An adaptive duty-cycle scheme for GPS scheduling in mobile location sensing applications *

Chao-Lin Wu†, Yu-Te Huang†‡, Cheng-Lung Wu†, Hao-hua Chu‡, Polly Huang‡, and Ling-Jyh Chen†
† Academia Sinica, Taiwan
† {clwu, ytls, nickwulul6178, cclljj}@iis.sinica.edu.tw
‡ National Taiwan University, Taiwan
‡ hchu@csie.ntu.edu.tw, phuang@cc.ee.ntu.edu.tw

ABSTRACT

Mobile location sensing applications (MLSAs) represent an emerging genre of applications that exploit Global Positioning System (GPS) technology and facilitate location-based services. The design of MLSAs must incorporate a tradeoff between information accuracy and energy efficiency because GPS technology is energy expensive and unaffordable for most MLSA platforms, which are battery-powered and therefore resource-constrained. In this study, based on our observation that the reception of GPS signals is spatially and temporally correlated, we propose a novel algorithm called the Adaptive Duty Cycle (ADC) scheme to exploit the spatio-temporal localities in the design of GPS scheduling algorithms. Using a comprehensive set of evaluations, as well as realistic hiker mobility traces, we evaluate the ADC scheme in terms of data granularity and power consumption. The results demonstrate that the scheme can achieve the Pareto optimum in all test cases. Moreover, the scheme is simple, effective, and generalizable to other mobile location sensing applications.

1. INTRODUCTION

Mobile location sensing applications (MLSAs) continue to permeate every part of our living environments. In addition to maintaining the functions of mobile applications (e.g., supporting mobility and battery-powered operations), MLSAs exploit Global Positioning System (GPS) technology and facilitate the provision of emerging location-based services (LBS). This new genre has already spawned a wide range of applications, such as mobile social networks [8], mobile urban sensing systems [1], and moving object tracking systems [4, 11].

The use of GPS technology in MLSAs raises a new technical challenge for two reasons. First, an MLSA requires timely and accurate location information, so it favors keeping GPS receivers in the ON mode continuously. Second, GPS receivers are widely regarded as power-hungry; thus, keeping them "always-on" is unaffordable for MLSAs because the lifespan of the latter is constrained by the battery capacity of the associated mobile devices. It is therefore necessary to find a tradeoff between information accuracy and energy efficiency. Several solutions to the problem have been proposed. Generally, they can be divided into two categories: *static duty-cycle* (SDC) approaches (e.g., [1, 11]), and *dynamic duty-cycle* (DDC) approaches (e.g., [5–7, 9, 10, 12]). SDC approaches turn GPS receivers ON and OFF at regular intervals, while DDC approaches adjust GPS duty cycles based on analytical models or events triggered by additional sensors.

However, existing DDC approaches are infeasible for generic MLSA scenarios for three reasons: 1) they require additional sensors (e.g., radio signal receivers and accelerometers) as triggers; 2) they need a prebuilt lookup table to handle the triggered events; and 3) when placed in outdoor environments without pre-learned radio patterns, they tend to keep GPS receivers in the ON mode continuously, and thus waste energy. On the other hand, although SDC approaches do not have the above limitations, they are operated in a "blind" manner, without considering the spatial and temporal correlation of GPS signal reception in duty-cycle scheduling.

In this study, we focus on the use of MLSAs in challenging environments, where the reception of GPS signals may be affected by the canopy cover, and the mobile devices are both resource- and power-constrained. Based on our observation that spatial and temporal localities exist in the reception of GPS signals, we propose the Adaptive Duty Cycle (ADC) algorithm, which is a DDC-based approach that exploits the localities in the design of GPS scheduling algorithms. Using a comprehensive set of simulations, as well as realistic traces of hiker mobility, we evaluate the ADC scheme in terms of information accuracy and energy efficiency. The results demonstrate that the scheme outperforms existing approaches in all test cases. Moreover, it is simple, practical, and effective for emerging mobile location sensing applications.

The remainder of this paper is organized as follows. Section 2 contains a review of GPS technology and existing power saving schemes in MLSAs. In Section 3, we describe the proposed ADC scheme; and in Section 4, we present a comprehensive set of experiment results, which we analyze and explain in detail. We then summarize our conclusions in Section 5.

