

RSSI Based Location-Aware PC Power Management

Zhong-Yi Jin and Rajesh K. Gupta
University of California, San Diego
{zhjin, rgupta}@cs.ucsd.edu

ABSTRACT

Many existing PC power saving features are under utilized since they may interfere with users' normal operations of computers. By taking into consideration the locations of computer users, location-aware PC power management techniques can overcome such limitations and offer both energy savings and non-intrusiveness. However, existing location-aware power management schemes either rely on complex location systems for accurate location information or lack the capabilities to support efficient PC power management. In this paper, we present NAPS, a zone based location-aware power management scheme that supports efficient PC power management yet does not require exact location information. In addition, we present our preliminary study on the feasibility of using the received signal strength indicator (RSSI) to estimate zone based locations in an indoor environment. Our experiments using sensor node radios indicate that the RSSI can be used for such purposes.

1. INTRODUCTION

Modern computer systems are designed with sophisticated power management features to reduce energy consumption. In principle, they help conserve energy by putting various components of a computer into the most energy efficient states based on real-time demands of computing resources. For example, the voltage of a processor can be scaled or an entire computer can be put into sleep (suspend to RAM) or hibernation mode (suspend to disk).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.



Figure 1: A location-aware scheme

However, many existing PC power management features are under utilized. According to a survey [6], 2/3 of desktops in office buildings are left on at night in the U.S.. This is due to the fact that applying the power management features can interfere with users' normal operations of computers. For example, once a computer is in sleep or hibernation mode, a user would have to wait for it to wake up before he/she is able to use it again. In addition, even when a power management feature is adopted, users often tend to trade energy for convenience. For example, it is common to choose a long idling period before automatically turning off the display or putting the computer into sleep/hibernation.

Location-aware power management techniques [5, 2, 1] are proposed to overcome some of the limitations of conventional power management schemes. By taking into consideration users' location information, less intrusive and more energy efficient power management decisions/policies can be made. For instance, a computer in hibernation mode consumes less energy than it does in sleep mode but requires more time to wake up. With the location-aware power management technique [1] illustrated in Figure 1, a computer can automatically be put into the most appropriate mode for both energy savings (up to 83% [1]) and non-intrusiveness based on the distance the user is away from the computer. By starting the wakeup process ahead of time once the user is within a certain range, the computer will be ready for use when the user returns to the computer, as if no power management is used.

The performance of location-aware power management schemes, defined in terms of energy savings and non-intrusiveness, depends on the accu-

racy of location information. Existing research on location-aware power management falls into two extreme categories according to its requirements or assumptions over the accuracy of location information. Research in the first category [2, 1] relies on highly accurate location information that can only be provided by sophisticated location systems such as ultrasonic systems [1]. While this type of work can achieve optimal performance, the needs of complex location systems limit its real world applications¹. The second category of work [2] does not require exact location information and exploits radio connectivity for relative locations. It identifies users’ relative locations by scanning the presence of their wireless devices such as Bluetooth enabled mobile phones [2]. While the latter type of work is more practical for actual deployment, just knowing whether users are in or out of communication ranges significantly limits its performance. For example, it cannot be used for the scenario illustrated in Figure 1.

In this paper, we present NAPS: a Non-intrusive location **A**ware **P**ower management **S**cheme designed specifically for PCs. It works without the need of exact location information as required by the first category of work but can achieve similar levels of performance. In addition, unlike existing location-aware schemes that blindly put a computer into sleep or hibernation mode based on locations alone, NAPS also monitors the resource consumption of each running application and manages power at different abstraction levels. This minimizes the interference to background applications such as compilers while reducing the resource consumption of user interactive applications such as web browsers.

NAPS works by dividing space into a few discrete virtual zones and making power management decisions based on that. Since zones in NAPS are relative, we propose and also present in this paper our preliminary experimental study on the feasibility of using the received signal strength indicator (RSSI) of sensor node radios to estimate the zone location of a user in an indoor environment. We choose sensor node radios for their low power (38mW for CC2420), relatively long range and low cost [4].

