



UNIVERSITY OF IOANNINA
SCHOOL OF INFORMATICS AND TELECOMMUNICATIONS
DEPARTMENT OF INFORMATICS AND TELECOMMUNICATIONS

PhD Thesis

**Diffuse energy utilization and power management for autonomous
cyber-physical microsystems**

Ioannis Masklavanos

Computer and Network Engineer

Arta, October 2025



Πανεπιστήμιο Ιωαννίνων

Σχολή Πληροφορικής και Τηλεπικοινωνιών

Τμήμα Πληροφορικής και Τηλεπικοινωνιών

Διδακτορική διατριβή

**Αποδοτική διαχείριση και αξιοποίηση διάχυτης μικροενέργειας για
την τροφοδοσία αυτόνομων κυβερνοφυσικών μικροσυστημάτων**

Ιωάννης Μασκλαβάνος

Μηχανικός Ηλεκτρονικών Υπολογιστών και Δικτύων

Άρτα, Οκτώβριος 2025

**DIFFUSE ENERGY UTILIZATION AND POWER MANAGEMENT
FOR AUTONOMOUS CYBER-PHYSICAL MICROSYSTEMS**

Ημερομηνία αίτησης του κου Μασκλαβάνου Ιωάννη: 25-09-2019

Ημερομηνία ορισμού Τριμελούς Συμβουλευτικής Επιτροπής: 23^η/23-10-2019

Μέλη Τριμελούς Συμβουλευτικής Επιτροπής:

Επιβλέπων:

Δουμένης Γρηγόριος, Επίκουρος Καθηγητής του Τμήματος Πληροφορικής και Τηλεπικοινωνιών του Πανεπιστημίου Ιωαννίνων

Μέλη:

Βαρτζιώτης Φώτιος, Επίκουρος Καθηγητής του Τμήματος Πληροφορικής και Τηλεπικοινωνιών του Πανεπιστημίου Ιωαννίνων

Τσάμης Χρήστος, Διευθυντής Ερευνών στο Εθνικό Κέντρο Ερευνών «ΔΗΜΟΚΡΙΤΟΣ», Ινστιτούτο Νανοεπιστήμης και Νανοτεχνολογίας, Εργαστήριο Συγκομιδής Ενέργειας & Αυτόνομων Αισθητήρων

Ημερομηνίας ορισμού θέματος διδακτορικής διατριβής: 23^η/23-10-2019

«Αποδοτική Διαχείριση και Αξιοποίηση Διάχυτης Μικροενέργειας για την Τροφοδοσία Αυτόνομων Κυβερνοφυσικών Μικροσυστημάτων»

ΟΡΙΣΜΟΣ ΕΠΤΑΜΕΛΟΥΣ ΕΞΕΤΑΣΤΙΚΗΣ ΕΠΙΣΤΡΟΦΗΣ: 97/09-10-2025

Δουμένης Γρηγόριος	Επίκουρος Καθηγητής του Τμήματος Πληροφορικής και Τηλεπικοινωνιών της Σχολής Πληροφορικής και Τηλεπικοινωνιών, με γνωστικό αντικείμενο «Ψηφιακή Επεξεργασία Σήματος – Επαναπροσδιοριζόμενα Ραδιοσυστήματα»
Τσάμης Χρήστος	Διευθυντής Ερευνών στο Εθνικό Κέντρο Ερευνών «ΔΗΜΟΚΡΙΤΟΣ», Ινστιτούτο Νανοεπιστήμης & Νανοτεχνολογίας, Εργαστήριο Συγκομιδής Ενέργειας & Αυτόνομων Αισθητήρων
Βαρτζιώτης Φώτιος	Επίκουρος Καθηγητής του Τμήματος Πληροφορικής και Τηλεπικοινωνιών της Σχολής Πληροφορικής και Τηλεπικοινωνιών, με γνωστικό αντικείμενο «Σχεδίαση και Βελτιστοποίηση Ψηφιακών Συστημάτων VLSI με έμφαση στον Έλεγχο Ορθής Λειτουργίας»
Αγγέλης Κωνσταντίνος	Καθηγητής του Τμήματος Πληροφορικής και Τηλεπικοινωνιών της Σχολής Πληροφορικής και Τηλεπικοινωνιών, με γνωστικό αντικείμενο «Μικροηλεκτρονικές Διατάξεις – Τηλεπικοινωνιακά Συστήματα»
Γιαννακέας Νικόλαος	Αναπληρωτής Καθηγητής του Τμήματος Πληροφορικής και Τηλεπικοινωνιών της Σχολής Πληροφορικής και Τηλεπικοινωνιών, με γνωστικό αντικείμενο «Σήματα και Συστήματα»

Ρεΐσης Διονύσης

Καθηγητής του Τμήματος Φυσικής του Εθνικού και Καποδιστριακού Πανεπιστημίου, με γνωστικό αντικείμενο «Προσβασιμότητα Συστημάτων Πληροφορικής και Τηλεπικοινωνιών – Τεχνολογίες Σήματος Φωνής»

Κιζήρογλου Μιχάλης

Καθηγητή του Τμήματος Μηχανικών Παραγωγής και Διοίκησης της Πολυτεχνικής Σχολής του Διεθνούς Πανεπιστημίου Ελλάδος, με γνωστικό αντικείμενο «Τεχνολογία Μικροηλεκτρονικής και Μικροσυστημάτων»

Έγκριση Διδακτορικής Διατριβής με βαθμό «ΑΡΙΣΤΑ» στις 15-10-2025

Αρτα 16-10-2025

**Ο ΠΡΟΕΔΡΟΣ ΤΟΥ ΤΜΗΜΑΤΟΣ
ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ
ΤΗΛΕΠΙΚΟΙΝΩΝΙΩΝ
Αλέξανδρος Τζάλλας
Καθηγητής**



Η Γραμματέας του Τμήματος

Ευαγγελία Χρήστου

Acknowledgements

This thesis would not have been possible without the support, guidance, and encouragement of many people and institutions, to whom I am deeply grateful.

First and foremost, I would like to express my sincere gratitude to my supervisor Dr Doumenis for his continuous guidance, insightful feedback, and unwavering support throughout the course of this work. His scientific rigor, patience, and enthusiasm for research have shaped both this thesis and my own way of thinking as an engineer and researcher. I am equally grateful to my co-supervisors and the members of my advisory and examination committees for their constructive comments, challenging questions, and valuable suggestions, which significantly improved the quality and clarity of this dissertation.

I would also like to thank my colleagues and collaborators at the Autonomous Systems Laboratory (ASL) for the many discussions, suggestions, and moments of support that contributed to this work.

This work has been supported by research projects T1EΔK-00360 PERPS, T2EΔK-EFOS and TAEΔK-06182 AMSA, which provided the resources, infrastructure, and experimental platforms needed to carry out the studies reported in this thesis. Their support is gratefully acknowledged.

On a more personal note, I would like to thank my friends for their companionship, understanding, and encouragement during the more demanding periods of this journey. Their presence and humor made the difficult moments more bearable and the good moments more enjoyable.

Finally, and most importantly, I owe my deepest gratitude to my family. Their unconditional love, patience, and trust have been a constant source of strength. Their support, in ways both visible and invisible, made it possible for me to pursue and complete this work. This thesis is dedicated to them.

Abstract

Autonomous cyber-physical microsystems (μ CPS) promise long-horizon sensing, inference, and actuation with minimal maintenance, yet their feasibility is governed less by average harvested power than by the dynamics and losses of the power path and by bursty, deadline-driven workloads observed through sparse telemetry. This thesis develops a principled framework that converts diffuse environmental energy into predictable service. First, it introduces macroscopic, loss-aware models that map incident source flux to usable electrical power through realistic conversion, storage, and protection, explicitly accounting for cold-start energy, quiescent drains, leakage, equivalent series resistance, and undervoltage hysteresis. These models define mission-level autonomy metrics—energy neutrality over a horizon and survivability within protection bounds—that guide both design and evaluation. Second, it presents DREAM (Dynamic Rest & Energy-Aware Management), a runtime that uses low-cost observables (primarily storage-voltage trajectories and discrete protection events) to orchestrate rest, energy aggregation, and execution of high-crest-factor tasks (sensing, computation, communication) so that deadlines are met without inducing brownouts. Third, it proposes a trace-driven sizing methodology that selects harvester technology/area, storage type/capacity, and converter modes to meet mission-level energy and service constraints over stated horizons with quantifiable confidence under realistic temporal variability in sources and networks. Validation across representative case studies demonstrates sustained service with materially fewer brownouts and improved predictability versus battery-only baselines and controllers that optimize average power without explicit deliverability or protection modeling. The results provide designers with a credible path from environmental energy fluxes to system-level performance guarantees for maintainable, sustainable, and trustworthy μ CPS.

Keywords: energy harvesting; energy-neutral operation; runtime power management; trace-driven harvester sizing; cyber-physical microsystems

Περίληψη

Η διατριβή πραγματεύεται τον σχεδιασμό και τη λειτουργία αυτόνομων κυβερνοφυσικών μικροσυστημάτων που τροφοδοτούνται από συγκομιδή ενέργειας, επιδιώκοντας να μετατρέψει στοχαστικές και χαμηλής στάθμης ροές ισχύος σε προβλέψιμη υπολογιστική και επικοινωνιακή υπηρεσία με τεκμηριωμένες εγγυήσεις ουδετερότητας ενέργειας και βιωσιμότητας. Η πρακτική πρόκληση δεν συνίσταται στον μέσο ρυθμό παροχής ισχύος, αλλά στη χρονική δομή της ενεργειακής αλυσίδας, στις απώλειες μετατροπής και διανομής, και στους περιορισμούς παραδοτέας ισχύος κατά τη στιγμή εκτέλεσης κορυφωμένων φορτίων. Ενεργοβόρες μεταβάσεις —όπως ενεργοποίηση υποσυστημάτων ραδιοεπικοινωνίας, προθέρμανση αισθητήρων και εγγραφές σε μη πτητική μνήμη— επιβάλλουν αυστηρές στιγμιαίες απαιτήσεις που δεν περιγράφονται επαρκώς από αναλύσεις μέσων τιμών. Ως εκ τούτου, υιοθετείται μια κόμβο-κεντρική και απωλειο-ευαίσθητη θεώρηση, όπου η ενέργεια αντιμετωπίζεται ως πρωτογενής πόρος υπό περιορισμούς απόδοσης, προστασίας και αποθήκευσης.

Η θεωρητική θεμελίωση βασίζεται σε μακροσκοπικά μοντέλα ροής ισχύος από την πηγή έως το φορτίο μέσω ρεαλιστικής μετατροπής, αποθήκευσης και προστασίας. Η ισχύς που διατίθεται στην είσοδο της αλυσίδας μετά την αντιστοίχιση και τη συλλογή παριστάνεται ως $P_H(t)$, ενώ το συνολικό ενεργειακό κόστος των λειτουργιών —λήψη μετρήσεων από τους αισθητήρες, τοπική επεξεργασία, επικοινωνία και ψευδοσταθερές καταναλώσεις— ως $P_L(t)$. Η αποθηκευμένη ενέργεια $E(t)$ παρακολουθείται μέσω της τάσης αποθήκευσης $V_{\text{store}}(t)$, ενώ οι αποδόσεις εισόδου και εξόδου, μαζί με όρους διαρροών και quiescent drains, ενσωματώνονται ρητά ώστε να αποδοθεί η πραγματική συμπεριφορά της αλυσίδας. Τα κυκλώματα προστασίας ορίζουν επιτρεπτή ζώνη λειτουργίας· η υπέρβαση των ορίων οδηγεί σε διακοπές (UVLO) και σε ανεπιθύμητες επανεκκινήσεις. Σε αυτό το πλαίσιο, ορίζονται δύο μετρικές αποστολής: η ουδετερότητα ενέργειας σε συγκεκριμένο ορίζοντα, όπου η συγκομιδή και η αποθήκευση ισοσκελίζουν το κόστος μετά τις απώλειες, και η βιωσιμότητα, δηλαδή η διατήρηση της τροχιάς $E(t)$ εντός των ορίων προστασίας με διασφαλισμένη δυνατότητα εκτέλεσης κορυφωμένων εργασιών χωρίς κατάρρευση τάσης. Η εστίαση μετατοπίζεται έτσι από την αθροιστική ενέργεια στην πραγματική δυνατότητα παράδοσης ισχύος τη δεδομένη στιγμή, η οποία αποδεικνύεται καθοριστική για την αξιοπιστία της υπηρεσίας.

Οι πηγές διάχυτης ενέργειας —φωτοβολταϊκά σε εσωτερικούς χώρους, θερμικές κλίσεις χαμηλού ΔT , μηχανικοί κραδασμοί και περιβαλλοντικό RF— χαρακτηρίζονται από υψηλή χρονική μεταβλητότητα με διαστήματα πλατό και υφέσεων. Η εργασία προσδιορίζει προσεκτικά τη χαρτογράφηση από τη φυσική ροή $\Phi(t)$ στη χρήσιμη ισχύ $P_H(t)$, λαμβάνοντας υπόψη ψυχρή εκκίνηση, περιορισμούς αντιστοίχισης/MPPT, δυσανεξία σε πολύ μικρές τάσεις ανοικτού κυκλώματος (ιδίως σε TEG) και βραχύβια παράθυρα ενέργειας σε κραδασμούς ή RF. Αναδεικνύεται ότι η πολυτροπική σύζευξη πηγών αυξάνει τη συχνότητα ευκαιριών συγκομιδής, υπό την προϋπόθεση ότι ελέγχεται η προσθήκη ψευδοσταθερών καταναλώσεων και οι αλληλεπιδράσεις ελέγχου μέσω προσεκτικού OR-ing, κοινής μετατροπής όπου είναι εφικτό και σωστής ρύθμισης υστερήσεων και κατωφλίων.

Στο επίπεδο αρχιτεκτονικής, προτείνεται ένας συμπαγής αλλά εκφραστικός αγωγός ισχύος: αντιστοίχιση/ανορθωτές/MPPT, μετατροπείς DC-DC, υποσύστημα αποθήκευσης (υπερπυκνωτές, μικρο-μπαταρίες ή υβριδικές λύσεις) και διανομή με προστασίες. Ο διαχωρισμός μονοπατιού ψυχρής εκκίνησης από το μονοπάτι κανονικής λειτουργίας επιτρέπει αφύπνιση υπό δυσμενείς συνθήκες και, στη συνέχεια, λειτουργία υψηλότερης απόδοσης μόλις συσσωρευθεί επαρκής ενέργεια. Η αποθήκευση αναλύεται ως μηχανισμός με χρονικά εξαρτώμενη δυνατότητα παράδοσης ισχύος: η στρογγυλή απόδοση, οι διαρροές και η ωμική συμπεριφορά (ESR) προσδιορίζουν την άμεση ικανότητα υποστήριξης φορτίων με υψηλό crest factor, ιδίως σε εκπομπές δεδομένων μεγάλης πυκνότητας.

Η διαχείριση ισχύος στο χρόνο εκτέλεσης στηρίζεται σε φειδωλή αλλά επαρκή παρατήρηση: μετρήσεις V_{store} σε χαμηλούς ρυθμούς και αξιοποίηση συμβάντων από τα κυκλώματα προστασίας. Για τον περιορισμό του κόστους τηλεμετρίας εφαρμόζονται σχήματα event-driven ενημέρωσης. Στο φορτίο υιοθετείται παρτιδοποίηση διεργασιών ώστε να μειώνεται ο αριθμός αφυπνίσεων και να ελαχιστοποιούνται μεταβατικές δαπάνες: προηγείται δειγματοληψία και τοπική επεξεργασία, ακολουθεί συγκέντρωση δεδομένων και έπειτα πραγματοποιείται ενιαία μετάδοση σε κατάλληλα διαμορφωμένα ενεργειακά παράθυρα. Οι πολιτικές επικοινωνίας ευθυγραμμίζονται με την ενεργειακή κατάσταση μέσω προσαρμογής του μεγέθους payload, των παραμέτρων MAC/PHY και του χρονισμού, ώστε να περιορίζεται ο on-air χρόνος για τη δεδομένη ποιότητα ζεύξης και να αποφεύγονται αποτυχημένες επαναλήψεις.

Κεντρική συμβολή αποτελεί το DREAM (Dynamic Rest & Energy-Aware Management), ένα ελαφρύ εκτελεστικό υπόστρωμα που μεταχειρίζεται την ενέργεια ως πόρο πρώτης τάξης. Το DREAM οργανώνει τη λειτουργία σε δύο καταστάσεις: ανάπαυση-συσσώρευση, κατά την οποία αυξάνεται το ενεργειακό απόθεμα, και συμπαγή εκτέλεση εργασιών (burst execution), όπου ολοκληρώνονται δεσμευμένες λειτουργίες αίσθησης, υπολογισμού και επικοινωνίας με σεβασμό στα όρια παραδοτέας ισχύος και στις προθεσμίες. Η προσέγγιση βασίζεται σε υδατοσήματα αποθήκευσης, τμηματοποίηση μεγάλων κορυφών σε υπο-bursts με ενδιάμεση επαναφόρτιση και ευθυγράμμιση με αναμενόμενα πλατώ συγκομιδής βάσει απλών χρονικών προτύπων. Η πειραματική αξιολόγηση καταδεικνύει σημαντική μείωση brownouts, σταθερότερη ποιότητα υπηρεσίας και εξάλειψη φαύλων κύκλων επανεκκινήσεων και ανεπιτυχών ραδιοδοκιμών.

Συμπληρωματικά, προτείνεται ιχνο-καθοδηγούμενη μεθοδολογία διαστασιολόγησης. Με αντιπροσωπευτικά ίχνη πηγών και φορτίων, τα μετρημένα μοντέλα μετατροπών και αποθήκευσης αναπαράγουν την πραγματική ροή ενέργειας, επιτρέποντας να προσδιοριστούν τεχνολογίες και εμβαδά συλλεκτών, τύποι και χωρητικότητες αποθήκευσης, κατώφλια και υστερήσεις προστασίας, καθώς και περιοχές λειτουργίας μετατροπών που ικανοποιούν τις μετρικές αποστολής με ποσοτικοποιημένη εμπιστοσύνη. Όταν τα ίχνη είναι περιορισμένα, χρησιμοποιούνται συνθετικά μοντέλα χαμηλής πολυπλοκότητας που διατηρούν κρίσιμες χρονικές στατιστικές (μήκη πλατώ και υφέσεων), αποφεύγοντας τα σφάλματα τόσο των αναλύσεων μέσω των τιμών όσο και των υπερσυντηρητικών περιβλημάτων.

Η εμπειρική επικύρωση διεξάγεται σε κόμβους φωτοβολταϊκής ενέργειας εσωτερικού χώρου και σε συστήματα θερμοηλεκτρικής συγκομιδής χαμηλού ΔT , καθώς και σε πρωτότυπα με αναλυτική τηλεμετρία της κατάστασης του DREAM. Τα αποτελέσματα επιβεβαιώνουν ότι, υπό αντίστοιχους μέσους προϋπολογισμούς, η προτεινόμενη προσέγγιση επιτυγχάνει βιώσιμη υπηρεσία με σαφώς λιγότερα επεισόδια υποτάσεως, βελτιωμένη προβλεψιμότητα εκτέλεσης εργασιών και ανθεκτικότητα σε παγίδες ψυχρής εκκίνησης. Αναδεικνύεται επίσης η σημασία μετατροπών πολύ χαμηλής τάσης για boot-strap εκκινήσεις και η χρησιμότητα συσσώρευσης πριν από κορυφωμένες μεταδόσεις δεδομένων.

Η προτεινόμενη προσέγγιση είναι σκόπιμα μακροσκοπική, ώστε να παραμένει πρακτική. Δεν απαιτεί επίπονη ηλεκτρική ταυτοποίηση εξαρτημάτων· επαρκούν αξιόπιστες μετρήσεις quiescent και διαρροών και ιχνοσύνολα με ρεαλιστική χρονική ποικιλία. Το τίμημα της

φειδωλής παρατήρησης είναι ευρύτερες αβεβαιότητες, οι οποίες αντιμετωπίζονται με ασφαλείς ζώνες λειτουργίας και προσαρμοστικά υδατοσχήματα. Σε εφαρμογές υψηλής κρισιμότητας (π.χ. ιατρικές), οι τελετουργίες ασφάλειας και οι ενημερώσεις πρέπει να ομαδοποιούνται σε ήδη αφυπνισμένα χρονικά παράθυρα και να προβλέπεται ενεργειακή εφεδρεία για λειτουργίες ασφαλούς αποτυχίας.

Συνοψίζοντας, η διατριβή εισφέρει ένα συνεκτικό πλαίσιο που συνδέει τις ροές περιβαλλοντικής ενέργειας με εγγυήσεις επιπέδου συστήματος για κυβερνοφυσικούς κόμβους χαμηλής ισχύος. Τα απωλειο-ευαίσθητα μοντέλα επιτρέπουν σχεδίαση με ρητές μετρικές ουδετερότητας και επιβιωσιμότητας· το DREAM ενορχηστρώνει, με ελάχιστη τηλεμετρία, την εναλλαγή συσσώρευσης και συμπαγών εκτελέσεων χωρίς επεισόδια υποτάσεως· και η ιχνο-καθοδηγούμενη διαστασιολόγηση επιλέγει συλλογή, αποθήκευση και προστασία που εξασφαλίζουν αυτονομία με μετρήσιμη εμπιστοσύνη. Η κεντρική θέση είναι σαφής: η στιγμιαία δυνατότητα παράδοσης ισχύος είναι ισοδύναμης σημασίας με το μακροπρόθεσμο ενεργειακό ισοζύγιο και, με προσεκτικά σχεδιασμένα παρατηρήσιμα σήματα, τα μικροσυστήματα μπορούν να λειτουργούν προβλέψιμα και αξιόπιστα σε απαιτητικά περιβάλλοντα για εκτεταμένους ορίζοντες λειτουργίας.

Λέξεις κλειδιά: συγκομιδή μικροενέργειας, δυναμική διαχείριση ισχύος, ενεργειακή ουδετερότητα, διαστασιολόγηση συλλεκτών, κυβερνοφυσικά μικροσυστήματα

Table of Contents

Acknowledgements	1
Abstract.....	2
Περίληψη	3
Table of Contents.....	7
List of tables.....	9
Nomenclature	12
1 Introduction.....	15
1.1 Background and motivation	15
1.2 Overview of diffuse energy sources.....	16
1.3 Power management techniques for embedded systems.....	16
1.4 Challenges and opportunities in diffuse energy utilization	17
1.5 Problem statement.....	18
1.6 Research objectives.....	19
1.7 Contributions.....	19
1.8 Novelty and positioning.....	20
1.9 Organization of the Thesis	23
2 Fundamentals of cyber-physical systems.....	24
2.1 From cyber-physical systems to microsystems	24
2.2 Node-level energy and timing model	27
2.3 Representative application domains.....	31
2.4 Summary and outlook.....	35
3 Diffuse energy sources	37
3.1 Definition of energy flux “ Φ ”.....	38
3.2 Characteristics and energy content of source classes Φ	47
3.3 Availability.....	56
3.4 Harvesting potential	57

4	Energy collection architectures and components	60
4.1	Introduction	60
4.2	Reference architectures and energy collection model	61
4.3	Energy harvesting and conversion systems	75
4.4	Energy storage technologies	89
4.5	Power distribution and management strategies	104
5	Energy-Aware runtime power management	110
5.1	Introduction	110
5.2	Power management algorithms	112
5.3	Energy-Aware runtime system management	119
6	Design methodology for energy autonomous embedded microsystems	136
6.1	Scope, assumptions and notation	136
6.2	Input traces and window selection	137
6.3	Harvester sizing from window neutrality	139
7	Case studies and experiments	150
7.1	Energy balance modeling techniques for embedded systems	150
7.2	Research projects	158
7.3	Low- ΔT thermoelectric generator evaluation	170
8	Conclusion and future directions	180
8.1	Summary of findings	180
8.2	Contributions to the field	181
8.3	Recommendations for future work	182
8.4	Emerging trends and future research directions	184
	Appendix A	186

List of tables

Table 3.1 Source summary of $\Phi(t)$ by class 55

Table 4.1 Battery characteristics per technology 95

Table 5.1 Energy management algorithms comparison table.....135

List of Figures

Figure 3.1 Taxonomy of diffuse environmental energy	37
Figure 4.1 Ideal architecture.....	62
Figure 4.2 Non-ideal architecture	63
Figure 4.3 Photovoltaic transducer equivalent model	77
Figure 4.4. Equivalent of thermoelectric module	78
Figure 4.5 Triboelectric generator equivalent circuit	81
Figure 4.6 Accumulator technologies Power density vs Energy density.....	90
Figure 4.7 Electric Double Layer Capacitor	97
Figure 5.1 Two-state operation.....	123
Figure 5.2 DREAM state diagram	126
Figure 5.3 DREAM flowchart.....	132
Figure 7.1 Simulation model (top level)	152
Figure 7.2 Energy Input block.....	153
Figure 7.3 Load simulation block	153
Figure 7.4 Microprocessor simulation block.....	154
Figure 7.5 (top) Capacitor voltage (bottom) Calculated sleeptime.....	156
Figure 7.6 (top) Capacitor voltage (bottom) Calculated sleeptime.....	157
Figure 7.7 Power supply architecture.....	159
Figure 7.8 PERPS board	159
Figure 7.9 (left) Device test in the lab (right) Demo device in situ	160
Figure 7.10 (left) EFOS concept design (right) Pilot device.....	162
Figure 7.11 EFOS power supply architecture	163
Figure 7.12 EFOS prototype	164
Figure 7.13 Installation of EFOS system	165
Figure 7.14 AMSA architecture.....	166
Figure 7.15 AMSA Harvester schematic	166
Figure 7.16 AMSA Board.....	167
Figure 7.17 AMSA pilot devices	168
Figure 7.18 Sandwich formation.....	171
Figure 7.19 3D printed case	172
Figure 7.20 Perfboard circuit.....	172
Figure 7.21 Assembled perfboard.....	173
Figure 7.22 Complete setup	173
Figure 7.23 TEM evaluator GUI.....	175
Figure 7.24 LabView block diagram.....	176

Figure 7.25 Marlow TG12-6 output in high ΔT	178
Figure 7.26 Marlow TG12-6 output in low ΔT	178
Figure 7.27 ATS-TEC40-39-004 output in low ΔT	179

Nomenclature

Term	Definition
A (area)	Effective collector aperture (for minimum-area sizing).
$A_{\phi_{\min}}(T)$	Availability of Φ above threshold $\phi_{\min}$ over window T (fraction).
Brownout	Hysteretic undervoltage region where loads are shed until recovery.
Cold-start	Bring-up from discharged store once (VCS, ECS) criteria are met.
CPS	Cyber-physical system.
DVFS	Dynamic Voltage and Frequency Scaling.
$E(t)$	Stored energy state.
ECS	Start-up energy budget to reach regulation.
$E_{e,\lambda}(\lambda, t, n)$	Spectral irradiance at the POA (optical).
E_{in}	Charge-packet energy used in RTE/dwell relations.
E_T	Environmental energy over window T (harvesting potential).
$h_{\text{eff}}(t)$	Effective heat-transfer coefficient (thermal coupling).
$I_{\text{leak}}(V, T)$	Leakage current model used to parameterize P_{leak} .
MPPT	Maximum Power Point Tracking (input control).
n	Unit normal of the plane of array.
OCV	Open circuit voltage.

Term	Definition
OVP	Over-Voltage Protection.
$P_H(t)$	Harvested power referred to the input/storage side (post input-path losses).
$P_L(t)$	Aggregate load power (sense + compute + comms + quiescent).
PMIC / PMU	Power Management IC / Unit.
POA	Plane-of-Array (irradiance at device plane).
$P_q(t)$	Quiescent power of conversion/protection blocks (loss).
$P_{leak}(t)$	Storage leakage/self-discharge power (loss).
$Q_p(\cdot)$	p-th percentile operator (e.g., Q5, Q50, Q95 of ET).
Rail sequencing / power-good	Enable/disable order and status signaling for survivable operation.
RF	Radio Frequency.
RTE	Round-trip efficiency.
SIMO	Single-Inductor Multiple-Output (regulator topology).
SoC	State of Charge.
T	Generic window horizon used for ET/availability/sizing.
$T_{L,cycle}$	Load cycle period in duty-cycled workload abstraction.
TSI	Total Solar Irradiance.

Term	Definition
$u(t)$	External stimulus to the harvester (e.g., irradiance level).
UVLO	Undervoltage Lockout (with hysteresis).
VCS	Cold-start Threshold Voltage.
$V_{\text{store}}(t)$	Storage voltage (proxy for stored energy).
$\{\tau_k\}$	Task set with per-task energy cost e_k and deadline d_k .
$\Delta T(t)$	Temperature difference across an interface (thermal).
Λ	Optical spectral band of interest (e.g., 300–1200 nm for PV).
$\Pi b(t), \Phi b(t)$	Band-integrated vibration/RF measures over band b .
$\Phi(t)$	Environmental energy flux at point of use (e.g., $\text{W} \cdot \text{m}^{-2}$ for optical).
Φ_{min}	Source-side threshold defining “usable” Φ (units of Φ).
$\eta_{\text{ch}}, \eta_{\text{dis}}$	Charge/discharge multipliers for DC-boundary RTE.
$\eta_{\text{dc}}(u, x)$	DC conversion efficiency vs. operating point.
$\eta_{\text{H}}(u, x)$	Harvester efficiency vs. stimulus u and internal state x .
$\eta_{\text{in}}(t)$	Input-path efficiency (source/front-end $\rightarrow$ store side).
$\eta_{\text{out}}(t)$	Output-path efficiency (store $\rightarrow$ rails).

1 Introduction

Autonomous cyber-physical microsystems (μ CPS) are increasingly embedded in natural, urban, and industrial environments to perceive context, execute local inference, and actuate with minimal human oversight. Their practical value—whether in long-horizon environmental monitoring, building automation, logistics, or clinical wearables—hinges on the ability to operate predictably under tight energetic constraints. While algorithmic advances and low-power hardware have reduced per-operation cost, field deployments routinely reveal that energy acquisition, storage, and delivery remain the dominant determinants of feasibility and quality of service. Battery-centric designs face recurring maintenance, supply-chain, and environmental burdens; simply appending harvesters is insufficient because ambient sources are variable, the power path is lossy, and high crest-factor tasks (e.g., radio bursts, sensor warm-ups) can induce brownouts despite adequate average energy. These observations motivate a holistic treatment of energetics that spans source characterization, loss-aware power-path modeling, storage dynamics, and runtime scheduling.

This thesis addresses that need by developing a loss-aware, node-centric framework that turns diffuse environmental energy into predictable service for sensing, computation, and communication. It combines: (i) a macroscopic, store-side model that captures conversion losses, storage dynamics, and protection behavior; (ii) a runtime scheduler that enforces window-based energy neutrality and survivability using minimal telemetry; and (iii) a trace-driven sizing methodology that maps mission requirements into concrete design choices for sources, storage, and thresholds.

These contributions are validated on representative μ CPS case studies spanning indoor photovoltaic, low- ΔT thermoelectric, and mixed-source deployments. Replay of measured and synthesized traces demonstrates that loss-aware modeling and runtime control substantially reduce brownouts and improve availability compared to baseline duty-cycling and average-power sizing, even under weak and intermittent sources.

1.1 Background and motivation

These microsystems integrate sensing, low-power compute, and short-range/LPWAN radios. In field deployments, feasibility and quality of service are primarily constrained by energetics rather than functionality. Variability in energy acquisition, conversion losses, storage leakage/quiescent draws, and high-crest-factor loads (e.g., radio bursts, sensor

warm-ups, flash writes) cause deadline misses and brownouts despite adequate average harvested power.

Scaling batteries or deepening sleep extends lifetime but preserves logistics, environmental cost, and safety risks. Adding a harvester is not a drop-in fix: indoor light, low-grade thermal, and weak vibration sources are intermittent and low density; interfaces introduce cold-start limits, MPPT overheads, diode/converter drops, leakage, and protection thresholds that can gate operation. Consequently, average-power reasoning is insufficient without accounting for the lossy power path and the timing structure of work.

At a high level, we take a node-centric, loss-aware, timing-sensitive view: each microsystem sees stochastic input power, a leaky store, a conversion/protection stack, and a bursty task set with deadlines. Mission outcomes are formalized in terms of energy balance and protection limits and achieving them requires co-design of hardware sizing and runtime policies that schedule activity in concert with energy availability. A precise abstraction and notation are given in the problem statement of this chapter and in Chapter 2.

High-fidelity telemetry and model-predictive control are often impractical at microwatt scales. Practical systems regulate on sparse observables—primarily storage-voltage trajectories and protection events—while shaping rest and activity to meet QoS with negligible overhead. This motivates a runtime that admits, defers, or aggregates work based on store state and recent protection behavior, and a trace-driven sizing flow that links environmental variability, power-path losses, and workload timing to concrete guarantees.

1.2 Overview of diffuse energy sources

At microsystem scale, environmental energy is low-density and intermittent. Indoors, optical flux is orders of magnitude below outdoor insolation and shaped by building geometry and occupancy; low- ΔT thermal gradients evolve slowly and are easily flattened by packaging; vibration sources are narrowband and site-specific. Rather than tying the thesis to any particular device, we treat the environment generically as a time-varying source $\Phi(t)$ and defer detailed transducer physics to Chapter 3, which supplies device-agnostic envelopes of windowed energy over missions of interest.

1.3 Power management techniques for embedded systems

Power management in μCPS spans four interacting layers: (i) sensing and actuation front ends, whose bias and conditioning costs set a non-trivial energy floor; (ii) the compute substrate, where deep-sleep states, DVFS, and peripheral-triggered wake reduce dynamic

and static drains; (iii) the communication stack, whose preambles, listen windows, and contention behavior often dominate payload energy; and (iv) the power subsystem, where conversion, protection, and storage choices shape the admissible power-path trajectories. Across these layers, control policies are either reactive—regulating on store proxies and protection events—or predictive—admitting work based on short- to medium-horizon forecasts. The thesis instantiates these ideas in DREAM, a lightweight runtime that enforces window-based energy neutrality and survivability using only storage-voltage trajectories and protection events.

1.4 Challenges and opportunities in diffuse energy utilization

Diffuse sources exhibit intermittency with structure: plateaus and lulls follow human or machine rhythms, so identical averages can yield different outcomes depending on deadline alignment. Feasibility is further constrained by a lossy, stateful power path—efficiency cliffs, quiescent drains, cold-start budgets, current limits, and protection hysteresis often cap deliverability for short bursts. At microwatt scales, observability is sparse because fine-grained metering can be net-negative; policies must remain stable using only a storage proxy and protection events. Finally, scale and integration matter: placement, packaging, and parasitics perturb the mapping from environmental flux to harvested power, while EMC and analog integrity interact with burst shaping, and aging drifts erode initial margins.

These constraints also expose opportunities. Workloads can be shaped—batching and coalescing sensing, compute, and communication during harvest plateaus—to reduce wake/sleep churn, amortize handshakes, and manage crest factor with planned headroom; edge inference can trade airtime for local MAC operations under scarcity. Reconfigurable power paths and hybrid stores (e.g., supercapacitor – microbattery) improve effective deliverability while bounding quiescent losses through staged thresholds and selective front-end enables. Trace-driven sizing and verification—replaying representative source and workload traces through measured converter and storage models—reveals cold-start traps, retry storms, and lull-induced depletions, enabling evidence-backed choices of harvester area, storage capacity, and thresholds. Finally, lean telemetry contracts and decision-relevant specifications (μW -scale efficiencies, quiescent currents, spectra and duty factors) reduce design margins and enable portable runtime control.

1.5 Problem statement

We seek to design autonomous cyber-physical microsystems (μ CPS) that run on diffuse, variable energy while providing predictable service under tight budgets. The core difficulty is not average harvested power, but the combination of a lossy, stateful power path and bursty, deadline-driven workloads observed only through sparse telemetry. The goal is a principled way to size hardware and schedule work so that nodes remain energy-neutral and survivable over explicit horizons, meeting application timing/quality constraints without high-overhead metering.

We adopt a node-centric abstraction. Let $\Phi(t)$ be incident source flux; $PH(t)$ the conditioned electrical power after conversion/protection; $E(t)$ the stored-energy state with measurable proxy $V_{\text{store}}(t)$; and $PL(t)$ the load power of sensing, computing, memory, and communication. Conversion/protection are captured by state-dependent efficiencies (η_{in} , η_{out}), quiescent drains, and thresholds/hysteresis that govern cold-start and UVLO. Storage contributes leakage and finite deliverability (ESR, converter current limits). The workload is a task sequence $\{\tau_k\}$ with energy e_k , deadlines d_k , and crest requirements (peak power or minimum headroom during execution). Success is defined by neutrality over a window T (harvested-after-losses $\geq$ costs), survivability ($E(t)$ within protection bounds with burst headroom), and predictability (each τ_k meets d_k without violating deliverability constraints).

The problem is to co-design: (i) harvester/storage sizing and key thresholds, and (ii) a runtime admission/timing policy that uses low-cost observables—primarily $V_{\text{store}}(t)$ and discrete protection/status events—to decide when to rest, aggregate charge, and execute bursts, explicitly handling cold-start, quiescent drains, converter operating regions, and instantaneous deliverability while honoring task deadlines.

Constraints narrow admissible solutions. Telemetry is scarce (high-rate metering/coulomb counting is off-budget). Availability is structured, not white noise (plateaus/lulls from occupancy, HVAC, machine timetables, network calendars). Models must be macroscopic yet loss-aware, exposing the phenomena that dominate at μ W scales while remaining tractable for design and control. Verification must be trace-grounded, replaying representative source and workload traces through measured converter/storage models to surface corner cases (e.g., dawn start-ups, retry storms, lull-induced depletions) that average-power analyses miss.

1.6 Research objectives

Objective 1 — Loss-aware node model. Formalize a macroscopic, tractable model that maps incident flux $\Phi(t)$ to conditioned input power $P_H(t)$ through realistic conversion/protection and storage, and represent workloads as deadline-bounded, bursty tasks. Establish mission metrics—energy neutrality over explicit horizons and survivability within protection bounds—and use them as design targets and evaluation criteria throughout the thesis.

Objective 2 — Decision-relevant source characterization. Provide parsimonious, μ CPS-scale descriptors of diffuse sources (plateau/lull statistics, availability windows) that preserve temporal structure and interface non-idealities shaping $\Phi \rightarrow P_H$, avoiding high-overhead metering.

Objective 3 — Minimal-telemetry runtime. Develop an admission/timing policy (DREAM) that regulates using low-cost observables—principally a store proxy and discrete protection events—to decide when to rest/accumulate and when to execute bursts, remaining stable across cold-start and efficiency cliffs and coordinating with comms to prevent retry storms.

Objective 4 — Trace-driven sizing. Translate autonomy targets into concrete hardware choices (harvester technology/area, storage type/capacity, thresholds/hysteresis, converter operating regions) by replaying representative input/load traces through measured macromodels to yield quantitative neutrality/survival guarantees and surface corner cases. Where traces are scarce, introduce lightweight synthetic generators that preserve mission-relevant statistics.

Objective 5 — Empirical validation. Demonstrate practicality on representative μ CPS case studies with instrumentation limited to low-cost observables, comparing against battery-only and average-power controllers that ignore deliverability/protection dynamics, and showing sustained predictable service with materially fewer brownouts.

1.7 Contributions

1. **Loss-aware node model (Chapters 2 and 4).** A macroscopic, store-side model for energy-autonomous μ CPS that captures harvested power, load power, storage dynamics, conversion efficiencies, and protection events. The model is explicitly parameterized by bench-measurable quantities and supports analysis in terms of windowed energy neutrality, survivability, and availability.

2. **DREAM runtime (Chapter 5).** A constant-overhead runtime scheduler that operates on fixed-length epochs using only a storage proxy and protection events. DREAM enforces window-based neutrality when energy is surplus and switches to a survival regime when energy is scarce, guaranteeing a user-defined survival horizon under model assumptions.
3. **Trace-driven sizing methodology (Chapter 6).** A design methodology that replays measured or faithfully synthesized input traces through the loss-aware model to determine minimum collector area, storage capacity, and protection thresholds that satisfy neutrality, survivability, and availability targets over explicit mission windows.
4. **Empirical validation and case studies (Chapter 7).** Model-driven and hardware-in-the-loop experiments on real boards and traces from diffuse sources (indoor PV, low- ΔT TEGs, etc.), demonstrating that the proposed framework improves autonomy and service availability compared to existing energy-neutral duty-cycling and naive sizing approaches.

1.8 Novelty and positioning

Energy-harvesting micro–cyber-physical systems (μ CPS) have been studied from a variety of angles, including energy-neutral duty cycling, harvester-aware control, intermittent-computing abstractions and storage sizing techniques. Existing approaches typically make at least one of the following assumptions: (i) the energy chain is ideal or only partially accounts for losses; (ii) controllers have access to rich telemetry (e.g. harvester output and per-subsystem power metering); or (iii) design-time sizing and runtime policies are treated as largely separate problems. This thesis departs from these assumptions by adopting a store-centric perspective and using it consistently across modelling, runtime control and design-time sizing.

- **1. Node-centric, loss- and protection-aware model**

A first contribution is a node-centric model of the energy chain that is explicitly **loss- and protection-aware** and is formulated at the **energy store interface**. Instead of assuming ideal harvest, conversion and storage, the model incorporates:

- conversion losses (harvester, DC–DC and regulation stages),

- storage non-idealities (charge/discharge efficiency, leakage, rate-dependent behavior),
- and protection mechanisms (UVLO, OVP and related protection thresholds and hysteresis).

By working at the store interface, the model aligns with what is practically observable and controllable on μ CPS nodes: storage voltage, energy content and protection events. This contrasts with much prior work that either abstracts away the energy chain behind an idealized battery model or focuses on harvester-side behavior without fully propagating losses and protections to the node-level energy budget. The model introduced here is therefore tailored to the realities of low-power μ CPS, where losses and protections are not second-order effects but dominant determinants of deliverable energy and system availability.

- **2. Minimal telemetry runtime (DREAM)**

A second contribution is the DREAM runtime, which implements energy-aware control using **only storage-side telemetry** and a constant, hardware-friendly overhead. DREAM does not assume access to detailed harvester measurements or per-subsystem power metering. Instead, it bases its decisions on:

- the instantaneous or filtered state of the energy store (e.g. storage voltage as a proxy for stored energy), and
- protection events and thresholds (e.g. UVLO triggers, minimum reserve margins).

Within this constrained telemetry budget, DREAM enforces **window-based neutrality and survivability objectives**, by modulating the timing and intensity of workload execution. The policy alternates between accumulation and active phases according to explicit energy windows and reserves, rather than relying on implicit safety margins or average-power arguments.

In contrast, many existing controllers either (i) require richer measurement infrastructure that is often unavailable or too costly on resource-constrained μ CPS nodes, or (ii) operate with simple threshold policies that do not tie back to a formal model of neutrality and survivability. The novelty here lies in demonstrating that such guarantees can be approached

using only store-centric measurements and a lightweight runtime that is compatible with typical embedded platforms.

- **3. Trace-driven sizing tied to runtime guarantees**

A third contribution is a **trace-driven sizing methodology** that is explicitly tied to the runtime's guarantees and uses the same store-centric, loss-aware formalism. Instead of sizing storage and protection thresholds based on average harvested power or static safety factors, the proposed methodology:

- uses measured or realistic $\Phi(t)$ traces of net inflow/outflow at the store,
- evaluates energy-neutrality and survivability over explicit time windows, and
- derives storage capacity and threshold configurations through inclusion-style arguments on these windows.

This trace-driven approach directly accounts for the temporal structure of harvested energy and load, the non-idealities of the energy chain and the protection behavior of the node. It thereby avoids common pitfalls such as underestimating cold-start requirements, ignoring protection-induced outages or overestimating effective capacity due to unmodelled losses.

While trace-based evaluation has been used in prior work, it is typically employed as a post hoc validation tool, or is not systematically linked to runtime policies and protection mechanisms. In this thesis, trace-driven sizing is positioned as a **design-time counterpart** to DREAM: both analysis and runtime decisions are expressed in terms of the same node-centric quantities and window-based metrics.

- **4. Unified perspective and scope**

The final aspect of novelty is the **unified perspective** that spans modelling, runtime and design-time analysis. The same store-centric, loss- and protection-aware formalism:

- defines the metrics of interest (neutrality, survivability, availability),
- underpins the DREAM runtime's control decisions,
- and drives the trace-based sizing and evaluation of real μ CPS nodes.

This avoids the fragmentation often seen in the literature, where idealized models are used for design, different abstractions are used for runtime control and yet another set of assumptions underlies evaluation. By maintaining a single, explicit view of the energy store and its losses and protections, the thesis provides a coherent framework in which guarantees, policies and sizing choices can be compared and reasoned about.

Within this scope—energy-harvesting, low-power μ CPS that operate from non-ideal energy chains and expose only store-side telemetry—the framework proposed here therefore occupies a distinct position: it offers a practically grounded, analytically explicit and implementation-ready path from measured $\Phi(t)$ to predictable service, under realistic constraints on losses, protections and observability.

1.9 Organization of the Thesis

Chapter 2 specializes canonical CPS concepts to microsystem scale and introduces the store-side energy and timing model and telemetry contract used throughout the thesis. Chapter 3 characterizes diffuse environmental energy sources via geometry-aware, device-agnostic flux models and windowed energy envelopes. Chapter 4 details energy-collection architectures and components, embedding non-idealities and protection behavior into the loss-aware node model and defining neutrality, survivability, and availability metrics. Chapter 5 presents DREAM, the runtime that enforces window-based autonomy using minimal telemetry. Chapter 6 develops the trace-driven sizing methodology based on the inclusion argument over input traces. Chapter 7 validates the framework on representative case studies using measured and synthesized traces. Chapter 8 summarizes the contributions, discusses limitations, and outlines directions for future work.

2 Fundamentals of cyber-physical systems

Chapter 2 develops the foundations for cyber-physical microsystems (μ CPS) at the node scale. The focus is deliberately narrow: a single device that senses or actuates, computes, communicates, and must do so from a constrained, often harvested, energy budget. We adopt a hardware-first, store-side view of energetics—harvesting, conversion/protection, storage, and distribution—and make explicit the signals and constraints that govern feasibility at μ W–mW power. Intermittent sources, lossy power paths, and protection dynamics (UVLO/OVP, inrush limits, ESR-limited droop) determine what work can be performed and when. To keep analysis aligned with what embedded runtimes can actually observe, we rely on minimal telemetry (store-state proxies and protection events) and on windowed autonomy metrics—neutrality, survivability, and availability—over explicit mission horizons.

This chapter has three roles. First, it specializes canonical CPS formulations to microsystem scale, fixing the node as the unit of analysis and clarifying how plant, network, and node interact. Second, it introduces a store-side energy and timing model that relates harvested power, load power, storage state, and protection logic. Third, it defines the telemetry contract between hardware and any runtime or design-time logic, emphasizing minimal observability.

2.1 From cyber-physical systems to microsystems

2.1.1 Definition and scope

Canonical definitions describe a cyber-physical system (CPS) as a tight integration of computation, communication, and physical processes, where embedded computers monitor and control physical dynamics, usually in feedback, subject to explicit timing and reliability constraints [1], [2], [3]. In this thesis we refine that view to the microsystem scale: a cyber-physical microsystem (μ CPS) is a single device that closes a feedback or inference loop with its local environment by combining:

- (i) concrete transducers and analog interfaces for sensing and/or actuation,
- (ii) a resource-constrained embedded computing platform,
- (iii) a protocol-bound communication interface, and
- (iv) a lossy power subsystem that harvests, converts, stores, and safeguards energy.

Conceptually, this is a specialization of the CPS picture—software interacting with a physical process in a closed loop—to a node whose behavior is bounded by embodiment (packaging, interfaces, conversion paths) and energetics (variability of input power, bursty loads, protection thresholds). This specialization is consistent with mainstream CPS

definitions and frameworks [1], [2], [3], and fixes the unit of analysis at the node level, which is where energy autonomy is ultimately decided.

Formally, let the plant P have state $x_p(t)$ and output $y(t)$. Let the controller C be realized on an embedded hardware–software stack

$$E = \langle HW, SW, I, O, \Pi \rangle \quad (2.1)$$

where HW and SW denote hardware and software, I/O the physical and logical interfaces, and Π the non-functional constraints (timing, memory, and—central here—energy). When a communication substrate N is present, transport delay and loss enter the closed loop. We therefore view a microsystem as the triple

$$\mathcal{C} = \langle C, P, N \rangle \quad (2.2)$$

together with its realization on E . This formulation aligns with CPS treatments that emphasize explicit modelling of time, implementation constraints, and the mapping from algorithms to embedded platforms [2], [3].

By qualifying CPS with microsystem, we restrict scope to devices whose size, mass, and energy budgets are tight enough that interface bias currents, radio ramp-up transients, conversion inefficiencies, storage leakage, and protection hysteresis materially shape feasible schedules. This emphasis reflects the trajectory from wireless sensor networks to today’s batteryless and energy-harvesting IoT nodes, where intermittent input power and high-crest-factor workloads dominate design choices. Contemporary surveys document the centrality of harvesting variability, conversion losses, and storage behavior to node-level feasibility and motivate the node-centered perspective adopted here [4], [5].

Within this thesis, the microsystem is the unit of analysis. The plant and any network are treated as environment coupled through I/O interfaces. Power may be supplied by ambient energy and/or a local store; when the thesis later asserts energy autonomy, it will do so over an explicit mission horizon and with respect to the energy required for sensing, computation, and communication on the controller platform. This node-wise stance is consistent with CPS frameworks that separate conceptual models from realization and assurance, and with recent work on intermittently powered, batteryless systems that formalizes autonomy and progress in the presence of power failures [3], [6].

2.1.2 Node-centric signals and proxies

To keep reasoning consistent from sources $\rightarrow$ power path $\rightarrow$ runtime $\rightarrow$ sizing, we fix a small set of signals and proxies that will be used throughout the thesis:

- $P_H(t)$: harvested and conditioned input power at the store interface (after harvester and front-end conversion).
- $P_L(t)$: aggregate load power, including sensing, compute, communication, and quiescent drains.
- $E(t)$: stored-energy state, with measurable proxy $V_{\text{store}}(t)$ (for capacitive stores the mapping is exact up to ESR; for batteries it is approximate and may require auxiliary estimation).
- $\eta_{\text{in}}(t), \eta_{\text{out}}(t)$: input/output conversion efficiencies of the power path, capturing macroscopic losses such as start-up overheads, quiescent currents, and finite round-trip efficiency.

Analyses are expressed on the store side of the power path using these variables, avoiding double counting and keeping non-idealities explicit. Over a finite analysis window T , autonomy is characterized by windowed metrics built from $P_H(t)$, $P_L(t)$, and $E(t)$:

- neutrality (harvested-after-losses $\geq$ expended energy over T),
- survivability (the trajectory $E(t)$ remains within protection bounds with headroom for bursts), and
- availability (the fraction of time the node can admit its workload while satisfying neutrality and survivability).

These notions follow system-level treatments of energy-harvesting microsystems and intermittent computing, which stress realistic windows, conversion and storage losses, and bursty workloads rather than long-horizon averages [4], [5], [7].

