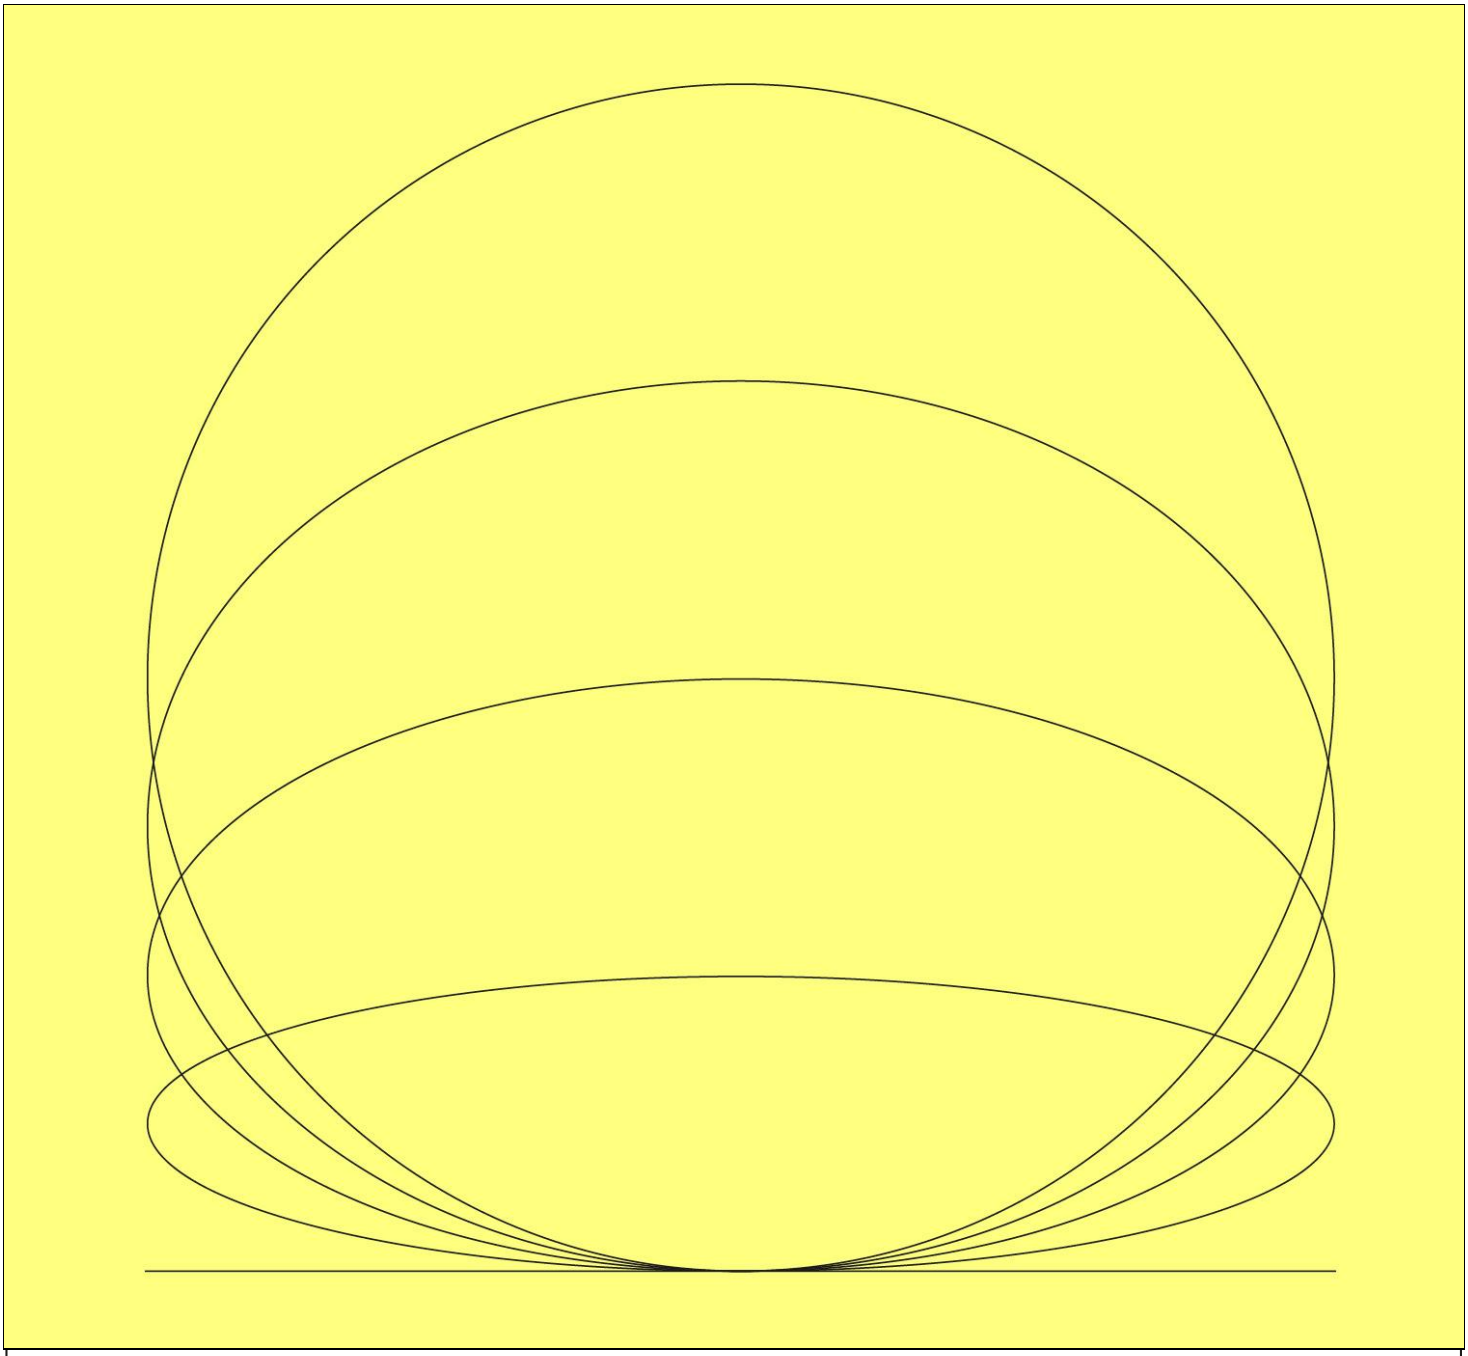


Report

H_2UpScale

Deliverable 3.3

Electrical system design

Version: 1.0
2026-07-28
HORIZON-JTI-CLEANH2-2024-03-02

H2UpScale

Clean Hydrogen Partnership

Co-funded by the European Union

Dissemination Level (Security)

Public	☒ **Public, fully open, e.g. web**
Sensitive	☐ Sensitive, limited under the conditions of the Grant Agreement. Restricted to the consortium and including the European Commission Services
Classified	EU classified (EUCI) under Decision No 2015/444: ☐ R-UE/EU-R (EU Restricted) ☐ C-UE/EU-C (EU Confidential) ☐ S-UE/EU-S (EU Secret)

Document verification and approval

	Name	**Beneficiary short name**	**Date** (DD-MM-YYYY)
Lead Author(s)	Theo Koch	RWTH	14-09-2026
Contributing Author(s)			
Quality Assurance Reviews the quality of contents in the deliverable	Sumit Roy	CU	08-09-2026
Approval Confirms that the deliverable conforms to the specific WP expectations/requirements	Federico Ribatti	TNO	15-09-2026

Project

Project acronym	H_2UpScale
Project title	High-power hydrogen fuel cell system balance of plant component up-scale and optimization
Project number	101192481
Work programme topic	HORIZON-JTI-CLEANH2-2024-03-02
Work Package	WP3 System Design and Optimisation
Deliverable number	D3.3
Deliverable type	Report
Dissemination level	Public
Due date	30-06-2026
Document date	14-09-2026
Keywords	Fuel Cell; Balance-of-Plant; upscale; system; architecture; design; optimisation

Revision history

Date	**Description**	**Author & organisation**
28-07-2026	First draft	Theo Koch (RWTH)
08-09-2026	Quality Assurance	Sumit Roy (CU)
14-09-2026	Final Draft	Theo Koch (RWTH)

Disclaimer

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Clean Hydrogen Partnership. Neither the European Union nor the granting authority can be held responsible for them.

All rights reserved

No part of this publication may be reproduced and/or published by print, photoprint, microfilm or any other means without the previous written consent of the H2UpScale Consortium.

Acknowledgement

The project is supported by the Clean Hydrogen Partnership and its members.

Table of abbreviations

Abbreviation	Meaning
APP-xxx-R/M/A	Application case: rated power in kW; R road, M marine, A aviation
BOP	Balance of Plant
BP1, BP2	Blueprint 1, Blueprint 2
CAC	Charge-Air Cooler
CAPEX	Capital Expenditure
DC	Direct Current
DC/DC	DC-to-DC converter
Dx.y	Deliverable x.y
EM	Electric Machine
EMS	Energy Management System
FBDF	Fixed-leading-coefficient Backward Differentiation Formula
FC	Fuel Cell
FCEV	Fuel Cell Electric Vehicle
FCS	Fuel Cell System
HT	High Temperature
HX	Heat Exchanger
ID	Identifier
ISA	International Standard Atmosphere
KPI	Key Performance Indicator
LHV	Lower Heating Value
MT	Medium Temperature
NSGA-II	Non-dominated Sorting Genetic Algorithm II
NTU	Number of Transfer Units
ODE	Ordinary Differential Equation
PEM	Proton Exchange Membrane
SciML	Scientific Machine Learning
SOC	State of Charge
UA	Product of overall heat-transfer coefficient and heat-transfer area
VECTO	Vehicle Energy Consumption Calculation Tool
WP	Work Package

Table of Contents

Table of abbreviations 3

Table of Contents 4

1 Introduction 6

1.1 Context 6

1.2 Connection to other WP and inputs used 6

1.3 Structure 7

2 Requirements, available inputs and assumptions 7

2.1 Application case 7

2.2 Duty cycles 8

3 The optimisation framework 11

3.1 Framework overview 12

3.2 Design vector 12

3.3 Fuel-cell system simulation model 15

3.4 Component modelling 16

3.5 Fidelity and symmetry reduction 17

3.6 Optimiser settings 17

3.7 Software implementation and numerical settings 17

3.8 Screening and confirmation 18

4 KPI, constraint and cost definitions 18

4.1 Objectives 18

4.2 Admissibility 19

4.3 Cost boundary 19

5 350 kW road platform (APP-350-R) 19

5.1 Blueprint architecture 19

5.2 Optimisation results 20

6 Architecture scaling method 25

7 700 kW marine platforms (APP-700-M) 26

7.1 Blueprint architectures 26

7.2 Optimisation results 28

8 1050 kW marine platforms (APP-1050-M) 33

8.1 Blueprint architectures 33

8.2 Optimisation results 36

9 1050 kW aviation platform (APP-1050-A) 41

9.1 Blueprint architecture 41

9.2 Optimisation results 42

10 Sensor positions and monitoring rationale 47

11 Comparison with state of the art **48**

12 Reproducibility **49**

13 Conclusions **49**

1 Introduction

1.1 Context

WP2 established what the system has to do. D2.1 defined the application-specific requirements, including the marine duty cycles. D2.2 consolidated the fuel-cell system requirements around a generic high-power stack model and fixed the modular node breakdown of 175 kW stacks. In WP3, D3.1 added a state-of-the-art review and a set of architecture blueprints for three power nodes (350, 700 and 1050 kW), and D3.2 provided should-cost functions for the BOP component families.

This deliverable takes those inputs, together with a dynamic fuel-cell system model built in Julia/SciML, and produces optimised BOP sizing per blueprint with the electrical and system-level evidence that supports it. The toolchain, the solver and its tolerances are stated in Section 3.7, and the availability of the model and data in the Reproducibility section.

The model itself and the generic component data behind it were developed in close cooperation with WP4, WP5, WP6 and WP7: WP4 for system simulation and the stationary and scalable fuel-cell stack and BOP component models, WP5 for the hydrogen-path components, WP6 for the air-path components, and WP7 for the coolant-path components. The simplified system model used here is the optimisation-oriented counterpart to WP4's more detailed simulation. It is fast enough to sit inside an outer optimisation loop that solves thousands of full missions, while remaining consistent with the component behaviour those work packages define.

1.2 Connection to other WP and inputs used

The stack characteristics, the node/module breakdown (2 × 175 kW stacks per 350 kW node) and the loss and efficiency conventions are taken from **D2.2**; the application requirements, performance targets and architecture blueprints from **D3.1**; and the BOP should-cost functions from **D3.2**. The marine duty cycles (700 and 1050 kW) come from the **D2.1** application-specific requirements, while the 350 kW road cycle is derived in WP3 from the **VECTO Long Haul** cycle applied to a 40 t FCEV longitudinal and battery/EMS model (§2.2). Generic component data and characteristics are supplied by the component work packages named above: **WP5** (ejector, recirculation blower), **WP6** (compressor, air filter, humidifier, exhaust/turbine) and **WP7** (heat exchangers, coolant), working with **WP4** on the stationary and scalable stack and BOP component models. The simplified dynamic system model and the optimisation framework around it are built in WP3 on top of these inputs.

Two points follow from this input set and hold throughout the deliverable.

First, **the stack is deliberately generic.** Consistent with D2.2, the project does not target a specific fuel-cell manufacturer; it uses a representative high-power stack so that the BOP components developed here remain compatible with any stack supplier operating within the defined performance envelope. Results are therefore valid as comparative engineering evidence for BOP sizing, and are not predictions of any particular commercial stack's behaviour.

Second, **the component data are generic by type**, consistent with that same principle and developed jointly with the component work packages. Each map is defined at a reference point. Where the optimiser moves a component away from that point (a compressor at a different rating, an electric machine at a different

power), the map is scaled by documented similarity or normalisation rules. Ratings far from their reference point therefore carry additional scaling uncertainty.

1.3 Structure

Section 2 states the requirements, inputs and assumptions the work is built on. Section 3 describes the optimisation framework, how the dynamic fuel-cell system simulation sits inside it, and how the individual components are modelled. Section 4 defines the objectives, admissibility gates and cost boundary that the framework evaluates against.

Section 5 reports the 350 kW road platform. Section 6 states the frozen rules by which that architecture is scaled to higher power nodes, and §7–§9 are the scaled platform chapters, each stating its architecture first and its optimisation results second: §7 the two 700 kW marine blueprints, §8 the two 1050 kW marine blueprints and §9 the 1050 kW aviation blueprint.

Sections 10–13 cover sensor positions and monitoring rationale, comparison with the state of the art, reproducibility, and the conclusions of this draft.

2 Requirements, available inputs and assumptions

2.1 Application case

Six blueprint columns are optimised across three power nodes and three applications. The 350 kW road column is the reference architecture from which the higher nodes are scaled (§6). At each marine node the same stacks are arranged in two competing ways, modular air/H_2 trains against a common rail, so the two columns of a pair can be compared directly.

Table 1. Blueprint overview.

Case ID	Application	Node	Architecture basis
APP-350-R	Heavy-duty road / truck	350 kW	Symmetric twin stack with shared BOP; fixed HT/MT topology
APP-700-M #1	Marine	700 kW	Four stack modules, two DC/DC converters, two modular air/H_2 trains
APP-700-M #2	Marine	700 kW	Same stacks/converters, one larger common-rail air/H_2 train
APP-1050-M #1	Marine	1050 kW	Six stack modules; the 350 kW column scaled linearly to three modular air/H_2 trains, with turbines for exhaust recuperation

Case ID	Application	Node	Architecture basis
APP-1050-M #2	Marine	1050 kW	Same six stacks; common-rail air/H_2 buffers supplied by two trains, sized as one high-flow and one low-flow unit
APP-1050-A	Aviation	1050 kW	Six stack modules, three modular air/H_2 trains

These six columns fix what is built. What each has to deliver over time is set by its duty cycle. The next section states those cycles, their provenance, and the ambient boundary conditions each application is evaluated against.

