

Harvest-Aware Real-Time Scheduling for Solar-Powered Edge Nodes under African Climatic Conditions

Raphael O. Ajah

ajahraphaelogah@gmail.com

University of Lagos, Nigeria

Research Article

Keywords: energy harvesting, real-time scheduling, mixed criticality, energy reserve, solar variability, extreme-value statistics, Africa

Posted Date: October 12th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11175504/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Harvest-Aware Real-Time Scheduling for Solar-Powered Edge Nodes under African Climatic Conditions

Anonymous author(s)

(identities withheld for double-anonymous review)

Abstract

Solar-powered edge nodes monitor mini-grids, farms and water systems across Africa, often where grid power is absent. Their deadlines are only as reliable as their energy supply, and sunshine varies with the night, the weather, the season and the year. We consider uniprocessor EDF nodes whose tasks are split into HI-criticality tasks, which must always meet their deadlines, and LO-criticality tasks, whose jobs may be rejected at release. We propose a reserve-based admission policy that admits a LO job only if the stored energy not yet committed to other LO jobs stays above a reserve. We prove that HI jobs, and all admitted LO jobs, then meet their deadlines whenever the reserve covers the worst harvest deficit. The result extends to seasonal reserves and to any energy manager filtered through the same test. The reserve of a harvest record is the maximum drawdown of a net-energy process and is computed exactly in linear time. Using 25 years (2001–2025) of NASA POWER data for six African sites, we find that the reserve for the same node differs by a factor of 3.6 across climates, and that reserves calibrated on one year of data are exceeded in 48% of later years. With a 1Wh store and reserves calibrated out of sample, the seasonal policy missed no HI deadline in 25 years at five sites. The baselines either missed HI jobs, on up to 2894 of 9131 days, or served 10–12 percentage points fewer LO jobs. Seasonal reserves doubled the LO work served at the most seasonal site.

Keywords: energy harvesting; real-time scheduling; mixed criticality; energy reserve; solar variability; extreme-value statistics; Africa

1 Introduction

In 2023 about 666 million people still lived without electricity, and roughly 85% of them were in sub-Saharan Africa [16]. Much of the new supply on the continent is decentralized: solar home systems, solar-powered water pumping, and village mini-grids. The equipment that monitors and protects these systems, meters their feeders, and reports faults is increasingly a small embedded *edge node* that is itself solar powered, because it sits exactly where grid power is absent or unreliable. Such nodes run a mix of activities with very different timing and importance. A fault check that must run every few seconds, a local alarm link, and interval telemetry to the operator are *critical*: missing them can leave a fault undetected or a mini-grid unbilled. Waveform analytics, on-device learning and bulk uploads are *useful* but can be postponed or skipped.

For these nodes, energy, not processor time, is the binding constraint. A few percent of processor utilization is typical, but the harvested energy varies on at least four time scales: the day–night cycle (fully predictable), weather over hours to days, the seasons, and the year-to-year variability of the seasons themselves. Across Africa the seasonal and interannual structure differs sharply between climate zones.

The West African monsoon brings weeks of thick cloud to the Guinea coast; the Sahel alternates between a dusty dry season and a short rainy season; equatorial highlands have two rainy seasons; the Mediterranean south-west of the continent has short, cloudy winter days. A node sized with a comfortable margin for the *average* day can still run out of energy in a bad week, and when it does, a naive scheduler loses the critical and the optional work alike.

The real-time community has developed a substantial theory of scheduling under energy-harvesting constraints. Lazy scheduling and its analysis with energy variability characterization curves [24], energy-aware DVFS [22], the optimal ED-H scheduler [7], fixed-priority analyses [1, 2], and weakly-hard and (m,k) -firm designs for energy-variable systems [28, 29] all take a model of the future harvest as *given*, typically as a lower curve on the energy harvested in any window. Energy-management work on sensor networks, in turn, adapts duty cycles to forecasts to reach energy neutrality [4, 18, 27]. Two practical questions sit between these bodies of work, and they are the subject of this paper:

1. Which work should be sacrificed, and how can the rest be guaranteed? When energy is insufficient, dropping jobs is unavoidable. We want the node to shed only degradable work, and to be able to *prove* that critical deadlines hold.

2. Where does the harvest model come from, and how much data does it need? A guarantee is only as good as its lower harvest curve. In practice curves are estimated from a pilot deployment or a “typical year”. The deficits that matter are rare multi-day events whose depth varies from year to year, so a curve estimated from a short record can be dangerously optimistic.

Approach. We consider a uniprocessor node scheduled by EDF whose tasks are split into HI-criticality tasks, whose jobs must always meet their deadlines, and LO-criticality tasks, whose jobs may be rejected at release when energy is scarce. We propose RES, a reserve-based admission policy: a LO job is admitted only if, after admitting it, the stored energy that is not yet committed to other LO jobs stays above a reserve level R . HI jobs are never rejected. The policy is a constant-time check per job release and needs no harvest forecast at run time; all knowledge about the climate is folded into R , which may vary with the season.

Contributions.

- *Guarantee.* We prove (Theorem 1) that under RES every HI job and every admitted LO job meets its deadline, and the node never browns out, whenever the reserve covers the worst deficit of the harvest relative to the HI energy demand. The proof handles storage overflow (harvest lost when the store is full) through an admission cap, and extends to season-dependent reserves (RES-S). The condition is tight up to a few job energies (Proposition 1).

- *Exact, linear-time reserve from long records.* We show that the reserve required by a harvest record is the maximum drawdown of a simple net-energy process, which the sequent-peak method of reservoir design [37] and Lindley’s recursion [20] compute in one pass. This makes it practical to evaluate decades of minute-resolution data, per season and per year, and to treat the reserve statistically: we fit extreme-value distributions [8] to annual requirements and choose R for a target return period.

- *A multi-decade African study.* Using 25 years (2001–2025) of satellite-derived daily irradiance, temperature and rainfall from NASA POWER for six sites that span the continent’s main climate zones (Table 2), we quantify (i) how much the required reserve differs between climates for the same node design, (ii) how often reserves calibrated on short records fail out of sample, (iii) what seasonal reserves recover in LO quality of service, and (iv) how RES compares with greedy, threshold and energy-neutral policies. For the same node, the required reserve differs by a factor

of 3.6 between sites, and a reserve calibrated on one year of data is exceeded in about half of the following years. With a 1Wh store and out-of-sample calibration, the seasonal policy misses no HI deadline in 25 years at five of the six sites while serving 93.9–98.9% of LO jobs. At the sixth, Cape Town, it misses on 4 days, three of them in a winter storm deeper than anything else in the record.

- *Validated tooling.* A minute-step simulator, cross-checked against an exact job-level EDF simulator with continuous energy dynamics, and the full data pipeline will be released as open source.

Organization. Section 2 reviews related work. Section 3 states the system model and Section 4 the analysis. Section 5 explains how reserves are computed and calibrated from climate records. Sections 6 and 7 present the experimental setup and results, Section 8 discusses implications for deployments in Africa and the limitations of the study, and Section 9 concludes.

2 Related work

Real-time scheduling with energy harvesting. Moser et al. introduced lazy scheduling for energy-harvesting sensor nodes and an admission test based on energy variability characterization curves, which bound the energy harvested in any window [24]. Liu, Qiu and Wu combined harvesting awareness with DVFS to stretch execution when stored energy is short [22]. Following early frameworks for harvesting-aware real-time scheduling [12], Chetto proposed ED-H, an EDF-based scheduler that decides when to idle and when to execute depending on slack time and slack energy, and proved it optimal for a model with a finite energy store [7]. For fixed priorities, Abdeddaïm et al. studied the optimality of PFPasap [1] and derived response-time analyses, showing that the synchronous release is not necessarily the worst case and that deadline-monotonic priorities are optimal for their model [2]. In their basic form, these works decide *when* to execute so that a given job set is feasible for a given harvest model; they do not choose *which* jobs to give up when the job set is not energy-feasible, and the harvest model is an input rather than an object of study.

Several works trade quality for sustainability. Moser, Chen and Thiele manage discrete service levels under harvesting [25]. Shirazi, Kargahi and Thiele schedule weakly-hard and (m,k) -firm tasks on energy-variable platforms by switching between performance levels [28, 29]. Our LO tasks are degradable in a similar spirit, but we give HI tasks a hard guarantee and make the harvest model, and the risk of getting it wrong, a first-class object of study.

Energy-neutral power management. In the sensor-network literature, Kansal et al. formalized energy-neutral operation and adapted duty cycles to an exponentially weighted moving-average (EWMA) harvest forecast [18]. Better short-term predictors followed, for example weather-conditioned moving averages [27]. Buchli et al. targeted long-term energy neutrality for solar systems over seasonal horizons [4], and Geissdoerfer et al. managed energy under time-varying utility [13]. See [32] for a survey. These controllers aim for good average utility. They do not provide per-job deadline guarantees, and their behaviour in rare deep deficits depends on how quickly the forecast and the battery feedback react. We use an EWMA energy-neutral controller as a baseline.