2.1.3 Microsystem architecture

At the scale of a single node, a cyber-physical microsystem is a concrete integration of four interacting subsystems:

- (i) sensing/actuation interfaces,
- (ii) a resource-constrained computing platform,
- (iii) a short-range communication interface, and
- (iv) a lossy power subsystem that conditions harvested or stored energy.

This partition is conventional in surveys of sensor/IoT node hardware and is the level at which implementation choices most strongly constrain feasibility under energy autonomy [4], [5], [8]. In subsequent sections we treat these subsystems through the lens of the node-centric signals above: sensing and actuation as contributors to $P_L(t)$; compute and communication as bursty, deadline-constrained workloads; and the power subsystem as the mechanism that maps environmental flux into $P_H(t)$, $E(t)$, and protection events that gate operation.

The remainder of the chapter builds on this foundation. Section 2.2 introduces the detailed node-level energy and timing model; Section 2.3 discusses telemetry and observability constraints; and Section 2.4 uses representative application domains to illustrate how these abstractions manifest in practice and motivate the runtime and sizing strategies developed in later chapters.

2.2 Node-level energy and timing model

The abstractions in § 2.1 fix the microsystem as the unit of analysis and identify the key store-side signals: harvested power $P_H(t)$, load power $P_L(t)$, stored-energy state $E(t)$ (with proxy $V_{\text{store}}(t)$), and conversion efficiencies $\eta_{\text{in}}(t)$, $\eta_{\text{out}}(t)$. This section turns those signals into an explicit energy–timing model for a node. The aim is not to capture every circuit-level detail, but to obtain a compact, implementation-aware contract that is realistic for energy-harvesting microsystems and batteryless devices [4], [5], [9], [10].

Under this model, feasibility and autonomy over a finite mission horizon are judged by whether the node can realize its workload schedule within the constraints of the energy balance, protection mechanisms, and limited observability.

2.2.1 Power and energy model

A cyber-physical microsystem can be viewed as an energy-processing system that transforms ambient or supplied power into useful work (sensing, computation, communication) while incurring conversion, leakage, and quiescent losses [4], [5], [9], [10].

At the store interface we write, at macroscopic time scales,

$$\frac{dE(t)}{dt} = \eta_{\text{in}}(t) P_H(t) - \frac{1}{\eta_{\text{out}}(t)} P_L(t) - P_{\text{leak}}(t) \quad (2.3)$$

where:

- $E(t)$ is the energy in the storage element(s);

- $P_H(t)$ is the conditioned input power delivered to the store after harvester and front-end conversion;
- $P_L(t)$ is the aggregate load power drawn by sensing, computing, communication, and quiescent drains;
- $P_{\text{leak}}(t)$ models self-discharge and parasitic currents (e.g., supercapacitor leakage, bias currents in front-end circuits) [4], [5], [11], [12], [13];
- $\eta_{\text{in}}(t)$ and $\eta_{\text{out}}(t)$ encode input and output path efficiencies, including start-up overheads, finite round-trip efficiency, and state-dependent behavior of converters [9], [10], [14], [15].

In practice, $E(t)$ is not directly measured. Instead, the runtime observes proxies such as capacitor voltage, ADC'd battery voltage, or coulomb-counted SoC [11], [12], [13], [16], [17]. For capacitive stores, $E(t)$ can be recovered from $V_{\text{store}}(t)$ up to ESR and non-idealities. For batteries or hybrid stores, SoC is estimated using model-based observers or coulomb counters, with non-negligible uncertainty [11], [12], [13], [15], [16], [17]. DREAM therefore treats $E(t)$ as a latent variable and uses $V_{\text{store}}(t)$ and event-level telemetry as stateful proxies, as detailed later.

On the load side, it is convenient to decompose

$$P_L(t) = P_{\text{sen}}(t) + P_{\text{cmp}}(t) + P_{\text{rad}}(t) + P_q(t) \quad (2.4)$$

where $P_{\text{sen}}(t)$ captures transducer and front-end consumption (including bias and acquisition), $P_{\text{cmp}}(t)$ the digital computing core, $P_{\text{rad}}(t)$ the radio, and $P_q(t)$ all quiescent and housekeeping drains (e.g., RTCs, reference circuits, always-on regulators). Surveys of node architectures and low-power MCU platforms stress that $P_q(t)$ and wake-up overheads can dominate long-term energy use if left unmanaged [8], [14], [18].

Protection constrains the admissible energy trajectory. Undervoltage lockout (UVLO), over-voltage protection (OVP), current limits, and brownout detection define a safe operating envelope

$$E_{\text{min}} \leq E(t) \leq E_{\text{max}} \quad (2.5)$$

within which control logic and data retention remain reliable [9], [10], [16], [19]. If $E(t)$ falls below the lower bound, the node is reset or held off; if it exceeds the upper bound, input is curtailed or dissipated to protect the store. These mechanisms effectively induce a state-

dependent on/off gating of $P_L(t)$, which is captured in later runtime and sizing analyses via inclusion arguments on admissible schedules.

Within a finite analysis window T , we use three windowed notions that recur throughout the thesis:

- **Neutrality:** net harvested energy after conversion minus net expended energy is non-negative over T .
- **Survivability:** the energy trajectory $E(t)$ respects protection bounds (and retains headroom for mandated bursts) throughout T .
- **Availability:** the fraction of time over T during which the node can admit its workload without violating neutrality or survivability.

These definitions echo system-level treatments of energy harvesting nodes and intermittent computing, which emphasize finite windows, conversion and storage losses, and bursty loads over simple long-term averages [4], [5], [6], [7], [9], [10].

2.2.2 Timing and workload characteristics

Time is the binding resource that couples computation, communication, and energetics. Even in small μW – mW nodes, tasks are typically constrained by deadlines, periods, and protocol windows rather than pure throughput [20], [21].

We model the microsystem workload as a set of tasks $\{\tau_k\}$, each characterised by:

- an arrival pattern (periodic, sporadic, or event-triggered),
- a worst-case execution-time and associated energy cost (C_k, e_k) ,
- a relative deadline d_k ,
- and, for communication tasks, a set of admissible transmission windows imposed by the link layer or MAC protocol.

Event-driven workloads are bursty by design: sensing and inference are concentrated into short intervals, often aligned with physical events or network schedules, followed by long idle intervals [20], [21]. Experiments on synchronous transmissions, low-power wireless stacks, and time-sensitive networking consistently show that these bursts have high crest factors (peak-to-average ratios), and their timing is tightly constrained by slot boundaries and rendezvous schedules [20], [21], [22], [23].

Intermittent power introduces a distinct temporal failure mode: power failures act as a non-preemptive, state-losing “scheduler” if not explicitly managed [6], [7]. Intermittent

computing work formalizes progress and correctness under such failures, typically by assuming an abstract power-failure model and providing software techniques (check-pointing, idempotent tasks, energy-aware scheduling) that tolerate them [6], [7]. In this thesis we adopt a complementary stance: we treat intermittent power as a property of the energy trajectory $E(t)$ and design runtimes and sizing methods that avoid or bound such failures within the hardware-imposed envelope (§ 3–§ 6), rather than absorbing them purely in software.

For the purposes of DREAM and the accompanying sizing framework, it is sufficient to assume that:

1. Each task τ_k has a well-defined energy cost envelope e_k under the target platform and operating points (e.g., DVFS settings).
2. Tasks interact with the power path only through their contribution to $P_L(t)$ and through induced timing choices (e.g., when radios are turned on).
3. Protocols such as TSCH, BLE, or LPWAN impose hard and soft time windows, which the scheduler must respect while ensuring that the necessary energy packets are available at the appointed instants [20], [21], [22], [23].

These assumptions allow us to couple the energy balance and the workload at the coarse time scales relevant to neutrality and survivability, while still capturing the tight coupling between event timing and energy availability.

2.2.3 Telemetry and operating constraints

Telemetry in a cyber-physical microsystem is sparse by necessity. Nodes rarely carry high-resolution energy metering; instead, they expose a small number of signals that reflect their energetics only indirectly: digitized store voltage, selected status bits (e.g., UVLO/OVP flags), possibly a coulomb-counter register, and a few performance counters [4], [5], [9], [15], [17]. These observables define the information contract between the physical node and any runtime or sizing logic.

We therefore partition observability into three classes:

1. Continuous proxies, such as $V_{\text{store}}(t)$ sampled at low rate, which provide coarse information about $E(t)$ but are affected by load-induced droop, ESR, and non-linear SoC relationships [11], [12], [13], [16].
2. Event-level telemetry, such as UVLO trips, brownouts, resets, and over-current flags, which encode safety violations and hard failures. These events are central to our

trace-driven sizing and validation methods because they directly reflect breaches of survivability at the store [9], [16], [19].

3. Self-metering mechanisms, such as smart coulomb counters and hybrid PMUs that can track net energy flow or SoC within known error bounds [15], [17]. These offer richer information but at the cost of area, power, and complexity, and are therefore not always present.

Any runtime that seeks to guarantee timing or energy properties must be correct under this limited observability and robust to the associated uncertainties [4], [5], [9]. In DREAM, this is addressed by (i) expressing guarantees with respect to observable proxies rather than unobservable internal variables, and (ii) designing sizing and runtime policies that operate with uncertainty bands on $E(t)$ derived from $V_{\text{store}}(t)$ and event statistics rather than requiring exact state knowledge.

When explicit self-metering is employed—for example, closing the loop between design-time sizing and deployment-time behavior must be integrated into this contract: its own consumption belongs in $P_q(t)$, its quantisation and drift define additional uncertainty, and any guarantees must hold under these observability constraints [15], [17].

Throughout this thesis, the runtime is assumed to have access to a deliberately minimal set of node-level telemetry signals: (i) the storage voltage $V_{\text{store}}(t)$, sampled no faster than the main control loop; (ii) discrete under- and over-voltage protection events originating from the power-path hardware; and, optionally, (iii) a low-rate, low-precision estimate of net charge flow into and out of the store. No per-subsystem power metering, high-rate current probing, or harvester-side sensing beyond these signals is assumed. All algorithms and guarantees in later chapters are enforced under this telemetry contract.

2.3 Representative application domains

The node-centric abstractions above are general, but concrete design choices are always shaped by application context. This section briefly reviews representative cyber-physical domains and explains how their domain-level requirements map to node-level energetics, timing, and autonomy constraints. The focus is on highlighting the energy-autonomy hook in each case, rather than providing an exhaustive survey.

2.3.1 Infrastructure-scale power systems

Power-system CPS bind physical grid dynamics to cyber functionality such as state estimation, protection, and control. Recent surveys emphasize multi-domain modelling that

simultaneously captures continuous-time power flows, discrete cyber logic, and the communication network [24]. Co-simulation frameworks such as HELICS and integrated transmission–distribution platforms enable detailed studies of coupled electromechanical and ICT dynamics, including latency, packet loss, and control interactions across scales [25], [26], [27].

From this system-level vantage point, the microsystem appears at the grid edge: micro-phasor measurement units (μ PMUs), substation monitors, and low-voltage IoT nodes push high-value measurements into the CPS [28], [29], [30], [31]. Their workloads are shaped by grid-side requirements (synchro phasor sampling, event detection, alarms) and by backhaul or LPWAN constraints. When these nodes are powered by constrained or harvested energy—e.g., in retrofit deployments or remote assets—the deliverability of time-synchronized measurements and alarms hinges on whether the lossy power path and storage can support bursty reporting and over-the-air updates within neutrality and survivability margins.

Thus, even though the overall CPS spans transmission to distribution, our analysis observes it through the lens of an edge microsystem whose energetics must be provisioned to uphold the timing and reliability contracts implied by protection and monitoring schemes.

2.3.2 Industrial and building systems

Industrial and building automation environments are canonical facility-scale CPS: they integrate occupancy sensing, HVAC control, lighting, safety, and security functions, often under explicit comfort and energy-efficiency objectives [32], [33]. Digital twin approaches, model-based control, and AI-driven optimization increasingly shape the supervisory layer, while the field layer is populated by large numbers of embedded nodes—access sensors, thermostats, IAQ monitors, asset trackers—connected via short-range wireless or fieldbuses [32], [33], [34].

The three-level view (field, automation, management) makes explicit how dynamics at each layer couple: physical processes (thermal mass, occupancy, daylight) evolve slowly but with significant variability; automation controllers enforce control policies subject to timing and quality-of-information constraints; management functions aggregate data and coordinate optimization across zones and systems [32], [33].

Microsystems at the field layer are where these guarantees must be realized under tight energy budgets. Battery-powered or energy harvesting nodes are expected to last years

without maintenance yet must honor strict wake-window schedules for BLE, 802.15.4e TSCH, or proprietary LPWAN links [22], [34]. Security standards such as ISA/IEC 62443 further require resilience under disturbance and attack, including fail-safe behavior and secure reconfiguration [32], [35]. All these requirements translate into constraints on $P_L(t)$, required duty cycles, and reserve energy for rare but critical events (alarms, rekeying, firmware updates).

In this context, the node-centric energy model provides a way to reason about whether proposed control and communication schedules are feasible under energy autonomy, rather than merely under idealized assumptions about battery capacity.

2.3.3 Medical & wearables

Clinical cyber-physical systems couple body-proximate sensing and actuation with data processing and clinical decision-making, often under stringent safety and regulatory constraints [36], [37], [38], [39], [40]. Wearable and implantable devices increasingly operate battery-free or with aggressively minimized batteries, relying on energy harvesting and wireless power transfer to achieve long-term autonomy while maintaining sensing fidelity and connectivity [36], [37].

At the system level, medical CPS integrate sensors, gateways, clinical backends, and human operators. At the microsystem level, feasibility reduces to whether a node can:

- acquire physiological signals with adequate sampling and quality,
- process and compress or classify data locally,
- transmit events, alarms, or streams within protocol-imposed windows,
- and uphold safety-critical behaviors (e.g., alarm delivery, actuator failsafe) under worst-case energy conditions [38], [41].

Regulatory guidance highlights cybersecurity, resilience, and lifecycle considerations for medical devices, including legacy systems [39], [40]. Many of these translate into additional energy demands—for secure communication, logging, or integrity checks—that must be budgeted alongside nominal sensing and communication.

The node-centric power and timing model allows us to integrate these requirements as additional tasks τ_k with their own energy and deadline parameters, and to ask whether the resulting schedule is sustainable under realistic $P_H(t)$, $E(t)$, and protection dynamics—rather than in an idealised, always-powered regime.

2.3.4 Environmental & agricultural sensing

Environmental and agricultural observatories are canonical long-lived deployments where energy autonomy is a primary design objective. Weather stations, soil-moisture nodes, and greenhouse monitors typically operate in remote or hard-to-service locations, often powered by small solar panels or indoor photovoltaics [4], [5], [42], [43]. Data are conveyed to backends over LPWANs such as Lora WAN, often under duty-cycle constraints and in the presence of interference [42], [43].

At the system level, these CPS support resource management, yield optimization, and environmental monitoring. At the microsystem level, feasibility again reduces to whether the node can maintain required sensing cadence and reporting freshness under highly variable $P_H(t)$, with sufficient reserve for maintenance operations such as key updates, firmware refreshes, and calibration [42], [43], [44].

Recent work on self-powered wireless sensor networks for sustainable agriculture emphasizes the coupling between deployment geometry, link budget, and energy budget: path loss and shadowing determine required transmit power and airtime; these, combined with sensing and processing costs, define the long-term energy demand that must be matched by harvested energy and storage [44].

The greenhouse and field deployments considered in later chapters are direct instantiations of this pattern: DREAM and the trace-driven sizing framework are used to validate that node-level neutrality, survivability, and availability can be guaranteed across a growing season under realistic light and temperature patterns.

2.3.5 Robotics & autonomous agents

Autonomous vehicles, UAVs, and mobile robots are cyber-physical systems in which perception, planning, and control are tightly coupled with physical motion and communication [63]. System-level studies focus on safety under corner cases, environmental shift, and complex interactions with infrastructure and other agents [45], [46], [47].

Within these systems, platform-level microsystems include perception nodes (e.g., camera or LiDAR processing modules), communication modules (e.g., 5G NR V2X OBU/RSU nodes), and embedded controllers for actuators. These nodes face stringent real-time and reliability requirements, but also increasing pressure to reduce size, weight, and power—especially in UAVs and small robots [45], [48], [49].

5G NR V2X and related technologies enable rich cooperative behaviors (platooning, cooperative perception, infrastructure-assisted driving), but they also introduce demanding communication and security workloads [64], [65]. Energy-aware path-planning and mission-planning schemes explicitly account for limited on-board energy stores and recharge opportunities, particularly for UAVs [48], [49].

In this context, the microsystem energy-timing model applies at two levels:

1. On-board nodes that must schedule compute- and communication-intensive tasks under limited local energy (e.g., in small UAVs or battery-constrained roadside units).
2. Infrastructure nodes (e.g., roadside V2X units) that may be partially or fully powered by harvested energy or constrained feeds yet must maintain high availability and predictable latency.

Our focus in this thesis remains on the second class—embedded nodes with constrained energy—but the same abstractions could be used to reason about energy-aware scheduling and autonomy margins in mobile platforms.

2.4 Summary and outlook

This chapter specialized general CPS notions to the microsystem scale, fixing the node as the unit of analysis. Starting from canonical CPS definitions, we introduced cyber-physical microsystems (μ CPS) as devices that couple sensing, computation, communication, and a lossy power subsystem to interact with their local environment under tight energy constraints. By treating plant and network as environment and focusing on the realization of the controller on an embedded platform, we obtained a node-centric view that makes embodiment and energetics explicit.

Building on this, we established a store-side energy model in terms of harvested power $P_H(t)$, load power $P_L(t)$, stored-energy state $E(t)$ (with proxy $V_{\text{store}}(t)$), and state-dependent conversion efficiencies $\eta_{\text{in}}(t)$, $\eta_{\text{out}}(t)$. Protection mechanisms such as UVLO, OVP, and current limits define a feasible energy envelope, within which the node can admit its workload. Over finite analysis windows, we expressed energy autonomy through neutrality, survivability, and availability, and linked these notions to the timing and structure of workloads: task sets $\{\tau_k\}$ with deadlines, bursty execution, and protocol-imposed communication windows.

We also clarified the telemetry contract between the physical node and any runtime or sizing logic. Rather than assuming rich metering, we deliberately constrained observability to realistic signals: coarse store-voltage traces, event-level telemetry (e.g., brownouts, UVLO trips), and, where available, lightweight self-metering. This sparse information is what later chapters will use for validation, adaptation, and trace-driven sizing, and it motivates an emphasis on robustness to uncertainty in $E(t)$.

Finally, representative CPS domains—power systems, industrial and building automation, medical and wearable devices, environmental and agricultural sensing, and robotics and autonomous agents—were revisited through this node-centric lens. In each case, system-level requirements such as latency, reliability, safety, or regulatory compliance manifest at the edge as constraints on node-level energetics and timing: required sensing and actuation cadences, communication obligations, and reserve energy for rare but critical events. These examples justify the focus on energy-autonomous microsystems and show that the adopted abstractions are broad enough to cover diverse application contexts.

The remainder of the thesis builds directly on this foundation. **Chapter 3** characterizes diffuse environmental energy Φ at the point of use, providing device-agnostic, statistically explicit envelopes that bound what is environmentally achievable over deployment windows. **Chapter 4** instantiates the loss-aware node model in concrete energy-collection architectures and components, making non-idealities and protection dynamics explicit and formalizing neutrality, survivability, and availability on the store side. **Chapter 5** introduces DREAM, an energy-aware runtime that regulates bursty workloads using minimal telemetry under this model. **Chapter 6** develops a trace-driven design methodology that maps mission-level autonomy targets into concrete choices of harvesting, storage, and thresholds. **Chapter 7** applies the framework to representative μ CPS case studies using measured and synthesized traces, and **Chapter 8** consolidates the results, positions them with respect to the state of the art, and outlines directions for future work.

3 Diffuse energy sources

Chapter 3 introduces diffuse environmental energy as the variable Φ , defined as the time-varying energy content available to a microsystem at its point of use. Here, Φ is treated as a stochastic field in time (and, where relevant, space) whose realizations capture both magnitude and intermittency of ambient sources such as indoor light, small thermal gradients, low-level vibrations, and opportunistic RF. Our focus is strictly on energy characteristics—not device behavior—so that Φ can serve as a neutral input to sizing, verification, and policy analyses independent of any particular transduction pathway.

We formalize Φ through descriptors that preserve their variability across operational horizons. Core measures include cumulative energy over a window $E_T = \int_t^{t+T} \Phi(\tau) d\tau$; intensity distributions and exceedance probabilities $P(\Phi > \phi)$; duty-weighted availability (fraction of time above a practical threshold); and temporal structure (autocorrelation, diurnal/weekly seasonality, calm-duration statistics, and rare-event tails). Because deployments often face strong context dependence, we also emphasize spatiotemporal heterogeneity (e.g., room-to-room or facade orientation effects), co-variates (temperature, occupancy, weather), and non-stationarity (long-term drifts). Throughout, units and normalization are chosen to reflect energy at the point of use (e.g., $\text{J} \cdot \text{m}^{-2}$ over a surface, J over a placement, or $\text{J} \cdot \text{m}^{-3}$ when volumetric flux is relevant), with clear aggregation windows that match operational duty cycles.




Figure 3.1 Taxonomy of diffuse environmental energy

To structure the discussion of Φ , we categorize diffuse environmental energy into mechanical, radiative, and thermal domains, each with context-dependent sub-classes whose temporal statistics govern availability and cumulative energy E_T . Figure 3.1 provides the taxonomy used throughout this chapter to report distributions, duty-time availability, calm durations, and other descriptors of Φ .

The remainder of the chapter provides empirically grounded envelopes for Φ in common settings (indoor/outdoor, static/mobile, human-occupied/unattended), specifies data-collection practices for constructing $\Phi(t)$ with minimal bias, and presents summary statistics (quantiles, reliability at specified percentiles, and worst-case calms) suitable for assessing energy availability and harvesting potential in a technology-agnostic manner. The result is a consistent, statistically explicit representation of environmental energy that can be consumed by subsequent methodology without presupposing any harvester characteristics.

These descriptors are chosen specifically so that the runtime and sizing methods in Chapters 5–7 can operate on storage-side energy envelopes and availability statistics, rather than device-specific harvester models.

3.1 Definition of energy flux “ Φ ”

Φ denotes the time-varying energy content available from diffuse environmental sources at the point of use. It is defined independently of any transduction mechanism or device characteristics. Unless otherwise specified, Φ is expressed as an energy flux with units appropriate to the domain of interest, e.g. $\text{J}\cdot\text{m}^{-2}\cdot\text{s}^{-1}$ for optical irradiance reaching a surface, $\text{J}\cdot\text{s}^{-1}$ for site-specific availability at a placement, or $\text{J}\cdot\text{m}^{-3}\cdot\text{s}^{-1}$ when volumetric fields are relevant. For a mission horizon T , the windowed environmental energy per unit area is

$$E_{\Phi}(T) = \int_{t_0}^{t_0+T} \Phi(t) dt \quad (3.1)$$

For periodic or cyclostationary processes, t_0 is chosen to align with a representative cycle; for non-stationary processes, $E_{\Phi}(T)$ is evaluated over shorter windows that tile the mission. These windowed energy envelopes are consumed directly in Chapter 6 for trace-driven sizing, where they define the environmental “budget” against which neutrality and survivability are judged.

3.1.1 Aggregation and normalization of Φ

In micro-systems, Φ is observed over centimeter-scale supports and short operational windows, so aggregation must preserve the fine temporal structure and the precise geometry of the measurement site. We define the cumulative micro-energy over a window T as $E_T(t) = \int_t^{t+T} \Phi(\tau) dt$, reported either per placement (J over T) or per area when a receiving surface is explicit ($\text{J}\cdot\text{cm}^{-2}$ over T). The associated window average $\bar{\Phi}_T(t) = E_T(t)/T$ is computed over windows aligned to the system’s duty rhythms—tens of seconds to minutes for task bursts, one day for diurnal structure, and multi-day horizons for resilience analysis. Minute-scale windows are recommended for indoor and human-activity-driven environments, where sub-hour variability is pronounced and coarse averaging inflates apparent availability; contemporary indoor-photovoltaic (IPV) and wearable-energy literature consistently documents this fine-scale fluctuation [50], [51], [52].

Normalization is stated with equal care. For radiative Φ , values are normalized to the projected area of the sensor plane and reported with the plane’s orientation and tilt; for vibration or motion-dominated fields, Φ is given per placement with the dominant axis of motion noted; and for RF, Φ we specify the monitored bands and report band-limited flux density so that comparisons are spectrally commensurate. Recent advances in IPV metrology emphasize that small changes in spectral content, geometry and measurement protocol can bias inferred energy content, which motivates explicit reporting of geometry, spectral window, and sampling settings alongside E_T and $\bar{\Phi}(t)$. In RF campaigns, city and home-scale studies show order-of-magnitude variation across bands and locations, reinforcing the need to publish the sweep configuration (RBW/VBW, dwell) used to estimate Φ . [51], [53], [54].

Data screening precedes aggregation to avoid bias in micro-campaigns. We remove specular spikes and shadowing gaps from indoor light records, enforce axis consistency and saturation checks in motion/vibration logs, and, for RF, ensure that repeated surveys use identical spectral settings so that differences in E_T reflect the environment rather than instrumentation. The IPV community has converged on protocols that control spectral invariance and temporal stability of the light source, reducing artefacts in the low-illuminance regime relevant to microsystems; adopting these controls when constructing Φ helps maintain comparability across sites and studies. [51], [55].

Because Φ often follows building schedules and human routines, we publish conditioned composites rather than a single unconditional average: diurnal profiles (hour-of-day means/quantiles) and schedule-conditioned statistics (e.g., business hours versus off-hours, occupied versus unoccupied). Wearable and kinetic-energy reviews report clustered activity with long calms, so we accompany E^T and $\bar{\Phi}(t)$ with calm-duration distributions and availability curves defined against a small context-specific threshold $\phi_{\min}$. This pairing—percentile envelopes of E_T over 24 h and distributions of calm lengths on the minute scale—adequately captures both the central tendency and the risk posed by low-energy intervals in micro-energy settings. [52], [56].

Finally, each dataset is accompanied by a succinct normalization record describing the window set TTT, sampling cadence and completeness rule, geometric conventions (area/volume/placement and orientation), and site metadata (height or body location, occlusions, nearby fixtures). Recent IPV inter-comparisons across metrology institutes illustrate how such metadata underpin reproducible cross-study synthesis at low power levels typical of microsystems [51], [55].

3.1.2 Availability, exceedance and calm durations

In micro-energy settings, availability expresses the fraction of time that the diffuse energy field $\Phi(t)$ remains above a practically meaningful threshold $\phi_{\min}$ within a window T . We take $\phi_{\min}$ to reflect usable energy at the point of use (e.g., a lux-to-irradiance level corresponding to non-negligible power density, or a band-limited RF flux floor), not any particular device threshold. Availability is then

$$A_{\phi_{\min}}(T) = \frac{1}{T} \text{meas} \{ \tau \in [t, t + T) : \Phi(\tau) \geq \phi_{\min} \} \quad (3.2)$$

reported alongside T and the sampling cadence. This metric is sensitive to the strong sub-hour structure typical of indoor light and human-activity-driven environments; minute-scale or finer sampling avoids biasing $A_{\phi_{\min}}(T)$ upward by smoothing short calms. Empirical indoor campaigns that map illuminance and spectra across offices, corridors and meeting rooms consistently show extended plateaus near zero interspersed with schedule- and occupancy-driven bursts—an archetype of availability that motivates percentile reporting rather than single averages [57], [58].

Because design decisions hinge on how often higher energy levels occur, we complement availability with exceedance probabilities $P(\Phi > \phi)$ and the quantile function $Q_p(\Phi)$. These capture rare but consequential periods, whether bursts or near-dark calms. Full-year indoor monitoring studies that log daily yields under realistic building use show exactly these tails: most days comfortably exceed modest energy needs, yet a small number of days each year fall short even for background overheads—evidence that low-percentile behavior, not the mean, governs survivability planning [57], [58].

For microsystems that experience long inactive stretches, the distribution of calm durations is an equally important descriptor. We define calms as maximal contiguous intervals with $\Phi(\tau) < \phi_{\min}$ and summarize their statistics via percentiles L_p^{calm} over the horizons of interest (e.g., minutes to hours). In body-worn and activity-proximate contexts, human-motion-linked sources exhibit clustered episodes with heavy-tailed gaps; reviews of wearable-context micro-energy repeatedly document this intermittency, reinforcing the need to publish calm-duration distributions alongside $A_{\phi_{\min}}(T)$ [52], [56].

RF-dominated environments display analogous behavior. Recent ambient electromagnetic wave surveys that measure power density across indoor, outdoor and basement spaces report order-of-magnitude swings with building use, attenuation and band selection; in such cases exceedance curves across bands (rather than a single aggregated number) more truthfully represent the resource available to microsystems. Complementary campus-scale measurements aimed at low-power applications reach similar conclusions and explicitly connect availability to feasible duty-cycled operation [59].

Finally, because $\Phi(t)$ is non-stationary and often exhibits “energy-on/energy-off” regimes, availability and exceedance naturally interface with intermittent-operation thinking: systems run opportunistically during on-periods and preserve state across off-periods. The intermittent-computing literature formalizes this reality and motivates reporting conservative, low-percentile envelopes of availability and calms so that algorithmic and sizing choices are anchored to the tails rather than the mean [60].

3.1.3 Temporal structure

The diffuse energy field $\Phi(t)$ encountered by microsystems is rarely white or stationary; it exhibits short-term correlation, multi-scale seasonality (daily, weekly, sometimes monthly), and clustered “on/off” regimes. We therefore characterize temporal structure with three

complementary views: (i) dependence at short lags, via the autocorrelation function and decorrelation time; (ii) periodic components, via seasonal–trend decomposition; and (iii) intermittency, via statistics of inter-event times and calm/burst clustering.

At short lags, we use the sample autocorrelation $\rho_\Phi(\tau)$ to estimate a decorrelation time τ_d beyond which $|\rho_\Phi(\tau)|$ falls below a small threshold (e.g., 0.1). This separates rapid fluctuations from slower drivers and guides the minimum sampling cadence needed to avoid aliasing in micro-energy records. Standard time-series texts formalize the link between $\rho_\Phi(\tau)$ and spectrum, allowing us to cross-check τ_d against observed power at diurnal or sub-diurnal frequencies [61]. In office interiors, light exposure often shows strong persistence during occupied periods with abrupt changes at schedule boundaries; post-occupancy light-dosimetry and illuminance studies confirm pronounced within-day structure and between-day variation shaped by human behavior and controls (Hartmeyer et al., 2024; Jang et al., 2024). Similarly, structural vibrations in buildings display persistent bands around natural modes excited by traffic or rail transients, with stability punctuated by timetable-driven events [62], [63].

To extract periodic structure, we decompose $\Phi(t)$ into seasonal, trend, and remainder components using robust STL-type procedures. For micro-energy, single-season STL is often insufficient because both daily and weekly cycles are present; modern multi-seasonal extensions (e.g., MSTL) and STL hybrids provide stable separation of these cycles even under volatility [64], [65]. The seasonal component $S(t)$ captures regular patterns (workday lighting schedules, commuter-driven vibration epochs), the trend $C(t)$ captures slow drift (e.g., daylight season), and the remainder $\varepsilon(t)$ aggregates idiosyncratic variability. Indoor environmental measurements in living labs demonstrate that even adjacent rooms can exhibit weak cross-correlation in illuminance and occupancy, underscoring the need to model seasonal structure per placement rather than at building level [66].

Beyond smooth seasonality, intermittency is a defining property at micro scales. We model event times—entries into an “energy-above-threshold” state—as a point process and analyze inter-event time distributions. Bursty temporal networks commonly exhibit fat-tailed inter-event times and residual correlation even after de-seasoning; these behaviors are well documented and provide principled metrics for quantifying clustering [67], [68], [69]. For Φ , this implies that high-energy episodes (e.g., meetings with lights at full output, trains passing beneath an over-track structure, or a Wi-Fi band becoming busy) tend to arrive in

trains separated by heavy-tailed calms—precisely the pattern observed in ambient vibration campaigns and indoor lighting usage studies [62], [70].

Ambient RF environments also vary across bands and time. Longitudinal multi-country assessments report statistically significant trends and intra-day variation in RF-EMF levels across microenvironments; within-band exceedance curves change with traffic and network configuration, reinforcing the need to analyze $\Phi(t)$ with band-resolved seasonality and burst metrics [71], [72], [73]. Where high-resolution time series are available, decorrelation times in the order of minutes are typical indoors, with extended tails during night-time or off-hours.

In reporting practice, we therefore pair (a) decorrelation time τ_d and spectral summaries to set defensible sampling cadences; (b) multi-seasonal decomposition to expose diurnal/weekly structure; and (c) burst metrics—inter-event distributions, burst-train counts, and calm-duration percentiles—to quantify intermittency. These descriptors remain deliberately technology-agnostic: they characterize the temporal scaffolding of Φ that any micro-scale methodology must respect.

3.1.4 Spatiotemporal heterogeneity and co-variates

At micro scale, $\Phi(r,t)$ is strongly conditioned by where a node sits and what the surrounding environment is doing. Even within a single room, illuminance and spectral composition change markedly with distance to fenestration, shading state, and the observer’s orientation; controlled field studies show that local viewing direction and window proximity significantly alter the incident light field, which explains the large between-placement spread routinely observed in office interiors [74], [75]. In open-plan offices the gradient from the facade to core zones persists even under modern shading or light-shelf strategies, and small differences in layout or luminous ceiling design can shift the entire distribution of indoor optical Φ at the desk scale [75], [76]. These spatial effects are not merely architectural artefacts; they induce order-of-magnitude differences in cumulative micro-energy over daily windows, so reporting Φ without a precise geometric description of placement risks non-reproducible comparisons.

Mechanical environments display analogous heterogeneity. Ambient vibration campaigns reveal location-dependent spectral signatures shaped by structural modes, transmission paths, and nearby excitations; floors above rail tunnels, for example, exhibit stable bands with episodic bursts synchronized to timetable events, while nominally similar residential

shells can show very different baseline levels once external noise sources and coupling are accounted for [62], [77]. Consequently, the same building can yield distinct Φ envelopes across rooms and storeys, and placement metadata (span, bay, distance to major masses or supports) becomes as critical for micro-energy characterization as orientation is for indoor light.

Band-limited electromagnetic fields also vary sharply across small distances and usage contexts. Multi-site surveys that stratify measurements by microenvironment—offices, transport hubs, homes, and outdoor streets—report strong spatial gradients and temporal modulation driven by user traffic and network configuration; summarizing RF Φ by band-resolved exceedance curves is therefore more faithful than a single aggregate level [78], [79]. From the perspective of micro-systems, this implies that a shift of only a few meters—or even a change from line-of-sight to partial shadow—can re-rank RF bands in terms of usable energy content over the same window T.

Heterogeneity is intertwined with co-variates that imprint structure on Φ over hours to seasons. In office and educational settings, lighting schedules, occupancy patterns, and shading control policies govern diurnal and weekly cycles, while exterior weather modulates the amplitude of the indoor optical field that reaches core areas; reviews of daylighting and window effects consistently link these drivers to measurable differences in local light exposure and its variability [80], [81]. For mechanical contexts, recurring human activities (footfall, doors, elevators) and transport timetables act as deterministic co-variates that increase predictability at specific frequencies; for RF, network load and technology mix (e.g., Wi-Fi vs. cellular bands) operate as time-of-day covariates that reshape exceedance behavior. Where buildings exploit active daylight harvesting, control set-points and sensor placement further modulate the residual indoor optical Φ , and recent analytical work highlights how these control choices propagate to the energy potentially accessible at micro scale [82].

In practice, characterizing $\Phi(r,t)$ for micro-energy therefore requires a site-explicit protocol: each record should pair the time series with placement descriptors (coordinates or a reproducible sketch, height, orientation, occlusions), note salient co-variates (occupancy schedule, shading state, HVAC or network schedules), and, where feasible, provide short co-located measurements at adjacent placements to quantify the spatial gradient. This

approach recognizes that variability across centimeters to meters is not noise but rather an intrinsic property of diffuse environmental energy at the scales relevant to microsystems.

3.1.5 Percentile envelopes and reliability

For micro-energy analysis, percentile envelopes condense uncertainty in $\Phi(t)$ and its windowed integral E_T . We report $E_p = Q_p(E_T)$ and $\bar{F}_p(T) = Q_p(\bar{\Phi}_T)$ over horizons T that match operation (minutes–hours for tasks; 24 h and 7 d for resilience). In indoor contexts the distribution is typically skewed, with a small fraction of days delivering most of the energy and a non-negligible tail of low-yield periods; envelopes at conservative percentiles (e.g., 5th–10th) therefore provide a more faithful basis for design than means or medians alone. Recent full-year monitoring of indoor energy conditions for IoT explicitly demonstrates this asymmetry, motivating percentile-based reporting rather than single “typical day” summaries [83].

We pair envelopes with a θ -reliable availability metric,

$$R_\theta(T) = \sup_{r \in [0,1]} P(A_{\phi_{\min}}(T) \geq r) \geq \theta \quad (3.3)$$

so, statements such as “ $R_{0.9}(24 \text{ h})=0.35$ ” can be read as: on at least 90% of days, Φ exceeds $\phi_{\min}$ for 35% of the time window. For band-structured fields (notably RF), reliability should be reported per band and then aggregated, because tails differ across spectra and micro-environments; recent surveys emphasize order-of-magnitude differences in available power density and duty over common indoor bands, which a single aggregate would hide [84], [85].

To make reliability risk-aware, we also publish tail-sensitive scalars alongside $E_p(T)$, adopting energy analogues of Value-at-Risk and Conditional-Value-at-Risk: $Var_\theta(E_T) = E_{1-\theta}(T)$ and $CVaR_\theta(E_T) = E[E_T | E_T \leq Var_\theta]$. Contemporary planning literature argues for CVaR-type metrics precisely because they penalize rare but severe shortfalls more than percentile cut-offs, a consideration that directly aligns with micro-systems that must survive infrequent, low-energy stretches [86], [87].

Finally, envelopes and reliability should be updated as the resource frontier shifts. Rapid progress in indoor photovoltaics—both device-level efficiency gains under 100–1000 lux and emerging metrology consensus for reproducible indoor testing—changes realistic expectations for usable Φ and thus the thresholds $\phi_{\min}$ used in availability definitions.

Publishing the chosen $\phi_{\min}$ with a cross-reference to current indoor-PV performance contexts preserves comparability across studies and time [88], [89], [90].

3.1.6 Uncertainty and estimation

Quantifying the uncertainty of Φ and its windowed integrals E_T is essential at micro scale, where centimeter-level placement, minute-level sampling, and local schedules dominate variability. Our approach is model-agnostic and emphasizes: dependence-aware resampling for distributional summaries; transparent propagation of measurement uncertainty; and reproducible handling of gaps and anomalies—framed for indoor optical, vibration/motion, and low-grade thermal contexts.

For time-dependent records, resampling must respect serial correlation and seasonality. We use block bootstrap variants (moving, circular, tapered) with block lengths chosen no shorter than the empirically estimated decorrelation time, yielding valid confidence intervals for $Qp(E_T)$, $A_{\phi_{\min}}(T)$ and calm statistics without assuming a parametric model [91]. When daily or weekly structure is pronounced, a seasonal/block hybrid (blocks aligned to diurnal/weekly cycles) preserves phase relationships and improves low-percentile estimates that drive survivability assessments [92]. In practice, we report the block selection rule, sensitivity to block length, and the resulting interval estimates alongside point summaries.

When Φ is derived from calibrated sensors, we propagate measurement uncertainty using Monte-Carlo techniques consistent with contemporary metrology guidance. Components include calibration and reference traceability, drift, resolution/quantization, alignment and cosine errors (for optical), and axis/fixture effects (for vibration). Combined standard uncertainty is then mapped through the normalization used for Φ (per placement or per area) to obtain uncertainty on E_T and availability. Recent indoor-photovoltaic metrology highlights the importance of spectral mismatch, source stability and spatial non-uniformity at low illuminance; adopting those reporting conventions materially improves uncertainty statements for indoor optical Φ [55]. For vibration/motion fields at building or wearable scale, uncertainty budgets should document sensor alignment, bandwidth and anti-aliasing, mounting repeatability, and temperature dependence, reflecting best practice in structural monitoring and kinetic-energy studies [56].

Data quality control precedes estimation. For indoor optical series we flag and, where appropriate, remove specular spikes and transient shadows; for vibration/motion we enforce

axis consistency, clipping checks and stationarity tests per window; for thermal gradients we track probe placement, contact quality and lag corrections. Window-level completeness thresholds (e.g., $\geq 90\%$ valid samples) are declared a priori; windows failing completeness are withheld from E_T and availability calculations. Missing values are handled with time-series-aware methods—state-space smoothing or LOESS/ARIMA-based imputation—paired with flagging so that uncertainty can be widened where imputation occurs [93], [94]. We avoid naive interpolation across long calms, which would bias low-percentile envelopes upward.

Finally, all estimates are accompanied by reproducible provenance: calibration certificates or references, QC rules and parameters, resampling configuration (block type and length), and scripts/notebooks that regenerate intervals and figures. This practice ensures that distributional summaries of Φ —especially the tails that govern risk—remain auditable as measurement campaigns evolve.

3.2 Characteristics and energy content of source classes Φ

We adopt a statistics-first view of $\Phi(t)$ to preserve intermittency and tail risk. For each source class we concisely report: (i) physical characteristics; (ii) energy content over mission windows T ; (iii) availability and calm distributions relative to a source-side threshold $\phi_{\min}$; and (iv) the harvesting potential as a conversion-agnostic upper bound.

3.2.1 Indoor and outdoor optical Φ

Throughout this thesis, $\Phi(t)$ denotes the plane-of-array (POA) irradiance incident at the exact plane and location where a harvester would sit, defined in a device-agnostic manner. Let $\mathbf{n}$ be the unit normal of that plane (a desk plane indoors; a module tilt/azimuth outdoors) and $E_{e,\lambda}(\lambda, t, \mathbf{n})$ the spectral irradiance [$\text{W} \cdot \text{m}^{-2} \text{nm}^{-1}$] measured at that plane. The source-side field is the broadband integral

$$\Phi(t) = \int_{\Lambda} E_{e,\lambda}(\lambda, t, \mathbf{n}) d\lambda [\text{Wm}^{-2}] \quad (3.4)$$

with Λ the application band (PV contexts typically $\approx 300\text{--}1200$ nm). This definition contains no device weighting (no spectral response, IV curve, thresholds, or MPPT behavior); those appear later in the conversion model of Chapter 4. For photon-centric arguments we may

also use the co-planar photon-flux density $\Psi(t) = \int_{\lambda} E_{e,\lambda}(\lambda, t, n)/(hc/\lambda)d\lambda$, but the object of analysis in Chapter 3 remains the broadband $\Phi(t)$.

Indoors, $\Phi(t)$ is the superposition of daylight transmitted through the fenestration system and electric-lighting contributions as they actually arrive at the plane of use. The field exhibits schedule coupling (occupancy, switching, stepped or continuous dimming) and pronounced centimeter-scale heterogeneity caused by luminaire photometry, furnishings, and facade-to-core gradients. Because geometry and context strongly condition outcomes, every statistic reported for $\Phi_{in}(t)$ is tied to an explicit plane (orientation/tilt), position, and any persistent occlude. Typical workplace illuminance targets (e.g., task-plane values around a few hundred lux, depending on task and standard) explain the quasi-plateaus commonly observed during occupied hours and the extended calms out of hours; these values are context, not device assumptions, and serve only to interpret the observed plateaus.

Outdoors, $\Phi(t)$ is governed by solar geometry and meteorology and is properly expressed as the sum of beam, diffuse, and ground-reflected components on the plane of array,

$$\Phi_{out}(t) = B_{POA}(t) + D_{POA}(t) + R_{POA}(t) \quad (3.5)$$

each depending on sun position, array orientation (tilt/azimuth or tracking), sky condition, and surface albedo; near- and far-shading modify these components as environmental effects at the plane, not as “device losses” in this chapter. In practice, POA is either measured with a co-planar pyranometer or modelled by transposition from horizontal or normal components using sky-diffuse formulations. A recent continuous form of the Perez diffuse sky model provides a modern, differentiable transposition that covers both forward and reverse cases, making it a suitable bibliographical anchor when we refer to model-based POA construction for $\Phi_{out}(t)$ [95].

For spectral context in the outdoor case, it is standard to reference the terrestrial AM1.5 spectrum when discussing device characterization; however, analyses in this chapter always use measured (or site-modelled) field spectra/irradiance at the actual plane. The broader extraterrestrial irradiance baseline itself has been updated in the satellite era; the SORCE/TIM results consolidated a lower total solar irradiance (TSI) value than pre-2000 conventions, which is relevant when connecting POA models to top-of-atmosphere

constraints and helps explain differences among legacy datasets. A concise, post-2010 anchor for that revision is Kopp and Lean (2011) [96].

Energy content is summarized consistently across settings using the windowed energy $E_T = \sum_{t \in T_T} \Phi(t) \Delta$ over mission-relevant horizons T . Because both indoor and outdoor fields are intermittent but in different ways—schedule-coupled plateaus and step changes indoors; diurnal cycles with sub-minute cloud-edge ramps and heavier upper tails outdoors—we report percentile envelopes $Q_{5,50,95}(E_T)$ rather than means, and we accompany them with availability $A_{\phi_{\min}}$ and calm-duration statistics relative to a source-side threshold $\phi_{\min}$. Evaluating $A_{\phi_{\min}}$ at one-minute cadence captures indoor control transitions and paired with minute-scale windows, resolves the short-timescale outdoor variability of interest for buffer sizing and control responsiveness. When synthetic or scenario-based $\Phi(t)$ traces are needed later (e.g., to stress-test scheduling sensitivity at hourly resolution before deploying minute-scale data), weather-conditioned stochastic models for solar radiation can be used as a supplement to measurement; such models have been formulated to reproduce hourly dynamics under varying weather regimes and are appropriate to reference in Chapter 6’s methodology section, not as a replacement for POA measurements in Chapter 3 [97].

This device-agnostic, geometry-aware definition of $\Phi(t)$ ensures that the brief envelopes reported here can be consumed directly by the loss-aware conversion model in Chapter 4 and, subsequently, by the DREAM scheduling and sizing analyses in Chapters 5–6. Outdoors, it aligns with contemporary POA measurement/modelling practice and sky-diffuse transposition; indoors, it preserves the centimeter-scale heterogeneity and schedule dependence that dominate feasibility at low power. By maintaining a single, consistent environmental quantity across both contexts, we avoid double-counting device effects and preserve a natural flow from source-side characterization to conversion and autonomy [95].

3.2.2 Low-grade thermal Φ

For low-grade thermal sources we keep the same source-side stance: $\Phi(t)$ denotes the thermal field at the point of use, independent of any converter. Let $\Delta T(t) = T_{\text{hot}}(t) - T_{\text{amb}}(t)$ be the temperature difference across the intended interface and let $h_{\text{eff}}(t)$ be the effective heat-transfer coefficient (lumped conduction–convection coupling through the mounting, clothing/insulation, and boundary layers). We define the thermal field as the heat-flux density incident on the interface,

$$\Phi_{th}(t) = h_{eff}(t)\Delta T(t)[Wm^{-2}] \quad (3.6)$$

and, when convenient, the windowed thermal energy per unit area over horizon T as

$$E_T^{(th)} = \sum_{t \in T_T} \Phi_{th}(t)\Delta \quad (3.7)$$

This definition is geometry-aware (through the actual contact and surroundings), device-agnostic (no Seebeck coefficients, internal resistances, or start-up thresholds), and aligns with the recommendation from recent reviews to characterize wearable and built-environment opportunities via ΔT and effective coupling rather than by device outputs [98], [99], [100].

In wearable/body-proximate contexts $\Delta T(t)$ is driven by metabolic heat and ambient conditions, while $h_{eff}(t)$ is governed by contact pressure, interface materials (e.g., textile, elastomer pads), local airflow, and posture. The dynamics are slow—minute-to-hour time constants—because tissue and packaging add thermal capacitance, producing quasi-steady plateaus during sustained wear interspersed with drift under activity or environmental change. Human-subject and mannequin studies consistently find modest but persistent driving ΔT (typically a few kelvin in temperate indoor conditions), with large variability attributable to micro-fit and insulation; hence placement repeatability dominates uncertainty [101], [102]. In equipment/building-interface contexts (ducts, motor enclosures, piping, electronics housings), ΔT can be larger and steadier, while h_{eff} depends on mounting geometry, insulation, and forced vs. natural convection; here, run-time duty cycles and control regimes (e.g., HVAC schedules) imprint a clear temporal structure on $\Phi_{th}(t)$ [98], [99].

Because $\Phi_{th}(t)$ is low-frequency and intermittency is shaped by usage cycles rather than fast transients, minute- and multi-hour windows are the natural scales for energy summarization. We therefore report $Q_{5,50,95}(E_T^{(th)})$ for $T \in \{1 \text{ h}, 4 \text{ h}, 24 \text{ h}\}$ to capture session-scale feasibility and daily survivability, and we accompany these with availability $A_{\phi_{min}}(T)$ and calm-duration statistics defined relative to a source-side threshold ϕ_{min} (e.g., a minimum heat-flux density deemed “usable” at the interface). Using percentiles rather than means is crucial, as low-percentile days and long calms—caused by cool ambient conditions, poor contact, or off-cycles—govern autonomy for ultra-low-power systems [98], [100].

Measurement and normalization remain strictly source-side. Differential sensors (thin-film thermopiles, paired thermistors) mounted across the interface record $\Delta T(t)$; co-recorded context (activity class, clothing/insulation, airflow or fan state) explains much of the variance. Estimating h_{eff} can be done by brief calorimetric steps or literature-based priors for comparable mountings; reporting both $\Delta T(t)$ and the assumed h_{eff} keeps $\Phi_{\text{th}}(t)$ reproducible across sites. Cadences of 10–60s are sufficient for dynamics without injecting noise; dependence-aware uncertainty (block bootstrap with blocks longer than the thermal time constant) should accompany Q_p and availability figures. Device-specific topics—Seebeck efficiency, electrical matching, cold-start energy and boost-converter thresholds—are deferred to Chapter 4, where the same $\Phi_{\text{th}}(t)$ drives a loss-aware conversion chain [99], [102].

Finally, design implications follow directly from the field statistics. In wearables, autonomy is set by the low-percentile $E_T^{(th)}$ over multi-hour windows and by the distribution of calm durations during low-activity or cool-ambient periods; micro-siting (body location) and contact engineering often yield larger gains than nominal increases in converter efficiency [101], [102]. In equipment/building interfaces, steadier ΔT supports predictable $E_T^{(th)}$, but maintenance and control cycles create structured calms that must be buffered.

3.2.3 Vibrations Φ

We retain a source-side stance. For mechanical environments, the field $\Phi(t)$ at the mounting point is most usefully represented by the vector acceleration time series $\mathbf{a}(t)=[a_x(t), a_y(t), a_z(t)]$ [ms^{-2}] together with its band-resolved spectral density. Device behavior (tuning, bandwidth, thresholds) is deferred to Chapter 4; here we characterize the environment itself.

