电子科技大学
格拉斯哥海南学院
Glasgow College Hainan,UESTC

Logbook

From: 03/10/2026 To: 27/10/2026

Month	**List the main activities** (only few words per activity)	**Interaction with the supervisor**			**Any other form of supervisory interaction** (second supervisor, industry, fellows etc.)
		Number of meetings	Mode of meeting (face- to-face, online e.g., Skype, WeChat etc.)	Number of emails exchanged	
October 2026	1. Supervisor meetings 2. Literature review 3. LP / TOU baselines 4. Specification report draft	2	Online (WeChat)	Not recorded	Not recorded

UESTCHN4008P · FINAL YEAR INDIVIDUAL PROJECT

Logbook

Daily records · 3 October – 27 October 2026

Student	Letian Xie
Student IDs	2959471 / 2023300901028
Project	SP2 Smart Charging Scheduling for Cost and Peak Reduction
1st Supervisor	Dr Amer Ghias
Compiled	27 October 2026

Contents

01 Monthly summary 3

02 Record basis and responsibility 4

03 3 October (Saturday): first online meeting 5

04 17 October (Saturday): second online meeting 6

05 20 October (Tuesday): first scheduler and tariff checks 7

06 21 October (Wednesday): working through the templates 9

07 22 October (Thursday): early energy requirement 11

08 27 October (Tuesday): review, form status and what comes next 13

Only six days between 3 and 27 October have records: 3, 17, 20, 21, 22 and 27 October. The days in between were not logged, and nothing has been written in to fill them.

Platforms, durations, timezones and email counts that were never recorded are written as “not recorded”, never as zero. Results from synthetic scenarios describe those scenarios only; they are not savings or annual returns for a real station.

01

Monthly summary

Page 1 of Log Book 1. The university header and the title come first, then my name, student number and the period covered, and below that the monthly table.

Month	List the main activities	Number of meetings	Mode of meeting	Number of emails exchanged	Any other form of supervisory interaction
October 2026	Online supervision; AI-assisted scheduling, tariff checks and report preparation.	2	Online	Not recorded	Not recorded
Total	Two online meetings: 3 and 17 October.	2	Online	Not recorded	Not recorded

A note on the entries

Only the two meetings I could confirm are counted: 3 October at 13:00 and 17 October at 14:00, both online. That gives two meetings and “Online” for the mode.

The email count is “Not recorded”, not zero. I never counted them, and the July application correspondence does not belong to this month. A zero would say there were none, which is not true. The last column is in the same position.

For activities I kept to four items: online supervision, AI-assisted scheduling work, tariff checks and report preparation. Anything longer would start describing things that did not happen.

The monthly progress form and the task book are separate items in the university system. This logbook does not stand in for either of them.

02

Record basis and responsibility

This is the opening of page 2. It sets out where the records come from and which parts I did myself.

Where these records come from

I compiled them on 27 October from two things: the two meetings I confirmed, and dated files on my computer.

The technical entries describe saved AI-assisted outputs. They are not a record of my reading time or working hours. I did not keep one, and I have not estimated numbers to fill the gap. The platform, length and decisions of the two meetings were never written down either. Email counts for the project were not recorded, and the July application correspondence is excluded.

Page index

Page 1 is the cover and monthly summary, page 2 the meetings, page 3 the initial software and tariff checks, page 4 the early energy example, and page 5 the review and next steps.

What I did and what the AI did

The mathematical reasoning, the code, the checks on the runs and the revision of each section are mine. OpenAI Codex helped with English expression, drafts of text and code, code organisation, and interpretation of results.

One thing has to be said plainly. This division does not mean that the manual checks I did not record are already done. I still have to check the work myself, and when I do, it gets recorded under its actual date.

03

3 October (Saturday): first online meeting

Amer and I had our first online meeting at one in the afternoon. We talked about how to divide the work, how the project should be designed, and how the timetable would run.

I did not note which platform we used, how long it lasted, what was decided, or what I was asked to do next. This entry was written on 27 October while putting the logbook together. The time is the 13:00 I was given; the timezone was never specified, so I have not added one.

It was the first of two meetings in October. The second was on 17 October.

What was not filled in later

If I had written two lines on the day, this entry would be easier. But going back over a meeting from memory is a good way to mistake what I know now for what was decided then, so the entry stops at “we discussed the division of work, design and timetable”.

What my supervisor suggested, what reading he asked for, when we would meet next — none of that is here, because I do not remember it.

What existed at the time

The project had only just started. The topic is SP2, cost and peak reduction for charging scheduling. The scheduler of 20 October and the early energy extension of 22 October both came after this meeting and cannot be counted into 3 October.

04

17 October (Saturday): second online meeting

Two in the afternoon, online again, and much the same ground: division of work, design, schedule.

What was decided, and what was assigned to me, was again not recorded separately. Platform, length and timezone are all missing. Like the 3 October entry, this one was reconstructed on 27 October.

