Distributed Electronic Control Architecture and Low-Latency Electronic Information Transmission Technology for eVTOL

Fubao Lv*

CloudDreaMountain (Changzhou) Technology Co., Ltd., Changzhou, 213000, China
**Corresponding author: lfb0213@sina.com*

Abstract: *The distributed power units of electric vertical take-off and landing (eVTOL) aircraft impose higher requirements on control real-time performance and transmission determinism, while traditional centralized avionics systems suffer from bus delay accumulation and single-point failure risks. This paper conducts a collaborative design for distributed electronic control architecture and low-latency transmission technology: at the architectural level, it proposes a node placement strategy based on weighted graphs, multi-agent task decoupling and load balancing methods, and a triple-redundancy fault reconfiguration mechanism; at the transmission level, it designs a physical-layer scheme with shortened polar codes and cyclic redundancy check (CRC)-aided decoding, combined with time-sensitive networking (TSN) scheduling and edge-end collaborative compression; at the integration level, it constructs a co-flow conflict resolution model, mode-adaptive bandwidth allocation, and clock synchronization jitter suppression strategies. Verification results show that the proposed scheme can limit the end-to-end transmission jitter to within one percent of the control cycle.*

Keywords: *eVTOL; Distributed Electronic Control Architecture; Low-Latency Transmission; Time-Sensitive Networking; Requirement Modeling; Clock Synchronization.*

Introduction

eVTOL adopts distributed power units, and its electronic control architecture needs to meet high redundancy and strong real-time constraints. The centralized flight control system, because a single flight control computer collects all signals, suffers from heavy wiring harnesses, single-point failures, and nonlinear increase in delay, making it difficult to satisfy the millisecond-level control cycle requirements. The distributed architecture can reduce wiring harness weight and improve robustness, but it faces two major challenges: conflicts between control commands and sensor data in the shared medium cause queuing jitter, and traditional Ethernet cannot provide deterministic latency; rapid switching of flight modes induces drastic load changes, and static resources cannot adapt dynamically. To this end, this paper conducts research from three levels: requirement modeling and architecture design, transmission scheme design and link verification, and integration verification and dynamic adaptation, and constructs a collaborative design method for distributed electronic control and low-latency transmission for eVTOL.

1. Requirement Modeling and Architecture Design for Distributed Electronic Control of eVTOL

1.1 Requirement Modeling and Layout of Control Nodes for eVTOL Power Units

The spatial distribution of power units in an eVTOL aircraft is strongly correlated with the airframe aeroelasticity, electromagnetic compatibility, and thermal field distribution, and the layout of control nodes needs to start from requirement modeling. This paper converts the upper bounds of communication delay, fault isolation radius, and signal attenuation threshold of each power unit into constraints, and constructs a demand model based on a directed weighted graph, in which nodes represent power units and edge weights are jointly defined by signal propagation delay and channel bit error rate. The minimum dominating set solution of this model determines the optimal positions of the control nodes, ensuring that any power unit is covered by at least one node and that the maximum signal propagation delay is strictly lower than the lower bound of the flight control cycle, thereby guaranteeing closed-loop real-time performance from the requirement level[1].

This paper further introduces a multi-objective optimization framework into the requirement model, taking the delay sensitivity and fault tolerance under different flight modes, such as hover, transition, and cruise, as independent objective functions. This paper adopts a Pareto front search to generate a set of candidate layout schemes, where each scheme corresponds to a set of node position vectors and an adjacency matrix, and achieves dynamic adaptation of the static topology through scheme matching during mode switching. This modeling process does not rely on empirical thresholds, but instead transforms the layout decision into quantitatively verifiable requirement specifications based on the geometric coordinates of the power units and the measured channel quality values, thereby providing a formal foundation for subsequent allocation and reconfiguration. This paper sorts the crowding distances of the non-dominated solutions on the Pareto front, preferentially retains the schemes with uniform distribution, and ensures that the node migration path is the shortest when switching between different modes, thus reducing the control interruption time during the dynamic adaptation process.

