

ARTIST Summer School in Europe 2009
Autrans (near Grenoble), France
September 7-11, 2009

Techniques for multiprocessor real-time scheduling

Invited Speaker: Sanjoy Baruah
The University of North Carolina

Techniques for multiprocessor real-time scheduling

Sanjoy Baruah

The University of North Carolina at Chapel Hill

Supported by the US National Science Foundation, the US Army Research Office, the US Air Force Office of Scientific Research, Northrop Grumman Corp., and Intel Corp.

Techniques for **multiprocessor** real-time scheduling

Why multiprocessors?

- provide greater **computing capacity**, at lower **cost**
- many real-time applications are **inherently parallelizable**
- uniprocessor systems are becoming **obsolete**

-(multicore CPU's)

Goal: A theory of multiprocessor real-time scheduling

Techniques for multiprocessor Deadline First scheduling

Earliest

Why multiprocessors?

- provide greater computing capacity, at lower cost
- many real-time applications are inherently parallelizable
- uniprocessor systems are becoming obsolete

-(multicore CPU's)

Goal: A theory of multiprocessor real-time scheduling

Why Earliest Deadline First (EDF)?

- more widely studied on multiprocessors
- analysis techniques (appear to) generalize

Techniques for multiprocessor Deadline First scheduling

Earliest

Overview of presentation

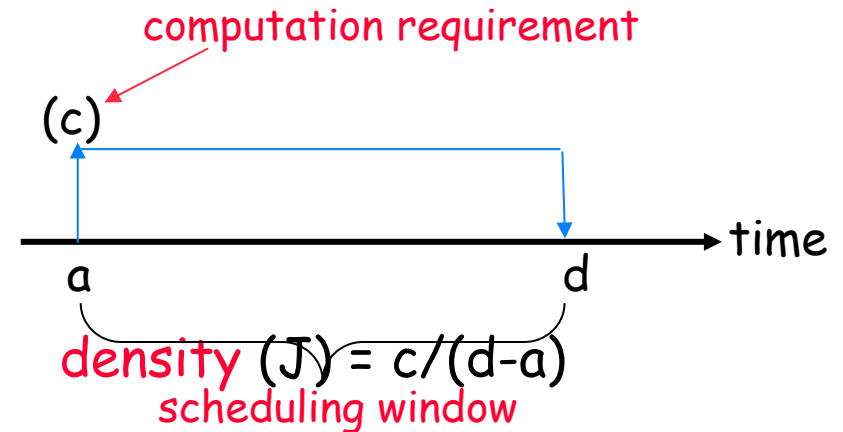
- * Task and machine model
 - Sporadic tasks; partitioned and global scheduling
- * The demand bound function (DBF)
- * Overview of theoretical results
 - Algorithms and lower bounds
- * Pragmatic considerations

Task model

Jobs executing on $m > 1$ identical processors

Job $J = (a, c, d)$

- Preemptable
- Not parallelizable



Recurring tasks or processes

- finite (a priori known) number of them
- generate the jobs
- represent code within an infinite loop
- different tasks are assumed independent

The sporadic task model

Task $\tau_i = (C_i, D_i, T_i)$

- worst-case execution requirement
- relative deadline
- minimum inter-arrival separation ("period")

Jobs

- first job arrives at any t
- consecutive arrivals $\geq T_i$
- each job has execution requirement C_i
- each job has its deadline D_i


NOTATION:

$$\tau = \{\tau_1, \tau_2, \dots, \tau_n\};$$

$(D_i \leq T_i \text{ for all } i)$

$$\text{dens}_{\max}(\tau) = \max_{\text{all } \tau_i \text{ in } \tau} \left(C_i / D_i \right)$$

Example: $\tau_i = (2, 3, 5)$



Global and partitioned scheduling

1. PARTITIONED

- Each task assigned to a processor

2. GLOBAL

- A job may execute on any processor
- A preempted job may resume on any processor

Global schedulability **dominates** partitioned schedulability

Global scheduling may have **higher run-time overhead**

The Demand Bound Function

$DBF(\tau_i, t) \equiv$ maximum cumulative execution requirement of jobs of sporadic task τ_i in any interval of length t

$$\text{load}(\tau) \equiv \max_{\text{all } t} \left(\sum_{\tau_i \in \tau} DBF(\tau_i, t) / t \right) \quad DBF(\tau_i, t) = c_i \times \max \left(0, \left\lfloor \frac{t - D_i}{T_i} \right\rfloor + 1 \right)$$

RESULT: Any sporadic task system τ is EDF-schedulable on a preemptive uniprocessor if $\text{load}(\tau) \leq 1$

Maximum total execution requirement by jobs of sporadic task system τ over any time-interval of length t

RESULT: Any L&L task system τ is RM-schedulable on a preemptive uniprocessor if $\text{load}(\tau) \leq \ln_e 2$


RESULT: Any sporadic task system τ is DM-schedulable on a preemptive uniprocessor if $\text{load}(\tau) \leq 0.567\dots$

Ω , the solution to the equation $x = \ln_e (1/x)$

The Demand Bound Function

$DBF(\tau_i, t) \equiv$ maximum cumulative execution requirement of jobs of sporadic task τ_i in any interval of length t

Baker and Cirinei. *A necessary and sometimes sufficient condition for the feasibility of sets of sporadic hard-deadline tasks.* RTSS 2006.




The Demand Bound Function

$\text{DBF}(\tau_i, t, s) \equiv$ maximum cumulative execution requirement of jobs of sporadic task τ_i in any interval of length t , **when executing on a speed- s processor** ($s \leq 1$)

Bonifaci, Marchetti-Spaccamola, and Stiller. *A constant-approximative feasibility test for multiprocessor real-time scheduling*. ESA 2008.

