

# Enabling the Usage of Multicore Platforms in Real-Time Safety-Critical Systems

James H. Anderson  
Kenan Professor, UNC Chapel Hill

CPS 1239135

CPS: Breakthrough: Collaborative Research: Bringing the Multicore  
Revolution to Safety-Critical Cyber-Physical Systems

PI: James H. Anderson, UNC Chapel Hill

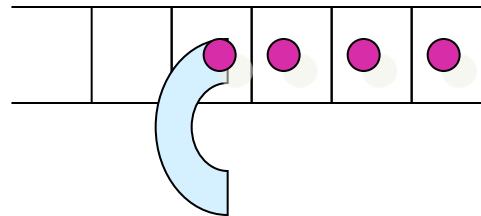
co-PI: Frank Mueller, NC State University

## Outline

- Real-time 101.
- Problems caused by multicore.
  - “The one-out-of-m problem.”
  - Why this is an important problem.
- Solution strategy.
- Evaluation.
- Project summary.

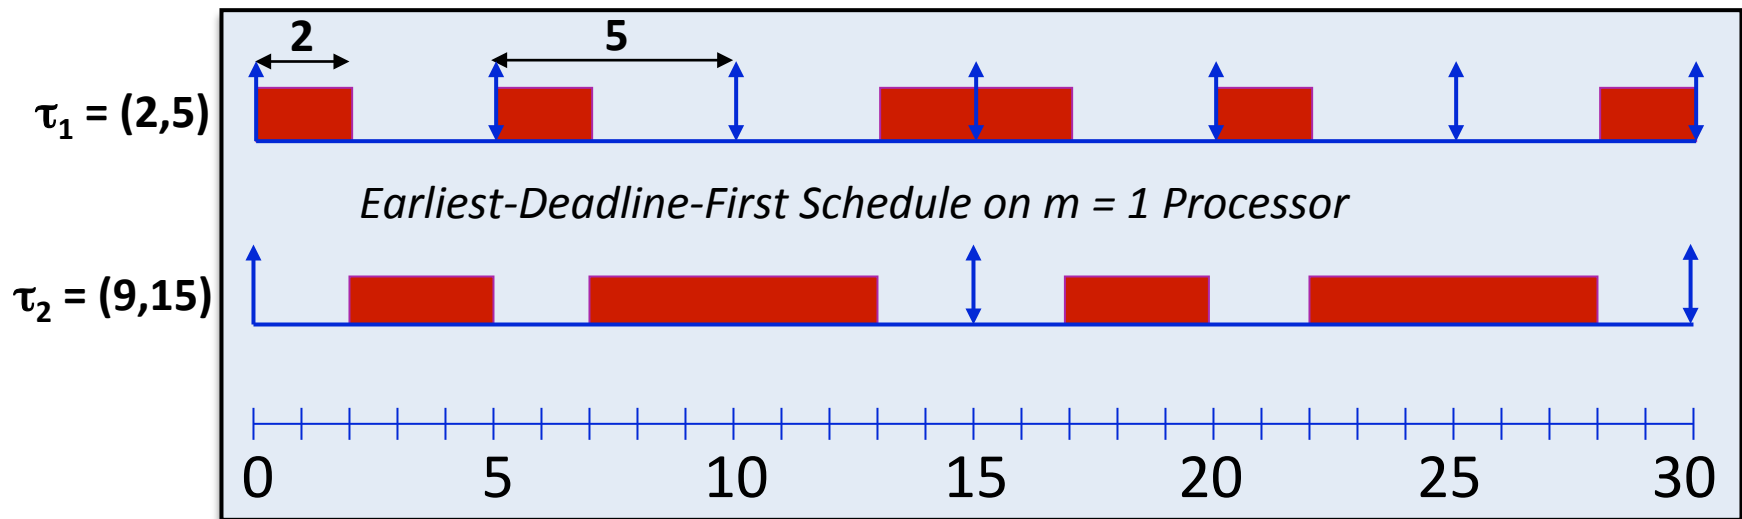
## What is a Real-Time System?

- A system with a dual notion of correctness:
  - *Logical correctness* (“it does the right thing”);
  - *Temporal correctness* (“it does it on time”).
- A system wherein *predictability* is as important as *performance*.
- **A Simple Example:** A robot arm picking up objects from a conveyor belt.



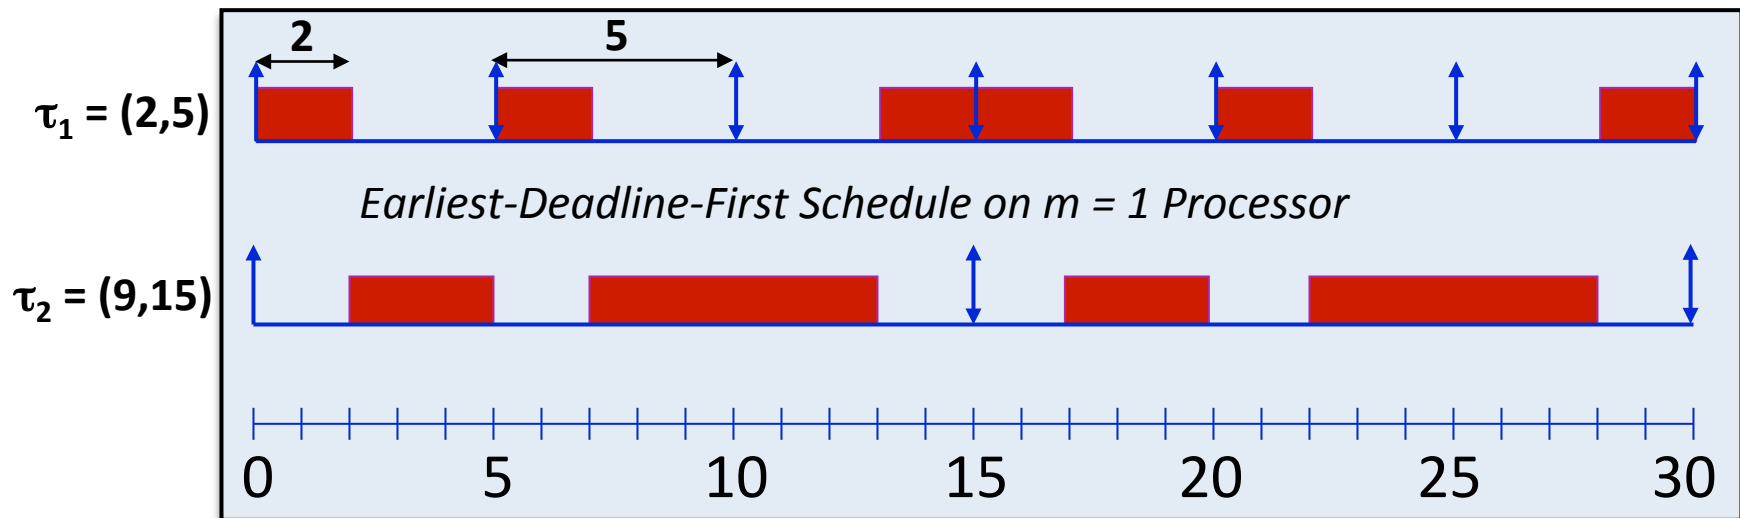
## Representing Real-Time Workloads: Periodic Task Systems

- Consider a set  $\tau$  of  $n$  **periodic tasks** on  $m$  processors:



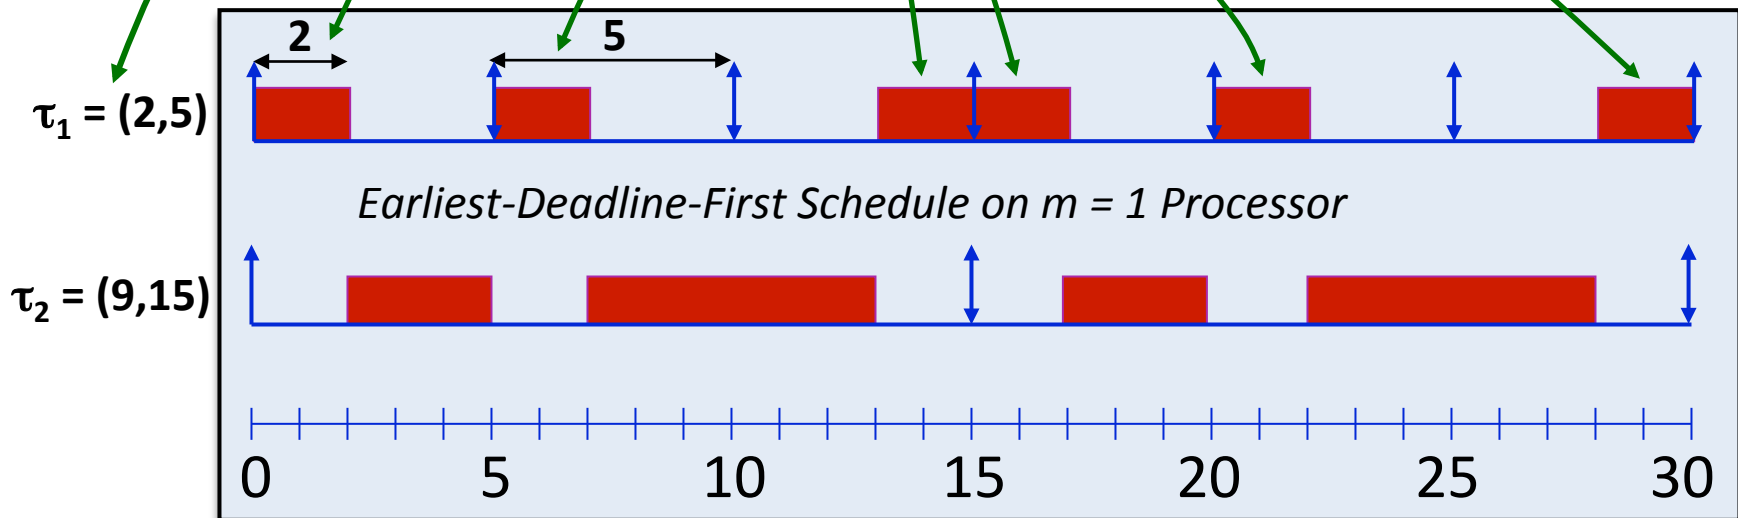
## Representing Real-Time Workloads: Periodic Task Systems

- Consider a set  $\tau$  of  $n$  **periodic tasks** on  $m$  processors:
  - Each task  $\tau_i = (C_i, T_i)$  releases a **job** with worst-case execution time (WCET)  $C_i$  at least  $T_i$  time units apart.
    - $\tau_i$ 's **utilization** is  $U_i = C_i/T_i$ .
    - Total utilization** is  $U(\tau) = \sum_i C_i/T_i$ .



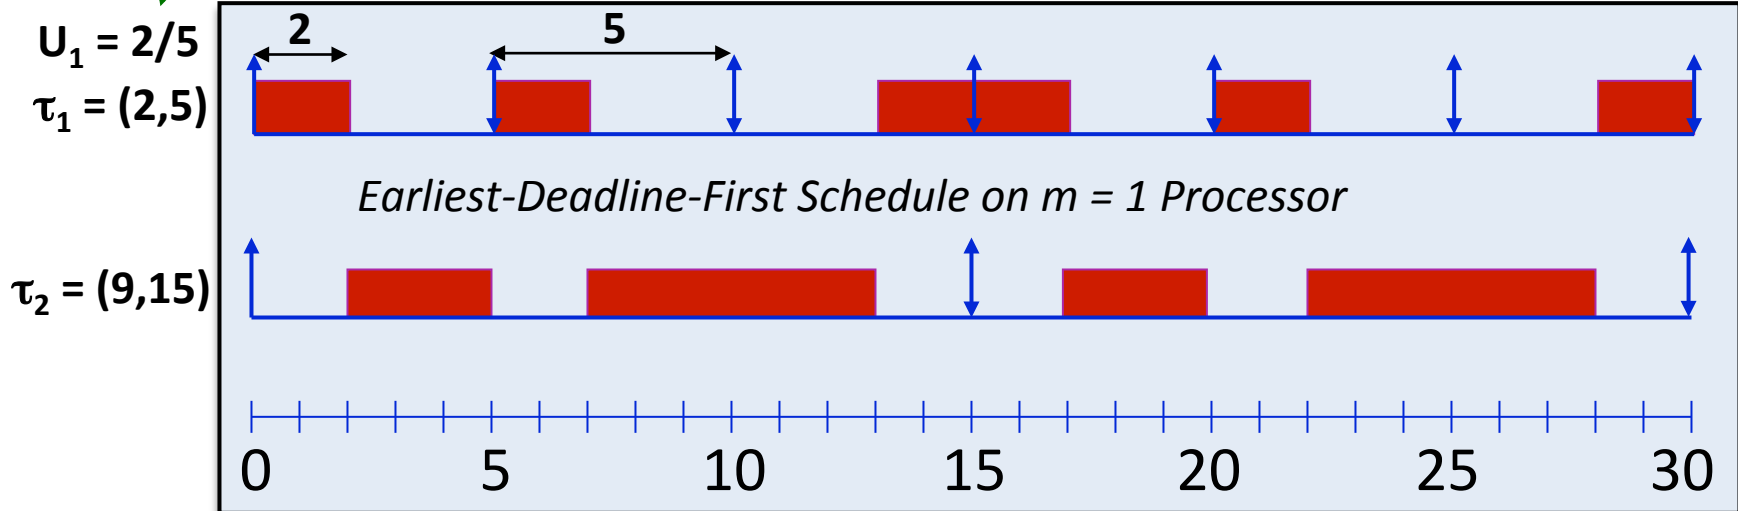
## Representing Real-Time Workloads: Periodic Task Systems

- Consider a set  $\tau$  of  $n$  **periodic tasks** on  $m$  processors:
  - Each task  $\tau_i = (C_i, T_i)$  releases a **job** with worst-case execution time (WCET)  $C_i$  at least  $T_i$  time units apart.
    - $\tau_i$ 's **utilization** is  $U_i = C_i/T_i$ .
    - Total utilization** is  $U(\tau) = \sum_i C_i/T_i$ .



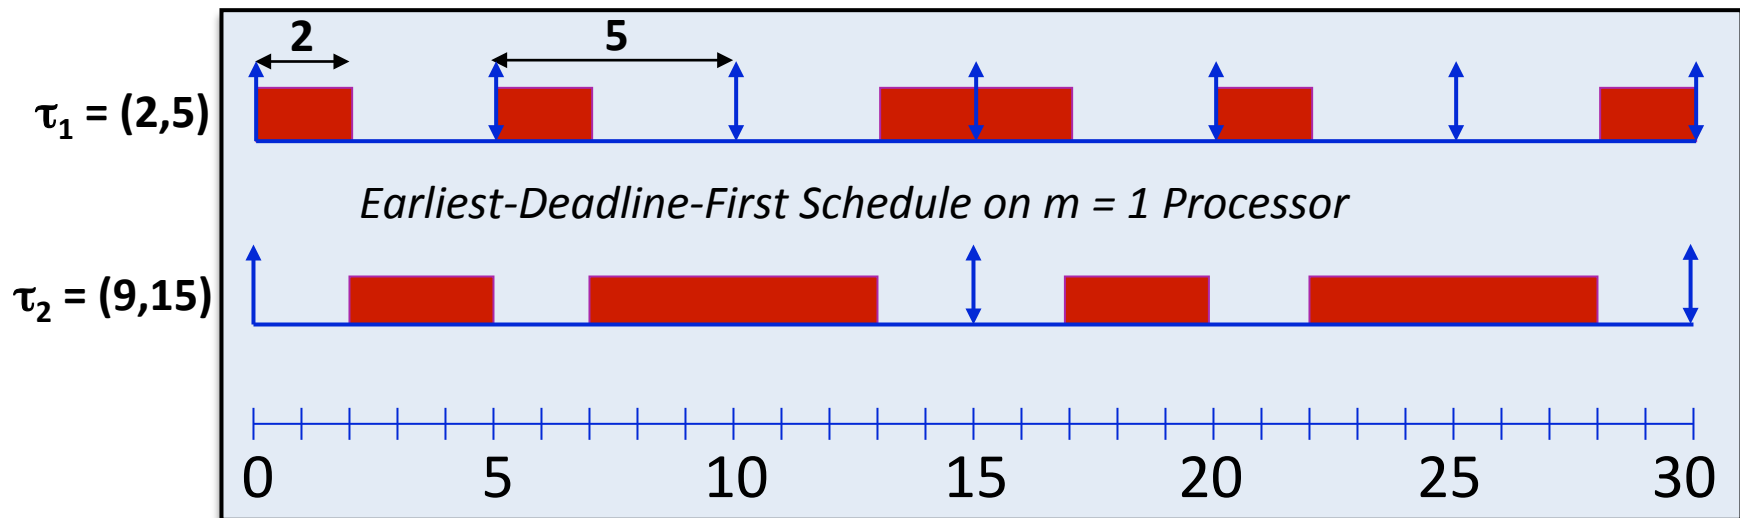
## Representing Real-Time Workloads: Periodic Task Systems

- Consider a set  $\tau$  of  $n$  **periodic tasks** on  $m$  processors:
  - Each task  $\tau_i = (C_i, T_i)$  releases a **job** with worst-case execution time (WCET)  $C_i$  at least  $T_i$  time units apart.
    - $\tau_i$ 's **utilization** is  $U_i = C_i/T_i$ .
    - Total utilization** is  $U(\tau) = \sum_i C_i/T_i$ .



## Representing Real-Time Workloads: Periodic Task Systems

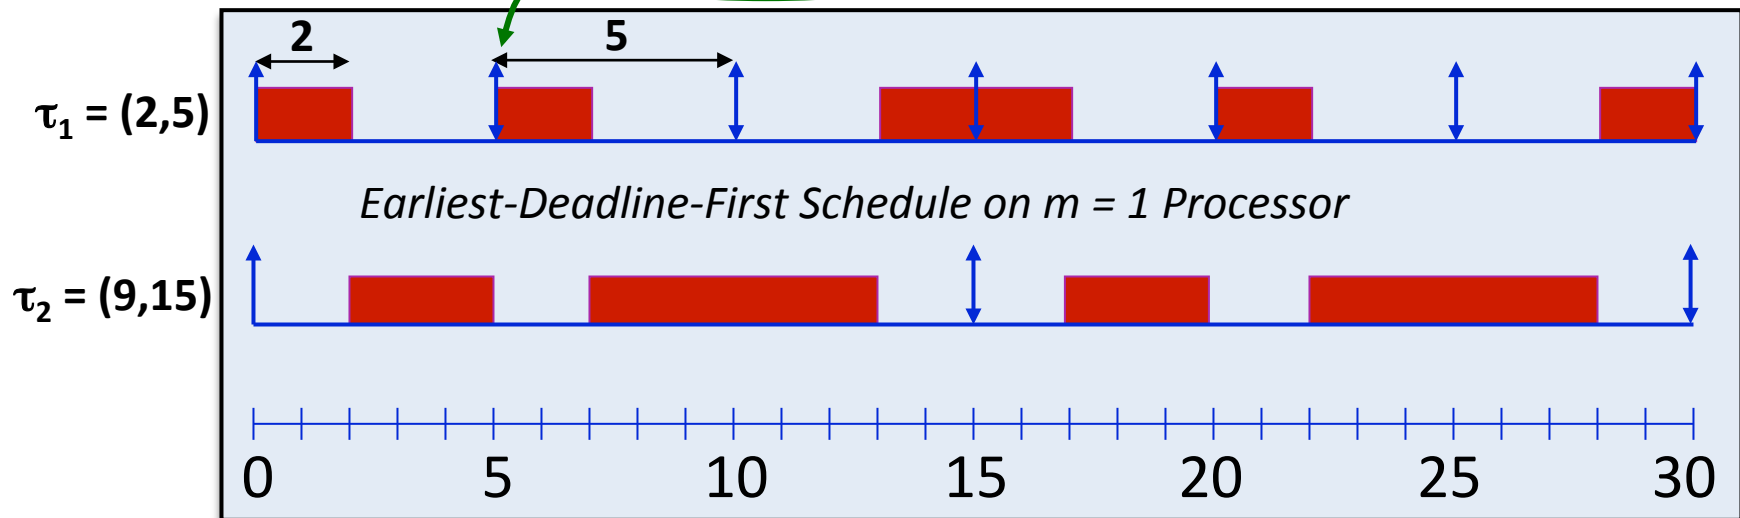
- Consider a set  $\tau$  of  $n$  **periodic tasks** on  $m$  processors:
  - Each task  $\tau_i = (C_i, T_i)$  releases a **job** with worst-case execution time (WCET)  $C_i$  at least  $T_i$  time units apart.
    - $\tau_i$ 's **utilization** is  $U_i = C_i/T_i$ .
    - Total utilization** is  $U(\tau) = \sum_i C_i/T_i$ .
  - Each job of  $\tau_i$  has a **relative deadline** given by  $T_i$ .





## Representing Real-Time Workloads: Periodic Task Systems

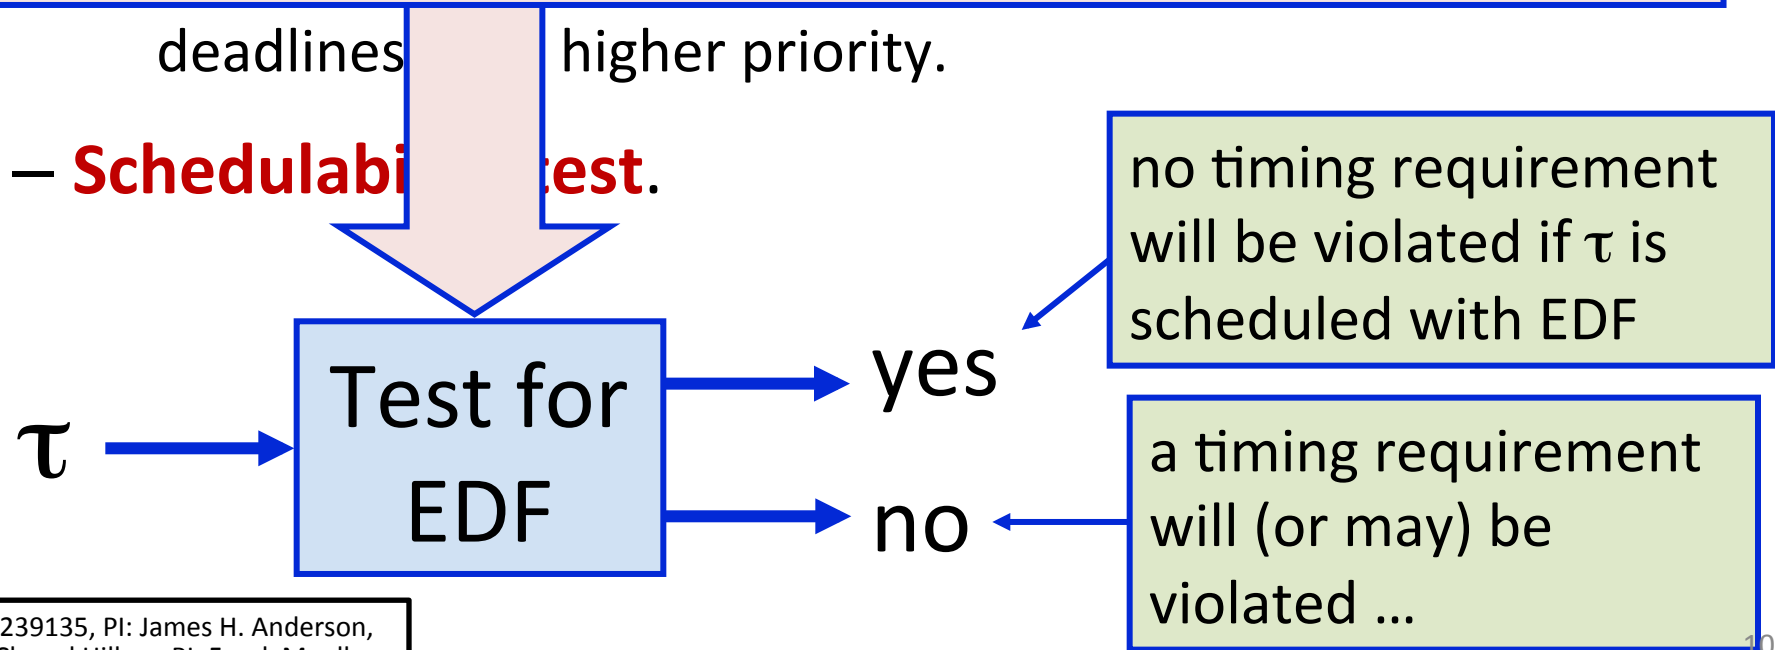
- Consider a set  $\tau$  of  $n$  **periodic tasks** on  $m$  processors:
  - Each task  $\tau_i = (C_i, T_i)$  releases a **job** with worst-case execution time (WCET)  $C_i$  at least  $T_i$  time units apart.
    - $\tau_i$ 's **utilization** is  $U_i = C_i/T_i$ .
    - Total utilization** is  $U(\tau) = \sum_i C_i/T_i$ .
  - Each job of  $\tau_i$  has a **relative deadline** given by  $T_i$ .



