

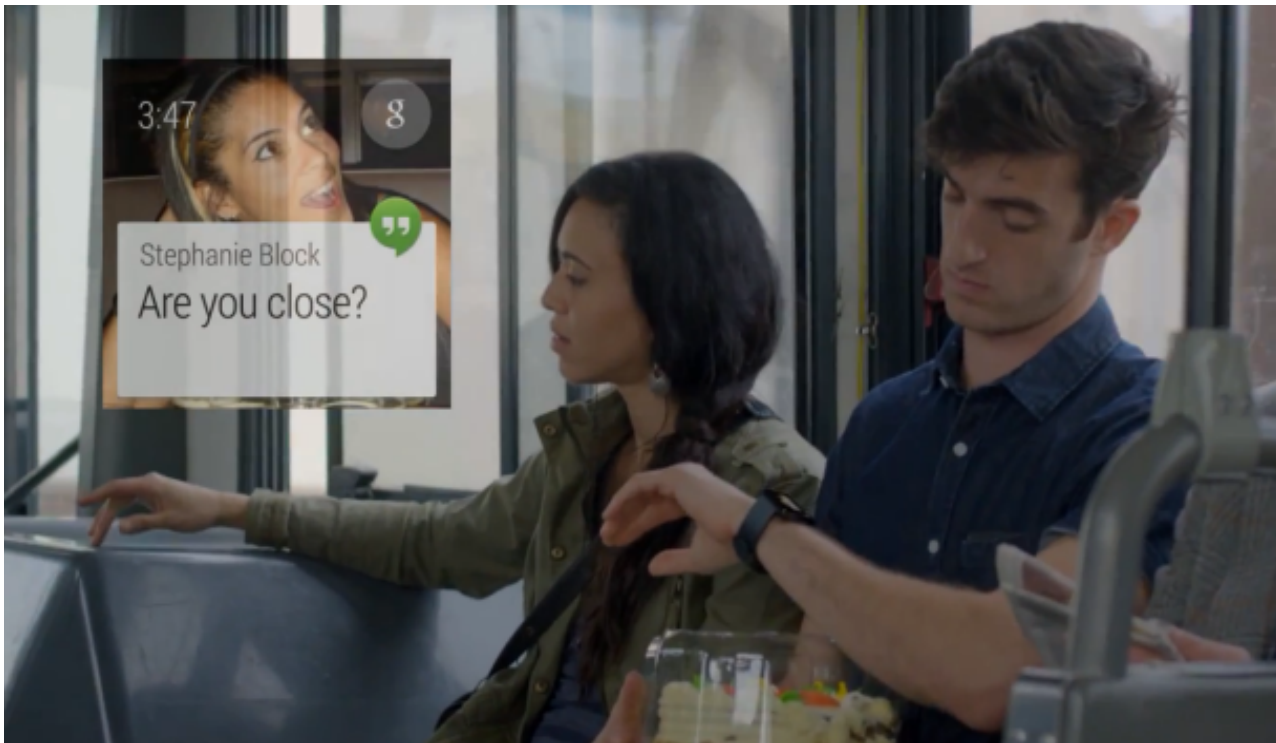
# Understanding the Characteristics of Android Wear OS

Renju Liu and Felix Xiaozhu Lin

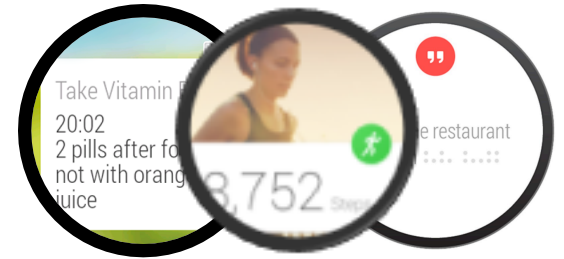
**Purdue ECE**

**PURDUE**  
ENGINEERING





## The Wearable stack



# Top questions

- Wearables should enjoy
  - Baremetal performance
  - Baremetal efficiency
- In this talk: **Android Wear**
  - Are we close to baremetal?
  - What is going on *inside*?
  - How should the OS evolve?

# Observation -- Symptoms

- The current performance & efficiency are far from baremetal
- **Pacing – inefficient**
  - face update: 400ms 88% busy

Clock face update




# Observation -- Symptoms


- The current performance & efficiency are far from baremetal

Launch “settings”

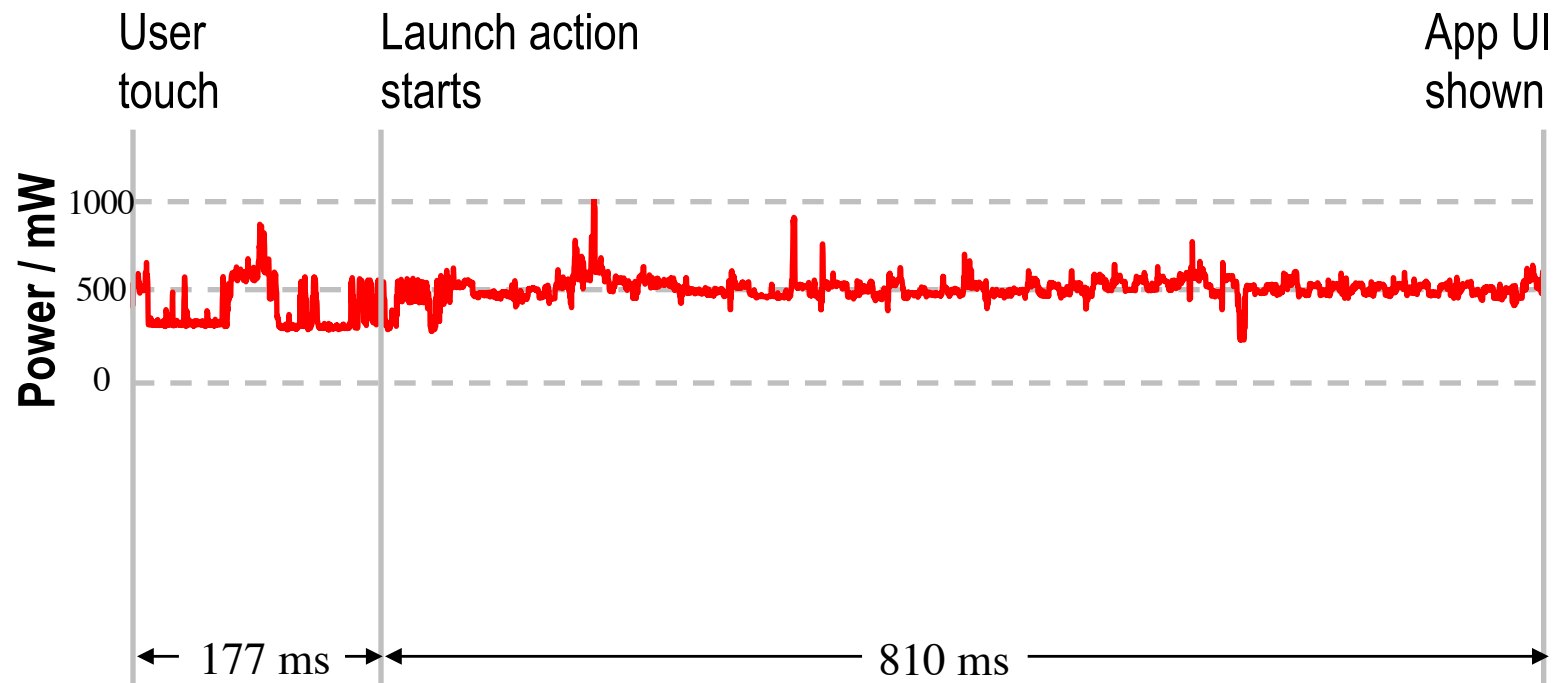
- **Pacing** – inefficient
  - face update: 400ms 88% busy
- **Racing** – slow
  - Launch an in-mem app: 1 sec



