

Uniprocessor Mixed-Criticality Scheduling with Graceful Degradation by Completion Rate

Zhishan Guo¹, Kecheng Yang², **Sudharsan Vaidhun**¹,
Samsil Arefin³, Sajal K. Das³, Haoyi Xiong⁴

¹Department of Electrical & Computer Engineering, University of Central Florida

²Department of Computer Science, Texas State University

³Department of Computer Science, Missouri University of Science & Technology

⁴Baidu Inc., Beijing, China

December 14, 2018

Mixed-Criticality (MC) Systems

Safety certification requirement

All safety critical tasks must **complete execution** within its **deadline**.


Non-safety critical tasks need not be guaranteed for safety certification, but rather by **mission certification**.

In general, different certification levels certify different levels of tasks, with safety critical being the highest.

Mixed-Critical (MC) Task Model

- System model of LO-criticality tasks:

$$\tau_i = \{C_i; T_i; D_i; \dots\}$$




Mixed-Critical (MC) Task Model

- System model of lo-criticality tasks:

$$i = f C_i; T_i; D_i; i g; 8 i 2 lo;$$


- System model of hi-criticality tasks:

$$i = f C_i^{lo}; C_i^{hi}; T_i; D_i; i g; 8 i 2 hi;$$



EDF-VD Algorithm

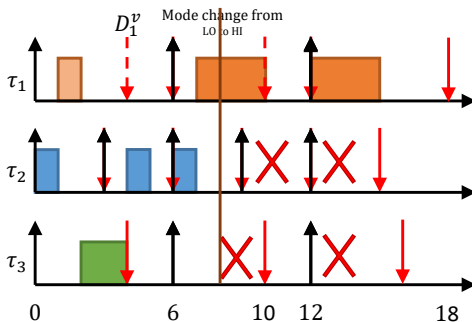
Earliest Deadline First with Virtual Deadline (EDF-VD) algorithm¹



¹S. Baruah et al. "The preemptive uniprocessor scheduling of mixed-criticality implicit-deadline sporadic task systems". In: *Proceedings of the 24th Euromicro Conference on Real-Time Systems*. 2012.

Example Schedule

Task ID	i	$C_i (C_i^{lo})$	C_i^{hi}	T_i	D_i	D_i^v
1	hi	1	3	6	6	4
2	lo	1	-	3	3	-
3	lo	2	-	6	4	-



Graceful Degradation

LO-criticality tasks are **not non-critical**.

Graceful Degradation

LO-criticality tasks are **not non-critical**.

Scenario 1 : In control systems, a reduced rate of sensing can still produce stable control.


Scenario 2 : In signal processing, losing few data packets can still produce decent sound/video.

Providing **degraded yet guaranteed service** is better than providing no guarantees.

Limitations of Existing Works on Graceful Degradation for Mixed Criticality

Utilization based degradation:²

- Reduced execution time - imprecise computation model

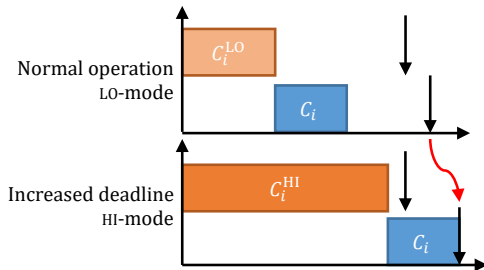


²Di Liu et al. "EDF-VD Scheduling of Mixed-Criticality Systems with Degraded Quality Guarantees". In: *Proceedings of the 37th IEEE Real-Time Systems Symposium*. 2016.

Limitations of Existing Works on Graceful Degradation for Mixed Criticality

Utilization based degradation:²

- Reduced execution time - imprecise computation model
- Increased deadline - delayed responses based on longer period



²Liu et al., "EDF-VD Scheduling of Mixed-Criticality Systems with Degraded Quality Guarantees".

Limitations of Existing Works on Graceful Degradation for Mixed Criticality

Asymptotic rate guarantees:^{3;4}

- No guarantee for short interval of time
- Example:
 - 1-out-of-4 instances scheduled and 1000-out-of-4000 instances scheduled the same percent of tasks (25%) **asymptotically**.
 - However, in the second case 3000 consecutive instances can be dropped.

³M. S. Branicky, S. M. Phillips, and Wei Zhang. "Scheduling and feedback co-design for networked control systems". In: *Proceedings of the 41st IEEE Conference on Decision and Control*. 2002.