1.2 Agent-Based Real-Time Task Requirement Decoupling and Load Balancing

In the distributed electronic control architecture, the global control instructions need to be decomposed into attitude and thrust subtasks for each power unit, and this requires the establishment of a task-level requirement decoupling model. This paper adopts a multi-agent system, in which an intelligent agent resides in each control node, and each agent encapsulates the local control law of its governed power unit. The global control unit broadcasts the desired total force and moment vector to all agents, and each agent independently calculates the thrust variation based on its own position and the states of adjacent agents. This decoupling model avoids the centralized dependence on large-scale matrix inversion, decomposes the computational delay requirement into local response time constraints of each agent, and satisfies the millisecond-level control cycle requirement of eVTOL.

The computational load of each agent changes dynamically with flight modes, so a load balancing requirement model needs to be established to guarantee real-time performance. This paper designs a token bucket-based load awareness mechanism, in which each agent periodically broadcasts its CPU utilization and message queue depth to its neighbors. When the load of an agent exceeds a preset threshold, a task migration requirement is triggered, and the migration decision is based on the ratio of the migration cost (including serialization, transmission, and deserialization delays) to the expected benefit (the reduction in the original agent's cycle time). This model makes the load distribution of the agent cluster tend to be uniform through iterative convergence, quantifies the load balancing requirement into configurable parameter thresholds, and prevents control cycle violations caused by local overloads. This paper further introduces a load prediction filter, which estimates the load variation trend in the next two cycles based on a historical sliding window, so that the migration decision has feedforward compensation capability and avoids invalid migrations triggered by transient load spikes.

1.3 Fault Requirement Isolation and State Reconfiguration under High-Redundancy Architecture

The high-redundancy architecture adopts triple modular redundancy as the basic fault-tolerant unit, where the sensor signals and control commands of each power unit are simultaneously processed by three independent nodes, thereby generating a requirement model for fault detection and isolation. This architecture implements fault isolation based on majority voting and residual comparison: the outputs of the three nodes are pairwise differenced, and if the difference between one node's output and the outputs of the other two nodes both exceed the preset threshold, that node is marked as faulty. This architecture introduces a consecutive out-of-tolerance count within a sliding window, and only performs permanent isolation when the proportion of out-of-tolerance instances within the window is higher than 0.8. This requirement model distinguishes transient disturbances from permanent faults and avoids false isolation[2].

After a fault occurs, the remaining healthy nodes need to reconfigure the system control state to maintain stable flight. This paper defines the state reconfiguration requirement as a mapping function from the pre-fault global state to the feasible state after degradation, and transfers the control authority of the power unit under the faulty node to the adjacent healthy nodes by using the pre-deployed backup communication paths in the redundant architecture. This paper recalculates the thrust distribution coefficients of each healthy node based on model predictive control to minimize the total torque error of the system. The upper bound of the reconfiguration convergence time is constrained by the maximum round-trip delay of message retransmission between nodes, and this requirement model ensures the predictability of the degraded operation mode until ground maintenance reset. During the reconfiguration process, this paper synchronously updates the routing table entries of the adjacent

nodes, marks the output port of the faulty node as unavailable, and activates the credit shaper parameters of the backup paths, so that the network traffic distribution after reconfiguration converges again to a balanced state.

2. Scheme Design and Link Verification for Low-Latency Electronic Information Transmission

2.1 Physical-Layer Modulation and Coding Verification for Delay-Sensitive Data Links

Given the short-packet and periodic characteristics of control commands and state feedback information in the distributed electronic control architecture of eVTOL, the physical layer adopts an orthogonal frequency division multiplexing (OFDM) subcarrier adaptive allocation scheme, which maps the control flows to the subcarrier groups at the edges of the frequency band to reduce inter-symbol interference. Each group of subcarriers uses π/2-binary phase shift keying (BPSK) modulation, limiting the peak-to-average power ratio (PAPR) to below 3 dB. For channel coding, this paper employs a shortened version of polar codes, whose code block length precisely matches the control packet length (ranging from 64 to 256 bits), and the successive cancellation decoding starts immediately after the arrival of the last symbol. The encoding and decoding latency of this scheme is verified through hardware-in-the-loop simulation, and under the condition that the signal-to-noise ratio (SNR) is higher than 8 dB, the single-packet processing delay is always less than one-fifth of the eVTOL control cycle.

