



Mobile and Ubiquitous Computing: Informed Mobile Prefetching

Brett Levasseur

*Computer Science Dept.
Worcester Polytechnic Institute (WPI)*






Introduction

- Where is data coming from?

- CPU Cache
- RAM
- Disk
- Networks
 - Optical
 - Copper
 - Wireless



- Prefetching data important to improve user experience



Introduction


- Common Prefetching applications
 - Databases
 - File systems
 - Distributed systems
 - HTML 5 comes with a prefetch Link type
- Mobile Device Prefetching
 - Fetches data from networks often
 - Normally use low bandwidth & high latency networks
 - Prefetching avoids network problems and latency with on-demand network use



Mobile Considerations

- Performance
 - Can't interfere with other user activity
 - Wireless conditions change cost
 - Class of data / Type of app
- Power aware
 - Network activity strong pull on battery
 - *"Majority of power consumption can be attributed to the GSM module and the display"* An Analysis of Power Consumption in A Smartphone
- Data Consumption
 - Extra charges for using too much data


What to do?




VS




VS






Proposal

- Add prefetch support to the mobile OS
- Informed Mobile Prefetching (IMP)
 - Library to support prefetching for mobile apps
 - Balance data fetched with resources available
 - Power Resources
 - Data Resources



Related Work

- Transparent Informed Prefetching (TIP)
 - Cost-benefit analysis informed fetching from disk arrays
- Intentional Networking
 - Label traffic type and network statistics inform choice on how to use the network
- Odyssey's Goal-Directed Adaptation
 - Applications modify behavior to conserve energy




Mobile Notes

- Performance
 - Measure benefit and impact costs
- Energy Use
 - Signal quality changes power use
 - WiFi uses less power than cellular network
- Cellular Data Usage
 - Cellular data limits
 - WiFi possible free data use



Methodology

- Adaptive management of budgeted resources
 - Conversion rates to compare power and network resources
 - Importance of a resource changes
 - Control loop changes conversion rate of budgeted resource
- Prefetch based on budget findings
 - Determine when and how to best retrieve data
 - Retrieve data in background
 - Does not interfere with other active applications



Cost/Benefit Decisions

- Inspired from TIP
 - App hints to indicate future data access
- Benefit dependent on
 - Size of data
 - Network conditions
- Cost without prefetch

$$T_{fetch} = \frac{S}{BW_{future}} + L_{future}$$



Fetch Cost

- Use past network data to approximate future conditions
 - Track average availability, latency and bandwidth
 - Uses active network measurement and passive measurements when data is prefetched or fetched
- Cost to fetch data over cellular and WiFi

$$T_{\text{fetch-WiFi}} * \text{Availability}_{\text{WiFi}} + T_{\text{fetch-cellular}} * (1 - \text{Availability}_{\text{WiFi}})$$



Prefetch Accuracy

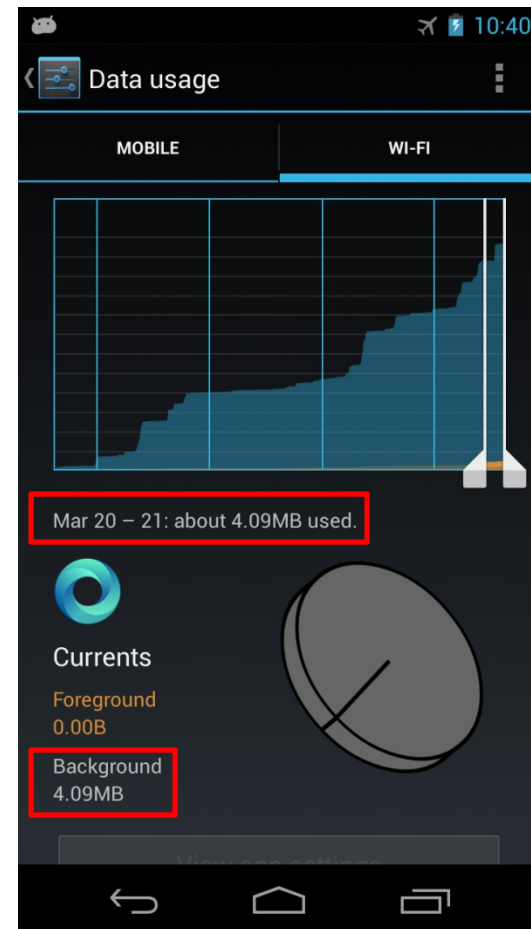
- Calculate accuracy of prefetch hints per app or classes within app


$$accuracy = hints_{consumed} / hints_{total}$$

- $hints_{total}$ incremented for each hint provided by app
- $hints_{consumed}$ incremented when app requests prefetch data
- Hints not prefetched tracked by checking if an app forces a fetch for data that was requested through prefetch but not yet retrieved

