Aalto University
School of Science
and Technology

Some results on power modeling of mobile devices

15.02.2011

Matti Siekkinen, Aalto University, Data communications software

Joint work with Yu Xiao, Rijubrata Bhaumik, Arto Karppanen, Petri Savolainen, Zhirong Yang, Antti Ylä-Jääski

Outline

- Introduction
- Practical power power modeling of data transmission over 802.11g
- System-level modeling for runtime power estimation on mobile devices
- Conclusions

Introduction

- Increasing gap between
  - mobile phone battery capacity and
  - energy consumption by typical usage
- Our focus: understand the energy consumption
  - Measurements
  - Modeling
- Why bother?
  - Energy-aware protocols and services
  - Smarter device power management

Methodology

- Power measurements
  - Software
  - Hardware
- Power modeling
  - Straightforward simplified and practical models
  - We care mostly about fairly rough numbers
    - Only significant energy savings make a difference
- Apply models
  - Energy aware application
  - Energy-efficient protocols

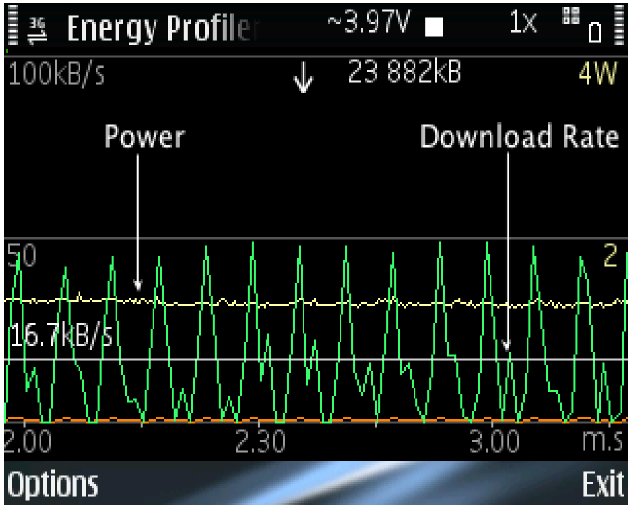

Power modeling of transfer over 802.11

- 802.11 Wireless network interface (WNI)
  - Power consumption depends on operating mode

Energy = Power(operating mode) Duration(operating mode)*

  - No common open API for getting this info

- Estimate the operating modes & durations from observed traffic patterns
  - Take into account 802.11 Power Saving Mode(PSM)

WNI mode	Average Power (W)		
	Nokia N810	HTC G1	Nokia N95
IDLE	0.884	0.650	1.038
SLEEP	0.042	0.068	0.088
TRANSMIT	1.258	1.097	1.687
RECEIVE	1.181	0.900	1.585

Per-burst computation

- TCP style bursty traffic naturally fits to this model
- Reduce computational work
  - If CBR -> use 1 pkt bursts
- Burst/bin definitions
  - Pkt interval < t
  - Bin rate $r = S_B/T = S_B/(T_B+T_I)$

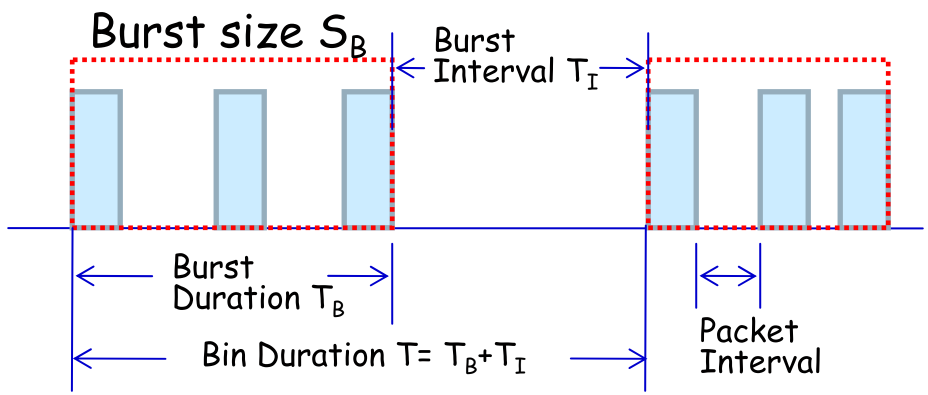

- PSM timeout effect
  - Threshold bin rate $r_c = S_B/(T_B + T_{timeout})$
- Scenario 1: {{r >= r_c } and {PSM is enabled}} or {PSM is not enabled}
  - No time to sleep in between bursts
  - Alternate between Rx/Tx and idle

