System energy consumption is a multi-player game

Mian Dong, Tian Lan and *Lin Zhong*

Energy accounting by software

How much energy does a process contribute given a time internal?

Modern mobile systems are multiprocessing

More cores, more types of cores, and more specialized cores

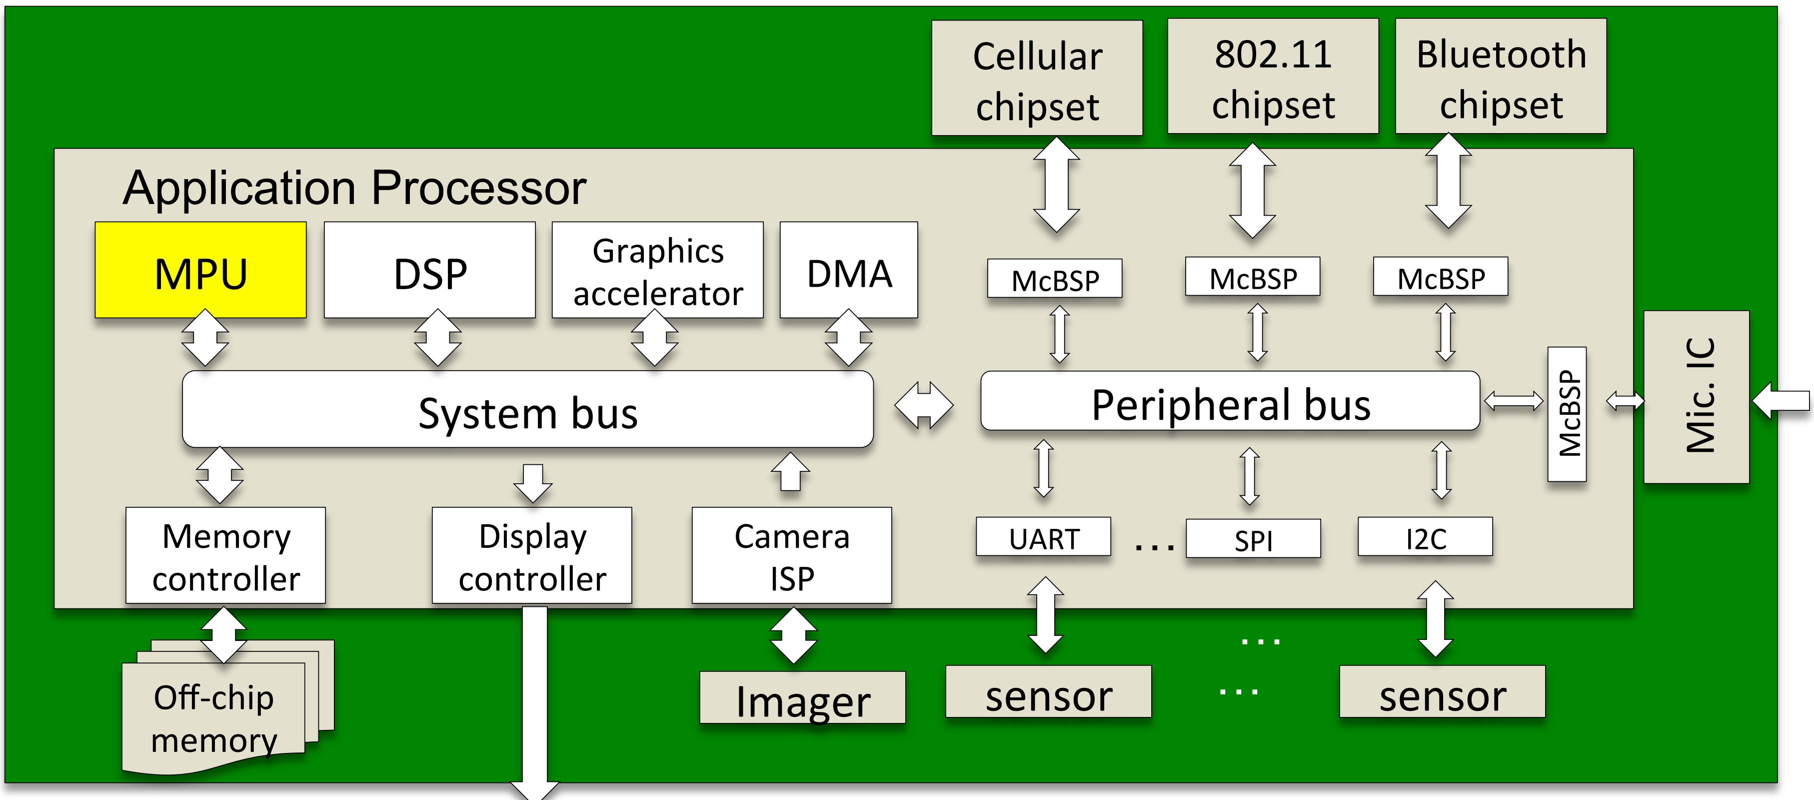

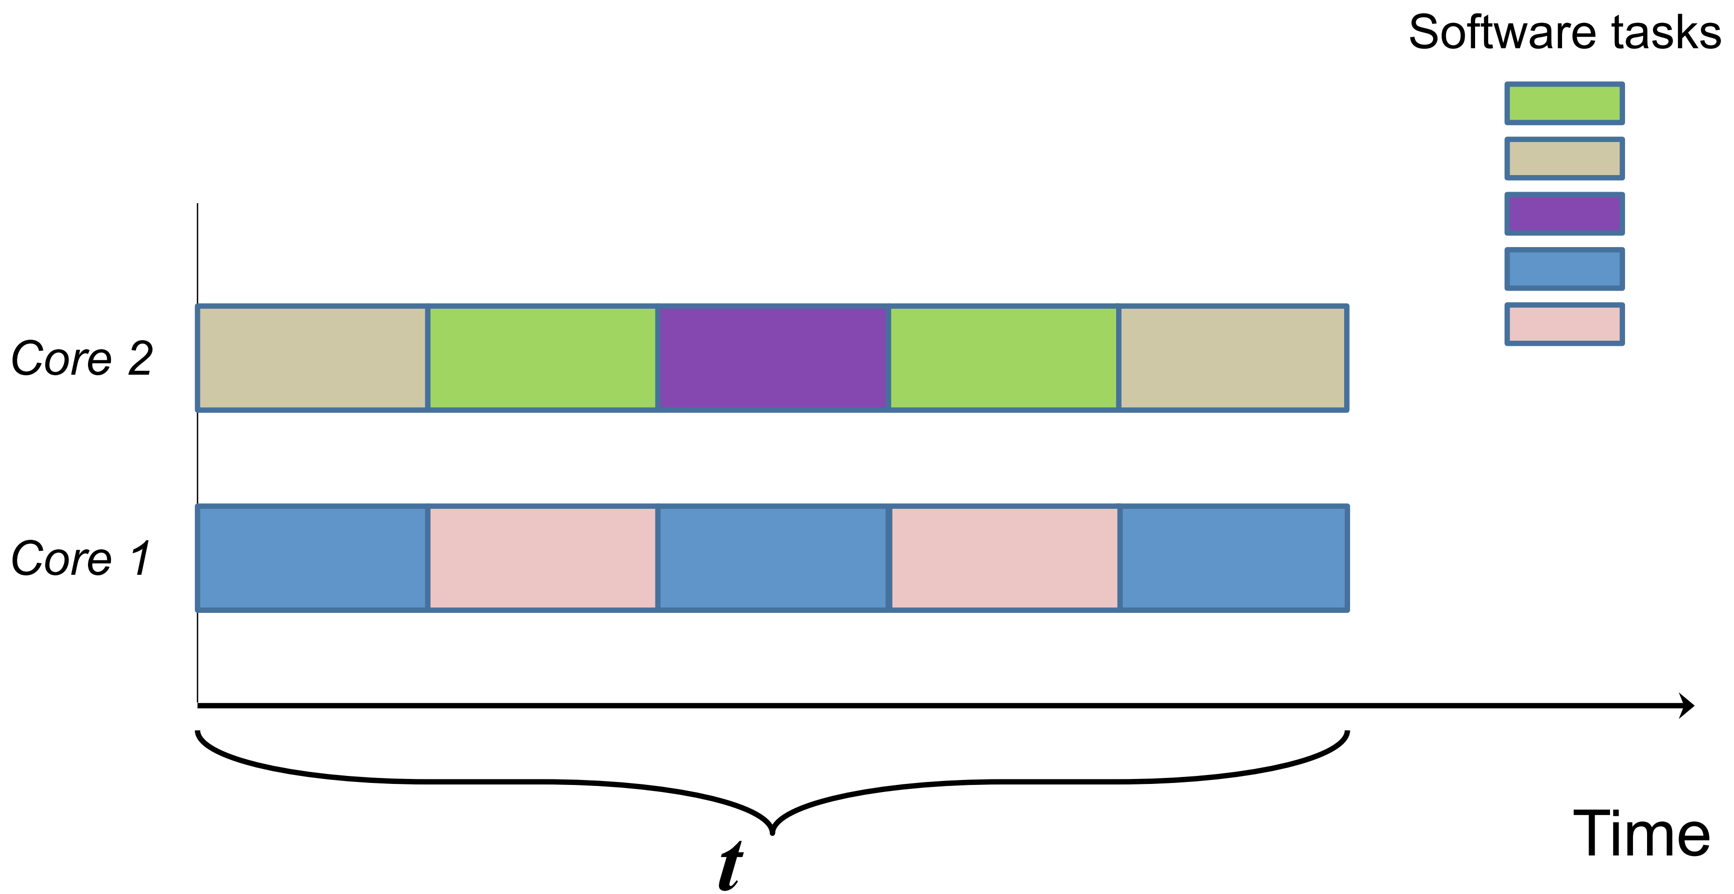
Software tasks
Core 2
Core 1
t
Time

Model-driven policy

$$E = \beta_0 + \beta_1 x_1 + \ldots + \beta_p x_p$$

Predictors x_i: System status variables

n tasks contribute to the predictors

$$x_1 = x_{1,1} + \ldots + x_{1,n}$$
$$\vdots \quad \vdots \quad \vdots$$
$$x_p = x_{p,1} + \ldots + x_{p,n}$$

Energy contribution by process i

$$\phi_i = \beta_1 x_{1,i} + \ldots + \beta_p x_{p,i}$$

H. Zeng *et al*, "ECOSystem: managing energy as a first class operating system resource," *ASPLOS'*02.
A. Kansal *et al*, "Virtual machine power metering and provisioning," *SoCC'*10.
A. Roy, "Energy Management in Mobile Devices with the Cinder Operating System". EuroSys 2011

Problems

$$E = \beta_0+\beta_1 x_1+\ldots+ \beta_p x_p$$

Predictors must be software accountable
Model must be linear
Constant factor (β_0)

Lone-wolf policy