# What happens underneath?



# What happens underneath?




# What happens underneath?





# What happens underneath?



# Four Aspects

**CPU busy?**


**CPU idle?**

**Thread-level parallelism (TLP)**


**Microarchitectural behaviors**


*Won't talk about our methodologies*



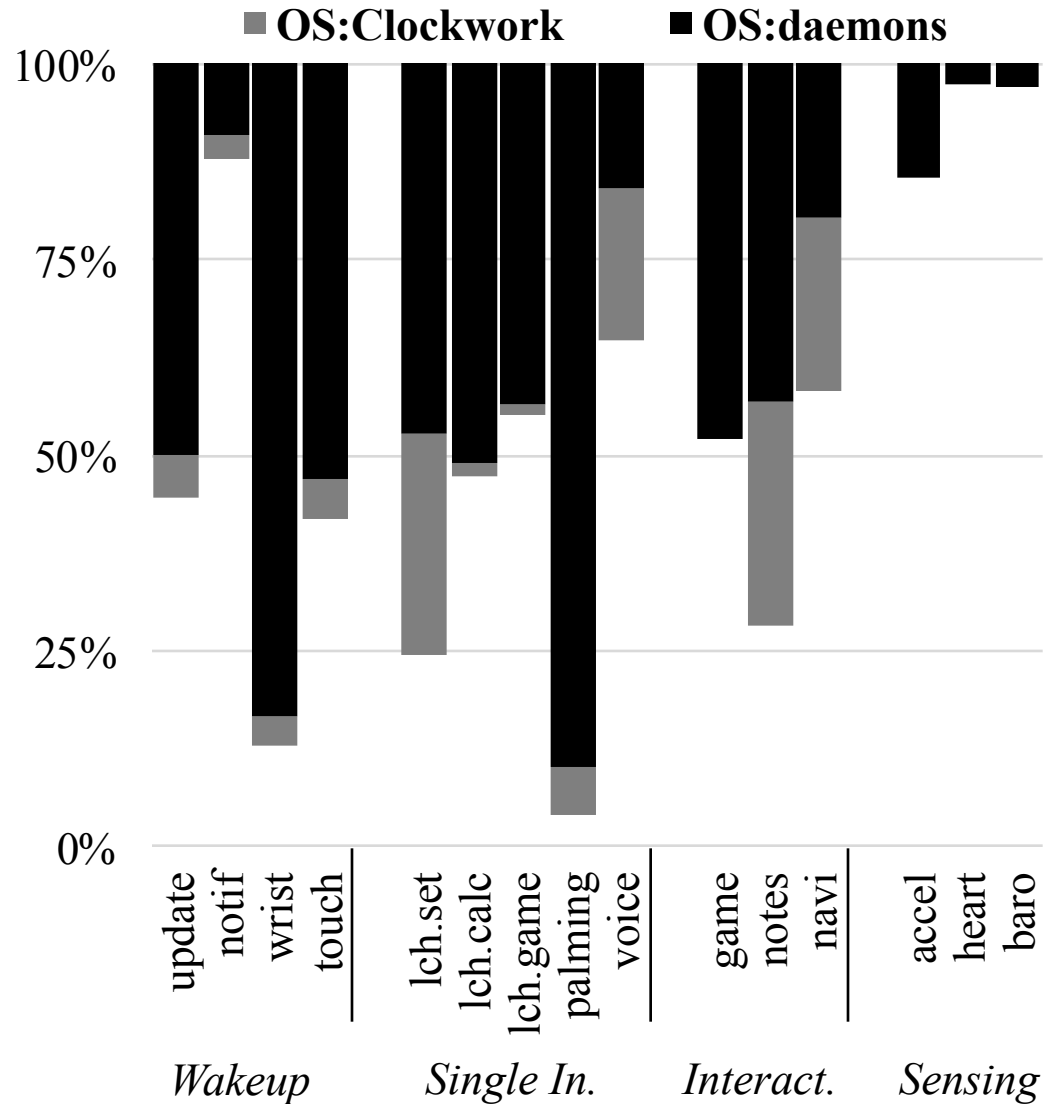



# OS execution dominates CPU usage.







# OS execution dominates CPU usage.







# OS execution dominates CPU usage.






# OS execution dominates CPU usage.






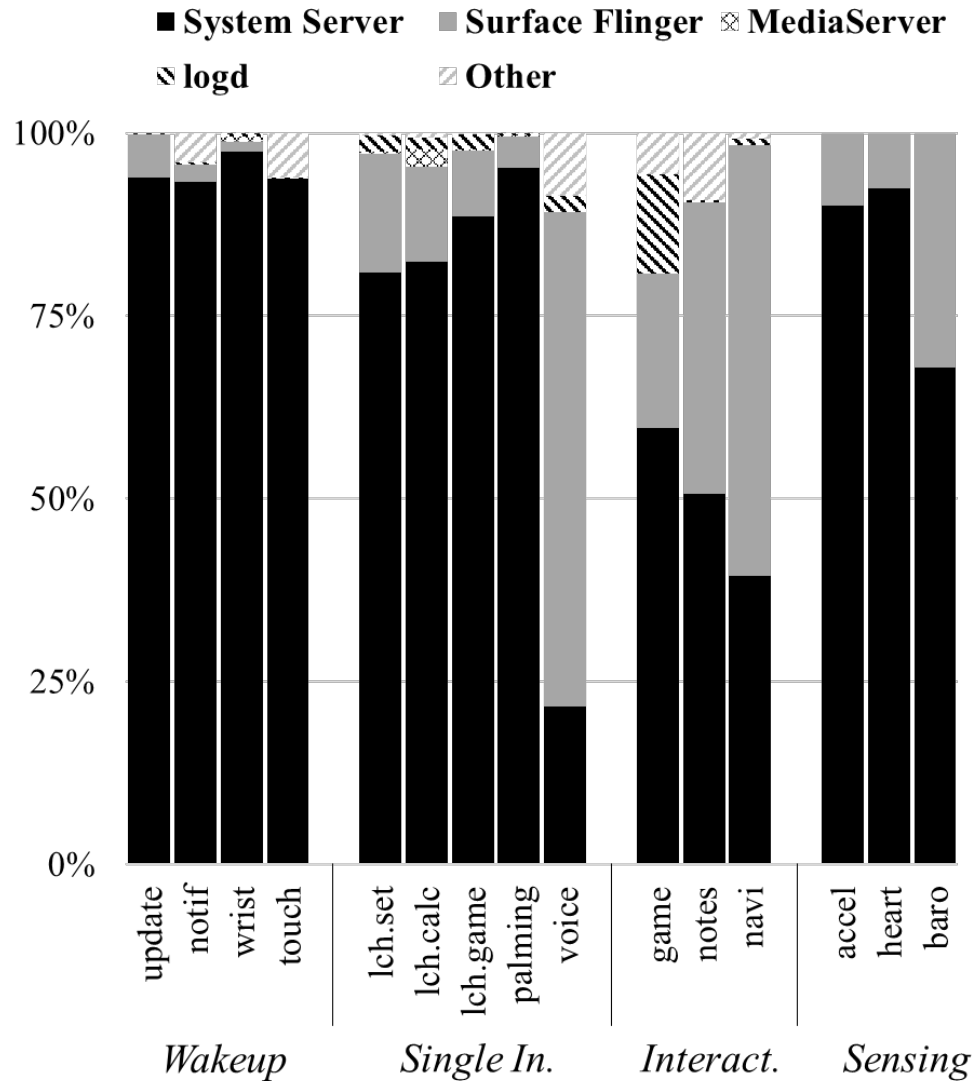
# OS execution dominates CPU usage.






