

Autonomic Trust Architecture for Secure Edge Intelligence in Cyber Physical Systems

Fully Funded PhD Studentship

School of Internet of Things | Xi'an Jiaotong–Liverpool University

Scholarship support

- Full tuition fee waiver.
- Living stipend of RMB 5,000 per month for 36 months, totalling RMB 180,000.
- Academic conference funding of RMB 16,500 in total for the entire PhD. This is not an annual allowance.

We invite applications from motivated candidates seeking doctoral research in this area. Scholarship administration and eligible conference expenditure are governed by the formal award terms and University procedures. The start date will be agreed in line with the relevant doctoral admissions intake.

Project overview

How can an edge intelligence system decide when to trust a device, when to restrict a workload, and how to recover after suspected compromise? This PhD will investigate an autonomic trust architecture that connects hardware-backed evidence, distributed attestation and runtime monitoring. The candidate will develop and evaluate methods for making bounded, auditable security decisions across heterogeneous edge devices. The research combines systems security with experimental implementation, focusing on a representative cyber-physical application rather than an entire production platform.

Research background and rationale

Edge AI places models close to sensors, actuators and users. Trust therefore depends on more than whether a model produces an accurate prediction: the platform, workload, data path and permissions must also be considered. A device may present valid startup evidence yet subsequently exhibit anomalous behaviour, while a runtime alert may be caused by a benign change rather than compromise.

The project will investigate how these different kinds of evidence can support defensible decisions over time. A hardware root of trust provides a basis for checking specified platform properties; remote attestation communicates evidence of those properties to a verifier. Neither mechanism alone proves that all application behaviour is safe. The research will make this distinction explicit when combining platform evidence with monitoring and response policies.

The scientific focus is the relationship between evidence freshness, trust decisions and operational cost. Candidate mechanisms will be compared with simpler alternatives to establish when adaptive verification and response offer a measurable benefit.

Research objectives and questions

The candidate will refine the research questions with the supervisory team following a critical literature review and initial feasibility experiments. The proposed questions are:

- How should evidence from platform attestation, workload identity and runtime monitoring be represented and combined without treating an anomaly score as proof of integrity?
- How can verification frequency and coordination adapt to changing risk while controlling communication cost, latency and energy consumption on constrained devices?
- How can a trust decision trigger a bounded response, such as restricting access or isolating a workload, while preserving an auditable route to recovery?

Research methods and evaluation

Threat model and experimental scope. Select a representative edge AI workload and a manageable set of supported devices. Define protected assets, trusted components and attacker capabilities, including compromised software, stale or replayed evidence and relevant input/output paths. Explicitly identify hardware and physical attacks outside the experimental scope.

Hardware anchored evidence and verification. Implement a baseline using available hardware isolation and attestation interfaces. Bind verification results to the device, workload and policy being evaluated. Investigate input/output protection where the selected platform supports it; document unprotected paths rather than assume CPU isolation extends to every peripheral.

Adaptive trust and response. Develop an evidence-combination and verification-scheduling method, then connect it to a small set of reversible response actions. Compare centralized and distributed coordination where relevant. Distributed ledgers or consensus mechanisms will be considered only if the threat model requires them.

Controlled evaluation. Use reproducible workloads, fault injection and authorized attack scenarios. Compare static and adaptive verification, and hardware-only, monitoring-only and combined approaches. Measure detection errors, stale-evidence acceptance, response and recovery time, overhead and scaling behaviour. Report repeated-run variability and ablation results; assess safety consequences in a controlled testbed.

Expected doctoral contribution

The intended doctoral contribution is a principled method for translating heterogeneous trust evidence into adaptive verification and response decisions, supported by a prototype and reproducible evaluation. The thesis should explain the security assumptions, failure modes and cost trade-offs under which the method is useful. New cryptographic algorithms, chip design and unrestricted autonomous control of safety-critical systems are outside the core scope.

Research training and environment

Close collaboration with our industry partner will be a central part of this PhD. The candidate will work with partner engineers to define use cases and validate the trust architecture with representative AI workloads. Experiments will examine security assumptions, attestation, response behaviour, performance overhead and system integration. Partner feedback will inform prototype refinement and analysis of the results.

Partner data and system access will follow agreed confidentiality, intellectual property, ethics and security arrangements. Approved benchmark and sandbox experiments will complement industry validation and support reproducibility.

Supervisory review will guide literature review, baseline experiments, method development and thesis evaluation, maintaining a focused, original doctoral contribution alongside industry collaboration.

Candidate profile

- A strong academic background in computer science, computer engineering, cybersecurity, electronic engineering or a closely related field, subject to the University's doctoral admission requirements.
- Sound knowledge of operating systems, networking and computer architecture, together with practical programming and debugging in C/C++, Rust, Go or Python.
- Evidence of research potential through a dissertation, substantial systems project or equivalent work, and the ability to design controlled experiments and explain technical decisions.
- Ability to read research literature, write clearly in English and work independently with academic guidance.

Desirable experience

Experience with trusted execution environments, embedded Linux, secure boot, remote attestation, distributed systems or edge AI is advantageous. Applicants are not expected to have prior expertise in every component.

Application and academic enquiries

Academic contact: Professor Dijiang Huang, School of Internet of Things, Xi'an Jiaotong–Liverpool University.

Email: Dijiang.huang@xjtlu.edu.cn

For an initial academic enquiry, please provide:

- A curriculum vitae describing your academic background and relevant research or technical experience.
- Academic transcripts and available degree information.
- A short statement of research interests explaining your fit with this project and a question you would like to investigate.
- Where available, a dissertation abstract, publication, technical report or repository showing your own contribution. Do not disclose confidential material.
- Your preferred start date and current study or employment commitments.

Suggested email subject: PhD Application — Autonomic Trust Architecture — Your Name

An academic enquiry does not replace the formal doctoral application. Admission and the scholarship award are subject to the University's formal selection and offer process. Detailed eligibility and application documentation should be confirmed through that process.

Research areas

Edge AI; AIoT security; confidential computing; hardware roots of trust; remote attestation; adaptive security