Let $S_a(f; t)$ denote the one-sided power spectral density (PSD) of the normal-to-mount acceleration over short analysis windows anchored at time t . Because harvestable mechanical energy for resonant and broadband devices is band-dependent, we will summarize the field via band-integrated vibration content

$$\Phi_B(t) \equiv \int_{f \in B} S_a(f; t) df [(m/s^2)^2 / \text{Hz} \cdot \text{Hz} = (m/s^2)^2] \quad (3.8)$$

for bands B centered on the dominant excitations (e.g., footfall 1–4 Hz; building services 10–80 Hz; vehicle chassis/engine bands 5–200 Hz). For time-domain summaries over a horizon T , we use windowed, band-specific integrals

$$E_T^{(B)} = \sum_{t \in T_T} \Phi_B(t) \Delta \quad (3.9)$$

and report percentile envelopes $Q_{5,50,95}(E_T^{(B)})$. This keeps the description device-agnostic (no assumed resonant frequency or damping), yet actionable for Chapter 4, where a transducer's effective band will be mapped onto B.

In building interiors, ambient vibration is typically narrow-band and location-specific, shaped by structural modes, services (HVAC, pumps), and human activity (footfall, elevators). Spectra show sharp lines or clusters at modal frequencies with sidebands from machinery, while the time series is bursty—clusters of events separated by long calms. Footfall-induced responses on floors and walkways concentrate energy between ~ 1 –4 Hz with higher harmonics, whereas machinery bays show peaks an order of magnitude higher [103]. These properties imply heavy-tailed event distributions and strong diurnal/weekly structure in $E_T(B)$ [104].

On vehicles, ride-induced chassis vibration and engine/road excitations populate low-to-mid bands (few Hz up to a few hundred Hz) whose centers and amplitudes vary with speed, surface roughness and powertrain order. Standard road roughness spectra (e.g., ISO 8608) explain the broad-band content transmitted to sprung masses; engine orders add narrow peaks whose frequency scales with RPM. Temporal patterns are sessional (trips/commutes), with relatively few calms during motion but extended calms between journeys [105], [106].

Walking, running, limb swing and incidental impacts generate clustered, low-frequency excitations with large peak-to-rms ratios; the dominant fundamental for gait is typically ~ 2 Hz, with substantial energy in the first few harmonics and sporadic higher-frequency bursts from impacts. Crucially, activity occurs in sessions, so availability is high during activity and near-zero between sessions [52], [107], [108].

Because useful content is band-concentrated, a single wideband scalar can be misleading. We therefore report, per context and placement, the set $\{Q_{5,50,95}(E_T^{(B)})\}_B$ over mission windows T (e.g., 1–5 min to capture bursts; 24 h for daily survivability), along with availability $A_{\varphi_{min}}^{(B)}(T)$ computed for a band-specific threshold $\varphi_{min}^{(B)}$ and calm-duration distributions. The low-percentile envelopes govern feasibility for autonomy at low duty

cycles, since heavy-tailed calms (nights, off-hours, between trips) dominate risk even when median activity is favorable [104], [109].

Mount tri-axial accelerometers at the intended harvester location and orientation; record axis alignment, mounting stiffness, nearby machinery, floor span/bay, and for vehicles the seat/frame position and typical speeds. Choose a cadence that resolves the highest band of interest with margin (sampling intervals 10–50 ms are typical; faster for powertrain orders). Compute PSDs on sliding windows with stated taper and segment length; report center frequencies and -3 dB bandwidths of dominant peaks to aid later matching. For human-motion measurements, note activity labels (walk, run, idle) and session timing. All statistics should include dependence-aware uncertainty (e.g., block bootstrap with blocks $\geq$ the longest correlation time in the band) to avoid overconfidence from strongly correlated windows [104], [109].

Since most energy arrives in bursts near specific bands, feasibility hinges on buffering and on aligning a device’s effective band with the environmental peaks. Micro-relocation—moving a few tens of centimeters on a floor slab or changing from frame to seat-rail in a vehicle—can change $Q_5(E_T^{(B)})$ by orders of magnitude due to modal shapes and transmission paths. Consequently, a short pre-deployment spectral survey across candidate placements is often more impactful than marginal improvements in converter efficiency [103], [105].

3.2.4 RF Φ

We define the RF field at the point of use as band-limited power-flux density: for a band $b=[f_1, f_2]$,

$$\Phi_b(t) = \int_{f_1}^{f_2} S(f, t, r) df [Wm^{-2}] \quad (3.10)$$

where $S(f, t, r)$ is the spectral power-flux density incident on the local plane. This is strictly source-side (no rectifier thresholds or matching). In practice, you either log $\Phi_b(t)$ with a calibrated antenna + receiver (reporting antenna factor, orientation, detector time constants) or infer band power at the receiver port and convert to flux with the stated aperture/gain. Minute-scale windows capture traffic bursts; 24-hour windows capture diurnal occupancy effects, so we report $Q_{5,50,95}(E_T^{(b)})$ and availability $A_{\varphi_{min}^{(\beta)}}(T)$ per band rather than a single

“wideband” number. This avoids masking tails, which differ markedly across Wi-Fi, cellular, and broadcast bands[78].

Three empirical facts shape design. First, levels vary by orders of magnitude across bands and micro-sites; moving the antenna tens of centimeters or rotating it can re-rank bands in the same room due to line-of-sight and multipath. Second, urban surveys consistently find ambient powers that are often sub- μ W at the antenna port but occasionally far higher near infrastructure—hence the need for per-band, placement-specific envelopes rather than generic claims. Third, traffic and scheduling imprint minute-scale bursts and clear daily structure, so low-percentile days and long calms often govern autonomy even when medians are attractive. These points are well documented in ambient-RF field studies and recent surveys [78], [110], [111].

Concise reading of the literature. Piñuela, Mitcheson and Lucyszyn provide a city-scale measurement study of ambient RF levels and harvesting practicality; Sherazi et al. give a comprehensive survey of RF energy harvesting models, measured powers, and applications; Sharma et al. review techniques for low-power WSNs and emphasize per-band variability and duty-cycle effects—exactly the rationale for our band-resolved envelopes.

Table 3.1 Source summary of $\Phi(t)$ by class

Source class	$\Phi(t)$	Typical magnitude (order)	Dynamics/ cadence	Availability signature	Windows for E_T
Optical (indoor)	Plane-of-array (POA) irradiance at device plane (daylight + electric light)	$\sim 0.1\text{--}10 \text{ W}\cdot\text{m}^{-2}$ in offices	Log at 1–5 s; report at 1-min	Moderate in hours, low overnight; micro-siting and shading dominate	1 min (bursts), 24 h (daily)
Optical (outdoor)	POA irradiance = beam + diffuse + ground on module plane	Up to $\sim 1000 \text{ W}\cdot\text{m}^{-2}$ clear sky; fast sub-min ramps under broken cloud	Log at 1–10 s; report at 1-min	Strong diurnal; short cloud-edge deficits/surpluses	1 min (ramps), 24 h (daily)
Thermal (low-grade ΔT)	Heat-flux density $\Phi_{th} = h_{eff}\Delta T$ at interface	Wearables $\sim 1\text{--}20 \text{ W}\cdot\text{m}^{-2}$; equipment interfaces often higher/steadier	Log at 10–60 s	Slow drifts; sessions & HVAC cycles imprint structure	1 h, 4 h, 24 h
Vibrations	Band-integrated vibration content $\Phi_B = \int_B S_a(f)df$	Buildings: $\sim 0.01\text{--}0.3 \text{ gRMS}$ in modal bands; vehicles/human motion: $\sim 0.1\text{--}>1 \text{ g}$ bursts	Log at 10–50 ms; report at 1-min	Bursty/sessional; heavy-tailed calms between events	1–5 min (bursts), 24 h
RF	Band power - flux density $\Phi_b = \int_{f_1}^{f_2} S(f)df$	Often sub- μW at antenna port; can be $\gg 100 \mu\text{W}$ near infrastructures	Log at 1 s; report at 1-min	Band- and site-specific; minute-scale traffic bursts; clear diurnals	1 min (bursts), 24 h

3.3 Availability

Availability quantifies how often the environmental field $\Phi(t)$ at the point of use exceeds a contextually meaningful source-side threshold $\phi_{\min}$. For a window T sampled every Δ with index set T_T , we compute

$$A_{\phi_{\min}}(T) = \frac{1}{|T_T|} \sum_{t \in T_T} \Phi(t) \geq \phi_{\min} \quad (3.11)$$

Unlike yield or autonomy (treated later), this chapter keeps $A_{\phi_{\min}}$ strictly device-agnostic: the threshold is expressed in the native units of $\Phi(t)$ on the same plane of use, and no conversion, thresholds, or control losses are folded in.

To keep $A_{\phi_{\min}}$ comparable across deployments, three items must be declared alongside the number: the geometry/band definition for $\Phi(t)$ (e.g., plane-of-array for optical; band limits for RF; analysis band for vibration; interface description for thermal), the evaluation cadence used to form T_T , and the horizon T . A one-minute cadence is a good default because it registers indoor lighting steps and outdoor cloud-edge fluctuations while smoothing second-scale burstiness in RF and mechanical logs that are originally captured faster; the raw data may—and often should—be recorded at higher rates, then aggregated to the common minute grid for the availability calculation [112].

Choosing $\phi_{\min}$ is an environmental decision, not a device setting. A defensible approach is to anchor $\phi_{\min}$ to a contextual percentile or use-case requirement that can be stated without invoking a converter—for example, a POA irradiance level associated with “usable light” in a space, a heat-flux density representing a meaningful body/equipment differential, a vibration band power integrated over a declared band, or an RF band power-flux over a licensed/unlicensed band. The chosen value must be reproducible from measurements alone (no IV curves, no rectifier sensitivities) and should remain fixed when comparing placements or days. When a study also reports device results in Chapter 4, that chapter may motivate a separate, device-aware threshold; the source-side $\phi_{\min}$ reported here is not retrofitted to match it.

Availability is estimated over at least two horizons. A short window (on the order of minutes) answers whether bursts can be supported without long gaps; a daily window ($T=24$) answers whether routine operation can be sustained through typical cycles. These two horizons make the statistic informative without duplicating the broader envelope work in §3.2. Uncertainty

should be attached to each estimate using dependence-aware resampling—moving-block bootstrap or related techniques with blocks longer than the dominant autocorrelation time—combined in quadrature with measurement components such as calibration and placement repeatability [113]. The result is an interval that reflects both sampling variability and measurement fidelity, rather than an over-confident single number.

Interpretation differs from mean-energy summaries. Two deployments can share similar E_T yet have very different $A_{\phi_{\min}}$: one with short, frequent excursions above $\phi_{\min}$, another with long calm spells punctuated by rare peaks. The former favors opportunistic policies; the latter demands buffering and deferral. This is precisely why availability is reported on a fixed cadence and horizon rather than inferred from averages or percentiles of E_T . In outdoor optical settings, for instance, sub-minute cloud dynamics can depress minute-scale availability despite ample daily energy [112]. In RF contexts, traffic scheduling and occupancy drive distinct minute-to-hour patterns that a daily mean would obscure [78]. In vibration environments, sessional activity creates bands of near-continuous availability interleaved with long calms; a band-resolved $A_{\phi_{\min}}$ captures this directly [104].

The deliverable of this section is a compact record: for each context and placement, the pair of availability values $A_{\phi_{\min}}(T)$ at minute-scale and at 24 h, each accompanied by its threshold definition, geometry/band metadata, cadence, and an uncertainty interval. These numbers—by design—are the only inputs Chapter 4 needs to impose device-side losses and thresholds, and they are the statistics that Chapters 5–6 use within DREAM to sequence work and size buffers against empirically observed gaps.

3.4 Harvesting potential

Harvesting potential is the environmental ceiling: the most energy that any harvester could plausibly extract on a given plane and over a stated time window, before device particulars enter the picture. We obtain it directly from the measured field $\Phi(t)$ by integrating over a horizon T ,

$$E_T = \sum_{t \in T_T} \Phi(t) \Delta \quad (3.12)$$

and then describing the distribution of E_T across comparable windows. Interpreted this way, E_T is not a promise of yield; it is a clean, device-agnostic bound that later chapters will reduce through spectral weighting, thresholds, tracking overheads and storage losses.

Because diffuse sources are intermittent and skewed, a single average obscures the design reality. We therefore summarize harvesting potential with percentiles of E_T —conservative, typical and favorable values—computed on two windows that matter for decision-making: a minute-scale window for burst feasibility and a 24-hour window for daily survivability. Availability from §3.3 complements this view. Availability asks how often the field clears a declared environmental level; harvesting potential asks how much energy accumulates when we look over the whole window. Two sites can share a similar median E_T and yet have very different availability if one suffers long calms and the other offers frequent short bursts; sizing and scheduling depend on knowing both.

The definition is uniform across sources, which keeps the chapter coherent. For light, indoors and outdoors alike, $\Phi(t)$ is the plane-of-array irradiance at the relevant surface. Outdoors, this naturally includes beam, diffuse and ground-reflected components and can be obtained either by co-planar measurement or by validated transposition from horizontal/normal components; using a modern Perez-family formulation maintains continuity across clear and overcast conditions without committing to any specific device model [95], [112]. For low-grade thermal contexts, $\Phi(t)$ is the incident heat-flux density $h_{\text{eff}}\Delta T$ at the interface, whose slow dynamics make multi-hour windows more informative; recent reviews emphasize reporting the environmental driver separately from thermoelectric conversion details, which we follow here [99]. For mechanical and RF environments, energy is intrinsically band-conditioned: we compute E_T on declared vibration or RF bands, preserving the narrow-band or channel structure that later determines which devices can effectively couple to the field [78], [104].

Uncertainty accompanies every percentile we report. It comes from two places: measurement (calibration, geometry and placement repeatability) and temporal dependence (adjacent samples are not independent). We therefore combine instrument uncertainties with dependence-aware resampling—for example, a moving-block bootstrap with blocks longer than the dominant correlation time—so the intervals reflect real variability rather than optimistic i.i.d. assumptions [113]. In outdoor optical cases, retaining minute-scale structure captures cloud-edge ramps that disproportionately shape the right tail of E_T ; in RF and mechanical cases, minute-scale aggregation smooths sub-second bursts while preserving the patterns that matter for buffering and task deferral.

The outcome of §3.4 is intentionally compact: for each plane and context we provide minute- and day-scale percentile summaries of E_T , with explicit geometry/band definitions and uncertainty. Chapter 4 then explains the gap between this environmental ceiling and deliverable energy by applying spectral response, start-up and threshold effects, maximum-power-point tracking, impedance and storage losses. Chapters 5–6 rely on both pieces together: harvesting potential sets what is even conceivable in a window, while availability and conversion losses determine how DREAM schedules work and sizes buffers to meet autonomy targets.

The brief envelopes derived here bound what any transducer could extract at the point of use, but they do not yet account for thresholds, start-up energy, impedance matching, storage inefficiencies, or tracking overheads. Chapter 4 makes these losses explicit by mapping $\Phi(t)$ through a calibrated conversion chain, yielding deliverable power-time profiles suitable for DREAM’s scheduling and sizing.

4 Energy collection architectures and components

4.1 Introduction

Energy-autonomous embedded systems depend on a carefully engineered energy-collection path that mediates between variable ambient sources and duty-cycled loads. The aim of this chapter is to lay the foundations—architectural and model-level—needed to reason about that path early in the design, before low-level hardware choices are fixed. We treat sources, storage, conversion, and telemetry as one coherent system so that later design decisions can be traced to first-order assumptions rather than hidden safety factors.

Designing for autonomy is difficult because available power is intermittent, non-stationary, and often weak, while power-management circuits impose cold-start thresholds, quiescent consumption, and MPPT/tracking overheads that erode the harvest. Storage contributes leakage and round-trip inefficiency, with usable capacity constrained by voltage limits and aging. Methods that match long-term average harvest to long-term average load routinely fail in deployment: they ignore variability and conversion losses, leading either to brownouts during realistic profiles or to unnecessary over-provisioning. Accordingly, this chapter adopts a window-based, loss-aware perspective. We characterize energy exchange over analysis windows that capture the dominant rhythms of both input and load (e.g., diurnal/seasonal PV, building schedules for indoor nodes), and we express non-idealities explicitly via macroscopic parameters (efficiencies and parasitic terms) that can be measured or conservatively bounded.

To anchor terminology and interfaces, Section 4.2 combines two elements: (i) reference architectures—an idealized view for upper-bound reasoning and a non-ideal, practice-driven view—and (ii) a unified energy-collection model that links those architectures to measurable quantities. Treating the architecture and model together avoids duplications and makes every diagrammatic choice map to a term in the state equation (e.g., input-path efficiency, converter quiescent power, storage leakage, output-path overheads). This coupling also clarifies which signals must be surfaced to software and how frequently they need to be observed.

A guiding principle throughout is minimal yet sufficient telemetry. The runtime should depend only on what the power path can expose with negligible overhead: a reliable store proxy, protection events (power-good/UVLO), and—optionally—one coarse source proxy

per harvester for mode selection. Oversensing and high-rate control loops can dissipate a non-trivial fraction of the harvested energy and complicate verification; a sparse, well-chosen contract decouples offline modelling from online scheduling while keeping the control surface stable across hardware revisions.

The scope is intentionally system-level. We do not attempt to catalogue every converter topology or chemistry-specific effect. Instead, we abstract their system impacts—efficiency, thresholds, leakage, and quiescent currents—so that components can be swapped without rewriting the theory. Likewise, load modelling is kept to the granularity needed to represent activity states and duty cycles; application specifics (sensor mix, radio protocol, firmware structure) are treated as inputs to the same framework. This abstraction preserves generality across use cases: indoor environmental tags with sporadic radio bursts, wearables with periodic sensing and BLE advertising, or outdoor nodes with diurnal harvesting and multi-day maintenance.

Because dimensioning entails project-specific objectives and acceptance criteria, all sizing rules and workflows are deferred to Chapter 5 (Design Methodology). There, the architecture-aware model from § 4.2 is operationalized into a procedure for parameter identification (efficiencies and parasitics), minimum collector aperture and storage capacity for a target availability, cold-start/brown-out evaluation, and trace-driven verification with sensitivity analysis. Finally, § 4.6 specifies runtime requirements—maintaining a safe store band during surplus, guaranteeing survival under energy poverty by adapting duty cycle at coarse cadence, tight energy/compute budgets, and clean interaction with protection—so that firmware and hardware implement a common, low-overhead contract.

Organization of the chapter:

§ 4.2 presents the Reference Architectures + Unified Model. §§ 4.3–4.5 retain the prior surveys under new numbering—harvesting & conversion interfaces, storage technologies, and power distribution/management. Chapter 6 builds on these foundations to provide the design methodology and dimensioning workflow.

4.2 Reference architectures and energy collection model

This section consolidates the architectural viewpoint and the analytical model into a single, consistent framework. The reference architectures provide the vocabulary for sources, store, conversion stages, and loads; the unified model maps each architectural element to explicit

terms in the state equation that governs stored energy. Dimensioning rules and workflows are deferred to Chapter 6; here we define the variables, assumptions, and performance metrics they rely on.

4.2.1 Reference architectures




Figure 4.1 Ideal architecture

Figure 4.1 (idealized architecture) portrays the system at its most abstract and lossless, making visible only the energy flow between at least two heterogeneous sources, a single storage element, a power path, and one or more controlled loads. Think of two different sources—say, an indoor PV cell and a vibration harvester—delivering electrical power into a common node without contention. The power path is represented by ideal switches that never dissipate energy and that instantaneously route power either into the store or toward the loads. The storage element is perfect: it has no leakage, no charge–discharge penalty, and no voltage-dependent usable capacity; its state can be described by either stored energy E or by the store voltage V_{store} . The loads are explicitly controlled rather than always-on: the runtime can enable or disable them without side effects, and there is only a small quiescent background to represent always-on supervision. Under these assumptions, the system’s macroscopic behavior reduces to a conservation statement: the rate of change of stored energy equals the sum of source powers minus the instantaneous load demand. Over any window of time, feasibility becomes a simple energy balance: if the time-integrated harvested energy from the two sources (and any additional sources) exceeds the time-

integrated demand of the controlled loads, the system is energy-neutral. This figure is deliberately silent about start-up, thresholds, and converter behaviors; its value is to set a clean vocabulary — “at least two sources → power path → store → controlled loads” — and to provide an upper bound against which all practical designs must underperform.




Figure 4.2 Non-ideal architecture

Figure 4.2 (non-ideal architecture) refines the same four constituents while keeping the viewpoint macroscopic. The two different sources are now terminated into distinct front ends, each represented as a single block that stands in for rectification, impedance presentation or MPPT, and voltage conversion. Rather than modelling transistors or control loops, the block is parameterized by an average efficiency and by two scalar overheads: a static quiescent draw and a tracking penalty associated with following the source’s operating point. The power path ceases to be ideal. It includes selection, protection, and regulation on both the input and output sides; in the macroscopic description these become an input efficiency for each source, an output efficiency that represents delivery from the store to the rails, and a small but finite quiescent consumption that persists even when the controlled loads are idle. The storage element acquires two aggregate properties that matter at system scale: a leakage term that drains energy slowly even when nothing else happens, and a round-trip efficiency that taxes every charge–discharge cycle. Finally, the controlled loads are still actuated by software, but their interaction with the power path is now realistic: some rails require sequencing, some loads present inrush or bursty demand, and all of them are bound by protection thresholds (e.g., UVLO) that define the safe operating region. Read this figure

as a contract: each box contributes either useful energy transfer or one of a small number of losses that can be measured once and reused across analyses.

From the macroscopic perspective, the two figures differ only in whether those aggregate loss and threshold terms are set to zero or to their measured values. In the ideal figure, all efficiencies are unity and all parasitics vanish; in the practical figure, each source's front end contributes an effective harvested power that is smaller than the raw transducer power by exactly the amount lost to quiescent and tracking overheads, and the output side inflates the apparent load by the inverse of its delivery efficiency plus its own quiescent draw. The store connects the two sides by accumulating the difference and slowly bleeding energy through leakage. Importantly, both figures presuppose at least two different sources feeding the same store through a power path—this is not an implementation detail but a modelling stance. It ensures that heterogeneity is explicit in the architecture and that arbitration across sources is part of the macroscopic picture rather than an afterthought.

Interpreting the diagrams this way yields immediate design consequences without resorting to device-level models. If one of the two sources is weak—for example, a low-illuminance PV cell—its front end's static draw can dominate unless it is gated by the power path; Figure 4.2 makes that interaction visible, and the macroscopic parameters make it quantifiable. If the controlled loads exhibit bursts (a radio transmission or a sensor heater), the output side's regulation and the store's round-trip penalty together determine how much stored energy must be available before a burst can be scheduled, even when average neutrality holds. If the system must start from an empty store, the presence of thresholds and a finite start-up ladder is implied by the non-ideal architecture and must be covered in the methodology; the ideal figure deliberately ignores this, which is why it serves as an upper bound rather than a design. In short, Figure 4.1 defines the topology and the invariants for “two sources $\rightarrow$ store $\rightarrow$ power path $\rightarrow$ controlled loads,” while Figure 4.2 attaches to the same topology just enough macroscopic physics—efficiencies, quiescent powers, leakage, and thresholds—to make predictions about neutrality, survivability, and availability that match field reality.

4.2.2 Unified macroscopic model

This subsection turns the two diagrams into a single state equation that governs the evolution of stored energy under heterogeneous inputs and controlled loads. Every system we consider contains at least two different sources, a storage element, a power path, and one or more loads that software can enable or disable. The modelling stance is macroscopic: we do not

track switch nodes or inductor currents; instead, we summarize each block in Figures 4.1–4.2 by a small set of efficiencies and parasitic terms that are measurable on the bench and stable across operating conditions of interest.

Consider $m \geq 2$ sources indexed by i and k loads indexed by ℓ . Each source has an ambient stimulus $\Phi_i(t)$ (irradiance, temperature gradient, vibration level, RF field strength) that produces a raw electrical power $P_i(t)$ at its transducer terminals. The power path presents each source to the store through a front end $N_{\text{conv},i}$ that encompasses rectification, impedance presentation or MPPT, and voltage conversion. Rather than modelling those circuits explicitly, we attribute to $N_{\text{conv},i}$ an average input-path efficiency $\eta_{\text{in},i}$, a static quiescent draw $P_{q,i}$, and a tracking penalty $P_{\text{trk},i}$ associated with extremum seeking or dithering. The effective harvested power from source i is therefore

$$P_{i,\text{eff}}(t) = \eta_{\text{in},i} P_i(t) - P_{q,i} - P_{\text{trk},i} \quad (4.1)$$

These parameters can be allowed to depend on operating point—weak-light versus bright-light, cold versus hot—without changing the form of the model; in practice they are identified as functions or piecewise-constant maps over the relevant ranges.

The loads together impose an application demand $P_L(t) = \sum_{l=1}^k P_{L,l}(t)$ that is seen by the store after the output side of the power path inflates it with delivery losses and adds its own housekeeping. The macroscopic representation of that effect is an output-path efficiency η_{out} and a quiescent term $P_{q,\text{out}}$. Storage leakage appears as an additional drain proportional to the instantaneous store voltage. The effective demand becomes

$$P_{\text{req}}(t) = \frac{P_L(t)}{\eta_{\text{out}}} + P_{q,\text{out}} + I_{\text{leak}} V_{\text{store}}(t) \quad (4.2)$$

Controller overhead can be treated either as part of $P_{q,\text{out}}$ (if it is tied to the power path) or absorbed into the load profile (if it scales with application activity). The distinction does not alter the algebra; it only changes where the overhead is booked.

Let $E(t)$ denote the energy stored. The storage element contributes round-trip losses, captured by charge and discharge efficiencies η_{ch} and η_{dis} . Round-trip efficiency (RTE) formalizes the net energy penalty of storing and later delivering energy at the same modelling boundary used throughout this subsection. Over a matched charge–discharge cycle, with

energies measured at the storage element's DC terminals $RTE = E_{out}/E_{in}$, where $E_{in} = \int_{t_c} v_{in}(t)i_{in}(t)dt$ and $E_{out} = \int_{t_d} v_{out}(t)i_{out}(t)dt$. Within the macroscopic model, DC-boundary RTE decomposes naturally into direction-specific factors η_{ch} and η_{dis} , while conversion inefficiencies and quiescent draws remain attributed to the input and output power paths; this separation avoids double counting when composing sources, storage, and loads and is consistent with the loss channels already identified for the non-ideal architecture.

Because duty cycles in CPS often include idle intervals between charge and discharge, it is useful to distinguish an intrinsic DC RTE, $\eta_{rt}^{int} = \eta_{ch}\eta_{dis}$, from an effective RTE over a cycle that includes a dwell of duration τ . If $P_{leak}(t)$ denotes the leakage power during the dwell (including redistribution effects for EDLC stacks), then for a charge packet of energy E_{in} followed by dwell and discharge,

$$RTE_{eff} = \eta_{dis} \left(\eta_{ch} - \frac{1}{E_{in}} \int_0^\tau P_{leak}(t) dt \right) \quad (4.3)$$

This expression makes explicit that long idle times degrade the apparent round-trip performance even when η_{ch} and η_{dis} are unchanged. Accordingly, parameter reporting and identification should state the measurement boundary, current profile, temperature, and dwell time so that neutrality and availability computations later in the chapter operate on effective energies rather than idealized ones, matching the window-based, loss-aware stance adopted in this chapter. In the regime of interest, a single multiplier η_{ch} on the incoming side suffices; if asymmetric behavior is significant, the cycle can be split into charge and discharge intervals with the corresponding multipliers applied to each. The state equation that results from the non-ideal architecture is

$$\frac{dE}{dt} = \sum_{i=1}^m \eta_{ch} P_{i,eff}(t) - P_{req}(t) \quad (4.4)$$

When all efficiencies are unity and all parasitic terms vanish, this reduces to the ideal law of Figure 4.1, $\dot{E} = \sum_i P_i - P_L$, which provides the well-known upper bound used for early feasibility checks.

To relate the energy state to a measurable quantity, we use the store voltage. For a supercapacitor of capacitance C , $E(t) = \frac{1}{2} CV_{store}^2(t)$ and $\dot{E} = CV_{store} \dot{V}_{store}$; the pair of

relations closes the model in terms of voltage. For batteries, $E(t)$ is tied to state-of-charge through the device's open-circuit-voltage curve; an internal-resistance term may be included when short bursts dominate the load, but for macroscopic energy accounting the OCV–SoC mapping and rate limits are sufficient. In either case, the model is driven by traces of $P_i(t)$ or, in modalities where an area or coupling factor is the primary design variable, by an effective stimulus representation $P_{in}(t) = A\eta_h\Phi(t)$ with A an equivalent aperture and η_h a lumped harvester efficiency. Collapsing two heterogeneous sources to a single effective input is legitimate for sizing if the downstream scheduler does not rely on their individual timing; where arbitration between sources matters, the two $P_{i,eff}$ terms should be retained separately.

Certain behaviors that are prominent in Figure 4.2 enter the model as constraints rather than continuous terms. Cold-start is represented by a threshold voltage V_{CS} and an associated start-up energy E_{CS} : the system cannot transition from “off” to “on” until $E(t)$ exceeds the energy required to bring up the front ends and the controller; below that boundary, loads are disconnected, and the only dynamics are those of trickle charging and leakage. Brown-out is represented by hysteretic undervoltage thresholds; when V_{store} falls below the lower threshold, the loads are shed and P_L collapses to its supervisory minimum, only to be re-enabled once the upper threshold is crossed. In trace-driven simulations this is implemented as event logic wrapped around the state equation; the differential model itself remains unchanged.

The identification of macroscopic parameters follows directly from the block meanings. Quiescent powers are measured with switching disabled at representative voltages and temperatures. Input- and output-side efficiencies are tabulated by sweeping operating points and taking the ratio of delivered to received power; the resulting maps are interpolated during simulation. Tracking penalties are obtained by comparing harvested power under MPPT with the power at a fixed operating point under matched conditions. Leakage is measured by holding the store at several voltages and recording the decay rate; an equivalent current model $I_{leak}(V,T)$ is fitted for use in P_{req} . Start-up energy is integrated along the cold-start ladder from a fully discharged state until regulation is established. None of these measurements require knowledge of switching waveforms; they are purposefully macroscopic, matching the level of the model.

The outcome is a compact description in which every box in the non-ideal diagram corresponds to a symbol in the equations and every symbol can be tied to a bench measurement. In § 4.2.3 the model is placed on a design window and accompanied by performance metrics—energy neutrality, survivability with respect to protection thresholds, and availability—that Chapter 5 turns into a sizing and verification procedure.

4.2.3 Design window, traces and signals

The model is exercised on a finite design window that is long enough to contain the dominant rhythms of both the heterogeneous sources and the controlled loads. For outdoor photovoltaic input the window must span seasonal variation; for indoor sources tied to building schedules or occupancy it should include multiple weekday–weekend cycles; for hybrid systems with at least two different sources the window must be long enough for their independent fluctuations to fully express. In practice the window is defined by a start time, a duration, and a uniform sampling interval.

$$W = [t_0, t_0 + T_W], \quad t_k = t_0 + k \Delta t, \quad k = 0, 1, \dots, N, \quad \Delta t = \frac{T_W}{N} \quad (4.5)$$

All signals—source stimuli or powers, store proxies, and load demand—are expressed on the same time base so that energy accounting is unambiguous.

Source traces can be provided either as raw stimuli (irradiance, temperature gradient, vibration level, RF field strength) or directly as electrical power at the transducer terminals. When stimuli are used, a calibrated transducer model converts them to electrical power before the front-end macromodel is applied. The two (or more) sources must remain distinct during this step so that policies and source-specific losses can be applied; only later, if the scheduler never arbitrates between sources, may an effective single input be formed for convenience. Load traces originate from the firmware schedule and from benchtop power profiling of each activity state. Bursty behavior such as radio transmissions is represented explicitly in the time series rather than smeared as a mean; this preserves the interaction between surge currents, output regulation efficiency, and store dynamics. All traces are resampled to a uniform cadence by energy-preserving averaging rather than point sampling, which avoids spurious neutrality errors when down sampling high-rate logs:

$$\bar{P}[k] = \frac{1}{\Delta t} \int_{t_k}^{t_{k+1}} P(\tau) d\tau \quad (4.6)$$

Cold-start and protection thresholds enter the simulation as events that gate the continuous dynamics rather than as additional differential terms. Below the cold-start energy the converters and loads remain disabled; once the store crosses the enable threshold, the front ends and rails become available, and the state equation governs the evolution. During brown-out the loads are shed, and only supervisory consumption remains until the upper hysteresis threshold is crossed again. This event logic is applied to the discrete-time traces with the same sampling interval used for integration. With the stored energy E_k at instant t_k , the forward update is:

$$E_{k+1} = E_k + \Delta t \left(\sum_{i=1}^m \eta_{ch} P_{i,eff}[k] - P_{req}[k] \right) \quad (4.7)$$

with the effective demand seen by the store given by

$$P_{req}[k] = \frac{P_L[k]}{\eta_{out}} + P_{q,out} + I_{leak} V_{store}[k] \quad (4.8)$$

Front-end enabling is likewise represented with a binary mask per source; The mask prevents the accumulation of quiescent losses for sources that are intentionally disconnected by the power path, which is critical when one of the two sources is weak.

$$P_{i,eff}[k] = u_i[k] (\eta_{in,i} P_i[k] - P_{q,i} - P_{trk,i}), \quad u_i[k] \in \{0,1\} \quad (4.9)$$

For capacitive storage, propagating the voltage directly is convenient; Using $E = \frac{1}{2} CV^2$, the consistent closed-form update is:

$$V_{store}[k+1] = \sqrt{V_{store}^2[k] + \frac{2\Delta t}{C} \left(\sum_{i=1}^m \eta_{ch} P_{i,eff}[k] - P_{req}[k] \right)} \quad (4.10)$$

The cold-start and hysteretic brown-out gates are applied as discrete events:

$$E_k < E_{CS} \Rightarrow u_i[k] = 0 \forall i, P_L[k] \rightarrow P_{supervisory},$$

$$V_{store}[k] \leq V_{UVLO}^- \Rightarrow P_L[k] \rightarrow P_{supervisory}, V_{store}[k] \geq V_{UVLO}^+ \Rightarrow \text{loads may be re-enabled} \quad (4.11)$$

Signal cadence serves two different purposes and must not be conflated. For modelling, the integration step is chosen small with respect to the store's time constant and the fastest burst

features in the load so that numerical error remains negligible; this step can be arbitrarily fine because it is computed offline. For runtime telemetry, the sampling rate is chosen coarsely enough that the average cost of sensing and decision making is negligible compared to the median harvested power. A single reliable store proxy and protection events suffice for most schedules; optional source proxies may be logged for diagnostics or mode hints, but they do not change the correctness of the energy accounting.

Missing data and outliers in measured traces are handled conservatively. When a source measurement is absent, interpolation is performed only over gaps short relative to the store time constant; longer gaps are filled with lower confidence bounds so that feasibility is not overstated. Load traces are always kept non-negative and aligned to activity start and end timestamps. All timestamps are normalized to a single time zone and clock epoch so that the composition of sources and loads is well-defined at every sample.

4.2.4 Definitions

The model is evaluated on a design window long enough to reveal the joint variability of at least two heterogeneous sources and the controlled loads. Within that window we use three system-level metrics that are defined purely at the macroscopic level and do not depend on circuit detail. The first is energy **neutrality**, which compares the energy effectively harvested after input-path and storage losses to the energy effectively demanded once output-path losses and leakage are accounted for. Writing the two window energies as integrals of the powers introduced in §4.2.2 gives

$$E_W^{in} = \int_W \sum_{i=1}^m \eta_{ch} P_{i,eff}(t) dt, E_W^{req} = \int_W P_{req}(t) dt \quad (4.12)$$

Neutrality holds if and only if the harvested energy meets or exceeds the requirement on the same window,

$$E_W^{in} \geq E_W^{req} \quad (4.13)$$

It is often convenient to track the instantaneous deficit—the algebraic difference between effective demand and effective harvest—because it reveals when the store must discharge. Using the same macroscopic terms, the deficit is

$$\delta(t) = P_{\text{req}}(t) - \sum_{i=1}^m \eta_{ch} P_{i,\text{eff}}(t) \quad (4.14)$$

Neutrality over the window corresponds to a non-positive time integral of $\delta(t)$; however, we postpone any use of the cumulative deficit for sizing to Chapter 5.

The second metric is **survivability**, which formalizes the requirement that the system remain within its protection limits throughout operation. In continuous time this demands that the store voltage never crosses the brown-out threshold and that any upper limits (for example, over-voltage protection on supercapacitors) are respected. Expressed as inequalities on the store proxy, the constraint reads

$$V_{\min} \leq V_{\text{store}}(t) \leq V_{\max} \quad \text{for all } t \in W \quad (4.15)$$

where $V_{\min}$ and $V_{\max}$ represent the effective lower and upper bounds implied by UVLO/OVP settings and safe operating limits. In the trace-driven replay this condition is enforced by the event logic of §4.2.3; here it serves as the formal definition against which trajectories are judged.

The third metric is **availability**, the fraction of the window during which every required load is powered without violating the survivability condition. Because the loads are controlled, availability is sensitive not only to energy balance but also to the timing of bursts relative to the store dynamics and to protection hysteresis. Using the indicator of successful service and safety, the definition over a window of duration $|W|$ is:

$$n_{av} = \frac{1}{|W|} \int_W 1 \text{ all required loads served and } V_{\min} \leq V_{\text{store}}(t) \leq V_{\max} dt \quad (4.16)$$

This metric is agnostic to the specific duty-cycling policy; it credits time only when the loads are enabled and the store remains inside the safe band. Chapter 5 turns these definitions into computable tests on measured or synthetic traces and shows how neutrality, survivability, and availability jointly determine the minimal collector aperture and storage capacity for a given design target.

4.2.5 Cold-start, Brown-out and survivability aspects

The unified model must be complemented by event constraints that capture how the system transitions between “off,” “start-up,” “normal,” and “survival” operation. These constraints

are macroscopic and apply to any system composed of at least two heterogeneous sources, a storage element, a power path, and controlled loads.

Cold-start is a boundary in state space rather than an ordinary loss term. While the store energy remains below the start-up requirement, front-end converters and application rails cannot be brought under control, and only passive charging through the weakest permissible path is possible. We represent this by a gating condition on energy and voltage: conversion and controlled loads may be enabled only once the stored energy exceeds the cold-start energy and the store voltage exceeds the enable threshold,

$$E(t) \geq E_{CS}, V_{\text{store}}(t) \geq V_{CS} \quad (4.17)$$

The cold-start requirement itself aggregates the work needed to charge the start-up buffer from its “dead” state to the regulator enable threshold and to power any digital bring-up; a convenient lower bound, useful for trace-based checks, is

$$E_{CS} \geq \frac{1}{2} C_{\text{buf}}(V_{\text{en}}^2 - V_{\text{dead}}^2) + E_{\text{boot,ctrl}} + E_{\text{boot,conv}} \quad (4.18)$$

In replay, the start-time is the first instant at which both inequalities are satisfied; prior to that moment the effective powers in the state equation are evaluated with all converters disabled and only leakage and passive input paths active.

Brown-out is the dual constraint that governs safe shut-down and prevents oscillation under energy poverty. The power path enforces a hysteretic undervoltage-lockout, with a lower threshold at which loads are disconnected and an upper threshold at which they may be re-enabled. We write this as

$$V_{\text{store}}(t) \leq V_{UVLO}^- \Rightarrow \text{loads off}, V_{\text{store}}(t) \geq V_{UVLO}^+ \Rightarrow \text{loads may be enabled} \quad (4.19)$$

Within the dead band between V_{UVLO}^- and V_{UVLO}^+ the scheduler must refrain from wake-ups that would immediately violate protection; the event logic in §4.2.3 enforces this during trace replay and underpins the survivability criterion defined in §4.2.4.

Survival operation formalizes the guarantee that, when inputs remain weak for an extended period, the system can continue in a minimal service mode for a prescribed horizon. Let $P_{\text{surv}}(t)$ denote the demand of this mode after delivery losses and housekeeping are accounted for. A sufficient condition for a survival horizon T_{sur} beginning at time t is that the usable

energy above the lower protection bound covers the cumulative requirement of the horizon. In continuous time this is expressed as:

$$\int_t^{t+T_{\text{sur}}} \left(\frac{P_{\text{surv}}(\tau)}{\eta_{\text{out}}} + P_{q,\text{out}} + I_{\text{leak}} V_{\text{store}}(\tau) \right) d\tau \leq E(t) - E_{\text{min}}, \quad (4.20)$$

where E_{min} is the energy corresponding to the lower voltage limit. For a capacitive store $E_{\text{min}} = \frac{1}{2} C V_{\text{min}}^2$ and $E(t) = \frac{1}{2} C V_{\text{store}}^2(t)$ so the right-hand side reduces to the band of usable energy between the instantaneous voltage and the protection floor. This condition does not prescribe how the scheduler chooses epoch lengths or task subsets; it simply states the macroscopic energy feasibility of surviving for T_{sur} without violating the hysteretic bounds. The design methodology in Chapter 5 turns this inequality into concrete rules for choosing survival cadence and for proving that the chosen cadence respects the protection hysteresis for the measured traces.

4.2.6 Telemetry and control interfaces

The unified model relies on a small set of runtime signals and actuators; the contract here states exactly what must be exposed by the power path so that software can operate correctly without incurring significant energy overhead. The goal is not to re-enumerate losses or restate the state equation, but to specify how measurements and commands map to the macroscopic quantities already defined, and how they should be timed and interpreted in a system that includes at least two heterogeneous sources, a storage element, a power path, and controlled loads.

At the core of the interface is a single store proxy. For capacitive storage this is the store voltage V_{store} sampled at a cadence slow relative to the store's time constant, yet fast enough to resolve the largest expected bursts of load activity. For batteries the proxy should be a temperature-aware SoC estimate derived from an OCV–SoC map or from coulomb counting corrected by periodic rest measurements; whichever proxy is used, the firmware must treat it as the state variable of the macroscopic model. The accuracy requirement is modest: absolute errors that translate into energy uncertainty small compared to the energy swing between the protection thresholds are sufficient, because the model's role is to steer timing and duty cycle, not to regulate a fixed voltage.

Protection events provide the second pillar of the interface. The power path must generate edge-triggered notifications when the store crosses the hysteretic undervoltage bounds and, where present, any over-voltage or temperature protection limits. These events have semantic priority over telemetry samples: they may immediately pre-empt scheduled activity by forcing loads off at the lower bound and authorizing re-enable at the upper bound. In trace-based replay they appear as instantaneous gates on the load series; at runtime they are delivered through an interrupt or wake-up source with debounce and monotonic sequence numbers so that no transition is missed when the controller is sleeping.

Optional source proxies complete the measurement side. Each enabled source may export a single scalar that is monotonically related to its instantaneous capability—fractional open-circuit voltage for PV, settled input current or equivalent Seebeck estimate for TEG, rectified envelope for RF, and stroke-synchronous charge for vibration. These proxies are not required for correctness of the energy accounting, which already incorporates input-side losses, but they do enable simple policies such as preferring the stronger source when both are available or throttling MPPT cadence under slowly varying stimuli. In multi-source systems these readings must be tagged with the identity of the source and with the enable state from the selection network so that the scheduler never mistakes a disabled input for a weak one.

On the actuation side the contract consists of rail enables, source selection, and coarse MPPT hints. Rail enables are level-sensitive controls that connect loads to their regulators; their semantics respect the protection gates, meaning asserts below the lower threshold remain latched off until the upper threshold is crossed. Source selection is expressed as a mask over the input network; it is the software-visible counterpart of the u_i mask used in the model and exists to avoid accumulating quiescent losses from inactive front-ends when one of the two sources is weak. MPPT hints are advisory toggles that reduce tracking cadence or fix the operating point during periods when the store is far from its limits; they are optional because their effect is captured at macroscopic level by the tracking penalty term yet exposing them allows the scheduler to trade a small loss in instantaneous harvest for a reduction in overhead when input change slowly.

Cadence and budgeting bind the contract to the energy model. Telemetry sampling must be paced so that its average power cost, plus the amortized cost of decision logic and housekeeping, is negligible compared to the median harvested power; in practice this means

sampling on the order of the store's electrical time constant in quiescent conditions and temporarily increasing cadence only around expected bursts or near protection thresholds. Anti-aliasing is achieved by averaging over windows commensurate with the sampling interval rather than by impulsive reads, aligning the runtime measurements with the energy-preserving resampling used offline. Slope limiting on the scheduling variables (for example, a maximum change in sleep interval per update) further prevents chatter induced by measurement noise without requiring any change to the underlying model.

To remain portable across hardware revisions, the interface must be versioned and self-describing. Each measurement record includes a timestamp on the common epoch used by trace replay, a source or rail identifier where applicable, and an indication of validity when the system is in start-up or brown-out states. Each actuator acknowledges with its effective state and, for selections that can be pre-empted by protection, with the reason code of the pre-emption. Temperature is logged at the same cadence as the store proxy whenever battery SoC estimation depends on it; otherwise, it can be sampled more sparingly and only when the controller is awake for other reasons.

Finally, the controller's own energy must be booked consistently. If rail enables and source selection are considered part of the power-distribution domain, their quiescent and dynamic costs belong with the output-side housekeeping; if they scale with application duty cycle, they may instead be folded into the load profile. Either choice is acceptable as long as it is applied uniformly in both runtime accounting and offline replay. With this contract, the architecture-model pair of § 4.2 can be exercised by software with minimal overhead, and Chapter 5 can layer a design methodology on top without depending on circuit-level detail.

4.3 Energy harvesting and conversion systems

4.3.1 Ambient energy sources

For design, PV sources behave approximately as current-sourced nonlinear devices whose MPP requires impedance presentation near R_{MPP} and low ripple at the source terminals; indoor operation additionally demands ultra-low-leakage VOC sampling and spectrum-aware MPPT. General irradiance characteristics are summarized in Chapter 3.

At electrical timescales a TEG is well-approximated by a Thevenin source (V_{oc} , R_i); however, optimal loading depends on the coupled thermal network, so MPPT cadence must

be slow relative to thermal time constants. General ΔT availability is summarized in Chapter 3.

Vibration and triboelectric harvesters. Mechanical harvesters produce AC outputs with high source impedance and a narrowband response centered at the mechanical resonance. Delivered power depends on the vibration spectral density and coupling. Rectification strategy and, in some regimes, synchronized charge extraction strongly affect the usable energy.

Far-field RF (special cases). At practical ranges, available power is typically in the micro-watt regime and depends on frequency allocation, antenna gain, and environment. RF is therefore considered only in proximity scenarios or as a supplemental source.

4.3.2 Transducer characteristics and source models

For design, each transducer can be represented by a compact equivalent model that captures the dominant physics over the intended operating range. Photovoltaic (PV) devices are well described by the single-diode I-V model with series and shunt parasitics; the maximum-power point (MPP) varies with irradiance and temperature and often lies near a fixed fraction of the open-circuit voltage under outdoor conditions [114], [115]. These characteristics motivate maximum-power-point tracking (MPPT) and place requirements on input impedance and ripple [116], [117]. Thermoelectric generators (TEGs) admit a Thevenin representation with open-circuit voltage proportional to the Seebeck coefficient and temperature difference, and an internal resistance set by leg geometry; the optimal electrical loading depends on the coupled thermal network rather than on simple resistance matching [118], [119]. Vibration harvesters are conveniently modelled as single-degree-of-freedom resonant systems with electromechanical coupling; piezoelectric devices appear as high-impedance AC current sources with parallel capacitance, while electromagnetic devices appear as AC voltage sources with series inductance and resistance [120], [121], [122], [123]. Interface circuits that employ synchronous charge extraction (e.g., SSHI/DSSH) modify the effective electrical damping and can increase harvested energy under resonant excitation [124]. These abstractions permit prediction of input power versus operating point and specify the constraints.

4.3.2.1 Photovoltaic transducers

A crystalline-silicon or thin-film photovoltaic module is well described by the single-diode model with series and shunt parasitics. In the following figure, the equivalent circuit is shown.




Figure 4.3 Photovoltaic transducer equivalent model

The output current of such a module is calculated as:

$$I = I_{ph} - I_s \left(e^{\frac{V+IR_s}{nV_T}} - 1 \right) - \frac{V + IR_s}{R_{sh}} \quad (4.21)$$

Where I and V are the terminal current and voltage, I_{ph} is photo-generated current (approximately proportional to irradiance and weakly dependent on temperature), I_s is diode saturation current, R_s and R_{sh} are series and shunt resistances, n is the ideality factor, and $V_T = kT/q$ the thermal voltage. The MPP (V_{MPP} , I_{MPP}) satisfies $\partial(VI)/\partial V = 0$ and varies with irradiance and temperature though I_{ph} , I_s and the parasitics. For crystalline modules outdoors, V_{mpp} is typically 70-80% of V_{oc} ; at low-lux indoor conditions this fraction is lower and more spectrum sensitive [115], [116], [117].

From input-converter perspective, the key requirement is to present an effective input impedance that places the module at, or sufficiently near, the MPP. A useful proxy is the MPP resistance

$$R_{mpp} \triangleq \frac{V_{mpp}}{I_{mpp}} = \left(\frac{dI}{dV} \right)^{-1} \quad (4.22)$$

Ripple on V or I reduces harvested power due to convexity of the P-V curve; hence, MPPT sampling and converter switching should be duty-cycled so that induced ripple and overheads remain negligible relative to harvested power [116], [117].

Parameter identification. The set $\{I_{sc}, V_{oc}, FF, \alpha T(\cdot)\}$ can be extracted from I-V sweeps under representative irradiance and temperature; when detailed first are required, one-diode parameters $\{I_s, n, R_s, R_{sh}\}$ can be obtained by two-slope or least-squares procedures [114]. The empirical fill-factor graphs and expressions of Green are often useful for quick checks and sensitivity analysis [115].

4.3.2.2 Thermoelectric generators

At electrical timescales, a TEG module is well represented by a Thevenin source




Figure 4.4. Equivalent of thermoelectric module

with open-circuit voltage and internal resistance:

$$V_{oc} = aN\Delta T, \quad R_i = R_{legs}(T) \quad (4.23)$$

where α is Seebeck coefficient per thermocouple, N the number of thermocouples, $\Delta T = T_h - T_c$ is the temperature difference across the module, and R_i the temperature-dependent electrical resistance of the legs. The terminal equation is:

$$V = V_{oc} - IR_i = aN\Delta T - IR_i \quad (4.24)$$

So that, for a load R_L ,

$$I = \frac{V_{oc}}{R_i + R_L}, \quad P_L = I^2 R_L = \frac{V_{oc}^2}{(R_i + R_L)^2} R_L \quad (4.25)$$

If ΔT were fixed, the electrical power delivered to a load R_L would be maximized at $R_L = R_i$ with

$$P_{L,max} = \frac{V_{oc}^2}{4R_i} \quad (4.26)$$

In practice, electrical loading perturbs the thermal gradient. A standard lumped model writes the heat rates at the hot and cold junctions as:

$$Q_h = \alpha T_h I - \frac{1}{2} I^2 R_i + K \Delta T \quad (4.27)$$

$$Q_c = \alpha T_c I - \frac{1}{2} I^2 R_i + K \Delta T \quad (4.28)$$

With thermal conductance K (including leg conduction and parasitics). Eliminating $Q_h - Q_c$ gives the operating-point dependence of the gradient

$$\Delta T = \Delta T_0 - \Theta_{th} \left(\Pi I - \frac{1}{2} I^2 R_i \right) \quad (4.29)$$

Where ΔT_0 is the open-circuit gradient set by boundary conditions, $\Pi = \alpha \bar{T}$ is the Peltier coefficient at average temperature $\bar{T} = (T_h + T_c)/2$, and Θ_{th} is the effective thermal resistance of the contacts and environment. Because ΔT now depends on I , the power-maximizing R_L generally deviates from the naïve match $R_L = R_i$; the optimum should be found numerically from measured K , Θ_{th} and α .