⁴Rupak Majumdar, Indranil Saha, and Majid Zamani. "Performance-aware Scheduler Synthesis for Control Systems". In: *Proceedings of the 9th ACM International Conference on Embedded Software*. ACM, 2011.

Graceful Degradation


Goals:

- Maintain limited functionality even under failure
- Provide some **guarantees** to LO tasks even in HI mode

Graceful Degradation

Goals:

- Maintain limited functionality even under failure
- Provide some **guarantees** to LO tasks even in HI mode



Guaranteed Graceful Degradation to I/O Criticality Tasks

System model:

$$\forall i \in \text{Io}: i = f(C_i; T_i; D_i; r_i; g_i)$$

$$\forall i \in \text{hi}: i = f(C_i^{\text{lo}}; C_i^{\text{hi}}; T_i; D_i; g_i)$$

r_i is the cumulative completion rate

Guaranteed Graceful Degradation to I O Criticality Tasks

System model:

$$\forall i \in \text{IO}: i = fC_i; T_i; D_i; r_i; ig;$$

$$\forall i \in \text{hi}: i = fC_i^{\text{lo}}; C_i^{\text{hi}}; T_i; D_i; ig;$$

r_i is the cumulative completion rate

Correctness criteria:

- In I O mode, all tasks will receive an execution time of C_i^{lo}

Guaranteed Graceful Degradation to I O Criticality Tasks

System model:

$$\forall i \in \text{Io}: i = f(C_i; T_i; D_i; r_i; ig_i)$$


$$\forall i \in \text{hi}: i = f(C_i^{\text{lo}}; C_i^{\text{hi}}; T_i; D_i; ig_i)$$

r_i is the cumulative completion rate

Correctness criteria:

- In I O mode, all tasks will receive an execution time of C_i^{lo}
- In hi mode, for any $N \in \mathbb{Z}^+$, **at least dr_i Ne-out-of- N I O jobs** will be guaranteed to receive a full execution time of C_i ; hi jobs will receive an execution time of C_i^{hi}


EDF-GVD Algorithm - I O-mode



Adapted from Earliest Deadline First with Virtual Deadline (EDF-VD) algorithm⁵.

⁵Baruah et al., "The preemptive uniprocessor scheduling of mixed-criticality implicit-deadline sporadic task systems".

EDF-GVD Algorithm - I O-mode



Adapted from Earliest Deadline First with Virtual Deadline (EDF-VD) algorithm⁵.


Virtual deadline setting

A simple heuristic is used to test the schedulability

$$D_i^v = \frac{C_i^{LO}}{C_i^{HI}} D_i; \quad \forall i \in \text{hi}$$

⁵Baruah et al., "The preemptive uniprocessor scheduling of mixed-criticality implicit-deadline sporadic task systems".

EDF-GVD Algorithm - I O-mode



Adapted from Earliest Deadline First with Virtual Deadline (EDF-VD) algorithm⁵.

Virtual deadline setting

If simple heuristic does not work, $q \in [0;1]$ is determined using binary search up to an error margin

$$D_i^v = q D_i; \quad \forall i \in \text{hi}$$


⁵Baruah et al., "The preemptive uniprocessor scheduling of mixed-criticality implicit-deadline sporadic task systems".

EDF-GVD Algorithm - hi mode

If any hi-criticality task does not signal completion by its C_i^{lo} , there is a **system wide mode switch**.

Example Schedule

Task ID	i	$C_i (C_i^{\text{lo}})$	C_i^{hi}	T_i	D_i	D_i^{v}	r_i	i
1	hi	1	3	6	6	4	-	-
2	lo	1	-	3	3	-	0.5	$f10g^1$
3	lo	2	-	6	4	-	0.4	$f10100g^1$



EDF-GVD Algorithm - hi mode

If any hi-criticality task does not signal completion by its C_i^{lo} , there is a **system wide mode switch**.

- hi-criticality tasks are scheduled for C_i^{hi} execution time.

EDF-GVD Algorithm - hi mode

If any hi-criticality task does not signal completion by its C_i^{lo} , there is a **system wide mode switch**.


