

LLAMA: An Adaptive Strategy for Performing Background Tasks on Mobile Devices

Denitsa Tilkidjieva, Nilanjan Banerjee[†], Maria Kazandjieva, Sami Rollins, Mark D. Corner[†]
Mount Holyoke College, [†]University of Massachusetts, Amherst
{dntilkid, makazand, srollins}@mtholyoke.edu, {nilanb, mcorner}@cs.umass.edu

1. Overview

Most computing applications are user-driven; a user must explicitly instruct a device to perform a particular task. To reduce latency and intrusiveness, some applications will act on a user's behalf even when the user is not actively using the device. For example, e-mail clients often periodically check for new mail and virus scanners can be scheduled to run during idle periods. This model is okay for devices that are always connected to a network and power supply, but it is not suitable for mobile devices that meet neither property.

When they are not actively used, mobile devices often go into a power-saving state. To perform periodic tasks on behalf of the user, applications would have to wake the device which would require large amounts of energy.

The goal of this work is to enable mobile devices to perform background tasks on behalf of the user by adaptively determining the amount of energy devices can spare. We consider the amount of energy remaining on the device when it begins to recharge, estimate the amount of energy that can be spent, and devote it to background tasks.

2. Analysis

This work presents Longer-Lifetime Algorithms for Mobile Applications (LLAMA), a middleware layer for mobile devices. LLAMA determines when a device should run background tasks based on previous battery usage and availability. At that time, it wakes the device from a power-saving state if necessary, allows applications to perform their tasks, and returns the device to the previous power state.

To test our adaptive strategy, we run a simulated experiment, using real battery traces collected from 13 devices over a two-month period (February 16 – April 15, 2006). We vary the confidence with which we can assure that we will not use more than the device's available energy and measure the total number of

background tasks performed as a percentage of the ideal number of background tasks that would be performed if we used a periodic algorithm (i.e. if we wake the device every five minutes).

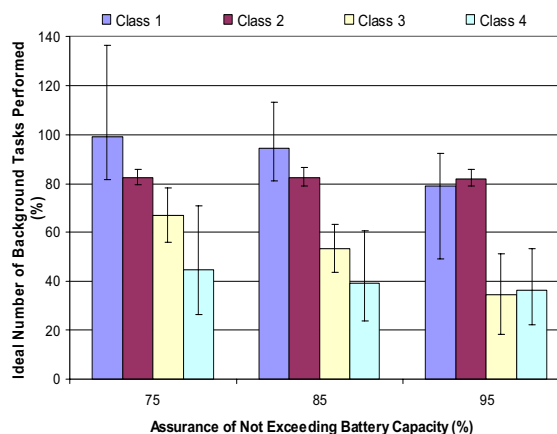


Figure 1. Performance of adaptive strategy on three confidence levels

Figure 1 plots the percentage of background tasks completed for three confidence levels. We group the 13 devices based their level of mobility, Class 1 being least mobile and Class 4 being most mobile. The average for each class is displayed; the error bars indicate the minimum and the maximum values for each class. We observe that, for all classes, increasing the confidence decreases the percentage of background tasks performed. This is expected; when the confidence is high, the adaptive strategy is more conservative, waking the device less often to perform tasks. All devices complete at least 18.37% of the ideal number of background tasks, compared to 0% if the adaptive strategy was not used.

Our strategy enables devices with different levels of mobility to adaptively use excess energy and support applications such as e-mail and RSS feed fetching, virus scan, file backup, and adjusting screen brightness.