

Energy- and Mobility-Aware Scheduling for Perpetual Trajectory Tracking

Philipp Sommer, Kai Geissdoerfer, Raja Jurdak, Branislav Kusy,
Jiajun Liu, Kun Zhao, Adam McKeown, David Westcott

Abstract—Energy-efficient location tracking with battery-powered devices using energy harvesting necessitates duty-cycling of GPS to prolong the system lifetime. We propose an energy- and mobility-aware scheduling framework that adapts to real-world dynamics to achieve optimal long-term tracking performance. To forecast energy, the framework uses an exponentially weighted moving average filter to compute a virtual energy budget for the remainder of the forecast period. The virtual energy budget is then used as input for our proposed information-based GPS sampling approach, which estimates the current tracking error through dead-reckoning and schedules a new GPS sample when the error exceeds a given threshold. In order to improve the long-term tracking performance, the threshold is adapted based on the current energy and movement trends to balance the expected information gain from a new GPS sample with its energy cost. We evaluate our approach on empirical traces from wild flying foxes and compare it to strategies that sample GPS using fixed and adaptive duty cycles and by using dead-reckoning with a fixed threshold. Our analysis shows that the proposed information-based GPS sampling strategy reduces the mean tracking error compared to existing methods and approaches the performance of the optimal offline sampling strategy.

Index Terms—Trajectory tracking, energy awareness, GPS, positioning, energy harvesting, scheduler

1 INTRODUCTION

Position tracking has become an essential building block for many applications, ranging from location-based services for mobile phones to wildlife tracking. Tracking can either be user-initiated, where users request location from the device, or autonomous, where the tracking devices collect user location information based on events of interest. Advanced tracking systems have used additional sensor inputs for autonomous detection of traffic conditions [25] or potholes [10] in a city, tracking and evaluation of out-patient health [28], or sharing cycling experiences [9].

Long-term location tracking, however, remains a challenge, with GPS requiring high power consumption to continuously deliver accurate positions. High energy consumption can be often traded off for lower localization accuracy, for example, by employing the signals of nearby Wi-Fi access points or cellular base stations to estimate the location instead of the GPS receiver. Energy improvements can be substantial [17] when continuously tracking the user context and selecting the most energy-efficient localization algorithm. Advanced techniques can use accelerometers to improve energy efficiency of location tracking [19], [26]

and magnetometers to improve trajectory tracking [18] by a factor of three or more.

Nearly all of the existing energy-efficient tracking algorithms focus on minimizing energy consumption subject to an application-specific position error bound. Because of the predictability of their energy budget (for instance, mobile phones are recharged every day), most methods track position within given accuracy constraints and aim to minimize energy consumption in the process. The accuracy constraints then implicitly define the energy budget.

In this paper, we are interested in autonomous long-term tracking that relies on energy harvesting. As opposed to having a fixed daily energy budget, we investigate algorithms that adapt to the changing nature of harvested energy. Our focus is to maximize the tracking accuracy given this dynamic energy budget, which is the reverse problem of most existing energy-efficient tracking algorithms. Adaptive tracking highlights the question of how to schedule GPS samples in a given future period on the basis of a dynamic energy budget. We build on the fact that the mobility of people [15] and many animal species [27] exhibits known recurring patterns, such as periodic preferential return to home base and time-correlated trip durations. Historical movement patterns can thus be used to forecast movement in the near future and to guide the choice of parameters for a given GPS sampling strategy. In a similar way, historical records of daily harvested energy can be used to predict future energy harvesting. While historical patterns can be used for sampling forecasts, it is important to re-evaluate the sampling strategy periodically in case the current behavior of the tracked object deviates from historical trends.

Building on these concepts, we propose energy- and mobility-aware tracking that minimizes tracking errors within the available energy budget. Our strategy uses low-

- P. Sommer is with ABB Corporate Research, Baden-Daettwil, Switzerland. E-mail: phsommer@acm.org
- K. Geissdoerfer is with TU Dresden, Center for Advancing Electronics Dresden, Germany. E-mail: kai.geissdoerfer@tu-dresden.de
- R. Jurdak and B. Kusy are with CSIRO, Data 61, Brisbane, QLD, Australia, E-mail: {raja.jurdak, brano.kusy}@csiro.au
- J. Liu is with Renmin University, Beijing, China, E-mail: jiajun-liu@ruc.edu.cn
- K. Zhao is with Suncorp Group, Brisbane, QLD, Australia.
- A. McKeown is with CSIRO, Ecosystem Sciences, Cairns, QLD, Australia, E-mail: adam.mckeown@csiro.au
- D. Westcott is with CSIRO, Ecosystem Sciences, Atherthon, QLD, Australia, E-mail: david.westcott@csiro.au

power inertial sensors to detect object motion and schedules GPS sampling within the object motion periods. We predict the available energy and the remaining movement duration until the end of the forecast period to determine suitable parameters for GPS sampling. Our focus is on tracked objects with periodic preferential return to one or few key locations, and temporally correlated trip duration. We use historical movement statistics coupled with instantaneous estimates of displacement from the most likely destination to forecast movement duration.

Both the available energy and the movement forecast are updated in real-time on the mobile node and the GPS sampling strategy is continuously adapted as new estimates become available. We propose an information-based GPS sampling strategy that considers expected movement duration, expected energy availability, and expected information gain of the next GPS sample. The ratio of the forecast movement duration and forecast remaining energy budget determines an adaptive threshold for the expected tracking error. The maximum tracking error is estimated through dead reckoning and, once it exceeds the adaptive threshold, a new GPS sample is triggered. This threshold balances the expected information gain from a new GPS sample, *i.e.*, how well the original trajectory can be approximated by taking into account that position, with its cost for longer-term tracking performance. Our simulation results using empirical data traces from wild flying foxes confirm that information-based tracking delivers significant reduction in tracking error compared to static and adaptive GPS sampling and performs close to an offline optimum algorithm.

Our previous work in [32] had proposed the information-based sampling strategy within the energy- and mobility-aware framework. That work included prediction of mobility patterns (the remaining moving time in a forecast period), with a fixed energy budget. The energy budget was fixed to target the motivating application of flying foxes, which move during the night, and roost in sunlight in trees during the day. This effectively resulted in a fixed energy budget for each night, based on energy harvested during the day. Fixed energy budgets are a special case of energy-awareness where we have full certainty of our energy budget. In this paper, we consider the more general scenario where the movement and energy harvesting periods can overlap, meaning the energy budget can vary over the course of the forecast period. The variable energy budget results in uncertainty regarding energy availability. Here, we use an exponentially weighted moving average (EWMA)-based energy budget prediction for online updating of the available energy budget for position sampling.

The remainder of the paper is organised as follows. Section 2 motivates long-term tracking and outlines its objectives. Section 3 introduces our task scheduling framework for long-term tracking. Section 4 presents our algorithms for online energy and mobility prediction within the framework, while Section 5 discusses online GPS sampling strategies using these predictions, as well as an offline oracle algorithm as benchmark. Section 6 details the setup of our trace-driven evaluations of the algorithms, while Section 7 presents the simulation results. Section 8 discusses related work, while Section 9 concludes the paper with a discussion of the results, limitations, and future work.

2 LONG-TERM MOBILITY TRACKING

Long-term tracking of small mobile entities is a challenging problem with high relevance in ecology, agriculture, and logistics. The key constraint here is energy. The very need to track mobile entities long-term implies that their location is unknown and not readily accessible, which limits opportunities for manually recharging the battery of tracking devices. An alternative approach is to support energy harvesting on tracking devices, such as through solar panels, to replenish energy supplies in situ. With energy harvesting, the available energy budget is subject to the amount of energy that has been harvested in the recent past and the expected energy to be harvested while the object needs to be tracked. Furthermore, decisions to acquire position samples have to be made autonomously and in real-time by the tracking device on the basis of both available energy and likely movement activity as communicating with the device for setting parameters manually might not be feasible.

2.1 Accuracy- vs. Energy-bound Tracking

Most recent work has focused on accuracy-bound tracking. Given an uncertainty or accuracy bound on location, the goal is to minimize energy consumption while operating within this bound to maximize lifetime. The accuracy bound de-facto defines the minimum required energy budget of the position tracking, given a typical motion pattern of the tracked object. If the available energy resources are smaller than the tracking budget, the device will run out of energy and the accuracy-bound tracking will fail to deliver the required performance.

With energy harvesting, a well-designed system can operate near-perpetually as long as it tailors its consumption to its available energy budget. Given a fixed energy budget, the focus becomes on how to maximize the location accuracy within a specific future period. We refer to this problem as energy-bound tracking. Solving this problem naturally depends on the sampling frequency of GPS and other supporting sensors (energy output), which are linked to the underlying movement dynamics of the tracked object and the harvesting energy dynamics of the environment (energy input). For energy forecasts, we use an EWMA filter on historical data to predict the harvested energy for a future period.

Objects that move rarely can be tracked reasonably well through a sparse scheduling strategy, while determining an energy-efficient schedule for tracking highly mobile objects with high accuracy is a challenging problem. For this reason, our work here not only considers the energy availability through harvesting sources, but also the mobility of the tracked objects. This paper specifically focuses on tracked objects that follow regular daily movement patterns, such as commuting trips to foraging or activity locations and return to one or few “home” locations [15], [27]. More specifically, we predict the remaining movement distance for a tracking interval on the basis of two estimates: (1) movement distance during the previous interval; and (2) current distance from the most likely destination. Each of these estimates proves to be more accurate in a specific time frame within the tracking interval.

