EMBEDDED SYSTEMS · FPGA · EMBEDDED AI

Mohamed Amine BELKACEM

I work where hardware meets software: VHDL on FPGA, analog front ends in SPICE, C down to the SIMD register, and the AI models that have to run on all of it. I treat every project as a measurement: a spec written before the first run, tests that are proven to fail when the code is broken, and numbers with their confidence intervals.

Available from March 2027 · 6 months · Work permit: student status

24	08	08
PROJECTS	COMPLETE	DOMAINS

Method

01
Spec first
Requirements and success criteria are written down before the first measurement.

02
Tests that can fail
I break the code on purpose to check that the tests catch it.

03
Honest numbers
Intervals, several seeds, and only what has been measured. Work in progress says so.

04
Reproducible
Every figure comes from a script in the repository and runs again in CI.

Path

2026 – 2027 — **M2 Communicating Electronic Systems Design (CSEC)**
University of Lorraine, Metz

2025 – 2026 — **M1 Communicating Electronic Systems Design (CSEC)**
University of Lorraine, Metz

2023 – 2025 — **Master's in Electronics, Embedded Systems**
USTHB, Algiers · thesis: connected health monitor

Jul – Aug 2024 — **Low-power LoRaWAN metering node**
CREDEG (Sonelgaz group) · ESP32 firmware for voltage and current measurement, deep-sleep state machine, LoRa-E5 module, MQTT

2020 – 2023 — **Bachelor's in Electronics**
USTHB, Algiers · thesis: smart-house solar supply

Skills

01

FPGA & digital

VHDL · Quartus Prime · Questa · GHDL · Cyclone IV · Basys 3 (Artix-7) · Gate-level simulation (SDF) · Fixed point

02

Analog & hardware

ngspice · Altium Designer · SKY130 PDK · Monte Carlo / corners · gm/Id · Current sensing

03

Embedded

C / C++ · ARM NEON · ESP32 · Arduino · Raspberry Pi · RP2040 · UART · SPI · I2C · LoRaWAN · POSIX threads

04

AI & data

PyTorch · scikit-learn · XGBoost · ONNX Runtime · int8 quantization · llama.cpp · wav2vec 2.0 · Flower / Opacus

05

Security & industrial

CPA / TVLA · Masking / hiding · OpenPLC · IEC 61131-3 ST · Modbus TCP · OPC UA

06

Method & tooling

Python · pytest · GitHub Actions · CMake · Git · Linux · Qt 6 · Flask · Flutter · MATLAB

LANGUAGES French · English · Arabic · Kabyle · Spanish

Reports & theses

2026–2027

INDIVIDUAL REPORT

FPGA co-design of a molecular dynamics simulator (bonded forces + PME)

M2 CSEC, University of Lorraine · in progress

2026

PROJECT REPORT

FP8 (E4M3) multiply-accumulate unit in VHDL for AI inference

M1 CSEC, University of Lorraine · report in French

2025

MASTER'S THESIS, DEFENDED JUNE 2025

Connected health monitor: ESP32 sensors, Raspberry Pi 5 server, web and mobile apps

USTHB, Algiers

2023

BACHELOR'S THESIS

Solar power supply for a smart-house model

USTHB, Algiers

Contact

Got a hardware problem worth measuring?

Open to internship offers in embedded systems, FPGA, electronics, embedded AI and test & validation.

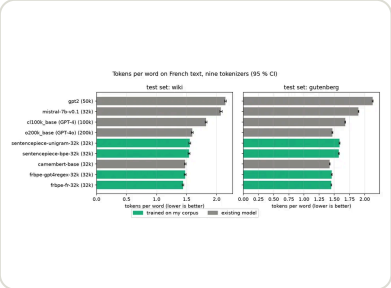
mohamed-amine.belkacem8@etu.univ-lorraine.fr · linkedin.com/in/mohamed-amine-belkacem-262582391 · github.com/amine-belkacem · amine-belkacem.github.io

All projects

AI & machine learning

● COMPLETE

French BPE tokenizer, from scratch to C



PROBLEM	How much do existing tokenizers penalise French, and can a French tokenizer run on a Raspberry Pi?
CONSTRAINT	Lossless and deterministic; the C encoder must give exactly the Python ids.
KEY RESULT	The same French text costs 27 % more tokens with GPT-4's tokenizer and 11 % more with GPT-4o's.
TOOLS	Python · C99 · tiktoken · SentencePiece · pytest · ASan / UBSan
PERIOD	Oct 2026
STATE	Complete · v1.0.1

