DIET: A Dynamic Energy Management Approach for Wearable Health Monitoring Devices

Nuzhat Yamin, Ganapati Bhat and Janardhan Rao Doppa
School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, 99163
Email: {nuzhat.yamin, ganapati.bhat, jana.doppa}@wsu.edu

***Abstract*—Wearable devices are becoming increasingly popular for health and activity monitoring applications. These devices typically include small rechargeable batteries to improve user comfort. However, the small battery capacity leads to limited operating life, requiring frequent recharging. Recent research has proposed energy harvesting using light and user motion to improve the lifetime of wearable devices. Most energy harvesting approaches assume that the placement of the energy harvesting device and sensors required for health monitoring are the same. However, this assumption does not hold for several real-world applications. For example, motion energy harvesting using piezoelectric sensors is limited to the knees and elbows, while a sensor for heart rate monitoring must be placed on the chest for optimal performance. To address this challenge, we propose a novel dynamic energy management approach referred to as DIET for wearable health applications enabled by multiple sensors and energy harvesting devices. The key idea behind DIET is to harvest energy from multiple sources and optimally allocate it to each sensor using a lightweight optimization algorithm such that the overall utility for applications is maximized. Experiments on real-world data from four users over 30 days show that the DIET approach achieves utility within 10% of an offline Oracle.**

I. Introduction

Wearable devices have the potential to transform health and activity monitoring by providing accurate and objective measures of users' health for enabling personalized treatments [1, 2]. Despite this potential, the adoption of wearable devices has been limited because of low battery lifetime and the need for frequent recharging [3]. Energy harvesting (EH) has been heralded as an effective method to increase the battery life of wearable devices. EH methods harvest energy from ambient sources, such as light and motion, to autonomously recharge the battery in wearable devices [4, 5]. Studies show that ambient light can provide up to 0.1 mW/cm^2 (indoor) – 100 mW/cm^2 (outdoor) power density [4], while user motion can generate up to 7.8 μJ per step when walking [6]. Ambient light and user motion are ideal sources of energy for wearable devices since they can be easily integrated as patches on clothes [5]. The wearable device can learn the activity patterns of the user to optimize both EH and consumption to meet user needs.

Energy harvesting and management algorithms for wearable devices typically assume that consumption and harvesting take place at the same location on the human body [4, 5]. However, this assumption is not true for many applications that require multiple sensors or devices for monitoring the user's health. For instance, the activity recognition approach in [7] places inertial motion unit on the chest, whereas motion energy is typically harvested from joints. Therefore, there is a strong need for algorithms that consider the energy harvested from multiple sources on the body and dynamically allocate it to a network of wearable sensors on the body.

This paper presents a general approach for **d**ynam**i**c **e**nergy managemen**t** (DIET) in wearable devices with multiple energy harvesters and consumers. The overall goal of DIET is to enable *recharge-free* operation of the system over a period of time while maximizing the utility to the user. The key challenge is that energy available from ambient sources is highly stochastic, and an effective algorithm should account for this inherent uncertainty. Therefore, we propose a novel *robust optimization* formulation to explicitly reason about the worst-case uncertainty and direct the energy available in each harvester to a battery or a sensor node as a function of the application requirements and target energy constraints. The primary components of the DIET approach are shown in Figure 1. We start by dividing each day into T equal intervals. At the beginning of each interval, the sensor nodes estimate their future energy requirements and send their requests to the harvesters. Similarly, each harvester estimates future energy availability and uncertainty to determine the energy available for distribution in the current interval such that the total energy consumption over a day equals the harvested energy. Using the energy available to distribute, each harvester solves a lightweight convex optimization problem to fulfill the requests by the sensors while maximizing the total system utility. The solutions obtained at the beginning of the interval are not optimal due to uncertainty in harvested and consumed energy. Therefore, at the end of each interval, we apply corrections to the energy allocations so that system constraints are met.

We evaluate the DIET approach on a wearable system with four sensor nodes and five EH devices with real-world activity data from four users over thirty days and calculate the overall utility to the user. We also implement a baseline approach presented in [8] and an offline Oracle that provides the optimal solutions. Comparisons to the baseline show that DIET achieves 189% higher utility for all users. Moreover, the utility achieved by DIET is within 10% of the optimal with less than 0.0003% overhead. The specific contributions of this paper are summarized below:

- A novel *robust optimization* formulation to manage the energy in a system with multiple harvesters and sensors
- A lightweight algorithm to distribute the harvested energy across multiple sensor nodes as a function of the user activity while maximizing the total system utility.
- Experimental evaluation on a real-world benchmark with thirty days of operation to show the robustness of DIET.

