Max Planck Institute for Software Systems

UPPSALA UNIVERSITET

Technische Universität Kaiserslautern

Quantifying the Effect of Period Ratios on Schedulability of Rate Monotonic

Mitra Nasri* Morteza Mohaqeqi Gerhard Fohler

RTNS, October 2016

Motivation

- Since 1973 it is known that <u>period ratio</u> affects the RM Schedulability

The hardest-to-schedule task set [Liu and Layland]

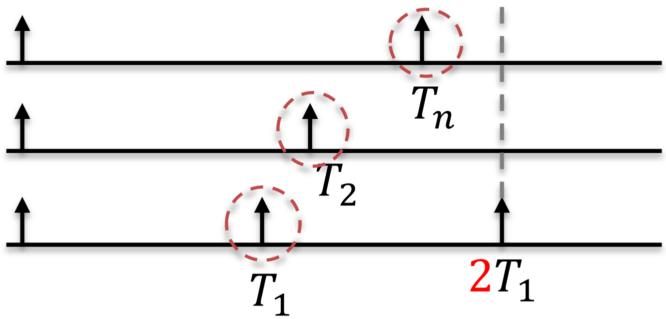

$$T_1 \le T_2 \le \cdots \le T_n \le 2T_1$$

$$1 \le \frac{T_2}{T_1} \le 2, \quad 1 \le \frac{T_3}{T_2} \le 2, \ldots$$

However, the **exact test** is very efficient in this case!

$$R_i \ge C_i + \sum_{j=1}^{i-1} \left\lceil \frac{R_i}{T_j} \right\rceil C_j$$

Small period ratios

Large period ratios

Schedulability ratio

The cost of analysis

NP-hard problem [Eisenbrand08]

Our Observation

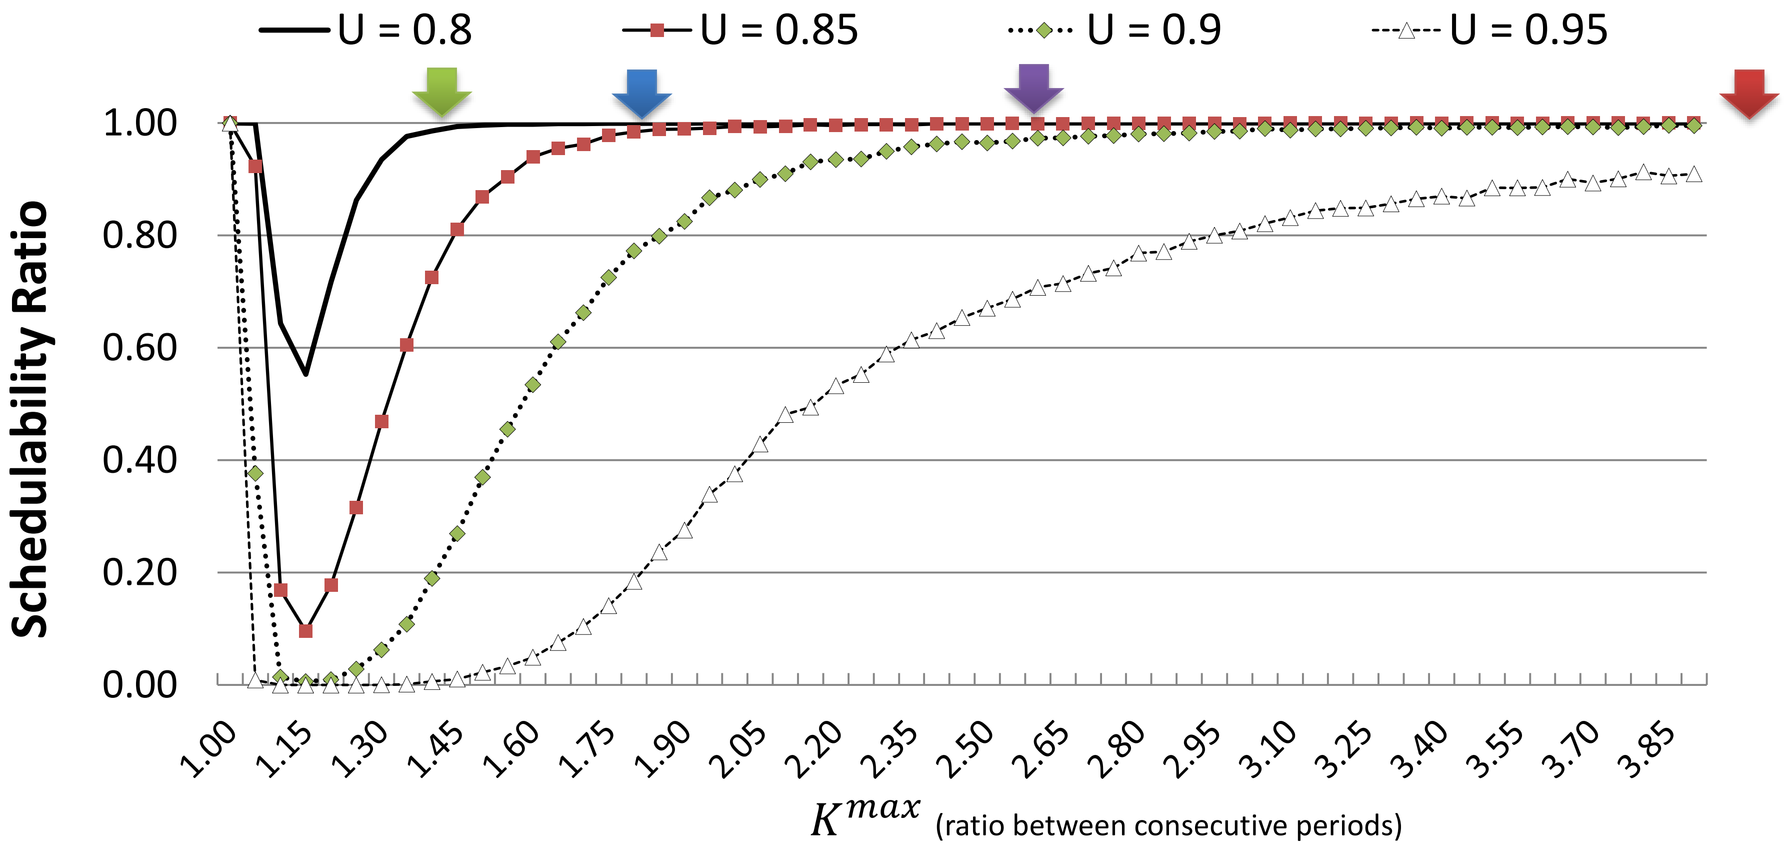

Periods: T_1 uniform distribution from [1, 10]

K_i uniform distribution from $[1, K^{max}]$

$T_i = K_i T_{i-1}$

Utilizations: uUniFast [Bini05]

WCETs: $C_i = u_i T_i$

Task sets satisfy the Liu and Layland assumptions

Why is it Important to Quantify this Effect?

A large number of tasks

Each task might have a set of configurations

Cost of schedulability analysis

System designer

Cost of re-configuration of the system if the previous configuration was not feasible

Identifying **RM-Friendly periods** reduces the costs

Why is it Important to Quantify this Effect?

- To **understand** what we get in the experiments!

Example:

Utilizations: uUniFast, U=0.9, Periods: uniform from [10, 1000], WCET: $C_i = u_i T_i$

What does it tell us?

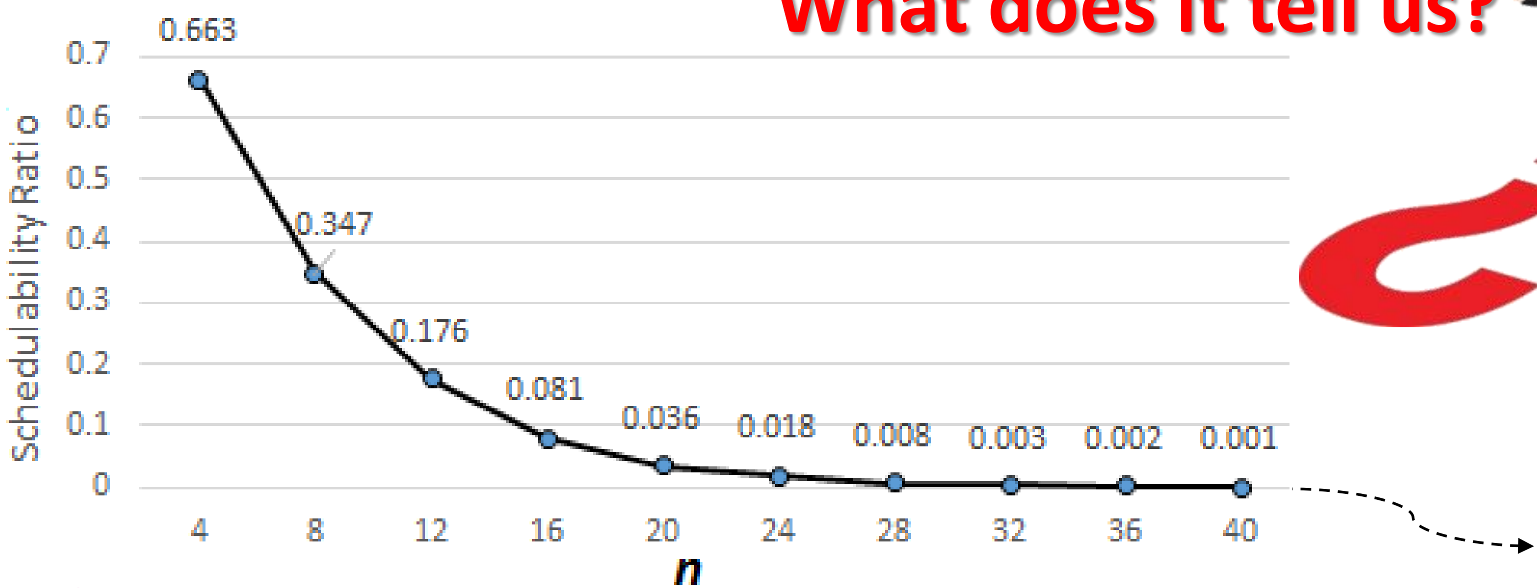

What does it NOT tell us?

