.

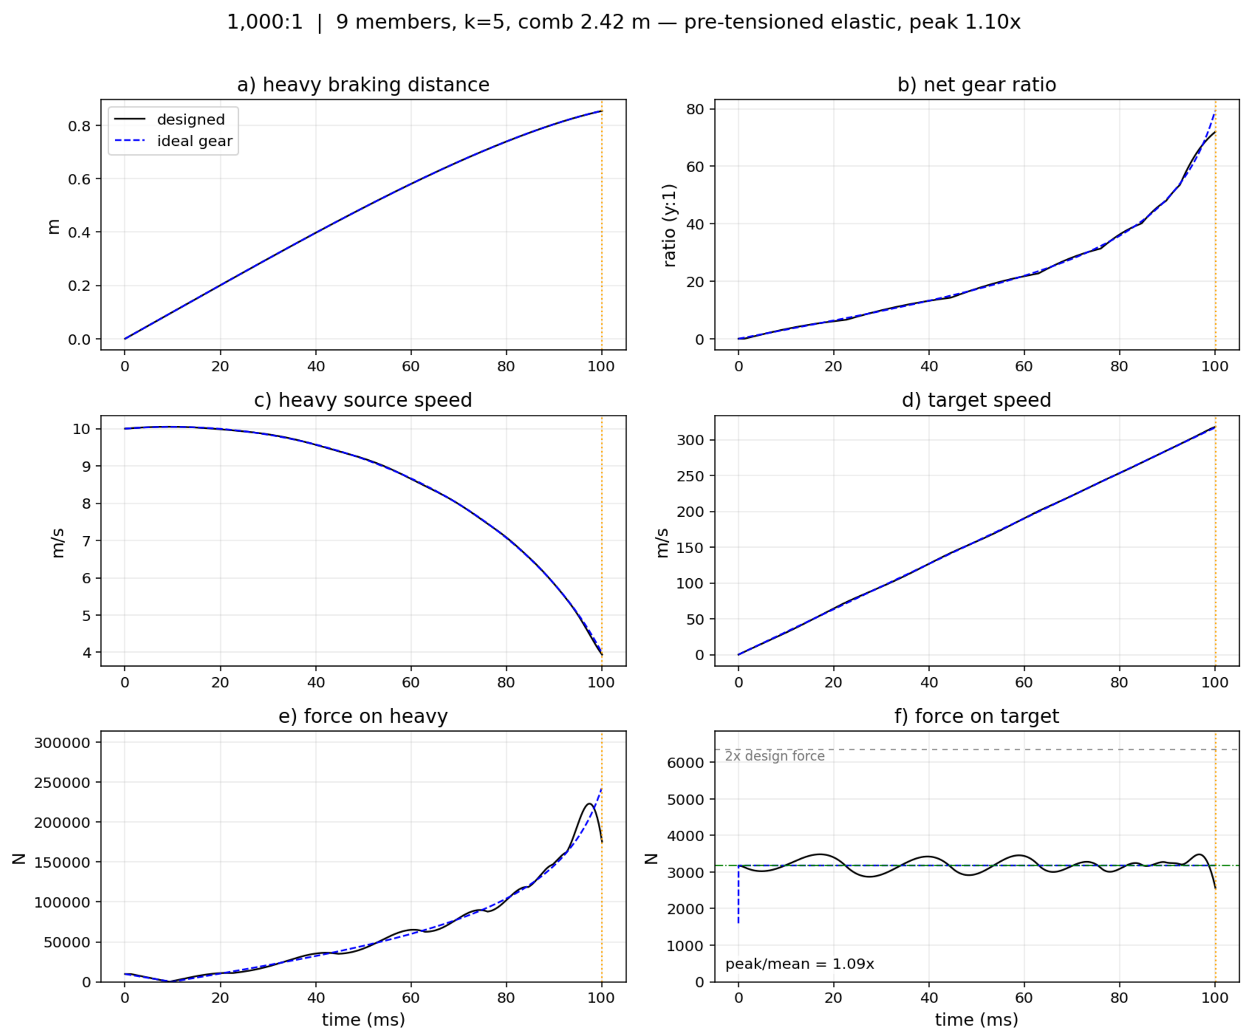

Figure B1. The same array with a compliant output member pre-tensioned to the working tension. Force ripples by ±2.4% about its mean and the member never unloads.

This result should not be read as a prediction that a real apparatus is free of force pulses. The model lumps all compliance into the output member, neglecting the array-side members, which are locally an order of magnitude stiffer and are where discrete engagement occurs; and it omits sheave rotational inertia, member mass and roller friction. Each of these is a plausible pulse mechanism, and their absence, rather than their suppression, is the more likely explanation for the clean force history obtained. The value of this appendix is in isolating the effect of pre-tension, not in predicting real force histories.

Appendix C — Verification of the ratio model

The closed form (4) was checked against rope length measured from an explicit polyline through support and engagement-member coordinates, finite-differenced with respect to carriage displacement. Across three configurations spanning 4 to 20 members, the worst relative disagreement was $1.5 \times 10^{-5}\%$ and cumulative take-up agreed to $7 \times 10^{-4}\%$.

Note that agreement between the *fitted* array profile and the target profile is not evidence of correctness: the fit produces that agreement by construction. The meaningful checks are the geometric verification above, and the agreement between simulated dynamics and the independently derived ideal velocity and force histories, which were not fitted to anything.

Appendix D — The 100:1 and 10,000:1 cases

Section 5 reports all three canonical cases numerically but illustrates only the 1,000:1 configuration. The remaining two are given here. They bracket the design space at either end, and the contrast between them is the clearest statement of where the difficulty lies.