## Checking Timing Constraints: Schedulability Tests

- W.r.t. scheduling, we actually care about two kinds

**Capacity loss** occurs when test requires conservative assumptions about task execution times, the manner in which tasks delay each other, etc.



## Outline

- Real-time 101.
- Problems caused by multicore.
  - “The one-out-of-m problem.”
  - Why this is an important problem.
- Solution strategy.
- Evaluation.
- Project summary.

## Schedulability Problems Caused by Multicore

- With multicore, the state of the art today in many safety-critical domains such as avionics is to ***disable all but one core.***
- Why?
  - On a machine with  $m$  cores, analysis pessimism can easily cause capacity loss of  $m - 1$  or greater.
  - We call this the **“one-out-of- $m$ ” problem:**
    - Given  $m$  cores, the capacity of the additional  $m - 1$  cores is wasted.*
    - *This is the most important open problem concerning “real-time on multicore” today.*

## The Root of the Problem

- This problem is caused by the presence of **shared hardware** (caches, buses, memory banks, etc.) that is not **predictably managed**.
- See the **FAA position paper CAST 32** for an extensive discussion of problems caused by shared-hardware interference.

## Outline

- Real-time 101.
- Problems caused by multicore.
  - “The one-out-of-m problem.”
  - Why this is an important problem.
- Solution strategy.
- Evaluation.
- Project summary.

## Our Solution Strategy

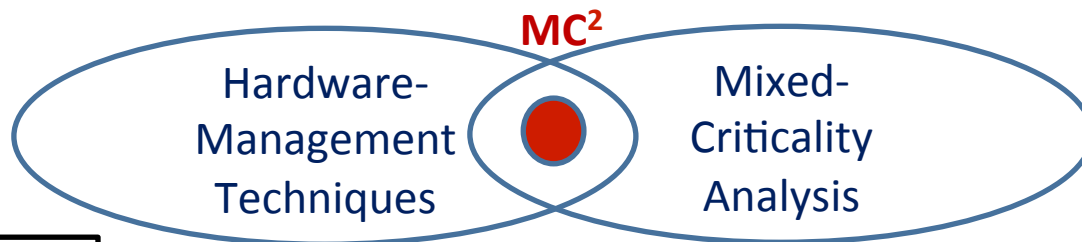
- W.r.t. lessening capacity loss generally (even on uniprocessors), two orthogonal approaches have been investigated previously:
  - **Hardware-management techniques** that reduce hardware interference.
  - **Mixed-criticality analysis techniques** that enable less critical tasks to be provisioned less pessimistically.

Hardware-  
Management  
Techniques

Mixed-  
Criticality  
Analysis

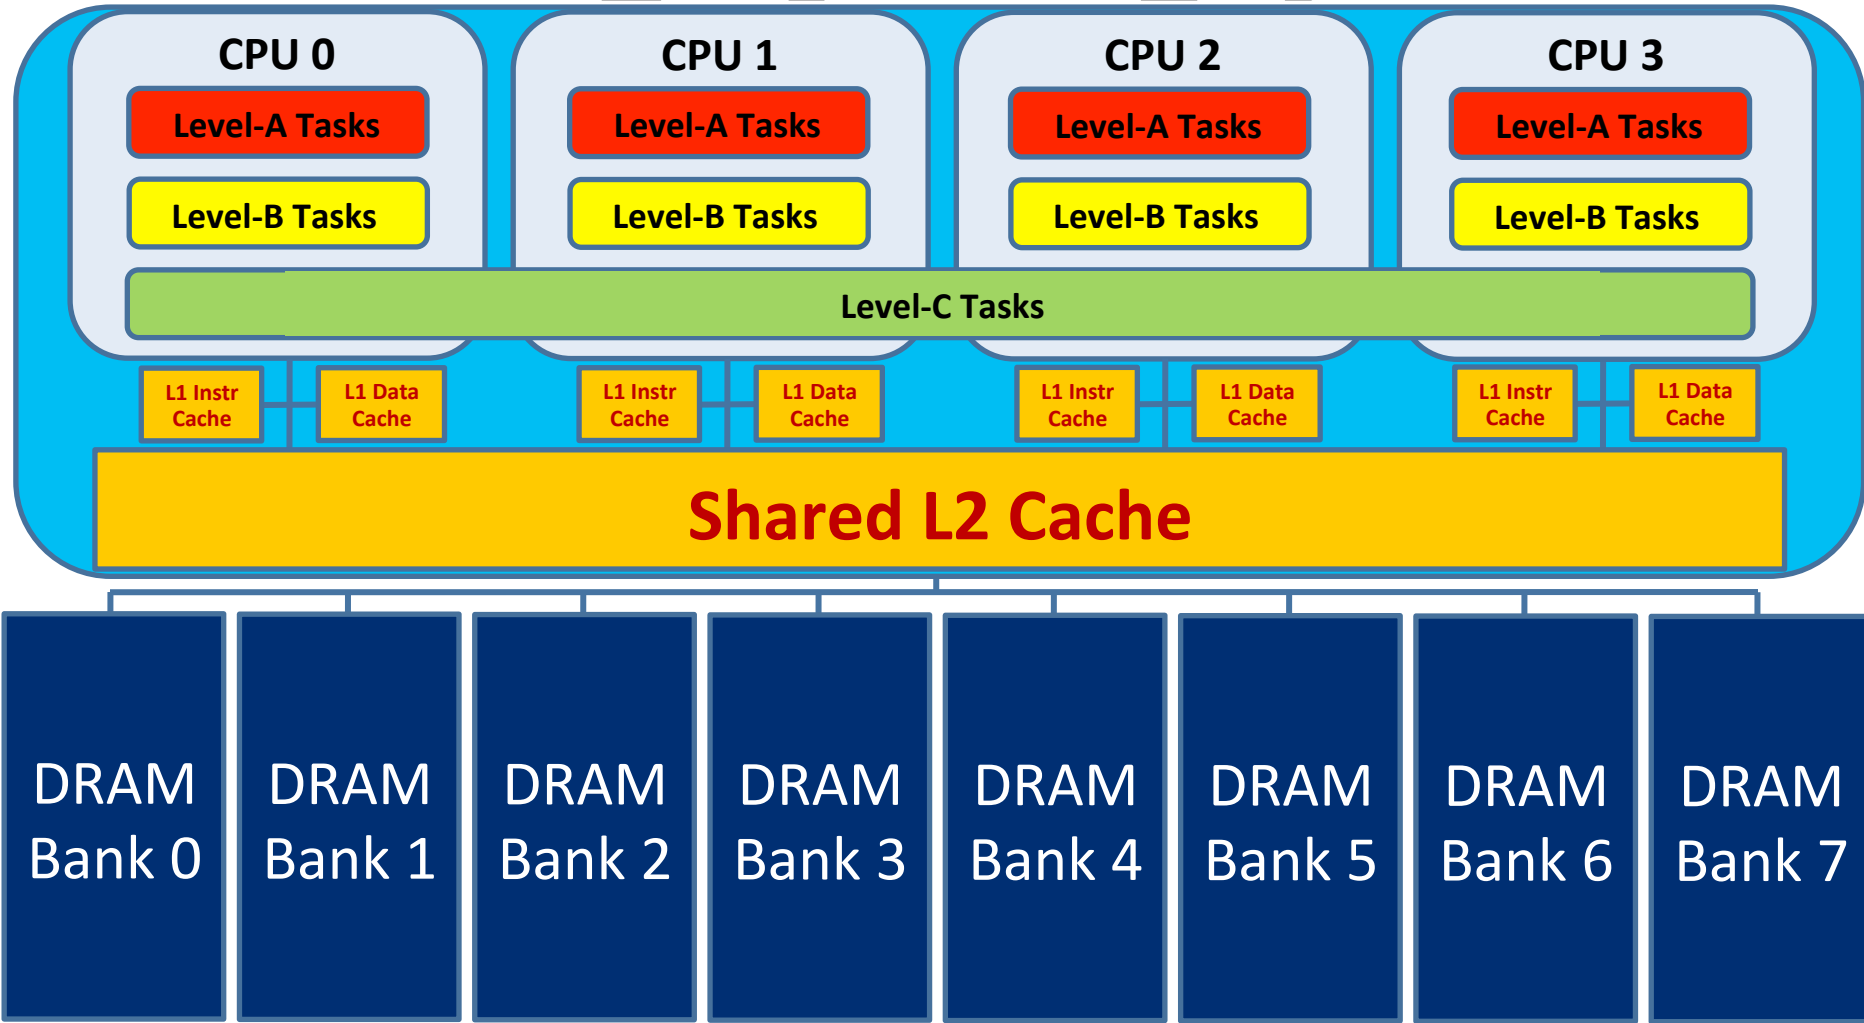
## Our Solution Strategy

- Our work focuses broadly on research questions that arise when applying both approaches together.
  - *Can better platform utilization be realized if hardware resources are managed differently at different criticality levels?*
  - *If so, how should resources be managed both within and across criticality levels?*
  - We are addressing such questions in the context of a resource-allocation and analysis framework developed by us called **MC<sup>2</sup> (mixed criticality on multicore).**



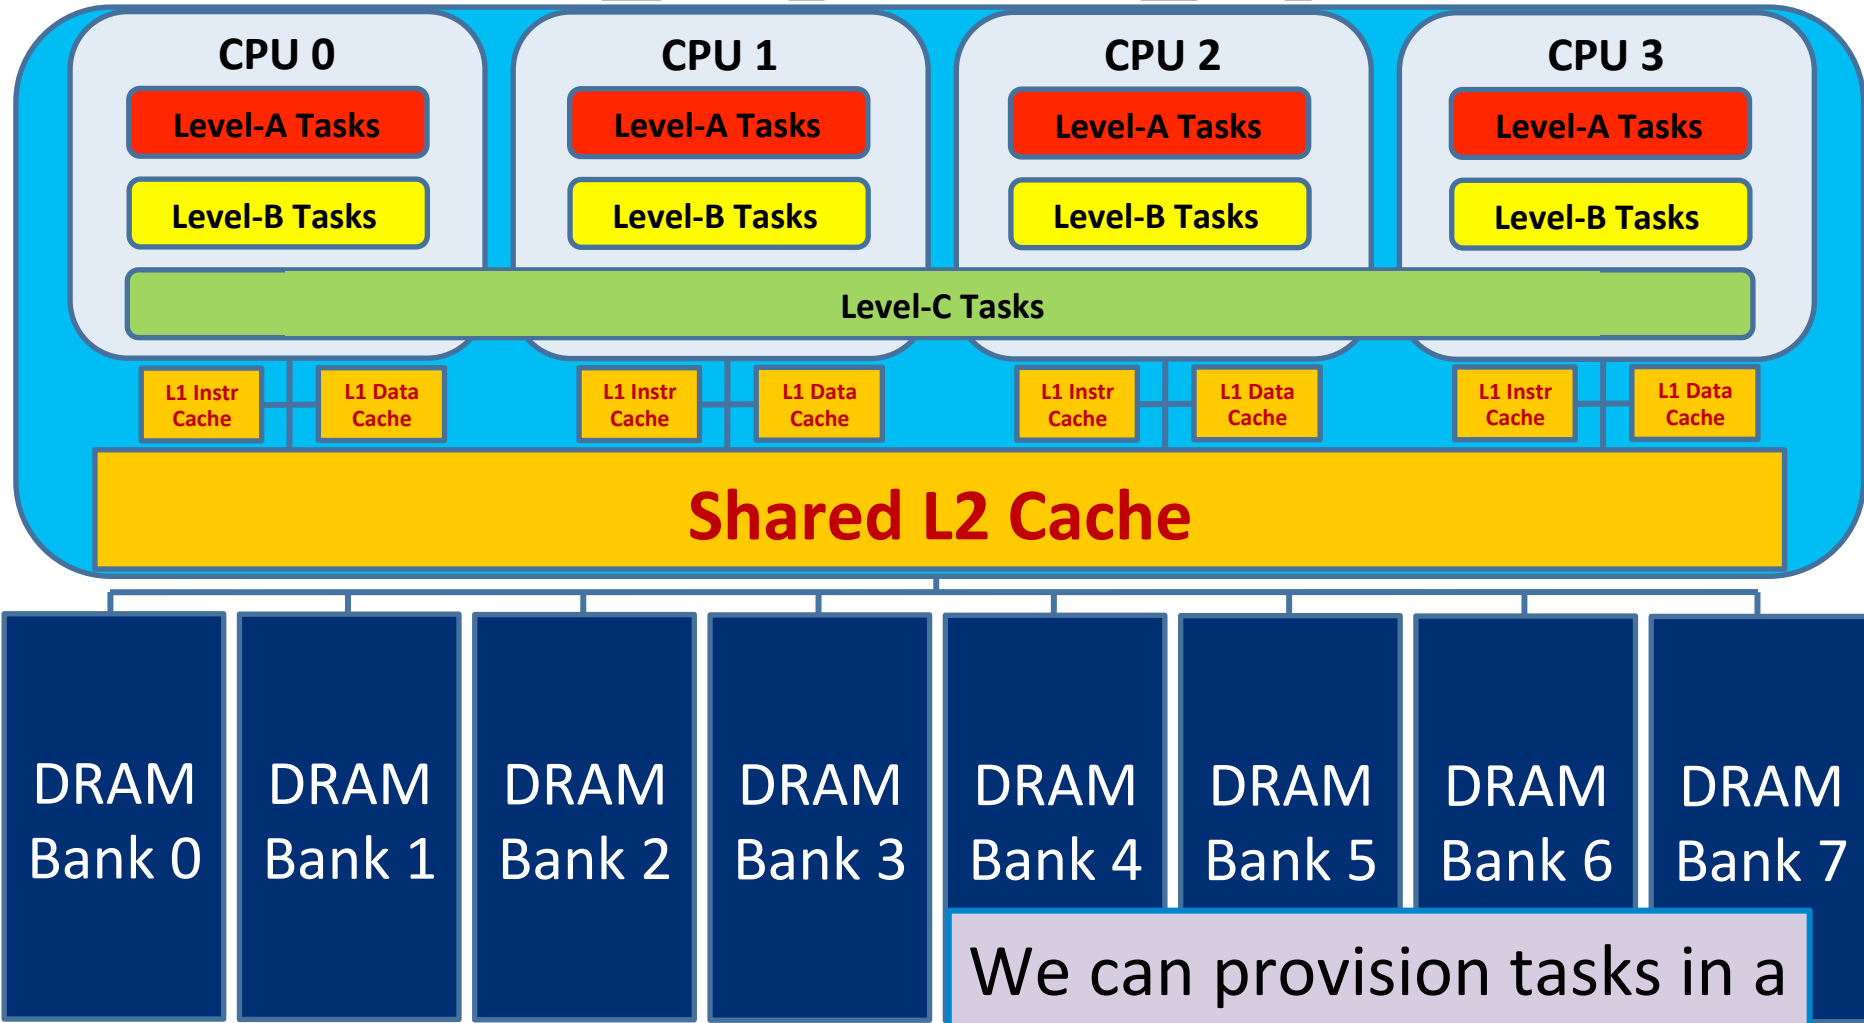


## MC<sup>2</sup> (Mixed-Criticality on Multicore)



CPS 1239135, PI: James H. Anderson,  
UNC Chapel Hill, co-PI: Frank Mueller,  
NC State University.

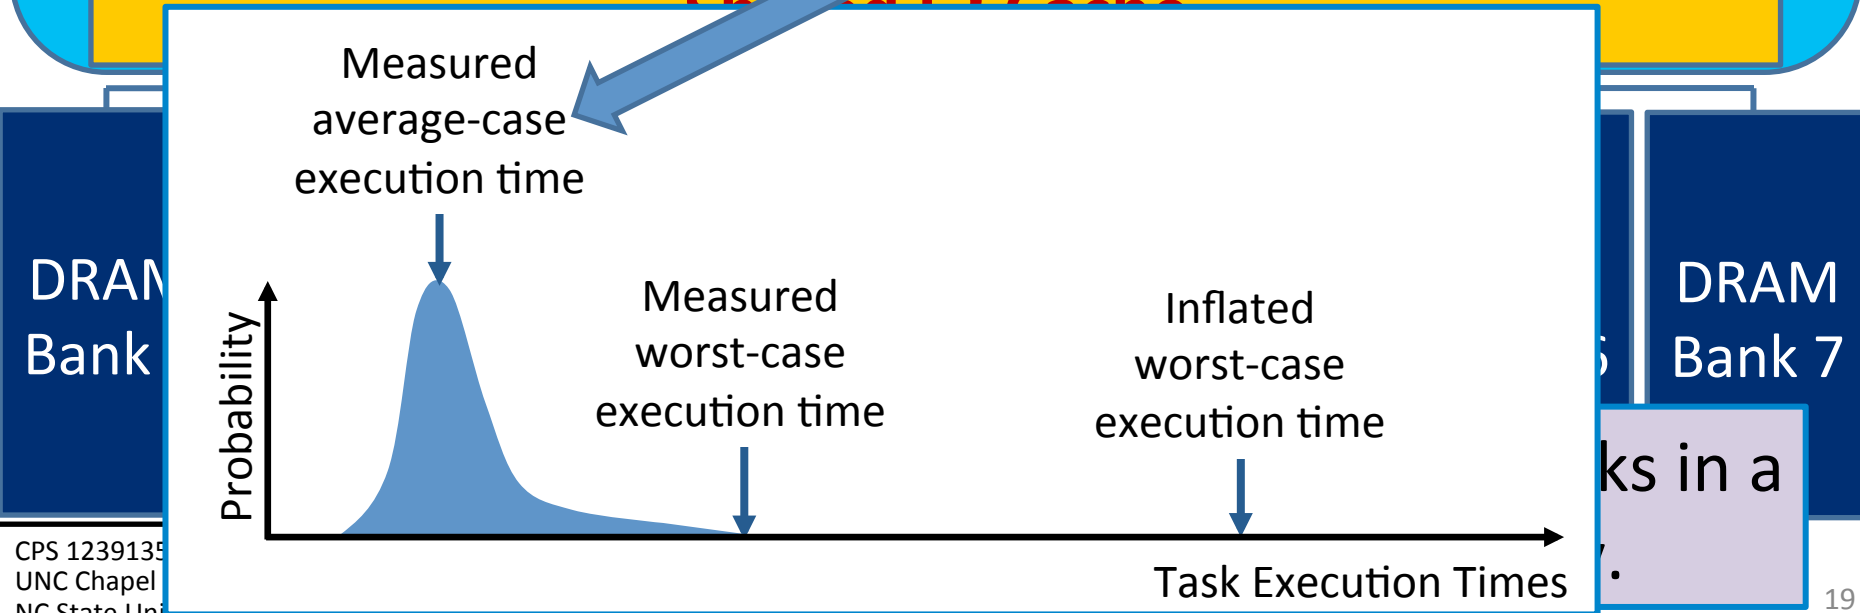
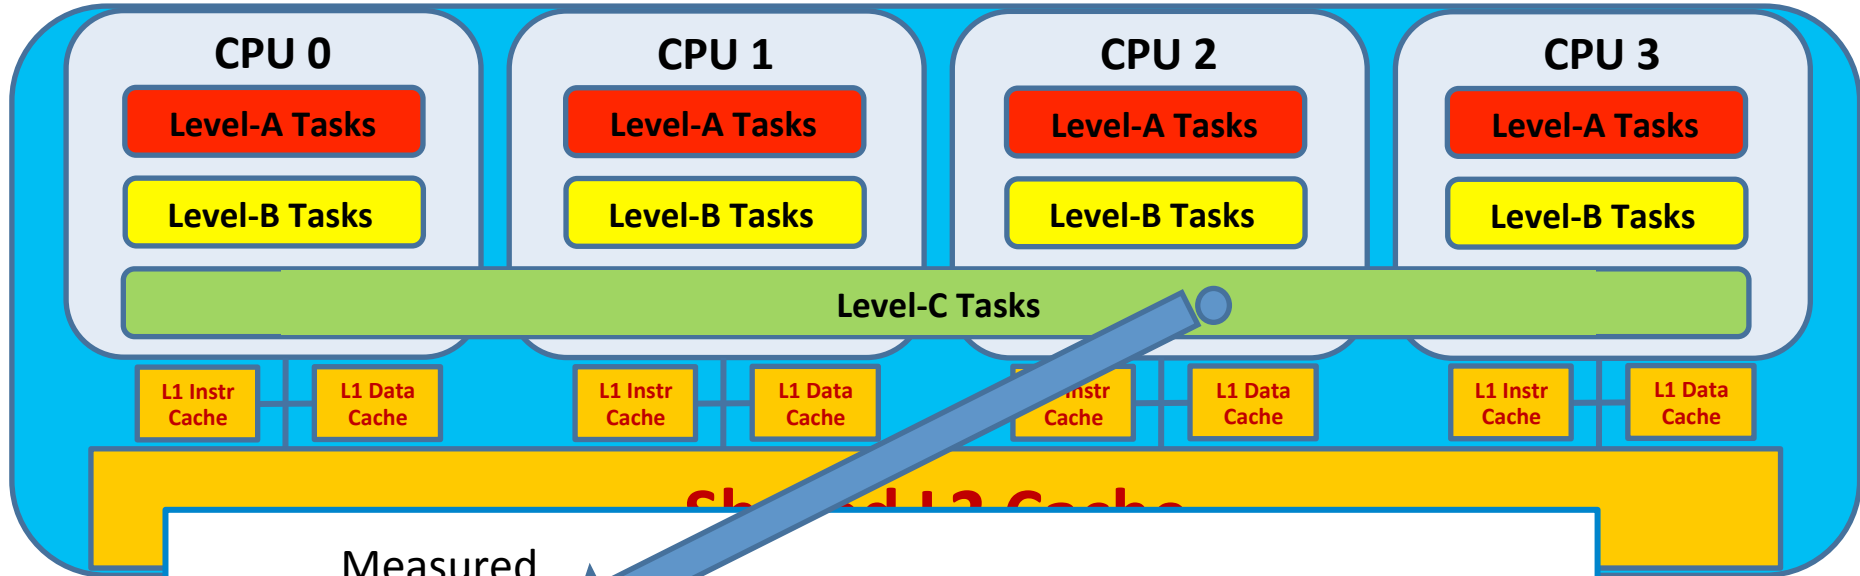
## MC<sup>2</sup> (Mixed-Criticality on Multicore)