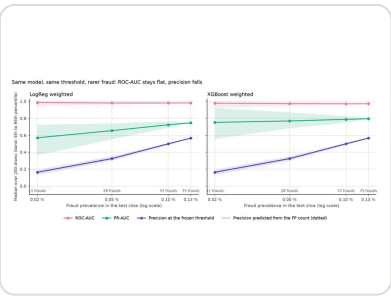
KEY RESULTS

- The same French text costs 27 % more tokens with GPT-4's tokenizer and 11 % more with GPT-4o's.
- A French-aware regex helps only a little: −2.45 % [−2.57, −2.33].
- C encoder identical to Python on 2,340 documents; 14.4 MB/s on the Pi, 8.5 MB peak memory.

github.com/amine-belkacem/french-bpe-tokenizer

● COMPLETE

Fraud detection under an alert budget



PROBLEM	Evaluate fraud models the way a fraud team uses them, at 0.17 % prevalence.
CONSTRAINT	Temporal split, thresholds fixed on validation, only 75 test frauds, so every key number has a bootstrap interval.
KEY RESULT	XGBoost catches 80 % [71, 88] of frauds at a 0.1 % false positive rate; PR-AUC 0.796.
TOOLS	Python · scikit-learn · XGBoost · imbalanced-learn · pytest
PERIOD	Oct 2026
STATE	Complete · v1.0.0

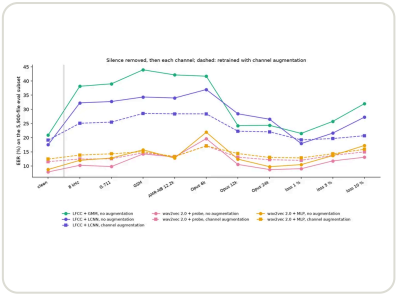
KEY RESULTS

- XGBoost catches 80 % [71, 88] of frauds at a 0.1 % false positive rate; PR-AUC 0.796.
- SMOTE and undersampling: no measurable gain, intervals centred on zero.
- A random split would have overstated PR-AUC by 0.07. All 6 pre-set success criteria met.

github.com/amine-belkacem/fraud-detection-alert-budget

● IN PROGRESS

Audio deepfake detection that generalises



PROBLEM	Do deepfake detectors still work on unseen generators, real fakes found online and phone lines?
CONSTRAINT	Model choices made on dev only, never on eval; the silence shortcut checked explicitly.
KEY RESULT	wav2vec 2.0 layer 3 + linear probe: 1.21 % EER, against 10.48 % for LFCC + GMM.
TOOLS	PyTorch · wav2vec 2.0 · scikit-learn · G.711 / Opus · Kaggle GPU
PERIOD	Oct 2026 → ongoing
STATE	In progress · Milestone 5 done

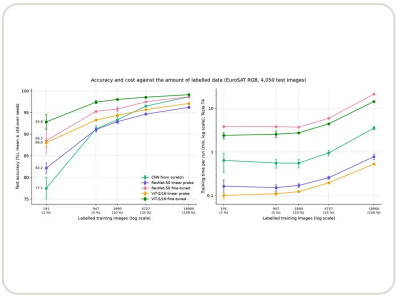
KEY RESULTS

- wav2vec 2.0 layer 3 + linear probe: 1.21 % EER, against 10.48 % for LFCC + GMM.
- Silence length alone gives 19.6 % EER; without it, a more honest figure is about 8 %.
- On real-world deepfakes (In-the-Wild) the best system rises to 22.4 %; through 10 simulated phone channels it averages 12 %.

github.com/amine-belkacem/audio-deepfake-generalization

● IN PROGRESS

EuroSAT: transfer learning for onboard triage



PROBLEM	Which transfer strategy classifies satellite tiles best with few labels and little compute on board?
CONSTRAINT	Same training images for every configuration at each data fraction; 3 seeds; Kaggle T4 for training.
KEY RESULT	75 training runs (5 configs × 5 data fractions × 3 seeds).
TOOLS	PyTorch · timm · Kaggle T4 · pytest
PERIOD	Oct 2026 → ongoing
STATE	In progress · Milestone 2 of 6