- hi-criticality tasks are scheduled for C_i^{hi} execution time.
- Lo-criticality tasks are scheduled based on an **admission control** procedure.
 - 1 Static admission control for rational values of r_i
 - 2 Dynamic admission control for irrational values of r_i

Static Admission Control

$$r_i = \frac{p}{q} \quad \text{pattern of } q \text{ length, with } p \text{ no. of '1's}$$

$$r_i = 0.4 = \frac{2}{5}, \quad i = f10100g^7$$

$$\text{actual } r_i = 1, \frac{1}{2}, \frac{2}{3}, \frac{1}{2}, \frac{2}{5}, \frac{1}{2}, \frac{3}{7}, \frac{1}{2}, \frac{4}{9}, \frac{2}{5},$$




Static Admission Control

$$r_i = \frac{p}{q} \quad \text{pattern of } q \text{ length, with } p \text{ no. of '1's'}$$

$$r_i = 0.4 = \frac{2}{5}, \quad i = f10100g^1$$

$$\text{actual } r_i = 1, \frac{1}{2}, \frac{2}{3}, \frac{1}{2}, \frac{2}{5}, \frac{1}{2}, \frac{3}{7}, \frac{1}{2}, \frac{4}{9}, \frac{2}{5}, \dots$$



Use cases:

- 1** Rational r_i value
- 2** Offline calculation

Dynamic Admission Control

Control the I O-tasks admission in hi mode **dynamically**

$$r_i = 1 = \rho_{\bar{2}} = 0.7071 \quad , \quad i = 11101101110$$

$$\text{actual } r_i = 1; 1; 1; \frac{3}{4}; \frac{4}{5}; \frac{5}{6}; \frac{5}{7}; \frac{3}{4}; \frac{7}{9}; \frac{4}{5}; \frac{8}{11};$$

Dynamic Admission Control

Control the low-criticality tasks admission in mode **dynamically**

$$r_i = 1 = \frac{p}{2} \quad \bar{r}_i = 0.7071 \quad , \quad \bar{c}_i = 11101101110$$

actual r_i | 1; 1; 1; $\frac{3}{4}$; $\frac{4}{5}$; $\frac{5}{6}$; $\frac{5}{7}$; $\frac{3}{4}$; $\frac{7}{9}$; $\frac{4}{5}$; $\frac{8}{11}$;

Use cases:

- 1 Irrational r_i value
- 2 Low memory - pattern need not be stored

Admission Control Patterns

S = Maximum number of consecutive drops

$$S(r_i) = \frac{1}{r_i} \quad 1:$$

Task ID	r_i	Pattern $_i$	S_i
1	0.4	f 10100g ¹	2
2	0.5	f 10g ¹	1
3	0.625	f 11011010g ¹	1
4	$1 = \frac{1}{2}$	11101101110	1

Provides the system designer control over maximum number of consecutive drops.

Backward Mode Switch

Task ID	i	$C_i (C_i^{lo})$	C_i^{hi}	T_i	D_i	D_i^y	r_i	i
1	hi	1	3	6	6	4	-	-
2	lo	1	-	3	3	-	0:5	f 10g ¹
3	lo	2	-	6	4	-	0:4	f 10100g ¹

Backward Mode Switch

Task ID	i	C_i (C_i^{lo})	C_i^{hi}	T_i	D_i	D_i^v	r_i	i
1	hi	1	3	6	6	4	-	-
2	lo	1	-	3	3	-	0:5	$f 10g^1$
3	lo	2	-	6	4	-	0:4	$f 10100g^1$

Backward Mode Switch

System switches back to lo mode when:

- 1 the processor idles in hi-criticality mode, and
- 2 the last instance of each hi-criticality task is finished within its lo -WCET

Backward Mode Switch

System switches back to lo mode when:

- 1 the processor idles in hi-criticality mode, and
- 2 the last instance of each hi-criticality task is finished within its lo -WCET

The maximum time after the last hi-critical task instance to return to lo -mode is

$$L_1 = \frac{\sum_{i \in \mathcal{P}_{hi}} (2C_i^{lo}) + \sum_{i \in \mathcal{P}_{lo}} (C_i + 2r_i C_i)}{\sum_{i \in \mathcal{P}_{hi}} (C_i^{lo} = T_i) + \sum_{i \in \mathcal{P}_{lo}} (r_i C_i = T_i)}$$

Schedulability Analysis

Demand bound function

$$g_i(t) = 0; \quad \text{dbf}_i(t) = \sum_{k=1}^{\lfloor t/T_i \rfloor} X_i$$

Schedulability Analysis

Demand bound function

$$g_i^B(0; \tau_i) = \sum_{i=2}^X dbf(i; \tau_i)$$

$dbf_A^B(i; \tau_i)$: dbf of A-criticality tasks in B-criticality mode

$dbf_{lo}^{do}(i; \tau_i)$; $dbf_{hi}^{do}(i; \tau_i)$; $dbf_{lo}^{hi}(i; \tau_i)$; $dbf_{hi}^{hi}(i; \tau_i)$