D.1 Fitted geometry

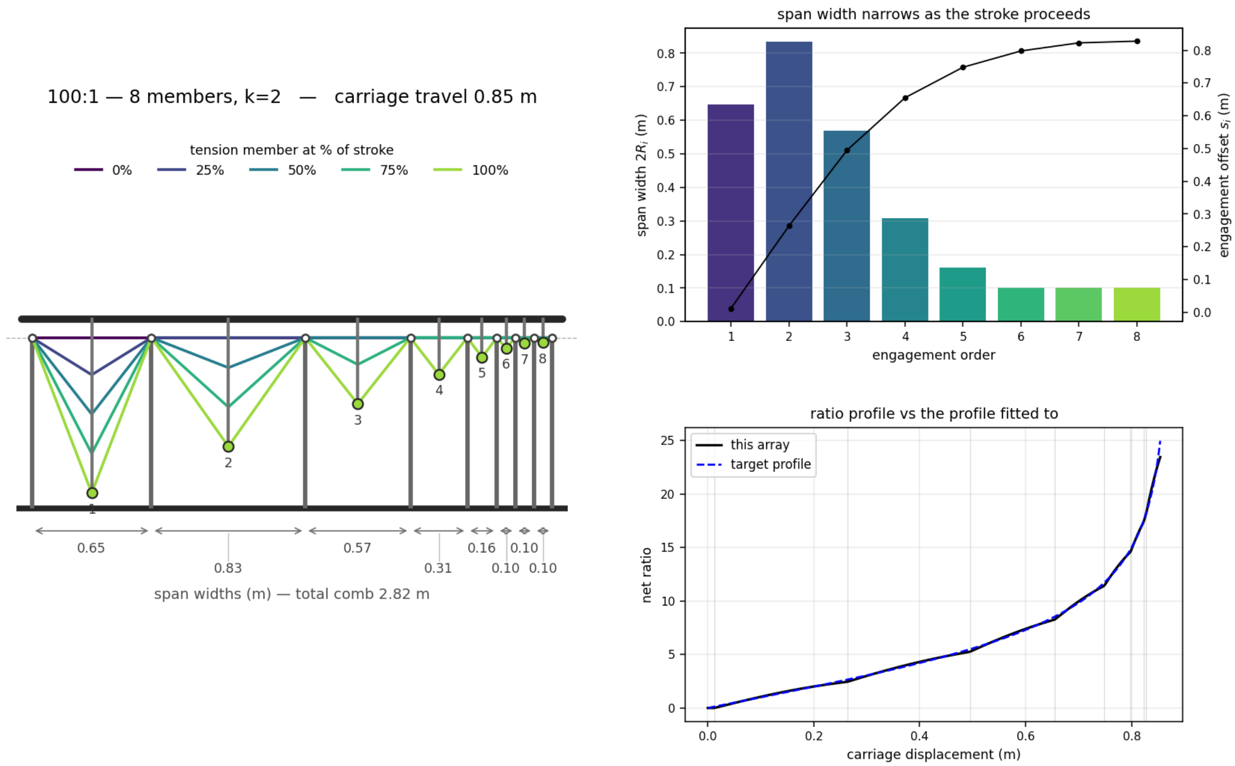

Figure D1. Fitted array for the 100:1 case: eight members in 2.82 m of array, spans of 65, 83, 57, 31, 16, 10, 10, 10 cm. A modest terminal ratio (25:1) spread over a full stroke demands a slow initial rise, and a slow rise demands wide spans, so this remains the widest of the three arrays, but at 2.82 m rather than the 10.21 m a width-blind selection returns.

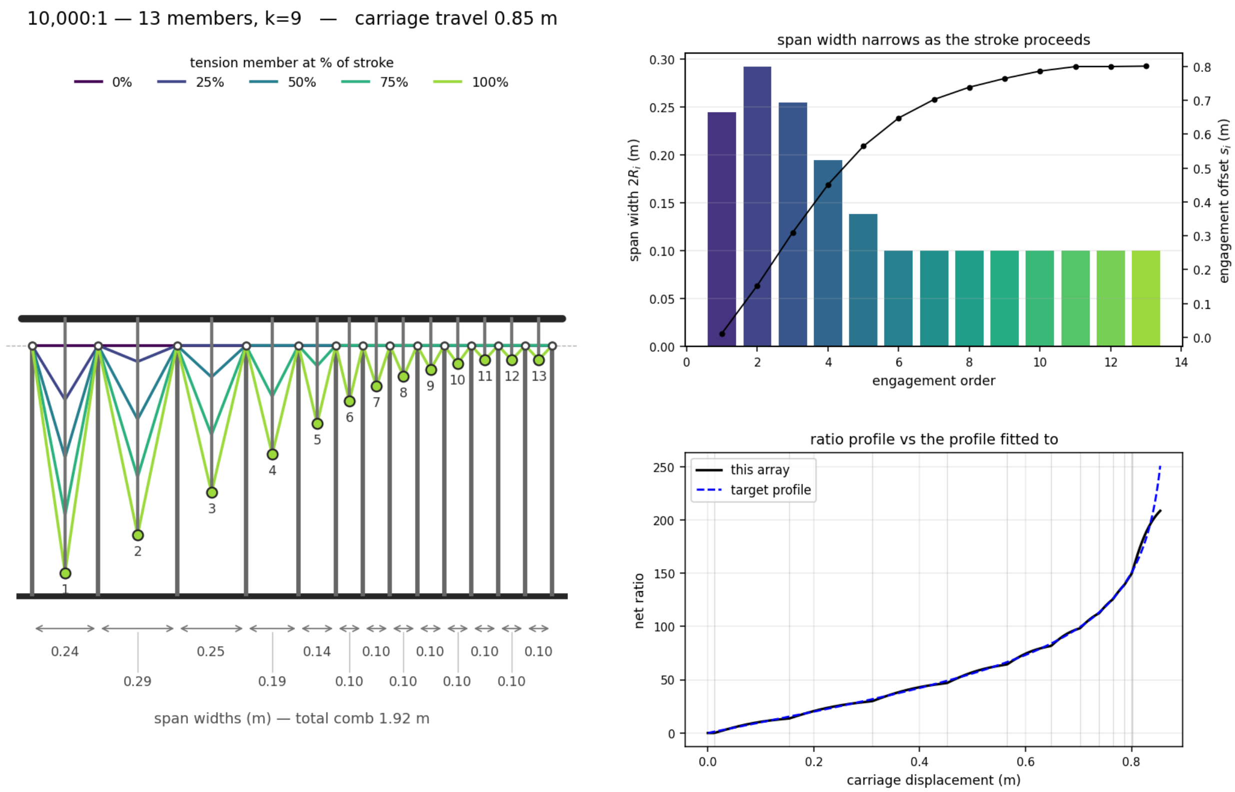

Figure D2. Fitted array for the 10,000:1 case: thirteen members in 1.92 m of array, spans of 24, 29, 25, 19, 14, 10, 10, 10, 10, 10, 10, 10, 10 cm. Although it requires a terminal ratio ten times that of the 100:1 case, this array is half as wide; the high-ratio case remains the geometrically cheap corner of the design space, though the margin narrows from a factor of five to about two once array width is priced into the selection (§4.7). Its fit residual, at 1.48%, is nonetheless the largest of the three.

