

The Opposed-Piston Free-Piston Linear Generator:

A First-Principles Case for Sub-Crankshaft Combustion and the Racing-Industrial Execution Model

Companion Working Paper to **Racing to Advanced Air Mobility: Project MAXIMUM BOOST** (SSRN Working Paper)

M. McAlpine

X1Racing, LLC

mac@x1racing.com

Working Paper — May 2026

Abstract

This paper develops a first-principles efficiency and mass analysis of the opposed-piston free-piston linear generator (OPFPLG) and answers, in terms a non-engineer regulator can verify, two questions left implicit in the companion paper *Racing to Advanced Air Mobility: Project MAXIMUM BOOST* [1]: (i) why the architecture is capable of brake thermal efficiency (BTE) in the 55–70% range that conventional crankshaft-based reciprocating engines cannot reach, and (ii) why the architecture has not been commercialized despite the underlying physics being well understood for decades. The efficiency case is built by decomposing the parasitic losses of a conventional four-stroke engine — crankshaft side-load friction, valvetrain friction, four-stroke pumping work, part-load throttling loss, and rotational-to-electrical conversion loss — and showing that each loss is structurally eliminated by the OPFPLG architecture rather than incrementally reduced. The analysis additionally shows that the architecture admits three efficiency-recovery mechanisms unavailable to conventional crankshaft engines: turbo-electric compounding of exhaust enthalpy onto the DC bus, pulse-width-modulated (PWM) firing at the single peak-BTE point, and charge-density modulation of boost pressure and per-stroke fuel mass to vary absolute power per stroke without leaving the peak-BTE operating point. The eVTOL use case that motivates the companion paper is shown to be a near-perfect match for this constant-output operating philosophy, because a VTOL aircraft has no idle requirement and a much narrower required output dynamic range than an automobile. Each gain is grounded in measured friction and conversion data from production engines, not in projected combustion improvements. The non-commercialization is shown to be an industry-structure problem, not a physics problem, traceable to four specific causes: emissions-only regulatory targeting from 1990 onward, the absence of a viable output-electrification platform until the last decade, the late convergence of enabling component technologies, and the collapse of internal-combustion R&D budgets following the industry's pivot to battery-electric. The paper concludes that the appropriate execution model to close the remaining BTE gap is rapid parallel iteration by elite motorsport engineering organizations operating under a shared-data contracting structure — the model formalized in the companion paper as Project MAXIMUM BOOST.

Keywords: free-piston linear generator; FPLG; opposed-piston engine; brake thermal efficiency; BTE; engine friction decomposition; turbo-electric compounding; MGU-H; pulse-width modulation; PWM; charge-density modulation; wastegate-controlled output; constant-output series hybrid; eVTOL idle-free

operation; motorsport-led innovation; Project MAXIMUM BOOST; industry inertia; emissions regulation; combustion efficiency.

JEL Codes: L62 (Automobiles; Aerospace); L94 (Electric Utilities); O31 (Innovation and Invention); O33 (Technological Change); Q42 (Alternative Energy Sources); Q55 (Environmental Technological Innovation).

Nomenclature

- **BTE** = Brake Thermal Efficiency
- **DOHC** = Dual Overhead Camshaft
- **DPF / SCR** = Diesel Particulate Filter / Selective Catalytic Reduction
- **FPLG** = Free-Piston Linear Generator
- **ICE** = Internal Combustion Engine
- **MGU-H** = Motor-Generator Unit, Heat (exhaust-driven electrical recovery)
- **NdFeB** = Neodymium-Iron-Boron (permanent magnet material)
- **OPFPLG** = Opposed-Piston Free-Piston Linear Generator
- **OEM** = Original Equipment Manufacturer
- **PMB** = Project MAXIMUM BOOST
- **PWM** = Pulse-Width Modulation

I. Introduction

The companion paper to this work [1] argues that a performance-based motorsport rulebook can pull next-generation combustion architectures into civil service a decade ahead of conventional certification pathways, and uses the free-piston linear generator (FPLG) as its central example. That paper assumes — but does not derive — that the OPFPLG architecture is capable of brake thermal efficiency well above the ~50% ceiling of the best contemporary production engines. The present paper supplies that derivation from first principles, in deliberately accessible terms, so that a federal regulator, program officer, or appropriations staff member can audit the efficiency claim without specialized combustion training.