2.2 Duty cycles

On every platform the optimiser's inner loop and the confirmation run use the same cycle, so screening and confirmation never see different missions. For the 350 and 700 kW platforms that cycle is the full-length reference profile. For the 1050 kW marine platform it is a shortened cycle derived from the supplied profile, for the reason given below.

Table 2. Overview of duty cycles for the different applications.

Cycle	Platform	Duration	Mean load	Source
road_350_reference	APP-350-R	4,503 s	121 kW (35 % of node)	VECTO Long Haul → 40 t FCEV model
maritime_700_reference	APP-700-M	7,080 s	401 kW (57 % of node)	D2.1 marine duty cycle
maritime_1050_reference	APP-1050-M	2,560 s	932 kW (89 % of node)	Derived from the 64,500 s D2.1 marine duty cycle
aviation_1050_reference	APP-1050-A	3,360 s	520 kW (50 % of node)	Derived in WP3 based on given data in D2.1

The road cycle is derived in WP3 from the VECTO Long Haul cycle. The route is passed through a 40 t heavy-duty longitudinal vehicle model to obtain DC-bus demand, and a battery and energy-management layer then splits that demand between the battery and the 350 kW fuel cell under charge-sustaining control. The resulting fuel-cell profile delivers 151.79 kWh over 4,503 s at a 34.7 % capacity factor, with 181 s at full 350 kW. The vehicle and EMS parameters are recorded with the cycle data. The marine cycles come from the D2.1 application-specific requirements. Both the road and the marine cases are run against the same boundary conditions: ambient air at 101.3 kPa and 298.15 K, with liquid-to-air heat rejection to a 298.15 K sink. Air

rejection is used for the marine columns as well, by project agreement, to increase modularity in radiator pack design across blueprints.

The supplied 1050 kW marine profile is the outlier in this set. It delivers 16,963 kWh over 17.9 h at a 90.2 % capacity factor, and sits at or above 99 % of nameplate for 51,791 s, which is 80 % of the mission. The structure is a sustained 1050 kW cruise broken by periodic shallow dips to ~700 kW and two deep manoeuvring or port excursions to ~285 kW. Where the road cycle is dominated by part load and the 700 kW marine cycle by a mid-band cruise, this one is close to a continuous rating case: the BOP is sized by sustained full-power operation, and thermal capacity has almost no low-load recovery periods to fall back on.

At 64,500 s that profile is roughly nine times the 700 kW mission, and a single blueprint column costs 4,100 evaluations (§3.6), so integrating it directly is not affordable. maritime_1050_reference is therefore derived from it: 2,560 s that keep the true 10 kW/s ramp rates, two shallow dips and one deep dip, reproducing the parent's 1050 kW peak and its ~90 % mean load at roughly 25× lower cost. This one cycle is used for both screening and confirmation, so the 1050 kW columns are treated the same way as the other four. The shortened derived cycle is long enough to cover potential long-term thermal drifts of the model under high steady loads.

The 1050 kW aviation profile is derived in WP3 based on the data given in D2.1. It is a 3,360 s single sortie in seven phases: taxi-out at 85 kW, a 1,050 kW take-off, an 860 s climb from 300 m to 7,620 m over which power tapers from 950 kW to 620 kW as inlet density falls, a 1,200 s cruise at 620 kW, a descent at 180 kW, approach and landing at 350 kW, and taxi-in at 85 kW. It delivers 485.4 kWh at a mean 520 kW, 49.5 % of nameplate, with 64 s at or above 99 % of nameplate. That places it between the road cycle at 34.7 % and the 700 kW marine cycle at 57 %, and well away from the continuous-rating character of the 1050 kW marine cycle.

It is also the only cycle in the set to carry a second trace. Altitude is integrated with the power profile rather than applied afterwards as a corner check, because the aviation case does not share the sea-level boundary conditions of the road and marine cases: its normal operating environment is cruise at 7,620 m (25,000 ft), where the ISA ambient is 37.6 kPa, about 37 % of sea-level pressure, and 238.7 K (−34.5 °C), and where the corrected air flow the compressor has to deliver rises by a factor of 2.45 against its sea-level value. That is a change of operating point rather than an offset on a sea-level mission, and it moves cycle efficiency, which is an optimiser objective.

All four cycles are shown below in the form the optimiser integrates them. Panel (a) of each figure is the time trace and panel (b) the load-duration curve, which plots power sorted descending against the cumulative share of the mission spent at or above it. The two panels answer different sizing questions. The trace shows *when* the load moves, which the thermal and air-path transients respond to. The duration curve shows *how long* the system sits at each load, which governs the ratings that survive screening.

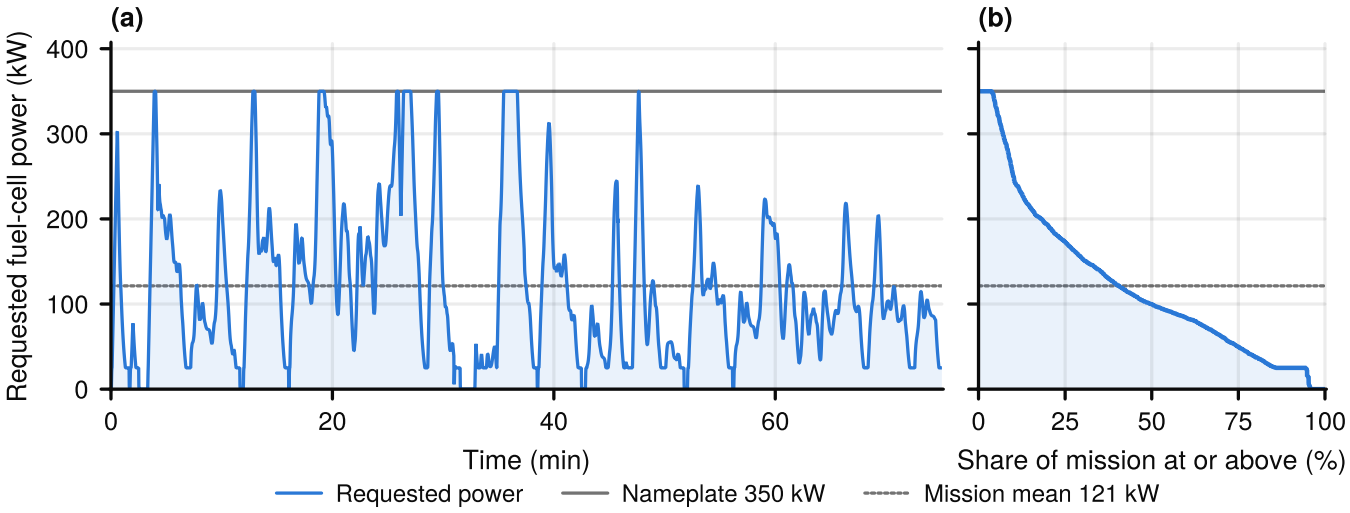


Figure 1. Duty cycle for the application APP-350-R.

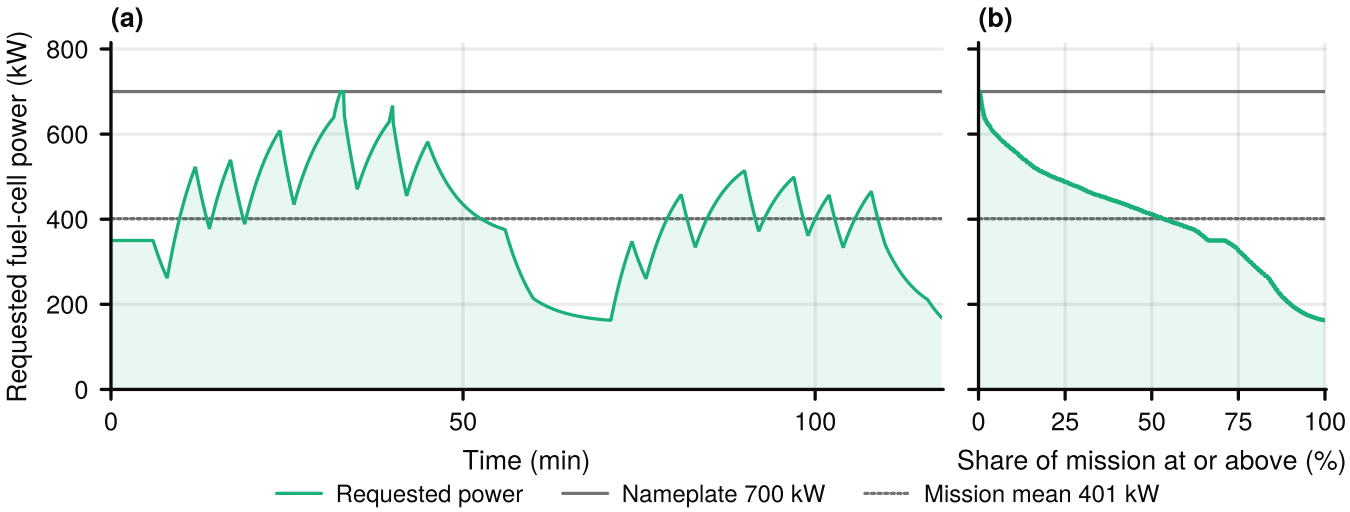


Figure 2. Duty cycle for the application APP-700-M.

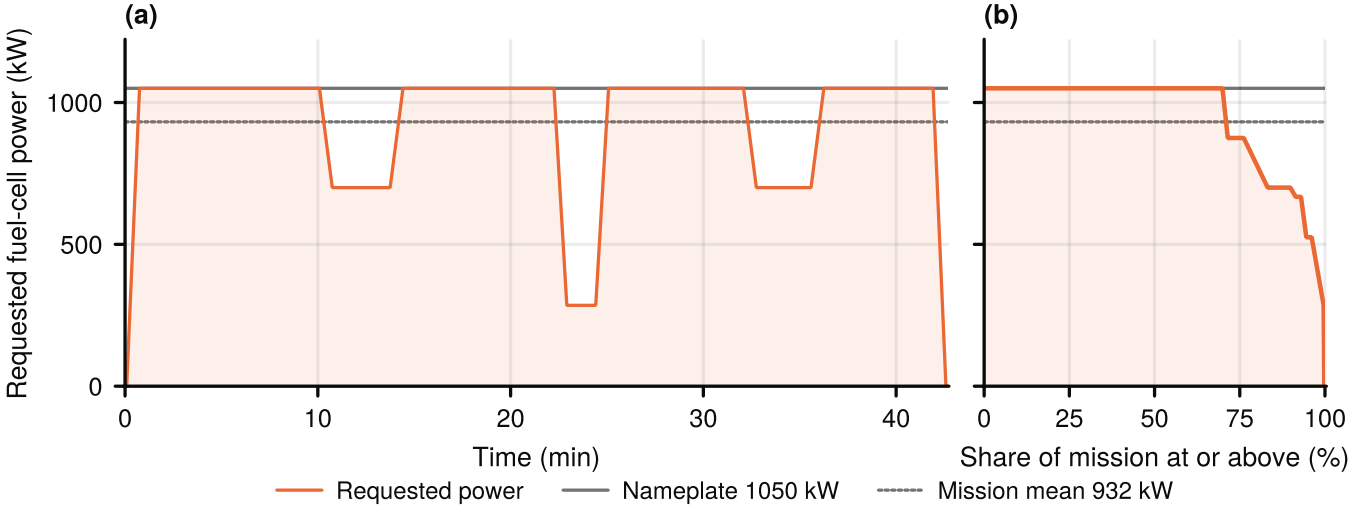


Figure 3. Duty cycle for the application APP-1050-M.

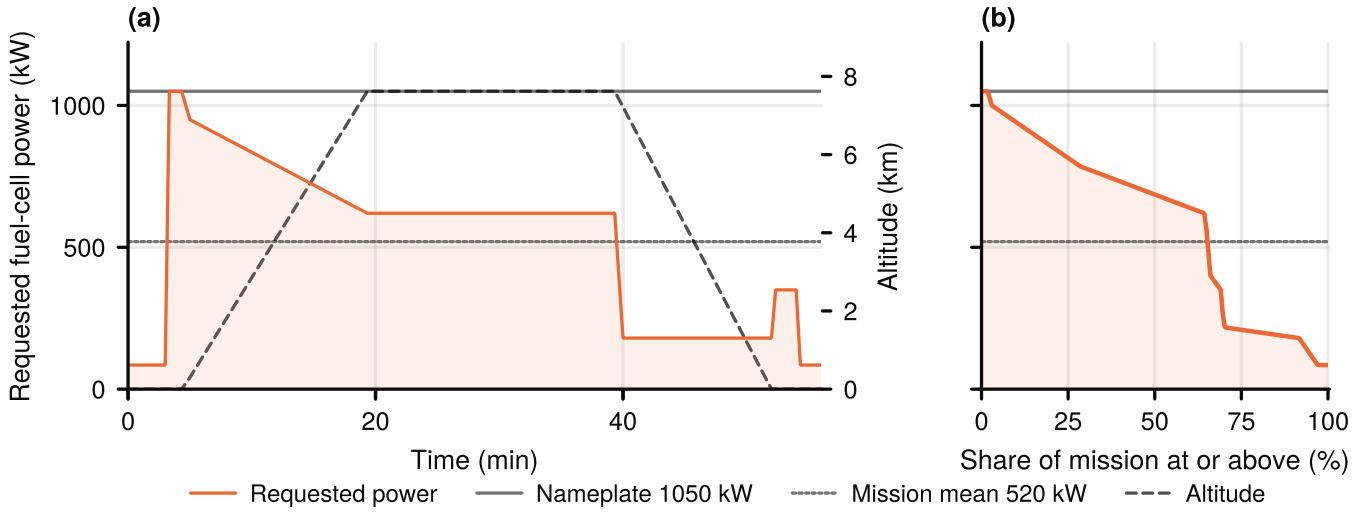


Figure 4. Duty cycle for the application APP-1050-A.

All four profiles represent battery-hybridised fuel-cell power demand, and each therefore embeds a specific hybridisation and ramp-limiting strategy. The sizing reported in this deliverable is conditioned on that strategy: a different battery size, SOC window or ramp policy would present the fuel cell with a different load signature and could change the optimised BOP sizing.

3 The optimisation framework

This section defines the framework that produces every reported number: what the optimiser controls (§3.2), how a dynamic mission simulation turns a single candidate design into evidence (§3.3), how the underlying physical components are modelled (§3.4), what reductions are applied to make the loop affordable (§3.5), and how the search and its confirmation stage are configured (§3.6–§3.8).

3.1 Framework overview

The framework is a closed loop between a multi-objective optimiser and a dynamic fuel-cell system simulation. The optimiser (NSGA-II) proposes a candidate design as eleven numbers: nine component-sizing variables and two operating set-points (step 1, Figure 5). Together with the derived BOP ratings defined in §3.2, these values build a dynamic plant model (step 2), and that model is run through the complete duty cycle (step 3). The resulting cycle efficiency, should-cost and constraint outcomes go back to the optimiser, which uses them to propose the next generation of designs (step 4).