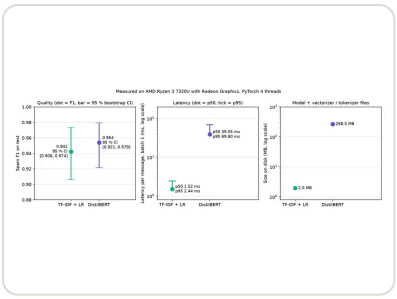
KEY RESULTS

- 75 training runs (5 configs × 5 data fractions × 3 seeds).
- With 191 labels: fine-tuned ViT 92.8 % vs scratch CNN 77.5 %. With all 18,900 the CNN catches up (98.63 %).
- A ViT linear probe trained in 6 s matches a fine-tuned ResNet-50 at 1 %.

github.com/amine-belkacem/eurosat-onboard-transfer

● IN PROGRESS

Spam filtering: classic model vs Transformer



PROBLEM	When is TF-IDF + logistic regression enough, and when is DistilBERT worth its cost?
CONSTRAINT	Duplicates removed before splitting; latency measured one message at a time on CPU.
KEY RESULT	Statistically tied on F1 (0.942 vs 0.954, difference interval contains 0).
TOOLS	scikit-learn · Hugging Face · DistilBERT · PyTorch
PERIOD	Oct 2026 → ongoing
STATE	In progress · Milestone 1 of 6 (PC)

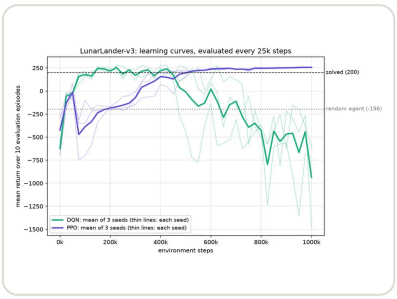
KEY RESULTS

- Statistically tied on F1 (0.942 vs 0.954, difference interval contains 0).
- DistilBERT is about 26× slower per message and 137× larger on disk.
- Duplicates removed before splitting (5,574 lines → 5,159 messages) so the test set does not reward memory.

github.com/amine-belkacem/spam-classic-vs-transformer-edge

● IN PROGRESS

DQN vs PPO on LunarLander



PROBLEM	DQN or PPO on LunarLander, judged on several seeds rather than one lucky run?
CONSTRAINT	Same 1 M-step budget for both algorithms; held-out evaluation episodes.
KEY RESULT	DQN reaches the 200 threshold 3 to 4 times sooner, then all three runs collapse below a random agent.
TOOLS	Gymnasium · PyTorch · DQN · PPO
PERIOD	Oct 2026 → ongoing
STATE	In progress · v0.1, 3 seeds

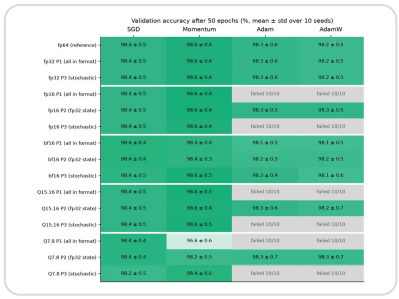
KEY RESULTS

- DQN reaches the 200 threshold 3 to 4 times sooner, then all three runs collapse below a random agent.
- PPO: 100 % success on 100 held-out episodes for every seed, 16 min vs 112 min of training.

github.com/amine-belkacem/lunarlander-dqn-vs-ppo

● IN PROGRESS

Optimizers in low precision



PROBLEM	What happens to SGD, Momentum, Adam and AdamW when the arithmetic drops to fp16, bf16 or fixed point?
CONSTRAINT	Optimizers written from scratch and matched against torch.optim before any experiment.
KEY RESULT	600 training runs, 6 number formats, 3 precision policies.
TOOLS	NumPy · PyTorch · pytest
PERIOD	Oct 2026 → ongoing
STATE	In progress · J4 of 5

KEY RESULTS

- 600 training runs, 6 number formats, 3 precision policies.
- Adam fails in every run with fp16 or fixed-point state: eps = 1e-8 rounds to 0. An fp32 state fixes it.
- AdamW's weight decay vanishes in every 16-bit format: 1 − lr·wd rounds to 1.

github.com/amine-belkacem/optimizers-low-precision

● JUST STARTED

Federated learning with differential privacy

PROBLEM	What does differential privacy cost each of 20 simulated hospitals, not just the average?
CONSTRAINT	Laptop CPU only (4 cores, 8 GB); images never leave their hospital.
KEY RESULT	Feasibility measured: DP-SGD epoch 1.344 s vs 0.560 s for plain SGD (2.4×), on a laptop CPU.
TOOLS	Flower · Opacus · PyTorch · MedMNIST
PERIOD	Oct 2026 → ongoing
STATE	Just started · J0: feasibility