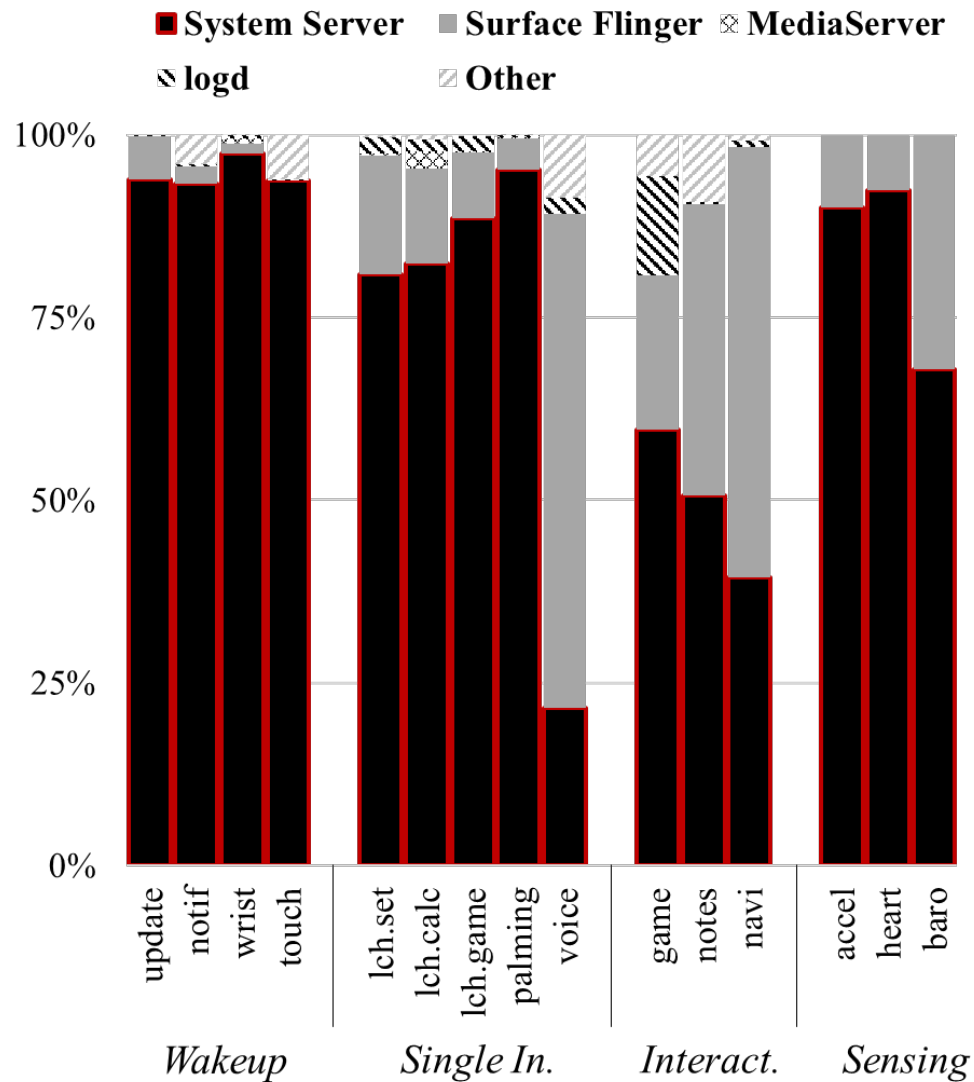



# OS execution dominates CPU usage.



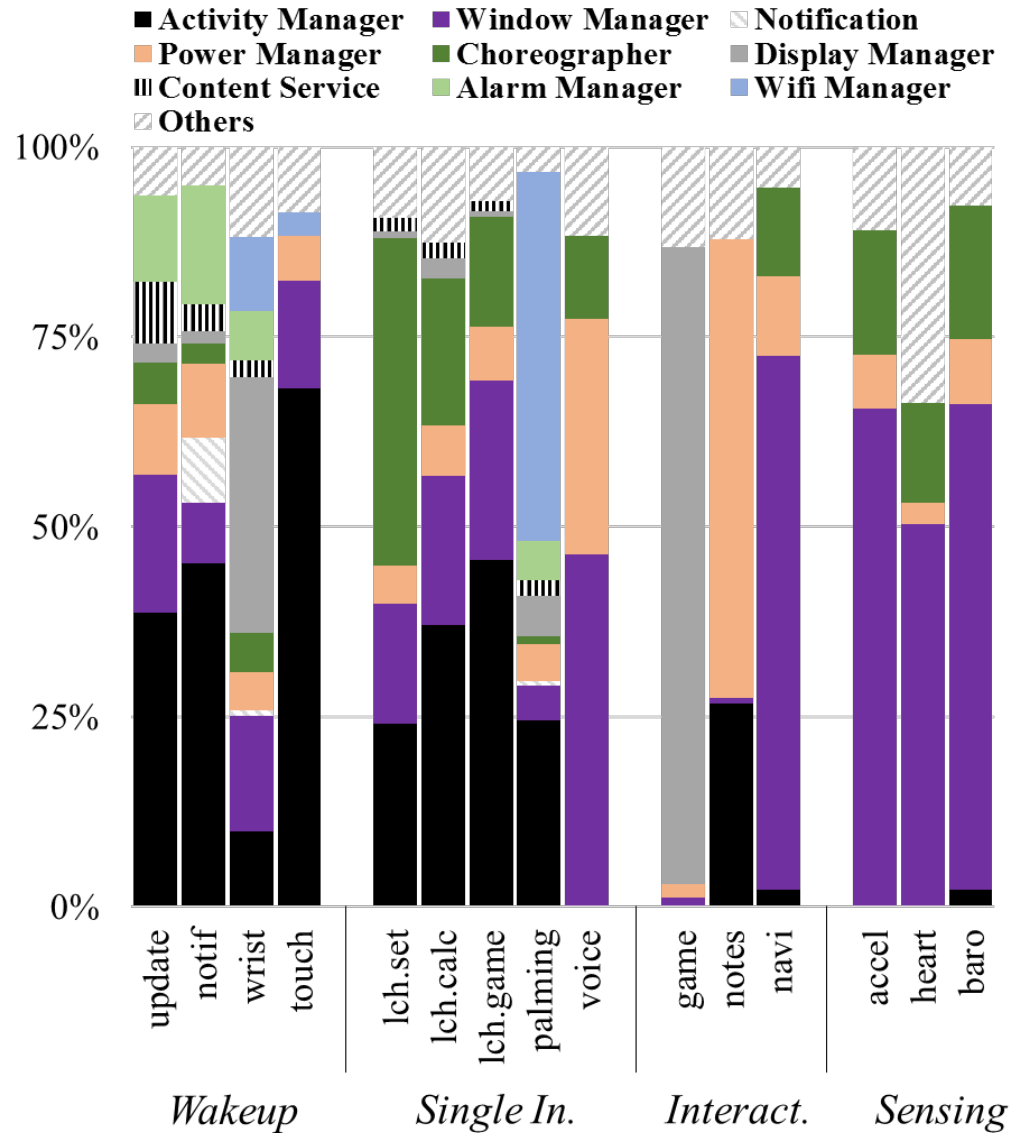



# OS execution dominates CPU usage.






# Costly OS services are ...





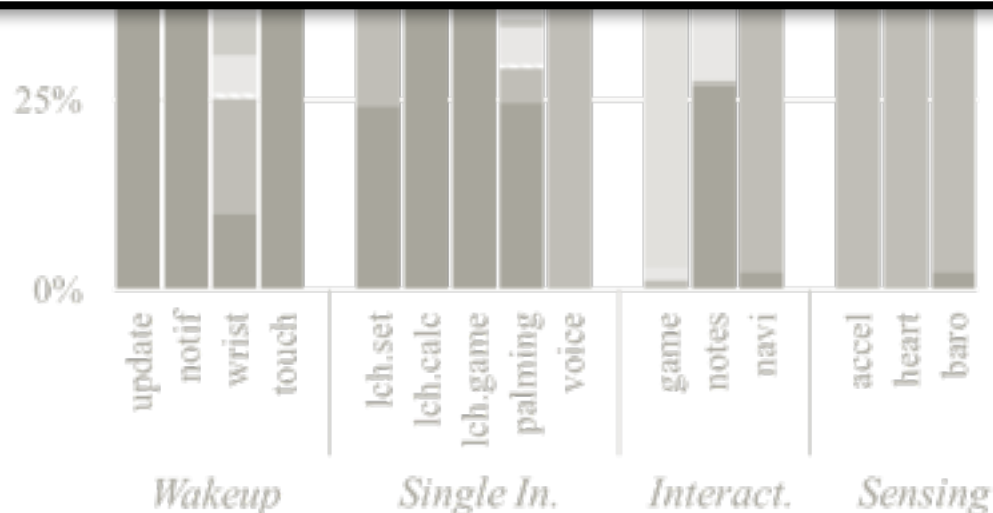
# Costly OS services are likely cruft.




**cruft** (krüft)

n.


1. Trash, debris, or other unwanted matter that accumulates over time.
2. Unnecessary digital information that accumulates over time, such as unneeded files or obsolete lines of code in software: *"By removing cruft, you can recover valuable disk space ... and reduce the chance of software conflicts" (Joe Kissell).*





## Hot functions: highly skewed distribution


Top 5  $\rightarrow$   $>20\%$  CPU cycles  
Top 50  $\rightarrow$   $>50\%$  CPU cycles



## Hot functions: highly skewed distribution

Top 5  $\rightarrow$   $>20\%$  CPU cycles  
Top 50  $\rightarrow$   $>50\%$  CPU cycles

Manipulating basic data structures  