2. RELATED WORK

Location or context aware power management techniques have been proposed for both home [5] and office [2] environments. The work in [5] builds a smart home using exact location information from RFIDs and pressure mats. The work in [1] studies

¹The focus of our paper is on indoor environments.

Table 1: Power management levels

Hierarchy	Level	Examples of Managed Components
Top	Application	Web browser
Middle	Device	CPU, Monitor, Disk
Bottom	System	PC (Sleep/Hibernation)

the potential energy savings in an office environment based on a large corpus of location data captured with an accuracy of about 3cm 95% of time. As mentioned in Section 1, NAPS is similar to [2] but uses zones for locations and systematically manages the power consumption of PCs at finer granularities. There is a large body of work on using RSSI for ranging based localization and we cover that briefly in Section 4. Our work is different as we use RSSI to estimate relative virtual locations or more importantly the relevant power states rather than exact physical locations or distance.

3. NAPS

An efficient power management scheme needs to take into account power management overheads defined in terms of energy and latency [3]. In NAPS, power management is carried out at three abstraction levels as shown in Table 1. At the application level, NAPS manages the resource consumptions of software programs. For example, a web browser can consume a significant number of CPU cycles in the background executing scripts or rendering dynamic content of downloaded web pages². If the user is not near the computer, there may be no need for such resource consumption and we can stop that by pausing the corresponding browser process(es). The list of programs that should be managed at the application level can be specified by the users or announced by individual programs at runtime based on their states. At the device level, NAPS manages the power of different hardware devices and at the system level the power of the entire computer as a whole. An example list of components that can be managed at each of the levels is shown in Table 1.

The three power management levels form a natural hierarchy since we cannot perform power management at a lower level without affecting power management decisions at higher levels. In other words, power management decisions at lower levels depend on power management decisions at higher levels. For example, if a compiler is running, then

²As an example, we create a web page at <http://cse.ucsd.edu/users/zhjin/powerdemo/demo1.html> that can consume 100% of a CPU/core.




Figure 2: NAPS Zones




Figure 3: NAPS state transition diagram

no power management should be performed on the compiler at the application level and as a result not at any other levels that could affect the performance of the compiler. Due to this natural hierarchy, both energy savings and power management overheads increase as the level goes down.

As shown in Figure 2, NAPS divides the physical space into virtual zones based on the relative distance between a user and a computer. The sizes of the zones are not explicitly defined using exact distance numbers. They are determined by the accuracy of the relative distance information provided to NAPS as well as the time that it takes to wake up the components of a PC from various low power states.

Three basic zones are defined in NAPS as shown in Figure 2. Usage-zone is an area where a user might have direct physical access to the computer. Nearby-zone is where a user does not have direct physical access but may quickly return to the computer. Away-zone is an area where a user is far away from the computer and it takes a relatively long period of time for the user to return to the computer. The gray regions are optional buffer areas set aside to prevent the zones from overlapping due to the accuracy of distance information. Note that the zones can be divided further based on the accuracy of relative distance information.

In NAPS, power management decisions are based on zones as well as power management levels which reflect energy savings and power management overheads. In other words, we want to maximize energy savings without interfering with users' normal operations of computers. In usage-zone, NAPS stops all location based power management actions. Application and device level power managements are carried out in nearby-zone. In away-zone, only system level power management decisions are made. It is important to note that NAPS is only respon-

sible for location related power management decisions and relies on existing power management schemes for other power management tasks. However, NAPS can significantly improve the performance of other power management schemes. For example, NAPS itself does not perform workload based dynamic voltage scaling (DVS) of CPUs but can improve the performance of any existing DVS schemes by pausing (*kill - STOP*) certain applications such as web browsers if the user is not in usage-zone. NAPS resumes (*kill - CONT*) paused applications automatically once the user returns to usage-zone for non-intrusive power management. A common state transition diagram of NAPS is shown in Figure 3.

4. RSSI FOR ZONES

A challenge in supporting NAPS is to provide the necessary relative distance information that can be used to establish the zones. In this paper, we propose to use RSSI for such information. RSSI has been studied intensively for ranging based localizations as summarized in [7]. The general conclusion is that RSSI alone is not sufficient to provide accurate and reliable range information in indoor environment due to multi-path. However, RSSI could be an ideal solution for NAPS as only coarse grain relative distance between a user and a computer is needed. In the following two sections, we present our preliminary experimental study on this topic.

