

Wishbone: Profile-based Partitioning for Sensornet Applications



- **Ryan Newton,**
Sivan Toledo, Lewis Girod,
Hari Balakrishnan,
Samuel Madden

+ Example Application: Locating Marmots

2

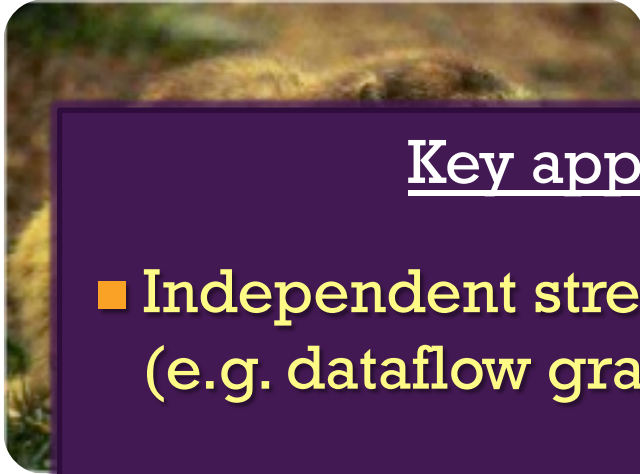


- Gothic, CO deployment August 2007
- Voxnet Platform
 - 2x PXA255, 64MB RAM, 8GB Flash, 802.11B, Mica2 supervisor, Li+ battery, Charge controller
 - Sensors: 4x48KHz audio, 3-axis accel, GPS, Internal temp

with Lewis Girod & UCLA Blumstein Lab

+ We target sensing applications

Animal localization




Pothole detection



Key application features

- Independent stream processing tasks (e.g. dataflow graphs)



- Predictable data rates

Challenges

- Heterogeneous platforms
- CPU & radio bottlenecks

Compu
Vis

Pipeline
leak detection

Speaker identification

EEG Seizure detection

+ Heterogeneous Platforms

Router
weak cpu,
strong radio



Smartphones
medium cpu,
strong radio

Brew

JavaME

Symbian




Android



iPhone SDK

Low power sensors
weak cpu/radio

TinyOS Contiki




Linux
microserver
Python

Java




Mix and Match!