$\phi_i = E(\boldsymbol{S} \cup \{i\}) - E(\boldsymbol{S})$ often ***S*** is an idle system

Problem: $\phi_i + \phi_j \neq E(\boldsymbol{S} \cup \{i,j\})$

Lin Zhong and Niraj K. Jha, "Graphical user interface energy characterization for handheld computers", in *Proc. Int. Conf. on Compilers, Architectures & Synthesis for Embedded Systems (CASES)*, Oct. 2003.

How can we evaluate an energy accounting policy?

How to split the utility bill?

How to split the profit?

$\boldsymbol{N}=\{1,2,..,n\}$	set of players
$\boldsymbol{S}$	subset of ***N*** (coalition)
$v(\boldsymbol{S})$	game surplus when ***S*** plays
$\phi_i(\boldsymbol{S})$	surplus received by player i

Shapley Value

How to determine the contribution of each individual player in a multi-player game?

Four axioms determines a unique distribution

L. S. Shapley, *A value for n-person games*: In Contributions to the Theory of Games, volume II, Annals of Mathematical Studies v. 28, pp. 307–317, Princeton Univ. Press, 1953.

System energy consumption (*E*): game surplus (v)

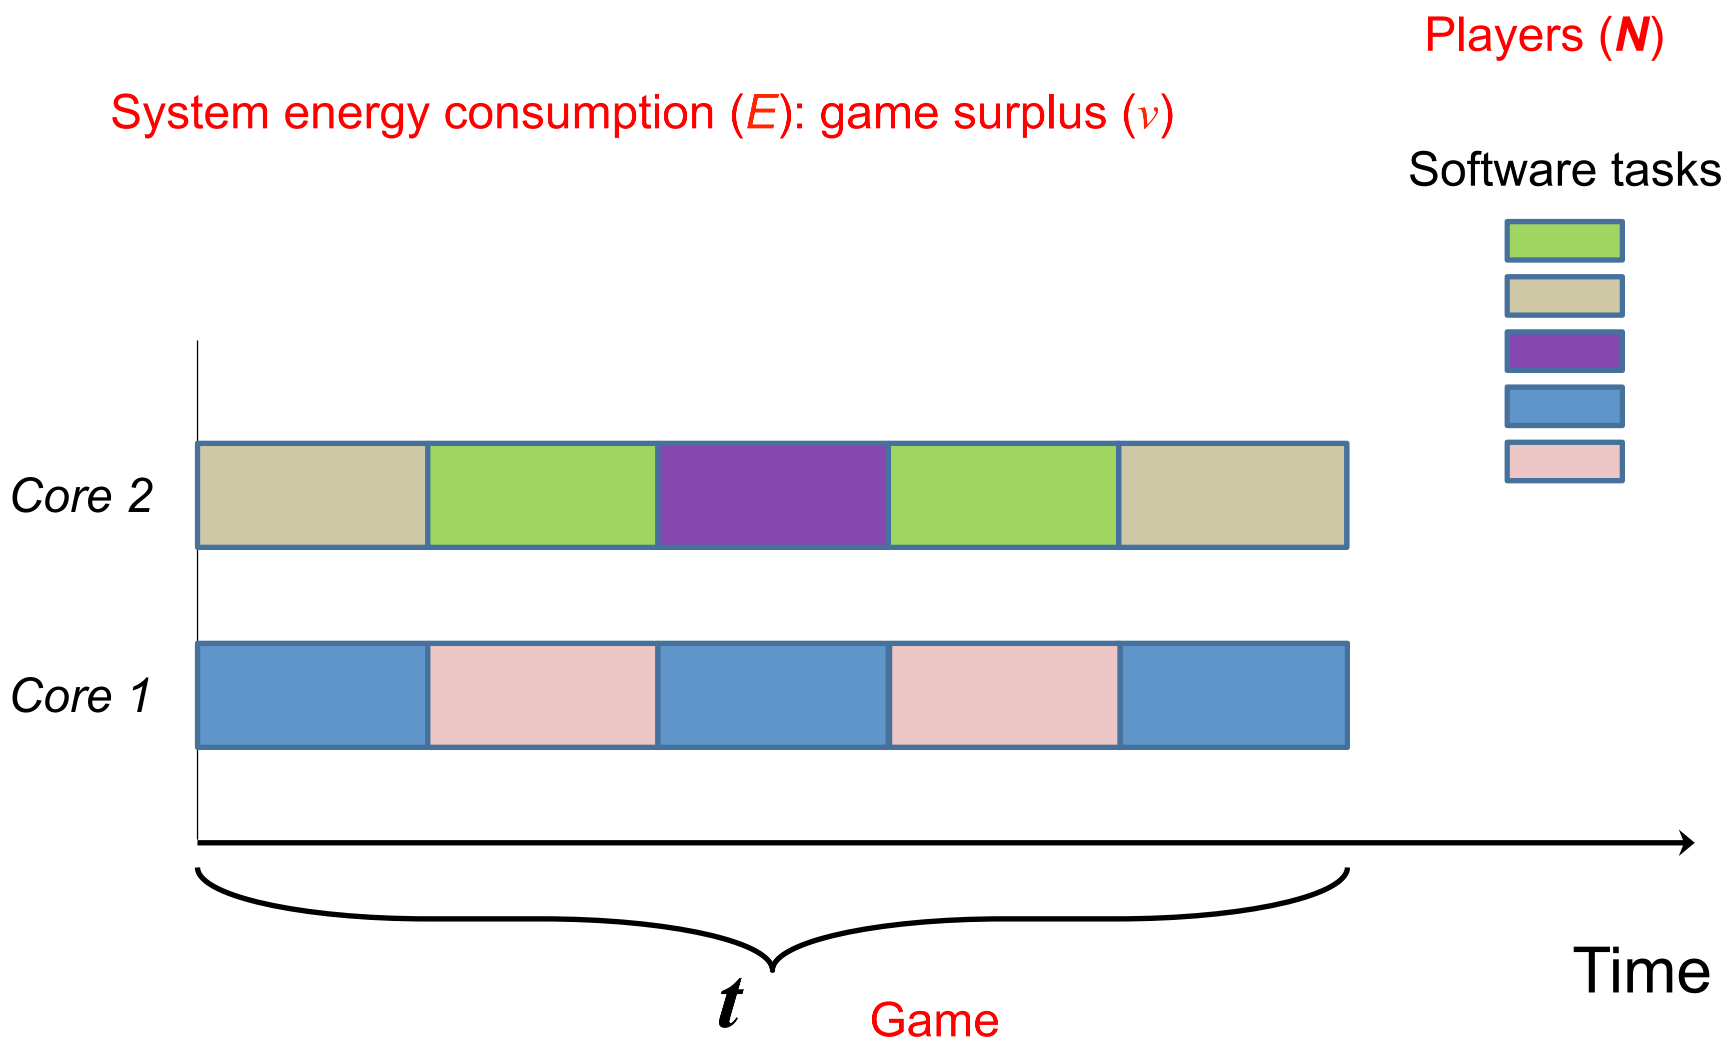

Axiom 1: Efficiency

The sum of the energy contributions by all tasks equals the system energy consumption.