D.2 Dynamics

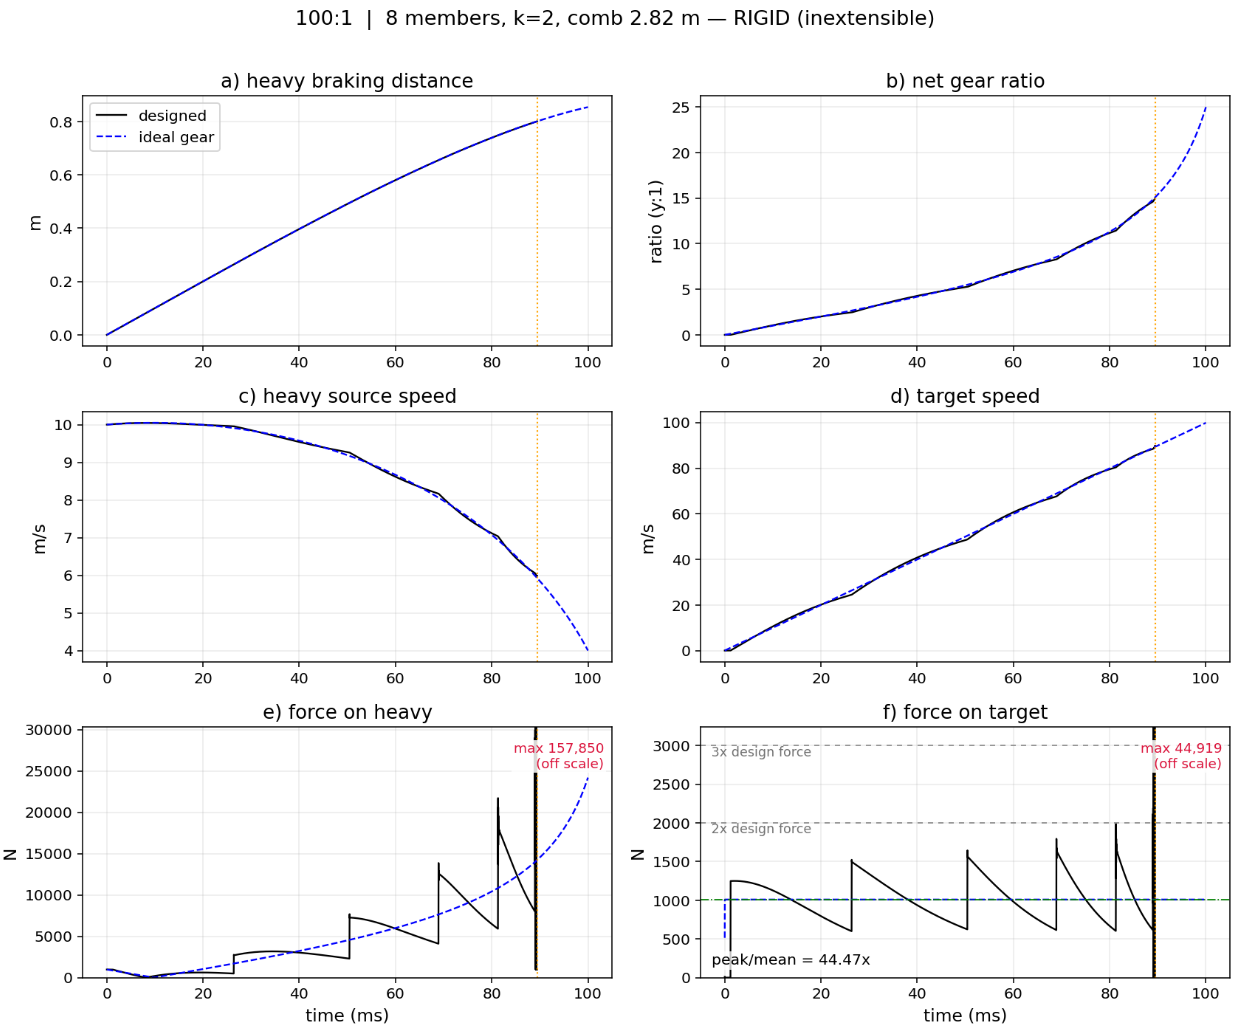

Figure D3. Rigid-body dynamics, 100:1. Exit velocity 89.8 m/s against an ideal 99.9 m/s. Panel (f) shows a sawtooth of 2.05 times the mean over the working stroke, rising to roughly 45 times the mean if the terminal traction-loss divergence is included. That multiple is not time-step converged (§6.4), and the release is specified to occur before the transient (§5.6, §6.3).