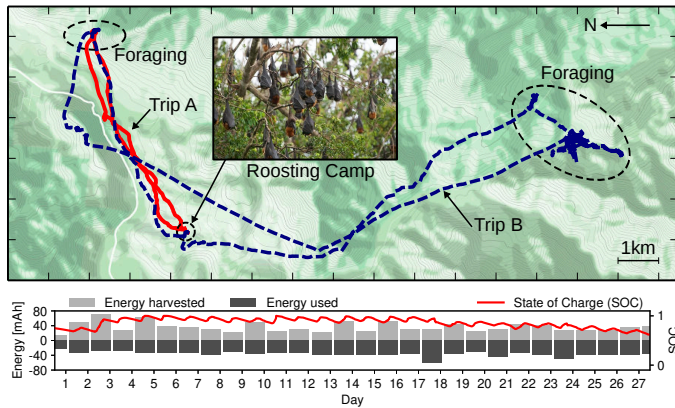


Fig. 1: Two GPS trajectories of a spectacled flying fox tagged with a mobile tracking device. The bar plot shows daily harvested solar energy, consumed energy of the tracking device, and battery state-of-charge (SOC).

2.2 Motivating Application

We focus our evaluation on a specific application of tracking flying foxes in their habitat across Australia. However, our approach can be applied to tracking any mobile objects that have regular motion patterns and periodic energy harvesting intervals, such as tracking commuting times of people to work, tracking people during leisure activities such as biking, roller-blading, or jogging, and tracking domestic, farm, or wild animals.

Flying foxes, also known as fruit bats, play an important ecological role as dispersers of pollen and seed in wet tropical areas of Africa and Asia, but are also vectors for a range of emerging infectious diseases which have serious consequences for human health, *e.g.*, Hendra virus in Australia, Nipah in Asia and Ebola in Africa. Flying foxes are highly mobile, with individuals able to fly hundreds of kilometers during nightly foraging. Key to understanding and managing these animals is to understand how they utilize landscapes and how they interact with other disease hosts, which requires a fine-grained understanding of their movement. However, their weight (600 g to 1 kg) limits the weight and size of tracking devices that can be placed on them [24], imposing a tight energy budget on sensing activities, such as GPS sampling.

Behavioral patterns. Flying foxes are nocturnal animals and during the daylight rest in large groups at sites called *camp*s or *roosts*. During the day they generally only move within the camp (within a radius of about 100-200 meters) and for most of the time they sleep at a single location. In contrast to the family of microbats, which are common in Europe and the US, flying foxes do not roost in caves, but hang upside down from the top branches of large trees exposed to direct sunlight. At dusk, they leave the camp and fly out to forage at sites that are usually tens of kilometers away. During the night they may change location several times before returning to the camp before sunrise, as shown in Figure 1.

2.3 Limitations of Static GPS Sampling

Current commercial motion trackers use a combination of motion-triggering and time-based duty cycling of GPS for

energy management. In particular, motion sensors are used to detect the start and end of motion events according to a given movement threshold and GPS samples are only taken during the motion events with a fixed duty cycle. We refer to this approach as static motion-based tracking, which effectively fixes the GPS sampling strategy for the duration of a typical motion event. This approach does not work well when the mobility dynamics or the available energy vary day-to-day. A static scheduler might miss key parts of a trajectory if it underestimates the length of the trajectory or overestimates the energy availability. Autonomous long-term tracking applications, therefore, need to adapt GPS sampling in order to reconstruct the original trajectory as accurately and efficiently as possible.

To illustrate the energy and motion duration dynamics, we present data from a flying fox tracking application in Figure 1. We show an example of two trips made by the same animal that vary in duration. We also note large differences in harvested energy for an individual animal on a day-to-day basis over a 4-week period of empirical data, illustrating the likelihood for differences in energy intake and movement duration among individuals and over time for the same individual.

3 TASK SCHEDULING FRAMEWORK

In this section, we introduce our framework for location tracking of mobile nodes that adapts to changing energy availability and motion patterns. At its core is the *Task Scheduler* component, which uses inputs from the *Energy-Awareness Layer* and *Mobility-Awareness Layer* to schedule sampling of different hardware components in the *Sensor Process Layer*, such as the GPS receiver or inertial sensors. The building blocks of the framework are depicted in Figure 2.

3.1 Task Scheduler

We assume that physical or remote access to the mobile node may not be feasible for prolonged time periods, so we develop mechanisms to support autonomous operation and reconfigurability. Specifically, the scheduler allows for adaptive execution of GPS sampling tasks based on the input from the energy- and mobility-awareness layers. This design principle allows us to adapt the GPS sampling rate in real-time based on the available energy and predicted duration of activity (see Section 5).

3.2 Sensor Process Layer

The sensor process layer contains software components that control the operation of the sensors implemented on the hardware platform (*e.g.*, GPS, accelerometer, magnetometer, etc.). This layer provides an abstraction of the underlying sensor hardware and allows to control and monitor the power state of the sensor, as well as to acquire sample values from the sensors. The sensing modalities used as input for the mobility-awareness layer depend on the capabilities of the hardware platform and on the application scenario.




Fig. 2: Building blocks of the energy- and mobility-aware scheduling framework.

3.3 Energy-Awareness Layer

The energy-awareness layer is responsible to keep track of the energy consumed and harvested by the system (energy monitoring) and to provide an estimate of the energy resources available in the near future (energy prediction).

Precise book-keeping of both energy consumption and energy harvesting is key to maintaining a realistic estimate of the energy available at any time. The *Energy Monitoring* component keeps track of the available energy stored in the batteries, the so called *State of Charge (SOC)*, the energy consumed by the node E_{used} , and the energy harvested E_H .

Energy Book-Keeping. We keep track of the energy consumption during runtime using a software-based book-keeping mechanism. The instantaneous overall consumption of a sensor node can be broken down into the power consumption of its individual hardware components, which includes the microcontroller, sensors, radio transceiver, and flash storage chip. The total instantaneous power consumption $P(t)$ of the sensor node can be calculated as the sum of the power consumption of its components at time t :

$$P(t) = \sum_i P_i(t) = \sum_i V_i(t) \cdot I_i(t) \quad (1)$$

where $P_i(t)$ is the instantaneous power consumption of component i , which depends on the supply voltage $V_i(t)$ and current consumption $I_i(t)$ of its operating state. We assume that the software is able to control and/or monitor the operating state of a component (e.g., active, idle or sleep mode) and that the current consumption in each state is known from device datasheets or measurements. Consequently, we can estimate the energy consumed by the node E_{used} from the beginning of the interval, by integrating the power consumption over the observation period. We have shown in lab experiments that software book-keeping can estimate the energy consumption within 10% even when using a simple model based on estimated average current draw and duration for each operating state of a component [31].

Energy Harvesting. The hardware platform is able to measure the voltage $V_H(t)$ and the charge current $I_H(t)$ into the battery while energy harvesting is in progress, which allows to estimate the harvested energy E_H during a time interval. Historical patterns of the amount of harvested energy are used by the *Energy Prediction* module to predict energy harvesting in the near future (see Section 4).

State of Charge Estimation. By combining energy book-keeping with the measurement of harvested energy, we can

calculate the estimated total available energy (SOC) at time t as follows:

$$SOC(t) = SOC(t_{\text{start}}) + \sum_{\tau=t_{\text{start}}}^t (E_H(\tau) - E_{\text{used}}(\tau)), \quad (2)$$

where $E_H(\tau)$ denotes the harvested energy and $E_{\text{used}}(\tau)$ is the total energy consumed by different system components at time τ . In this paper, the time interval begins at noon of the current day (t_{start}) and lasts until noon of the next day (t_{end}), at which point it resets.

3.4 Mobility-Awareness Layer

The mobility-awareness layer uses the input of different sensing modalities to capture the underlying mobility dynamics of the tracked objects (activity detection) and predicts the duration of future movements accurately (motion prediction). As these dynamics are largely dependent on the type of the tracked object, we focus here on our motivating application of long-term flying fox tracking.

Activity Detection. In order to facilitate activity detection based on low-power sensor inputs, many mobile tracking platforms feature inertial sensors. For example, the Camazotz hardware platform (see Section 6.1) is equipped with a combined 3-axis accelerometer and 3-axis magnetometer. The accelerometer provides samples of the acceleration along its three axes at a sampling frequency between 1 Hz and several kHz. Furthermore, the accelerometer provides a programmable interrupt for detection of movement events, e.g., when the animal changes from roosting to flying. This functionality can unburden the micro-controller from having to continuously read samples from the accelerometer's output buffer, so that it can remain in sleep mode for most of the time. The magnetometer measures the strength of the magnetic field at a sampling rate between 0.75 Hz and 220 Hz along three axes, which can detect heading changes.

Classification of Activity. We show a real-world example for activity detection using inertial sensors on a flying fox tagged with a Camazotz node in Figure 3. GPS data is used to provide ground truth on the animal activity state, movement duration, and heading. In our application scenario, continuous sampling of the accelerometer at 10 Hz provides enough information to detect animal movement with high accuracy. Specifically, our activity classifier accumulates acceleration changes along the z-axis in each 5-second window. The animal is classified as active in the time window if the average accumulated acceleration change is above a learned threshold. For groundtruth we use time windows in which both GPS and accelerometer readings are




Fig. 3: Accelerometer readings (top) and magnetometer heading (bottom) versus GPS speed/heading recorded on a flying fox. The animal starts flying at $t=30$ s.

present. The animal is considered truly active if the average GPS speed in the time window is more than 1 m/s.

The process of getting an adaptive threshold for the classification is as follows: first we use a fixed population-level threshold as the initial threshold, which we choose based on a high-frequency dataset with both GPS and accelerometer data that we collected in the initial phase of the animal tracking project. Then this initial threshold is adjusted adaptively in a dedicated learning phase, with both accelerometer and GPS readings accessible to the algorithm regardless of the animal’s motion status. In the learning phase, the algorithm classifies each time window using the current threshold, and validates the result with the actual activity state from the GPS. If a false negative (active classified as non-active) occurs, the threshold is reduced by a percentage (i.e. 30% in the experiments) of its current value. Otherwise when a false positive (non-active classified as active) occurs, the threshold is increased by the same percentage. The threshold is then fixed and used at all times when the GPS is not active on the node.