116 Mixed criticality and energy. Mixed-criticality scheduling originated with Vestal’s model of execution 117
time assurance levels [33]; Burns and Davis survey the field [5]. Energy has been brought into mixed-
118 criticality analysis for battery-less systems with energy uncertainty [3], for response-time analysis
119 of energy-harvesting systems in the adaptive mixed-criticality framework [35], for mixed-criticality 120
scheduling of energy-harvesting systems more generally [34], and for heterogeneous multicores
powered

121 by harvesting [36]. In our model the uncertain resource is energy rather than execution time:
criticality 122 determines which jobs are protected when the harvest falls short. The policy is
deliberately simple. Our 123 focus is on making its single parameter, the reserve, trustworthy for real
climates.

124 Intermittent computing. Batteryless devices that compute through frequent power failures are an 125
active area [15, 23], including deadline-aware scheduling of computation and harvesting [17]. We target 126
nodes with a rechargeable store (a small battery or supercapacitor bank) that must avoid brownouts 127
altogether, which is the norm for infrastructure monitoring.

128 Storage sizing, queues and network calculus. Our reserve is the energy analogue of a backlog
129 bound. The maximum of cumulative demand minus cumulative supply over all windows is the quantity




Figure 1: Solar-powered node. The storage level $B(t)$ integrates harvested power minus consumption. HI jobs are always released; a LO job is admitted only if the uncommitted energy after its admission stays at or above the (possibly seasonal) reserve $R(t)$.

Table 1: Main notation.

Symbol	Meaning
$\tau_i = (C_i, T_i, P_i, \chi_i)$	WCET, period (= relative deadline), board power while running, criticality $\in \{\text{HI}, \text{LO}\}$
P_b, P_l	baseline board power while on; storage self-discharge and charger quiescent draw
$e_i = (P_i - P_b)C_i$	worst-case incremental energy of a job of τ_i
$p_h(t), H(s, t)$	power delivered into storage; energy delivered in $[s, t]$
$B(t), B_{\max}$	stored energy; storage capacity
$\pi(t)$	energy committed to admitted, unfinished LO jobs
$B'(t) = B(t) - \pi(t)$	uncommitted stored energy
$\alpha_{\text{HI}}(\Delta)$	upper bound on HI + baseline energy demand in any window of length Δ
ρ, σ	slope and burst of the affine bound $\alpha_{\text{HI}}(\Delta) \leq \rho\Delta + \sigma$

$\beta_c(\Delta)$	lower bound on $H(s, s+\Delta)$ for window starts s in class c
$R(t), R_c, R^*$	reserve at time t ; reserve of class c ; $R^* = \sup_t R(t)$

130 computed by the sequent-peak algorithm for reservoir sizing in hydrology [37], by Lindley’s recursion
for 131 queues [20], and by the vertical deviation between arrival and service curves in network calculus
[9, 19].

132 Stand-alone photovoltaic systems are traditionally sized for a target loss-of-load probability [11]. We 133
connect this machinery to criticality-aware admission, prove the resulting deadline guarantee, and use 134
extreme-value statistics [8] to reason about the year-to-year variability of the requirement.

135 African context and data. Energy access is concentrated in sub-Saharan Africa [16], and dust
is a 136 known cause of photovoltaic losses in West Africa, with the harmattan season a particular concern
in 137 Nigeria [6]. We use NASA POWER data [26, 30], whose surface irradiance from 2001 onwards
comes 138 from the CERES SYN1deg product [10]. Evaluations in the scheduling literature typically
rely on short 139 traces from a few locations. We are not aware of a study of how the energy reserve
needed for real-time 140 guarantees varies across African climate zones, or of its year-to-year statistics.

141 3 System model

142 Figure 1 summarizes the model and Table 1 the notation.

143 Platform and workload. A single processor executes a set $\Gamma = \Gamma_{\text{HI}} \cup \Gamma_{\text{LO}}$ of sporadic tasks under 144
preemptive EDF. Task τ_i has worst-case execution time C_i , minimum inter-arrival time T_i equal to its 145
relative deadline, and criticality χ_i . While the node is on, the board draws a baseline power P_b (sleeping
146 processor, always-on sensing front end, regulators). While a job of τ_i executes, the board draws at
most

147 $P_i \geq P_b$, so a job consumes at most $e_i = (P_i - P_b)C_i$ above baseline. Radio transmissions performed by 148 a
task are part of its execution and power. The store additionally loses P_l to self-discharge and charger 149
quiescent current.

150 Assumption 1. $U = \sum_{\tau_i \in \Gamma} C_i/T_i \leq 1$.

151 Under Assumption 1, EDF meets every deadline of any subset of the jobs of Γ as long as the processor
152 is powered [21]. Low utilization is typical for monitoring nodes; the reference workload in Section 6
has 153 $U = 0.108$. Timing therefore fails only through energy starvation, and that is what the rest of the
model 154 addresses.

155 Energy source and store. Let $p_h(t) \geq 0$ be the power delivered into the store after the panel,
maximum-

156 power-point tracking and charging losses, and $H(s, t) = \int_s^t p_h$. The store holds $B(t) \in [0, B_{\max}]$. With

157 $p_c(t)$ the node’s consumption, $B' = p_h - p_c$ while $0 < B < B_{\max}$. When $B = B_{\max}$, surplus harvest is lost.

158 When $B = 0$ and $p_c > p_h$, the node *browns out*: its supply collapses, the running and pending jobs are 159
lost, and it stays off until the store is recharged to a restart threshold.

160 Criticality semantics. Every HI job must complete by its deadline. A LO job may be *rejected* at its 161 release,
in which case it is dropped; once admitted, it must complete by its deadline. The quality of 162 service of LO
tasks is the fraction of LO jobs that complete (LO QoS). This is a firm-admission model:

shedding happens at the boundary of a job, never in the middle of one, which suits tasks such as uploads that are useless if interrupted.

Harvest curves. The analysis needs lower bounds on harvested energy.

Definition 1 (Lower harvest curves). Partition time into classes $c(t)$ (one class for a static analysis; for example, calendar months for a seasonal one). A function $\theta_c: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a lower harvest curve for class c if $H(s, s+\Delta) \geq \theta_c(\Delta)$ for every $\Delta \geq 0$ and every window start s in the *closure* of class c .

Using closures means that a window starting exactly on a class boundary is covered by the curves of both neighbouring classes; this matters for time-varying reserves. For a static analysis, θ is the energy variability characterization lower curve of Moser et al. [24]. Section 5 computes the curves, or rather the quantities the analysis needs, from multi-decade records, and Section 7 measures how well curves estimated from the past hold in the future.

4 Reserve-based admission and its guarantee

4.1 HI energy demand

Lemma 1 (HI demand bound). *Consider any schedule of Γ in which every executed job meets its deadline.*

For any interval of length $\Delta \geq 0$, the energy consumed by the baseline, the store's self-discharge and the HI jobs in the interval is at most

$$\alpha_{\text{HI}}(\Delta) = (P_b + P_l)\Delta + \sum_{\tau_i \in \Gamma_{\text{HI}}} \left(\left\lceil \frac{\Delta}{T_i} \right\rceil + 1 \right) e_i \leq \rho \Delta + \sigma, \quad (1)$$

where $\rho = P_b + P_l + \sum_{\tau_i \in \Gamma_{\text{HI}}} e_i / T_i$ and $\sigma = 2 \sum_{\tau_i \in \Gamma_{\text{HI}}} e_i$.

Algorithm 1 RES run-time admission (per LO job release).

Require: fuel-gauge estimate B ; committed energy π ; reserve table R_c ; $R^{\wedge}$; $B_{\max}$

```

1: procedure ONLORELEASE(job  $J$  of  $\tau_j$ , time  $t$ )
2:    $r \leftarrow R_{c(t)}$   $\triangleright$  e.g. month-of-year lookup
3:   if  $B - \pi - e_j \geq r$  and  $\pi + e_j \leq B_{\max} - R^{\wedge}$  then
4:      $\pi \leftarrow \pi + e_j$ ; EDFINSERT( $J$ )
5:   else
6:     drop  $J$ 
7:   end if
8: end procedure
9: procedure ONLOPROGRESS(energy  $\varepsilon$  consumed by a LO job above baseline)
10:   $\pi \leftarrow \max(0, \pi - \varepsilon)$   $\triangleright$  on completion, release the job's remaining commitment
11: end procedure

```

Proof. The baseline and self-discharge consume at most $(P_b + P_l)\Delta$. Consider the interval $[t_1, t_1 + \Delta)$.

A job of τ_i that executes in it must have been released after $t_1 - T_i$ (otherwise its deadline, and hence its completion, precedes t_1) and before $t_1 + \Delta$. The open interval $(t_1 - T_i, t_1 + \Delta)$ has length $\Delta + T_i$, so it contains at most $\lceil (\Delta + T_i)/T_i \rceil = \lceil \Delta/T_i \rceil + 1$ releases of a sporadic task with minimum inter-arrival time T_i . Each job consumes at most e_i above the baseline. The affine bound follows from $\lceil x \rceil + 1 \leq x + 2$. $\square$

The consumption of LO jobs above the baseline is deliberately left out of α_{HI} : the admission policy accounts for it separately.

4.2 The RES policy

Let the reserve be a piecewise-constant, right-continuous function $R(t) = R_{c(t)}$, with $R^\wedge = \max_c R_c$. The committed energy $\pi(t)$ is the sum, over admitted LO jobs that have not completed, of e_j minus the incremental energy each has already consumed; when a LO job completes, its unused commitment is released.