$$P(r_d) = P_I+[T_d(P_R-P_I)+T_u(P_T-P_I)]r_d/S_B$$

- Scenario 2: { r < r_c } and {PSM is enabled}

$$P(r_d) = P_S+[T_d(P_R-P_S)+T_u(P_T-P_S)+T_{timeout}(P_I-P_S)]r_d/S_B$$

Validation

- Setup a local testbed
- Experiments with a few devices
  - Nokia N95, N810
  - HTC G1 (Dream)

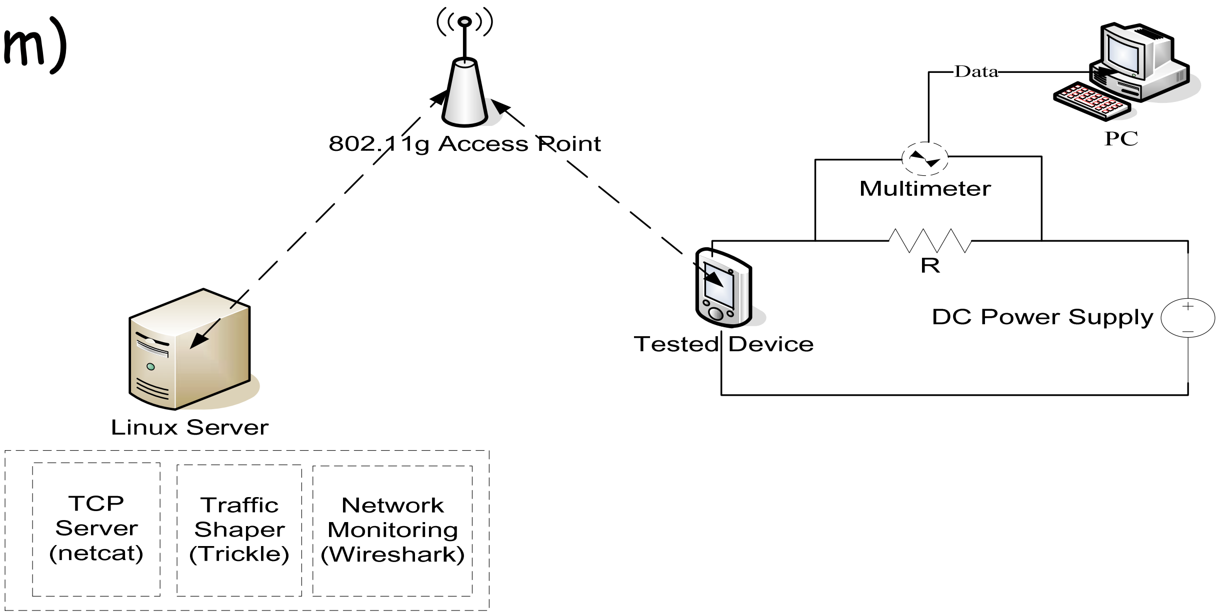

Results

Downloads: MAPE less than 7%

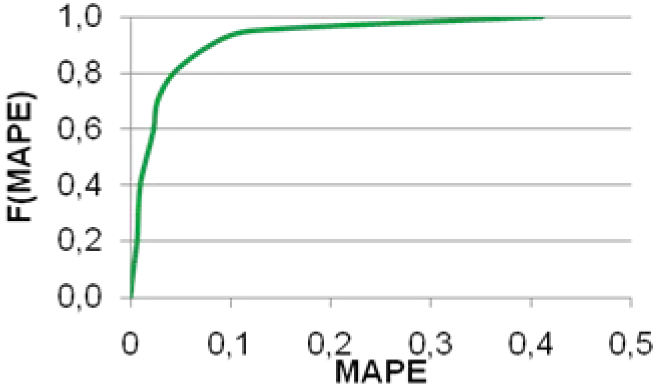

Uploads: MAPE less than 6%

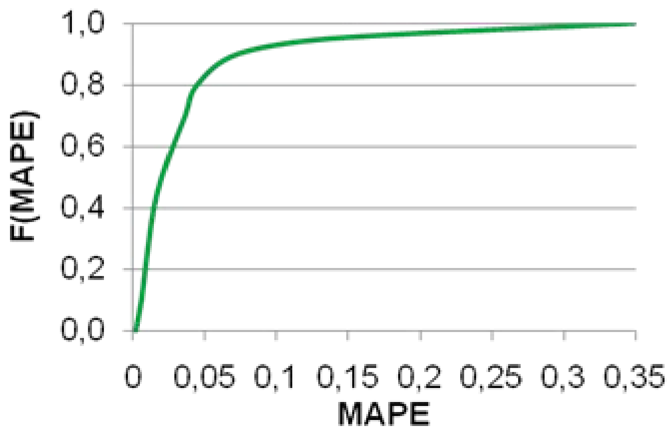

- Also tried run time estimation
  - YouTube download through public AP
  - MAPE about 11%
    - Most likely due to additional e.g. broadcast traffic

Limitations and possibilities

- Model does not “see” L2 traffic
  - Has an effect on power consumption
- Model does not separately consider computation caused by packet processing
  - Implicitly included in baseline measurements
- Could be incorporated to network simulators
  - Useful for protocol developers
- Enable energy-aware applications
  - E.g. energy-efficient peer/mirror server selection

System-level power modeling on mobile devices

- Cannot always directly measure power
  - Software unavailable
  - Hardware measurements impractical
- A system-level power model
  - Include main components
    - Processors, wireless network interfaces, display
  - Build using system-level measurements
    - Difficult to measure components separately
  - Generic model
    - Not tied to specific application
- We apply linear regression

Linear Regression with Nonnegative Coefficients

- Linear regression widely used for processor power modeling
