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 unit of the electric vertical take-off and landing (eVTOL) aircraft should have good real-time control performance and transmission determinism, and the traditional centralised avionics system will be prone to bus delay accumulation and single points of failure. Collaborative Design for a Distributed Electronic Control Architecture and Low-Latency Transmission Technology: Weighted graphs are used to determine the position of nodes in the architecture, multi-agent task decoupling and load-balancing schemes are introduced, a triple-redundancy fault reconfiguration mechanism has been established; at the transmission level, a physical-layer scheme with shortened polar codes and CRC-aided decoding has been developed, and it is combined with Time-Sensitive Networking (TSN) scheduling and edge-end collaborative compression; finally, at the integration level, a co-flow conflict resolution model, mode-adaptive bandwidth allocation and clock synchronization jitter suppression strategies have been developed. Based on the above verification results, the total transmission jitter of the proposed scheme will be less than 1% of the control period.*

Keywords: *eVTOL; Distributed Electronic Control Architecture; Low-latency Transmission; Time-sensitive Networking; Requirement Modeling; Clock Synchronization*

Introduction

EVTOLs have distributed power units, so the electronic control system needs to be very reliable and have good real-time performance. The central flight control system needs to collect all the signals in one place; otherwise, there will be many wires, single points of failure and a large increase in delay, and it will not meet the requirement of millisecond-level control cycle. Although the distributed structure can reduce the weight of wiring harnesses and improve robustness, there are still two main problems: conflicts between control commands and sensor data in the shared medium cause queuing jitter, and traditional Ethernet cannot guarantee deterministic latency; rapid switching of flight modes results in a large fluctuation in load, and static resources cannot be adjusted dynamically. Therefore, the three stages of research in this paper will be requirement modelling and architecture design; transmission scheme design and link verification; and integration verification and dynamic adaptation, and together they will form a collaborative design method for distributed electronic control and low-latency transmission of 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 position of the power unit in an eVTOL aircraft is related to the aeroelasticity of the airframe, electromagnetic compatibility and heat distribution; thus, the placement of control nodes should be determined according to the requirement model. The upper bounds of the communication delay, fault isolation radius and signal attenuation threshold for each power unit are converted into constraints, and a demand model based on a directed weighted graph is constructed; the nodes are the power units, and the edge weights are jointly defined by the signal propagation delay and the channel bit error rate. The minimum dominating set solution of this model finds the best location for the control node; that is to say, all power units must be within the range of at least one node, and the maximum signal propagation delay should be smaller than the lower bound of the flight control cycle to meet the requirements for closed-loop real-time performance at[1].

Add a multi-objective optimisation framework to the requirement model in this paper and set the delay sensitivity and fault tolerance in the different flight modes of hover, transition and cruise as separate objective functions. Pareto front search is used to find a set of candidate layout schemes, and each one is associated with a set of node position vectors and an adjacency matrix; dynamic adaptation of the static topology is achieved by scheme matching in mode switching. There is no empirical threshold; based on the geometric coordinates of the power units and the measured channel quality values, the layout decision will be converted into quantifiable requirement specifications and formally approved for the next allocation and reconfiguration. Sort the crowding distances of the non-dominated solutions on the Pareto front, prefer schemes with a uniform distribution, make the node migration path as short as possible when switching between different modes, and thus reduce the control interruption time during dynamic adaptation.

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

All the control commands in a distributed electronic control system need to be divided into attitude and thrust subtasks for each power unit, and a decoupling model should be set up at the task level. A Multi-agent System is employed in this paper, and at each control node, an intelligent agent stores the local control law of the power unit it is responsible for. The World Control Unit distributes the total force and moment vector required by all units, and then each unit independently decides how much thrust to produce based on its position and the state of nearby units. Large-scale matrix inversion does not need to be carried out in the decoupling scheme; instead, the computation delay requirement will be distributed among the local response times of all agents to meet the millisecond-level control cycle demand of eVTOLs.

The computation demands of all the agents are not the same in different flight modes, so at the same time, a load-balancing scheme needs to be introduced. Design a Token Bucket Load Awareness Mechanism in this paper, and have each agent periodically publish its CPU utilisation and message queue length to the neighbouring nodes. If the load of an agent exceeds a certain threshold, then task migration needs to be carried out; otherwise, only if the migration cost (serialization, transmission and deserialisation delay) is less than the expected gain (reduction in the original agent's cycle time) will it be done. Gradually spread the load of the agent cluster more evenly among the models, quantify the necessity of load balancing based on the threshold of the preset parameter, and avoid control-cycle violations caused by local overload. A load prediction filter is also employed in this paper to forecast the trend of load fluctuation in the next two cycles based on a historical sliding window, and thus the migration decision will be pre-compensated for; migration failure caused by a temporary increase in load will also be avoided.