Task sets satisfy the Liu and Layland assumptions

Why is it Important to Quantify this Effect? (cont.)

- To understand what we get in the experiments!
- Example:
  - Utilizations: uUniFast, U=0.9   Periods: uniform from [10, 1000]   WCET: $C_i = u_i T_i$

What does it NOT tell us?   The effect of period ratio!

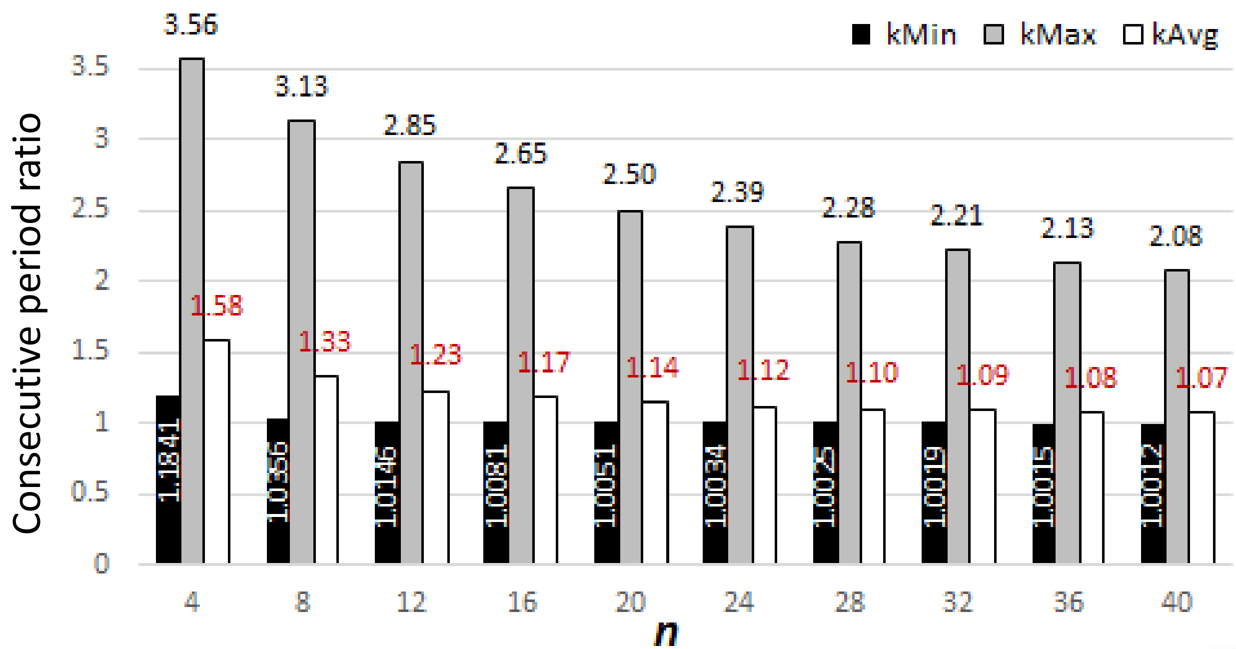

Average period ratio is 1.07

Maximum per-task utilization for n=40 is 0.09.

Contributions

- We quantify the effect of **period ratio** on RM schedulability

Base period ratio
(integer part of period ratio)

Period residual
(fractional part of period ratio)

and
Task utilization

- We derive a set of design hints

- We present a necessary schedulability test for RM based on period ratios

Read it in the paper

Agenda

- **Related work**
- System model and definitions
- Quantifying the effect of period ratios
- Evaluation
- Conclusion

Related Work

- When the **period ratio** approaches to infinity,
  the maximum schedulable utilization reaches to 1 [Lehozcky89]

- If periods are **harmonic**, and $U \leq 1$, the task set will be schedulable by RM [Han97]

- Davis et al., showed that if periods are selected randomly by **log-uniform distribution**, RM schedulability increases [Davis08, Emberson10].

- Wei et al., presented an efficient schedulability bound for RM based on the **ratio** between **the smallest and the largest periods** and utilization of the tasks [Wei08].

- Bini presented a utilization-based schedulability test in which the minimum value of consecutive period ratios is used too [Bini15].

Task Model and Definitions

Assumptions

- Preemptive periodic or sporadic tasks
- Implicit deadline
- No dependency or self-suspension

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

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

Tasks are indexed by their periods

Period ratio of two tasks

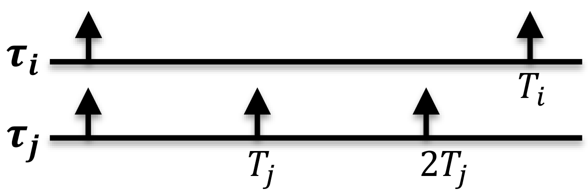

Base period ratio (integer part)

Period residual (fractional part)

$$K_{i,j} = \frac{T_i}{T_j} = k_{i,j} + \gamma_{i,j} \qquad \begin{cases} k_{i,j} \in \mathbb{N} \\ 0 \le \gamma_{i,j} < 1 \end{cases}$$

Example

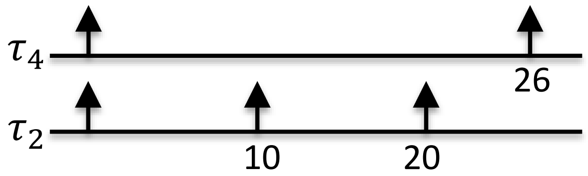

$$K_{4,2} = 2.6 \qquad \begin{cases} k_{4,2} = 2 \\ \gamma_{4,2} = 0.6 \end{cases}$$

Agenda

- Related work
- System model and definitions
- **Quantifying the effect of period ratios**
- EValuation
- Conclusion

Our Solution

- We start from a sufficient schedulability test for τ_i

$$t \geq C_i + \sum_{j=1}^{i-1} \left\lceil \frac{t}{T_j} \right\rceil C_j$$

We evaluate the WCRT equation at $t = T_i$

$$T_i \geq C_i + \sum_{j=1}^{i-1} \left\lceil \frac{T_i}{T_j} \right\rceil C_j$$

(some arithmetic operations)

$$\sum_{\substack{1=j \\ \frac{T_i}{T_j} \notin \mathbb{N}}}^{i} \frac{1-\gamma_{i,j}}{k_{i,j}+\gamma_{i,j}} u_j + \sum_{j=1}^{i} u_j \leq 1$$

$K_{i,j} = \frac{T_i}{T_j} = k_{i,j} + \gamma_{i,j}$ $\quad$ $k_{i,j} \in \mathbb{N}, \quad 0 \leq \gamma_{i,j} < 1$

Understanding the Effect of Period Ratios

$$\sum_{j=1}^{i} \frac{1-\gamma_{i,j}}{k_{i,j}+\gamma_{i,j}} u_j + \sum_{j=1}^{i} u_j \le 1$$

$\frac{T_i}{T_j} \notin \mathbb{N}$

1- The effect of periods that are **not-harmonic with T_i**

2- The contribution of **task utilization**

3- The contribution of **period ratio**

4- The contribution of **period residual**

Design hints:

1- Only the tasks with non-harmonic period with T_i have an adverse effect on the schedulability of τ_i

2- If you have a highly utilized task in the system, try to force other periods to be harmonic with its period.

3- Try to have either large period ratios or low utilization for the tasks that are not harmonic with T_i.

4- Force the period residual of highly utilized tasks to be large with respect to T_i.

Note: $\gamma_{i,j}$ shows how close is a period to be harmonic with another period. $\gamma_{i,j}\sim 0$ or $\gamma_{i,j}\sim 1$ are almost harmonic.

Remember: $K_{i,j} = \frac{T_i}{T_j} = k_{i,j} + \gamma_{i,j}$ $\quad k_{i,j} \in \mathbb{N}, \quad 0 \le \gamma_{i,j} < 1$

Evaluation

Main Questions of Our Experiments

- What is the effect of period residual?
- How different schedulability tests react towards an increase in
  - The maximum value of consecutive period ratios
  - The period residual

- **Schedulability Tests:**
  - Linear approximation
  - DCT [Han97] based on harmonic periods
  - Park 2014: verify the WCRT inequality at the latest releases of the high priority tasks

