QPC Aerospace Pilot

Phase 1 — Problem Definition

Three-Manufacturer Competitive Strategy Optimization — Airbus, Boeing, COMAC

156-Qubit Polycontextural QAOA Execution on IBM Fez

Horizon: 2026 – 2029

Quantum Polycontextural Computation (QPC)

Dr. Ilan Kreitmann · Ra'anana, Israel

Document status: Phase 1 specification — pre-data, pre-code

Version 1.0 · April 2026

1. Executive Summary

This document defines Phase 1 of the QPC Aerospace Competitive Strategy Pilot: a quantum optimization study implementing a three-manufacturer, three-segment competitive dynamics model as a 156-qubit polycontextural quantum computation executed on IBM Fez. It is the reference specification for all subsequent project phases (data collection, circuit implementation, hardware execution, results analysis, and industrial case study report).

The pilot addresses a genuinely hard computational problem that no single-context quantum or classical optimizer handles natively: the simultaneous strategic positioning of three manufacturers — Airbus, Boeing, and COMAC — competing across three commercial aerospace market segments (narrowbody, widebody, regional), with different objectives per manufacturer, different constraints, different addressable geographies, and market-share dynamics that couple the three manufacturers' outcomes together. QPC's polycontextural architecture represents each manufacturer as an independent logical contexture, with transjunctional coupling between contextures encoding the zero-sum nature of aircraft orders.

The pilot is built exclusively on public data. It is framed as an industrial case study applicable to Airbus, Boeing, any major airline, any aircraft lessor, and any aerospace defense prime. It does not claim to be a named-customer engagement. The case study deliverable is designed for use as a sales asset in subsequent approaches to industrial prospects.

The pilot is scoped to the horizon 2026 through 2029, because (i) order books are largely locked through this window, (ii) production capacity decisions are already made, and (iii) beyond 2029 the analysis becomes speculative in ways no optimization result can defend.

2. Strategic Objective

Each manufacturer seeks to maximize, over the 2026–2029 window, the sum of (i) new aircraft contracts won and (ii) new airline customer relationships established, subject to financial, operational, certification, and geopolitical constraints.

The optimization is expressed, per manufacturer context i, as a cost function to be minimized:

$$J_i = -\left[\text{expected\_contracts\_won}_i + w_i \cdot \text{expected\_new\_airlines}_i \right] + \lambda \cdot \text{constraint\_violations}_i$$

The three manufacturer costs are combined with transjunctional coupling into a single polycontextural objective:

$$J_{total} = \Sigma_i J_i + \mu \cdot \text{inconsistency\_penalty}$$

where the inconsistency penalty discourages configurations in which two manufacturers both claim the same airline contract. This term is the mathematical embodiment of the zero-sum market-share dynamic and is implemented physically by transjunctional CZ gates between corresponding qubit positions across manufacturer contexts.

3. Decision Variables (Strategic Levers)

Each manufacturer has five strategic levers per market segment. Each lever is encoded as a single qubit with binary intensity (baseline vs. aggressive action). This choice keeps the decision space interpretable while preserving the ability to discriminate strategy configurations.

3.1 The Five Levers

#	Lever	Baseline (qubit = 0)	Aggressive action (qubit = 1)
L1	Price positioning	List price, standard discounts	Strategic discount ≥10% on target campaigns
L2	Delivery reliability	Historical delivery performance	Committed on-time delivery program with reserved capacity
L3	Quality & rework reduction	Current production defect rate	Active quality-recovery program; reduced in-service rework
L4	Fleet portfolio adaptation	Current product mix	Accelerated model variants aligned to demand shift (range, capacity, freighter)
L5	Customer service & support	Standard service network	Enhanced MRO, spares, and airline-facing engineering support

The choice of binary encoding (rather than 2-bit or 3-bit per lever) is deliberate: it keeps the per-manufacturer decision space tractable at 5 × 3 = 15 qubits per context, and the analytical output — "should Airbus discount aggressively in narrowbody, yes or no" — is directly interpretable by an industrial reader. Finer granularity is available via Phase-2 extensions once the binary-level signal is established.

3.2 Market Segment Participation

The three manufacturers do not compete in every segment. Participation is constrained to published reality:

Segment	Airbus	Boeing	COMAC
Narrowbody	A320neo family, A220	737 MAX family	C919
Widebody	A330neo, A350	787, 777X	— (CR929 indefinitely postponed)
Regional (under 150 seats)	A220-100 (partial presence)	— (no current regional jet)	ARJ21

Where a manufacturer does not compete (Boeing regional, COMAC widebody), the corresponding qubits are forced to the zero state by the encoding and contribute zero to the objective. This is honest: it preserves the 156-qubit register while faithfully representing the competitive structure.