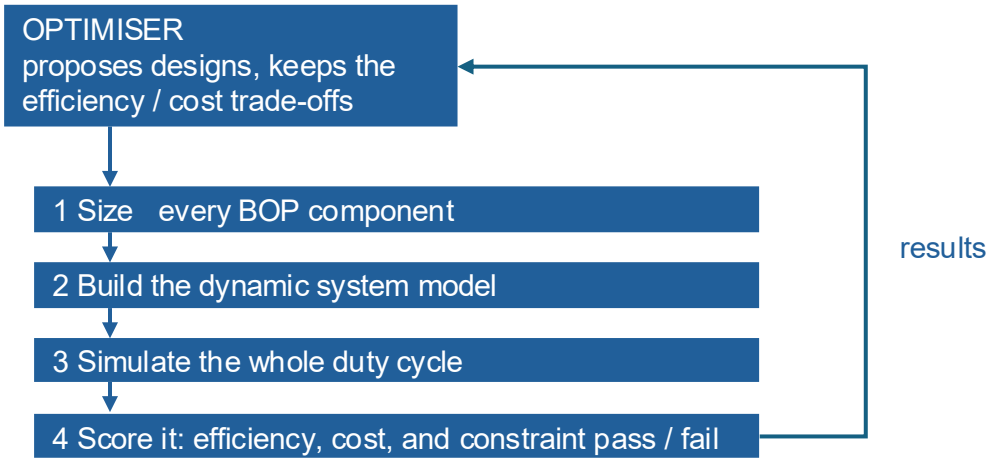


Figure 5. Optimisation framework overview.

3.2 Design vector

Eleven continuous variables are searched, identically on every platform except APP-1050-A, which adds a twelfth (§9.2):

Table 3. Design variables.

Variable	Meaning	Type	Domain
comp_kW	Compressor rating (per air train)	sizing	energetic
comp_pr	Compressor design pressure ratio	sizing	energetic
turb_kW	Backpressure turbine rating	sizing	energetic
pump_HT_kW	HT coolant pump rating	sizing	thermal (D3.4)

Variable	Meaning	Type	Domain
pump_MT_kW	MT coolant pump rating, per branch	sizing	thermal (D3.4)
fan_kW	Aggregate radiator-fan pack rating	sizing	thermal (D3.4)
blow_kW	Anode recirculation blower rating	sizing	energetic
stoich	Cathode stoichiometry target	operating set-point	energetic; moves thermal load
bp_kPa	Cathode backpressure target	operating set-point	energetic; moves thermal load
radHT_kW	HT radiator rating	sizing	thermal (D3.4)
radMT_kW	MT radiator rating	sizing	thermal (D3.4)

Nine are component ratings. The remaining two, *stoich* and *bp_kPa*, are operating set-points, which makes this a combined sizing and operating co-design. The domain column records whether the variable belongs to the electric or thermal system (D3.4) domain. Five variables size the thermal system and are analysed in D3.4, and the two set-points act on both domains at once. All eleven are searched together in one optimisation.

The eleven variables do not enumerate the bill of material. They fix the ratings of the components they name. Every remaining BOP item (filter, humidifier, separator, resonator, ejector, hydrogen heat exchanger, charge-air cooler and the air-handling electric machine) is derived instead of searched: its rating follows by frozen rules from the flows and duties that the searched ratings and set-points produce over the mission. The two groups are therefore coupled. Because *stoich* and *bp_kPa* change how the plant is run, they move the derived ratings as well as the searched ones.

Both radiator ratings drive loop UA as well as cost, and the two coolant pumps size their loops' flow independently. Design bounds are physical and headroom ranges chosen independently of the D3.2 cost-node spans. Where an optimised rating falls outside a tabulated D3.2 span, its cost is read by extrapolating the D3.2 law.

The search box at the 350 kW node is given below. These are the bounds the optimiser samples within.

Table 4. Design-variable search bounds at the 350 kW node.

Variable	Lower	Upper	Unit
comp_kW	40	75	kW
comp_pr	1.5	2.6	-
turb_kW	8	28	kW

Variable	Lower	Upper	Unit
pump_HT_kW	0.5	4.0	kW
pump_MT_kW	0.05	0.5	kW
fan_kW	10	50	kW
blow_kW	0.5	3.5	kW
stoich	1.5	2.4	-
bp_kPa	60	130	kPa
radHT_kW	300	500	kW
radMT_kW	34	130	kW

The higher-power columns do not re-derive these numbers, they scale them. Per-train variables (comp_kW, turb_kW, blow_kW) scale with the per-train power of the blueprint, so a common-rail column feeding two electrical groups from one air train carries twice the box of a modular column at the same node. System-level variables (pump_HT_kW, pump_MT_kW, fan_kW, radHT_kW, radMT_kW) scale with node power. The three intensive variables (comp_pr, stoich, bp_kPa) do not scale at all. For symmetric columns, two factors set the scaled bounds: s_sys equals P_node divided by 350 kW for the system-level variables, and s_air = s_sys divided by the number of air trains for the per-train variables. The asymmetric APP-1050-M #2 column is treated separately below.

Table 5. Design-bound scale factors per blueprint column.

Column	Air trains	s_air	s_sys
APP-350-R	1	1.0	1.0
APP-700-M #1 (modular)	2	1.0	2.0
APP-700-M #2 (common rail)	1	2.0	2.0
APP-1050-M #1 (linear)	3	1.0	3.0
APP-1050-M #2 (common rail)	2	1.5	3.0
APP-1050-A	3	1.0	3.0

Four of the six columns have s_air = 1.0, so their per-train box is the 350 kW box unchanged, and only the two common-rail columns differ. For APP-1050-M #2, the tabulated s_air = 1.5 is the nominal mean per-train

scaling implied by two trains at a 3× node. One per-train variable is searched at that scaling; the asymmetry is applied afterwards as a declared 2:1 ratio, so the high-flow train carries 4/3 of the searched rating and the low-flow train 2/3, leaving the installed total unchanged.

3.3 Fuel-cell system simulation model

The modelled system comprises the cathode air path (filter, compressor, charge-air cooler, humidifier, stack cathode, water separator, turbine); the anode path (pressure regulator and pressure-control valve, hydrogen heat exchanger, passive ejector with assist blower, anode volume, separator, intermittent purge); the stack; the DC/DC conversion stage; and two coolant loops (HT and MT) with radiators, pumps and one aggregate fan pack.

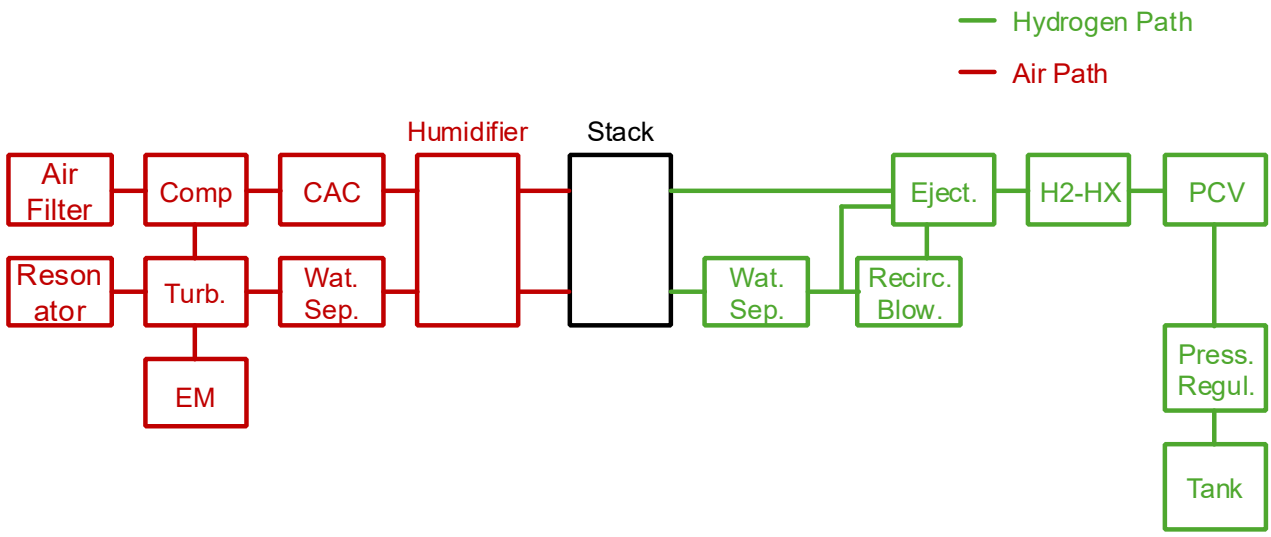


Figure 6. Fuel-cell system simulation model overview.

The plant is a differential-algebraic-free ODE system of 66 states, structured as:

Table 6. ODE system and its states.

Block	Count	States each	Content
Gas plena (energy-tracking)	10	5	4 species masses + internal energy
Stack anode / cathode volumes	2	4	4 species masses, isothermal
Liquid-water inventories	4	1	separator, cathode, anode, recirculation
Anode water ledger	1	1	signed running balance
Purge-cycle controller	1	1	hybrid controller phase

Block	Count	States each	Content
Coolant temperatures	2	1	HT and MT lumped nodes

The ten gas plena are the hydrogen tank, regulator plenum, pressure-control-valve plenum and ejector motive plenum on the anode supply; the filter and compressor plena, supply and return manifolds on the cathode; and the anode recirculation manifold and blower plenum. Every gas volume carries a four-species mixture (O_2, N_2, H_2O, H_2) as an ideal-gas mixture with species mass and energy conserved.

The exogenous command is DC/DC output power taken from the duty cycle, mirroring real supervisory control. A power-following cascade resolves it: an inner current loop solves, at each state, for the stack current whose DC/DC output equals the request by bracketed bisection on [0, *I_ceiling*], with the design current as the anti-windup cap. Compressor speed is a feed-forward schedule from the O_2 stoichiometry requirement. The anode pressure controller, thermal controllers and intermittent purge close around that.

3.4 Component modelling

Components are represented either by physics-based models or by characteristic maps. The physical models share a small set of governing correlations, which keeps behaviour consistent across the plant. Flow restrictions throughout the system (valves, the pressure-control valve, purge orifice, turbine and ejector primary nozzle) use the same compressible-orifice formulation, with component-specific geometry and discharge characteristics.

The fuel-cell stack calculates cell voltage from the Nernst potential, accounting for activation (Tafel), ohmic and concentration losses. Membrane conductivity follows the Springer water-content correlation. Concentration losses come from mass-transfer limits set by binary diffusion through the gas-diffusion layer. The stack is isothermal, which supports system-level heat-duty and component-sizing analyses but not spatial temperature distributions or hotspot predictions (§13).

A scaled manufacturer map gives compressor speed, pressure ratio, mass flow and efficiency. The turbine combines a compressible-flow restriction with an isentropic-efficiency model to determine mass flow and recovered power. The ejector is passive and stateless: its primary nozzle meters fresh H_2 from the prevailing upstream and downstream conditions, and primary-jet momentum generates the mixing-section pressure reduction that entrains recirculated anode gas.

The humidifier uses a two-sided membrane model with an NTU-effectiveness formulation, transferring water vapour and sensible enthalpy from the wet exhaust to the dry supply stream. Radiators and the charge-air cooler use cross-flow NTU-effectiveness models, so their ratings set heat-transfer conductance instead of imposing a fixed heat-duty limit. The water separator assumes equilibrium condensation to saturation at the outlet conditions, and removes liquid water at a prescribed separation efficiency.

The DC/DC converter and the recirculation blower are treated as regulated devices with constant or rating-scaled efficiencies. Coolant pump power follows from the required hydraulic duty and a prescribed efficiency. Radiator-fan power is calculated from the required airflow and the cooling-pack pressure drop, using prescribed aerodynamic and motor/inverter efficiencies. These reduced-order representations are deliberate. At system level what matters is how much parasitic power these devices draw and what that does to net efficiency, not how their internal losses arise.

3.5 Fidelity and symmetry reduction

The physical 350 kW architecture is a symmetric twin of two 500-cell modules (1,000 cells total, 330 cm^2 active area, 660 A ceiling) with shared BOP, that is, the two 175 kW modules per 350 kW node defined in D2.2. Production optimisation and confirmation use the behaviourally equivalent 1 × 1000-cell single-stack reduction, whose aggregate behaviour agrees with the explicit twin to within 0.01 %. The modular 700 kW blueprint is evaluated through one symmetric 350 kW representative electrical group at half system request, with repeated-group extensive quantities doubled.

These reductions are valid for aggregate, symmetric studies. They are not valid for module mismatch, flow maldistribution, isolation, degradation, or fault behaviour, and no such claim is made.

3.6 Optimiser settings

Multi-objective NSGA-II uses deterministic seeds and a population of 100. The standard higher-power runs use 20 generations. The implementation produces 2N offspring per generation, so the evaluation budget is

100 (initial population) + 20 × 200 (offspring) = 4,100 evaluations

per blueprint column. APP-350-R used an extended 40-generation budget, giving 8,100 evaluations (§5.2.1). The standard initial population is sampled uniformly across the design bounds; APP-700-M #1 is the stated exception and was warm-started from an earlier archive, whereas APP-700-M #2 used the standard initialization. Evaluation is threaded and deterministic; the parallel path reproduces the serial Pareto front exactly.

3.7 Software implementation and numerical settings

The system model, the optimisation driver and the evaluation scripts are written in Julia. Every mission is solved with the FBDF implicit integrator using finite-difference Jacobians, at a relative tolerance of 1×10^{-6}, absolute tolerances scaled per state group, and a ceiling of 10^6 steps. FBDF was selected after a cross-solver qualification that ran the same configurations under the explicit Tsit5 and the stiff Rodas5P integrators.

Table 7. Software used to produce the reported results.

Package	Version	Role
Julia	1.12.6	Language and runtime
OrdinaryDiffEqBDF	1.26.0	FBDF integrator; solves every mission
OrdinaryDiffEqTsit5	1.12.0	Explicit integrator; solver qualification only
OrdinaryDiffEqRosenbrock	1.31.1	Rodas5P; second qualification solver
SciMLBase	2.155.1	Problem and solution interfaces
Metaheuristics	3.5.0	NSGA-II multi-objective search

Package	Version	Role
ADTypes	1.22.2	Finite-difference Jacobian selection

These are the versions the reported results were produced with. What each retained result set records, and the availability of the model, the data and the scripts, are stated in the Reproducibility section.

3.8 Screening and confirmation

Screening produces a Pareto set per column. Three report roles are then selected deterministically from the admissible Pareto pool (*low-cost*, *high-efficiency*, and an equal-weight *compromise*) and re-evaluated at finer sampling on the same mission. Final report-role KPIs and blueprint comparisons come from this confirmation stage. Screening statistics are reported separately and explicitly labelled as screening results.

The three roles do not rank the designs. They are the two bookends of the trade-off surface plus its centre, retained so that the shape of the surface stays visible instead of collapsing into a single recommendation.

4 KPI, constraint and cost definitions

4.1 Objectives

The search optimises two objectives:

Table 8. Optimiser objectives.

Objective	Definition	Direction
Cycle efficiency η_cycle	Net energy ÷ fresh-H_2 LHV energy over the cycle	maximise
BOP should-cost C	Sum of the D3.2 component cost lines at the design's actual ratings	minimise

Tracking is a constraint, not an objective. Earlier revisions of this work minimised the tracking shortfall *J_track* as a third objective. That did not work. The optimised designs meet the duty essentially exactly, so *J_track* settles several orders of magnitude below its own admissibility threshold (~10^{-8} against a normalised threshold of 1). NSGA-II then preserved Pareto membership on differences that were numerically meaningless, inflating the front with designs that differed only in solver noise. At 700 kW this collapsed all three report roles onto a single row: the genuine cost/efficiency surface held only one design while the front carried 10–14 points.

J_track is therefore now enforced through the admissibility gate of §4.2, which already contained the three tracking checks (shortfall, maximum deficit, maximum overshoot). A design that fails any of them is dominated. The tracking shortfall is still recorded for every evaluation as a diagnostic; it no longer competes with cost and efficiency for front membership. The reported Pareto set is consequently a two-dimensional cost/efficiency trade-off surface.