+ Contributions



+ Contributions

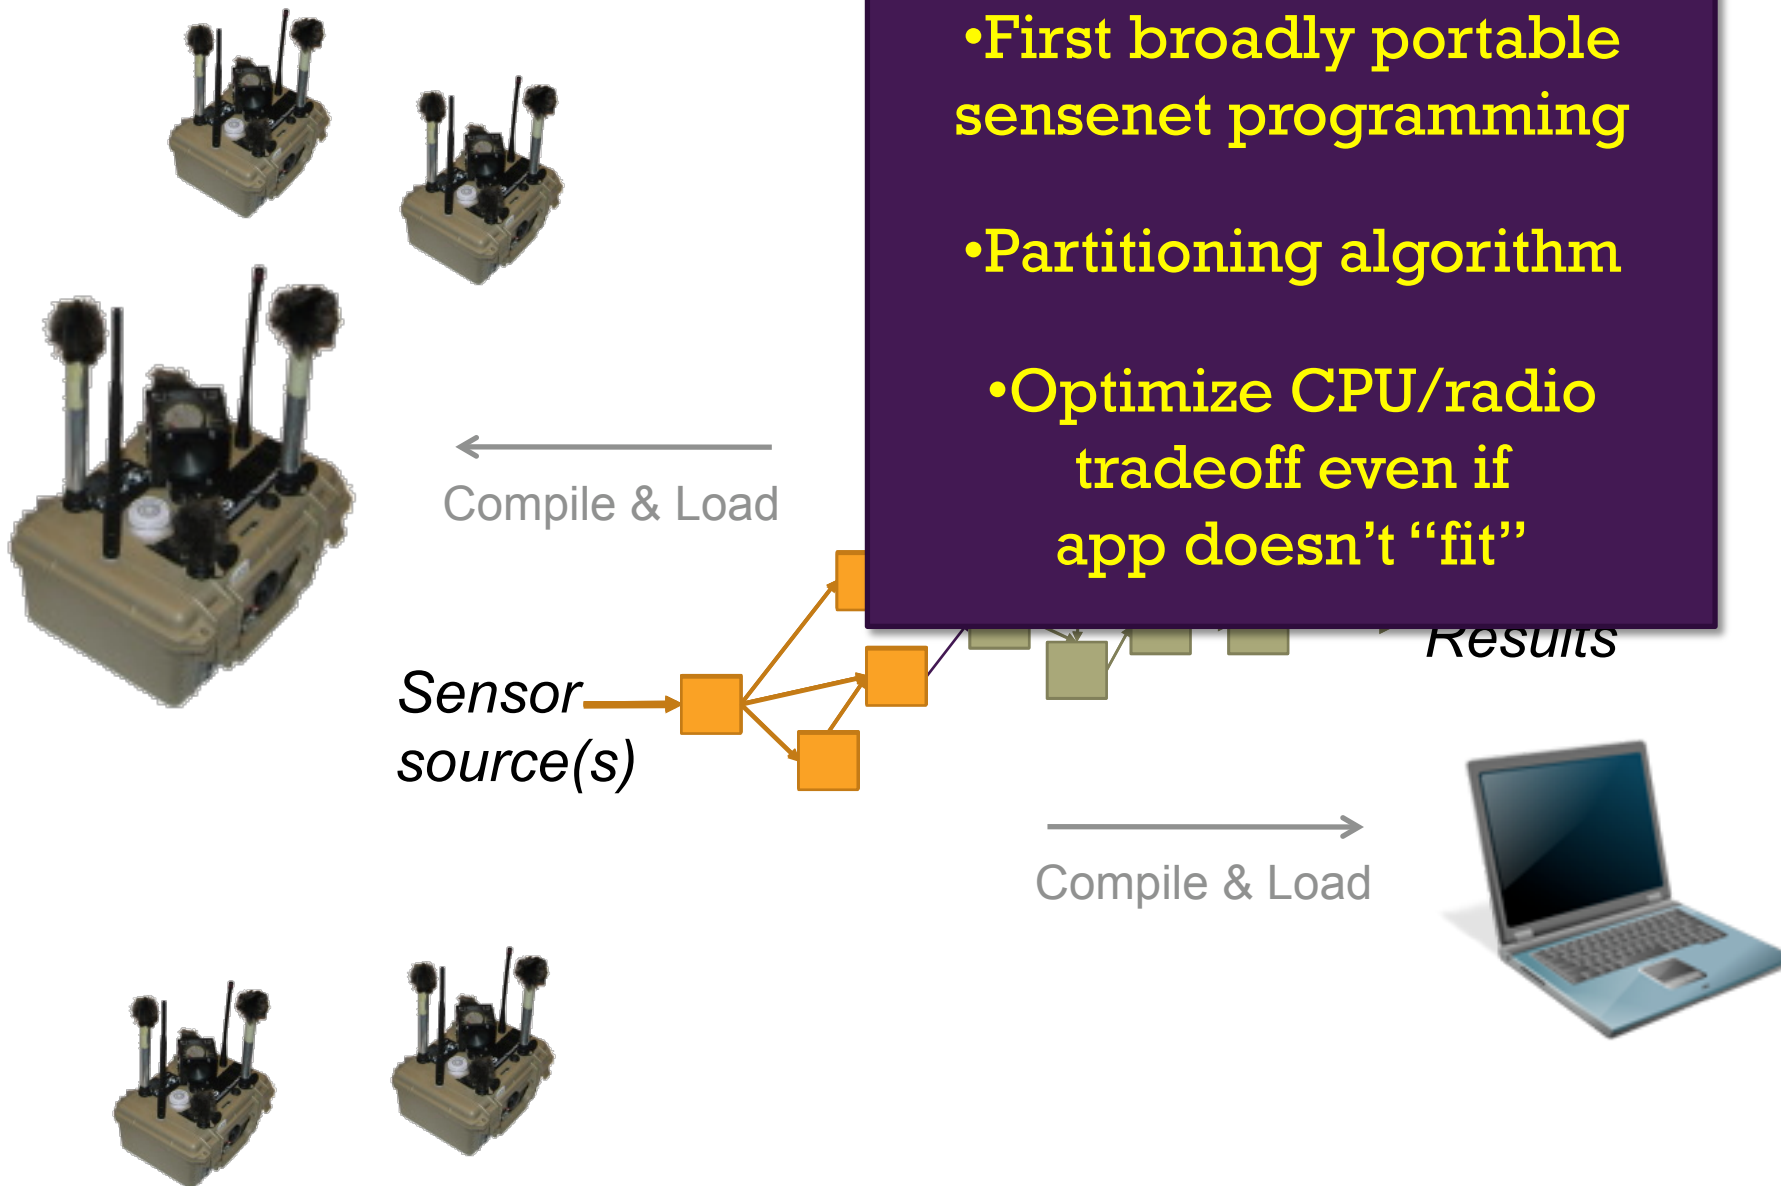


+ Contributions

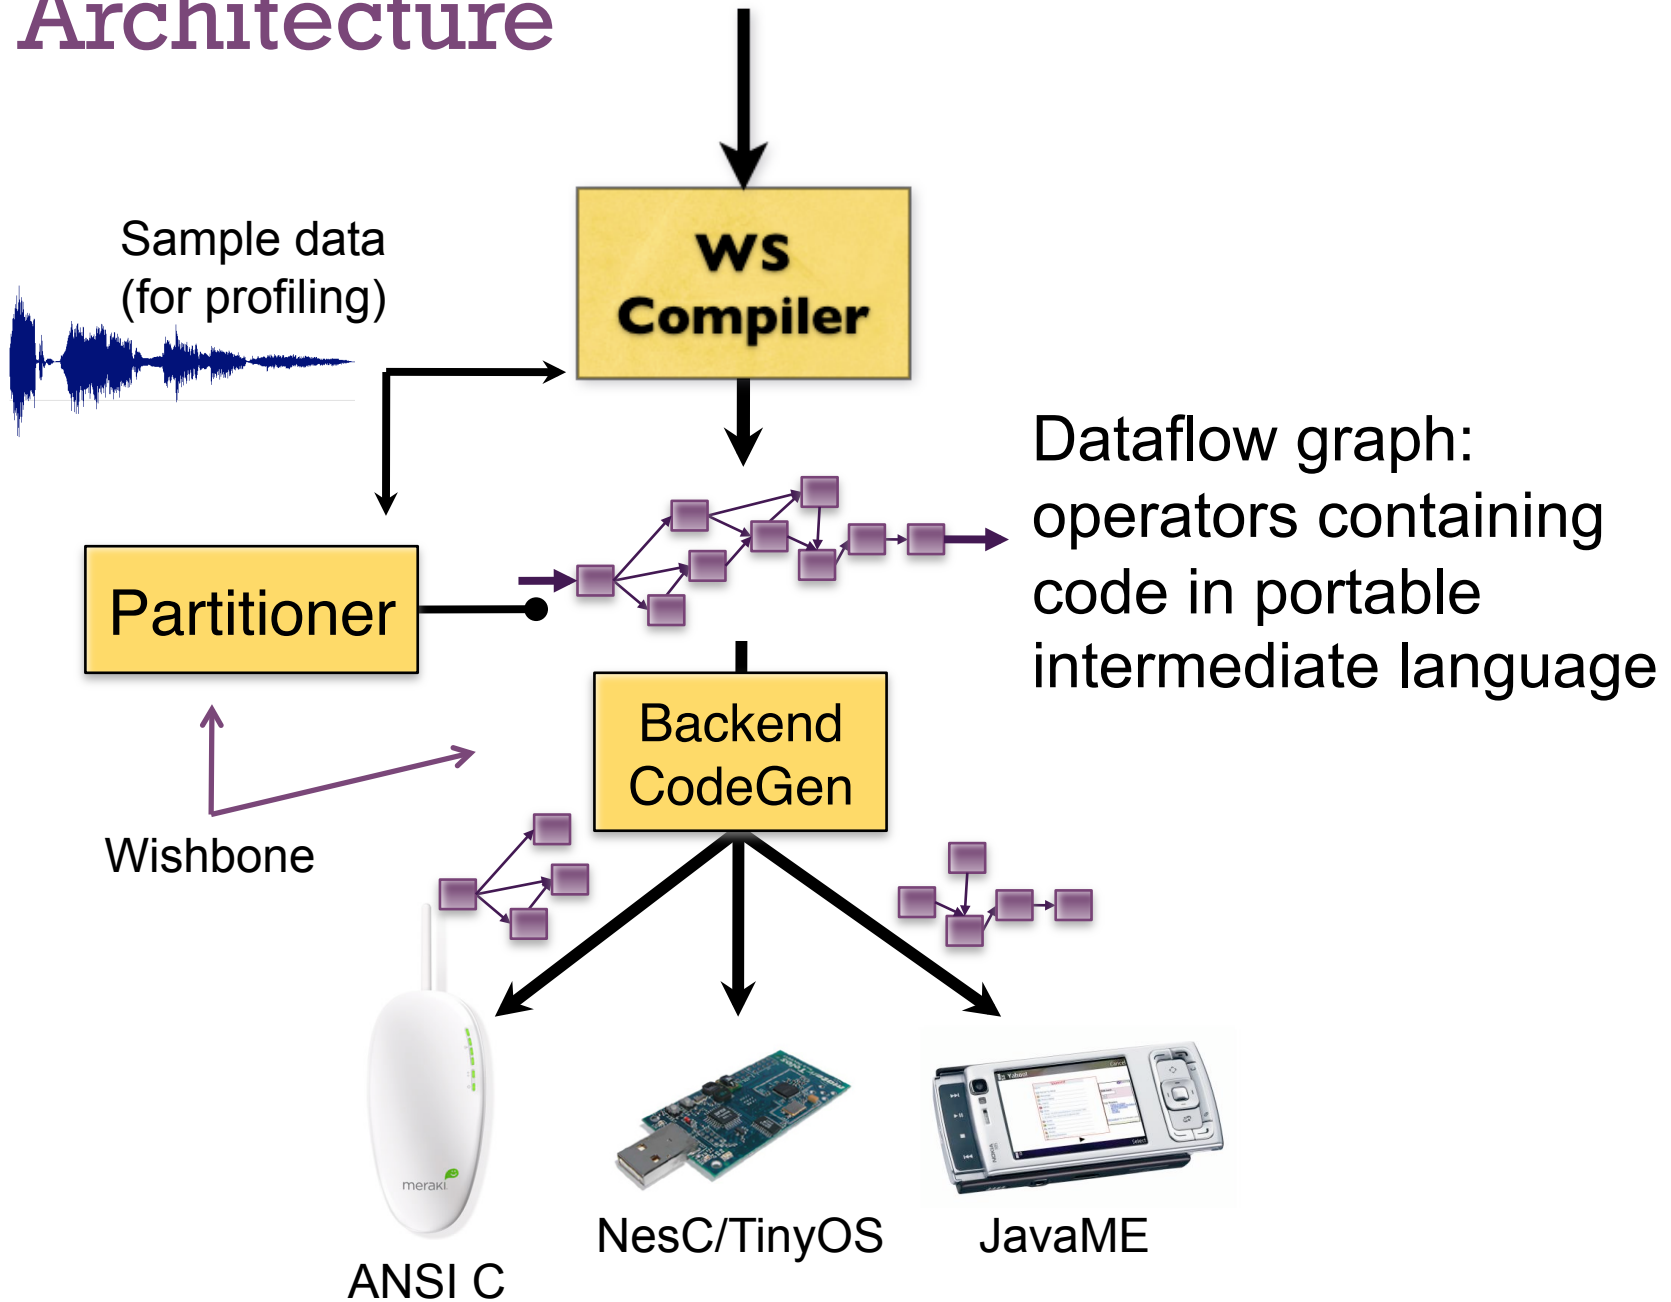
Contributions

- First broadly portable sensenet programming
- Partitioning algorithm
- Optimize CPU/radio tradeoff even if app doesn't "fit"

7



+ Architecture



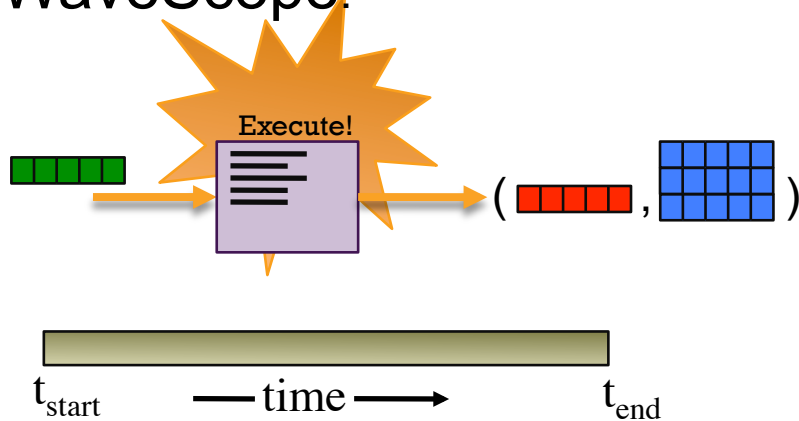
+ Targeting TinyOS

Task granularity, messaging model




- 16 bit microcontroller⁹
- 10K RAM
- No mem. protection
- No threads

WaveScope:



TinyOS:



*Profile-directed
Cooperative