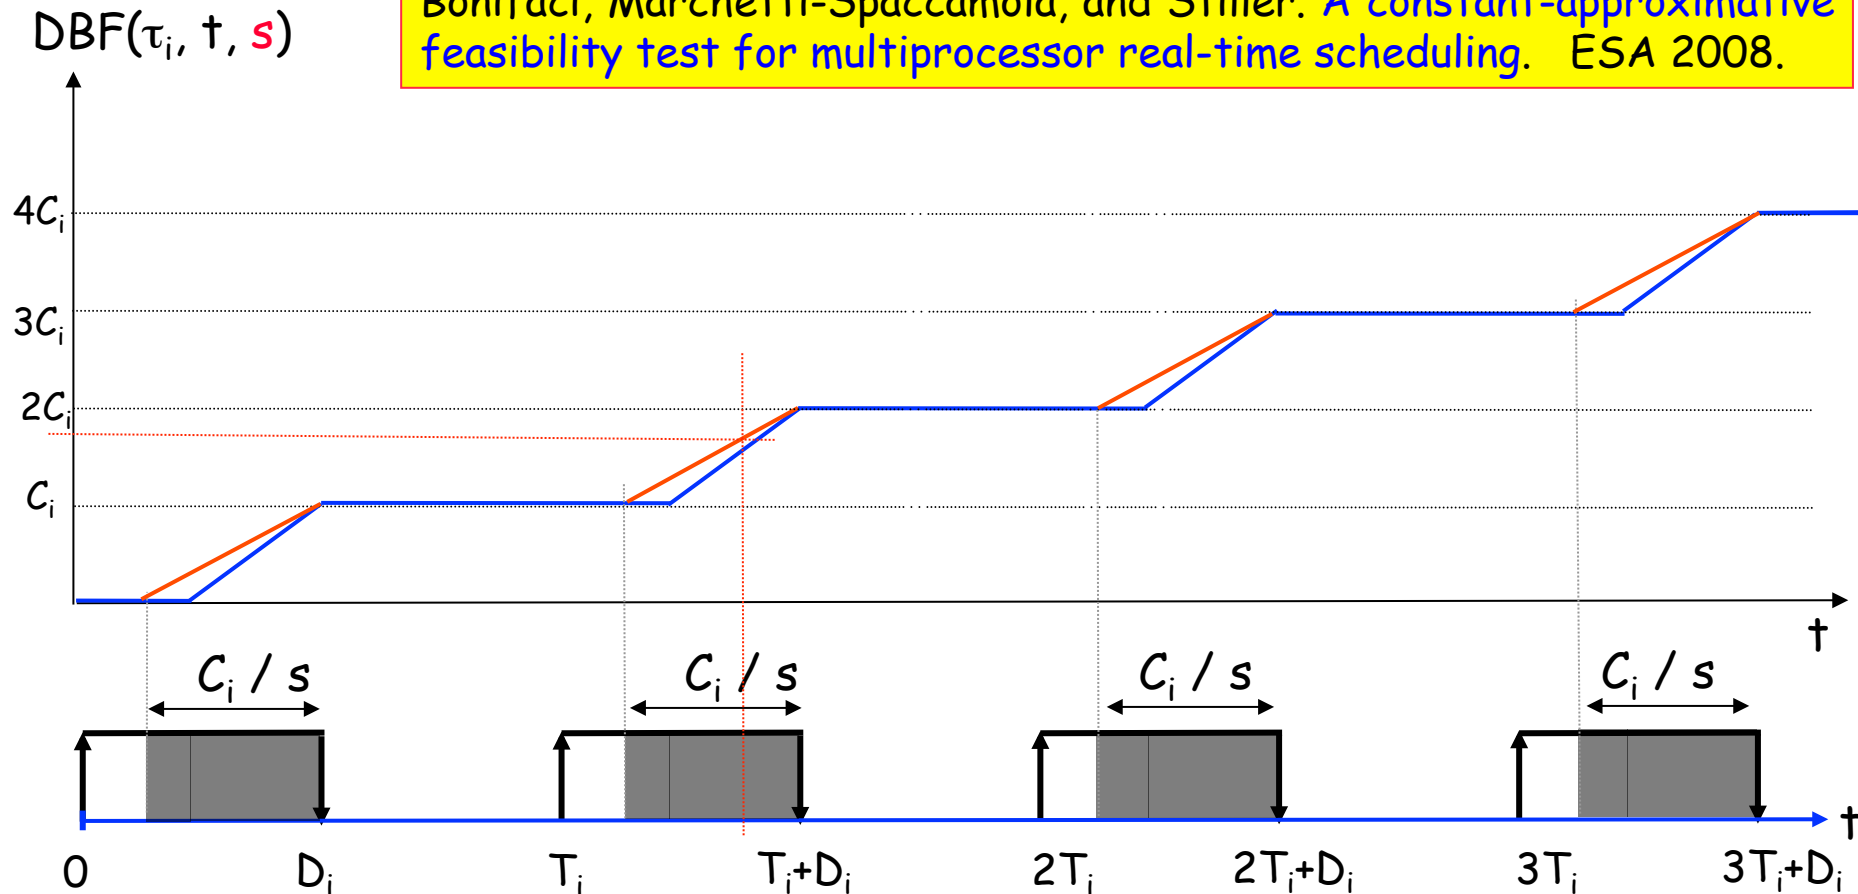
$\text{DBF}(\tau_i, t, 1)$



The Demand Bound Function

$\text{DBF}(\tau_i, t, s) \equiv$ maximum cumulative execution requirement of jobs of sporadic task τ_i in any interval of length t , **when executing on a speed- s processor** ($s \leq 1$)


Bonifaci, Marchetti-Spaccamola, and Stiller. *A constant-approximative feasibility test for multiprocessor real-time scheduling*. ESA 2008.



The Demand Bound Function

$\text{DBF}(\tau_i, t, s) \equiv$ maximum cumulative execution requirement of jobs of sporadic task τ_i in any interval of length t , **when executing on a speed- s processor** ($s \leq 1$)

Bonifaci, Marchetti-Spaccamola, and Stiller. *A constant-approximative feasibility test for multiprocessor real-time scheduling*. ESA 2008.