Validation. We perform cross-validation on the GPS and accelerator samples that are collected from two sensor nodes and span 137 minutes. The dataset consists of 2,283 GPS samples and 30,972 matching accelerometer samples from continuous time frames for each of the nodes. One node is mostly stationary and the other mostly active. In total, we have 2,277 valid samples in two segments, with 1,081 in moving state and 1,196 in stationary state. We split each node’s data into five consecutive segments, and iteratively combine the training and test set from each segment into a unified training set and a unified test set. Therefore there are 25 repetitions. For each repetition, we use the adaptive algorithm on the training set to learn the classification threshold, and then perform the detection on the test set to evaluate the prediction performance. Note that all folds after the split are still consecutive samples to avoid the over-estimation effect mentioned in [34]. Finally we use the average precision and recall from the 25 repetitions as the overall validation results. Empirical result for the average precision reaches 98.2%, and 98.7% for the average recall, with the learned threshold fluctuating around 0.7~0.9.

4 ENERGY AND MOBILITY PREDICTION

In this section, we describe the models used in our scheduling framework to predict the amount of energy harvested

in the near future and the duration of the remaining activity within the forecast period. Prediction of the harvested energy allows us to consider cases when the activity of the tracked object and energy harvesting overlap in time, such as animals that are active during the day. Note that our previous work [32] assumed the energy harvesting and motion/activity phases to be separate, which is the case in the flying fox application scenario. Separating the harvesting and motion periods effectively results in a deterministic energy budget. In a prediction context, this provides a perfect prediction with zero uncertainty. By allowing overlap of harvesting and motion periods, we explore the more general case where energy budgets change over time and need to be predicted into a future forecast period, which inherently involves uncertainty.

4.1 Energy Prediction

Our approach assumes that the harvested energy on a day is subject to seasonal variations and that it is similar to the average energy harvested during the recent past. We adopt the exponentially weighted moving average (EWMA) filter approach [16] for short-term energy harvesting prediction to predict the expected energy that will be harvested from the start (denoted as t_{start}) until the end of the forecast period (denoted as t_{end}), as shown in Algorithm 1. We choose EWMA because it works with only historical values of harvested energy and has low memory and computational demands, allowing us to run it on our resource constrained platform.

We set the maximum forecast period to one full day (starting and ending at noon), and split it into fixed 30 min slots. We introduce the time slot index $i = (t \bmod 1d)/30m$. $E_H(i)$ stands for the energy harvested during slot i . We maintain an array of the exponentially weighted moving average of energy harvested $\bar{E}_H(i)$ during slot i . For each slot, we update $\bar{E}_H(i)$ as the weighted average ($\alpha = 0.5$) of the historical average and the latest observation of energy harvested during this slot:

$$\bar{E}_H(i) = \alpha \bar{E}_H(i) + (1 - \alpha) E_H(i). \quad (3)$$

We forecast the energy harvested in a time slot for one day to be equal to the average calculated on the previous day. Consequently, the predicted total energy to be harvested until the end of the day $\hat{E}_{H,tot}(i)$ is calculated as an aggregate of the estimated energy harvested in each future time slot until the end of the forecast period t_{end} :

$$\hat{E}_{H,tot}(i) = \sum_{j>i}^{t_{\text{end}}} \bar{E}_H(j). \quad (4)$$

Note that $\hat{E}_{H,tot}(i)$ can be calculated ahead for later time slots on the same day. We assume the energy input to be arriving at a constant rate during each time slot. At any time t , the estimated energy input until the end of the day is then given as the linear interpolation between the values at the beginning of the current slot i and the beginning of the next time slot $i + 1$ as follows:

$$\hat{E}_{H,tot}(t) = \hat{E}_{H,tot}(i) + \frac{\hat{E}_{H,tot}(i+1) - \hat{E}_{H,tot}(i)}{30m} \cdot (t \bmod 30m). \quad (5)$$

Algorithm 1 Energy Prediction Algorithm

Require: $\bar{E}_H(i)$ \triangleright Recent Avg. harvested energy in slot i
Require: $E_H(t)$ \triangleright Harvested energy from book-keeping
1: **procedure** PREDICT_ENERGY_INPUT(i)
2: $\hat{E}_{H,tot}(i) = \sum_{j>i}^{t_{end}} \bar{E}_H(j)$ \triangleright Predict energy input
3: $\bar{E}_H(i) = \alpha \bar{E}_H(i) + (1 - \alpha) E_H(i)$ \triangleright Update average

Algorithm 2 Virtual Budgeting Algorithm

Require: $E_{balance}(t)$ \triangleright Accumulated energy balance
Require: $\hat{E}_{H,tot}(i)$ \triangleright Energy prediction beginning of slot
Require: $\hat{E}_{H,tot}(i+1)$ \triangleright Energy prediction end of slot
Require: $E_H(t)$ \triangleright Harvested energy from book-keeping
Require: $E_{used}(t)$ \triangleright Used energy from book-keeping
1: **procedure** CALCULATE_VIRTUAL_BUDGET(t)
2: $E_{balance}(t) += (E_H(t) - E_{used}(t))$
3: $\hat{E}_{H,tot}(t) = \hat{E}_{H,tot}(i) + \frac{\hat{E}_{H,tot}(i+1) - \hat{E}_{H,tot}(i)}{30m} \cdot (t \bmod 30m)$
4: $E_{budget}(t) = \hat{E}_{H,tot}(t) + E_{balance}(t)$

The scheduler can over- or under-spend the available energy if the actual harvested energy or animal mobility deviate from their respective predictions. This deviation $E_{balance}$ is accumulated from all past values of harvested and spent energy since the beginning of the forecast period t_0 , as reported by the energy book-keeping module:

$$E_{balance}(t) = \sum_{\tau=t_0}^t (E_H(\tau) - E_{used}(\tau)). \quad (6)$$

The scheduler compensates for this deviation in an on-line fashion. Specifically, we introduce the concept of the virtual energy budget E_{budget} , which is the energy that is available to the sensor node for the rest of the forecast period at time t . We calculate E_{budget} from the energy we predict will be harvested for the rest of the forecast period and then compensate for the prediction errors made during the current and previous forecast periods:

$$E_{budget}(t) = \hat{E}_{H,tot}(t) + E_{balance}(t). \quad (7)$$

The virtual budget is used by the scheduler for sampling decisions of the GPS algorithm (see Section 5). We refer the reader to a detailed description of the virtual budget calculation in Algorithm 2. Note that the algorithm is designed to run after an initial training period, during which only steps 2 (tracking the energy balance for each time step) and 3 (updating the average energy harvested for each slot) are executed. The training period needs to be at least one forecast period, so that the algorithm can populate $\bar{E}_H(i)$ for all i .

4.2 Mobility Prediction

Optimal scheduling depends both on the available energy budget and the amount of energy that will be used for tracking movements. Therefore, it is important that the GPS process can estimate for how long the node will be moving within a certain time interval. Clearly, the accuracy of our prediction will directly influence the energy consumption of the tracking algorithm. If the actual total tracking time exceeds the predicted time, we risk running out of energy early and we sacrifice accuracy of tracking towards the end of the observation period. On the other

Algorithm 3 Mobility Prediction Algorithm

Require: $d_{camp}(t)$ \triangleright Current distance to camp
Require: d_{prev_max} \triangleright Previous max. distance from camp
Require: Δt_{motion} \triangleright Estimate for remaining movement
Require: v_{avg} \triangleright Average flying speed
Require: $d_{total}(d)$ \triangleright Estimate of total flight distance
Require: t_{start}, t_{end} \triangleright Observation interval
procedure PREDICT_REMAINING_DISTANCE(t)
if $d_{camp}(t) < d_{prev_max}$ **then**
if $\Delta t_{motion} \cdot v_{avg} < d_{camp}(t)$ **then**
 $\gamma \leftarrow \frac{t - t_{start}}{t_{end} - t_{start}}$ \triangleright Weight factor
 $d_{remaining}(t) \leftarrow (1 - \gamma) \cdot \Delta t_{motion} \cdot v_{avg} + \gamma \cdot d_{camp}(t)$
else
 $d_{remaining}(t) \leftarrow d_{camp}(t)$
else
 $d_{remaining}(t) = d_{total}(d_{camp}(t))$

hand, overestimating the total tracking time will result in conservative budgeting of energy resources and may lead to larger localization errors. The implementation of the mobility prediction module needs to be tailored to the designated application scenario and might take into account input from several sensors, last location from GPS, the current time/date information and historical data.

Prediction of Movement Duration. For our animal tracking application, we developed a hybrid approach to predict animal mobility based on individual- and population-based models and real-time estimation of the remaining motion duration based on the remaining distance to the nearest roosting camp (see Algorithm 3). We start by using an initial estimate for the total flight distance based on the previous night or an individual-based model. If we detect that the animal has exceeded the maximum distance from camp observed in the previous night d_{prev_max} , we switch to a population-based model to estimate the total flight distance based on historical trajectories.

Mobility models. We assume that we are given a set of historical trajectories of an object of interest with a certain level of regularity of motion that allows us to extract simple motion features. Note that this motion regularity is shared by many other species including humans [33]. An estimate for the motion duration of an animal per night Δt_{motion} can be calculated by observing typical animal behavior in historical data (see Section 6). Tracking of new animals can be done using a population-level motion model which maps the current distance from the camp $d_{camp}(t)$ to the total flight distance d_{total} , while a more accurate individual-level model can be derived over time as more data becomes available.