978-3-9819263-6-1/DATE22/©2022 EDAA

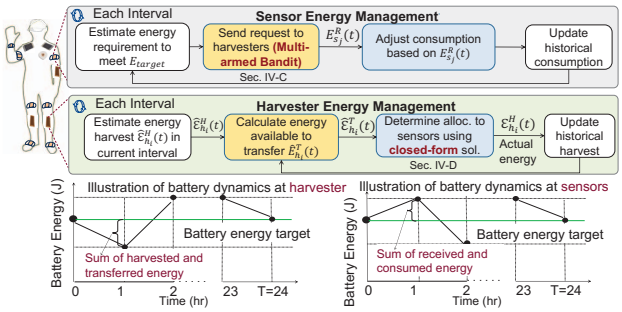

Fig. 1: High-level overview of the DIET approach.

II. Related Work

Wearable devices have gained prominence recently due to their applicability in health applications such as Parkinson's disease diagnosis and rehabilitation [2, 9]. These devices must have a small form-factor for user comfort [3]. As a result, the battery capacity of wearable devices is severely constrained, which leads to a low operating lifetime [10].

EH has been proposed as an effective method to increase the lifetime of wearable devices [4–6, 11]. EH from ambient sources must be managed optimally to ensure that sufficient energy is available for the applications. For effective management, it is important to predict the energy available in the future [11, 12]. In a wearable device, the future energy depends on the activities and location of the user. Therefore, activity prediction is crucial to have an effective energy management system. Several approaches have been proposed to predict future activities of the user [13, 14]. In particular, Bayesian Networks (BNets) have been used extensively in activity prediction. Following the success of BNets in prior work, we adopt them in the DIET approach to predict future activities.

Several algorithms have been proposed for runtime management of EH devices [8, 11, 15–17]. The work in [8] proposes an energy allocation scheme for sensor nodes with multi-source EH. The approach first assigns an energy budget from multiple EH sources and then uses it to determine the energy allocation of tasks in the sensor node. However, this approach suffers from two key limitations. First, the energy allocation decisions are made by a central manager, which is not feasible in a real-world scenario with multiple harvesters and sensors. Second, it does not consider the maximization of the application utility, which is one of the critical components for the quality of service. Furthermore, future application needs and EH potential are not taken into account. To address these limitations, the DIET approach uses a finite horizon stochastic optimization to maximize the utility of the application. Specifically, it obtains energy using multiple harvesters on the body and allocates it to the sensors while maximizing the utility.

III. Overview and Preliminaries

We consider a system with M sensing nodes with processors and N energy harvesters mounted on the body, as shown in Figure 1. The set of sensors and harvesters is denoted by $\mathcal{S}$ and $\mathcal{H}$, respectively. The M sensors are used to monitor the activities of the user and report them to a host device, such as a smartphone, for long-term analysis. The energy harvesters obtain energy from ambient light and user motion to power the sensors used in the application. We also consider that all sensors and harvesters have dedicated batteries to store energy.

User activities and EH typically follow a pattern with day-to-day variations. To account for these variations, we perform the energy management with a one-day horizon. The one-day horizon is further divided into T intervals of equal length to account for intra-day variations in activity and EH patterns. The set of intervals in a horizon is denoted by $\mathcal{T}$. With this setup, our goal is to ensure that the sensors operate on harvested energy without the need to recharge their batteries manually.

Battery dynamics at each harvester: The battery energy of a harvester $h_i \in \mathcal{H}$ at the beginning of interval t is denoted as $E^B_{h_i}(t)$ $(t \in \mathcal{T})$. In each interval, the harvested energy at harvester h_i is given by the random variable $\mathcal{E}^H_{h_i}(t)$, while the energy transferred by the harvester to sensor s_j is given by $E^T_{h_i,s_j}(t)$ $(s_j \in \mathcal{S}, h_i \in \mathcal{H})$. Using this, we can express the battery level at the beginning of interval $t+1$ as:

$$E^B_{h_i}(t+1) = E^B_{h_i}(t) + \eta \mathcal{E}^H_{h_i}(t) - \sum_{j=1}^{M} E^T_{h_i,s_j}(t), h_i \in \mathcal{H}, t \in \mathcal{T}$$

where η is the EH efficiency. Note that the battery level at the end of a horizon (day) is equal to the battery level at the beginning of the next horizon, i.e. $E^B_{h_i}(T) = E^B_{h_i}(0)$. To achieve operation on harvested energy, sufficient energy must be present in the battery at the end of the horizon so that the next horizon can start without manual recharging. Therefore, we introduce a target energy constraint on each harvester as: $E^B_{h_i}(T) = E^B_{h_i}(0) \geq E_T,\ h_i \in \mathcal{H}$, where E_T is the battery target.