Multitasking:*

```

iterate x in S {
  f(); yield();
  for(i=...) {
    ... if (i==387) yield(); ...
  }
  g(); yield(); g();
}

```

Same goal as
Protothreads

+ Profiling Streams and Operators

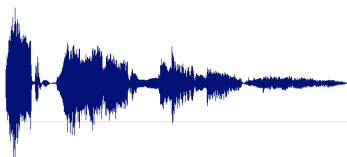
- Every sensor source is paired with sample data



- Includes timing info

- Measure rates, execution times

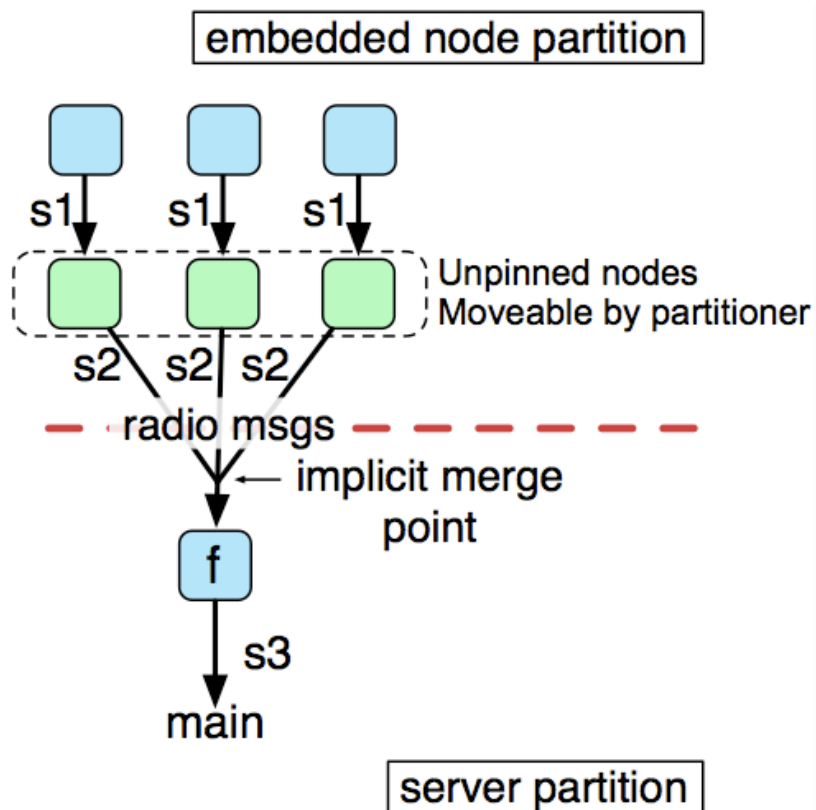
- Separately: profile network channel in deployment environment

- → per-node send rate

audioStream = 
 IFPROF(readFile("foo8kHz",
 readSensor()))