Legacy/improper OS designs




## Hot functions: highly skewed distribution

Top 5  $\rightarrow$   $>20\%$  CPU cycles  
Top 50  $\rightarrow$   $>50\%$  CPU cycles

Manipulating basic data structures  
Legacy/improper OS designs

Anecdotes

Backlight UI layout low-mem killer




# Idle episodes: plentiful and of various lengths

Time (ms)	Pct. Overall	Episodes	Pct. Explained
614.1	17.1%	376	100.0%
843.3	50.5%	352	100.0%
722.6	50.9%	205	99.9%
185.2	25.6%	110	92.9%
153.6	15.6%	120	91.4%
16.8	10.6%	6	100.0%
223.0	61.2%	44	100.0%
2173.0	52.80%	912	100.0%
4035.6	86.80%	277	100.0%

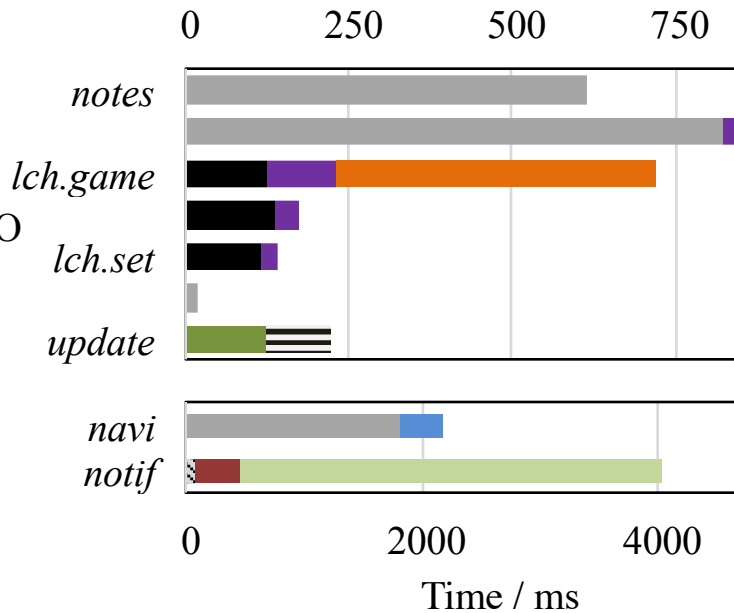
*notes*  
*voice*  
*lch.game*  
*lch.calc*  
*lch.set*  
*touch*  
*update*  
  
*navi*  
*notif*






# Idle anomalies are caused by ...

- Device suspend
- Voice UI
- Cont. interaction
- Cont. interact.+NetI/O
- Storage I/O
- User think
- Bluetooth tail time
- = OS shell policy
- App policy