Battery dynamics at each sensor: Let us assume that the battery energy of sensor s_j at the beginning of interval t is $E^B_{s_j}(t)$ $(s_j \in \mathcal{S}, t \in \mathcal{T})$. At the beginning of each interval, the sensors receive energy from the harvesters based on the requests made by the sensors. We denote the total energy received by sensor s_j at interval t as $E^R_{s_j}(t)$. The energy consumed by the sensor in interval t is given by $E^c_{s_j}(t)$. Consequently, the battery energy at the beginning of the next interval is:

$$E^B_{s_j}(t+1) = E^B_{s_j}(t) + \eta E^R_{s_j}(t) - E^c_{s_j}(t), s_j \in \mathcal{S}, t \in \mathcal{T} \quad (1)$$

Similar to the harvesters, the battery level at the end of a horizon is equal to the energy at the beginning of the next horizon, i.e $E^B_{s_j}(T) = E^B_{s_j}(0)$. We also impose target energy constraints on the sensor battery levels as: $E^B_{s_j}(T) = E^B_{s_j}(0) \geq E_T,\ s_j \in \mathcal{S}$.

Application and utility function: DIET is applicable for any wearable application such as vital sign monitoring and activity monitoring. The energy required by the application in each interval is a function of the user activities. For instance, if the user is jogging, we must sample the sensors with high frequency to capture the activities and heart rate accurately. We model the change in sampling frequency of sensor s_j as a function of activity with duty ratio $\rho_{s_j}(t)$, i.e., the percentage of time in an interval that the sensors are active. We can calculate the energy required in each interval for each sensor as:

$$E^{Req}_{s_j}(t) = [\rho_{s_j}(t)[P^a_{s_j} + P^{com}_{s_j}] + (1 - \rho_{s_j}(t))P^{idle}_{s_j}]T_P \quad (2)$$

where T_P is the length of the interval, and $P^a_{s_j}$, $P^{com}_{s_j}$, and $P^{idle}_{s_j}$

are the active, communication, and idle power consumption. If the energy consumed by the sensor s_j in an interval t is equal to the required energy $E^{Req}_{s_j}(t)$, it means that the user gets the required quality of service (QoS). However, the QoS is reduced if the energy available to consume is lower than the required energy. We define a utility function to model the QoS as:

$$U_{s_j}(t) = \ln\left(\frac{E^c_{s_j}(t)}{E^{Req}_{s_j}(t)} * 100\right), \quad s_j \in \mathcal{S}, t \in \mathcal{T} \tag{3}$$

The utility function achieves maximum when the energy requirements for an interval are met.

IV. Robust Optimization Approach

A. Problem Formulation

The overall goal of DIET is to maximize the utility to the user over each horizon while satisfying the battery dynamics by reasoning about the uncertainty in EH. Therefore, we formulate a robust optimization problem as shown below:

$$\max \quad \mathcal{O} = \mathbb{E}\left[\sum_{t=0}^{T-1}\sum_{j=1}^{M} \ln\left(\frac{E^c_{s_j}(t)}{E^{Req}_{s_j}(t)} * 100\right)\right] \tag{4}$$

$$\text{s. t.} \quad E^B_{h_i}(t+1) = E^B_{h_i}(t) + \eta\mathcal{E}^H_{h_i}(t) - \sum_{j=0}^{M} E^T_{h_i,s_j}(t) \tag{5}$$

$$E^B_{s_j}(t+1) = E^B_{s_j}(t) + \eta E^R_{s_j}(t) - E^c_{s_j}(t) \tag{6}$$

$$E^B_{s_j}(t) \geq E_T, E^B_{h_i}(t) \geq E_T, h_i \in \mathcal{H}, s_j \in \mathcal{S} \tag{7}$$

$$E^c_{s_j}(t) \leq E^{Req}_{s_j}(t), \quad s_j \in \mathcal{S} \tag{8}$$

$$E^B_{s_j}(t) \geq 0, \quad E^B_{h_i}(t) \geq 0, h_i \in \mathcal{H}, s_j \in \mathcal{S} \tag{9}$$

The objective function maximizes the total utility of all sensors over the finite horizon. The first three constraints account for the battery dynamics described in the previous section. Equation 8 constrains the energy consumption in each interval to the energy required in the interval. Finally, Equation 9 requires that all battery levels be non-negative at all times.

Key Challenges: The robust optimization problem outlined above has a concave objective and linear constraints. Solving this problem optimally at runtime is infeasible for the following key reasons: **1)** true values of the uncertain EH for the current and future intervals within the finite horizon are unknown; **2)** wearable devices are constrained by computing and energy resources needed to use a commercial solver to obtain the solution. Therefore, DIET employs a two-stage approach to efficiently solve the problem at runtime by reasoning about the uncertainty in the EH. Our approach is guided by two key insights: 1) We can use historical data about user activity and energy harvest to design supervised learning algorithms to predict the future expected EH and the corresponding uncertainty. 2) The predicted values can then be used in a lightweight optimization problem at each energy harvester to determine the energy transfer to each sensor. Similarly, each sensor can decide its energy consumption based on the predicted activities and available energy. In the following, we detail these steps.