Accuracy Counts

- Currents – Google news reader & aggregator
- App was **not used** from March 20th-21st
- 4.09MB downloaded
- Rarely use app





Energy Use

- Compare energy needed to prefetch now with fetching later on demand
- T_{prefetch} calculated like T_{fetch} but with current conditions for each network (cell and WiFi)

$$T_{\text{prefetch}} = (S / BW_{\text{now}}) + L_{\text{now}}$$

- PowerTutor used to calculate energy cost of prefetch and fetch
 - Specific to hardware and carrier



Energy Use Cont.

- WiFi – Uses power coefficient $P_{\text{WiFi-xmit}}$ or power to send and receive on WiFi

$$E_{\text{prefetch}} = P_{\text{WiFi-xmit}} * T_{\text{prefetch}}$$


- 3G – Stays in high power state after transmission completes

$$E_{\text{prefetch}} = (P_{\text{3G-xmit}} * T_{\text{prefetch}}) + E_{\text{tail}}$$

$$E_{\text{tail}} = \begin{cases} P_{\text{tail}} \times T_{\text{tail}} & \text{if transition occurs and tail begins} \\ P_{\text{tail}} \times T_{\text{inactivity}} & \text{if tail time is extended} \end{cases}$$

- Net cost to prefetch

$$E_{\text{prefetch}} - (E_{\text{fetch}} * \text{Accuracy})$$



Data Consumed

- Estimate the cost to fetch data on cell plan

$$D_{fetch} = S * (1 - Availability_{WiFi})$$

- If WiFi available $D_{fetch} = 0$ and if not $D_{fetch} = S$
- Net cost to prefetch

$$D_{prefetch} - (D_{fetch} * Accuracy)$$



Compare Measurements

- Calculation values in seconds, Joules and bytes
- Odyssey's goal-directed adaptation adjusts conversion rates for these metrics
 - Once a sec remaining supply of resource checked
 - Subtract 5% of remaining and 1% of original

$$C_{adjustment} = \frac{estimated_demand}{reduced_supply} \quad C_{new} = C_{old} * C_{adjustment}$$

- Used to calculate conversions for data and energy



Decision

- Each network calculates benefit vs cost

$$T_{fetch} \times Accuracy - (c_{energy} \times (E_{prefetch} - Accuracy \times E_{fetch}) + c_{data} \times (D_{prefetch} - Accuracy \times D_{fetch}))$$

- Prefetch over the network with a positive value or if both positive prefetch over either



Implementation



- IMP implemented as an Android Java library

Application → IMP	IMP → Application
<code>prefetch(Fetcher)</code> → Future <code>Future.get()</code> → Result <code>Future.cancel()</code> <code>Fetcher.setSize(int)</code> <code>Fetcher.setPrefetchClass(int)</code>	<code>Fetcher.fetch(Labels)</code> → Result

- Hints provided through prefetch call
- Calling “get” retrieves the data
 - If prefetched it is available
 - If not then IMP makes the call on demand



Evaluation Apps


- K9 email client 
 - Used IMAP proxy to intercept traffic to server
 - Proxy downloads email headers
 - Decides which emails to prefetch and issues hints
- OpenIntents News Reader 
 - Atom/RSS feed reader
 - Modified Apache HTTPComponents
 - Prefetch link contents from feed summary
 - Made version with and without prefetch classes



Evaluation Hardware

- Apps run on Nexus One running Android 2.3.4 over AT&T
 - Modified Android to allow using either WiFi or cellular
 - Added Intentional Networking
- Used isolated WiFi and private Cisco MicroCell
- All traffic passes through computer to emulate network conditions
- Used private servers for the app data (email, news articles)