Axiom 2: Symmetry

If replacing one with another will not change the system energy consumption under any circumstances, two tasks should have the same energy contributions.

Axiom 3: Null Player

If adding a task under any circumstances does not increase system energy consumption, this task should have zero energy contribution.

Axiom 4: Additivity

The same energy attribution policy should work for all the time intervals.

Shapley Value

$$\phi_i(E(\mathbf{N})) = \sum_{\mathbf{S} \subseteq \mathbf{N} \setminus \{i\}} \frac{E(\mathbf{S} \cup \{i\}) - E(\mathbf{S})}{(|\mathbf{N}| - |\mathbf{S}|)\binom{|\mathbf{N}|}{|\mathbf{S}|}}$$

Systems challenges

- $E(\boldsymbol{S})$ is highly random
- $E(\boldsymbol{S})$ not available for many $\boldsymbol{S}$
- $E(\boldsymbol{S})$ depends on hardware configuration

Challenge I: Shapley value only cares about IF a player participates in a game but not HOW

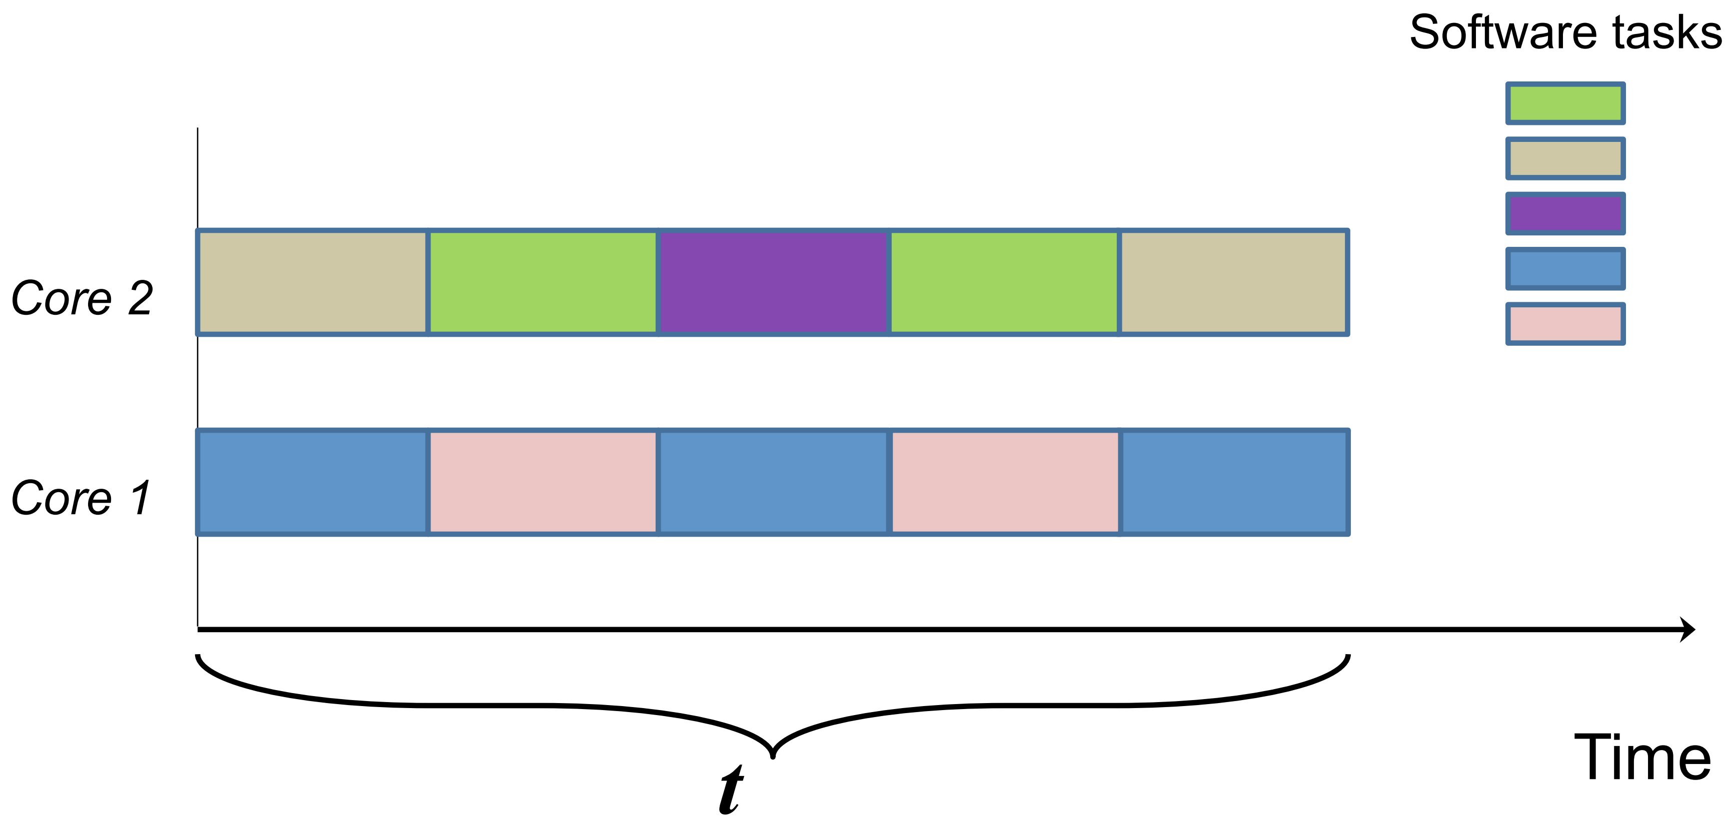

Challenge I: Shapley value only cares about IF a player participates in a game but not HOW