B. Machine-Learning Based Activity and Energy Prediction

The energy availability is a function of the user's activity, location, and the duration of the activity. Therefore, we develop an activity prediction model as the first step in DIET.

Activity Prediction: Activity prediction is a supervised learning problem where input features of the system state s, $\phi(s)$ are used to make predictions of the activities. We use the ground truth activity values A^* to train the supervised learning algorithm. Prior studies have shown that there exists a conditional dependence between consecutive activities of a user [13, 14]. Following this, we use the following features for the activity prediction algorithm: {Current activity (A_n), Current user location (L_n), Time of the day (τ_n) and Day of the week (d_n)}. These features are used to train a Dynamic Bayesian networks (DBN) for activity prediction. We use a DBN since they have been successfully used to predict activities [13] and they provide a lightweight solution to predict activities on the wearable device. The output of the DBN prediction is the future activity A_{n+1} at any given time instant n.

The final consideration in the activity prediction is to predict the duration of the next activity once it starts. To this end, we use a decision tree to obtain the duration of the next activity. The features used by the decision tree include the current activity, location, time, day, next activity, and next location.

Estimation of Future Energy Harvest: The DIET approach uses piezoelectric sensors and photovoltaic (PV) cells to power the wearable health application. The energy harvested by piezoelectric sensors depends on the level of movement of the joints where the sensor is attached [6]. Using this insight, we can assign an intensity for each user activity and use it to calculate the energy available from the piezoelectric sensors. The motion intensity is then combined with the piezoelectric energy model proposed in [6] to estimate the energy available in the future and make energy management decisions.

The energy harvested by PV cells depends on the irradiance of the light, which in turn depends on the location of the user. Therefore, we use the DBN described in the previous section to first predict the location of the user. The location is then used to predict the light irradiance in either outdoor or indoor conditions [18–20]. The irradiance is used along with the current-voltage characteristics of the employed PV cell and activity duration to estimate energy available in the future.

C. Energy Management at Sensor Nodes

The energy management problem at each sensor node involves estimating future energy requirements, sending energy requests to harvesters, and deciding the energy consumption in the current interval. We start by estimating the energy required in an interval t using the predicted activity, location, and the duration in Equation 2. However, the estimation of energy required in the current interval alone is not sufficient to satisfy the target energy requirements at the end of the interval. Therefore, we maintain an exponentially weighted moving average (EWMA) of energy requirements in each interval over all observed horizons. Specifically, the historical average for each interval $E_{hist,s_j}(t)(s_j \in \mathcal{S}, t \in \mathcal{T})$ in a sensor is given by:

$$E_{hist,s_j}(t) = \alpha\bar{E}_{hist,s_j}(t) + (1-\alpha)E^c_{s_j}(t) \tag{10}$$

where $\alpha < 1$ is a weighing factor and $\bar{E}_{hist,s_j}(t)$ is the EWMA of the consumption up to the current horizon. Using

the historical average, the energy required in the rest of the horizon, $\hat{E}_{s_j}(t)(s_j \in \mathcal{S}, t \in \mathcal{T})$, to maintain the battery target is:

$$\hat{E}_{s_j}(t) = E_{target} - \left[E^B_{s_j}(t) - \hat{E}^c_{s_j}(t) - \sum_{k=t+1}^{T-1} E_{hist,s_j}(k)\right] \quad (11)$$

If the value of $\hat{E}_{s_j}(t)$ is greater than zero, it means that the current battery level is insufficient to meet future energy needs. Therefore, the sensor must request energy from the harvesters.

It is crucial to choose the appropriate harvester when requesting energy since the energy availability in each harvester varies as a function of the user activity and requests the harvester receives. We model the problem of choosing the harvester for energy requests as a multi-armed bandit and use the ϵ-greedy policy for bandits to choose the best harvester in each interval [21]. Specifically, each sensor maintains a numerical preference for each harvester based on the energy received from the harvester in the past. Then, the harvester with the highest preference is selected with a probability of 1 - ϵ, while a random harvester is chosen with a probability of ϵ. The ϵ-greedy approach ensures that all harvesters are explored by the sensor. At the end of each interval, the harvester preferences are updated to reflect the energy received from the harvester.