The argument proceeds in four steps. Section II identifies three structural taxes that every conventional reciprocating internal-combustion engine pays regardless of combustion quality — the crankshaft, the four-stroke cycle, and the valvetrain — and a fourth tax that any non-direct-drive engine pays at the mechanical-to-electrical conversion step. Section III decomposes the friction and mass cost of each tax using measured data from production engines, shows that the OPFPLG architecture structurally eliminates each one, and identifies two additional efficiency-recovery mechanisms — turbo-electric compounding of exhaust enthalpy and pulse-width-modulated peak-point operation — that the architecture admits at no structural cost and that conventional crankshaft engines cannot easily admit. Section IV addresses the second question: given that the physics has been understood for decades, why has no major automaker, aerospace prime, or government program brought an OPFPLG to production? Section V argues that the appropriate execution model to close the remaining gap is rapid parallel iteration by elite motorsport engineering organizations, formalized as Project MAXIMUM BOOST in [1].

The paper makes no speculative thermodynamic claims. Every efficiency gain enumerated below is the direct elimination of a measured loss in production engines, not a projected improvement in combustion. The OPFPLG's combustion event itself is broadly comparable to a modern high-pressure compression-ignition or premixed-charge compression-ignition cycle; the architectural gain comes from what the architecture removes, not from what it adds.

II. The Three Structural Taxes of Conventional Combustion

Every conventional reciprocating internal-combustion engine, from a 1908 Model T to a 2026 heavy-duty diesel, pays three structural taxes that have nothing to do with the quality of its combustion event:

- **The crankshaft tax.** The engine converts linear piston motion into rotational shaft motion through a connecting rod and crankshaft. The connecting rod transmits force at an angle to the cylinder axis, which forces the piston laterally into the cylinder wall. This side-load is the largest single source of mechanical friction in any piston engine.
- **The four-stroke tax.** The engine produces one power stroke for every two crankshaft revolutions. Half the strokes are dedicated to moving air rather than producing work, and the throttling and valve-overlap losses required to manage that air exchange consume an additional fraction of indicated power.
- **The valvetrain tax.** The engine carries camshafts, valves, springs, lifters, and a timing drive to gate air exchange. Overcoming valve-spring force is not free; the valvetrain consumes a measurable fraction of total engine friction and contributes substantial mass.

A series-hybrid powertrain — combustion engine driving an electrical generator that powers an electric drive — adds a fourth tax: a rotational-to-electrical conversion step that is, even in the best modern alternators, an order of magnitude lossier than direct linear-to-electrical conversion. And every conventional engine, regardless of architecture, discards a substantial fraction of fuel energy as exhaust enthalpy and as part-load throttling loss when operated below its peak operating point.

The OPFPLG architecture eliminates all four structural taxes and addresses both residual loss categories. Two opposed pistons translate linearly along a common axis, with no crankshaft, no connecting rods, and therefore no side-load. Combustion gas pushes the pistons apart; a compressed gas spring (or, in some designs, a second combustion event on the opposing face) returns them. The pistons cover and uncover intake and exhaust ports in the cylinder wall on the same stroke, eliminating the valvetrain. Permanent magnets on the moving piston assemblies pass through a linear stator, generating electricity directly from linear motion. Exhaust energy is recovered electrically by a turbo-electric compounding stage feeding the same DC bus. And because the generator has no mechanical link to the wheels — or, in the eVTOL case, to the rotors — the engine can be designed to run only at its single best operating point, with output modulated by pulse-width modulation of firing events and charge-density modulation of boost pressure and per-stroke fuel mass rather than by throttling.