5. EXPERIMENTAL SETUP

In our experiments, we use Chipcon CC2420 radios on TelosB Motes [4] to study the feasibility of establishing the three basic zones from RSSI in an indoor environment. On the CC2420, RSSI is calculated over 8 symbol periods and stored in the RSSI_VAL register. RSSI in dBm can be computed as $RSSI_VAL - 45$. Since the TelosB uses no more than 41mW in sending or receiving [4], it is ideal for our power management applications.

Our experiments use 2 TelosB nodes, a base node and a mobile node. The base node is connected to a stationary computer via USB and the mobile node is moved around to measure RSSI values at different locations. The experiments are carried out by having the base node broadcast a probe packet periodically. Once the mobile node receives the probe, it immediately replies with an ack packet, which includes the RSSI of the previously received probe packet. If the ack packet is received by the base node, it sends the RSSI values of both the probe and ack packets to a data collection program running on the computer. To help identify locations




Figure 4: Experiment Setup

in the analysis of our experimental data, the ack packet also contains a location number which is increased by one for each button press on the mobile node. The transmission time of both the probe and ack packets are also collected as part of the process.

The experiments are conducted in and around our office (7.24×7.09 meters) and the floorplan is shown in Figure 4. The office is fully furnished and divided with cubical walls and bookcases. For all the experiments, the base node is attached to a laptop sitting on a desk at location number 0 in Figure 4. The base node sends a probe packet every 250ms. Both base and mobile nodes send packets at the maximum allowable transmission power of CC2420. The two nodes have a direct line of sight only when the mobile node is at location 0 and 1. Experiments are carried out at night with 1-2 people in the office. The hallway has people walking by frequently.

6. EXPERIMENTS

6.1 At fixed locations inside office

Our first experiment is used to establish a general RSSI profile of the office by measuring RSSI at each of the 9 numbered locations in Figure 4. Data are collected by leaving the mobile node on a chair at each of the 9 locations in sequence (alphabetic order). The mobile node is moved to the next location after RSSI values stabilize on our data collection program which plots the data in real time. The location number of the mobile node is updated after each movement. The results of this experiment are shown in Figure 5 which plots the values of RSSI and a 5.25s moving average of that. Note that since the values of probe-RSSI and ack-RSSI are very close in all our experiments, we only plot the probe-RSSI values in the figures.

As shown in Figure 5, once the mobile node is left on the chair, the variation of RSSI is only around 10dBm. By following the moving average, we can see that the RSSI values fall into a few clearly separated tiers with the highest values at location 0 and lowest at location 5. The RSSI moving average values at location 1 are larger than those at location 7 because the nodes are in direct line of sight in




Figure 5: RSSI at fixed locations in the office




Figure 6: RSSI when moving inside the office

the former case. These tiers suggest the potential of using RSSI for relative distance.

We can also see in Figure 5 that the variation of RSSI is very significant while the mobile node is being moved to different locations. This is mainly caused by the change of multi-path at different locations. However, by using a 5.25s moving average, we can remove most of the outliers and the smallest RSSI value becomes -68dBm.

6.2 Walking inside office

Our second experiment is designed to study the change of RSSI while the mobile node is moving continuously inside the office. This experiment is similar to the first one except the mobile node is carried inside a pocket of a person who walks around the office by following the numbered locations three times while changing speed and orientation.

The results of the second experiment are shown in Figure 6. We can see that the variation of RSSI is more significant compared to the first experiment in Figure 5 due to random speed and orientation changes. However, after applying the 5.25s moving average, the smallest RSSI is -65dBm, very close to that in the first experiment. By following the moving average, we can see three peaks at time 400, 630 and 900 (in units of 250ms). These are RSSI values that correspond to the three times the person returns to position 0. The small peak at time 850 corresponds to location 7. There are similar




Figure 7: RSSI when walking away from and back into the office

peaks before time 400 and 630 according to the raw RSSI values but they are filtered out by the moving average since the mobile node is moved quickly across location 7 in the first 2 loops. However, besides these peaks, the rest of RSSI values cannot be clearly separated into tiers as in the first experiment.