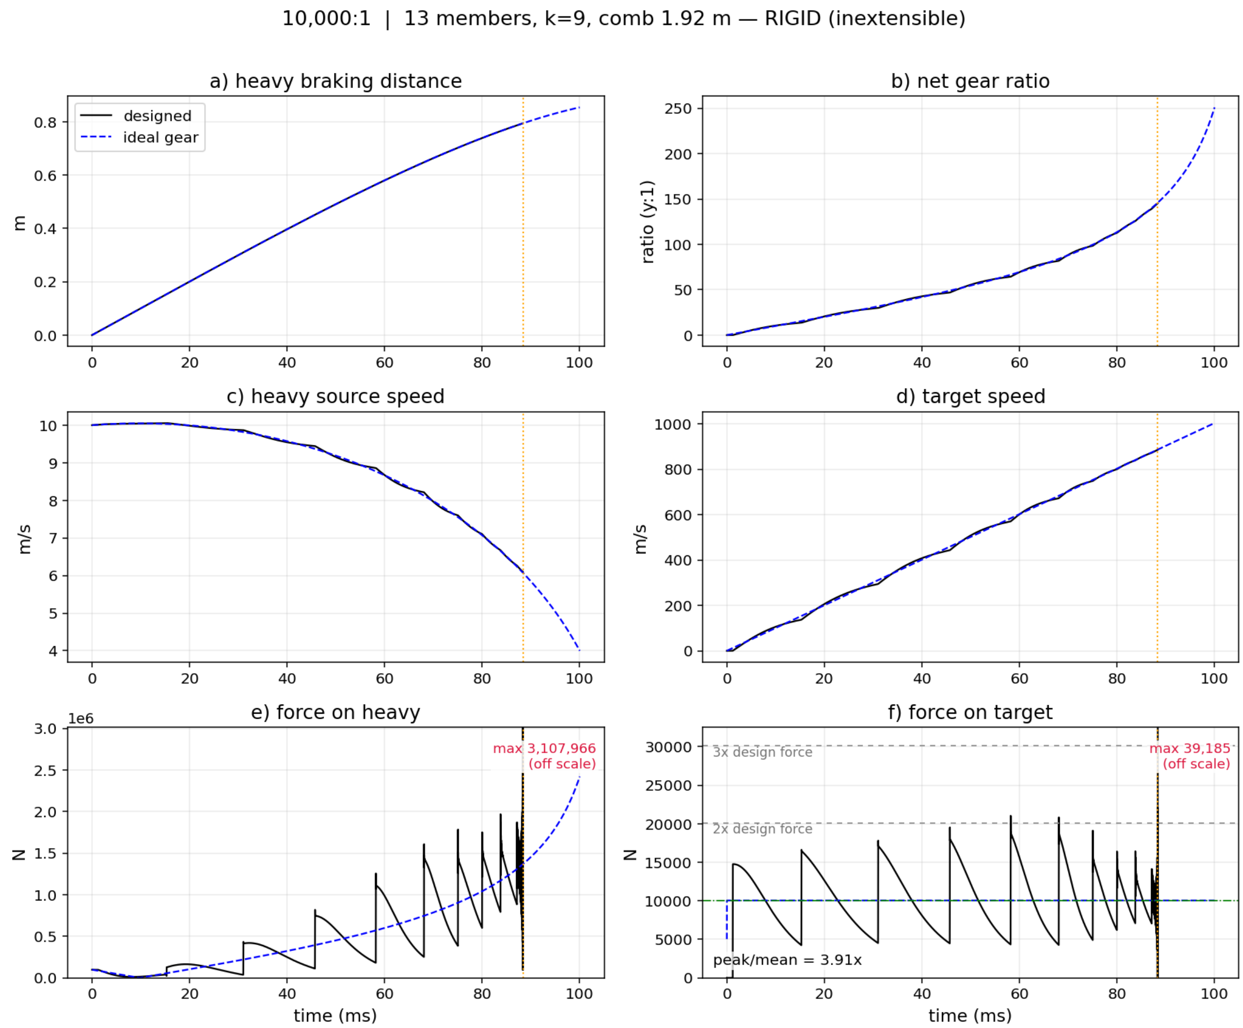

Figure D4. Rigid-body dynamics, 10,000:1. Exit velocity 885.9 m/s against an ideal 1003.5 m/s. Peak payload force over the working stroke is 2.10 times the mean, against 2.05 times at 100:1 (Figure D3), comparable rather than markedly better, which is the revised finding of §6.2.

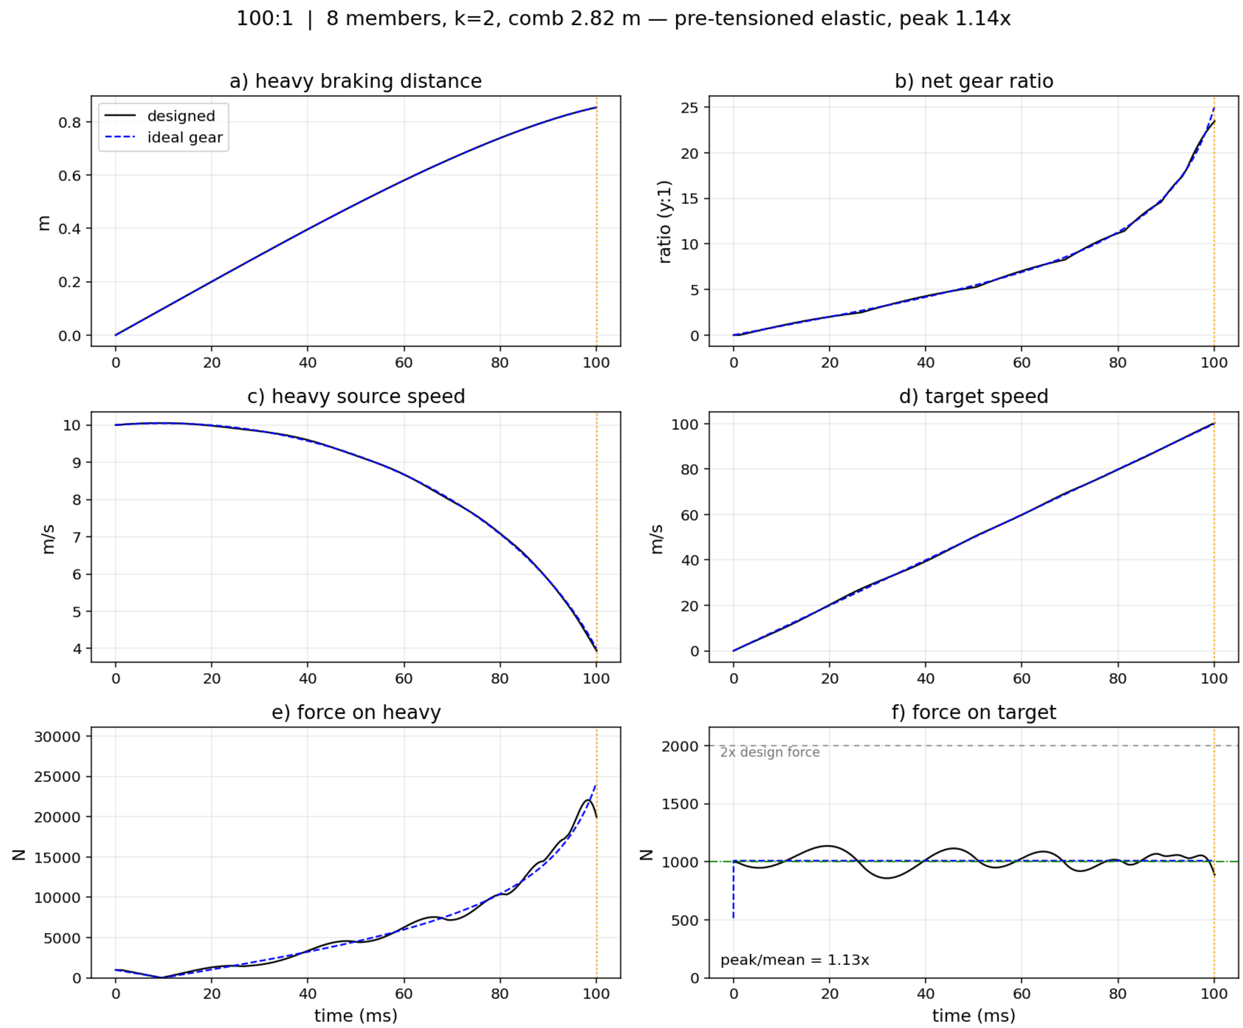

Figure D5. The 100:1 array with a compliant, pre-tensioned output member: 100.2 m/s at 1.14 times the design force. Comparison with Figure D3 shows compliance closing the remaining gap from 2.05 to 1.12 times the mean. The correction is a factor of 1.82 here, 1.90 at 1,000:1 and 1.76 at 10,000:1: compliance does the same amount of work in all three cases, which was not true before N and k were selected jointly.