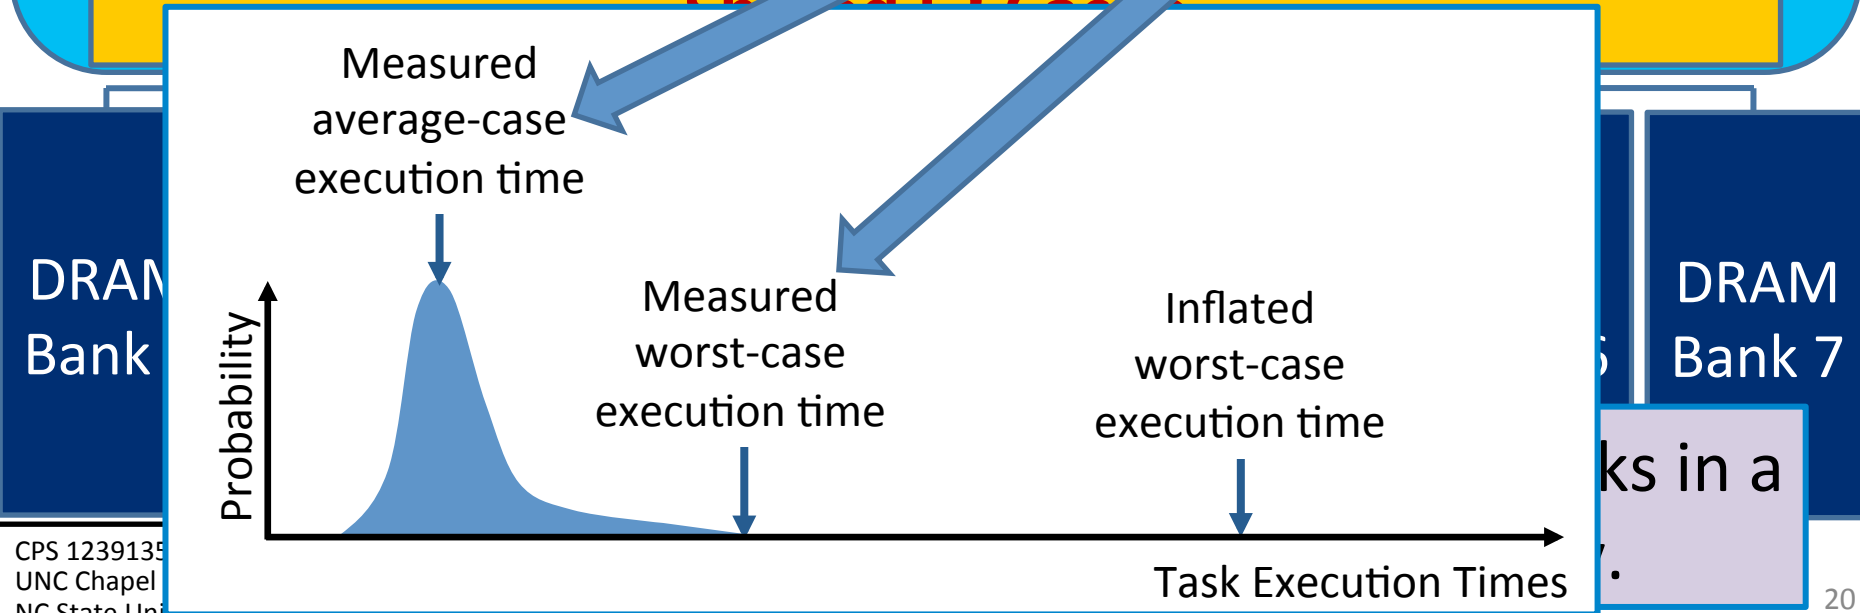
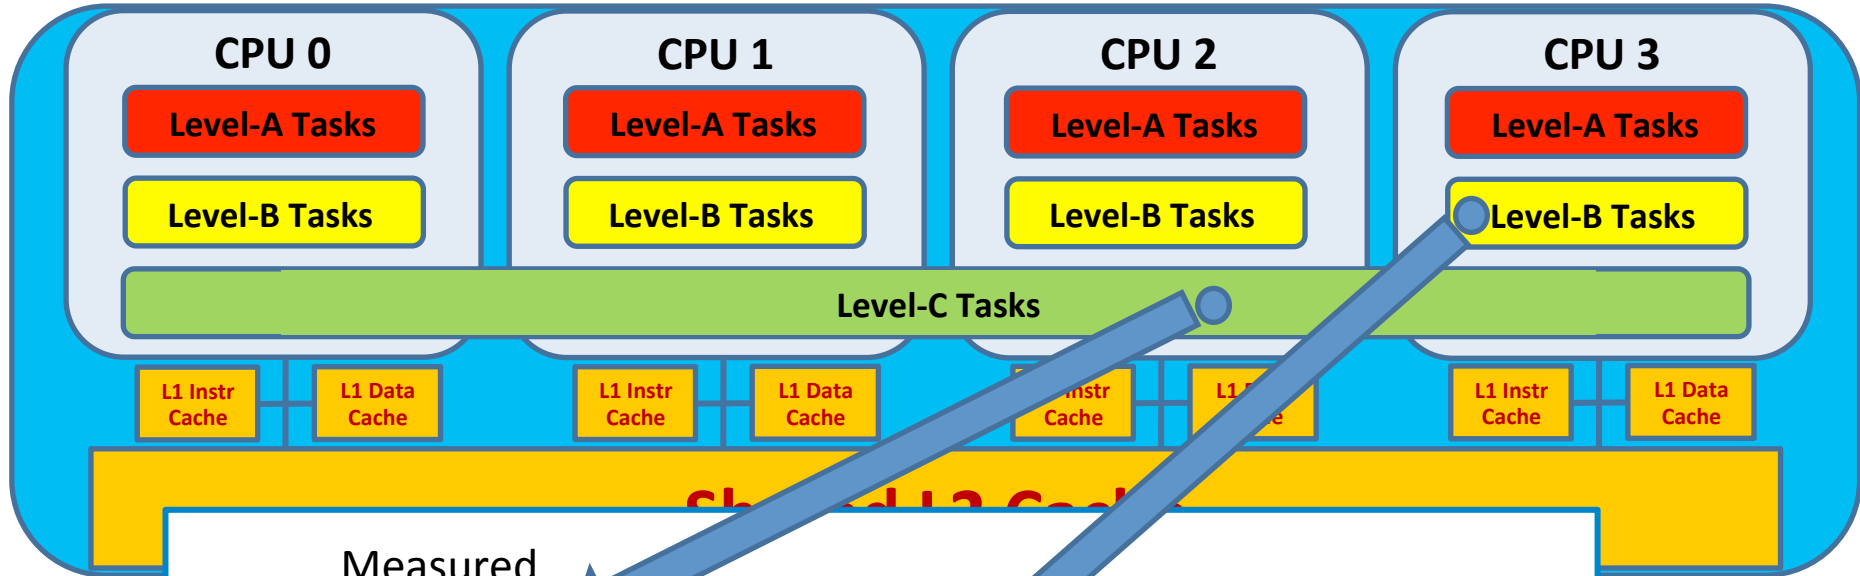
We can provision tasks in a **criticality-aware** way.

CPS 1239135, PI: James H. Anderson, UNC Chapel Hill, co-PI: Frank Mueller, NC State University.

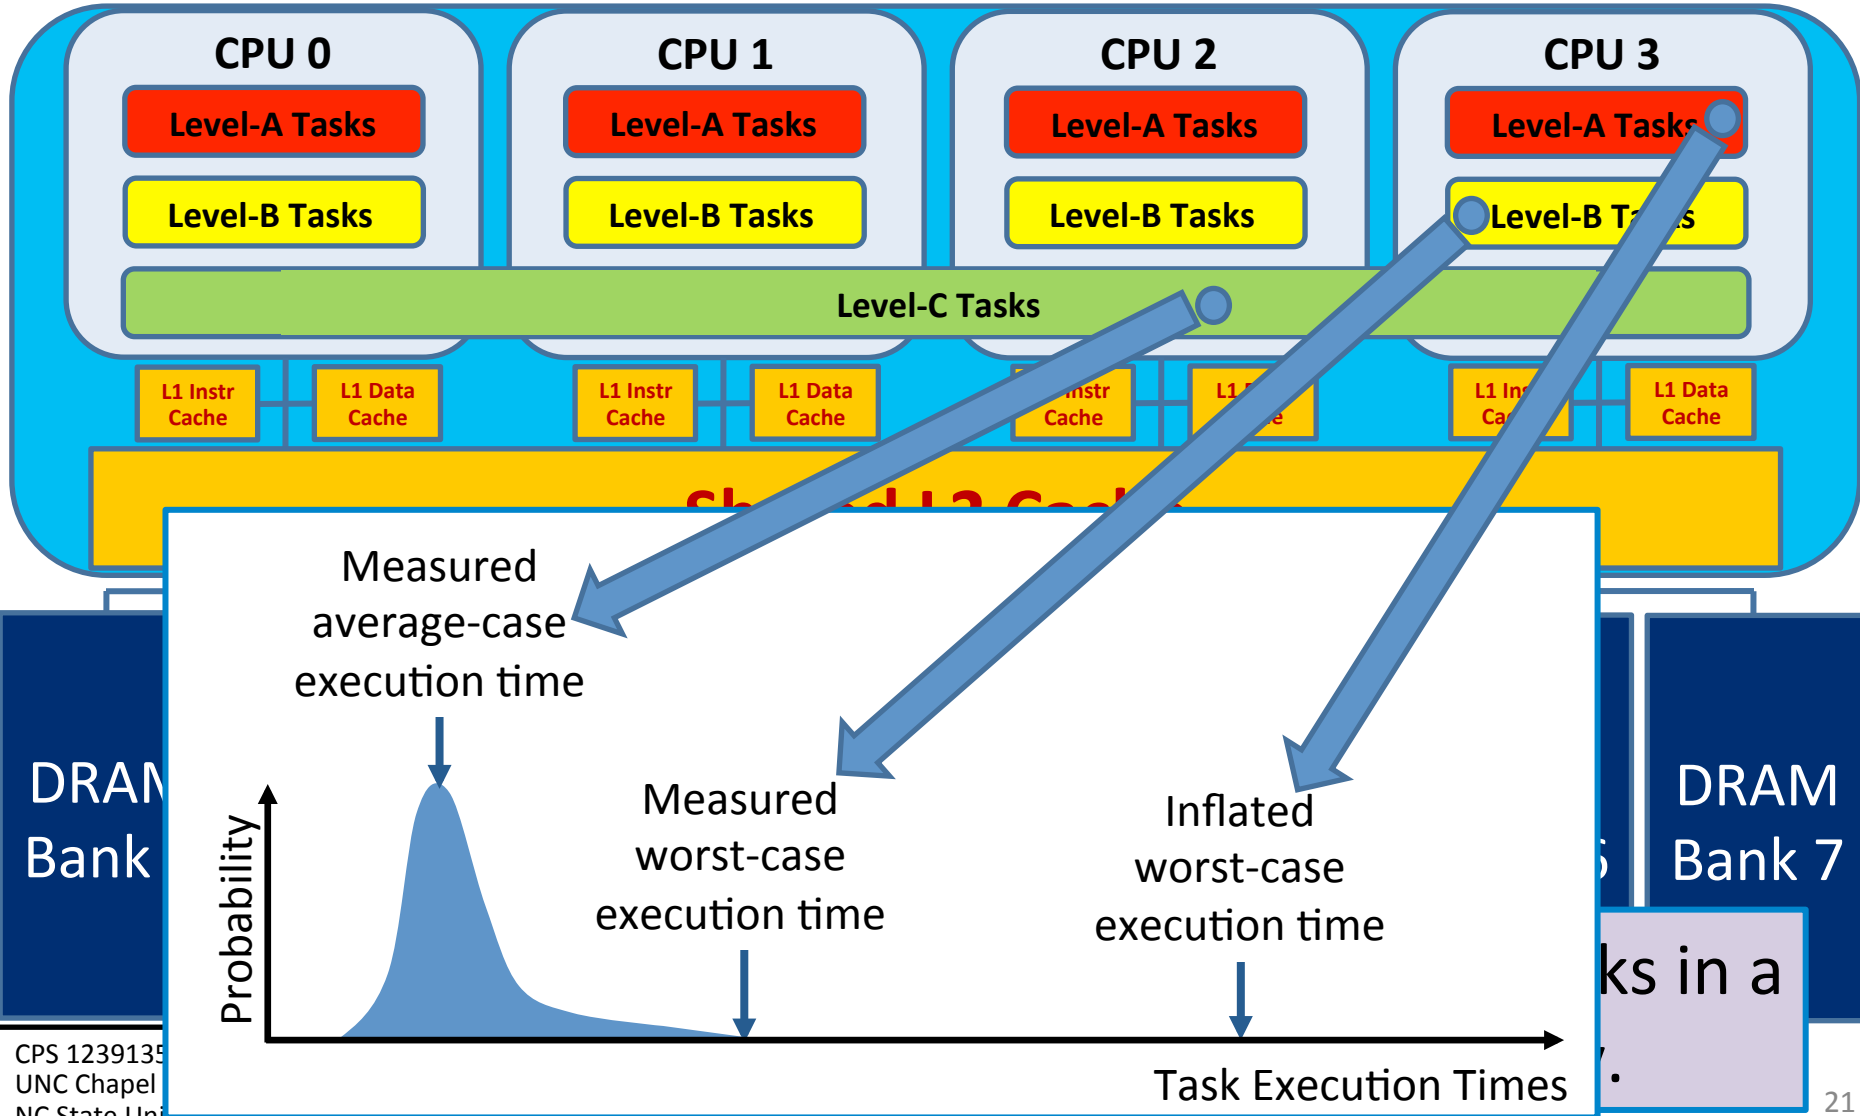
# MC<sup>2</sup> (Mixed-Criticality on Multicore)



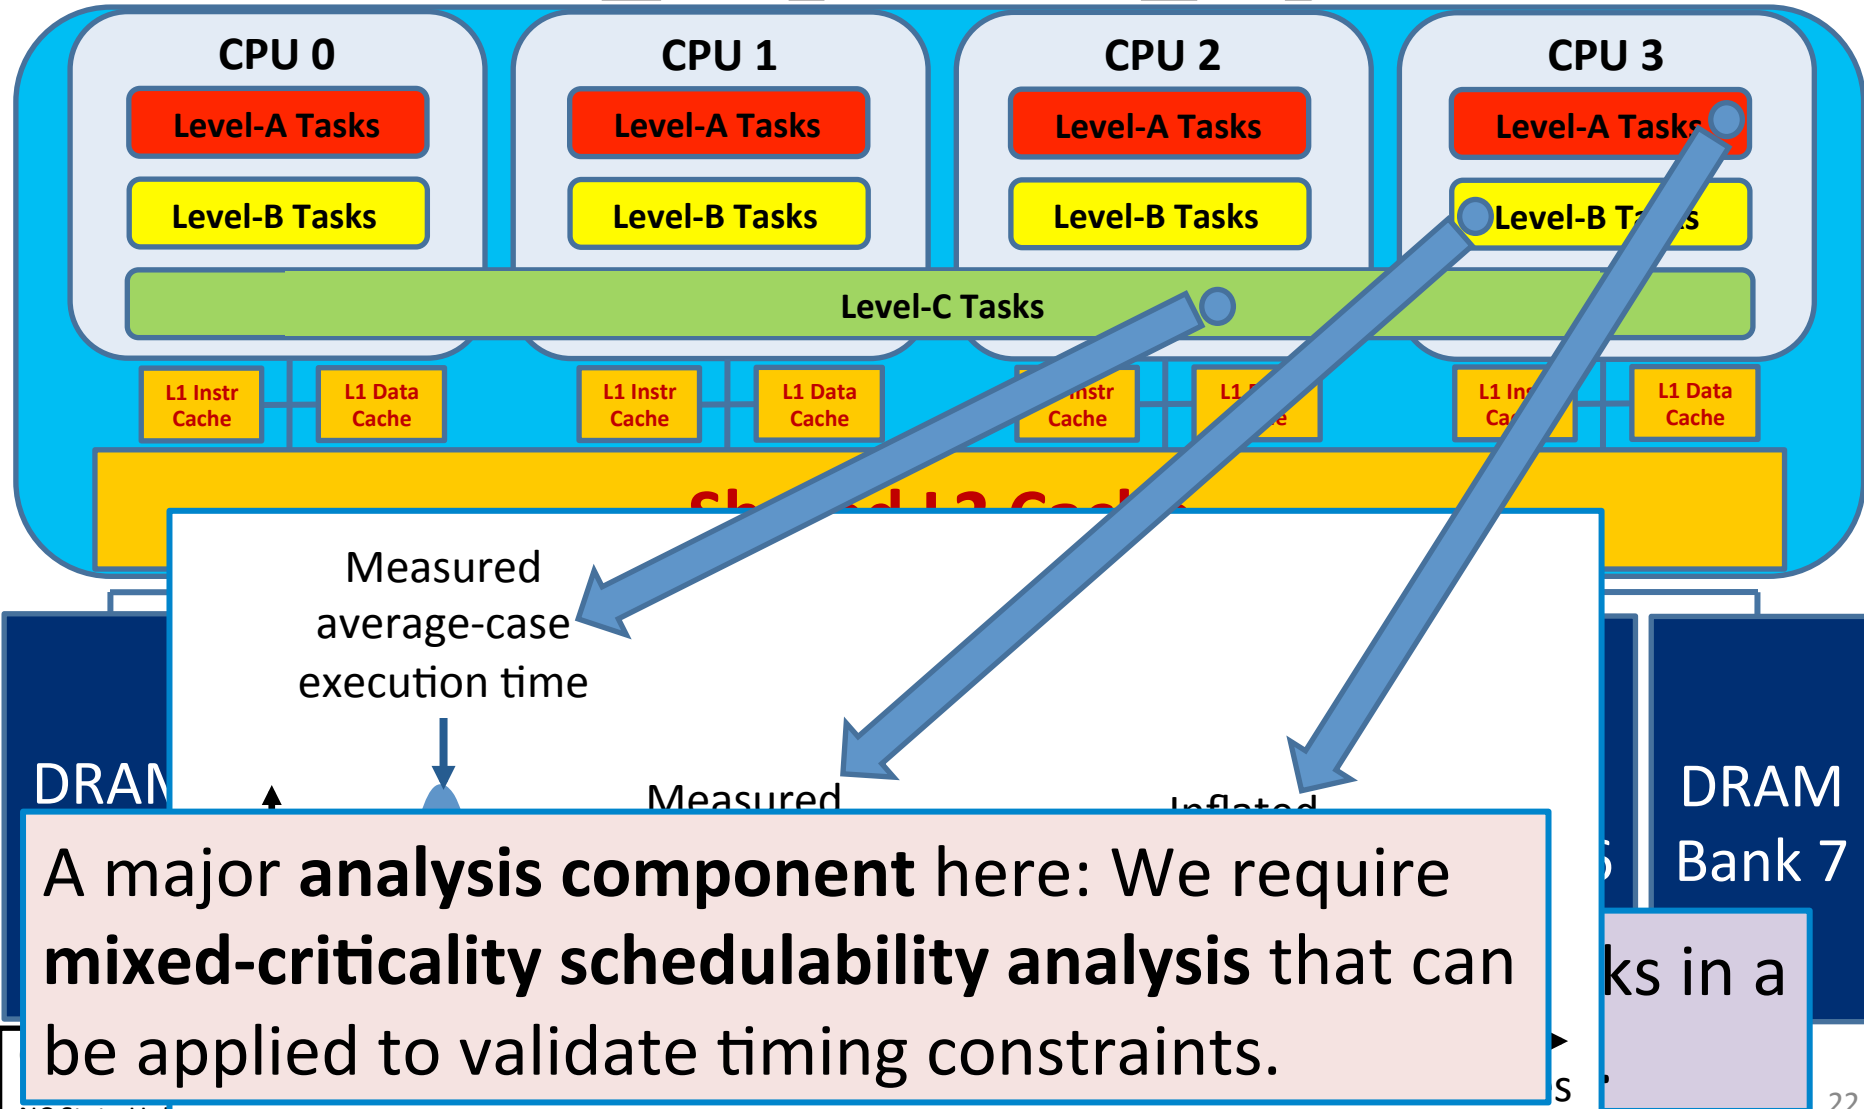
# MC<sup>2</sup> (Mixed-Criticality on Multicore)



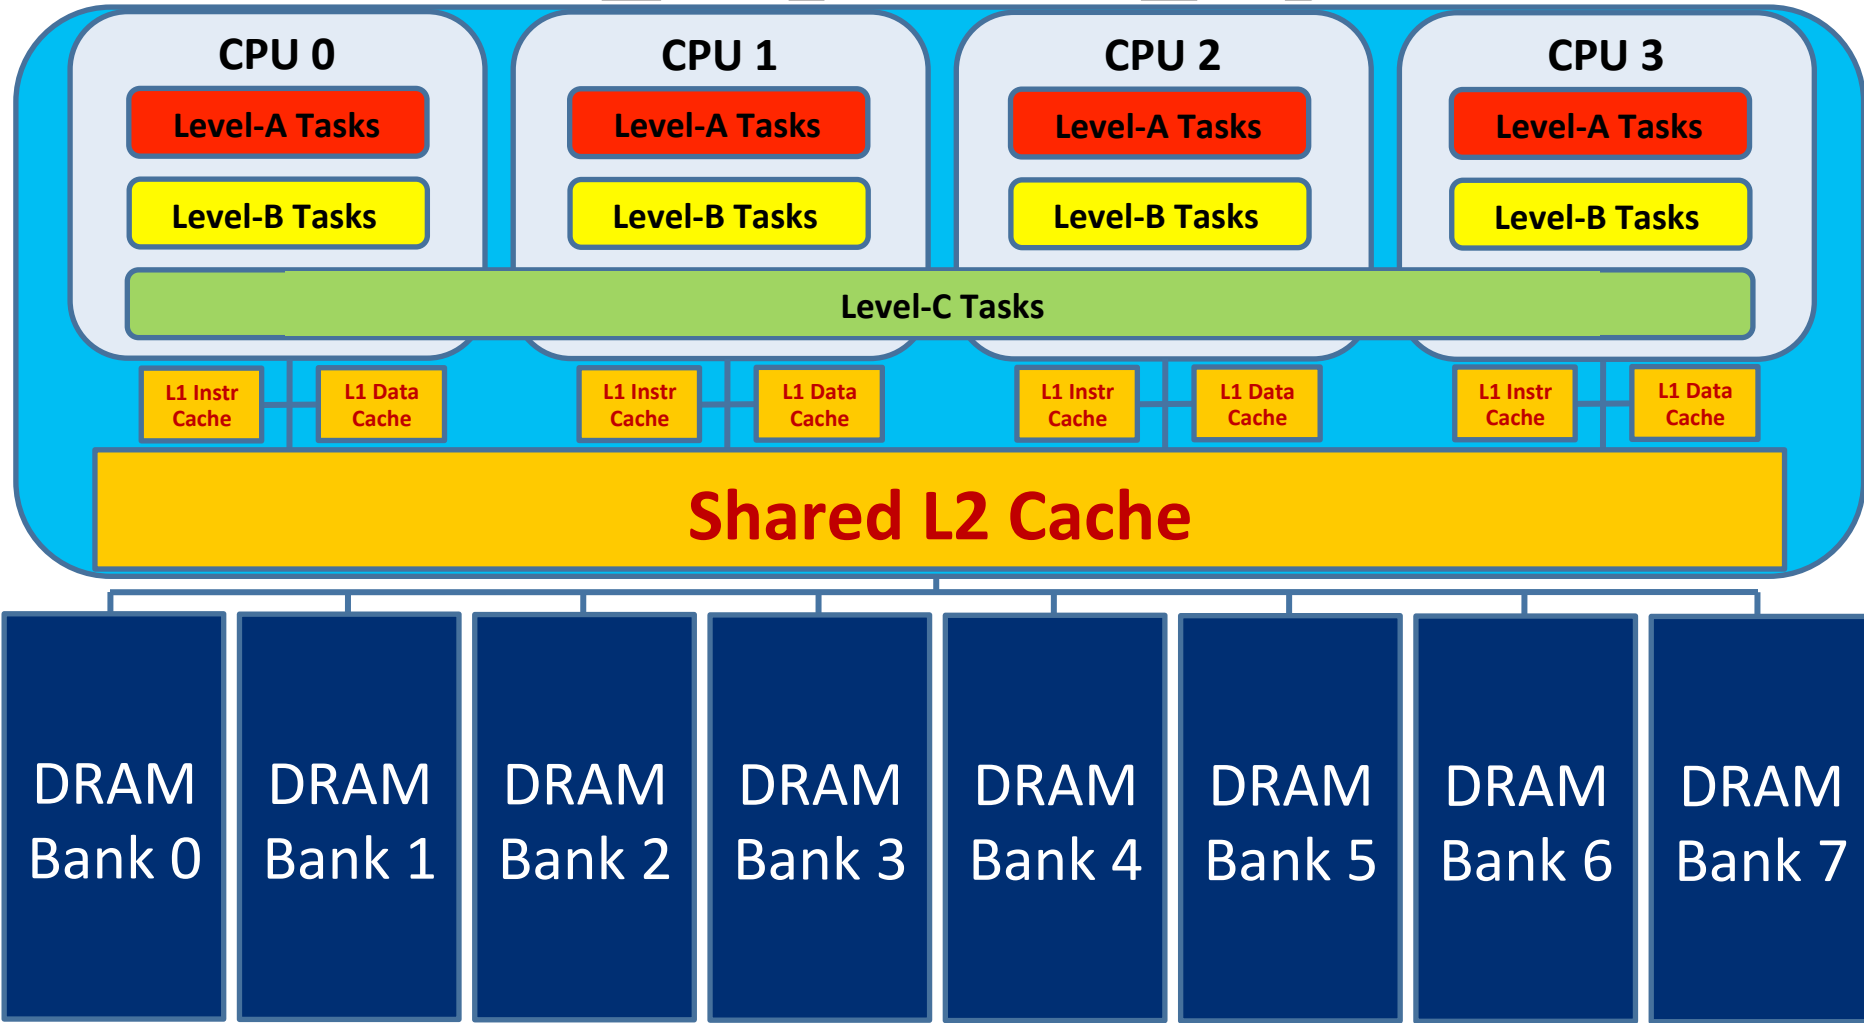
# MC<sup>2</sup> (Mixed-Criticality on Multicore)



