

# An Asymptotically Optimal Real-Time Locking Protocol for Clustered Scheduling under Suspension-Aware Analysis

12/5/2012  
RTSS'12 WIP



Max  
Planck  
Institute  
for  
Software Systems

Björn B. Brandenburg  
[bbb@mpi-sws.org](mailto:bbb@mpi-sws.org)

# Background: Suspension-Based Multiprocessor Real-Time Locking

<p><b>Blocking Optimality</b> [— &amp; Anderson, 2010]</p>	<p><b>suspension oblivious</b></p>	<p><b>suspension aware</b></p>
<p>How are suspensions analyzed?</p>	<p>CPU demand is <b>over-approximated</b></p>	<p>CPU demand is modeled <b>accurately</b></p>
<p>Lower bound on <u>maximum priority inversion blocking</u> <math>\max_i \{B_i\}</math></p>	<p><math>\Omega(m)</math> <math>m = \#CPUs</math></p>	<p><math>\Omega(n)</math> <math>n = \#tasks</math></p>

[— & Anderson, 2010] *Optimality Results for Multiprocessor Real-Time Locking*, RTSS 2010.

# Asymptotically Optimal Locking Protocols

<p>JLFP = job-level fixed-priority</p>	<p><u>Suspension Oblivious</u> Any JLFP Scheduler</p>	<p><u>Suspension Aware</u> EDF w/ Implicit Deadlines</p>	<p><u>Suspension Aware</u> Any JLFP Scheduler</p>
<p><b>Partitioned</b> (no migrations)</p>			
<p><b>Global</b> (jobs migrate freely)</p>			
<p><b>Clustered</b> (jobs migrate only among subset of processors)</p>			

[Block et al., 2007] *A Flexible Real-Time Locking Protocol for Multiprocessors*, RTCSA 2007.  
 [— & Anderson, 2010] *Optimality Results for Multiprocessor Real-Time Locking*, RTSS 2010.  
 [— & Anderson, 2011] *Real-Time Resource-Sharing under Clustered Scheduling: Mutex, Reader-Writer, and k-Exclusion Locks*, EMSOFT 2011.  
 [—, 2011] *Scheduling and Locking in Multiprocessor Real-Time Operating Systems*, PhD thesis, UNC, 2011.

# Asymptotically Optimal Locking Protocols

JLFP = job-level fixed-priority	<u>Suspension Oblivious</u> Any JLFP Scheduler	<u>Suspension Aware</u> EDF w/ Implicit Deadlines	<u>Suspension Aware</u> Any JLFP Scheduler
<b>Partitioned</b> (no migrations)	<b>P-OMLP</b> ✓ [— & Anderson, 2010]		
<b>Global</b> (jobs migrate freely)	<b>G-OMLP</b> ✓ [— & Anderson, 2010]		
<b>Clustered</b> (jobs migrate only among subset of processors)	<b>C-OMLP</b> ✓ [— & Anderson, 2011]		

[Block et al., 2007]

*A Flexible Real-Time Locking Protocol for Multiprocessors*, RTCSA 2007.

[— & Anderson, 2010]

*Optimality Results for Multiprocessor Real-Time Locking*, RTSS 2010.

[— & Anderson, 2011]

*Real-Time Resource-Sharing under Clustered Scheduling: Mutex, Reader-Writer, and k-Exclusion Locks*, EMSOFT 2011.

[—, 2011]

*Scheduling and Locking in Multiprocessor Real-Time Operating Systems*, PhD thesis, UNC, 2011.

# Asymptotically Optimal Locking Protocols

JLFP = job-level fixed-priority	<u>Suspension Oblivious</u> Any JLFP Scheduler	<u>Suspension Aware</u> EDF w/ Implicit Deadlines	<u>Suspension Aware</u> Any JLFP Scheduler
<b>Partitioned</b> (no migrations)	<b>P-OMLP</b> ✓ [— & Anderson, 2010]	<b>SPFP</b> (asymptotical tightness) ✓ [— & Anderson, 2010] <b>FMLP+</b> (practical protocol) ✓ [—, 2011]	
<b>Global</b> (jobs migrate freely)	<b>G-OMLP</b> ✓ [— & Anderson, 2010]		
<b>Clustered</b> (jobs migrate only among subset of processors)	<b>C-OMLP</b> ✓ [— & Anderson, 2011]		

[Block et al., 2007]

[— & Anderson, 2010]

[— & Anderson, 2011]

[—, 2011]

*A Flexible Real-Time Locking Protocol for Multiprocessors*, RTCSA 2007.

*Optimality Results for Multiprocessor Real-Time Locking*, RTSS 2010.

