

Energy Demand in Stereolithography Manufacturing: Equipment- and Build-Level Evidence, with Throughput and Allocation as Open Requirements

Industrial economics and deployment

Juan Seguro

3Dresyns by Resyner Technologies S.L.

© 2026 Juan Seguro. All Rights Reserved.

Correspondence: info@3Dresyns.com · www.3Dresyns.com

ER-475 · Engineering Review · Version 0.1 · 12 October 2026

Abstract

The evidence available for this review supports energy-demand analysis in stereolithography at equipment, build and desktop-case-study level, but it does not establish factory-scale throughput or a general allocation framework [1,2]. Source [1] addresses energy-consumption modelling for stereolithography-based additive manufacturing. Source [2] addresses life-cycle impacts of desktop stereolithography using case-specific parameters including product mass, support mass and building time. Source [1] was checked against the openly available full article during this review; source [2] was checked against its publisher and institutional-repository abstracts. Accordingly, throughput and allocation are treated as open requirements rather than measured or validated outputs. The review identifies the minimum distinctions needed to prevent equipment energy, build duration, throughput and allocation from being conflated. Synthetic energy accounts distinguish attempted and accepted output, rework and shared-load allocation while keeping measured totals separate from allocation conventions.

Keywords: energy demand · stereolithography · equipment energy · build energy · throughput · allocation

1. Scope and ownership

ER-475 concerns energy demand in stereolithography manufacturing at the levels directly represented by the available evidence: equipment, build and desktop case study. It does not claim factory-level energy mapping across multiple machines, shifts or sites. It does not claim a validated throughput model or a general allocation method.

2. Evidence base and source-access depth

Source [1] is a stereolithography energy-consumption modelling study and supports equipment/build energy analysis and sensitivity to process variables [1]. Source [2] is a desktop-stereolithography life-cycle assessment that uses product mass, support mass and building time as case parameters [2]. The full text of [1] was inspected for this appraisal of its subsystem model and experimental design. Source [2] was inspected at publisher- and institutional-abstract depth, so detailed numerical findings not supported there are not reproduced or inferred.

3. Equipment energy is not factory energy

Energy measured or modelled for a stereolithography build is not automatically equivalent to the energy demand of a production factory. A factory-level result would need to distinguish machine states, idle time, ancillary equipment, post-processing, failed builds, rework and site services where those are inside the declared system boundary. The present evidence does not establish such a complete multi-machine boundary.

4. Build time is not throughput

Building time can influence energy use and can be a useful parameter in a case study [2]. It is not, by itself, a throughput metric. Throughput would require a declared production denominator, such as accepted parts or accepted functional output per unit time, together with the treatment of failed builds, queueing, setup and post-processing. This review names throughput as an open requirement and does not report it as a result.

5. Allocation remains an open requirement

When one build or one machine serves multiple parts, products or jobs, a production analysis must decide how shared energy is assigned. Possible denominators can differ materially, and the present two records do not establish a generally transferable allocation rule for stereolithography manufacturing. Source [2] provides case-specific parametrisation but is not treated here as a validated general allocation framework.

6. Minimum reporting set

A bounded energy-demand record should identify the machine or equipment state, build duration, build contents, accepted output, failed or repeated work if included, post-processing boundary, ancillary equipment boundary and the denominator used for any normalised energy result. If throughput or allocation is not measured, it should be stated as not measured rather than reconstructed from build time.

7. Meter the energy boundary before allocating it

The proposed calculation separates energy measured for a build from energy assigned to its output. A time-resolved meter record can distinguish the declared operating interval from preceding or subsequent idle periods. Washing, post-cure, ventilation and other shared operations should enter only under the stated boundary, with their allocation rules visible. A machine-only result and a complete-workflow result should not share an unqualified label.

Consider a synthetic batch for which the declared printer interval consumes 4.0 kWh and the assigned share of post-processing consumes 2.0 kWh. If eight of ten attempted parts are accepted, the measured-and-assigned total is 6.0 kWh and the ratio is 0.75 kWh per accepted

part. Dividing by ten would instead give 0.60 kWh per attempted part. Both calculations can be reproduced, but they answer different questions. Neither is an observed VPP benchmark.

If the same accepted output also requires rework consuming 0.8 kWh within the chosen boundary, the accepted-part ratio becomes $6.8/8 = 0.85$ kWh. A comparison that omits rework in one route while including it in another would be inconsistent even if each meter reading were correct.