1.3 Fault Requirement Isolation and State Reconfiguration in a High-Redundancy Architecture

First, we can use the concept of Triple Modular Redundancy (TMR); that is to say, for each power unit, both the sensor signal and the control command are obtained at the same time from the three independent modules to meet the fault detection and isolation requirement model. The structure is fault isolation based on majority voting and residual comparison; that is to say, 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 a preset threshold, then that node is considered faulty. There is a sliding window to count consecutive out-of-tolerance items, and only when the proportion of out-of-tolerance instances in the window exceeds 0.8 will they be permanently excluded. The Model of Need distinguishes between temporary interruptions and prolonged malfunctions, but it does not make this division[2].

If there is a failure, the remaining good nodes need to switch to a different system control mode and maintain the stability of the aircraft. The purpose of the state remapping in this paper is to map the pre-fault global state to the feasible state after degradation, and at the same time, it is necessary to transfer the control authority of the power unit at the faulty node to the adjacent healthy nodes via pre-deployed backup communication paths in the redundant architecture. With the help of Model Predictive Control, we will recompute the thrust distribution coefficients for all healthy nodes and reduce the total torque error in the system. The upper bound of the reconfiguration convergence time is the maximum round-trip delay of message retransmission among all nodes, and this requirement model guarantees that the degraded operation mode will be stable before ground maintenance is carried out. At the same time, during reconfiguration, this paper will update the routing table entries of the neighbouring nodes, mark the output port of the faulty node as unavailable, and enable the credit shaper parameters for the backup paths; thus, after reconfiguration, the distribution of network traffic will be balanced again.

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

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

Given that the control commands and state feedback information in the distributed electronic control architecture of eVTOLs have a short packet size and are periodic, an OFDM subcarrier adaptive allocation scheme is used at the physical layer to map the control flow onto subcarrier groups at the edges of the frequency band and reduce inter-symbol interference. Each group of subcarriers is modulated by $\pi/2$-BPSK, and the peak-to-average power ratio (PAPR) is less than 3 dB. Polar codes are shortened for channel coding in this paper, and the code block length is the same as that of the control packet (64-256 bits); successive cancellation decoding starts immediately after the arrival of the last symbol. Based on Hardware-in-the-Loop simulation, it has been determined that the encoding and decoding latency of this scheme is relatively small, and when the signal-to-noise ratio (SNR) is greater than 8 dB, the single-packet processing delay will always be less than one-fifth of the eVTOL control period.

To reduce the burst errors caused by electromagnetic interference from the motor on board, a CRC-assisted list decoding scheme will be added to the polar code in this paper. At the time of the bit decision, the decoder considers the two remaining paths and prunes one only if the difference between the two log-likelihood ratios exceeds an adaptive threshold; this threshold is inversely proportional to the estimated signal-to-noise ratio of the channel. Construct an FPGA-based Physical-Layer Test and Measurement Platform and place the measured sequence of on-board electromagnetic noise in it in this paper. According to the results of the above tests, the proposed scheme can reduce the backtracking latency by about 60% compared with that of the standard list decoding, keep the decoding delay within the range of five symbol periods, and achieve a bit error rate that meets the reliability requirements for eVTOL control commands. Softly combine the weights of the residual paths after pruning the decoding paths in this paper to achieve a coding gain of about 0.5 dB without increasing the critical path delay and thus improve the link margin in the presence of electromagnetic interference.

2.2 Verification of Priority Queue Scheduling for Deterministic Ethernet Onboard Transmission

The distributed electronic control nodes will be based on Time-Sensitive Networking (TSN) and will use it for deterministic transmission. The 8 priority queues of an Ethernet switch are divided into two levels, and there is strict priority and credit-based shaping. Add the control commands to the high-priority queue in this paper, and the gate control list will pre-calculate the opening time windows of each output port according to the flight control cycle to ensure that control frames are forwarded without blocking and the queuing delay is limited to a fixed upper bound. The sensor data should be in the second highest-priority queue, and a credit-based shaper is used to limit the burst length. All sensor flows start with a credit; one credit is used per transmitted byte, and transmission is not allowed if the credit reaches zero before the refresh in the next period[3].