*Real-Time Resource-Sharing under Clustered Scheduling: Mutex, Reader-Writer, and k-Exclusion Locks*, EMSOFT 2011.

*Scheduling and Locking in Multiprocessor Real-Time Operating Systems*, PhD thesis, UNC, 2011.

# Asymptotically Optimal Locking Protocols

JLFP = job-level fixed-priority	<u>Suspension Oblivious</u> Any JLFP Scheduler	<u>Suspension Aware</u> EDF w/ Implicit Deadlines	<u>Suspension Aware</u> Any JLFP Scheduler
<b>Partitioned</b> (no migrations)	<b>P-OMLP</b> ✓ [— & Anderson, 2010]	<b>SPFP</b> (asymptotical tightness) ✓ [— & Anderson, 2010] <b>FMLP+</b> (practical protocol) ✓ [—, 2011]	
<b>Global</b> (jobs migrate freely)	<b>G-OMLP</b> ✓ [— & Anderson, 2010]	<b>FMLP</b> ⚠ [Block et al., 2007]	?
<b>Clustered</b> (jobs migrate only among subset of processors)	<b>C-OMLP</b> ✓ [— & Anderson, 2011]	?	?

[Block et al., 2007]

*A Flexible Real-Time Locking Protocol for Multiprocessors*, RTCSA 2007.

[— & Anderson, 2010]

*Optimality Results for Multiprocessor Real-Time Locking*, RTSS 2010.


[— & Anderson, 2011]

*Real-Time Resource-Sharing under Clustered Scheduling: Mutex, Reader-Writer, and k-Exclusion Locks*, EMSOFT 2011.

[—, 2011]

*Scheduling and Locking in Multiprocessor Real-Time Operating Systems*, PhD thesis, UNC, 2011.

# Asymptotically Optimal Locking Protocols

JLFP = job-level fixed-priority	<u>Suspension Oblivious</u> Any JLFP Scheduler	<u>Suspension Aware</u> Any JLFP Scheduler
<b>Partitioned</b> (no migrations)	<b>P-OMLP</b> ✓ [— & Anderson, 2010]	<h2>This Work</h2> <h3>The Generalized FMLP+</h3> 
<b>Global</b> (jobs migrate freely)	<b>G-OMLP</b> ✓ [— & Anderson, 2010]	
<b>Clustered</b> (jobs migrate only among subset of processors)	<b>C-OMLP</b> ✓ [— & Anderson, 2011]	

[Block et al., 2007]

[— & Anderson, 2010]

[— & Anderson, 2011]

[—, 2011]

*A Flexible Real-Time Locking Protocol for Multiprocessors*, RTCSA 2007.

*Optimality Results for Multiprocessor Real-Time Locking*, RTSS 2010.

*Real-Time Resource-Sharing under Clustered Scheduling: Mutex, Reader-Writer, and k-Exclusion Locks*, EMSOFT 2011.

*Scheduling and Locking in Multiprocessor Real-Time Operating Systems*, PhD thesis, UNC, 2011.

# The Generalized **FMLP**<sup>+</sup>

**F**IFO **M**ultiprocessor **L**ocking **P**rotocol

## The Goal

- $\max_i \{B_i\} = O(n)$  maximum priority inversion blocking
  - $n = \#tasks$
- Use a **F**IFO queue!