8. Make shared-load allocation a sensitivity, not a hidden fact

A mixed build creates an allocation problem. Equal assignment per item, assignment by occupied area, by volume or by another declared basis can produce different product-level results from the same measured build energy. The meter cannot decide which accounting convention is appropriate. The article therefore proposes reporting the build total as the primary observation and the allocated product result as a derived quantity.

Where the decision changes with the allocation convention, that dependence is itself a result. A study could report a range across justified conventions rather than implying that one untested rule is physically exact. Idle energy requires similar care: including an entire weekend in one small build's denominator may be appropriate for an operational period account but not for a narrowly defined process comparison.

Repeated builds should preserve loading, duration, accepted output and the operating boundary. This permits examination of whether a change reduces energy per build, per attempted part or per accepted part. The analysis ends at energy accounting. Conversion to environmental impact requires a separate lifecycle boundary and appropriate factors, and is not supplied by an electricity measurement alone.

8.1. Comparative appraisal of the two direct evidence anchors

Yang et al. [1] develop and experimentally assess a subsystem-based energy model for mask-image-projection stereolithography. Their analysis links configurable build parameters to electrical demand and also examines part surface quality. This is direct evidence that process parameter optimization can be evaluated against both an energy objective and a product-quality constraint. It does not, by itself, demonstrate an accepted-part energy intensity across production shifts, a shared post-processing network or different printer architectures. The distinction matters because the measured experimental machine and program define the scope of the model.

Mele et al. [2] address a different question: cradle-to-gate life-cycle impacts for a bottom-up desktop stereolithography case, parameterized by product mass, support mass and build time. That assessment includes material production, transport, transformation and post-processing. It therefore provides a broader impact boundary than a printer electricity meter, yet remains anchored to its case equipment, location and inventory choices. A result expressed as environmental impact per printed part cannot be treated as interchangeable with electrical energy per accepted part unless functional unit, inventory coverage and yield are reconciled.

The papers are complementary rather than competing measurements of the same quantity: [1] probes equipment-level energy modeling and process-variable effects; [2] addresses a life-cycle inventory and impact question. Their boundaries, endpoints and operating contexts should be compared before any numerical values are pooled. Neither paper establishes a universal industrial energy benchmark for all vat-photopolymerization platforms.

8.2. Measurement architecture and uncertainty

A qualification campaign can use at least three nested meter boundaries: the printer input, the cell including wash and post-cure, and the factory operation including shared auxiliaries. Report each boundary separately rather than replacing the first with a silently allocated estimate of the third. The sampling protocol should capture standby and warm-up, layerwise exposure or scanning, separation motion and post-build operations insofar as they are observable on the instrument. Meter sampling frequency, resolution, power-factor treatment and calibration status are part of the traceability record; their suitability depends on the transient behaviour of the tested machine.

For a build with variable power $P(t)$, electrical energy E is the integral of $P(t)$ over the declared time interval. If only interval-average readings are available, the approximation is the sum of average power multiplied by interval duration, with the sampling method disclosed. The analyst should retain the raw time series or export, the build-file identifier, firmware version and resin/process configuration. Differences in operating state are potential confounders: an apparent resin effect may instead arise from layer count, separation time, tray loading, control logic or post-cure recipe.

Measurement uncertainty and accounting uncertainty are distinct. The former concerns instrument accuracy, sampling and repetition; the latter concerns allocation of idle and shared loads and assignment to accepted output. A sensitivity analysis should vary both sets independently. Where accepted-part intensity is reported, retain numerator and denominator, the acceptance criteria and any exclusions. If failure rates are estimated rather than observed, the output is a scenario, not a measured benchmark.

8.3. Review-level comparison and research gaps

An adequate cross-study comparison needs a common functional unit, explicit process scope, machine architecture, geometry and orientation, layer schedule, batch packing, material consumption and support strategy, accepted yield, electricity boundary and geography/time relevant to impact conversion. A missing field does not make a paper unusable, but it precludes the comparison that depends on that field. When such fields are unavailable, narrative comparison of study design is more defensible than pooled energy-intensity averages.

The most consequential gap between these two sources and industrial decision-making is the lack of directly comparable, replicated measurements covering equipment, post-processing, accepted yield and utilization under one transparent production boundary. A future study should predefine the primary functional unit, stratify by printer architecture and loading, measure a representative run schedule, and show how conclusions change under alternative allocation rules. Environmental impact conclusions would additionally require independently documented electricity and material inventories; reduced electricity alone does not establish a reduced cradle-to-gate impact.