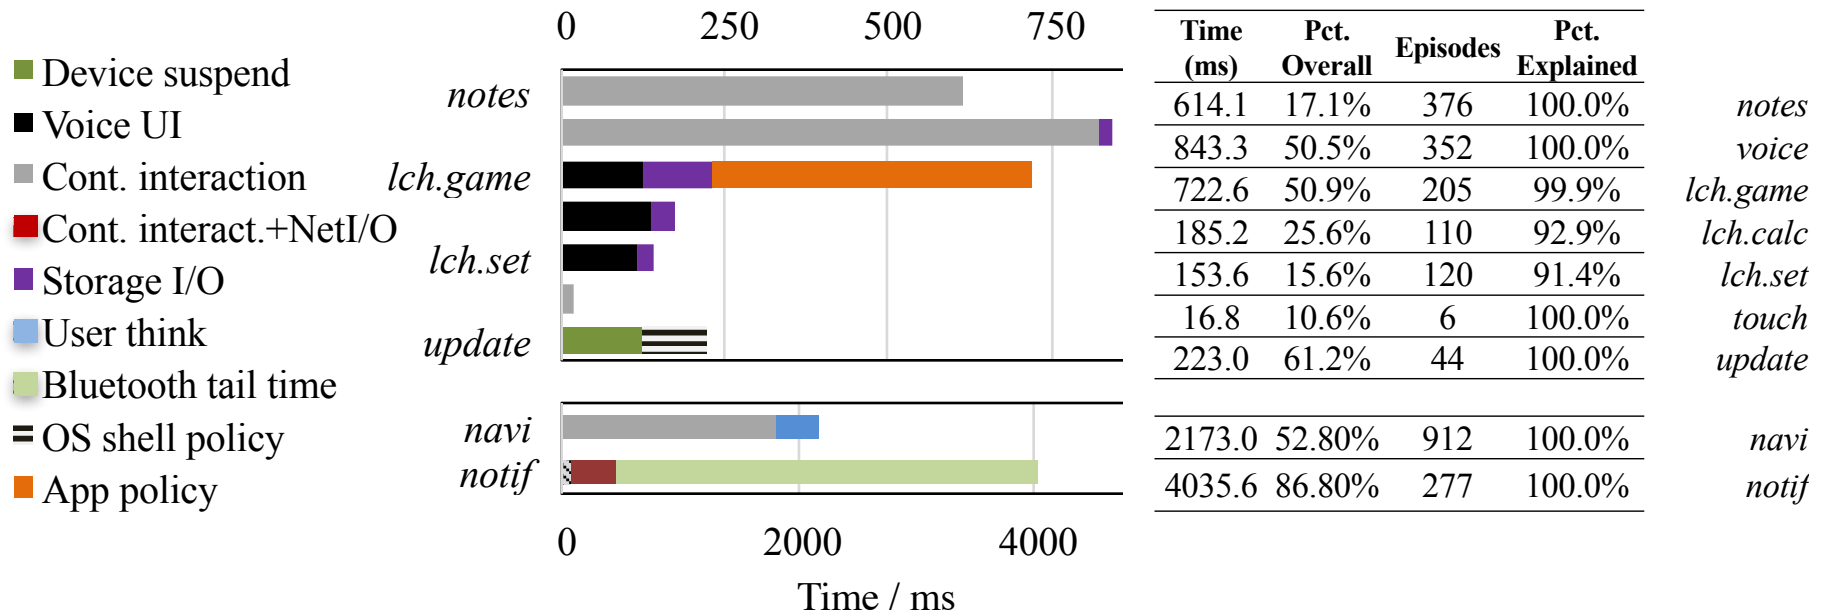


Time (ms)	Pct. Overall	Episodes	Pct. Explained
614.1	17.1%	376	100.0%
843.3	50.5%	352	100.0%
722.6	50.9%	205	99.9%
185.2	25.6%	110	92.9%
153.6	15.6%	120	91.4%
16.8	10.6%	6	100.0%
223.0	61.2%	44	100.0%
2173.0	52.80%	912	100.0%
4035.6	86.80%	277	100.0%

notes  
voice  
lch.game  
lch.calc  
lch.set  
touch  
update  
navi  
notif




# Idle anomalies are caused by ...



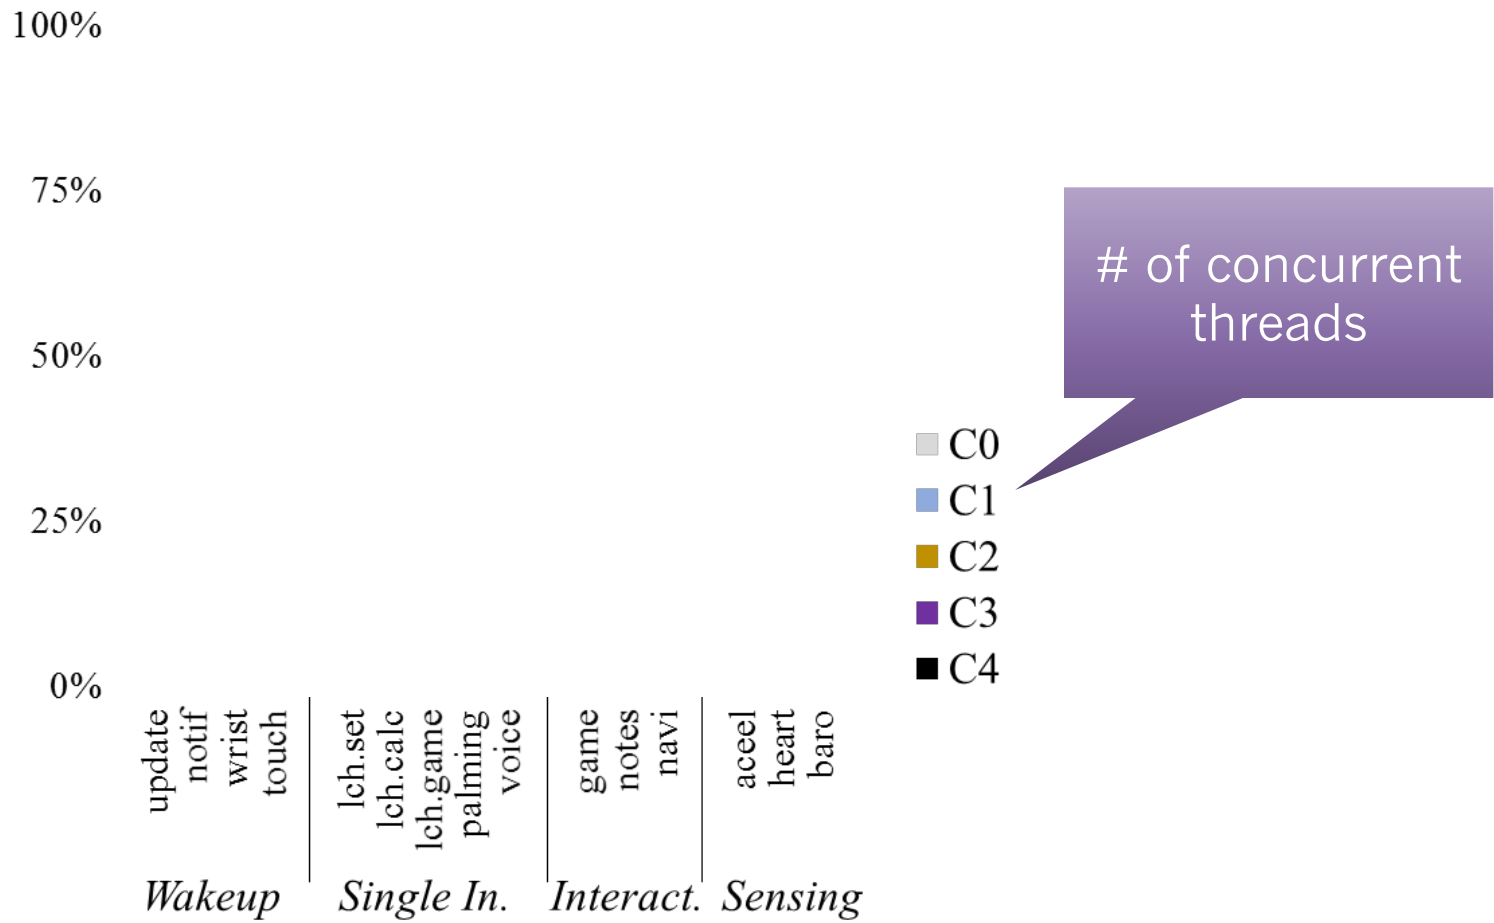
Legacy/improper OS designs  
Performance overprovisioning