2. BACKGROUND AND RELATED WORK

The Global Positioning System (GPS) is a global navigation satellite system (GNSS) that provides reliable location and time infor-

*This research was supported in part by the National Science Council of Taiwan under Grants: NSC 100-2219-E-001-001 and NSC 100-2219-E-001-002.

mation anywhere anytime worldwide [2]. Each GPS satellite continually transmits messages that contain the precise time a message was sent, as well as the ephemeris, and the almanac. Since messages are transmitted at the speed of light, the GPS receiver determines the distance to each GPS satellite by calculating the delay between the time a message is sent and the time it is received. Then, using *trilateration* algorithms, the GPS receiver calculates its own location after receiving messages from more than three GPS satellites.

However, the effectiveness of GPS localization may degrade when the number of GPS satellites in view of the receiver decreases due to the cover created by surrounding environments [2]. Moreover, the reception of GPS signals may be affected by the canopy cover (e.g., the terrain, obstacles, and clouds), which leads to errors in GPS localization calculations. Thus, it is recommended that four or more satellites should be used to calculate the receiver's location [11].

One of the most challenging aspects of GPS technology is its power consumption, which is expensive and renders it unaffordable for emerging mobile and battery-powered platforms, such as smart phones and wireless sensor networks. To prolong the lifespan of built-in GPS systems, several *static duty-cycle* (SDC) schemes reduce the power consumption of GPSs by turning them ON and OFF periodically. For instance, the BikeNet [1] and ZebraNet [11] projects use SDC-based approaches and turn their GPS receivers ON and OFF according to a fixed duty cycle. The former project turns the GPS receivers on for one minute every 8 minutes; while the latter turns the receivers on for 150 seconds every 1,000 seconds. In contrast, *dynamic duty-cycle* (DDC) approaches adjust the GPS duty cycle based on sophisticated models [5, 9, 10] or events triggered by add-on sensors. For example, they adjust duty cycles based on whether people are in outdoor/indoor environments using radio signal recognition [6], or whether the tracked objects have been moving over a certain distance [7, 12]).

3. ADAPTIVE DUTY CYCLE (ADC)

3.1 The Algorithm

We propose a novel Adaptive Duty Cycle (ADC) algorithm to facilitate the efficient use of GPS power in mobile location sensing applications. The rationale behind ADC is based on the observation that the failures of GPS lock attempts are usually spatially and temporally correlated; hence, ideally, a GPS scheduling algorithm should consider the spatial and temporal localities of failures.

We define a GPS lock attempt as a *success* if the GPS receiver receives messages from at least four satellites within $T_{attempt}$ time; otherwise, it is deemed a *failure*. Let W_i be the time interval between the $(i-1)$-th and the i-th GPS lock attempts; and let W_{min} and W_{max} be the lower- and upper-bound of W_i respectively ($T_{attempt} < W_{min} < W_{max}$). There are two cases where the ADC scheme calculates the value of W_i after the $(i-1)$-th GPS lock attempt:

- **Case 1:** if the $(i-1)$-th GPS lock attempt is successful, ADC derives W_i by Equation 1, which multiplies W_{i-1} by a scaling factor of the *information accuracy* ($0 < \alpha < 1$) to increase the frequency of GPS lock attempts, thereby improving the timeliness of location sensing.

$$W_i = max(W_{i-1} \times \alpha, W_{min}) \quad (1)$$

- **Case 2:** if the $(i-1)$-th GPS lock attempt fails, ADC multiplies W_{i-1} by a scaling factor of the *energy efficiency* ($1 < \beta < 2$), which saves energy by delaying subsequent GPS lock attempts, as shown in Equation 2.

$$W_i = min(W_{i-1} \times \beta, W_{max}) \quad (2)$$

Using Equations 1 and 2, the ADC scheme strategically updates the frequency of its GPS lock attempts by exploiting the localities of the GPS lock results. As a result, it can better accommodate the tradeoff between information accuracy and power consumption.