The instantaneous electrical output and efficiency are:

$$P_e = IV = I(\alpha N \Delta T - I R_i), \quad \eta = \frac{P_e}{Q_h} \quad (4.30)$$

Using the standard material figure of merit

$$Z = \frac{\alpha^2}{R_i K}, \quad ZT = Z \bar{T} \quad (4.31)$$

An upper-bound (idealized contacts) for conversion efficiency is

$$\eta_{max} = \frac{\Delta T}{T_h} \frac{\sqrt{1 + Z \bar{T}} - 1}{\sqrt{1 + Z \bar{T}} + \frac{T_c}{T_h}} \quad (4.32)$$

With $\bar{T} = (T_h + T_c)/2$. Current density that maximizes local efficiency follows the “compatibility factor” formulations; 1-D,

$$s = \frac{\sqrt{1 + ZT} - 1}{\alpha T} \quad (4.33)$$

Which guides spatially varying designs and indicates that constant-property legs are generally sub-optimal over large temperature spans. [118], [119].

Converter implications.

Two practical consequences follow for the input power stage:

- i. At small ΔT , the available Voc may lie below the cold-start threshold of the downstream converter, requiring staged start-up or charge-pump assistance
- ii. Because the effective source resistance seen by the converter depends on ΔT and contact quality, aggressive perturb-and-observe tracking should be bandwidth-limited to avoid exciting thermal transients.

The pair (Voc, Ri) can be measured by open-circuit and small-signal tests across controlled ΔT ; the Seebeck coefficient follows the slope of Voc versus ΔT . Thermal contact resistances can be estimated from transients to heat-flow steps.

4.3.2.3 Vibration and triboelectric transducers

A broad class of vibration harvesters is well approximated by a single-degree-of-freedom mass–spring–damper system with electromechanical coupling [120], [121], [122], [123]. For base acceleration $y(t)$, relative displacement x , mass m , mechanical damping c_m , stiffness k , and coupling coefficient θ , the linearized dynamics yield

$$m\ddot{x} + c_t\dot{x} + kx = -m\ddot{y}, \quad \omega_0 = \sqrt{\frac{k}{m}}, \quad Q = \frac{m\omega_0}{c_t} \quad (4.34)$$

where $c_t = c_m + c_e$ includes electrical damping c_e induced by the transducer and load. Under stationary base excitation with acceleration spectral density $S_a(f)$, the average extractable power increases with both Q and the product of mechanical and electrical damping; maximum occurs when electrical damping matches mechanical damping (load matching) [120]. Piezoelectric harvesters are well represented electrically by an AC current source in parallel with a capacitance C_p ; electromagnetic harvesters map to an AC voltage source with series inductance and resistance [121], [122]. Rectification (full-bridge vs synchronous) and

synchronous charge-extraction (SSHI/DSSH) alter the effective c_e and can increase energy per cycle [122], [124].

Triboelectric nanogenerators (TENGs). TENGs convert mechanical motion to electricity via contact electrification and electrostatic induction. Four canonical modes exist—contact-separation, lateral sliding, single-electrode, and freestanding—and device output scales with interfacial area, surface charge density and relative speed. TENGs present ultra-high internal impedance ($M\Omega$ – $G\Omega$), high voltages and low currents, producing bursty waveforms; consequently, the front-end must minimize leakage, rectify with very low bias, and avoid loading that collapses surface charge density.




Figure 4.5 Triboelectric generator equivalent circuit

For rotational stimuli, rotating TENGs (R-TENGs) exploit sliding or non-contact operation to trade off charge density against wear. A comprehensive review by Segkos et al details operating principles, impedances, waveform characteristics and design strategies, including hybrid EMG–TENG combinations and charge-pumping interfaces that are directly relevant to impedance presentation and matching at the power path[125]. Building on this, Bardakas et al demonstrate a rotational–linear TENG (RL-TENG) that converts rotation into contact–separation motion using ZnO-on-Kapton tribolayers, documenting load-sweep methodology and measured outputs that can be reused for sizing and verification[126]. In parallel, Anastasopoulos et al. report flexible triboelectric tactile sensors, underscoring the dual role

of TENG structures as both harvesters and self-powered sensors and the need for high-impedance readout/rectification for repeatability[127].

4.3.2.4 *Far-field RF rectennas*

An RF energy harvester (rectenna) comprises an antenna, an impedance-matching network, a nonlinear rectifier, and a smoothing/output network. In line-of-sight, narrowband conditions the received power is well approximated by Friis' transmission relation[128]:

$$P_{rx} = P_t G_t G_r \left(\frac{\lambda}{4\pi R} \right)^2 L_{pol} L_{sh} \quad (4.35)$$

where P_t is transmit power, G_t and G_r are the transmitter/receiver gains, λ is the wavelength, R the separation, and L_{pol}, L_{sh} capture polarization/mismatch and shadowing/fading.

At the rectifier input, the RF-to-DC efficiency depends strongly on input level in the nanowatt–microwatt regime relevant to ambient harvesting [129]:

$$\eta_{rect}(P_{rf}) = \frac{P_{dc}}{P_{rf}} \quad (4.36)$$

A small-signal expansion of a diode around a bias V_0 (self-bias if no DC source) with $v_d(t)=V_0+v_{rf}(t)$ yields

$$i_d(t) \approx I_s e^{\frac{V_0}{nV_T}} \left(1 + \frac{v_{rf}(t)}{nV_T} + \frac{v_{rf}^2(t)}{2(nV_T)^2} \right) - I_s \quad (4.37)$$

so that the rectified DC component (time-average over a carrier period) scales as

$$I_{dc} \approx I_s e^{\frac{V_0}{nV_T}} \frac{\langle v_{rf}^2 \rangle}{2(nV_T)^2} \quad (4.38)$$

This square-law behavior at low power motivates (i) impedance matching at the intended level to maximize v_{rf} ; (ii) exploiting waveforms with higher crest factor where permissible; and (iii) careful co-design with the subsequent DC–DC stage, whose input impedance fixes the rectifier's operating point[129], [130].

Because the rectifier input impedance $Z_{in}(P_{rf}, T)$ varies with power and temperature, narrowband high-Q matching maximizes peak efficiency only over a limited dynamic range. Broadband networks reduce sensitivity to level/frequency drift at some cost in peak

efficiency. The impact of mismatch can be expressed through the power transmission coefficient [131]:

$$L_{\Gamma} = 1 - |\Gamma|^2, \quad \Gamma = \frac{Z_{\text{in}} - Z_0}{Z_{\text{in}} + Z_0} \quad (4.39)$$

So even modest errors in Z_{in} near threshold translate into superlinear losses in P_{dc} via the quadratic dependence on v_{rf} .

Ambient field surveys provide realistic P_{rx} distributions in indoor/urban environments—typically tens of nW to a few μW at the antenna port—together with system-level efficiencies achieved with practical antennas and rectifiers. These data are essential for budgeting cold-start, storage pre-charge, and duty-cycling policies[130]. When the rectified node must cold-start a downstream converter, the rectenna’s smoothed output must reliably exceed the converter’s start-up threshold VCS under worst credible incident fields; staged start-up or charge-pump assistance may be required in the sub- μW regime [129], [130].

4.3.3 Input power-conditioning and maximum power point tracking

Input power conditioning (IPC) mediates between an intermittent transducer and an energy store or load by shaping electrical boundary conditions at the source terminals, providing rectification and voltage/current conversion, and executing a maximum power point tracking (MPPT) policy that is net-energy positive under ultra-low-power constraints. In autonomous nodes, the IPC must (i) maintain the source near its instantaneous maximum-power point (MPP), (ii) guarantee cold-start from the weakest credible stimulus, and (iii) limit quiescent overheads so that control consumes significantly less energy than it enables. This section formalizes these roles, then tailors MPPT choices to photovoltaic (PV), thermoelectric (TEG), vibration/triboelectric, and far-field RF harvesters.

Roles, losses and control cadence

The power–voltage characteristic $P(v) = v i(v)$ of most harvesters is locally concave around the MPP; hence ripple in terminal voltage produces a Jensen-type penalty that depresses average power. Input networks and control laws should therefore be designed to constrain the ripple magnitude Δv relative to the local curvature of $P(v)$, while recognizing that further attenuation increases component size and quiescent loss. Because environmental dynamics vary widely—milliseconds for vibrational strokes, seconds for PV irradiance ramps, minutes

for thermal drifts—the tracker’s update period must be chosen from measured source time constants, rather than fixed a priori. As a rule-of-thumb for low-power nodes, the MPPT update period is set between one-fifth and one times the dominant source time constant, with additional event-driven acceleration when a fast proxy (e.g., a rectified envelope) indicates a step change.

Cold-start can dominate feasibility. Typical switch-mode converters require hundreds of millivolts to oscillate; dim-light PV, small- ΔT TEGs, and weak RF produce less. Practical IPCs therefore adopt staged start-up (passive rectification into a reservoir followed by a self-oscillating step-up that hands over to a synchronous converter once a UVLO threshold is crossed) or specialized cold-start ICs that exploit transformers and charge-pumps designed for tens of millivolts.

MPPT strategies: selection and net-power optimality

While many algorithms converge to the MPP on paper, few are competitive when measurement, computation, and actuation energy are accounted for. Four strategies are broadly useful:

- i. Fractional open-circuit/short-circuit sampling. For PV, regulating $v \approx k V_{OC}$ or $i \approx k I_{SC}$ (with k characterized for the cell and spectrum) yields good results indoors and under quasi-stationary conditions. The method requires sparse, low-leakage sampling of V_{OC} or I_{SC} and minimal arithmetic, which benefits tiny budgets [132], [133], [134].
- ii. Perturb-and-observe (P&O) with adaptive step size. Classical hill-climbing is attractive for its simplicity; adaptive step rules reduce limit-cycle loss and accelerate recovery after disturbances [135].
- iii. Incremental conductance. Estimating di/dv and solving $dP / dV=0$ provides better behavior during irradiance ramps than fixed-step P&O, at the cost of additional sensing [132], [134].
- iv. Proxy-based extremum seeking. When direct $v-i$ measurement is expensive or ill-posed (e.g., in high-impedance triboelectric and RF rectifiers), a correlated proxy—charge per stroke, rectified DC level, envelope power—can be optimized by slow hill-climbing across discrete settings [111], [136], [137].

In all cases, the net-power condition must hold: average harvested power with MPPT minus without MPPT must exceed the tracker's own consumption integrated over its duty cycle. When this is violated (very low input, very slow dynamics), freezing a recently optimal set-point or reverting to fractional sampling is preferable.

Photovoltaic harvesters (outdoor and indoor)

A voltage-programmed boost or buck-boost converter that enforces a panel-side set-point decouples the PV from the storage bus and provides a natural handle for MPPT. Outdoors, incremental-conductance or P&O with adaptive steps at 1–5 Hz is sufficient to follow temperature drift and cloud edges without exciting undesirable ripple. Indoors, where spectral composition and irradiance are comparatively stable but absolute power is tiny, fractional- V_{OC} with very sparse sampling (e.g., 0.1–0.5 Hz) is often net-positive relative to more elaborate schemes; the sampling network must exhibit picoampere leakage to avoid biasing V_{OC} on small cells [133], [134].

Commercial power-management ICs provide useful reference designs. Ultralow-power PV front-ends (e.g., TI BQ25570, ADI ADP5091, ST SPV1050) integrate programmable MPPT set-points, V_{OC} sampling timers, and cold-start paths; their data emphasize the practical trade-offs among sampling cadence, input ripple, and quiescent current, which closely mirror the analytical considerations above [138], [139], [140].

Thermoelectric generators

Unlike PV, a TEG's electrical loading perturbs the thermal field through Peltier and Joule effects; the optimal operating point depends on the coupled electro-thermal network rather than a fixed internal resistance. Accurate control therefore benefits from input-current regulation (which sets electrical damping) combined with slow MPPT (periods of tens of seconds) so that measurements reflect the settled thermal state. Analytical and experimental studies show that the MPP occurs at a load ratio that varies with ΔT and heat-sink characteristics, not simply $R_L=R_i$ [141], [142]. In very low ΔT regimes, cold-start is the bottleneck; transformer-based step-ups designed for ~20–30 mV enable self-start and then hand over to a higher-efficiency synchronous path [143]. A PV-style fast hill-climber is inappropriate here—it injects oscillations into the thermal plant and overestimates power by sampling during transients. Instead, use averaged measurements and conservative step sizes, or lock-in techniques that estimate the gradient with tiny, slow dithers [142].

Vibration and triboelectric harvesters

Resonant piezoelectric harvesters resemble high-impedance AC current sources shunted by the piezo capacitance. The rectifier and IPC define the electrical damping seen by the mechanical resonator; “tracking” is therefore equivalent to matching electrical to mechanical damping. Beyond simple diode-bridge + DC–DC, synchronous switch harvesting families (e.g., SSHI) flip the piezo voltage at displacement extrema, boosting extracted charge per cycle and lifting conversion efficiency at low excitation [136], [144]. In practice, the controller adjusts a small set of discrete parameters—switching instants, effective input resistance—to maximize the average rectified current over hundreds of cycles. Because stroke energy is packetized, stroke-synchronous updates (every few to few-dozen cycles) are effective and energy-cheap.

Triboelectric nanogenerators (TENGs) exhibit mega- to giga-ohm source impedances and very high open-circuit voltages. IPC must therefore minimize leakage, use charge-conserving interfaces, and optimize charge-per-stroke rather than instantaneous power. MPPT via classical $v-i$ sensing is rarely meaningful; timing control and surface-charge preservation dominate performance. Reviews of TENG physics and power conditioning emphasize that a small set of discrete interface choices (full-wave vs. Bennet doubler, synchronous charge extraction, buffered charge pumps) cover most operating regimes; trackers need only explore across a slow schedule of those choices [145].

Far-field RF rectennas

At nanowatt-to-microwatt inputs, rectifier efficiency and input impedance are strongly amplitude-dependent. The effective “MPP” is therefore inseparable from impedance matching and waveform design. Two practical levers are (i) adaptive matching networks that hill-climb among a small set of impedances using the rectified DC as a proxy, and (ii) multisine excitations (when the transmitter is controllable) that exploit diode nonlinearity to raise conversion efficiency at low input levels [137]. In ambient scenarios with unknown, fading sources, low-rate (sub-Hz) proxy optimization is sufficient because channel coherence times are long relative to IPC dynamics; faster updates only waste energy. Measurements in urban/office environments quantify realistic RF power densities and demonstrate that rectifier-aware matching dominates yield at the tail of the distribution [111].

4.3.4 Multi-source energy collection architectures

Multi-source energy collection (MSEC) addresses the intermittency and low power density of individual transducers by combining heterogeneous ambient sources—typically light (PV), thermal gradients (TEG), mechanical vibration (piezo/tribo), and RF—into one power-delivery subsystem. Properly architected MSEC can increase average harvested power, extend autonomy during source outages, and improve reliability through source diversity and intelligent power routing (e.g., source selection, time-sharing, or concurrent harvesting). Recent surveys emphasize that benefits accrue only if the electronics decouple each source’s optimum operating point (e.g., MPPT set-points) from the load’s instantaneous needs, while minimizing quiescent losses and cross-regulation between channels [146], [147], [148].

Architectural families

i. Front-end combining with shared conversion

The lightest-weight MSEC class uses passive or low-overhead front-end “combining” (e.g., diode OR, switched networks) to funnel multiple rectified sources into a shared DC-DC stage and storage element. This minimizes silicon area and control complexity but sacrifices independent MPPT and can incur diode/switch conduction losses. For ultra-low-power nodes where every nanowatt matters, quiescent draw in the shared stage often dominates, so simple combiners are preferred only when source impedances and dynamics are similar [147].

ii. Multi-input power converters

Modern MSEC increasingly uses multi-input converter topologies that provide explicit ports per source with coordinated control. Two broad categories are common:

Single-inductor, time-multiplexed interfaces (MISIMO/SI-MISO). These share one inductor across multiple inputs/outputs using fast scheduling. Notable techniques include event-driven MPPT with decoupled load regulation [149], cycle-by-cycle source tracking with adaptive peak-inductor-current control [150], and reversely-polarized energy recycling with automatic source selection for triple-source operation [151]. Such controllers reduce switch/inductor count and silicon area and have demonstrated high peak efficiency across μW – mW regimes in scaled CMOS; however, they require careful arbitration to suppress cross-regulation and preserve each source’s MPPT window [147], [149], [151].

Multi-port (often isolated) converters. Where power levels and voltage ratios are higher—or where bidirectionality to hybrid storage is desired—multi-port boost/bridge topologies dominate. These allow simultaneous energy flow from several sources and to multiple storage elements (battery/supercapacitor), while maintaining port-level control. Recent designs show modular multi-input bidirectional converters for microgrids and EV fast-charging DC links, highlighting trade-offs among device stress, component count, and control complexity [152], [153], [154]. Although aimed at higher power, their principles (port decoupling, concurrent MPPT, prioritization logic) directly inform low-power MSEC [146].

iii. PMIC-centric dual-source managers

Commercial PMICs increasingly expose two dedicated harvester inputs with built-in MPPT and a regulated application rail. Analog Devices’ LTC3331 integrates an energy-harvesting AC/DC + buck front-end with a battery-fed buck-boost, an input prioritizer, and a 10-mA shunt charger, enabling seamless switchover between harvested energy and a rechargeable cell at sub- μA quiescent levels [155]. More recently, e-peas’ AEM13920 adds two independent harvester ports (e.g., PV, TEG, RF, pulsed kinetic) with concurrent energy extraction, cold-start near 275 mV/5 μW , and a regulated buck output—substantially lowering integration effort for dual-source IoT nodes[156].

Control and power-path strategies

Source arbitration and MPPT. The controller determines whether to *time-share* a single inductor among sources (MISIMO) or to *truly* combine sources simultaneously (multi-port). Event-driven scheduling and per-port hysteretic comparators[149] allow fast context switches with minimal cycle overhead, while CCST [150] locks the active port to its instantaneous MPPT. For triple-source systems, automatic mode selection and energy-recycling paths expand usable input ranges[151]. These mechanisms are key to preventing the “tug-of-war” that otherwise occurs when sources with disparate impedances share passives.

Hybrid storage and load support. MSEC often complements a rechargeable battery with a supercapacitor to buffer burst loads and improve cycle life. A recent ultra-low-power CMOS supercapacitor storage unit demonstrates how careful regulation around an unregulated DC-DC output can sustain operation during source-poor intervals and protect

storage elements [157]. In higher-power contexts, bidirectional multi-port converters manage energy exchange between batteries and supercaps while tracking multiple sources [152], [153].

Minimal telemetry for modelling and control. The energy-collection subsystem exposes only three state proxies used later by §4.5 and Chapter 5: (i) a store-state proxy (the storage voltage V_{store} and, where applicable, a temperature-compensated SoC estimate); (ii) a per-source input-power proxy suitable for extremum seeking (e.g., fractional V_{OC} sampling for PV, settled input current for TEG, stroke-synchronous rectified charge for vibration/TENG, and rectified DC level for RF); and (iii) power-good/UVLO events that govern cold-start/brown-out arbitration. Telemetry cadence follows source dynamics (sub-Hz for PV/TEG/RF; stroke-synchronous for vibration/TENG) and is duty-cycled such that average telemetry power remains below $\sim 1\text{--}5\%$ of the median harvested power over the design window. These signals form the contract consumed by the runtime scheduler in Chapter 5.

4.4 Energy storage technologies

This chapter surveys the electrochemical storage options that underpin autonomy and reliability in cyber-physical systems. It focuses on the two workhorse classes—secondary batteries and electrochemical capacitors (supercapacitors)—linking their mechanisms to practical performance, temperature limits, ageing, and safety, and establishing a comparative basis for later design choices [158].

4.4.1 Secondary batteries vs supercapacitors

Electrochemical storage is the constraining element of most cyber-physical systems (CPS): it sets the bounds on autonomy, burst capability, lifetime, and safety. In practice, CPS rely on two technology classes with complementary strengths: secondary batteries, which provide high specific/volumetric energy for sustained operation, and electrochemical capacitors (supercapacitors), which provide low-impedance buffering for short, high-crest-factor events. Section 4.4.1 examines these classes from a mechanistic and performance standpoint—how their physics maps to system-level behavior—before later sections address hybridization and state estimation.




Figure 4.6 Accumulator technologies Power density vs Energy density

For consistency, devices are compared on: (i) energy and power metrics ($\text{Wh} \cdot \text{kg}^{-1}$, $\text{Wh} \cdot \text{L}^{-1}$, $\text{W} \cdot \text{kg}^{-1}$), (ii) impedance/ESR and burst response, (iii) temperature windows and charge constraints, (iv) ageing modes (calendar vs cycle; self-discharge for capacitors), and (v) safety envelopes under abuse. Emphasis is placed on chemistries and formats commonly used in CPS—notably NMC/NCA, LFP, and LTO-anode lithium-ion for batteries, and EDLC/pseudocapacitive devices for capacitors. Subsection 4.4.1.1 details the secondary batteries and their characteristics, while 4.4.1.2 treats supercapacitors; overlap is minimized so the two can be read as complementary parts of a single design framework.

4.4.1.1 Secondary batteries

This subsection surveys rechargeable (secondary) battery technologies that are actually used in cyber-physical systems (CPS)—embedded sensor/actuator nodes, wearables, industrial controllers and gateways—focusing on operating mechanism, practical performance (energy/power density, efficiency), usable temperature windows, ageing behavior and safety. In contemporary CPS, four chemistries dominate in practice: lithium-ion layered-oxide systems (NMC/NCA, typically in pouch/“Li-polymer” or cylindrical formats), lithium-iron-phosphate (LFP), lithium-titanate-oxide (LTO-anode) variants, and nickel–metal hydride (NiMH). Sealed lead–acid (VRLA/AGM) remains common at the infrastructure edge (cabinets/gateways with float charge). Emerging options (e.g., sodium-ion) are noted briefly for context but are not yet the default choice in deployed CPS. In

practice, round-trip efficiency (RTE) is not a constant but varies with current rate, voltage window or state-of-charge, temperature, and ageing, and its architectural impact is coupled to the non-ideal power path. For lithium-ion cells operated at moderate rates and within their nominal thermal window, coulombic efficiency is high, yet the energy RTE is lower because discharge occurs at a reduced average voltage relative to charge; under high-crest-factor workloads the combination of output-side delivery losses and the store's own round-trip penalty sets the minimal pre-burst energy required even when average neutrality holds.

By contrast, electrochemical double-layer capacitors (EDLCs) can be highly efficient over short bursts dominated by ESR, but their effective RTE over application-relevant intervals becomes explicitly time-dependent because idle leakage and charge redistribution reduce the recoverable energy; consequently, deep-sleep autonomy is leakage-limited rather than cycle-wear-limited. Any comparative statement must therefore specify both the measurement boundary (DC-level storage versus bus-level conversion included) and the cycle's time structure to avoid conflating storage physics with converter housekeeping and to ensure that sizing and scheduling decisions reflect the true macroscopic losses used elsewhere in the model.

Lithium-ion layered-oxide systems

Graphite and layered-oxide couples, notably $\text{LiNi}_x\text{Mn}_y\text{Co}_z\text{O}_2$ (NMC) and LiNiCoAlO_2 (NCA), dominate compact CPS (wearables, portable nodes) because they offer high gravimetric/volumetric energy and are available in thin pouches ("Li-polymer") and cylindrical 18650/21700 cells. The faradaic intercalation mechanism provides cell-level $\approx 150\text{--}260 \text{ Wh kg}^{-1}$ with chemistry- and loading-dependent specific power; the stable operating window is typically 2.7–4.2 V per cell [159], [160], [161].

In CPS, these cells are usually charged CC–CV to 4.20 V (sometimes 4.10–4.15 V to extend life), with end-of-charge taper cut-offs of $C/10\text{--}C/20$. Internal resistance and kinetic polarization rise with C-rate and at low temperature, so high-crest-factor bursts (radios, motors) are often current-limited or buffered (by decoupling capacitors or, where appropriate, supercapacitors—treated later in § 4.4.1.2). Round-trip coulombic efficiency is high at moderate rates, but even 0.1 % inefficiency integrates to significant loss over long sleep horizons, which matters for maintenance intervals [161].

Calendar ageing accelerates at high SoC and high temperature via SEI/CEI growth and electrolyte oxidation; cycle ageing intensifies with deep/high-rate cycling and at low-temperature charge where lithium plating can occur. Storage near mid-SoC and charge-voltage ceilings below 4.2 V are widely used levers to extend life (Keil & Jossen, 2017; Birkel et al., 2017; Attia et al., 2020). Low-temperature operation (sub-zero) depresses power and usable capacity; charge at ≤ 0 °C is often prohibited unless pre-heated [162].

Ni-rich layered oxides achieve high energy but narrow abuse margins at high SoC; meticulous protection (cell-level guardianship ICs, pack fusing, thermal design) is standard. Comparative tests show higher thermal-runaway onset and lower peak pressures for LFP than for NMC; layered oxides therefore demand conservative limits [160], [163].

Where compactness and autonomy dominate (wearables, mobile nodes), NMC/NCA are the default due to energy density and ubiquitous supply chains. Designers trade a fraction of nameplate energy for life by lowering the charge ceiling and by smoothing bursts.

Lithium-ion layered-oxide systems

LiFePO₄ couples a graphite (or occasionally hard-carbon) anode to an olivine phosphate cathode. The flat 3.2–3.3 V plateau, excellent thermal stability, and long cycle life make LFP a workhorse for industrial/outdoor CPS where safety and life outrank maximal energy [160], [161].

Typical cell-level energy density is ≈ 90 – 160 Wh kg⁻¹, lower than Ni-rich layered oxides, but power capability is strong, and the plateau simplifies regulation. Standard charge limits at 3.60–3.65 V per cell reduce oxidative stress; discharge cut-offs around 2.5–2.8 V preserve life.

LFP is comparatively tolerant of elevated temperature and overcharge; the cathode's benign oxygen chemistry reduces thermal-runaway propensity. Low-temperature power is still limited by graphite kinetics; cycle life can exceed 2000–5000 cycles at moderate depth of discharge with appropriate thermal management [160], [161].

In fielded industrial nodes/gateways (smart metering, roadside cabinets, remote telemetry) where abuse tolerance, predictable voltage, and long life trump the last joule of energy density, LFP is often preferred.

Lithium-titanate-oxide (LTO-anode) lithium-ion

Replacing graphite with $\text{Li}_4\text{Ti}_5\text{O}_{12}$ (LTO) raises anode potential to ≈ 1.55 V vs Li/Li^+ , eliminating most lithium-plating risk and enabling “zero-strain” insertion with minimal SEI growth. The result is exceptional cycle life, fast-charge acceptance, and low-temperature power, at the cost of lower energy density (cell nominal ≈ 2.3 – 2.4 V with LFP or LMO cathodes) [159], [161].

LTO cells tolerate high C-rates and maintain usable power down to -20 °C or below, where graphite-anode LIBs struggle. Practical energy density is often 50 – 90 Wh kg^{-1} , but cycle life can reach >10 k cycles at moderate depth of discharge. Elevated anode potential and robust interphases yield excellent abuse tolerance.

In harsh-environment CPS (cold climates, high-rate pulsed actuation, frequent charge-discharge from intermittent harvesters) where service life and cold-start are decisive, LTO is a strong fit despite lower energy.

Nickel-metal hydride

NiMH pairs a NiOOH positive with a hydrogen-absorbing alloy negative in an aqueous KOH electrolyte. Gravimetric energy (≈ 60 – 120 Wh kg^{-1} at cell level) trails LIB, but NiMH offers robust safety, good low-temperature discharge power, and benign failure modes (no flammable organic electrolyte) [164].

Self-discharge is higher than LIB, though “low-self-discharge” (LSD) variants mitigate it. Charge management commonly uses $-\Delta V$, $\Delta T/\text{dt}$, and temperature caps rather than CC–CV to a defined potential. Cycle life of 500 – 1000 cycles is typical under moderate depth of discharge.

NiMH persists in safety-critical or cost-sensitive CPS where aqueous chemistry and simple charging are desirable, and in modular packs require 1.2 V cells for straightforward series stacking.

Sealed lead-acid (VRLA/AGM)

VRLA (AGM/gel) immobilizes the sulfuric acid electrolyte and recombinants gases, enabling sealed operation. Although gravimetric energy is low (≈ 30 – 50 Wh kg^{-1}), VRLA

remains common in infrastructure CPS (edge servers, cabinets, telecom gateways) where float charging, surge currents, and low cost dominate [165], [166].

Life is sensitive to temperature (grid corrosion) and depth of discharge (sulfation); AGM supports better high-rate performance; gel variants improve deep-discharge tolerance. Mature recycling is a sustainability advantage.

For stationary CPS power buffers/UPS with float service, VRLA provides predictable behavior and lowest installed cost per Wh, with straightforward logistics.

Sodium-Ion

Sodium-ion batteries preserve the “rocking-chair” architecture of Li-ion while substituting Na^+ carriers to improve materials security and cost resilience. Techno-economic analysis indicates that—despite lower gravimetric energy than state-of-the-art Li-ion—SIBs can be cost-competitive in stationary and other volume-tolerant applications under plausible learning curves and commodity-price trajectories, which aligns well with CPS roles such as edge gateways, cabinets and UPS-backed controllers [167].

On the anode side, hard carbon (HC) remains the workhorse. Recent synthesis connects HC microstructure (closed-pore fraction, micro/mesoporosity, precursor chemistry) to initial Coulombic efficiency, plateau capacity and rate capability, providing design rules (e.g., heteroatom doping, pore-architecture control) that are now delivering application-level gains [168]. Electrolyte engineering (solvent/salt blends and SEI-forming additives) co-optimized with HC has materially improved sub-zero kinetics, shrinking SIB’s historic cold-temperature penalty—salient for outdoor CPS [169].

For cathodes, two families dominate. Prussian blue/white analogues (PBA/PW) combine open 3D frameworks with manufacturability; current work targets vacancy and water management to stabilize capacity and enable fast charging [170]. In parallel, layered oxides (P2/O3, including high-entropy and anion-redox designs) are advancing through compositional and structural tuning to mitigate phase transitions and oxygen-redox instabilities, pushing SIB performance toward LFP-class use cases [171].

Industrialization has accelerated and now spans grid storage, mobility pilots and manufacturing build-outs. In Qianjiang, Hubei, the first 50 MW/100 MWh SIB BESS entered service in July 2024, evidencing bankable stack integration [172]. CATL launched

the Naxtra sodium-ion brand with mass production slated for December 2025, signaling tier-1 supply readiness [173]. In mobility, Yiwei/JAC began serial production (Jan 2024) and exports (Feb 2024) of sodium-ion EVs—modest volumes, but proof of pack-level BMS/thermal feasibility [174]. In Europe, Northvolt + Altris validated $\sim 160 \text{ Wh kg}^{-1}$ SIB cells for ESS (Nov 2023), while in North America Natron Energy announced a US\\$1.4 bn, $\sim 24 \text{ GWh yr}^{-1}$ plant in North Carolina targeting data-center/industrial power niches [175], [176]. Collectively, these signals reduce adoption risk for CPS in stationary/edge roles and foreshadow broader availability across prismatic (and emerging cylindrical) formats.

Quantum batteries

Quantum batteries are prototype devices that store energy in quantum states (e.g., excitations of molecules, spins or superconducting circuits) and may exhibit collective quantum effects—such as entanglement and superabsorption—that in principle allow faster charging power than classical analogues. Recent reviews and proof-of-concept experiments (e.g., superabsorption in organic microcavities) show rapid theoretical progress and first demonstrations, but present platforms require specialized photonic/cavity or cryogenic hardware, suffer decoherence-limited retention, and are not electrochemical drop-ins for CPS. Consequently, they are out of scope for deployed CPS design; at most, they merit a brief outlook note rather than inclusion alongside secondary batteries or supercapacitors[177], [178], [179].

Table 4.1 Battery characteristics per technology

Chemistry	Nominal voltage (Volt)	Voltage range (Volt)	Gravimetric energy (Wh kg^{-1})	Cycle life to 80%	Operating temperature ($^{\circ}\text{C}$)
Li-Po	3.6-3.7	2.7-4.2	150-260	500-1500	-20+60
LiFePO ₄	3.2-3.3	2.5-3.65	90-160	2000-5000	-20+60
LTO	2.3-2.4	1.8-2.8	50-90	5000-10000	-30+60
NiMH	1.2	1.0-1.5	60-120	500-1000	-20+60
VRLA/AGM	2	1.75-2.45	30-50	300-700	-20+50
SIB	3	2.0-3.8	90-160	1000-3000	-20+60

4.4.1.2 Supercapacitors

Supercapacitors employed in CPS fall into two principal categories, distinguished by their charge-storage mechanism and, consequently, by voltage limits, leakage behavior, ageing modes and frequency response. Presenting them separately avoids conflating purely electrostatic storage with surface-redox contributions.

Electrochemical double-layer capacitors (EDLC)

EDLCs store charge electrostatically by forming an electric double layer at the interface between an electronically conductive porous electrode (typically activated or carbide-derived carbon) and an ionically conductive electrolyte [180]. Capacitance scales with accessible surface area and pore architecture; equivalent series resistance (ESR) is set by ionic transport in the electrolyte, electronic pathways in the carbon/tabs, and interconnects. Organic electrolytes (acetonitrile or propylene carbonate based) provide $\sim 2.7\text{--}3.0$ V per cell and thus dominate compact CPS; aqueous systems are limited to ~ 1 V per cell but offer excellent low-temperature behavior and benign safety; ionic liquids extend the voltage window (≥ 3.2 V per cell) and high-temperature stability at the cost of low-temperature transport [181], [182].

Over an operating window $V_{\min} \rightarrow V_{\max}$, stored energy is $E = \frac{1}{2} C (V_{\max}^2 - V_{\min}^2)$, so state of energy is quadratic in voltage; under a step current, terminal voltage exhibits an ohmic jump $I R_{\text{tot}}$ followed by a linear ramp I/C . In the frequency domain, EDLCs show near-constant-phase behavior across decades before rolling off at a corner set by ion transport in the smallest pores; compact distributed-RC or fractional-order models reproduce this cheaply in embedded code [182].

Because storage is electrostatic, cycle endurance is exceptional ($10^5\text{--}10^6$ shallow cycles are routine). Reliability is dominated by self-discharge (open-circuit voltage decay) and slow drifts in ESR/effective capacitance. OCV decay is a superposition of charge redistribution in hierarchical pores (multi-/stretched-exponential kinetics) and activation-controlled leakage through side reactions/impurities, with clear Arrhenius temperature acceleration [183]. Engineering strategies such as ion-transport confinement in gel/solid-state electrolytes demonstrably suppress leakage and extend hold-time, trading some rate capability [184].

Sub-nanometer pores commensurate with ion size can boost capacitance via partial dissolution and confinement; simulations and in-situ studies link high capacitance to these

pore-level effects, underscoring why pore-size distribution matters as much as total surface area for CPS parts targeting minimal volume [182], [185].




Figure 4.7 Electric Double Layer Capacitor

Pseudocapacitors

Pseudocapacitors add fast, surface-confined faradaic reactions (or shallow intercalation) that remain capacitive on practical time-scales, increasing areal capacitance while preserving low impedance. Representative materials include MnO_2 , Nb_2O_5 , conducting polymers, and MXenes (2D carbides/nitrides such as $\text{Ti}_3\text{C}_2\text{T}_x$). While devices are often paired with the same electrolytes as EDLCs, their differential capacitance $C_d(V)$ can vary with potential as redox states are accessed; in MXenes, finite electronic density of states introduces quantum-capacitance limitations near the Fermi level [186], [187], [188], [189].

The quadratic energy relation still holds with the effective $C(V, T)$, but voltage-dependent C_d means hold-time prediction in CPS benefits from characterizing C over the intended window and temperature. Frequency response can remain quasi-capacitive over a broader band than EDLCs due to rapid surface redox, which is advantageous for burst spectral content extending into higher frequencies.

Compared with EDLCs, pseudocapacitive electrodes introduce material-specific degradation (e.g., dissolution/structural change for MnO_2 , volumetric strain for polymers,

surface termination evolution for MXenes). Device designs mitigate these via coatings, conductive scaffolds and electrolyte optimization; nevertheless, leakage and parameter drift remain the primary constraints on long idle horizons in CPS [186], [187].

In both categories, C and ESR must be specified together because transient droop scales with both; converter efficiency vs input voltage should be characterized across $V_{\min} \rightarrow V_{\max}$; and, in series stacks, balancing is essential because leakage/capacitance dispersion drives unequal cell voltages and over-voltage can trigger electrolyte decomposition and gas evolution. Empirical studies on micropower energy-harvesting supplies confirm that, when operated within a constrained window, supercapacitors deliver high burst efficiency with minimal conversion overhead, while leakage dictates deep-sleep autonomy [182], [190].

4.4.2 Hybrids and selection criteria

Hybridizing a secondary battery with a supercapacitor lets cyber-physical systems (CPS) decouple average-energy supply from burst-power delivery, lowering battery RMS and peak current (hence temperature and ageing drivers) while preserving high crest-factor capability and efficient energy-harvesting interfaces [182], [191]. In practice the battery services the low-frequency (mean) component of the load and sets autonomy; the supercapacitor services the high-frequency transients and accepts intermittent inflows with minimal loss, so battery C-rate, overpotential and heat generation during bursts are suppressed [190], [191].

Three interconnection patterns occur. Passive hybrids place a supercapacitor across the DC bus: they are simple and reliable, but the battery still sees some high-frequency current and energy use in the capacitor is not actively scheduled. Semi-active hybrids insert a bidirectional DC–DC converter on one branch (commonly the supercapacitor), enabling power-split control, pre-charge, and bus regulation over a wide capacitor voltage swing. Fully active hybrids use converters on both branches to optimize battery stress and conversion efficiency at the cost of complexity and quiescent overhead [191], [192]. These topologies—well studied at traction scale—translate cleanly to embedded CPS provided converter quiescent current and supercapacitor leakage are explicitly budgeted [182], [192].

Sizing follows the load spectrum. The supercapacitor stack is dimensioned from the burst; the battery from the mean energy and permitted C-rate. For a burst energy ΔE delivered while the capacitor swings from V_1 to V_2 , the required capacitance is

$$C_{\text{req}} = \frac{2\Delta E}{V_1^2 - V_2^2} \quad (4.40)$$

and the ESR-induced step droop is $\Delta V_{\text{ESR}} = I_{pk} R_{\text{tot}}$. For near-constant-power bursts P of duration τ , the terminal voltage evolves as $V(\tau) \approx \sqrt{V_1^2 - 2P\tau/C}$, from which admissible duty cycles follow for a given droop budget [182]. Series cell count is set by the bus requirement and per-cell ceiling V_{max} : $N_s = \lceil \alpha V_{\text{max,bus}}/V_{\text{max,cell}} \rceil$ with a derating $\alpha > 1$ for balancing tolerance and ageing. Once the burst is assigned to the supercapacitor, the battery can be selected at a lower design C-rate, reducing heat and ageing [191].

Energy-management strategies range from frequency-shaping (low-pass the battery current; high-pass to the capacitor) and rule-based peak-shaving (reservoir thresholds with hysteresis) to optimal control (Pontryagin/MPC) that trades converter loss against battery degradation. At CPS power levels, rule-based schemes often suffice when converter efficiency vs input voltage is characterized across the full capacitor window, while structured loads (e.g., periodic radio TX windows) benefit from predictive pre-charge of the capacitor to minimize bus droop and idle loss [191], [192].

A practical constraint for embedded nodes is that deep-sleep autonomy is leakage-limited rather than cycle-wear-limited. Supercapacitor OCV decay comprises charge-redistribution in hierarchical pores plus activation-controlled leakage with clear Arrhenius temperature acceleration; both must be included in hold-time budgets along with converter quiescent current [182], [183]. Empirical micropower studies show that, inside a constrained voltage window with low ESR, the burst path is highly efficient and battery stress is reduced, but the idle horizon is ultimately set by leakage and quiescent drains—hence the importance of electrolyte choice (organic vs aqueous vs ionic liquid) and, where suitable, gel/solid-state confinement to suppress leakage [184], [190].

When series-stacking capacitors to reach bus voltage, cell-to-cell dispersion in leakage and capacitance leads to unequal voltage sharing; balancing (passive for simplicity or active for efficiency) is therefore mandatory to prevent over-voltage decomposition, gas evolution and swelling [181], [182]. Converter interaction also matters: wide input swings require characterizing efficiency vs input voltage across $V_{\text{min}} \rightarrow V_{\text{max}}$, and undervoltage thresholds should reflect the quadratic SoE of capacitors to avoid marooning useful energy at the low end [182].

An alternative to explicit two-device hybrids is the lithium-ion capacitor (LIC), which pairs a battery-like anode (often pre-lithiated carbon) with an EDLC-type cathode. LICs extend energy beyond EDLCs while retaining low impedance and fast charge acceptance, simplifying boards where moderate energy and high bursts are required; current reviews cover pre-lithiation methods, electrode choices and performance envelopes [193], [194].

4.4.3 State-of-charge (SoC) estimation

This subsection gives a concise, implementation-oriented account of SoC (and, for capacitors, SoE) estimation suited to cyber-physical systems (CPS) with duty-cycled workloads, modest microcontroller resources, and long idle intervals. We separate (i) rechargeable batteries—where the key challenge is tracking charge and slowly drifting model parameters—and (ii) supercapacitors, where energy is quadratic in voltage and long-idle behavior is set by redistribution and leakage. A short note maps these methods to fuel-gauge PMICs commonly used in real designs.

Batteries: practical SoC estimation with light models

A robust baseline for single-cell Li-ion in CPS is coulomb counting corrected by an OCV model: integrate current to track charge, then re-anchor to a temperature-indexed open-circuit-voltage curve during quiescent intervals to arrest drift from sensor offset and clock error. This baseline is widely recommended in the BMS literature because it is computationally light yet markedly improves long-idle accuracy when compared with pure coulomb counting, provided the OCV(SoC, T) map is accurate and voltage plateaus (e.g., LFP) are handled carefully. [195], [196], [197]

For bursty loads and varying temperature, a 1-RC equivalent-circuit model (ECM) with an extended/unscented Kalman filter (EKF/UKF) materially improves fidelity at MCU-friendly cost. The filter estimates SoC while adapting ohmic resistance R_0 and a polarization branch; adding slow states for usable capacity Q_{use} (SoH) and R_0 captures ageing. Scheduling small identification pulses (tens to hundreds of milliseconds) during idle windows stabilizes parameter tracking. Numerous studies document EKF/UKF performance across chemistries and profiles, and recent reviews synthesize best practice and pitfalls (noise tuning, observability on flat OCV regions, temperature coupling). [196], [197], [198], [199], [200]

For resource-constrained CPS and energy harvesting nodes, simple yet well-calibrated voltage-based SoC approximations can be competitive when paired with temperature

awareness and empirically derived curves. Buchli’s work in WSNs shows that lightweight OCV-based estimators—augmented by occasional current/temperature—achieve >95 % accuracy under realistic workloads without heavy models, and his doctoral thesis places such estimators in an end-to-end energy-neutral design framework [201], [202]. These results justify a tiered approach: start simple; add ECM+EKF only when bursts/temperature dynamics demand it.

Minimal recipe (batteries). (1) High-resolution voltage, current, temperature sensing with bias calibration; (2) coulomb counter + temperature-indexed OCV re-anchoring after defined rests; (3) optional EKF/UKF on a 1-RC ECM for bursty profiles; (4) slow states for $Q_{\text{usc}}(t)$ and $R_0(t)$ to follow ageing; (5) explicit charge derating below 0 °C to avoid plating artefacts that corrupt the model. [196], [197]

Supercapacitors: SoE-first with leakage-aware tracking

For EDLC-type supercapacitors operated between $V_{\min}$ and $V_{\max}$,

$$E = \frac{1}{2} C (V_{\max}^2 - V_{\min}^2) \quad (4.41)$$

So, state-of-energy (SoE)—not SoC—is the natural quantity; it is quadratic in voltage. In practice, effective capacitance $C(V,T)$ and ESR $R(T)$ should be identified online, and long-idle behavior must include charge redistribution (multi-exponential) and activation-controlled leakage (Arrhenius). Lightweight observers combine: (i) small scheduled pulses to estimate C and ESR; (ii) idle-period fitting of $V(t)$ to update leakage parameters, yielding accurate hold-time predictions for deep sleep. When leakage dominates autonomy, gel/solid-state electrolytes help at some cost to low-temperature transport. [182], [183], [184]

Where capacitor dynamics are non-ideal, Kalman-family observers on distributed-RC models have been demonstrated to track internal states and SoE with good accuracy over wide operating windows, reinforcing that estimator complexity can remain modest while capturing key physics.[203], [204]

Minimal recipe (supercapacitors). (1) Compute SoE from voltage; (2) identify C /ESR via brief, safe excitations; (3) maintain a two-time-scale leakage model from sleep telemetry to forecast hold-time; (4) if series-stacked, include cell-level monitoring/balancing to avoid over-voltage divergence during charge or thermal excursions.[182], [183]

Fuel-gauge PMICs: mapping methods to practice

It is appropriate to cite representative fuel-gauge ICs to show how these estimators are deployed in silicon:

- Analog Devices MAX17260 (ModelGauge™ m5 EZ): single-cell Li-ion gauge combining coulomb counting with a dynamic battery model and temperature compensation; designed for very low quiescent current and minimal host configuration while reporting SoC/remaining time/SoH. [205]
- Texas Instruments BQ27441-G1 (system-side Impedance Track™): adapts an internal model to ageing and temperature, providing SoC/SoH with low host burden and clear integration guidance in the TRM. [206]

Mentioning such parts (briefly) grounds the section: both families implement the “coulomb-count + model correction” paradigm described above, with embedded parameter adaptation to mitigate drift and ageing.

4.4.4 Battery management system requirements

This subsection consolidates practical design rules for integrating secondary batteries and supercapacitors in cyber-physical systems (CPS), with emphasis on protection, thermal/voltage derating, balancing, leakage budgeting, and verification. The aim is to translate the mechanisms and estimation methods of §§4.4.1–4.4.3 into implementable requirements that meet safety and compliance constraints while preserving autonomy and burst capability.

To support verification and field diagnostics, the battery-management and supervisory firmware should accumulate direction-tagged charge and discharge energies, $\int v(t) i(t) dt$, at the DC storage boundary and co-log temperature and the idle interval between directions. Pairing a charge packet with its subsequent discharge after a dwell of duration τ yields an in-situ estimate of the effective round-trip efficiency and the dwell leakage; these quantities furnish acceptance checks against the chapter’s windowed neutrality and survivability criteria and expose impedance growth as a systematic RTE drift under fixed operating conditions. The required instrumentation is already implied by the platform interface—a single store proxy sampled at coarse cadence together with hysteretic protection events—and can be implemented with negligible overhead, because measurements occur at the cadence used by the epoch-based runtime (DREAM) rather than continuously; the same

logs are consumed by the scheduler to keep field behavior aligned with the macroscopic accounting used for design.

The first principle is operating-point derating matched to chemistry. For lithium-ion batteries, sustained operation near the upper cut-off voltage accelerates SEI/CEI growth and electrolyte oxidation, so it is common to cap charge at 4.10–4.15 V for layered-oxide cells and 3.60–3.65 V for LiFePO₄, trading a few percent of energy for reduced calendar kinetics and thermal stress [160], [161]. Low-temperature charge should be curtailed or preceded by pre-heat because graphite-anode systems approach the plating regime below 0 °C, which both degrades the cell and corrupts SoC inference; LTO-anode variants are more tolerant but still merit current derating (Hu *et al.*, 2019). For supercapacitors, voltage derating per cell is the primary life lever: running organic-electrolyte EDLCs below their nominal ceiling reduces the rate of over-voltage-triggered side reactions and gas evolution; temperature derating further limits leakage growth, which otherwise dominates deep-sleep autonomy [182].

Protection architectures should enforce over-voltage/under-voltage, over-current and over-temperature limits in hardware. Single-cell Li-ion packs typically combine a dedicated protector (or gauge-with-protections) with a host-side supervisor that honors charge/discharge inhibit signals and derates set-points with temperature. In multi-cell stacks (more typical for gateways and cabinets), cell balancing prevents divergence: passive shunts are simple and predictable for long-life, low-C-rate systems, whereas active balancing can recover energy at the cost of control complexity. The same principle extends to series-stacked EDLCs, where dispersion in leakage and capacitance yields unequal voltage sharing; balancing is not optional—without it, one cell will over-voltage under charge or at elevated temperature. Because EDLC hold-time is intrinsically leakage-limited, balancing networks and supervisory electronics must be included in the quiescent-drain budget [182], [184].

Thermal design couples directly to safety and lifetime. Ni-rich layered-oxide cells deliver high energy density but exhibit narrower abuse margins at high state of charge; comparative sealed-enclosure experiments show earlier runaway and higher pressures for NMC than for LFP under equivalent stimuli, reflecting the latter’s benign oxygen chemistry [160], [163]. Practical CPS enclosures should therefore provide localized thermal paths away from heat-sensitive components, avoid insulating foams around cells or capacitors, and provide venting or pressure relief where applicable. For supercapacitors, although thermal-runaway

mechanisms of lithium-ion do not apply, over-voltage and high-temperature storage still cause gas generation, dry-out and ESR rise; mechanical allowances for swelling and controlled ambient exposure are prudent [182].

A verification plan suitable for CPS should combine electro-thermal characterization, burst-response testing, and long-idle leakage/redistribution profiling. Electro-thermal characterization measures OCV–SoC maps, R_{0R_0R0} and one-RC polarization time constants (for batteries) and CCC/ESR vs temperature (for EDLCs), preferably in a controlled chamber across the intended range. Burst-response tests replay real duty cycles (sensor warm-up, RF TX/RX windows, actuator inrush) and confirm that bus droop and thermal rise respect budgets; hybrids should show reduced battery RMS and peak currents when the supercapacitor branch is active [191]. Long-idle profiling measures open-circuit decay at two or more temperatures over hours–days; fitting a two-time-scale model (short-term redistribution plus temperature-activated leakage) yields predictive hold-time parameters for deployment [183], [184]. These validated models should then be embedded into the SoC/SoE estimator (§4.4.3) so that supervisory decisions (pre-charge, derating, duty-cycle scheduling) use measured device behavior rather than generic assumptions.