Definition 2 (RES admission). HI jobs are always released. A LO job of τ_j released at time t is admitted if and only if

$$B(t) - \pi(t) - e_j \geq R(t), \quad (2)$$

$$\pi(t) + e_j \leq B_{\max} - R^\wedge. \quad (3)$$

Admitted jobs are scheduled together with the HI jobs by EDF. Rejected jobs are dropped.

Condition (2) keeps the uncommitted energy $B' = B - \pi$ at or above the reserve at every admission. Condition (3) bounds the energy tied up in pending LO work so that, when B' is below the reserve, the store cannot be full and no harvest is lost (Step 3 of the proof below). With a single class, (3) follows from (2) and $B \leq B_{\max}$. Algorithm 1 gives the run-time procedure; it performs a constant number of arithmetic operations per LO release and needs only a fuel-gauge reading of $B(t)$.

4.3 Guarantee

Theorem 1 (HI and admitted-LO guarantee). *Let Assumption 1 hold and let β_c be lower harvest curves (Definition 1). Suppose that for every class c*

$$R_c \geq R_c^* := \sup_{\Delta \geq 0} [\alpha_{HI}(\Delta) - \beta_c(\Delta)], \quad (4)$$

that $B_{\max} \geq R$, and that $B(0) \geq R(0)$. Then, under EDF with RES admission, the node never browns out, every HI job meets its deadline, and every admitted LO job meets its deadline.

Proof. Virtual system. Let V be identical to the real system (same releases, same harvest, admission evaluated on V 's own state, same overflow rule) except that its store may become negative instead of browning out. The two systems coincide until V 's store first becomes negative, if ever. In V the processor is always powered, so by Assumption 1 EDF meets the deadline of every HI job and every admitted LO job, and Lemma 1 applies to V on every interval. It therefore suffices to prove $B(t) \geq 0$ in V for all t . Since $\pi \geq 0$, we prove the stronger statement $B'(t) = B(t) - \pi(t) \geq 0$. All quantities below refer to V . **Step 1: committed energy is bounded.** π increases only at admissions, and by (3) it is at most $B_{\max} - R^\wedge$ right after each; between admissions it does not increase. Hence $\pi(t) \leq B_{\max} - R^\wedge$ for all t .

Step 2: the last time above the reserve. Assume, for contradiction, that $B'(t_2) < 0$ for some t_2 . Let

214 $S = \{t \in [0, t_2] : B'(t) \geq R(t)\}; 0 \in S$ because $B'(0) = B(0) \geq R(0)$. Let $t_1 = \sup S$. The functions B' and R
 215 are right-continuous; B jumps down only at admissions and up only when a LO job completes with
 unused

216 commitment, and R jumps only at class boundaries. We claim that $B'(t_1) \geq \tilde{R} := \min\{R(t_1^-), R(t_1)\}$
 .

217 If $t_1 \in S$ this is immediate. Otherwise there are $t_n \uparrow t_1$ in S , so $B'(t_1^-) \geq R(t_1^-)$. A downward jump
 218 of B' at t_1 can only be an admission, which by (2) would leave $B'(t_1) \geq R(t_1)$ and put t_1 in S . Hence
 219 $B'(t_1) \geq B'(t_1^-) \geq R(t_1^-)$. Since $R \geq 0 > B'(t_2)$, we also have $t_1 < t_2$.

220 *Step 3: what can happen in $(t_1, t_2]$.* For every $t \in (t_1, t_2]$, $B'(t) < R(t) \leq R^*$. (i) No LO job is
 221 admitted in $(t_1, t_2]$: an admission at t would leave $B'(t) \geq R(t)$ and put t in S . (ii) No harvest is lost:
 222 $B(t) = B'(t) + \pi(t) < R^* + (B_{\max} - R^*) = B_{\max}$ by Step 1. (iii) While an admitted LO job executes, B and
 223 π decrease by the same amount, so B' is unchanged; a completion can only increase B' .
 Consequently

$$B'(t_2) \geq B'(t_1) + H(t_1, t_2) - E_{\text{HI}}(t_1, t_2),$$

224 where $E_{\text{HI}}(t_1, t_2)$ is the energy consumed by the baseline, self-discharge and HI jobs in $(t_1, t_2]$.

225 *Step 4: contradiction.* Let C' be the class whose reserve equals R^* , i.e. $C' = c(t_1)$ or, if t_1 is a
 226 class boundary, possibly $C' = c(t_1^-)$. In both cases t_1 lies in the closure of C' , so Definition 1 gives 227
 $H(t_1, t_2) \geq \beta_{C'}(\Delta)$ with $\Delta = t_2 - t_1$, and Lemma 1 gives $E_{\text{HI}}(t_1, t_2) \leq \alpha_{\text{HI}}(\Delta)$. Therefore

$$B'(t_2) \geq R_{C'} + \beta_{C'}(\Delta) - \alpha_{\text{HI}}(\Delta) \geq R_{C'} - R_{C'}^* \geq 0,$$

228 contradicting $B'(t_2) < 0$. Hence $B' \geq 0$ and $B \geq 0$ in V at all times; V and the real system coincide, the
 229 real system never browns out, and all HI and admitted LO jobs meet their deadlines. $\square$

230 **Corollary 1 (Static reserve and minimum storage).** *With a single class, RES guarantees all HI
 deadlines if*

231 $R \geq R^* = \sup_{\Delta} [\alpha_{\text{HI}}(\Delta) - \beta(\Delta)]$, $R \leq B_{\max}$ and $B(0) \geq R$; moreover $R^* \leq \sigma + \sup_{\Delta} [\rho \Delta - \beta(\Delta)]$. Because 232 the
 static curve is the pointwise minimum of the class curves, $R^* = \max_C R_C^*$: a seasonal reserve never 233
 reduces the storage needed for the HI guarantee, but it lets LO jobs use more of the store in benign 234
 seasons.

235 **Remark 1 (RES as a safety floor).** The proof uses only two properties of the admission rule: every
 236 admitted LO job satisfies (2) and (3), and HI jobs are never rejected. Any LO admission policy that 237
 admits a LO job only when both tests hold in the current state, such as an energy-neutral controller, a
 238 learned policy or a utility-maximizing planner whose decisions are filtered through the two tests,
 inherits 239 Theorem 1. RES can therefore serve as a certified floor under an arbitrary energy manager.
 Section 7.4 240 evaluates one such combination.

241 Theorem 1 does not require any forecast at run time, and it holds for any LO workload, including
 LO 242 jobs that consume less than e_j or that arrive in bursts. Admitted LO jobs are protected too, which is
 what 243 makes firm admission meaningful: a bulk upload is either not started or completed.

244 **Proposition 1 (Near-tightness).** *Let the HI jobs of periodic tasks execute for their WCET at power P_i*

and let the harvest be a trace h with $D^* = \sup_{s \leq t} [\rho(t-s) - H(s, t)]$. If the store holds $B(s^*) = x$ with no pending LO work at the start s^* of a window attaining D^* , and $x < D^* - \sigma$, then a HI job is lost to brownout in that window. Hence no reserve below $D^* - \sigma$ can guarantee HI deadlines in general, while $D^* + \sigma$ suffices (Corollary 1).

Proof sketch. In a window of length Δ each periodic HI task completes at least $\lfloor \Delta/T_i \rfloor - 1 \geq \Delta/T_i - 2$ full jobs, so HI and baseline consumption is at least $\rho\Delta - \sigma$. As long as the store stays below the reserve (no LO admissions), $B(s^* + \Delta) \leq x + H - (\rho\Delta - \sigma) < D^* - \rho\Delta + H$, which is negative for the window attaining D^* . A negative store is impossible, so the node browns out and at least one HI job is lost. RES produces the precondition $B(s^*) \approx R$ whenever LO demand exceeds the harvest before s^* , because it admits LO work until B reaches the reserve. $\square$

4.4 Exact linear-time computation of the reserve

The supremum in (4) ranges over all window lengths and positions. For a recorded trace it can be evaluated exactly and cheaply.

Proposition 2. Let p_h be piecewise constant on a grid of step δ over $[0, N\delta]$, let $X_k = H(0, k\delta) - \rho k\delta$ for $k = 0, \dots, N$, and let the classes be unions of grid cells. Then

$$\sup_{0 \leq s \leq t \leq N\delta} [\rho(t-s) - H(s, t)] = \sigma + \max_{0 \leq j \leq k \leq N} (X_j - \min_{k \geq j} X_k) \quad (5)$$

$$R_{\text{aff}c} = \max_{s \in c, t \geq s} X(t) - \min_{j: j\delta \in c} X_j \quad (6)$$

and both are computed in $O(N)$ time. $R_{\text{aff}c} \geq R_c^*$ and $R_{\text{aff}c} - R_c^* \leq \sum_{i \in \Gamma_{\text{HI}}} e_i$ when the harvest record is used as β_c .