KEY RESULTS

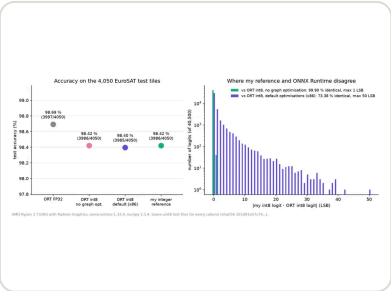
- Feasibility measured: DP-SGD epoch 1.344 s vs 0.560 s for plain SGD (2.4×), on a laptop CPU.
- No accuracy results yet.

github.com/amine-belkacem/dp-federated-hospitals

Embedded AI

● IN PROGRESS

Hand-written int8 inference engine (C / NEON)



PROBLEM	What really happens inside an int8 inference engine, down to the SIMD register?
CONSTRAINT	Bit-exact with an integer reference, layer by layer, before any speed claim; no malloc at inference.
KEY RESULT	Bit-exact: 89,100 / 89,100 op outputs identical to the reference over 4,050 test tiles (98.42 % accuracy).
TOOLS	C11 · ARM NEON · ONNX Runtime · CMake · QEMU · ASan / UBSan
PERIOD	Nov 2026 → ongoing
STATE	In progress · J3: Pi 5 benchmarks next

KEY RESULTS

- Bit-exact: 89,100 / 89,100 op outputs identical to the reference over 4,050 test tiles (98.42 % accuracy).
- Static 256 KB arena, no malloc at inference (checked by a test on the object files).
- Found why ONNX Runtime drifts up to 50 LSB on AVX2: int16 pair saturation in VPMADDUBSW.

github.com/amine-belkacem/int8-neon-inference-pi5

● JUST STARTED

What breaks first below 2 bits? LLMs on a Pi 5

PROBLEM	Which capability breaks first below 2 bits per weight, and is the speed gained worth it?
CONSTRAINT	Method and statistics frozen before the first run; measured on a Raspberry Pi 5.
KEY RESULT	Method frozen before the first run: paired McNemar + Holm, bootstrap, "broken" = retention upper bound below 90 %.
TOOLS	llama.cpp · GGUF · Python · Raspberry Pi 5
PERIOD	Nov 2026 → ongoing
STATE	Just started · J0: method and tooling

KEY RESULTS

- Method frozen before the first run: paired McNemar + Holm, bootstrap, "broken" = retention upper bound below 90 %.
- 84 tests; no model evaluated yet.

github.com/amine-belkacem/llm-below-2-bits-pi5

JUST STARTED

What quantization costs a small LLM on a Pi 5

PROBLEM	What does 2- to 8-bit quantization really cost a small LLM on a Raspberry Pi 5?
CONSTRAINT	Every number from a measurement on the board; llama.cpp version pinned.
KEY RESULT	Tooling and tests in place (Wilson intervals, KV-cache memory model, bandwidth ceiling); no results published yet.
TOOLS	llama.cpp · Python · shellcheck · Raspberry Pi 5
PERIOD	Nov 2026 → ongoing
STATE	Just started · J0: tooling

KEY RESULTS

- Tooling and tests in place (Wilson intervals, KV-cache memory model, bandwidth ceiling); no results published yet.

`github.com/amine-belkacem/pi5-llm-quant-map`

FPGA & digital design

IN PROGRESS

Molecular dynamics on FPGA (HW/SW co-design)

PROBLEM	Compute bonded forces and the reciprocal PME chain of molecular dynamics in fixed-point hardware.
CONSTRAINT	Portable VHDL (target board not fixed yet); results compared against an OpenMM reference.
KEY RESULT	Python reference model matches OpenMM to 1e-13, PME included.
TOOLS	VHDL · Python · OpenMM · Quartus · Questa
PERIOD	2026–2027 (M2) → ongoing
STATE	In progress · Step 2: fixed-point bonded forces

KEY RESULTS

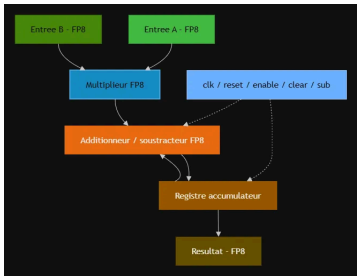
- Python reference model matches OpenMM to 1e-13, PME included.
- Bit-exact fixed-point model of the bond forces, 536 test vectors; formats chosen from measured error (Q8.24 positions, Q16.48 forces).
- Found issues in the spec: Q8.24 cannot hold k_b, Q16.48 saturates for Δr > 0.0708 nm.