## MC<sup>2</sup> (Mixed-Criticality on Multicore)

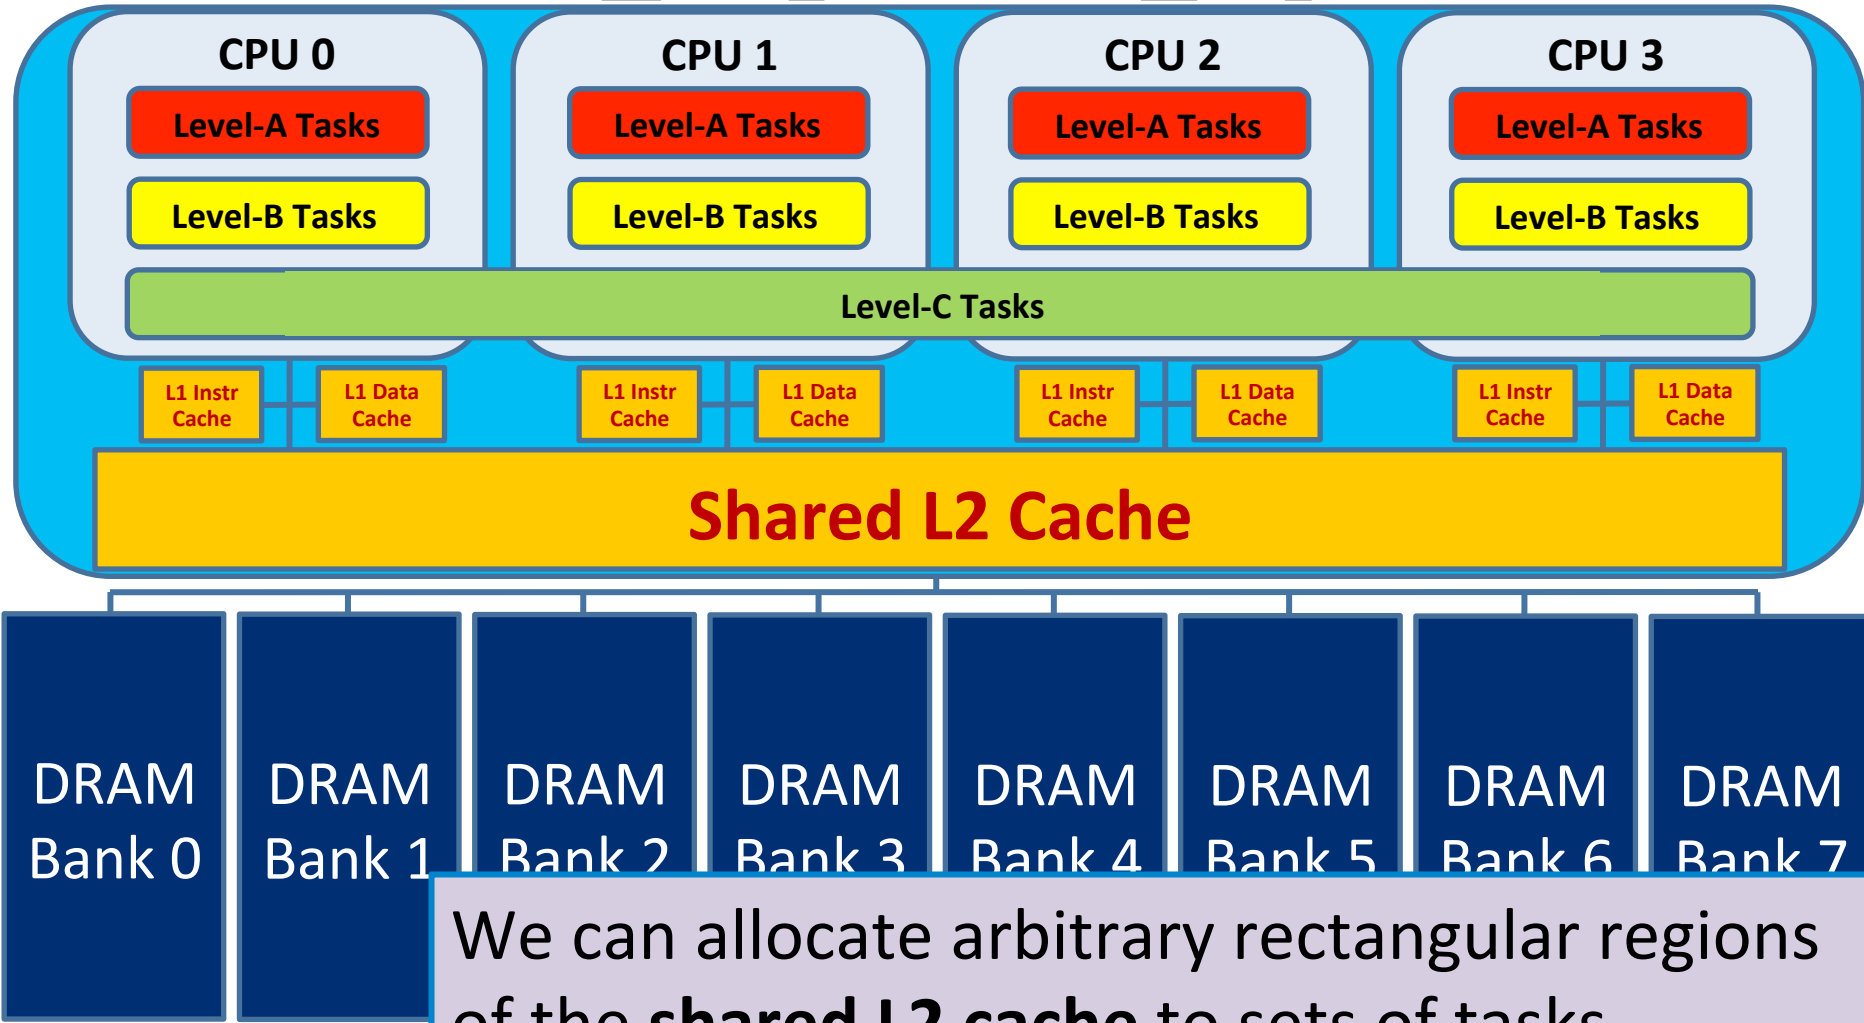


## MC<sup>2</sup> (Mixed-Criticality on Multicore)



CPS 1239135, PI: James H. Anderson,  
UNC Chapel Hill, co-PI: Frank Mueller,  
NC State University.

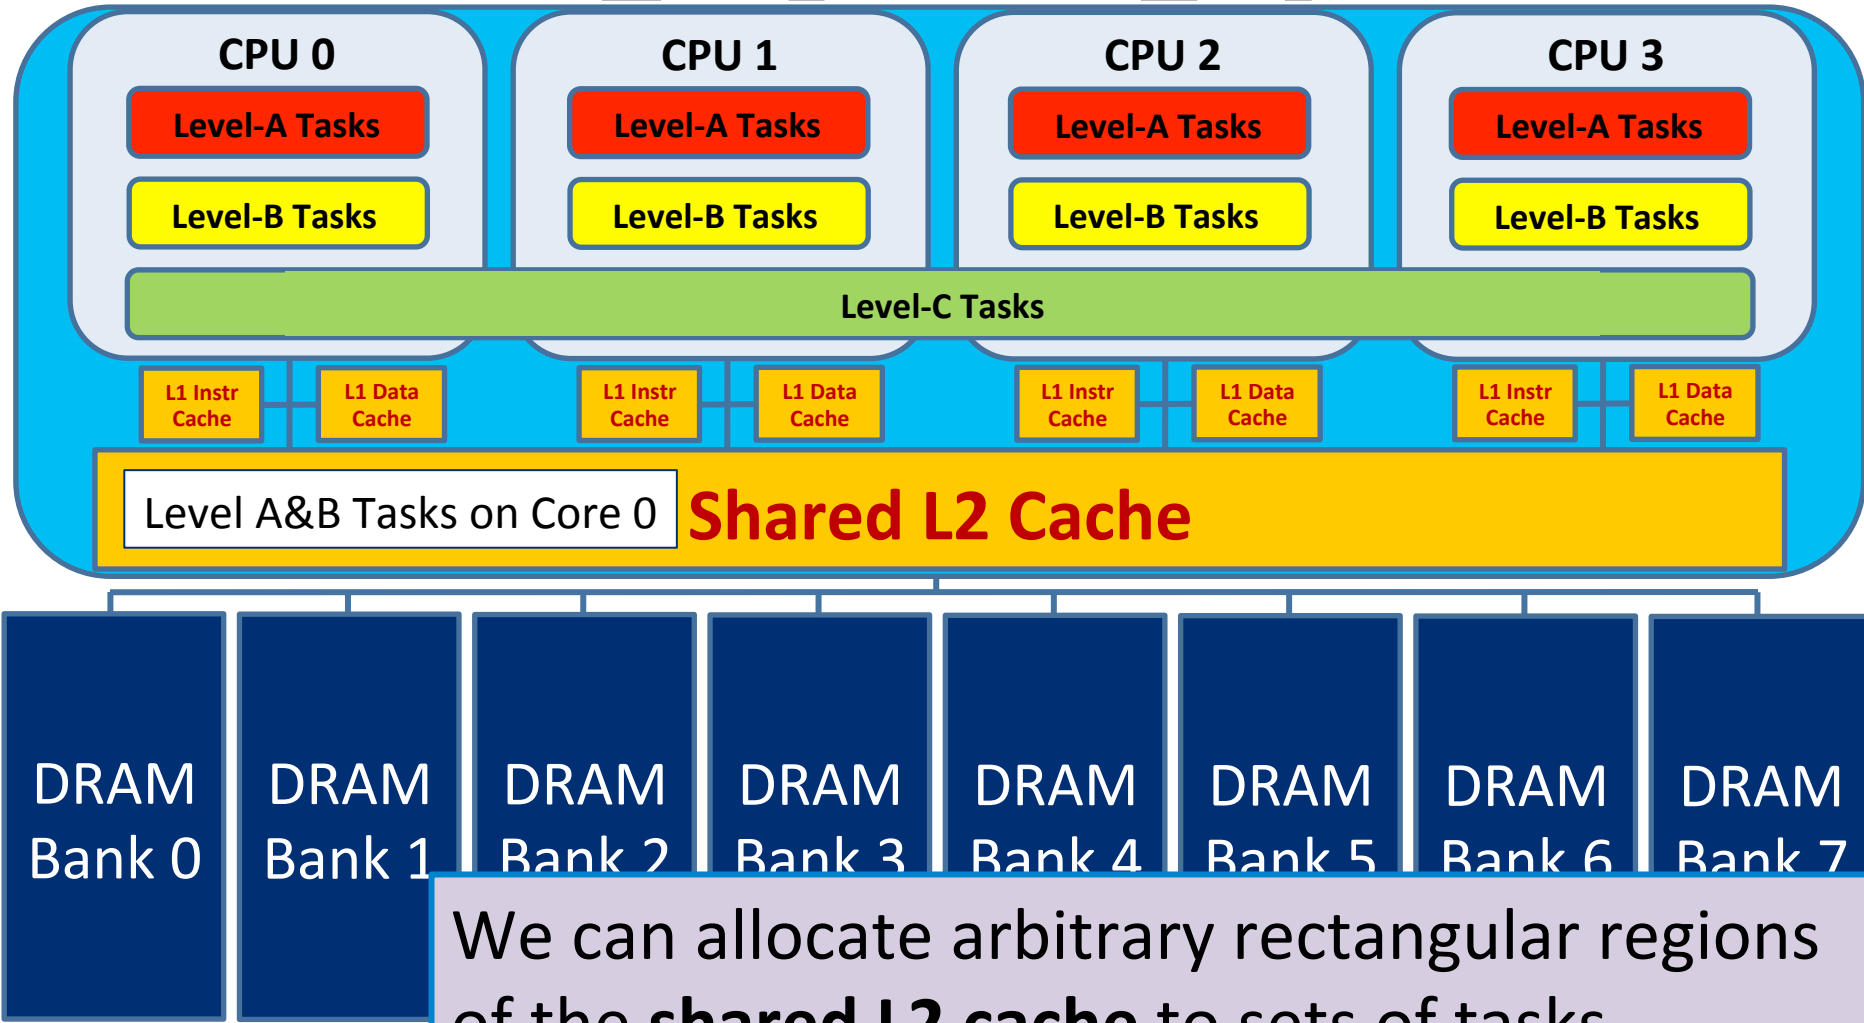
## MC<sup>2</sup> (Mixed-Criticality on Multicore)



We can allocate arbitrary rectangular regions of the **shared L2 cache** to sets of tasks.

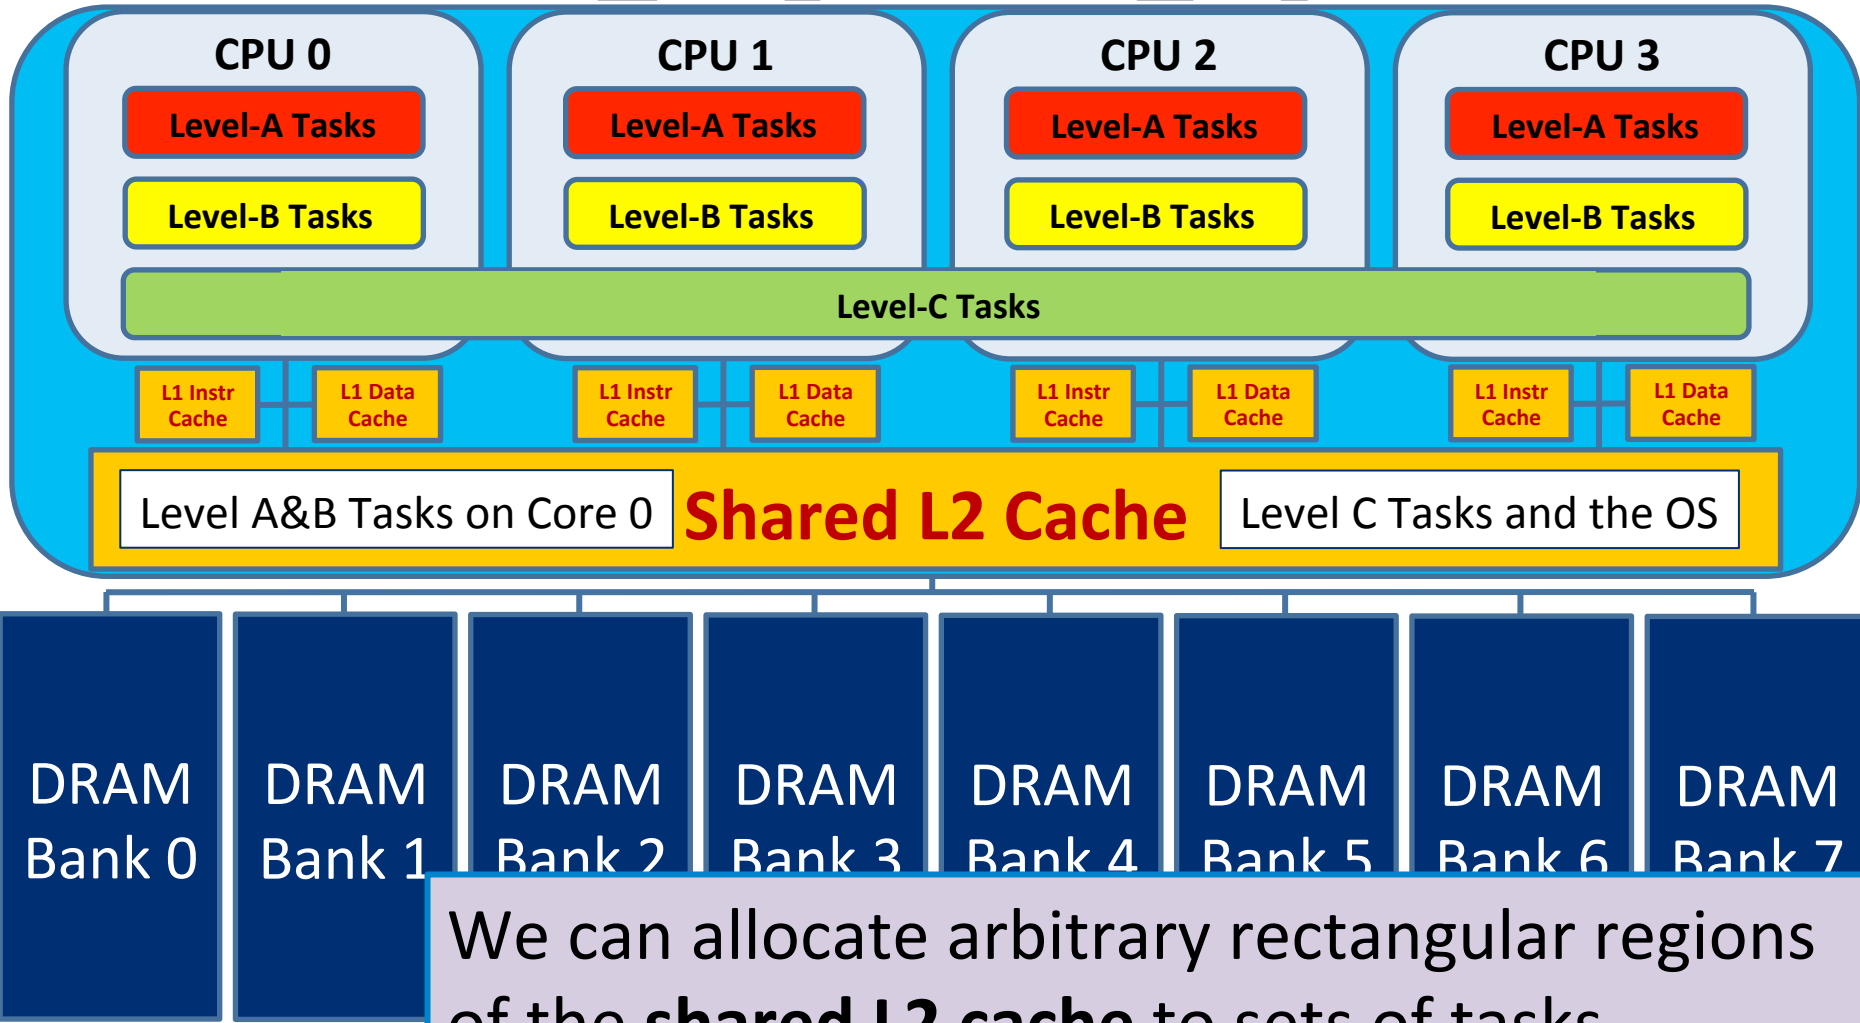


## MC<sup>2</sup> (Mixed-Criticality on Multicore)



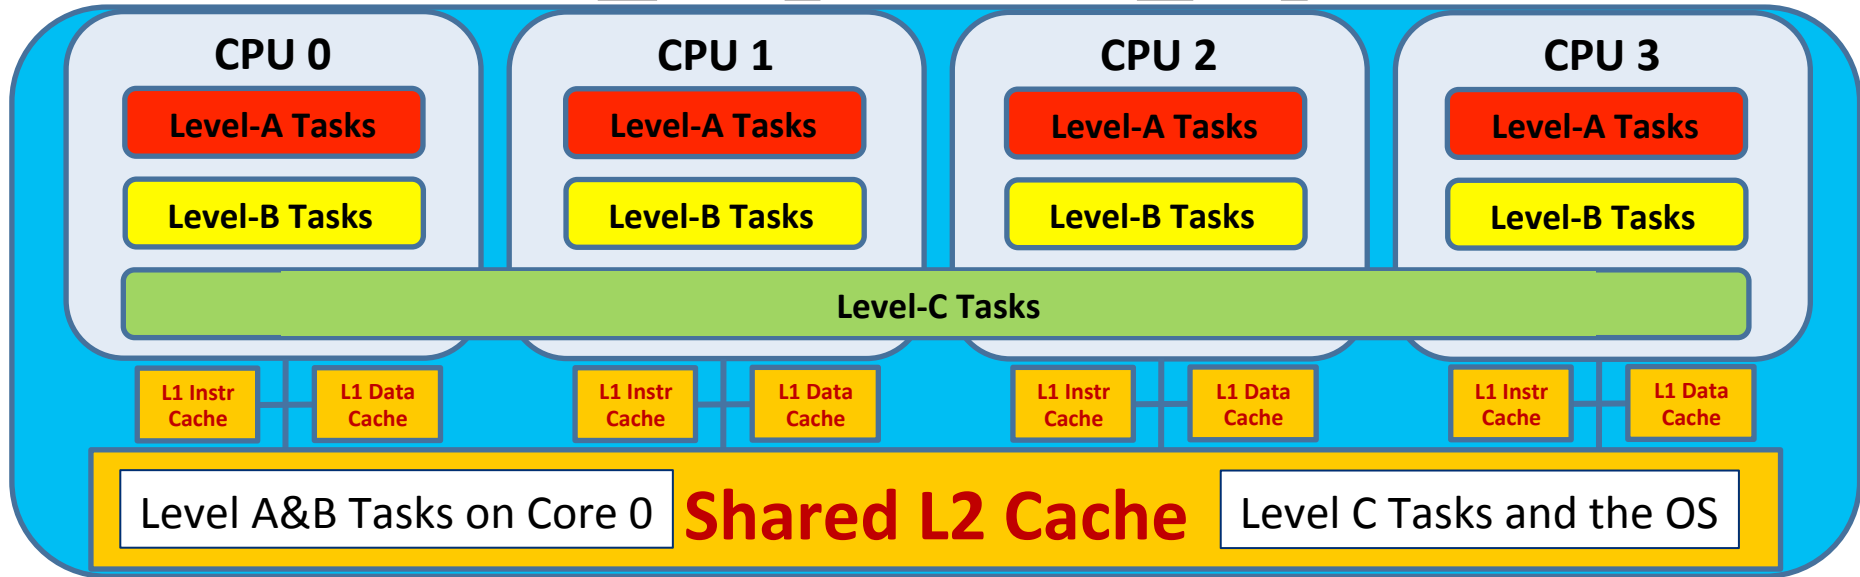
We can allocate arbitrary rectangular regions of the **shared L2 cache** to sets of tasks.

# MC<sup>2</sup> (Mixed-Criticality on Multicore)



We can allocate arbitrary rectangular regions of the **shared L2 cache** to sets of tasks.

## MC<sup>2</sup> (Mixed-Criticality on Multicore)



We have an **optimization framework** that can automatically produce such allocations.