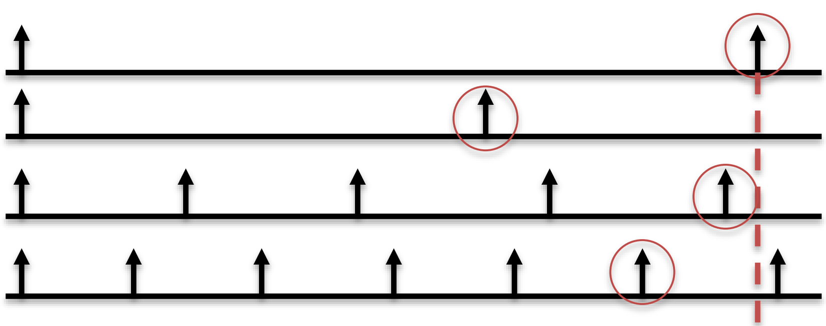

$$t \geq C_i + \sum_{j=1}^{i-1} \left\lceil \frac{t}{T_j} \right\rceil C_j$$

The Effect of Period Residual

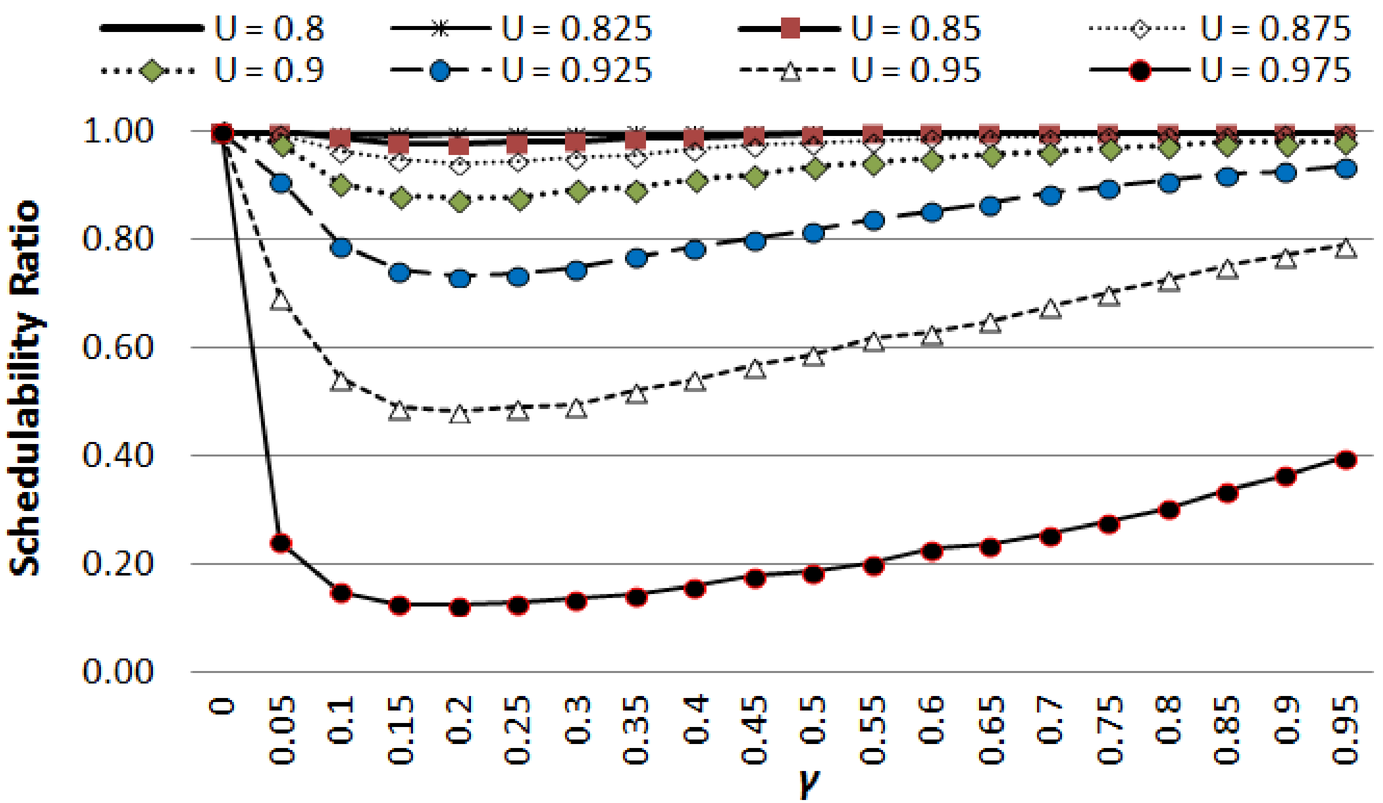

Periods: T_1 uniform distribution from [1, 10]

$k_{i,i-1}$ uniform distribution from $\{1, 2, 3\}$

$\gamma_{i,i-1}$ uniform distribution from $[0, \gamma]$

$T_i = (k_{i,i-1} + \gamma_{i,i-1})T_{i-1}$

Utilizations: uUniFast [Bini05]

WCETs: $C_i = u_i T_i$

The Impact of Different Schedulability Tests

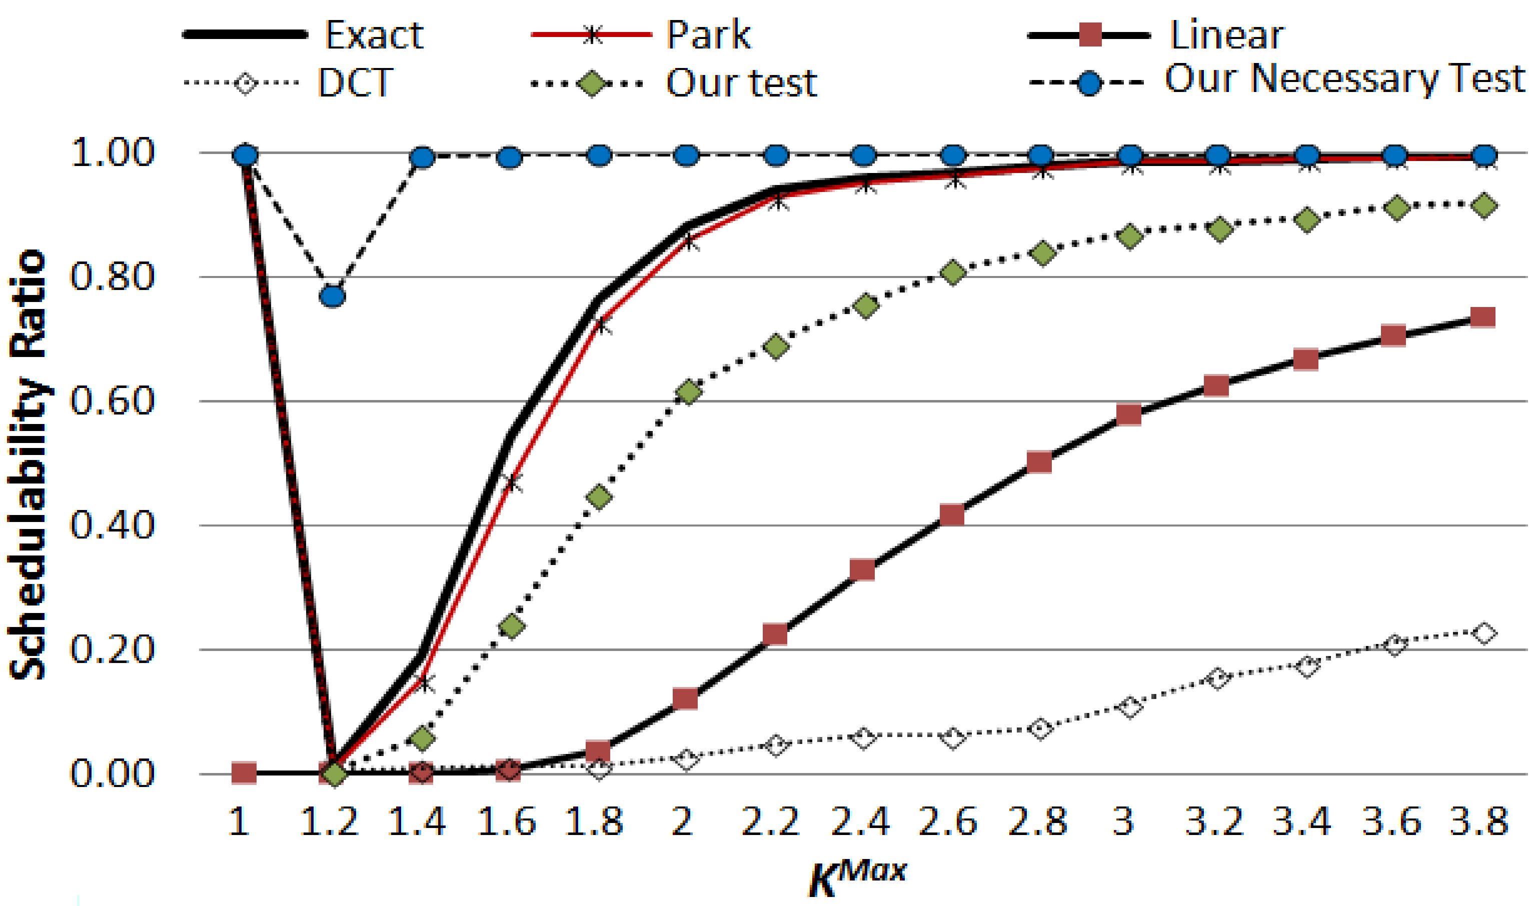

Periods: T_1 uniform distribution from [1, 10]

K_i uniform distribution from $[1, K^{max}]$

$T_i = K_i T_{i-1}$

Utilizations: uUniFast [Bini05]

WCETs: $C_i = u_i T_i$

The Impact of Different Schedulability Tests, cont.