4.2 Admissibility

A design is report-admissible only if it passes every gate:

- solver success with finite states and KPIs;
- inventory and thermodynamic domain validity;
- minimum loaded cell voltage and maximum current density;
- tracking within threshold — shortfall, maximum deficit and maximum overshoot against the requested DC/DC output (§4.1);
- compressor and EM map-envelope feasibility;
- commanded pump/fan/blower peak draw within design rating;
- HT and MT loop temperatures within their assumed limits (the limits, their basis and the margins achieved: described in D3.4);
- component rating utilisation ≤ 1.00, for components whose rating represents a physical capacity gate

Thresholds are engineering policy. "Admissible" means the design passes this comparative screen.

4.3 Cost boundary

Cost is the D3.2 BOP should-cost only. It excludes the stack, hydrogen storage, enclosure, wiring, controls, integration labour and manufacturing scale-up. Deliverable-level CAPEX targets from D3.1 Table 5 are *not* comparable with these subtotals, and no such comparison is made.

5 350 kW road platform (APP-350-R)

5.1 Blueprint architecture

Two symmetric 500-cell stack modules (1,000 cells, 330 cm² active area) share one BOP set: a single air train, a single anode train, one DC/DC converter, and one HT and one MT coolant loop rejecting through a shared liquid-to-air radiator pack. The architecture and its reporting boundary are Figure 6 (§3.3); it is the source architecture from which the 700 kW and 1050 kW blueprints of §7–§9 are built under the count rules of §6.

5.2 Optimisation results

5.2.1 Screening outcome

Screening on road_350_reference evaluated 8,100 designs. 5,292 were report-admissible (65.3 %), the retained Pareto set holds 20 points, and 18 evaluations failed in the solver. A further 214 evaluations reached the per-evaluation wall-clock guard and were scored as failures. The admissible design space spans €3,554–6,818 in BOP should-cost and 0.197–0.599 in cycle efficiency.

The retained Pareto set is degenerate. All 20 points carry the same air path to within rounding: comp_kW 42.4–42.5 kW, comp_pr 2.49–2.51, stoich 1.63–1.66, and bp_kPa at the 130 kPa(g) ceiling on every point. Only the cooling hardware varies, with radHT_kW between 366 and 470 kW and fan_kW between 16.4 and 20.6 kW. Efficiency is therefore flat across the whole cost range: binning the admissible designs by cost gives a best efficiency of 0.5978 at €3,500–3,750 and 0.5971 at €4,000–4,250, a spread of 0.18 % against a 21 % spread in cost.

This is a property of the platform rather than of the search. Cathode backpressure sets cell voltage through the Nernst term, and bp_kPa is capped at 130 kPa(g) by the D2.2 stack operating envelope (260 kPa(a) cathode inlet, 280 kPa(a) anode inlet); the confirmed designs reach 242.5 kPa(a) at the cathode with the anode about 30 kPa above. Compressor pressure ratio, the other lever on efficiency, is capped at 2.6. At 2.9 the stack pins the 2.00 A/cm² current-density gate and the compressor map check fails. With both levers at their limits the optimiser can no longer buy efficiency, so selection pressure falls entirely on cost. The 350 kW road optimum is cathode-pressure limited. No efficiency–cost trade-off curve exists for this platform, and none is reported.

Table 9. APP-350-R confirmed report roles against the baseline design.

Role	BOP cost (€)	η_cycle	Peak HT (K)	Peak MT (K)	Admissible
Baseline	5,092.53	0.5775	355.47	328.91	yes
Low-cost	3,689.64	0.5970	355.40	328.93	yes
Compromise	3,700.35	0.5974	355.51	328.87	yes
High-efficiency	3,774.73	0.5983	355.63	328.80	yes

Confirmation solved the full 4,503 s reference cycle uncapped at 4×10^6 iterations, a more permissive policy than screening, and it reproduced the screening result. The screening archive's best efficiency was 0.5982 at €3,852, while the confirmed high-efficiency role returns 0.5983 at €3,775; this close agreement does not indicate a favourable screening bias for the road case. All three optimised roles pass every gate.

The three roles span €3,690–3,775 and 0.5970–0.5983, so the last €85 of BOP cost buys 0.13 efficiency points. All three are cheaper and more efficient than the baseline design: +3.6 % efficiency for −27.5 % cost at the low-cost role.

The optimised designs carry almost no margin. Every sized component on the low-cost design runs between 93 % and 100 % of its rating — EM inverter 0.998, radiator fan 0.988, compressor 0.962, HT radiator 0.953 —

and five design variables sit at or next to a bound. This follows from asking for the cheapest admissible design: the optimiser buys exactly the duty and nothing more. A design at 99.8 % of its inverter rating will not survive a supplier derating or a change to the duty cycle. The compromise and high-efficiency roles cost 0.3–2.3 % more and leave usable radiator margin. The compromise design is used as the reported baseline below because it adds margin at only 0.3 % more BOP cost than the low-cost role.

5.2.2 Design vector of the report roles

Table 10. Design vectors of the confirmed APP-350-R report roles.

Variable	Low-cost	Compromise	High-efficiency	Unit
comp_kW	42.462	42.509	42.460	kW
comp_pr	2.489	2.490	2.513	—
turb_kW	19.630	19.655	19.393	kW
pump_HT_kW	0.546	0.628	0.523	kW
pump_MT_kW	0.051	0.061	0.053	kW
fan_kW	20.578	17.556	16.717	kW
blow_kW	0.500	0.506	0.505	kW
stoich	1.643	1.626	1.630	—
bp_kPa	129.994	129.713	130.000	kPa gauge
radHT_kW	365.701	396.729	470.154	kW
radMT_kW	35.228	34.577	40.519	kW

The three designs share one machinery configuration. Compressor power, pressure ratio, turbine, blower and stoichiometry agree to within a few percent, and all three sit at the cathode backpressure ceiling. They separate on cooling: the HT radiator spans 366–470 kW and the fan 16.7–20.6 kW. Buying more radiator does not raise efficiency here, it buys hot-ambient margin, so the cost difference between the roles purchases robustness rather than performance.

Six of the eleven variables sit at or beside a bound. bp_kPa, blow_kW and pump_MT_kW are pinned; pump_HT_kW, stoich and comp_kW are within a few percent of their floors. Only bp_kPa is binding in the sense that matters: it caps efficiency, and it is set by the D2.2 stack envelope. The floors on the blower and the pumps are reached because the optimiser is minimising cost and those components have no cheaper setting, not because the design wants to go below them.

5.2.3 Component sizing, utilisation and cost

The inventory below is the compromise design, recommended above as the reported baseline. Ratings are the purchased values, and the platform carries one unit of every line apart from the two stack modules. The radiator, fan, coolant-pump, charge-air-cooler and hydrogen-heat-exchanger lines appear here because the cost objective is evaluated over one bill of material; their thermal sizing basis, duties and capacity margins are reported in D3.4.

Table 11. APP-350-R BOP inventory of the confirmed compromise design.

Component	Rating	Unit	Utilisation (%)
PEM stack module	500	cells/module	—
Compressor	42.51	kW	95.5
Turbine	19.65	kW	92.6
Air-handling EM + inverter	25.14	kW	99.3
HT coolant pump	0.628	kW	62.1
MT coolant pump	0.061	kW	31.1
Radiator fan pack	17.56	kW	99.1 (A)
Recirculation blower	0.506	kW	100.0
HT radiator	396.73	kW	87.5
MT radiator	34.58	kW	30.6
Charge-air cooler	66.32	kW	31.5
Water separator	320.04	g/s	100.0 (B)
Cathode air filter	320.04	g/s	100.0 (B)
Resonator	320.04	g/s	100.0 (B)
Humidifier	320.04	g/s	100.0 (B)
H_2 ejector	8.00	g/s	79.6
H_2 heat exchanger	0.50	m^2	(318) (C)

Component	Rating	Unit	Utilisation (%)
Total			

Three utilisation figures in this table are conventions rather than sizing results. A: The radiator fan is rated against the 40 °C hot-ambient sizing draw and not against the mission peak: the mission runs at a 25 °C sink with highway ram air, and its peak fan draw is roughly a third of what a stopped 40 °C condition demands. B: The water separator, air filter, resonator and humidifier are sized to the peak flow they pass, which makes their utilisation structural. C: The hydrogen heat exchanger is sized for thermal adequacy. The 0.5 m² unit brings fresh hydrogen to within about 2.4 K of the HT coolant, and the bracketed figure comes from a flow-area correlation that does not apply to that sizing basis, so it is a diagnostic and not a capacity gate.

Two lines are closely sized and are the ones to watch under any assumption change: the air-handling machine and inverter at 99.3 % utilisation, and the HT radiator at 87.5 %. The two coolant pumps sit at 62.1 % and 31.1 %, but those peaks come from a 35 %-mean-load profile against a 25 °C sink and do not show that the pumps are oversized for a hotter or heavier duty.

5.2.4 Cycle KPIs and parasitics

Table 12. APP-350-R confirmed cycle KPIs over the full road mission.

KPI	Baseline	Low-cost	Compromise	High-efficiency
Requested net-bus energy (kWh)	151.79	151.79	151.79	151.79
Delivered net-bus energy (kWh)	151.78	151.78	151.79	151.79
DC/DC output (kWh)	168.05	165.50	165.39	165.41
Parasitic energy, DC/DC output less net bus (kWh)	16.27	13.72	13.60	13.62
Parasitic share of DC/DC output (%)	9.68	8.29	8.23	8.24
Fresh H_2 (kg)	7.883	7.625	7.620	7.609
H_2 utilisation (%)	99.887	99.894	99.916	99.926
Purge loss over fresh H_2 (%)	0.034	0.087	0.085	0.085
Cycle efficiency (%)	57.75	59.70	59.74	59.83
Min loaded cell voltage (V)	0.660	0.678	0.679	0.684
Max current density (A/cm²)	1.807	1.740	1.734	1.718

KPI	Baseline	Low-cost	Compromise	High-efficiency
Min loaded cathode stoichiometry	1.805	1.561	1.556	1.588

The optimised designs gain efficiency in two ways. Parasitic energy falls from 16.27 kWh on the baseline design to 13.60–13.72 kWh, or from 9.7 % to 8.2 % of DC/DC output, mainly because the air-handling machine is sized to the net shaft demand and the coolant pumps run far below the baseline rating. Fresh hydrogen falls from 7.883 kg to 7.609–7.625 kg for the same energy delivered to the bus.

Fresh hydrogen and passive ejector recirculation are credited first in the anode balance, and the blower supplies the remaining stoichiometry deficit, so both contribute through most loaded operation. Purge is normally closed and intermittent, and its hydrogen loss is included in the utilisation figures above. The model reports purge loss but does not predict nitrogen crossover concentration, because no validated crossover correlation exists for this stack.

5.2.5 Comparison with the D3.1 truck cruise target

D3.1, through D2.1, sets a 55 % efficiency target at approximately 51 % load for truck applications. The road profile spends 520 s in the 45–57 % load band at a mean request of 179.1 kW, or 51.2 % of the node, very nearly the target point. Over that band the reported design delivers 63.74 % net-bus over LHV efficiency, integrated from the retained cycle-response trace of the confirmed compromise design

6 Architecture scaling method

The 350 kW platform of §5 is the source architecture. The higher power nodes of §7–§9 are built by repeating source equipment behind shared infrastructure, under frozen rules that keep counts auditable:

- One authoritative air-train count drives the filter, compressor, compressor motor/inverter, charge-air cooler, humidifier, separator, resonator, turbine, ejector and hydrogen heat exchanger. Installed stages cannot carry inconsistent counts within one blueprint. (See Table 7 in D3.1)
- DC/DC-converter count is independent of air-train count. Stack modules are grouped in series behind each converter; converter outputs are paralleled at the system DC bus.
- HT source count is the physical stack-branch count plus one paired H_2-HX → CAC coolant branch per air train, all merging before one shared radiator.
- The MT main path splits through X parallel DC/DC cold plates, merges and mixes, then splits through Y parallel compressor-EM cold plates and merges again, where X is the converter count and Y the air-train count. The aggregate radiator-fan drive sits on its own parallel branch.
- Each system has exactly one functional HT radiator, one functional MT radiator and one aggregate fan pack. Radiators may vary in frontal area or thickness. Multiple physical cores or fans inside those assemblies do not change the architecture count.

These rules are what make the blueprint columns of §7–§9 comparable to each other and to the 350 kW column: every column is the same component set at a different, explicitly recorded count

Extension to multi-megawatt systems

The scaling rules above can be extended to fuel-cell systems up to approximately 6 MW without changing the 175 kW stack module that forms the basis of the present architectures. Higher system power is obtained primarily by increasing the number of stack modules and electrical groups rather than by increasing the power of the individual stack. In the present framework, two 175 kW stack modules form one 350 kW power unit; repeating 17 such units would therefore provide 5.95 MW from 34 stack modules.

At multi-megawatt scale, the architecture becomes hierarchical. Stack modules are grouped into repeated power units with their associated DC/DC conversion and air and hydrogen handling, while selected infrastructure can be shared at system level. Air, hydrogen and cooling equipment need not be consolidated into single 6 MW components; several parallel BOP trains can instead serve groups of stack modules and connect to common electrical, hydrogen and thermal interfaces. The existing scaling rules therefore remain applicable, but the number of repeated functional groups increases with system power.

This extension represents an architectural extrapolation of the present 350–1050 kW framework. Introduction of a higher-power stack module would constitute a new reference architecture and would require the stack and component scaling assumptions to be redefined rather than simply extending the present count-based rules.

7 700 kW marine platforms (APP-700-M)

7.1 Blueprint architectures

Both 700 kW blueprints are built from the 350 kW source architecture of §5 under the count rules of §6. They share four 500-cell stack modules (2,000 cells, 330 cm² active area) arranged in two electrical groups, two DC/DC converters paralleled at the system DC bus, one shared HT radiator, one shared MT radiator and one aggregate liquid-to-air fan pack. They differ only in the air and hydrogen path.

Blueprint 1 — modular (D3.1 700 kW #1): two independent air trains and two hydrogen trains, one per electrical group. Because each group has its own air train, this blueprint may de-stage to one running group below a 320 kW system request.

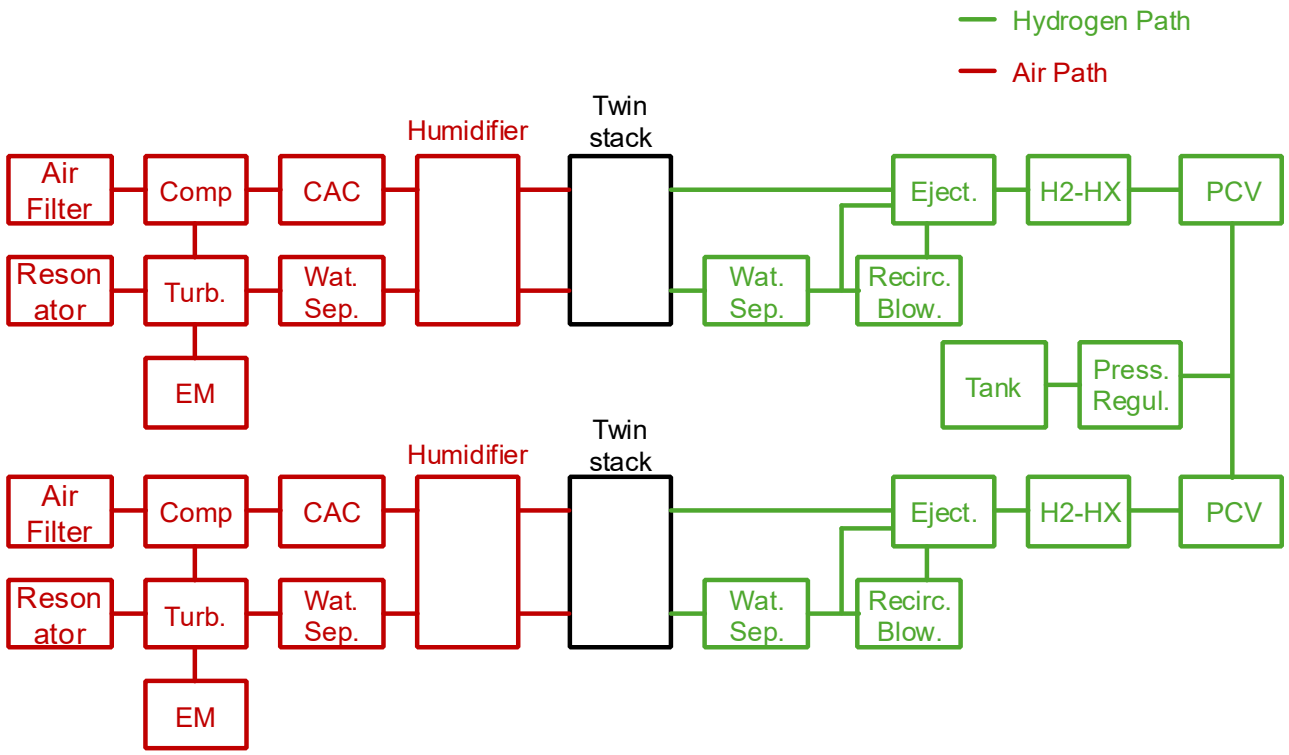


Figure 7. Fuel-cell system simulation model overview for blueprint 700 kW #1.