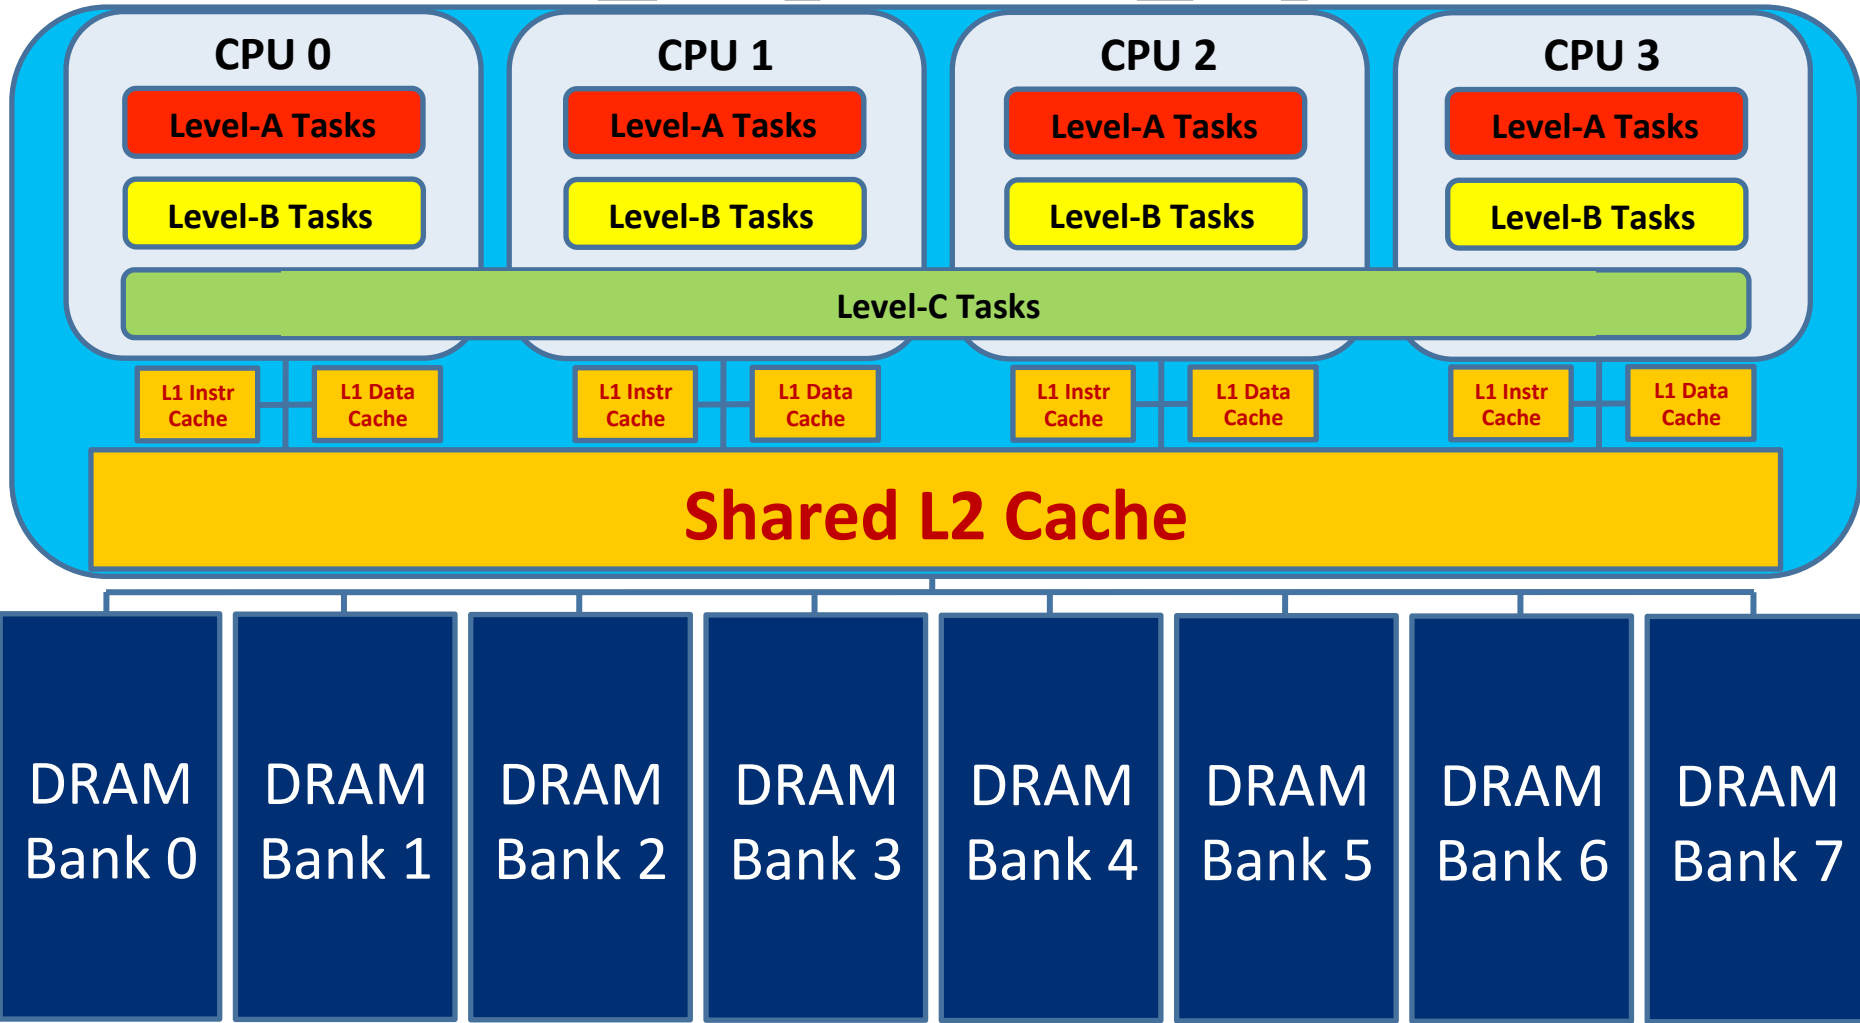
Bank 0 | Bank 1 | Bank 2 | Bank 3 | Bank 4 | Bank 5 | Bank 6 | Bank 7

DRAM

DRAM

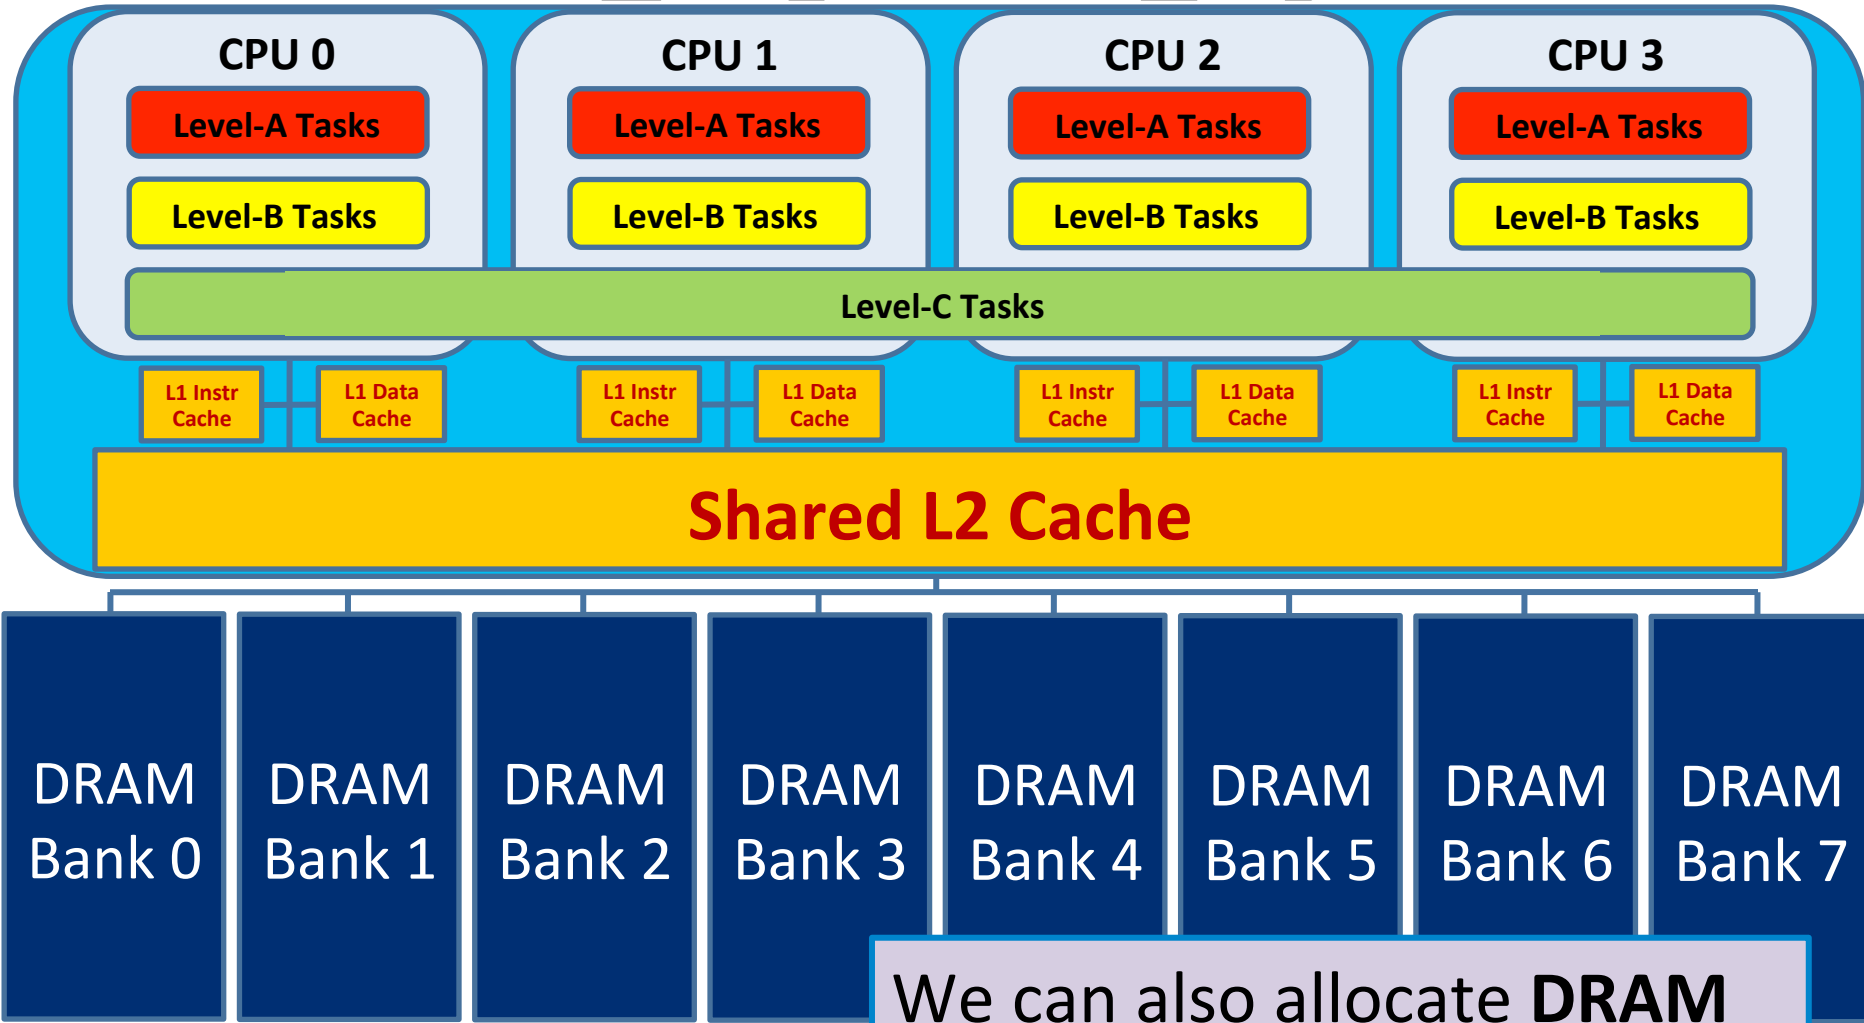
We can allocate arbitrary rectangular regions of the **shared L2 cache** to sets of tasks.

# MC<sup>2</sup> (Mixed-Criticality on Multicore)



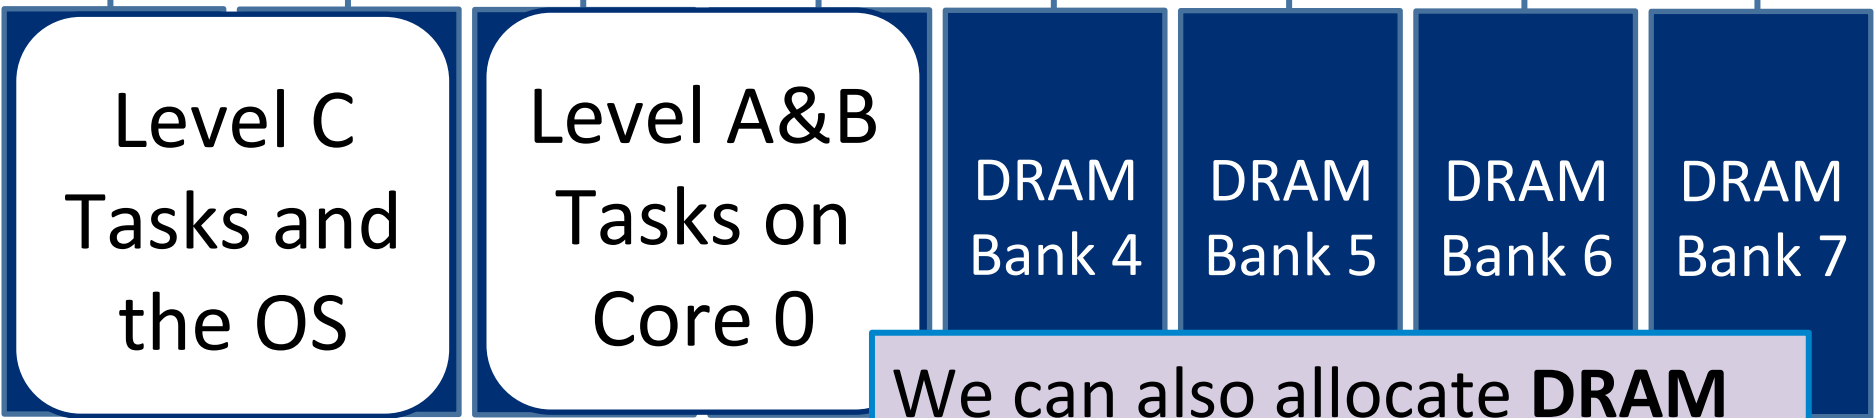
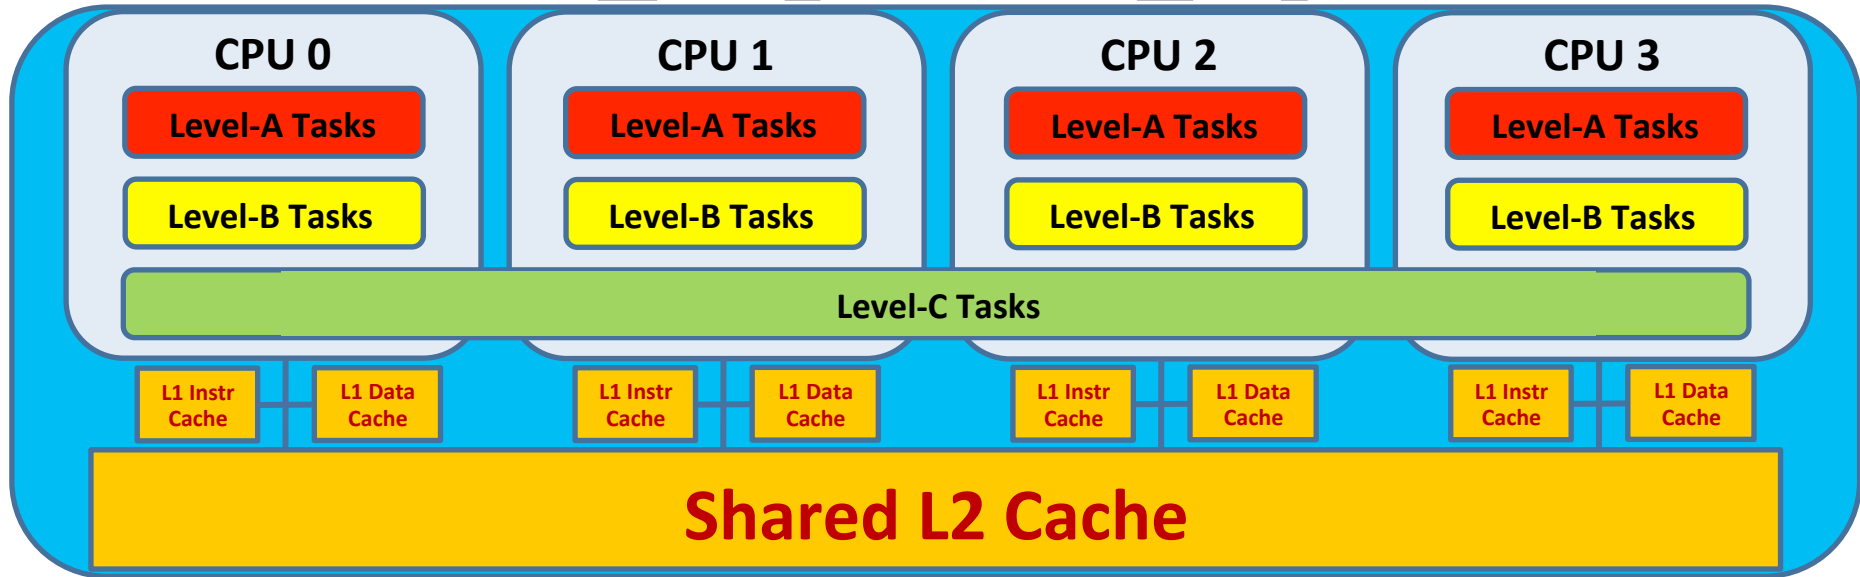
CPS 1239135, PI: James H. Anderson,  
UNC Chapel Hill, co-PI: Frank Mueller,  
NC State University.

## MC<sup>2</sup> (Mixed-Criticality on Multicore)



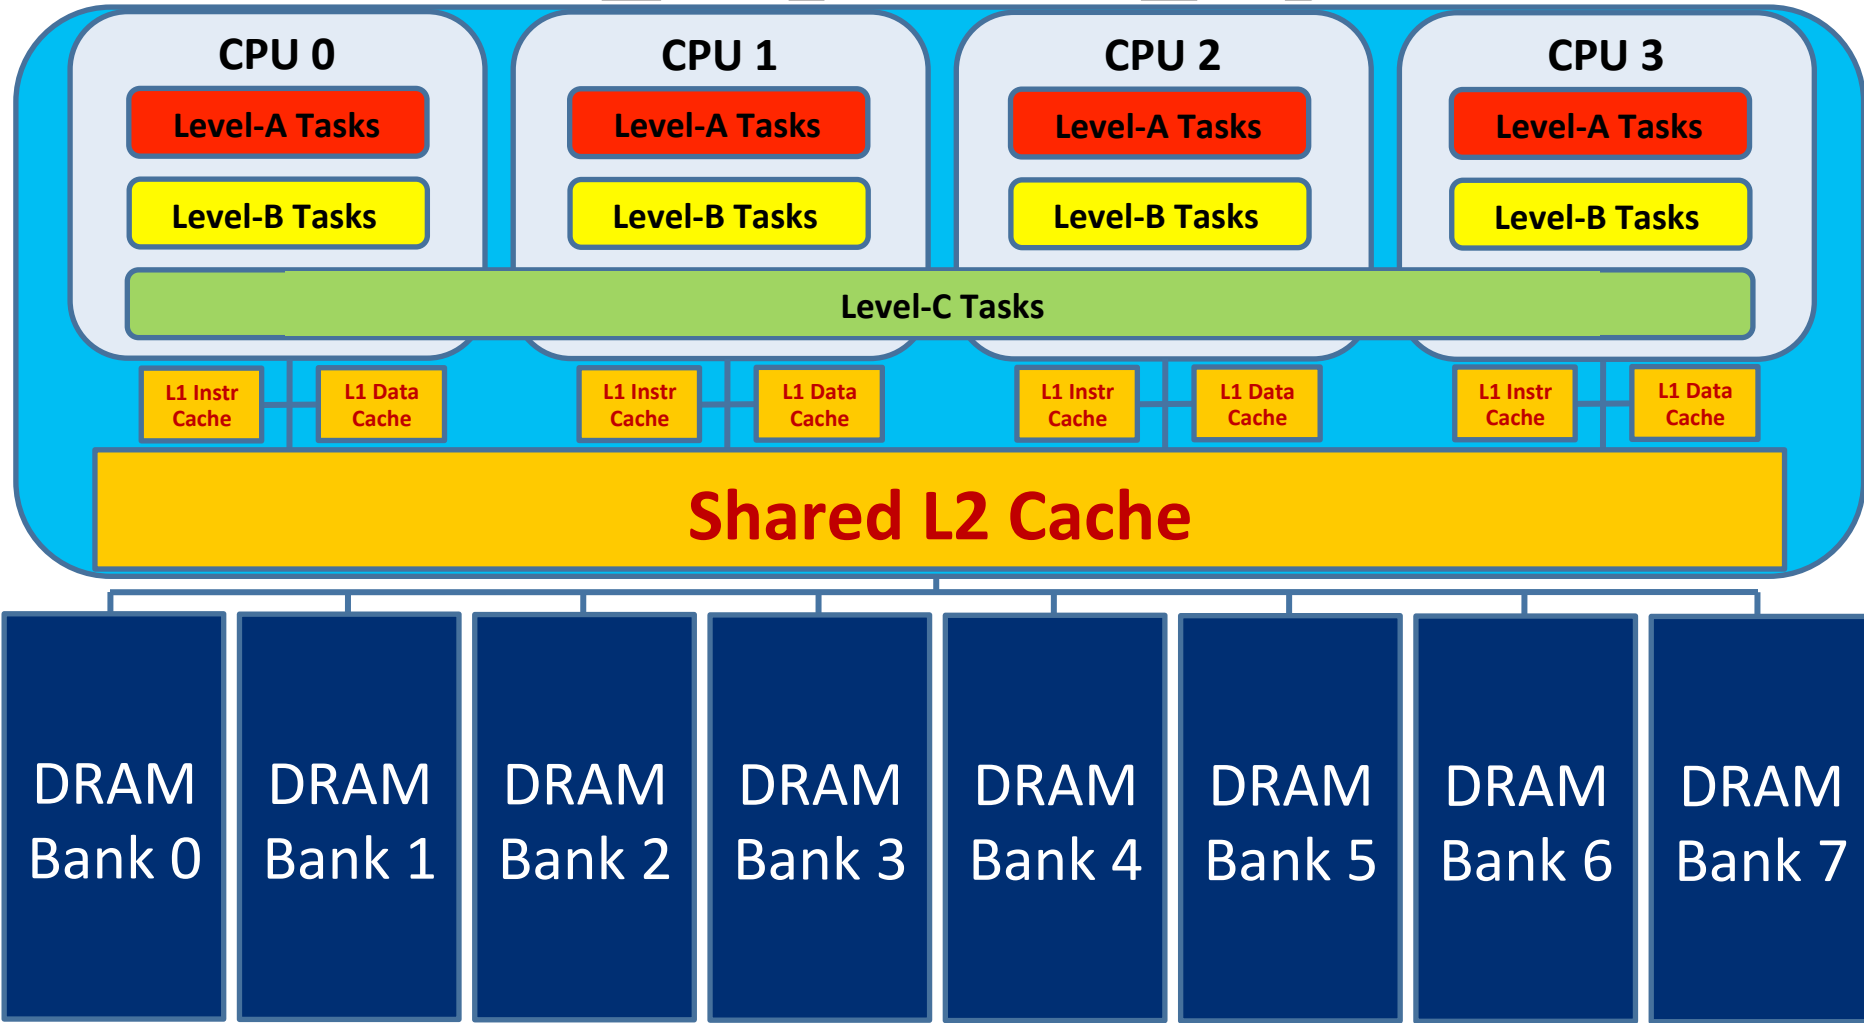
We can also allocate **DRAM banks** to certain sets of tasks.

## MC<sup>2</sup> (Mixed-Criticality on Multicore)



We can also allocate **DRAM banks** to certain sets of tasks.