The generation of the schedule is based on the periodic time-triggered flow set of the eVTOL flight mission. All flows are modelled as a directed acyclic graph (DAG), and then a satisfiability modulo theory (SMT) solver is used to search the space and find the gate opening times. To verify the determinism of scheduling, an onboard network testbed consisting of four switching nodes and eight end systems has been built, and mixed traffic in hover, transition and cruise modes has been introduced. Based on the above measurements, the end-to-end queuing jitter of control commands is less than 1.5 microseconds, and the deviation between the maximum delay of high-priority traffic and the theoretical schedule is less than 0.3 microseconds; therefore, it can be confirmed that the priority queue scheduling scheme meets the hard real-time requirements of eVTOL. Therefore, compressed storage is employed to group the time periods in which there are the same gating actions in adjacent time slots, reducing the total occupied space of several pre-stored schedules in flash memory by 40% and enabling faster switching among a larger number of flight modes.

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

Compress the high-frequency vibration sensor data (at a sampling rate of 2 kHz or higher) and the motor state vectors of each power unit in the eVTOL before transmission to reduce the bandwidth demand of the link. At the edge end, Differential Coding is employed with exponential Golomb entropy coding; that is to say, the difference between the current sampled value and the previous sampled value is taken, this sequence of differences is mapped to the amplitude domain by a non-linear quantizer, and the quantization step size is adjusted according to the dynamic range of the signal. Compression does not need to save dictionaries or train models, and the compression delay per sample is fixed at two clock cycles. Based on the actual measurement at the rotor test bench, the compression ratio is 3.2:1 and the signal-to-noise ratio of the reconstructed signal is still over 35 dB.

The Decompression Process will be carried out at the control station. A local decoding buffer of 32 samples is used for the differential-coded data stream, and then the original values are obtained by cumulative decoding of the difference sequence. If there is packet loss, this paper does not request retransmission but rather uses a third-order linear prediction model to estimate the missing data based on the prediction coefficients obtained from the Yule-Walker equations of the previous complete window. Construct an edge-end collaborative processing verification system and simulate a channel condition of 5% packet loss in this paper. Based on the test results above, the root mean square error of the extrapolation between the sensor data and the original data is smaller than the quantization noise of the sensor; therefore, it can be concluded that the collaborative compression scheme meets the accuracy requirements for attitude computation in the presence of limited packet loss. At the same time, this paper proposes a bidirectional negotiation mechanism for compression parameters; that is, according to the channel bit error rate, the edge side will dynamically change the quantization step size and prediction order to find the best trade-off between compression ratio and restoration accuracy under different channel conditions[4].

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

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

If the physical transmission medium of the forward control command and the reverse sensor data is the same, there will be a problem of bidirectional co-flow competition; thus, this paper will build a conflict resolution model based on a time-expanded graph for joint verification. The input and output buffer states of each switching node at different times are taken as the discrete-time nodes in this paper, and the command flow and data flow are shown by directed edges of different colours; a conflict occurs if both types of flows try to use the same output port at the same time slot. A mixed-criticality scheduling algorithm will be used in this paper to give the command flow a high-criticality label and grant it full preemption rights. After each preemption, this paper will save the context and add a preemption cost function to quantify how many bytes have been transmitted in the preempted flow so far.

A test platform with six control nodes and three switching units has been constructed in this paper, the typical control command flow and sensor data flow in hover mode have been presented, and tests of the conflict resolution mechanism have also been carried out. Based on the above test results, if the preemption cost exceeds the remaining deadline margin of the command flow, a variable-length codeword replacement strategy will be automatically selected by the scheduler; that is to say, the command flow will be divided into several micro-instruction slices and these slices will be inserted transparently in the gaps of the sensor data frame for transmission. Therefore, there will be no conflict at the physical layer, the end-to-end command transmission success rate will be above 99.97%, and protocol stack re-encapsulation is not required; thus, it has been confirmed that the co-flow conflict resolution model can be applied to high-load eVTOL scenarios. Fixed-length (8-byte) identification and cyclic sequence number identification are used to recognise the micro-instruction slices, and at the receiving end, the complete instruction is reconstructed based on the sequence number; thus, the control law can be delivered incrementally without waiting for all slices to arrive, and the reassembly delay is reduced[5].