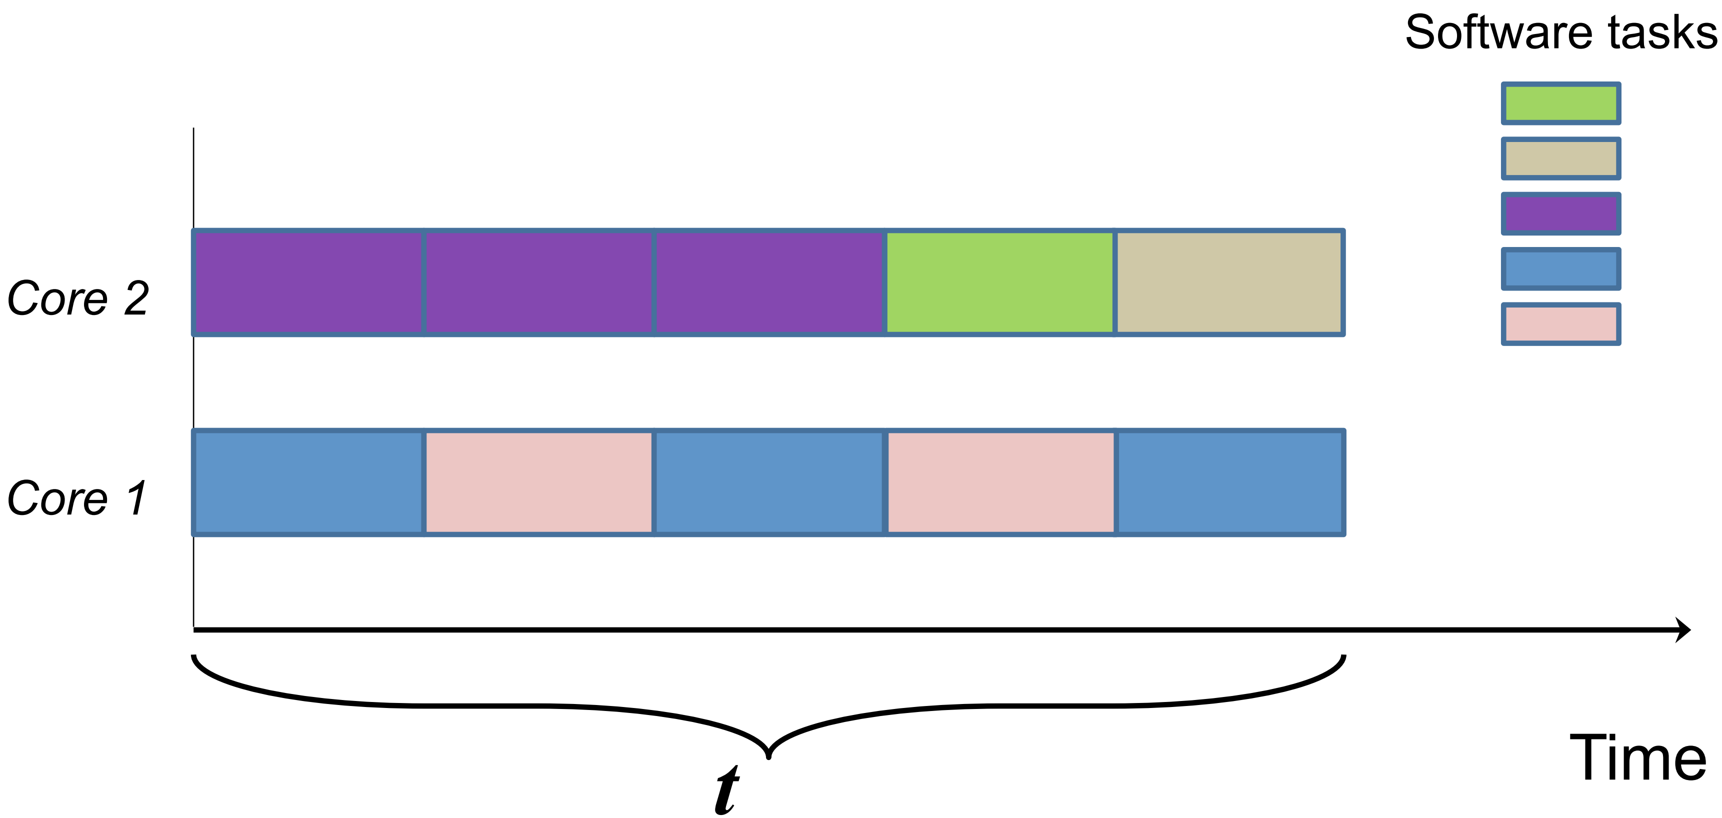

Challenge I: Shapley value only cares about IF a player participates in a game but not HOW

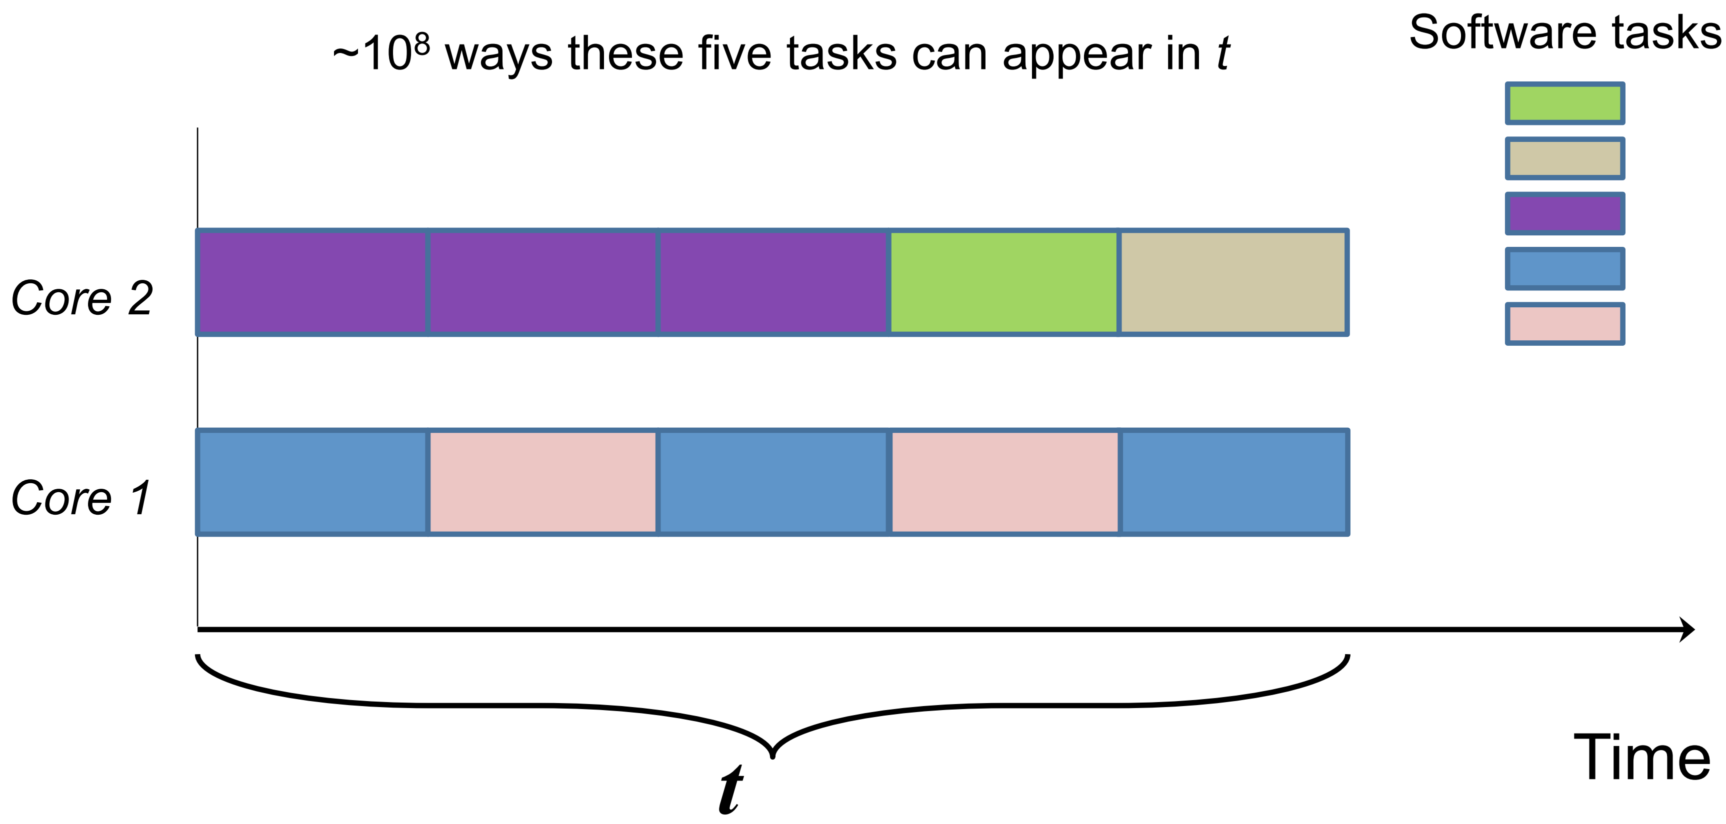

Challenge II: Not all combinations of tasks have been observed

To distribute system energy to the five tasks, Shapley value requires $E(\boldsymbol{S})$ for every subset $\boldsymbol{S}$ of the five tasks