The Demand Bound Function

$DBF(\tau_i, t, s)$ \equiv maximum cumulative execution requirement of jobs of sporadic task τ_i in any interval of length t , **when executing on a speed- s processor** ($s \leq 1$)

$$\text{load}(\tau, s) \equiv \max_{\text{all } t} \left(\sum_{\tau_i \in \tau} DBF(\tau_i, t, s) / t \right)$$

RESULT: A necessary condition for τ to be [EDF-]schedulable on m speed- s processors: $\text{load}(\tau, s) \leq m s$

GLOBAL EDF SCHEDULING

On $m > 1$ processors

At each instant, schedule m active jobs with the **earliest deadlines**

- (fewer active jobs than processors: idle remaining processors)

A global EDF schedulability test

RESULT: A sufficient condition for τ to be EDF-schedulable on

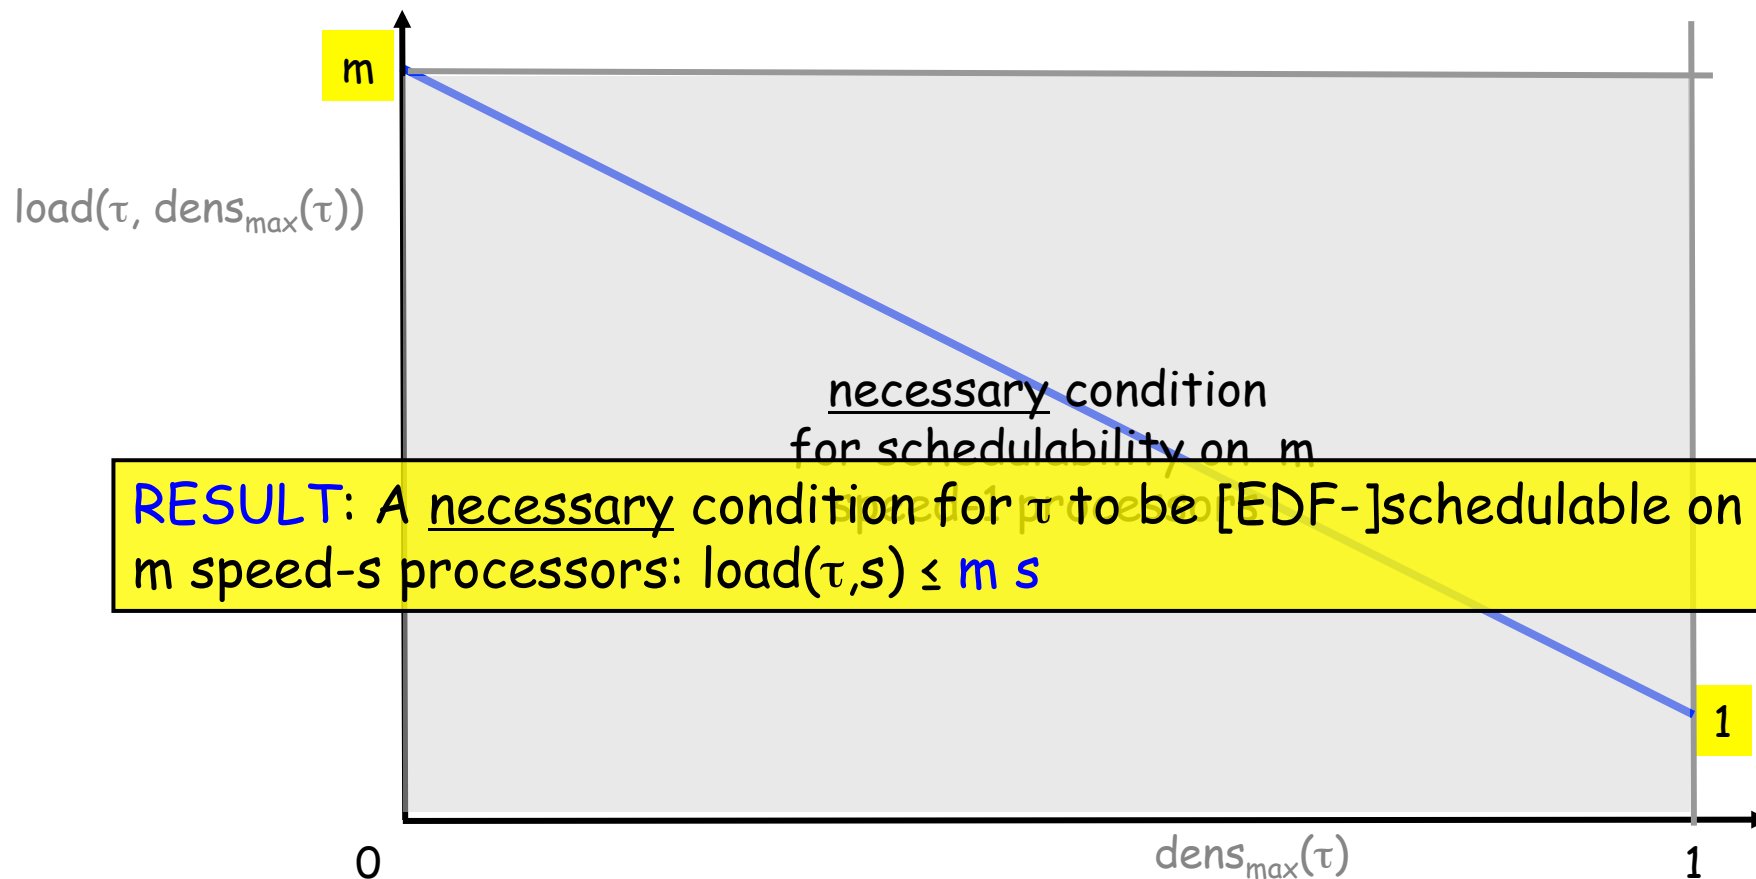
m speed-1 processors: $\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq [m - (m-1) \times \text{dens}_{\max}(\tau)]$

A diagram consisting of a yellow rectangular box with a black border. Inside the box is the mathematical expression $\max_{\text{all } \tau_i \text{ in } \tau} \left[C_i / D_i \right]$. A vertical arrow points upwards from the top center of the box to a horizontal line segment positioned above the $\text{dens}_{\max}(\tau)$ term in the inequality above.