4. Qubit Allocation (156Q on IBM Fez)

The full register of 156 qubits is allocated deliberately. Every qubit has an assigned role; none are included as padding.

4.1 Per-Manufacturer Decision Qubits (98 Qubits Total)

Manufacturer	Narrowbody (5L)	Widebody (5L)	Regional (varies)	Manufacturer total
Airbus	5	5	5 (A220)	15 logical + 23 extended = 38
Boeing	5	5	0	15 logical + 15 extended = 30
COMAC	5	0	5 (ARJ21)	15 logical + 15 extended = 30
Subtotal	15	10	10	98

Extended qubits within each manufacturer context carry geographic market specialization (see Section 5) and intermediate circuit structure required for faithful encoding of the objective function.

4.2 Coupling Ancillas (20 Qubits)

Transjunctional coupling between manufacturer contexts is implemented via 20 ancilla qubits positioned to mediate entanglement between homologous (lever, segment) positions across pairs of manufacturers. Using ancillas rather than direct long-range coupling is the single most important noise-mitigation architectural choice: it allows deeper logical coupling with shallower transpiled depth on the Heron lattice.

The 20 ancillas are allocated: 10 for Airbus–Boeing coupling (both compete in narrowbody and widebody), 5 for Airbus–COMAC narrowbody coupling, 5 for Boeing–COMAC narrowbody coupling. Regional segment coupling is deferred to Phase-2 extensions.

4.3 Market-Condition Ancillas (18 Qubits)

Exogenous market conditions affect all three manufacturers simultaneously. Encoding them as shared quantum variables lets the circuit respond to market context rather than assuming the manufacturers operate in a vacuum. Six exogenous variables at three qubits each:

- Gulf carrier demand cycle (3Q) — Emirates, Etihad, Qatar, Saudia widebody orders
- Chinese domestic aviation growth rate (3Q) — Big 3 + regional Chinese carrier demand
- Post-COVID global traffic recovery trajectory (3Q) — load factor and fleet renewal urgency
- Jet fuel price regime (3Q) — shapes narrowbody-vs-widebody economics
- Low-cost carrier capacity expansion (3Q) — IndiGo, Ryanair, Wizz, VietJet order activity
- Regulatory and certification climate (3Q) — EASA, FAA, CAAC pace and reciprocity

4.4 Reserved Qubits (20)

The final 20 qubits are reserved. During transpilation, Qiskit's router uses unused qubits as routing ancillas when two logical qubits need to interact across the Heron coupling graph. By reserving 20 qubits explicitly rather than packing the logical encoding into all 156, we reduce the probability that the transpiler corrupts the logical layout by conscripting logical qubits into routing roles.

4.5 Total Allocation

Category	Qubits	% of 156
Per-manufacturer decision qubits (3 × ~32)	98	62.8%
Transjunctional coupling ancillas	20	12.8%
Shared market-condition ancillas (6 × 3Q)	18	11.5%
Reserved for transpilation routing	20	12.8%
Total	156	100%

5. Geographic Constraint Structure

Each manufacturer has a realistic addressable market in the 2026–2029 window. The optimization respects these constraints via hard penalties in the objective:

Manufacture	Addressable markets (yes)	Constrained or excluded (no)
Airbus	Europe; North America; Latin America; Middle East; Southeast Asia; India; Africa; Australia-NZ; limited China (historical)	Certain Chinese orders now routed to COMAC; selected sanctioned markets
Boeing	North America; Europe; Middle East; Southeast Asia; Japan; Korea; Latin America; Australia-NZ; India	Meaningful China market access constrained post-2018; Russia market sanctioned
COMAC	China domestic; Belt-and-Road Southeast Asia (GallopAir Brunei, TransNusa Indonesia); selected African and Central Asian markets	EASA-certified routes; FAA-certified routes; all markets requiring bilateral certification with Western regulators

This constraint structure is encoded in the objective function through region-specific penalty terms that activate when a manufacturer's decision configuration would imply winning an airline in a market it cannot realistically serve. The penalty amplitude is large enough to exclude infeasible configurations from the optimization's effective search space.

6. Polycontextural Circuit Architecture

The circuit follows the standard three-layer QPC architecture previously validated on IBM Torino (PRCBS, Crash Detection, Boundaries Test) and IBM Fez (PQST-64, Holographic Memory). The implementation is native-gate targeted for the Heron processor to minimize compiled depth.

6.1 Layer 1 — Kenogrammatic Preparation

Hadamard gates applied across all 98 decision qubits and the 18 market-condition ancillas, placing the register in an equal superposition across the full decision space. Coupling ancillas are prepared in |0⟩ state. Reserved qubits are untouched.

