Reducing Energy Consumption of a Modem via Selective Packet Transmission Delaying

Daedong Park, Philkyue Shin, Seongsoo Hong
Department of Electrical and Computer Engineering Seoul National University
Seoul, Republic of Korea
{ddpark, pkshin, sshong}@redwood.snu.ac.kr

***Abstract*—Most practical energy-saving mechanisms for a modem rely on a packet transmission delaying technique where the smartphone attempts to reduce the energy consumption of the modem by delaying transmissions of delay-tolerant packets and piggybacking them onto later packets. However, unconditional packet transmission delaying may lead to unanticipated energy loss at the modem since the exact radio resource control state of the modem is not considered. In order to address this problem, we propose a mechanism that selectively delays the packet transmissions only if such delays are expected to achieve energy savings. Our mechanism consists of three key components, which are *deferrable packet identifier*, *pattern-based next packet transmission predictor* and *packet transmission time designator*. In order to demonstrate the effectiveness, we have implemented all three components into Android 4.4 KitKat running on a Google Nexus 5 smartphone and then performed extensive experiments. Our experiments show that the proposed mechanism reduces the energy consumption of an LTE modem by up to 22.5 percent when compared to that of a legacy LTE modem.**

Keywords-modem; smartphone; energy saving; delayed transmission

I. INTRODUCTION

It is of great importance to reduce the energy consumed by a modem in order to prolong the operating time of a battery powered smartphone. In practice, a modem is reported to consume 31 to 47% of the total energy used in a smartphone [1]. Furthermore, this amount is expected to steadily increase as a result of the rapid growth in mobile data traffic [2].

The radio resource control (RRC) protocol often have a great impact on the amount of energy consumed to transmit a given amount of data. When a modem attempts to transmit a packet, it should first acquire the radio resources from the base station by using the RRC protocol. After the completion of the packet transmission, a modem releases the radio resources. If a modem needs to perform both the radio resource acquisition and the release processes whenever a packet is transmitted, the modem and the base station surely suffer overhead. In order to reduce these overhead, the RRC protocol delays the radio resource release process for a predefined time interval after completing the packet transmission while anticipating a subsequent packet transmission request. This time interval is referred to as the tail time of a modem. In the case of burst packet transmissions where multiple packets are transmitted within a short period of time, the tail time effectively reduces such overhead by avoiding repetitive radio resource acquisition and release.

On the other hand, when there is no packet transmission request within the tail time, the modem wastes energy by unnecessarily holding the radio resources during the tail time. The amount of energy wasted comprises up to 60% of the modem's energy consumption [3].

In order to rectify this problem, quite a few researchers have proposed various approaches to reducing the amount of energy wasted while retaining the advantage of the tail time [3, 4, 5, 6]. They can be classified into two categories. Approaches in the first category prematurely terminate the tail time if a subsequent packet transmission request is expected to occur only after the pre-specified time [4, 5]. However, these approaches are impractical since modems can cause severe signal overhead on base station due to frequent radio resource acquisition and release [7]. Approaches in the second category is to adjust the packet transmission times of multiple packets so as to transmit them altogether if possible [3, 6]. For this, most existing approaches attempt to identify delay-tolerant packets and to defer their transmissions. However, they simplistically delay packet transmissions without considering the exact radio resource usage state of the modem, which may lead to unanticipated energy loss of the modem.

In this paper, we propose a new packet transmission delaying mechanism that selectively delays the transmissions of deferrable packets only if such delays are expected to accrue energy savings. Our mechanism consists of three key components: (1) *deferrable packet identifier*, (2) *pattern-based next packet transmission predictor* and (3) *packet transmission time designator*. We have implemented the proposed mechanism and evaluated the energy savings on a Google Nexus 5 smartphone running Android 4.4 KitKat. The results of the experiment demonstrate that our mechanism can reduce the energy consumption of the LTE modem by up to 22.5% compared to that of a legacy LTE modem, while incur only negligible run-time overhead.

The remainder of this paper is organized as follows. In Section II, we model our target system and state our problem that needs to be solved. Section III provides the technical treatment of our approach. Section IV reports on the experimental evaluation. Finally, Section V concludes the paper.

II. System Model and Problem Description

In this section, we explain the target system model and state our problem at hand

A. System Model

We begin by modeling the packet transmission in a mobile device. In general, a modern mobile device has a strictly layered architecture. Tasks running in the application layer issue the packet transmission requests to a network stack in the kernel layer by invoking system calls. In turn, the network stack creates a header for each packet and delivers the complete packet to the underlying modem device driver. The modem device driver partitions the packets into frames and finally, the modem transmits the frames to the external network.