3.2 Evaluation Metrics

Let δ_i denote the result of the i-th GPS lock attempt ($\delta_i = 1$ if it is a success, and $\delta_i = 0$ otherwise). We design two performance metrics to evaluate the performance of the ADC scheme:

1. To evaluate the *information accuracy* aspect, we measure the **data granularity** (Φ), i.e., the average time interval between two consecutive successes of GPS lock attempts, as shown in Equation 3. Intuitively, the smaller the value of Φ, the more accurate the location information provided.

$$\Phi = \frac{\sum_{\forall i} W_i}{\sum_{\forall i} \delta_i}. \quad (3)$$

2. For the *energy efficiency* aspect, we measure the **power consumption** (Ψ), i.e., the percentage of time that the GPS receiver is turned ON. The larger the value of Ψ, the more energy consumed by the GPS receiver. Specifically, if the i-th GPS lock attempt is successful (i.e., $\delta_i = 1$), it takes T_{hot} or T_{warm} time depending on whether the lock attempt is a *hot start* or a *warm start*[1]. On the other hand, a failed GPS lock attempt takes $T_{attempt}$ time. Let θ_i denote whether or not the i-th GPS lock attempt is a successful hot start ($\theta_i = 1$ if it is a hot start; otherwise, $\theta_i = 0$). We obtain the value of Ψ by

$$\Psi = \frac{\sum_{\forall i} (\delta_i(\theta_i T_{hot} + (1-\theta_i)T_{warm}) + (1-\delta_i)T_{attempt})}{\sum_{\forall i} W_i}. \quad (4)$$

4. EVALUATION

We evaluate the proposed ADC scheme via both trace-based simulations and real-world experiments. The traces used for simulations were collected on the *Yushan Peak Trail* in Yushan National Park on May 27 and 28, 2009. The trail is a 10.9km long and runs from the park entrance to the summit of the *Yushan Mountain* (3,952m - the highest mountain in North East Asia) with a 1302m altitudinal shift. We recruited 20 volunteers and gave each of them a GPS tracker (Model: Genie GT-31[2]). On the first day, the participants

[1]The GPS receiver initiates a *hot start* for a GPS lock attempt if it is less than 60 minutes since the last successful GPS lock attempt, and the almanac and ephemeris stored in its memory are still valid. It initiates a *warm start* if only the almanac stored in its memory is still valid; otherwise, a *cold start* is initiated. In general, a hot start takes less than five seconds, and a warm start takes about 30 seconds. The cold start is the slowest type, because it has to download the almanac and ephemeris data, which usually takes about 45 seconds. Note that the exact time required for GPS hot/warm/cold starts may vary with different chipsets, but the order remains the same.

[2]Genie GT-31/BGT-31: http://www.locosystech.com

hiked from the entrance to the Paiyun Lodge, which is about 2.4km beneath the summit and provides very basic overnight accommodation for hikers. The next day, they climbed to Yushan Peak, and then returned to the entrance. Using the SDC approach, each GPS tracker recorded the timestamp, latitude, and longitude information of the hiker's position, as well as the number of GPS satellites sensed by the tracker, every minute. Each hiker contributed a 32-hour trace approximately. The traces provided the ground truth of the hikers' mobility on the trail, together with the results of each tracker's GPS lock attempts, which we subsequently analyzed to calculate the canopy cover rates of the trail.

In addition, we implemented the ADC and SDC schemes on the Android OS platform, and conducted the real-world experiment using four HTC Hero cellphones on the *Chihsing Peak Trail* in Yangmingshan National Park on July 21 and 25, 2011. The trail is 5km long and runs from the *Lengshueikeng*, through the summit of the *Chihsing Mountain* (1,120m - the highest mountain in Great Taipei City area with a 400m altitudinal shift), and to the *Siaoyukeng*. To eliminate the influence of uncontrolled variables during experimentation, two of the cellphones run the ADC scheme on July 21 and the SDC scheme on July 25; and the other cellphones run the SDC scheme on July 21 and the ADC scheme on July 25. Each cellphone contributed an approximately 2-hour trace on each day.