Blueprint 2 — common rail (D3.1 700 kW #2): one larger air train and one hydrogen train feeding both groups through common rails with buffer volumes. It cannot de-stage; both groups always run.

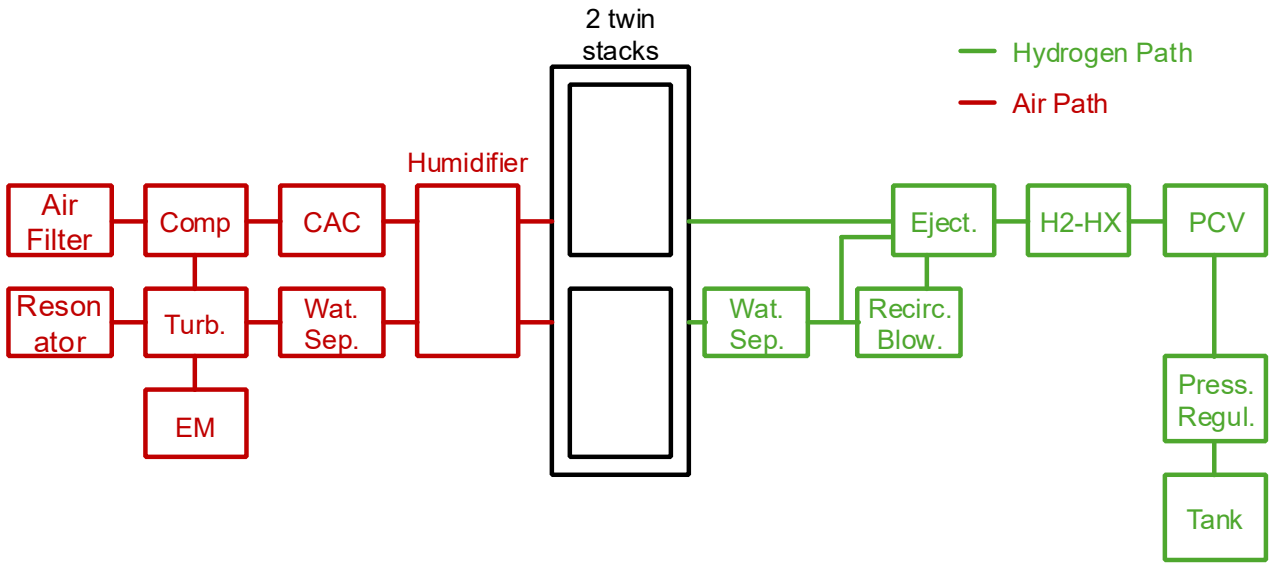


Figure 8. Fuel-cell system simulation model overview for blueprint 700 kW #2.

Table 13. 700 kW blueprint comparison.

Architecture count	Blueprint 1 — modular	Blueprint 2 — common rail
Stack modules	4	4
Electrical groups	2	2
DC/DC converters	2	2
Air trains (filter, compressor, EM/inverter, CAC, humidifier, separator, resonator, turbine)	2	1
Hydrogen trains (ejector, H_2 heat exchanger, recirculation blower)	2	1
Common air / H_2 rails with buffer volume	no	yes
HT radiator / MT radiator / aggregate fan pack	1 / 1 / 1	1 / 1 / 1
HT source branches (stack branches + one H_2-HX → CAC branch per air train)	4 + 2	4 + 1
De-stage capability	yes, below 320 kW	no

The comparison the deliverable is set up to make on this node is therefore narrow and clean: two air trains sized for half flow each, against one air train sized for full flow, on identical stacks, electrical architecture and heat-rejection architecture.

7.2 Optimisation results

7.2.1 Screening outcome

Screening the modular blueprint on maritime_700_reference evaluated 4,100 designs: 2,070 report-admissible (50.5 %), 12 Pareto points and 21 solver failures (0.5 %). A further 378 evaluations (9.2 %) reached the 1,800 s per-evaluation wall-clock guard and were scored as failures, and 1,612 were rejected on physical constraints. The admissible space spans €7,302–11,594 in BOP should-cost and 0.544–0.599 in cycle efficiency at screening fidelity.

The retained Pareto set separates into two groups by cathode backpressure. Five points sit at 82.5–84.6 kPa(g) with comp_pr near 2.20 and efficiency 0.573–0.575. Seven sit at 127.3–128.6 kPa(g) with comp_pr near 2.48 and efficiency 0.595–0.599. The rest of the design vector is almost constant across the front: comp_kW varies by 2.0 %, turb_kW by 3.9 %, stoich by 4.3 %, and every point holds blow_kW at its 0.5 kW floor. This front is a cathode-pressure trade with the cooling and pumping hardware adjusting around it.

The marine platform therefore behaves differently from the road platform of §5, where all 20 Pareto points sat at the 130 kPa(g) ceiling and no efficiency–cost trade-off existed. Here the cheaper half of the front gives up pressure to save cost, so a trade-off curve does exist. The high-efficiency half still works against the same ceiling, which encodes the D2.2 stack pressure envelope, so efficiency at that end of the front is limited by the stack's rated pressure and not by the search.

Blueprint 2 yielded only five Pareto points, and they are degenerate. Across the whole front comp_kW varies by 0.2 % (97.70–97.89 kW), comp_pr by 0.9 % and bp_kPa by 2.9 % (115.2–118.6 kPa gauge); only radHT_kW (796–873 kW) and radMT_kW (68.5–77.2 kW) move. Every point carries one air path. The role-selection rule of §3.8 consequently returns the same design for the compromise and the high-efficiency role, so blueprint 2 offers two distinct report designs rather than three. This echoes the weak cost-efficiency trade-off at 350 kW (§5), but is more extreme because two report roles collapse to the same design; blueprint 1 instead splits into two pressure clusters.

Table 14. APP-700-M confirmed report roles on the full 7,080 s marine mission.

Role	BOP cost (€)	η_cycle	Min thermal margin (K)	Admissible
BP1 modular — seed baseline	9,855.98	0.5722	2.46	yes
BP1 modular — low-cost	7,302.41	0.5728	2.70	yes
BP1 modular — compromise	7,441.97	0.5865	2.75	yes
BP1 modular — high-efficiency	7,660.18	0.5864	2.59	yes
BP2 common rail — seed baseline	8,724.30	0.5798	2.49	yes
BP2 common rail — low-cost	6,853.52	0.5912	2.74	yes
BP2 common rail — compromise	6,985.80	0.5917	2.67	yes
BP2 common rail — high-efficiency	6,985.80	0.5917	2.67	yes

The minimum thermal margin is the smaller of the HT margin to 358.15 K and the MT margin to 343.15 K. Blueprint 2's compromise and high-efficiency roles resolve to the same archive row. All eight designs are admissible on the full mission, and three results need stating before the numbers are used.

The roles reorder at confirmation fidelity. On the modular blueprint the compromise design confirms at 0.5865 against the high-efficiency design's 0.5864, and costs €218 less. The selection rule picks roles from the screening archive, which is solved at 10 s sampling, whereas confirmation solves the same designs at 1 s. The ordering the rule relied on does not survive the finer solve. For this column the compromise design dominates the high-efficiency design on both objectives, so the third role adds nothing and the reported pair is low-cost and compromise.

Screening overstates efficiency, and the error is not uniform. Confirmation moves the modular high-efficiency design from 0.5992 to 0.5864 and the compromise design from 0.5979 to 0.5865, losses of 1.28 and 1.14 points. The low-cost design does not move at all, 0.5728 at both fidelities. The two high-pressure designs lose

about a point and the low-pressure design loses nothing, so the 10 s screening solve resolves the high-pressure operating point less well than the low-pressure one. Because the error tracks the design rather than offsetting all of them equally, it also reorders the front, which is what produces the role inversion above. Final role comparisons below therefore use only confirmed efficiencies.

Blueprint 2 loses a comparable amount, 0.6034 at screening against 0.5917 confirmed, or 1.17 points. Its designs run at 115–119 kPa(g), well below the pressure ceiling blueprint 1 works against, so the error is not confined to high-pressure operation. On this mission, the 10 s screening solve therefore introduces a design-dependent bias that cannot be assumed to cancel across the 700 kW front.

7.2.2 Design vector of the report roles

Table 15. APP-700-M design vectors of the report roles.

Variable	BP1 low-cost	BP1 compromise	BP1 high-efficiency	BP2 compromise	Unit
comp_kW	46.09	45.75	46.39	97.70	kW
comp_pr	2.209	2.483	2.513	2.475	—
turb_kW	21.64	21.90	21.14	48.21	kW
pump_HT_kW	1.897	1.480	2.145	1.096	kW
pump_MT_kW	0.315	0.323	0.366	0.111	kW
fan_kW	51.47	88.14	67.26	68.83	kW
blow_kW	0.508	0.509	0.507	1.127	kW
stoich	1.576	1.645	1.638	1.583	—
bp_kPa	82.49	127.26	127.68	118.60	kPa gauge
radHT_kW	808.54	844.92	854.29	872.86	kW
radMT_kW	72.27	82.97	106.09	72.26	kW

Ratings are per installed unit. pump_MT_kW is the nameplate of one MT branch pump and the loop carries two; fan_kW, radHT_kW and radMT_kW are system-level, because each system has one functional HT radiator, one functional MT radiator and one fan pack.

The three modular roles differ mainly in cathode pressure and in cooling hardware. The low-cost design runs 82.5 kPa(g) at comp_pr 2.21 and buys the smallest radiator and fan; the other two run about 127 kPa(g) at

comp_pr near 2.5. Compressor and turbine ratings agree to within 1.4 % across all three, and every role holds the recirculation blower at its 0.5 kW box floor.

The optimiser cuts the HT coolant pump hard, from the 4.94 kW seed to 1.48–2.14 kW. The coolant loops carry explicit hydraulic resistance, so pump nameplate buys real parasitic power, and surplus nameplate is charged twice: once in cost and once in the parasitic ledger below.

7.2.3 Component sizing, count, utilisation and cost

Table 16. APP-700-M BOP inventory of the confirmed compromise design of each blueprint.

Component	Count BP1/BP2	Rating BP1	Rating BP2	Util BP1 (%)	Util BP2 (%)
PEM stack module	4 / 4	500 cells/mod	500 cells/mod	—	—
Compressor	2 / 1	91.49 kW	97.70 kW	91.9	87.1
Turbine	2 / 1	43.81 kW	48.21 kW	84.4	75.8
Air-handling EM + inverter	2 / 1	52.45 kW	54.43 kW	99.8	99.0
DC/DC converter	2 / 2	800.0 kW	800.0 kW	95.9	96.3
HT coolant pump	1 / 1	1.48 kW	1.10 kW	65.7	71.1
MT coolant pump	2 / 2	0.65 kW	0.22 kW	14.7	16.8
Radiator fan pack	1 / 1	88.14 kW	68.83 kW	39.3	58.1
Recirculation blower	2 / 1	1.02 kW	1.13 kW	100.0	100.0
HT radiator	1 / 1	844.92 kW	872.86 kW	84.5	84.2
MT radiator	1 / 1	82.97 kW	72.26 kW	25.9	30.3
Charge-air cooler	2 / 1	133.70 kW	128.12 kW	31.9	32.3
Water separator	2 / 1	780.60 g/s	750.98 g/s	84.1	89.9
Cathode air filter	2 / 1	780.60 g/s	750.98 g/s	84.1	89.9
Resonator	2 / 1	780.60 g/s	750.98 g/s	84.1	89.9
Humidifier	2 / 1	780.60 g/s	750.98 g/s	84.1	89.9

Component	Count BP1/BP2	Rating BP1	Rating BP2	Util BP1 (%)	Util BP2 (%)
H_2 ejector	2 / 1	17.00 g/s	17.00 g/s	83.7	80.5
H_2 heat exchanger	2 / 1	1.00 m²	1.00 m²	356 (diag.)	342 (diag.)
Total					

Ratings are installed totals across the count shown, so they are the per-unit values of the design-vector table multiplied by that count. Three entries need reading with care.

The hydrogen heat exchanger shows utilisation above 300 % on both blueprints. Its rating is a flow-throughput proxy in which heat-transfer area is forced to scale with hydrogen flow, and it is a reported diagnostic rather than a capacity gate: the component is thermally adequate at this area, since a measured solve warms fresh hydrogen to 2.4 K of the coolant at 0.5 m² per unit. The number does not indicate an undersized exchanger.

The radiator fan is rated against the 40 °C hot-ambient sizing draw, so its utilisation column is that sizing condition and not a mission peak. Comparing it with the pumps, which are on measured mission peaks, would misread the fan as oversized.

7.2.4 Cycle KPIs and parasitics

Table 17. APP-700-M confirmed cycle KPIs over the full 7,080 s marine mission.

KPI	BP1 low-cost	BP1 compromise	BP2 low-cost	BP2 compromise
Requested net-bus / delivered DC/DC output (kWh)	789.35 / 837.07	789.35 / 844.74	789.35 / 843.05	789.35 / 843.58
Canonical net-bus energy (kWh)	789.38	789.46	789.42	789.41
Aggregate parasitic energy (kWh)	47.69	55.28	53.63	54.17
Parasitic share of DC/DC output (%)	5.70	6.54	6.36	6.42
Fresh H_2 (kg)	41.33	40.37	40.05	40.02
H_2 utilisation (%)	99.96	99.88	99.89	99.87
Purge loss over fresh H_2 (%)	0.045	0.104	0.082	0.085
Cycle efficiency (%)	57.28	58.65	59.12	**59.17**
Min loaded cell voltage (V)	0.637	0.672	0.654	0.664

KPI	BP1 low-cost	BP1 compromise	BP2 low-cost	BP2 compromise
Max current density (A/cm²)	1.876	1.764	1.826	1.795

Both designs track the mission. Net-bus energy matches the 789.35 kWh request to within 0.03 kWh, and the tracking shortfall stays near 1×10^{-4} against a threshold of 1.

Parasitic power is 5.7–6.5 % of DC/DC output. The compromise design spends 7.6 kWh more than the low-cost design, mostly on compression, and recovers considerably more than that at the stack. Hydrogen utilisation is above 99.8 % in every case, with purge losses between 0.045 % and 0.104 % of fresh hydrogen; these follow from the intermittent-purge proxy and are not a measured crossover result. The number of de-stage events and the time spent de-staged are not written by the confirmation artefact, and recording them requires an addition to the KPI writer rather than a new run.