Distance to home base. Tracked objects may exhibit large variations in motion patterns across the population, or across time (see Figure 1). Prediction of motion duration that is based on simple statistics only, such as the mean duration of motion, will therefore be inaccurate. We observe that flying foxes usually return to the same roosting camp after the nightly foraging (similarly to humans commuting to/from home and work). The current location of the animal can thus be used as a lower bound on the remaining flight distance. This provides us with a lower bound on the remaining motion time, which in turn determines the lower bound on energy required to capture the return trip to the camp. Clearly, this approach can only provide a good estimate after the animal has reached the maximum dis-

tance from the camp, which usually happens in the second half of the night. We therefore employ a linear weighting factor γ which applies more weight to the individual- or population-based mobility prediction in the beginning of the observation period, while the distance to base becomes more important towards the end of the observation period, or as soon as it exceeds the initial distance estimate.

Having defined the scheduling framework and prediction algorithms, the next 3 Sections focus on our approach to use these predictions to drive GPS sampling in a real world tracking scenario. Section 5 introduces three online GPS sampling strategies that use some or all of the predictions, as well as an offline oracle algorithms that we use a performance benchmark. Section 6 presents the hardware platform assumptions and empirical traces that we use for comparing the performance of these GPS sampling strategies. Section 7 evaluates our prediction accuracy for energy and mobility and the performance of our proposed online GPS sampling algorithm against existing online approaches and the oracle approach.

5 GPS SAMPLING STRATEGIES

In this section, we present four online and one offline strategies to configure sampling of the GPS receiver. Our scheduling framework is agnostic of the actual GPS sampling strategy, and it assumes that each strategy will tune its sampling parameters based on the given constraints. Energy harvesting and mobility prediction described in Section 4 provide two key inputs to the GPS scheduler, as they define energy inputs and outputs that need to be balanced out. This introduces an additional time-varying constraint on GPS sampling as the used energy should be smaller than the harvested energy at any point in time. We assume that the battery capacity is substantially larger than the energy harvested per day and therefore the system can operate around a target SoC that allows to buffer surplus harvested energy and balance energy deficits. The GPS schedulers can then operate with a virtual energy budget E_{budget} and aim to achieve energy neutral operation at the end of the forecast period, rather than at any point in time.

5.1 Motion-based GPS Sampling

The activity detection module (see Section 3.4) can be configured to enable the GPS receiver while movement of the tracked object is detected. During the movement period (Δt_{motion}), periodic GPS samples are acquired. Between position samples the GPS receiver remains in sleep mode to reduce the power consumption. The total energy consumption in the time interval $\Delta t_{\text{interval}} = t_{\text{end}} - t_{\text{start}}$ is proportional to the GPS scheduling period (T_{sampling}) and can be calculated as follows:

$$E_{\text{used}} = \Delta t_{\text{interval}} \cdot P_{\text{baseline}} + k \cdot T_{\text{hotstart}} \cdot P_{\text{tracking}} \quad (8)$$

P_{baseline} denotes the baseline power consumption of the system, P_{tracking} is the additional power used during tracking operation, and T_{hotstart} is the average duration of a GPS hotstart. $k = \lfloor \frac{\Delta t_{\text{motion}}}{T_{\text{sampling}}} \rfloor$ is the total number of recorded GPS samples.

Strategy #1: Static Motion-based. The static motion-based strategy employs a fixed GPS sampling interval during




Fig. 4: Dead-reckoning using heading estimation based on magnetometer data. The current position is extrapolated from the last GPS position and speed v_{gps} .

movement periods. Specifically, the optimal GPS sampling interval can be calculated using the following inequality: $E_{\text{used}} \leq E_H$. Performance of the static GPS scheduler critically depends on our initial estimates of the energy budget and mobility patterns of the tracked object, which determine the constant GPS sampling interval T_{sampling} . We can get estimates for both the harvested energy and motion duration by studying typical behavior of a population of animals. However, any individual deviations from the typical behavior, for example, due to weather or changes in food availability will result in a suboptimal GPS sampling strategy. This strategy is used by many commercial motion-based GPS trackers that operate at a fixed GPS duty cycle.

Strategy #2: Adaptive Motion-based. In contrast to static movement-based tracking that sets a fixed T_{sampling} for the whole time window, the adaptive approach continually updates the GPS sampling interval to meet the energy constraints. By distributing the *remaining* GPS samples $k(t)$ uniformly over the predicted *remaining* duration of motion $\Delta t_{\text{motion}}(t)$, we obtain the *updated* GPS sampling interval $T_{\text{sampling}}(t)$ and the corresponding duty cycle $DC(t)$ as follows:

$$T_{\text{sampling}}(t) = \frac{\Delta t_{\text{motion}}(t)}{k(t)} \quad DC(t) = \frac{k(t) \cdot T_{\text{hotstart}}}{\Delta t_{\text{motion}}(t)} \quad (9)$$

5.2 Information-based GPS Sampling

Periodic duty-cycling of GPS during movement does not take into account the intrinsic characteristics of the trajectory, for example sharp turns, to identify key points that will result in small errors when used as input for trajectory interpolation. Therefore, our aim is to obtain GPS samples at a minimal number of key locations that provide detailed information about the shape of the trajectory.

Sampling inertial sensors *while* the node is in motion can provide a rough estimate of the current heading while using only a small fraction of the power consumption of the GPS (see Figure 3). By employing *dead-reckoning* using the magnetometer heading, the current position can be extrapolated using the last known GPS position and speed v_{gps} . In each time step, we calculate the tracking errors $e[t]$ between the dead-reckoning-based position estimates at time t and the corresponding interpolated positions on a straight line between the last known GPS position and the most recent position estimate by dead-reckoning (see Figure 4).

If the interpolated trajectory based on the previous GPS sample and dead-reckoning exhibits small estimated tracking errors $e[t]$, there is little benefit in spending energy to take another GPS sample. On the other hand, a large estimated tracking error increases the value of taking another GPS sample to avoid incurring even larger errors in

the future. Therefore, as soon as the estimated maximum tracking error $e_{\max} = \max_t e[t]$ exceeds the threshold R another GPS sample is acquired and the dead-reckoning phase is restarted.

Strategy #3: Static Information-based. In this strategy, a *fixed* error threshold R for dead-reckoning is defined at the start of the observation window based on the expected movement duration and the available energy budget. This strategy is conceptually similar to the heading-aware sampling approach of EnTracked_T [18], where a fixed error threshold can be configured. Note that the trajectory error of EnTracked_T is calculated as the orthogonal distance from the straight line defined by the initial heading, while the information-based trajectory error is calculated based on a straight line between the initial position and the most recent position estimate based on dead-reckoning.

Strategy #4: Adaptive Information-based. In the adaptive version of the information-based GPS tracking, we calculate a *dynamic* threshold $R(t)$ for the tolerated tracking error based on the current duty cycle $DC(t)$ using the following equation:

$$R(t) = R_0 \cdot \left(\frac{1}{DC(t)} - 1 \right)^\beta. \quad (10)$$

The intuition is to continually update the GPS sampling policy based on the relationship between the expected *remaining* movement duration and the *remaining* energy budget. The ratio of movement duration and energy budget expressed in remaining GPS sample duration is effectively the inverse of the duty cycle $DC(t)$ in Eq. 9. As a result, the threshold $R(t)$ will increase when the remaining energy budget decreases, which will yield fewer GPS samples and a more conservative sampling policy. On the other hand, the calculated threshold decreases if the remaining energy budget and the expected motion duration are similar (when the expected movement duration roughly matches the budgeted GPS samples), allowing for more aggressive GPS sampling. R_0 and β are tuning parameters for calibrating the scale of the threshold. In case of $\lim_{\beta \rightarrow 0} R(t)$ converges to a fixed threshold R_0 .

The parameter settings for R_0 and β depend on the characteristics of the trajectories. Given a dataset with historical trajectories, we can obtain the resulting duty cycle for specific choices of the threshold parameter R by means of simulation and calculate the most suitable setting for R_0 and β by curve fitting. Alternatively, the scheduler can start with an initial estimate and optimize the parameters used to calculate the threshold R in an online fashion based on the measured resulting duty cycle.

5.3 Optimal Offline Scheduling of GPS

As a performance benchmark, we consider an offline oracle-based optimal strategy that provides a lower bound on the tracking errors for a given energy budget by finding the optimal sampling points. The offline optimal strategy cannot be implemented on a mobile device since it requires knowledge of the complete trajectory.

Strategy #5: Offline GPS Scheduling. Let a sequence $P = \{1, 2, \dots, N\}$ denote the points of the ground truth trajectory. Suppose we take GPS samples at the first and last points of the trajectory by default, and the energy

budget only allows us to take at most k additional GPS samples to form an estimated trajectory. The optimal estimated trajectory is a subsequence of P , denoted by $S = \{1, i_1, \dots, i_k, \dots, i_{|S|-2} \leq k, N\}$, which has the minimum worst-case tracking error $e_{opt} = \operatorname{argmin}_S e_{\max}(S)$. This problem is equivalent to the min- ϵ problem in curve approximation and can be solved by a graph-search approach in polynomial time [5]. The basic idea is to construct a directed graph $G(V, E)$ with the vertices set $V = P$ and edges set $E = \{(i, j) | d_{ij} \leq \epsilon, i < j\}$ and use G to find the optimal estimated trajectory. Here d_{ij} is the tracking error between the estimated trajectory \hat{i}_j to the ground truth trajectory. A shortest path from vertex 1 to vertex N in G then represents the optimal estimated trajectory with tracking error ϵ .