D. Energy Management at Harvesters

The energy management problem at each harvester is to allocate energy to the sensors based on the available energy and requests from sensors. To this end, we first find the expected energy available in the current interval, $\hat{\mathcal{E}}^H_{h_i}$ using the activity and location predictions. Similar to the sensors, each harvester maintains EWMA of the historical energy harvest to account for the uncertainty in the future harvest, $\mathcal{E}^H_{hist,h_i}(t)(h_i \in \mathcal{H}, t \in \mathcal{T})$, in each interval of the finite horizon. The historical energy harvest is then used to determine the energy available to transfer while satisfying the target battery constraint as:

$$\hat{\mathcal{E}}^T_{h_i}(t) = E^B_{h_i}(t) + \hat{\mathcal{E}}^H_{h_i}(t) + \sum_{k=t+1}^{T-1} \mathcal{E}^H_{hist,h_j}(k) - E_T \quad (12)$$

where $\hat{\mathcal{E}}^T_{h_i}(t)$ is the energy available to allocate and $\hat{\mathcal{E}}^H_{h_i}(t)$ is the estimated energy in the current interval. The historical harvest is updated at the end of each interval based on the actual EH so that future allocations use the feedback from the actual EH.

The available energy $\hat{\mathcal{E}}^T_{h_i}(t)$ must be optimally allocated to the sensors such that the total utility of the system is maximized. As described in Equation 3, utility depends on the ratio of sensor energy allocation and the actual energy requirements in an interval. In turn, the actual energy consumption depends on the activity of the user and the sensors most critical to monitor the activity accurately. In order to model the importance of sensors, the harvester assigns a priority to each sensor $\gamma_{s_j}(A_{n+1})$ as a function of the predicted activity. Finally, each harvester $h_i \in \mathcal{H}$ solves the following optimization problem to determine the energy allocations for each sensor:

$$\max \quad \mathcal{O} = \sum_{j=1}^{M} \gamma_{s_j}(A_{n+1}) \ln\left(\frac{E^T_{h_i,s_j}(t)}{\hat{E}_{s_j}(t)}\right) \quad (13)$$

$$\text{s. t.} \quad \sum_{j=1}^{M} E^T_{h_i,s_j}(t) \leq \hat{\mathcal{E}}^T_{h_i}(t) \;\&\; 0 \leq E^T_{h_i,s_j}(t) \leq \hat{E}_{s_j}(t) \quad (14)$$

The first constraint specifies that the total energy allocation to sensors does not exceed the available energy, while the second constraint ensures that all energy allocations are non-negative and less than the sensor request.

Lightweight Solution: We consider two cases for solving the above convex optimization problem with linear constraints. In the first case, the harvester has sufficient energy to fulfill all the sensor requests. In this case, the solution is trivial, and the values of $E^T_{h_i,s_j}(t)$ are set to the respective requests. In the second case, the total available energy is lower than the requests. We obtain a closed-form solution to the second case by using the Karush-Kuhn-Tucker conditions [22]. Specifically, the energy transferred to each sensor is given by:

$$E^T_{h_i,s_j}(t) = \frac{\gamma_{s_j}(A_{n+1})}{\sum_{j=1}^{M} \gamma_{s_j}(A_{n+1})} \hat{E}^T_{h_i}(t) \quad (15)$$

The solution specifies that the energy transferred is determined as a function of the normalized priorities. Intuitively, this ensures that the most important sensor to the activities receives the highest energy, which in turn improves the application utility. At the same time, lower priority sensors are not starved of energy so that they contribute to the application utility as well. We omit the proof of the solution due to space constraints. In summary, the closed-form solution ensures that the energy management has negligible overhead.

V. Experimental Evaluation

A. Experimental Setup

Wearable device model: We use a health monitoring device with four sensors and five energy harvesters. Specifically, we consider Texas Instruments (TI) CC2652 microcontroller [23] with inertial motion units (IMU), one stretch sensor [24], and one ECG sensor for monitoring the user activities. We characterize the power consumption of the TI CC2652 microcontroller using a digital multimeter. The stretch sensor is passive in nature and consumes energy only when it is communicating data to the processor. Therefore, we only consider communication power for the stretch sensor. We obtain the power consumption values of the ECG sensors from the study reported in [25].

The wearable device used also includes five EH modules. Specifically, we include piezoelectric sensors on the knees and elbows (a total of four locations) so that they can harvest energy when the user's legs or hands are moving. The system also includes an SP3-37 [26] flexible PV-cell to obtain energy from ambient light. The PV-cell is mounted on the user's shoulder to allow maximum exposure to ambient light.

User activity benchmark and pre-processing: We use the ARAS activity data benchmark [27] to obtain traces of user activities. The ARAS benchmark contains activity data collected from two houses with four users over 30 days. It includes a total of 27 activities collected in a smart home environment. We first pre-process the activity labels to merge the redundant activities. For example, the ARAS benchmark classifies cooking and eating breakfast, lunch, and dinner separately. We merge these labels as cooking and eating, respectively. At the end of pre-processing, we have nine activity labels: {*Active, cook, eat, clean, sleep, leisure, work, care, and other*}. The activities are

then used in the EH model presented in Section IV-B to obtain the energy available for each of the activities.