How it differs from the first

Honestly, both meetings left the same amount of information: a direction, no conclusions. The only thing the second one adds is the time, 14:00 rather than 13:00. It would be easy to give the later meeting details the first one does not have, such as “we settled on the LP” or “we agreed to submit at the end of October”, but nothing supports that.

What followed

The first working scheduler came out on 20 October, with three strategies compared and 23 tests passing. On 22 October the ready_by constraint was added. Those are real artefacts, but they were not decided at this meeting and should not be mixed into it.

05

20 October (Tuesday): first scheduler and tariff checks

The first working Python module came out of this day. It compares three strategies: capacity-limited immediate charging, a cheapest-slot TOU heuristic, and a linear programme.

Each session is described by four quantities: when it arrives, when it departs, how much grid-side energy it requests, and its power limit. The LP minimises energy cost plus a weighted hourly peak. Every feasible schedule has to deliver the full request inside the connection window. If it cannot, it is reported as infeasible. Partial delivery is not treated as a valid answer.

SP1 supplies aggregate hourly demand, which shows neither when individual vehicles leave nor how much energy is flexible. The demonstration therefore has to invent its own sessions: 50 of them, explicitly synthetic, with a 100 kW site limit and 1007 kWh of total demand. One available connection is assumed per session. Charger assignment, PV, background load, battery losses and V2G are not modelled.

Results

Strategy	Energy kWh	Cost USD	Peak kW
Immediate	1007	123.66937	100.00000
TOU heuristic	1007	107.65837	100.00000
LP	1007	107.65837	77.46154

All three deliver 1007 kWh, so the costs are directly comparable. TOU and the LP cost the same, 107.65837. The difference is in the peak: the LP brings 100 kW down to 77.46 kW. That comes from the peak term, not from cheaper electricity. It should not be read as the LP buying at a better tariff.

What the result does and does not show

The saved verification record reports 23 passing tests, covering tariffs, feasibility, conservation and exported costs. The peak weight of 0.5 USD/kW in the LP is an engineering preference I set, not a utility demand charge. The two should not be conflated, and the lower peak holds only for this fixture.

The cost comparison holds only at equal energy. It says nothing about what a real charging station would save, and it cannot be turned into an annual return.

Tariff and interface problems

The benchmark uses a fixed Xcel summer 2023 residential-rate snapshot. It is a benchmark, not the contract the station actually signed. Earlier years are incompletely covered, and some are missing entirely.

The Palo Alto tiered example was calculated separately, and the result is slightly counter-intuitive. Within one billing period, equal total energy gives equal energy charges regardless of when the charging happens.

Exported CSV schedules use whole-hour UTC timestamps. The proposed USD fields are still a proposal; the team has not agreed them yet.

What was kept

`scheduler.py` and `sp2-kpi-v2.json` are in the SP5 handover folder. The test record is `tests_2026-10-20.txt`.

One point needs stating. These are AI-assisted outputs with a saved record. The date I ran the code myself is not recorded. When I do run it, it gets its own entry under that date, and should not be folded into 20 October.

06

21 October (Wednesday): working through the templates

Most of this day went on documents. I went through the report template and the form templates issued by Glasgow College Hainan one by one, and revised the English wording along the way.

Two things were held to. The report still describes software only. It does not extend to physical charge control, and there is no laboratory testing or charger hardware in it. And every tariff figure carries its year and source rather than being written as a generic number.

Which sentences changed, and which revision they reached, was not recorded. I cannot reconstruct it now, so I have not tried.

Files in hand

File	Purpose
`2959471X_Xie_UESTCHN4008P_Specification_Report_26-27.docx`	Report body, 8 pages, 15 references, 4 SMART objectives, block diagram and Gantt
`2959471X_Xie_Risk_Assessment_to-be-signed.docx`	Risk assessment form
`2959471X_Xie_ECD_to-be-signed.docx`	Ethical consideration declaration
`2959471X_Xie_Resources_Request_pending-approval.docx`	Resource request, 0 RMB additional budget
`2959471X_Xie_Logbook1_Cover_notes-pending.docx`	Logbook cover, for editing only, not submitted on its own

These were finalised on 27 October. The 21st was the template check; the two dates should not be merged.

What is not written here

No meeting took place that day and nothing went back and forth with the supervisor. So there is no “the supervisor has reviewed it” or “all comments are addressed” — the supervisor signature fields were still blank on the 27th.

What each form said before and after was not recorded separately, so it is not listed line by line. Setting it out in detail now would mean inventing it.

What came next

The review on 27 October built on this. The two meetings and the early energy extension went into the logbook and the report. After that: 28 October for checking the report section by section and running the code, and 29 October for submission.

07

22 October (Thursday): early energy requirement

This day added two fields, `ready_by` and `minimum_energy_kwh`.