## MC<sup>2</sup> (Mixed-Criticality on Multicore)



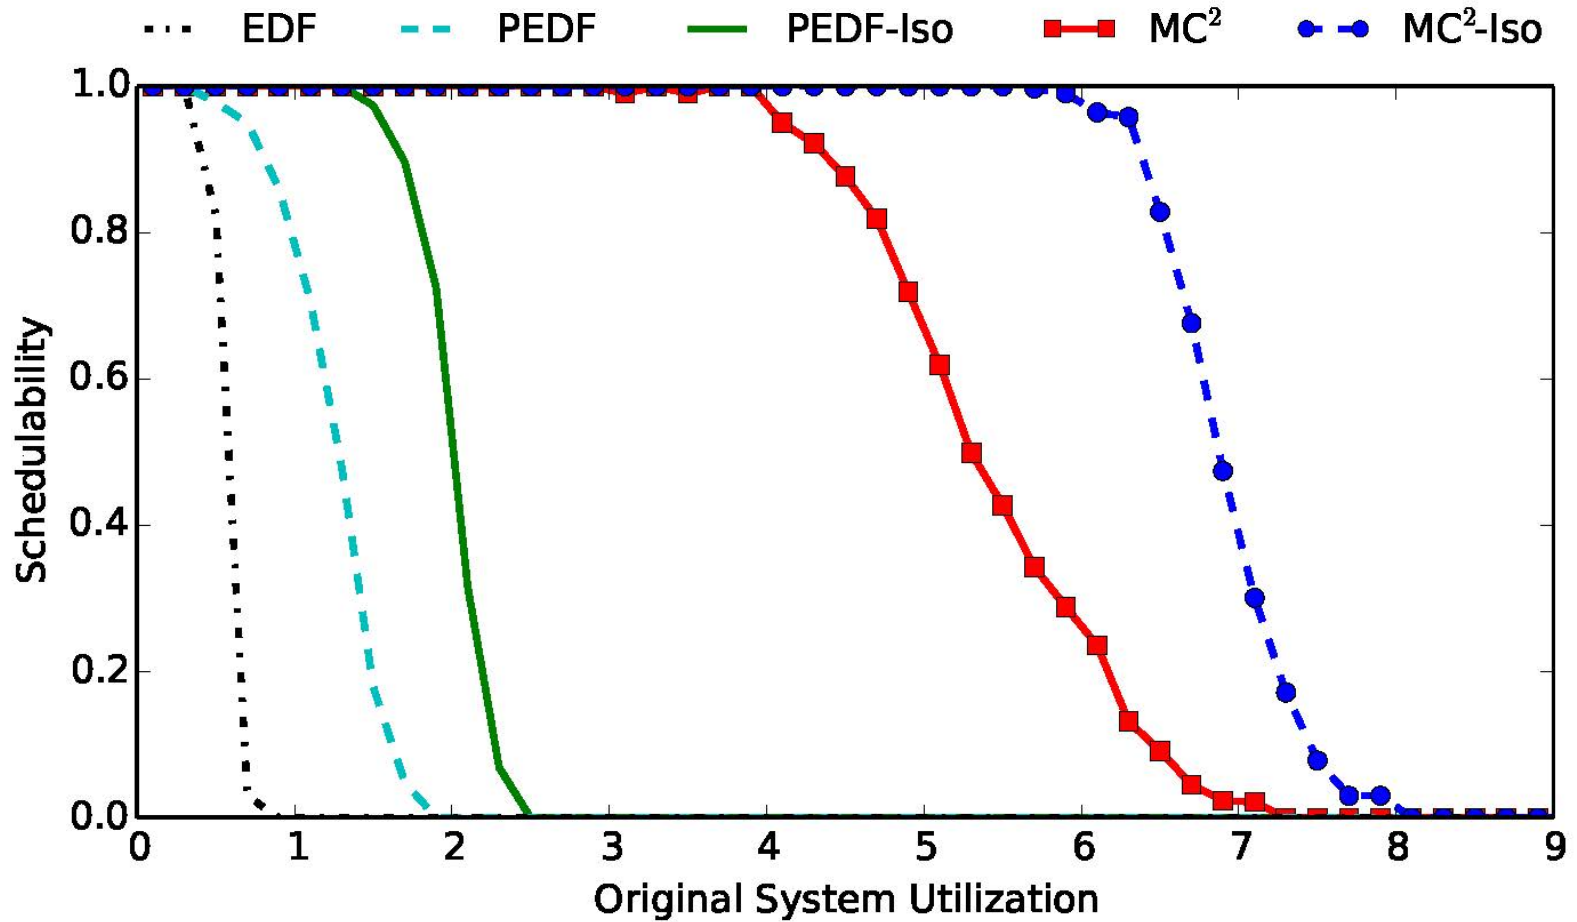
CPS 1239135, PI: James H. Anderson,  
UNC Chapel Hill, co-PI: Frank Mueller,  
NC State University.

## Outline

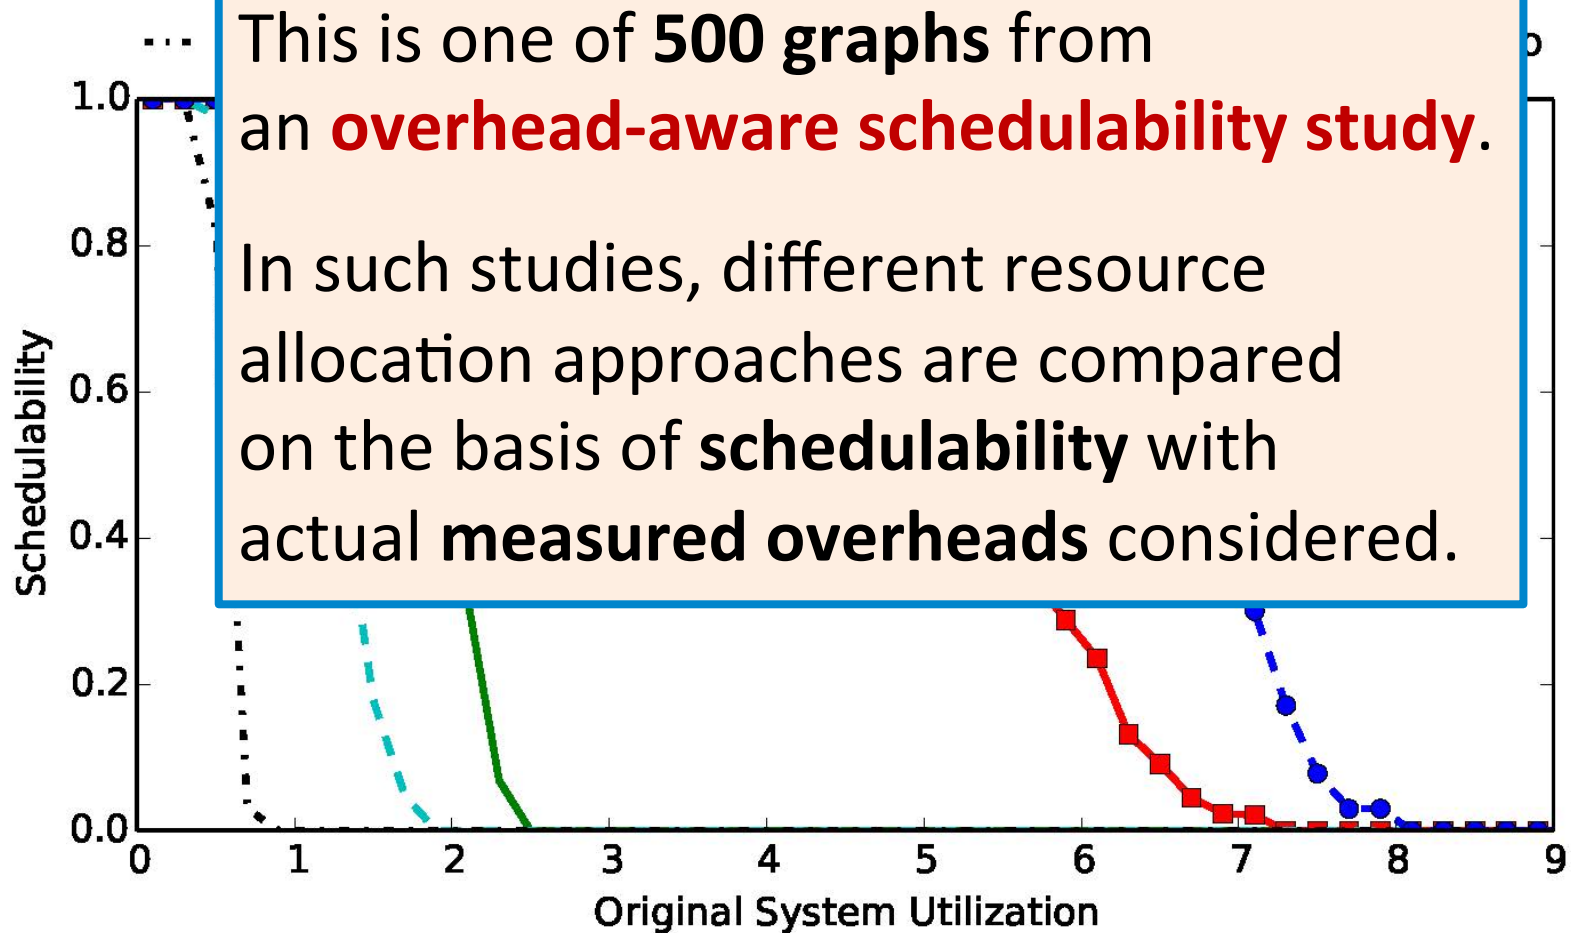
- Real-time 101.
- Problems caused by multicore.
  - “The one-out-of-m problem.”
  - Why this is an important problem.
- Solution strategy.
- Evaluation.
- Project summary.



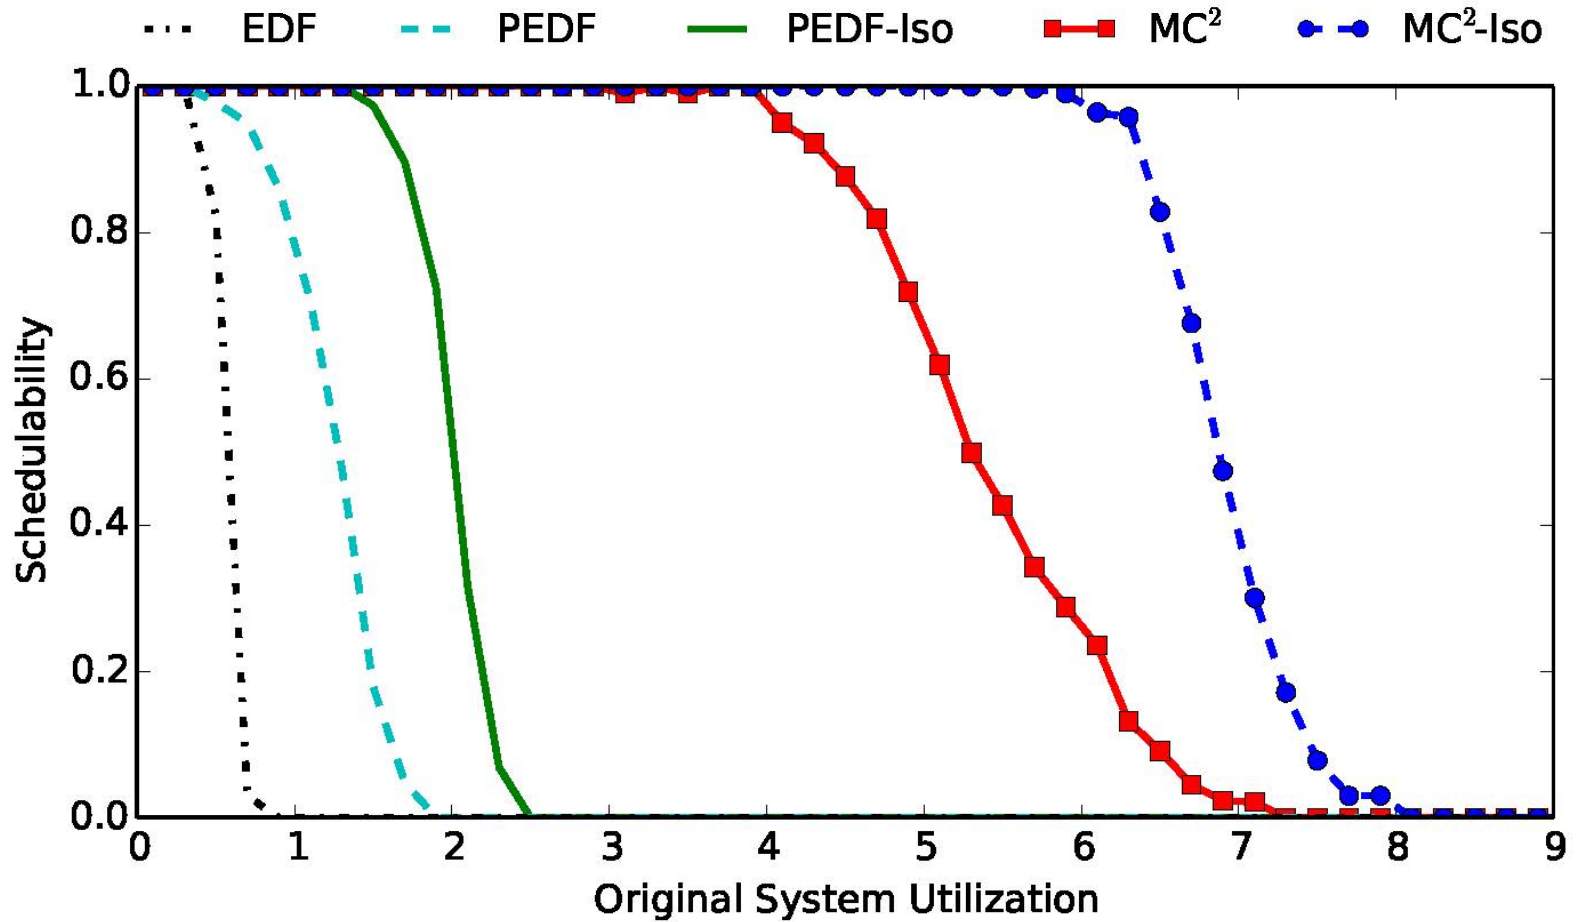
## A Little Experimental Evidence in Favor of our Approach



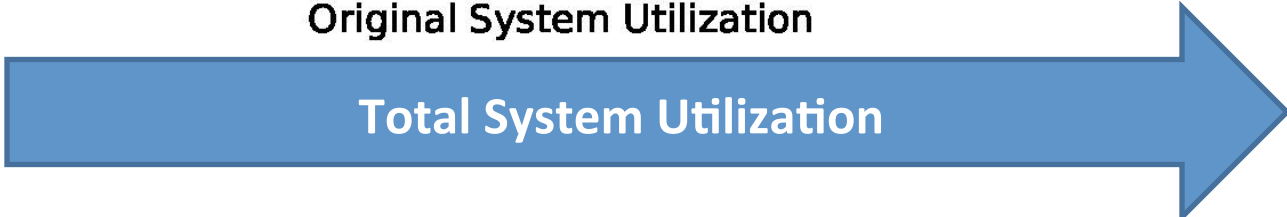
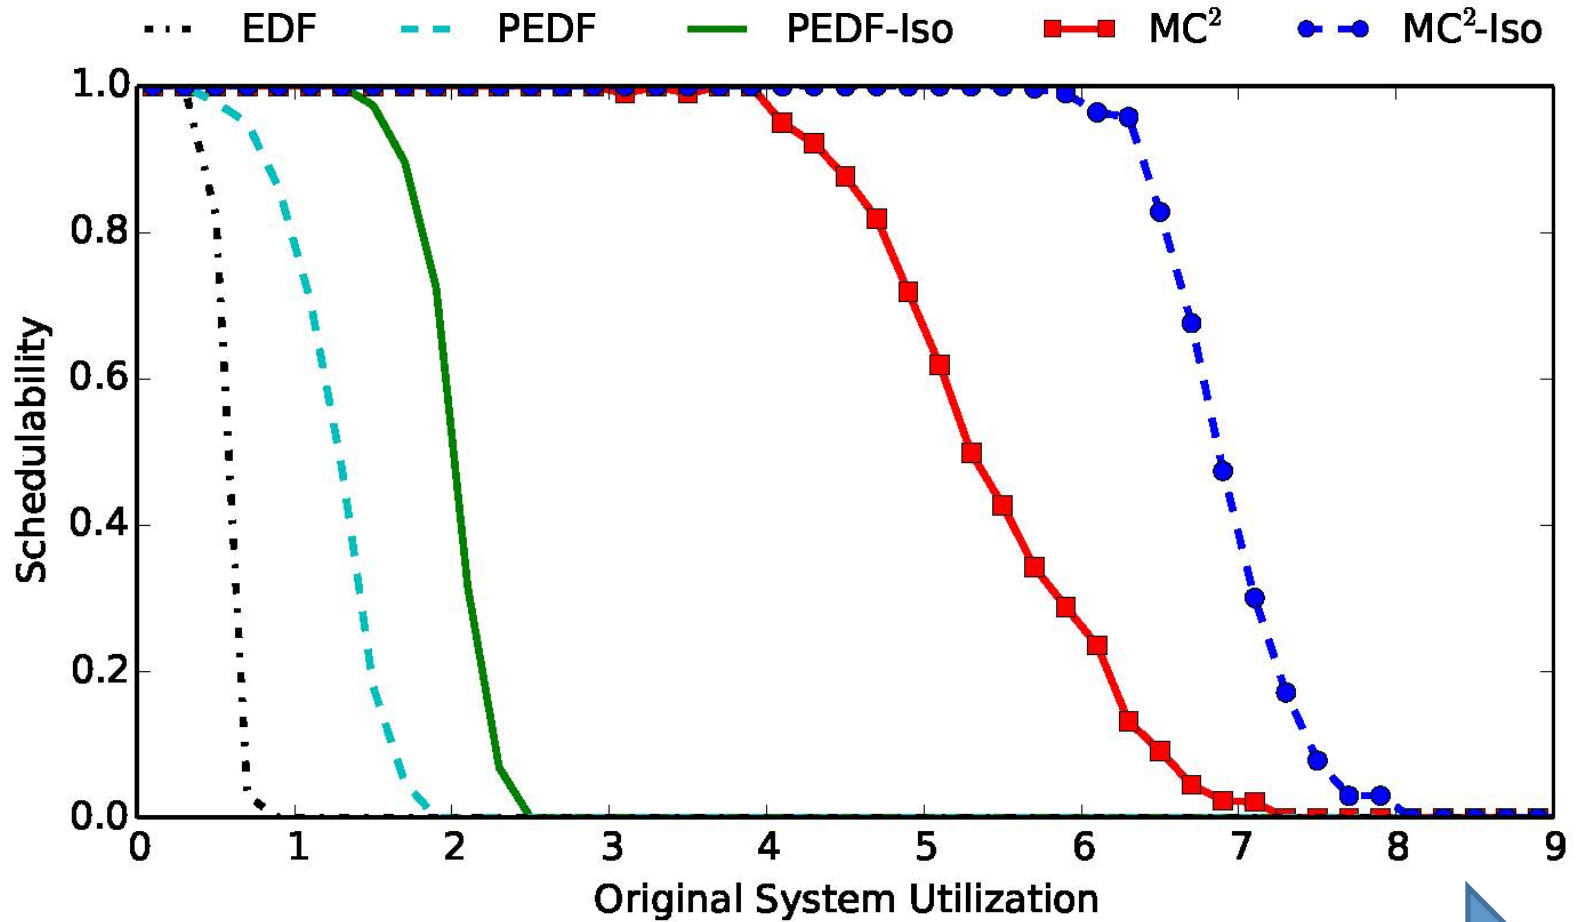
## A Little Experimental Evidence in Favor of our Approach



## A Little Experimental Evidence in Favor of our Approach



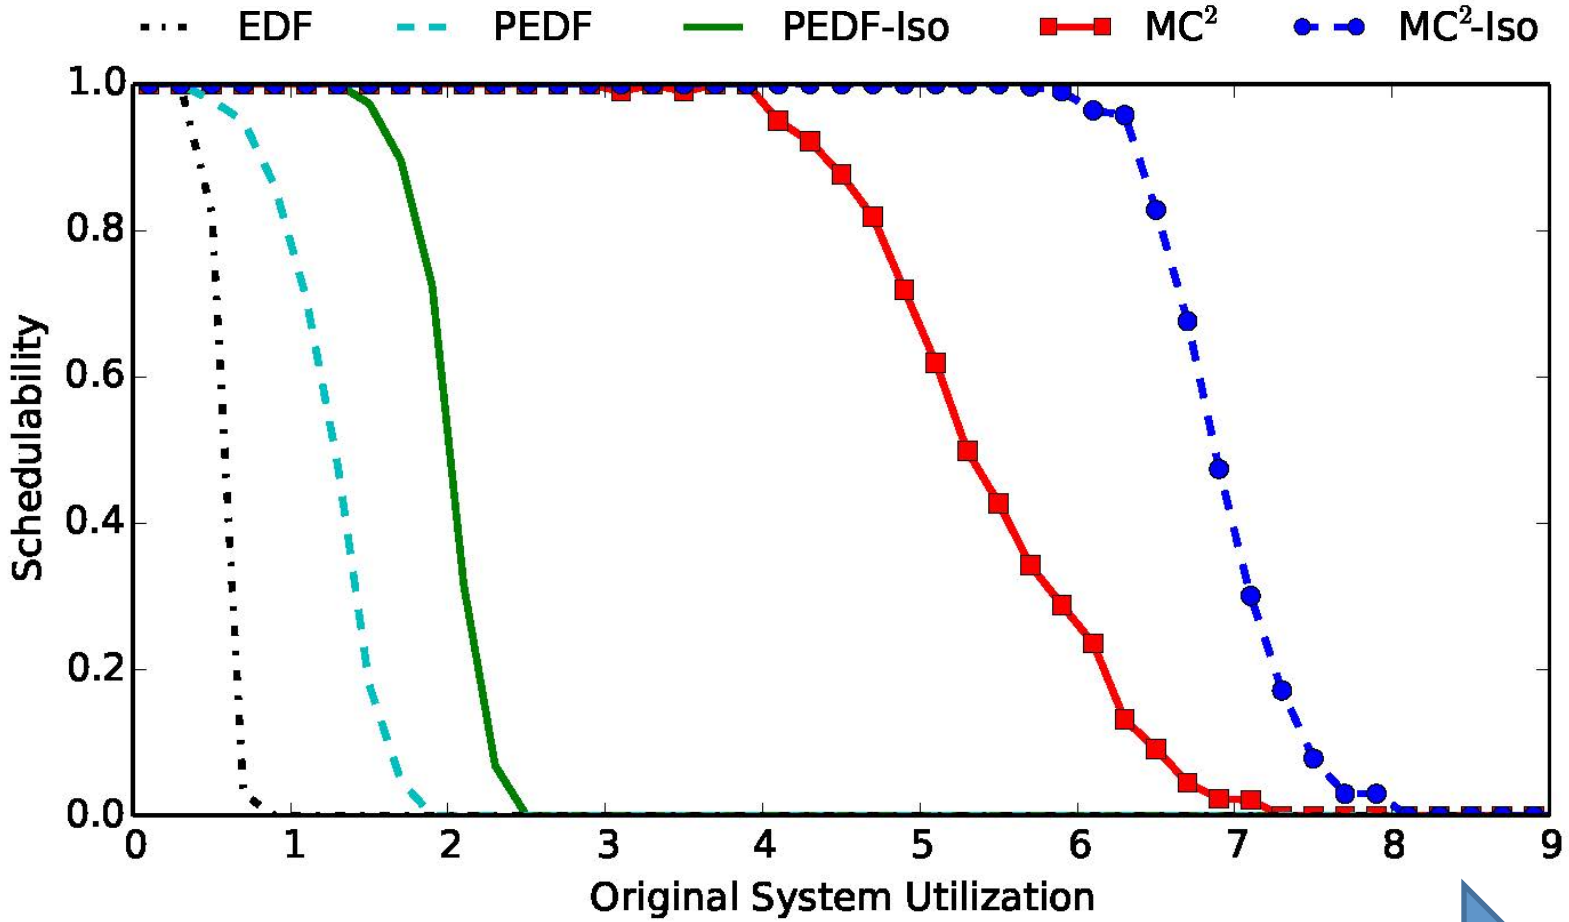
## A Little Experimental Evidence in Favor of our Approach



CPS 1239135, PI: James H. Anderson,  
UNC Chapel Hill, co-PI: Frank Mueller,  
NC State University.