DIET parameters: The hyperparameters for DIET include the weight α for the EWMA filter and the learning rate for the multi-armed bandit. We set α to 0.5 since it has shown good results in prior work [11]. The learning rate for the multi-armed bandit is set to 10^{-3} so that sensors can adapt to changes in EH patterns. The time horizon is set to 24 hours, and each interval is 1 hour. The battery target is set to 5 J.

Oracle solution: We use the CVX [28] package to obtain an Oracle solution to the optimization problem in Equations 4–9. The Oracle uses actual energy harvest and consumption over the entire horizon, which is not feasible in a runtime solution due to the lack of future information.

Baseline solution: We use the approach proposed in [8] as a baseline for comparisons. The baseline approach first obtains an energy budget from multiple energy sources using the algorithm in [15]. The energy budget is then allocated to tasks on the sensor nodes such that the amount of data transmitted is maximized while avoiding power failures. While not practical due to the central energy manager, it provides a useful baseline to analyze the effectiveness of the solutions obtained by DIET.

B. Validation of Activity and Energy Prediction

To evaluate the accuracy of the activity and duration prediction models, we first use 14 days of data to train the initial DBN for each user. The remaining 16 days are used to perform the initial accuracy analysis before updating the DBN at runtime. With the initial training, we achieve 30% activity prediction accuracy for user 1, which is insufficient for effective energy management. Therefore, we continuously update the DBN as new data is observed at runtime. Table I shows the activity prediction accuracy for all the users after completing the online updates. As we can see, the accuracy for user 1 has improved to 56%, while the accuracy for other users is about 62%. The accuracy for user 1 is lower since the activities of the user have a higher degree of randomness compared to other users. We also note that the accuracy of activity prediction achieved in this work is comparable to the prediction accuracy of prior approaches [13]. The second column of the table also shows that the duration prediction approach achieves more than 85% accuracy for all the users. The high accuracy ensures that we make accurate predictions of energy harvest.

The predicted activity and duration are used to predict the energy available in an interval. The last three columns of Table I show the mean absolute error of energy prediction for all the users. We observe that the error in motion EH is lower than 0.5 J for all the users. The error in light energy prediction is higher due to the higher magnitude of difference between outdoors and indoors energy. For instance, user 4 has a higher error because the user spends most of their time outside, which leads to higher variations due to weather changes. Overall, the higher error for light energy is acceptable since we apply corrections to the energy allocations at the beginning of each interval as a function of the actual energy.

TABLE I: Activity and duration prediction validation accuracy

User	Act. Acc.	Dur. Acc.	Light Energy MAE	Motion Energy Hand (MAE)	Motion Energy Leg (MAE)
1	56%	88	0.70	0.05	0.03
2	63%	87%	10.47	0.15	0.12
3	68%	88%	6.06	0.15	0.10
4	73%	86%	19.41	0.33	0.27

C. Utility Comparison with Baseline and Oracle

One of the primary objectives of DIET is to enable the long-term operation of the device on harvested energy. To this end, we simulate DIET over a period of 30 days for each user. Then, we compare the utility obtained by the Oracle, DIET and the baseline, as shown in Figure 2. The DIET approach closely follows the utility of the Oracle for users 2, 3, and 4. In contrast, the baseline approach has a significantly lower utility. This is because the baseline is unable to learn the user's activity and EH patterns. The utility of DIET is lower in the first few days because it has not learned the activity and EH patterns of the users. Once DIET learns the activity patterns of the user, it quickly adapts and achieves utility that is equal to the Oracle. User 1 experiences higher variation in the utility due to the highly stochastic nature of the activities. Moreover, the total EH for user 1 is lower than other users, due to which DIET is unable to distribute the energy as effectively as the Oracle. The baseline approach achieves higher utility than DIET on days 4, 5, and 6 for user 1. This is because the baseline allocates lower energy on the first three days with lower utility and uses the saved energy on days 4, 5, and 6. However, it fails to adapt to the changes in the EH pattern after day 6. The average utility achieved by DIET for user 1 is within 35% of the optimal, which shows that DIET works effectively for users with stochastic activities and energy harvest. In summary, the utility of the DIET is within 10% of the optimal for all users.

D. Evaluation of Energy Allocations by DIET

In this section we compare the performance of DIET with the Oracle in terms of the energy consumption. Figure 3 shows a comparison of the energy consumption of the four sensors for user 1 on a specific day. The blue dotted line shows the energy requirement in each hour, while the red and black (dashed) lines show the Oracle and DIET, respectively. We observe that the Oracle follows the required energy level for the stretch sensor and the ECG. The DIET approach has a lower energy consumption in the early hours because it does not have knowledge of the future and conserves energy. During