Proof. $X(t) = H(0, t) - \rho t$ is continuous and piecewise linear with breakpoints on the grid, and $\rho(t-s) - H(s, t) = X(s) - X(t)$. For fixed t , $X(s)$ is linear on each cell, so its maximum over $s \leq t$ is attained at a breakpoint or at $s = t$; symmetrically for t . This gives (5), the maximum drawdown of X , computed with one running maximum. For (6), a start s inside a cell $[j\delta, (j+1)\delta]$ satisfies $X(s) - \min_{t \geq s} X(t) \leq \max\{d_j, d_{j+1}\}$ with $d_j = X_j - \min_{k \geq j} X_k$, and both cell ends lie in c because classes are unions of cells; the suffix minima take one backward pass. The last claim follows from $\rho\Delta + \sum e_i \leq \alpha_{\text{HI}}(\Delta) \leq \rho\Delta + \sigma$. $\square$

Equation (5) is the sequent-peak algorithm of reservoir design and Lindley's recursion $W_k = \max(0, W_{k-1} + \rho\delta - h_k)$ in disguise [20, 37]: $W_k = \max_{j \leq k} (X_j - X_k)$ is the energy the store must have supplied since the last time the harvest caught up with the HI demand. In network-calculus terms [9, 19],

R_* is the vertical deviation between the HI arrival curve α_{HI} and the harvest service curve β . The contribution here is not the formula but its use: Proposition 2 makes it cheap to evaluate reserves over decades of minute-resolution data, for every season, every year and every cross-validation fold, which is what the statistical treatment in Section 5 requires.

275 5 From climate records to reserves

276 Theorem 1 turns the question “will HI deadlines hold?” into “does the reserve cover the worst harvest 277 deficit the node will meet?”. This section explains how we obtain harvest traces from public climate 278 records, how we compute static and seasonal reserves, and how we treat the reserve statistically, since the 279 worst deficit of the *future* is unknown.

280 5.1 Harvest traces

281 Data. We use daily all-sky global horizontal irradiation (GHI), clear-sky GHI, 2-m air temperature 282 and bias-corrected precipitation from the NASA POWER daily point API [30]. From 2001 onwards 283 the irradiance comes from CERES SYN1deg on a 1° grid [10, 26]; temperature and bias-corrected 284 precipitation are MERRA-2 products, as stated in the headers of the downloaded files. Days are in local 285 solar time. We use complete calendar years 2001–2025 (25 years). Missing irradiance values, if any, are 286 replaced by the *minimum* recorded for the same day of the year, so that gap filling can never make a site 287 look sunnier than its record; none of the six series had gaps.

288 Intra-day shape. Each day’s irradiance is spread over the day at one-minute resolution following
289 a clear-sky diurnal shape, $G(t) \propto \cos \vartheta_z e^{-0.057/\cos \vartheta_z}$ [14], with the solar zenith angle ϑ_z from the site 290 latitude and the declination series of Spencer [31]. This preserves every daily total and the exact length of 291 each night, but not intra-day cloud variability. The reserves in this study range from 921J to 3347J, while 292 the most a minute of sunshine delivers to the store is 17J at any site, so the reserves are governed by night 293 length and by multi-day totals, both of which the disaggregation preserves. We return to it in Section 8.3.

294 Panel and charger. The DC power of a panel of rating W_p is

$$p_{pv}(t) = W_p \frac{G(t)}{1000} \frac{\eta_{dc} (1 + \gamma (T_{cell} - 25)) (1 - s(d))}{T_{cell} = T_{air}}, \quad + \frac{NOCT - 20}{800} G(t)$$

295 with DC-side losses $\eta_{dc} = 0.85$, temperature coefficient $\gamma = -0.4\%/K$, NOCT = 45 °C and the day’s 296 mean air temperature T_{air} . The power reaching the store is $p_h = \eta p_{pv}$ with charging and round-trip 297 efficiency $\eta = 0.85$. $s(d)$ is an optional soiling loss that grows by a fixed rate on dry days, is halved by a 298 day with at least 1mm of rain and is reset by 5mm or more, capped at 30%. The soiling rate is uncertain 299 and site-specific [6], so we treat it as a sensitivity parameter (Section 8.1) and set $s = 0$ elsewhere.

300 5.2 Static and seasonal reserves

301 Given a trace h_k (joules delivered in minute k), the static reserve is $R^{\text{aff}} = \sigma + \max_{j \leq k} (X_j - X_k)$ (Proposi-
302 tion 2). For the seasonal policy Res-S we use calendar months as classes and compute, for every minute
303 j , the per-start deficit $d_j = X_j - \min_{k \geq j} X_k$; the monthly reserve is $R_m = \sigma + \max d_j$ over the starts whose 304 day of the year lies in month m extended by seven days on either side. The extension implements the
305 closures of Definition 1 (and slightly more), so that a deficit that starts on the last day of a month is also 306 covered by the next month’s reserve. Deficits are allowed to run past the end of the month or year in
307 which they start.

A calendar-based class is a simple, forecast-free way to exploit seasonality. Any partition that the node can evaluate at run time (for example a class switched by a rain gauge or by a coarse weather forecast) fits Theorem 1 equally well, provided the curves are computed for the same partition.

5.3 Statistical calibration

A reserve computed from a record is guaranteed for that record only. For deployment, what matters is how likely a future year is to contain a deficit deeper than the reserve. Let R_y be the reserve needed by the windows that *start* in calendar year y . It is the analogue of an annual maximum in hydrology. We consider three ways of setting the reserve for a year y that was not used for calibration:

- *Empirical, leave-one-year-out (LOYO)*. R is the static (or monthly) reserve of the other 24 years. Windows are not allowed to cross the held-out year. This reproduces a designer who has a long record, and asks how often the maximum of 24 years is exceeded in a new year.
- *Short records*. R is computed from the $k \in \{1,2,3,5,10\}$ years immediately preceding y (rolling origin, strictly past data). This reproduces a designer who calibrates on a pilot deployment or on the most recent years of data.

• *Return levels*. A generalized extreme-value (GEV) distribution is fitted by maximum likelihood to the annual requirements $\{R_y\}_{y=y}$ [8], and R is set to the T -year return level, the $(1-1/T)$ -quantile of the fitted distribution, for $T \in \{5,10,25,50,100\}$. If the fitted shape parameter is implausible (magnitude above 1), the Gumbel special case is used. Under stationarity, a year then exceeds its reserve with probability $1/T$.

Exceedance of the reserve is a necessary but not a sufficient condition for a HI miss, because the store is often above the reserve when a deficit begins. We therefore report both the calibration of the reserve (how often $R_y > R$) and the actual HI misses in simulation.

Table 2: Sites and their 2001–2025 climate from NASA POWER. GHI: mean daily global horizontal irradiation ($\text{kWhm}^{-2} \text{d}^{-1}$); CV: coefficient of variation of daily GHI; worst month: lowest monthly mean GHI (month in brackets); worst 5-d: lowest 5-day mean GHI as % of the site mean; run: longest spell of days below 50% of the day-of-year mean; rain: mean annual precipitation (mm); W_p : panel rating at $S = 1.3$.

Site	Climate	Lat	Lon	GHI	CV	Worst month	Worst 5-d	Run	Rain	W_p (W)
Kano	Sudano-Sahelian savanna	12.00	8.52	5.99	0.15	5.51 (D)	53%	2 d	676	0.35
Lagos	Humid tropical coast	6.52	3.38	4.60	0.25	3.81 (J)	42%	2 d	1559	0.45
Douala	Equatorial monsoon	4.05	9.70	4.25	0.23	3.24 (A)	45%	2 d	2405	0.49
Nairobi	Equatorial highland	-1.29	36.82	5.80	0.21	4.75 (J)	45%	2 d	581	0.36
Cape Town	Mediterranean	-33.92	18.42	5.56	0.42	2.61 (J)	20%	3 d	522	0.37
Aswan	Hot desert	24.09	32.90	6.52	0.21	4.45 (D)	52%	2 d	2	0.33

Table 3: Reference workload (assumed, representative values). $P_b = 4.0\text{mW}$, $P_l = 0.5\text{mW}$. Average power is e_i/T_i .

Task	χ	C_i (s)	T_i (s)	P_i (mW)	e_i (J)	avg. (mW)
Fault check	HI	0.015	2	35	0.0005	0.23
LoRa status	HI	1.5	300	160	0.234	0.78
Interval telemetry (cellular)	HI	12	900	900	10.75	11.95
Power-quality analytics	LO	0.6	10	40	0.022	2.16
Bulk waveform upload (cellular)	LO	30	1800	1500	44.88	24.93
Load-forecast update	LO	20	3600	45	0.820	0.23

330 6 Experimental setup

331 6.1 Sites

332 Table 2 lists the six sites. They were chosen before looking at the data to cover the main climate zones
333 of the continent with one large city each: a Sudano-Sahelian savanna site with a long dusty dry season
334 (Kano), the humid Guinea coast (Lagos), an equatorial monsoon site, the wettest of the six (Douala),
335 an equatorial highland with two rainy seasons (Nairobi), a Mediterranean site with winter rains and
the 336 shortest winter days of the set (Cape Town), and a hot desert as a best case (Aswan).

337 6.2 Node and workload

338 Table 3 lists the reference workload, a solar-powered monitoring node on a rural mini-grid feeder. HI 339
tasks are a fault check every 2s, a LoRa status message to a local operator every 5min, and interval
340 telemetry over a cellular link every 15min (the usual metering interval). LO tasks are power-quality 341
analytics with an on-device classifier every 10s, a bulk waveform upload every 30min, and an hourly
342 load-forecast update. Power figures are representative of a low-power microcontroller with LoRa
and an 343 LTE-M/Cat-1 modem. They are *assumed* order-of-magnitude values, not measurements,
and every result 344 scales with them. The HI tasks, including the baseline, average 17.5mW and the
LO tasks 27.3mW, so 345 HI is 39% of the full demand of 1.07Wh per day. $U = 0.108$.