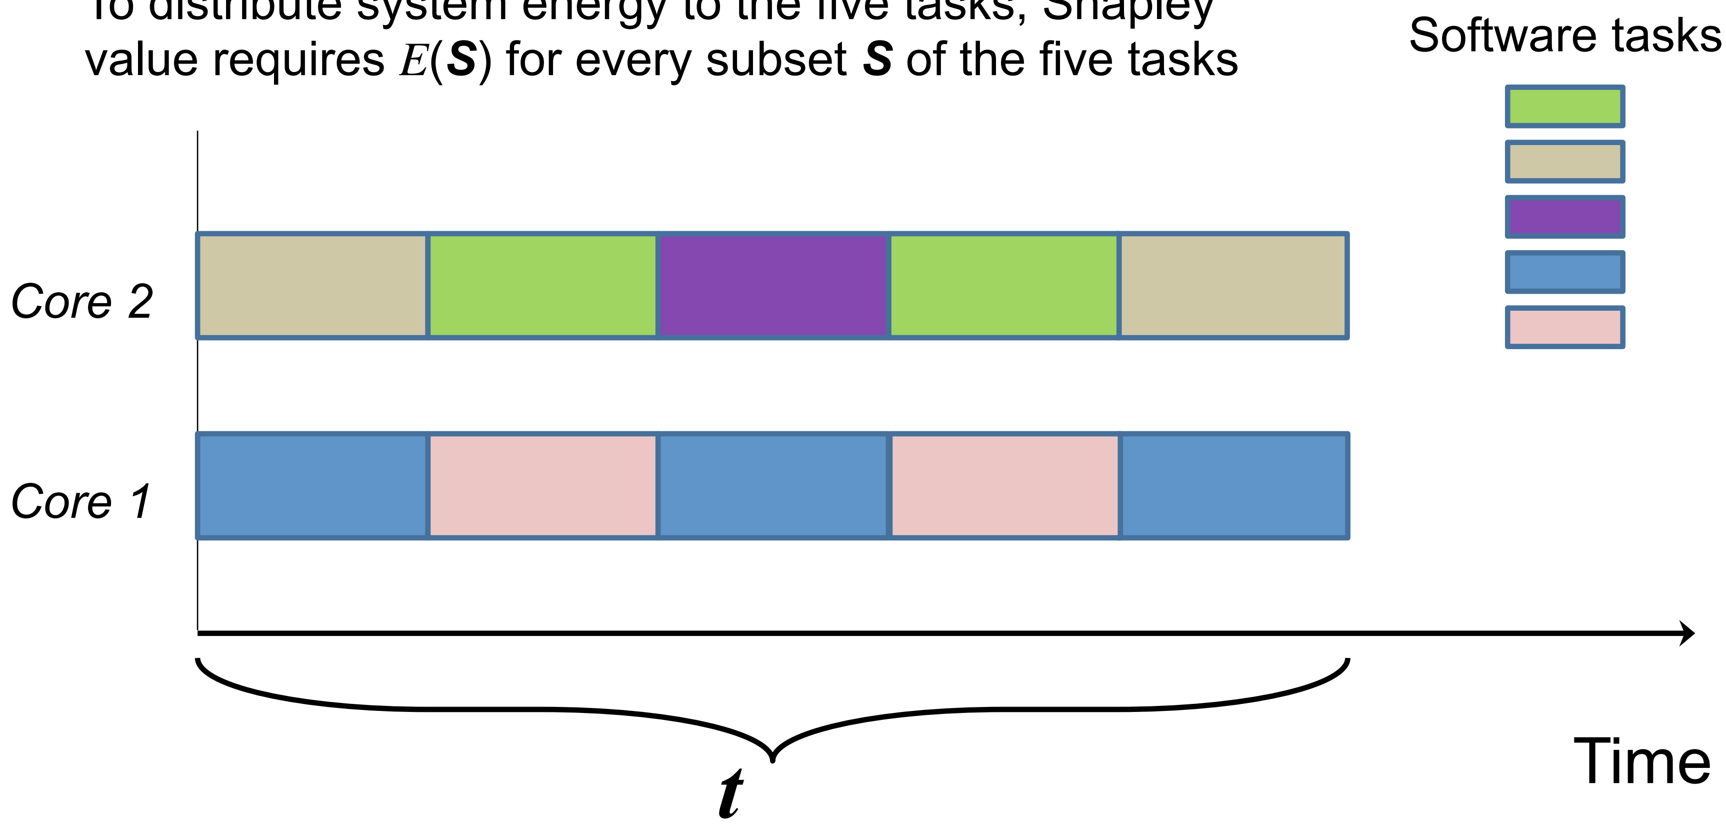

Solutions

- $E(\boldsymbol{S})$ is highly random
- $E(\boldsymbol{S})$ not available for all $\boldsymbol{S}$
  - Estimate E for short time interval (10 ms)
- E(S) depends on hardware configuration

Challenge I: Shapley value only cares about IF a player participates in a game but not HOW

~10^8 ways these five tasks can appear in *t*

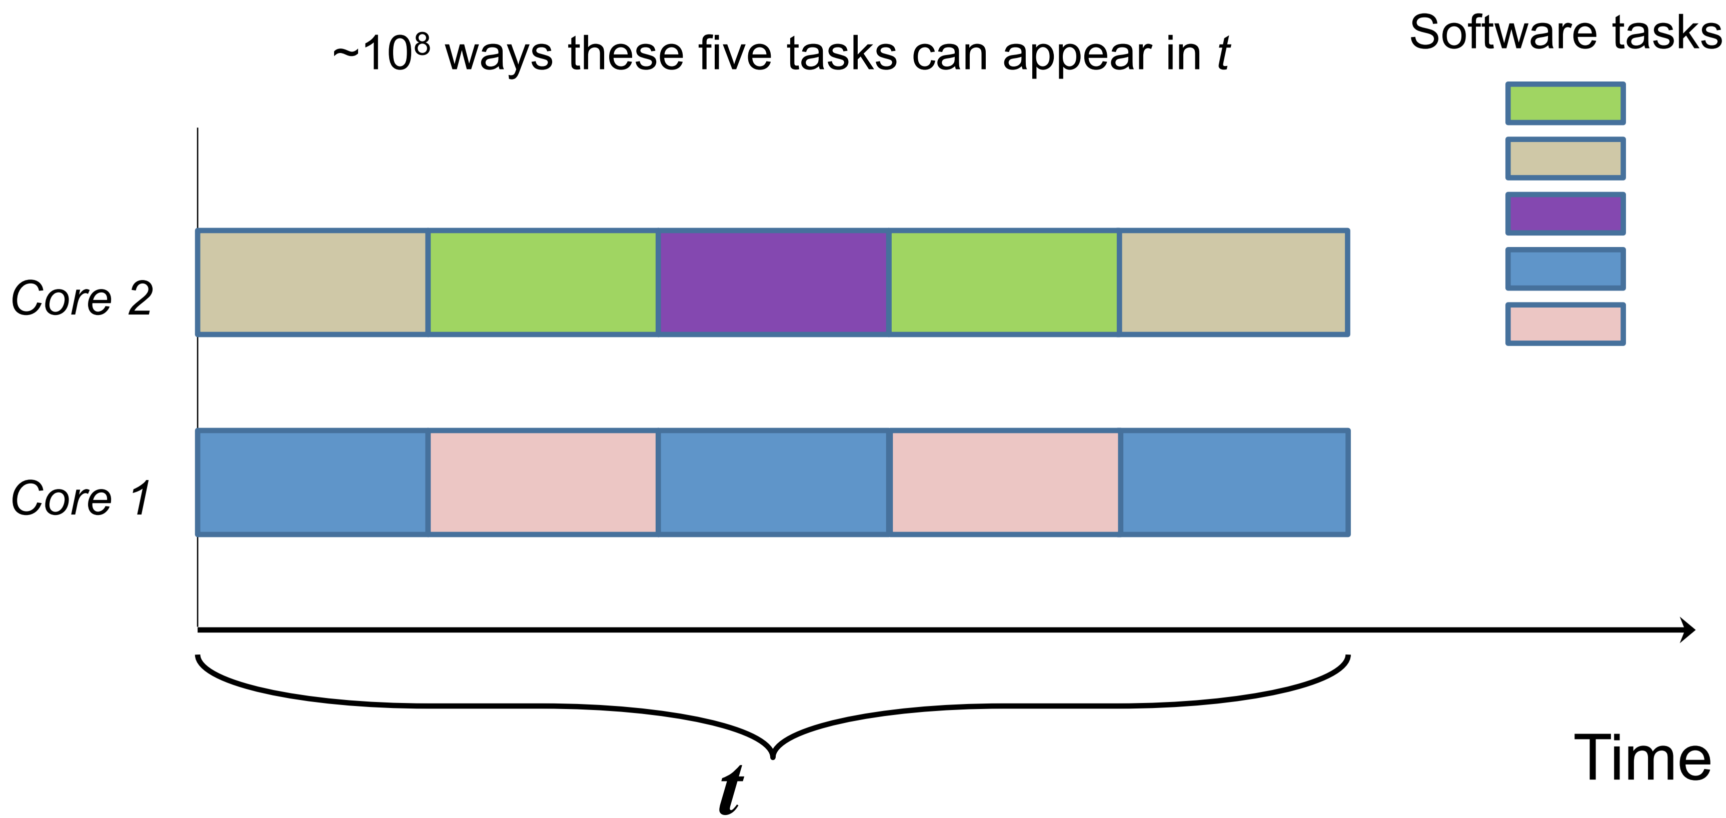

Challenge I: Shapley value only cares about IF a player participates in a game but not HOW