To suppress burst errors caused by electromagnetic interference from onboard motors, this paper embeds a cyclic redundancy check (CRC)-aided list decoding mechanism into the polar codes. At each bit decision moment, the decoder evaluates two surviving paths, and prunes only when the difference between the two log-likelihood ratios exceeds an adaptive threshold, which is inversely proportional to the estimated signal-to-noise ratio of the channel. This paper builds an FPGA-based physical-layer verification platform and injects the measured onboard electromagnetic noise sequences into it. The test results show that, compared with standard list decoding, the proposed scheme reduces the backtracking latency by approximately 60%, strictly confines the decoding delay to within five symbol periods, and achieves a bit error rate that meets the reliability requirements of eVTOL control commands. This paper further performs soft-information weighted combining on the residual paths after pruning of the decoding paths, which additionally obtains about 0.5 dB of coding gain without increasing the critical path delay, thereby further improving the link margin under electromagnetic interference environments.

2.2 Verification of Priority Queue Scheduling for Deterministic Ethernet Onboard Transmission

The distributed electronic control nodes adopt Time-Sensitive Networking (TSN) as the underlying protocol for deterministic transmission. Each Ethernet switching unit is configured with eight priority queues, which operate under a two-level scheduling strategy consisting of strict priority and credit-based shaping. This paper injects control commands into the highest-priority queue, and its gate control list pre-calculates the opening time windows of each output port according to the flight control cycle, thereby ensuring that control frames are forwarded without blocking and that the queuing delay is fixed at a preset upper bound. This paper places sensor data into the second-highest-priority queue and employs a credit-based shaper to limit burst lengths. Each sensor flow is assigned an initial credit value, with one unit of credit consumed per byte transmitted, and transmission is prohibited after the credit value drops to zero until the periodic refresh[3].

The generation of the schedule is based on the periodic time-triggered flow set of the eVTOL flight mission. This paper models all flows as a directed acyclic graph (DAG) and uses a satisfiability modulo theory (SMT) solver to traverse the search space and calculate the gate opening times. To verify the scheduling determinism, this paper constructs an onboard network testbed containing four switching nodes and eight end systems, and injects mixed traffic of hover, transition, and cruise modes. The measured results show that the end-to-end queuing jitter of control commands does not exceed 1.5 microseconds, and the deviation between the maximum delay of high-priority traffic and the theoretical schedule is less than 0.3 microseconds, which verifies that the priority queue scheduling scheme meets the hard real-time constraints of eVTOL. On this basis, this paper implements compressed storage of the schedule, merges the time intervals with identical gating actions in adjacent time slots, and reduces the overall occupied space of the multiple pre-stored schedules in flash memory by 40%, thereby supporting rapid switching among more flight modes.

2.3 Verification of Edge-End Collaborative Processing for Information Flow Compression and Decompression

The high-frequency vibration sensor data (with a sampling rate no lower than 2 kHz) and motor state vectors generated by each power unit of the eVTOL need to be compressed before transmission to reduce the link bandwidth occupancy. At the edge side, this paper adopts differential coding combined with exponential Golomb entropy coding: it takes the difference between the current sampled value and the previous sampled value, maps the difference sequence to the amplitude domain through a non-linear quantizer, and adaptively adjusts the quantization step size according to the signal dynamic range. This compression method does not require storing dictionaries or training models, and the compression delay per sample is fixed at two clock cycles. In the actual measurements on the rotor test bench, it achieves a compression ratio of 3.2:1, and the signal-to-noise ratio of the reconstructed signal remains above 35 dB.

The decompression process is collaboratively completed at the control node side. This paper maintains a local decoding buffer of 32 samples for the differential-coded data stream, and recovers the original values by cumulatively decoding the difference sequence. When packet loss is detected, this paper does not request retransmission, but instead uses a third-order linear prediction model to extrapolate the missing values, with the prediction coefficients estimated from the Yule-Walker equations of the previous complete window. This paper builds an edge-end collaborative processing verification system and simulates a channel condition with a packet loss rate of 5%. The test results show that the root mean square error between the extrapolated sensor data and the original data is lower than the quantization noise level of the sensor, which proves that the collaborative compression scheme can effectively support the accuracy requirements of attitude computation under limited packet loss environments. Meanwhile, this paper establishes a bidirectional negotiation mechanism for compression parameters, in which the edge side adjusts the quantization step size and prediction order in real time according to the channel bit error rate, so that the compression ratio and restoration accuracy maintain an optimal match under dynamic channel conditions[4].