Finally, component choices should be estimator-aware. Battery fuel-gauge PMICs that implement coulomb-count + model correction (e.g., MAX17260 ModelGauge m5 EZ; TI bq27441-G1 Impedance Track) reduce firmware burden and provide temperature-aware SoC and remaining-time estimates at very low idle current; however, their configuration must reflect the measured OCV map and capacity at the chosen charge ceiling, and their quiescent draw must be included in idle budgets. For supercapacitors, the supervisory MCU can maintain SoE from voltage while periodically re-identifying CCC/ESR via safe micro-pulses and updating leakage parameters from sleep telemetry; if series-stacked, simple per-cell monitors (divider + MUX) greatly improve observability and balancing effectiveness.

4.5 Power distribution and management strategies

Energy-autonomous microsystems must convert irregular, low-level harvested inputs into stable, application-grade supply rails with minimal overhead. Building on the loss-aware architecture of §§4.2–4.4, this section formalizes how the power path is realized in practice (source selection, protection, start-up, and delivery to the store) and how regulation and voltage domains are organized for mixed digital/analog/radio loads. The emphasis is on

quiescent-power budgets, survivability under brown-out, and rail sequencing that preserves both energy efficiency and signal integrity.

Recent advances enable multi-rail delivery at ultra-low power without duplicating magnetics, notably single-inductor multiple-output (SIMO) and single-input/multi-output variants with control schemes that suppress cross-regulation and support deterministic power-up sequencing—features that map directly to the availability and survivability metrics defined earlier. These architectures will be referenced for selecting regulators and scheduling rail enables in energy harvesting nodes [147], [207].

On the load side, domain partitioning is paired with dynamic techniques (power gating, DVFS, and duty-cycling) to minimize both static and dynamic consumption while respecting analog/RF noise constraints and protection hysteresis; this section states implementable rules so that runtime (cf. Chapter 5, DREAM) can actuate them with negligible telemetry overhead [208], [209].

Finally, because many harvesters operate near converter thresholds, we summarize cold-start and UVLO/OVP strategies that guarantee controlled bring-up from a discharged store and safe fallback during energy poverty, along with selection guidance for EH-oriented PMICs. The discussion is grounded in recent reviews of IoT energy harvesting power paths and cold-start circuits [210], [211].

4.5.1 Power path

The power path mediates energy between heterogeneous sources, the storage element, and regulated rails while enforcing cold-start and brown-out constraints with negligible overhead. Behaviors are formalized as event gates around the system state (enable/disable governed by explicit inequalities on store energy/voltage and hysteretic protection thresholds), and the runtime consumes a minimal contract—one store proxy and per-rail power-good/UVLO signals—to avoid measurement overhead and policy oscillations.

At the μW – mW scale, passive diode OR-ing is inefficient; MOSFET-based ideal-diode and power-multiplexer controllers reduce forward drop to tens of millivolts, provide active reverse-current blocking, and implement deterministic break-before-make source selection, which directly supports survivability and availability under source switchover. These controllers are therefore preferred for source OR-ing or prioritization ahead of downstream regulation [212], [213], [214], [215], [216], [217].

Cold start is treated as a boundary condition: until the stored energy and voltage exceed start-up requirements, only a passive, current-limited charging path is permitted; supervisory logic and the first regulation stage draw from a small start-up buffer, and the main conversion is enabled only once explicit inequalities are satisfied with hysteresis to prevent chatter. Recent reviews emphasize the centrality of cold-start circuits and tabulate strategies and budgets for ultra-low-power harvesters [210]. Concrete implementations that embed this behavior appear in monolithic PMUs for indoor PV—with start-up, OVP and output switching integrated [15]—and in system-level PMCS designs for solar-assisted embedded vision, where a dedicated input power switch and low-power sleep path realize controlled bring-up [218].

Undervoltage lockout and over-voltage protection require explicit hysteresis so that a lower trip disconnects loads and disables nonessential conversion, and an upper re-arm authorizes recovery only after the store replenishes. Application notes quantify how divider networks and comparator hysteresis avoid bounce under finite source impedance, and how soft-start limits inrush to protect the cold-start budget [19], [219].

Finally, because multi-rail embedded nodes demand deterministic rail sequencing and minimal cross-regulation, the regulation stage coupled to the power path should support ordered power-up and low quiescent draw; recent SIMO architectures explicitly implement power-up sequencing for IoT multi-rail loads [207], aligning the regulator’s behavior with the power-path enable logic described above.

4.5.2 Regulation and voltage domains

Regulation and voltage domains in a CPS must deliver application-grade rails with deterministic behavior under the event gates defined earlier, because control loops, sensing chains, and radio links impose hard bounds on start-up ordering, rail droop, phase noise, and timing jitter. In practice this means enables are slaved to hysteretic protection thresholds from the power path, while the runtime (DREAM) only observes a store proxy and per-rail power-good/UVLO events; all other policy remains local so that deadlines and quality-of-service (QoS) guarantees are preserved even under energy scarcity. Recent CPS-oriented studies link energy management to real-time guarantees and reliability, reinforcing this separation of concerns [220].

Partitioning follows noise sensitivity and task criticality. Digital compute and radio baseband typically use high-efficiency switching conversion, whereas precision analog front-ends, PLL/VCOs, and RX paths require rails with high power-supply rejection and low integrated noise; modern guidance and measurements show that correctly-compensated LDOs provide broadband PSRR when paired with low-parasitic capacitors and verified on dedicated benches. These choices localize conducted and radiated noise away from time-critical sensing and estimation tasks in CPS [221], [222].

Rail sequencing must be explicit so that actuators and radios never see back-bias or latch-up during bring-up; processor and memory vendors prescribe order, ramp rates, and inter-rail delays, and platform notes show that improper sequencing causes inrush, reliability degradation, and timing faults—failure modes that manifest as missed deadlines in closed-loop CPS. Practical patterns use regulator enable chains and power-good interlocks, or PMIC sequencers, to generate deterministic edges consumed by the runtime [223].

Quiescent current dominates deep-sleep autonomy in event-driven CPS. Always-on domains should collapse to one housekeeping rail and use ultra-low-IQ buck conversion; recent devices such as the TPS62840 achieve ~60 nA IQ (≈ 120 nA in 100 % mode) while maintaining clean switching for radios when preceded or followed by modest filtering or a short post-LDO. This allows sensors and time-based to remain alive without compromising availability metrics [224].

Sharing magnetics through single-inductor multi-output regulation reduces size and idle loss for multi-rail nodes, provided cross-regulation is controlled and power-up is deterministic. Contemporary SIMO designs for IoT introduce constant charge-transfer and built-in sequencing, while cross-regulation analyses quantify coupling mechanisms and mitigation in SISO/SIMO topologies—properties directly mapped to CPS survivability because they bound rail interference during task preemption and radio bursts. Where an RF rail demands tighter noise than SIMO can deliver, a dedicated converter with a post-LDO remains appropriate [207], [225].

Dynamic techniques operate within domains under real-time constraints. DVFS reduces dynamic power roughly with V^2f and is effective when transitions are scheduled at phase boundaries of the control workload; recent embedded/CPS literature reports significant energy savings when DVFS is coordinated with timing analysis, while also noting reliability

and scheduling implications that justify keeping protection events outside of software control. For intermittently powered devices, DVFS policies must account for buffer voltage and convert efficiency, a point highlighted in recent work on energy-harvesting CPS [208], [220], [226].

For noise-critical rails, a measurement-driven recipe is preferred: choose regulators with documented high-frequency PSRR and low output-noise density; verify with vendor test setups; and, if a switching precursor is used, adopt low-noise buck stages (e.g., Silent Switcher® families or low-ripple architectures) plus gentle slew/soft-start so that bring-up does not inject broadband energy into sensing bands. These practices preserve estimation accuracy and RF sensitivity in CPS while avoiding static losses associated with stacked LDOs [227], [228].

It is useful to anchor the partitioning to common rail levels in contemporary platforms. FPGA/processor cores operate near 0.9–1.2 V (e.g., Artix-7 VCCINT $\approx$ 0.9–1.0 V with explicit sequencing constraints), auxiliary rails around 1.8 V, and I/O banks select among LVCMOS standards at 1.2/1.5/1.8/2.5/3.3 V as per JESD8; DDR5 memories standardize VDD = VDDQ = 1.1 V with a 1.8 V pump rail (VPP); mainstream MCUs remain single-supply 1.8–3.6 V. In Wi-Fi/BT SoCs such as the ESP32 family, a 3.3 V external rail typically feeds on-chip regulators that derive $\sim$ 1.1 V core and segregated always-on domains for RTC and wake logic; this naturally suggests a “3.3 V system + local core generation” pattern with strict brown-out thresholds and RTC retention during deep sleep. Ultra-low-power MCUs such as TI’s MSP430 are natively single-supply ($\approx$ 1.8–3.6 V) with internal core regulation and supply-voltage supervisors; frequency limits track VDD, so DVFS reduces to discrete operating-point choices constrained by that curve. Across these devices, the prevailing levels motivate a practical topology: an efficient buck (or SIMO stage) to generate 3.3 V and 1.8 V trunks, selective post-LDOs for noise-critical analog/RF rails, and explicit sequencing so core-domain rails precede or track I/O and memory rails. This mapping keeps CPS sensing, control, and radio chains within their noise and timing budgets while minimizing static loss [229], [230], [231], [232], [233], [234], [235].

The resulting regulator set, and domain plan expose to software only the small interface already defined: a coarse-cadence store proxy and per-rail events for enable/UVLO. With this contract, DREAM can schedule domain enables, power-gating, and DVFS points in

constant time while protection gates retain priority, ensuring that control-loop deadlines and CPS QoS are satisfied even when harvested power fluctuates.

The conversion model provides deliverable energy trajectories and constraint sets (start-up, ramp, storage limits). Chapter 5 embeds these into DREAM, which schedules rest phases against E_T , $A_{\phi_{\min}}$ and R_{θ} to meet autonomy targets with stated confidence and derives minimal storage and buffer requirements for each deployment scenario.

5 Energy-Aware runtime power management

5.1 Introduction

Runtime power management for energy-autonomous cyber-physical microsystems departs fundamentally from classical low-power design. In battery-dominated systems the manager's role is to throttle consumption against a fixed energy budget; under diffuse and intermittent ambient inputs, it must instead orchestrate consumption and opportunistically align it with time-varying harvest while preserving survivability. This chapter develops a runtime viewpoint and translates it into concrete scheduling policies and system contracts suitable for untethered microsystems operating on heterogeneous sources (e.g., photovoltaic, thermoelectric, vibrational, RF). Building directly on the loss-aware architecture and macroscopic energy model established in Chapter 4, we treat the energy path (sources $\rightarrow$ conversion $\rightarrow$ storage $\rightarrow$ distribution) and the software scheduler as a single closed loop with explicit constraints and measurable proxies. In particular, the runtime will reason over the stored-energy state, cold-start/brown-out thresholds, and the neutrality-survivability-availability triad of performance metrics defined previously, rather than ad-hoc voltage or duty-cycle heuristics.

The operating context is characterized by three features. First, stochastic supply: ambient power exhibits strong non-stationarity (diurnal/seasonal PV, occupancy-driven indoor light, slow thermal drifts, bursty mechanical excitation). Second, packetized demand: application workloads are a mix of quiescent supervision and crest-factor bursts (sensing, computation, radio), whose timing can often be deferred or coalesced without violating task-level constraints. Third, non-ideal power paths: conversion inefficiencies, quiescent draws, storage leakage, and protection hysteresis impose state-dependent losses and gating events that materially affect feasibility. A correct runtime must therefore (i) maintain operation within protection limits under uncertainty, (ii) adapt load timing and intensity to the instantaneous and forecasted energy state, and (iii) do so with overheads that are provably negligible relative to the harvested power on the design window.

We organize the design space into two algorithmic families that will be made precise in §5.2: reactive managers, which regulate on measured store proxies and protection events with minimal look-ahead, and predictive managers, which incorporate short- to medium-horizon forecasts of harvest and/or load to schedule discretionary work. Reactive policies offer robustness and tiny overheads and are indispensable near cold-start or under prolonged

energy poverty. Predictive policies can materially increase task availability and quality-of-service when inputs or workloads exhibit exploitable temporal structure, provided that forecasting and planning remain net-energy positive [236], [237], [238].

To unify these strands, §5.3 introduces DREAM (Dynamic Rest & Energy-Aware Management), a runtime framework that couples a minimal telemetry contract to a two-mode scheduler. In survival mode, DREAM enforces invariants that guarantee adherence to the brown-out hysteresis while servicing a minimal task set at a provable cadence. In opportunistic mode, it exploits surplus by stretching sleep intervals down, coalescing bursts to reduce delivery losses, and admitting optional tasks under an energy-neutrality test computed on a sliding window. Mode transitions are triggered by thresholds on a single store proxy (e.g., V_{store} for capacitive stores or SoC for batteries), augmented by edge-triggered protection events from the power path; no high-rate sensing is required. The framework is parameterized only by macroscopic quantities identified once per platform—conversion efficiencies, quiescent currents, leakage, start-up energy—so it remains portable across hardware revisions.

Methodologically, this chapter makes three contributions. First, it formalizes the runtime problem against the unified model of Chapter 4: the state equation for stored energy, event logic for cold-start/brown-out, and window-based metrics (neutrality, survivability, availability) become explicit constraints for scheduling. Second, it provides net-power-positive decision rules for both reactive and predictive regimes, including sufficiency conditions under which forecasting is beneficial once its sensing/compute overheads are debited to the energy budget [238], [239], [240]. Third, it defines a telemetry and control interface—one reliable store proxy, protection events, and optional coarse source proxies—that is sufficient to implement DREAM with bounded overhead and without application-specific tuning.

The remainder of the chapter proceeds as follows. Section 5.2 surveys and systematizes reactive and predictive power-management algorithms for environmentally powered nodes, clarifying their assumptions, stability properties, and overheads in the ultra-low-power regime [236], [241]. Section 5.3 instantiates these concepts in DREAM, states operating assumptions and notation, and derives policy tests that map directly to measurements available at runtime. Collectively, these results close the loop between the loss-aware energy

path and software scheduling, enabling principled autonomy for diffuse-energy microsystems rather than heuristic tuning.

5.2 Power management algorithms

This section systematizes the algorithmic landscape for runtime energy management in energy-harvesting microsystems. We adopt a two-axis taxonomy: (i) reactive controllers that regulate on measured energy-store proxies and protection events using thresholding, hysteresis, or event-driven duty-cycling; and (ii) predictive controllers that incorporate short-horizon forecasts of harvest and/or load to solve admission, pacing, and burst-coalescing decisions under an energy-neutrality constraint. The survey emphasizes stability, and survivability guarantees near cold-start, the net-energy positivity of any estimation or planning overheads, and the minimal telemetry required to implement each class in ultra-low-power regimes. We link these design choices to hardware realities—conversion inefficiencies, leakage, and quiescent losses—and to application-level QoS via admission control and task slack management. Representative exemplars span perpetual-system runtimes, energy-aware admission/duty-cycling, and information-theoretic formulations that clarify when look-ahead is provably beneficial. The resulting synthesis provides the comparative backdrop against which §5.3 instantiates DREAM’s two-mode scheduler and its sufficiency conditions for safe opportunism [242], [243], [244], [245].

5.2.1 Reactive

- **Hibernus (reactive save–restore on brown-out).**

Balsamo et al. introduce *Hibernus*, a minimalist, event-driven supervisor that guarantees brown-out–safe operation by coupling dual store-voltage watermarks with a single-shot state checkpoint and deep sleep on imminent undervoltage. By triggering exclusively on low-rate store proxies and protection events (comparator interrupts), Hibernus avoids continuous metering and complex forecasts, yielding microcontroller-class overheads that are negligible on harvested budgets. Its principal contribution is a *reactive invariants* design: pre-empt before protection trips, resume only above a hysteretic upper bound, and touch non-volatile memory at most once per outage. The approach sets a clear baseline for survivability near cold-start and during prolonged energy poverty, though it is conservative under surplus because it lacks admission granularity across optional tasks. [246].

- **QuarkOS (ultra-fine fragmentation under intermittent power).**

Zhang et al. argue that forward progress on micro-powered sensors is best secured by fragmenting software into ultra-short, idempotent quanta that fit within volatile energy “packets.” QuarkOS pioneers this *micro-tasking* stance, combining aggressive duty-cycling with power-failure-tolerant execution that simply resumes at the next boot. The key insight is that when outages are frequent and unpredictable, *reacting* with fine granularity and atomic commits outperforms heavier checkpointing or prediction. Limitations include restricted expressiveness for complex I/O atomicity and potential overheads from very fine partitioning, which later systems address [247].

- **Intermittent computation without hardware support (OSDI’16)**

Van der Woude and Hicks formalize intermittent execution on commodity microcontrollers without specialized hardware or restart recovery. They present a software-only runtime that ensures crash consistency by constraining how state is read/written across *task boundaries*, and by committing only at well-defined points that tolerate abrupt power loss. The contribution is a practical recipe for *reactive robustness*—use brown-out events to pre-empt, and structure all progress as idempotent task completion—together with evidence that correctness and acceptable overheads are achievable in real applications. The work catalyzed a family of intermittent-operating systems and clarified where reactive control ends and programming model support must begin [248].

- **Alpaca (task-atomic intermittent execution without checkpoints)**

Maeng, Colin and Lucia remove explicit checkpoints by framing the programming model around *idempotent, task-atomic* units whose side effects are committed only upon successful completion; partial work is discarded automatically on the next power-up. Alpaca exemplifies reactive design at the language/runtime boundary: it needs no forecasts, relies on store proxies and protection hysteresis to pace progress, and achieves low energy/time overhead by avoiding repeated NVM writes. Its guarantee—monotone progress in the presence of arbitrary outages—makes it a strong fit for battery-less nodes, though the strictness of task idempotence requires disciplined application structuring [249].

- **Coala (dynamic task-based intermittent execution).**

Majid et al. extend the above by *reactively* adapting task granularity at run time—splitting or fusing tasks based on recent energy availability and power-on durations. This keeps

execution quanta aligned with the ambient energy packets without any explicit harvest forecasting. The system also prioritizes tasks under energy scarcity while preserving forward progress guarantees. The main value is empirical: dynamic granularity reduces wasted partial work and improves effective throughput on harvested power; the trade-off is added runtime complexity that must remain net-energy positive [250].

- **ReplayCache (volatile caches for energy-harvesting systems).**

Zeng et al. show that conventional volatile caches can be made *reactively persistent* by logging cache-line fills/evictions and replaying them after a power failure, thereby avoiding wholesale checkpointing of memory hierarchies. The method triggers on power events and low-rate proxies (no forecasting), bringing a key subsystem into the intermittent-computing fold with modest area/energy overheads. This expands the reactive toolkit beyond scheduling into memory-system design, though, as with any logging, pathological burst patterns can stress the log unless sized with care [251].

- **Ocelot (automatic freshness and I/O consistency).**

Surbatovich et al. address a subtle failure mode in intermittent systems: *stale or inconsistent inputs/outputs* after outages. Ocelot wraps I/O and dependent computations in automatically inferred *atomic regions*, revalidating input freshness and either committing or rolling back upon restoration. The runtime remains purely reactive—no prediction, only event-driven checks tied to power-loss recovery—and provides strong semantic guarantees with low programmer burden. Its benefits are most pronounced in sensing/actuation pipelines; the cost is additional metadata and validation work that must be kept tiny under scarce energy [252].

- **Write-Light Cache (reducing writing energy in intermittent contexts).**

Choi et al. propose a *write-light* volatile cache design that reduces the energy and time overhead of persistence-related writes—critical when power cycles are frequent. While architectural in focus (ISCA), the contribution is squarely reactive in spirit: anticipate brown-outs and minimize the footprint of state that must be stabilized across them. The design complements ReplayCache-like logging by cutting write amplification, thereby improving the viability of more capable memory subsystems on harvested budgets [253].

- **Wake-up receivers + harvesting (radio/sensing gating as reactive policy).**

Ait Aoudia et al. empirically demonstrate that combining *wake-up receivers (WURs)* with energy harvesting enables large reductions in idle energy by *reactively* gating high-draw radios until an external trigger arrives or the store surpasses a watermark. This is a quintessential reactive peripheral policy: it requires only coarse store telemetry and event interrupts, yet it unlocks order-of-magnitude lifetime improvements in duty-cycled IoT nodes. The approach is immediately practical and orthogonal to core execution models but relies on WUR sensitivity/selectivity commensurate with the application environment [254].

- **Capacity over capacitance (store-aware invariants and sizing).**

Jackson, Adkins and Dutta argue—via analysis and SenSys evaluation—that *right-sizing energy capacity and thresholds* in view of leakage, conversion losses and protection hysteresis can improve reliability more than naively adding capacitance. The work reframes reactive watermarks as *store-aware invariants* derived from macroscopic platform parameters (quiescent currents, delivery efficiency, leakage), providing a principled basis for admission/hibernate levels. It also cautions against over-provisioning that bloats quiescent losses—an insight we adopt when parameterizing thresholds [255].

- **Reactive task schedulers for battery-less IoT (per-task watermarks).**

Sabovic et al. present an energy-aware task scheduler for *battery-less* devices that attaches per-task admission tests to a store proxy (voltage/SoC) and re-queues tasks when conditions degrade. The design is resolutely reactive—no look-ahead—yet achieves high service ratios by prioritizing essential tasks and exploiting momentary surpluses. Their results quantify a sweet spot where simple gating and priority rules deliver most of the benefit of heavier planners, provided the workload is packetized and dependencies are encoded [256].

- **DVFS + duty-cycling as reactive knobs; survey evidence.**

Khriji et al. synthesize post-2012 evidence that dynamic voltage/frequency scaling combined with classical duty-cycling remains a *reactive mainstay* for low-power nodes, provided control cadence is slow and telemetry minimal. The review’s relevance here is twofold: it highlights how *tiny* control loops must be to stay net-energy positive on harvested budgets, and it enumerates pitfalls (e.g., chatter near thresholds) that our own policies avoid through hysteresis and slope limiting [209].

5.2.2 Predictive

- **Online EH prediction with sliding-window models.**

Cammarano, Petrioli and Spenza develop lightweight online predictors tailored to environmentally powered nodes, combining persistence/EMA baselines with seasonal components and adaptive windowing to track diurnal structure. They embed the predictor into an energy-neutral controller that scales activity, so planned consumption matches expected inflow over a near-term horizon. The key result is that even simple, well-tuned models yield material availability gains over purely reactive policies, provided sensing and model-updating costs are kept net-energy positive [240].

- **ASARIMA: adaptive seasonal ARIMA for harvested power**

Li and Han propose ASARIMA, an adaptive scheme that re-weights and re-parameters SARIMA models to cope with changing irradiance regimes (clear vs. cloudy) and slow seasonal drift. On representative solar traces, ASARIMA improves short-horizon forecasts and stabilizes duty-cycling decisions by producing less biased next-interval energy estimates. The contribution is a principled statistical alternative to black-box learning that remains interpretable and MCU-friendly [257].

- **Q-learning for prediction under regime shifts**

Kosunalp frames energy prediction as a reinforcement-learning problem in which a Q-learner selects among candidate predictors (or parameterizations) based on recent error and reward signals. This meta-predictor adapts across regimes without explicit weather labels, enabling controllers to be more aggressive when forecasts are reliable and conservative otherwise. The approach highlights that adaptivity in the choice of predictor can matter as much as the model class itself [258].

- **Edge LSTM for short-term irradiance**

Pi, Jin and Chen implement multichannel LSTM models that fuse local irradiance and auxiliary signals to forecast minute-scale solar power on edge hardware. Compared to linear baselines, the LSTM reduces RMSE and enables higher admission rates for deferrable tasks without increasing brown-out risk—so long as model evaluation costs are budgeted. The

work shows that compact deep models can be viable on microcontroller-class or embedded edge platforms [259].

- **Cloud-cover–assisted harvest prediction**

Renner demonstrates that incorporating coarse cloud-cover forecasts into local time-series models reduces day-ahead and intra-day error, particularly during fast-changing conditions. This exogenous-data assimilation is low overhead (few bytes per update) yet yields clearer admission signals for forecast-aware schedulers. Methodologically, it motivates blending global meteorological cues with local energy traces in simple controllers [260].

- **Energy management patterns across predictive controllers**

Ashraf et al. survey prediction-informed energy managers and distill design patterns: match control cadence to source periodicity, use sliding energy budgets, and explicitly account for sensing/compute overheads. Their synthesis provides actionable guidance on when look-ahead yields net benefit—e.g., under diurnal solar with moderate variability—and when reactive fallbacks dominate. These patterns underpin robust implementations on real EH nodes [261].

- **Robustness analysis and design guidelines**

Stricker, Gupta and Thiele analyze how forecast error, storage inefficiency, and model mismatch propagate through predictive controllers, offering robustness conditions and tuning rules. They recommend conservative baselines plus opportunistic “spending” only within confidence bounds, effectively bridging predictive and safety-first control. The result is a principled recipe for uncertainty-aware duty-cycling on intermittent supplies [262].

- **Uncertainty-aware prediction and control on battery-less IoT**

Yamin et al. present a pipeline that quantifies prediction uncertainty and feeds it into admission control for capacitor-powered devices. By executing optional work only when confidence intervals indicate surplus harvest, they improve availability while keeping brown-out incidence low. The study provides empirical evidence that explicit uncertainty treatment outperforms point-forecast controllers under volatility [263].

- **Optimal scheduling for battery-less devices with forecasts**

Delgado, Famaey and Mercelis formulate forecast-informed task scheduling as an optimization problem that enforces energy-neutral operation with minimal deadline violations. They demonstrate that small-horizon solvers (dynamic programming/MILP variants) are feasible on EH-class platforms when problem structure is exploited, and they quantify the useful look-ahead beyond which added horizon brings little benefit. This clarifies when MPC-like planning pays off [264].

- **Harvest prediction for WSN deployment planning**

Janković, Žemaitis and Navickas evaluate practical predictors for node-level harvested energy and show which features and horizons matter most for WSN duty-cycling. Their results inform both runtime policy (horizon selection) and pre-deployment sizing (panel/storage choices) when long local traces are unavailable. It is a deployment-focused complement to embedded prediction work [265].

- **Predictive energy-aware routing (PEAR) over 6TiSCH**

Jecan et al. move beyond single nodes and use per-node energy forecasts to shape routing, MAC duty-cycling, and transmission pacing. By steering traffic toward nodes predicted to have surplus and throttling drains on predicted-poor nodes, PEAR improves network lifetime and latency under ENO constraints. This generalizes node-level prediction to network-level resource allocation [266].

- **Predictive energy-management units (EMU) on real devices**

Rajappa et al. build a microcontroller-class EMU that executes daily solar prediction and sets next-day duty targets, reporting negligible energy overhead and stable operation across seasons. Their engineering contribution is a reproducible, end-to-end controller that integrates prediction, budgeting, and actuation with clear interfaces to the application. It demonstrates practicality on industrial IoT hardware [267].

- **LT-ENO: long-term energy-neutral operation**

Buchli et al. (SenSys'14) combine capacity planning with long-horizon predictive control to maintain energy neutrality across seasons, markedly reducing duty-cycle variance over 11-year solar traces. The runtime scales activity based on astronomical irradiance models and

storage constraints, explicitly accounting for round-trip inefficiency and self-discharge. LT-ENO is the canonical long-horizon predictive baseline for solar-powered nodes [268].

- **Minimum energy-utilization guarantees**

Follow-up work establishes optimal policies that guarantee minimum energy utilization under uncertainty, extending LT-ENO's guarantees and clarifying how to allocate energy while preserving neutrality. These results give theoretical footing to budget-based admission policies commonly used in practice. The emphasis is on provable guarantees rather than heuristic tuning [269], [270].

- **Capacity planning without long local traces**

Buchli et al. (EWSN'14) show how to size harvesting and storage using astronomical models and coarse location inputs, eliminating the need for site-specific long-term traces. This planning step conditions the effectiveness of any predictive runtime by ensuring that neutrality targets are actually attainable. It couples naturally with LT-ENO controllers deployed afterward [271].

- **PreAct: predictive control under time-varying utility**

Geissdoerfer et al. extend predictive energy management to time-varying utility: they track desired utility profiles (e.g., weekend priorities) while staying energy-neutral, using a low-cost predictor and PID-style controller. PreAct explicitly incorporates inefficiencies and uncertainty and demonstrates microjoule-scale daily overhead on MSP-class MCUs. It evidences that modest look-ahead plus uncertainty handling can yield tangible QoS gains [272].

5.3 Energy-Aware runtime system management

5.3.1 DREAM - Dynamic Rest & Energy-Aware Management

DREAM is the runtime duty-cycle scheduler that implements, in the deployed node, the control logic implied by the unified energy model introduced earlier in this chapter. Its task is to choose the next sleep interval at the end of each epoch so that, when harvested power is sufficient, the storage element exhibits no long-term drift over representative windows, and, when conditions deteriorate, the node remains operative for a configured survival horizon without breaching protection limits. DREAM therefore sits exactly at the boundary

between modelling and hardware: it reads the storage proxy and protection events exported by the power path and emits an actuation decision—principally the next sleep time and the enable state of the rails—that is consistent with the macroscopic accounting used by the chapter’s analysis.

The algorithm operates on an epoch abstraction. Each epoch has duration $T_e(n) = T_{act} + T_s(n)$, where T_{act} is the fixed active time determined by the application, and $T_s(n)$ is the sleep time selected by DREAM for the subsequent epoch. Active and sleep powers are treated as profiled constants of the deployed firmware image; this mirrors the chapter’s deliberate separation between offline characterization and online control. By confining itself to these coarse-grained quantities, the scheduler matches the scale at which energy accumulates in storage and avoids the measurement overheads and modelling fragility that accompany fine-grain telemetry.

The only mandatory measurement is a store-state proxy sampled once per epoch. For capacitive storage this is the terminal voltage, mapped to energy through $E = \frac{1}{2}CV^2$; for batteries it is an SoC estimate that may incorporate temperature compensation. DREAM does not depend on dense source or load sensing, nor on predictions of future input. Instead, it enforces the chapter’s windowed criteria using the most recent proxy sample, the fixed active time, and the profiled powers. Protection is represented as event-gated hysteresis—cold-start and undervoltage/over-voltage thresholds—that pre-empt ordinary operation; the scheduler’s decisions are always subordinate to these gates. In particular, it must not schedule a wake that would immediately re-enter a forbidden region, and it only re-arms loads once the upper threshold is crossed.

Two complementary control objectives guide the selection of T_s . In energy-surplus conditions the goal is neutrality: the algorithm corrects the store deviation measured at the end of the current epoch by adjusting the next sleep time so that the expected net change over the following epoch cancels that deviation. Under the steady-input-between-epochs approximation used throughout the chapter, this rule admits a closed form that depends only on the last proxy measurement and the profiled parameters. In energy-poor conditions the goal is survival: if the incoming energy cannot amortize the cost of the next active phase, DREAM switches policy and chooses a cadence that guarantees operation for at least a target horizon T_{sur} within the usable energy band defined by the lower protection threshold. This

survival cadence can be computed once upon entry to scarcity and, if desired, re-evaluated each epoch to exploit incidental input without departing from constant-time operation.

The metrics established earlier—energy neutrality on windows, survivability with respect to protection, and availability as the fraction of time the load is safely powered—provide the acceptance criteria against which DREAM’s behavior is judged. Neutrality requires that the storage trajectory close each representative window without negative drift; survivability requires that the proxy remain within the admissible band; availability is maximized subject to these constraints. By basing decisions on a single proxy and event-gated protection, DREAM renders these metrics checkable both in offline trace replay and online, while keeping runtime overhead negligible relative to typical harvested power levels.

Adopting an epoch-based controller is not merely an implementation convenience. It aligns the controller’s temporal granularity with the physics of the storage element and with the communication and sensing patterns of the application, for which work is naturally organized in bursts separated by quiescent intervals. This choice also yields deterministic computational cost. The algorithm maintains only two mutable quantities—the last sleep time and the last energy estimate—and applies constant-size arithmetic at each epoch boundary. No forecasting, optimization loops, or sliding histories are required. The simplicity of the state and update law, limits failure modes eases formal reasoning about behavior under protection events and facilitates deployment on microcontrollers with severe energy and memory budgets.

Practical considerations arising from evaluation inform two further aspects of the design. First, measurement noise at the proxy and small, rapid fluctuations in the store can induce oscillatory adjustments to the sleep time if unconstrained. DREAM therefore applies bounded updates and recommends platform-level clamps for the minimum and maximum sleep interval; these clamps are chosen to respect watchdog servicing, data freshness requirements, and connectivity needs of the application. Second, while the algorithm is indifferent to the presence or absence of per-source telemetry, it allows such proxies to be ingested for diagnostics or to break ties when multiple schedules satisfy the neutrality or survival objectives; correctness, however, never depends on them.

Finally, the scope of DREAM is deliberately narrow. It does not perform maximum-power-point tracking, source arbitration, or calendar-time scheduling online. Those choices are

captured in the chapter’s unified model via efficiencies, masks, and the definition of representative windows, and they are exercised during offline sizing and verification. This separation of concerns preserves the clarity of the runtime policy, keeps its costs fixed and small, and ensures that the algorithm’s guarantees—neutrality under surplus and a quantified survival horizon under scarcity—rest on a small, explicit set of assumptions that are satisfied by the platform interface defined in the earlier sections of the chapter.

Concretely, **DREAM** pursues the following design objectives:

1. **Energy neutrality (surplus):** drive the expected store drift over representative windows toward zero by selecting T_s consistent with the measured epoch-level energy change and an online estimate of input power.
2. **Survivability (scarcity):** when input is insufficient to amortize the active load, schedule sleeps to meet a user-defined minimum survival time without violating lower energy bounds.
3. **Safety under protection:** honor cold-start enable and brown-out hysteresis at all times, including deep-sleep re-arming and safe re-entry to normal operation.
4. **Minimal telemetry and compute:** rely on a single proxy sample per epoch and constant-time arithmetic; avoid dense sensing and long-horizon forecasting.
5. **Portability:** express control in macroscopic quantities (epoch energies/powers) so that characterization and operation remain robust across hardware revisions and storage technologies.

5.3.2 Operating assumptions and notation

DREAM operates on a minimal runtime contract: a single, reliable store-state proxy sampled once per epoch (supercapacitor voltage or battery SoC) and edge-triggered protection events (cold-start enable and hysteretic UVLO/OVP). Decisions are taken at epoch boundaries using macroscopic energy accounting consistent with the chapter’s loss-aware model, avoiding dense telemetry or long-horizon forecasting.

Operating assumptions. Execution proceeds in epochs composed of a fixed active interval and a sleep interval chosen by the runtime.



Figure 5.1 Two-state operation

Active and sleep powers are treated as profiled constants of the firmware build. The proxy sample is the only mandatory measurement; optional per-source signals are advisory and do not affect correctness. Protection events pre-empt scheduling and gate rail enables; the scheduler must not issue a wake that would immediately violate hysteresis. This stance keeps per-epoch compute and sensing overheads negligible on MCU-class platforms.

Notation. Let n index epochs and let DREAM select the next sleep $T_s(n+1)$ at the end of epoch n . The epoch duration is:

$$T_e(n) = T_{\text{act}} + T_s(n) \quad (5.1)$$

For supercapacitor stores the energy state is derived from the measured terminal voltage; batteries use a temperature-aware SoC map. The supercapacitor relation is:

$$E(n) = \frac{1}{2} C V^2(n) \quad (5.2)$$

Per-epoch accounting follows the discrete balance:

$$\Delta E(n) = E(n) - E(n-1) = E_{\text{in}}(n) - E_{\text{out}}(n) \quad (5.3)$$

with harvested and consumed aggregates:

$$E_{\text{in}}(n) = P_{\text{in}}(n) T_e(n) \quad (5.4)$$

$$E_{\text{out}}(n) = E_{\text{act}} + P_{\text{sleep}} T_s(n), \quad E_{\text{act}} = P_{\text{act}} T_{\text{act}} \quad (5.5)$$

Combining these equations gives, for supercapacitors, the measurable form used operationally:

$$\Delta E(n) = \frac{1}{2} C [V^2(n) - V^2(n-1)] \quad (5.6)$$

An online estimate of the average harvested power over epoch n follows by rearranging (Eq. 5.1)–(Eq. 5.5); this estimator is evaluated once per epoch from boundary measurements (Eq. 5.7):

$$P_{in}(n) = \frac{\Delta E(n) + E_{act} + P_{sleep} T_s(n)}{T_{act} + T_s(n)} \quad (5.7)$$

Protection and bands are represented by admissible store limits $[E_{LB}, E_{UB}]$ (or equivalently $[V_{min}, V_{max}]$) and hysteretic thresholds for cold-start and brown-out ($V_{CS}, V_{UVLO}^{\downarrow}, V_{UVLO}^{\uparrow}$). All scheduling decisions are subordinate to these event gates; crossings authorize, or forbid load enables, and the controller remains quiescent inside the dead band.

The symbols and identities above are the sole primitives consumed by the controller. They map one-to-one to the flowchart's formulas and state transitions and will be invoked directly in §5.3.3 to derive the neutrality law and in §5.3.4 to construct the survival cadence under scarcity.

5.3.3 Neutrality control law (energy surplus)

Neutrality seeks a zero net drift of the store between consecutive epochs. Using the per-epoch balance with the harvested-power estimate, the standard one-step persistence $P_{in}(n+1) \approx P_{in}(n) = P_{in}$, the sleep for the next epoch is chosen so that $\Delta E(n+1)=0$. The resulting closed-form update is :

$$T_s(n+1) = \frac{(P_{in} - P_{act})T_{act}}{P_{in} - P_{sleep}} \quad (5.8)$$

This expression is evaluated at the end of epoch n using P_{in} . It is feasible only when the denominator is positive, i.e., when the average harvested power exceeds the background draw P_{sleep} . In practice, $T_s(n+1)$ is bound to the platform's admissible range $[T_s^{min}, T_s^{max}]$ and constrained to be non-negative; if $P_{in} \leq P_{sleep}$ or the denominator is near zero, neutrality is deemed infeasible and the scarcity policy of §5.3.4 is invoked.

To suppress limit cycles due to measurement noise and quantization, a first-order smoothing and clamping step is applied to the raw update T_s^* from (Eq. 5.8). With $w \in (0,1)$, the implementation uses:

$$T_s(n+1) \leftarrow \text{clip}(wT_s^* + (1-w)T_s(n), T_s^{\min}, T_s^{\max}) \quad (5.9)$$

All scheduling remains subordinate to protection gates (cold-start enable and UVLO hysteresis): if a gate asserts, the controller refrains from issuing wakes until the re-arm threshold is crossed. Computational cost is constant - time per epoch, as the realization depends only on the two boundary proxy samples, the profiled load parameters, and the last sleep value.

5.3.4 Survival control law (scarcity)

When the estimated input cannot amortize the next active phase, neutrality is infeasible, and the scheduler must guarantee operation for at least a configured horizon T_{sur} while respecting the lower admissible store limit E_{LB} . Let E_0 denote the store energy at the instant of entering scarcity; epochs retain the same fixed active duration T_{act} , and the runtime seeks a sleep T_s that satisfies the survival constraint over the window $[0, T_{\text{sur}}]$.

To guard against optimistic measurement noise and short-lived bursts, the controller uses a conservative input-power surrogate obtained by scaling the online estimate with a confidence factor $\rho \in (0, 1]$:

$$P_{\text{in}}^{\text{sur}} = \rho P_{\text{in}}(n) \quad (5.10)$$

Under the macroscopic balance, the total energy after N equal epochs is

$$E(N) = E_0 + N[(P_{\text{in}}^{\text{sur}} - P_{\text{sleep}})T_s + (P_{\text{in}}^{\text{sur}} - P_{\text{act}})T_{\text{act}}] \quad (5.11)$$

with $N = \frac{T_{\text{sur}}}{T_{\text{act}} + T_s}$. Enforcing $E(N) \geq E_{\text{LB}}$ and solving at equality for the smallest admissible T_s yields the closed-form survival cadence:

$$T_s^* = -T_{\text{act}} \frac{(E_0 - E_{\text{LB}}) + T_{\text{sur}}(P_{\text{in}}^{\text{sur}} - P_{\text{sleep}})}{(E_0 - E_{\text{LB}}) + T_{\text{sur}}(P_{\text{in}}^{\text{sur}} - P_{\text{act}})} \quad (5.12)$$

The implementation clamps T_s^* to the platform bounds $[T_s^{\min}, T_s^{\max}]$ and to non-negativity. Feasibility depends on the sign of the denominator in (Eq. 5.12): if $(E_0 - E_{\text{LB}}) + T_{\text{sur}}(P_{\text{in}}^{\text{sur}} - P_{\text{sleep}}) > 0$, the left-hand side of the survival inequality increases with T_s , and the minimal feasible sleep is precisely T_s^* after clamping. If the denominator is non-positive, increasing sleep does not buy safety; the controller must then select the boundary that maximizes survivability (typically $T_s = T_s^{\max}$) and check the residual deficit. When

even continuous sleep cannot sustain the store — i.e., $P_{in}^{sur} \leq P_{sleep}$ and $(E_0 - E_{LB}) < T_{sur}(P_{sleep} - P_{in}^{sur})$ — the policy escalates to deep sleep; the expected time-to-threshold under sleep-only drain is:

$$T_{deep} \approx \frac{E_0 - E_{LB}}{P_{sleep} - P_{in}^{sur}} \quad \text{for } P_{in}^{sur} < P_{sleep} \quad (5.13)$$

Survival mode is re-evaluated at each epoch boundary using updated $E_0 \leftarrow E(n)$ and P_{in}^{sur} , so that incidental improvements in input shorten the cadence back toward neutrality as soon as feasible. All actuation remains subordinate to protection gates: if undervoltage asserts during scarcity, wake-ups are inhibited until the re-arm threshold is crossed, at which point (Eq. 5.12) is re-applied.

5.3.5 Mode selection and protection logic

At the end of each epoch the scheduler chooses the next operating mode by combining the store state with the online input-power estimate, while always honoring protection gates.




Figure 5.2 DREAM state diagram

Neutrality is attempted only when it is both feasible and safe. Feasibility is encoded by the predicate $\mathcal{F}_{neu}(n)$, which requires harvested power to exceed the background draw and the store to lie within admissible limits:

$$\mathcal{F}_{neu}(n) : (P_{in}(n) > P_{sleep}) \wedge (E_{LB} < E(n) \leq E_{UB}) \quad (5.14)$$

Safety is enforced through a protection gate predicate $\mathcal{G}(n)$ that captures cold-start enable and hysteretic undervoltage; when this gate is asserted, no wake is permitted regardless of computed cadences:

$$\mathcal{G}(n) : (V(n) \leq V_{UVLO}^{\downarrow}) \vee (V(n) < V_{CS}) \quad (5.15)$$

Given these predicates, the mode for the next epoch is selected as follows: deep sleep when protection asserts; survival when neutrality is unsafe or infeasible and the estimated input cannot amortize the active phase; neutrality otherwise:

$$mode(n+1) = \begin{cases} DeepSleep, & \mathcal{G}(n), \\ Survival, & \neg \mathcal{G}(n) \wedge (E(n) \leq E_{SoC,des} \vee \neg \mathcal{F}_{neu}(n)) \wedge (P_{in}(n) \leq P_{act}), \\ Neutrality, & \neg \mathcal{G}(n) \wedge \mathcal{F}_{neu}(n) \end{cases} \quad (5.16)$$

Re-entry from a protected state requires clearing the hysteresis margin and cold-start threshold; only then is the computed cadence applied, with platform bounds encoding watchdog servicing and data-freshness constraints:

$$Re - arm: V(n) \geq V_{UVLO}^{\uparrow} \wedge V(n) \geq V_{CS}, \quad T_s(n+1) = clip(T_s^*(n+1), T_s^{min}, T_s^{max}) \quad (5.17)$$

This logic ensures that scheduling decisions remain energy-aware and protection-safe: neutrality holds the store steady when possible; survival guarantees a configured horizon when input is insufficient; and deep sleep preserves the store through enforced brown-out or cold-start conditions.

5.3.6 Online parameter identification

The runtime requires two load aggregates: the active-phase energy E_{act} for the fixed application workload, and the low-power draw P_{sleep} that accrues during rest. Both can be identified on the target platform using brief experiments that read only the store proxy at epoch boundaries.

In the supercapacitor case, operate in a short window where harvested input is negligible, so that per-epoch energy decrements reflect the load alone. Run three consecutive epochs with fixed T_{act} and two distinct sleep durations X and Y in epochs n and $n+1$. The discrete balances become:

$$E_{\text{act}} + P_{\text{sleep}} X = \frac{1}{2} C [V^2(n-1) - V^2(n)] \quad (5.18)$$

$$E_{\text{act}} + P_{\text{sleep}} X = \frac{1}{2} C [V^2(n) - V^2(n+1)] \quad (5.19)$$

Eliminating E_{act} yields a closed-form estimate for the background draw,

$$P_{\text{sleep}} = \frac{\frac{1}{2} C [V^2(n-1) - 2V^2(n) + V^2(n+1)]}{Y - X} \quad (5.20)$$

after which the active energy follows directly from any one balance, e.g.

$$E_{\text{act}} = \frac{1}{2} C [V^2(n-1) - V^2(n)] - P_{\text{sleep}} X \quad (5.21)$$

With P_{sleep} known, a short run with $T_s=0$ provides E_{act} per epoch from a single energy drop and thus the active power,

$$P_{\text{act}} = \frac{E_{\text{act}}}{T_{\text{act}}} \approx \frac{\frac{1}{2} C [V^2(k) - V^2(k+1)]}{T_{\text{act}}} \quad \text{with } T_s = 0 \quad (5.22)$$

On battery platforms, replace the supercapacitor mapping with the SoC-to-energy relation; when the estimator exposes usable state of energy or a calibrated capacity E_{max} , use

$$\Delta E(n) \approx E_{\text{max}} [\text{SoC}(n) - \text{SoC}(n-1)] \quad (5.23)$$

and substitute $\Delta E(\cdot)$ for the voltage-squared differences in (5.18) – (5.19). If truly negligible input cannot be guaranteed, alternate the two sleep settings in quick succession so that average input remains approximately constant over the pair; subtracting the two balances then cancels the common input term to first order. Repeating the triplet experiment for several (X,Y) pairs and averaging the resulting estimates reduces noise. Leakage and regulator housekeeping are, by construction, absorbed into P_{sleep} ; when leakage varies with

voltage, conduct the experiment inside the intended operating band so that the identified values reflect the region where the scheduler will regulate.

5.3.7 Complexity, memory footprint, and overhead

DREAM is engineered for constant-time operation per epoch. At each epoch boundary the runtime executes a fixed, bounded arithmetic pipeline: one proxy acquisition; one energy-difference computation (for supercapacitors, a squared-voltage difference scaled by $\frac{1}{2}C$; for batteries, a SoC difference scaled by capacity/SoE map); one harvested-power estimate using the rational form of (Eq. 5.7); one sleep-time update from the neutrality or survival law (Eqs. 5.9/5.12); optional first-order smoothing and clamping; and a small set of comparisons for mode selection and protection gates. None of these steps depends on history length or input rate, so the time complexity is $O(1)$ per epoch. Even with modest noise-handling—e.g., a two-sample moving average or a single-epoch outlier guard—the window is fixed and the bound remains $O(1)$.

The arithmetic cost is dominated by at most two divisions along the hot path (one for the input-power estimate, one for the sleep update); all other operations are additions, multiplications, and comparisons. This makes latency predictable on MCU-class targets. If a hardware FPU is unavailable, the formulas admit straightforward fixed-point (Q-format) implementations with pre-scaled constants and saturation, preserving numerical stability within the operating band while further tightening worst-case latency.

The mode-selection logic of §5.3.5 adds only a handful of predicate evaluations for feasibility and protection; this is branch-only work with no iterative search. The online identification procedure in §5.3.6 is also constant-time when invoked (two discrete balances and a closed-form 2×2 solve). Because identification is run infrequently (e.g., at commissioning or on long intervals to track drift), its cost is amortized over many epochs and does not affect steady-state timing.

The memory footprint is similarly bounded. The controller maintains only the last proxy sample and energy estimate, the current T_s , and the identified/profiler parameters (E_{act} , P_{sleep}), plus platform constants (capacitance or SoE mapping handles, admissible energy/voltage limits, watchdog and freshness bounds, hysteresis thresholds). This is a constant number of scalars, implying space complexity $O(1)$ with no dynamic allocation, buffers, or queues.

Energy overhead is limited to the proxy read and the brief end-of-epoch compute. Because the sampling rate is one read per epoch and the arithmetic is constant-size, the control energy is negligible relative to application activity and background leakage across both surplus and scarcity conditions. Overhead remains stable as conditions vary: scarcity changes only the scheduled cadence, not the amount of computation per boundary. Latency and real-time behavior are governed by two external tunables — T_{act} (the minimum service quantum) and the clamping bounds $[T_s^{\min}, T_s^{\max}]$. Provided these are respected, DREAM’s timing is deterministic: the update executes once at the boundary, and the device then remains in a low-power state for the scheduled rest without intermediate intervention.

Finally, portability follows from the macroscopic formulation. All quantities are either directly measured at epoch boundaries or identified in situ with brief experiments; no converter-level waveforms, current sensors, or vendor-specific APIs are required. The same implementation carries across storage technologies and hardware revisions by adjusting only configuration constants and the proxy-to-energy relation at build time.

5.3.8 Flowchart-based realization

This subsection instantiates the control laws of §§5.3.3–5.3.5 as a compact, per-epoch routine. At each epoch boundary, the controller samples the store proxy, forms the epoch energy increment, estimates the average harvested power, selects the mode (neutrality or survival) subject to protection gates, computes a candidate sleep T_s^* , applies smoothing and platform clamps, and commits the schedule only if re-arm thresholds are cleared. The routine is single-pass and constant-size, inheriting the $O(1)$ time/space properties of §5.3.7.

DREAM ALGORITHM

Input: T_{act} , E_{act} , P_{sleep} , bounds $\{T_s^{\min}, T_s^{\max}\}$,
thresholds $\{V_{\text{CS}}, V_{\text{UVLO}\uparrow}, V_{\text{UVLO}\downarrow}\}$ or $\{E_{\text{LB}}, E_{\text{UB}}\}$,
smoothing weight $w \in (0,1)$, survival horizon T_{sur}
State: last proxy sample (V_{prev} or SoC_{prev}), last energy E_{prev} ,
last sleep T_s , current mode.

- 1 Acquire proxy sample at end of epoch n : V or SoC
- 2 Compute ΔE from proxy difference; update $E(n)$
- 3 Estimate $P_{\text{in}}(n)$ using ΔE , E_{act} , P_{sleep} , T_{act} , T_s
- 4 Evaluate protection gate $G(n)$. If asserted $\rightarrow$ DeepSleep; skip to step 9
- 5 Test neutrality feasibility: $P_{\text{in}}(n) > P_{\text{sleep}}$ and $E_{\text{LB}} < E(n) \leq E_{\text{UB}}$.
If true $\rightarrow$ compute T_s^* via neutrality (ΔE target 0); go to step 7
- 6 Scarcity path:
- If $P_{\text{in}}(n) \leq P_{\text{act}}$ or $E(n) \leq$ desired band $\rightarrow$ compute T_s^* via survival
(with conservative P_{in} surrogate if configured).
- Else $\rightarrow$ neutrality as in step 5

- 7 Smooth and clamp: $T_s \leftarrow \text{clip}(w \cdot T_s^* + (1 - w) \cdot T_s, T_s^{\min}, T_s^{\max})$
 - 8 If protection re-arm conditions not met (e.g., $V < V_{\text{UVLO}\uparrow}$ or $V < V_{\text{CS}}$), force DeepSleep
 - 9 Persist minimal state: {proxy sample, E(n), T_s , mode};
 - 10 Schedule sleep T_s
-




Figure 5.3 DREAM flowchart

5.3.9 Limitations and extensions

DREAM is deliberately macroscopic and reactive: it regulates via a single store proxy and per-epoch aggregates, forecasting only one-step persistence of harvested power. This minimizes sensing/compute cost but bounds performance. With structured inputs (e.g., diurnal cycles or periodic loads), neutrality can converge conservatively and survival may select safe, sub-optimal cadences. Averaging over power, rather than fine-grain currents, also misses sub-epoch transients (e.g., brief regulator sweet spots or input surges).