6.2 Layer 2 — Morphogrammatic Encoding

Context-specific single-qubit RZ rotations encode each manufacturer's lever-specific strategic weight. The rotation angles are parameterized by the classical cost-function coefficients derived from the public data (financial strength, production capacity, order backlog, certification status). Each manufacturer block receives its own rotation pattern — Airbus does not use Boeing's coefficients — which is the mechanism by which the three contexts encode different objectives.

Short-range CZ entanglement within each manufacturer block encodes correlations between adjacent levers (for example, aggressive price discounting requires corresponding production capacity commitment). These intra-block gates are native on Heron and add minimal depth.

6.3 Layer 3 — Transjunctional Coupling

Transjunctional CZ gates, mediated through the 20 coupling ancillas, link homologous (lever, segment) positions across manufacturer pairs where they actually compete. When Airbus chooses aggressive narrowbody pricing and Boeing simultaneously chooses aggressive narrowbody pricing, the transjunctional coupling produces interference that redistributes amplitude away from configurations where both manufacturers simultaneously claim the same incremental market share. This is the physical embodiment of the zero-sum airline-order dynamic.

6.4 QAOA Optimization Loop

COBYLA classical optimizer drives variational parameters across iterations. Per-iteration shots: 2048 initially, increasing to 4096 for final verification runs. QAOA depth p = 1 for initial execution; p = 2 only if the p = 1 result passes quality gates (mean pairwise correlation, uniqueness ratio, top-1 bitstring share). The quality-gate thresholds are the same as those used in the Boundaries Test Report to ensure methodological continuity.

7. Noise Mitigation Strategy

At 156 qubits with the circuit depth required for three-context transjunctional coupling, noise is the central engineering risk. The Boundaries Test Report demonstrated that compiled depths above roughly 1000 layers on IBM Fez begin to fail strict quality gates even when the hardware still returns usable bitstring data. The pilot addresses this through five specific techniques applied from the start:

1. Hardware-native gate set with minimal decomposition. Use CZ (native on Heron) for all transjunctional coupling rather than CX (decomposes to two ECR gates plus single-qubit gates). Estimated compiled-depth reduction: 30–40% on transjunctional layers.
2. Transpilation at optimization_level = 3 with pre-specified initial layout. Contexts are mapped to physically adjacent qubits on the Heron coupling graph so that transjunctional gates require minimal SWAP insertion. At 156Q, SWAP inflation is typically the dominant source of depth growth.
3. Dynamical decoupling on idle qubits. Qubits not actively gated during a given layer receive XY4 or CPMG decoupling sequences, extending effective coherence time. Enabled via the Sampler primitive's built-in dynamical_decoupling option.
4. Readout error mitigation via measurement calibration. Post-processing applies the backend's published calibration matrix to the measured distribution. QPC's own Noise Reducer toolkit is applied on the aggregated counts for consistency with prior QPC publications.
5. Staged execution with graceful fallback (Section 8). If the 156Q run produces noise-dominated output, the smaller-scale runs provide a scaling curve that remains a defensible deliverable.

8. Staged Execution Plan

The pilot is executed in five staged runs of increasing complexity. This produces five data points rather than one, provides graceful fallback if noise dominates at full scale, and turns the scaling behavior itself into a commercial argument.

Stage	Configuration	Qubits	What it proves	Expected quality
S1	Airbus-only, three segments	38	Per-manufacturer baseline optimization	High (standard QPC envelope)
S2	Boeing-only, two segments	30	Second independent baseline, different constraints	High
S3	COMAC-only, two segments	30	Third independent baseline, geographic restrictions	High
S4	Three contexts + coupling, no market ancillas	98 + 20	Transjunctional coupling effect is isolated and measurable	Moderate, watched carefully
S5	Full 156Q configuration	156	Complete polycontextural showcase with market context	Uncertain; quality gates applied

Each stage produces a JSON artifact containing job ID, backend name, counts, computed objective, feasibility statistics, and the circuit depth / 2-qubit gate count after transpilation. The five artifacts together form the pilot's primary evidence package.

If S5 produces a signal that is indistinguishable from random sampling, the pilot deliverable shifts its emphasis: the scaling story from S1 through S4 becomes the primary result, and S5 is reported honestly as a noise-limit observation rather than a strategic finding. This fallback is built into the plan — we do not need it to become real in order to have something to deliver.

9. Data Sources (Public Only)

The pilot uses public data exclusively. No proprietary sources, no NDA-protected datasets, no confidential industry reports. This is a commitment, not a fallback — it ensures the case study is fully reproducible and can be shown openly to any industrial prospect without disclosure-control friction.