Measurement architecture for a comparable energy-intensity metric

Machine-rated power and measured electric energy are not interchangeable. A reproducible study needs an interval-defined measurement of build, standby, thermal control, auxiliary motion and post-processing loads, with a stated meter boundary. For bottom-up photopolymerisation, layer exposure time alone omits separation, recoating, waiting and ancillary power draw. Normalisation per printed volume is useful for process studies but not necessarily for comparing economic function: a rejected or overcured part consumes energy without contributing an accepted unit. Energy per conforming part should therefore accompany any energy per build or per unit cured volume.

System-boundary differences must be treated as uncertainty sources rather than silently normalised away. The cited energy studies investigate specific equipment and operational settings, and their reported trends cannot be extended directly to all stereolithography, DLP and LCD devices. Reported sensitivity should include build utilisation, part orientation, thickness, total build duration and post-processing duty cycles. This is a minimum reporting recommendation for future controlled comparisons, not a pooled estimate of process energy demand.

Separating machine electricity from accepted-output energy intensity

Electricity metering at the printer inlet does not necessarily include wash, drying, extraction, curing and environmental conditioning. A comparable assessment should state which peripherals are inside the system boundary and how shared services are allocated to the accepted production batch. Background idle consumption and standby time also require explicit treatment when the manufacturing schedule is intermittent.

Energy per nominal build and energy per accepted part answer different questions. A print that completes optically but fails dimensional qualification has consumed electricity without creating accepted output. Consequently the denominator should preserve reprints and post-processing rejects, while reporting a supplementary build-only metric for process diagnosis. Use measured electrical power profiles and timestamps where feasible rather than assuming nameplate power remains constant.

Claims about comparative sustainability need additional boundaries for upstream resin production, consumables and disposal. An equipment-level electricity result cannot establish life-cycle superiority over another manufacturing route. The review should classify each study by measured power scope, functional unit, yield treatment and allocation approach before making inter-study comparisons.

9. Functional-unit and measurement design

Electricity consumption should be decomposed into machine idle, thermal conditioning, exposure, mechanical movement, cleaning, post-curing and auxiliary loads. These stages have different dependencies on part geometry, batch size, environmental settings and utilization. An energy number measured at the printer plug cannot silently stand for whole-site electricity; conversely, allocation of building ventilation or shared post-processing to one build requires an explicit allocation rule.

The functional unit changes the interpretation of energy intensity. Energy per build is appropriate for equipment comparison under a fixed program, whereas energy per accepted part should include utilization, rejection and any rework; energy per dimensional volume can conceal geometry- and quality-dependent differences. For economically relevant comparisons, the accepted-part denominator is generally the safer reporting target, accompanied by its raw numerator and the acceptance definition. Neither the number of printed objects nor their mass alone guarantees functional equivalence.

A reproducible metering study would record sampling interval, instrument calibration, start and end criteria, machine state, build file, layers, loading fraction, resin temperature and each ancillary operation. Replicated builds with different loading levels would help separate fixed baseload from incremental exposure and motion demands. A factorial comparison must avoid presenting machine-to-machine differences as resin-only effects when optics, slicing, vat condition or firmware also change.

Uncertainty analysis should compare alternative shared-load allocations, rejection assumptions and utilization scenarios rather than reporting a single false-precision estimate. The literature may motivate this experimental design, but without boundary-matched photopolymer-specific measurements it cannot establish a universal energy intensity or a ranking against molding or other additive routes.

10. Transfer limits

The available evidence is direct for stereolithography but narrow in scale. It does not justify claims about all photopolymer manufacturing routes, all factory configurations or all allocation practices. The two records do not support extension to multi-machine or factory-level analysis.

11. Conclusions

The present evidence supports equipment- and build-level energy analysis for stereolithography and a desktop case-study view of life-cycle impacts [1,2]. It does not establish throughput or a general allocation framework. Those two elements therefore remain explicit qualification requirements for later work rather than outputs of this review.

References

1. Yang Y, Li L, Pan Y, Sun Z. Energy Consumption Modeling of Stereolithography-Based Additive Manufacturing Toward Environmental Sustainability. *Journal of Industrial Ecology*. 2017;21(S1):S168-S178. <https://doi.org/10.1111/jiec.12589>.
2. Mele M, Campana G, D'Avino G. Life cycle impact assessment of desktop stereolithography. *Journal of Cleaner Production*. 2020;244:118743. <https://doi.org/10.1016/j.jclepro.2019.118743>.

Competing interests

The author is Founder, CEO and CTO of 3Dresyns by Resyner Technologies S.L. The company develops and supplies photoresins and related process materials. No proprietary formulation, manufacturing sequence, process window or comparative product recommendation is disclosed.

Data availability

No new experimental dataset was generated.

Copyright

© 2026 Juan Segurolo. All Rights Reserved.