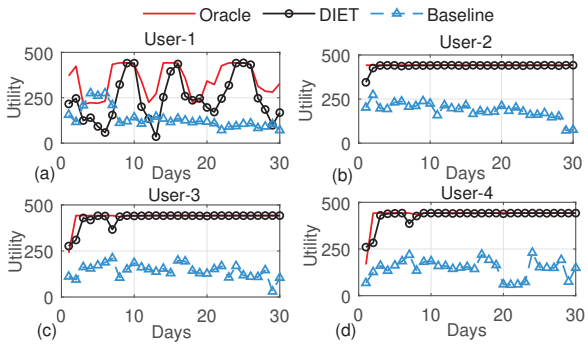

Fig. 2: Utility comparison of proposed algorithm to optimal and baseline for a) user 1, b) user 2, c) user 3 and d) user 4 over 30 days

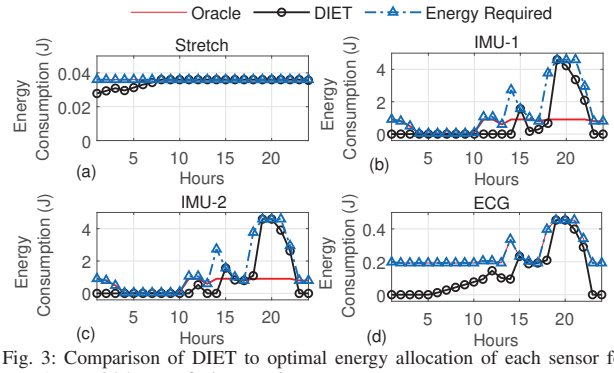

Fig. 3: Comparison of DIET to optimal energy allocation of each sensor for user 1 over 24 hours of Time Horizon

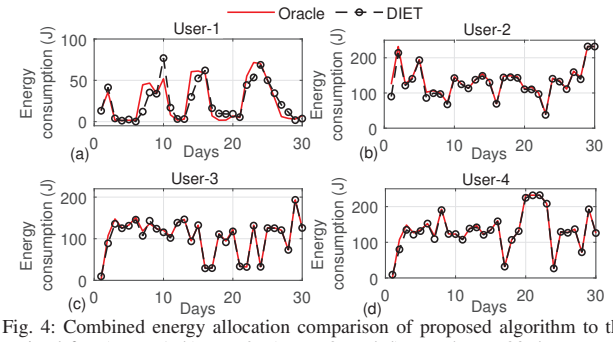

Fig. 4: Combined energy allocation comparison of proposed algorithm to the optimal for a) user 1, b) user 2, c) user 3, and d) user 4 over 30 days

the later hours of the day, the sensors have sufficient energy to follow the energy requirements. The IMU sensors have a higher energy requirement due to their higher power consumption. The Oracle is unable to meet the energy requirement in all intervals because sufficient energy is not available in the system. The DIET approach also has to consume lower energy during the initial parts of the day due to insufficient energy. However, there is sufficient energy in the battery during the later parts of the day to meet the energy requirements.

Finally, we compare the total energy consumption for all four users for 30 days in Figure 4. This comparison allows us to verify whether the total energy consumption of the DIET approach matches the Oracle. The total energy consumption of the DIET approach closely follows the Oracle, except for few days. The energy consumption of DIET is higher than the Oracle on some days because of the fact that DIET consumed lower energy in the previous days. The surplus energy is used in the following days to achieve a higher utility. In summary, these results show that DIET closely follows the energy consumption of the Oracle for all four users.

E. Implementation Overhead

We implement the DIET approach on the CC2652 processor to characterize its execution time and energy overhead. We observe that the lightweight optimization algorithm to allocate energy used by each harvester takes about 0.05 ms to execute, which is negligible compared to the interval length of one hour. Similarly, the bandit update algorithm takes about 0.09 ms for each sensor. The energy consumption for the two algorithms is 0.60 μJ and 1.18 μJ, respectively. The energy overhead is less than 0.0003% of the typical EH value.

VI. Conclusion

Wearable devices are transforming activity and health monitoring with the ability to record user data continuously. However, their widespread usage is still constrained by energy limitations. Recent work has proposed ambient energy harvesting as an effective solution to augment the battery of wearable devices. However, most energy harvesting methods assume that the placement of the harvester and sensors is the same. The DIET approach addresses this issue and performs dynamic energy management with multiple harvesters and sensors. Specifically, we harvested energy from multiple sources and allocated it to the sensors while maximizing the utility. Experiments on a real-world dataset showed that DIET achieves utility that is within 10% of the optimal Oracle solution.

Acknowledgements: This work was supported in part by the Washington State University New Faculty Seed Grant.