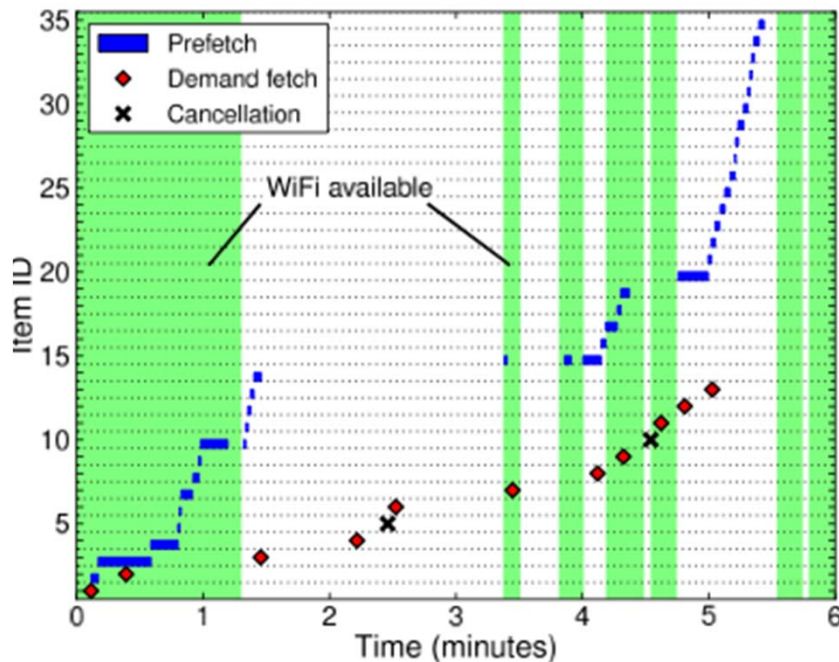
Evaluation Schemes & Measurements



- Compare IMP to other schemes
 - Never-prefetch, Always-prefetch, Size-limit, WiFi-only
 - Other schemes allowed to use Intentional Networking
- Measure cellular data usage with Linux sysfs interface
- Measure power use with PowerTutor model for Nexus One
- Collected example conditions through driving and walking traces



Example Trace



Email prefetch and demand fetch activity for one run of IMP, energy-constrained, on the driving trace. The horizontal bars show issued prefetches. The diamonds show demand fetches. The shaded regions show periods when WiFi is available. Each canceled prefetch hint is marked with an X. Since prefetching finishes early in the experiment, only the first six minutes are shown.

- IMP with data constraint
- Example fetches, prefetches, some canceled
- Set of batch prefetches at end




Evaluation Test Data


- Email
 - Day long email traces
 - 35 emails, 28 read, 7 deleted
 - 32 KB threshold
- News Reader
 - 25 articles over 5 feeds
 - Read rate varies by feed to a total of 64% of articles read
 - 128 KB threshold
- 20 minute benchmarks




Email Driving Trace



Average fetch time for 28 emails fetched over 20 minutes. The left set of bars shows results for simple heuristic-based prefetching strategies; the right set of bars shows results for IMP when energy, cellular data, or both are constrained. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.





Total energy usage for the email benchmark. The left set of bars shows results for simple heuristic-based prefetching strategies; the right set of bars shows results for IMP when energy, cellular data, or both are constrained. Where there is an energy budget set for IMP, it appears as a solid line above the bar. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.




Total cellular data usage for the email benchmark. The left set of bars shows results for simple heuristic-based prefetching strategies; the right set of bars shows results for IMP when energy, cellular data, or both are constrained. Where there is a data budget set for IMP, it appears as a solid line above the bar. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.

Energy Limit: 300 Joules
Data Limit: 2 MB
Both: 325 Joules & 2 MB


Email Walking Trace



Average fetch time for 28 emails fetched over 20 minutes. The left set of bars shows results for simple heuristic-based prefetching strategies; the right set of bars shows results for IMP when energy, cellular data, or both are constrained. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.



Total energy usage for the email benchmark. The left set of bars shows results for simple heuristic-based prefetching strategies; the right set of bars shows results for IMP when energy, cellular data, or both are constrained. Where there is an energy budget set for IMP, it appears as a solid line above the bar. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.

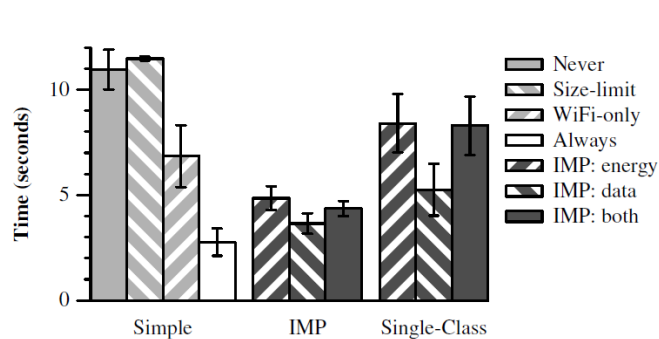


Total cellular data usage for the email benchmark. The left set of bars shows results for simple heuristic-based prefetching strategies; the right set of bars shows results for IMP when energy, cellular data, or both are constrained. Where there is a data budget set for IMP, it appears as a solid line above the bar. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.