3.2 Verification of Bandwidth Adaptive Allocation Method for eVTOL Flight Modes

The frame rate of the control commands and the generation frequency of the sensor data are not the same in the four flight modes (hover, transition, cruise and manoeuvre), so a bandwidth-adaptive allocation algorithm and integrated verification need to be developed in this paper. Based on the results of offline simulation and online measurement in this paper, a mode-bandwidth mapping matrix is constructed. Offline: According to the aerodynamic model, calculate the theoretical demand for the control bandwidth at all operating conditions, and in the online stage, count the actual link load rate in a sliding window to obtain residual values. An exponentially weighted moving average filter is used on the residuals, and then the mapping matrix is updated; thus, the bandwidth allocation will gradually move to an optimal operating point, and the working cycle of the algorithm will be in sync with the cycle of flight mode switching.

Inject mixed traffic of the four modes on the hardware-in-the-loop simulation platform and verify the bandwidth reallocation process: After detecting the mode switching signal, freeze the transmission permissions of all current links, look up the target bandwidth allocation vector in the table according to the new mode, and then resume transmission via a step-wise expansion window. Set the initial width of the expansion window to one-fourth of the original window, and after each successful acknowledgment, double the width of the window until the target bandwidth is reached. Based on the above measured data, the maximum instantaneous queue overflow caused by the bandwidth switching process is less than 15% of the buffer capacity, and the convergence time is less than two control cycles; therefore, it can be concluded that the algorithm has good fast-tracking ability for mode change and is anti-congestion. For the intermediate transition states that are not in the mode mapping matrix, linear interpolation is employed in this paper to obtain temporary bandwidth allocation coefficients; these are referred to as transient entries, and after reaching a steady state, they are replaced by online measured values to achieve continuous adaptation in the mode space[6].

3.3 Verification of Transmission Jitter Suppression under Distributed Clock Synchronization Mechanism

Clock deviation and drift at the distributed electronic control nodes are the root causes of end-to-end transmission jitter; therefore, this paper will determine how much jitter can be reduced by the clock synchronization mechanism in integrated verification. Precision Time Protocol (PTP) is used in this paper to obtain hardware timestamping of the clock at each node, synchronization message exchange is carried out over all links between the control node and the power unit, and the relative clock deviation and rate ratio are calculated. The sending end records the local sending time at the time it encapsulates the transmission packet, and the receiving end, based on its local time and the known clock difference, converts this into the corresponding global time, then stores the packet in a jitter buffer with adaptive depth control and does not directly deliver it to the upper layer.

A clock-synchronisation test environment has been set up for the four control nodes and the four power-unit simulators, and the transmission jitter in all kinds of dynamic conditions has been measured. Based on the test results above, when the estimated radial acceleration causes a dynamic shortening of the synchronization period, the clock deviation among nodes can be limited to ± 200 nanoseconds, the depth of the jitter buffer will be adjusted according to long-term statistical data and the variance of load fluctuations, and the end-to-end transmission jitter will be less than one per cent of the control cycle. As shown in the verification results above, the proposed jitter-suppression method can provide a stable time reference for all nodes in the distributed electronic control system and meet the requirement of transmission determinism in high-dynamic flight conditions. The trigger threshold for the dynamic adjustment of the synchronisation period is the variance of the sliding window of radial acceleration. If the variance is too large, we will reduce the synchronisation period from 1ms to 0.25ms and then gradually increase it after the variance has stabilised to control the bandwidth overhead of the synchronisation protocol while maintaining the synchronisation accuracy.

Conclusion

Briefly introduce the general structure of the distributed electronic control system and the low-latency electronic information transmission technology for eVTOLs in this paper. Graph theory is used to model the requirements and design the architecture; the distribution of the organisation of computing resources is done by placing nodes according to multi-objective Pareto search, agent-based real-time task decoupling and load balancing are employed, and under triple modular redundancy, fault isolation and state reconfiguration are carried out. Design of the Low-Latency Transmission Scheme and Link Verification at the Physical Layer is presented in this paper. Shortened Polar Codes with CRC-assisted list decoding have been used to apply to the short-packet control flow, priority queue scheduling and schedule compression storage mechanism of Time-Sensitive Networking (TSN), and at the edge end, collaborative differential coding compression and linear predictive extrapolation recovery techniques have also been applied. The problem of joint verification and dynamic adaptation is introduced in this paper. A co-flow transmission conflict resolution model, a mode-driven bandwidth-adaptive allocation algorithm and a jitter suppression strategy under clock synchronization have been proposed, and their actual effects in typical eVTOL flight conditions have been confirmed by hardware-in-the-loop simulation and test platforms. In the future, research will be carried out on algorithms for topology dynamic reconfiguration of large-scale eVTOL clusters, adaptive modulation and coding schemes based on channel state learning and prediction, collaborative optimisation methods that consider the dynamics of transmission delay in the closed-loop design of control laws, etc.

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.