5.4 Comparison of GPS Tracking Strategies

We provide a case study trajectory from our flying foxes dataset comparing the behaviour of different GPS sampling strategies given the same energy budget (see Figure 5). Static sampling strategies employ a *population-based* historical estimate of movement duration. For the adaptive sampling strategies, the node starts out with an *individual-based* historical estimate of movement duration, and as the time progresses the weight of this estimate progressively decreases, while the weight of the current distance from base increases. We have chosen a challenging example trajectory, where the tracked animal explores additional foraging areas resulting in a considerably longer travel distance than the average nightly foraging distance.

Clearly, the information-based and offline approaches capture the original trajectory well. In contrast, the static motion-based approach underestimates the flight time as this animal is flying longer than the initial estimate would suggest, so the node depletes its energy resources early during the night and suffers high position errors from that point on. The remaining flight trajectory back to the roosting camp can only be interpolated. The example traces demonstrate the benefits of the information-based scheduling strategy, which also takes into account the shape of the trajectory based on the dead-reckoning approach. As a result, the sampled trajectory clearly shows non-uniform sampling patterns, where long straight segments are sampled infrequently and more tortuous segments are sampled with higher frequency. Adaptive information-based tracking continuously adapts the error threshold for sampling, and thereby the sampling schedule, according to updated estimates of remaining flight time and energy budget. Static information-based scheduling approach, which is similar to EnTracked_T, shows similar performance for the first part of the trajectory. However, as the error threshold is not adapted based on the remaining movement duration, it slightly overspends on energy which allows to take fewer samples towards the end of the trajectory, which results in a larger tracking error.

6 EXPERIMENTAL SETUP

In this section, we describe the experimental setup for evaluating our tracking framework using data gathered during a deployment on wild flying foxes.

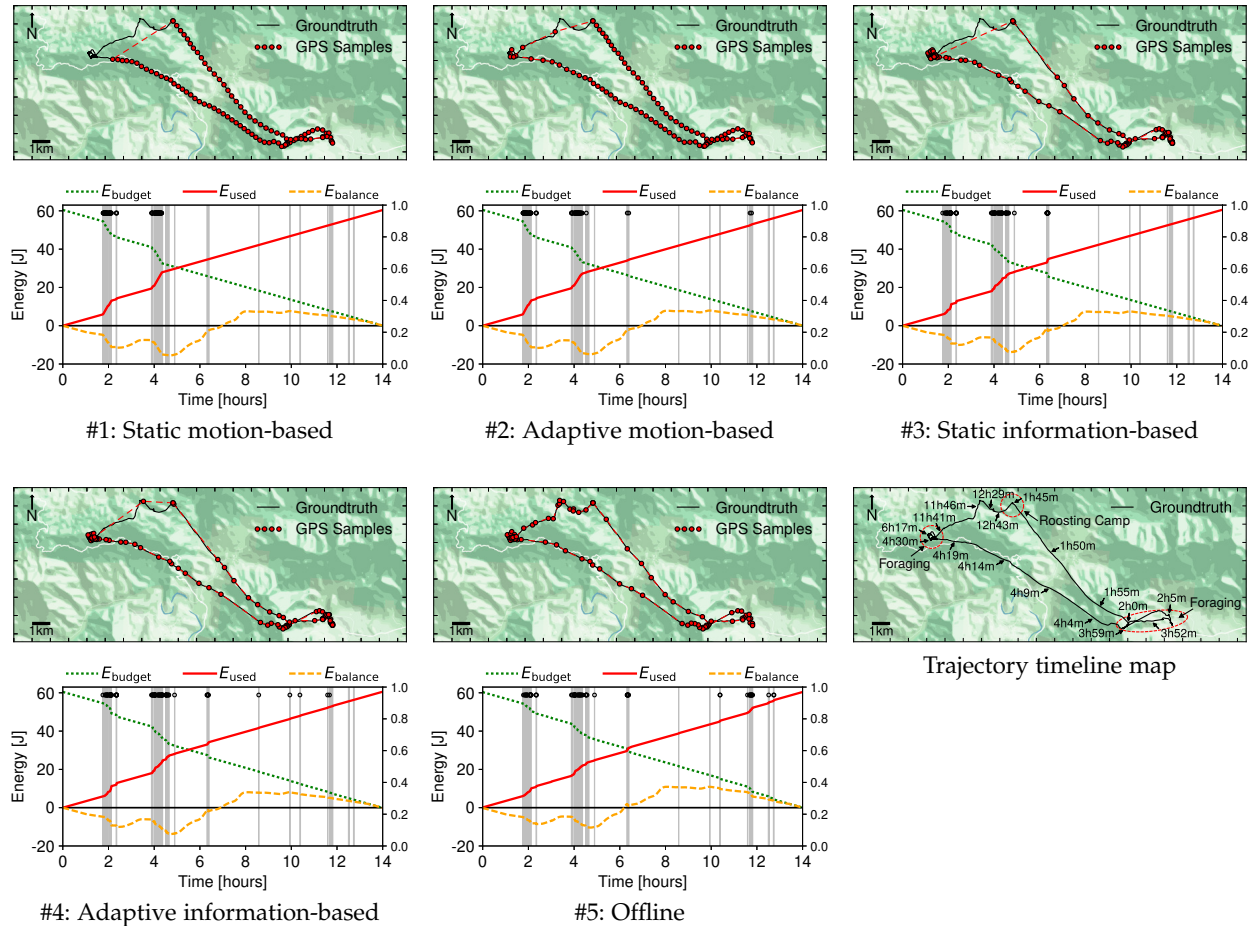


Fig. 5: GPS sampling points and energy profile for different GPS sampling strategies on a flying fox trajectory. Maps show the groundtruth trajectory and the selected sampling points for each GPS sampling strategy. The grid size of the map is 1 kilometre. The plot shows the energy budget $E_{budget}(t)$, energy consumption $E_{used}(t)$ and energy balance $E_{balance}(t)$ (see Equation 6). Plot areas shaded in gray indicate activity periods and GPS samples are indicated with a black circle. Time is relative to the start of the observation period at 5 pm local time.

6.1 Hardware Platform

In this work, we employ the Camazotz platform [14], a low-power sensor platform for GPS location tracking, which is optimized for small size and low weight (see Figure 6). The platform is based on the Texas Instruments CC430F5137 system-on-chip, which combines a microcontroller with a sub-GHz short-range radio transceiver. Besides a GPS receiver, the Camazotz platform provides various sensing modalities, such as accelerometer, magnetometer, pressure and temperature sensor, and a microphone. In this work, we employ the inertial sensors (accelerometer and magnetometer) as input for activity detection, while the other sensors are not used and remain in sleep mode. Full details of the platform’s energy supply and harvesting, GPS, sensors and wireless communication capabilities are available in [32].

6.2 Empirical Traces

Our empirical data set is based on tracking and energy harvesting data from 5 free-living flying foxes that were tagged with a Camazotz tracking device¹. Details of the

1. Ethical approval for experiments with animal subjects has been granted.

GPS task configuration are available in [32]. Based on an observation period of several weeks, we select a dataset with 51 trips containing more than 140,000 GPS data points with high temporal resolution, which allows reconstruction of the corresponding trajectories with only minor gaps (e.g., due to GPS hotstarts or temporary signal degradation). We further employ the empirical traces to learn the population-wide parameters used by the information-based tracking algorithm (see Equation 10).

Energy Consumption and Harvesting. We estimate both energy consumption E_{used} and harvested solar energy E_H for each 24-hour period starting at noon, as shown in Figure 7. For energy harvesting, the charge current from the solar panel is sampled every 10 minutes. We obtain the solar power by multiplying with the system voltage and up-sample to 1 s by linear interpolation. Other than the energy consumption of the GPS receiver, we also account for the baseline consumption of the microcontroller, inertial sensors, flash storage chip, and radio transceiver during idle listening, contact logging and data offload to the base station. During the data collection phase, we report an average energy consumption of 33.0 mAh and an average harvested energy of 27.3 mAh daily. It is important to note that the




Fig. 6: Camazotz collars (left) and a spectacled flying fox (*Pteropus conspicillatus*) with collar (right).




Fig. 7: Daily energy consumption and harvested solar energy for the Camazotz nodes on wild flying foxes.

GPS task configuration is chosen to provide consecutive GPS samples to serve as a high resolution ground truth data set for our animal tracking case study. However, the energy consumption of such a scenario, which is mainly dominated by the GPS receiver, will exceed the amount of daily harvested solar energy. It can therefore not guarantee long-term tracking operation, which clearly shows the need for an energy management framework that can balance energy inputs and outputs.

Inertial sensors. As collecting raw samples from the inertial sensors is not feasible due to storage constraints on the mobile nodes, we simulate magnetometer heading samples based on GPS heading angle affected by a random jitter from a normal distribution with standard deviation of 10.0° . This value has been obtained by comparing magnetometer samples with heading information provided by GPS obtained during experiments on wild flying foxes.

7 EVALUATION

In this section, we present an evaluation of the individual building blocks of our task scheduling framework presented in Section 3. Furthermore, we evaluate the tracking performance for the GPS sampling strategies from Section 5 with our proposed framework when applied to a real-world animal tracking dataset and application that are detailed in Section 6.

7.1 Mobility Prediction

Distance from camp. We calculate the maximum distance from the camp and the total flight distance within that day (see Figure 8) for a dataset consisting of 314 trips by




Fig. 8: Maximum distance from the camp for two consecutive days (left), and maximum distance from the base camp and the total flight distance (right).