References

[1] W. Maetzler, J. Klucken, and M. Horne, "A Clinical View on the Development of Technology-Based Tools in Managing Parkinson's Disease," *Movement Disorders*, vol. 31, no. 9, pp. 1263–1271, 2016.
[2] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A Survey," *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.
[3] A. Ozanne *et al.*, "Wearables in Epilepsy and Parkinson's disease—A Focus Group Study," *Acta Neurologica Scandinavica*, vol. 137, no. 2, pp. 188–194, 2018.
[4] A. Valenzuela, "Energy Harvesting for No-Power Embedded Systems," 2008, https://bit.ly/3fnA6Vm, accessed 3/28/2021.
[5] D. Matsouka *et al.*, "Piezoelectric Textile Fibres for Wearable Energy Harvesting Systems," *Materials Research Express*, vol. 5, no. 6, p. 065508, 2018.
[6] Y. Tuncel *et al.*, "Towards Wearable Piezoelectric Energy harvesting: Modeling and Experimental Validation," in *Proc. of ISLPED*, 2020, pp. 55–60.
[7] D. Castro *et al.*, "Wearable-Based Human Activity Recognition Using an IoT Approach," *J. Sensor and Actuator Netw.*, vol. 6, no. 4, p. 28, 2017.
[8] P.-D. Gleonec *et al.*, "Energy Allocation for LoRaWAN Nodes with Multi-Source Energy Harvesting," *Sensors*, vol. 21, no. 8, p. 2874, 2021.
[9] F. Terraneo, A. Leva, and W. Fornaciari, "A High-Performance, Energy-Efficient Node for a Wide Range of WSN Applications," in *Proc. EWSN*, 2016, pp. 241–242.
[10] Y.-K. Chen, "Challenges and Opportunities of Internet of Things," in *ASPDAC*, 2012, pp. 383–388.
[11] A. Kansal *et al.*, "Power Management in Energy Harvesting Sensor Networks," *ACM Trans. Embedd. Comput. Syst.*, vol. 6, no. 4, p. 32, 2007.
[12] N. Yamin and G. Bhat, "Online solar energy prediction for energy-harvesting internet of things devices," in *Proc. ISLPED*, 2021, pp. 1–6.
[13] E. Nazerfard and D. J. Cook, "CRAFFT: An Activity Prediction Model Based on Bayesian Networks," *JAIHC*, vol. 6, no. 2, pp. 193–205, 2015.
[14] Z.-H. Wu *et al.*, "A Bayesian Network Based Method for Activity Prediction in a Smart Home System," in *Proc. SMC*, 2016, pp. 001 496–001 501.
[15] C. M. Vigorito, D. Ganesan, and A. G. Barto, "Adaptive Control of Duty Cycling in Energy-Harvesting Wireless Sensor Networks," in *IEEE SECON*, 2007, pp. 21–30.
[16] Y. Xiao *et al.*, "Dynamic Energy Trading for Energy Harvesting Communication Networks: A Stochastic Energy Trading Game," *IEEE J. on Selected Areas in Comm.*, vol. 33, no. 12, pp. 2718–2734, 2015.
[17] D. Hussein, G. Bhat, and J. R. Doppa, "Adaptive Energy Management for Self-Sustainable Wearables in Mobile Health," in *Proc. AAAI*, 2022.
[18] Sandia National Laboratories, "Sandia's Ephemeris Model," 2017, https://bit.ly/2XCqhMJ, accessed 9/18/2021.
[19] P. Ineichen and R. Perez, "A New Airmass Independent Formulation for the Linke Turbidity Coefficient," *Solar Energy*, vol. 73, no. 3, pp. 151–157, 2002.
[20] X. Ma, S. Bader, and B. Oelmann, "Characterization of Indoor Light Conditions by Light Source Classification," *IEEE Sensors J.*, vol. 17, no. 12, pp. 3884–3891, 2017.
[21] R. S. Sutton and A. G. Barto, *Introduction to Reinforcement Learning*, 2nd ed. MIT Press, 2018.
[22] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge Univ. Press, 2004.
[23] Texas Instruments Inc., "CC2652R Microcontroller," [Online] https://www.ti.com/product/CC2652R, accessed 1 November 2020, 2018.
[24] B. O'Brien, T. Gisby, and I. A. Anderson, "Stretch Sensors for Human Body Motion," in *Proc. Electro. Polymer Actuators and Dev.*, vol. 9056, 2014, p. 905618.
[25] M. Altini *et al.*, "An ECG Patch Combining a Customized Ultra-Low-Power ECG SoC with Bluetooth Low Energy for Long Term Ambulatory Monitoring," in *Proc. Conf. on Wireless Health*, 2011, pp. 1–2.
[26] FlexSolarCells, "SP3-37 Datasheet," 2013, https://bit.ly/3dcJ0lK, accessed 3/28/2021.
[27] H. Alemdar *et al.*, "ARAS Human Activity Datasets in Multiple Homes with Multiple Residents," in *Int. Conf. Perv. Comput. Tech. Health. and Work.*, 2013, pp. 232–235.
[28] M. Grant and S. Boyd, "CVX: Matlab Software for Disciplined Convex Programming, V. 2.1," 2014, http://cvxr.com/cvx, accessed 9/14/2021.