III. Friction, Mass, and Conversion Decomposition

This section quantifies each of the four taxes using measured data from production engines, then identifies what the OPFPLG architecture removes.

A. Eliminating the Crankshaft

In a conventional engine, piston-and-ring friction — driven primarily by crankshaft-induced side-load — accounts for roughly 40–50% of total mechanical friction. Crank and rod bearing friction adds a further 25–30%. Removing the crankshaft removes both. A reasonable engineering estimate is a 40–55% reduction in total mechanical friction from this change alone.

The mass implications are equally significant. A crankshaft on a comparable-displacement engine weighs approximately 20–40 kg, and a flywheel — which exists only to smooth the torque pulses a crank inherently produces — adds a further 10–20 kg. A free-piston layout eliminates both. The reciprocating assembly also becomes substantially lighter because there is no need to oversize the piston skirt to resist side-load wear; the piston needs only to seal against combustion pressure.

B. Two-Stroke Operation

A four-stroke engine produces one power stroke per two crankshaft revolutions. A two-stroke fires every cycle. For the same delivered power, a two-stroke engine requires roughly half the displacement and half the reciprocating mass of a four-stroke. The 5–10% of indicated power that a four-stroke loses to throttling and valve-overlap pumping is also eliminated.

Two-stroke operation has historically been associated with poor emissions and oil consumption in carbureted small engines. Modern direct-injection two-stroke combustion at production injection pressures avoids both problems: fuel is injected after the exhaust port closes, so there is no short-circuiting of unburned charge, and the lubricant is not mixed with the fuel.

There is a further architectural consequence specific to the series-hybrid use case. A conventional automotive engine spends most of its operating life at part load, where brake thermal efficiency collapses — a modern gasoline engine that achieves 38% BTE at its best operating point typically delivers 18–25% BTE under real-world drive-cycle conditions because it is throttled, off-cam, and pumping against vacuum for most of the cycle. The OPFPLG, as a series-hybrid generator with no mechanical connection to the wheels, has no such constraint. It can be designed to run only at its single best operating point — the combustion condition that maximizes BTE — and modulate its average output to meet demand by two complementary mechanisms operating together. First, pulse-width modulation (PWM): firing strokes at the optimum point and either stopping between cycles or running short bursts at full output, with the buffer battery and the linear-generator inertia absorbing the resulting power pulses. Second, charge-density modulation: varying the boost pressure delivered by the compounding turbine, the wastegate setting, and the corresponding fuel injection mass per stroke so that each stroke runs at the same BTE-optimal combustion phasing and equivalence ratio but delivers a tunable absolute power per stroke. Together the two mechanisms give the system a continuous output range from a few percent of rated power up to peak, with every fired stroke occurring at or very near the peak-BTE operating point. The engine is either at its peak-efficiency point or it is off. There is no part-load regime. This recovers, at the system level, the 10–20 percentage-point gap between peak-point BTE and drive-cycle-average BTE that the conventional automotive use case throws away. PWM and charge-density modulation also dramatically simplify the combustion control problem, since the engine no longer needs to maintain a stable combustion event across a wide operating envelope — it only needs to be very good at one operating point.

The advantage is even more pronounced in the eVTOL use case that motivates the companion paper [1]. An automobile spends large fractions of its operating life at idle — at stoplights, in traffic, in parking — which is precisely the operating regime where a conventional engine is least efficient and where the PWM strategy is most valuable but also most stressed by frequent restart cycles. A VTOL aircraft has effectively no idle requirement: the vehicle is either powered down on the ground or it is airborne under continuous power demand (takeoff, hover, climb, cruise, descent, landing). The PWM and charge-density-modulation envelope therefore needs to span only the airborne power band — typically a 3-to-1 to 5-to-1 ratio between minimum continuous and maximum peak demand — rather than the 20-to-1 or greater dynamic range that defines automotive duty. This narrower required envelope further simplifies the combustion design and makes the constant-output, peak-efficiency operating philosophy a near-perfect match for the eVTOL mission profile. The OPFPLG's structural advantage over conventional engines is largest precisely where the companion paper places its first civil application.