Repository published once the course project is graded.

`Repository coming soon`

● COMPLETE

FP8 multiply-accumulate unit in VHDL



PROBLEM	Build the basic operator of a neuron, acc ± a × b, in the FP8 format used for AI inference.
CONSTRAINT	Small entry-level Cyclone IV FPGA; every block verified by a self-checking testbench.
KEY RESULT	Full MAC in 360 logic elements, about 6 % of an EP4CE6.
TOOLS	VHDL · Quartus Prime · Questa · GHDL
PERIOD	2025–2026 (M1)
STATE	Complete · M1 project

KEY RESULTS

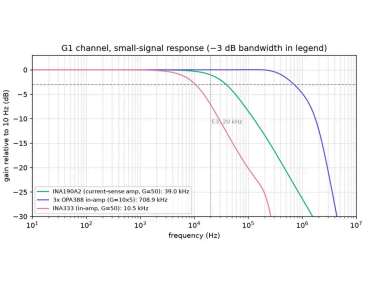
- Full MAC in 360 logic elements, about 6 % of an EP4CE6.
- Saturation to ±448, flush to zero; testbenches run with GHDL in CI on every push.

github.com/amine-belkacem/fp8-mac-vhdl

Analog & electronics

● IN PROGRESS

Autorange current profiler (Pi 5 HAT)



PROBLEM	Measure 1 µA to 500 mA without losing a sample when the range switches.
CONSTRAINT	€0 until the very end: every block validated in ngspice before building; burden under 100 mV.
KEY RESULT	Four shunt ranges, every range under 4 % of the error limit before calibration; burden 61.95 mV at 500 mA.
TOOLS	ngspice · Altium Designer · Python · C · RP2040 · Monte Carlo
PERIOD	Nov 2026 → ongoing
STATE	In progress · J7 of 10: Altium schematic

KEY RESULTS

- Four shunt ranges, every range under 4 % of the error limit before calibration; burden 61.95 mV at 500 mA.
- Range switching in 20 to 90 ns with comparators, latches and gate drivers (transient simulation).
- Monte Carlo on 200 instruments, 0–50 °C: 100 % yield on the error spec.
- Firmware core in C tested against the Python reference: 0 divergence on 5,660 samples. 105 tests in CI.

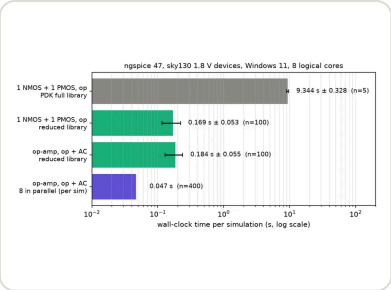
Simulation only so far; the real board comes at the end.

github.com/amine-belkacem/autorange-current-profiler

JUST STARTED

Can an optimiser size an op-amp better than me?

PROBLEM	Can an optimiser size an op-amp better than a hand design, and does its circuit survive process variation?
CONSTRAINT	Equal number of ngspice simulations for every method; hypotheses frozen before any optimisation.
KEY RESULT	Reduced PDK library identical to the full one to the 12th digit: 0.18 s per simulation instead of 9.3 s of loading.
TOOLS	ngspice · SKY130 PDK · Python · Monte Carlo
PERIOD	Nov 2026 → ongoing
STATE	Just started · J0 of 7



KEY RESULTS

- Reduced PDK library identical to the full one to the 12th digit: 0.18 s per simulation instead of 9.3 s of loading.
- Mismatch spread within 2.2 % of my prediction over 1,000 draws. Hypotheses frozen before any optimisation.

github.com/amine-belkacem/sky130-opamp-sizing-bench

JUST STARTED

Analog in-memory computing: an error budget

PROBLEM	How much hardware imperfection can a neural network take when its products run inside a resistive crossbar?
CONSTRAINT	PC only; every expected test value computed by hand first.
KEY RESULT	Differential weight mapping, DAC/ADC models and energy blocks; 45 hand-checked tests. No accuracy claim yet.
TOOLS	Python · NumPy · SciPy · ngspice
PERIOD	Nov 2026 → ongoing
STATE	Just started · J0: foundations