Voice UI

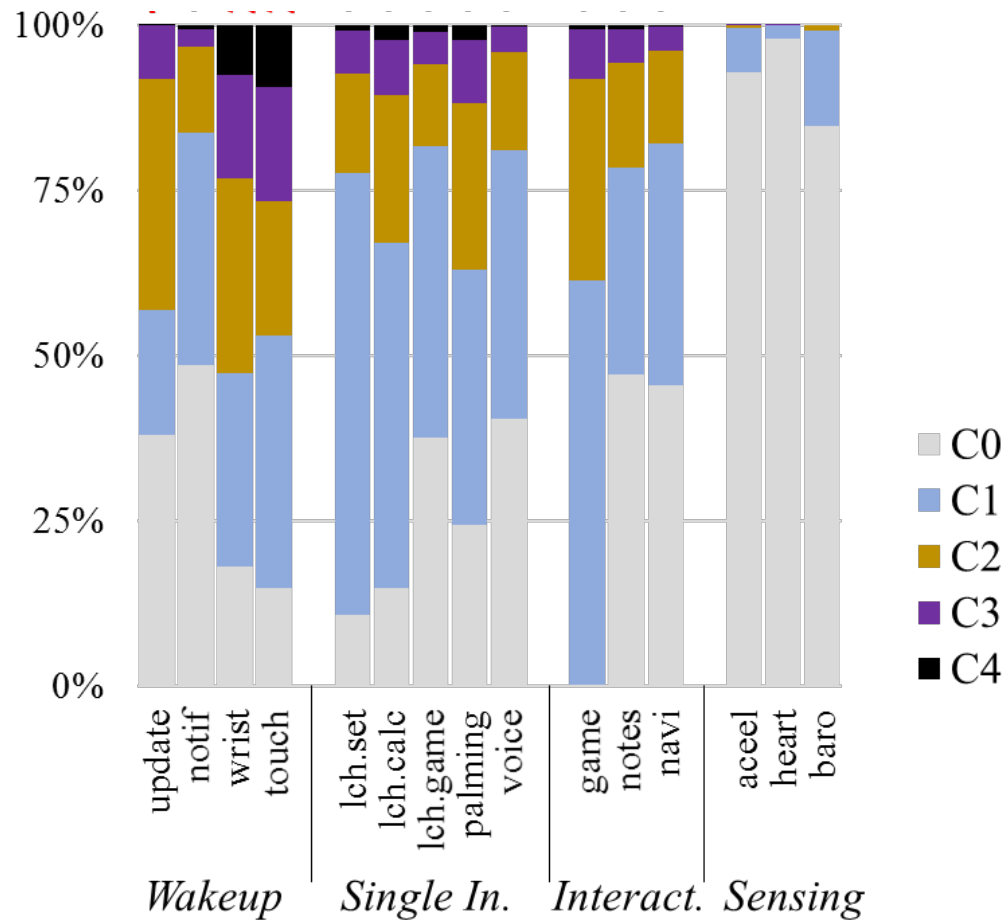



# Substantial TLP on a par with desktop



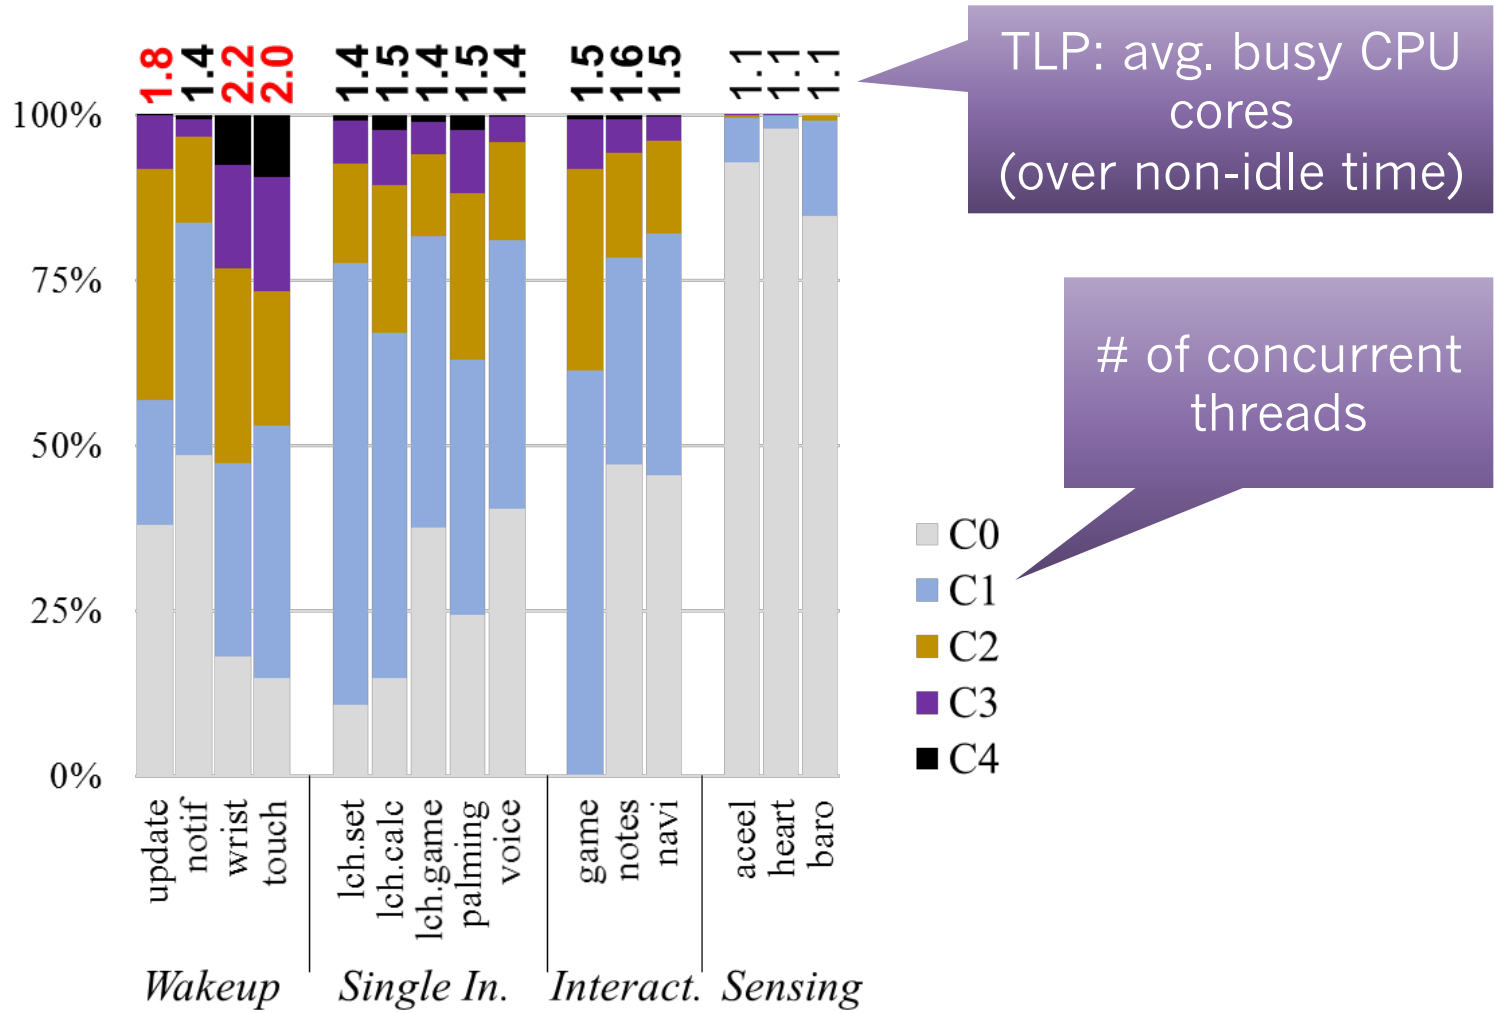



# Substantial TLP on a par with desktop







# Substantial TLP on a par with desktop







# ...due to short interactions.






# Apps are mostly single-threaded; OS contributes to TLP significantly.






# Wearable suffers from uArch inefficiency

**Cycles-per-instruction** (lower is better)  
**2 -- 5** (high!)








# Wearable suffers from uArch inefficiency


**Cycles-per-instruction** (lower is better)  
**2 -- 5** (high!)

## Smartphone as a comparison

1.3 -- 2.5 web rendering

<2 SPEC INT







# Wearable suffers from uArch inefficiency

**Cycles-per-instruction** (lower is better)  
**2 -- 5** (high!)

Smartphone as a comparison

1.3 -- 2.5 web rendering  
<2 SPEC INT







# Wearable suffers from uArch inefficiency

**Cycles-per-instruction** (lower is better)  
**2 -- 5** (high!)


Smartphone as a comparison

1.3 -- 2.5 web rendering  
<2 SPEC INT