## The Problem [—, 2011]

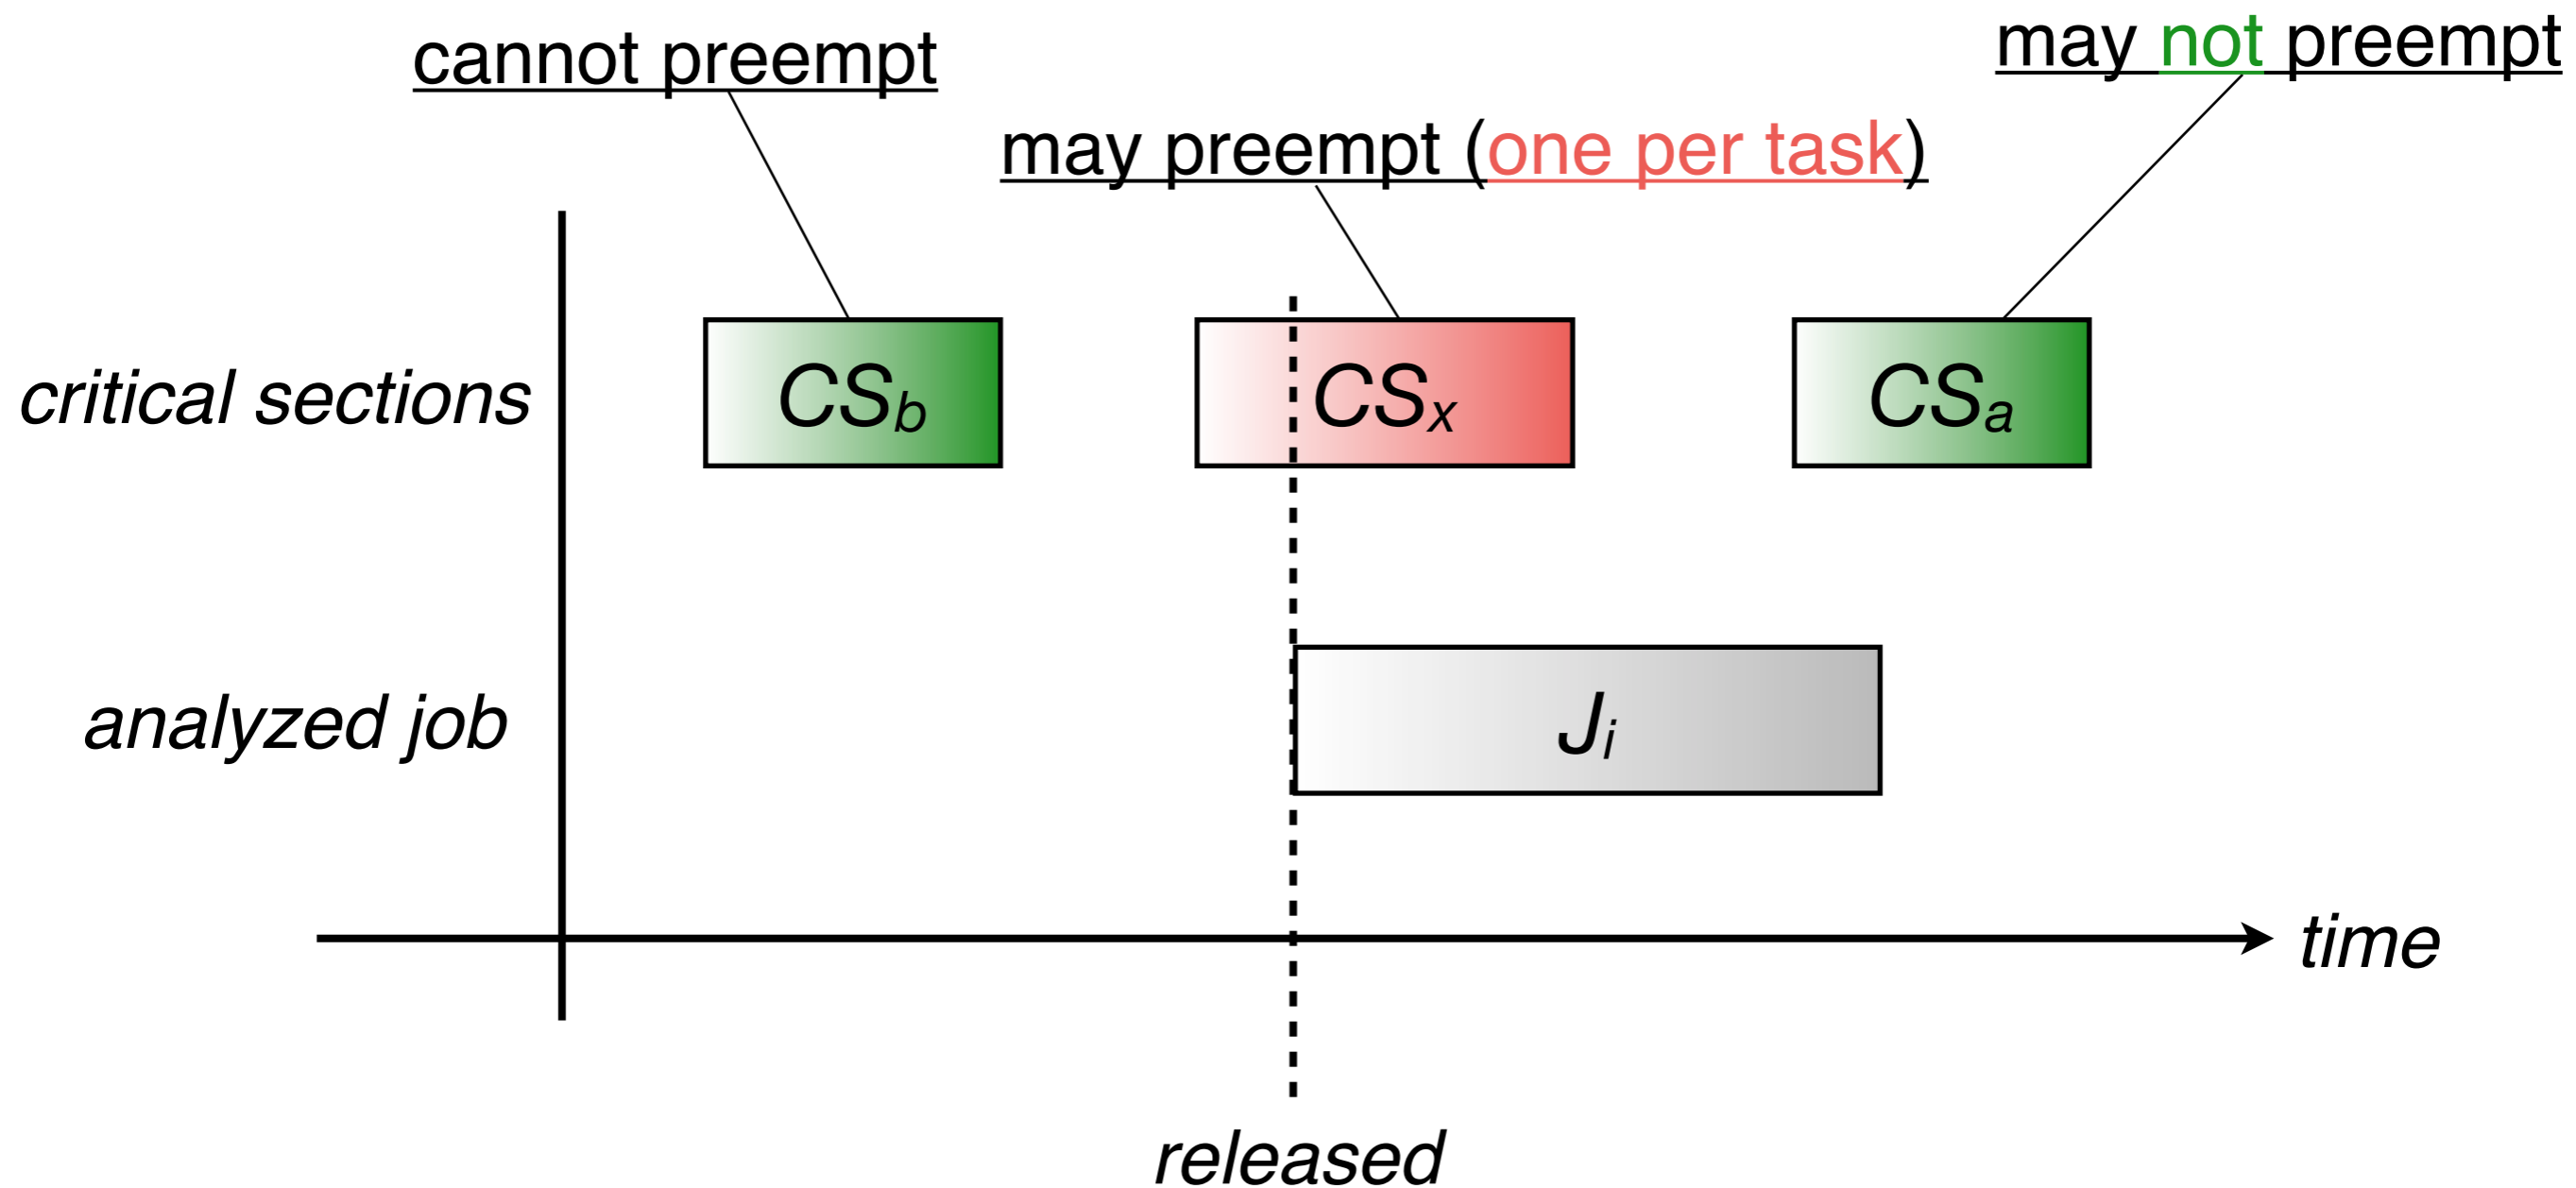
- Priority inheritance leads to  $\Omega(\phi)$  blocking.
  - $\phi = \text{ratio largest to shortest period}$ , unbounded in general
- Priority boosting also leads to  $\Omega(\phi)$  blocking...

[—, 2011] *Scheduling and Locking in Multiprocessor Real-Time Operating Systems*, PhD thesis, UNC, 2011.



# New Solution: **FIFO Boosting**

*prioritize by order of lock request & release times*



→  $O(n)$  preemptions

# Thanks!

*I'll be happy to answer your questions on Friday...*