Only two tasks can appear in *t*
Two ways any two tasks can appear in *t*

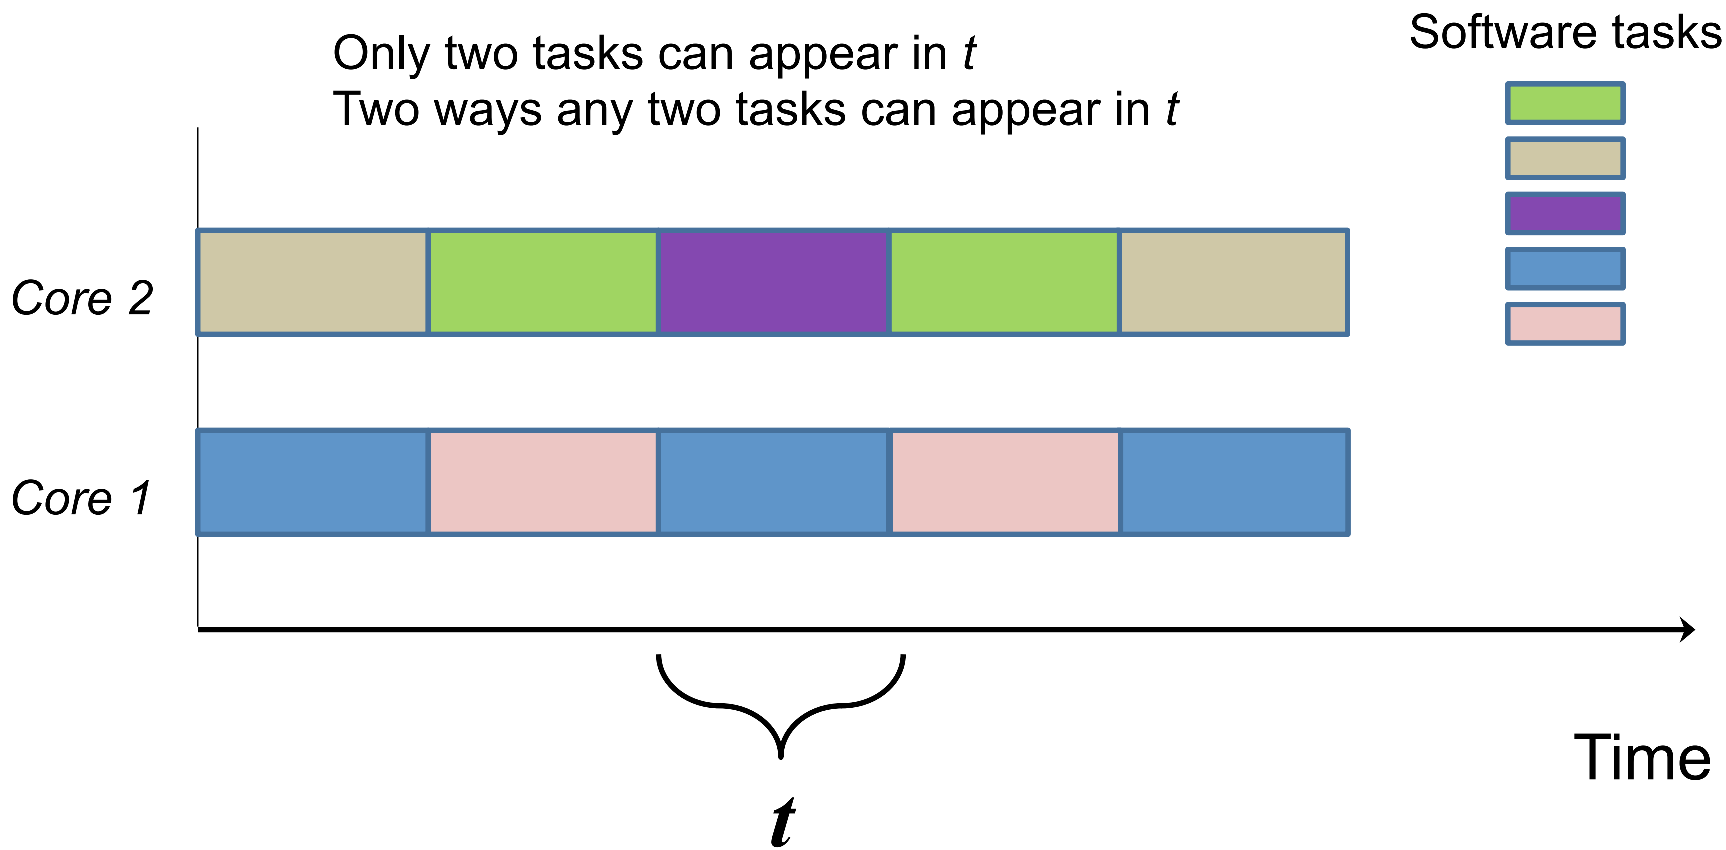

Challenge II: Not all combinations of tasks have been observed

To distribute system energy to the five tasks, Shapley value requires $E(\boldsymbol{S})$ for any subset $\boldsymbol{S}$ of the ***five*** tasks

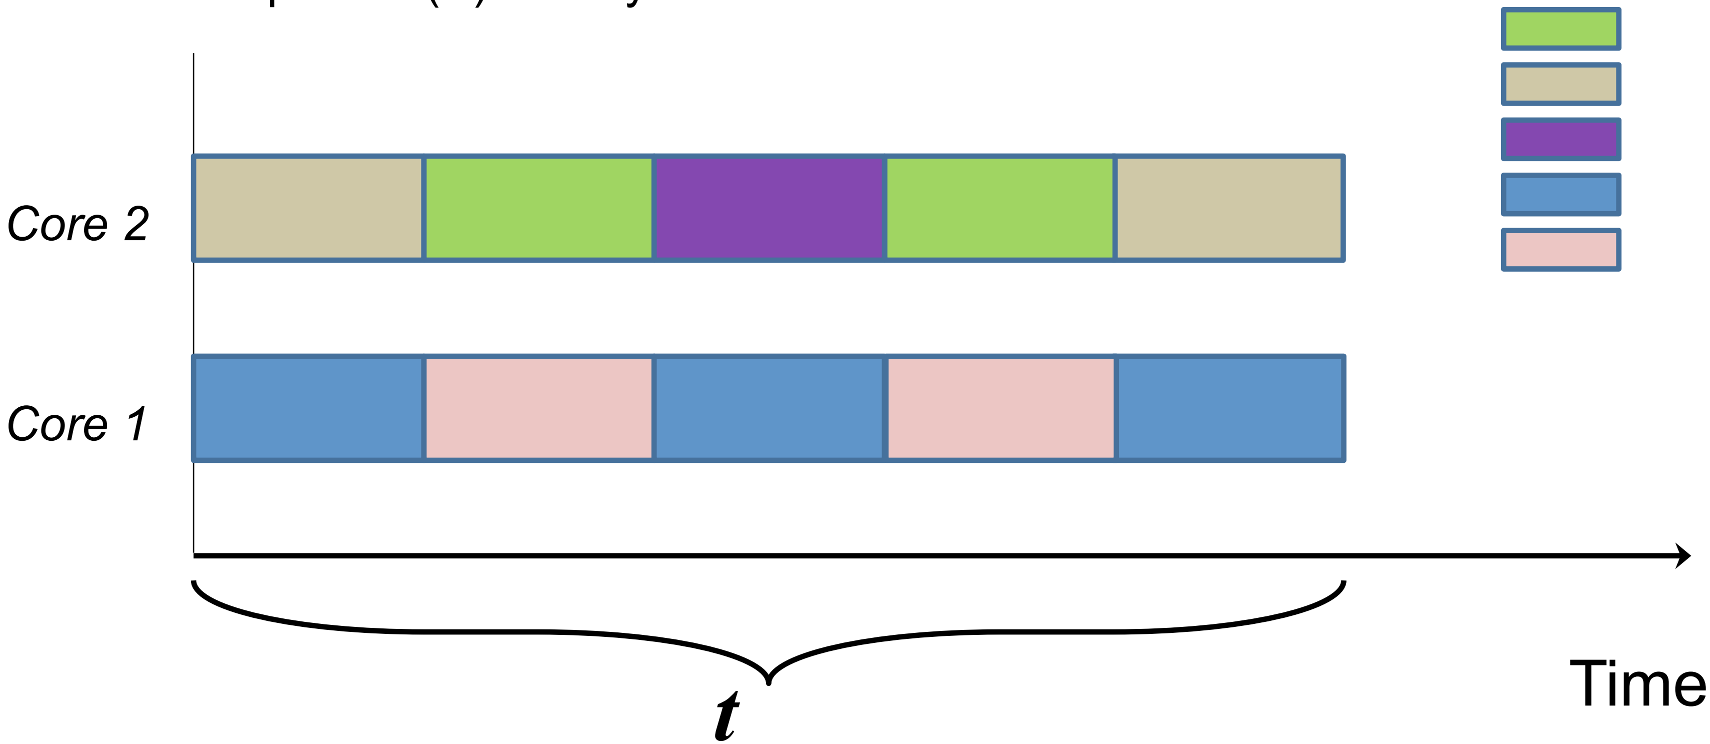

Challenge II: Not all combinations of tasks have been observed

To distribute system energy to the five tasks, Shapley value requires $E(\boldsymbol{S})$ for any subset $\boldsymbol{S}$ of the ***two*** tasks