# A Little Experimental Evidence in Favor of our Approach

Fraction of Task Systems Deemed Schedulable

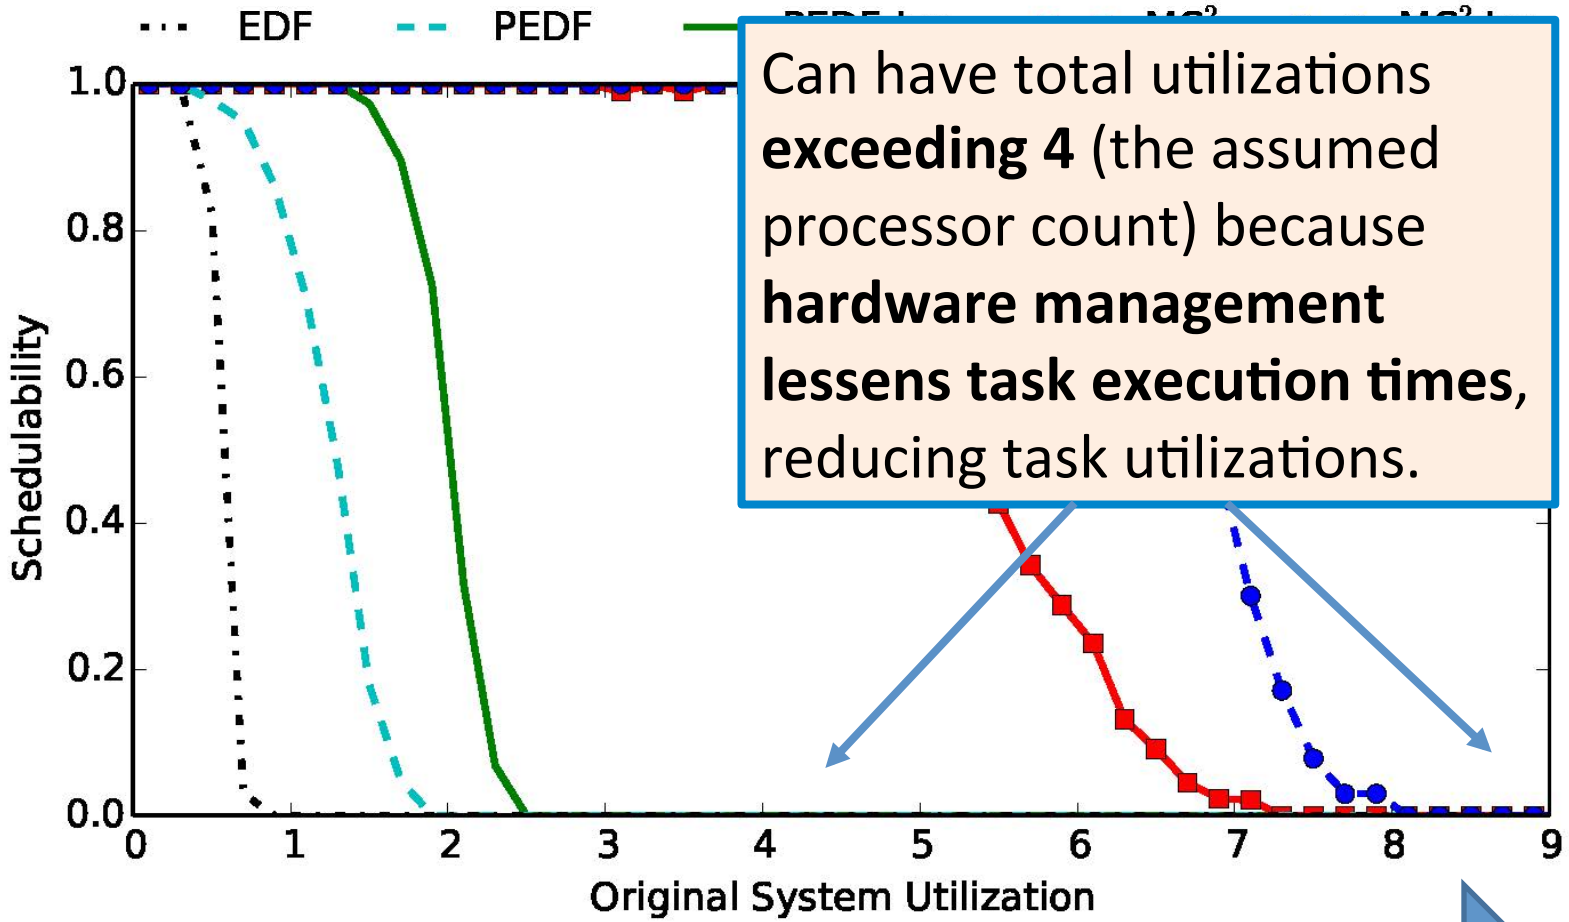


Total System Utilization

CPS 12: James H. Anderson, UNC Charlotte, co-PI: Frank Mueller, NC State University.

# A Little Experimental Evidence in Favor of our Approach

Fraction of Task Systems Deemed Schedulable



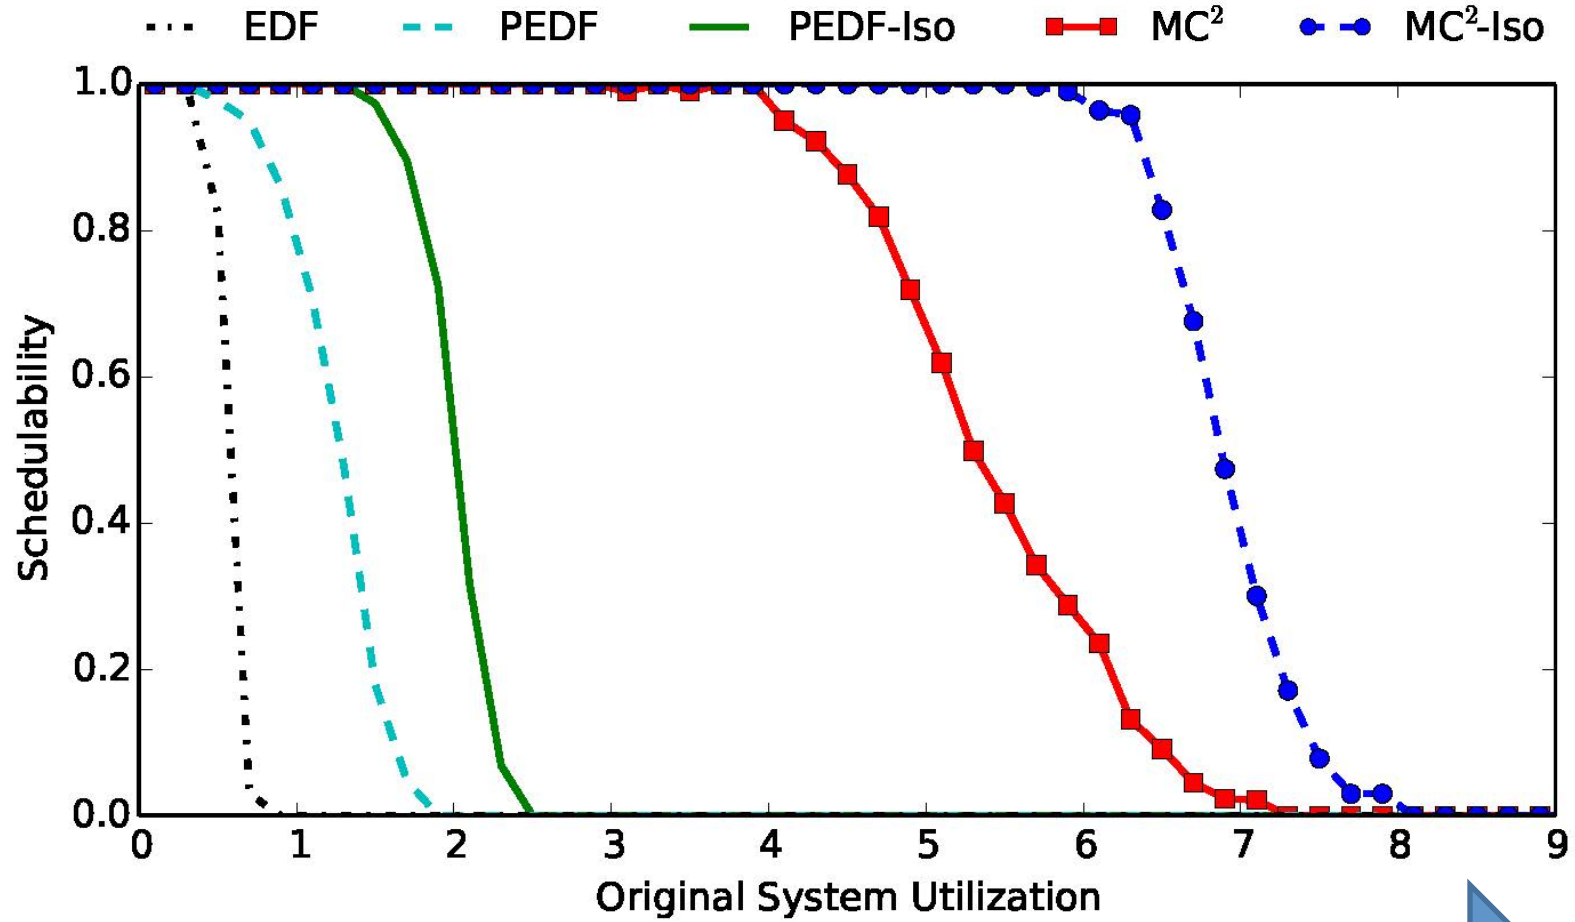
Can have total utilizations **exceeding 4** (the assumed processor count) because **hardware management lessens task execution times**, reducing task utilizations.

Total System Utilization

CPS 12: James H. Anderson, UNC Charlotte, co-PI: Frank Mueller, NC State University.

# A Little Experimental Evidence in Favor of our Approach

Fraction of Task Systems Deemed Schedulable

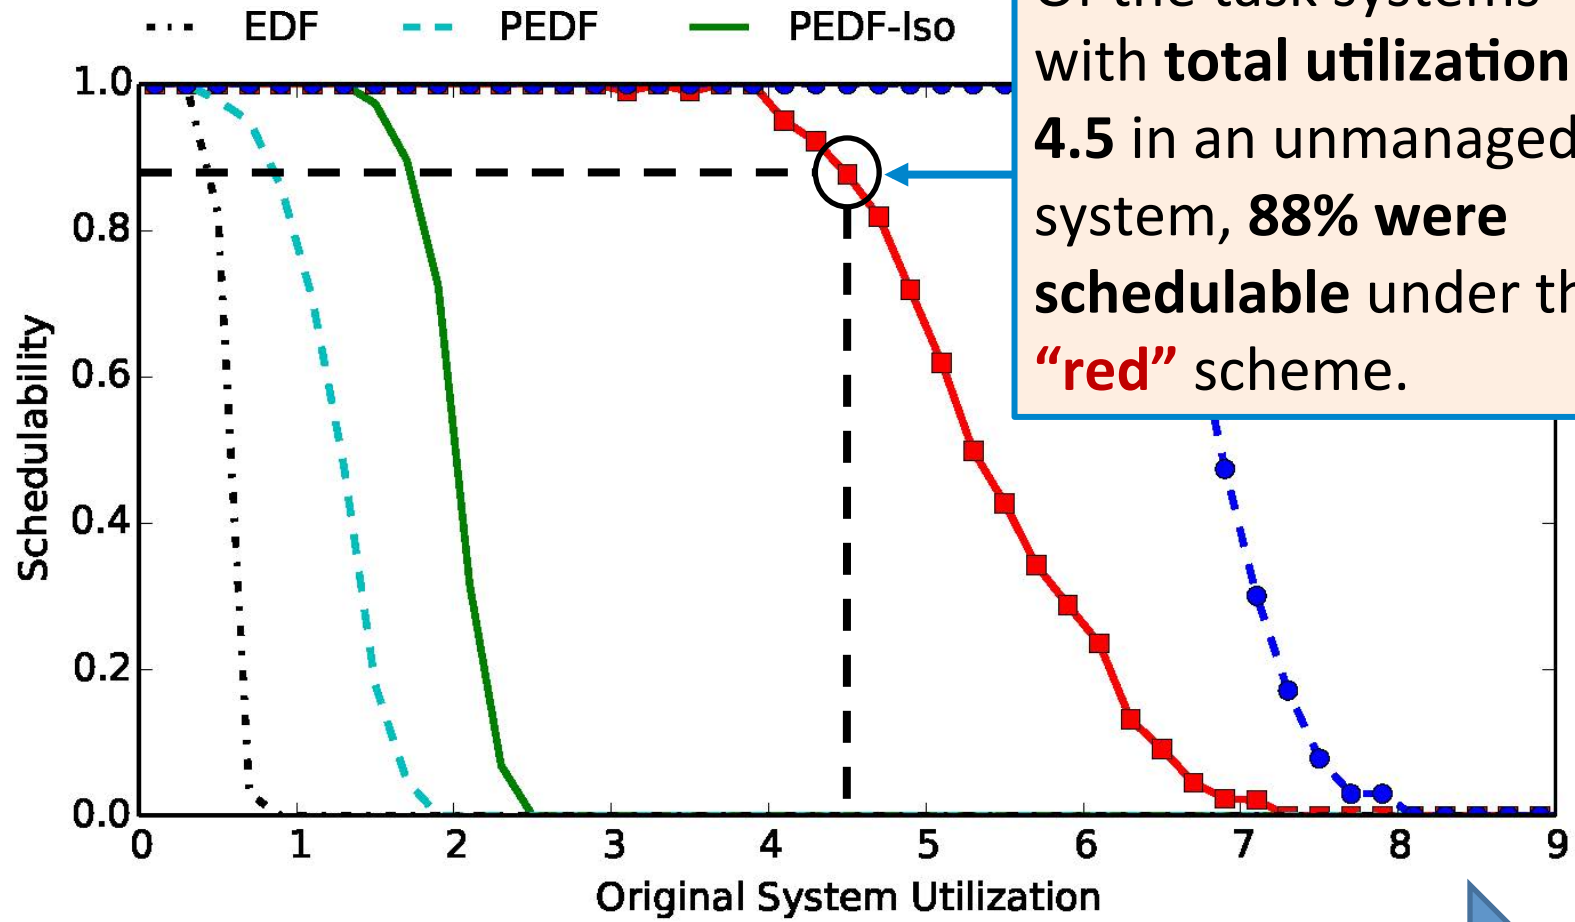


Total System Utilization

CPS 12: James H. Anderson, UNC Charlotte, co-PI: Frank Mueller, NC State University.

# A Little Experimental Evidence in Favor of our Approach

Of the task systems with **total utilization 4.5** in an unmanaged system, **88% were schedulable** under the **“red”** scheme.



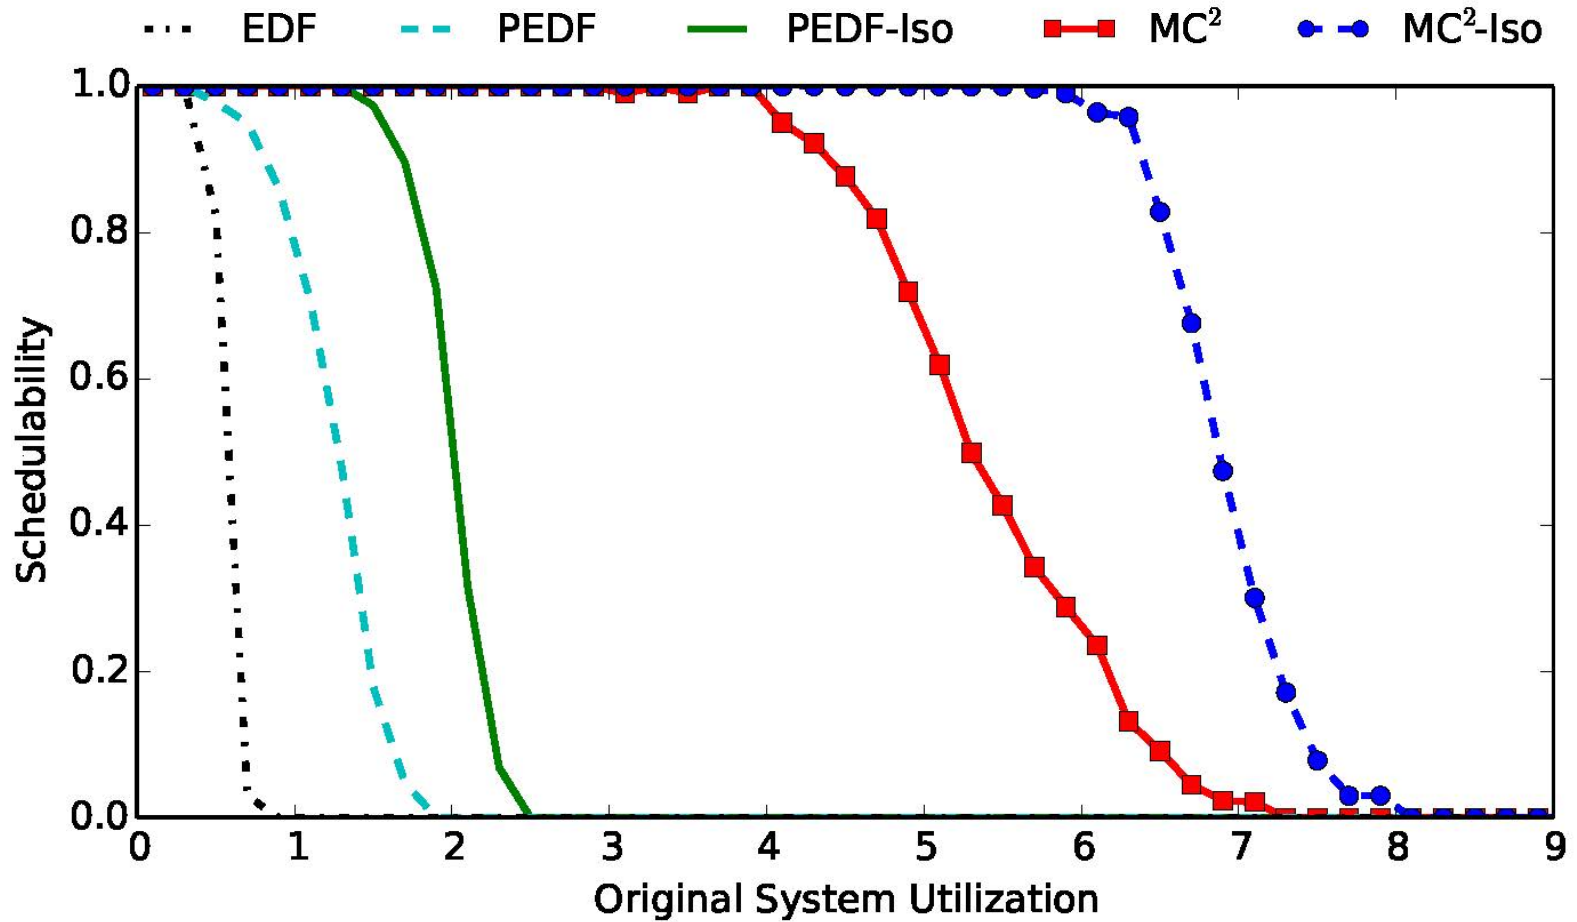
Fraction of Task Systems Deemed Schedulable

Total System Utilization

CPS 12: James H. Anderson, UNC Charlotte, co-PI: Frank Mueller, NC State University.

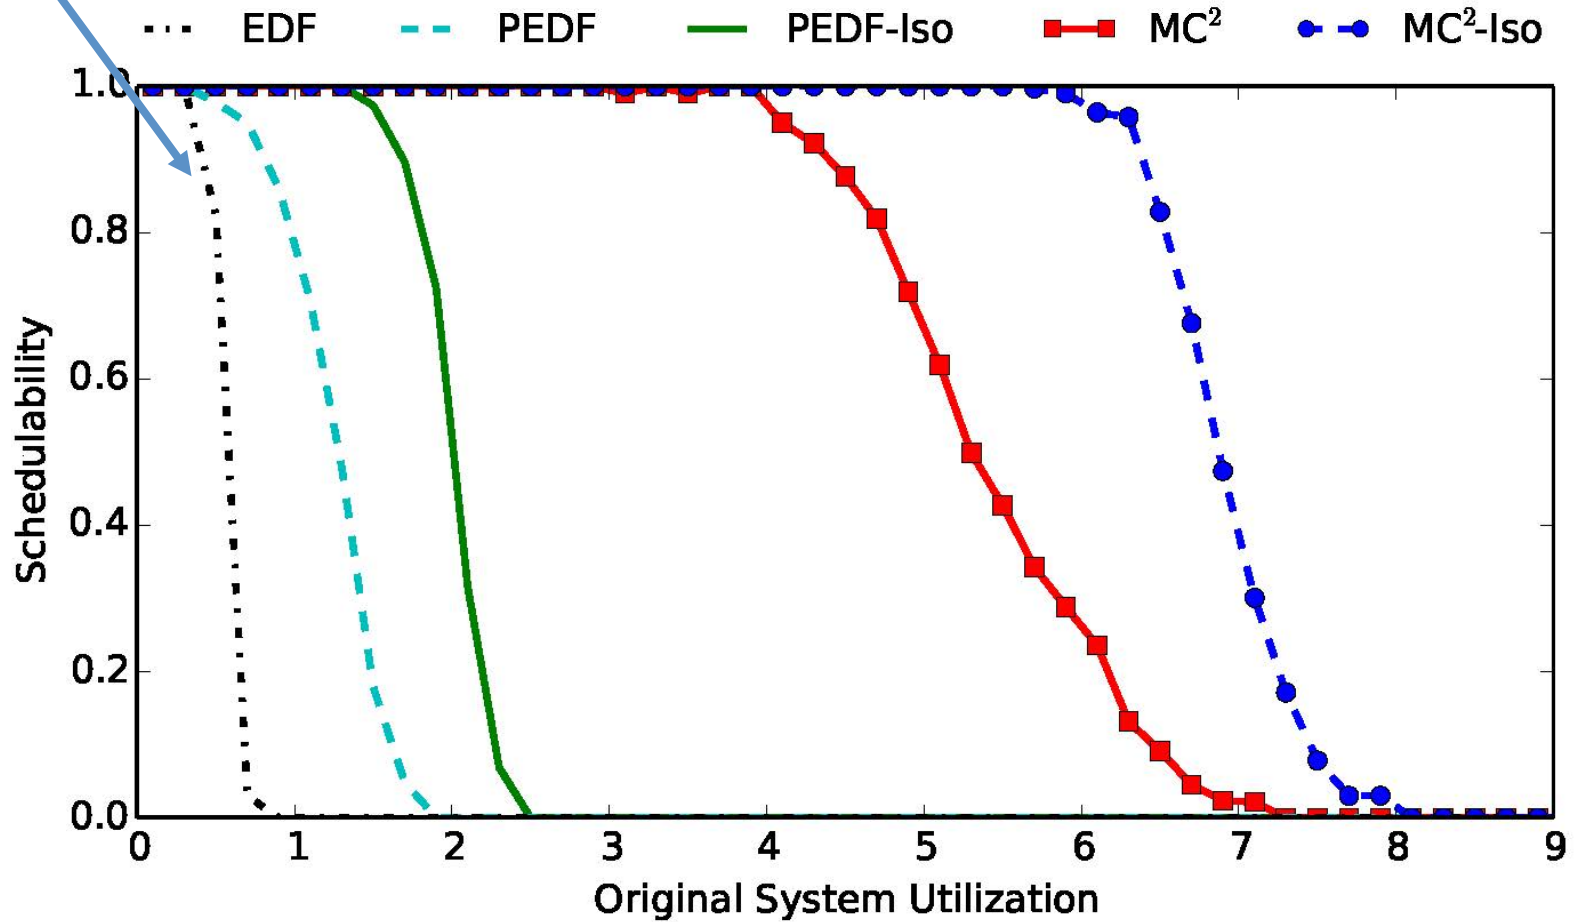


## A Little Experimental Evidence in Favor of our Approach



Uniprocessor  
EDF (the current  
de facto standard)

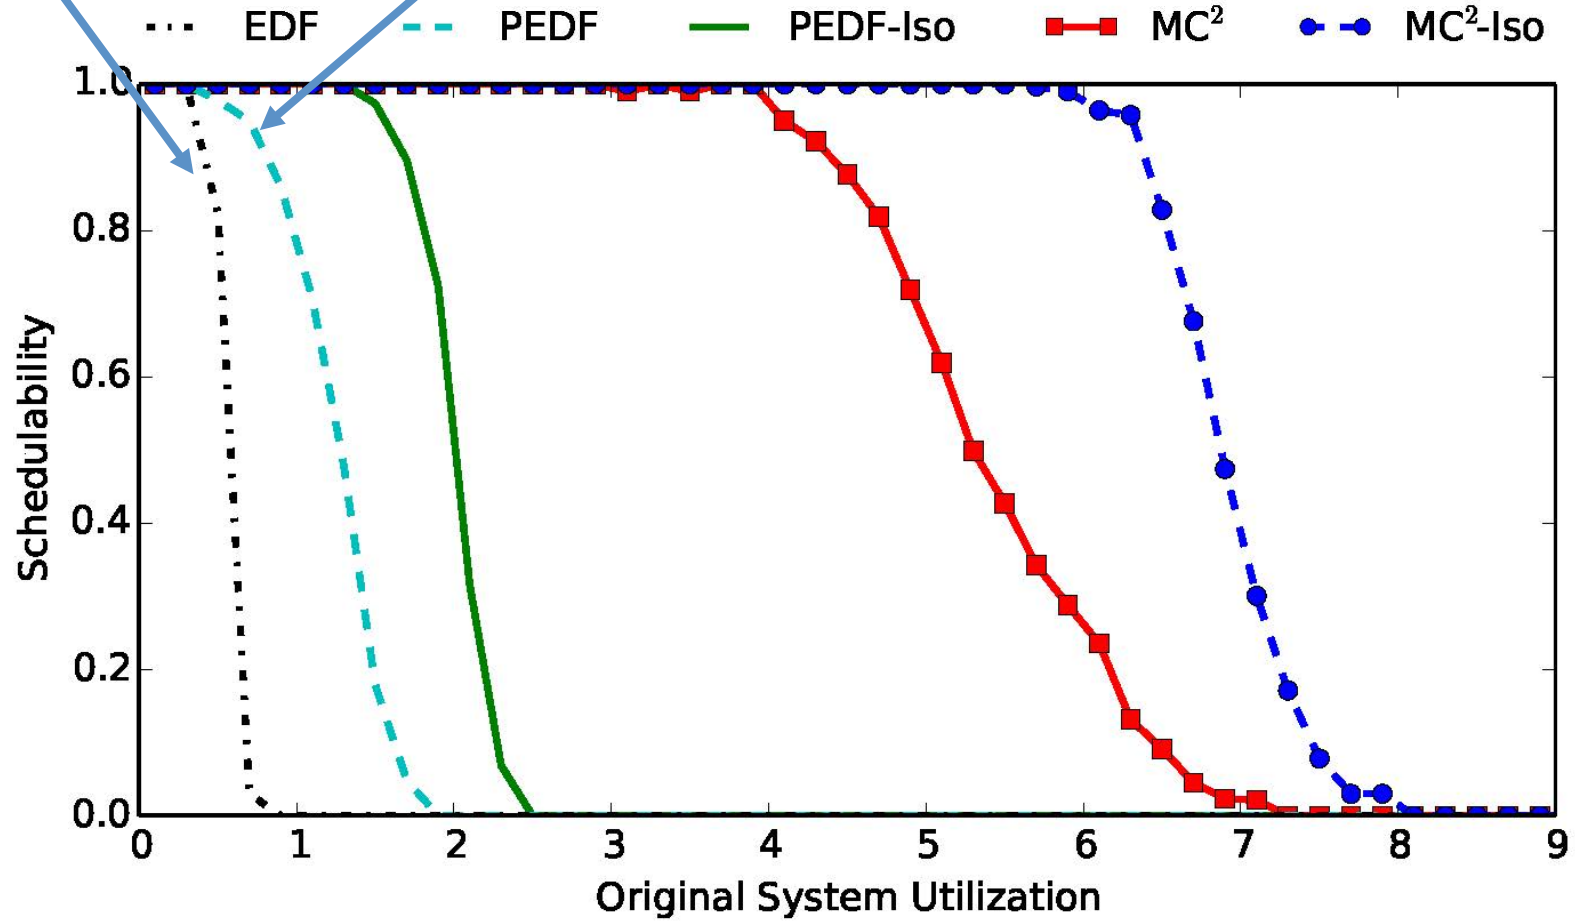
# Experimental Evidence in Favor of our Approach



Uniprocessor  
EDF (the current  
de facto standard)

Partitioned EDF

# Experimental Evidence in Favor of our Approach



CPS 1239135, PI: James H. Anderson,  
UNC Chapel Hill, co-PI: Frank Mueller,  
NC State University.