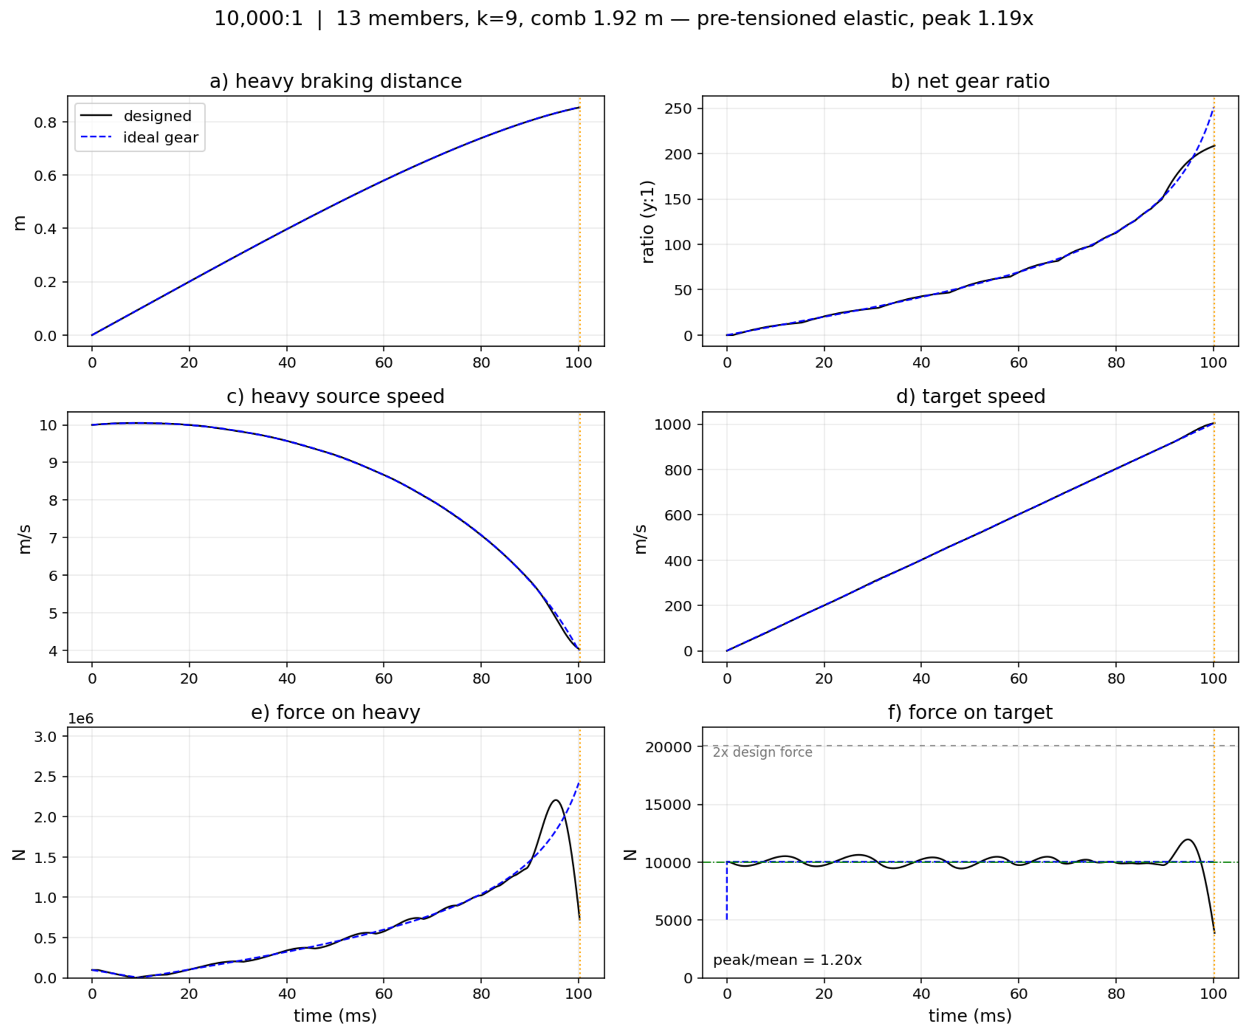

Figure D6. The 10,000:1 array, compliant: 1004.8 m/s at 1.19 times the design force, 100% of the ideal exit velocity. As in Figures 13 and D5, both force panels are drawn from zero so the traces can be read against the absolute scale and compared across cases.

D.3 Instability

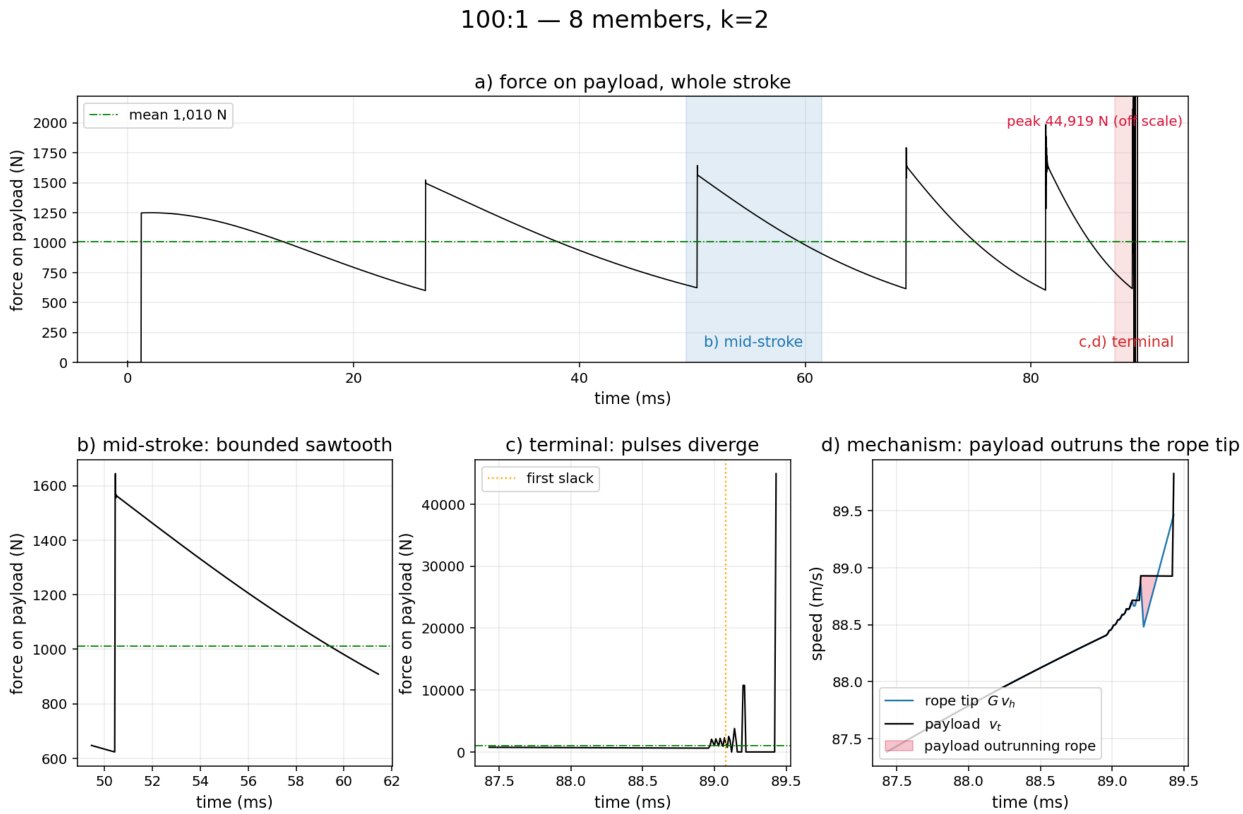

Figure D7. Traction-loss instability, 100:1. Formerly the most severe of the three by a wide margin; under joint selection of N, k and array width the three cases lie within 3% of one another.