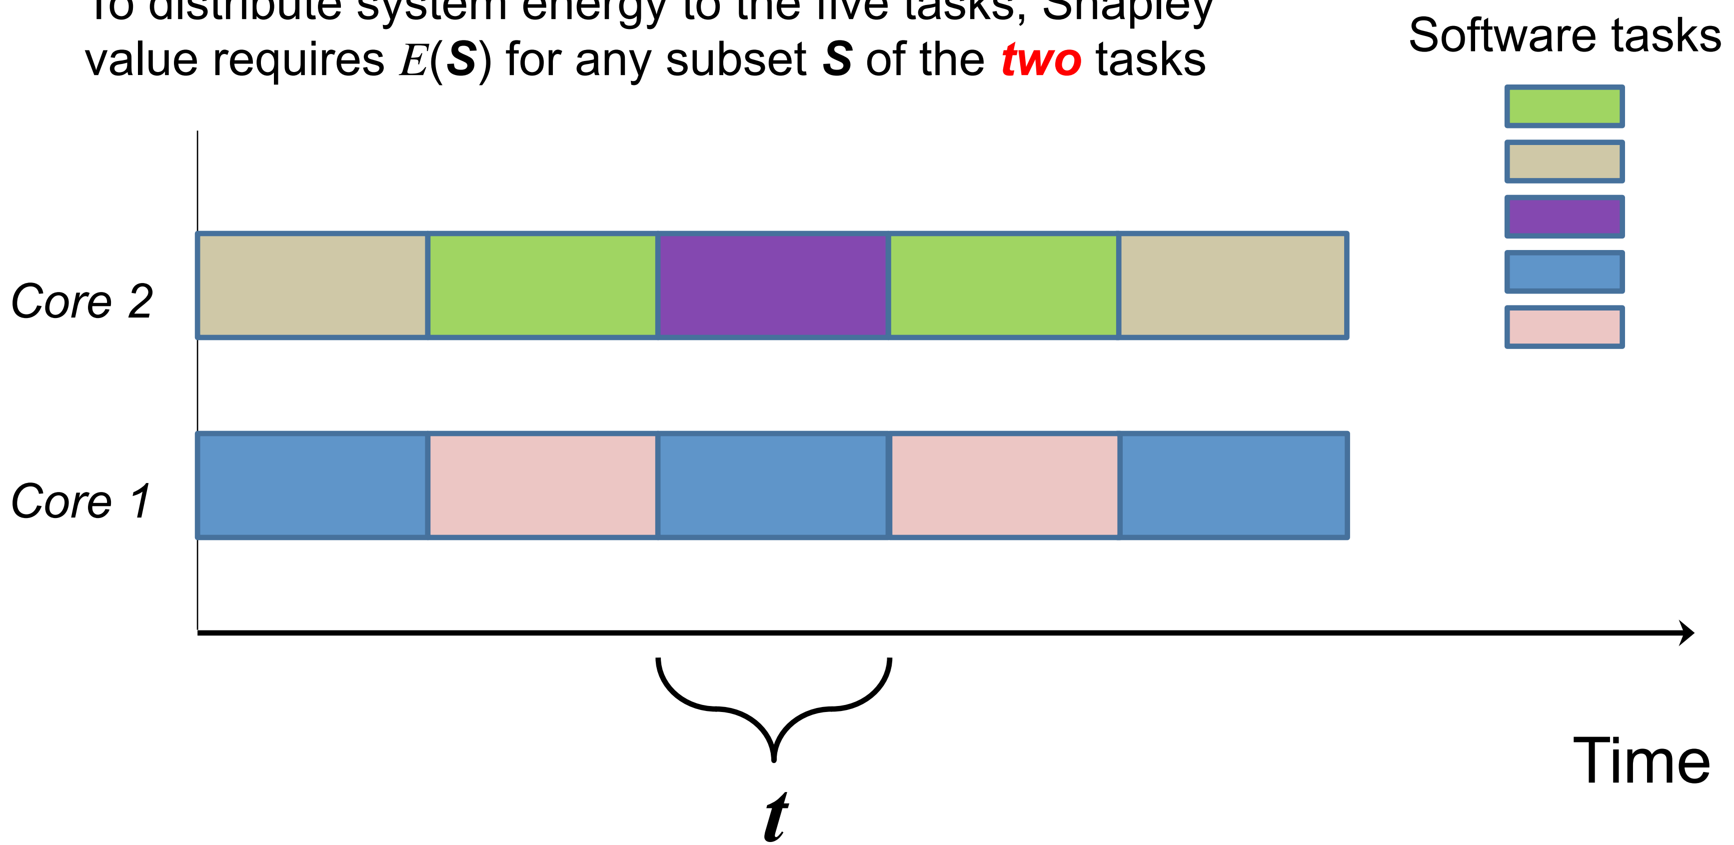

Solutions

- $E(\boldsymbol{S})$ is highly random
- $E(\boldsymbol{S})$ not available for all $\boldsymbol{S}$
  - Estimate E for short time interval (10 ms)
  - Estimate E *in situ*
- E(S) depends on hardware configuration

Smart battery interface

- Machine learning techniques (~80%)
  - *Dong and Zhong (MobiSys 2011)*

- Improved hardware/software (~95%)

Mian Dong and Lin Zhong, "Self-constructive, high-rate energy modeling for battery-powered mobile systems," in *Proc. ACM/USENIX Int. Conf. Mobile Systems, Applications, and Services* (*MobiSys*), June 2011.

Solutions

- E(***S***) is highly random
- E(***S***) not available for all ***S***
  - Estimate E for short time interval (10 ms)
  - Estimate E *in situ*
  - Use the crowd
- E(S) depends on hardware configuration

Solutions

- $E(\boldsymbol{S})$ is highly random
- $E(\boldsymbol{S})$ not available for all $\boldsymbol{S}$
  - Estimate E for short time interval (10 ms)
  - Estimate E *in situ*
  - Use the crowd
  - Extend Shapley Value framework
- E(S) depends on hardware configuration

Extend Shapley Value framework

- Approximate unknown $E(\boldsymbol{S})$ by their per-task energy cost allocations
- Example
  - Tasks: $\{1, 2, 3\}$
  - $E(\boldsymbol{S})$ known for $\boldsymbol{S} = \{1,2,3\}, \{2,3\}, \{1\}, \{2\}$
  - $E(\boldsymbol{S})$ can be recursively estimated for others
    - $\hat{E}(\{3\}) = E(\{1,2,3\}) - E(\{1\}) - E(\{2\})$
    - $\hat{E}(\{1,2\}) = E(\{1\}) + E(\{2\})$
    - $\hat{E}(\{1,3\}) = E(\{1\}) + E(\{3\})$