8 1050 kW marine platforms (APP-1050-M)

8.1 Blueprint architectures

Both 1050 kW marine blueprints carry six 500-cell stack modules (3,000 cells) and reject through one shared HT radiator, one shared MT radiator and one aggregate liquid-to-air fan pack, assembled under the same §6 count rules. They differ in how the air and hydrogen paths are scaled and with them the electrical grouping.

Blueprint 1 — linear scaling (D3.1 1050 kW maritime): the 350 kW column repeated three times behind shared heat rejection. Three modular air trains and three hydrogen trains, one per two-module electrical group, each with a turbine for exhaust recuperation, and the DC/DC converters paralleled at the system bus. This is the most directly comparable column in the set: it is the source architecture at 3× count, so any deviation from three times the 350 kW result is an effect of the mission, of the shared heat rejection, or of the cost model's size scaling.

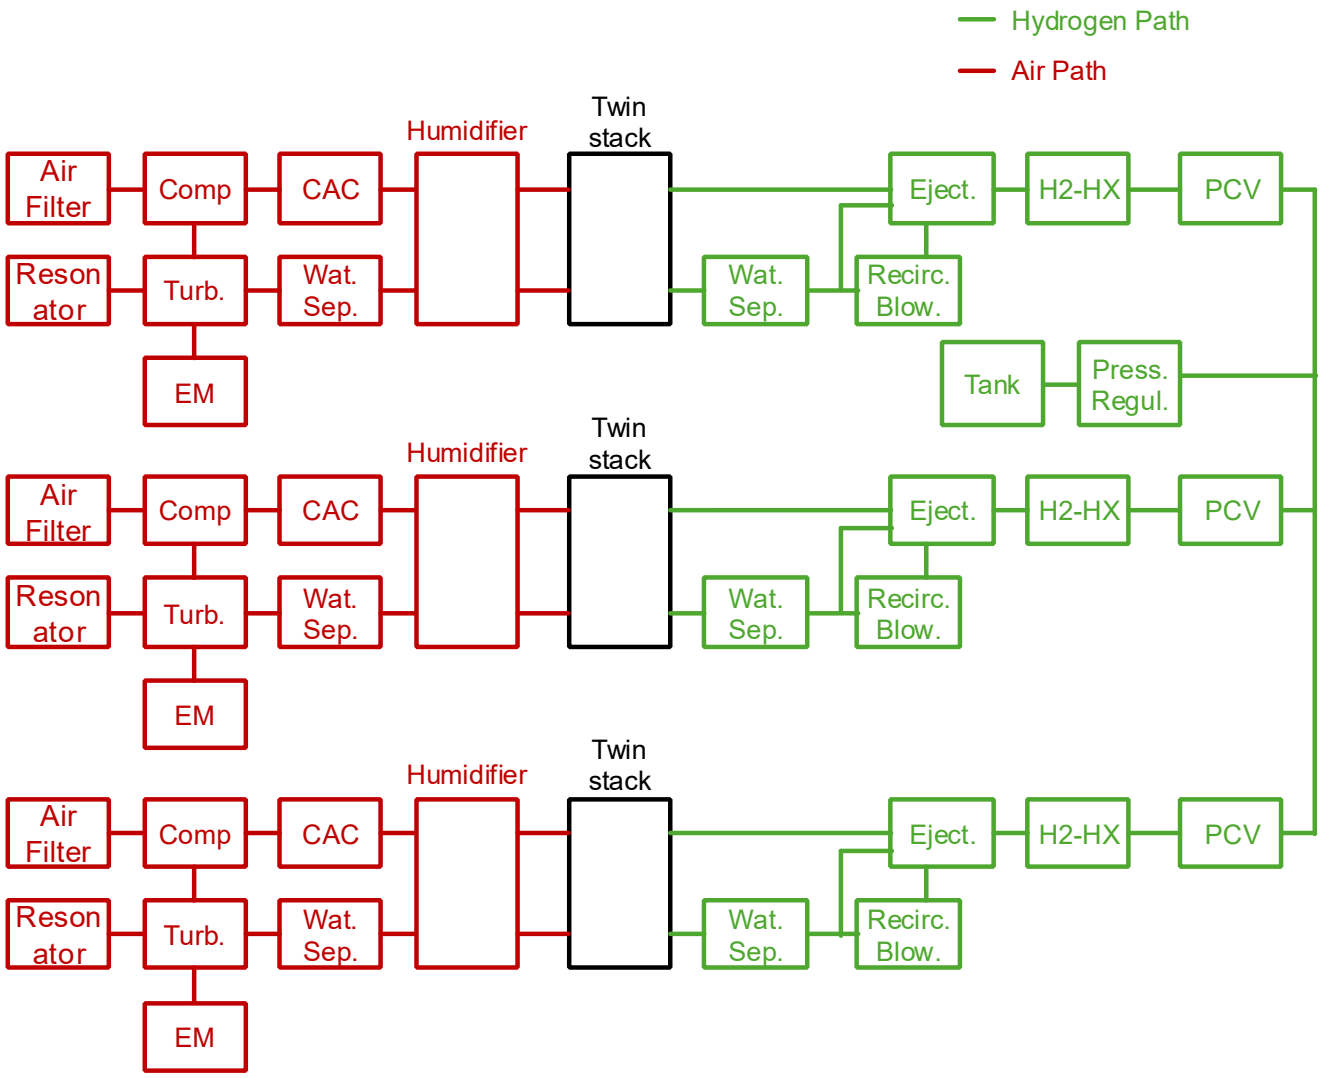


Figure 9. Fuel-cell system simulation model overview for blueprint 1050 kW #1.

Blueprint 2 — common rail (D3.1 1050 kW #2): the same six stack modules fed from air and hydrogen common rails with buffer volumes, supplied by two trains that are deliberately not identical: one high-flow and one low-flow unit. The split is carried as sizing and cost at a declared 2:1 ratio rather than as an operating strategy: the two trains discharge into a common rail, so every stack draws from one buffered plenum and the model sees a single aggregate air supply. Which unit runs when is not simulated. On this mission that omission is conservative, because only 6.5 % of the duration falls below half load, which is where a high/low split pays.

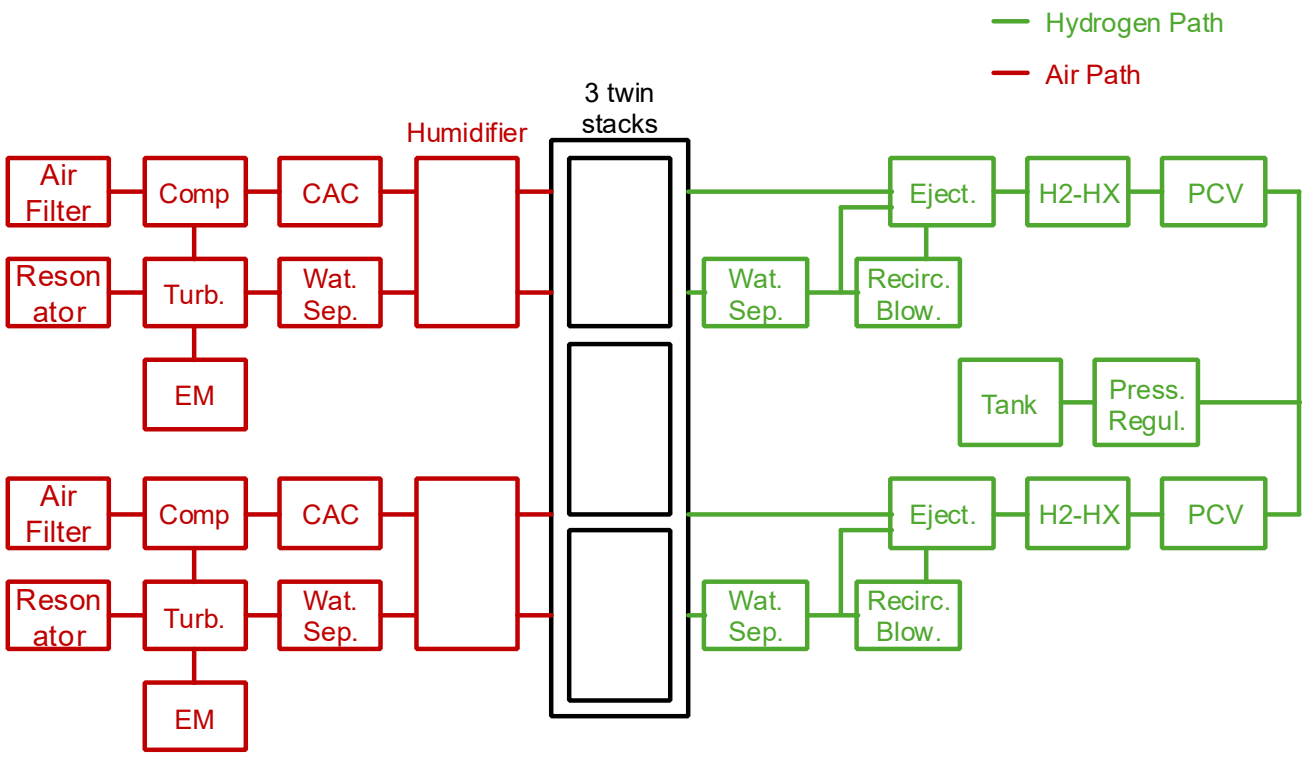


Figure 10. Fuel-cell system simulation model overview for blueprint 1050 kW #2.

Table 18. 1050 kW blueprint comparison.

Architecture count	BP1 — linear scaling	BP2 — common rail
Stack modules (500 cells each)	6	6
Electrical groups	2	2
DC/DC converters	3	3
Air trains	3 (identical)	2 (1 high-flow + 1 low-flow)
Hydrogen trains	3 (identical)	2 (identical)
Turbine (exhaust recuperation)	3	2
Common air / H_2 rails with buffer volume	no	yes
HT radiator / MT radiator / aggregate fan pack	1 / 1 / 1	1 / 1 / 1
HT source branches (stack branches + one H_2-HX → CAC branch per air train)	6 + 3	6 + 2

Architecture count	BP1 — linear scaling	BP2 — common rail
De-stage capability	yes, per group	no; the rail aggregates and staging is not modelled

8.2 Optimisation results

8.2.1 Screening outcome

Screening blueprint 1 evaluated 4,100 designs: 2,070 report-admissible (50.5 %), 86 Pareto points, 1,811 rejected on physical constraints and 56 on performance thresholds. A further 163 evaluations (4.0 %) reached the per-evaluation wall-clock guard and were scored as failures. The admissible space spans €11,144–17,010 in BOP should-cost and 0.4575–0.5053 in cycle efficiency at screening fidelity.

Blueprint 2 evaluated the same budget: 2,484 admissible (60.6 %), 16 Pareto points, 1,407 physical-constraint and 31 performance-threshold rejections, and 178 timeouts (4.3 %). Its admissible space spans €10,012–15,686 and 0.4398–0.5086.

Cathode backpressure separates the two fronts. All 16 blueprint 2 points sit at 128.3–130.0 kPa(g), against the 130 kPa(g) bound that encodes the D2.2 stack pressure envelope of 260 kPa(a) cathode inlet. The efficiency end of that front is set by the stack's rated pressure and not by the search. Blueprint 1 spans 115.4–129.1 kPa(g) with a mean of 122.5, so a pressure trade remains available across its range. The same ceiling limits the 350 kW column of §5 and the high-efficiency half of the 700 kW front in §7.

Table 19. APP-1050-M confirmed report roles on the 2,560 s marine mission.

Role	BOP cost (€)	η_cycle	Min thermal margin (K)	Admissible
BP1 modular — seed baseline	14,195.33	0.4798	2.79	yes
BP1 modular — low-cost	11,143.77	0.4922	2.98	yes
BP1 modular — compromise	11,628.14	0.5025	2.80	yes
BP1 modular — high-efficiency	12,590.94	0.5043	2.67	yes
BP2 common rail — seed baseline	13,052.64	0.4832	2.79	yes
BP2 common rail — low-cost	10,012.18	0.4942	3.00	yes
BP2 common rail — compromise	10,643.43	0.5072	2.82	yes
BP2 common rail — high-efficiency	10,951.29	0.5070	2.80	yes

The minimum thermal margin is the smaller of the HT margin to 358.15 K and the MT margin to 343.15 K. All eight designs are admissible on the full mission.

The HT loop binds on every design. HT margins run 2.67–3.00 K while MT margins run 12.93–13.58 K, leaving about five times more room on the MT loop. This repeats the 350 kW result of §5 and the 700 kW result of §7 at a third power node, with one HT and one MT radiator now serving six stack modules. D3.4 carries the heat duties and loop temperatures behind these margins.

The roles reorder at confirmation fidelity on blueprint 2. Its compromise design confirms at 0.5072 against the high-efficiency design's 0.5070, and costs €308 less. Roles are selected from the screening archive at 10 s sampling and confirmed at 1 s, and the ordering the selection rule relied on does not survive the finer solve. The same inversion appears on the modular 700 kW column in §7.2, so it is a property of the two-fidelity procedure rather than of this architecture. For blueprint 2 the compromise design dominates the high-efficiency design on both objectives, and the reported pair is low-cost and compromise.

8.2.2 Design vector of the report roles

Table 20. APP-1050-M design vectors of the confirmed report roles.

Variable	BP1 low-cost	BP1 compromise	BP1 high-efficiency	BP2 low-cost	BP2 compromise	BP2 high-efficiency	Unit
comp_kW	51.62	50.94	50.95	74.89	72.20	72.29	kW
comp_pr	2.489	2.519	2.471	2.530	2.565	2.560	-
turb_kW	24.36	22.70	21.16	38.01	31.94	32.05	kW
pump_HT_kW	2.436	2.467	2.432	1.853	2.433	3.401	kW
pump_MT_kW	0.366	0.183	0.387	0.222	0.235	0.198	kW
fan_kW	78.31	145.48	73.54	92.00	102.45	93.29	kW
blow_kW	0.556	0.584	0.733	0.761	0.755	0.996	kW
stoich	1.645	1.753	2.041	1.554	1.688	1.724	-
bp_kPa	117.01	124.74	126.07	129.99	129.82	129.32	kPa gauge
radHT_kW	1,199.96	1,415.23	1,480.99	1,199.08	1,385.71	1,339.44	kW
radMT_kW	144.65	102.99	110.03	119.44	118.18	114.85	kW

Ratings are per installed unit. pump_MT_kW is the nameplate of one MT branch pump and the loop carries two; fan_kW, radHT_kW and radMT_kW are system-level, because each system has one functional HT radiator, one functional MT radiator and one fan pack. For APP-1050-M #2 the two installed air trains differ in size: the tabulated value is the nominal per-train rating, and the declared 2:1 split gives a high-flow train at 4/3 of it and a low-flow train at 2/3. Across the low-cost, compromise and high-efficiency roles that is 99.85/49.93,

96.27/48.13 and 96.39/48.19 kW of compressor and 50.68/25.34, 42.59/21.29 and 42.73/21.37 kW of turbine. The split applies to the air-train lines only; hydrogen-path ratings, blow_kW among them, are uniform.

Blueprint 2's air path is sized for two trains rather than three, so its per-unit compressor and turbine ratings are about 45 % larger than blueprint 1's while the installed totals stay within 6 %. Its three roles hold cathode pressure at 129.3–130.0 kPa(g) and vary mainly in cooling hardware, which is the degenerate pattern §5 reports at 350 kW. Blueprint 1 keeps a wider pressure range, 117.0–126.1 kPa(g).

Blueprint 1's high-efficiency role is the one design in either column that buys efficiency with air rather than pressure. It runs stoich 2.041 against 1.55–1.75 everywhere else, and pairs that with the largest HT radiator in the set at 1,481 kW. The compromise role reaches almost the same efficiency, 0.5025 against 0.5043, for €963 less.

8.2.3 Component sizing, count, utilisation and cost

Table 21. APP-1050-M BOP inventory of the confirmed compromise design of each blueprint.