# The major cause: complex OS code (L1 icache, iTLB, and branch predictor)



**The major cause: complex OS code  
(L1 icache, iTLB, and branch predictor)**

**uArch problem will NOT be gone with  
future wearable CPUs**

# Four Aspects

## CPU busy

- ◆ OS dominates
- ◆ Lots of cruft
- ◆ Skewed hot functions
- ◆ Legacy bottlenecks

## CPU idle

- ◆ Anomalous
- ◆ OS flaws
- ◆ Too much performance

## Thread-level parallelism

- ◆ Desktop-like
- ◆ OS-contributed

## Microarchitectural behaviors

- ◆ Mismatch
- ◆ OS code complexity

# Repair, don't overhaul (yet)

## CPU busy

- ◆ OS dominates ◆ Lots of cruft
- ◆ Skewed hot functions ◆ Legacy bottlenecks

## CPU idle

- ◆ Anomalous ◆ OS flaws
- ◆ Too much performance

## Thread-level parallelism

- ◆ Desktop-like ◆ OS-contributed

## Microarchitectural behaviors

- ◆ Mismatch ◆ OS code complexity

How about after that?  
(i.e. “next-gen wearable OS”)

We probably will reach a point when OS overhaul/redesign is justified.



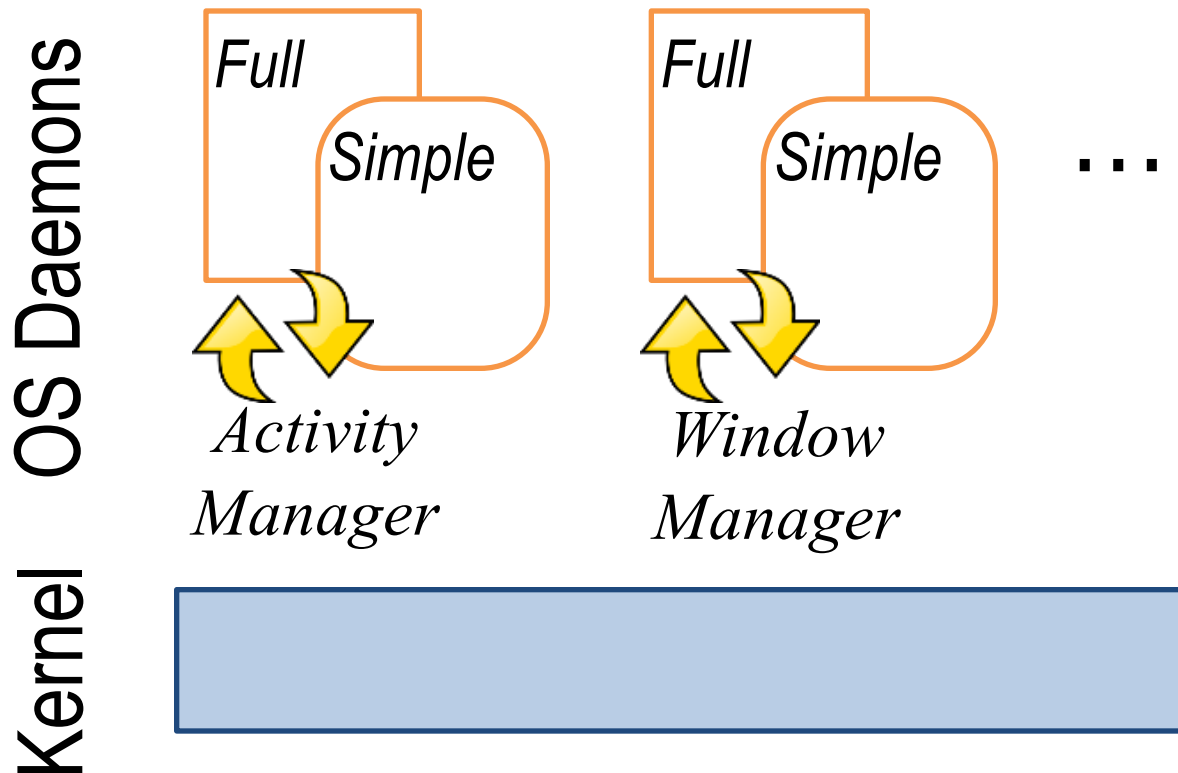
**Specializing OS**  
for common, single-app scenarios