Based on the RSSI moving average of the first two experiments, we can establish two potential boundaries for estimating NAPS zones in this office environment. The first boundary is at -50dBm (about 2.5s ahead of the three peaks in the second experiment). The RSSI value of this boundary is only exceeded in location 0, 1 and 7 in both experiments. The second boundary is at -68dBm , the minimum RSSI value of the two experiments.

6.3 Walking away from or back into office

Our last experiment is used to study the changes of RSSI values when the mobile node leaves or enters the office. In this experiment, a person carrying the mobile node leaves location 0 and walks into the hallway following the path shown as a dotted line in the hallway part of Figure 4 before returning to the base node. The far right end of the path is where the base node and mobile node are out of communication range. That can be easily detected since the mobile node flashes an LED for each received probe packet. For this experiment, the same walk is repeated twice at normal walking speed and the results are shown in Figure 7.

By following the moving average in Figure 7, we can see that for both walks, the RSSI value closely correlates to the distance between the two nodes. This suggests that the -68dBm boundary we identified from the first two experiments can be used to tell whether a person is inside the office or not.

We can also see in Figure 7 that the smallest raw RSSI value is -93dBm and it takes about 30s of walking time to move the nodes out of communication range. This is enough time to wake up most computers from the sleep mode when -93dBm is used as the away-zone boundary in NAPS.

6.4 Discussion

Based on the experiments, we identify four RSSI boundaries: -50dBm , -68dBm , -93dBm and $-\infty$ (out of communication range). With these boundaries, we can establish four zones for our office environment: $(+\infty, -50\text{dBm}]$ for usage-zone, $(-50\text{dBm}, -68\text{dBm}]$ for nearby-zone, $(-68\text{dBm}, -93\text{dBm}]$ and $(-93\text{dBm}, -\infty)$ for away-zone. By subdividing the away-zone, sleep and hibernation modes can be used separately for system level power management.

7. CONCLUSION AND FUTURE WORK

We present NAPS, a zone based location-aware power management scheme that takes into account application-level information. While our experimental results are promising and validate the use of RSSI for zones in an indoor environment, it opens up a number of interesting challenges. A particular challenge is pre-wake strategy needed. Computers in hibernation can be woken up using WOL by forming networks with base nodes on active computers or gateway nodes[3], or alternatively using the base node as a power switch. Furthermore, identification of zones and their mappings to power states needs to be done for diverse radio and geographical environments. While we assume a typical office environment where machines are bounded to individuals, multi-user scenarios can be supported since RSSI values are collected from individual packets that contain the MAC addresses of the senders. A prototype of NAPS is under development and our current version can manage the power of displays.

8. REFERENCES

- [1] R. K. Harle and A. Hopper. The potential for location-aware power management. In *UbiComp '08*, pages 302–311. ACM, 2008.
- [2] C. Harris and V. Cahill. Exploiting user behaviour for context-aware power management. *WiMob '05*, 4:122–130 Vol. 4, Aug. 2005.
- [3] Z. Jin, C. Schurgers, and R. K. Gupta. A gateway node with duty-cycled radio and processing subsystems for wireless sensor networks. *ACM Trans. Design Autom. Electr. Syst.*, 14(1), 2009.
- [4] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. *IPSN '05*, pages 364–369, April 2005.
- [5] A. Roy, S. Das Bhaumik, A. Bhattacharya, K. Basu, D. Cook, and S. Das. Location aware resource management in smart homes. *PerCom '03*, pages 481–488, March 2003.
- [6] C. Webber, J. Roberson, M. McWhinney, R. Brown, M. Pinckard, and J. Busch. After-hours power status of office equipment in the usa. Technical Report LBNL-57470, LBNL, 2005.
- [7] K. Whitehouse, C. Karlof, and D. Culler. A practical evaluation of radio signal strength for ranging-based localization. *SIGMOBILE Mob. Comput. Commun. Rev.*, 11(1):41–52, 2007.