Component	Count BP1/BP2	Rating BP1	Rating BP2	Util BP1 (%)	Util BP2 (%)
PEM stack module	6 / 6	500 cells/mod	500 cells/mod	-	-
Compressor	3 / 2	152.81 kW	144.40 kW	90.4	97.1
Turbine	3 / 2	68.10 kW	63.88 kW	86.9	95.5
Air-handling EM + inverter	3 / 2	93.19 kW	88.57 kW	94.4	99.7
DC/DC converter	3 / 2	1,200.0 kW	1,280.0 kW	96.4	90.4
HT coolant pump	1 / 1	2.47 kW	2.43 kW	63.0	60.5
MT coolant pump	2 / 2	0.37 kW	0.47 kW	17.6	17.1
Radiator fan pack	1 / 1	145.48 kW	102.45 kW	29.8	41.2
Recirculation blower	3 / 2	1.75 kW	1.51 kW	100.0	69.1
HT radiator	1 / 1	1,415.23 kW	1,385.71 kW	74.8	75.0
MT radiator	1 / 1	102.99 kW	118.18 kW	32.1	28.1
Charge-air cooler	3 / 2	217.70 kW	214.18 kW	38.9	36.8
Water separator	3 / 2	1,248.16 g/s	1,201.57 g/s	85.8	89.8

Component	Count BP1/BP2	Rating BP1	Rating BP2	Util BP1 (%)	Util BP2 (%)
Cathode air filter	3 / 2	1,248.16 g/s	1,201.57 g/s	85.8	89.8
Resonator	3 / 2	1,248.16 g/s	1,201.57 g/s	85.8	89.8
Humidifier	3 / 2	1,248.16 g/s	1,201.57 g/s	85.8	89.8
H_2 ejector	3 / 2	20.30 g/s	21.90 g/s	94.1	87.2
H_2 heat exchanger	3 / 2	1.50 m²	1.00 m²	318 (diag.)	478 (diag.)
Total					

Ratings are installed totals across the count shown, so they are the per-unit values of the design-vector table multiplied by that count. Four entries need reading with care.

The hydrogen heat exchanger again shows utilisation far above 100 %. As in §7.2, its rating is a flow-throughput proxy in which heat-transfer area is forced to scale with hydrogen flow, and it is reported as a diagnostic rather than a capacity gate. The component is thermally over-adequate at any area in this range, so the number does not indicate an undersized exchanger.

The radiator fan is rated against the 40 °C hot-ambient sizing draw, not a mission peak, so its utilisation column is not comparable with the pump and radiator columns beside it. Blueprint 1's compromise role buys a 145 kW fan and uses 29.8 % of that sizing draw, which is the largest fan in either column and the least used.

Blueprint 1's recirculation blower sits at exactly 100 % of its rating. The optimiser drives blow_kW down until the anode loop saturates the blower, so the search sets this rating at the point where it is exactly consumed and no engineering margin remains. Blueprint 2's blower runs at 69.1 %. The two blowers also price perversely against each other, blueprint 1's three units costing less than blueprint 2's two, because both sit below the first tabulated anchor of the D3.2 curve where cost is read by extrapolation.

The air-handling machine is the largest single BOP line on both blueprints, and it runs at 94.4 % and 99.7 % of its rating. Blueprint 2's motor and inverter have almost no headroom on the mission peak.

8.2.4 Cycle KPIs and parasitics

Table 22. APP-1050-M confirmed cycle KPIs over the 2,560 s marine mission.

KPI	BP1 low-cost	BP1 compromise	BP1 high-efficiency	BP2 low-cost	BP2 compromise	BP2 high-efficiency
Requested energy (kWh)	662.7	662.7	662.7	662.7	662.7	662.7
Delivered DC/DC (kWh)	727.1	726.3	729.8	724.7	726.4	726.5

KPI	BP1 low-cost	BP1 compromise	BP1 high-efficiency	BP2 low-cost	BP2 compromise	BP2 high-efficiency
Canonical net-bus energy (kWh)	662.0	662.3	662.4	661.7	662.3	662.3
Aggregate parasitic energy (kWh)	65.1	64.0	67.4	63.0	64.1	64.1
Parasitic share of DC/DC output (%)	9.0	8.8	9.2	8.7	8.8	8.8
Fresh H2 (kg)	40.33	39.53	39.40	40.16	39.17	39.18
H_2 utilisation (%)	99.9	99.9	99.9	99.9	99.9	99.9
Purge loss / fresh H2 (%)	0.08	0.08	0.08	0.09	0.09	0.08
Cycle efficiency (%)	49.22	50.25	50.43	49.42	50.72	50.70
Min loaded cell voltage (V)	0.6687	0.6836	0.6896	0.6674	0.6900	0.6897
Max current density (A/cm²)	1.789	1.745	1.734	1.786	1.728	1.729

Parasitic energy is the difference between DC/DC output and the canonical net-bus energy, and it takes 8.7–9.2 % of DC/DC output on every design. The spread across six designs is 4.4 kWh, so the parasitic ledger is not what separates them. Cell voltage does the rest: the minimum loaded cell voltage rises from 0.667 V on the cheapest designs to 0.690 V on the most efficient, and peak current density falls from 1.789 to 1.728 A/cm², inside the 2.00 A/cm² gate.

Hydrogen utilisation is 99.9 % on all six designs and purge loss stays below 0.1 % of fresh hydrogen. The anode loop is not a differentiator at this node. The model reports purge loss but does not predict nitrogen crossover concentration, because no validated crossover correlation exists for this stack.

The mission mean request is 932 kW, 88.7 % of the node, and the profile spends most of its length near nameplate. The mission-integrated cycle efficiency is therefore close to a direct reading of the D3.1 marine target condition of 50 % at 82–90 % load. Four of the six reported designs meet it, 50.25 % and 50.43 % on blueprint 1 and 50.72 % and 50.70 % on blueprint 2, and the two low-cost designs fall just short at 49.22 % and 49.42 %. The comparison carries the same qualifications as §5: it is a simulated result on an assumption-based model, and D3.1 does not fully define the auxiliary boundary or the heating-value basis of its target.

9 1050 kW aviation platform (APP-1050-A)

9.1 Blueprint architecture

The aviation column carries six 500-cell stack modules (3,000 cells) in three electrical groups behind the DC/DC converters, with three modular air trains and three hydrogen trains, one per group, and one shared HT radiator, one shared MT radiator and one aggregate liquid-to-air fan pack. Topologically it is blueprint 1 of §8.1 at the same count. What separates the two columns is the operating environment and the fuel supply, not the equipment list.

Ambient conditions vary along the mission rather than sitting at a fixed corner. On the ground the model uses 101.3 kPa and 298.15 K; at the 7,620 m cruise the ISA values are 37.6 kPa and 238.7 K. The corrected air flow the compressor must deliver rises by a factor of 2.45 between the two, so this is a change of operating point rather than a derating applied to a sea-level result. Four consequences follow for the equipment.

The compressor uses a different map. EK40C-2429 is a higher-pressure-ratio, lower-flow trim of the same Rotrex frame, with a pressure-ratio envelope reaching 3.12 against EK40C-3425's 2.594. The marine map does not contain the cruise operating point, and off-map operation is not a result. The trade is design-point efficiency, 0.732 against 0.755.

The turbine is retained. D3.1 Table 7 removes it from the aviation blueprint on space grounds, and that is overridden here. At the same commanded backpressure the expansion ratio available at cruise is larger than at sea level, so the turbine recovers more shaft work in the environment where the blueprint proposed to delete it.

The hydrogen path is modelled in two parts. The gas-side chain upstream of the pressure regulator uses the compressed-gas supply of the other five columns. The cryogenic vaporiser the blueprint specifies is carried separately, as a commanded heat sink on the MT coolant loop, at about 3 bar tank delivery with 25 K saturation, 390 kJ/kg latent heat and 13.0 kJ/(kg K) mean gas specific heat. That is roughly 4.0 MJ/kg, of which boiling is about a tenth. The vaporiser is therefore a cold sink on the MT loop rather than a heat load, and §9.2 shows it in the MT temperature margins.

De-staging is restricted to the ground. A staging threshold on a flying mission has to sit below the lowest in-flight power request, so this column uses 150 kW against the 320 kW of the marine columns.

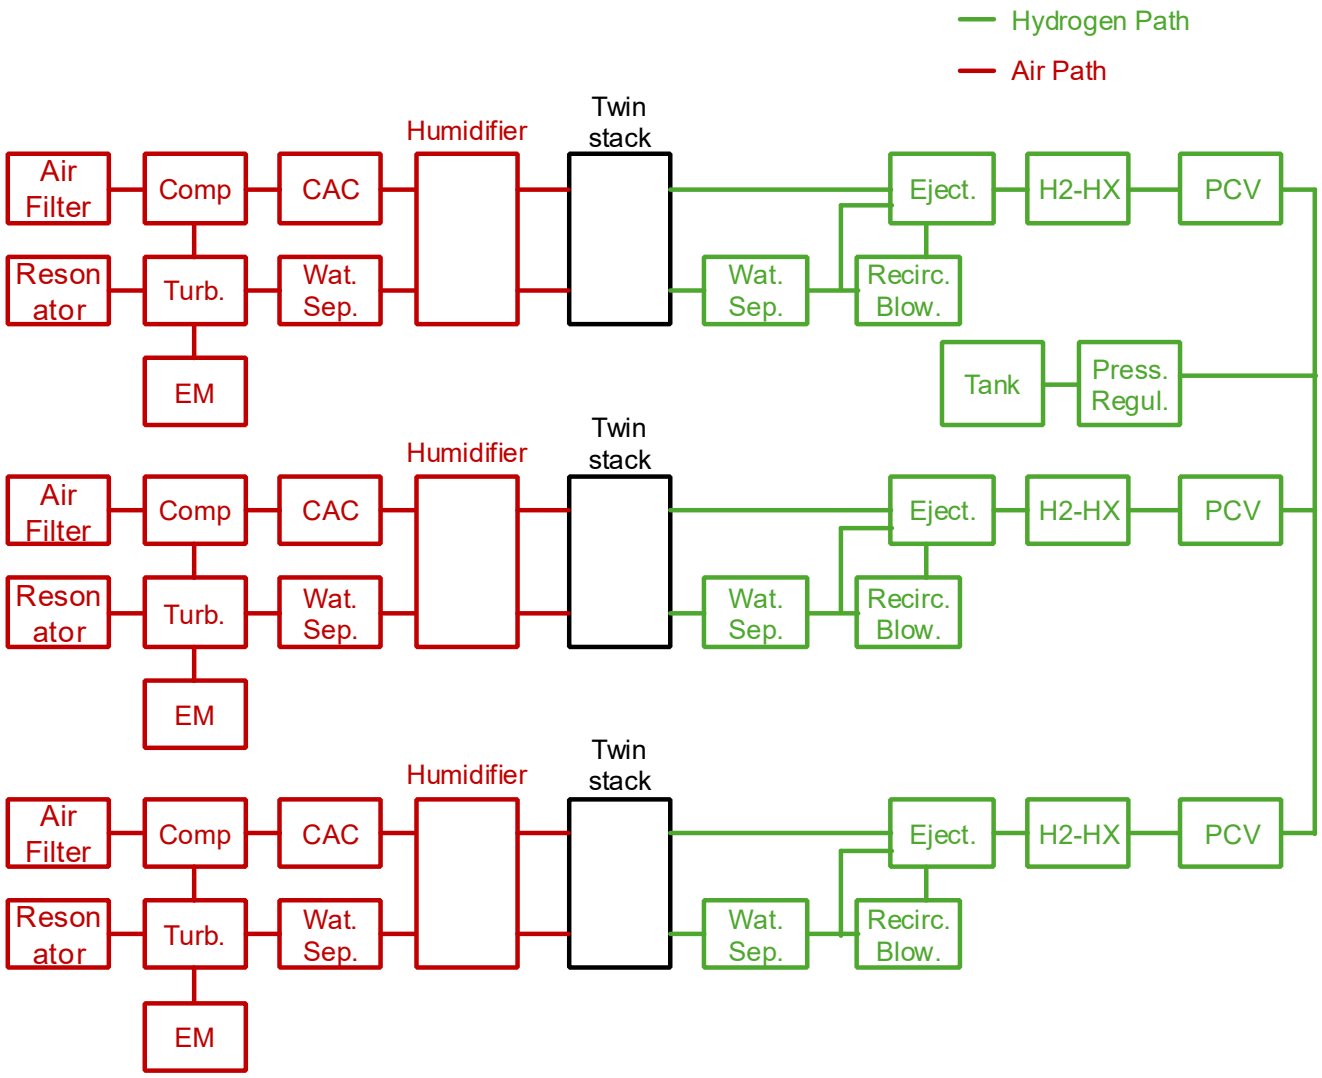


Figure 11. Fuel-cell system simulation model overview for blueprint APP-1050-A.

9.2 Optimisation results

9.2.1 Screening outcome

Screening evaluated 4,100 designs on the 3,360 s aviation mission: 2,819 report-admissible (68.8 %), 21 Pareto points, 1,280 rejected on physical constraints and one lost to the per-evaluation wall-clock guard. The admissible space spans €10,613–14,842 in BOP should-cost and 0.5345–0.5510 in cycle efficiency at screening fidelity. The admissible fraction is the highest of the six columns, and the solver behaved better here than on the marine node, where 163 and 178 evaluations timed out.

The retained front is narrow in every variable except two. Across all 21 points comp_pr varies by 0.9 % (2.303–2.323), stoich by 3.3 %, turb_kW by 2.3 %, and cathode backpressure sits at 128.9–129.7 kPa(g) on every point. The two variables that do move are em_margin, 1.114–1.247, and the cooling hardware. This is a pressure-limited front with an electric-machine sizing trade running through it.

The pressure limit is the same one that binds at 350 kW and at the efficiency end of the 700 kW and 1050 kW marine fronts: the 130 kPa(g) bound encoding the D2.2 stack envelope of 260 kPa(a) cathode inlet. Under the

altitude scheduling of §9.1 that bound is reached on the ground, so it is take-off and not cruise that fixes the cathode pressure of every design in the front.

em_margin is a twelfth design variable carried only on this column. The air-handling machine shares one shaft with the compressor and turbine and is sized to the net demand, em_kW = em_margin × (comp_kW − turb_kW). On the other five columns that factor is fixed at 1.1. At altitude the compressor needs more shaft torque at low corrected speed, and with the factor fixed the optimiser had no actuator for the resulting EM torque constraint. Releasing it as a search variable is what makes admissible designs reachable on this column.

Table 23. APP-1050-A confirmed report roles on the 3,360 s aviation mission.

Role	BOP cost (€)	η_cycle	Min thermal margin (K)	em_map margin	Admissible
Seed baseline	14,200.70	0.5178	1.65	-0.179	no
Low-cost	10,613.19	0.5449	1.87	+0.0012	yes
Compromise	10,637.78	0.5452	1.94	-0.0014	no
High-efficiency	11,248.90	0.5447	1.92	+0.0996	yes

The minimum thermal margin is the smaller of the HT margin to 358.15 K and the MT margin to 343.15 K. Two of the four designs are inadmissible, and in both cases the sole failing gate is em_map, the air-handling machine's torque envelope. Every other physical check passes on all four.

The front has converged onto that gate. The low-cost design clears it by 0.0012 in normalised torque margin and the compromise design misses it by 0.0014, a band of about 0.3 % straddling the constraint. Only the high-efficiency design carries real headroom, 0.0996. This is the expected consequence of making em_margin a searched variable: torque headroom is bought with motor and inverter cost, so the optimiser purchases as little as the constraint allows. It also means the 10 s screening grid and the 1 s confirmation grid disagree on admissibility rather than only on efficiency, because a torque peak resolved at 1 s is not resolved at 10 s.