Practical proxies add error: supercapacitor voltage and battery SoC drift with temperature, bias, leakage, and housekeeping currents; online identification (§5.3.6) reduces but does not remove bias, so operating near protection limits can force conservative scheduling. DREAM further assumes workload stability across epochs; abrupt changes (firmware, rare heavy tasks) stale Eact, temporarily degrading neutrality until re-identified.

Protection handling is simple and event-driven—safe and deterministic, but excluding aggressive recoveries (staged rail re-enables, speculative wake-ups under rising input). DREAM omits MPPT, multi-source arbitration, and task-level admission control, presuming these are handled by the power path and that active workload quanta have externally defined freshness/latency.

Extensions that retain $O(1)$ control include: (i) opportunistic, budgeted prediction (bounded-order smoothing for P_{in} or time-of-day bias with RTC); (ii) adaptive identification (periodic/triggered two-point updates on ΔE to track temperature/aging drift); (iii) multi-band objectives (bias neutrality toward a preferred sub-band for QoS headroom via a small offset); and (iv) a coarse load-shedding tier before deep sleep for heterogeneous loads, using constant-size comparisons.

In sum, DREAM trades predictive optimality for robustness, portability, and negligible overhead. Where needed, the above extensions add performance while preserving its essentials: constant-time epoch arithmetic, a single reliable proxy, and strict adherence to protection gates.

5.3.10 Summary

This subsection consolidated a lightweight runtime for energy harvesting nodes that regulates at epoch boundaries using a single store proxy and macroscopic energy accounting. After stating provenance and the control problem, §5.3.2 fixed the operating assumptions

and notation; §5.3.3 derived the neutrality law that nulls store drift under surplus; §5.3.4 constructed a survival cadence that guarantees a configured horizon under scarcity; §5.3.5 formalized mode selection and protection gating; §5.3.6 detailed an on-device identification of the load aggregates E_{act} and P_{sleep} ; §5.3.7 established constant-time and constant-space bounds; and §5.3.8 mapped the formalism to an executable flowchart. §5.3.9 discussed limitations and extensions, outlining bounded-complexity enhancements that preserve determinism and minimal overhead.

Collectively, these elements yield a controller whose decisions are traceable to measurable quantities, whose compute and memory costs are fixed and small, and whose safety is enforced by explicit protection predicates. The separation of concerns—loss-aware modelling at design time and constant-size control at runtime—enables portable deployment across storage technologies and hardware revisions while meeting neutrality and survivability objectives with negligible overhead. To situate the proposed runtime within the literature, Table 5.1 compares representative reactive and predictive managers against DREAM. We report only salient characteristics and the most informative deltas, using one umbrella row where multiple papers form a coherent family.

Table 5.1 Energy management algorithms comparison table

Work	Year	Type	Core characteristics	Differences vs. DREAM	Advantages of DREAM
Hibernus	2015	Reactive	Brown-out–safe save/restore with dual watermarks; single checkpoint on undervoltage	Coarse global gating; no energy-neutral windowing or survival horizon	Prediction-free epoch control with explicit neutrality in surplus and guaranteed survival cadence under scarcity
Alpaca	2017	Reactive	Task-atomic, idempotent execution without checkpoints	Programming/runtime model; not an energy scheduler	Drops cleanly beneath DREAM’s epoch timing; DREAM supplies neutrality/survival with constant-time updates
Coala	2020	Reactive	Dynamic split/merge of tasks based on recent energy	Adapts task granularity, not sleep timing vs. budgets	Regulates epoch sleep to meet budgets; can co-exist to exploit energy packets more efficiently
Ocelot	2021	Reactive	Automatic I/O freshness and crash-consistent atomic regions	Semantics layer; no energy budgeting	Provides the energy policy while Ocelot safeguards I/O correctness within epochs
Per-task watermark scheduler	2022	Reactive	Event-driven per-task thresholds; priority and re-queue	Lacks global neutrality test; local thresholds only	Adds explicit neutrality window + survival mode with a single store proxy
Capacity over capacitance	2019	Reactive	Store sizing and invariant thresholds beat raw capacitance	Design guidance; not an online scheduler	Operationalizes these invariants into per-epoch decisions with tiny, fixed overhead
Cammarano et al.	2016	Predictive	Sliding-window/seasonal predictors with ENO control	Requires forecast maintenance / tuning	Attains neutrality/survival without forecasting or model upkeep
ASARIMA	2022	Predictive	Adaptive seasonal ARIMA for harvested power	Statistical model fitting and parameter adaptation on device	Forecast-free, interpretable control; deterministic MCU cost per epoch
LT-ENO family	2014–2019	Predictive	Long-horizon neutrality via astronomical models; theoretical utilization guarantees	Day-scale planning loop and offline sizing assumptions	Purely online, solver-free neutrality correction; portable across platforms with minimal telemetry
PreAct	2019	Predictive	Utility-aware predictive control with uncertainty handling (PID-style)	Needs forecasts, utility profiles, and uncertainty pipeline	Delivers robust autonomy with no exogenous inputs and one store-proxy measurement per epoch

6 Design methodology for energy autonomous embedded microsystems

6.1 Scope, assumptions and notation

This chapter builds directly on the unified, loss-aware node model already introduced: the environmental input is represented as a time series $\Phi(t)$; conversion through the input path yields conditioned electrical power at the storage interface; the storage state is tracked as energy $E(t)$ with a store-voltage proxy $V_{\text{store}}(t)$; the demand is the aggregate load power $P_L(t)$; and non-idealities enter through macroscopic parameters—conversion efficiencies, quiescent drains, leakage, equivalent series resistance (ESR), and protection thresholds with hysteresis. No new symbols are introduced unless explicitly defined locally. Throughout, we treat inputs and workloads as traces rather than relying on closed-form irradiance formulae, and we fold device-level effects into the effective input power so that sizing proceeds without double counting losses. This stance is intentionally technology-agnostic and aligns with the loss-aware, window-based perspective established earlier.

Feasibility is enforced over a design window $W = [t_0, t_0 + T_W]$ chosen to capture the dominant rhythms of both input and load (e.g., diurnal/seasonal structure for outdoor deployments or occupancy-driven patterns indoors). Long-term viability requires that stored energy does not drift across repeated windows, which is expressed by an energy-neutrality condition on the same window once conversion losses, quiescent terms, and leakage are included. The prior chapter formalized this window selection process and its neutrality check; we adopt that formalism here without re-introducing new notation.

Design targets are expressed through three mission-level metrics used both for sizing and for verification. First, neutrality compares harvested-after-losses to demanded-after-losses on W and precludes cumulative deficits. Second, survivability requires that trajectories remain within the admissible store band implied by UVLO/OVP and safe operating limits; in practice, this is enforced as inequalities on the store proxy. Third, availability measures the fraction of W during which required loads are served while survivability holds; its value depends not only on energy balance but also on the timing of bursts relative to storage dynamics and protection hysteresis. These metrics were defined precisely in the modelling chapters and are used here as constraints and evaluation criteria.

Two additional considerations are treated explicitly in this methodology because they dominate behavior at the microwatt–milliwatt scale. The first is transient deliverability: burst events must not violate rail-droop limits set by storage ESR and converter current capability, a constraint that couples burst envelopes to the admissible operating band $[V_L, V_H]$. The second is cold-start and brown-out behavior: bring-up from a discharged store and fallback under energy poverty are governed by hysteretic thresholds and start-up energy budgets; these appear as state-space boundaries rather than ordinary loss terms and are checked alongside neutrality and availability during trace replay.

Finally, all guarantees are grounded in evidence from trace-driven replay through calibrated macromodels. Modelling cadence is chosen small relative to the storage time constant and the fastest burst features, whereas runtime telemetry remains sparse by design to maintain negligible overhead while preserving correctness of the energy accounting. The remainder of this chapter operationalizes these assumptions into sizing rules for the harvester and the accumulator, deliverability and protection checks, and a reproducible workflow with sensitivity analysis.

6.2 Input traces and window selection

All sizing and verification in this chapter proceed from time-series traces rather than closed-form irradiance models. The environmental input is supplied as an electrical-power trace at the storage interface (i.e., after conversion/protection), and the workload is supplied as a load-power trace that preserves burst structure; macroscopic non-idealities (conversion efficiencies, quiescent draws, leakage) are folded into these traces so that energy accounting is performed on the same side of the power path without double counting. This node-centric stance matches the architecture established earlier—harvested power $P_H(t)$, load power $P_L(t)$, stored energy $E(t)$ with a measurable proxy $V_{\text{store}}(t)$, and state-dependent efficiencies—keeping the model lightweight while capturing the losses that dominate at the μW – mW scale.

The design window $W=[t_0, t_0+T_W]$ is chosen to capture the slowest rhythms that materially shape deficits. When input and load periods are commensurate, T_W is taken as their least common multiple; when they are not, T_W spans a long representative horizon so that all patterns affecting feasibility are included. For outdoor photovoltaics, pronounced seasonality makes the maximum repetition time $T_{\text{cycle}} \approx 1$ year, and sizing should cover a full year to include worst-case seasons. Regardless of the domain, the fundamental requirement

is no drift in stored energy across repeated windows: each W should close with a non-negative change in the state of charge when losses are included, and in steady operation the expected drift is set to zero.

Neutrality on a window is evaluated by the input-side balance already defined in the thesis: harvested-after-losses versus demanded-after-losses over W . Using the macroscopic terms introduced previously, the window energies are written as integrals of the corresponding powers; neutrality holds if and only if the harvested energy meets or exceeds the requirement on the same window. This alignment of sides (both referred to the input path) prevents double counting and ensures that storage terms enter once, via the accumulator contribution.

Because feasibility depends on temporal structure, not just averages, traces must be prepared with care. For modelling/replay, the integration step is chosen small relative to the storage time constant and the shortest burst features, so numerical error remains negligible; for runtime telemetry, a much coarser cadence is admissible to keep sensing overhead negligible. Measured traces are screened and synchronized: timestamps are normalized to a single time zone and epoch; short gaps (relative to the store constant) may be interpolated, while long gaps are conservatively filled with lower confidence bounds so feasibility is not overstated. Load traces remain non-negative and aligned to activity boundaries. These practices make the ensuing neutrality and availability checks evidence-bearing rather than optimistic.

This trace-driven stance is deliberate. Replaying representative input and load windows through measured converter and storage macromodels surfaces cold-start traps, lull-induced depletions, and retry storms that average-power analyses miss; it also yields defensible harvester area/gain, storage capacity, and threshold placements. Where measured traces are scarce, lightweight synthetic generators that preserve plateau lengths, lull distributions, and other mission-relevant temporal statistics can still guide design effectively—provided their windowing matches the chosen T_w .

Finally, because the thesis already formalizes $\Phi(t)$ and the conversion from incident flux to input-side electrical power, Chapter 6 does not re-derive early irradiance equations; instead, it consumes the input-side power trace $P_{in}(t)$ produced by those definitions and applies the window-based neutrality test above. This preserves consistency with your earlier modelling while keeping the methodology focused on design windows, replay, and sizing.

6.3 Harvester sizing from window neutrality

This section determines the minimum collector aperture A_{min} that closes the energy balance over a representative window $W = [t_0, t_1]$, while also exposing the instantaneous power deficit and cumulative energy deficit used later in verification.

Methodology

This section formalizes a window-based, source-agnostic model for energy-autonomous embedded systems. The model couples (i) a compact description of the ambient input power, (ii) a possibly intermittent application load, and (iii) macroscopic conversion and storage efficiencies. From these ingredients we derive closed-form conditions for energy-neutral operation and sizing rules for both the harvester (collector area) and the accumulator (storage capacity) under an availability target.

Approach for harvester sizing

i) Operating point dependence

In the macroscopic, loss-aware model, efficiencies may depend on the operating point rather than being constant. We therefore write

$$\eta_H = \eta_H(u(t), x(t)), \quad \eta_{dc} = \eta_{dc}(P_{rail}(t), V_{in}(t), V_{out}(t), x(t)) \quad (6.1)$$

where $u(t)$ denotes the external stimulus seen by the harvester (e.g., intensity, frequency, gradient, field strength) and $x(t)$ collects internal conditions (e.g., temperature, bias, mode). The effective powers are

$$P_{in,eff}(t) = \eta_H(\cdot)P_{raw}(t), \quad P_{req,store}(t) = \frac{P_L(t)}{\eta_{dc}(\cdot)} + P_{q,out} \quad (6.2)$$

Harvesters and conversion paths often exhibit different efficiencies across operating regimes (e.g., low vs. high stimulus level, off-resonance vs. on-resonance, light-load vs. heavy-load converter modes). Representing η_H and η_{dc} as operating-point-dependent functions captures these effects without changing the form of the energy balance.

ii) Application load

Let $P_L(t)$ denote the instantaneous application power. When the workload is structured as repeating intervals (P_{Li}, T_{Li}), $P_L(t)$ is periodic with cycle $T_{L,cycle}$. The mean load over a cycle is obtained by averaging $P_L(t)$ on $T_{L,cycle}$ (Eq. (6.3)). This abstraction accommodates duty-

cycled sensing, sporadic communications, and burst compute phases without fixing implementation details:

$$P_{L,avg} = \frac{1}{T_{L,cycle}} \int_0^{T_{L,cycle}} P_L(t) dt = \frac{\sum_i P_{Li} \cdot T_{Li}}{\sum_i T_{Li}} \quad (6.3)$$

iii) Heteroperiodic inputs and analysis window

When input power and load have different periods (e.g., diurnal PV vs. weekly maintenance bursts), they must be treated as independent time series. Energy balance is then enforced over a window W that captures both rhythms. If the periods are commensurate, W is set to their least common multiple; otherwise, W spans a sufficiently long horizon so that all patterns that materially affect energy are included. For outdoor PV, strong seasonality dictates $T_{cycle} \approx 1$ year; therefore, W should cover a full year for sizing with respect to worst-case seasons. The fundamental requirement is that the store's state-of-charge (SoC) at the end of each W does not drift downward, ensuring long-term survivability.

iv) Electrical power at the input stage

The harvested electrical power delivered by the input stage is modeled as:

$$P_{in}(t) = \eta_H \cdot A \cdot \Phi(t) \quad (6.4)$$

where A is the collector area and $\eta_H \in (0,1)$ is the macroscopic input-side efficiency (transducer + front-end). For $\Phi(t)$ one may use Eq. (3) (or the detailed Eq. (1)/(2) when needed). This abstraction folds ripple and tracking losses into η_h for sizing purposes and keeps the model lightweight.

v) System lifetime

The total lifetime of the system, T_{life} , defines the total time specified for the system to be operational. The energy analysis will be performed in a smaller (more representative) window. However, it is important that the system should not allow an energy deficit to accumulate over time because at some point the system will shut down.

vi) Maximum repetition time

T_{cycle} is defined as the largest repeating cycle that affects the available input $\Phi(t)$. The choice of T_{cycle} is influenced by the form of energy at the system input.

In the case of outdoor photovoltaics, the power flow has a 24-hour form. However, the power flow changes significantly due to seasonality. Therefore, in this case, $T_{\text{cycle}} = 1$ year is selected.

vii) Output energy calculation

The energy required for the load (at its input) is defined as:

$$E_{out,W} = \int_0^W P_L(t) dt = P_{L,avg} \cdot W \quad (6.6)$$

The energy required by the system, considering the efficiency of the individual systems, is defined as:

$$E_{req,W} = \int_0^W \frac{P_L(t)}{\eta_{dc}} dt = \frac{P_{L,avg} \cdot W}{\eta_{dc}} \quad (6.7)$$

Where η_{dc} is defined as the total power efficiency of the system. ορίζεται η συνολική απόδοση διαδρομή της ισχύος στο σύστημα.

$$\eta_{dc} = \eta_{PMU} \times \eta_{PP} \times \eta_{store} (\leq 1) \quad (6.8)$$

viii) Energy balance

Define the output-side demand referred to the input side as

$$P_{out}(t) \triangleq \frac{P_L(t)}{\eta_{dc}(t)} + P_{q,out}(t) \quad (6.9)$$

and define the charge power seen at the input side of the storage path as

$$P_{charge}(t) \triangleq \frac{1}{\eta_{ch}(t)} \left(\dot{E}(t) + P_{leak}(V_{store}(t)) \right) \quad (6.10)$$

Then the node balance at the input side is

$$P_{in}(t) \geq P_{out}(t) + P_{charge}(t) \quad (6.11)$$

(with equality when there is neither protection-driven shedding nor deliberate dumping).

Over any analysis window $W=[t_0, t_1]$, define

$$E_{in} \triangleq \int_{t_0}^{t_1} P_{in}(t) dt, \quad E_{load} \triangleq \int_{t_0}^{t_1} \left(\frac{P_L(t)}{\eta_{dc}(t)} + P_{q,out}(t) \right) dt \quad (6.12)$$

The accumulator (input-side) term is

$$E_{acc} \triangleq \int_{t_0}^{t_1} P_{charge}(t) dt = \frac{1}{\eta_{ch}^{eff}} \left(\Delta E + \int_{t_0}^{t_1} P_{leak}(V(t)) dt \right) \quad (6.13)$$

with $\Delta E = E(t_1) - E(t_0)$. The window balance becomes

$$E_{in} \geq E_{load} + E_{acc} \quad (6.14)$$

which guarantees the storage level over W does not decrease (or increases by ΔE if targeted).

ix) Calculation of minimum collector surface area

Assume the raw input scales with the area A , so $P_{raw}(t) = A \phi(t)$. Over a window $W=[t_0, t_1]$ with $\Delta E = E(t_1) - E(t_0)$, enforcing energy neutrality with storage losses gives the required minimum area below.

$$A_{min} \geq \frac{\Delta E_{target} + \int_{t_0}^{t_1} \left(\frac{P_L(t)}{\eta_{dc}(t)} + P_{q,out} + P_{leak}(V(t)) \right) dt}{\eta_{ch} \int_{t_0}^{t_1} \eta_H(t) \phi(t) dt} \quad (6.15)$$

Approach for accumulator sizing

i) Period separation

We divide the period under consideration T_s into:

- Operational (up) time T_{op} -> System is operating
- Blackout time T_{blk} -> System is not operating

With

$$T_s = T_{op} + T_{blk} \quad (6.16)$$

We define

$$n_{av} = \frac{T_{op}}{T_s} \in [0,1] \quad (6.17)$$

ii) Energy balance equation

Defining $R(t)$ as

$$R(t) = \frac{P_L(t)}{\eta_{dc}} \quad (6.18)$$

We consider the accumulated energy $E(t) \in [0, C]$ (J), where C the maximum usable capacity.

$$\frac{dE}{dt} = \begin{cases} \eta_{ch}(P_{in}(t) - R(t)), & \text{av } P_{in} \geq R \text{ (charging)} \\ \frac{P_{in}(t) - R(t)}{\eta_{dis}}, & \text{av } P_{in} < R \text{ (discharging)} \end{cases} \quad (6.19)$$

Where η_{ch} and η_{dis} the charging and discharging performance of the energy storage device. The system reaches saturation at the limits $E=0$ and $E=C$. (When the storage device is full, the system rejects the excess energy).

iii) Calculating E_{in} (for given A)

On the collector's side, it is calculated as:

$$E_{in} = \int_{t_0}^{t_0+\Delta} P_{in}(t) dt = \int_{t_0}^{t_0+\Delta} \eta_H \cdot A \cdot \Phi(t) dt \quad (6.20)$$

The energy required by the system (on the same side) is:

$$E_{req}(\Delta) = \int_{t_0}^{t_0+\Delta} R(t) dt = \int_{t_0}^{t_0+\Delta} \frac{P_L(t)}{\eta_{dc}} dt \quad (6.21)$$

The local deficit is calculated as:

$$E_{def} = \max\{0, E_{req} - E_{in}\} \quad (6.22)$$

iv) Minimum C to avoid blackouts (with empty accumulator)

We define the net balance at time t as :

$$\delta(t) = P_{in}(t) - R(t) \quad (6.23)$$

- $\delta(t) > 0$: surplus $\rightarrow$ accumulator charging
- $\delta(t) < 0$: deficit $\rightarrow$ accumulator discharging

We define the cumulative difference as:

$$S(t) = \int_0^t (P(t) - P_{in}(\tau))d\tau, \quad S'(t) = R(t) - P_{in}(t), \quad S(0) = 0 \quad (6.24)$$

- When the system has a deficit ($R > P_{in}$), $S(t)$ rises
- When the system has a surplus ($R < P_{in}$), $S(t)$ falls

With initial charge $SoC(0)=0 \Rightarrow E(0)=0$, the stored energy is :

$$E(t) = \Pi_{[0,C]}(E(0) - S(t)) = \Pi_{[0,C]}(-S(t)) \quad (6.25)$$

Where $\Pi_{[0,C]}(x) = \min\{C, \max\{0, x\}\}$ is the saturation at $[0, C]$.

- If at that moment $S(t)$ is positive, then $-S(t) \leq 0 \Rightarrow E(t) = 0$ (empty accumulator)
- If $S(t)$ is negative, then $-S(t) > 0 \Rightarrow E(t) > 0$

Definition of blackout

The blackout index is defined as:

$$b(t) = 1(E(t) = 0 \wedge R(t) > P_{in}(t)) = 1(E(t) = 0 \wedge S'(t) > 0) \quad (6.26)$$

In other words, we have a blackout when:

- The accumulator is empty $E(t)=0$

And

- At that moment there is energy deficit ($R > P_{in}$, so $S' > 0$)

The total blackout time T_s is:

$$T_{blk} = \int_0^{T_s} b(t)dt \quad (6.27)$$

And the availability:

$$n_{av} = 1 - \frac{T_{blk}}{T_s} \quad (6.28)$$

We want the following to always hold true for some initial $E(0)=E_0$:

$$0 \leq E(t) = E_0 - S(t) \leq C \quad \forall t \in [0, W] \quad (6.28)$$

This is equivalent to two inequalities:

$$E_0 \geq \max_t S(t) \quad (\text{not to fall below zero}) \quad (6.29a)$$

$$C \geq E_0 - \min_t S(t) \quad (\text{not to overflow over } C) \quad (6.29b)$$

For minimum C , we take the smallest possible E_0 , i.e. $E_0^* = \max S$ then:

$$C_{min} = \max S - \min S \quad (6.30)$$

With this choice, the minimum value of $E(t)$ becomes exactly zero (when $S = \max S$) and the maximum becomes C_{min} (when $S = \min S$).

$$C_{min} = \max_k S[k] - \min_k S[k] \quad (6.31)$$

$$E_0^* = \max_k S[k] \quad (6.32)$$

Distinct form (discrete time, steps of Δt)

Let the analysis horizon be $[0, T_s]$ and define discrete time steps

$$t_k = k \Delta t, k = 0, 1, \dots, K, T_s = K \Delta t \quad (6.33)$$

Let

- $P_{in}[k]$ be the input (harvested) power referred to the storage side,
- $R[k]$ be the required power at the same side (load demand plus losses), and
- C be the maximum usable energy of the accumulator.

1. Instantaneous net balance (power deficit/surplus)

$$\delta[k] = R[k] - P_{in}[k] \quad (6.34)$$

- $\delta[k] > 0$: deficit $\Rightarrow$ accumulator discharging
- $\delta[k] < 0$: surplus $\Rightarrow$ accumulator charging

2. Cumulative energy difference (surplus/deficit)

$$S[k] = \sum_{n=0}^k \delta[n] \Delta t \quad (6.35)$$

When the system has a deficit ($R[k] > P_{\text{in}}[k]$), $S[k]$ increases; when it has a surplus, $S[k]$ decreases.

3. Stored energy with saturation

For an initial energy $E[0] = E_0$, the stored energy is

$$E[k] = \Pi_{[0,C]}(E_0 - S[k]) \quad (6.36)$$

where

$$\Pi_{[0,C]}(x) = \min\{C, \max\{0, x\}\} \quad (6.37)$$

enforces the physical limits of the accumulator.

In the special case of an initially empty store, $E_0 = 0$, so

$$E[k] = \Pi_{[0,C]}(-S[k]) \quad (6.38)$$

4. Blackout indicator and availability

A blackout occurs at step k when

- the accumulator is empty: $E[k] = 0$, and
- There is an energy deficit: $\delta[k] > 0$.

Define the blackout indicator

$$b[k] = \begin{cases} 1, & \text{if } E[k] = 0 \text{ and } \delta[k] > 0 \\ 0, & \text{otherwise} \end{cases} \quad (6.39)$$

The total blackout time over the horizon is

$$T_{\text{blk}} = \sum_{k=0}^K b[k] \Delta t \quad (6.40)$$

and the availability is

$$n_{\text{av}} = 1 - \frac{T_{\text{blk}}}{T_s} \quad (6.41)$$

Ignoring any upper capacity limit, the stored energy evolves as $E[k] = E_0 - S[k]$. To avoid blackout we need $E[k] \geq 0$ for all k , which implies

$$E_0 \geq S[k] \forall k \Rightarrow E_0^* = \max_k S[k] \quad (6.42)$$

When the accumulator has a finite usable capacity C and the initial energy is not fixed, the stored energy must remain in $[0, C]$. This requires that the entire excursion of the cumulative deficit can be embedded within the interval $[0, C]$, which leads to the minimum capacity

$$C_{\min} = \max_k S[k] - \min_k S[k] \quad (6.43)$$

Intuitively, E_0^* captures how much pre-charge is needed to survive the worst cumulative deficit, while $C_{\min}$ captures how much capacity is needed to ride out both deficits and surpluses without hitting empty or full.

The expressions for E_0^* and $C_{\min}$ characterize, respectively, the minimum initial energy and the minimum capacity required to avoid blackout over a given horizon when both can, in principle, be chosen freely. In practice, however, many deployments start from a nearly empty store and rely on future harvested energy to build up their state of charge. It is therefore useful to specialize these general results to the important case of an initially empty accumulator and study how the blackout-free regime and availability depend on the chosen capacity.

Minimum C to avoid blackouts (empty accumulator)

We define the energy in the accumulator when it is full as (from equation 23):

$$E(t) = \Pi_{[0,C]}(C - S(t)) \quad (6.44)$$

Solving C as a function of a goal n_{av}^*

For $C_2 \geq C_1$ applies:

$$\Pi(C_2 - S) = \Pi(C_1 - S) \Rightarrow E_{C_2}(t) \geq E_{C_1}(t) \forall t \quad (6.45)$$

So, the $n_{av}(C)$ is a non-decreasing function of C .

We define:

$$C^*(n_{av}^*; A) = \inf \{C \geq 0 : n_{av}(C; A) \geq n_{av}^*\} \quad (6.46)$$

Due to monotonicity, you can reliably find it by bisecting in C .

If the goal is $n_{av}=1$ in a full window W , so that it never empties:

$$E(t) = C - S(t) > 0 \forall t \in [0, W] \Leftrightarrow S(t) < C \forall t \quad (6.47)$$

Therefore, it is necessary:

$$C \geq S_{\max}, \quad S_{\max} = \max_{t \in [0, W]} S(t) \quad (6.48)$$

Energy upper limit on availability (for any C)

The maximum energy that can be given to the load in T_s is $\leq C + \int_0^{T_s} P_{in} dt$

The energy required for 100% uptime is $\int_0^{T_s} R dt$

So:

$$n_{av}(C) \leq \min \left(1, \frac{C + \int_0^{T_s} P_{in} dt}{\int_0^{T_s} R dt} \right) \quad (6.49)$$

Effect of A

Increase in A $\Rightarrow$ increase in $P_{in} \Rightarrow$ decrease in $S'(t) = R - P_{in} \Rightarrow$ the $S(t)$ is lower everywhere
 $\Rightarrow$

- Smaller $S_{\max} \Rightarrow$ smaller C for $n_{av} = 1$
- Smaller fluctuations $\Rightarrow$ smaller $C^*(n_{av}^*; A)$ for any target

Is 2 x C sufficient:

It depends on the value of C calculated in d.iv:

- If we used range $C_{iv} = \max S - \min S$:

With a full start, the smallest number $S_{\max} \leq C_{iv}$ is sufficient. 2 x C is conservative.

- If we used the maximum deficit from zero start $D_{\max} = \max S$:

For a full start, $C \geq S_{\max} = D_{\max}$ for $n_{av} = 1$ in the cycle is sufficient. The «2 x $D_{\max}$ » is not a general rule. In some profiles it is an exaggeration, in others it may not be enough when considering a larger horizon or when energy neutrality is required in successive cycles.

Best practice:

- For $n_{av} = 1$ in a full cycle with $\text{SoC}(0) = 100\%$:

$C_{\text{req}} = S_{\text{max}}$ and we also check that $S(W) \leq 0$ to ensure that the system is viable in repetitions.

- For a general target $n_{av}^* < 1$:

We find $C^*(n_{av}^*; A)$ by dichotomy (since $n_{av}(C)$ is monotonic)

7 Case studies and experiments

Chapter 7 grounds the thesis in model-driven and empirical evidence. Section 7.1 develops a parameterized Simulink model of an energy-harvesting autonomous sensor that is instantiated from an actual system and was used to evaluate the operation of DREAM, capturing harvester, storage, and load dynamics under realistic inputs. Section 7.2 presents three research projects as case studies, demonstrating how the methodology transfers across distinct hardware, workloads, and deployment contexts, and analyzing the resulting performance and trade-offs. Section 7.3 reports a targeted experiment on thermoelectric generators at low temperature differentials (ΔT), characterizing conversion efficiency, start-up behavior, and stability, and relating the findings back to the model and controller behavior. Together, these components provide convergent validation of DREAM’s practicality and limits, and yield guidance for sizing and control in real deployments.

7.1 Energy balance modeling techniques for embedded systems

7.1.1 Purpose

This subsection introduces a parameterized MATLAB/Simulink model that emulates the end-to-end operation of an energy-harvesting autonomous sensor and embeds the DREAM controller in the loop. The model is instantiated from an actual system: all key parameters (harvester/converter characteristics, protection thresholds, storage limits, and load profiles) are derived from measurements on the reference platform. By mirroring the real power path—harvester $\rightarrow$ conversion & protection $\rightarrow$ storage $\rightarrow$ load—the model enables controller-in-the-loop studies in which DREAM’s decisions act on the same signals (e.g., power-good, UVLO/OVLO, store proxy) available on hardware.

The scope is twofold. First, the model supports design-space exploration and sensitivity analysis without obscuring effects from cycle-accurate device physics: converters are captured with macroscopic efficiency, cold-start behavior, and quiescent drains; storage follows an energy-centric representation $E = \frac{1}{2}CV^2$; the workload is profiled as active quanta and radio bursts. Second, it provides a repeatable evaluation harness for DREAM, exercised under controlled ambient traces and perturbations to examine neutrality targets, survival behavior, and protection compliance before (and alongside) laboratory trials.

7.1.2 Reference hardware and measured parameters

The Simulink model is anchored to the evaluation setup exactly as it was built: the energy input is a measured power trace $P_{in}(t)$ in milliwatts/microwatts; there is no modeled power path or converter, and thus no efficiencies, quiescent drains, cold-start, or UVLO behavior. Storage is a single ideal supercapacitor of capacitance C ; leakage, ESR, voltage derating, and temperature effects are intentionally excluded. The load is realized as a resistor network that the MCU switches between two discrete operating states—ACTIVE and LPM—so that instantaneous consumption follows $P_{load}(t) = V^2(t)/R$ with R equal to the branch in use. The DREAM controller runs on the MCU; at each epoch boundary the MCU samples the capacitor voltage $V(t)$ as the sole proxy for stored energy, updates DREAM's recurrence, schedules a fixed-duration active burst, and then sleeps for the computed interval before repeating.

All parameters used in simulation are taken directly from this setup unless explicitly varied for sensitivity analysis. The input block consumes the recorded $P_{in}(t)$ at its native sampling interval Δt and reproduces the trace verbatim. The storage block integrates net power into the capacitor state using the relations $E(t) = \frac{1}{2} CV^2(t)$ and $\frac{dE}{dt} = P_{in}(t) - P_{load}(t)$, which yield $\frac{dV}{dt} = [P_{in}(t) - P_{load}(t)]/[CV(t)]$. The load block implements the two modes by selecting either a low resistance for ACTIVE or a much higher resistance for LPM; if the low-power mode was effectively off in hardware, the model sets the branch open so that $P_{lpm} = 0$. The controller block executes DREAM with tunable survival horizon H , smoothing weight w (if used), a minimum sleep floor, and an epoch timer that defines the cadence of active bursts and sleep intervals; these timings mirror the firmware configuration.

For reproducibility, the baseline instantiation records the provenance of each quantity: the power trace is measured under the intended environmental conditions and imported unchanged; the capacitance C is taken from the component used on the bench (optionally verified by a simple charge–discharge slope); the resistances defining ACTIVE and LPM are the installed values as measured with a DMM at room temperature; the active-burst duration and epoch timing are those executed by the MCU firmware during the experiments. Because the model omits any conversion stage, protection appears only as reporting thresholds $[V_{min}, V_{max}]$ against the capacitor voltage rather than enforced cutoffs. This minimalist mapping isolates the interaction between DREAM's scheduling decisions and the store–load dynamics, allowing the evaluation to focus on controller behavior under a

known input power profile without confounding effects from conversion losses or protection hysteresis.

7.1.3 Model architecture

The model consists of four blocks connected by one power signal and one voltage signal, plus a discrete control line; it is organized for inspection and traceability rather than device-level detail. The top-level diagram wires the four elements of the minimal system. Energy In (DC Bus) injects the measured harvested power trace onto the bus. The Charge block is the ideal supercapacitor connected to the same node. The MCU reads the capacitor/bus voltage and runs DREAM; on each epoch it sets the digital enable line EN that drives Load, which applies the resistor corresponding to ACTIVE or LPM to the bus (V_{in}). The single bus node therefore closes the loop: input power charges the capacitor; the MCU observes the resulting voltage and schedules ACTIVE/LPM; the load's draw, in turn, shapes the bus voltage trajectory.




Figure 7.1 Simulation model (top level)

This subsystem feeds the model with the measured harvested power trace and enforces any optional physical caps specific to the transducer. The trace enters at left and is routed to a scope (“Power, V_{bus} ”) for verification. The function block “Power_limit_by_TEG” accepts the commanded operating voltage v (bus voltage) and the raw input u (trace power) and outputs a bounded input power consistent with low- ΔT TEG behavior (e.g., non-negative, limited by the module’s $V-I$ characteristic around its operating point). The selected output then drives the DC bus through a controlled source, with the bus voltage fed back to the limiter and logged. This block is where the recorded $P_{in}(t)$ enters the simulation and, if

desired, is constrained by device-level feasibility before being integrated by the storage model.




Figure 7.2 Energy Input block

The diagram shows the ohmic load that the MCU enables via a single digital line EN. The supply node “Load V_{CC}” fans out to two grounded branches: RLPM1 (low-power mode) and RTx (active mode). A controlled switch selects which branch is connected; when EN=0 the circuit sinks through RLPM1 (very high resistance, emulating sleep draw or OFF if taken $\rightarrow \infty$), and when EN=1 it sinks through RTx (low resistance, emulating ACTIVE power). Instantaneous consumption follows $P_{\text{load}}(t) = V_{\text{load}}^2/R$ with $R \in \{R_{\text{LPM}}, R_{\text{act}}\}$. This block provides the only path from the storage node to ground and therefore sets the epoch’s active-burst power versus the background draw in sleep. Parameters extracted from this schematic are R_{act} , R_{LPM} (or OPEN), and the EN polarity.




Figure 7.3 Load simulation block

This subsystem implements the discrete controller that schedules ACTIVE/LPM. The central MATLAB Function block (tsleep_energy_management_fcn) receives the current

capacitor voltage V_{cap} (via the “uC I/O” block), the previous sleep duration (t_{sleep_prev}), and the configured survival horizon ($t_{survival}$). An interrupt trigger marks epoch boundaries; on each trigger the function computes the next sleep time t_{sleep_int} and updates internal state outputs (V_{old_out} , $t_{survival_out}$). A Pulse generator produces the fixed-duration ACTIVE burst (period/dc) and drives the digital control line to the load. “Memory register #1–#3” hold the minimal history DREAM needs (previous sleep, previous voltage, last decision), ensuring constant-time updates. This block therefore realizes the controller-in-the-loop evaluation: read $V_{cap} \rightarrow$ update DREAM $\rightarrow$ emit next sleep T_s and the ACTIVE pulse.




Figure 7.4 Microprocessor simulation block

7.1.4 Simulation configuration and run procedure

Each simulation run is defined by four user-set parameters: capacitor start voltage V_{init} , survival time $T_{survival}$ (the horizon DREAM must guarantee), the input-power profile $P_{in}(t)$ (selected trace and any scaling), and the simulation duration T_{sim} . After setting these, the simulation is started with these initial values and DREAM enabled on the MCU block. The controller samples the capacitor voltage at epoch boundaries, computes the next sleep interval, and the load follows the issued ACTIVE/LPM schedule. During the run we observe the system in real time via scopes (voltage trajectory, mode flag, and sleep sequence) to qualitatively assess behavior under the chosen configuration.

To ensure the reliability of our experiment, we conducted several trials under varying conditions to evaluate the algorithm’s performance. During the trials, we monitored the

board's behavior and recorded the data generated by the system. The results of our experiment demonstrated that the algorithm performed efficiently, and the board functioned reliably under different scenarios.

7.1.5 Simulation examples

Reporting thresholds $V_{\min} = 2.2V$ and $V_{\max} = 3V$ are monitored (not enforced), with controller bounds $T_s^{\min} = 20s$ to avoid oversampling and $T_s^{\max} = 5000s$ for deep sleep.

- Example A (Charging circle)

The capacitor is initialized just above $V_{\min}$. At $t = 0$, the measured input power is present, so the store voltage $V_C(t)$ rises monotonically with a concave-down envelope. Because DREAM samples V_C as the sole proxy for state of charge, the very low initial voltage yields a large, computed sleep time $T_s(0)$ to conserve energy. As V_C increases, DREAM progressively shortens T_s , allowing more frequent ACTIVE bursts while preserving the configured survival horizon T_{survival} . After a brief transient, the system reaches a quasi-steady regime in which average input and average load powers balance: $V_C(t)$ flattens, and T_s converges to the configured minimum sleep (a small floor used to avoid oversampling and unnecessary wakeups). No threshold crossings or chatter are observed, indicating stable operation in a charging/neutral state with a fixed duty cycle.

- Example B (Discharge/Survive circle)

The capacitor begins near 50% SoC and the system charges normally until $t = 10$ ks. During this interval, $V_C(t)$ rises and DREAM correspondingly decreases T_s , mirroring Example A. At $t = 10$ ks the input power is forced to zero and remains off until $t = 30$ ks. Immediately after the cut, DREAM reacts by lengthening T_s to maintain the survival horizon H ; ACTIVE events become sparse and $V_C(t)$ decays smoothly but remains above $V_{\min}$ for the required time. When input resumes at $t = 30$ ks, the controller shortens T_s again to rebuild headroom while still guaranteeing that the node can survive at least 25 ks should another outage occur. As the run proceeds, the accumulated deficit and conservative pacing eventually drive the system to the controller's maximum sleep limit; by $t = 65$ ks the node enters deep sleep (LPM/OFF held), and $V_C(t)$ asymptotically approaches its floor without oscillation.




Figure 7.5 (top) Capacitor voltage (bottom) Calculated sleeptime




Figure 7.6 (top) Capacitor voltage (bottom) Calculated sleeptime

Simulation outputs were compared against laboratory recordings on the same setup. Agreement is assessed on (i) the capacitor-voltage envelope $V_C(t)$, (ii) timing and duration of ACTIVE bursts, and (iii) the sleep-time sequence $\{T_s\}$. Runs are accepted when $V_C(t)$ stays within ± 20 mV of the measured envelope, ACTIVE timing matches within one epoch, and $\{T_s\}$ follows the same trend (monotone decrease/ increase and final plateau). (§7.2.1)

The model deliberately omits converter/power-path losses and protection hysteresis and treats the supercapacitor as ideal (no leakage/ESR). The load is a two-state resistive abstraction. These simplifications isolate DREAM’s scheduling decisions and the store–load interaction under a known $P_{in}(t)$.

7.2 Research projects

This section evaluates the proposed power management techniques together with the design methodology introduced in Chapter 6 across three prototypes that represent distinct operating envelopes: (i) a state-of-the-art power-supply module designed and implemented for energy-autonomous nodes, (ii) an autonomous weather station capable of reading the nearby electric field, and (iii) an autonomous steel-stress sensor for naval applications. Using the experimental setup of §7.1, we quantify energy neutrality and survivability (time outside protection bounds). For each prototype we perform trace-driven replay and in-situ tests, compare against battery-only and naïve harvesting baselines, and monitor the contribution of the runtime policies. The Chapter 6 methodology is referenced throughout to clarify architectural choices and implementation trade-offs.

7.2.1 Perpetual power supply-PERPS

Our objective in PERPS was to design and construct a low-power energy-harvesting power supply for embedded systems. The power supply should be capable of harvesting energy from different sources with high efficiency. PERPS served as the seed project for the work in this thesis: it informed the design methodology of Chapter 6 and provided the first deployment target for the DREAM runtime algorithm, which was later generalized to the other prototypes.

Designing an ultra-low-power energy-harvesting supply is intrinsically complex due to cold-start requirements, widely varying source impedances/voltages, leakage and self-consumption constraints, and the need to buffer burst loads while maintaining regulated delivery. Figure 7.7 summarizes the chosen architecture: a single ambient source—

photovoltaic (PV), thermoelectric module (TEM), triboelectric nanogenerator (TENG), or piezoelectric (PZ)—is connected to the harvester front-end. Harvested energy is stored in the accumulator, which buffers short-term power imbalance, while the harvester regulates the output node V_{out} to a configurable setpoint. Only one source is present at the input at any given time; source selection is therefore external to the module. This partitioning isolate source-specific conversion from energy storage and from the regulated output, simplifying control while preserving flexibility in the target application.




Figure 7.7 Power supply architecture




Figure 7.8 PERPS board

Figure 7.8 shows the PERPS board. It follows an Arduino-Uno “shield-style” layout so it can sit on top of a common microcontroller board or be used on its own. The connectors are arranged to make wiring straightforward: one input for the chosen energy source, terminals for the accumulator, and a regulated V_{out} for the load. Key points are labeled for easy probing during experiments, but detailed component information is intentionally omitted.

We validated PERPS in a representative deployment powering a low-power XBee sensing/telemetry node. The energy source was a small indoor photovoltaic (PV) cell placed under the room’s lamps; the power path followed the architecture of Figure 7.1, with the harvester regulating V_{out} and buffering surplus energy in the accumulator. In the original demonstration, the chosen PV cell provided ample headroom and the system operated perpetually (energy-neutral with positive storage drift over a 24-hour cycle). Despite the absence of a formal methodology at the time, perpetual operation was achieved because both the PV cell and the storage capacitor were deliberately oversized.




Figure 7.9 (left) Device test in the lab (right) Demo device in situ

We apply retroactively the Chapter 6 window-neutrality method to a concrete indoor-light, always-on scenario with the following parameters (all powers referred to the input/storage side of the power path):

- **Lighting:** 150–250 lux, approximately constant (24/7).

- **Load (periodic):** Active power $P_A = 42.9 \text{ mW}$ for $t_A = 0.19 \text{ s}$ once every $T = 600 \text{ s}$; sleep power $P_S = 0.009 \text{ mW}$ otherwise.
- **Harvester:** a-Si PV with an empirically measured maximum-power output $P_{mpp} = 0.432 \text{ mW}$ at 200 lx for an aperture area $A = 41.6 \text{ cm}^2$
- **Power path:** effective DC-path efficiency $\eta_{dc} \in [0.70, 0.85]$ over the realized operating region.
- **Accumulator:** supercapacitor 5 F , operated in the window $2.7 \text{ V} \rightarrow 2.2 \text{ V}$.

The mean load over one 10-minute cycle is:

$$\bar{P}_L = \frac{P_A t_A + P_S (T - t_A)}{T} = \frac{42.9 \cdot 0.19 + 0.009 \cdot (600 - 0.19)}{600} \text{ mW} = 22.58 \mu\text{W}$$

The energy per cycle: $E_{\text{cycle}} = \bar{P}_L T = 0.01355 \text{ J}$. The incremental TX-over-sleep energy is $E_{\text{burst}} = (P_A - P_S) t_A = 8.15 \text{ mJ}$.

Usable supercapacitor energy in the $2.7 \rightarrow 2.2 \text{ V}$ window is $E_{\text{usable}} = \frac{1}{2} C (C_{\text{max}}^2 - C_{\text{min}}^2) = \frac{1}{2} \cdot 5 (2.7^2 - 2.5^2) = 6.125 \text{ J}$. The small-signal voltage droop during the 0.19 s transmission is $\Delta V = \frac{E_{\text{burst}}}{C V_{\text{avg}}} = \frac{8.15}{5 \cdot 2.45} = 0.66 \text{ mV}$, which is negligible for deliverability.

From the 200 lux measurement $p_{PV,200} = P_{mpp}/A_{pv} = 0.432 \text{ mW}/41.6 \text{ cm}^2 = 10.385 \mu\text{W}/\text{cm}^2$ at MPP.

With constant lighting, neutrality reduces to mean powers:

$$A_{\text{min}} = \frac{\bar{P}_L}{\eta_{dc} p_{PV}} = \frac{22.58}{10.38 \cdot 0.8} = 2.71 \text{ cm}^2$$

The time to raise the supercapacitor from $2.2 \text{ V} \rightarrow 2.5 \text{ V}$ while operating:

$$t = \frac{E_{\text{usable}}}{P_{\text{surplus}}} = \frac{6.125 \text{ J}}{0.35 \text{ mW}} = 6.1 \text{ h}$$

Applying the methodology clarifies the sizing margins and shows that the original design was conservatively over-provisioned. Using the measured P_{mpp} density at 200 lx (about $10.4 \mu\text{W}/\text{cm}^2$) and accounting for η_{dc} , the same accumulator would permit a PV **area roughly ~16× smaller** while maintaining perpetual operation under the same lighting and usage conditions. If ambient illumination were reduced or became more variable, the

methodology would either select a modestly larger PV device or retain the smaller cell and adapt the telemetry cadence via DREAM to remain energy-neutral.

7.2.2 Electric Field over Seas-EFOS

EFOS targeted the development of an autonomous coastal weather station capable of measuring the ambient electric field using a rotating field-mill sensor.




Figure 7.10 (left) EFOS concept design (right) Pilot device

The principal objective was the realization of a resilient power subsystem that could harvest energy from heterogeneous sources, integrate a state-of-the-art triboelectric nanogenerator (TENG), and supply a regulated high-voltage path for the field-mill motor while maintaining stable rails for digital control and low-noise analog sensing. Deliberate right-sizing of the harvesters and the accumulator was outside the scope of this demonstrator; the emphasis was placed on architectural integration, cold-start capability, and survivability under intermittency.

The power supply follows a modular, source-decoupled topology.




Figure 7.11 EFOS power supply architecture

Each energy source—photovoltaic, TENG, thermoelectric, and vibration—interfaces to the system through a dedicated harvester front-end that performs rectification, impedance presentation, and, where appropriate, simple MPPT-like regulation. By isolating sources at the front-end, the design prevents back-drive and contention when one source is dominant or when others are quiescent. The conditioned outputs of the harvesters are then combined through a common buffer node that absorbs bursty or high-crest-factor inflows without imposing cross-coupling penalties. This buffer feeds a single battery-charger stage that manages the electrochemical store, exposing clear thresholds for cold-start, charge acceptance, and over-/under-voltage protection.

Downstream of the store, three independent DC/DC converters generate the rails required by the station. The first rail supplies the microcontroller and communications, prioritizing low quiescent consumption and good transient response during duty-cycled activity. The second rail biases the sensors and analog front-end of the electric-field readout chain; this path is treated as noise-sensitive, with sequencing and local decoupling to prevent coupling from digital bursts and motor events. The third rail is dedicated to the field-mill motor/driver and is allowed to operate at a higher instantaneous power level, with inrush control and undervoltage lockout to avoid brown-outs that could corrupt ongoing measurements. This separation of rails is central to maintaining measurement integrity while guaranteeing deliverability for the motor’s spin-up, which otherwise represents the most demanding transient in the system.




Figure 7.12 EFOS prototype

Additive manufacturing was used to design and fabricate a custom enclosure that integrated three outdoor-rated photovoltaic cells, each mounted on a distinct facet so that the array sampled different azimuths and captured incident light under varying ship headings and seasons. The printed geometry incorporated precise cell recesses, cable management channels, and gasketed interfaces for weather protection, while allowing rapid iteration of tilt and spacing to reduce mutual shading and improve structural rigidity with minimal mass. This enclosure was installed as part of the EFOS prototype on a research vessel operating in the Mediterranean Sea, where it experienced continuous vibration, salt-laden air, and wide temperature and irradiance swings. The assembly remained fully functional for more than six months at sea, after which it was removed intact and returned to the lab for post-deployment evaluation.



Figure 7.13 Installation of EFOS system

Unfortunately, the sizing methodology outlined in Chapter 6 could not be applied to EFOS. The deployment did not yield the continuous, source-resolved power measurements required to construct representative duty-cycle traces and energy windows; as a result, neither harvester apertures nor accumulator capacity could be right-sized from data. Any sizing would therefore have been speculative, and the system is reported as an integration prototype rather than an optimized design.

7.2.3 Autonomous Multifunctional Stress Assessment Sensors for Naval Applications-AMSA

The AMSA project addresses long-duration monitoring of stress in ship steel structures by integrating a proprietary anisotropic magneto resistive sensor into a compact node designed for autonomous operation. Energy autonomy shaped every architectural choice so that sensing fidelity and reporting cadence could be preserved without manual intervention.

The delivered platform adopts a harvesting path centered on a single power-management IC that accepts multiple energy inputs concurrently and manages storage and regulation. Two independent sources feed this controller so that available ambient energy can be aggregated rather than alternated, while the storage element at the controller's output provides both energy buffering and burst current. The functional organization is summarized in Figure 7.15, which illustrates a **single-PMIC, dual-source** topology capable of **simultaneous** harvesting into a common accumulator and regulated output.




Figure 7.14 AMSA architecture

Electrical realization follows directly from this organization. The schematic implements per-input conditioning for the two sources, the battery interface with protection and chemistry-appropriate charge termination, and the regulated system rail that powers the sensing front-end and communications. Care is taken to minimize quiescent draw in the quiescent state while ensuring adequate transient response during active intervals. This connectivity is shown in Figure 7.13, which preserves the one-PMIC philosophy while exposing clear test points for source characterization and storage health.