346 Sizing. Each site gets its own panel, sized so that the mean daily energy delivered to the store over 347
2001–2025 equals $S = 1.3$ times the full daily demand. This is how installers commonly size small 348 systems
(from the annual mean with a margin), and it removes the trivial effect of mean irradiance so 349 that
differences between sites reflect *variability*. The resulting panels range from 0.33W to 0.49W.

350 Section 8.1 varies S from 1.0 to 2.0. Storage capacities range from 0.25 to 8Wh (900J to 28.8kJ),
from a 351 small supercapacitor bank to a single small lithium cell. The node starts with a full store
and restarts after 352 a brownout once the store holds 2% of its capacity plus a 1J boot energy.

353 6.3 Policies

354 All policies release every HI job and differ only in LO admission.

- 355 • GREEDY: admit every LO job (energy-agnostic EDF).
- 356 • THRESH- x : admit a LO job if $B - \pi - e_j \geq xB_{\max}$, for $x \in \{25\%, 50\%\}$. This is the “low-battery 357 mode”
common in practice.
- 358 • ENO-EWMA: an energy-neutral controller in the style of [18]. An EWMA ($\alpha = 0.5$) predicts 359 the harvest
of each half-hour slot of the day. At midnight the fraction of LO jobs to admit is set to
360 $\min\{1, \max\{0, (\hat{E}_{\text{day}} - E_{\text{HI}} - E_b + (B - 0.5B_{\max})/3)/E_{\text{LO}}\}\}$, balancing predicted harvest against 361 HI, baseline and
LO demand with a three-day battery feedback, and LO jobs are admitted at that
362 rate.
- 363 • RES and RES-S: Definition 2 with a static or a monthly reserve. Unless stated otherwise the
reserve 364 for year y is calibrated *out of sample* by LOYO (Section 5.3). In-sample versions,
calibrated on all 365 25 years, are reported to check Theorem 1.
- 366 • ENO+RES-S: ENO-EWMA admission, but a LO job is admitted only if the RES-S test also passes 367
(Remark 1).

6.4 Metrics

The *HI miss ratio* is the fraction of HI jobs that miss their deadline, and *HI-miss days* counts the days with at least one such miss. *LO QoS* is the fraction of released LO jobs that complete. We also report brownouts, hours spent browned out, and the fraction of harvest lost to a full store.

6.5 Simulators

Twenty-five years at one-minute resolution are 13.1 million steps per site and configuration, so the experiments use a minute-step simulator compiled with Numba. At the start of each minute it releases the minute's jobs (all periods divide or are multiples of 60s), applies the admission policy with the current B and π , and balances harvest against baseline, HI and admitted LO energy. If the store cannot cover the minute's demand, the node browns out: the fraction of the minute's jobs that the available energy can pay for completes, the rest miss, and the node stays off (drawing only self-discharge) until it reaches the

restart threshold. Because $U \approx 0.108$ and every WCET is well below a minute, EDF completes every admitted job within the minute of its release whenever energy is available, so energy starvation is the only source of misses.

We validated this simulator against an exact job-level simulator that we wrote independently: it runs preemptive EDF with synchronous periodic releases, integrates the storage level in continuous time between events (releases, completions, minute boundaries, the store becoming empty or full), and applies the same admission rules with the exact $B(t)$ and $\pi(t)$ at each release. Table 5 (Section 7.3) compares the two on 21-day windows around the deepest deficits of Cape Town and Kano for four policies and three storage sizes; the minute-step simulator reproduces HI miss ratios within 0.94 percentage points and LO QoS within 0.96 points.

7 Results

7.1 How variable is the harvest?

Table 2 and Figure 2 summarize the six records. Mean daily GHI ranges from 4.25 kWh m^{-2} in Douala to 6.52 kWh m^{-2} in Aswan, so the panels sized for $S = 1.3$ range from 0.33W to 0.49W. Variability differs more than the means do. The coefficient of variation of daily GHI is 0.15 in Kano and 0.42 in Cape Town.




Figure 2: Monthly climatology of daily GHI, 2001–2025: mean, 10th–90th percentile band of individual days, and the minimum day of each month.

Table 4: Reserves required by the HI workload (Wh). R^* : in-sample static reserve over 2001–2025 (Proposition 2); critical window: length and start of the deficit that attains it; R_m : range of monthly reserves; R_y : range of annual requirements; R_T : GEV return levels fitted to all 25 annual requirements.

Site	R^*	Critical window		R_m min–max	R_y min–max	R_{10}	R_{50}	R_{100}
			start					
Kano	0.481	1.5	2020-08-05	0.223–0.481	0.239–0.481	0.298	0.335	0.351
Lagos	0.479	1.5	2025-06-13	0.241–0.479	0.238–0.479	0.445	0.543	0.585
Douala	0.346	1.5	2010-02-12	0.234–0.346	0.232–0.346	0.320	0.367	0.387
Nairobi	0.256	1.6	2025-08-17	0.231–0.256	0.232–0.256	0.241	0.249	0.254
Cape Town	0.930	3.7	2008-06-17	0.200–0.930	0.342–0.930	0.648	0.893	1.024
Aswan	0.256	0.6	2021-12-27	0.210–0.256	0.253–0.256	0.255	0.257	0.258

The darkest five consecutive days in 25 years average 20% of the site mean in Cape Town, 42–45% in Lagos, Douala and Nairobi, and 52–53% in Aswan and Kano. On the Guinea coast the lowest monthly means occur in the monsoon (July in Lagos, August in Douala). In Kano the lowest monthly mean is in December, when the sun is lowest and the dusty harmattan season is under way [6], but the darkest individual days fall in the July–August rains (dashed minima in Figure 2). The month with the lowest average is not necessarily the month that sets the reserve. At every site, runs of days below half of the day-of-year mean are short: two or three days at most (Table 2).

7.2 How large must the reserve be?

Table 4 gives the in-sample reserves for the reference node. For the same node, sized to the same multiple of each site’s mean harvest, the reserve ranges from 0.26Wh (Aswan) to 0.93Wh (Cape Town), a factor of 3.6. The deficit that sets it (Figure 3) is short. It is one long night plus a dim day in Aswan (0.6days, December), about a day and a half of rain or cloud in Kano, Lagos, Douala and Nairobi, and a 3.7-day run of winter storms in Cape Town (June 2008). With HI demand at 30% of the mean harvest, the reserve bridges nights and short dim spells, not seasons.

The year-to-year picture (Figure 4) separates the sites into two groups. In Aswan all 25 annual requirements R_y lie within 5% of the smallest; in Nairobi 23 and in Kano 18 do. At these sites the




Figure 3: HI demand minus the lower harvest curve, $\alpha_{HI}(\Delta) - \beta(\Delta)$, over 2001–2025. The maximum (dot) is the reserve R^* of Corollary 1; the saw-tooth is the day–night cycle.

requirement has a floor set by the longest night and rises above it only in occasional years. In Kano the floor years have their worst window in the long nights of December and January, while the seven years above it (2011, 2014, 2019–2022 and 2024) all have deficits starting between late July and mid-September, in the rains; August 2020 doubled the requirement. In Lagos (2 of 25 years near the floor), Douala (9) and Cape Town (3), most years contain a deficit deeper than a night. Not all of them fall in the monsoon:

in Lagos the worst deficit of the year started in January–March in 6 of the 25 years, and in Douala in February–April in 7. Cape Town’s June 2008 event stands apart: it needed 1.44 times the second-largest annual requirement in the record (0.64Wh).

Seasonality is equally site-specific (Figure 5). The monthly reserve in Cape Town ranges from 0.20Wh in summer to 0.93Wh in June. In Kano it is elevated only in July and August (and slightly in September), and in Lagos it is highest in April, June and July. In Nairobi and Aswan it barely moves, following only the length of the night. As Corollary 1 predicts, the largest monthly reserve equals the static one at every site.

When the HI load grows. The short critical windows are a property of the reference workload. Scaling the HI job energies (Section 8.1) shows how the nature of the reserve changes. At HI demand equal to 52% of the mean harvest, the reserve rises to 0.71–1.23Wh at five sites, but to 7.26Wh in Cape Town, where the critical window becomes a 68.6-day winter season. At 75%, it ranges from 2.00Wh in Kano (critical window 6.6days) to 38.33Wh in Cape Town (139.6days). Kano, with the least variable sunshine, is the last to need seasonal storage.

7.3 Guarantee and simulator checks

Across all sites and storage sizes there are 58 configurations in which the in-sample reserve fits in the store ($R \leq B_{\max}$). In all of them RES and RES-S missed no HI deadline in 25 simulated years, as Theorem 1 requires. When the reserve exceeds the capacity (every site at 0.25Wh, and Cape Town at 0.5Wh), the theorem does not apply and every policy misses HI jobs. Planning the storage from R^* is therefore the

first design step.

Table 5 compares the minute-step and exact simulators on 21-day windows around the deepest deficits

of Cape Town and Kano. The largest difference in HI miss ratio is 0.94 percentage points (ENO-EWMA,

Cape Town, 1Wh), the largest difference in LO QoS is 0.96 points, and brownout counts differ by at most four. For RES the HI miss ratios agree to within 0.001 percentage points.