C. Eliminating the Valvetrain

A modern dual-overhead-camshaft cylinder head carries two camshafts, sixteen or more valves, springs, lifters, and a timing drive (chain or belt). The valvetrain consumes 10–15% of total engine friction in a contemporary four-valve-per-cylinder engine — overcoming valve-spring force is the dominant component — and adds 30–70 kg per head. An opposed-piston engine has no cylinder head and no valvetrain; the pistons themselves uncover the intake and exhaust ports as they approach bottom dead center. The friction is removed; the mass is removed; and a substantial source of cycle-to-cycle variability (valve timing, valve lash, spring degradation) is removed with it.

The elimination of the cylinder head has a second-order thermodynamic benefit: the cylinder head is the largest single heat-loss path in a conventional engine, because its surface area is large relative to the combustion volume. An opposed-piston combustion chamber, bounded by two pistons and the cylinder wall, has a substantially lower surface-to-volume ratio at top dead center. Less heat is lost to the structure; more heat is converted to work.

D. Direct Linear-to-Electrical Conversion

A conventional series hybrid runs the energy path combustion → linear piston motion → rotational shaft motion → electrical power. The mechanical-to-rotational conversion through the crankshaft is approximately 85–90% efficient under best conditions. The OPFPLG architecture eliminates the rotational step entirely. A linear alternator with a Halbach-array NdFeB translator and an amorphous-metal (Metglas) stator converts linear piston motion directly into electrical power at 96–98% conversion efficiency. This is a 7–10 percentage-point gain on the conversion path alone, before any thermodynamic improvement in the combustion event is considered.

E. Turbo-Electric Compounding

A conventional engine discards a substantial fraction of its fuel energy as exhaust enthalpy — hot, pressurized gas that leaves the cylinder still carrying significant thermodynamic potential. A turbocharger recovers a small portion of this energy mechanically, but only to compress intake air, and only across the narrow speed range where the turbo and the engine match. Excess exhaust energy is dumped through the wastegate.

Because the OPFPLG produces electrical output rather than shaft torque, exhaust energy can be recovered electrically without a mechanical coupling to the engine. A turbo-electric compounding stage — an exhaust-driven turbine spinning a high-speed permanent-magnet alternator, with the resulting electrical power fed onto the same DC bus as the main linear generator — converts otherwise-wasted exhaust enthalpy directly into useful electricity. The recovered power adds to system output without burning additional fuel. Modern turbo-compounded heavy-duty diesels recover 3–5% of fuel energy this way using mechanical compounding; an electrical compounding stage, freed from the requirement to gear-match a crankshaft, can be optimized purely for thermodynamic recovery and routinely delivers 5–8% additional system efficiency.

This is the same architectural principle that Formula 1 power units have used since 2014, where the MGU-H (motor-generator unit, heat) recovers exhaust energy electrically and either feeds the traction battery or drives the compressor on demand. F1 has demonstrated that the technology is production-mature and reliable across season-length operational cycles. The OPFPLG architecture admits MGU-H-class compounding natively — the DC bus is already there, the power electronics are already there, the generator-side control loop is already there — at no architectural cost. A conventional crankshaft engine cannot easily admit electrical compounding because it has no DC bus to feed into without adding the rotational-to-electrical conversion stage that the OPFPLG has structurally removed.

F. Cumulative Result

Table 1 summarizes the four taxes plus the compounding recovery and what the OPFPLG architecture delivers from each.

Table 1. Decomposition of Conventional ICE Losses vs. OPFPLG Architecture.