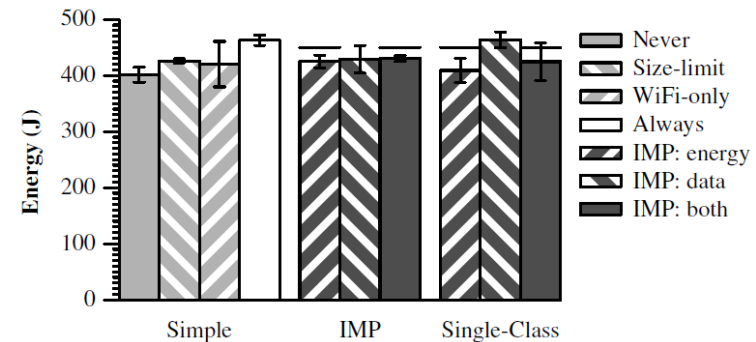
Energy Limit: 150 Joules
Data Limit: 2 MB
Both: 150 Joules & 2 MB



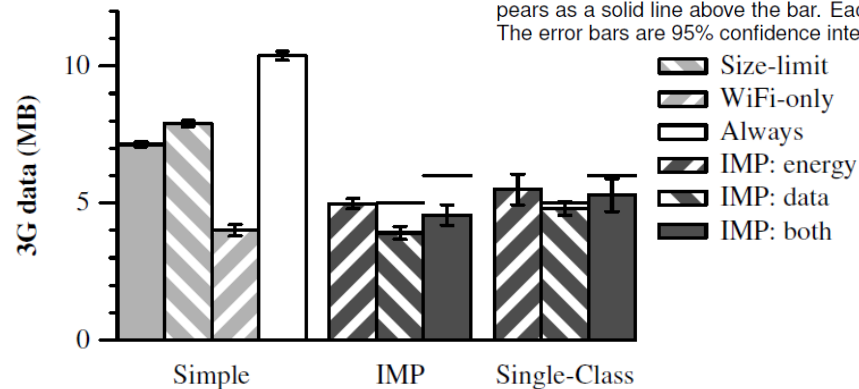
News Reader Driving Trace



Average news article fetch time for 16 articles fetched over 20 minutes. The left set of bars shows results for simple heuristic-based prefetching strategies. The middle set of bars shows results for IMP when energy, cellular data, or both are constrained. The right set of bars shows results for IMP with the same constraints when IMP does not separate prefetch hints by class. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.



Total energy usage for the news reader benchmark. The left set of bars shows results for simple heuristic-based prefetching strategies. The middle set of bars shows results for IMP when energy, cellular data, or both are constrained. The right set of bars shows results for IMP with the same constraints when IMP does not separate prefetch hints by class. Where there is an energy budget set for IMP, it appears as a solid line above the bar. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.

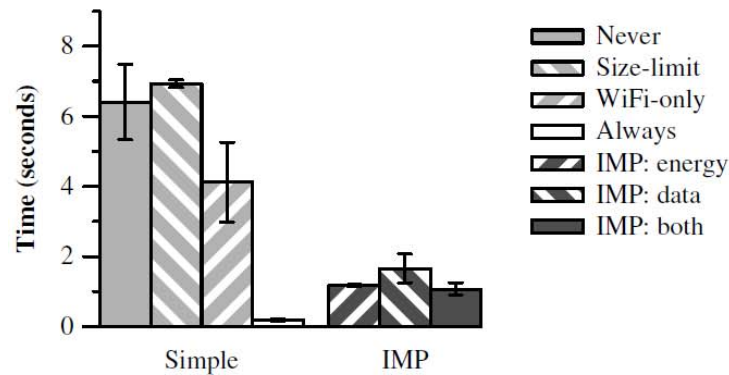


Total cellular data usage for the news reader benchmark. The left set of bars shows results for simple heuristic-based prefetching strategies. The middle set of bars shows results for IMP when energy, cellular data, or both are constrained. The right set of bars shows results for IMP with the same constraints when IMP does not separate prefetch hints by class. Where there is a data budget set for IMP, it appears as a solid line above the bar. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.