In order to model the packet transmission delaying mechanism, we need to clearly segregate the class of packets that become the target of the mechanism. To do so, we classify the packets using the characteristics inherited from their transmitting tasks. We first classify the packets into real-time and non-real-time packets. We further divide the non-real-time packets into interactive and non-interactive packets. A real-time packet is a packet whose transmission is subject to a timing constraint. An interactive packet is a non-real-time packet whose transmission latency affects the user-perceived response. Finally, a non-interactive packet is a packet that is free from any timing or performance constraints. We thus refer to the third type of packet as a deferrable packet. Clearly, our mechanism can safely delay the transmissions of only the deferrable packets.

B. Problem Description

In our approach, we attempt to reduce the energy that is wasted during the tail time of a modem by using a new packet transmission delaying mechanism. In order to formulate the problem that needs to be solved, we first define a metric to evaluate the mechanism. We then state our problem.

Let $S = < p_1, p_2, \ldots, p_n >$ be a sequence of packets arranged in increasing order of their request times. For the packet transmission delaying mechanism D, we define $E_D(S)$ as the amount of energy to transmit S, consumed by a modem that implements D. Similarly, we define $E_L(S)$ as the energy consumed by a legacy LTE modem that does not adopt any packet transmission delaying mechanism. Throughout this paper, we evaluate the packet transmission delaying mechanism D with $\Delta E(S)$ where $\Delta E(S) = E_L(S) - E_D(S)$. Clearly, our objective is to maximize $\Delta E(S)$ by selectively delaying the transmissions of deferrable packets.

Contrary to intuition, existing mechanisms that unconditionally postpone packet transmissions cannot guarantee $\Delta E(S) \geqq 0$ for all possible packet sequences. The reason is that they do not consider the RRC state of a modem. In order to avoid such energy loss, a new packet transmission delaying mechanism must be able to selectively delay the transmission of a deferrable packet according to the estimated value of ΔE. Thus, the sub-problems we address in this paper can be described as follows: For a currently requested packet p_c and the most recently transmitted packet p_{prev}, our packet transmission delaying mechanism (1) determines whether p_c is deferrable; (2) predicts the next packet transmission request time of a mobile device; and (3) selectively delays the transmissions of p_c to maximize ΔE.

III. Selective Packet Transmission Delaying Mechanism

This section presents the proposed packet transmission delaying mechanism that selectively delays the packet transmissions. We first explain the details of the deferrable packet identifier (DPI) and the learning and online phase of the pattern-based next packet transmission predictor (PNP). We then present our packet transmission time determination algorithm for the packet transmission time designator (PTD).

A. Deferrable Packet Identification

When a packet transmission request occurs, the DPI determines whether the requested packet is deferrable or not by using the attribute of the requesting task. In our approach, a real-time task is clearly defined as a task which has a timing constraint. An interactive task is a non-real-time task that belongs to a user-interactive task chain. A user-interactive task chain, which we borrow from [8], is a sequence of task executions that begin with a task handling an input event from a user-interactive device and end with a task writing output to a user-interactive device. The rest are defined as deferrable tasks. The DPI conservatively classifies the packet so as to avoid a false positive classification with respect to the deferrable packet class.

B. Pattern-based Next Packet Transmission Prediction

The PNP predicts the next packet transmission request time of the mobile device on online phase from the packet transmission patterns derived on learning phase.

On learning phase, the PNP first collects a series of packet transmission requests of a task and creates groups for them. In order to cover the diverse packet transmission characteristics of a task, we define three classes of patterns: (1) *periodic*, (2) *geometric* and (3) *other*. A periodic pattern is a sequence of equivalent groups where the time interval between any two consecutive elements is fixed. Similarly, a geometric pattern is a pattern where the equivalent groups occur at exponentially increasing time intervals. An other pattern is a sequence of equivalent groups occur irregularly in time. We have borrowed key ideas for the pattern detection from the exhaustive periodicity searching algorithm proposed in [6].

On the online phase, the PNP predicts the next packet transmission request time of each group found in packet requests by using the patterns that are derived, and it then selects the earliest time among them as the result of the prediction.

C. Estimated Energy Gain-based Packet Transmission Time Determination

The PTD calculates the estimated energy gain and determines whether to delay the packet or not according to the value of the gain. We first present out estimated energy gain model. For the currently requested packet p_c, the estimated energy gain model captures the energy consumed by the modem in transmitting the three packets in series: the previously transmitted packet p_{prev}, p_c itself and the subsequent packet p_{next} onto which the current packet may be piggybacked.

We now explain how to determine whether to delay the p_c or not according to the value of the gain. When a packet transmission request occurs, it first checks to see if p_c is deferrable. If so and the estimated energy gain is positive, it puts p_c into a queue. Otherwise, it transmits all delayed packets in the queue with p_c.

IV. Implementation and Experimental Evaluation

We first describe the implementation and experimental setup. We then present the experimental results and analyze them.