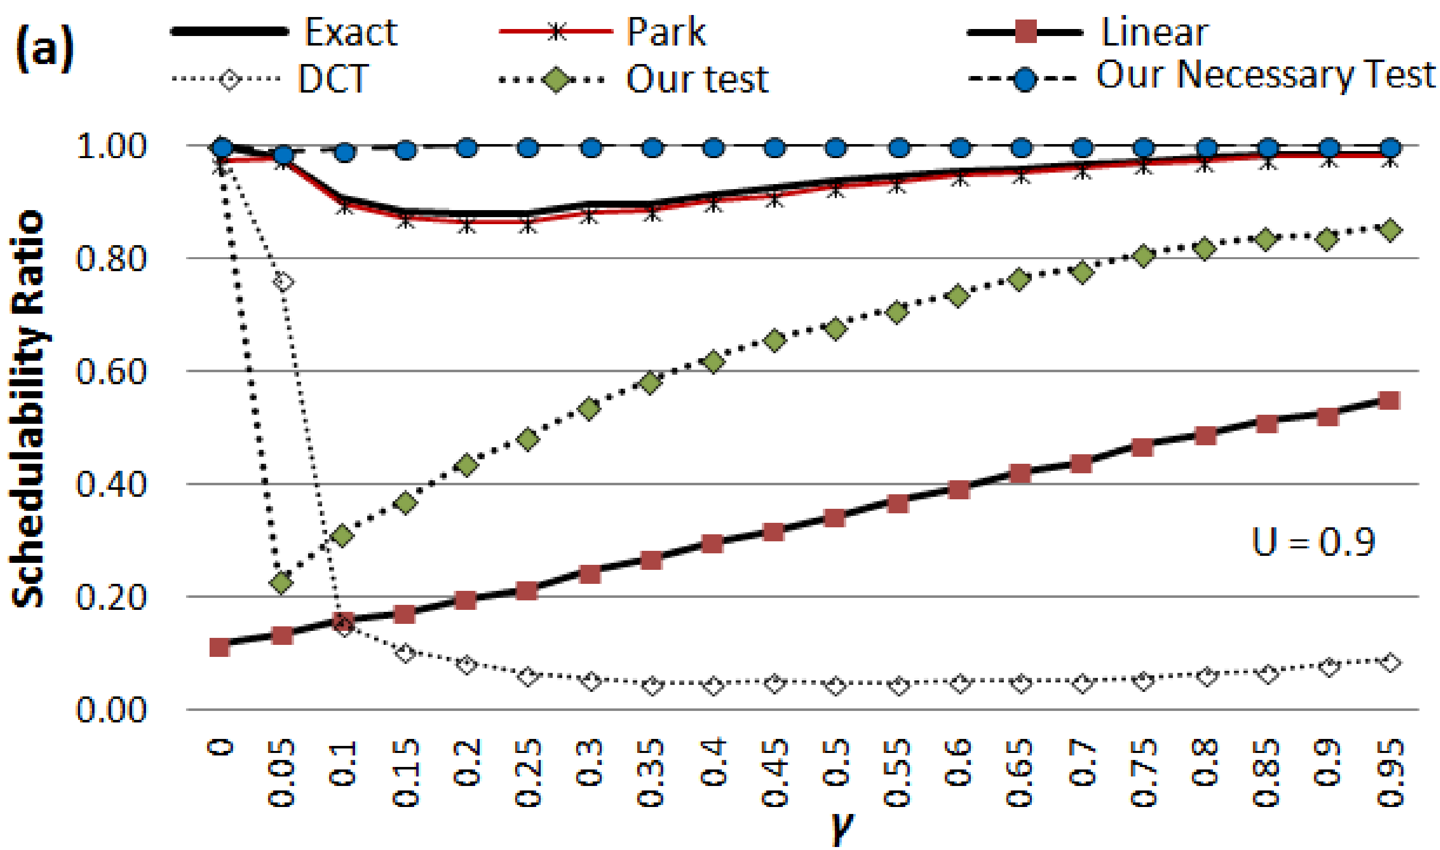

Periods: T_1 uniform distribution from [1, 10]

$k_{i,i-1}$ uniform distribution from $\{1, 2, 3\}$

$\gamma_{i,i-1}$ uniform distribution from $[0, \gamma]$

$T_i = (k_{i,i-1} + \gamma_{i,i-1})T_{i-1}$

Utilizations: uUniFast [Bini05]

WCETs: $C_i = u_i T_i$

Agenda

- Related work
- System model and definitions
- Quantifying the effect of period ratios
- A necessary schedulability test
- Experiments
- **Conclusion**

Conclusion

- We quantified the effect of period ratio
- We have considered the effect of
  - Base period ratio
  - Period residual
  - Utilization of each task

It helps designers to create RM-Friendly task sets

It helps us to understand the experimental results

It helps us to design fair experiments

- **Future work**
  - Designing an efficient task partitioning algorithm based on RM-Friendly tasks
  - Considering tasks with constrained or arbitrary deadlines
  - Using our result to build a parameter assignment tool for systems with a set of configurations

Thank you.

Now you are behind the scene!

Agenda

- **A necessary schedulability test**
- More interesting experiments
- The proof for small period residual