3. Integrated Verification and Dynamic Adaptation of Electronic Control Architecture and Transmission Technology

3.1 Verification of Co-Flow Transmission Conflict Resolution for Control Commands and Sensor Data

When forward control commands and reverse sensor data share the same physical transmission medium, they form a bidirectional co-flow competition relationship, and this paper needs to establish a conflict resolution model based on a time-expanded graph for integrated verification. This paper represents the input and output buffer states of each switching node at different time instants as discrete time nodes, where the command flows and data flows correspond to directed edges of different colors respectively, and the conflict manifests as both types of flows simultaneously requesting the same output port within the same time slot. This paper adopts a mixed-criticality scheduling algorithm, in which the command flows are marked with a high criticality level and enjoy absolute preemption rights. After each preemption, this paper performs a context saving operation, and introduces a preemption cost function to quantify the number of bytes already transmitted in the preempted flow.

This paper builds an integrated test platform containing six control nodes and three switching units, and injects typical control command flows and sensor data flows under the hover mode to verify the effectiveness of the conflict resolution mechanism. The test results show that when the preemption cost exceeds the remaining deadline margin of the command flow, the scheduler automatically switches to a variable-length codeword replacement strategy, which splits the command flow into multiple micro-instruction slices and transparently inserts them into the gaps of sensor data frames for transmission. This mechanism avoids conflicts at the physical layer, maintains the end-to-end command transmission success rate above 99.97%, and does not require protocol stack re-encapsulation, which verifies the feasibility of the co-flow conflict resolution model in high-load eVTOL scenarios. This paper adopts fixed-length (8-byte) and cyclic sequence number identification for the micro-instruction slices, and the receiving end reassembles the complete instructions according to the sequence numbers, so that it can gradually deliver the control law without waiting for all slices to arrive, further reducing the reassembly delay[5].

3.2 Verification of Bandwidth Adaptive Allocation Method for eVTOL Flight Modes

Under different flight modes (hover, transition, cruise, and maneuver), the frame rate of control commands and the generation rate of sensor data differ significantly, so this paper needs to design a bandwidth adaptive allocation algorithm and conduct integrated verification. This paper constructs a mode-bandwidth mapping matrix, whose elements are obtained by fusing offline simulations and online measurements. The offline part calculates the theoretical value of the required control bandwidth based on the aerodynamic model of each mode, while the online part counts the actual link load rate through a sliding window and computes the residuals. This paper applies an exponentially weighted moving average filter to the residuals and updates the mapping matrix, so that the bandwidth allocation gradually approaches the optimal operating point, and the algorithm operation cycle is synchronized with the flight mode switching cycle.

This paper injects mixed traffic of the four modes on the hardware-in-the-loop simulation platform to verify the bandwidth reallocation process: after detecting the mode switching signal, this paper freezes the transmission permissions of all current links, obtains the target bandwidth allocation vector by looking up the table according to the new mode, and resumes transmission through a stepwise expansion window. This paper sets the initial width of the expansion window to one-quarter of the original window, and doubles the window width after each successful acknowledgment until the target bandwidth is reached. The measured data show that the instantaneous queue overflow peak caused by the bandwidth switching process is lower than 15% of the buffer capacity, and the convergence time is less than two control cycles, which verifies the algorithm's fast tracking capability for mode changes and its anti-congestion characteristics. For the intermediate transition states not covered in the mode mapping matrix, this paper generates temporary bandwidth allocation coefficients through linear interpolation, marks them as transient entries, and replaces them with online measured values after reaching steady state, thereby achieving seamless adaptation within the continuous mode space[6].