$$\max_{\text{all } \tau_i \text{ in } \tau} \left[C_i / D_i \right]$$

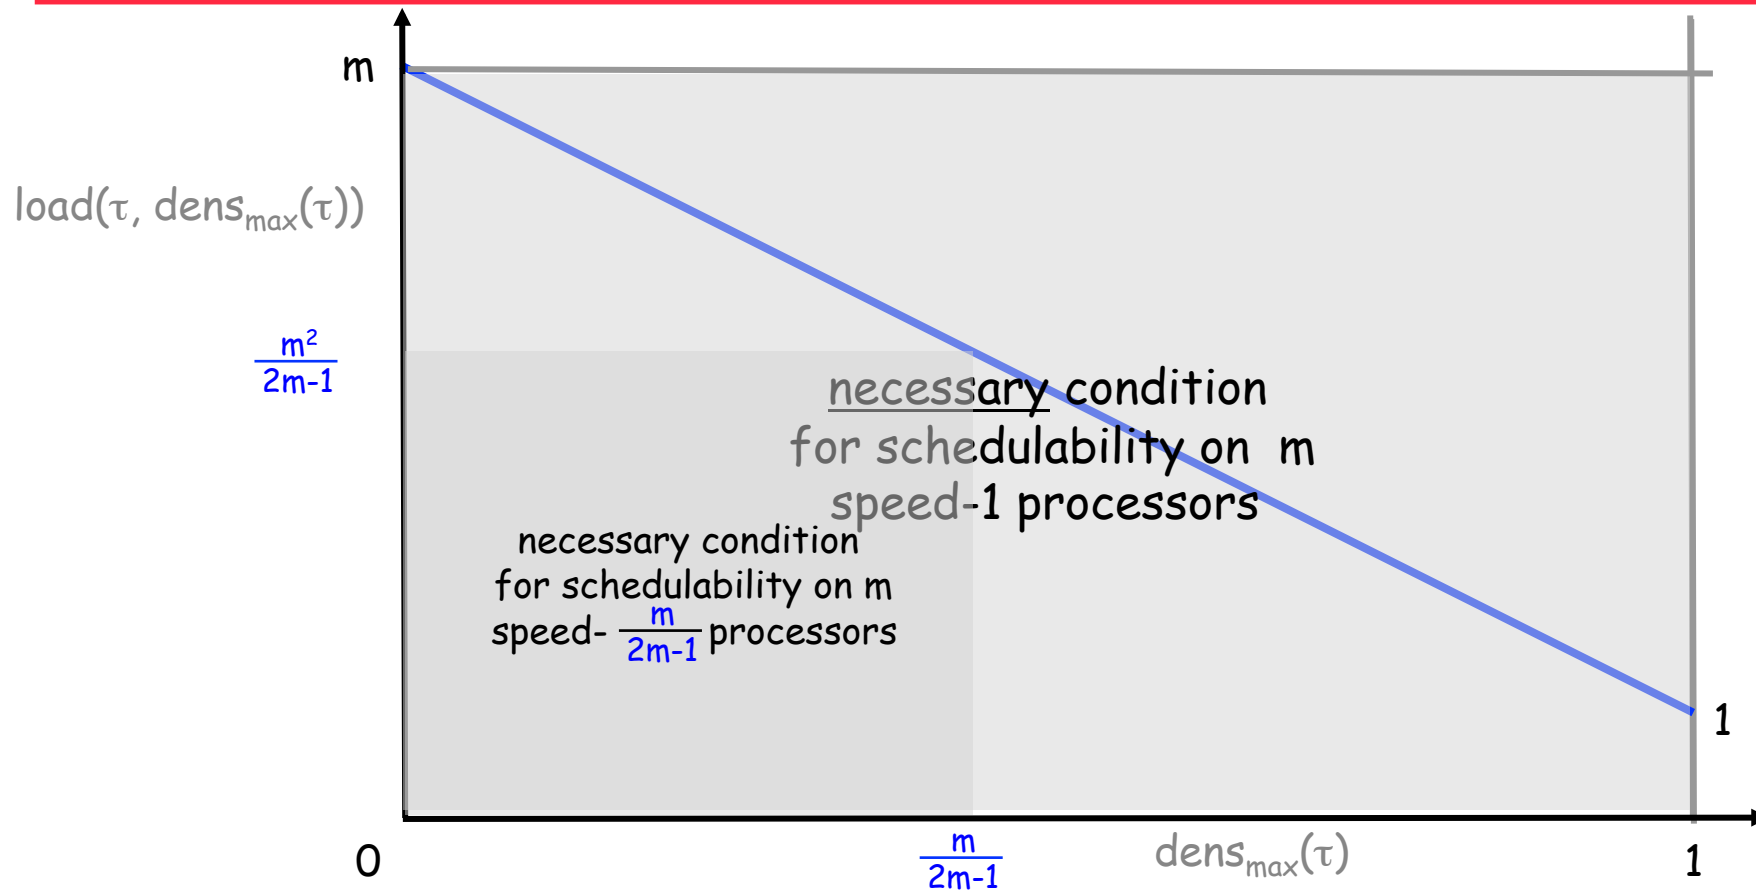
A global EDF schedulability test

RESULT: A sufficient condition for τ to be EDF-schedulable on m speed-1 processors: $\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq [m - (m-1) \times \text{dens}_{\max}(\tau)]$