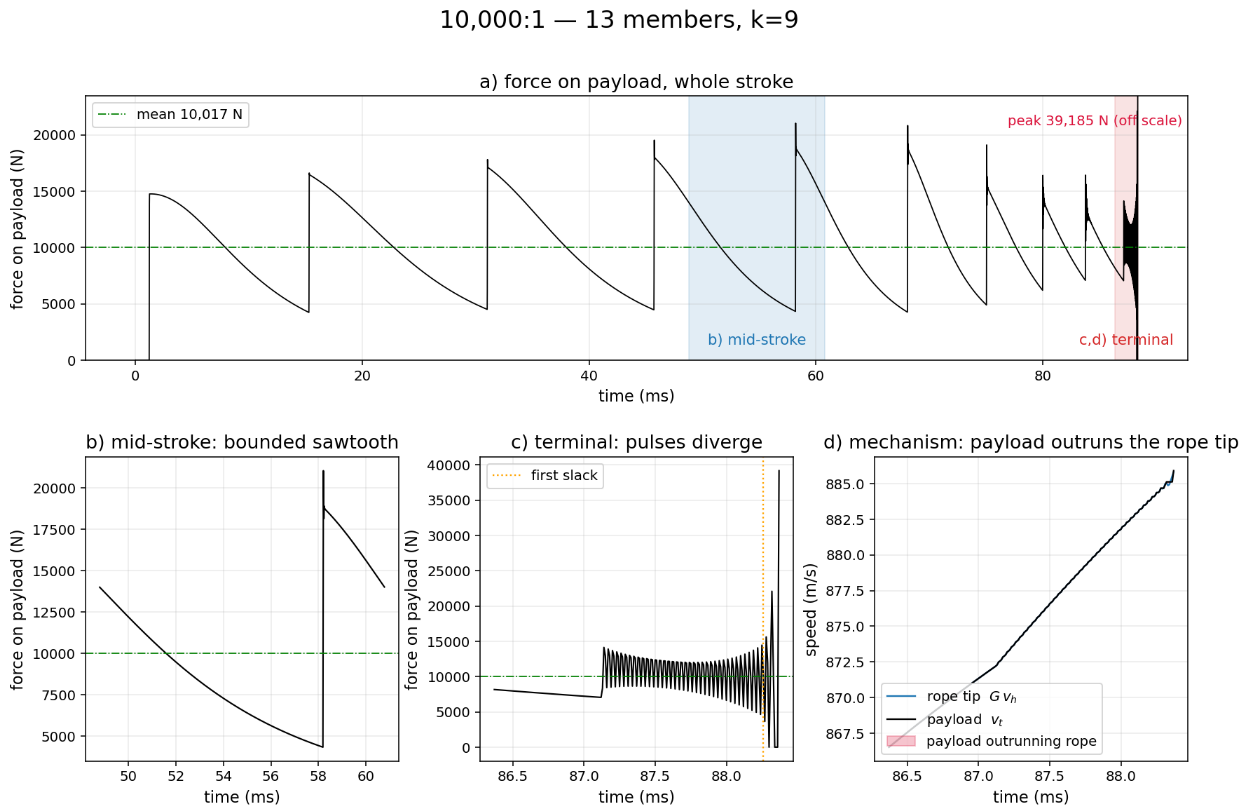

Figure D8. Traction-loss instability, 10,000:1. Comparable in severity to the 100:1 case of Figure D7 rather than an order of magnitude milder, which is the evidence for the revised §6.2: once N and k are selected jointly, mass ratio no longer orders severity.

D.4 Summary

Table 14 summarises the three canonical designs alongside the critical-speed margin of §7.1.

Table 14: Summary of the three canonical designs, with the critical-speed margin of §7.1.

	100:1	1,000:1	10,000:1
members built	8	9	13
fixed-stage ratio	2	5	9
array width	2.82 m	2.42 m	1.92 m
narrowest span	10 cm	10 cm	10 cm
exit velocity, fraction of ideal	100%	100%	100%
peak-to-mean, rigid	2.05x	2.07x	2.10x
peak-to-mean, compliant	1.12x	1.09x	1.19x
v/c at safety factor 3	0.09	0.29	0.92

The two binding constraints run in opposite directions. Geometry becomes **harder** as mass ratio rises (the fit residual grows from 0.57% to 1.48% and the rigid peak-to-mean force from 2.05x to 2.10x) while the tension member simultaneously approaches its critical speed. What improves with mass ratio is array width, not fidelity. At 10,000:1 the payload leaves at 0.53 of the member's critical speed, and at a safety factor of 3 the working point is 0.92 of it, past the 0.577 optimum of §7.1. Geometry has ceased to be the constraint there.

Appendix E — Implementation variants

The apparatus drawn in Figures 1, 4 and 7 was chosen for clarity of exposition rather than for efficiency. Two variants are given here because each addresses a limit identified in the body of the paper, and because neither changes the transmission's ratio profile: the geometry of §3.2 and the design procedure of §4.7 apply unaltered.