Loss Source	Conventional 4-Stroke	OPFPLG
Mechanical friction (piston, bearings, valvetrain)	10–15% of fuel energy	3–5%
Pumping & throttling losses	5–10%	≈ 0%
Cylinder-head heat loss	High (large head surface area)	Eliminated (no head; lower surface-to-volume ratio)
Mechanical-to-electrical conversion	85–90% (rotational alternator)	96–98% (direct linear)
Exhaust enthalpy recovery	0–5% (mechanical turbocompounding, limited)	5–8% (turbo-electric compounding onto DC bus)
Part-load penalty (drive-cycle average vs. peak BTE)	10–20 percentage-point collapse under throttled operation	≈ 0 (PWM + charge-density modulation, peak-point only)
Dry mass per continuous horsepower	≈ 0.8–1.2 kg/hp	≈ 0.35 kg/hp

Stacking these gains is what produces the OPFPLG design targets of approximately 67% BTE at 1,500 hp continuous and 55% BTE at 2,000 hp peak in a 533 kg package referenced in [1] and Exhibit B [3]. The current best heavy-duty production diesel reaches approximately 50% BTE under optimum conditions; the on-road operational average is 42–46%. The Travis Air Force Base 250 kW linear-generator pilot deployed in March 2026 [2] delivers 46% in federal service even when optimized for fuel flexibility

rather than peak efficiency, providing an existing federal data point that brackets the lower bound of the achievable range.

IV. Why Has This Not Been Built Before?

The first-principles case in Section III is not new. The Pescara free-piston compressor of the 1920s and the Junkers Jumo opposed-piston aviation diesel of the 1930s established the underlying mechanical principles a century ago. Toyota Central R&D Labs demonstrated a two-stroke 10 kWe FPLG prototype at 42% indicated thermal efficiency in 2014 [4]; the German Aerospace Center (DLR) ran a parallel program [5]; Volvo, GM, Ford, and others built running FPLG prototypes between 2000 and 2015. None reached production. The reasons are not in physics. They are in industry structure, and they decompose into four causes.

A. Regulators Set the Wrong Target

From roughly 1990 onward, U.S. and EU emissions regulation focused almost exclusively on criteria pollutants: oxides of nitrogen (NO_x), particulate matter, unburned hydrocarbons, and carbon monoxide. The regulatory architecture rewarded aftertreatment — three-way catalysts, diesel particulate filters, selective catalytic reduction — bolted onto existing crankshaft-based engine architectures. Thermal efficiency and CO₂ were treated as side concerns under separate, weaker frameworks (CAFE in the United States, fleet-average CO₂ targets in the EU). The result is that for thirty years, manufacturer engineering budgets were directed at compliance with the criteria-pollutant framework rather than at the architectural redesign that would have delivered the largest absolute efficiency gain. There was no commercial pressure to redesign the engine itself.

B. The Architecture Has No Output Shaft

An FPLG produces electricity, not torque. It is fundamentally incompatible with a transmission-based powertrain. Until battery and power-electronics costs fell sufficiently in the 2015–2025 window to make series-hybrid architectures commercially viable, there was no production vehicle platform that could actually use a free-piston engine. The OEM FPLG prototypes built in the 2000–2015 window were shelved not because the engines failed — most ran successfully on dynamometers — but because they did not fit the transmission tunnels and final-drive architectures the manufacturers were tooled to build. The architecture was orphaned by the powertrain it did not match.

C. Enabling Component Technologies Matured Separately

Four distinct technology streams had to mature before a commercially-viable OPFPLG could be assembled from off-the-shelf components: Halbach-array NdFeB permanent magnets at production cost and temperature stability; amorphous-metal (Metglas) stator cores at production yield; gigapascal-class direct-injection fuel systems; and real-time, sub-millisecond piston-position control electronics with the computational capacity to manage two free pistons against combustion-pressure disturbances. Each of these matured in the 2010–2025 window, driven by demand from unrelated markets (consumer electronics, wind turbines, common-rail diesel passenger cars, edge-AI inference hardware). Building an OPFPLG in 2005 required inventing four enabling technologies simultaneously. Building one in 2026

means assembling existing components. The window for cheap construction has opened only in the last 36 months.