The seed baseline misses by 0.179, and the search says why. Its air-handling machine is sized at the fixed 1.1 factor of the other columns, while the converged designs need 1.114–1.247. The seed's EM is undersized for this environment by roughly 14 %, which is a result about the blueprint rather than about the search.

The reported design is the low-cost role. It is admissible, it is the cheapest of the four, and at 0.5449 against 0.5447 it is also marginally more efficient than the high-efficiency role, which it therefore dominates on both objectives. The high-efficiency role remains the only design with meaningful em_map headroom, and a build would be better served by it than by the €636 saving.

The HT loop binds, as on every other column, but more tightly. HT margins are 1.65–1.94 K against 2.67–3.00 K on the 1050 kW marine column and 2.46–2.75 K at 700 kW. MT margins run the other way, 15.57–16.14 K against about 13 K on the marine columns, because the cryogenic vaporiser removes heat from the MT loop instead of adding it. The cold sink buys MT margin the marine columns do not have, and it does not help the loop that actually binds.

9.2.2 Design vector of the report roles

Table 24. APP-1050-A design vectors of the confirmed report roles.

Variable	Low-cost	Compromise	High-efficiency	Unit
comp_kW	49.86	49.33	51.03	kW
comp_pr	2.321	2.322	2.303	-
turb_kW	24.76	24.69	25.03	kW
pump_HT_kW	1.709	1.693	1.660	kW
pump_MT_kW	0.156	0.253	0.362	kW
fan_kW	47.05	53.05	61.72	kW
blow_kW	0.502	0.513	0.538	kW
stoich	1.715	1.732	1.746	-
bp_kPa	128.87	129.44	129.42	kPa gauge
radHT_kW	1,082.73	1,046.41	1,111.16	kW
radMT_kW	108.60	114.26	108.95	kW
em_margin	1.138	1.148	1.247	-

Ratings are per installed unit. pump_MT_kW is the nameplate of one MT branch pump and the loop carries two; fan_kW, radHT_kW and radMT_kW are system-level. bp_kPa is a scheduling coefficient rather than a fixed gauge offset, so the commanded cathode pressure it produces falls with ambient along the mission.

The three roles are close to one another in the air path and separate mainly in em_margin and cooling hardware. Compressor and turbine ratings agree within 3.4 %, cathode pressure within 0.4 %, and every role holds the recirculation blower near its 0.5 kW box floor. The high-efficiency role buys 9.6 % more EM headroom and a 31 % larger fan than the low-cost role, and returns 0.0002 in cycle efficiency for €636.

9.2.3 Component sizing, count, utilisation and cost

Table 25. APP-1050-A BOP inventory of the confirmed low-cost design.

Component	Count	Rating	Measured peak	Utilisation (%)
PEM stack module	6	500 cells/mod	-	-
Compressor	3	149.59 kW	125.17 kW	83.7
Turbine	3	74.27 kW	56.26 kW	75.8
Air-handling EM + inverter	3	85.72 kW	76.67 kW	89.4
DC/DC converter	3	1,200.0 kW	1,128.16 kW	94.0
HT coolant pump	1	1.71 kW	0.91 kW	53.4
MT coolant pump	2	0.31 kW	0.05 kW	15.3
Radiator fan pack	1	47.05 kW	44.63 kW	94.9
Recirculation blower	3	1.51 kW	1.51 kW	100.0
HT radiator	1	1,082.73 kW	966.34 kW	89.3
MT radiator	1	108.60 kW	30.76 kW	28.3
Charge-air cooler	3	191.67 kW	55.26 kW	28.8
Water separator	3	1,220.90 g/s	1,007.71 g/s	82.5
Cathode air filter	3	1,220.90 g/s	1,007.71 g/s	82.5
Resonator	3	1,220.90 g/s	1,007.71 g/s	82.5
Humidifier	3	1,220.90 g/s	1,007.71 g/s	82.5
H_2 ejector	3	22.10 g/s	19.10 g/s	86.4
H_2 heat exchanger	3	1.50 m²	-	318 (diag.)

Ratings are installed totals across the count shown. Three entries need reading with care, and one is specific to this column.

The radiator fan runs at 94.9 % of its rating. On the other five columns the fan utilisation column reports the 40 °C hot-ambient sizing draw, which is a ground condition; here the same check is applied and the low-cost design leaves only 5 % of it unused. That is the tightest fan margin in the deliverable, and it sits on the design the search selected for cost.

The hydrogen heat exchanger again reports utilisation far above 100 %, for the reason given in §7.2: its rating is a flow-throughput proxy and a reported diagnostic, not a capacity gate. On this column the exchanger also carries the vaporiser duty, which is modelled as the commanded MT sink rather than through this rating.

The recirculation blower again sits at exactly 100 % of a rating the box floor sets rather than the search. Its cost is the same D3.2 extrapolation artefact described in §8.2.3 and should not be read as a real price for small blowers.

9.2.4 Cycle KPIs and parasitics

Table 26. APP-1050-A confirmed cycle KPIs over the 3,360 s aviation mission.

KPI	Seed baseline	Low-cost	Compromise	High-efficiency
Requested energy (kWh)	485.4	485.4	485.4	485.4
Delivered DC/DC (kWh)	536.0	517.0	517.1	517.9
Canonical net-bus energy (kWh)	485.9	485.4	485.4	485.4
Aggregate parasitic energy (kWh)	50.1	31.5	31.6	32.5
Parasitic share of DC/DC output (%)	9.3	6.1	6.1	6.3
Fresh H_2 (kg)	28.14	26.72	26.71	26.73
H_2 utilisation (%)	99.9	100.0	100.0	99.9
Purge loss / fresh H_2 (%)	0.01	0.02	0.02	0.02
Cycle efficiency (%)	51.78	54.49	54.52	54.47
Min loaded cell voltage (V)	0.6756	0.6976	0.6983	0.6995
Max current density (A/cm²)	1.764	1.667	1.669	1.665

The optimised designs cut parasitic energy from 50.1 kWh to 31.5–32.5 kWh, from 9.3 % of DC/DC output to 6.1–6.3 %. That is the largest parasitic reduction against a seed in the deliverable, and most of it is air-path work: the seed runs a larger compressor at a lower cathode pressure, and the scheduled backpressure lets the optimised designs hold a higher pressure ratio without paying for it at altitude.

Cycle efficiency of 54.5 % is the highest reported in this deliverable, and it is a property of the mission before it is a property of the architecture. The aviation profile delivers a mean 520 kW, 49.5 % of the node, against

88.7 % for the 1050 kW marine mission, so the stack spends most of the sortie at a load where cell voltage is high. The two columns share an equipment list and cannot be ranked against each other on this number.

10 Sensor positions and monitoring rationale

Following the dedicated internal WP3 sensor-positioning review (T. Dankers, C. Bekdemir, A. Balaji and F. Ribatti, System architectures and sensor positioning in PEM fuel cell systems, TNO Powertrains, 10 April 2026), sensor placement uses a non-intrusive, boundary-based approach. Measurements are concentrated at subsystem interfaces and stack inlets and outlets, where they support control, protection and diagnostics without requiring sensors inside the electrochemical stack.

Table 27. Sensor positions and monitoring rationale.

Subsystem and position	Measurement	Monitoring rationale
Anode, each stack inlet and outlet	Pressure; inlet temperature	Anode-pressure control, purge/recirculation supervision and cathode-anode differential-pressure protection
Anode outlet	Outlet temperature and H_2 composition, where required	Experimental thermal monitoring and architecture-dependent purge/bleed control
Active hydrogen recirculation blower	Rotational speed	Blower control and parasitic-power assessment; passive ejectors require no internal sensing
Compressor inlet and outlet	Pressure; compressor speed and electrical power	Operating-envelope protection, actuator control and efficiency monitoring
Cathode, each stack inlet	Pressure and temperature; humidity where required	Local cathode boundary conditions, pressure/stoichiometry control and water management
Cathode, each stack outlet	Pressure, where required	Pressure-drop monitoring and diagnostics
Shared air buffer, where fitted	Buffer pressure	Primary feedback for buffer-pressure control
Coolant, stack/loop inlet and outlet	Temperature	Thermal control, over-temperature protection and heat-balance monitoring
Coolant loop, where required	Flow, pressure and conductivity	Thermal-balance, fault and coolant-health monitoring

Subsystem and position	Measurement	Monitoring rationale
Stack and package boundary	Stack voltage and current; enclosure H_2 detection	Power control, electrical protection, leak detection and emergency shutdown

In parallel multistack systems, pressure measurements are replicated at equivalent locations for every stack: cathode inlet and anode inlet/outlet. A central pressure measurement alone cannot reveal stack-to-stack imbalance. Per-stack mass flow and gas composition sensors are generally avoided; these quantities are preferably estimated from pressure measurements and system models to limit sensor count. Direct flow or composition sensing remains application dependent.

11 Comparison with state of the art

D3.1 Table 6 lists eight high-power fuel-cell systems with reported power, gravimetric density and peak efficiency. Setting the results of §5 to §9 beside them requires care on one point above all. D3.1 Table 6 reports a peak, and D3.1 records that those peaks are typically reached at 20–30 % load. Every efficiency in this deliverable is an average over a complete duty cycle whose mean load is higher. The two quantities answer different questions, and no ranking against a named product is made here.

On scale the modelled nodes sit inside the reported field rather than beyond it. D3.1 Table 6 spans 120 kW to 1500 kW, and the columns here are 350, 700 and 1050 kW. The maritime entries cluster at 200–225 kW with a single 1500 kW system above them, so the 700 kW and 1050 kW marine blueprints populate a gap in that field instead of extending its upper end.

Where the figures can be placed side by side, the deliverable's are averages taken at a higher load than the peaks they would be compared against, so the comparison runs against the modelled systems rather than favouring them. The road column averages 59.70–59.83 % over a mission at 35 % mean load, against the 60 % peak reported for the 120 kW Ballard FCmove. The 700 kW marine column averages 58.65–59.17 % at 57 % mean load, against 53.5–54 % peaks for the 200–225 kW maritime entries. The 1050 kW marine column, nearest in scale to the 1500 kW HDF system and its 49 % peak, averages 50.25–50.72 % on a mission that holds 89 % of nameplate throughout.

The aviation pair cannot be compared at all. ZeroAvia's ZA600 is reported at 55–60 %, but D3.1 records its parameters at stack level, so that figure excludes the balance-of-plant draw which the 54.47–54.52 % of §9 includes. Correcting one boundary onto the other would need a stack-level accounting this deliverable does not produce, and none is attempted.

Three quantities are left uncompared. Gravimetric density is the least comparable column in D3.1 Table 6, spanning 0.014 to 1.4 kW/kg, and no stack, BOP, coolant or enclosure mass model exists here, so no kW/kg figure is offered. Gravimetric density is nevertheless pursued indirectly. In the component models a higher rating means a higher should-cost, and for BOP components a higher rating generally also means a heavier unit, so minimising should-cost drives the ratings, and with them BOP mass, towards the duty actually required. The cost boundary of §4.3 is the D3.2 BOP should-cost and excludes stacks, storage and integration,

so it cannot be set against a system CAPEX figure. No degradation or ageing model exists either, so completing a mission is evidence of capacity rather than of durability.

What the results do support is internal comparison: design against design along one blueprint's trade-off surface, and blueprint against blueprint at the same power node, evaluated under one mission and one search policy.

12 Reproducibility

The simulation model, the duty-cycle data and the optimisation scripts are not released. This document is the public output of the task. The restriction follows from the inputs rather than from the code: the should-cost model is built directly on the D3.2 component cost curves, which are partner-confidential, and the component maps are supplier data provided to the project under the same terms. Publishing the scripts without those inputs would not let anyone reproduce a number, and publishing them with the inputs is not permitted.

Reproducibility is therefore carried by the written record. The design vector and its search bounds are in Section 3.2, the duty cycles in Section 2.2, the optimiser settings and seed policy in Section 3.6, the software versions, solver and tolerances in Section 3.7, and the objectives, admissibility gates and cost boundary in Section 4. Each retained result set records the design vector, the derived component sizing, the cycle definition, the solver policy and seed, the KPI and constraint outcomes, the cost breakdown, the assumptions in force and the software revision that produced it.

13 Conclusions

This deliverable optimises six blueprint columns across four application nodes on one framework: a coupled fuel-cell and thermal system model, a multi-objective search over the design vector of §3.2, and a confirmation stage that re-solves every reported design at 1 s sampling on the mission it was optimised against. That is the full supplied mission on every column except APP-1050-M, where it is the 2,560 s cycle derived from the supplied 64,500 s profile (§2.2). All report-role designs are admissible except the aviation compromise design, which misses the air-handling machine’s torque envelope by 0.0014 in normalised margin against the low-cost design’s 0.0012 clearance (§9.2).

Balance-of-plant should-cost runs €3,690–3,775 at 350 kW, €6,854–7,660 at 700 kW, €10,012–12,591 at 1050 kW marine and €10,613–11,249 on the aviation column. Cycle efficiency runs 59.70–59.83 %, 57.28–59.17 %, 49.22–50.72 % and 54.47–54.52 % in the same order. Efficiency tracks the mission's mean load rather than the power node: the 1050 kW marine mission holds 89 % of nameplate throughout and is the least efficient, the road mission at 35 % mean load the most.

One result repeats on every column. The high-temperature loop binds and the medium-temperature loop does not. HT margins run 1.65–3.00 K against the 358.15 K limit while MT margins run 12.9–16.1 K against 343.15 K, so the MT loop carries roughly five times more room at every node. The admissibility gate compares the lumped loop temperature directly against these limits. With two lumped coolant nodes and an isothermal stack (§3.4), margins of a few kelvin are therefore not a hardware reserve: they are reliable for ranking designs and for identifying the HT loop as the binding one, while absolute sizing is to be confirmed against the detailed WP4 model and the limit basis in D3.4. The optimiser also cuts the MT radiator sharply from each seed, because surplus MT core is charged twice: once when it is bought, and again when it obstructs the shared frontal area ahead of the binding HT core.

A second constraint recurs. Cathode backpressure is bounded at 130 kPa(g) by the D2.2 stack envelope, and on most columns that bound rather than the search sets the efficiency end of the front. The 350 kW front sits on it entirely, as do the high-efficiency half of the 700 kW front, all sixteen points of the 1050 kW common-rail front and the whole aviation front. A front pinned there is limited by the stack's rated pressure: a bounded optimum, not a search that stopped early.

Consolidating air trains changes cost and parasitic draw but not the thermal bottleneck. At 700 kW the common rail undercuts the modular blueprint by €456 at the compromise role and is 0.52 points more efficient, while peak HT temperature differs by 0.08 K (§7.2). At 1050 kW both blueprints are reported without a ranking; their HT margins agree within 0.02 K and their HT radiators within 2 %.

The aviation column is the exception, and its binding constraint is electrical rather than thermal. Its front converges onto the air-handling machine's torque envelope: two of the four confirmed designs miss that gate alone, and the two that pass straddle it within 0.13 %. The machine sized at the fixed margin used on the other five columns falls about 14 % short of what the altitude environment needs, a result about the blueprint rather than about the search.

Five limits qualify the above. Screening overstates cycle efficiency by roughly 1.2 points against confirmation, so none is reported. That error is design-dependent, so the 10 s screening front is not the 1 s front: the reported designs are the designs the two-stage procedure identified and confirmed rather than a demonstrated 1 s optimum, and re-evaluating each column's full Pareto pool at 1 s to reconstruct the confirmed front is the first item of further work (§3.8). The 1050 kW marine designs are confirmed on the 2,560 s cycle derived from the supplied 64,500 s mission rather than on that mission itself. Loop time constants of a few seconds rule out thermal drift over the longer mission (§2.2), but the surrogate reproduces the parent's statistics rather than its particular sequence of excursions, and solving the reported designs on the full profile is the second item of further work. The 40 °C hot-ambient capacity check is algebraic rather than a second mission solve. The comparisons that can and cannot be drawn against the state of the art are set out in §11.