Schedulability Test A - ~~folo~~ -mode

Condition A:

$$\begin{aligned}
 & \sum_{i=0}^{\infty} \left(\sum_{j=2}^{\infty} \text{dbf}_{lo}^j(i; \tau) \right) \\
 & + \sum_{i=2}^{\infty} \left(\sum_{j=2}^{\infty} \text{dbf}_{hi}^j(i; \tau) \right) \leq \tau
 \end{aligned}$$

⁶Pontus Ekberg and Wang Yi. "Bounding and shaping the demand of generalized mixed-criticality sporadic task systems". In: *Real-Time Systems* 50 (2014), pp. 48{86.

Schedulability Test A - for I O-mode

Condition A:

$$g_i \leq 0; \quad \sum_{i \in \text{lo}} \text{dbf}_{i,0}^{\text{lo}}(\tau_i) + \sum_{i \in \text{hi}} \text{dbf}_{i,0}^{\text{hi}}(\tau_i) \leq c$$

Lemma 3: If $\sum_{i \in \text{lo}} (C_i = T_i) + \sum_{i \in \text{hi}} (C_i^{\text{lo}} = T_i) \leq c$ and $c < 1$
Then, Condition A⁶ holds, if it is true for all τ such that

$$\tau_i < \frac{c}{1-c} \max_{i \in \text{lo}} f T_i \quad D_i; \max_{i \in \text{hi}} f T_i \quad D_i^{\text{lo}}; g_i$$

⁶Ekberg and Yi, "Bounding and shaping the demand of generalized mixed-criticality sporadic task systems".

Schedulability Test B - for hi-mode

Condition B:

$$g \cdot 0; \quad \times \quad \text{dbf}_{l_0}^{\text{hi}}(i; \tau) \\ + \quad \times \quad \text{dbf}_{\text{hi}}^{\text{hi}}(i; \tau) \quad \tau;$$

Lemma 4: If $\prod_{i \in l_0} (r_i \cdot C_i = T_i) \leq c_1$, $\prod_{i \in \text{hi}} (C_i^{\text{hi}} = T_i) \leq c_2$, $c_1 < 1$, $c_2 < 1$, and $c_1 + c_2 < 1$.

Then, Condition B holds, if it is true for all τ such that

$$\tau < \frac{c_1 \max_{i \in l_0 \wedge r_i > 0} fT_i \cdot D_i + \frac{T_i}{r_i} g + c_2 \max_{i \in \text{hi}} fT_i \cdot D_i + D_i^y g}{1 - c_1 - c_2};$$

Experiment via Simulation

Task set generation:


- $P_{hi} = 0.5$: probability of hi-tasks.
- $C_i^{lo} \sim U[1;10]$
- $R_{hi} = 4$: ratio of C_i^{hi} to C_i^{lo}
- $T_{max} = 200$: $T_i \sim U[C_i^{lo}; T_{max}]$.
- $minDR \sim U[0.1;0.9]$: $D_i = \frac{C_i^{lo}}{T_i}$; where D_i is uniformly generated from the range $[minDR; 1]$.

Tasks are generated in a way so that each task set have a targeted average utilization U_{avg} .

Simulation results

Variation of acceptance ratio with average utilization

- Constrained deadline tasks - deadline randomly chosen in the range $[0.5; 1]$ period.
- 1000 tasksets each.

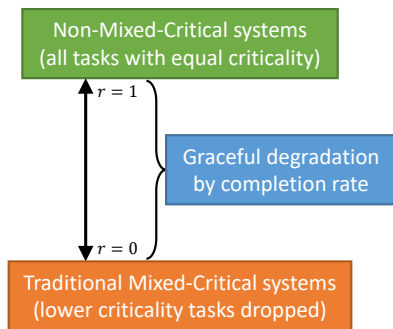


Conclusions

- **Stronger correctness criteria** - guaranteed service rate to non-safety-critical tasks in hi mode.

Conclusions

- **Stronger correctness criteria** - guaranteed service rate to non-safety-critical tasks in hi mode.
- Bridges mixed-critical scheduling and non-mixed-critical scheduling



Conclusions

- **Stronger correctness criteria** - guaranteed service rate to non-safety-critical tasks in hi mode.
- Bridges mixed-critical scheduling and non-mixed-critical scheduling
- The variable S_j is provided to aid the system designers in choosing the r_j value based on the tolerance on consecutive dropped jobs.

Thank You

Thank you for your time!