Figure 4: Annual requirement R_y (reserve needed by deficits starting in year y), with the record maximum R^* and the 10- and 50-year GEV return levels fitted to all 25 years.

7.4 Policy comparison

Figures 6 and 7 show HI-miss days and LO QoS over the 25 years for six storage sizes, and Tables 6 and 7 give the values for 1 and 2Wh. RES and RES-S are calibrated out of sample (LOYO) throughout.

Energy-agnostic and energy-neutral control miss HI deadlines at night. With a 0.5Wh store, GREEDY missed HI jobs on 9130 of 9131 days at every site: LO work drains the store in the evening and the node browns out before dawn. ENO-EWMA balances energy over a day but not over the day-night cycle, and was almost as bad at 0.5Wh (8265–9117 days). At 1Wh it still missed HI jobs on 34 (Aswan) to 429 (Lagos) days. GREEDY still missed HI deadlines at all six sites with a 2Wh store, from 3 days in Kano to 2620 in Cape Town.

A fixed threshold is right at some sites and wrong at others. THRESH-25% avoided HI misses at 1Wh in Aswan, Douala and Nairobi, whose requirements are close to its 0.25Wh threshold, but missed on 252 days in Cape Town, whose requirement is 0.93Wh. THRESH-50% was safe at five sites from 1Wh and in Cape Town from 2Wh, but kept half the store idle. At 1Wh its LO QoS was 83.9–87.3% at the five sites, against 93.9–98.9% for RES-S.

RES and RES-S. With out-of-sample reserves and a 1Wh store, RES and RES-S missed no HI deadline in 25 years at five sites. In Cape Town RES missed on 3 days and RES-S on 4. Three of these days are in the June 2008 event, which lies outside the range of the other 24 years (Section 7.6); the fourth, for RES-S, is in 2015. The price of the static reserve is paid in LO QoS when the store is small: at 1Wh in Cape Town RES delivered only 38.6% of LO jobs, because it holds back the winter reserve all year. The seasonal reserve recovers most of this: RES-S raised LO QoS at 1Wh from 38.6% to 81.2% in Cape Town, from 89.3% to 97.9% in Kano and from 86.3% to 93.9% in Lagos. From 2Wh upwards the store is large enough that the difference disappears.

The reserve as a safety floor. ENO+RES-S admits LO jobs at the ENO-EWMA rate, but only when the RES-S test also passes. Every job it admits passes the RES-S tests, so Theorem 1 covers it (Remark 1). It




Figure 5: Monthly reserves R_m (in-sample) against the static reserve R^* .

lost less than one point of LO QoS relative to RES-S at five sites at 1Wh (97.7% vs 97.9% in Kano) and 79.0% vs 81.2% in Cape Town. It had the fewest HI misses in Cape Town: 1 day at 1Wh and none at 2Wh.

7.5 How much data does a reserve need?

Table 8 and Figure 8 calibrate the static reserve on the k years before each test year. Pooled over the six sites and the 90 test site-years, a reserve computed from the previous year alone was exceeded in 43 test years (48%). With two years of data the figure was 37%, with three 28%, with five 19%, and with ten 13%. With the other 24 years (LOYO) it was 4%, close to the $1/25$ chance that a given year holds the record maximum.

Exceeding the reserve is necessary for a HI miss but not sufficient. When a deficit is short, the store is rarely at the reserve at its start, because it has been refilled during the preceding day. HI misses therefore occurred only in Cape Town, where long winter deficits first drain the store to the reserve through LO admissions and then exhaust it. There, one-year calibration missed HI jobs on 12 days in 4 of 15 test years, and calibration on ten or more years never did. The same pattern holds at 2Wh and 8Wh.

Return levels. Table 9 evaluates reserves set to GEV return levels fitted out of sample. At $T = 5$ the fit is well calibrated: 29 exceedances in 150 site-years against 30.0 expected. Further into the tail it becomes optimistic: 19 exceedances against 15.0 at $T = 10$, and 5 against 1.5 at $T = 100$. Two features of the data explain this. At night-dominated sites most annual requirements sit on a hard floor with rare jumps, a shape that a three-parameter GEV fitted to 24 values does not capture. And the Cape Town 2008 event is far outside the rest of its record. With $T \geq 10$, HI misses at 1Wh were confined to a single site-year, Cape Town 2008.

7.6 Anatomy of the worst deficit

Figure 9 follows the deepest deficit in the record, Cape Town in June 2008, with a 2Wh store. On the darkest day the store received 0.11Wh, a quarter of the HI demand alone. GREEDY missed HI jobs on all 30 days of the window, because LO work empties the store every evening. The out-of-sample June reserve of RES-S, computed from the other 24 years, was 0.64Wh, and it was not enough: RES-S

missed HI jobs on 3 days. ENO-EWMA, whose forecast reacts to a run of dim days by cutting LO work, Table 5: Minute-step simulator vs. exact job-level EDF simulator on 21-day windows around the deepest deficit. HI miss ratio and LO QoS in %; brownouts as counts.

Site	Policy	$B_{\max}$ (Wh)	HI miss (%)		LO QoS (%)		Brownouts	
			minute	exact	minute	exact	minute	exact
Cape Town	Greedy	0.5	47.83	47.70	52.2	52.3	66	66
Cape Town	Thresh-25%	0.5	21.23	21.22	37.8	37.8	16	16
Cape Town	ENO-EWMA	0.5	14.88	14.73	23.3	23.5	12	12
Cape Town	Res (in-sample)	0.5	5.00	5.00	0.0	0.0	3	3
Cape Town	Greedy	1	44.54	44.51	55.5	55.5	26	26
Cape Town	Thresh-25%	1	8.67	8.67	34.8	34.8	7	7
Cape Town	ENO-EWMA	1	4.24	5.18	27.4	27.7	3	5
Cape Town	Res (in-sample)	1	0.00	0.00	12.4	12.4	0	0
Cape Town	Greedy	2	39.82	39.76	60.2	60.2	20	21
Cape Town	Thresh-25%	2	4.89	4.88	36.9	36.9	3	3
Cape Town	ENO-EWMA	2	0.00	0.00	30.8	29.9	0	0
Cape Town	Res (in-sample)	2	0.00	0.00	30.4	30.4	0	0
Kano	Greedy	0.5	18.14	18.20	81.9	81.8	39	41
Kano	Thresh-25%	0.5	4.13	4.13	66.0	66.0	4	4
Kano	ENO-EWMA	0.5	15.06	14.95	73.2	73.2	35	39
Kano	Res (in-sample)	0.5	0.00	0.00	30.3	30.3	0	0
Kano	Greedy	1	6.50	6.50	93.5	93.5	5	5
Kano	Thresh-25%	1	0.00	0.00	88.0	88.0	0	0
Kano	ENO-EWMA	1	3.26	3.19	84.2	84.3	2	2
Kano	Res (in-sample)	1	0.00	0.00	78.0	77.9	0	0
Kano	Greedy	2	0.74	0.74	99.3	99.3	1	1
Kano	Thresh-25%	2	0.00	0.00	95.7	95.7	0	0
Kano	ENO-EWMA	2	0.00	0.00	89.8	89.4	0	0
Kano	Res (in-sample)	2	0.00	0.00	96.1	95.8	0	0

happened to keep more energy in the store and missed on 1 day. ENO+RES-S missed on none, delivering 20% of LO jobs in the window against 24% for RES-S.

8 Discussion

8.1 Sensitivity

HI load. Section 7.2 showed that the reserve changes character as HI demand grows relative to the harvest. It bridges nights and 1–4-day spells at 30% of the mean harvest, covers multi-day spells of 2.5–4.6 days at five sites (but a 69-day winter in Cape Town) at 52%, and becomes seasonal storage at 75%. When a node's critical function is a large share of its energy budget, the design problem shifts from real-time admission to seasonal energy planning of the kind addressed by long-term energy-neutral methods [4]. RES still provides the guarantee, but the store it requires is then measured in watt-hours to tens of watt-hours.

Panel sizing. Sizing the panel to exactly the mean full demand ($S = 1.0$) raises the in-sample reserve

to between 0.31Wh (Aswan) and 1.37Wh (Cape Town). GREEDY then misses HI jobs on 1211–3951 days even with a 2Wh store. RES with out-of-sample reserves missed none at five sites and 5 days in Cape Town. Doubling the panel ($S = 2.0$) lowers the reserve to 0.23–0.64Wh; Cape Town still needs the largest reserve by far.

Soiling. With the panel sized on the clean record, a soiling rate of 0.2–0.5% per dry day removes 29.2–29.7% of Aswan's harvest (almost no rain ever cleans the panel, so the loss saturates at the 30%




Figure 6: Days with at least one HI deadline miss over 2001–2025 (symmetric-log scale) against storage capacity.

Table 6: Policies over 2001–2025 with $B_{\max} = 1\text{Wh}$. Each cell: HI-miss days / LO QoS (%). RES and RES-S are calibrated out of sample (LOYO).