4.1 Locality of GPS Signal Reception

First, we analyzed the Yushan traces and verified that localities existed in the reception of GPS signals. We consider each hiker's trace as a sequence of *success* and *failure* periods; and each *success* and *failure* period may consist of one or more back-to-back GPS lock successes and failures respectively. Figure 1 shows the distribution of the lengths of the *success* and *failure* periods as cumulative distribution function (CDF) curves. We observe that approximately 40% of the GPS lock failures and 75% of the GPS lock successes appear successively (i.e., the length of the period is greater than one), which confirms our argument that the results of GPS lock attempts are spatially and temporally correlated (i.e., localities do exist).

In addition, suppose each successful GPS lock attempt takes T_{hot} time and each GPS lock failure takes $T_{attempt}$ time. Figure 2 shows the distribution of the GPS receiver's power consumption when $T_{hot} = 2$ seconds and $T_{attempt} = 50$ seconds in the Yushan traces. We observe that 13.8% of the energy was consumed by successful GPS lock attempts (i.e., the length of successive failures was zero); and 66.3% of the energy was consumed by successive GPS lock failures (i.e., the length of successive failures was greater than or equal to two). In other words, *successive GPS lock failures consume an excessive amount of energy and must be avoided*.

4.2 Evaluation Under Synthetic Cover Rates

Next, we compare the performance of the ADC and SDC schemes in terms of information accuracy and energy efficiency under various canopy cover rates. Using the Yushan traces, we employ the *Trimble Planning Software*[3] to calculate *the number of GPS satellites that were in view of each hiker's GPS receiver* on the trail under 20%, 40%, and 60% cover rates. We consider that a GPS lock attempt is successful if the receiver can see four or more satellites; otherwise, it is deemed a failure.

Based on the results derived by the *Trimble Planning Software*,

[3]Trimble; http://www.trimble.com/planningsoftware.html

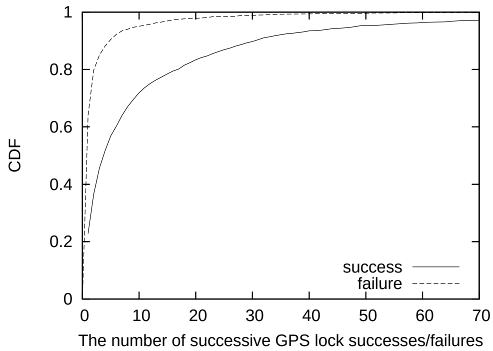

Figure 1: The CDF curves of the lengths of GPS lock *success* and *failure* periods in the Yushan traces

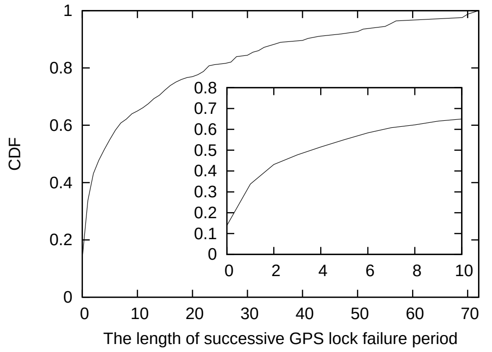

Figure 2: The CDF curves of the power consumption of the GPS receivers in the Yushan traces when $T_{hot} = 2$ seconds and $T_{attempt} = 50$ seconds

we evaluate the ADC and SDC schemes using trace-based simulations. Specifically, when the ADC scheme is used, we fix the values of W_{max} and W_{min} at 15 minutes and 1 minute respectively. Moreover, we vary the value of α as 0.5, 0.6, 0.7, 0.8, and 0.9; and the value of β as 1.1, 1.2, 1.3, ..., and 2.0 (i.e., there are $5 \times 10 = 50$ combinations of α and β pairs). Meanwhile, under the SDC scheme, the time intervals between two GPS lock attempts are set at 1, 2, 3, ..., and 15 minutes respectively.

Using Equations 3 and 4, we calculate the data granularity (Φ) and power consumption (Ψ) of each hiker's GPS receiver under each configuration of the ADC and SDC schemes. In total, there are 50 $\times$ 20 = 1,000 instances for the ADC scheme, and 15 $\times$ 20 = 300 instances for the SDC scheme[4]. We average the evaluation results of each configuration under the two schemes, and calculate their *Pareto frontiers* [3], as shown in Figure 3, respectively.