The reason is practical. A vehicle may nominally stay until 22:00, but the driver may need it at 18:00. If the scheduler puts all the charging after 18:00, the vehicle leaves short. The earlier model had no such constraint, so its “optimal” schedule was of little use in that case.

How the constraint works

Before the declared readiness time, cumulative grid-side energy must not fall below `minimum_energy_kwh`. At departure the original full request is still due. The minimum is conditional: it assumes the schedule is feasible and that the vehicle stays connected until `ready_by`.

It constrains grid-side delivered energy, not battery state of charge. It should not be written as a guarantee of range or of a particular charge level. It guarantees neither.

An eight-vehicle example

Item	Value
Sessions	8 (synthetic)
Connection window	14:00 to 22:00
Request per vehicle	14 kWh, 7 kW limit
Site limit	28 kW
Tariff	Historical Xcel snapshot
Peak penalty	0, switched off deliberately to isolate energy cost

Both LPs use the same secondary rule: among cost-optimal schedules, maximise the smallest fraction of energy delivered by readiness.

Measure	Original LP	Early energy LP
Minimum per vehicle by 18:00	3.5 kWh	7.0 kWh
Full-plan energy	112 kWh	112 kWh
Full-plan cost	USD 14.21532	USD 18.69840
Maximum hourly site power	28 kW	28 kW

The minimum delivered by 18:00 rises from 3.5 to 7.0 kWh, exactly half the request. Total energy is unchanged at 112 kWh, but the cost goes from 14.21532 to 18.69840, or 4.48308 USD more. The peak does not move; both cases sit at 28 kW. The money bought earlier availability and nothing else.

Fifteen requests

Fifteen constructed requests were tried: 13 feasible, 2 not. The two failures happen because the site cannot supply enough energy before `ready_by`, not because the algorithm failed.

The prototype passes five tests.

One more thing. If a vehicle does leave early, the charging scheduled after that point has to be removed and the bill recalculated from what was actually delivered. Less energy delivered is not an efficiency gain; it is a shortfall.

What was kept

`readiness_demo.py`, `readiness_results.json`, `verification_2026-10-22.json` and `test_readiness.py` are in the early_energy folder of the SP5 handover.

The prototype and the analysis were produced with AI assistance. There are no real drivers, no questionnaires and no physical charging trials behind them.

08

27 October (Tuesday): review, form status and what comes next

This was the wrap-up day. The two October meetings went into the logbook, email counts were set to “not recorded”, and the early energy extension was folded into the report — its question, objectives and evaluation plan. The report still describes software only, and the tariff assumptions still carry their dates.

Where the forms stand

Item	On my side	Still missing
Risk assessment	Signed by hand on 27 October; training, COSHH and supervision boxes completed	Supervisor’s judgement items and signature
ECD	Signed by hand	Supervisor signature and date
Resources request	Signed by hand	Supervisor approval
Specification report	Body complete, 8 pages	Originality declaration, which I have to read and sign myself
Log Book 1	5 pages complete	No signature required

The two items in the risk form — whether training is adequate and whether risks are adequately controlled — are my supervisor’s judgement. I have not filled in a conclusion on his behalf.

The handwritten annotations already in the PDFs have to be kept. Re-exporting a Word file and converting it again would wipe the signatures.

What I worked out this week

The first priority is to keep service requirements explicit. A schedule can look excellent on cost and still serve a driver badly when that driver leaves early. The early energy constraint makes the requirement testable, at the price of a smaller feasible set: two of fifteen requests became impossible. Evaluation therefore has to report per-vehicle service and the number of infeasible requests, not just the cost and peak reduction.

A few concepts also need keeping apart. An aggregate forecast from SP1 is not a set of vehicle requests; it has totals but no individual behaviour. Grid-side energy is not battery state of charge, since the meter sits outside the battery and losses are in between. A lower electricity

price is not lower emissions; the two are not linked. Later comparisons have to give every method the same information, or the comparison means nothing.

Plan

Date	Action	Completion evidence
28 Oct	Check each report section, run the code myself, record the real results and the corrections; follow up on the supervisor's comments, the risk judgement items and the signatures	Report signed by me, supervisor items complete, review record traceable in the log
29 Oct	Open every signed file page by page and submit the four PDFs on both platforms, keeping the receipts	Submission confirmed on both, file names, times and receipts stored
13 Nov	Agree the session readiness inputs and the SP4/SP5 service summary fields with the team; audit data and tariff coverage	Fields and audit record settled
11 Dec	Complete the four-method comparison and the constraint checks	Comparison results and check record

The four PDFs are the Specification Report, the Risk Assessment, the ECD and the complete Log Book 1, submitted once on each platform. The resource request goes separately by email to the supervisor. It is not a fifth platform item; the additional budget is 0 RMB, but it still needs approval.

What is not written here

Reading dates, development hours, meeting decisions and email totals were not recorded, and none of them has been estimated to fill the space. If I do run the code on 28 October, that gets its own entry under the 28th rather than being folded back into this one.