Site	Greedy	Thresh-25%	Thresh-50%	ENO-EWMA	Res	Res-S	ENO+Res-S
Kano	107 / 99.8	1 / 99.3	0 / 86.3	68 / 99.4	0 / 89.3	0 / 97.9	0 / 97.7
Lagos	672 / 98.7	2 / 96.5	0 / 83.9	429 / 96.8	0 / 86.3	0 / 93.9	0 / 93.1
Douala	567 / 99.0	0 / 97.4	0 / 85.2	218 / 97.5	0 / 96.6	0 / 96.9	0 / 96.2
Nairobi	396 / 99.3	0 / 98.3	0 / 86.4	102 / 98.2	0 / 98.3	0 / 98.3	0 / 97.8
Cape Town	2894 / 91.2	252 / 85.7	6 / 79.2	417 / 85.0	3 / 38.6	4 / 81.2	1 / 79.0
Aswan	712 / 99.3	0 / 98.9	0 / 87.3	34 / 98.7	0 / 98.9	0 / 98.9	0 / 98.7

cap) and 9.6–13.8% of Kano's, against at most 5.7% at the other sites. GREEDY suffers accordingly: at 2Wh its HI-miss days rise from 321 to 4379 in Aswan and from 3 to 2002 in Kano. The effect on the reserve depends on when the binding deficit occurs. Kano's reserve is unchanged (0.48Wh), because its binding deficits fall in the rainy season, when rain has just cleaned the panel. Aswan's rises from 0.26Wh to 0.34Wh, because its binding deficit is a winter night and the panel is then at its dirtiest. If the reserve accounts for soiling, RES misses no HI deadline at any soiling rate outside Cape Town. If it is calibrated on the clean record and the panel is never cleaned, it misses HI jobs on 3 days in

Aswan. ⁵¹⁵ Soiling therefore belongs in the harvest model used to compute reserves at arid sites, or panels must be ⁵¹⁶ cleaned on a schedule.

8.2 Implications for deployments

Plan storage from the site's own long record. The reserve differs by a factor of 3.6 across the six sites ⁵¹⁹ for the same node sized to the same multiple of its site's mean harvest. The mean irradiance is a poor ⁵²⁰ guide: Cape Town's mean is higher than those of Lagos and Douala but its reserve is by far the largest, ⁵²¹ and Kano needs nearly twice Aswan's reserve because of a few days of monsoon cloud. Twenty-five ⁵²² years of NASA POWER data are free for any location on the continent, and Proposition 2 evaluates them ⁵²³ in seconds.




Figure 7: LO quality of service over 2001–2025 against storage capacity.

Table 7: Policies over 2001–2025 with $B_{\max} = 2\text{Wh}$. Each cell: HI-miss days / LO QoS (%). RES and RES-S are calibrated out of sample (LOYO).

Site	Greedy	Thresh-25%	Thresh-50%	ENO-EWMA	Res	Res-S	ENO+Res-S
Kano	3 / 100.0	0 / 100.0	0 / 99.8	0 / 99.8	0 / 100.0	0 / 100.0	0 / 99.8
Lagos	69 / 99.8	0 / 99.5	0 / 98.3	0 / 98.5	0 / 99.5	0 / 99.5	0 / 98.5
Douala	192 / 99.6	0 / 99.2	0 / 98.5	0 / 98.6	0 / 99.3	0 / 99.3	0 / 98.6
Nairobi	100 / 99.8	0 / 99.5	0 / 99.0	0 / 99.0	0 / 99.6	0 / 99.6	0 / 99.0
Cape Town	2620 / 91.6	6 / 86.6	0 / 86.2	1 / 85.8	3 / 86.3	4 / 86.7	0 / 85.8
Aswan	321 / 99.6	0 / 99.3	0 / 99.1	0 / 99.1	0 / 99.4	0 / 99.4	0 / 99.1

⁵²⁴ Do not calibrate on a pilot year. A reserve estimated from one year of data was exceeded in about ⁵²⁵ half of the following years, and from three years in more than a quarter. Whether an exceedance ⁵²⁶ causes misses depends on the length of the deficit, and hence on the site. Nodes at sites with ⁵²⁷ long deficits, such as Cape Town's winter, suffered from short-record calibration, while sites with ⁵²⁸ short deficits were forgiving in our experiments. A designer cannot know which case applies without the long record.

529 Leave a margin beyond the record. Even 24 years were not enough for Cape Town in 2008, and 530 extreme-value extrapolation to $T = 50$ or $T = 100$ years was optimistic in the tail (Table 9). The
531 requirement's floor, set by night length, and its occasional large jumps are poorly described by a GEV 532
fitted to two dozen values. A peaks-over-threshold model of the excess above the night floor [8], or 533
simply an engineering margin on the record maximum, is more defensible. The margin is cheap when the
534 storage can hold it and the seasonal policy keeps LO QoS high.

535 Use seasonal reserves and combine with forecasting. With small stores, monthly reserves doubled 536 LO
QoS in Cape Town at no cost in HI guarantees, and they need nothing at run time but a calendar. 537 Forecast-
based control and the reserve are complementary. The forecast reacts to a developing dim spell
538 and preserves energy early, while the reserve provides a floor that holds even when the forecast is wrong.
539 With a 1Wh store their combination had the fewest HI misses of all policies at the most demanding site.

Table 8: Static reserves calibrated on the k years before each test year (rolling origin), test years 2011–2025, $B_{\max} = 1\text{Wh}$. Each cell: test years whose requirement R_y exceeded the calibrated reserve / test years with at least one HI miss (out of 15 per site).

Site	1 y	2 y	3 y	5 y	10 y	24 y (LOYO)
Kano	8 / 0	5 / 0	4 / 0	3 / 0	3 / 0	1 / 0
Lagos	8 / 0	5 / 0	4 / 0	2 / 0	2 / 0	1 / 0
Douala	5 / 0	4 / 0	4 / 0	1 / 0	0 / 0	0 / 0
Nairobi	8 / 0	7 / 0	6 / 0	6 / 0	4 / 0	1 / 0
Cape Town	7 / 4	6 / 2	3 / 1	2 / 1	0 / 0	0 / 0
Aswan	7 / 0	6 / 0	4 / 0	3 / 0	3 / 0	1 / 0
All sites	43 / 4	33 / 2	25 / 1	17 / 1	12 / 0	4 / 0




Figure 8: Short calibration records, test years 2011–2025, $B_{\max} = 1\text{Wh}$. Left: fraction of test years whose requirement exceeded the calibrated reserve. Right: HI-miss days (symmetric-log scale); only Cape Town has any.

540 8.3 Limitations and threats to validity

541 Harvest model. Our traces spread daily satellite totals over a clear-sky diurnal shape. This keeps
daily
542 totals and night lengths exact, but it omits intra-day cloud transients, which matter for small stores.
The 543 satellite product represents a 1° cell, which smooths local cloud variability, so point deficits may
be

deeper than the cell averages we used. Both effects would make real reserves larger than ours, not smaller. Validating the reserves against ground radiometer data and hourly records is the most important next step.

Node model. Task powers are assumed values, not measurements from a specific board. The storage model has a constant charging efficiency and self-discharge, and ignores capacity fade, temperature effects on usable capacity, and the error of a fuel gauge. The last is easily absorbed by adding the gauge’s worst-case error to R . Results scale with the workload’s energy, so a measured workload can be substituted directly.

Scheduling model. We assume a single processor, implicit deadlines and low utilization, so that timing fails only through energy starvation. The energy argument of Theorem 1 does not depend on the timing analysis beyond the requirement that admitted jobs meet their deadlines while the processor is powered. Extending it to constrained deadlines or partitioned multicores therefore mainly requires replacing Assumption 1 by the corresponding schedulability test, and Lemma 1 by the matching demand bound.

Table 9: Return-level reserves fitted out of sample (LOYO GEV), pooled over all sites and years; HI misses simulated with $B_{\max} = 1\text{Wh}$. Expected exceedances = N/T under stationarity.

T (years)	site-years N	expected	observed exceedances	years with HI miss	mean R (Wh)
5	150	30.0	29	4	0.338
10	150	15.0	19	1	0.366
25	150	6.0	9	1	0.404
50	150	3.0	5	1	0.436
100	150	1.5	5	1	0.471

Stationarity. All statistics assume that the next 25 years resemble the last 25. Trends in cloudiness and aerosol, and changes in the satellite products, could violate this assumption; NASA POWER cautions against trend analysis across changes of source data [26]. Periodic recalibration with a rolling long window is a practical safeguard.

9 Conclusion

Solar-powered edge nodes in Africa face harvest deficits whose depth, duration and season depend strongly on the climate zone. We showed that a simple reserve-based admission rule gives HI tasks a hard deadline guarantee on such nodes, provided the reserve covers the worst deficit. We also showed that this reserve can be computed exactly and cheaply from decades of freely available satellite data. The experiments turned the guarantee into design guidance. Reserves must be computed per site rather than inferred from mean irradiance, from long records rather than pilot deployments, and with a margin beyond the record maximum. Seasonal reserves recover most of the optional work that a static reserve gives up on small stores, and the reserve works well as a safety floor under a forecast-based energy manager.

Three extensions would strengthen the results. The first is a hardware prototype, deployed at one of the Nigerian sites, to replace the assumed task energies with measurements and to test the policy with a real fuel gauge. The second is validation of the reserves with hourly and ground-station irradiance, which include the intra-day and local variability that our daily satellite data smooth out. The third is a peaks-over-threshold treatment of the excess above the night floor, which should give better-calibrated return levels than the block-maximum fits evaluated here.