Energy Limit: 450 Joules
Data Limit: 5 MB
Both: 450 Joules & 6 MB



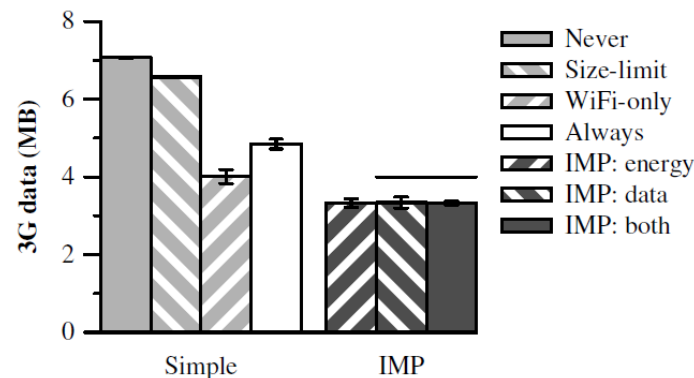
News Reader Walking Trace



Average news article fetch time for 16 articles fetched over 20 minutes. The left set of bars shows results for simple heuristic-based prefetching strategies. The right set of bars shows results for IMP when energy, cellular data, or both are constrained. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.




Total energy usage for the news reader benchmark. The left set of bars shows results for simple heuristic-based prefetching strategies. The right set of bars shows results for IMP when energy, cellular data, or both are constrained. Where there is an energy budget set for IMP, it appears as a solid line above the bar. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.



Total cellular data usage for the news reader benchmark. The left set of bars shows results for simple heuristic-based prefetching strategies. The right set of bars shows results for IMP when energy, cellular data, or both are constrained. Where there is a data budget set for IMP, it appears as a solid line above the bar. Each bar is the mean of 5 trials. The error bars are 95% confidence intervals.


Energy Limit: 200Joules
Data Limit: 4 MB
Both: 200 Joules & 4 MB



Conclusions

- Always-Prefetch best during walking with energy constraints for the News Reader
- All other cases IMP is best

Test	Constraints	Avg Fetch to Allways-Prefetch (within)	Avg Fetch to Never, Size and WiFi Only Prefetch Strategies	Energy Reduction	3G Data Reduction
Email Driving	Energy	200ms	2-8x	21-43%	NA
	Data	410ms	2-7x	NA	NA
	Both	240ms	2-8x	9-38%	3x
Email Walking	Energy	40-150ms	NA	30-65%	NA
	Data	40-150ms	NA	NA	2-4x
News Driving	Energy	NA	29-58%	NA	NA
	Data (single-class)	NA	47-68%	NA	45-62%
	Data (multi-class)	NA	42-47% (multi-class better than single)	NA	NA
	Both	NA	36-62%	NA	NA
News Walking	Energy	NA	2-6x	25-35%	NA
	Data	NA	2-6x	NA	17-53%



Future Work

- Pay as you go data plans
 - Different structure to determine network constraints
- Cache space on device
 - Assumed unlimited here but could be a potential issue
- Network throttling
 - Can't be detected my checking network strength



References

- Berjon, R., Leithead, T., Navara, E. D., O'Connor, E., Pfeiffer, S. HTML5 A vocabulary and associated APIs for HTML and XHTML. <http://www.w3.org/TR/html5/>, December 17, 2012.
- Carroll, A., Heiser, G. *An Analysis of Power Consumption in A Smartphone*. In Proc. Usenix 2010
- Flinn, J., Satyanarayanan, M. Managing battery lifetime with energy-aware adaptation. *ACM Transactions on Computer Systems (TOCS)* 22, 2 (May 2004), 137–179.
- Higgins, B. D., Flinn, J., Gluli, T. J., Noble, B., Peplin, C., Watson, D. Informed Mobile Prefetching, *In MobiSys'12*, June 25–29, 2012.



References Cont.

- Higgins, B. D., Reda, A., Alperovich, T., Flinn, J., Gluli, T. J., Noble, B., and Watson, D. Intentional networking: Opportunistic exploitation of mobile network diversity. In *Proceedings of the 16th International Conference on Mobile Computing and Networking* (Chicago, IL, September 2010), pp. 73–84.
- Patterson, R. H., Gibson, G. A., Ginting, E., Stodolsky, D., and Zelenka, J. Informed prefetching and caching. In *Proceedings of the 15th ACM Symposium on Operating Systems Principles* (Copper Mountain, CO, December 1995), pp. 79–95.