

Empirical Study of Energy Minimization Issues for Mixed-Criticality Systems with Reliability Constraints

Zheng Li, Xiayu Hua, Chunhui Guo and Shangping Ren
Email: {zli80, xhua, cguo13}@hawk.iit.edu, ren@iit.edu

Illinois Institute of Technology

November 3, 2014

Outline

- 1 Introduction
- 2 System Model
- 3 Problem and Analysis
- 4 Algorithm
- 5 Evaluation
- 6 Conclusion

Mixed-Criticality (MC) System

To reduce size, weight, power consumption and cost, tasks at different criticality levels share the same platform, e.g., Unmanned Aerial Vehicle (UAV).

- HI-criticality task: flight control
- LO-criticality task: photo capturing

MC Research Issues

- Schedulability: EDF-VD [Baruah, 2012], demand bound function analysis [Ekberg, 2012]
- Reliability: recovery [Zhao, 2012]
- Energy efficiency: Dynamic Voltage and Frequency Scaling (DVFS) [Zhu, 2004]

Processor Model

Processor: single DVFS-enabled processor

Available frequencies: $F_a = \{f_1, \dots, f_q\}$ in descending order where $f_{\max} = f_1 = 1, f_{\min} = f_q$

Task Model

Mixed-criticality task: $\tau_i = (C_i, T_i, L_i)$

- $C_i = \{C_i(LO), C_i(HI)\}$: worst-case execution times under f_{\max}
- T_i : period, which is equal to deadline
- $L_i = \{LO, HI\}$: criticality level

HI-criticality task: $C_i(LO) \leq C_i(HI)$

LO-criticality task: $C_i(LO) = C_i(HI)$

System Model

Fault Recovery Model

Fault detection: at end of job execution, cost is counted in C_i

Backward recovery: re-executed under f_{\max} , deadline is same with the corresponding job

Reliability Model

Transient fault rate: $\lambda(f_i) = \hat{\lambda}_0 10^{-\hat{\alpha}f_i}$

Job-level reliability: $R_i(f_i) = e^{-\lambda(f_i) \cdot \frac{C_i}{T_i}}$

Task-level reliability within a hyper-period H :

$\Phi_i(f_i, r_i, k_i) = R_i(f_i)_i^k + \sum_{j=1}^{r_i} \binom{k_i}{j} (1 - R_i(f_i))^j R_i(f_i)^{k_i-j}$ where r_i is the recovery number and $k_i = H/T_i$ is the number of jobs

Energy Model

Energy consumption of a job: $E(f_i, C_i) = P_{\text{ind}} \frac{C_i}{f_i} + C_{\text{ef}} C_i f_i^{\theta-1}$

Expected energy consumption of recovery: ignored, as probability of fault occurrence is very small (e.g., $\leq 10^{-6}$)

Expected energy consumption of the system within a hyper-period H :

$EC(F) = \sum_{\tau_i \in \Gamma} (H/T_i) \cdot E(f_i, C_i)$ where $F = \{f_1, \dots, f_i, \dots, f_{|\Gamma|}\}$ is the

frequency assignment of task set Γ

Assumption: all jobs of the same task have the same execution frequency

System Model

- The system has two running modes, i.e., LO-mode and HI-mode, and the system runs at the LO-mode initially.
- In the LO-mode, each task τ_i can run up to $\frac{C_i(LO)}{f_i}$ within each period. If any task τ_i executes beyond $\frac{C_i(LO)}{f_i}$ in a period, the system enters into the HI-mode.
- In HI-mode, all LO-criticality tasks are removed from the system and every HI-criticality task τ_i 's maximum execution time is $\frac{C_i(HI)}{f_{\max}}$.
- It is worth pointing out that fault recovery time is not counted toward task execution time for triggering system mode switch.

Problem Formulation

Problem Formulation

Given a DVFS-enabled processor with q different processing frequencies $F_a = \{f_1, \dots, f_q\}$ and a mixed-criticality task set $\Gamma = \{\Gamma_{LO}, \Gamma_{HI}\}$, develop an algorithm to **assign an execution frequency** to every task so that the system's expected energy consumption is minimized and the system's schedulability constraint and reliability requirement are guaranteed.

Reliability requirement: $\forall 1 \leq i \leq |\Gamma| \quad \Phi_i(f_i, r_i, k_i) \geq \Phi_i(f_{max}, 0, k_i)$

Assumption: task execution frequency does NOT change once it is assigned.

Strategy

- LO-mode: tasks run at lower frequencies
- HI-mode: all HI-criticality tasks run at f_{max}

Energy Consumption Analysis

- Lower frequency reduces system's energy consumption.
- Lower frequency results in longer execution time and lower reliability, the system needs more recoveries to maintain reliability requirement. Both longer execution time and more recoveries negatively impact schedulability.
- There is a tradeoff between energy consumption and schedulability and reliability constraints.

- In the LO-mode, at least r_i recoveries are needed to maintain reliability requirement under frequency f_i . Given the reliability requirement, we can calculate r_i .
- In the HI-mode, as the system runs under f_{\max} , no recovery is needed.