10 animals. Time periods where we do not have continuous high-resolution GPS trajectories available are interpolated based on hourly GPS positions. We can observe that today's and yesterday's maximum distance from the camp are weakly correlated (Pearson correlation = 0.35), confirming that yesterday's distance alone is not sufficient for predicting today's distance. Furthermore, the maximum distance from the camp is also related to the total flight distance on that day (Pearson correlation 0.86). We use these observations to build a model for the total flight distance d_{total} based on the current distance from camp $d_{camp}(t)$ (see Algorithm 3).

Motion Duration. We present a summary of motion statistics based on the high-resolution GPS dataset collected with five tagged animals in Table 1. It can be observed that daily motion duration can vary greatly for individual animals. For each animal, we evaluate the performance of trip duration prediction based on individual statistics. Specifically, we calculate the prediction error of a trip as the difference between the trip duration and the mean duration of all other trips in the dataset for the same animal.

Prediction Error. To examine the behaviour of our flight time prediction algorithm more closely, we adopt the mean absolute prediction error (MAPE). The MAPE metric calculates the absolute error between the estimated total flight time and the actual flight time at any time of the day, averaged over N trips:

$$\hat{T}_{motion}(t) = t_{motion}(t) + \Delta t_{motion}(t)$$

$$MAPE(t) = \frac{100}{N} \sum \left| \frac{T_{motion} - \hat{T}_{motion}(t)}{T_{motion}} \right|$$

At any time t , $t_{motion}(t)$ is the known, past movement duration, $\Delta t_{motion}(t)$ the predicted remaining motion duration and T_{motion} is the total movement duration of the current trip.

Figure 9 shows the MAPE metric for activity duration averaged over 48 trajectories of our data set, after removing three outliers. The flight duration prediction is based on experiences from previous trips and if the current trip differs strongly from previous trips, the prediction can fail. One example could be when an individual moves farther from the roosting camp than usual to access more remote sources of food. A more sophisticated algorithm might help to mitigate the effect of such outliers on prediction performance.

Animal identifier (number of trips)		1 (10)	2 (8)	3 (9)	4 (18)	5 (6)	All (51)
Motion duration/night [minutes]	Minimum	20.3	52.4	43.7	8.7	50.3	8.7
	Average	48.2	73.1	57.2	16.6	69.1	52.8
	Maximum	110.5	109.0	64.6	23.9	95.0	110.5
Avg. prediction error [minutes]	Individual	31.9	14.7	6.2	3.8	18.3	15.0
	Population	28.7	28.7	12.7	29.0	24.6	24.7

TABLE 1: Actual and predicted motion duration for 5 animals and 51 daily trajectories in the flying foxes dataset.




Fig. 9: Mean absolute prediction error (MAPE) for energy harvesting and movement duration.

However we only observed deviations in 3 out of 51 trips, indicating that the proposed method works in a majority of cases. The prediction error of total flight time at the beginning of the observation period is around 33% and decreases afterwards. A small prediction error remains in cases where the animal has changed its roosting location, which cannot be foreseen with the current mobility prediction algorithm.

7.2 Energy Prediction

To evaluate the performance of our energy prediction algorithm based on an exponential moving average (EWMA) filter (see Section 4), we use the MAPE, similar to the previous evaluation of total flight time. In this case, the MAPE metric calculates the absolute error between the estimated total energy budget and the actual total energy budget for each time slot i averaged over N trips:

$$\text{MAPE}(t) = \frac{100}{N} \sum \left| \frac{\sum_i E_H(i) - (E_H(t) + \hat{E}_{H,tot}(t))}{\sum_i E_H(i)} \right|$$

Figure 9 shows the mean absolute prediction error over all trajectories of our data set. It can be observed that the prediction error of the total available energy budget at the beginning of the observation period is around 15% and that the uncertainty will decrease later on.

7.3 Energy Consumption

We report the energy consumption of the individual system components in Table 2 when executing the different tracking strategies. The system baseline consumption includes the task scheduler, which periodically wakes up the microcontroller from sleep mode, as well as additional permanent consumers (e.g., voltage regulators, disabled sensors in sleep mode, leakage currents). The motion detection module requires the inertial sensors to be powered on continuously to

detect activity of the tracked object and signal a hardware interrupt to the microcontroller, which accounts for around 30% of the total energy consumption. Furthermore, the GPS receiver consumes roughly 6% of the energy while in standby mode, while active GPS tracking accounts for 14-15% of the total energy budget.

7.4 Emulation with Real-World GPS Traces

We use the previously described 51 empirical GPS traces from 5 free living flying foxes to emulate the behaviour of our scheduling framework with different tracking strategies under identical realistic energy constraints. As the magnetometer readings used as input by the information-based algorithm are modelled with a random jitter in our simulation framework (see Section 6), we calculate the mean results of 10 simulation runs for each trajectory.

Metrics. We focus on two metrics when evaluating tracking strategies, namely the average tracking error and the power consumption. The chosen sampling strategy determines when a GPS position sample is taken. The location of the tracked object between GPS sampling points is interpolated. Given a sampling strategy, the average tracking error metric is used to measure the accuracy of the obtained trajectories, and the power consumption characterizes its energy efficiency. The instantaneous tracking error $e[t_i]$ at the discrete time instant t_i is defined as the Euclidean distance between the ground truth location (x, y) and our estimated location (\hat{x}, \hat{y}) . We then calculate the average tracking error e_{avg} in the interval $[0, T]$ as follows:

$$e_{avg} = \frac{1}{T} \sum_{i=0}^T e[t_i]. \quad (11)$$

Energy Prediction Modes. In our emulation scenario, we consider the following two scenarios for energy harvesting:

- **Separated harvesting/tracking:** Separate periods of harvesting (day) and tracking (night), such as in the flying foxes application, results in a fixed energy budget for the tracking period.
- **EWMA-based prediction:** This scenario overlays harvesting and tracking synthetically and sets an energy budget for the tracking phase that is based on a EWMA-based prediction of harvested energy (see Section 4).

7.5 Tracking Performance

We evaluate our framework with the following GPS sampling strategies: *static motion-based* (Strategy #1), *adaptive motion-based* (Strategy #2), and *adaptive information-based* (Strategy #4). Furthermore, we compare the performance

Algorithm	Baseline (P=0.54 mW)		GPS tracking (P=141 mW)		GPS standby (P=0.07 mW)		Inertial (P=0.33 mW)	
	Duty Cycle	Energy	Duty Cycle	Energy	Duty Cycle	Energy	Duty Cycle	Energy
#1: static-motion	100%	49.02%	0.12%	15.04%	99.88%	5.99%	100%	29.95%
#2: adaptive-motion	100%	49.02%	0.12%	15.03%	99.88%	5.99%	100%	29.96%
#3: static-information	100%	49.58%	0.11%	14.06%	99.89%	6.06%	100%	30.30%
#4: adaptive-information	100%	49.42%	0.11%	14.34%	99.89%	6.04%	100%	30.20%
#5: offline	100%	49.20%	0.11%	14.72%	99.89%	6.01%	100%	30.07%

TABLE 2: Energy breakdown by system components for the different tracking strategies.

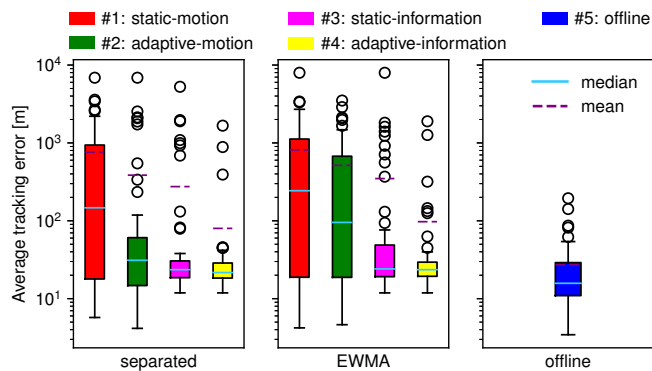


Fig. 10: Tukey boxplot of the average daily tracking error for separated harvesting and tracking, EWMA-based energy prediction, and the offline scenario.

of these algorithms with the *static information-based* strategy (Strategy #3), which is conceptually similar to the heading-aware strategy of EnTracked_T [18], and an optimal *offline* schedule calculated with full access to the trajectory (Strategy #5). Figure 10 shows the Tukey boxplot of the average daily tracking error for the four online strategies (see Section 5) with EWMA-based energy prediction, separated harvesting and tracking phases, and the optimal offline strategy with perfect energy prediction. The boxplot exposes the non-normal distribution of daily average tracking errors, highlighting the impact of variations in trace length and shape on the mean tracking error.

Motion-based GPS sampling. The *static motion-based* strategy (Strategy #1), which sets a fixed duty cycle at the beginning of the observation period based on a population-based estimate for movement duration, exhibits a median error of 146.9 m for separated energy harvesting and tracking phases and of 243.3 m for EWMA-based energy prediction. Employing online estimates for remaining movement duration, as in the *adaptive motion-based* strategy (Strategy #3), results in a smaller median tracking error of 31.1 m for separated harvesting and tracking phases, and in a tracking error of 95.3 m for EWMA-based energy prediction.

Information-based GPS sampling. While the *optimal offline* strategy (Strategy #5) expectedly has the lowest median (15.8 m) and the smallest spread of daily average tracking errors, the *adaptive information-based* sampling strategy (Strategy #4) yields a significantly lower mean daily tracking error with less outliers than the other online approaches. Since the adaptive information-based approach adapts the dead-reckoning error threshold to both energy and mobility forecasts in real-time and tunes its sampling decisions on

trajectory features rather than a fixed duty cycle setting, it is able to vary its performance to instantaneous changes in both movement and energy profiles.

Although the *static information-based* strategy (Strategy #3) achieves the next lowest median tracking error (separated: 23.6 m, EWMA: 24.1 m) after the adaptive information-based approach (separated: 21.6 m, EWMA: 23.6 m), daily tracking errors still have high variation due to the fixed error threshold. This strategy performs well for days where the movement duration and energy budget match well with the fixed error threshold. Otherwise, the threshold is mismatched due to longer than expected movement or less than expected energy availability, resulting in higher daily tracking error in many instances.