E.1 Minimal fast path

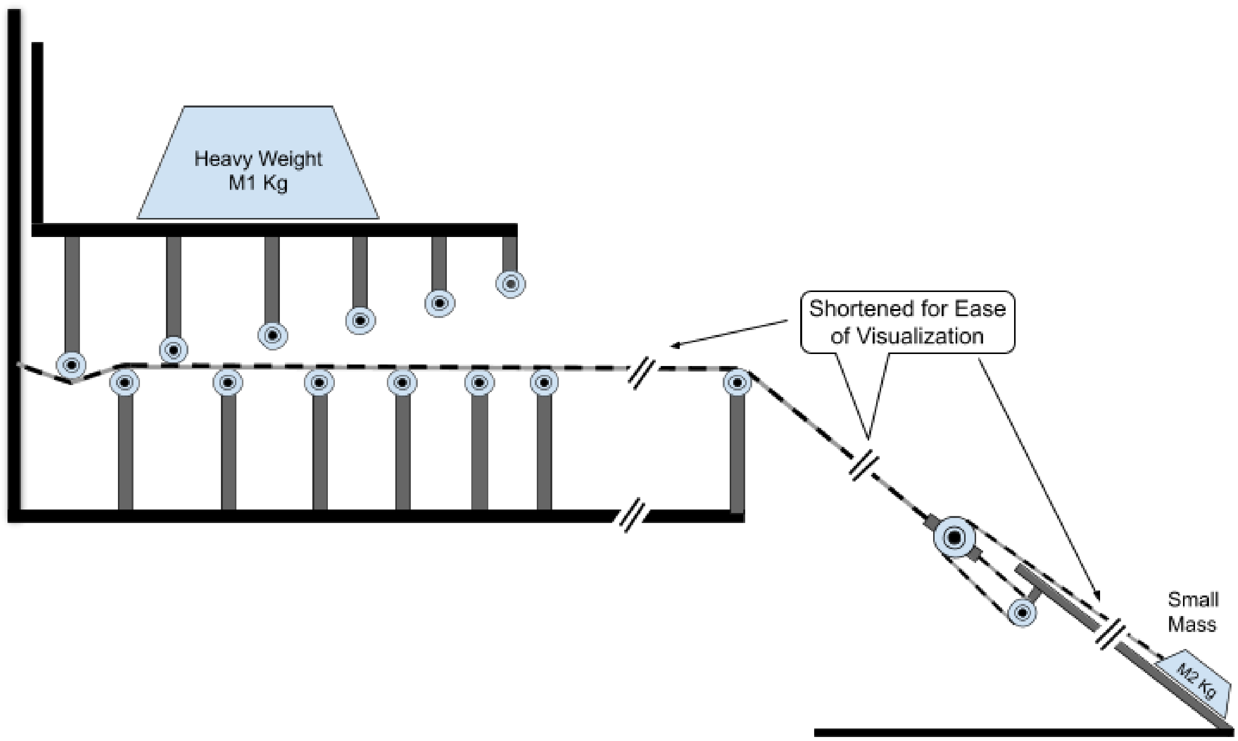

Figure E1. A twelve-member array with a 3:1 second stage, arranged so that the output member passes over one fewer sheave at payload speed than the arrangement of Figure 7. Breaks in the rope runs are drawn for compactness and do not represent joints.

§7.2 identifies the rotational energy of high-speed sheaves as the largest term omitted from the models in this paper. A sheave turning at payload rim speed stores $\frac{1}{2}m_s v^2$, the translational kinetic energy of its own mass at payload velocity, but only fixed-axle sheaves on the output member do so: the movable blocks of the fixed-ratio stage turn at $1/k$ of payload speed and store $1/k^2$ as much. Two sheaves run at full speed in the arrangement of Figure 7, the orientation sheave and the first sheave taking the output member. Counting the remainder through n_{eff} of (11) gives $n_{\text{eff}} \approx 2$ to 3, so 100 g sheaves store 10 to 15 kJ against a 50.6 kJ output at 1,000:1, and 101 to 151 kJ against 505 kJ at 10,000:1. In both cases that is 20 to 30% of the energy delivered to the payload, equivalently 17 to 23% of the energy extracted from the source.

Removing a single full-speed sheave therefore recovers about 10% of the output energy, independent of mass ratio, because both the stored and delivered energies scale as v^2. The orientation sheave is a convenience of layout rather than a requirement, and removing it takes n_{eff} from about 2 to about 1. It costs nothing in ratio profile: the fixed-stage ratio k is unchanged, only the number of direction changes imposed on the output member.

The same scaling has a consequence for the exit velocities reported in §5.3, and it should be stated plainly. Because sheave energy and payload energy both go as v^2, fast-path inertia acts exactly as added payload mass, and $v = \sqrt{2E/(m + n_{\text{eff}} m_s)}$, which is (11) rearranged. Against a 1 kg payload, $n_{\text{eff}} \approx 2$ to 3 is a 20 to 30% mass increase and therefore a **9 to 12% reduction in exit velocity** (Table 15):

Table 15: Exit velocity corrected for the rotational inertia of sheaves in the fast path, which acts as added payload mass. $n_{\text{eff}} = \sum_i (v_i/v)^2$ counts each sheave by the square of its rim speed.

	§5.3 as reported	$n_{\text{eff}} = 3$	$n_{\text{eff}} = 2$	$n_{\text{eff}} = 1$ (E.1)
100:1	100.2 m/s	87.9	91.5	95.5
1,000:1	318.2 m/s	279.1	290.5	303.4
10,000:1	1004.8 m/s	881.3	917.2	958.0

The exit velocities of §5.3 are therefore optimistic by roughly 9 to 12% on this account alone, and each full-speed sheave removed recovers about 4% of velocity. This is a larger effect than any difference between the candidate geometries of §5.3, whose exit velocities all lie within 1% of the ideal. It does not change the peak-to-mean force figures, which are ratios and largely insensitive to a uniform added mass, but it does mean the velocity column should be read as an upper bound.

The general rule follows directly: **minimise the number of elements the output member touches, not the number of elements in the machine.** Sheaves on the array side travel at $1/k$ of payload velocity and store $1/k^2$ of the energy, so they are nearly free; a single sheave moved out of the fast path is worth k^2 of them.

E.2 Symmetric array

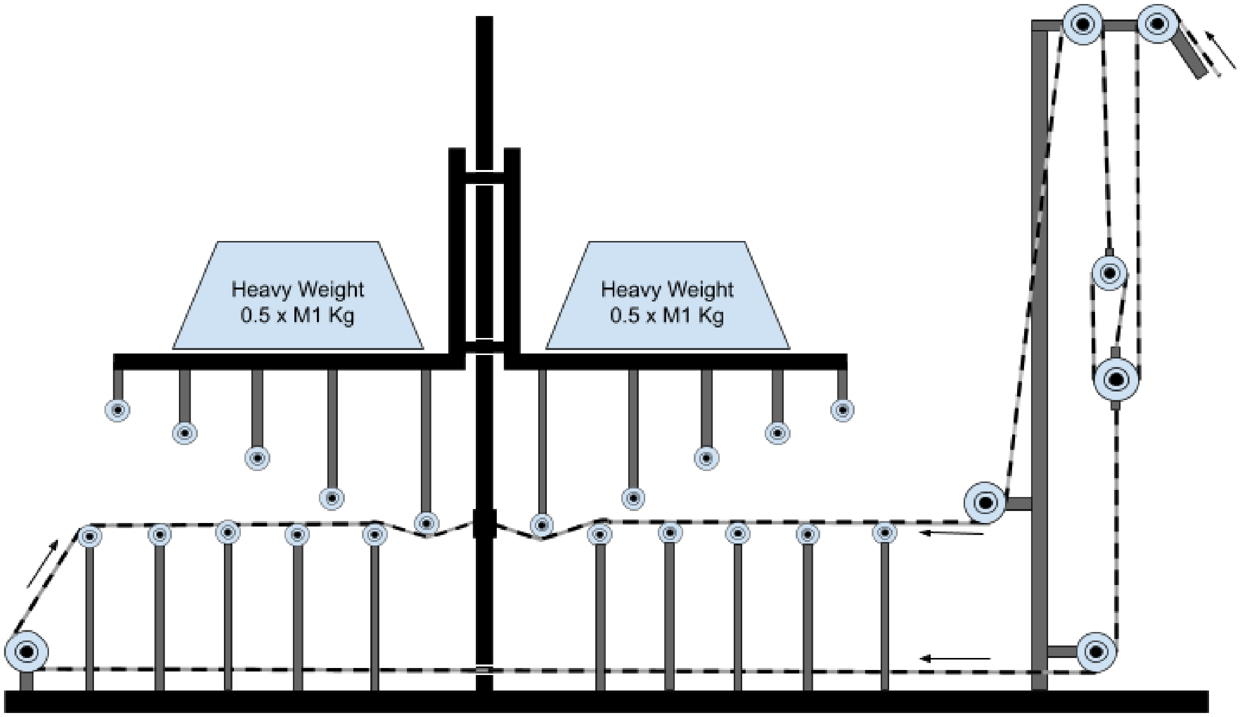

Figure E2. A sixteen-member symmetric arrangement with a 9:1 second stage. The source mass is divided into two halves of $M/2$ either side of a central guide, each driving its own half-array, with the two ropes combining at the second stage.