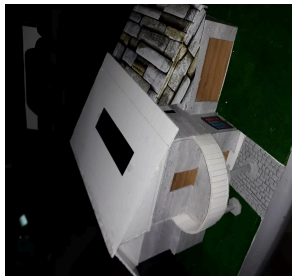
KEY RESULTS

- Differential weight mapping, DAC/ADC models and energy blocks; 45 hand-checked tests. No accuracy claim yet.

github.com/amine-belkacem/analog-imc-error-budget

● COMPLETE

Solar power supply for a smart-house model



PROBLEM	Power an Arduino smart-house model from the sun.
CONSTRAINT	Cheap off-the-shelf parts; the TP4056 charger accepts 8 V at most.
KEY RESULT	Measured 585 mA charge current at 5.72 V panel voltage; sizing for 9,281 Wh/day, driven by the fridge.
TOOLS	Arduino · TP4056 · XL6009 · Python · pytest
PERIOD	2022–2023 (Bachelor's thesis)
STATE	Complete · Bachelor's thesis, USTHB

KEY RESULTS

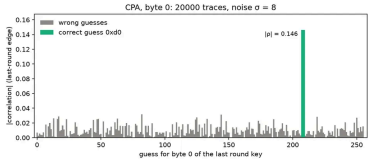
- Measured 585 mA charge current at 5.72 V panel voltage; sizing for 9,281 Wh/day, driven by the fridge.

`github.com/amine-belkacem/smart-house-solar-energy`

Security

● IN PROGRESS

AES-128 side-channel lab



PROBLEM	Does a countermeasure that looks secure at register level still hold once the glitches of the real netlist count?
CONSTRAINT	Simulation only, no board: traces from GHDL and Questa, area and Fmax from Quartus.
KEY RESULT	Unprotected core: 4,890 LE, 105.7 ± 2.7 MHz on Cyclone IV; CPA recovers the full 128-bit key, TVLA flags leakage from 43 traces.
TOOLS	VHDL · GHDL · Questa · Quartus · Python · NumPy · CPA / TVLA
PERIOD	Nov 2026 → ongoing
STATE	In progress · Milestone 6 of 9

KEY RESULTS

- Unprotected core: 4,890 LE, 105.7 ± 2.7 MHz on Cyclone IV; CPA recovers the full 128-bit key, TVLA flags leakage from 43 traces.
- Same tooling breaks a real FPGA from public oscilloscope traces: full key in 373 traces.
- Random dummy rounds: 28× more traces for an unprofiled attacker, only 3.7× for a profiled one.
- Gate-level timing simulation of the netlist: 73 % of LUT toggles are glitches.

`github.com/amine-belkacem/aes128-sidechannel-lab`

Industrial automation

JUST STARTED

Is a soft-PLC on a Pi 5 deterministic enough?

PROBLEM	Is a soft-PLC on a stock Linux ARM board deterministic enough to replace a classic PLC?
CONSTRAINT	Success criterion fixed in advance: max jitter ≤ 10 % of a 10 ms cycle over 1 h under load.
KEY RESULT	Plant physics, buffer tank with a PI level loop, IEC 61131-3 ST programs with unit tests; 68 tests in CI.
TOOLS	OpenPLC · Structured Text · Modbus TCP · OPC UA · PREEMPT_RT · Python
PERIOD	Nov 2026 → ongoing
STATE	Just started · J0: bench bring-up

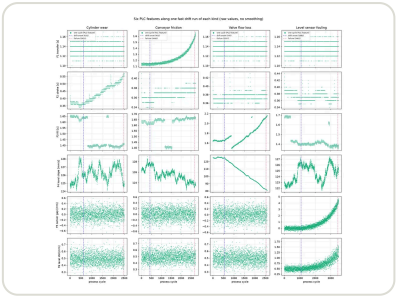
KEY RESULTS

- Plant physics, buffer tank with a PI level loop, IEC 61131-3 ST programs with unit tests; 68 tests in CI.
- No jitter results yet; success criterion fixed in advance (max jitter ≤ 10 % of the 10 ms cycle over 1 h under load).