The Idea of our Necessary Schedulability Test for RM

- Example of a feasible schedule

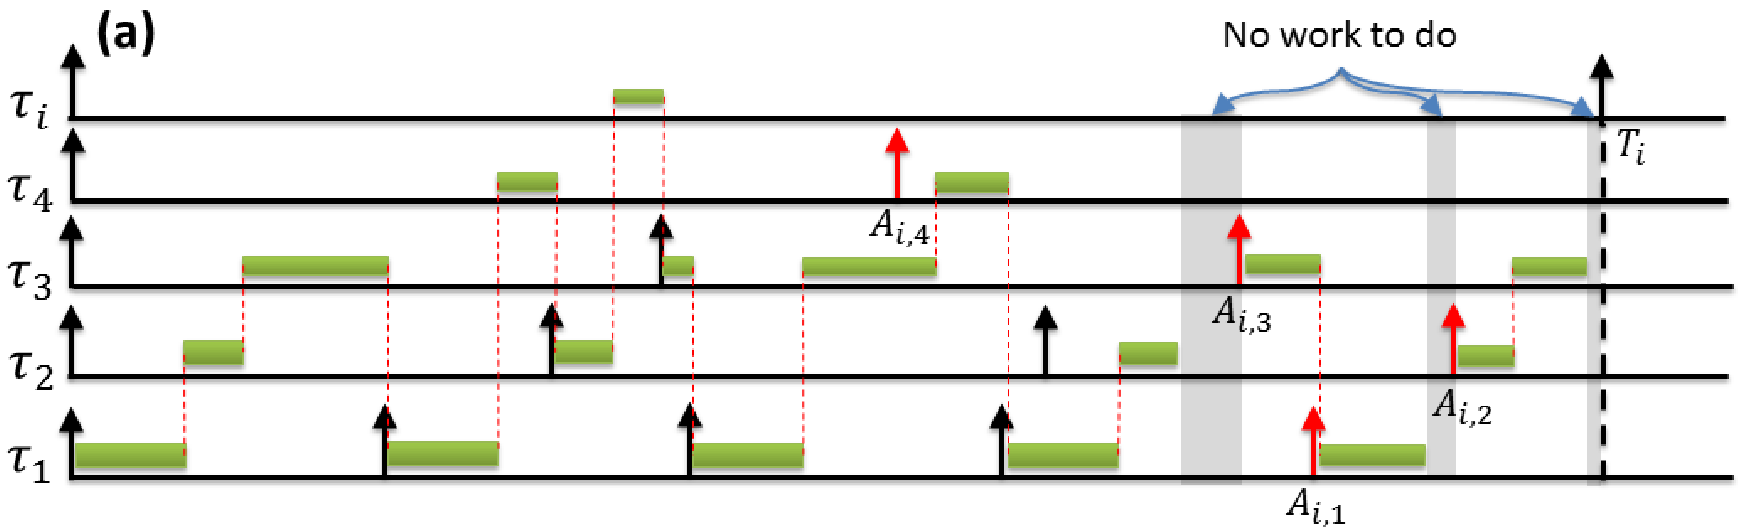

The Idea of our Necessary Schedulability Test for RM

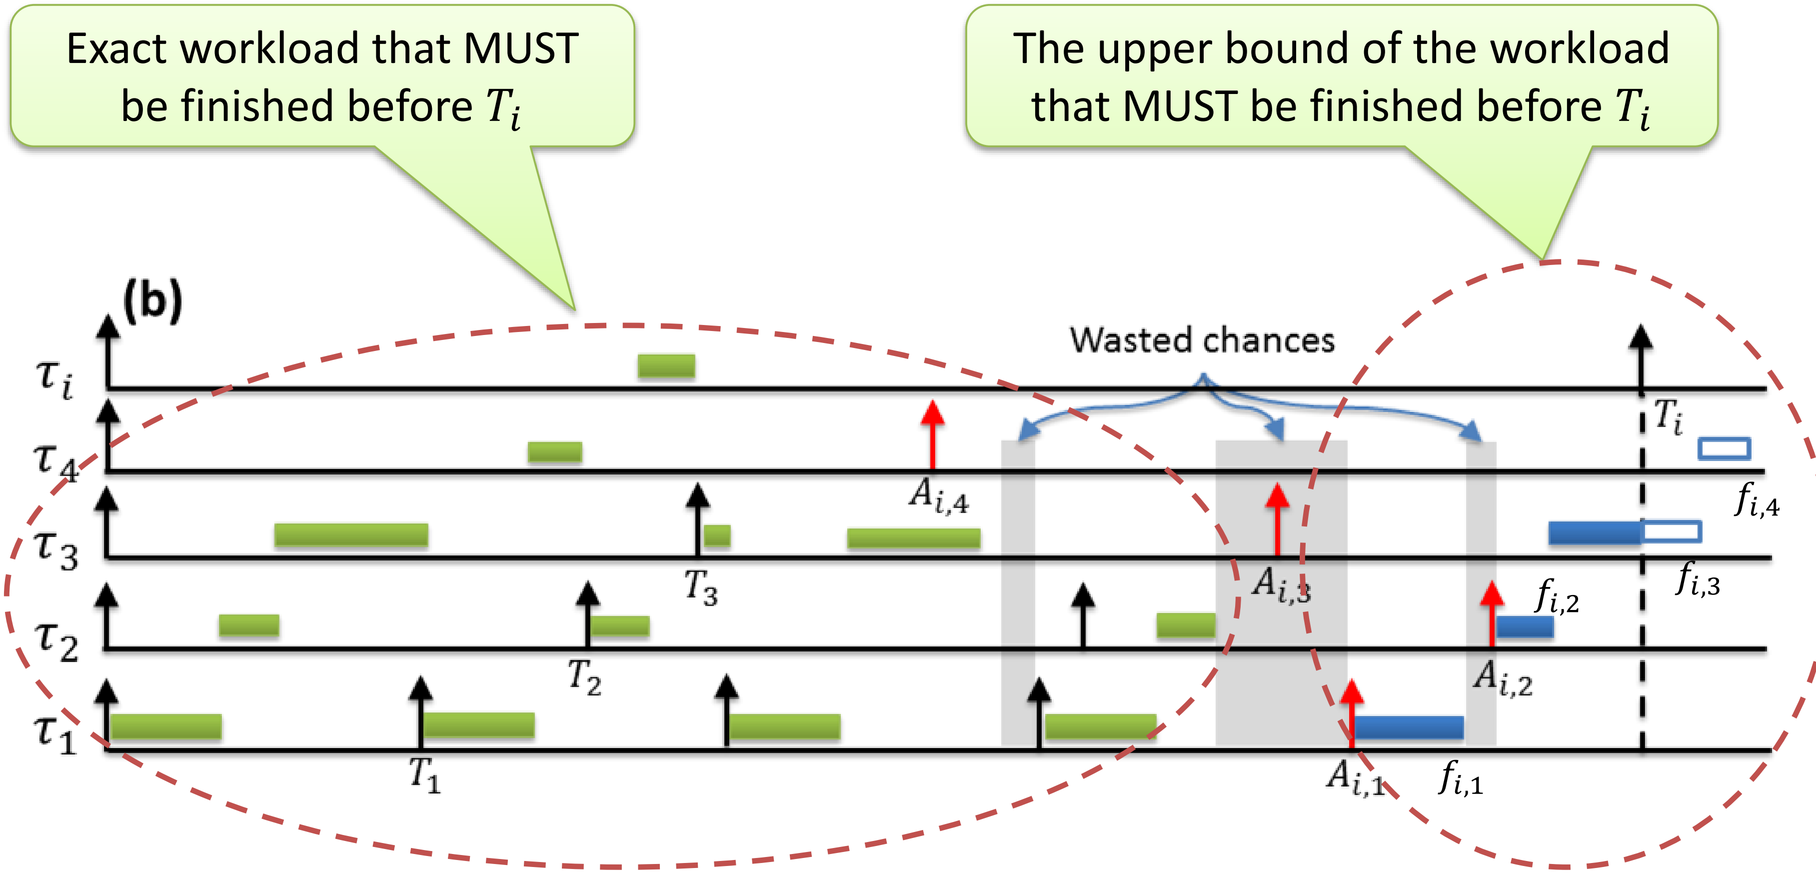

$$T_i \geq \sum_{1 \leq j \leq i} \left\lfloor \frac{T_i}{T_j} \right\rfloor C_j + \sum_{1 \leq j < i} (f_{i,j} - s_{i,j})$$

Agenda

- A necessary schedulability test
- **More interesting experiments**
- The proof for small period residual

Why it is Important to Quantify this Effect? (cont.)

- To understand what we get in the experiments!
- Another example:
  - Utilizations: uUniFast, U=0.9, Periods: log-uniform from [1, 1000], WCET: $C_i = u_i T_i$

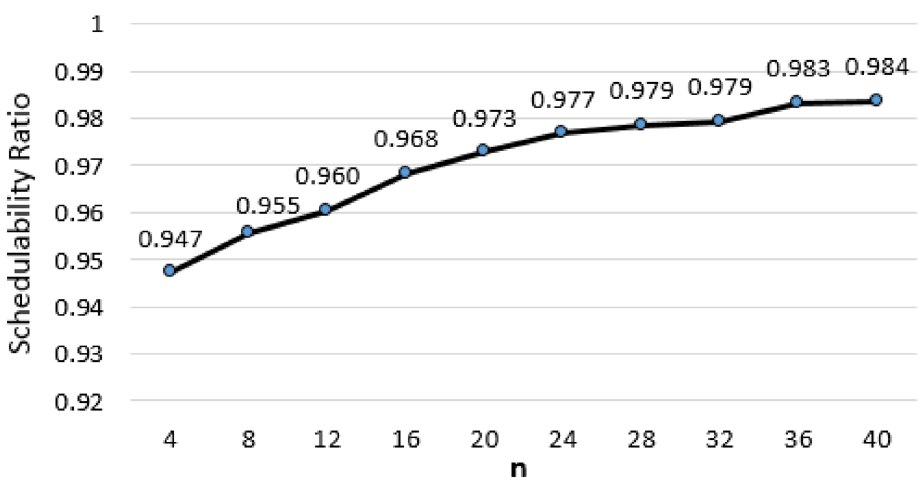

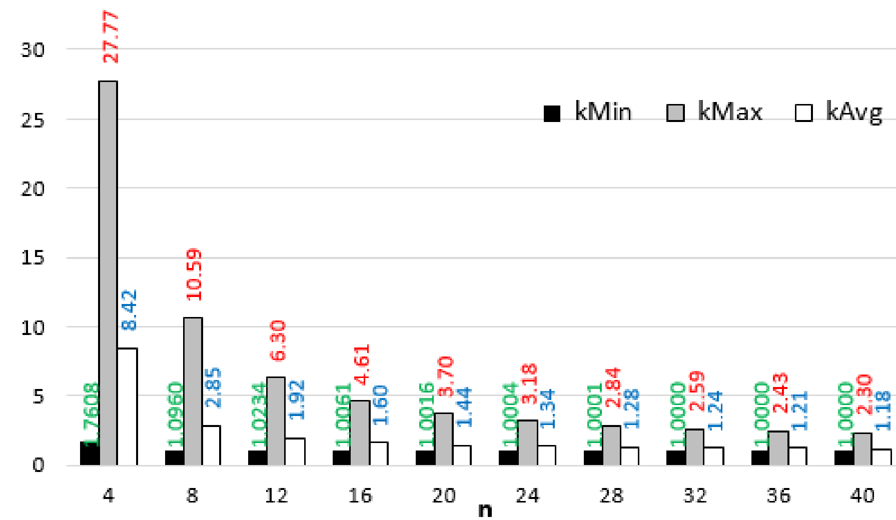

Average period ratio: 1.20

Maximum per-task utilization: 0.09

Why it is Important to Quantify this Effect?

- To understand what we get in the experiments!
- Example:
  - Utilizations: uUniFast, U=0.9 Periods: uniform from [10, 1000] WCET: $C_i = u_i T_i$

So how about the effect of task utilizations? Will it not help?

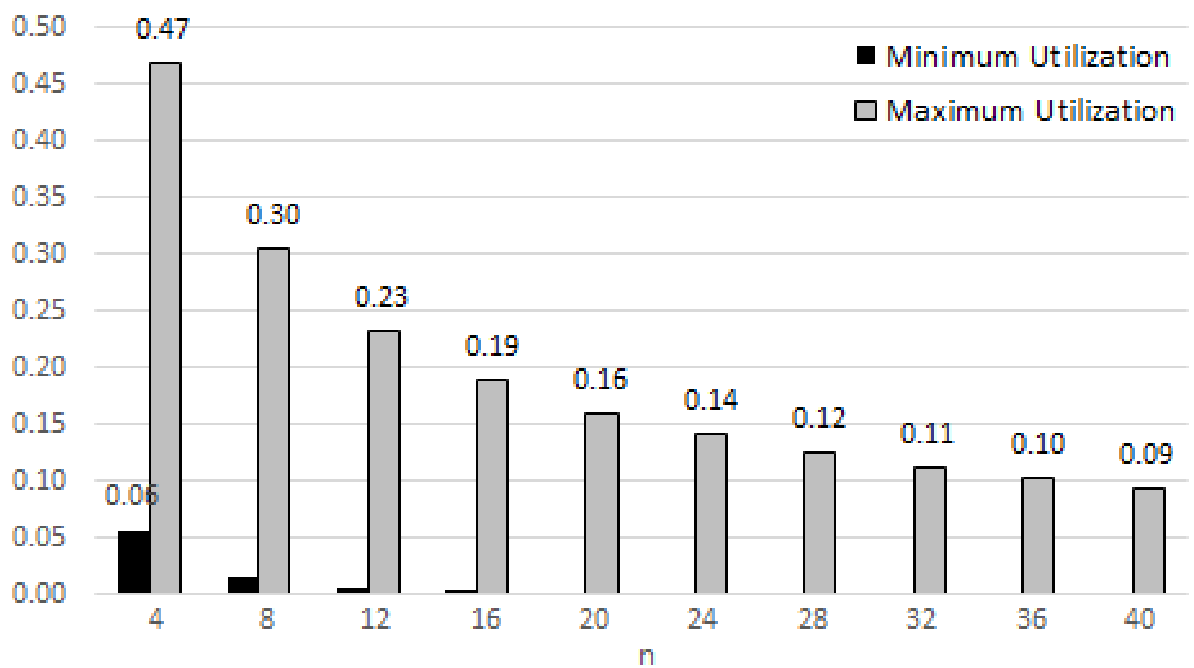

Maximum per-task utilization: 0.09

Evaluating the Point-Reduction Method (HET) of [Bini04]

$$\mathcal{TS}(\tau_i) \doteq \mathcal{P}_{i-1}(D_i)$$

Initial Value

Recursive step

$$\mathcal{P}_j(t) = \begin{cases} \{t\} & j = 0 \\ \mathcal{P}_{j-1}\left(\left\lfloor \frac{t}{T_j} \right\rfloor T_j\right) \cup \mathcal{P}_{j-1}(t) & \text{otherwise} \end{cases}$$