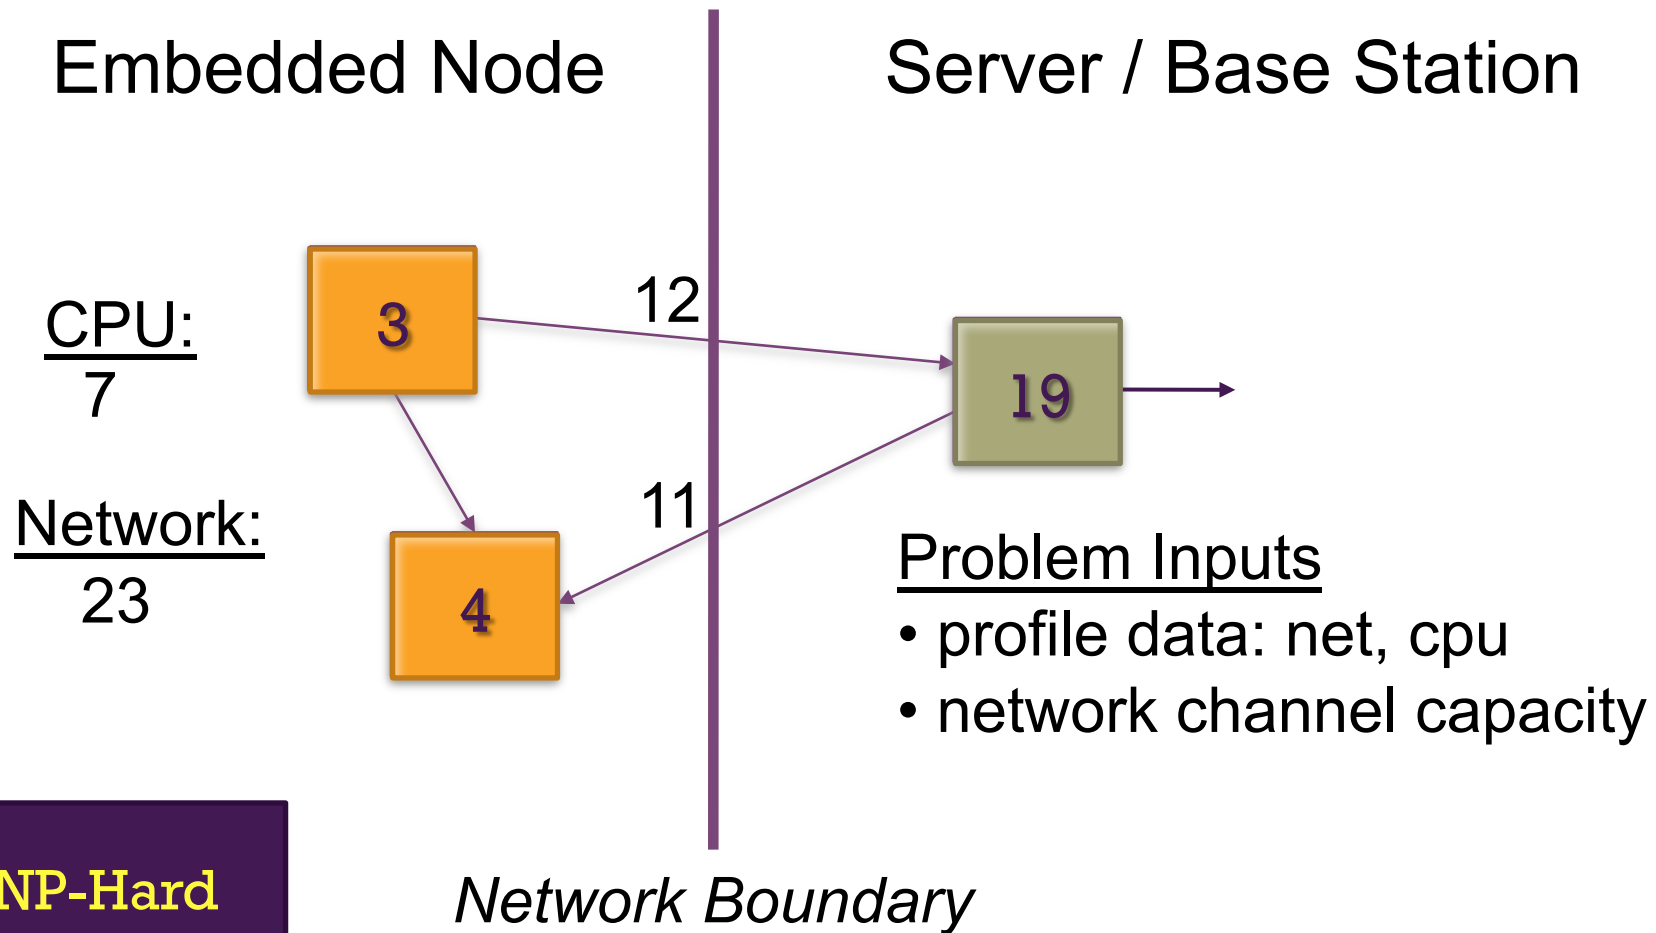
+ State, Replication, and Pinning



Pinning Constraints

- All stateless ops: **unpinned**
- Stateful replicated ops: **unpinned**
- Stateful global ops: ***pinned to server – don't distribute!***

+ Problem Scenario



+ Partitioning Algorithm: Integer linear program formulation

Tricky bit (see paper):