Figure 7.15 AMSA Harvester schematic




Figure 7.16 AMSA Board

The pilot devices were deployed on representative steel structures and benefited from enclosures produced via 3D printing, which enabled protective cases with complex geometries tailored to the mounting surfaces and cable exits. This approach allowed tight integration of the sensing head and harvesting surfaces while providing mechanical robustness and environmental shielding appropriate to shipboard conditions.




Figure 7.17 AMSA pilot devices

A supervisory routine governs the node's duty cycle at runtime, adapting the sleep interval to the available energy and the state of storage so that operation remains self-sustaining without altering the sensing chain or communications stack.

In this case we deployed the design methodology retroactively as well. We consider the following parameters for the system.

- **Lighting:** 200 lux, approximately constant (24/7).
- **Load (periodic):** Active power $P_A = 163 \text{ mW}$ for $t_A = 1.825 \text{ s}$ once every $T = 600 \text{ s}$; sleep power $P_S = 0.4125 \text{ mW}$ otherwise.
- **Harvester:** a PV cell with an measured output power of $10.4 \mu\text{W}/\text{cm}^2$
- **Power path:** effective DC-path efficiency $\eta_{dc} \in [0.8]$ over the realized operating region.

Accumulator: a LiFEPO4 battery with rated capacity of 1800mAh, operating in the window $3.0 \text{ V} \rightarrow 3.6 \text{ V}$.

By applying the equation 6.15 for the PV sizing, we calculate the required aperture of the PV cell. For a 10-minute cycle we set $\Delta E_{\text{target}} = 0$, $P_{q,\text{out}} = 0$, and take leakage as already counted inside $P_L(t)$. Indoor PV density is constant $\Phi(t) = \Phi = 10.4 \mu\text{W}/\text{cm}^2$; combined harvester/charging efficiency is $\eta_{\text{ch}}\eta_H = 0.8$.

Active: $P_a = 163 \text{ mW}$, $t_a = 1.825 \text{ s} \rightarrow E_a = P_a t_a = 0.082632 \text{ mWh}$.

Sleep: $I_s = 0.15 \text{ mA}$ at $V \approx 3.3 \text{ V} \rightarrow P_s = 0.495 \text{ mW}$; sleep time $T - t_a = 598.175 \text{ s} \rightarrow$

$$E_s = 0.082249 \text{ mWh.}$$

So,

$$\int P_L(t)dt = E_{cyc} = E_a + E_{sleep} = 0.164881 \text{ mWh} = 0.59357 \text{ J}$$

$$\begin{aligned} \eta_{ch} \int \eta_H \Phi(t)dt &= 0.8 \cdot \Phi \cdot T = 0.8 \cdot (10.4 \times 10^{-6} \text{ W/cm}^2) \cdot 600 \text{ s} \\ &= 4.992 \times 10^{-6} \text{ J/cm}^2 \end{aligned}$$

So,

$$A_{min} \geq \frac{0.59357 \text{ J}}{4.992 \times 10^{-3}} = 118.9 \text{ cm}^2$$

For the accumulator sizing, we size the accumulator for (1) TX burst support and (2) blackout autonomy. Depth-of-discharge limit DoD = 80% and regulated rail $\approx 3.3 \text{ V}$.

The burst energy in the active interval:

$$E_{burst} = E_\alpha = 0.082632 \text{ mWh} = 0.297 \text{ J}$$

The average burst current is

$$I_{burst,avg} = \frac{P_\alpha}{V} = \frac{163 \text{ mW}}{3.3 \text{ V}} = 49 \text{ mA}$$

For $t_a=1.825 \text{ s}$. A LiFEPO4 cell easily sources this. The minimum theoretical battery capacity set purely by burst energy is tiny.

$$C_{min,burst} = \frac{E_{burst}}{V} \cdot \frac{1}{3600} \cdot \frac{1}{DoD} = \frac{0.000082632 \text{ Wh}}{3.3 \text{ V}} \cdot \frac{1}{0.8} = 31 \mu\text{Ah.}$$

In practice, if we just ensure adequate low-ESR decoupling at the rail to limit droop during the $\sim 1.8 \text{ s}$ burst, the battery will handle the energy.

For the blackout autonomy, we consider a 10-minute cadence with average load power

$$\bar{P}_L = \frac{E_{cyc}}{T} = 0.989 \text{ mW}$$

The daily average energy $E_{\text{day}} = 0.989\text{mW} \times 24\text{h} = 0.02374\text{Wh}$. The daily charge at 3.3V is:

$$Q_{\text{day}} = \frac{0.02374}{3.3} = 7.2\text{mAh/day}.$$

The required accumulator for N dark days can be calculated by:

$$C_{\text{req}}(N) = \frac{Q_{\text{day}}N}{0.8}$$

e.g. $C_{\text{req}}(1) = 9.0\text{ mAh}$, $C_{\text{req}}(3) = 27.0\text{ mAh}$, $C_{\text{req}}(7) = 63.0\text{ mAh}$, $C_{\text{req}}(30) = 270\text{ mAh}$.

In the final prototype the installed PV area was slightly below the minimum predicted by Eq. (6.16), placing the node in a mild energy-deficit regime at the original 10-minute cadence. Rather than alter the hardware, the supervisory algorithm (DREAM) compensated in situ by lengthening the sleep interval to 13.24 minutes. This adaptive duty-cycle change reduced the average load to match the available harvest, restoring energy neutrality and enabling perpetual operation under the same indoor illumination.

7.3 Low- ΔT thermoelectric generator evaluation

7.3.1 Objective and scope

The objective of this work is to design and validate an automated experimental platform for the low- ΔT characterization of thermoelectric modules (TEMs), with controlled temperature differentials to approximately as low as 2 °C across the device. Unlike the extensively studied large- ΔT regime, small (to very small) temperature differences remain comparatively underexplored; the present setup targets this gap by enabling precise thermal control and high-fidelity electrical measurements relevant to near-isothermal operation.

The scope of Section 7.3 is to describe the mechanical, thermal, electrical, and software elements of this setup; detail the operating procedure used to generate stable $\Delta T \geq 2\text{ °C}$; and define the measurement outputs— V_{OC} , I_{SC} , R_{int} , and P_{MPP} —that will be used to evaluate TEM performance in the low- ΔT regime.

7.3.2 Experimental apparatus

The thermal subsystem was designed to establish and maintain small, well-controlled temperature differentials across the TEM while minimizing lateral gradients and ambient drift. Heating of the hot face was provided by a bank of four 10 W MICA ceramic resistors powered by a laboratory DC supply. Their distributed placement and low thermal inertia

enabled fine power trimming and uniform heat flux at the module interface, which is essential when operating in the 2–8 °C differential range. The cold face was actively refrigerated using a Peltier device operated conservatively to support low- ΔT studies. Because the Peltier’s effective thermal resistance is low and its operation adds heat to its own hot junction, a forced-convection heat sink was attached to that side to evacuate the additional load and stabilize the cold-side temperature during steady operation. The assembled sandwich of spreaders–TEM–spreaders is shown outside the enclosure in the following image.




Figure 7.18 Sandwich formation

To promote isothermality, 0.5 cm-thick aluminum plates were used as spreaders on both sides of the TEM. These plates increased in-plane thermal conductance and suppressed local hot and cold spots, improving the fidelity of the applied differential at the active junctions. Industrial-grade PT100 RTDs (readable to two decimal places) were embedded in shallow, machined recesses in the plates so that, upon assembly, the plates mated flush with the TEM in a sandwich configuration without sensor-induced gaps. The measured face temperatures are denoted T_{hot} and T_{cold} , and the working differential is defined as

$$\Delta T = T_{\text{hot}} - T_{\text{cold}}$$

To reduce parasitic losses to the environment and improve repeatability, the assembled sandwich was housed in a custom 3D-printed enclosure, shown in the following figures.




Figure 7.19 3D printed case

The case was fabricated from PETG, selected for its higher thermal durability and dimensional stability under modest temperature rise compared with common PLA filaments. By shielding the stack from ambient drafts and limiting direct radiative and convective exchange, the enclosure helped maintain stable boundary conditions at low ΔT while accommodating cable feedthroughs and mounting features for the spreaders and heat sink.

The electrical output of the TEM was routed to a custom perfboard that consolidated measurement and loading functions. The circuit topology of the perfboard—comprising the open-circuit voltage path to a precision DMM, the short-circuit current path through a calibrated ammeter/shunt, and a variable load implemented via a high-power potentiometer—is depicted in Figure 7.19, while the assembled perfboard hardware is shown in Figure 7.20. This arrangement enabled measurement of V_{OC} and I_{SC} and a continuous load sweep to locate the maximum power point (MPP) by trimming the potentiometer while logging voltage and current.




Figure 7.20 Perfboard circuit




Figure 7.21 Assembled perfboard

All hardware—power supplies for the resistor bank and the Peltier, measurement instruments, and safety interlocks—was controlled by a LabVIEW supervisory program implementing dual PID loops. The hot-side loop modulated the resistor power to track $T_{\text{hot,set}}$, and the cold-side loop modulated the Peltier input current to track $T_{\text{cold,set}}$. Within this configuration the platform can establish and hold $\Delta T \approx 2^\circ\text{C}$ at the lower limit; the results reported in this study focus on $\Delta T = 2\text{-}8^\circ\text{C}$. A photograph of the complete experimental arrangement, including the enclosure, heatsink, wiring, and instrumentation, is provided in Figure 7.21.




Figure 7.22 Complete setup

7.3.3 Control and user interface

All thermal and electrical subsystems were coordinated through a LabVIEW™ supervisory program that provides closed-loop temperature regulation and integrated data acquisition. The graphical user interface (GUI), shown in Figure x, allows the operator to set independent temperature setpoints for the hot and cold faces. Setpoints can be adjusted on-the-fly while the system is running; the control loops track the updated targets without requiring a restart. The GUI continuously displays the key electrical observables at the TEM—voltage, current, instantaneous power, and inferred resistance—alongside real-time plots of the measured hot- and cold-side temperatures. For completeness, the voltages delivered by the laboratory power supplies to the heater bank and to the Peltier stage can also be visualized, enabling the user to correlate actuator effort with thermal response.

The underlying control logic, summarized in Figure x, implements two independent feedback loops. The hot-side loop reads the corresponding PT100 sensor and modulates the resistor-bank supply to track $T_{\text{hot,set}}$, while the cold-side loop reads the cold-face PT100 and modulates the Peltier input to track $T_{\text{cold,set}}$. Temperature, actuator commands, and TEM electrical measurements are time-stamped and logged for post-processing and for constructing the $V-\Delta T$ and $P-\Delta T$ curves reported later




Figure 7.23 TEM evaluator GUI




Figure 7.24 LabView block diagram

7.3.4 Metrics

The platform reports both thermal-control quality indicators and device-level electrical figures of merit, each computed from the synchronized data streams acquired during steady operation and during controlled load sweeps. Thermal performance is summarized by the commanded face temperatures T_{hot} and T_{cold} and by their difference $\Delta T = T_{\text{hot}} - T_{\text{cold}}$.

Electrical characterization begins with the open-circuit voltage V_{OC} and the short-circuit current I_{SC} , obtained by switching between the corresponding measurement paths. In the low- ΔT linear regime, the effective module Seebeck coefficient is estimated as $\alpha_{\text{eff}} = V_{\text{OC}}/\Delta T$; when multiple ΔT levels are available, α_{eff} is more robustly obtained from the slope of a least-squares fit of V_{OC} versus ΔT . The internal electrical resistance of the device is inferred from the Thevenin relation as $R_{\text{int}} \approx V_{\text{OC}}/I_{\text{SC}}$. An independent estimate is obtained from the slope of the I - V characteristic recorded during load sweeps, and the two values are compared as a consistency check.

Load sweeps are performed by trimming the potentiometer while logging voltage and current at the module terminals; instantaneous electrical power is computed as $P = VI$ and mapped to the corresponding load resistance $R_L = V/I$. The maximum power point is defined by $(V_{\text{MPP}}, I_{\text{MPP}})$ with $P_{\text{MPP}} = V_{\text{MPP}}I_{\text{MPP}}$. Under an ideal Thevenin model, R_L at MPP equals R_{int} , and $V_{\text{MPP}} \approx V_{\text{OC}}/2$, $I_{\text{MPP}} \approx I_{\text{SC}}/2$, yielding $P_{\text{MPP}} \approx V_{\text{OC}}I_{\text{SC}}/4$. Departures from these equalities quantify non-idealities such as contact resistances, temperature drift during the sweep, or weak nonlinearity at the smallest ΔT .

We report $V_{\text{OC}}(\Delta T)$, $I_{\text{SC}}(\Delta T)$, and $P_{\text{MPP}}(\Delta T)$; when R_{int} is approximately constant over 2-8°C, P_{MPP} exhibits the expected near-quadratic scaling with ΔT . Each curve includes uncertainty bands obtained by propagating the measurement errors of temperature, voltage, and current, and repeatability is summarized by the spread of repeated runs at fixed ΔT . Together, the principal metrics are R_{int} , $V_{\text{OC}}(\Delta T)$, $I_{\text{SC}}(\Delta T)$, and $P_{\text{MPP}}(\Delta T)$.

7.3.5 Results

The high- ΔT reference (commercial module) demonstrates substantially larger terminal levels than those observed under near-isothermal operation: both the voltage at maximum power (V_{MPP}) and the corresponding maximum power (P_{MPP}) are markedly higher when the temperature differential is large. This reference behavior, shown in Figure 7.25, provides scale for interpreting the low- ΔT results that follow.




Figure 7.25 Marlow TG12-6 output in high ΔT

In contrast, when the temperature differential is constrained to 2–8°C, the measured V_{MPP} and P_{MPP} are greatly reduced but remain well behaved and monotonic with ΔT . For the first device (TEM-A), V_{MPP} increases steadily with ΔT , and P_{MPP} rises smoothly across the range; the same qualitative trends are observed for the second device (TEM-B), with absolute levels shifted relative to TEM-A. These datasets are presented in Figure 7.26 (TEM-A) and Figure 7.27 (TEM-B).




Figure 7.26 Marlow TG12-6 output in low ΔT




Figure 7.27 ATS-TEC40-39-004 output in low ΔT

Despite the much lower outputs than in high- ΔT operation, both devices deliver sufficient power in the $\Delta T = 2\text{--}8^{\circ}\text{C}$ window to run an embedded node under an appropriate energy budget (e.g., duty-cycled sensing, intermittent radio). In this regime the open-circuit voltage and the operating voltage at maximum power are very small, so practical deployments require ultra-low-voltage energy harvesters—of the type discussed in the previous chapter—capable of cold start and efficient boost at millivolt-level inputs. Consistent with the manual, conversion efficiency drops sharply as ΔT decreases (5.03% at $\Delta T = 180$, 2.42% at $\Delta T = 60$), and our end-to-end measurements on the target platform show only 0.41% at $\Delta T = 8$, reflecting realistic coupling and conversion losses. These figures tighten the neutrality margin at modest ΔT and increase the likelihood of survival-mode operation unless workloads are derated or additional headroom is provisioned.

8 Conclusion and future directions

Chapter 8 consolidates the thesis’ technical results into a coherent narrative and translates them into design guidance. We first recapitulate the main findings—environmental energy characterization, loss-aware power-path modelling, the DREAM runtime, and trace-driven sizing—emphasizing the guarantees each provides and the evidence supporting them. We then articulate the contributions relative to the state of the art, discuss practical implications and limitations, and chart specific directions where the approach can be extended. The goal is to leave the reader with a concise, actionable picture of what is now possible for energy-autonomous μ CPS and what remains to be done.

8.1 Summary of findings

This thesis advanced a node-centric, loss-aware framework that makes energy-autonomous operation of micro-scale cyber-physical systems (μ CPS) predictable rather than best-effort. Beginning from a principled definition and characterization of diffuse environmental energy Φ as observed at the point of use, we established decision-relevant descriptors (windowed energy envelopes, availability above source-side thresholds, calm-duration statistics) that are independent of any particular transduction pathway and preserve the temporal structure that matters for design and control. These descriptors provide conservative, evidence-bearing bounds for what is even conceivable in a deployment window.

Building on this environmental ceiling, we introduced a unified energy-collection architecture and macroscopic model that explicitly represent non-idealities across the power path—converter operating regions and efficiencies, cold-start energy, protection thresholds, storage leakage and round-trip penalties, and output-path overheads. Treating sources, conversion, storage, and controllable loads as a closed loop with measurable proxies yields tractable state equations in the space of stored energy and enables reasoning on the same side of the power path for both harvest and demand. This stance avoids double counting and aligns design-time modelling with runtime observability.

From this model, the thesis delivered two complementary artifacts. First, DREAM (Dynamic Rest & Energy-Aware Management), a runtime that regulates work at epoch boundaries using sparse, low-cost telemetry (primarily a reliable store proxy and protection events). DREAM provides: (i) an explicit neutrality law that nulls drift under surplus; (ii) a survival cadence that guarantees a configured horizon under scarcity; (iii) mode selection and

protection gating that bound behavior across cold-starts and brown-outs; and (iv) constant-time, constant-space control contracts that are portable across storage technologies and hardware revisions. Collectively, these properties replace ad-hoc duty-cycle heuristics with interpretable, evidence-tracing decisions.

Second, a trace-driven sizing methodology operationalized the model into concrete hardware choices—harvester technology and area, storage type and capacity, thresholds and hysteresis, and converter operating regions. The method replays representative input and load traces through measured macromodels to close the energy balance over deployment-relevant windows, quantify survivability and autonomy margins, and surface corner cases (e.g., dawn start-up, lull-induced depletion, retry storms) that average-power analyses miss. Where measured traces are scarce, lightweight synthetic generators that preserve mission-relevant temporal statistics offer a practical surrogate without biasing neutrality checks.

The approach was validated through case studies and experiments spanning simulation and hardware prototypes. Results demonstrated that, even under weak and intermittent sources (e.g., low- ΔT thermoelectric operation), the proposed architecture and runtime can sustain predictable service with materially fewer brown-outs than battery-only baselines or managers that ignore conversion losses and protection dynamics. The overall outcome is a principled path from environmental fluxes to system-level guarantees, enabling maintainable, sustainable μ CPS deployments.

The guarantees provided in this thesis remain bounded by the traces and models used: they assume reasonably stationary statistics of the harvested energy over the mission horizon, do not account for rare catastrophic faults or hardware aging, and do not claim optimality of DREAM beyond the invariants we prove for windowed neutrality and survival.

8.2 Contributions to the field

The thesis contributes:

1. **A device-agnostic formalization of diffuse environmental energy Φ for μ CPS design.** The definition and accompanying statistics (percentile envelopes of windowed energy, availability above practical thresholds, calm-duration distributions) create a neutral interface between environmental measurement and subsequent sizing and control, independent of specific transducers or circuits.

2. **A unified, loss-aware energy-collection model and architectural vocabulary.** By coupling reference architectures with a macroscopic state equation in stored-energy space, the work makes explicit the role of cold-start, thresholds, leakage, and efficiency plateaus, and identifies the minimal telemetry that the power path must expose for sound runtime control.
3. **DREAM: a constant-overhead runtime with neutrality and survival guarantees.** DREAM’s epoch-based regulation uses only low-cost observables, achieves neutrality under surplus, and enforces a survival cadence under scarcity, providing interpretable control with bounded compute/memory costs suitable for microcontrollers.
4. **A trace-driven, window-neutrality sizing methodology with quantitative guarantees.** The method closes the energy balance on the input/storage side using measured or faithfully synthesized traces, producing concrete harvester area, storage capacity, and threshold settings that meet autonomy targets with stated confidence while revealing critical corner cases.
5. **Evidence of practicality across representative μ CPS.** Through simulation models, hardware prototypes, and comparative baselines, the thesis shows that a loss-aware, telemetry-sparse approach improves autonomy and predictability without burdensome sensing or complex forecasting.

Taken together, these contributions reframe energy autonomy from a component-level efficiency pursuit to a system-level guarantees problem with clear contracts between hardware, firmware, and measurement practice.

8.3 Recommendations for future work

The results point to several avenues where focused effort can yield outsized gains:

1. **Tighter integration of conversion control and runtime policy.** Present implementations treat MPPT/boost/buck control as largely independent of DREAM’s epoch timing. Joint design—in which converter set-points and mode changes are co-scheduled with epoch decisions—could recover additional margin, especially in millivolt-level thermoelectric and low-illuminance indoor-PV regimes.

2. **Deliverability-aware sizing beyond energy neutrality.** The current methodology guarantees energy balance and survival over windows; future work should embed burst deliverability constraints explicitly (e.g., radio TX/RX spikes, sensor warm-ups) in the sizing objective, capturing ESR, converter current limits, and transient impedance effects so that worst-case crest factors are provably admissible.
3. **Formal robustness under non-stationarity.** While percentile envelopes and calm statistics address variability, systematic drifts (seasonality shifts, occupancy changes, aging) motivate robust counterparts—e.g., distributionally robust neutrality tests, adaptive survival horizons, and online identification of efficiency/ leakage parameters that preserve guarantees under slow change.
4. **Multi-source coordination and hybrid storage.** Extending the model to cover concurrent heterogeneous harvesters with shared constraints (e.g., indoor PV + vibration) and hybrid stores (supercapacitor + rechargeable micro-battery) would enable cross-domain arbitration, prioritize sources by marginal round-trip efficiency, and mitigate cold-start traps.
5. **End-to-end verification toolchain.** A practical outcome is a reference implementation that ingests environmental/ load traces, identifies macromodel parameters from lab measurements, runs window-neutrality/survivability checks, and emits firmware-level configurations (epoch length, thresholds, survival cadence) together with guarantees. Packaging this flow as an open tool would accelerate adoption and facilitate reproducible evaluation.
6. **Telemetry contracts and standards.** The runtime’s minimal-telemetry stance suggests a small, standard interface—store proxy characteristics, protection event semantics, and optional source proxies—between power-management hardware and firmware. Defining and standardizing this interface would improve portability across vendors and platforms.
7. **Application-specific studies at scale.** Larger deployments in buildings, logistics, and environmental monitoring can probe model fidelity under operational diversity, quantify maintenance savings versus battery-centric designs, and surface new edge cases for the sizing and runtime policies.

8.4 Emerging trends and future research directions

Several technological and methodological currents align with and extend the thesis:

- **Ultra-low-voltage conversion and cold-start innovations.** Advances in sub-50 mV cold-start, nano-ampere quiescent controllers, and SIMO architectures can materially expand the feasible envelope for low- ΔT thermoelectrics and indoor-PV. Co-design with DREAM can translate these gains into longer survival and smaller stores.
- **Energy-proportional radios and sensing front-ends.** Radios that expose energy-aware MACs, hardware pacing for preamble sampling, and wake-on-pattern receivers reduce crest factors and align communication bursts with harvest. Similar trends in sensors (e.g., duty-proportional analog front-ends) simplify deliverability guarantees.
- **Learning-light controllers with guarantees.** Lightweight, certification-friendly learning modules (e.g., bandit-style epoch tuning, robust regression on efficiency proxies) can adapt parameters without violating safety predicates. The objective is not forecast accuracy per se, but guarantee-preserving adaptation under slow drift.
- **Trace ecosystems and benchmarks.** Public corpora of μ CPS-scale environmental and load traces—measured with protocols that preserve spectral/temporal fidelity and accompanied by geometry/context metadata—would enable fair comparisons of sizing and runtime policies, and de-risk transfer across sites.
- **Sustainable materials and circular design.** As autonomy improves, the bill of materials can pivot away from large primary cells toward recyclable stores and small secondary chemistries, aided by right-sizing from window neutrality. Life-cycle analyses under realistic duty/harvest profiles can quantify environmental benefits.
- **Security and trust under intermittent power.** Guarantee-preserving power contracts can simplify secure boot, key retention, and privacy-sensitive sensing under brown-out dynamics, motivating research at the intersection of energy-aware control and system security.

By elevating environmental characterization, loss-aware modelling, and minimal-telemetry runtime control to first-class citizens—and by validating the approach with trace-driven sizing and representative prototypes—the thesis turns ambient energy from a convenience into a design foundation. The proposed framework equips designers to build μ CPS that are not merely efficient but **predictable** in the wild, and it opens clear paths for community contributions in modelling, tooling, hardware co-design, and large-scale evaluation.

Appendix A

Author's publications

1. Naskari, V.; Doumenis, G.; Masklavanos, I. Irradiance Non-Uniformity in LED Light Simulators. *Information* 2023, 14, 316. <https://doi.org/10.3390/info14060316>
2. I. Masklavanos, V. Naskari, C. Koutsos, F. Vartziotis, G. Doumenis, S. Siskos, A. Bardakas, A. Segkos, C. Tsamis, C. Papakis and G. Koukas, "Design of an Autonomous Wireless Electric Field sensor for maritime applications: the EFOS project", *Edge Intelligence 2023, Emerging Tech Conference*
3. G. Doumenis, V. Naskari, E. Hristoforou, P. Pattakos, G. Stamou, C. Papakis and I. Masklavanos, "Design of an Autonomous, Multi-functional Stress Assessment Sensor for Naval Applications: the AMSA project", *Edge Intelligence 2023, Emerging Tech Conference*
4. K. Tsiapali, S. Katsikas, A. Sakellariou, G. Doumenis and I. Masklavanos, "Self-powered IoT devices," *Edge Intelligence 2022, Emerging Tech Conference*, <https://doi.org/10.63438/GNHK6931>
5. G. Doumenis, I. Masklavanos and K. Tsiapali, "Lightweight operation scheduling for self-powered IoT devices," *2022 7th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM)*, 2022, pp. 1-7, doi: 10.1109/SEEDA-CECNSM57760.2022.9932933
6. V. Naskari, G. Doumenis, C. Koutsos and I. Masklavanos, "Design and implementation of an indoors light simulator," *2022 7th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM)*, 2022, pp. 1-6, doi: 10.1109/SEEDA-CECNSM57760.2022.9932899.
7. I. Masklavanos, C. Koutsos and G. Doumenis, "Efficiency of SuperCap storage in energy harvesting micropower supplies," *2022 7th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM)*, 2022, pp. 1-6, doi: 10.1109/SEEDA-CECNSM57760.2022.9932904.
8. "Autonomous energy harvesting system for maritime applications", V. Gogolou, K. Kozalakis, S. Siskos, A. Segkos, A. Bardakas, V. Vamvakas, I. Masklavanos, C. Tsamis, K. Tsiapali, A. Sakellariou, S. Katsikas, 19th International Conference on Nanosciences & Nanotechnologies (NN22), 5-8 July 2022, Thessaloniki, Greece
9. "Improvement of the performance of triboelectric generators by plasma induced surface roughness", A. Segkos, A. Bardakas, V. Vamvakas, I. Masklavanos, A. Zeniou, E. Gogolides, C. Tsamis, 19th International Conference on Nanosciences & Nanotechnologies (NN22), 5-8 July 2022, Thessaloniki, Greece

References

- [1] R. (Raj) Rajkumar, I. Lee, L. Sha, and J. Stankovic, "Cyber-physical systems: the next computing revolution," in *Proceedings of the 47th Design Automation Conference*, in DAC '10. New York, NY, USA: Association for Computing Machinery, June 2010, pp. 731–736. doi: 10.1145/1837274.1837461.
- [2] R. Alur, *Principles of Cyber-Physical Systems*. Cambridge, MA, USA: MIT Press, 2023.
- [3] E. R. Griffor, C. Greer, D. A. Wollman, and M. J. Burns, "Framework for Cyber-Physical Systems: Volume 1, Overview," *NIST*, June 2017, Accessed: Sept. 21, 2025. [Online]. Available: <https://www.nist.gov/publications/framework-cyber-physical-systems-volume-1-overview>
- [4] F. K. Shaikh and S. Zeadally, "Energy harvesting in wireless sensor networks: A comprehensive review," *Renew. Sustain. Energy Rev.*, vol. 55, pp. 1041–1054, Mar. 2016, doi: 10.1016/j.rser.2015.11.010.
- [5] J. Singh, R. Kaur, and D. Singh, "Energy harvesting in wireless sensor networks: A taxonomic survey," *Int. J. Energy Res.*, vol. 45, no. 1, pp. 118–140, 2021, doi: 10.1002/er.5816.
- [6] M. Surbatovich, B. Lucia, and L. Jia, "Towards a formal foundation of intermittent computing," *Proc ACM Program Lang*, vol. 4, no. OOPSLA, p. 163:1-163:31, Nov. 2020, doi: 10.1145/3428231.
- [7] S. Umesh and S. Mittal, "A survey of techniques for intermittent computing," *J. Syst. Archit.*, vol. 112, p. 101859, Jan. 2021, doi: 10.1016/j.sysarc.2020.101859.
- [8] F. Karray, M. W. Jmal, A. Garcia-Ortiz, M. Abid, and A. M. Obeid, "A comprehensive survey on wireless sensor node hardware platforms," *Comput. Netw.*, vol. 144, pp. 89–110, Oct. 2018, doi: 10.1016/j.comnet.2018.05.010.
- [9] "The Internet of Batteryless Things – Communications of the ACM." Accessed: Sept. 21, 2025. [Online]. Available: <https://cacm.acm.org/research/the-internet-of-batteryless-things/>
- [10] Q. Lian, P. Han, and N. Mei, "A Review of Converter Circuits for Ambient Micro Energy Harvesting," *Micromachines*, vol. 13, no. 12, p. 2222, Dec. 2022, doi: 10.3390/mi13122222.
- [11] M. Ceraolo, G. Lutzemberger, and D. Poli, "State-Of-Charge Evaluation Of Supercapacitors," *J. Energy Storage*, vol. 11, pp. 211–218, June 2017, doi: 10.1016/j.est.2017.03.001.
- [12] P. Saha, S. Dey, and M. Khanra, "Accurate estimation of state-of-charge of supercapacitor under uncertain leakage and open circuit voltage map," *J. Power Sources*, vol. 434, p. 226696, Sept. 2019, doi: 10.1016/j.jpowsour.2019.226696.
- [13] K. Movassagh, A. Raihan, B. Balasingam, and K. Pattipati, "A Critical Look at Coulomb Counting Approach for State of Charge Estimation in Batteries," *Energies*, vol. 14, no. 14, p. 4074, Jan. 2021, doi: 10.3390/en14144074.
- [14] L. Lin, Z. Tang, N. Tan, and X. Xiao, "Power Management in Low-Power MCUs for Energy IoT Applications," *J. Sens.*, vol. 2020, no. 1, p. 8819236, 2020, doi: 10.1155/2020/8819236.
- [15] A. Lopez-Gasso, A. Beriain, H. Solar, and R. Berenguer, "Power Management Unit for Solar Energy Harvester Assisted Batteryless Wireless Sensor Node," *Sensors*, vol. 22, no. 20, p. 7908, Jan. 2022, doi: 10.3390/s22207908.
- [16] "7 μ A 1-Cell Fuel Gauge with ModelGauge m5 EZ".
- [17] L. Steinfeld, J. Oreggioni, D. A. Bouvier, C. A. Fernández, and J. Villaverde, "Smart Coulomb Counter for Self-Metering Wireless Sensor Nodes Consumption," *J. Low Power Electron.*, vol. 11, no. 2, pp. 236–248, June 2015, doi: 10.1166/jolpe.2015.1370.
- [18] B. Baker, J. Woods, M. J. Reed, and M. Afford, "A Survey of Short-Range Wireless Communication for Ultra-Low-Power Embedded Systems," *J. Low Power Electron. Appl.*, vol. 14, no. 2, p. 27, June 2024, doi: 10.3390/jlpea14020027.
- [19] N. Smith, "Understanding Undervoltage Lockout in Power Devices," 2016.
- [20] M. Zimmerling, L. Mottola, and S. Santini, "Synchronous Transmissions in Low-Power Wireless: A Survey of Communication Protocols and Network Services," *ACM Comput Surv*, vol. 53, no. 6, p. 121:1-121:39, Dec. 2020, doi: 10.1145/3410159.

- [21] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele, "Low-power wireless bus," in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, in SenSys '12. New York, NY, USA: Association for Computing Machinery, Nov. 2012, pp. 1–14. doi: 10.1145/2426656.2426658.
- [22] D. De Guglielmo, S. Brienza, and G. Anastasi, "IEEE 802.15.4e: A survey," *Comput. Commun.*, vol. 88, pp. 1–24, Aug. 2016, doi: 10.1016/j.comcom.2016.05.004.
- [23] T. Sayjari, R. Melo Silveira, and C. Borges Margi, "Application-Aware Scheduling for IEEE 802.15.4e Time-Slotted Channel Hopping Using Software-Defined Wireless Sensor Network Slicing," *Sensors*, vol. 23, no. 16, p. 7143, Jan. 2023, doi: 10.3390/s23167143.
- [24] M. Abdelmalak, V. Venkataramanan, and R. Macwan, "A Survey of Cyber-Physical Power System Modeling Methods for Future Energy Systems," *IEEE Access*, vol. 10, pp. 99875–99896, 2022, doi: 10.1109/ACCESS.2022.3206830.
- [25] T. Hardy, B. Palmintier, P. Top, D. Krishnamurthy, and J. Fuller, "HELICS: A Co-Simulation Framework for Scalable Multi-Domain Modeling and Analysis," *IEEE Access*, vol. 12, pp. 24325–24347, 2024, doi: 10.1109/ACCESS.2024.3363615.
- [26] V. R. Motakatla *et al.*, "Integrated Transmission and Distribution Co-Simulation Platform for Demonstration of Bulk Grid Services Using Distributed Energy Resources," *Energies*, vol. 17, no. 13, p. 3215, June 2024, doi: 10.3390/en17133215.
- [27] L. N. H. Pham, V. Rosero-Morillo, A. Shukla, F. Gonzalez-Longatt, and V. Meza-G, "Real-Time Cyber-Physical Power System Testbed for International Electrotechnical Commission 61850 Generic Object-Oriented Substation Event Transfer Time Measurements," *Eng. Proc.*, vol. 77, no. 1, p. 17, 2024, doi: 10.3390/engproc2024077017.
- [28] B. Deepa, S. Hampannavar, and S. Mansani, "Phasor estimation using micro-phasor measurement unit (μ PMU) in distribution network for situational awareness," *J. Electr. Syst. Inf. Technol.*, vol. 11, no. 1, p. 52, Nov. 2024, doi: 10.1186/s43067-024-00179-5.
- [29] B. H. Hameed and S. Kurnaz, "Secure low-cost photovoltaic monitoring system based on LoRaWAN network and artificial intelligence," *Discov. Comput.*, vol. 27, no. 1, p. 36, Oct. 2024, doi: 10.1007/s10791-024-09475-0.
- [30] A. W. Indrawan *et al.*, "LoRa-based mesh network for distribution substation monitoring," *AIP Conf. Proc.*, vol. 3140, no. 1, p. 040021, July 2024, doi: 10.1063/5.0221081.
- [31] X. Wang, W. Zhao, and X. Niu, "Low voltage user power internet of things monitoring system based on LoRa wireless technology," *Energy Inform.*, vol. 8, no. 1, p. 12, Jan. 2025, doi: 10.1186/s42162-025-00472-1.
- [32] G. Li *et al.*, "A critical review of cyber-physical security for building automation systems," *Annu. Rev. Control*, vol. 55, pp. 237–254, Jan. 2023, doi: 10.1016/j.arcontrol.2023.02.004.
- [33] T. W. Kang and Y. Mo, "A comprehensive digital twin framework for building environment monitoring with emphasis on real-time data connectivity and predictability," *Dev. Built Environ.*, vol. 17, p. 100309, Mar. 2024, doi: 10.1016/j.dibe.2023.100309.
- [34] Y. Zhao *et al.*, "A Comprehensive Evaluation of Bluetooth Low Energy Mesh," in *2024 IEEE 30th International Conference on Parallel and Distributed Systems (ICPADS)*, Belgrade, Serbia: IEEE, Oct. 2024, pp. 544–551. doi: 10.1109/ICPADS63350.2024.00077.
- [35] "ISA/IEC 62443 Series of Standards - ISA," isa.org. Accessed: Sept. 24, 2025. [Online]. Available: <https://www.isa.org/standards-and-publications/isa-standards/isa-iec-62443-series-of-standards>
- [36] J. V. Vaghasiya, C. C. Mayorga-Martinez, and M. Pumera, "Wearable sensors for telehealth based on emerging materials and nanoarchitectonics," *Npj Flex. Electron.*, vol. 7, no. 1, p. 26, June 2023, doi: 10.1038/s41528-023-00261-4.
- [37] G. Maroli *et al.*, "Wearable, battery-free, wireless multiplexed printed sensors for heat stroke prevention with mussel-inspired bio-adhesive membranes," *Biosens. Bioelectron.*, vol. 260, p. 116421, Sept. 2024, doi: 10.1016/j.bios.2024.116421.
- [38] H. Jia *et al.*, "UR2M: Uncertainty and Resource-Aware Event Detection on Microcontrollers," Mar. 12, 2024, *arXiv*: arXiv:2402.09264. doi: 10.48550/arXiv.2402.09264.
- [39] C. for D. and R. Health, "Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions." Accessed: Sept. 24, 2025. [Online]. Available:

- <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-system-considerations-and-content-premarket-submissions>
- [40] “Principles and Practices for the Cybersecurity of Legacy Medical Devices | International Medical Device Regulators Forum.” Accessed: Sept. 24, 2025. [Online]. Available: <https://www.imdrf.org/documents/principles-and-practices-cybersecurity-legacy-medical-devices>
 - [41] P. Busia, M. A. Scrugli, V. J.-B. Jung, L. Benini, and P. Meloni, “A Tiny Transformer for Low-Power Arrhythmia Classification on Microcontrollers,” *IEEE Trans. Biomed. Circuits Syst.*, vol. 19, no. 1, pp. 142–152, Feb. 2025, doi: 10.1109/TBCAS.2024.3401858.
 - [42] Y.-T. Ting and K.-Y. Chan, “Optimising performances of LoRa based IoT enabled wireless sensor network for smart agriculture,” *J. Agric. Food Res.*, vol. 16, p. 101093, June 2024, doi: 10.1016/j.jafr.2024.101093.
 - [43] E. Bicomumakuba, E. Habineza, Samsuzzaman, M. N. Reza, and S.-O. Chung, “IoT-enabled LoRaWAN gateway for monitoring and predicting spatial environmental parameters in smart greenhouses: A review,” *Precis. Agric. Sci. Technol.*, vol. 7, no. 1, pp. 28–46, Mar. 2025, doi: 10.22765/PASTJ.20250003.
 - [44] X. Xiao, “Sustainable Agriculture with Self-Powered Wireless Sensing,” *Agriculture*, vol. 15, no. 3, p. 234, Jan. 2025, doi: 10.3390/agriculture15030234.
 - [45] D. Fernández Llorca *et al.*, “Testing autonomous vehicles and AI: perspectives and challenges from cybersecurity, transparency, robustness and fairness,” *Eur. Transp. Res. Rev.*, vol. 17, no. 1, p. 38, July 2025, doi: 10.1186/s12544-025-00732-x.
 - [46] M. Tabassum, F. H. Bastos, A. Oliveira, and A. Klautau, “NR Sidelink Performance Evaluation for Enhanced 5G-V2X Services,” *Vehicles*, vol. 5, no. 4, pp. 1692–1706, Dec. 2023, doi: 10.3390/vehicles5040092.
 - [47] E. Bitsikas and A. Ranganathan, “Security Analysis of 5G NR Device-to-Device Sidelink Communications,” Feb. 23, 2025, *arXiv*: arXiv:2502.16650. doi: 10.48550/arXiv.2502.16650.
 - [48] G. Ahmed and T. Sheltami, “Novel Energy-Aware 3D UAV Path Planning and Collision Avoidance Using Receding Horizon and Optimization-Based Control,” *Drones*, vol. 8, no. 11, p. 682, Nov. 2024, doi: 10.3390/drones8110682.
 - [49] P. Thantharate, A. Thantharate, and A. Kulkarni, “GREENSKY: A fair energy-aware optimization model for UAVs in next-generation wireless networks,” *Green Energy Intell. Transp.*, vol. 3, no. 1, p. 100130, Feb. 2024, doi: 10.1016/j.geits.2023.100130.
 - [50] H. Zhang *et al.*, “Harvesting Inertial Energy and Powering Wearable Devices: A Review,” *Small Methods*, vol. 8, no. 1, p. 2300771, 2024, doi: 10.1002/smt.202300771.
 - [51] D. E. Parsons, G. Koutsourakis, and J. C. Blakesley, “Performance measurements for indoor photovoltaic devices: Classification of a novel light source,” *APL Energy*, vol. 2, no. 1, p. 016110, Mar. 2024, doi: 10.1063/5.0186028.
 - [52] J. Gołabek and M. Strankowski, “A Review of Recent Advances in Human-Motion Energy Harvesting Nanogenerators, Self-Powering Smart Sensors and Self-Charging Electronics,” *Sensors*, vol. 24, no. 4, p. 1069, Jan. 2024, doi: 10.3390/s24041069.
 - [53] A. Chakraborty *et al.*, “Photovoltaics for indoor energy harvesting,” *Nano Energy*, vol. 128, p. 109932, Sept. 2024, doi: 10.1016/j.nanoen.2024.109932.
 - [54] S. Zeiske, P. Meredith, A. Armin, and G. Burwell, “Importance of spectrally invariant broadband attenuation of light in indoor photovoltaic characterization,” *APL Energy*, vol. 1, no. 2, p. 026103, Sept. 2023, doi: 10.1063/5.0159289.
 - [55] A. Eghbali *et al.*, “Results of an International Comparison of Indoor Photovoltaic Measurements Among Seven Metrological Institutes,” *Sol. RRL*, vol. 9, no. 9, p. 2400855, 2025, doi: 10.1002/solr.202400855.
 - [56] G. Qu *et al.*, “Current Research Status and Future Trends of Vibration Energy Harvesters,” *Micromachines*, vol. 15, no. 9, p. 1109, Sept. 2024, doi: 10.3390/mi15091109.
 - [57] K. Seunarine *et al.*, “Light power resource availability for energy harvesting photovoltaics for self-powered IoT,” *J. Phys. Energy*, vol. 6, no. 1, p. 015018, Jan. 2024, doi: 10.1088/2515-7655/ad1764.

- [58] Y. Cao *et al.*, “Human-motion adaptability enhancement of wearable electromagnetic vibration energy harvesters toward self-sustained body sensor networks,” *Cell Rep. Phys. Sci.*, vol. 5, no. 9, p. 102117, Sept. 2024, doi: 10.1016/j.xcrp.2024.102117.
- [59] I. S.o and O. K.o, “Measurement of Ambient Radio Frequency Energy for Low Powered Applications,” *J. Energy Technol. Environ.*, vol. 7, no. 3, pp. 181–193, Sept. 2025, doi: 10.37933/jete/7.3.2025.1946.
- [60] B. Lucia, V. Balaji, A. Colin, K. Maeng, and E. Ruppel, “Intermittent Computing: Challenges and Opportunities,” in *2nd Summit on Advances in Programming Languages (SNAPL 2017)*, B. S. Lerner, R. Bodík, and S. Krishnamurthi, Eds., in Leibniz International Proceedings in Informatics (LIPIcs), vol. 71. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017, p. 8:1-8:14. doi: 10.4230/LIPIcs.SNAPL.2017.8.
- [61] “Time series analysis and its applications: with R examples.” Accessed: Sept. 28, 2025. [Online]. Available: https://link.springer.com/content/pdf/10.1007/0-387-36276-2_6.pdf
- [62] H. Di, G. Su, J. Yu, Y. Xu, P. Xu, and S. Zhou, “Field measurement and evaluation of vibrations inside buildings above metro tunnels,” *Soil Dyn. Earthq. Eng.*, vol. 166, p. 107767, Mar. 2023, doi: 10.1016/j.soildyn.2023.107767.
- [63] J. Pettersson and T. Edström, “An Evaluation of the Dynamic Behavior from Ambient Vibrations in a High-Rise Timber Building”.
- [64] A. Sohrabbeig, O. Ardakanian, and P. Musilek, “Decompose and Conquer: Time Series Forecasting with Multiseasonal Trend Decomposition Using Loess,” *Forecasting*, vol. 5, no. 4, pp. 684–696, Dec. 2023, doi: 10.3390/forecast5040037.
- [65] M. Wang, Y. Meng, L. Sun, and T. Zhang, “Decomposition combining averaging seasonal-trend with singular spectrum analysis and a marine predator algorithm embedding Adam for time series forecasting with strong volatility,” *Expert Syst. Appl.*, vol. 274, p. 126864, May 2025, doi: 10.1016/j.eswa.2025.126864.
- [66] T. Ramos, S. Dedesko, J. A. Siegel, J. A. Gilbert, and B. Stephens, “Spatial and Temporal Variations in Indoor Environmental Conditions, Human Occupancy, and Operational Characteristics in a New Hospital Building,” *PLOS ONE*, vol. 10, no. 3, p. e0118207, Mar. 2015, doi: 10.1371/journal.pone.0118207.
- [67] M. Karsai, K. Kaski, A.-L. Barabási, and J. Kertész, “Universal features of correlated bursty behaviour,” *Sci. Rep.*, vol. 2, no. 1, p. 397, May 2012, doi: 10.1038/srep00397.
- [68] H.-H. Jo, M. Karsai, J. Kertész, and K. Kaski, “Circadian pattern and burstiness in mobile phone communication,” *New J. Phys.*, vol. 14, p. 013055, Jan. 2012, doi: 10.1088/1367-2630/14/1/013055.
- [69] S. Unicomb, G. Iñiguez, J. P. Gleeson, and M. Karsai, “Dynamics of cascades on burstiness-controlled temporal networks,” *Nat. Commun.*, vol. 12, no. 1, p. 133, Jan. 2021, doi: 10.1038/s41467-020-20398-4.
- [70] S. L. Hartmeyer, A. Davies, and M. Andersen, “Post-Occupancy Evaluation of Office Lighting – What Can We Learn from Light-Dosimetry?,” *LEUKOS*, vol. 0, no. 0, pp. 1–24, doi: 10.1080/15502724.2024.2434976.
- [71] L. Belácková *et al.*, “Temporal change of outdoor RF-EMF levels in four European countries: a microenvironmental measurement study,” *Environ. Res.*, vol. 285, p. 122315, Nov. 2025, doi: 10.1016/j.envres.2025.122315.
- [72] A. F. Veludo *et al.*, “Assessing radiofrequency electromagnetic field exposure in multiple microenvironments across ten European countries with a focus on 5G,” *Environ. Int.*, vol. 200, p. 109540, June 2025, doi: 10.1016/j.envint.2025.109540.
- [73] R. Ramirez-Vazquez, I. Escobar, E. Arribas, and G. A. E. Vandenbosch, “Systematic Review of Exposure Studies to Radiofrequency Electromagnetic Fields: Spot Measurements and Mixed Methodologies,” *Appl. Sci.*, vol. 14, no. 23, p. 11161, Jan. 2024, doi: 10.3390/app142311161.
- [74] J. van Duijnhoven, M. P. J. Aarts, and H. S. M. Kort, “The importance of including position and viewing direction when measuring and assessing the lighting conditions of office workers,” *WORK*, vol. 64, no. 4, pp. 877–895, Dec. 2019, doi: 10.3233/WOR-193028.

- [75] M. Brzezicki, "An Evaluation of Useful Daylight Illuminance in an Office Room with a Light Shelf and Translucent Ceiling at 51° N," *Buildings*, vol. 11, no. 11, p. 494, Nov. 2021, doi: 10.3390/buildings11110494.
- [76] A. Jamrozik *et al.*, "Access to daylight and view in an office improves cognitive performance and satisfaction and reduces eyestrain: A controlled crossover study," *Build. Environ.*, vol. 165, p. 106379, Nov. 2019, doi: 10.1016/j.buildenv.2019.106379.
- [77] D. Caballol, Á. P. Raposo, F. Gil Carrillo, and M. Morales-Segura, "Measurement of ambient vibration in empty buildings and relation to external noise," *Appl. Acoust.*, vol. 186, p. 108431, Jan. 2022, doi: 10.1016/j.apacoust.2021.108431.
- [78] H. H. R. Sherazi, D. Zorbas, and B. O'Flynn, "A Comprehensive Survey on RF Energy Harvesting: Applications and Performance Determinants," *Sensors*, vol. 22, no. 8, p. 2990, Jan. 2022, doi: 10.3390/s22082990.
- [79] N. Loizeau *et al.*, "Comparison of ambient radiofrequency electromagnetic field (RF-EMF) levels in outdoor areas and public transport in Switzerland in 2014 and 2021," *Environ. Res.*, vol. 237, p. 116921, Nov. 2023, doi: 10.1016/j.envres.2023.116921.
- [80] N. G. Vasquez, R. F. Rupp, R. K. Andersen, and J. Toftum, "Occupants' responses to window views, daylighting and lighting in buildings: A critical review," *Build. Environ.*, vol. 219, p. 109172, July 2022, doi: 10.1016/j.buildenv.2022.109172.
- [81] N. Ghaeili Ardabili, J. Wang, and N. Wang, "A systematic literature review: building window's influence on indoor circadian health," *Renew. Sustain. Energy Rev.*, vol. 188, p. 113796, Dec. 2023, doi: 10.1016/j.rser.2023.113796.
- [82] Z. Murphy, S. Kapp, and J.-K. Choi, "Analytical study on daylight harvesting potential for manufacturing facilities: Techno-economic-environmental perspective," *Energy Build.*, vol. 320, p. 114565, Oct. 2024, doi: 10.1016/j.enbuild.2024.114565.
- [83] M. Pirc, Ž. Ajdič, D. Uršič, M. Jošt, and M. Topič, "Indoor Energy Harvesting With Perovskite Solar Cells for IoT Applications—A Full Year Monitoring Study," *ACS Appl. Energy Mater.*, vol. 7, no. 2, pp. 565–575, Jan. 2024, doi: 10.1021/acsaem.3c02498.
- [84] S. N. Arinze, E. R. Obi, S. H. Ebeniwa, and A. O. Nwajana, "RF Energy-Harvesting Techniques: Applications, Recent Developments, Challenges, and Future Opportunities," *Telecom.*, vol. 6, no. 3, p. 45, Sept. 2025, doi: 10.3390/telecom6030045.
- [85] S. Singh, M. Kumar, and R. Kumar, "Powering the future: A survey of ambient RF-based communication systems for next-gen wireless networks," *IET Wirel. Sens. Syst.*, vol. 14, no. 6, pp. 265–292, 2024, doi: 10.1049/wss2.12094.
- [86] D. Pan, L. Zhang, B. Wang, J. Jia, Z. Song, and X. Zhang, "Multi-objective planning of integrated energy system based on CVaR under carbon trading mechanism," *Front. Energy Res.*, vol. 12, Feb. 2024, doi: 10.3389/fenrg.2024.1310301.
- [87] D. Stenclik, "New Resource Adequacy Criteria for the Energy Transition: Modernizing Reliability Requirements," NREL/SR--6A40-89334, 2372882, MainId:90113, Mar. 2024. doi: 10.2172/2372882.
- [88] R. L. Z. Hoye *et al.*, "Reaching a consensus on indoor photovoltaics testing," *Joule*, vol. 0, no. 0, Sept. 2025, doi: 10.1016/j.joule.2025.102127.
- [89] "Indoor photovoltaics with over 40 percent efficiency developed - Fraunhofer ISE," Fraunhofer Institute for Solar Energy Systems ISE. Accessed: Sept. 28, 2025. [Online]. Available: <https://www.ise.fraunhofer.de/en/press-media/news/2025/indoor-photovoltaics-with-over-40-percent-efficiency-developed.html>
- [90] Y. Olzhabay *et al.*, "Evaluation of Indoor Power Performance of Emerging Photovoltaic Technology for IoT Device Application," *Energies*, vol. 18, no. 5, p. 1118, Jan. 2025, doi: 10.3390/en18051118.
- [91] J.-P. Kreiss and S. N. Lahiri, "1 - Bootstrap Methods for Time Series," in *Handbook of Statistics*, vol. 30, T. Subba Rao, S. Subba Rao, and C. R. Rao, Eds., in Time Series Analysis: Methods and Applications, vol. 30., Elsevier, 2012, pp. 3–26. doi: 10.1016/B978-0-444-53858-1.00001-6.
- [92] A. E. Dudek, J. Leśkow, E. Paparoditis, and D. N. Politis, "A Generalized Block Bootstrap for Seasonal Time Series," *J. Time Ser. Anal.*, vol. 35, no. 2, pp. 89–114, 2014, doi: 10.1002/jtsa.12053.