Schedulability Analysis

Schedulability Condition

$$\forall l \in [0, H] : \sum_{\tau_i \in \Gamma} \text{dbf}_{\text{LO}}(\tau_i, l) \leq l$$

$$\forall l \in [0, H] : \sum_{\tau_i \in \Gamma_H} \text{dbf}_{\text{HI}}(\tau_i, l) \leq l$$

Worst Case

r_i recoveries take place in the first r_i jobs of τ_i

Virtual Deadline

HI-criticality task: $\text{VD}_i \leq T_i$

LO-criticality task: $\text{VD}_i = T_i$

Schedulability Analysis




Figure: LO-mode Job Execution

LO-mode Demand Bound Function

$$\text{dbf}_{\text{LO}}(\tau_i, l, f_i) = \begin{cases} \left[\left(\left\lfloor \frac{l - \text{VD}_i}{T_i} \right\rfloor + 1 \right) \cdot \left(1 + \frac{f_{\max}}{f_i} \right) \cdot C_i(\text{LO}) \right]_0 & \text{if } l \leq r_i \cdot T_i \\ r_i \cdot \left(1 + \frac{f_{\max}}{f_i} \right) \cdot C_i(\text{LO}) + \left[\frac{f_{\max}}{f_i} \cdot \left(\left\lfloor \frac{l - r_i \cdot T_i - \text{VD}_i}{T_i} \right\rfloor + 1 \right) \cdot C_i(\text{LO}) \right]_0 & \text{otherwise} \end{cases}$$

where $\lceil x \rceil_0 = \max\{x, 0\}$

Schedulability Analysis




Figure: Carry-over Job

HI-mode Demand Bound Function

$$dbf_{HI}(\tau_i, l) = \left[\left\lfloor \frac{l - (T_i - VD_i)}{T_i} \right\rfloor + 1 \right] \cdot C_i(HI)$$

Heuristic Search based Energy Minimization (HSEM) Algorithm

- 1 Initially, all the tasks are running under the f_{max} .
- 2 Heuristically select a task to scale down its frequency without violating the reliability and schedulability constraints.
- 3 Repeat step (2) until no task in step (2) is available.

Schedulability Test

Use GREEDY algorithm [Ekberg, 2012] to determine virtual deadlines and test schedulability.

Algorithm

Metric (Heuristic Criteria)

The ratio between energy consumption change and execution time demand change

$$ED(F_i, F'_i) = \frac{EC(F) - EC(F')}{\text{Gap}(F) - \text{Gap}(F')}$$

where $F_i = \{f_1, \dots, f_i, \dots, f_{|\Gamma|}\}$

$F'_i = \{f_1, \dots, f'_i, \dots, f_{|\Gamma|}\}$

$\text{Gap}(F) = \min_{l \in [0, H]} \{g(l) \mid g(l) \neq l\}$

$g(l) = l - \sum_{\tau_i \in \Gamma} \text{dbf}_{\text{LO}}(\tau_i, l)$




Figure: Demand Bound Function

Task Selection Policy (TSP)

$$ED = \begin{cases} \max_{i \in [1, |\Gamma|]} \{ED(F_i, F'_i) \mid ED(F_i, F'_i) < 0\} & \text{if } \exists ED(F_i, F'_i) < 0 \\ \max_{i \in [1, |\Gamma|]} \{ED(F_i, F'_i)\} & \text{otherwise} \end{cases}$$

ALGORITHM 1: HSEM ($\Gamma, F = \{f_{\max}, \dots, f_{\min}\}$)

```
1 int[ $|\Gamma|$ ]  $FI = \{0, \dots, 0\}$ ;  
2 while  $\exists i : FI[i] < |F|$  do  
3   | find  
   |  $I = \{i | \text{CHECK}(\Gamma, F, FI, i) == \text{TRUE} \wedge 0 \leq i \leq |\Gamma| - 1\}$   
4   | if  $I$  is not empty then  
5   |   | find  $\tau_k$  based on TSP  
6   |   |  $FI[k] = FI[k] + 1$ ;  
7   |   | end  
8   |   | else  
9   |   | break;  
10  |   | end  
11 end  
12 return  $FI$ ;
```

ALGORITHM 2: CHECK (Γ, F, FI, i)

```
1 int[ $|\Gamma|$ ] A={0,...,0};
2  $FI[i] = FI[i] + 1$ 
3 for ( $j = 0; j < |\Gamma|; i++$ ) do
4   |  $A[j] =$  Obtain from MRT with  $F[FI[j]]$  and MRT;
5 end
6 if ( $GREEDY(\Gamma, F, FI, A) == SUCCESS$ ) then
7   | return TRUE;
8 end
9 return FALSE;
```


Comparisons

- LUF (Largest Utilization First): chooses the task with **largest** HI-mode utilization to scale down its execution frequency as low as it does not violate the reliability and schedulability constraints.
- SUF (Smallest Utilization First): chooses the task with **smallest** HI-mode utilization to scale down its execution frequency as low as it does not violate the reliability and schedulability constraints.


Performance Evaluation

Experiment Setting

3 LO-criticality tasks, 3 HI-criticality tasks



(a) $U_L(\Gamma_L)$ with $U_H(\Gamma_H) = 0.3$



(b) $U_H(\Gamma_H)$ with $U_L(\Gamma_L) = 0.3$

Figure: Normalized Energy Consumption (%)

Conclusion

- Analyze the resource demand of a mixed-criticality task set with both deadline and reliability constraints under a given frequency assignment.
- Propose HSEM algorithm to find a frequency assignment which can minimize system's energy consumption without violating system's reliability and schedulability constraints.
- Compared with LUF and SUF, the HSEM algorithm performs better with respect to energy saving.
- Future work is to extend our developed demand analysis and algorithm to address the problem that recovery is needed even under f_{\max} .

Thank You