From the results, we observe that the performance of ADC and SDC are comparable when the canopy cover rate is 20% because they have the same Pareto frontier, as shown in Figure 3-a. The reason is that the signal reception of GPS receivers is good when the cover rate is low. In this case, the ADC scheme behaves in a similar

[4]The performance of the ADC and SDC schemes may vary substantially under different parameter settings. To compare the two schemes, we use an exhaustive set of configurations for evaluation and observe their Pareto frontiers. We discuss the impacts of the parameter settings on the simulation results further in later subsections.

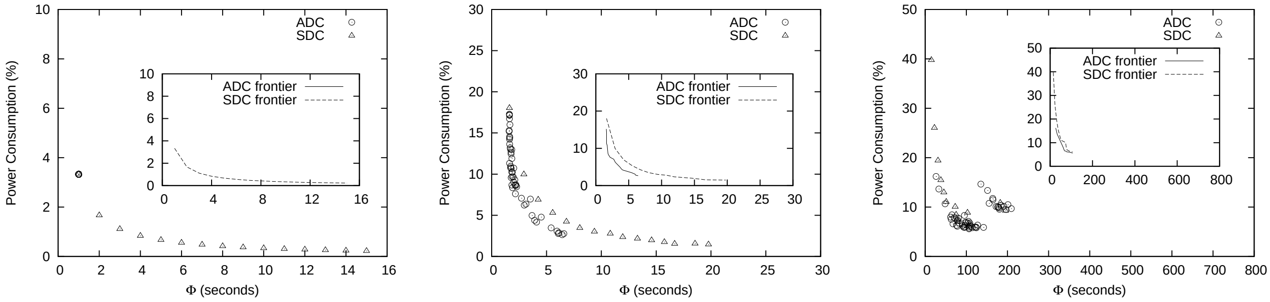

(a) 20% cover rate (b) 40% cover rate (c) 60% cover rate

Figure 3: The evaluation results of the ADC and SDC schemes under different synthetic canopy cover rates.

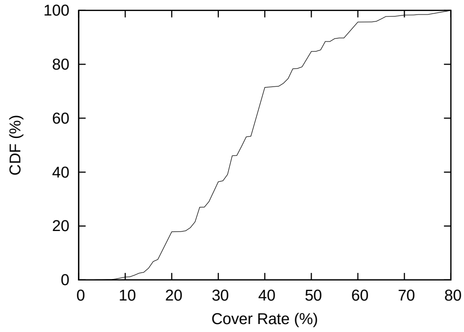

Figure 4: The CDF curves of realistic canopy cover rates in the Yushan traces

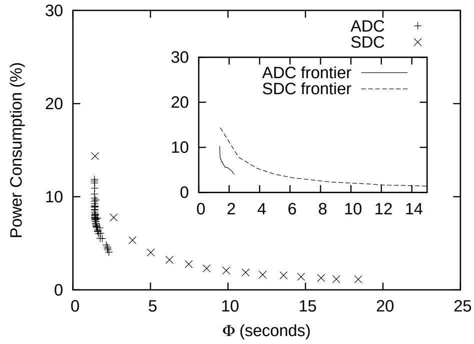

Figure 5: The evaluation results of the ADC and SDC schemes under realistic canopy cover rates

way to SDC with a nearly fixed time interval for W_{min} between every two consecutive GPS lock attempts. However, when the cover rate is above 20%, the ADC scheme outperforms the SDC scheme and it always achieves the "Pareto optimum" [3]. In other words, there are no instances on SDC's frontier that have finer data granularity and lower power consumption than the instances on ADC's frontier.

4.3 Evaluation Under Realistic Cover Rates

We also calculate the realistic canopy cover rates of each hiker along the trail by *reverse engineering* approaches; that is, we use a brute-force search and the *Trimble Planning Software* to determine the cover rate at each location for each timestamp in a hiker's mobility trace. Figure 4 shows the distribution of the obtained canopy cover rates as a CDF curve.

Using the same parameter settings, we evaluate the performance of the two schemes in terms of data granularity and power consumption under realistic canopy cover rates. From the results shown in Figure 5, we observe that, once again, the ADC scheme outperforms the SDC scheme because it achieves the *Pareto optimum* in all test cases.