Relating f and g while
staying linear

$f_u \in \{0,1\}$ where

$g_{uv} \in \{0,1\}$ where

ounds

3 Parameters

C, N, α

■ $cpu < C$ where $cpu = \sum_u f_u (compute_u)$

■ $net < N$ where $net = \sum_{uv \text{ Edges}} g_{uv} (data_{uv})$

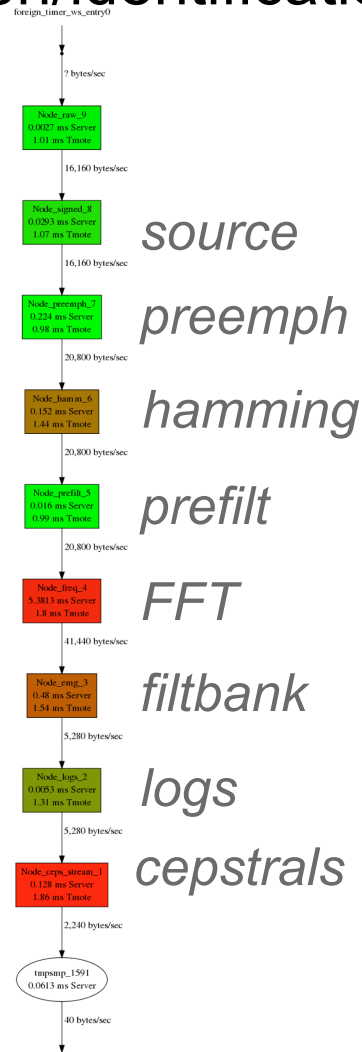
■ Minimize objective function

$$\min(\alpha cpu + net)$$

Proxy for
Energy

+ Evaluation: Two Applications

Human speech
detection/identification



source

preemph

hamming

prefilt

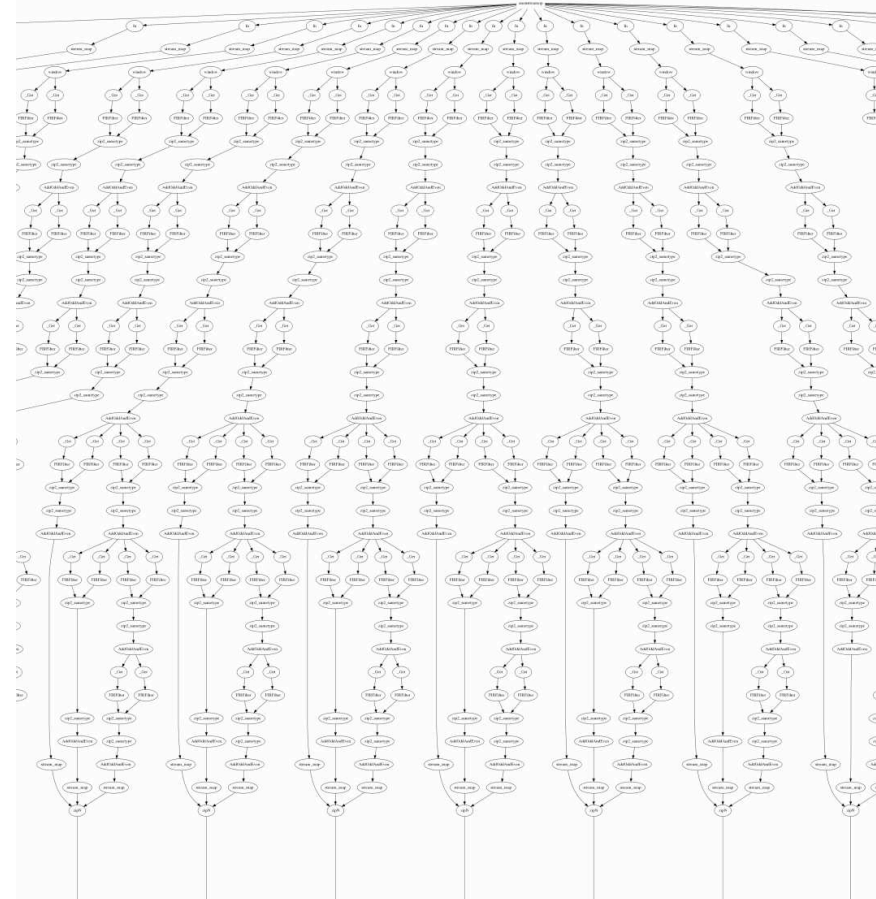
FFT

filtbank

logs

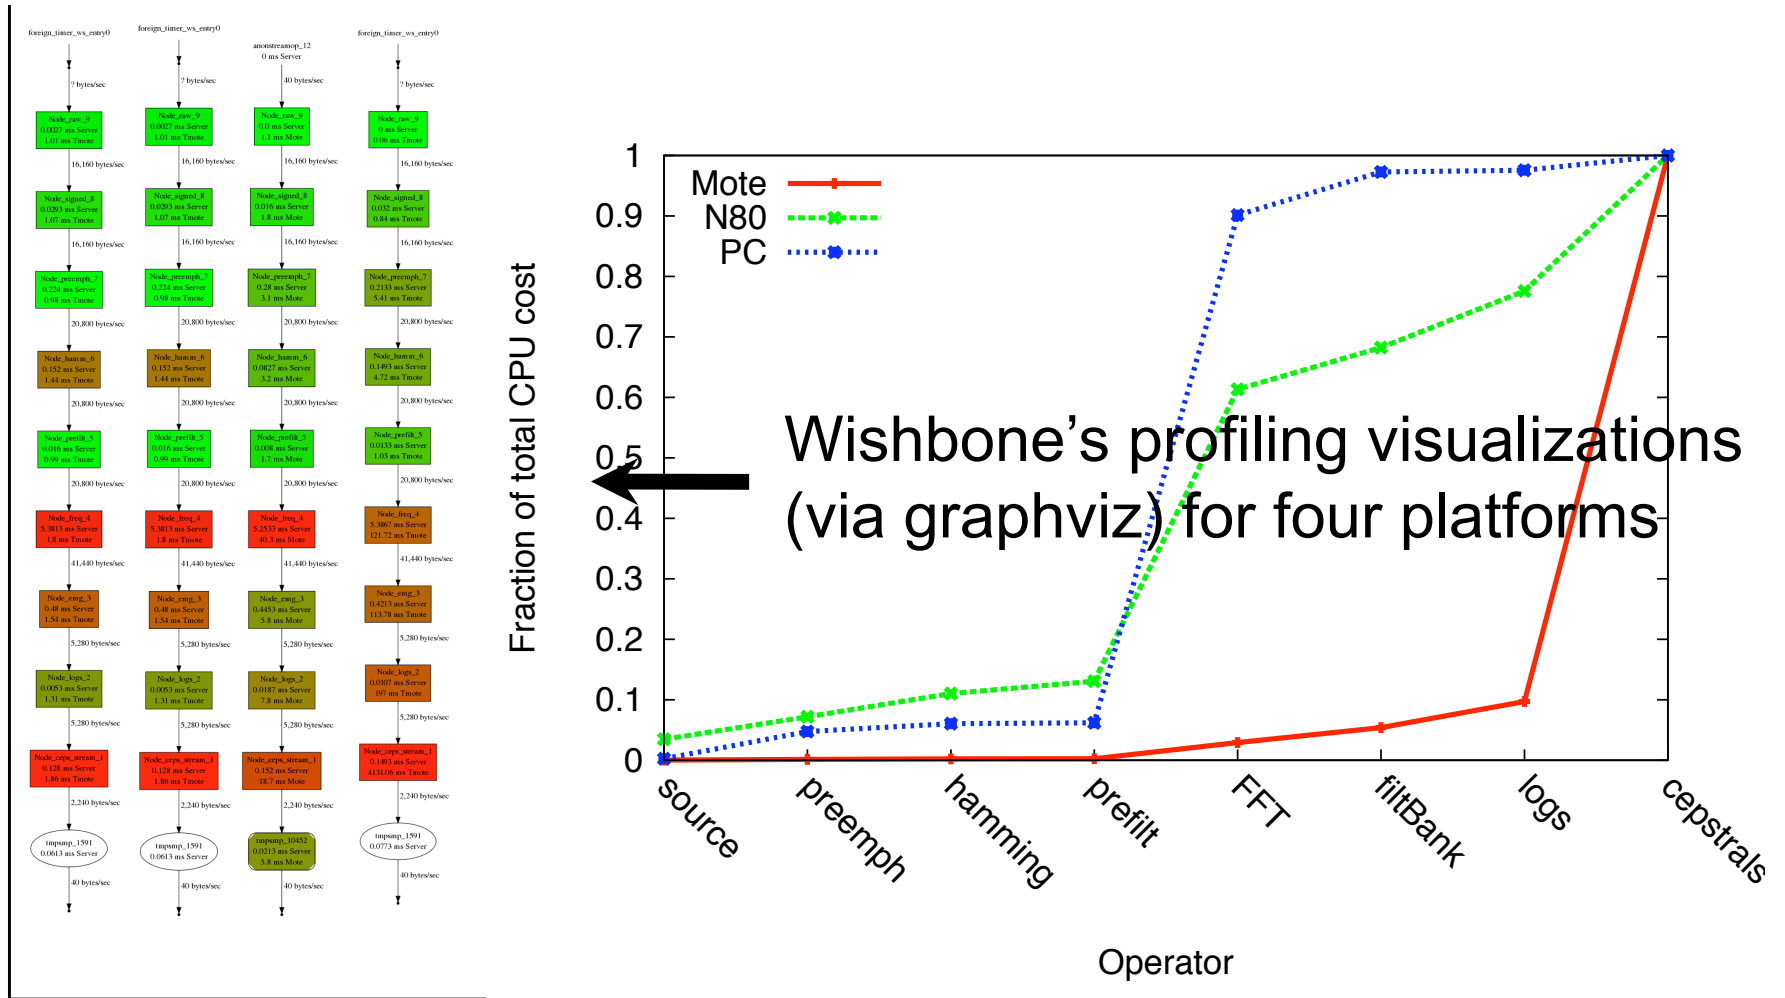
cepstrals

EEG-based seizure
onset detection