The force opposing the carriage reaches roughly **23 times the weight of the source mass** in all three designs of §5.3: 23.4 kN at 100:1, 228 kN at 1,000:1 and 2.09 MN at 10,000:1. In the single-sided arrangement that load is reacted at an offset from the guide axis, so the guide bearings carry a moment as well as the vertical load, and they carry it while sliding at up to 10 m/s. Bearing friction under such a side load is not modelled anywhere in this paper and would appear directly as lost stroke energy.

Dividing the source mass and the array symmetrically about the guide cancels that moment. Equal tension either side is the natural state rather than something to be engineered: the two halves are mirror images driven by the same carriage through the same displacement, so they draw in rope at the same rate by construction.

The arrangement also degrades gracefully. The vertical load on the guide is unchanged (it is the sum of the two halves), but the lateral reaction is the *difference*, so any imbalance short of total still improves on the single-sided case, and improves on it by the reciprocal of the mismatch (Table 16):

Table 16: Lateral load on the guide: single-sided against symmetric with a tension mismatch between the two half-arrays.

	single-sided	symmetric, 10% mismatch	symmetric, 20% mismatch
100:1	23.4 kN	2.3 kN	4.7 kN
1,000:1	228 kN	22.8 kN	45.7 kN

	single-sided	symmetric, 10% mismatch	symmetric, 20% mismatch
10,000:1	2.09 MN	209 kN	419 kN

A 20% mismatch, far larger than anything a built machine should exhibit, still reduces the lateral load fivefold. There is no configuration in which the symmetric arrangement is worse.

The tolerance required to stay well inside that is undemanding, because the array members are compliant. A difference δ in free length between the two ropes produces a fractional tension imbalance $\delta/(\varepsilon L_r)$, where ε is the working strain and L_r the rope length. At 1% working strain on 10 m ropes, 1 mm of length error gives 1% imbalance and 10 mm gives 10%. The compliance that §5.7 shows to be necessary for force uniformity thus also makes load sharing self-correcting, and this is why the arrangement is not analogous to tandem rigid actuators, where a small length error transfers load abruptly.

The practical consequences are lower bearing friction, more uniform descent, and a guide that can be sized for a nearly axial load.

The symmetric arrangement carries two further advantages that follow from the same geometry. The frame sees a balanced load rather than an overturning moment, which reduces the mass of structure required to react it. And with two ropes in parallel the tension in each is halved, which by the relation $T = kF/N_r$ of §3.3 permits either a smaller cross-section or a higher safety factor in the array members. This does *not* improve the critical-speed margin of §7.1: that limit binds on the output member, which carries the full payload force F and travels at payload velocity regardless of how many array members feed it, and $c = \sqrt{\sigma/\rho}$ falls rather than rises if working stress is reduced at fixed cross-section. The array members travel at $1/k$ of payload velocity and are nowhere near their critical speed in any case.

E.3 What the variants do not change

Neither variant alters $G_{\text{array}}(d)$, the target profile, the fitting procedure, or any result in §5. The array geometry is a property of the span widths and engagement offsets alone. What the variants change is the parasitic mass in the fast path and the load path through the structure, both of which sit outside the models used here and both of which are, on the estimates above, larger effects than the differences between the candidate geometries of §5.3.

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

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

[4] General Atomics Aeronautical Systems. EMALS: Electromagnetic aircraft launch system. Product datasheet, 2024.

[5] Naval Sea Systems Command. USS Gerald R. Ford (CVN 78) EMALS testing. Technical report, Naval Sea Systems Command, 2017.

[6] 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.

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

Code and data availability

The rigid-body and compliant simulators, the target solver, the least-squares array fit and the design procedure of §4.7 are available as a documented Python package at https://github.com/ramimahdi/ropecomb, released under the MIT licence. It ships the three canonical geometries of §5.3 and a regression suite that reproduces their exit velocities, peak-to-mean forces and fit residuals, together with the geometry verification of Appendix C. The figure-generation scripts are not included.

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

A United States provisional patent application, No. 64/124,997, directed to aspects of the transmission described in this article, was filed before this article was made public. The author has a financial interest in its commercialisation. 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

Large language models (Anthropic Claude, Google Gemini) were used in preparing this manuscript: for copy-editing and structural review; for LaTeX formatting and plot generation; for transcribing and packaging the author's simulation and fitting code into the released Python library, and for running the parameter searches with it; and for algebraic checking and cross-verification of reported values against the underlying data. Figure 1 is a stylised line rendering, produced with AI assistance, of a three-dimensional model built by the author. The invention, the simulation models, the fitting formulation, the specification of the design problem, all design decisions, and the interpretation of the results are the author's own. The author has verified every claim reported here and takes full responsibility for the content of this article.

Correspondence

Rami N. Mahdi, Interesting Machines LLC, Sacramento, California, USA. ramimahdi@interestingmachines.com

ORCID: https://orcid.org/0009-0000-0817-3887

Figures

Table 17 indexes every figure in this paper against the file that produces it.

Table 17: Index of figures.

#	file	content
1	`fig00_hero_line.png`	the apparatus, pictorial
2	`fig02_concept.png`	energy transfer concept
3	`fig03_target_profiles.png`	target ratio profiles
4	`fig01_howitworks.png`	apparatus annotated with part names
5	`fig04_single_member.png`	single member, and its ratio
6	`fig05_configurations.png`	four arrangements, and their profiles
7	`fig06_apparatus.png`	one-rope and two-rope apparatus
8	`fig07_stage_options.png`	second-stage options
9	`FIG_search_vs_fit.png`	search versus fitting
10	`FIG_comb_1000to1.png`	fitted geometry, 1,000:1
11	`FIG_target_selection.png`	why the source is stopped at 4 m/s
12	`FIG_rigid_1000to1.png`	dynamics, rigid baseline
13	`FIG_dynamics_1000to1.png`	dynamics, compliant and pre-tensioned
14	`FIG_instability_1000to1.png`	traction-loss instability

Appendix D reproduces the geometry, dynamics and instability figures for the 100:1 and 10,000:1 cases as Figures D1–D8.