4.4 Evaluation with Different Parameters

Next, we evaluate the impacts of the parameter settings (α, β, W_{min}, and W_{max}) on the information accuracy and energy efficiency under the ADC scheme. Using the Yushan traces, as well as the realistic canopy cover rates derived in the previous subsection, we configure the parameters by their default values ($\alpha = 0.7$, $\beta = 1.3$, $W_{min} = 1$ minute, and $W_{max} = 15$ minutes), and vary them one by one in the experiment. From the results shown in Figure 6, we observe that the different values of the four parameters may cause the behavior of the overall system to vary. Therefore, the parameters should be configured properly based on a combination of several factors, such as the balance of the tradeoff between information accuracy and energy efficiency, the minimum lifespan required by the system, and the maximum data granularity allowed by the system.

4.5 Real-world Experiment

Finally, we present the results of the real-world experiment in Figure 7. The duty cycle used in the SDC scheme is set to 1 minute, and the parameters used in the ADC scheme are $\alpha = 0.7$, $\beta = 1.3$, $W_{min} = 1$ minute, and $W_{max} = 15$ minutes. The results in Figure 7 show that the ADC scheme consumes less power (3% approximately in average) but SDC has finer data granularity (about 20 seconds in average). Note that, in this study, the experiment was conducted with a clear sky, in which the SDC scheme was benefited because most of GPS lock attempts were succeeded. However, if the canopy cover (e.g., the terrain, obstacles, and clouds) was common on the trail, successive GPS lock failures were inevitable and the information accuracy of the SDC scheme would degrade significantly. Meanwhile, the ADC scheme is more favorable, because it is more responsive and better able to preserve the information accuracy while keeping the power consumption minimal.

5. CONCLUSION

In this paper, we propose an Adaptive Duty-Cycle (ADC) scheme for GPS scheduling in mobile location sensing applications. The scheme exploits the localities of GPS signal reception, and strategically adapts the duty cycle schedule of GPSs to accommodate the performance criteria of the information accuracy and energy efficiency aspects. Using real data collected from hikers' mobility traces, we first verify the spatial and temporal correlation of GPS