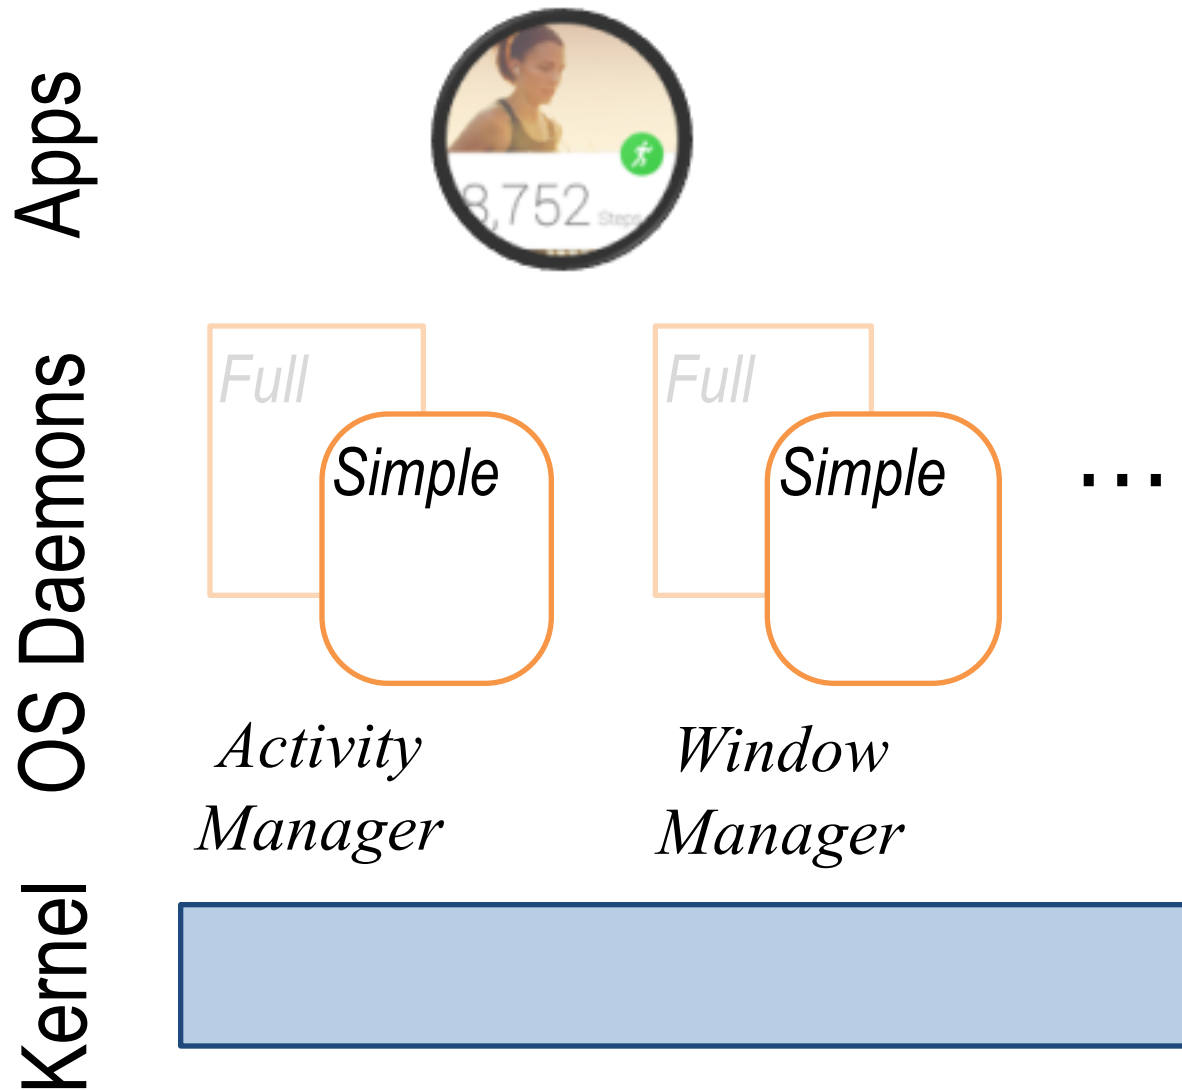
# Restructuring OS for Wearable



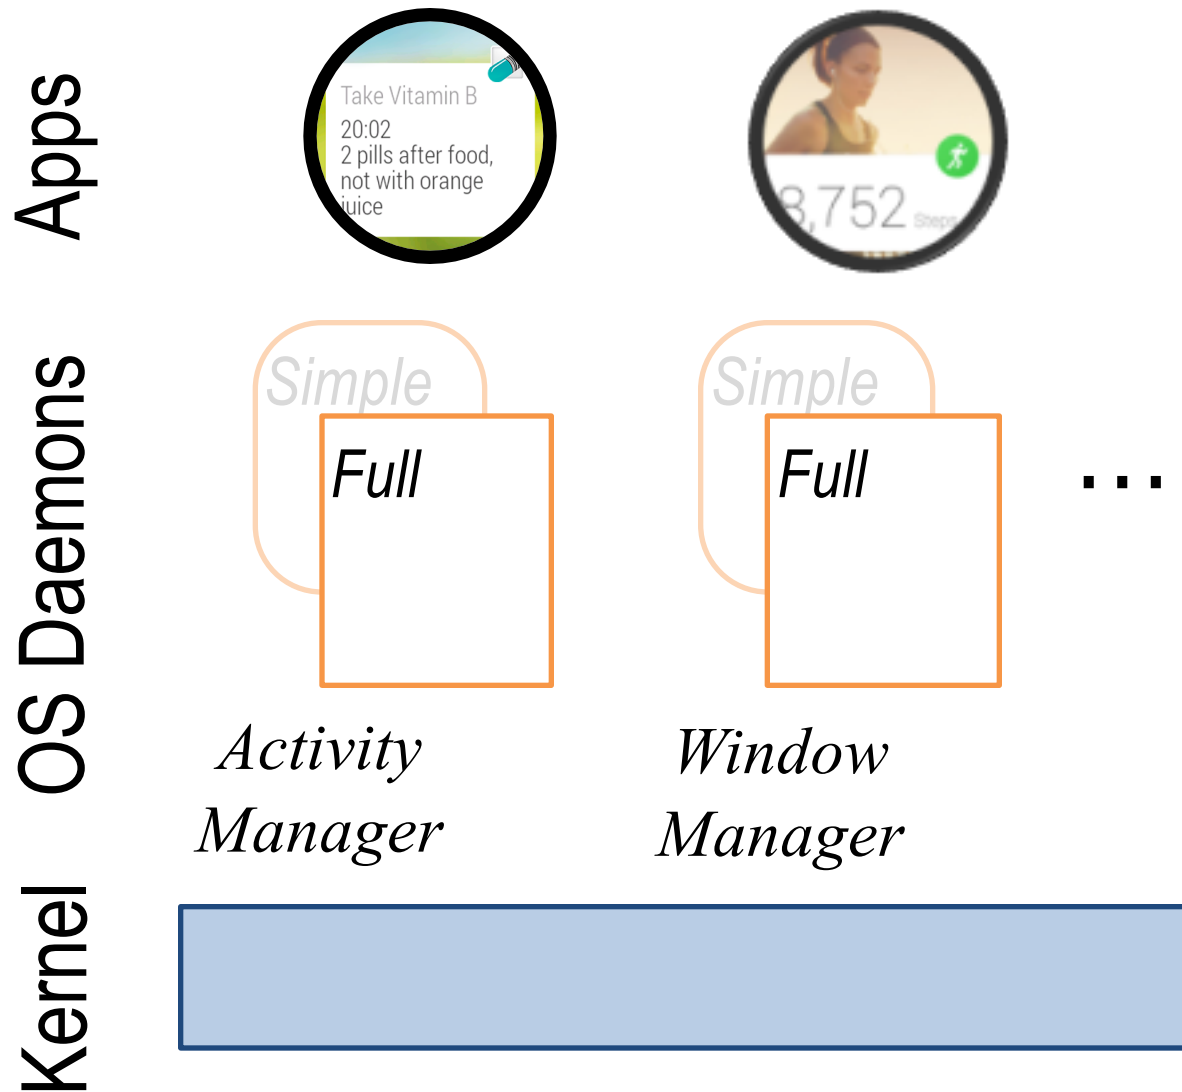
Specializing OS  
for common, single-app scenarios



# Restructuring OS for Wearable



# Restructuring OS for Wearable



# Final takeaway

- Wearables: unique usage and hardware
- Many mobile OS tradeoffs are invalid
  - efficiency v.s. flexibility & programming ease
- Immediate actions: fixing individual OS components
- Future: OS specialization may be needed



Tools, data, and benchmark videos

**[xsel.rocks/p/wear](https://xsel.rocks/p/wear)**

# FAQ

- You forgot Apple Watch or Samsung Tizen.
- Isn't your discovery just some oversight of Google engineers?
- Aren't these things easy to fix?
- Doesn't multicore wearable sound crazy?
- Power! I want to learn about power.
- I bet the Android Wear team already fixed these!

**`xsel.rocks/p/wear`**

# Has Android Wear improved?

## Android Wear 2.0 Developer Preview

### New User Interface



- Material design for wearables
- Expanded notifications
- Darker UI

### Standalone Apps



- Direct network access to cloud
- Apps run on watch even when your phone (Android or iOS) isn't with you

### Watch Face



Complications API:  
any watch face can  
show data from  
any app

### Messaging



New input methods:  
handwriting, keyboard,  
Smart Reply, and 3rd  
party IMEs

### Fitness



Google Fit platform:  
automatic activity  
recognition and  
data API