576 Statements and declarations

577 Data availability. All irradiance, temperature and precipitation series are public NASA POWER data
578 (daily point API, community “RE”, 2001–2025) for the coordinates in Table 2.

579 Code availability. The simulator, analysis scripts and the scripts that produce every table and figure
will 580 be released in a public repository on acceptance (link withheld for anonymous review).

581 Competing interests. [To be completed by the authors.]

582 Funding. [To be completed by the authors (withheld for review).]

583 References

584 [1] Yasmina Abdeddaïm, Younès Chandarli, and Damien Masson. The optimality of PFPasap algorithm 585
for fixed-priority energy-harvesting real-time systems. In *Euromicro Conference on Real-Time 586 Systems*
(ECRTS), 2013.

587 [2] Yasmina Abdeddaïm, Younès Chandarli, Robert I. Davis, and Damien Masson. Response time 588
analysis for fixed priority real-time systems with energy-harvesting. *Real-Time Systems*, 52(2): 589 125–160,
2016. doi: 10.1007/s11241-015-9239-7.




Figure 9: Cape Town, 30 days from 2008-06-11, $B_{\max} = 2\text{Wh}$. Top: daily harvest delivered to the store, against HI and full daily demand. Middle: stored energy (hourly) and the out-of-sample monthly reserve of RES-S. Bottom: share of HI jobs missed per day.

590 [3] Sedigheh Asyaban, Mehdi Kargahi, Lothar Thiele, and Morteza Mohaqeqi. Analysis and scheduling 591
of a battery-less mixed-criticality system with energy uncertainty. *ACM Transactions on Embedded 592*
Computing Systems, 16(1):1–26, 2016. doi: 10.1145/2964201.

- [4] Bernhard Buchli, Felix Sutton, Jan Beutel, and Lothar Thiele. Dynamic power management for long-term energy neutral operation of solar energy harvesting systems. In *ACM Conference on Embedded Network Sensor Systems (SenSys)*, 2014.
- [5] Alan Burns and Robert I. Davis. A survey of research into mixed criticality systems. *ACM Computing Surveys*, 50(6), 2017. doi: 10.1145/3131347.
- [6] Yusuf N. Chanchangi, Aritra Ghosh, Senthilarasu Sundaram, and Tapas K. Mallick. Dust and PV performance in Nigeria: A review. *Renewable and Sustainable Energy Reviews*, 121:109704, 2020. doi: 10.1016/j.rser.2020.109704.
- [7] Maryline Chetto. Optimal scheduling for real-time jobs in energy harvesting computing systems. *IEEE Transactions on Emerging Topics in Computing*, 2(2):122–133, 2014.
- [8] Stuart Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer, 2001.
- [9] Rene L. Cruz. A calculus for network delay, Part I: Network elements in isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, 1991.
- [10] David R. Doelling, Norman G. Loeb, Dennis F. Keyes, Michele L. Nordeen, Daniel Morstad, Cathy Nguyen, Bruce A. Wielicki, David F. Young, and Moguo Sun. Geostationary enhanced temporal interpolation for CERES flux products. *Journal of Atmospheric and Oceanic Technology*, 30(6):1072, 2013. doi: 10.1175/JTECH-D-12-00136.1.
- [11] M. Egido and E. Lorenzo. The sizing of stand alone PV-systems: A review and a proposed new method. *Solar Energy Materials and Solar Cells*, 26(1–2):51–69, 1992.
- [12] Hussein El Ghor, Maryline Chetto, and Rafic Hage Chehade. A real-time scheduling framework for embedded systems with environmental energy harvesting. *Computers & Electrical Engineering*, 37(4):498–510, 2011.
- [13] Kai Geissdoerfer, Raja Jurdak, Brano Kusy, and Marco Zimmerling. Getting more out of energy-harvesting systems: Energy management under time-varying utility with PreAct. In *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2019. doi: 10.1145/3302506.3310393.
- [14] Bernhard Haurwitz. Insolation in relation to cloudiness and cloud density. *Journal of Meteorology*, 2(3):154, 1945.
- [15] Josiah Hester and Jacob Sorber. The future of sensing is batteryless, intermittent, and awesome. In *ACM Conference on Embedded Network Sensor Systems (SenSys)*, 2017. doi: 10.1145/3131672.3131699.
- [16] IEA, IRENA, UNSD, World Bank, and WHO. Tracking SDG 7: The energy progress report 2025. Technical report, World Bank, Washington, DC, 2025.
- [17] Bashima Islam and Shahriar Nirjon. Scheduling computational and energy harvesting tasks in deadline-aware intermittent systems. In *IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2020.

- [18] Aman Kansal, Jason Hsu, Sadaf Zahedi, and Mani B. Srivastava. Power management in energy harvesting sensor networks. *ACM Transactions on Embedded Computing Systems*, 6(4), 2007. doi: 10.1145/1274858.1274870.
- [19] Jean-Yves Le Boudec and Patrick Thiran. *Network Calculus: A Theory of Deterministic Queuing Systems for the Internet*, volume 2050 of *Lecture Notes in Computer Science*. Springer, 2001.
- [20] D. V. Lindley. The theory of queues with a single server. *Mathematical Proceedings of the Cambridge Philosophical Society*, 48(2):277–289, 1952.
- [21] C. L. Liu and James W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of the ACM*, 20(1):46–61, 1973.
- [22] Shaobo Liu, Qinru Qiu, and Qing Wu. Energy aware dynamic voltage and frequency selection for real-time systems with energy harvesting. In *Design, Automation and Test in Europe (DATE)*, 2008.
- [23] Brandon Lucia, Vignesh Balaji, Alexei Colin, Kiwan Maeng, and Emily Ruppel. Intermittent computing: Challenges and opportunities. In *Summit on Advances in Programming Languages (SNAPL)*, volume 71 of *LIPICs*, pages 8:1–8:14, 2017. doi: 10.4230/LIPICs.SNAPL.2017.8.
- [24] Clemens Moser, Davide Brunelli, Lothar Thiele, and Luca Benini. Real-time scheduling for energy harvesting sensor nodes. *Real-Time Systems*, 37(3):233–260, 2007. doi: 10.1007/s11241-007-9027-0.
- [25] Clemens Moser, Jian-Jia Chen, and Lothar Thiele. Power management in energy harvesting embedded systems with discrete service levels. In *ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED)*, 2009. doi: 10.1145/1594233.1594338.
- [26] NASA Langley Research Center. POWER docs: Data sources. <https://power.larc.nasa.gov/docs/methodology/data/sources/>, 2026. Accessed 26 September 2026.
- [27] Joaquín Recas Piorno, Carlo Bergonzini, David Atienza, and Tajana Simunic Rosing. Prediction and management in energy harvested wireless sensor nodes. In *International Conference on Wireless Communication, Vehicular Technology, Information Theory and Aerospace & Electronic Systems Technology (Wireless VITAE)*, 2009.
- [28] Mahmoud Shirazi, Mehdi Kargahi, and Lothar Thiele. Resilient scheduling of energy-variable weakly-hard real-time systems. In *International Conference on Real-Time Networks and Systems (RTNS)*, pages 297–306, 2017. doi: 10.1145/3139258.3139282.
- [29] Mahmoud Shirazi, Mehdi Kargahi, and Lothar Thiele. Performance maximization of energy-variable self-powered (m,k)-firm real-time systems. *Real-Time Systems*, 56(1):64–111, 2020. doi: 10.1007/s11241-020-09344-1.
- [30] Adam H. Sparks. nasapower: A NASA POWER global meteorology, surface solar energy and climatology data client for R. *Journal of Open Source Software*, 3(30):1035, 2018. doi: 10.21105/joss.01035.
- [31] J. W. Spencer. Fourier series representation of the position of the sun. *Search*, 2(5):172, 1971.
- [32] Sujesha Sudevalayam and Purushottam Kulkarni. Energy harvesting sensor nodes: Survey and implications. *IEEE Communications Surveys & Tutorials*, 2011.

- ⁶⁶⁷ [33] Steve Vestal. Preemptive scheduling of multi-criticality systems with varying degrees of execution ⁶⁶⁸ time assurance. In *IEEE Real-Time Systems Symposium (RTSS)*, 2007. doi: 10.1109/RTSS.2007.35.
- ⁶⁶⁹ [34] Kankan Wang and Qingxu Deng. Mixed-criticality scheduling of energy-harvesting systems. In ⁶⁷⁰ *IEEE Real-Time Systems Symposium (RTSS)*, pages 435–446, 2022.
- ⁶⁷¹ [35] Kankan Wang, Yuhan Lin, and Qingxu Deng. Response time analysis for energy-harvesting mixed-
⁶⁷² criticality systems. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, ⁶⁷³ pages 1293–1298, 2022. doi: 10.23919/DATE54114.2022.9774646.
- ⁶⁷⁴ [36] Yi Xiang and Sudeep Pasricha. Mixed-criticality scheduling on heterogeneous multicore systems ⁶⁷⁵ powered by energy harvesting. *Integration, the VLSI Journal*, 61:114–124, 2018. doi: 10.1016/j.⁶⁷⁶ vlsi.2017.11.007.
- ⁶⁷⁷ [37] G. K. Young. Storage yield: Extending the sequent peak algorithm to multiple reservoirs. *Water* ⁶⁷⁸ *Resources Research*, 5(5):1110, 1969. doi: 10.1029/WR005i005p01110.