Prototype implementation

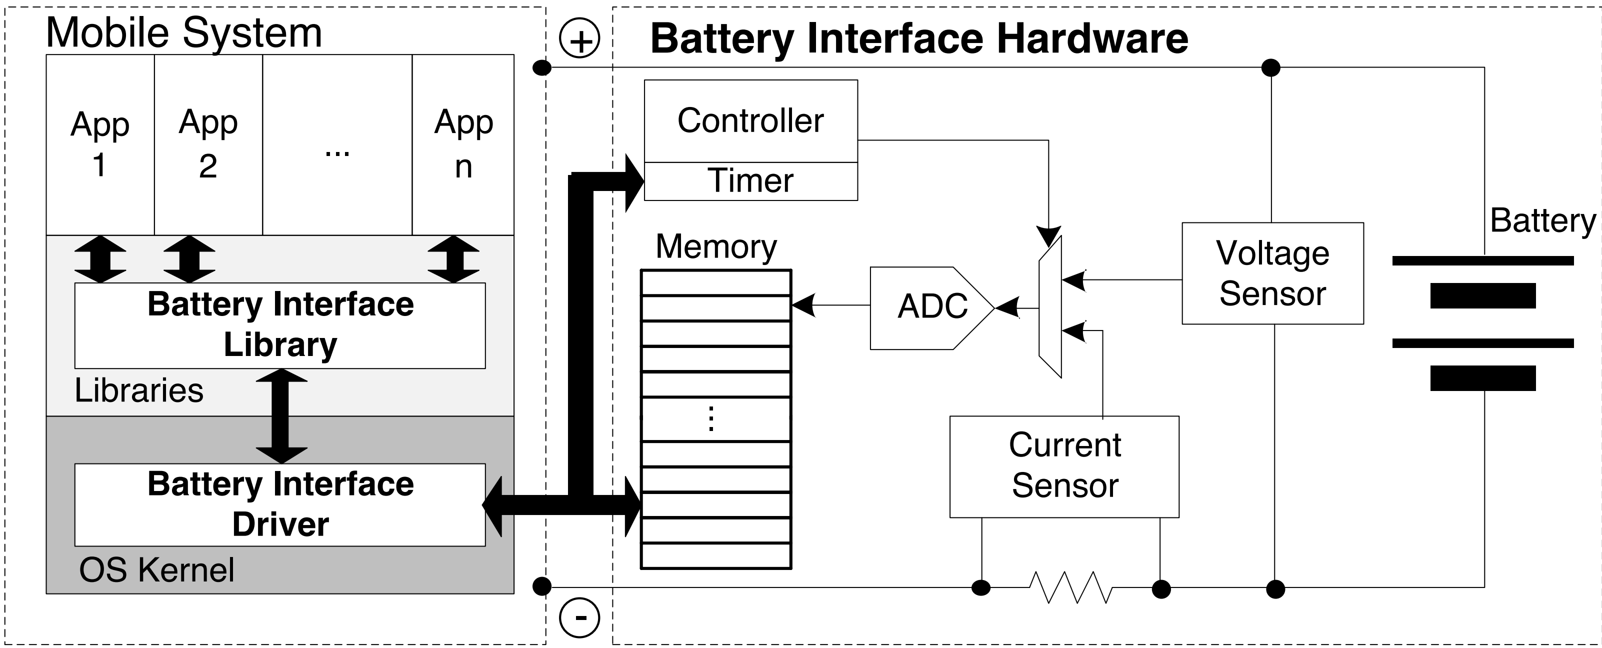

Texas Instruments Pandaboard (OMAP4430) with Android

MAXIM DS2756 battery fuel gauge

Evaluating existing policies

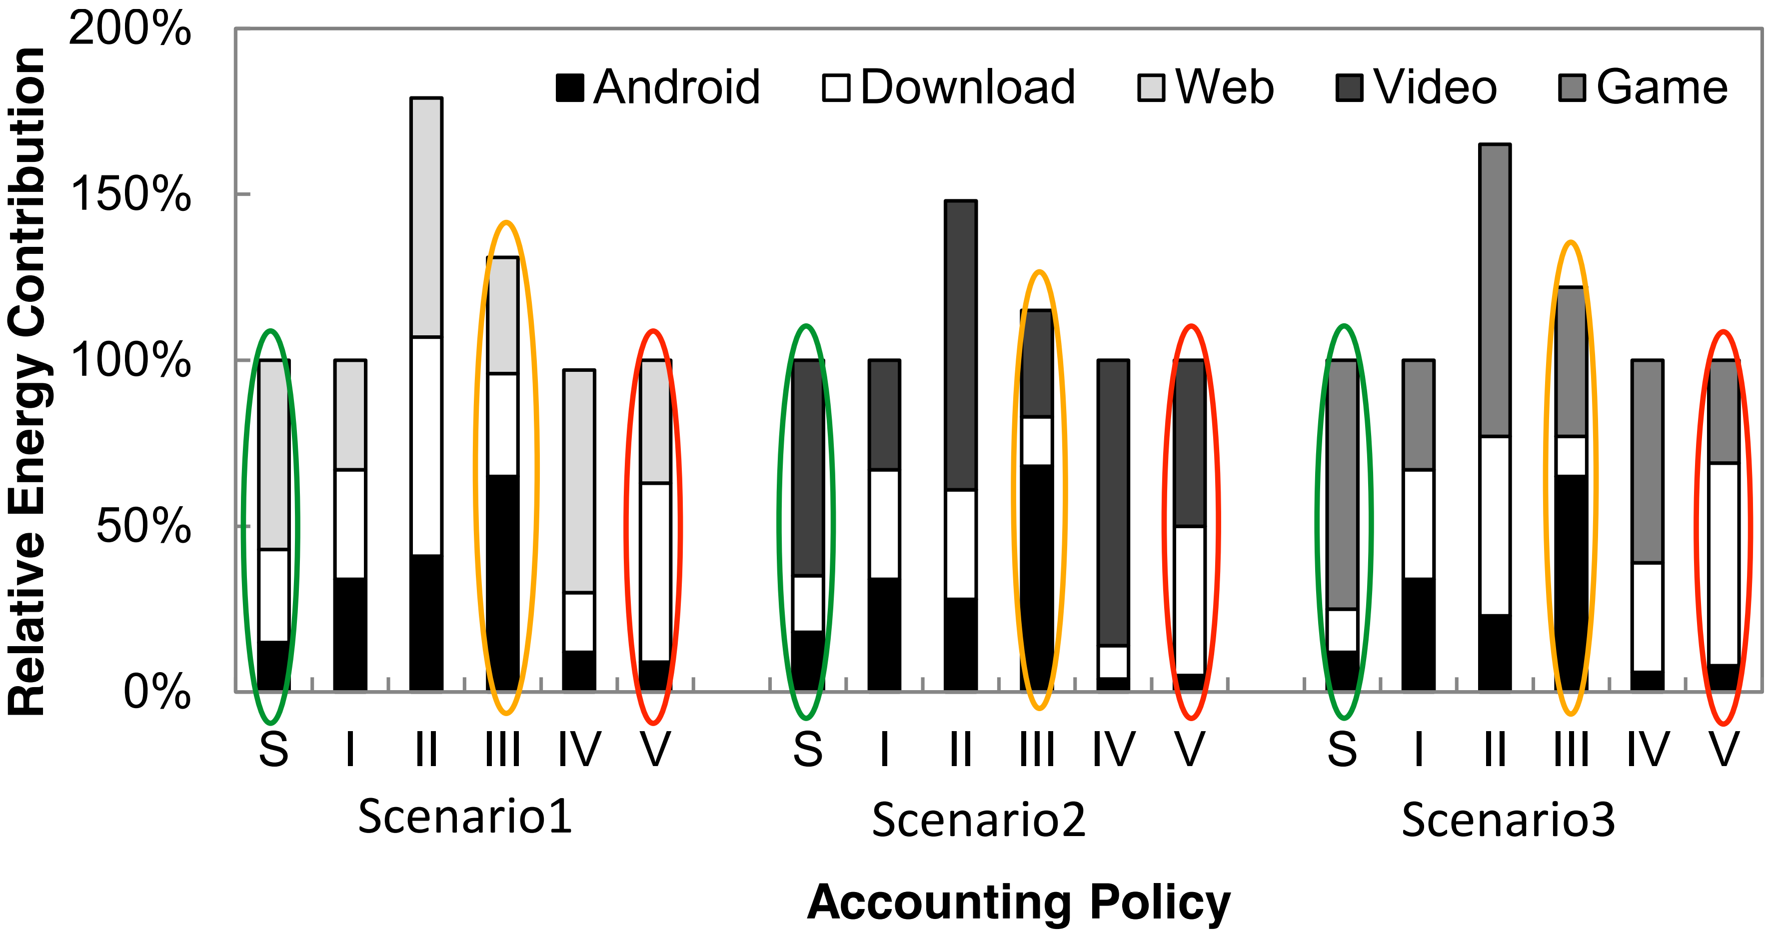

III = Lone-Wolf V = Model-driven

Conclusions

- Shapley value as ground truth for energy accounting
- System challenges can be addressed