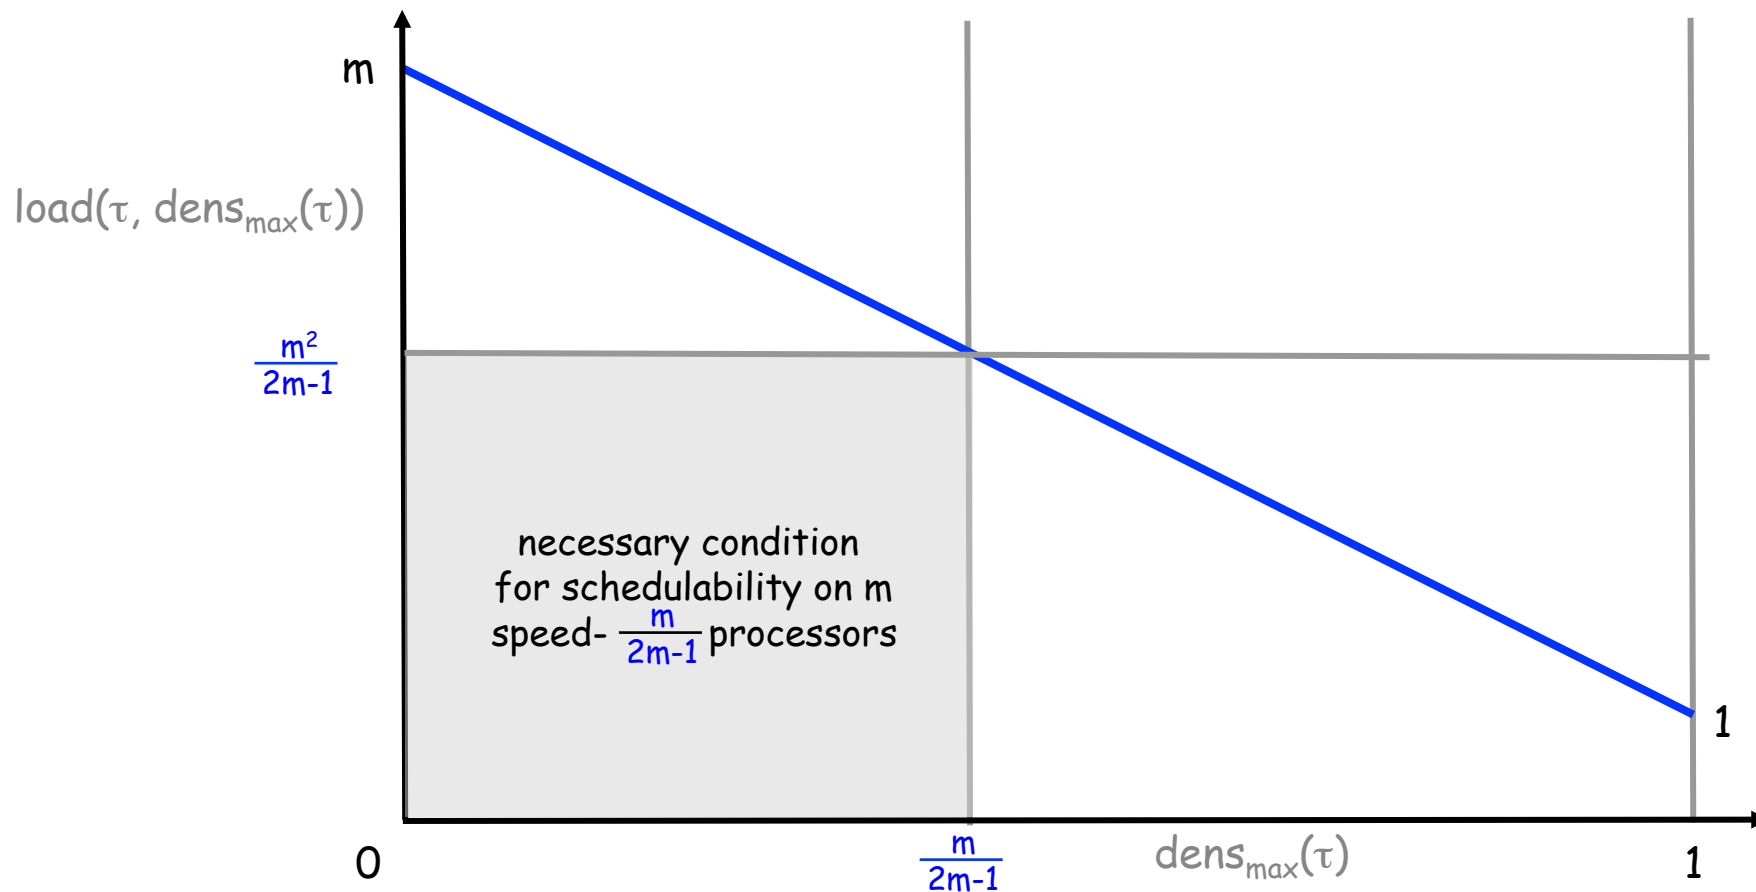
A global EDF schedulability test

Any sporadic task system schedulable upon m processors is global EDF schedulable on m processors that are each $(2 - 1/m)$ times as fast




A global EDF schedulability test - Proof

RESULT: A sufficient condition for τ to be EDF-schedulable on m speed-1 processors: $\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq [m - (m-1) \times \text{dens}_{\max}(\tau)]$



A global EDF schedulability test - Proof

RESULT: A sufficient condition for τ to be EDF-schedulable on m speed-1 processors: $\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq [m - (m-1) \times \text{dens}_{\max}(\tau)]$




Job j_2 :

- arrives at $t_2 < t_1$
- is active at t_1
- has executed for $< (t_1 - t_2) \times \text{dens}_{\max}(\tau)$ by t_1

\Rightarrow All m procs are busy for $> (t_1 - t_2) \times (1 - \text{dens}_{\max}(\tau))$ over $[t_2, t_1)$
 \equiv Total idled capacity over $[t_2, t_1) < (m-1) \times (t_1 - t_2) \times \text{dens}_{\max}(\tau)$

A global EDF schedulability test - Proof

RESULT: A sufficient condition for τ to be EDF-schedulable on m speed-1 processors: $\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq [m - (m-1) \times \text{dens}_{\max}(\tau)]$




Job j_i :

- arrives at $t_i < t_{i-1}$
- is active at t_{i-1}
- has executed for $< (t_{i-1} - t_i) \times \text{dens}_{\max}(\tau)$ by t_{i-1}

\Rightarrow Total idled capacity over $[t_i, t_{i-1}) < (m-1) \times (t_{i-1} - t_i) \times \text{dens}_{\max}(\tau)$

A global EDF schedulability test - Proof

RESULT: A sufficient condition for τ to be EDF-schedulable on m speed-1 processors: $\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq [m - (m-1) \times \text{dens}_{\max}(\tau)]$




Repeat until no further such jobs

- Any job arriving at $t_k < t_i$ has executed for at least $(t_i - t_k) \times \text{dens}_{\max}(\tau)$ over $[t_k, t_i]$

A global EDF schedulability test - Proof

RESULT: A sufficient condition for τ to be EDF-schedulable on m speed-1 processors: $\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq [m - (m-1) \times \text{dens}_{\max}(\tau)]$



Repeat until no further such jobs

1. total exec. requirement over $[t_i, t_0)$ is $\leq \text{dbf}(\tau, L, \text{dens}_{\max}(\tau))$
2. total work done over $[t_i, t_0)$ is $= mL - \text{total idle time over } [t_i, t_0)$
 $\geq mL - (m-1) \times L \times \text{dens}_{\max}(\tau)$

$$\text{dbf}(\tau, L, \text{dens}_{\max}(\tau)) > [m - (m-1) \times \text{dens}_{\max}(\tau)] \times L$$