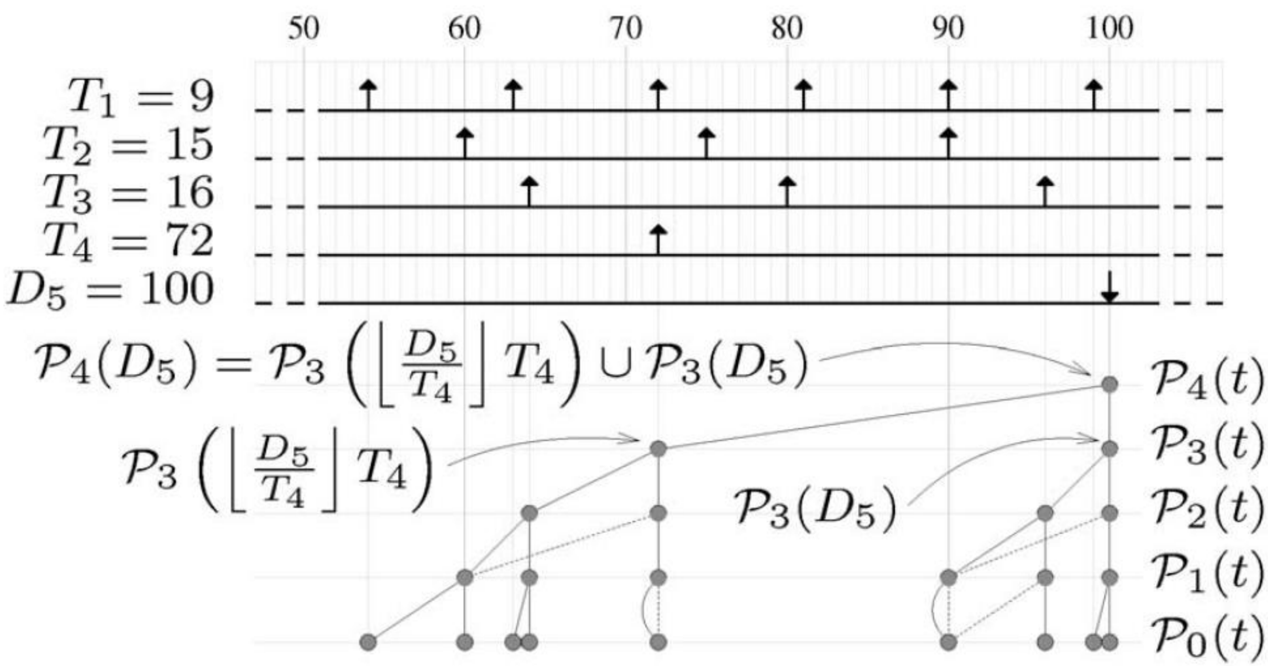

Evaluating the Point-Reduction Method (HET) of [Bini04]
Period range: [1, 100]

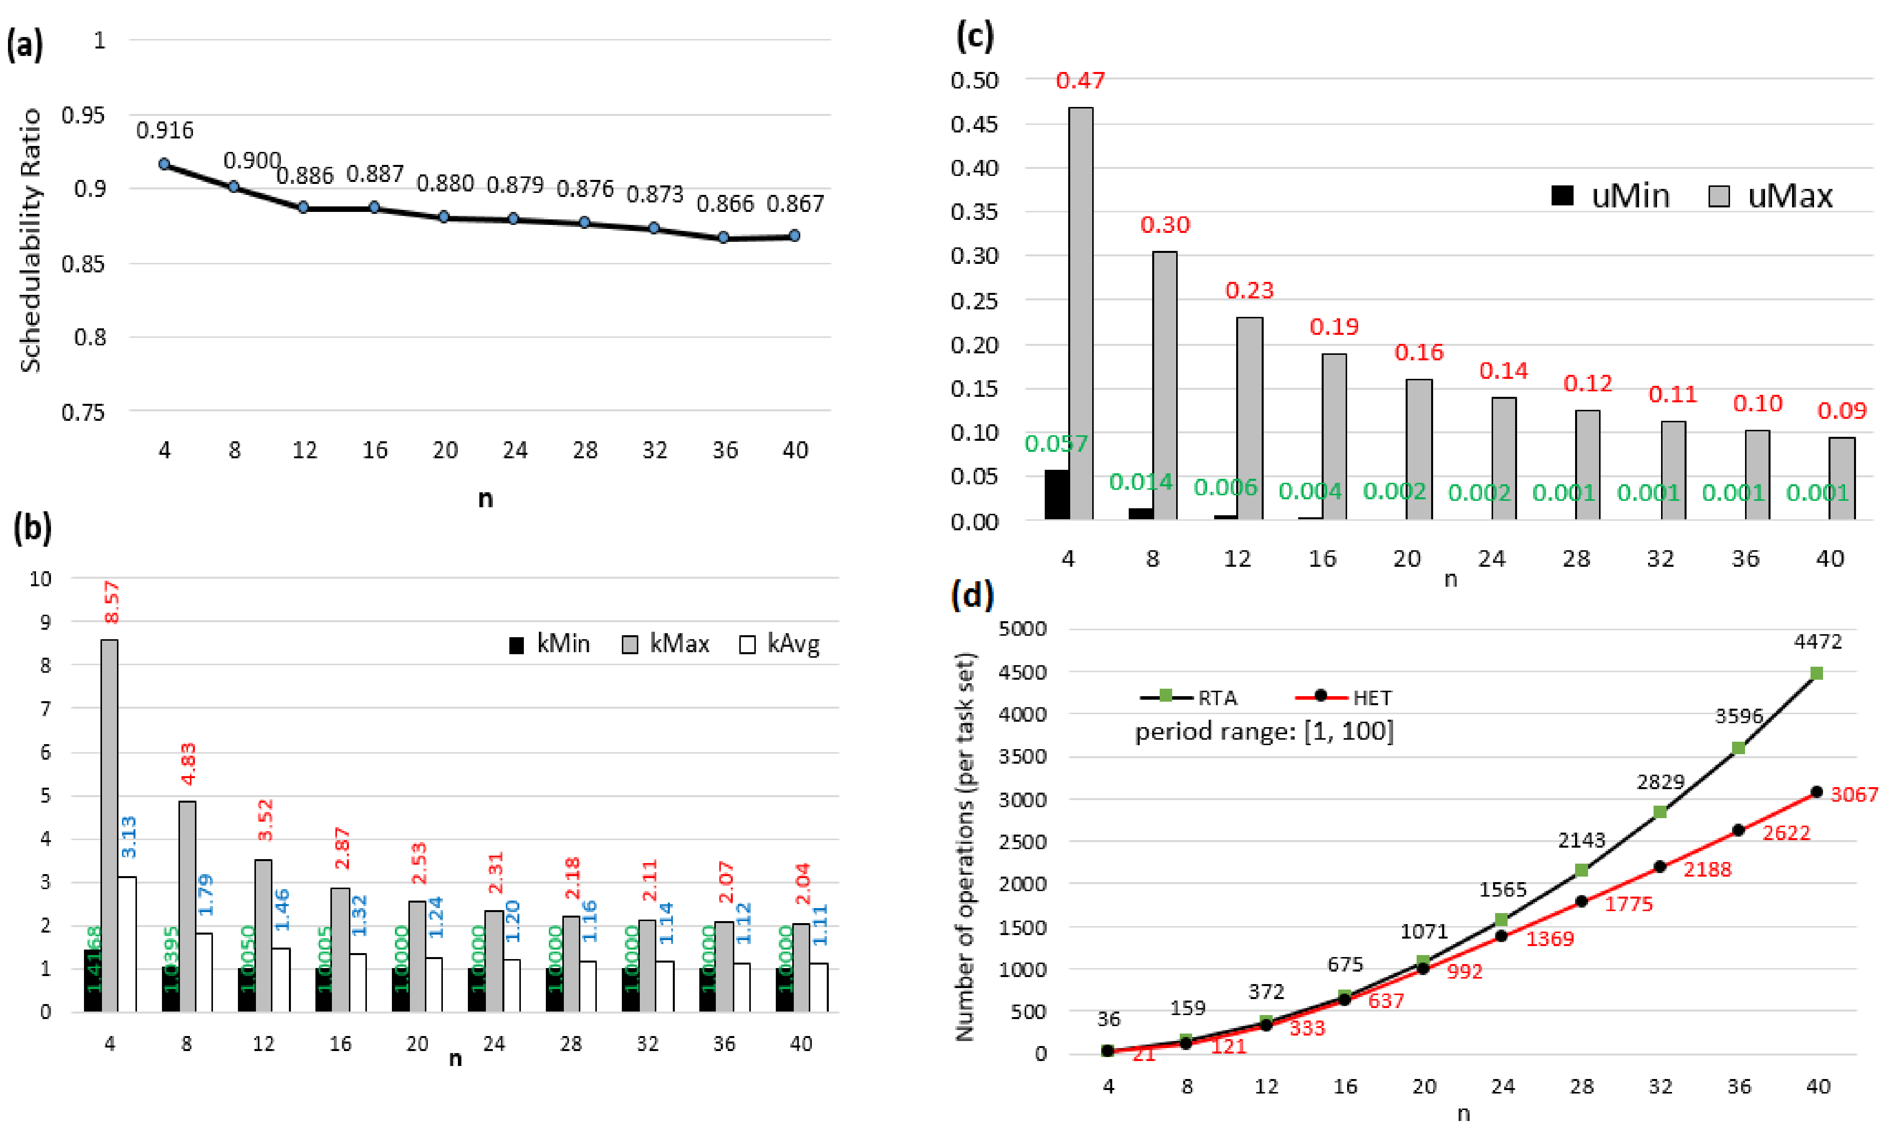

Evaluating the Point-Reduction Method (HET) of [Bini04]
Period range: [1, 1000]