  - Build a linear relationship between the *p* predictor variables and power consumption based on *n* observations

$$f(y_i) = \beta_0 + \sum_{j=1}^{p} \beta_j g_j(x_{i,j})$$

  - Hardware performance counters (HPC) as variables
- Challenge 1: mobile devices can only monitor a subset of HPCs at a time
  - E.g., 3 out of 17 HPCs in Nokia N810
  - Can we still apply this method with reasonable accuracy?
    - Yes, reduce the set of HPCs using nonnegative coefficients constraint
- Challenge 2: Network and display also consume power
  - We include them into the model

Regression variables

- Target: reflect resource consumption of a mobile application
- Local computing
  - 17 HPC-based event rates
    - CPU activity, memory access
- Network I/O
  - Download data rate
  - Upload data rate
  - CAM switch
    - Effect of 802.11 power saving mechanism
- Display
  - Brightness level

Benchmarks & model fitting

- Goals:
  - Stress all selected variables
  - Explore the space of their cross product
- Nokia Internet Tablet N810
- Different categories of workloads:
  - Idle with different brightness levels
  - Audio/video players/recorders
  - File download/upload at different (limited) transfer rates
  - Audio/video streaming
- Model fitting with non-negative least squares: $S(\beta_0, \ldots, \beta_p) = \sum_{i=1}^{n} (y_i - f(y_i))^2$
  - Matlab function lsqnonneg
- Outcome:

$$Power(W) = 0.7655 + 0.2474 \times g_0(x_0) + 0.0815 \times g_1(x_1) + 0.0606 \times g_2(x_2) + 0.0011 \times g_{17}(x_{17}) + 0.0015 \times g_{18}(x_{18}) + 0.3822 \times g_{19}(x_{19}) + 0.125 \times g_{20}(x_{20}).$$