1400 operators

+ Observation: Relative cost varies by platform

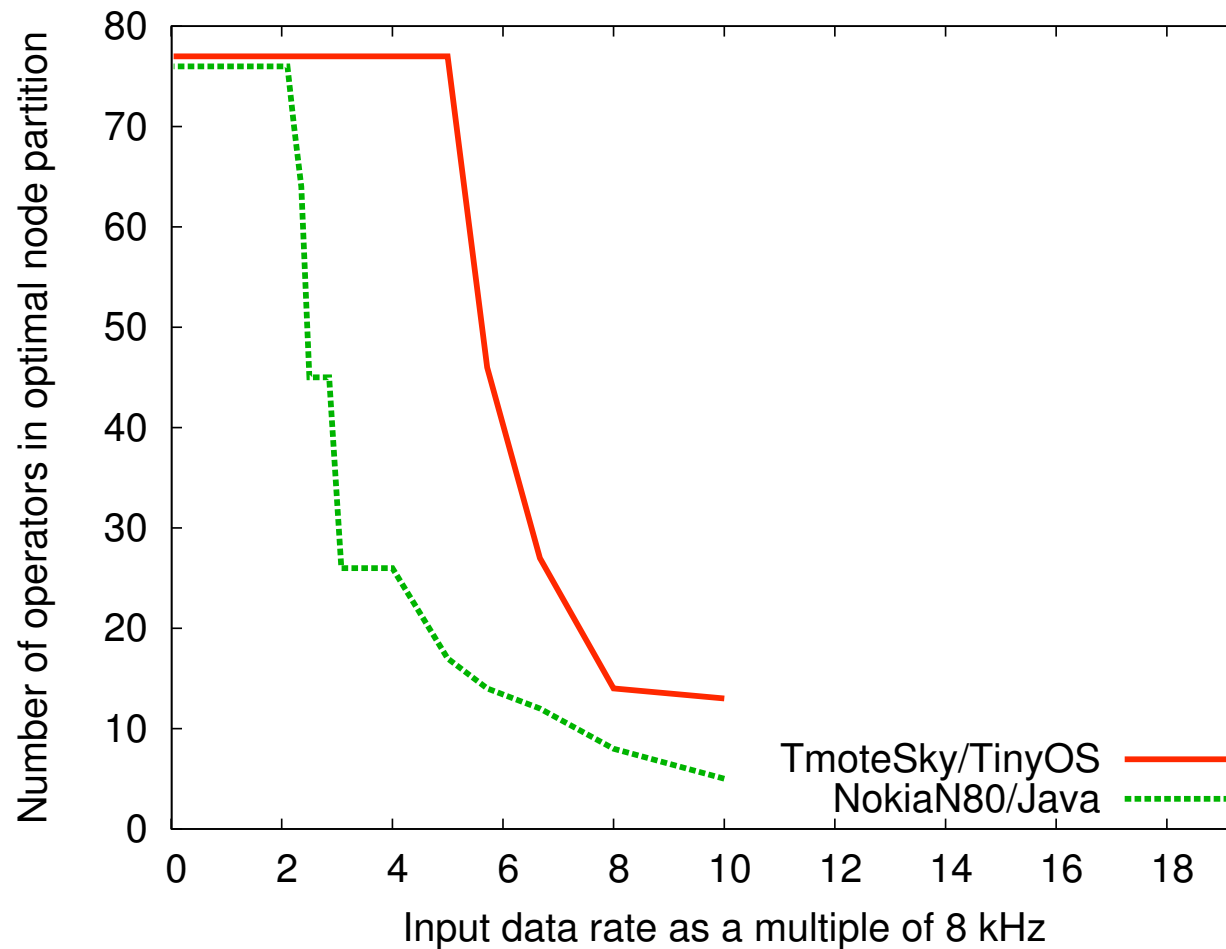


+ Visualizing Profile Data: Bandwidth vs. Compute





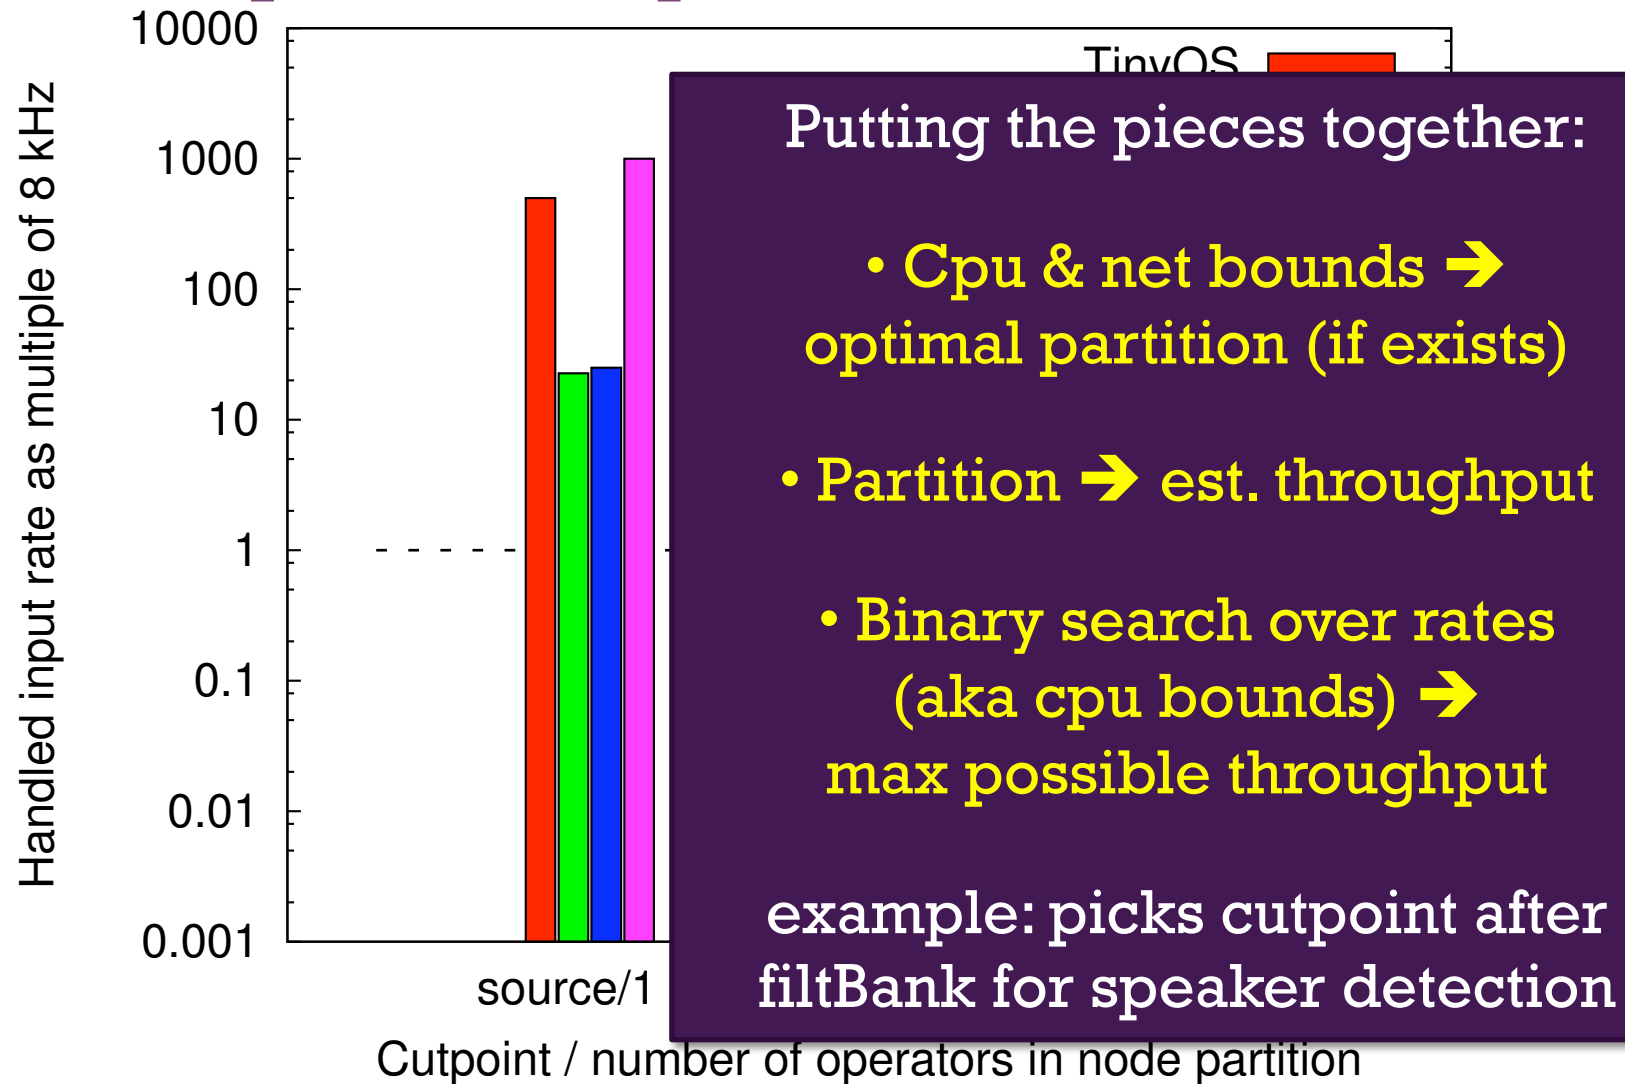
Optimal partitions across platforms



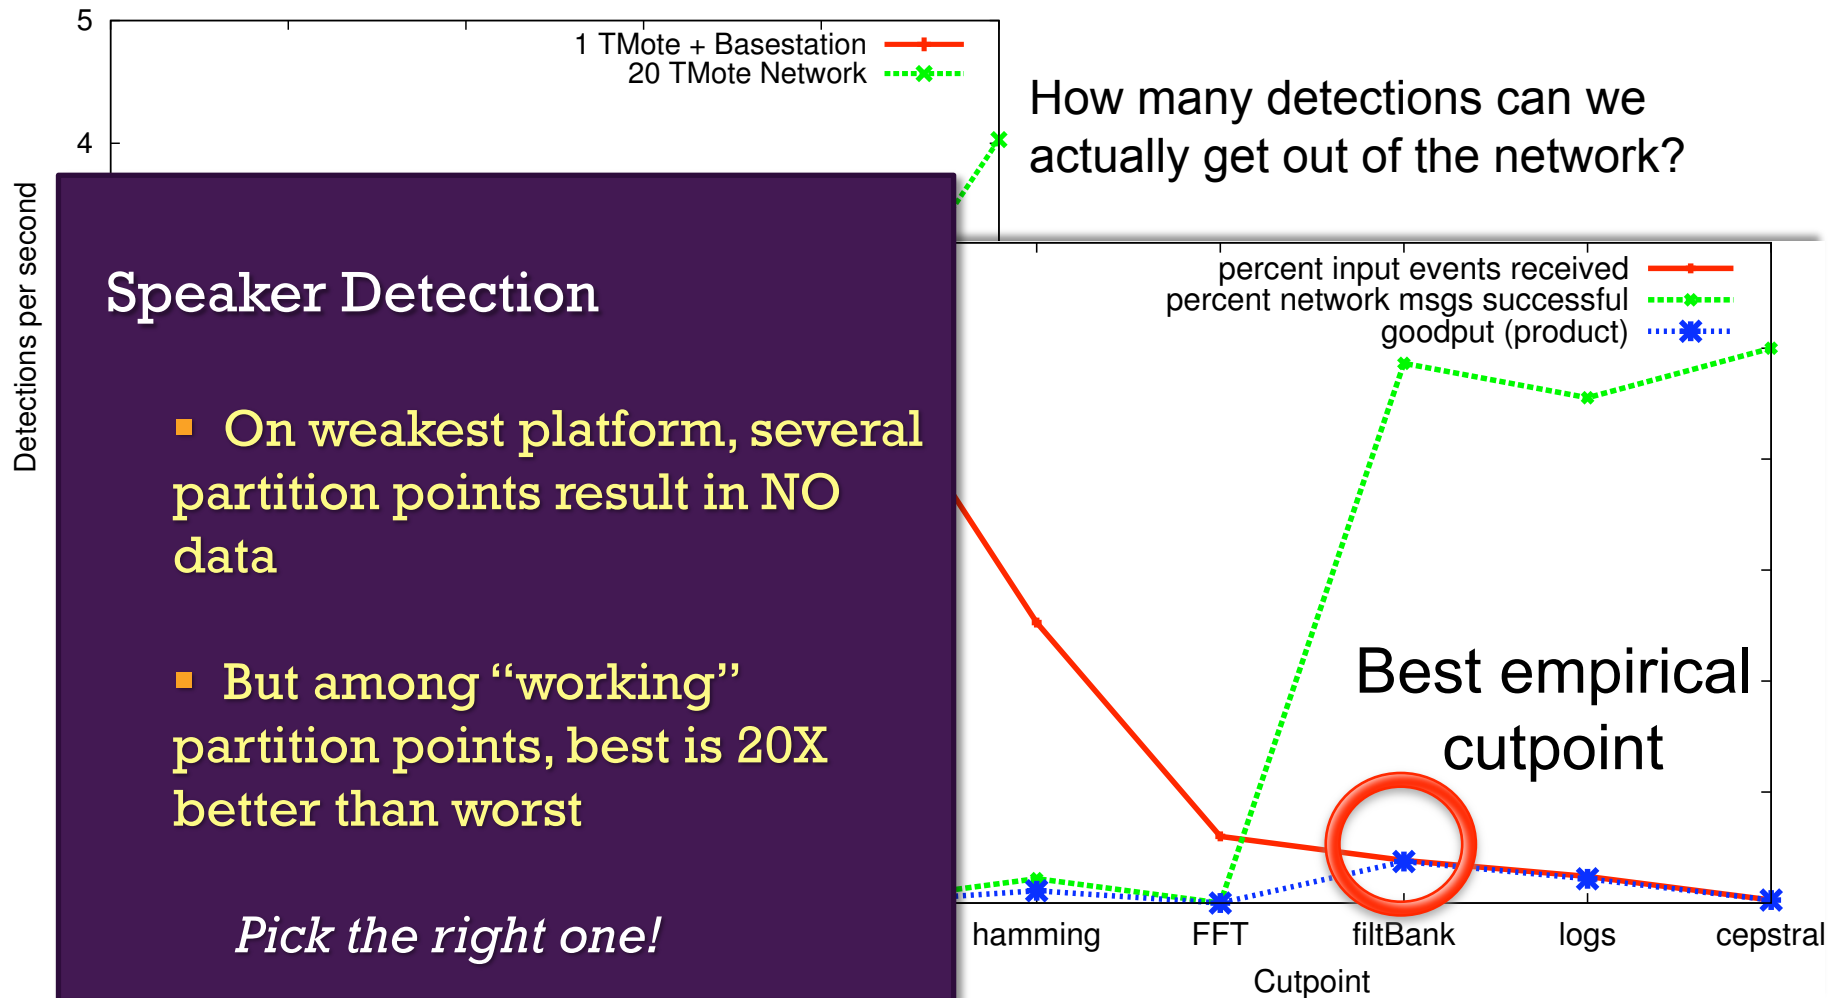
EEG Application (1 of 22 channels)

Each line represents 2100 partitioner-runs

+ Speaker Detection: CPU performance across partitions/platforms



+ Groundtruth: Testbed deployment, 20 motes



+ Related Work

- Graph partitioning for scientific codes
 - balanced, heuristic – e.g. Zoltan
- Task scheduling, commonly list scheduling