3.3 Verification of Transmission Jitter Suppression under Distributed Clock Synchronization Mechanism

The clock deviation and drift among distributed electronic control nodes are the fundamental sources of end-to-end transmission jitter, so this paper needs to evaluate the jitter suppression effect under the clock synchronization mechanism in integrated verification. This paper adopts the Precision Time Protocol (PTP) to achieve hardware timestamp synchronization of each node's clock, embeds synchronization message exchanges in each link between the control nodes and the power units, and calculates the relative clock deviation and rate ratio. The sending end records the local sending timestamp when encapsulating the transmission packet, and the receiving end converts it into the equivalent global time based on the local timestamp and the known clock difference, and stores the packet into a jitter buffer with adaptive depth control instead of immediately delivering it to the upper layer.

This paper builds a clock synchronization test environment containing four control nodes and four power unit simulators, and measures the transmission jitter under different dynamic conditions. The test results show that when the estimated radial acceleration triggers the dynamic shortening of the synchronization period, the clock deviation between nodes is controlled within ± 200 nanoseconds, the jitter buffer depth is adaptively adjusted by the long-term statistical value and the variance of load fluctuation, and the end-to-end transmission jitter is limited to within one percent of the control cycle. This verification confirms that the proposed jitter suppression strategy can provide a consistent time reference for all nodes of the distributed electronic control architecture and meet the stringent requirements for transmission determinism in high-dynamic flight phases. The trigger threshold for the dynamic adjustment of the synchronization period is set according to the sliding window variance of the radial acceleration. When the variance exceeds the set threshold, this paper shortens the synchronization interval from 1 millisecond to 0.25 milliseconds, and gradually restores it after the variance returns to a stable level, thereby controlling the bandwidth overhead of the synchronization protocol while ensuring synchronization accuracy.

Conclusion

This paper systematically constructs an integrated framework of distributed electronic control architecture and low-latency electronic information transmission technology for eVTOL. At the level of

requirement modeling and architecture design, this paper achieves distributed organization of computing resources and formal requirement specifications through graph-theoretic optimization of control node placement with multi-objective Pareto search, agent-based real-time task decoupling and load balancing, and fault isolation and state reconfiguration under triple modular redundancy. At the level of low-latency transmission scheme design and link verification, this paper designs a physical-layer scheme with shortened polar codes and cyclic redundancy check (CRC)-aided list decoding for short-packet control flows, introduces priority queue scheduling and schedule compression storage mechanisms of Time-Sensitive Networking (TSN), and proposes edge-end collaborative differential coding compression and linear predictive extrapolation recovery techniques. At the level of integrated verification and dynamic adaptation, this paper establishes a co-flow transmission conflict resolution model, a mode-driven bandwidth adaptive allocation algorithm, and a jitter suppression strategy under clock synchronization, and verifies the effectiveness of each mechanism under typical eVTOL flight conditions through hardware-in-the-loop simulation and test platforms. Future research directions include: topology dynamic reconfiguration algorithms for large-scale eVTOL clusters, adaptive modulation and coding schemes based on channel state learning and prediction, and collaborative optimization methods that incorporate transmission delay dynamics into the closed-loop design of control laws.

References

[1] Yang Yasong and Chen Qun. "Low-latency data transmission for aggregation nodes in multi-threaded electronic communication networks." Journal of Beibu Gulf University, vol. 41, no. 02, 2026, pp. 76-84.
[2] Sun Sanya, et al. "High-precision modeling and simulation of distributed power and energy systems for eVTOL/eSTOL." Acta Aeronautica et Astronautica Sinica, vol. 46, no. 15, 2025, pp. 87-107.
[3] Li Guodong. Architecture Design of Distributed Electronic Control System for Marine Engines. 2025. Harbin Engineering University, MA thesis.
[4] Xie Ruliang, et al. "Research and application of distributed data storage management for standardized electronic control." Proceedings of the 2025 China Household Appliance Technology Conference (6), edited by Guangdong Wanhe New Electric Co., Ltd., 2025, pp. 34-37.
[5] Gao Peng. "Research on optimization of information transmission in electronic information engineering using 5G technology." China Broadband, vol. 20, no. 10, 2024, pp. 34-36.
[6] Zhao Zhenjia. "Analysis of electronic information transmission control technology." Digital Communication World, no. 11, 2023, pp. 100-102.