➔ $\text{load}(\tau, \text{dens}_{\max}(\tau)) > [m - (m-1) \times \text{dens}_{\max}(\tau)]$

PARTITIONED EDF SCHEDULING

1. partitioning algorithm *Must be defined*
2. uniproc. EDF scheduling
Can reuse results for uniproc EDF


DBF*(τ_i, t): an approximation to DBF(τ_i, t, s)

DBF(τ_i, t, s) \equiv maximum cumulative execution requirement of jobs of sporadic task τ_i in any interval of length t , when executing on a speed- s processor ($s \leq 1$)

DBF(τ_i, t, s)

1. $DBF^*(\tau_i, t) \geq DBF(\tau_i, t, s)$

2. $DBF^*(\tau_i, t) < 2 \times DBF(\tau_i, t, s)$



$$2 - \left(\frac{C_i}{(T_i * s)} \right)$$

A partitioning algorithm

Sporadic task system $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$, on m speed-1 processors

assume $D_i \leq D_{i+1}$ for all i

for $i := 1$ to n

- assign τ_i to any processor j such that

τ_i "fits" on processor j

A partitioning algorithm

Sporadic task system $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$, on m speed-1 processors


assume $D_i \leq D_{i+1}$ for all i

for $i := 1$ to n

- assign τ_i to any processor j such that

$$D_i \left\{ \sum_{\tau_k \text{ assigned to proc } j} \text{DBF}^*(\tau_k, D_i) \geq C_i \right.$$

(Upper bound on) execution
already allocated over $[t, t + D_i)$



A partitioning algorithm


Sporadic task system $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$, on m speed-1 processors

assume $D_i \leq D_{i+1}$ for all i

for $i := 1$ to n

- assign τ_i to any processor j such that

$$D_i \left[\sum_{\tau_k \text{ assigned to proc } j} \text{DBF}^*(\tau_k, D_i) \geq C_j \right]$$



A partitioning algorithm

Sporadic task system $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$, on m speed-1 processors

assume $D_i \leq D_{i+1}$ for all i

for $i := 1$ to n

- assign τ_i to any processor j such that

$$D_i \left[\sum_{\tau_k \text{ assigned to proc } j} \text{DBF}^*(\tau_k, D_i) \geq C_i \right]$$

- if no such processor j , then return failure

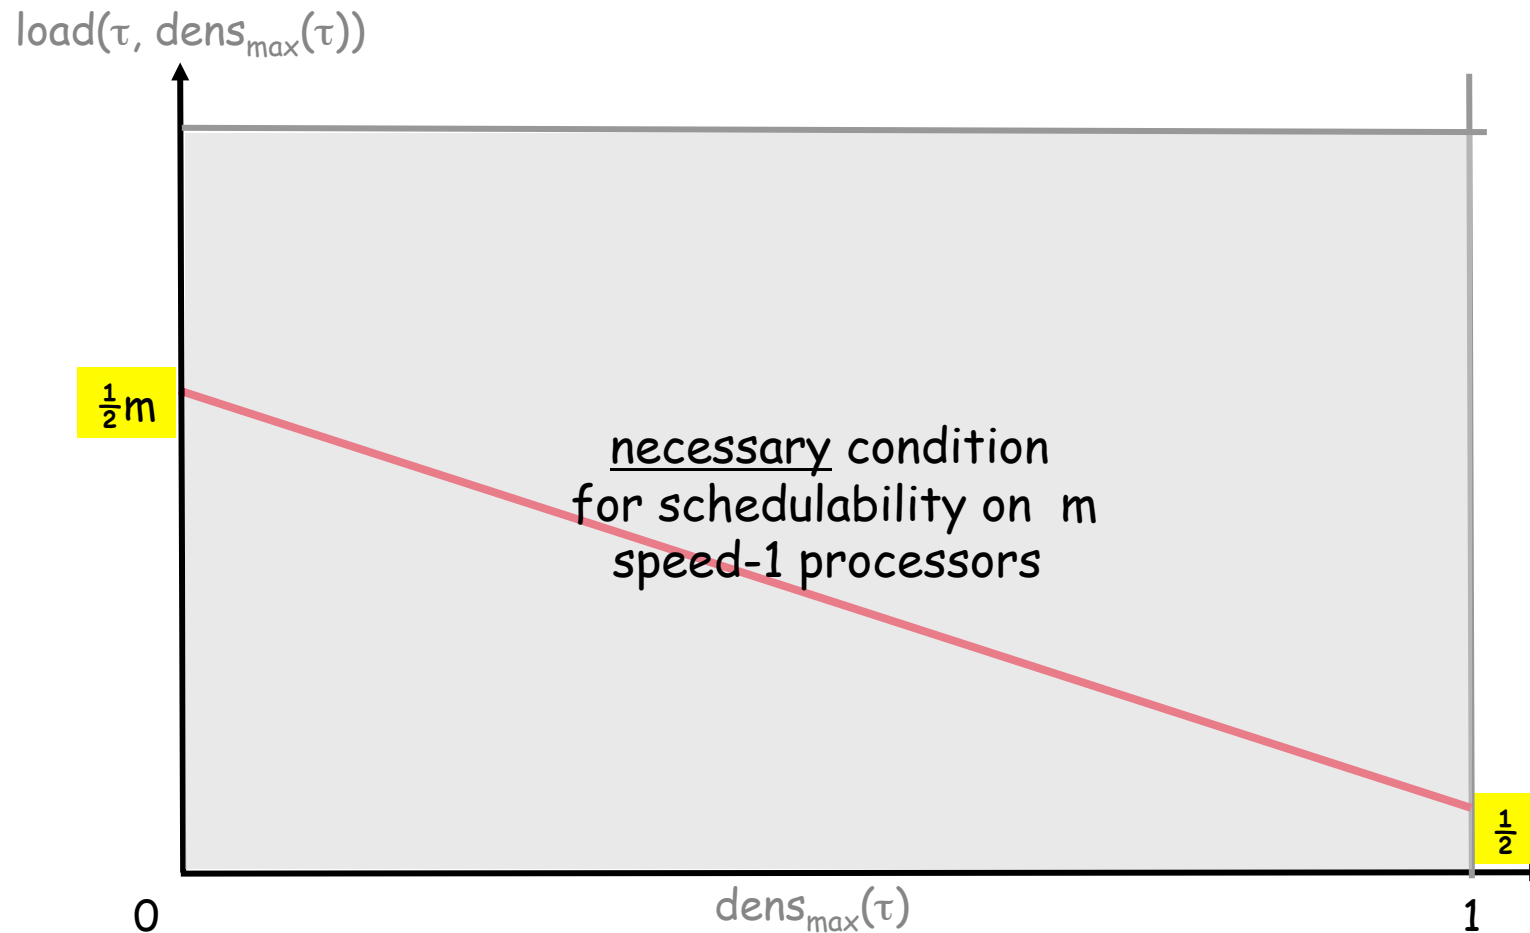
Runtime complexity: $O(n^2)$

Schedulability test: Run the partitioning algorithm!