- [93] R. J. Hyndman and G. Athanasopoulos, *Forecasting: principles and practice*. OTexts, 2018.
- [94] S. Moritz and T. Bartz-Beielstein, “imputeTS: Time Series Missing Value Imputation in R,” *R J.*, June 2017, [Online]. Available: <https://digitalcommons.unl.edu/r-journal/453>
- [95] A. Driesse, A. R. Jensen, and R. Perez, “A continuous form of the Perez diffuse sky model for forward and reverse transposition,” *Sol. Energy*, vol. 267, p. 112093, Jan. 2024, doi: 10.1016/j.solener.2023.112093.
- [96] G. Kopp and J. L. Lean, “A new, lower value of total solar irradiance: Evidence and climate significance,” *Geophys. Res. Lett.*, vol. 38, no. 1, 2011, doi: 10.1029/2010GL045777.
- [97] T. Sung, S. Y. Yoon, and K. C. Kim, “A Mathematical Model of Hourly Solar Radiation in Varying Weather Conditions for a Dynamic Simulation of the Solar Organic Rankine Cycle,” *Energies*, vol. 8, no. 7, pp. 7058–7069, July 2015, doi: 10.3390/en8077058.
- [98] R. A. Kishore and S. Priya, “A Review on Low-Grade Thermal Energy Harvesting: Materials, Methods and Devices,” *Materials*, vol. 11, no. 8, p. 1433, Aug. 2018, doi: 10.3390/ma11081433.
- [99] J. He, K. Li, L. Jia, Y. Zhu, H. Zhang, and J. Linghu, “Advances in the applications of thermoelectric generators,” *Appl. Therm. Eng.*, vol. 236, p. 121813, Jan. 2024, doi: 10.1016/j.applthermaleng.2023.121813.
- [100] D. Champier, “Thermoelectric generators: A review of applications,” *Energy Convers. Manag.*, vol. 140, pp. 167–181, May 2017, doi: 10.1016/j.enconman.2017.02.070.
- [101] V. Leonov, “Thermoelectric Energy Harvesting of Human Body Heat for Wearable Sensors,” *IEEE Sens. J.*, vol. 13, no. 6, pp. 2284–2291, June 2013, doi: 10.1109/JSEN.2013.2252526.
- [102] S. J. Kim, J. H. We, and B. J. Cho, “A wearable thermoelectric generator fabricated on a glass fabric,” *Energy Environ. Sci.*, vol. 7, no. 6, pp. 1959–1965, May 2014, doi: 10.1039/C4EE00242C.
- [103] J. M. W. Brownjohn *et al.*, “Footbridge system identification using wireless inertial measurement units for force and response measurements,” *J. Sound Vib.*, vol. 384, pp. 339–355, Dec. 2016, doi: 10.1016/j.jsv.2016.08.008.
- [104] L. Tang, Y. Yang, and C. K. Soh, “Toward Broadband Vibration-based Energy Harvesting,” *J. Intell. Mater. Syst. Struct.*, Dec. 2010, doi: 10.1177/1045389X10390249.
- [105] L. Zuo and P.-S. Zhang, “Energy Harvesting, Ride Comfort, and Road Handling of Regenerative Vehicle Suspensions,” *J. Vib. Acoust.*, vol. 135, no. 011002, Feb. 2013, doi: 10.1115/1.4007562.
- [106] “ISO 8608:2016,” ISO. Accessed: Sept. 29, 2025. [Online]. Available: <https://www.iso.org/standard/71202.html>
- [107] M. Pozzi and M. Zhu, “Plucked piezoelectric bimorphs for knee-joint energy harvesting: modelling and experimental validation,” *Smart Mater. Struct.*, vol. 20, no. 5, p. 055007, Apr. 2011, doi: 10.1088/0964-1726/20/5/055007.
- [108] J. Jiang, S. Liu, L. Feng, and D. Zhao, “A Review of Piezoelectric Vibration Energy Harvesting with Magnetic Coupling Based on Different Structural Characteristics,” *Micromachines*, vol. 12, no. 4, p. 436, Apr. 2021, doi: 10.3390/mi12040436.
- [109] S. Nabavi and L. Zhang, “Portable Wind Energy Harvesters for Low-Power Applications: A Survey,” *Sensors*, vol. 16, no. 7, p. 1101, July 2016, doi: 10.3390/s16071101.
- [110] P. Sharma and A. K. Singh, “A survey on RF energy harvesting techniques for lifetime enhancement of wireless sensor networks,” *Sustain. Comput. Inform. Syst.*, vol. 37, p. 100836, Jan. 2023, doi: 10.1016/j.suscom.2022.100836.
- [111] M. Piñuela, P. D. Mitcheson, and S. Lucyszyn, “Ambient RF Energy Harvesting in Urban and Semi-Urban Environments,” *IEEE Trans. Microw. Theory Tech.*, vol. 61, no. 7, pp. 2715–2726, July 2013, doi: 10.1109/TMTT.2013.2262687.
- [112] C. A. Gueymard and J. A. Ruiz-Arias, “Extensive worldwide validation and climate sensitivity analysis of direct irradiance predictions from 1-min global irradiance,” *Sol. Energy*, vol. 128, pp. 1–30, Apr. 2016, doi: 10.1016/j.solener.2015.10.010.
- [113] D. S. Wilks, *Statistical Methods in the Atmospheric Sciences*. Academic Press, 2011.
- [114] W. De Soto, S. A. Klein, and W. A. Beckman, “Improvement and validation of a model for photovoltaic array performance,” *Sol. Energy*, vol. 80, no. 1, pp. 78–88, Jan. 2006, doi: 10.1016/j.solener.2005.06.010.

- [115] M. Green, "Solar cell fill factors: General graph and empirical expressions," *Solid-State Electron.*, vol. 24, pp. 788–789, 1981.
- [116] T. Esram and P. L. Chapman, "Comparison of Photovoltaic Array Maximum Power Point Tracking Techniques," *IEEE Trans. Energy Convers.*, vol. 22, no. 2, pp. 439–449, June 2007, doi: 10.1109/TEC.2006.874230.
- [117] V. Salas, E. Olías, A. Barrado, and A. Lázaro, "Review of the maximum power point tracking algorithms for stand-alone photovoltaic systems," *Sol. Energy Mater. Sol. Cells*, vol. 90, no. 11, pp. 1555–1578, July 2006, doi: 10.1016/j.solmat.2005.10.023.
- [118] S. B. Riffat and X. Ma, "Thermoelectrics: a review of present and potential applications," *Appl. Therm. Eng.*, vol. 23, no. 8, pp. 913–935, June 2003, doi: 10.1016/S1359-4311(03)00012-7.
- [119] G. J. Snyder and T. S. Ursell, "Thermoelectric Efficiency and Compatibility," *Phys. Rev. Lett.*, vol. 91, no. 14, p. 148301, Oct. 2003, doi: 10.1103/PhysRevLett.91.148301.
- [120] C. B. Williams and R. B. Yates, "Analysis of a micro-electric generator for microsystems," *Sens. Actuators Phys.*, vol. 52, no. 1, pp. 8–11, Mar. 1996, doi: 10.1016/0924-4247(96)80118-X.
- [121] P. D. Mitcheson, E. M. Yeatman, G. K. Rao, A. S. Holmes, and T. C. Green, "Energy Harvesting From Human and Machine Motion for Wireless Electronic Devices," *Proc. IEEE*, vol. 96, no. 9, pp. 1457–1486, Sept. 2008, doi: 10.1109/JPROC.2008.927494.
- [122] S. R. Anton and H. A. Sodano, "A review of power harvesting using piezoelectric materials (2003–2006)," *Smart Mater. Struct.*, vol. 16, no. 3, p. R1, May 2007, doi: 10.1088/0964-1726/16/3/R01.
- [123] S. Roundy, P. K. Wright, and J. Rabaey, "A study of low level vibrations as a power source for wireless sensor nodes," *Comput. Commun.*, vol. 26, no. 11, pp. 1131–1144, July 2003, doi: 10.1016/S0140-3664(02)00248-7.
- [124] A. Badel, D. Guyomar, E. Lefeuvre, and C. Richard, "Piezoelectric Energy Harvesting using a Synchronized Switch Technique," *J. Intell. Mater. Syst. Struct.*, vol. 17, no. 8–9, pp. 831–839, Sept. 2006, doi: 10.1177/1045389X06057533.
- [125] A. Segkos and C. Tsamis, "Rotating Triboelectric Nanogenerators for Energy Harvesting and Their Applications," *Nanoenergy Adv.*, vol. 3, no. 3, pp. 170–219, Sept. 2023, doi: 10.3390/nanoenergyadv3030010.
- [126] A. Bardakas, A. Segkos, and C. Tsamis, "Zinc Oxide-Based Rotational–Linear Triboelectric Nanogenerator," *Appl. Sci.*, vol. 14, no. 6, p. 2396, Jan. 2024, doi: 10.3390/app14062396.
- [127] A. Anastasopoulos, V. Zacharia, A. Bardakas, and C. Tsamis, "Development of Triboelectric Devices for Human–Machine Interface Applications," *Proceedings*, vol. 97, no. 1, p. 7, 2024, doi: 10.3390/proceedings2024097007.
- [128] H. T. Friis, "A Note on a Simple Transmission Formula," *Proc. IRE*, vol. 34, no. 5, pp. 254–256, May 1946, doi: 10.1109/JRPROC.1946.234568.
- [129] C. R. Valenta and G. D. Durgin, "Harvesting Wireless Power: Survey of Energy-Harvester Conversion Efficiency in Far-Field, Wireless Power Transfer Systems," *IEEE Microw. Mag.*, vol. 15, no. 4, pp. 108–120, June 2014, doi: 10.1109/MMM.2014.2309499.
- [130] H. J. Visser and R. J. M. Vullers, "RF Energy Harvesting and Transport for Wireless Sensor Network Applications: Principles and Requirements," *Proc. IEEE*, vol. 101, no. 6, pp. 1410–1423, June 2013, doi: 10.1109/JPROC.2013.2250891.
- [131] J. A. Hagerty, F. B. Helmbrecht, W. H. McCalpin, R. Zane, and Z. B. Popovic, "Recycling ambient microwave energy with broad-band rectenna arrays," *IEEE Trans. Microw. Theory Tech.*, vol. 52, no. 3, pp. 1014–1024, Mar. 2004, doi: 10.1109/TMTT.2004.823585.
- [132] K. H. Hussein, I. Muta, T. Hoshino, and M. Osakada, "Maximum photovoltaic power tracking: an algorithm for rapidly changing atmospheric conditions," *IEE Proc. - Gener. Transm. Distrib.*, vol. 142, no. 1, pp. 59–64, Jan. 1995, doi: 10.1049/ip-gtd:19951577.
- [133] D. P. Hohm and M. E. Ropp, "Comparative study of maximum power point tracking algorithms," *Prog. Photovolt. Res. Appl.*, vol. 11, no. 1, pp. 47–62, 2003, doi: 10.1002/pip.459.

- [134] B. Subudhi and R. Pradhan, "A Comparative Study on Maximum Power Point Tracking Techniques for Photovoltaic Power Systems," *IEEE Trans. Sustain. Energy*, vol. 4, no. 1, pp. 89–98, Jan. 2013, doi: 10.1109/TSTE.2012.2202294.
- [135] N. Femia, G. Petrone, G. Spagnuolo, and M. Vitelli, "Optimization of perturb and observe maximum power point tracking method," *IEEE Trans. Power Electron.*, vol. 20, no. 4, pp. 963–973, July 2005, doi: 10.1109/TPEL.2005.850975.
- [136] G. K. Ottman, H. F. Hofmann, A. C. Bhatt, and G. A. Lesieutre, "Adaptive piezoelectric energy harvesting circuit for wireless remote power supply," *IEEE Trans. Power Electron.*, vol. 17, no. 5, pp. 669–676, Sept. 2002, doi: 10.1109/TPEL.2002.802194.
- [137] "Maximizing DC power in energy harvesting circuits using multi-sine excitation." Accessed: Sept. 04, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/5973371>
- [138] "BQ25570 data sheet, product information and support | TI.com." Accessed: Sept. 04, 2025. [Online]. Available: <https://www.ti.com/product/BQ25570>
- [139] "ADP5091 | data sheet and product info Ultralow Power Energy Harvester PMU with MPPT and Charge Management | Analog Devices." Accessed: Sept. 04, 2025. [Online]. Available: <https://www.analog.com/en/products/adp5091.html>
- [140] "SPV1050 | Product - STMicroelectronics." Accessed: Sept. 04, 2025. [Online]. Available: <https://www.st.com/en/power-management/spv1050.html>
- [141] S. Lineykin and S. Ben-Yaakov, "Modeling and Analysis of Thermoelectric Modules," *IEEE Trans. Ind. Appl.*, vol. 43, no. 2, pp. 505–512, Mar. 2007, doi: 10.1109/TIA.2006.889813.
- [142] A. Montecucco and A. R. Knox, "Maximum Power Point Tracking Converter Based on the Open-Circuit Voltage Method for Thermoelectric Generators," *IEEE Trans. Power Electron.*, vol. 30, no. 2, pp. 828–839, Feb. 2015, doi: 10.1109/TPEL.2014.2313294.
- [143] "LTC3108 Datasheet and Product Info | Analog Devices." Accessed: Sept. 04, 2025. [Online]. Available: <https://www.analog.com/en/products/ltc3108.html>
- [144] E. Lefeuvre, A. Badel, C. Richard, L. Petit, and D. Guyomar, "A comparison between several vibration-powered piezoelectric generators for standalone systems," *Sens. Actuators Phys.*, vol. 126, no. 2, pp. 405–416, Feb. 2006, doi: 10.1016/j.sna.2005.10.043.
- [145] Z. L. Wang, "Triboelectric Nanogenerators as New Energy Technology for Self-Powered Systems and as Active Mechanical and Chemical Sensors," *ACS Nano*, vol. 7, no. 11, pp. 9533–9557, Nov. 2013, doi: 10.1021/nn404614z.
- [146] H. Liu, H. Fu, L. Sun, C. Lee, and E. M. Yeatman, "Hybrid energy harvesting technology: From materials, structural design, system integration to applications," *Renew. Sustain. Energy Rev.*, vol. 137, p. 110473, Mar. 2021, doi: 10.1016/j.rser.2020.110473.
- [147] Z. Gao, S. Wang, Y. Li, and M. Chen, "Review of the Multi-Input Single-Inductor Multi-Output Energy Harvesting Interface Applied in Wearable Electronics," *Front. Electron.*, vol. 2, p. 793780, Dec. 2021, doi: 10.3389/felec.2021.793780.
- [148] V. Gogolou, T. Noulis, and V. F. Pavlidis, "Multi-Source Energy Harvesting Systems Integrated in Silicon: A Comprehensive Review," *Electronics*, vol. 14, no. 10, p. 1951, Jan. 2025, doi: 10.3390/electronics14101951.
- [149] S. S. Amin and P. P. Mercier, "MISIMO: A Multi-Input Single-Inductor Multi-Output Energy Harvesting Platform in 28-nm FDSOI for Powering Net-Zero-Energy Systems," *IEEE J. Solid-State Circuits*, vol. 53, no. 12, pp. 3407–3419, Dec. 2018, doi: 10.1109/jssc.2018.2865467.
- [150] C. W. Liu, H. H. Lee, P. C. Liao, Y. L. Chen, M. J. Chung, and P.-H. Chen, "Dual-Source Energy-Harvesting Interface with Cycle-by-Cycle Source Tracking and Adaptive Peak-Inductor-Current Control," *IEEE J. Solid-State Circuits*, vol. 53, no. 10, pp. 2741–2750, Oct. 2018, doi: 10.1109/JSSC.2018.2844358.
- [151] P.-H. Chen, H.-C. Cheng, and C.-L. Lo, "A Single-Inductor Triple-Source Quad-Mode Energy-Harvesting Interface With Automatic Source Selection and Reversely Polarized Energy Recycling," *IEEE J. Solid-State Circuits*, vol. 54, pp. 2671–2679, Oct. 2019, doi: 10.1109/JSSC.2019.2917549.
- [152] B. Zhu, H. Hu, H. Wang, and Y. Li, "A Multi-Input-Port Bidirectional DC/DC Converter for DC Microgrid Energy Storage System Applications," *Energies*, vol. 13, no. 11, p. 2810, Jan. 2020, doi: 10.3390/en13112810.

- [153] A. Almutairi, K. Sayed, N. Albagami, A. G. Abo-Khalil, and H. Saleeb, "Multi-Port PWM DC-DC Power Converter for Renewable Energy Applications," *Energies*, vol. 14, no. 12, p. 3490, Jan. 2021, doi: 10.3390/en14123490.
- [154] S. Farajdadian, A. Hajizadeh, and M. Soltani, "Recent developments of multiport DC/DC converter topologies, control strategies, and applications: A comparative review and analysis," *Energy Rep.*, vol. 11, pp. 1019–1052, June 2024, doi: 10.1016/j.egyr.2023.12.054.
- [155] "LTC3331 Datasheet and Product Info | Analog Devices." Accessed: Sept. 04, 2025. [Online]. Available: <https://www.analog.com/en/products/ltc3331.html>
- [156] e-peas, "e-peas introduces first energy harvesting PMIC to handle two independent energy sources at the same time." Accessed: Sept. 04, 2025. [Online]. Available: <https://www.prnewswire.com/news-releases/e-peas-introduces-first-energy-harvesting-pmic-to-handle-two-independent-energy-sources-at-the-same-time-301910276.html>
- [157] V. Gogolou, K. Kozalakis, E. Koutroulis, G. Doumenis, and S. Siskos, "An Ultra-Low-Power CMOS Supercapacitor Storage Unit for Energy Harvesting Applications," *Electronics*, vol. 10, no. 17, p. 2097, Jan. 2021, doi: 10.3390/electronics10172097.
- [158] Y. Shao *et al.*, "Design and Mechanisms of Asymmetric Supercapacitors," *Chem. Rev.*, vol. 118, no. 18, pp. 9233–9280, Sept. 2018, doi: 10.1021/acs.chemrev.8b00252.
- [159] N. Nitta, F. Wu, J. T. Lee, and G. Yushin, "Li-ion battery materials: present and future," *Mater. Today*, vol. 18, no. 5, pp. 252–264, June 2015, doi: 10.1016/j.mattod.2014.10.040.
- [160] R. Schmich, R. Wagner, G. Hörpel, T. Placke, and M. Winter, "Performance and cost of materials for lithium-based rechargeable automotive batteries," *Nat. Energy*, vol. 3, no. 4, pp. 267–278, Apr. 2018, doi: 10.1038/s41560-018-0107-2.
- [161] A. Masias, J. Marcicki, and W. A. Paxton, "Opportunities and Challenges of Lithium Ion Batteries in Automotive Applications," *ACS Energy Lett.*, vol. 6, no. 2, pp. 621–630, Feb. 2021, doi: 10.1021/acsenenergylett.0c02584.
- [162] A. Belgibayeva *et al.*, "Lithium-ion batteries for low-temperature applications: Limiting factors and solutions," *J. Power Sources*, vol. 557, p. 232550, Feb. 2023, doi: 10.1016/j.jpowsour.2022.232550.
- [163] T. H. Dubaniewicz, T. L. Barone, C. B. Brown, and R. A. Thomas, "Comparison of thermal runaway pressures within sealed enclosures for nickel manganese cobalt and iron phosphate cathode lithium-ion cells," *J. Loss Prev. Process Ind.*, vol. 76, p. 104739, May 2022, doi: 10.1016/j.jlp.2022.104739.
- [164] X. Yi, C. Kirk, and N. Robertson, "Nickel hydroxide-based energy storage devices: nickel-metal hydride batteries vs. nickel hydroxide hybrid supercapacitors," *Carbon Neutrality*, vol. 3, no. 1, p. 39, Dec. 2024, doi: 10.1007/s43979-024-00114-7.
- [165] G. J. May, A. Davidson, and B. Monahov, "Lead batteries for utility energy storage: A review," *J. Energy Storage*, vol. 15, pp. 145–157, Feb. 2018, doi: 10.1016/j.est.2017.11.008.
- [166] N. Vangapally *et al.*, "Lead-acid batteries and lead-carbon hybrid systems: A review," *J. Power Sources*, vol. 579, p. 233312, Sept. 2023, doi: 10.1016/j.jpowsour.2023.233312.
- [167] A. Yao, S. M. Benson, and W. C. Chueh, "Critically assessing sodium-ion technology roadmaps and scenarios for techno-economic competitiveness against lithium-ion batteries," *Nat. Energy*, vol. 10, no. 3, pp. 404–416, Mar. 2025, doi: 10.1038/s41560-024-01701-9.
- [168] C. Wu *et al.*, "Hard carbon for sodium-ion batteries: progress, strategies and future perspective," *Chem. Sci.*, vol. 15, no. 17, pp. 6244–6268, May 2024, doi: 10.1039/D4SC00734D.
- [169] M. Li *et al.*, "Low-temperature performance of Na-ion batteries," *Carbon Energy*, vol. 6, no. 10, p. e546, 2024, doi: 10.1002/cey2.546.
- [170] P. Hong, C. Xu, C. Yan, Y. Dong, H. Zhao, and Y. Lei, "Prussian Blue and Its Analogues for Commercializing Fast-Charging Sodium/Potassium-Ion Batteries," *ACS Energy Lett.*, vol. 10, no. 2, pp. 750–778, Feb. 2025, doi: 10.1021/acsenenergylett.4c02915.
- [171] Y.-J. Guo *et al.*, "Sodium layered oxide cathodes: properties, practicality and prospects," *Chem. Soc. Rev.*, vol. 53, no. 15, pp. 7828–7874, July 2024, doi: 10.1039/D4CS00415A.
- [172] C. Murray, "World's largest sodium-ion BESS comes online in China as it seeks to diversify away from lithium," *Energy-Storage.News*. Accessed: Sept. 13, 2025. [Online]. Available:

- <https://www.energy-storage.news/first-half-world-largest-200mwh-sodium-ion-project-comes-online-china/>
- [173] “China’s CATL launches new sodium-ion battery brand,” *Reuters*, Apr. 21, 2025. Accessed: Sept. 13, 2025. [Online]. Available: <https://www.reuters.com/technology/chinese-battery-maker-catl-launches-second-generation-fast-charging-battery-2025-04-21/>
 - [174] “First sodium-ion battery EVs go into serial production in China - electrive.com.” Accessed: Sept. 13, 2025. [Online]. Available: <https://www.electrive.com/2024/01/02/first-sodium-ion-battery-evs-go-into-serial-production-in-china/>
 - [175] S. Dareen and S. Dareen, “Natron plans \$1.4 bln sodium-ion battery plant in North Carolina,” *Reuters*, Aug. 15, 2024. Accessed: Sept. 13, 2025. [Online]. Available: <https://www.reuters.com/business/energy/natron-plans-14-bln-sodium-ion-battery-plant-north-carolina-source-says-2024-08-15/>
 - [176] “Northvolt develops state-of-the-art sodium-ion battery.” Accessed: Sept. 13, 2025. [Online]. Available: <https://northvolt.com/articles/northvolt-sodium-ion/>
 - [177] F. Campaioli, S. Gherardini, J. Q. Quach, M. Polini, and G. M. Andolina, “Colloquium: Quantum batteries,” *Rev. Mod. Phys.*, vol. 96, no. 3, p. 031001, July 2024, doi: 10.1103/RevModPhys.96.031001.
 - [178] J. Q. Quach *et al.*, “Superabsorption in an organic microcavity: Toward a quantum battery,” *Sci. Adv.*, vol. 8, no. 2, p. eabk3160, Jan. 2022, doi: 10.1126/sciadv.abk3160.
 - [179] J. Q. Quach, G. Cerullo, and T. Virgili, “Quantum batteries: The future of energy storage?,” *Joule*, vol. 7, no. 10, pp. 2195–2200, Oct. 2023, doi: 10.1016/j.joule.2023.09.003.
 - [180] H. Ahuja, P. Dhapola, N. Sahoo, V. Singh, and P. Singh, “Ionic liquid (1-hexyl-3-methylimidazolium iodide)-incorporated biopolymer electrolyte for efficient supercapacitor,” *High Perform. Polym.*, vol. 32, pp. 220–225, Mar. 2020, doi: 10.1177/0954008319897763.
 - [181] A. Brandt, S. Pohlmann, A. Varzi, A. Balducci, and S. Passerini, “Ionic liquids in supercapacitors,” *MRS Bull.*, vol. 38, no. 7, pp. 554–559, July 2013, doi: 10.1557/mrs.2013.151.
 - [182] P. Simon and Y. Gogotsi, “Perspectives for electrochemical capacitors and related devices,” *Nat. Mater.*, vol. 19, no. 11, pp. 1151–1163, Nov. 2020, doi: 10.1038/s41563-020-0747-z.
 - [183] A. Y. Rychagov *et al.*, “Self-Discharge Processes in Symmetrical Supercapacitors with Activated Carbon Electrodes,” *Materials*, vol. 16, no. 19, p. 6415, Jan. 2023, doi: 10.3390/ma16196415.
 - [184] Z. Wang *et al.*, “Extremely low self-discharge solid-state supercapacitors via the confinement effect of ion transfer,” *J. Mater. Chem. A*, vol. 7, no. 14, pp. 8633–8640, Apr. 2019, doi: 10.1039/C9TA01028A.
 - [185] C. Merlet *et al.*, “On the molecular origin of supercapacitance in nanoporous carbon electrodes,” *Nat. Mater.*, vol. 11, no. 4, pp. 306–310, Apr. 2012, doi: 10.1038/nmat3260.
 - [186] V. Augustyn, P. Simon, and B. Dunn, “Pseudocapacitive oxide materials for high-rate electrochemical energy storage,” *Energy Environ. Sci.*, vol. 7, no. 5, pp. 1597–1614, Apr. 2014, doi: 10.1039/C3EE44164D.
 - [187] M. R. Lukatskaya, B. Dunn, and Y. Gogotsi, “Multidimensional materials and device architectures for future hybrid energy storage,” *Nat. Commun.*, vol. 7, no. 1, p. 12647, Sept. 2016, doi: 10.1038/ncomms12647.
 - [188] B. Anasori, M. R. Lukatskaya, and Y. Gogotsi, “2D metal carbides and nitrides (MXenes) for energy storage,” *Nat. Rev. Mater.*, vol. 2, no. 2, p. 16098, Jan. 2017, doi: 10.1038/natrevmats.2016.98.
 - [189] T. Brousse, D. Bélanger, and J. W. Long, “To Be or Not To Be Pseudocapacitive?,” *J. Electrochem. Soc.*, vol. 162, no. 5, p. A5185, Mar. 2015, doi: 10.1149/2.0201505jes.
 - [190] I. Masklavanos, C. Koutsos, and G. Doumenis, “Efficiency of SuperCap storage in energy harvesting micropower supplies,” in *2022 7th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM)*, Sept. 2022, pp. 1–6. doi: 10.1109/SEEDA-CECNSM57760.2022.9932904.
 - [191] H. Rezaei, S. E. Abdollahi, S. Abdollahi, and S. Filizadeh, “Energy management strategies of battery-ultracapacitor hybrid storage systems for electric vehicles: Review, challenges, and

- future trends,” *J. Energy Storage*, vol. 53, p. 105045, Sept. 2022, doi: 10.1016/j.est.2022.105045.
- [192] A. Urooj and A. Nasir, “Review of Hybrid Energy Storage Systems for Hybrid Electric Vehicles,” *World Electr. Veh. J.*, vol. 15, no. 8, p. 342, Aug. 2024, doi: 10.3390/wevj15080342.
- [193] D. Karimi, H. Behi, J. Van Mierlo, and M. Bercebar, “A Comprehensive Review of Lithium-Ion Capacitor Technology: Theory, Development, Modeling, Thermal Management Systems, and Applications,” *Molecules*, vol. 27, no. 10, p. 3119, Jan. 2022, doi: 10.3390/molecules27103119.
- [194] J. Ding, W. Hu, E. Paek, and D. Mitlin, “Review of Hybrid Ion Capacitors: From Aqueous to Lithium to Sodium,” *Chem. Rev.*, vol. 118, no. 14, pp. 6457–6498, July 2018, doi: 10.1021/acs.chemrev.8b00116.
- [195] G. L. Plett, *Battery management systems*. in Artech House Power engineering series. Boston: Artech house, 2016.
- [196] M. A. Hannan, M. S. H. Lipu, A. Hussain, and A. Mohamed, “A review of lithium-ion battery state of charge estimation and management system in electric vehicle applications: Challenges and recommendations,” *Renew. Sustain. Energy Rev.*, vol. 78, pp. 834–854, Oct. 2017, doi: 10.1016/j.rser.2017.05.001.
- [197] X. Hu, F. Feng, K. Liu, L. Zhang, J. Xie, and B. Liu, “State estimation for advanced battery management: Key challenges and future trends,” *Renew. Sustain. Energy Rev.*, vol. 114, p. 109334, Oct. 2019, doi: 10.1016/j.rser.2019.109334.
- [198] R. Xiong, X. Gong, C. C. Mi, and F. Sun, “A robust state-of-charge estimator for multiple types of lithium-ion batteries using adaptive extended Kalman filter,” *J. Power Sources*, vol. 243, pp. 805–816, Dec. 2013, doi: 10.1016/j.jpowsour.2013.06.076.
- [199] S. Yuan, H. Wu, and C. Yin, “State of Charge Estimation Using the Extended Kalman Filter for Battery Management Systems Based on the ARX Battery Model,” *Energies*, vol. 6, no. 1, pp. 444–470, Jan. 2013, doi: 10.3390/en6010444.
- [200] J. Xing and P. Wu, “State of Charge Estimation of Lithium-Ion Battery Based on Improved Adaptive Unscented Kalman Filter,” *Sustainability*, vol. 13, no. 9, p. 5046, Jan. 2021, doi: 10.3390/su13095046.
- [201] B. R. Buchli, “Power Subsystem Design and Management for Solar Energy Harvesting Embedded Systems,” ETH Zurich, 2015. doi: 10.3929/ETHZ-A-010476716.
- [202] B. Buchli, D. Aschwanden, and J. Beutel, “Battery State-of-Charge Approximation for Energy Harvesting Embedded Systems,” in *Wireless Sensor Networks*, P. Demeester, I. Moerman, and A. Terzis, Eds., Berlin, Heidelberg: Springer, 2013, pp. 179–196. doi: 10.1007/978-3-642-36672-7_12.
- [203] A. Nadeau, G. Sharma, and T. Soyata, “State-of-charge estimation for supercapacitors: A Kalman filtering formulation,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Florence, Italy: IEEE, May 2014, pp. 2194–2198. doi: 10.1109/ICASSP.2014.6853988.
- [204] C. Liu, Q. Li, and K. Wang, “State-of-charge estimation and remaining useful life prediction of supercapacitors,” *Renew. Sustain. Energy Rev.*, vol. 150, p. 111408, Oct. 2021, doi: 10.1016/j.rser.2021.111408.
- [205] “MAX17260 Datasheet and Product Info | Analog Devices.” Accessed: Sept. 11, 2025. [Online]. Available: <https://www.analog.com/en/products/max17260.html>
- [206] “BQ27441-G1 data sheet, product information and support | TI.com.” Accessed: Sept. 11, 2025. [Online]. Available: <https://www.ti.com/product/BQ27441-G1>
- [207] A. Chakraborty, A. K. Jha, A. Deo, A. Maity, and A. Patra, “A Scalable Single-Inductor Multiple-Output DC-DC Converter With Constant Charge-Transfer and Power-Up Sequencing for IoT Applications,” *IEEE Trans. Circuits Syst. Regul. Pap.*, vol. 71, pp. 2964–2975, Jan. 2024, doi: 10.1109/TCSI.2024.3349404.
- [208] J. Zidar, T. Matić, I. Aleksi, and Ž. Hocenski, “Dynamic Voltage and Frequency Scaling as a Method for Reducing Energy Consumption in Ultra-Low-Power Embedded Systems,” *Electronics*, vol. 13, no. 5, p. 826, Jan. 2024, doi: 10.3390/electronics13050826.

- [209] S. Khriji, R. Chéour, and O. Kanoun, “Dynamic Voltage and Frequency Scaling and Duty-Cycling for Ultra Low-Power Wireless Sensor Nodes,” *Electronics*, vol. 11, no. 24, p. 4071, Jan. 2022, doi: 10.3390/electronics11244071.
- [210] X. Shi, M. Cai, and Y. Jiang, “Key Role of Cold-Start Circuits in Low-Power Energy Harvesting Systems: A Research Review,” *J. Low Power Electron. Appl.*, vol. 14, no. 4, p. 55, Dec. 2024, doi: 10.3390/jlpea14040055.
- [211] S. Naifar, O. Kanoun, and C. Trigona, “Energy Harvesting Technologies and Applications for the Internet of Things and Wireless Sensor Networks,” *Sensors*, vol. 24, no. 14, p. 4688, Jan. 2024, doi: 10.3390/s24144688.
- [212] “Primer on PowerPath Controllers, Ideal Diodes & Prioritizers | Analog Devices.” Accessed: Sept. 14, 2025. [Online]. Available: <https://www.analog.com/en/resources/technical-articles/primer-on-powerpath-controllers-ideal-diodes-prioritizers.html>
- [213] “Ideal Diode Controller Eliminates Energy Wasting Diodes in Power OR-ing Applications | Analog Devices.” Accessed: Sept. 14, 2025. [Online]. Available: <https://www.analog.com/en/resources/technical-articles/ideal-diode-controller-eliminates-energy-wasting-diodes-in-power-or-ing-applications.html>
- [214] “Basics of Power MUX (Rev. A).” Accessed: Sept. 14, 2025. [Online]. Available: <https://www.ti.com/lit/an/slvae51a/slvae51a.pdf>
- [215] “LTC4412 - Low Loss PowerPath Controller in ThinSOT.” Accessed: Sept. 14, 2025. [Online]. Available: <https://www.analog.com/media/en/technical-documentation/data-sheets/4412fb.pdf>
- [216] “Integrated Ideal Diode Source Selector and eFuse Improve System Robustness | Analog Devices.” Accessed: Sept. 14, 2025. [Online]. Available: <https://www.analog.com/en/resources/analog-dialogue/articles/integrated-ideal-diode-source-selector.html>
- [217] “Basics of Ideal Diodes (Rev. B).” Accessed: Sept. 14, 2025. [Online]. Available: <https://www.ti.com/lit/an/slvae57b/slvae57b.pdf>
- [218] M. Balle, W. Xu, K. F. Darras, and T. C. Wanger, “A Power Management and Control System for Portable Ecosystem Monitoring Devices,” *IEEE Trans. AgriFood Electron.*, vol. 3, no. 1, pp. 134–143, Mar. 2025, doi: 10.1109/TAFE.2024.3472493.
- [219] “Adding Hysteresis for Smooth Undervoltage and Overvoltage Lockout | Analog Devices.” Accessed: Sept. 14, 2025. [Online]. Available: <https://www.analog.com/en/resources/analog-dialogue/articles/adding-hysteresis-for-smooth-undervoltage-and-overvoltage-lockout.html>
- [220] Y. Alishzadeh, S. J. Jassbi, and M. Reshadi, “Energy management of fault-tolerant real-time embedded systems through switching-activity-based techniques,” *Microprocess. Microsyst.*, vol. 102, p. 104929, Oct. 2023, doi: 10.1016/j.micpro.2023.104929.
- [221] “A Comprehensive Guide to LDO Regulators: Navigating Noise Compromise Applications and Trends | Analog Devices.” Accessed: Sept. 15, 2025. [Online]. Available: <https://www.analog.com/en/resources/analog-dialogue/articles/a-comprehensive-guide-to-ldo-regulators.html>
- [222] “AN-CM-399 Enhancing LDO PSRR and Noise Performance Measurements Using Capacitors with Low Parasitic Parameters.” Accessed: Sept. 15, 2025. [Online]. Available: https://www.renesas.com/en/document/apn/cm-399-enhancing-ldo-psrr-and-noise-performance-measurements-using-capacitors-low-parasitic?srsId=AfmBOoppxrRNOQdNeLE0Q_kGZLB1nraDW8hDn90RlgZ44ZKB5BaDup-p
- [223] S. P. Singh, “Simple power rail sequencing solutions for complex multi-rail systems”.
- [224] “TPS62840 1.8-V to 6.5-V, 750-mA, 60-nA IQ Step-Down Converter datasheet (Rev. D).” Accessed: Sept. 15, 2025. [Online]. Available: <https://www.ti.com/lit/ds/symlink/tps62840.pdf>
- [225] Y. Wang, J. Xu, and G. Zhou, “A Cross Regulation Analysis for Single-Inductor Dual-Output CCM Buck Converters,” *J. POWER Electron.*, vol. 16, no. 5, pp. 1802–1812, Sept. 2016.
- [226] A. Maioli, K. A. Quinones, S. Ahmed, M. H. Alizai, and L. Mottola, “Dynamic Voltage and Frequency Scaling for Intermittent Computing,” *ACM Trans. Sens. Netw.*, vol. 21, no. 2, pp. 1–34, Mar. 2025, doi: 10.1145/3714470.

- [227] “Benefits of Ultralow Noise Switching Regulator in Noise-Sensitive RF Applications | Analog Devices.” Accessed: Sept. 15, 2025. [Online]. Available: <https://www.analog.com/en/resources/technical-articles/benefits-of-ultralow-noise-switching-regulator-in-noise-sensitive-rf-applications.html>
- [228] S. Schnier, “Minimize noise and ripple with a low-noise buck converter,” 2023.
- [229] “MSP430FR5969 data sheet, product information and support | TI.com.” Accessed: Sept. 15, 2025. [Online]. Available: <https://www.ti.com/product/MSP430FR5969>
- [230] “esp32-c6_datasheet_en.pdf.” Accessed: Sept. 15, 2025. [Online]. Available: https://www.espressif.com/sites/default/files/documentation/esp32-c6_datasheet_en.pdf
- [231] “AM62x Sitara™ Processors datasheet (Rev. B).” Accessed: Sept. 15, 2025. [Online]. Available: <https://www.ti.com/lit/ds/symlink/am623.pdf>
- [232] V. John, “Discrete Power Solution for AM62x,” 2022.
- [233] “IMX8MMCEC.” Accessed: Sept. 15, 2025. [Online]. Available: <https://www.nxp.com/docs/en/data-sheet/IMX8MMCEC.pdf>
- [234] “Artix 7 FPGAs Data Sheet: DC and AC Switching Characteristics,” 2024.
- [235] “2.3.8 2.5V LVCMOS.” Accessed: Sept. 15, 2025. [Online]. Available: <https://onlinedocs.microchip.com/oxy/GUID-DAFD9EF4-C157-47D1-BD7C-1934D43D152B-en-US-2/GUID-DB356DD0-0871-4AA2-8694-6D9F71193C4B.html>
- [236] L. Benini, A. Bogliolo, and G. De Micheli, “A survey of design techniques for system-level dynamic power management,” *IEEE Trans. Very Large Scale Integr. VLSI Syst.*, vol. 8, no. 3, pp. 299–316, June 2000, doi: 10.1109/92.845896.
- [237] C. Moser, L. Thiele, D. Brunelli, and L. Benini, “Adaptive Power Management for Environmentally Powered Systems,” *IEEE Trans. Comput.*, vol. 59, no. 4, pp. 478–491, Apr. 2010, doi: 10.1109/TC.2009.158.
- [238] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava, “Power management in energy harvesting sensor networks,” *ACM Trans Embed Comput Syst*, vol. 6, no. 4, pp. 32-es, Sept. 2007, doi: 10.1145/1274858.1274870.
- [239] C. M. Vigorito, D. Ganesan, and A. G. Barto, “Adaptive Control of Duty Cycling in Energy-Harvesting Wireless Sensor Networks,” in *2007 4th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, June 2007, pp. 21–30. doi: 10.1109/SAHCN.2007.4292814.
- [240] A. Cammarano, C. Petrioli, and D. Spenza, “Pro-Energy: A novel energy prediction model for solar and wind energy-harvesting wireless sensor networks,” in *2012 IEEE 9th International Conference on Mobile Ad-Hoc and Sensor Systems (MASS 2012)*, Oct. 2012, pp. 75–83. doi: 10.1109/MASS.2012.6502504.
- [241] J. Recas Piorno, C. Bergonzini, D. Atienza, and T. Simunic Rosing, “Prediction and management in energy harvested wireless sensor nodes,” in *2009 1st International Conference on Wireless Communication, Vehicular Technology, Information Theory and Aerospace & Electronic Systems Technology*, May 2009, pp. 6–10. doi: 10.1109/WIRELESSVITAE.2009.5172412.
- [242] O. Ozel and S. Ulukus, “Achieving AWGN Capacity Under Stochastic Energy Harvesting,” *IEEE Trans. Inf. Theory*, vol. 58, no. 10, pp. 6471–6483, Oct. 2012, doi: 10.1109/TIT.2012.2204389.
- [243] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. Srivastava, “Design considerations for solar energy harvesting wireless embedded systems,” in *IPSN 2005. Fourth International Symposium on Information Processing in Sensor Networks, 2005.*, Apr. 2005, pp. 457–462. doi: 10.1109/IPSN.2005.1440973.
- [244] J. Sorber, A. Kostadinov, M. Garber, M. Brennan, M. D. Corner, and E. D. Berger, “Eon: a language and runtime system for perpetual systems,” in *Proceedings of the 5th international conference on Embedded networked sensor systems*, in SenSys ’07. New York, NY, USA: Association for Computing Machinery, Nov. 2007, pp. 161–174. doi: 10.1145/1322263.1322279.
- [245] S. Sudevalayam and P. Kulkarni, “Energy Harvesting Sensor Nodes: Survey and Implications,” *IEEE Commun. Surv. Tutor.*, vol. 13, no. 3, pp. 443–461, 2011, doi: 10.1109/SURV.2011.060710.00094.

- [246] D. Balsamo, A. S. Weddell, G. V. Merrett, B. M. Al-Hashimi, D. Brunelli, and L. Benini, "Hibernus: Sustaining Computation During Intermittent Supply for Energy-Harvesting Systems," *IEEE Embed Syst Lett*, vol. 7, no. 1, pp. 15–18, Mar. 2015, doi: 10.1109/LES.2014.2371494.
- [247] P. Zhang, D. Ganesan, and B. Lu, "{QuarkOS}: Pushing the Operating Limits of {Micro-Powered} Sensors," presented at the 14th Workshop on Hot Topics in Operating Systems (HotOS XIV), 2013. Accessed: Sept. 16, 2025. [Online]. Available: <https://www.usenix.org/conference/hotos13/session/zhang>
- [248] J. V. D. Woude and M. Hicks, "Intermittent Computation without Hardware Support or Programmer Intervention," presented at the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), 2016, pp. 17–32. Accessed: Sept. 22, 2025. [Online]. Available: <https://www.usenix.org/conference/osdi16/technical-sessions/presentation/vanderwoude>
- [249] K. Maeng, A. Colin, and B. Lucia, "Alpaca: intermittent execution without checkpoints," *Proc ACM Program Lang*, vol. 1, no. OOPSLA, p. 96:1-96:30, Oct. 2017, doi: 10.1145/3133920.
- [250] A. Y. Majid *et al.*, "Dynamic Task-based Intermittent Execution for Energy-harvesting Devices," *ACM Trans. Sens. Netw.*, vol. 16, no. 1, pp. 1–24, Feb. 2020, doi: 10.1145/3360285.
- [251] J. Zeng *et al.*, "Enabling Volatile Caches for Energy Harvesting Systems," in *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, Oct. 2021, pp. 170–182. doi: 10.1145/3466752.3480102.
- [252] M. Surbatovich, L. Jia, and B. Lucia, "Automatically enforcing fresh and consistent inputs in intermittent systems," in *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, Virtual Canada: ACM, June 2021, pp. 851–866. doi: 10.1145/3453483.3454081.
- [253] J. Choi, J. Zeng, D. Lee, C. Min, and C. Jung, "Write-Light Cache for Energy Harvesting Systems," in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, in ISCA '23. New York, NY, USA: Association for Computing Machinery, June 2023, pp. 1–13. doi: 10.1145/3579371.3589098.
- [254] F. Ait Aoudia, M. Gautier, M. Magno, O. Berder, and L. Benini, "Leveraging Energy Harvesting and Wake-Up Receivers for Long-Term Wireless Sensor Networks," *Sensors*, vol. 18, no. 5, p. 1578, May 2018, doi: 10.3390/s18051578.
- [255] N. Jackson, J. Adkins, and P. Dutta, "Capacity over capacitance for reliable energy harvesting sensors," in *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*, in IPSN '19. New York, NY, USA: Association for Computing Machinery, Apr. 2019, pp. 193–204. doi: 10.1145/3302506.3310400.
- [256] A. Sabovic, A. K. Sultania, C. Delgado, L. D. Roeck, and J. Famaey, "An Energy-Aware Task Scheduler for Energy-Harvesting Batteryless IoT Devices," *IEEE Internet Things J.*, vol. 9, no. 22, pp. 23097–23114, Nov. 2022, doi: 10.1109/JIOT.2022.3185321.
- [257] L. Li and C. Han, "ASARIMA: An Adaptive Harvested Power Prediction Model for Solar Energy Harvesting Sensor Networks," *Electronics*, vol. 11, no. 18, p. 2934, Jan. 2022, doi: 10.3390/electronics11182934.
- [258] S. Kosunalp, "A New Energy Prediction Algorithm for Energy-Harvesting Wireless Sensor Networks With Q-Learning," *IEEE Access*, vol. 4, pp. 5755–5763, 2016, doi: 10.1109/ACCESS.2016.2606541.
- [259] M. Pi, N. Jin, D. Chen, and B. Lou, "Short-Term Solar Irradiance Prediction Based on Multichannel LSTM Neural Networks Using Edge-Based IoT System," *Wirel. Commun. Mob. Comput.*, vol. 2022, no. 1, p. 2372748, 2022, doi: 10.1155/2022/2372748.
- [260] C. Renner, "Solar harvest prediction supported by cloud cover forecasts," in *Proceedings of the 1st International Workshop on Energy Neutral Sensing Systems*, in ENSSys '13. New York, NY, USA: Association for Computing Machinery, Nov. 2013, pp. 1–6. doi: 10.1145/2534208.2534210.

- [261] N. Ashraf, M. Faizan, W. Asif, H. K. Qureshi, A. Iqbal, and M. Lestas, “Energy management in harvesting enabled sensing nodes: Prediction and control,” *J. Netw. Comput. Appl.*, vol. 132, pp. 104–117, Apr. 2019, doi: 10.1016/j.jnca.2019.01.015.
- [262] N. Stricker, R. Da Forno, and L. Thiele, “Robustness of predictive energy harvesting systems: Analysis and adaptive prediction scaling,” *IET Comput. Digit. Tech.*, vol. 16, no. 4, pp. 106–124, 2022, doi: 10.1049/cdt2.12042.
- [263] N. Yamin and G. Bhat, “Uncertainty-aware Energy Harvest Prediction and Management for IoT Devices,” *ACM Trans Autom Electron Syst*, vol. 28, no. 5, p. 87:1–87:33, Sept. 2023, doi: 10.1145/3606372.
- [264] C. Delgado and J. Famaey, “Optimal Energy-Aware Task Scheduling for Batteryless IoT Devices,” *IEEE Trans. Emerg. Top. Comput.*, vol. 10, no. 3, pp. 1374–1387, July 2022, doi: 10.1109/TETC.2021.3086144.
- [265] S. Jankovic and L. Saranovac, “Prediction of Harvested Energy for Wireless Sensor Node,” *Elektron. Ir Elektrotehnika*, vol. 26, no. 1, pp. 23–31, Feb. 2020, doi: 10.5755/j01.eie.26.1.23807.
- [266] E. Jecan, C. Pop, O. Ratiu, and E. Puschita, “Predictive Energy-Aware Routing Solution for Industrial IoT Evaluated on a WSN Hardware Platform,” *Sensors*, vol. 22, no. 6, p. 2107, Jan. 2022, doi: 10.3390/s22062107.
- [267] J. Anuj *et al.*, “An Energy Management Unit for Predictive Solar Energy Harvesting IoT;,” in *Proceedings of the 8th International Conference on Internet of Things, Big Data and Security*, Prague, Czech Republic: SCITEPRESS - Science and Technology Publications, 2023, pp. 39–50. doi: 10.5220/0011839500003482.
- [268] B. Buchli, F. Sutton, J. Beutel, and L. Thiele, “Dynamic power management for long-term energy neutral operation of solar energy harvesting systems,” in *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*, in SenSys ’14. New York, NY, USA: Association for Computing Machinery, Nov. 2014, pp. 31–45. doi: 10.1145/2668332.2668333.
- [269] B. Buchli, P. Kumar, and L. Thiele, “Optimal Power Management with Guaranteed Minimum Energy Utilization for Solar Energy Harvesting Systems,” in *2015 International Conference on Distributed Computing in Sensor Systems*, June 2015, pp. 147–158. doi: 10.1109/DCOSS.2015.9.
- [270] R. Ahmed, B. Buchli, S. Draskovic, L. Sigrist, P. Kumar, and L. Thiele, “Optimal Power Management with Guaranteed Minimum Energy Utilization for Solar Energy Harvesting Systems,” *ACM Trans Embed Comput Syst*, vol. 18, no. 4, p. 30:1–30:26, June 2019, doi: 10.1145/3317679.
- [271] B. Buchli, F. Sutton, J. Beutel, and L. Thiele, “Towards Enabling Uninterrupted Long-Term Operation of Solar Energy Harvesting Embedded Systems,” in *Wireless Sensor Networks*, B. Krishnamachari, A. L. Murphy, and N. Trigoni, Eds., Cham: Springer International Publishing, 2014, pp. 66–83. doi: 10.1007/978-3-319-04651-8_5.
- [272] K. Geissdoerfer, R. Jurdak, B. Kusy, and M. Zimmerling, “Getting more out of energy-harvesting systems: energy management under time-varying utility with PreAct,” in *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*, in IPSN ’19. New York, NY, USA: Association for Computing Machinery, Apr. 2019, pp. 109–120. doi: 10.1145/3302506.3310393.

[Οπισθόφυλλο. Κενή σελίδα]