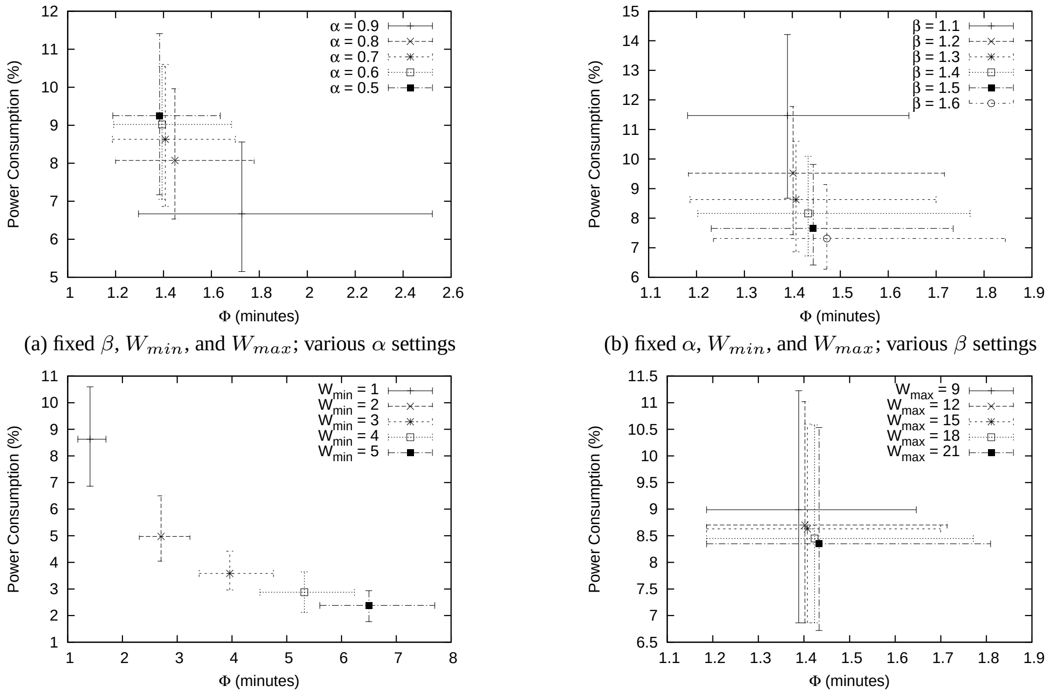

(a) fixed β, W_{min}, and W_{max}; various α settings

(b) fixed α, W_{min}, and W_{max}; various β settings

(c) fixed α, β, and W_{max}; various W_{min} settings

(d) fixed α, β, and W_{min}; various W_{max} settings

Figure 6: The evaluation results of the ADC scheme with different parameter values. The default values of the parameters are: $\alpha = 0.7$**,** $\beta = 1.3$**,** $W_{min} = 1$ **minute, and** $W_{max} = 15$ **minutes.**

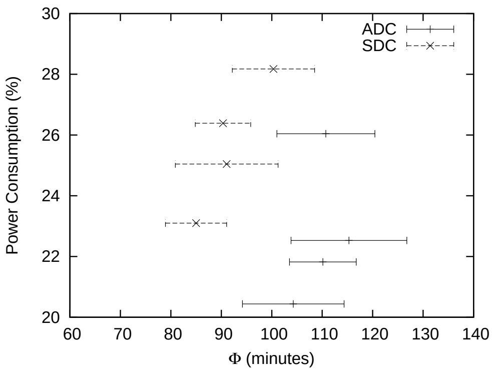

Figure 7: The evaluation results of the ADC and SDC schemes in the Yangmingshan trace

lock results. Then, we compare the ADC scheme with the SDC scheme in terms of data granularity and power consumption. The results demonstrate that the ADC scheme can achieve the Pareto optimum in all test cases. We also examine the impacts of different parameter settings on the performance of the proposed scheme, and suggest that the configuration of the ADC scheme should be based on a combination of several factors, including the balance of the tradeoff between information accuracy and energy efficiency, the minimum lifespan required by the system, and the maximum data granularity allowed by the system. Finally, we implement the proposed ADC scheme on the Android OS platform, and verify its performance using real-world experiments. The proposed ADC scheme is simple, practical, effective, and shows promises for future mobile location sensing applications.

References

[1] 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 Transactions on Sensor Networks*, 6(1):6:1–6:39, December 2009.

[2] J. A. Farrell and M. Barth. *The Global Positioning System & Inertial Navigation*. McGraw-Hill Professional, Dec. 1998.

[3] D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, Auguest 1991.

[4] J.-H. Huang, S. Amjad, and S. Mishra. Cenwits: A sensor-based loosely coupled search and rescue system using witnesses. In *ACM SenSys*, 2005.

[5] R. Jurdak, P. Corke, D. Dharman, and G. Salagnac. Adaptive GPS duty cycling and radio ranging for energy-efficient localization. In *ACM SenSys*, 2010.

[6] D. H. Kim, Y. Kim, D. Estrin, and M. B. Srivastava. SensLoc: sensing everyday places and paths using less energy. In *ACM SenSys*, 2010.

[7] M. B. Kjargaard, J. Langdal, T. Godsk, and T. Toftkjar. Entracked: energy-efficient robust position tracking for mobile devices. In *ACM SenSys*, 2009.

[8] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell. Sensing meets mobile social networks: The design, implementation and evaluation of the CenceMe applications. In *ACM SenSys*, 2008.

[9] J. Paek, J. Kim, and R. Govindan. Energy-efficient rate-adaptive GPS-based positioning for smartphones. In *ACM MobiSys*, 2010.

[10] S. Sundaresan, I. Koren, Z. Koren, and C. M. Krishna. Event-driven adaptive duty-cycling in sensor networks. *International Journal of Sensor Networks*, 6(2):89–100, 2009.

[11] P. Zhang, C. M. Sadler, S. A. Lyon, and M. Martonosi. Hardware design experiences in ZebraNet. In *ACM SenSys*, 2004.

[12] Z. Zhuang, K.-H. Kim, and J. P. Singh. Improving energy efficiency of location sensing on smartphones. In *ACM SenSys*, 2010.