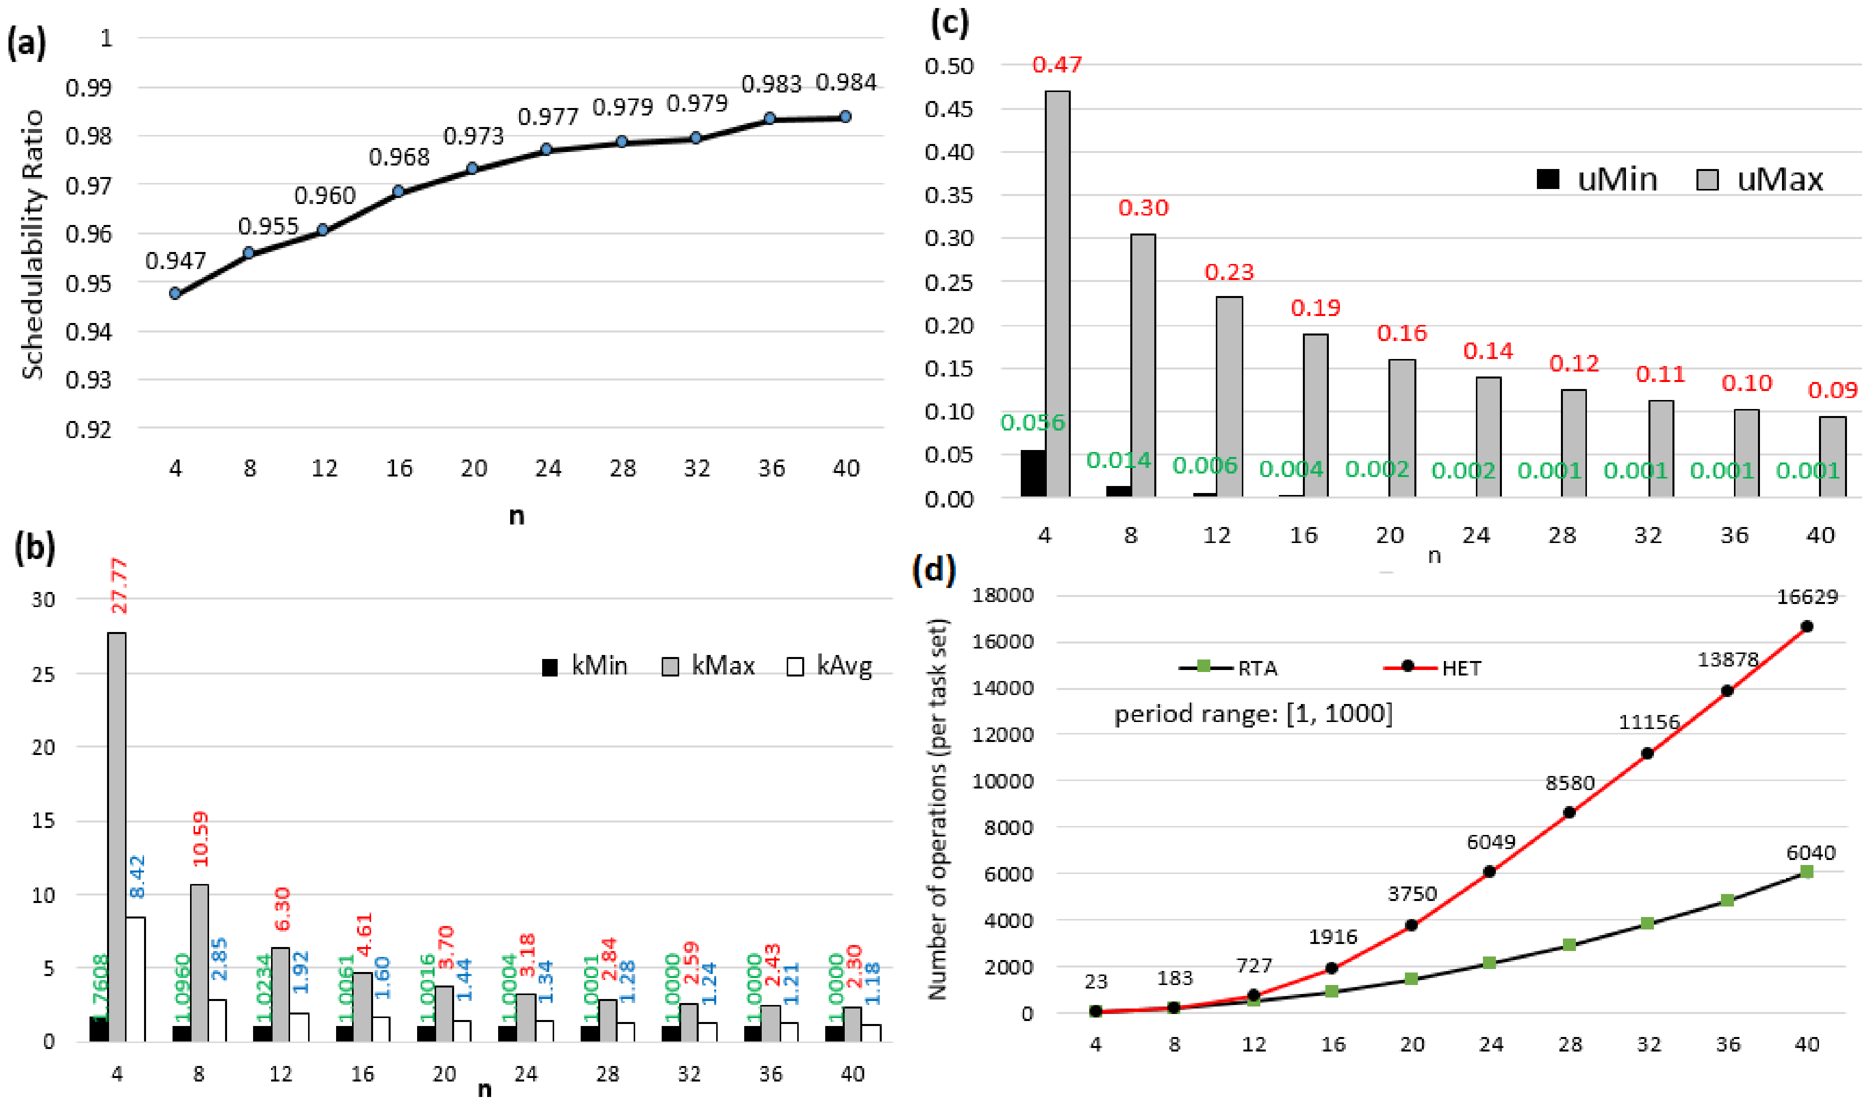

Evaluating the Point-Reduction Method (HET) of [Bini04]
Period range: [1, 10000]