- Dynamic: Map-reduce, Condor, etc.


- Sensor network context: Tenet and Vango
 - Linear pipeline of operators
 - Manual partition
 - Run TinyOS code on both server and sensor



CONCLUSION

- **WS/Wishbone System:** (*available, open source*)
 - Convenient, efficient way to run one program on many devices
 - Dataflow graph profiling and visualization extremely useful even outside of auto partitioning
- **Contributions**
 - Algorithm for partitioning streaming programs
 - Backend(s) enable high level programming even on motes
- Techniques for node/server partitioning could be extended to handle heterogeneous cores intra-device.

Contact: Ryan Newton - newton@mit.edu



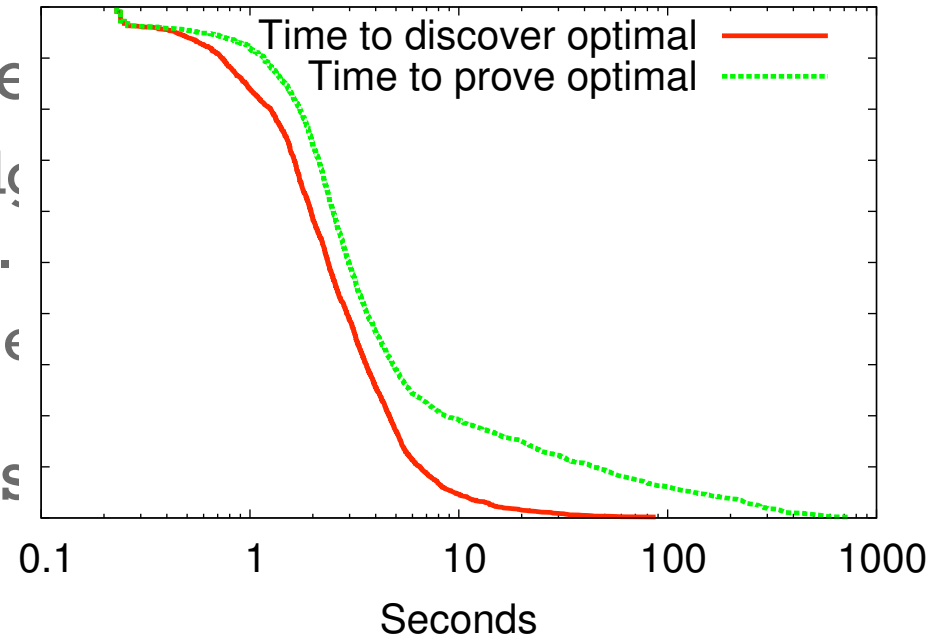
+ Partitioning: Algorithm Runtime

■ Graph Preprocessing step


- Merge vertices until all edges are monotonically decreasing.
- Eliminates the majority of ϵ

■ Even without preprocessing

- 8000 runs,
- partitioning the 1400-node EEG dataflow graph,
- with different CPU budget,
- took under 10 seconds 95% of the time.
- But there is a long tail... luckily ILP solvers produce approximate solutions as well!



+ Motivating Example



*Unstable optimal partition.
Flips between horizontal and vertical partition.*