Consider the separated energy prediction scenario in Figure 10, which provides a priori definition of the energy budget. To highlight the performance gain of the proposed *adaptive information-based* strategy, we calculate the reduction of the median tracking over a comparison algorithm as $G = \frac{e_{comp} - e_{adaptive-information}}{e_{comp}}$. The reduction of the median tracking error by 8.2%-85.2% over the other online approaches can be attributed to linking the dynamic movement prediction component to the adaptive error threshold $R(t)$ in dead reckoning.

7.6 Prediction of Harvested Energy

The EWMA scenario highlights the combined effect of considering both real-time movement and energy predictions, further decreasing the median error of adaptive information-based sampling by 2.2% compared to static information-based, by 75.2% compared to adaptive motion-based, and by 90.3% compared to static motion-based.

The adaptive motion-based strategy performance degrades in the EWMA prediction scenario as it directly computes a sampling duty cycle based on movement and energy forecasts, which both include uncertainty. Although the same forecasts are also used by the adaptive information-based strategy to compute the dynamic error threshold, employing dead-reckoning to estimate the *instantaneous* trajectory error to decide on sampling yields significant benefits.

Interestingly, EWMA prediction yields a slightly higher median error for adaptive information-based sampling than for the scenario with separated harvesting and tracking. We attribute this effect to hidden constructive interaction between the uncertainties in mobility and energy prediction. Furthermore, the variance in the daily tracking error for information-based sampling in the EWMA scenario is slightly higher than for the separated scenario, highlighting the effect of uncertainty in the energy budget on the tracking performance.




Fig. 11: Daily average tracking error versus energy carry-over to the following day without and with EWMA-based prediction of energy harvesting.

Benefits of Energy Prediction. To shed further light on the effects of energy prediction on the tracking performance, we compare the adaptive information-based strategy in a scenario without and with EWMA-based energy prediction. Figure 11 shows the average error for all 51 traces against the remaining energy at the end of each day. In both cases, energy harvesting traces are overlaid with tracking activity in time synthetically.

Without a mechanism to predict the amount of harvested energy, information-based sampling initially has a very small available energy budget, based on what has been harvested since the beginning of the simulation, leading to higher tracking errors and an underspending on energy. As time progresses, the early energy savings translate into a lower error threshold and higher accuracy tracking for trajectory segments with significant information content. With EWMA-based energy prediction, the information-based strategy uses its energy forecast effectively to set lower tracking error thresholds early on, leading to relatively small average tracking errors for 49 out of the 51 traces. It also results in a smaller positive energy carryover than without prediction, indicating it is able to better utilize the harvested energy.

7.7 Benefits of Information-based GPS Sampling

We conclude that adaptive information-based sampling outperforms other online strategies as it adapts gracefully to uncertainty in harvested energy predictions when the tracking activity overlaps with the energy harvesting period. Figure 12 shows the performance of the adaptive information-based strategy for the different energy harvesting scenarios. As a benchmark, the optimal offline tracking strategy, which has full visibility into future movement and energy input, yields an average tracking error of 27.3 m. In comparison, adaptive information-based sampling without energy prediction leads to an average error of 1000.9 m. With EWMA prediction, information-based sampling reduces the mean tracking error by 90.3% to 97.4 m compared to the scenario without prediction. When the energy budget is deterministic, such as in the scenario with separated harvesting and tracking periods, information-based sampling achieves an average tracking error of around 79.9 m. Thus, the uncertainty in harvested energy prediction leads to an increase of the mean error by 21.9% over deterministic energy prediction, yet still achieves roughly a nine-fold reduction in error compared to the scenario without prediction. In terms of




Fig. 12: Tracking performance of the adaptive information-based strategy for the different energy prediction modes. Error bars represent one standard deviation.

median performance, adaptive information-based sampling achieves a daily average tracking error of 23.6 m (EWMA) and 21.6 m (separated), compared to 15.8 m with the optimal offline strategy.

8 RELATED WORK

Localisation of mobile devices have been studied extensively in the past motivated by applications in navigation, tracking and rescue. Location awareness is also a key requirement for emerging mobile applications such as location-based search, turn-by turn navigation, social media, or participatory sensing [9], [10], [25].

Wildlife tracking. GPS technology is being used to track increasingly smaller animals, thanks to miniaturisation of GPS modules. However, certain wildlife tracking applications cannot rely on GPS due to the application scenario, *e.g.*, tracking badgers underground [8], and thus employ alternative localization techniques such as RFID tags or radio contact logging [4], [21]. With our focus on tracking in remote areas with limited cellular network coverage, we rely on GPS only to determine the current position of tagged animals. Although the power consumption of GPS receivers has been continually improving, it remains a major part of the overall energy budget of mobile devices. Liu et al. [22] proposed to reduce the power consumption by offloading the GPS signal processing to the cloud. However, this requires to store and transfer raw data, which creates another bottleneck for resource-constrained applications.

GPS Duty-Cycling. Continuous tracking applications exhibit a trade-off between accuracy of localization and the amount of energy spent to do so. Therefore, mechanisms to adapt the duty cycle of the GPS receiver and other localization modalities have been extensively studied in previous work [9], [35]. Paek et al. [26] use the location-time history of a user and a blacklist based on cell tower signal strength to decide when to activate the GPS and to avoid unavailability of GPS indoors. SensLoc [17] uses algorithms to detect the user context, such as frequently visited places, and movement detection to adapt the GPS duty cycle. Jurdak et al. [13] propose to use short-range radio contact-logging to bound the position uncertainty when duty-cycling the GPS. SmartDC [6] provides adaptive GPS duty-cycling using prediction mechanisms for regularities in user mobility. eNav [12] employs dead-reckoning using the acceleration in smartphones to enable the GPS only when the user reaches the next navigation waypoint.

Energy Management. Adaptive localization algorithms, such as EnLoc [7], EnTracked [2], [18], [19], and a-Loc [20], dynamically select the localization modality that provides the most energy-efficient position estimate within provided uncertainty bounds. Another approach to duty-cycling the GPS is to rely on low-power sensors such as accelerometer or compass to trigger GPS sampling only when the device is moving [17], [26] or when it changed its heading [18]. Jigsaw [23] manages battery life of mobile phones by balancing power consumption of sensing processes with quality requirements of the application and user context.

Unlike most adaptive localization algorithms, we focus on the inverse problem of determining the most accurate GPS sampling strategy given energy constraints. Consequently, instead of sampling GPS location when the uncertainty of the current location estimate increases above a certain limit, we constrain the GPS sampling subject to the available energy. Our approach will not strictly enforce that all location estimates are within the error bounds, but will strive to achieve best performance overall given the energy constraints.

Energy-aware process scheduling frameworks have been proposed to adjust the long-term system performance based on prediction of harvested energy [3]. While such frameworks consider long-term operation of solar energy harvesting systems over years at a guaranteed minimum performance level, the goal of this paper is to manage short-term energy usage under daily varying harvesting and mobility.

9 CONCLUSIONS AND DISCUSSION

This paper has presented an energy- and mobility-aware scheduling framework for perpetual tracking applications. At the core of this framework is an adaptive information-based tracking strategy that maximizes the tracking information return subject to a given energy budget. Our strategy uses historical motion statistics and distance from the destination to forecast activity duration to strategically select the best GPS sampling policy. It also predicts energy availability to set a virtual energy budget for a future sampling period. The algorithm then updates its forecast and adapts its acceptable error threshold in response to mobility forecast and energy availability. We evaluated our approach on empirical traces from wild flying foxes and showed that our approach can significantly increase positioning accuracy for a dynamic energy budget.

While our work was motivated by animal tracking, we believe that our proposed energy- and mobility-aware scheduling framework is applicable to a broad range of tracking applications. For instance, most living beings, including humans, exhibit very strong preferential return behavior to one or two locations, typically home and work or study [11], [15], and exhibit time-of-day correlations with the most likely destination (e.g., going to work at 8 am or returning home at 5 pm [1]). As such, the concepts for scheduling of GPS by predicting motion duration on the basis of previous days, population-based statistics, or current distance from the most likely destination are transferable to tracking people and indeed many other living species [30]. A related area is asset tracking in logistics where ground assets often need to be tracked over long

durations. Using historic motion patterns, whether for the asset group or individual assets, and energy availability to schedule position sampling can provide more accurate asset tracking.

An interesting direction for future work is to improve the prediction of motion duration and energy harvesting for more accurate scheduling of GPS samples by applying more advanced algorithms or a combination of several methods [29]. For the flying fox application, it will be important to detect if the animal is flying to a previously unknown location. In addition, temperature and solar exposure inputs can be used as proxies for weather conditions that may better predict the extent of movement and energy demand for a single forecast period. Finally, it will be important to evaluate the information-based tracking strategy for other application scenarios, such as for tracking people, different animal species, or movable assets in logistics. We expect our work to be an important step towards highly autonomous tracking with energy harvesting that delivers detailed trajectory data near-perpetually for a broad range of applications.

10 ACKNOWLEDGMENTS

This work was supported by CSIRO's Sensor and Sensor Networks Transformation Capability Platform.