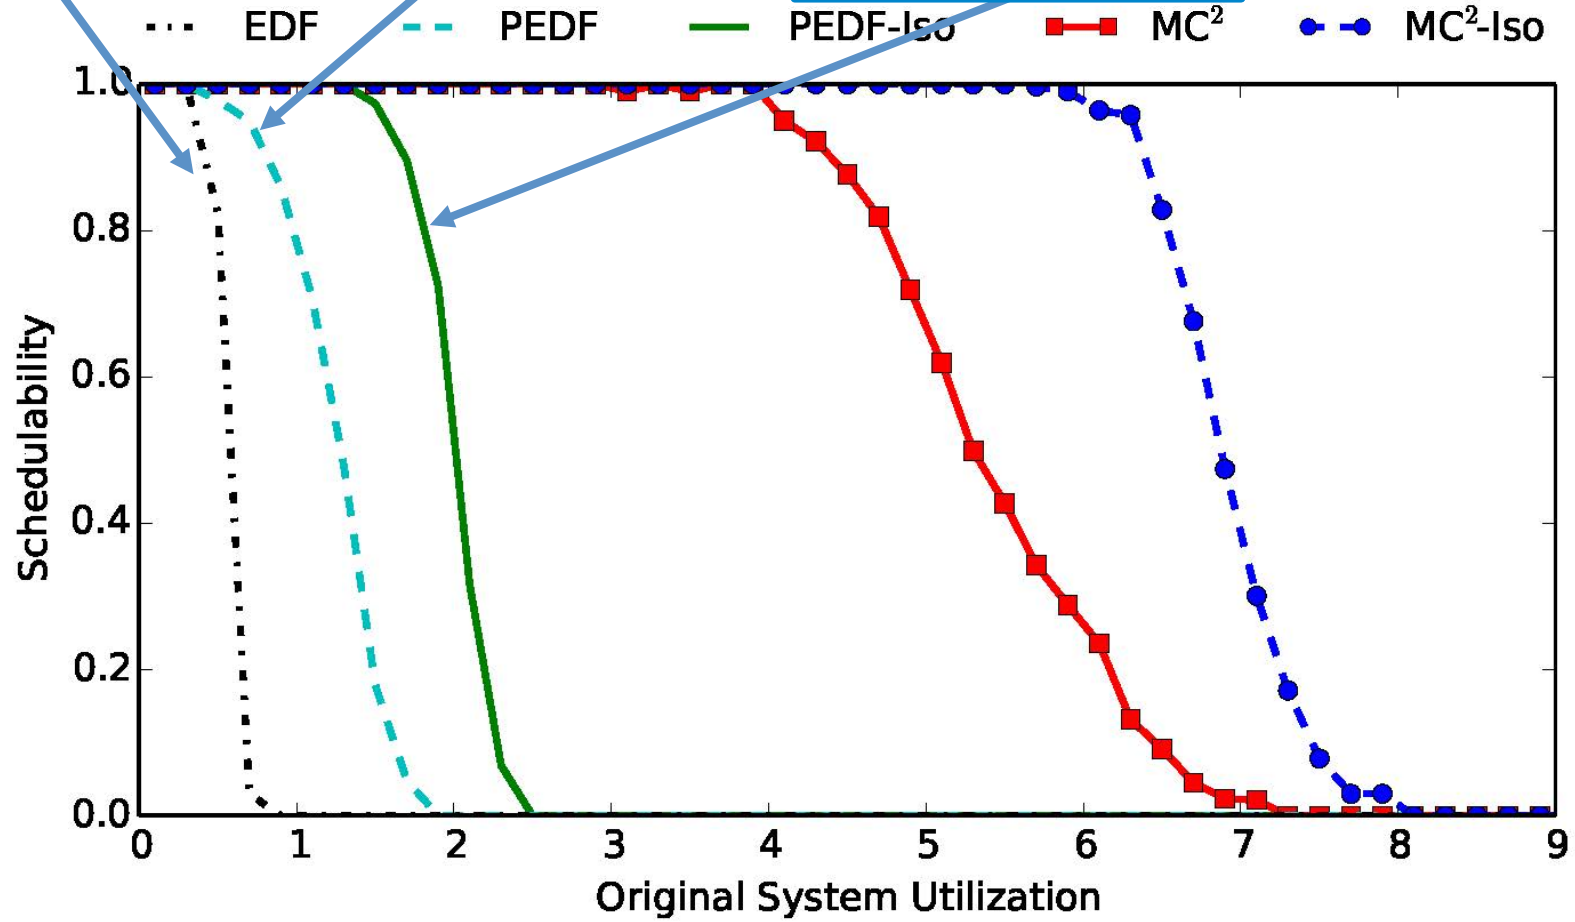
Uniprocessor EDF (the current de facto standard)

Partitioned EDF

Partitioned EDF with hardware management

### Experimental Evidence

### Approach



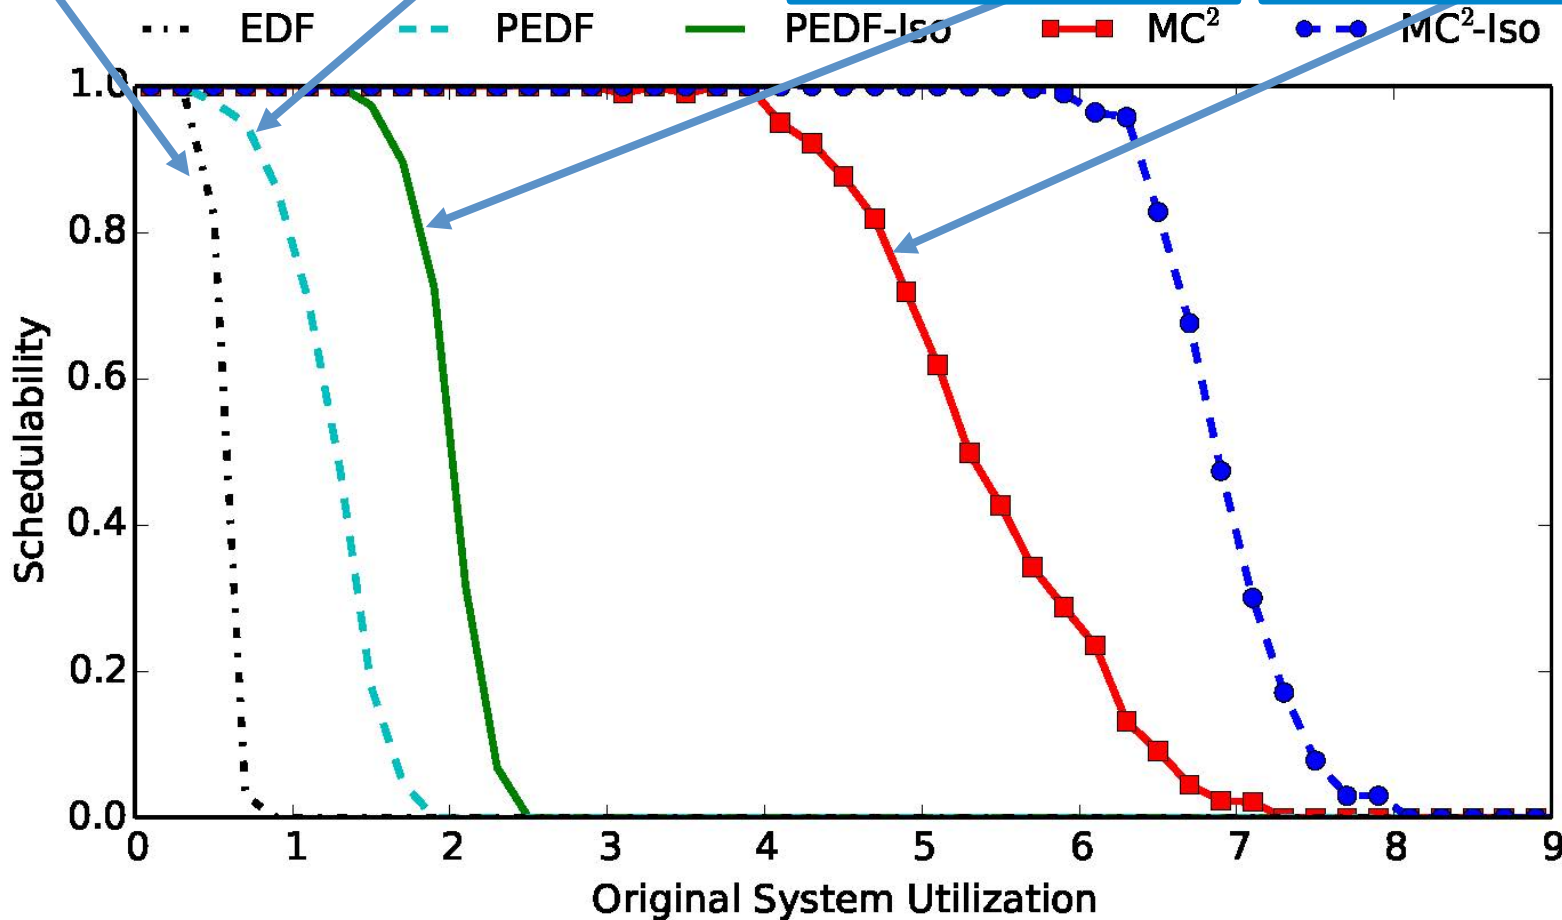
CPS 1239135, PI: James H. Anderson, UNC Chapel Hill, co-PI: Frank Mueller, NC State University.

Uniprocessor EDF (the current de facto standard)

Partitioned EDF

Partitioned EDF with hardware management

MC<sup>2</sup> without hardware management



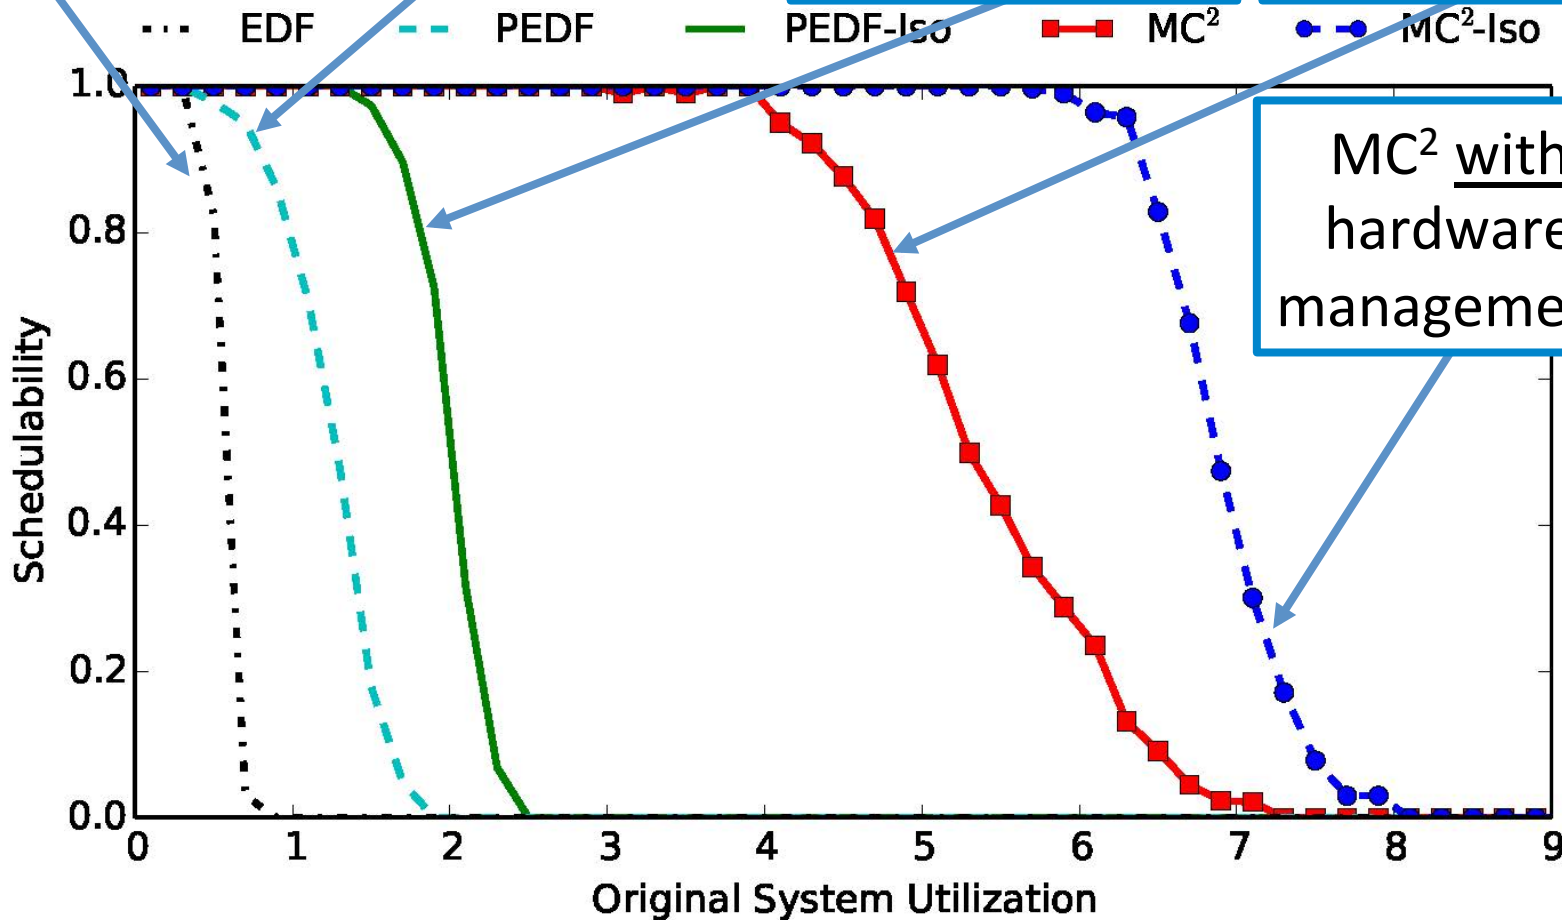
Uniprocessor EDF (the current de facto standard)

Partitioned EDF

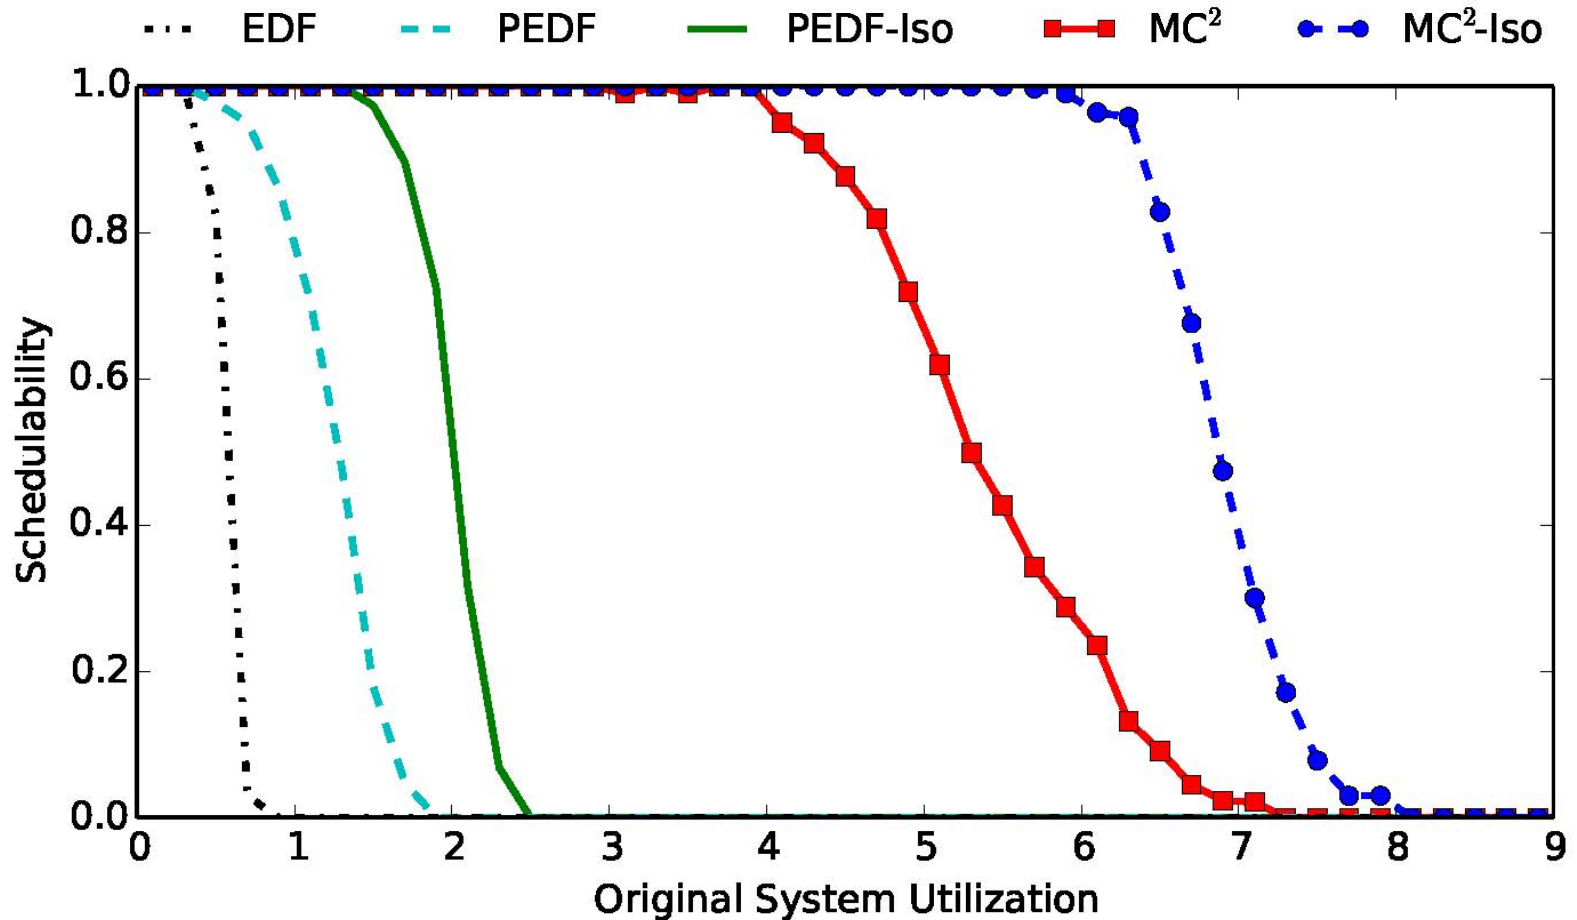
Partitioned EDF with hardware management

MC<sup>2</sup> without hardware management

MC<sup>2</sup> with hardware management



## A Little Experimental Evidence in Favor of our Approach



## Project Summary

# CPS: Breakthrough: Collaborative Research: Bringing the Multicore Revolution to Safety-Critical Cyber-Physical Systems

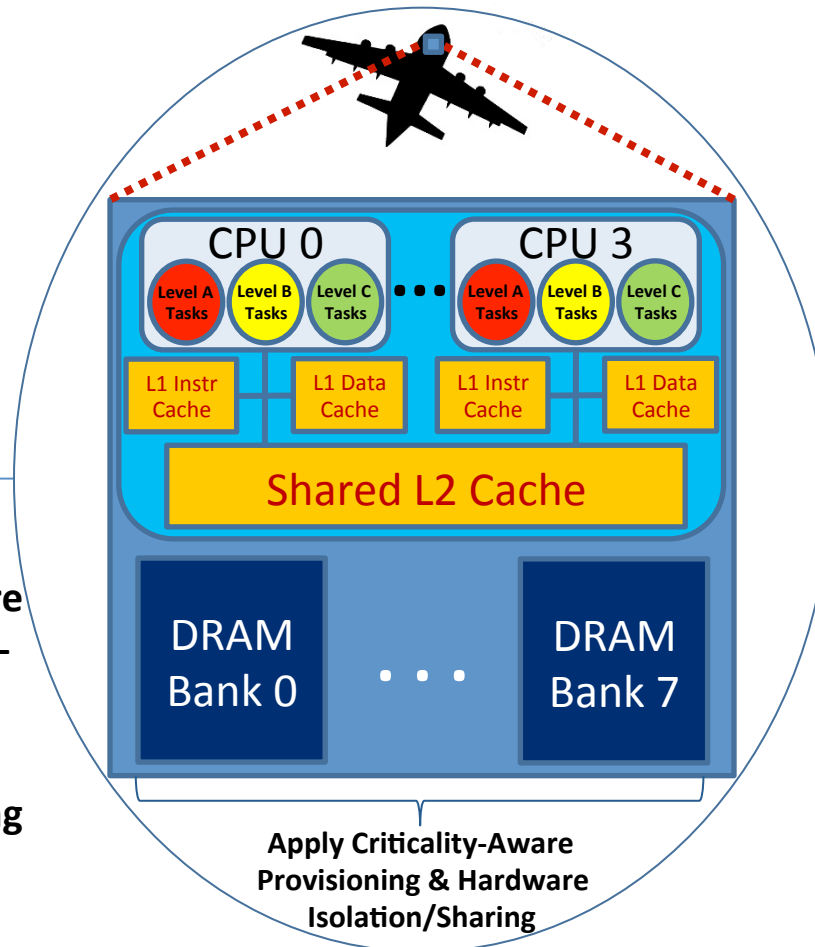
### Challenge:

- To enable the usage of **multicore platforms** in safety critical domains, particularly avionics.
- On multicore, **cross-core interactions** are hard to predict.

### Solution:

- Combine **shared hardware management** with **mixed-criticality provisioning**.
- Enforce **isolation** where needed but enable **sharing** where possible.

CPS 1239135, PI: James H. Anderson, UNC Chapel Hill, co-PI: Frank Mueller, NC State University.



### Scientific Impact:

- Proposed framework can be generally applied to support **computationally intensive real-time workloads on multicore**.
- Workloads **4-8x “larger”** were enabled in experiments.

### Broader Impact:

- Adopted approach devised in consultation with **avionics industry** colleagues.
- Research directly targets concerns raised in the FAA’s **CAST 32** report.
- Two new **CPS courses** created.