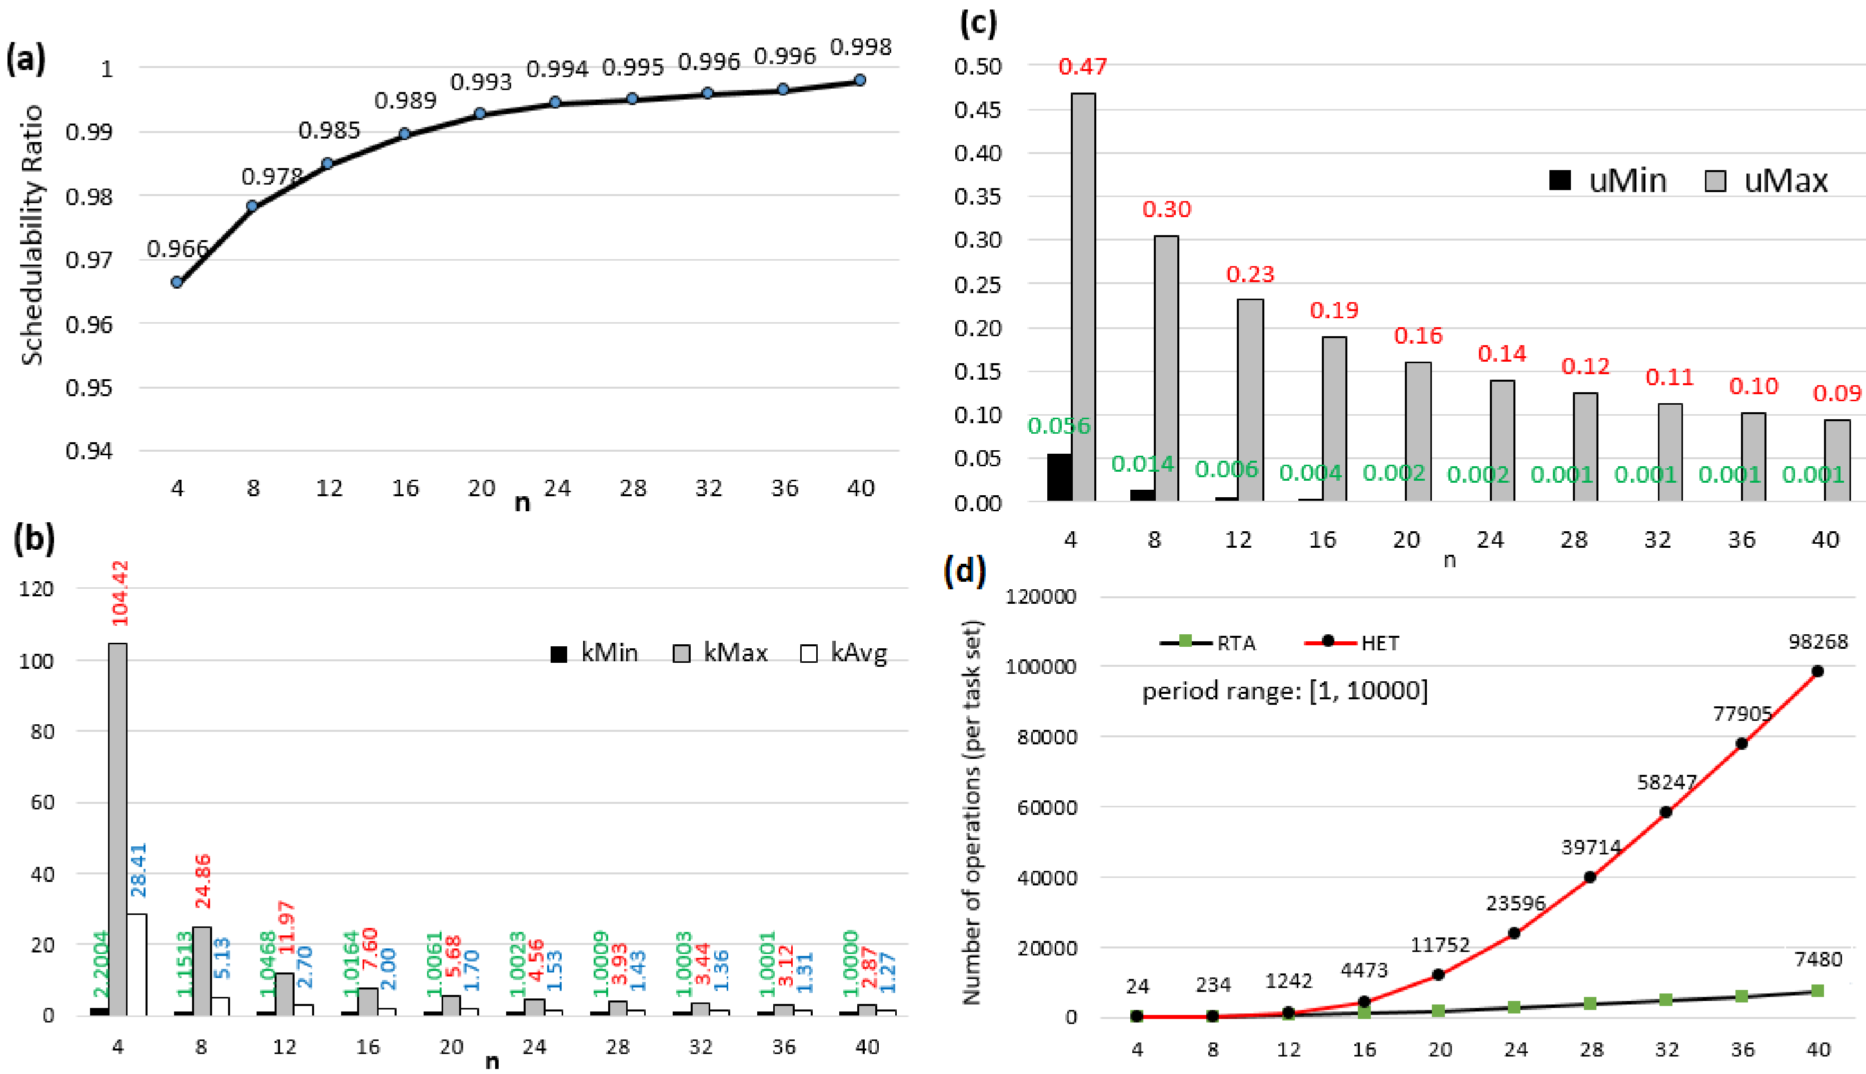

Evaluating the Point-Reduction Method (HET) of [Bini04]
Period range: [1, 100000]

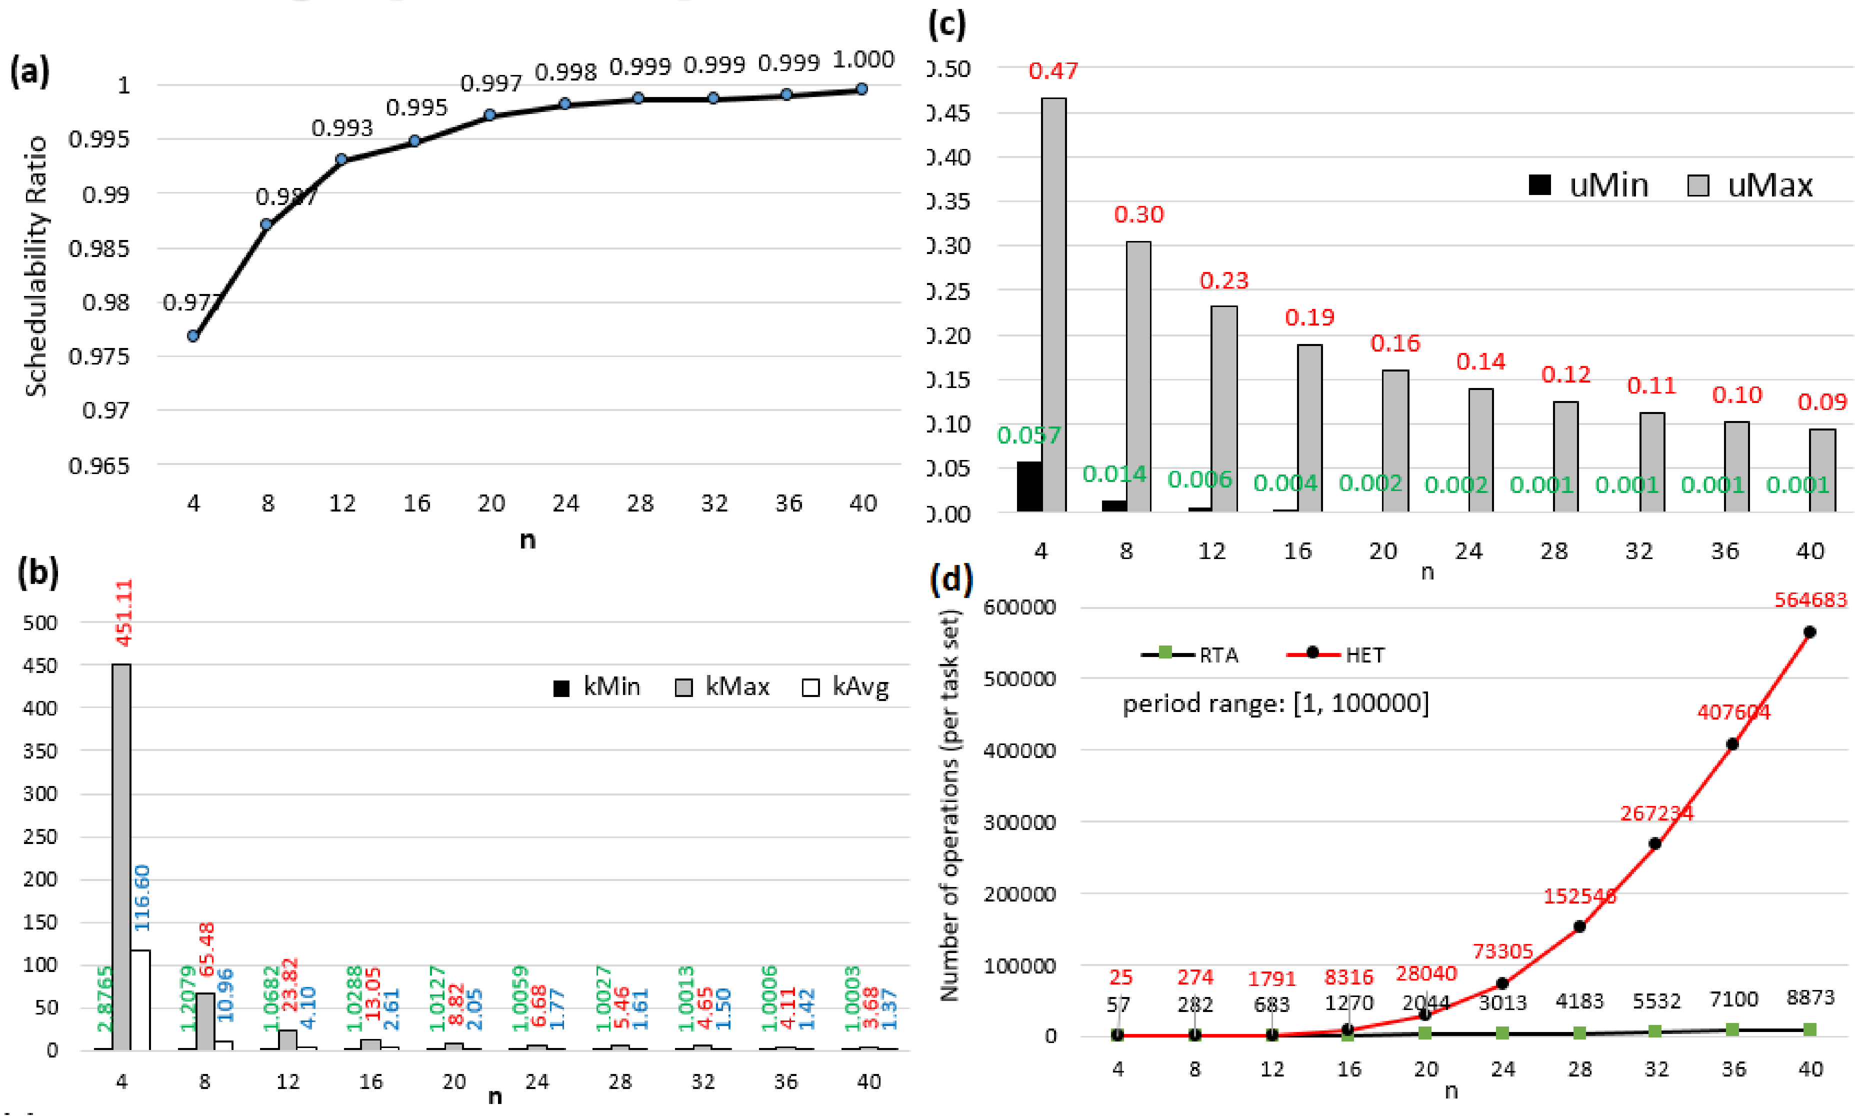

Our results is consistent with [Emberson10]

Algorithm	Orders of magnitude spanning task periods 1	2	3	4	5	6
RTA	1247	1652	2050	2462	2847	3297
HET	380	1326	5642	23014	81636	197642
HET/RTA	0.3	0.8	2.75	9.35	28.7	60.0

Agenda

- A necessary schedulability test
- More interesting experiments
- **The proof for small period residual**

Small Period Residual

$$\frac{k_{i,m}}{k_{i,m}+\gamma_{i,m}} + \sum_{1\le j<i} \frac{\gamma_{i,j}}{k_{i,j}+\gamma_{i,j}} u_j \ge \sum_{1\le j\le i} u_j$$

$$\gamma_{i,j} \to 0 \Rightarrow 1 \ge \sum_{j=1}^{i} u_j$$