9.1 Airbus Sources

- Airbus SE Annual Report and Registration Document (Euronext Paris — AIR.PA)
- Airbus monthly orders & deliveries bulletin (company investor relations site)
- Airbus Global Market Forecast (20-year outlook, updated annually)
- EASA Type Certificate data sheets for A320neo family, A330neo, A350, A220

9.2 Boeing Sources

- Boeing Company 10-K and 10-Q filings (NYSE — BA)
- Boeing Orders & Deliveries monthly data (boeing.com investor relations)
- Boeing Commercial Market Outlook (20-year forecast, updated annually)
- FAA Type Certificate data sheets for 737 MAX family, 787, 777X

9.3 COMAC Sources

- AVIC / COMAC annual reports and press releases (publicly distributed)
- Cirium and Flightglobal C919 and ARJ21 orderbook trackers
- CAAC (Civil Aviation Administration of China) certification registry
- Chinese state media and industry press for fleet-in-service figures, cited explicitly as vendor-stated

9.4 Industry-Wide Sources

- Aviation Safety Network (flightsafety.org) — accident and incident history
- IATA Annual Report and World Air Transport Statistics
- ICAO traffic and fleet statistics
- Individual airline annual reports for fleet composition and order status

All data used in the pilot will be tagged with source, retrieval date, and reliability classification (audited regulatory filing / company-reported / industry-tracker / vendor-stated). The methodology section of the final report will expose this provenance openly.

10. Phase-1 Deliverables and Project Phases

This document (Phase 1) is the specification. Subsequent phases deliver concrete artifacts:

Phase	Duration	Deliverable	Status
Phase 1	Week 1	Problem definition document (this document)	Draft in review
Phase 2	Week 2	Structured JSON dataset for all three manufacturers	Pending
Phase 3	Weeks 3–4	Python QPC circuit and QAOA loop for all five stages	Pending
Phase 4	Week 5	Five JSON execution artifacts from IBM Fez runs	Pending
Phase 5	Week 6	Industrial case-study report (Crash Detection format)	Pending

Estimated effort from the owner: 4–6 hours per week. Estimated IBM Fez runtime: approximately 8 minutes of compute across five staged runs plus queue wait time (typically several hours of wall clock).

11. Framing and Honesty Commitments

The final report will state clearly that this pilot is a public-data industrial case study, not a named-customer engagement. The intended subtitle is: "Automotive and Aerospace Competitive Strategy Analysis — Airbus, Boeing, COMAC 2026–2029 — Industrial Case Study Using Public Data, Applicable to Any OEM, Lessor, Airline, or Aerospace Defense Prime."

The report will distinguish three categories of result rigorously:

- **Measured result:** any figure, number, or ranking produced directly from the JSON execution artifact from IBM Fez.
- **Derived result:** any finding computed by post-processing the measured bitstrings through the QPC reducer or classical interpretation layer.
- **Interpretation:** strategic readings offered by the author on the basis of the measured and derived results. These are clearly labeled and kept separate from the measured evidence.

No figure in the report will be synthetically interpolated. If the hardware execution does not produce a given curve or statistic, the report says so explicitly, and the corresponding figure is either omitted or presented as a fallback visualization with clear labeling. This commitment is a direct lesson from the GIC 2026 pre-submission review and is non-negotiable.

The report will also state the pilot's limitations openly: QAOA is a heuristic search without global-optimality certificate; shot noise produces run-to-run variation; the three-context snapshot optimization is an equilibrium analysis rather than a dynamic game simulation; the 2026–2029 horizon is as far as defensible forecasting extends with current orderbook data.

12. Go / No-Go Gate for Proceeding to Phase 2

Before data collection (Phase 2) begins, this document should be reviewed and explicitly approved. The approval gate asks four questions:

1. Does the three-manufacturer, three-segment, five-lever structure capture the competitive reality as the owner understands it, or do decision variables need to be added or removed?
2. Is the 156Q allocation acceptable, or should a different balance be chosen between decision qubits, coupling ancillas, market ancillas, and reserved qubits?
3. Are the geographic constraint structure and the public-data commitment acceptable as framed?
4. Are the staged-execution plan (S1 through S5) and the honesty commitments (no synthetic figures, explicit measured / derived / interpretation separation) agreed as binding?

Upon affirmative answer to all four, Phase 2 begins. If any answer is negative, this document is revised to version 1.1 before data work commences. No data collection, no code, no hardware submission starts until Phase 1 is locked.

— End of Phase 1 Problem Definition, Version 1.0 —

Quantum Polycontextural Computation (QPC) · readytogo@quantumpolycontextural.ai