`github.com/amine-belkacem/softplc-pi5-determinism`

JUST STARTED

Digital twin and drift detection next to a soft-PLC



PROBLEM	Catch a drifting part before it fails, on the same Pi as the PLC, without disturbing its cycle.
CONSTRAINT	Detectors compared at the same false-alarm rate; the twin is never copied from the simulator.
KEY RESULT	Seeded drift datasets and evaluation metrics (delay, Garwood and Clopper-Pearson intervals); 42 tests. Success criteria written before any measurement.
TOOLS	Python · NumPy · scikit-learn · OpenPLC
PERIOD	Nov 2026 → ongoing
STATE	Just started · J1: data generator

KEY RESULTS

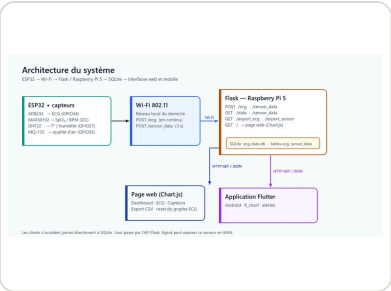
- Seeded drift datasets and evaluation metrics (delay, Garwood and Clopper-Pearson intervals); 42 tests. Success criteria written before any measurement.

`github.com/amine-belkacem/plc-drift-twin`

Embedded & IoT

● COMPLETE

Connected health monitor (ESP32 + Pi 5)



PROBLEM	Follow a patient's vital signs at home without sending medical data to a third-party cloud.
CONSTRAINT	One ESP32 and one Raspberry Pi 5 on the home network.
KEY RESULT	Whole chain working on the real setup; readings every 3.07 s (median) against a 3 s timer.
TOOLS	ESP32 · Arduino C++ · Flask · SQLite · Flutter · Chart.js
PERIOD	2024–2025 (Master's thesis)
STATE	Complete · Master's thesis, USTHB

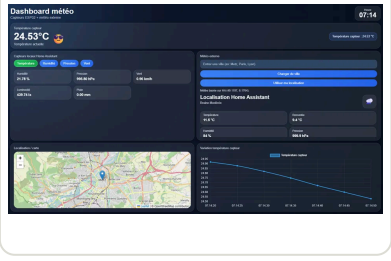
KEY RESULTS

- Whole chain working on the real setup; readings every 3.07 s (median) against a 3 s timer.
- Going back through the code, I found and fixed a bug that truncated the ECG to whole volts.

`github.com/amine-belkacem/health-monitor-iot`

● COMPLETE

Connected weather station dashboard



PROBLEM	Show live ESP32 sensor values next to official weather data, reliably.
CONSTRAINT	No backend of my own; the access token never written in the code.
KEY RESULT	13 tests on the unchanged code run in CI and caught one real bug.
TOOLS	JavaScript · Home Assistant · REST API · ESP32 · Chart.js
PERIOD	2025–2026 (M1 lab)
STATE	Complete · M1 lab, 10 students

KEY RESULTS

- 13 tests on the unchanged code run in CI and caught one real bug.

`github.com/amine-belkacem/home-assistant-weather-station`

Software & tools

● COMPLETE

Qt6 UART API, PC ↔ FPGA



PROBLEM	Let a PC application exchange data with an FPGA board without handling serial-port details.
CONSTRAINT	Project under NDA: only my own part, the PC-side API, is published.
KEY RESULT	QTest tests run in CI on virtual serial ports.
TOOLS	C++17 · Qt 6 · QSerialPort · CMake · Basys 3
PERIOD	2025–2026 (M1)
STATE	Complete · M1 project

KEY RESULTS

- QTest tests run in CI on virtual serial ports.
- My part of a project under NDA, published with my supervisor's agreement.

`github.com/amine-belkacem/qt6-uart-fpga-api`

● COMPLETE

Multithreaded railway traffic simulator

PROBLEM	Let seven trains share single tracks without ever being on the same section.
CONSTRAINT	POSIX threads, mutexes, semaphores, pipes and fork only.
KEY RESULT	Two strategies: deadlock detection (Holt graph reduction, victim backs off) and prevention (all-or-nothing route reservation).
TOOLS	C · POSIX threads · semaphores · pipes · fork · TSan
PERIOD	Course project
STATE	Complete · Course project

KEY RESULTS

- Two strategies: deadlock detection (Holt graph reduction, victim backs off) and prevention (all-or-nothing route reservation).
- Safety invariant checked independently by the control process; ThreadSanitizer build.

`github.com/amine-belkacem/simu_ferroviaire`