D. EV Capital Displaced ICE R&D

Once major OEMs committed to battery-electric vehicle platforms from approximately 2018 onward, internal-combustion innovation budgets collapsed across the industry. Engineering headcount that would historically have been deployed against efficiency was redirected to battery chemistry, traction motors, and charging infrastructure. The FPLG was orphaned a second time, by an industry that decided the correct answer to combustion's limitations was "no engine at all" rather than "a fundamentally better engine." That decision is now visibly under stress as grid capacity, critical-mineral supply chains, and consumer adoption rates fall short of the assumptions on which the BEV pivot was based, but the engineering capacity to revisit the combustion question inside the legacy OEMs is no longer present at the scale required.

The combined effect of these four causes is that the OPFPLG sits in an industry gap that legacy automotive cannot close, because closing it requires admitting that thirty years of efficiency improvement was bounded not by the physics of combustion but by an architecture the industry chose to keep.

V. The Racing-Industrial Execution Model

The architecture has been understood at a first-principles level for decades. What has been missing is an organization with the freedom to build it from scratch rather than fitting it into a transmission tunnel — and the right execution model to actually do so at speed. Project MAXIMUM BOOST, formalized in the companion paper [1] and elaborated in Exhibit B [3], supplies that model. Its execution premise is that the cheapest, fastest, and best path from a 46% federal-baseline FPLG to a 60–70% championship-level FPLG is rapid parallel iteration by the world's best racing engine experts, operating outside the constraints of legacy OEM development cycles.

A. Why Racing

Racing engine organizations have, for a century, treated each dyno session as a learning loop rather than a milestone gate. They routinely take an engine concept from clean-sheet design to running hardware in months. They instrument every test to the millisecond. They run multiple build variants in parallel as a matter of course because the season schedule does not permit sequential development. And — uniquely among engine-development organizations — they are paid to extract maximum work from every gram of fuel and every gram of mass, because grams of fuel and grams of mass directly determine race outcomes. The combustion knowledge that the legacy OEM engineering function lost in the BEV pivot has been preserved, intact and continuously refined, inside the motorsport engineering community.

B. Parallel Iteration Under Shared-Data Contracting

Project MAXIMUM BOOST applies the racing execution model directly to the OPFPLG. Under its contracting structure, tens of variants are developed in parallel — for the first time ever in the history of free-piston engine development — each variant exploring a different bore/stroke ratio, port-timing geometry, injection strategy, translator-mass schedule, or gas-spring configuration. All variants share instrumentation data, dyno results, and failure-mode learnings under a common contracting structure that

prevents the proprietary siloing that has historically slowed engine development inside the OEM environment. No competing program in the history of FPLG development has had this many independent expert teams working the same architecture simultaneously with a shared data pipeline.

C. Why This Is the Cheapest, Fastest, and Best Path

It is the cheapest path because parallel small-team iteration burns a fraction of the overhead of a sequential corporate development cycle. Racing organizations are structurally lean — the per-engineer overhead of an F1 or LMDh powertrain group is a small fraction of the equivalent inside a global automaker. Spreading the development across many such organizations multiplies the number of independent experiments while keeping total program overhead bounded by the contracting structure rather than by any single organization's burden rate.

It is the fastest path because the learning rate scales with the number of independent experiments running concurrently. A sequential development program learns one variant per development cycle; a parallel program of n teams learns n variants per cycle. With shared data, every team learns from every other team's results. The expected time to convergence on the optimum design falls roughly as $1/n$, bounded below by the irreducible dyno-test cycle time, which is itself short by the standards of conventional engine development.

It is the best path because the people who already know how to extract maximum work from every gram of fuel and every gram of mass are the ones who win races for a living. The combustion expertise required to close the gap from 46% to 60%+ BTE is, at this point in history, concentrated in motorsport rather than in the legacy automotive engineering function. Project MAXIMUM BOOST is the contracting instrument that puts that expertise to work on the architecture where it can deliver the largest absolute efficiency gain available to internal combustion.