A. *Implementation and Experimental Setup*

We have run our implementation on a Google Nexus 5 smartphone running Android 4.4 KitKat with the 4G LTE carrier network provided by Korea Telecom (KT). We have performed our measurement with a Monsoon Power Meter to measure the power consumption between a smartphone and its battery. As a metric for evaluating our approach, we use the energy gain $\Delta E(S)$ of our mechanism.

We have conducted three different experiments: two to evaluate the energy savings and one to evaluate the run-time performance overhead. In the first experiment, we simultaneously ran ten applications and measured the energy consumption of the target smartphone with and without our mechanism. In the second experiment, we simulated the state transition and energy consumption of a modem by using packet transmission logs that are collected during the real usage of a smartphone. In the third experiment, we measured the blocking time of the system call handlers during the experiment comparing the results with those of a legacy smartphone.

B. *Experimental Results*

In our first experiment, we measured the energy consumption of the target smartphone for 20 minutes of the online phase after two hours of the learning phase. TABLE I summarizes the data measured in this experiment. It shows that our mechanism saves 22.5% of energy.

For the second experiment, we collected logs from two male and one female participants aged between 25 and 30 in daily life. The first half of the packet transmission logs was used for pattern detection, and the second half to calculate $\Delta E(S)$. TABLE II summarizes the results. It shows that our mechanism saves 12.8% of the energy on average.

In the last experiment, we used the same workload and the tunable parameters that were used in the first experiment. During 20 minutes of an online phase, the target smartphone was blocked for 1.99s while the legacy smartphone was blocked for 1.97s. The run-time performance overhead that was incurred by executing our algorithms was only up to 1.0%.

TABLE II. Summary of the first experiment

Applications	Facebook, BBC News, Bloomberg, ESPN, Naver Line, Naver Map, Naver Webtoon, Naver Music, Daum tvPot, Kweather	
Energy consumption of modem (mWh)	W/ our approach	W/o our approach
	141.4	182.4
$\Delta E(S)$ (mWh)	41.0	
Energy savings (%)	22.5	

TABLE I. Summary of the second experiment

	Hours of use (h)	Energy consumption of modem (J)		$\Delta E(S)$ (J)	Energy savings (%)
		W/ our approach	W/o our approach		
#1	97.6	30,718	35,007	4,289	12.3
#2	159.4	13,340	15,508	2,168	14.0
#3	7.9	4,463	5,070	607	12.0

V. Conclusion

In this paper, we have proposed an energy-saving mechanism for a modem that exploits selective packet transmission delaying. The proposed mechanism was composed of the three key components: DPI, PNP and PTD. When a packet transmission request occurs, the DPI checks whether the requested packet is deferrable or not. For a deferrable packet, the PTD calculates the estimated energy gain by using the next packet transmission request time that is predicted by the PNP. The PTD finally determines the transmission time of the packet by using the gain. In order to show the effectiveness of the proposed mechanism, we evaluated it with extensive experiments. The results indicate it achieved up to 22.5% savings in energy consumption for an LTE modem relative to legacy LTE modems.

The present research can be further extended into new research directions. We are looking to elaborate on the packet transmission patterns of a task and the detection algorithm of the patterns using machine learning techniques. We also plan to extend our pattern-based prediction approach to take into account the incoming packets as well as the outgoing ones.

Acknowledgment

This work was supported by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (No. 2016-0-00170, Development of wide area driving environment awareness and cooperative driving technology which are based on V2X wireless communication) and the Ministry of Science, ICT and Future Planning (MSIP), Korea, under the Information Technology Research Center (ITRC) support program (IITP-2017-2013-0-00717) supervised by the Institute for Information & communications Technology Promotion (IITP).

References

[1] A. Carroll and G. Heiser, "An Analysis of Power Consumption in a Smartphone," USENIX annual technical conference, Vol. 14, 2010.

[2] Cisco, “Cisco visual networking index: Global mobile data traffic forecast update,” 2014-2019, Feb 2015.

[3] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, "Energy consumption in mobile phones: a measurement study and implications for network applications," Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference, 2009.

[4] P. K. Athivarapu, et al., "Radiojockey: mining program execution to optimize cellular radio usage," Proceedings of the 18th annual international conference on Mobile computing and networking, 2012.

[5] K. Yeseong, J. Boyeong, and K. Jihong, "A Personalized Network Activity-Aware Approach to Reducing Radio Energy Consumption of Smartphones," IEEE Transactions on Mobile Computing 15.3, pp. 544-557, 2016.

[6] F. Qian, et al., "Periodic transfers in mobile applications: network-wide origin, impact, and optimization," Proceedings of the 21st international conference on World Wide Web, 2012..

[7] 3GPP, "System impact of poor proprietary fast dormancy," 3gpp discussion and decision notes rp-090941 2009.

[8] H. Sungju, Y. Jonghun, and H. Seongsoo, "Improving interactivity via vt-cfs and framework-assisted task characterization for linux/android smartphones," Embedded and Real-Time Computing Systems and Applications (RTCSA), 2012.