REFERENCES

- [1] A. Bamis and A. Savvides. Lightweight Extraction of Frequent Spatio-Temporal Activities from GPS Traces. In *Proc. RTSS*, 2010.
- [2] S. Bhattacharya, H. Blunck, M. B. Kjaergaard, and P. Nurmi. Robust and energy-efficient trajectory tracking for mobile devices. *IEEE Trans. Mobile Comput.*, 14(2):430–443, Feb 2015.
- [3] B. Buchli, F. Sutton, J. Beutel, and L. Thiele. Dynamic Power Management for Long-term Energy Neutral Operation of Solar Energy Harvesting Systems. In *Proc. SenSys*, 2014.
- [4] B. Cassens, S. Ripperger, M. Hierold, F. Mayer, and R. Kapitza. Automated encounter detection for animal-borne sensor nodes. In *Proc. EWSN*, 2017.
- [5] W. S. Chan and F. Chin. Approximation of polygonal curves with minimum number of line segments or minimum error. *Int. J. Comput. Geom. Appl.*, 6(01):59–77, 1996.
- [6] Y. Chon, E. Talipov, H. Shin, and H. Cha. Smartdc: Mobility prediction-based adaptive duty cycling for everyday location monitoring. *IEEE Trans. Mobile Comput.*, 13(3):512–525, 2014.
- [7] I. Constandache, S. Gaonkar, M. Saylor, R. R. Choudhury, and L. Cox. EnLoc: Energy-Efficient Localization for Mobile Phones. In *Proc. INFOCOM*, 2009.
- [8] V. Dyo, S. A. Ellwood, D. W. Macdonald, A. Markham, C. Mascolo, B. Pásztor, S. Scellato, N. Trigoni, R. Wohlers, and K. Yousef. Evolution and Sustainability of a Wildlife Monitoring Sensor Network. In *Proc. SenSys*, 2010.
- [9] S. B. Eisenman, E. Miluzzo, N. D. Lane, R. A. Peterson, G.-S. Ahn, and A. T. Campbell. Bikenet: A mobile sensing system for cyclist experience mapping. *ACM Trans. Sen. Netw.*, 6(1):6:1–6:39, 2010.
- [10] J. Eriksson, L. Girod, B. Hull, R. Newton, S. Madden, and H. Balakrishnan. The Pothole Patrol: Using a Mobile Sensor Network for Road Surface Monitoring. In *Proc. MobiSys*, 2008.
- [11] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196), 2008.
- [12] S. Hu, L. Su, S. Li, S. Wang, C. Pan, S. Gu, M. T. Al Amin, H. Liu, S. Nath, R. R. Choudhury, and T. F. Abdelzaher. Experiences with eNav: A Low-power Vehicular Navigation System. In *Proc. UbiComp*, 2015.
- [13] R. Jurdak, P. Corke, D. Dharman, and G. Salagnac. Adaptive GPS Duty Cycling and Radio Ranging for Energy-Efficient Localization. In *Proc. SenSys*, 2010.
- [14] R. Jurdak, P. Sommer, B. Kusy, N. Kottege, C. Crossman, A. Mckeown, and D. Westcott. Camazotz: Multimodal Activity-based GPS Sampling. In *Proc. IPSN*, 2013.

- [15] R. Jurdak, K. Zhao, J. Liu, M. AbouJaoude, M. Cameron, and D. Newth. Understanding human mobility from twitter. *PLOS ONE*, 10(7):1–16, 07 2015.
- [16] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava. Power management in energy harvesting sensor networks. *ACM Trans. on Embedded Computing Sys.*, 6:32, 2007.
- [17] D. H. Kim, Y. Kim, D. Estrin, and M. B. Srivastava. SensLoc: Sensing Everyday Places and Paths using Less Energy. In *Proc. SenSys*, 2010.
- [18] M. B. Kjaergaard, S. Bhattacharya, H. Blunck, and P. Nurmi. Energy-efficient trajectory tracking for mobile devices. In *Proc. MobiSys*, 2011.
- [19] M. B. Kjaergaard, J. Langdal, T. Godsk, and T. Toftkjaer. EnTracked: Energy-Efficient Robust Position Tracking for Mobile Devices. In *Proc. MobiSys*, 2009.
- [20] K. Lin, A. Kansal, D. Lymberopoulos, and F. Zhao. Energy-Accuracy Trade-Off for Continuous Mobile Device Location. In *Proc. MobiSys*, 2010.
- [21] A. Lindgren, C. Mascolo, M. Lonergan, and B. McConnell. Seal-2-Seal: A Delay-Tolerant Protocol for Contact Logging in Wildlife Monitoring Sensor Networks. In *Proc. MASS*, 2008.
- [22] J. Liu, B. Priyantha, T. Hart, H. S. Ramos, A. A. F. Loureiro, and Q. Wang. Energy Efficient GPS Sensing with Cloud Offloading. In *Proc. SenSys*, 2012.
- [23] H. Lu, J. Yang, Z. Liu, N. D. Lane, T. Choudhury, and A. T. Campbell. The Jigsaw Continuous Sensing Engine for Mobile Phone Applications. In *Proc. SenSys*, 2010.
- [24] A. McKeown and D. Westcott. Assessing the accuracy of small satellite transmitters on free-living flying-foxes. *Australian Ecology*, 37, 2012.
- [25] P. Mohan, V. N. Padmanabhan, and R. Ramjee. Nericell: Rich Monitoring of Road and Traffic Conditions using Mobile Smartphones. In *Proc. SenSys*, 2008.
- [26] J. Paek, J. Kim, and R. Govindan. Energy-Efficient Rate-Adaptive GPS-based Positioning for Smartphones. In *Proc. MobiSys*, 2010.
- [27] F. Papi. *Animal homing*. Springer Science & Business Media, 2012.
- [28] J. Ryder, B. Longstaff, S. Reddy, and D. Estrin. Ambulation: A Tool for Monitoring Mobility Patterns over Time Using Mobile Phones. In *Proc. CSE*, 2009.
- [29] F. C. Sangogboye and M. B. Kjaergaard. Promt: predicting occupancy presence in multiple resolution with time-shift agnostic classification. *Computer Science - Research and Development*, 33(1):105–115, 2018.
- [30] S. Scellato, M. Musolesi, C. Mascolo, V. Latora, and A. Campbell. NextPlace: A Spatio-temporal Prediction Framework for Pervasive Systems. In *Proc. Pervasive*. 2011.
- [31] P. Sommer, B. Kusy, and R. Jurdak. Energy Estimation for Long-Term Tracking Applications. In *Proc. ENSys*, 2013.
- [32] P. Sommer, J. Liu, K. Zhao, B. Kusy, R. Jurdak, A. McKeown, and D. Westcott. Information Bang for the Energy Buck: Towards Energy- and Mobility-Aware Tracking. In *Proc. EWSN*, 2016.
- [33] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási. Limits of predictability in human mobility. *Science*, 327(5968):1018–1021, 2010.
- [34] A. Stisen, H. Blunck, S. Bhattacharya, T. S. Prentow, M. B. Kjaergaard, A. Dey, T. Sonne, and M. M. Jensen. Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition. In *Proc. Sensys*, 2015.
- [35] P. Zhang, C. M. Sadler, S. A. Lyon, and M. Martonosi. Hardware Design Experiences in ZebraNet. In *Proc. SenSys*, 2004.




Kai Geissdoerfer received his B.Sc. and M.Sc. in Electrical Engineering from Technische Universität Berlin in 2015 and 2017 respectively. During his master studies he's been working as a student research assistant with the Telecommunication Networks Group at TU Berlin and the Distributed Sensing Systems Group at CSIRO, Australia. He's now a PhD student at Technische Universität Dresden. His main interests are embedded communication systems, energy harvesting and battery-less sensing systems.




Raja Jurdak is a Principal Research Scientist and leads the Distributed Sensing Systems Group at CSIRO. He received the B.E. degree from the American University of Beirut in 2000, and the MSc and Ph.D. degrees from the University of California Irvine in 2001 and 2005 respectively. His current research interests are around energy and mobility in sensor networks. He is an Adjunct Professor at Macquarie University, and Adjunct Associate Professor at UNSW Australia.



Branislav Kusy is a Principal Research Scientist at CSIRO and an Adjunct Associate Professor at the University of Queensland. He obtained his PhD degree from Vanderbilt University in 2007. His research interests include embedded networked systems and algorithms, with a specific focus on in-situ sensing and intelligence to allow IoT systems to respond to real-world challenges in an autonomous way.



Jiajun Liu is an Associate Professor at Renmin University of China. He received his PhD and BEng from The University of Queensland, Australia and from Nanjing University, China in 2012 and 2006 respectively. Before joining Renmin University he has been a Postdoctoral Fellow at the CSIRO of Australia from 2012 to 2015. From 2006 to 2008 he also worked as a Researcher/Software Engineer for IBM China R&D Labs. His main research interests are in multimedia and spatio-temporal data management and mining.



Kun Zhao is a senior data scientist in Suncorp Group Australia. He was a OCE Postdoctoral Fellow at CSIRO Data 61 Australia. He obtained his PhD and BS in Physics from Northeastern University USA and Nankai University China respectively. His research interests include data science, complex networks and mobility data analysis in sensor networks.



Adam McKeown is a Movement and Population Ecologist working for CSIRO, Australia on a variety of projects. He has worked on a wide range of wildlife over the last 25 years, with a large amount of the last 10 years spent tracking flying-foxes around Australia.



Philipp Sommer is a Senior Research Scientist at ABB Corporate Research in Switzerland. He received MSc and PhD degrees in Electrical Engineering from the ETH Zurich in 2007 and 2011 respectively. He has been a postdoctoral fellow at the Autonomous System lab at CSIRO, Australia between Sept. 2011 and Oct. 2014. His research interests include a broad range of topics in the field of wireless sensor networks, distributed computing and embedded systems.



David Westcott is a Senior Principal Research Scientist at CSIRO, Australia. He is a zoologist whose research focuses on the ecology and behaviour of tropical species. He received the MSc and Ph.D. degrees from the University of British Columbia in 1991 and 1995 respectively.