A property: this algorithm schedules any system satisfying

$$\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq \frac{1}{2} [m - (m-1) \times \text{dens}_{\max}(\tau)]$$

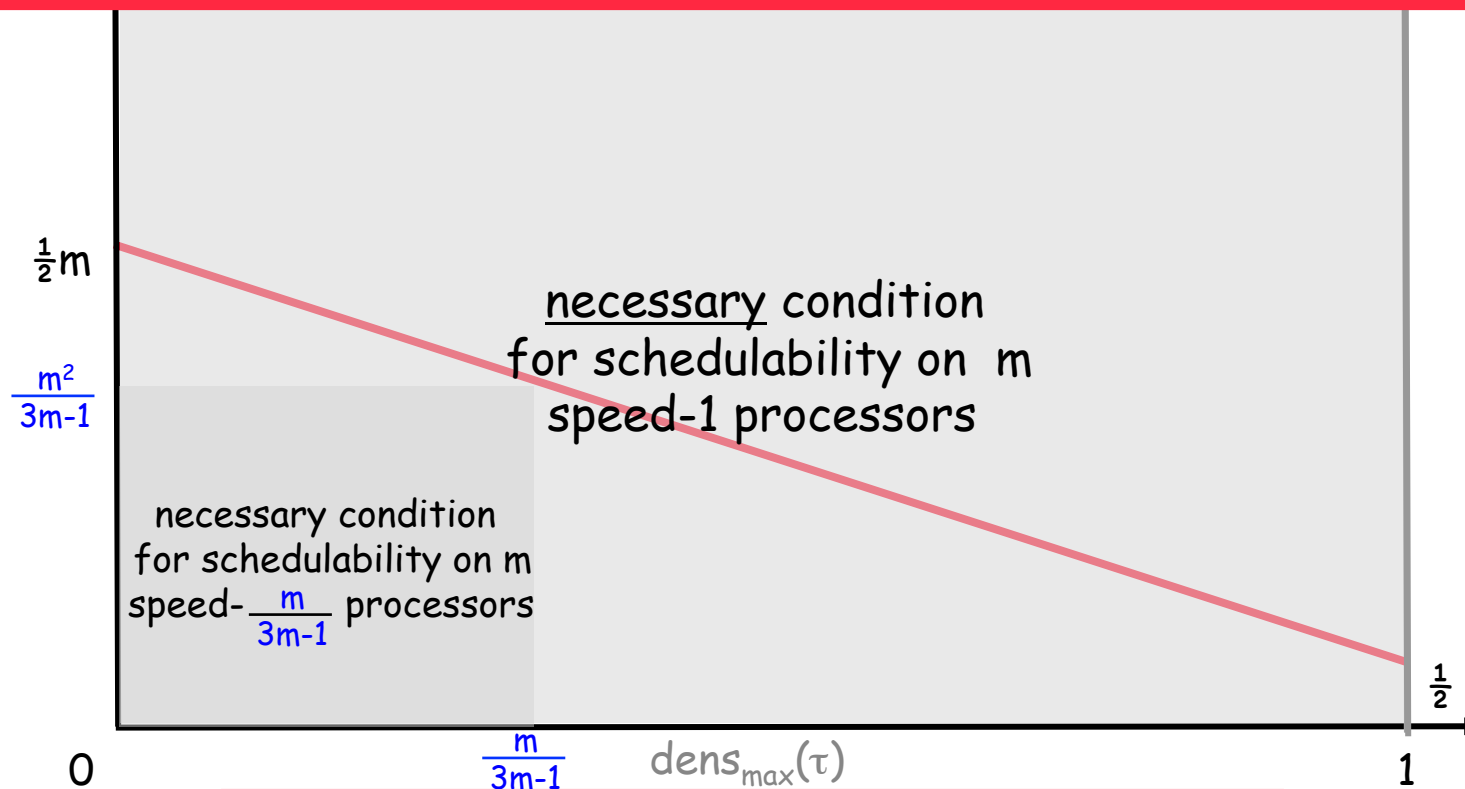
A partitioning algorithm



$$\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq \frac{1}{2} [m - (m-1) \times \text{dens}_{\max}(\tau)]$$

A partitioning algorithm

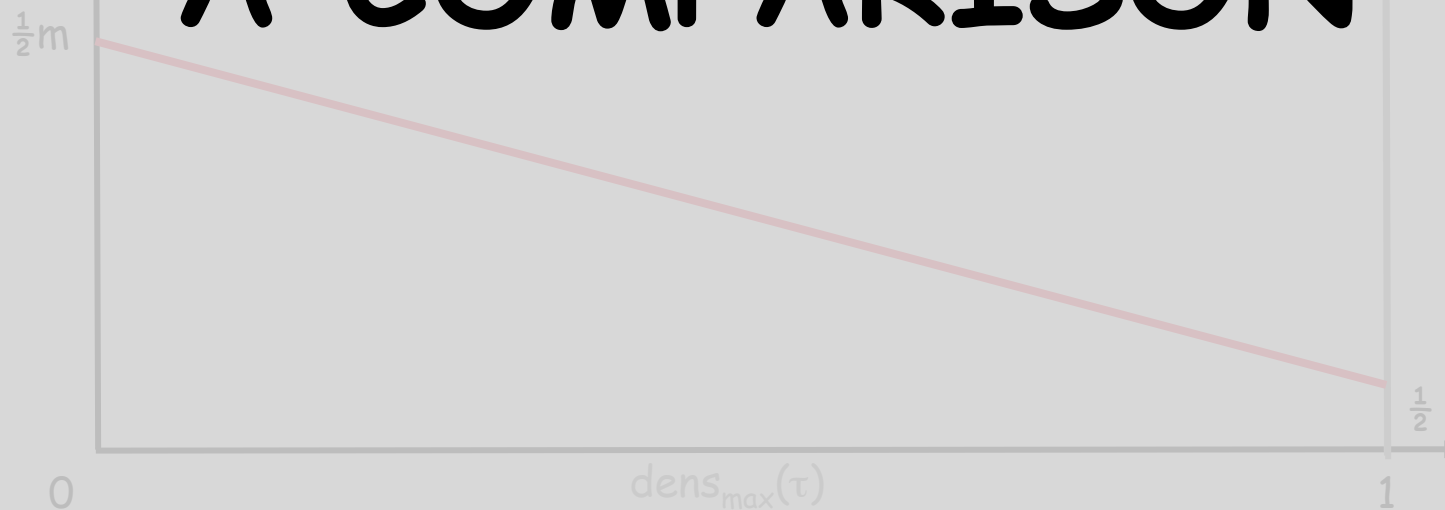
Any sporadic task system schedulable upon m processors is partitioned-EDF schedulable on m processors that are each $(3 - 1/m)$ times as fast