VI. Conclusion

The OPFPLG efficiency case is not a speculation about future combustion improvements. It is the structural elimination of measured losses in production engines: crankshaft side-load friction, valvetrain friction, four-stroke pumping work, cylinder-head heat loss, and rotational-to-electrical conversion loss. Each loss is documented in the production-engine literature; each is removed by the OPFPLG architecture as a direct consequence of geometry rather than as a projected improvement.

The non-commercialization of the architecture is similarly structural. It reflects thirty years of emissions-only regulatory targeting, an architectural incompatibility with transmission-based powertrains, the late convergence of enabling components, and the collapse of internal-combustion R&D budgets inside the legacy OEMs. None of these obstacles is a physics constraint, and the first three are no longer binding.

The fourth — the loss of internal-combustion engineering capacity inside the OEMs — is binding, and it is the specific obstacle that Project MAXIMUM BOOST is designed to dissolve. Rapid parallel iteration by elite motorsport engineering organizations, under a shared-data contracting structure, is the cheapest, fastest, and best execution model available to close the 17-point BTE gap between today's best production engines and the physical ceiling that first principles indicate the OPFPLG architecture can reach. The

required timeline is 24–36 months, not a decade, because the architecture, the components, the regulatory on-ramp, and the engineering talent are all present and available today.

The United States need not invent new combustion physics, new battery chemistry, or new aviation statute to recover the efficiency that the architectural choices of the twentieth century left on the table. It needs a contracting structure that aligns motorsport engineering capacity with national infrastructure outcomes. That is what the companion paper proposes, and what this paper supplies the technical first-principles case for.

References

- [1] McAlpine, M., "Racing to Advanced Air Mobility: Project MAXIMUM BOOST," SSRN Working Paper, abstract id 6621460, 2026.
- [2] U.S. Department of the Air Force, Air Force Office of Energy Assurance, Tradewinds Solutions Marketplace contract award to Mainspring Energy, Inc. for 250 kW linear generator deployment, Travis Air Force Base, California, March 2026.
- [3] X1Racing, LLC, "Exhibit B: American Energy Sovereignty — Project MAXIMUM BOOST," 2026. Available: x1racing.com/x-files.
- [4] Goto, S., Moriya, K., Kosaka, H., Akita, T., Hotta, Y., Umeno, T., and Nakakita, K., "Development of Free Piston Engine Linear Generator System Part 2 — Investigation of Control System for Generator," SAE Technical Paper 2014-01-1193, Toyota Central R&D Labs, 2014.
- [5] Kock, F., Haag, J., and Friedrich, H. E., "The Free Piston Linear Generator — Development of an Innovative, Compact, Highly Efficient Range-Extender Module," SAE Technical Paper 2013-01-1727, German Aerospace Center (DLR), 2013.
- [6] Mikalsen, R., and Roskilly, A. P., "A Review of Free-Piston Engine History and Applications," Applied Thermal Engineering, Vol. 27, Nos. 14–15, 2007, pp. 2339–2352.
- [7] Heywood, J. B., Internal Combustion Engine Fundamentals, 2nd ed., McGraw-Hill, 2018. (Reference text for engine friction decomposition data cited in Section III.)
- [8] Richman, R. M., and Reynolds, W. C., "A Computer-Controlled Poppet-Valve Apparatus for Two-Stroke Engine Air Exchange Studies," SAE Technical Paper 840378, 1984. (Historical baseline on pumping-loss decomposition.)
- [9] Achatés Power, Inc., "Opposed-Piston Two-Stroke Engine Development Program," technical publications, 2014–2025.
- [10] McAlpine, M., "Nothing Rolls Coal Like an EV: The Case for Super-Efficient Fossil," LinkedIn, 2025. Available: x1racing.com/x-files.

— END WORKING PAPER —