CPU_CYCLES
DCACHE_WB
TLB_MISS
dl rate
ul rate
CAM switch
brightness level

Model evaluation

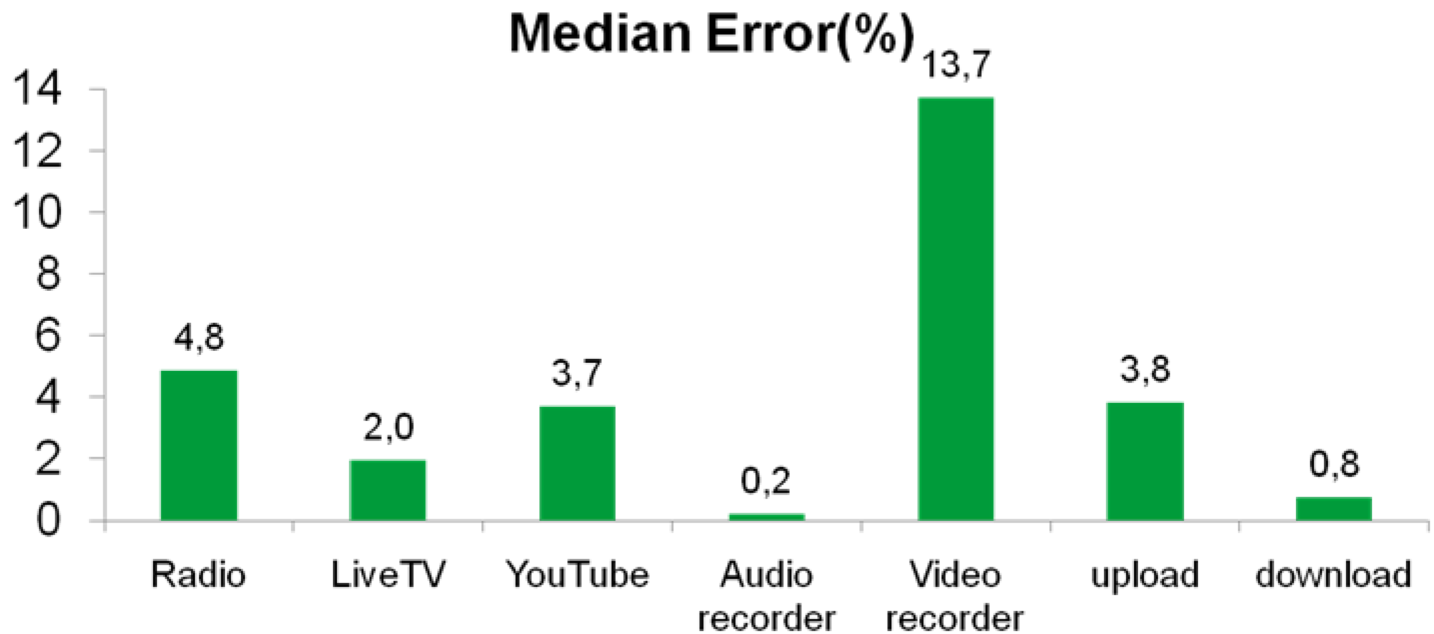

- Evaluation based on similar data used for model fitting
  - But not the same
- Similar (or even better) results with another independent data set
  - Different set of applications

Limitations and further work

- Can we extend to a *per-process model?*
  - HPCs can be monitored per process
  - Traffic can be analyzed for each process
  - But power can not be measured for each process separately
    - Necessary for validation purposes
- Could we incorporate this model into simulators?
  - E.g. enhanced node model for network simulators

Conclusions

- Different levels of modeling
  - Complementary work
- We believe this is important activity
  - Enable energy-aware applications
  - Facilitate energy-efficient protocol development
- More future work
  - 3G (WCDMA) suffers from long wasted “tail” energy, how about upcoming LTE?
  - Apply the models