$$\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq \frac{1}{2} [m - (m-1) \times \text{dens}_{\max}(\tau)]$$

Any sporadic task system schedulable upon m processors is partitioned-EDF schedulable on m processors that are each $(3 - 1/m)$ times as fast

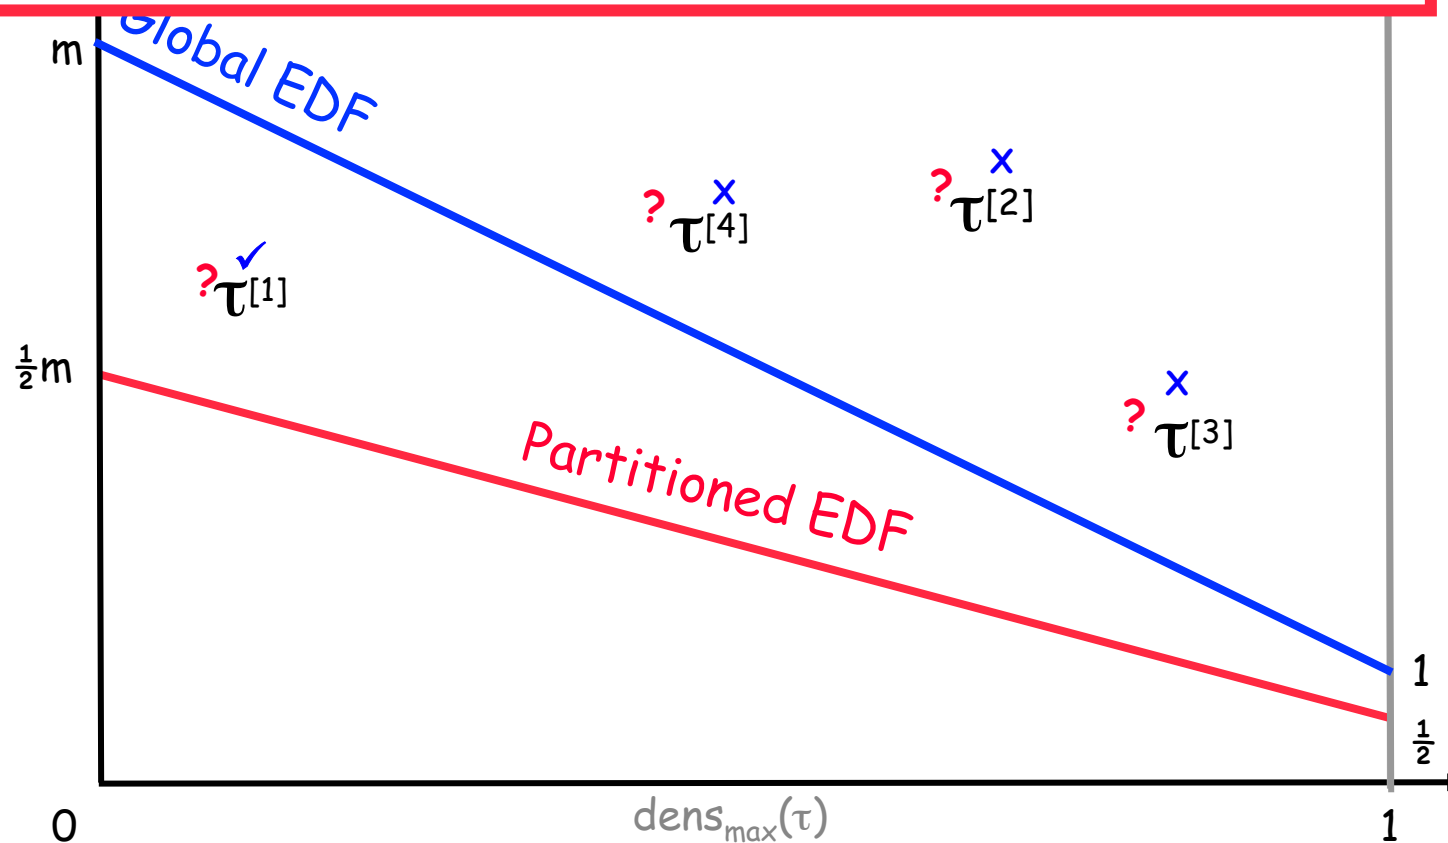
A COMPARISON



$$\text{load}(\tau, \text{dens}_{\max}(\tau)) \leq \frac{1}{2} [m - (m-1) \times \text{dens}_{\max}(\tau)]$$

Comparing Global and Partitioned EDF

BOTTOM LINE: Partitioned EDF and Global EDF offer comparable schedulability (to random tasksets)



Context and conclusions

Multiprocessor systems are increasingly important

⇒ need a theory of multiprocessor RT scheduling

Some breakthroughs recently...

- sporadic tasks on identical multiprocessors, scheduled using EDF
- other models, other algorithms

Multiprocessor RT scheduling theory today

≈

Uniprocessor RT scheduling theory in mid-1980's

⇒ significant progress soon??