

Workload-Intensity-Sensitive Timing Behavior Analysis for Distributed Multi-User Software Systems

Matthias Rohr^{1,2}, André van Hoorn², Wilhelm
Hasselbring^{2,3}, Marco Lübcke⁴, and Sergej Alekseev^{5,6}

¹ BTC Business Technology Consulting AG, Germany,

^{*}, ² Graduate School TrustSoft, University of Oldenburg, Germany,

³ Software Engineering Group, University of Kiel, Germany,

⁴ CeWe Color AG & Co. OHG, Oldenburg, Germany,

⁵ Nokia Siemens Networks GmbH, Berlin, Germany,

⁶ Hochschule Mittweida, University of Applied Sciences, Mittweida, Germany


January 29, 2010

Joint WOSP/SIPEW International Conference on Performance Engineering,
San Jose, California, USA



^{*} This work is supported by the German Research Foundation (DFG), grant GRK 1076/1

Software operation response times



Motivation

Foundations

Approach


Evaluation

Related Work

Conclusions
and future
work

- Workload-intensity can be a major influence to timing behavior in enterprise information systems
- Varying workload-intensity can cause high variance in timing behavior
- High variance can make it difficult to draw statistical conclusions
 - E.g., proper threshold determination for anomaly detection

Software operation response times



Motivation

Foundations

Approach


Evaluation

Related Work

Conclusions
and future
work

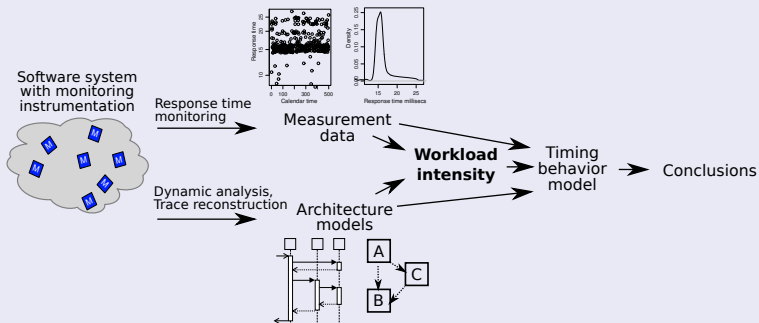
- Workload-intensity can be a major influence to timing behavior in enterprise information systems
- Varying workload-intensity can cause high variance in timing behavior
- High variance can make it difficult to draw statistical conclusions
 - E.g., proper threshold determination for anomaly detection

Without considering workload intensity



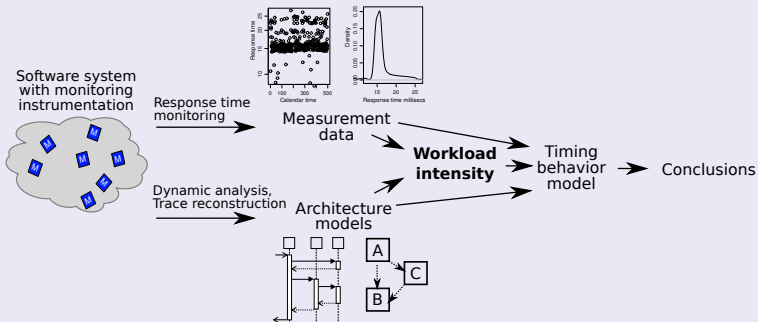
Our approach

With considering workload intensity



Our approach

With considering workload intensity



Our approach

- Goal: “Reduce” variation for statistical timing behavior analysis
- Categorization based on workload-intensity levels
- Requires only light-weight common monitoring infrastructure

- 1 Motivation
- 2 Foundations**
- 3 Workload-intensity-sensitive timing behavior analysis
- 4 Empirical evaluation
- 5 Related work
- 6 Conclusions and future work

Motivation

Foundations

Approach

Evaluation

Related Work

Conclusions
and future
work

- System architecture and implementation:
 - Hardware design
 - Software design
 - Middleware [?]
- **System usage:**
 - Workload-intensity
 - Concurrent service requests [Happe et al. 2008]
 - Number of active users [?]
 - Individual request characteristics
 - Parameter values and parameter size [?]
 - Caller identity / stack content [?]
- State:
 - Cache content
 - Load balancer state
 - Software application state
 - Other active processes on same platform
 - Database content

Motivation

Foundations


Approach

Evaluation

Related Work

Conclusions
and future
work

Relation between response times and workload intensity



(Schematic illustration based on ?)

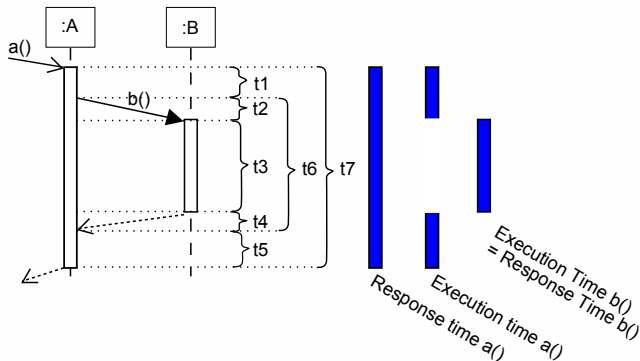


Figure: Response times and execution times.

- 1 Motivation
- 2 Foundations
- 3 Workload-intensity-sensitive timing behavior analysis**
- 4 Empirical evaluation
- 5 Related work
- 6 Conclusions and future work

Motivation

Foundations

Approach

Evaluation

Related Work

Conclusions
and future
work

1. Monitoring

- Recording of:
 - **Response times**: Time between start and end of software operation executions
 - **Execution sequences** corresponding to a user request
 - Host identifier
- Reconstruction of Traces and Dependency Graphs
- Kieker framework^a [?]

^a<http://kieker.sourceforge.org>

2. Computation of workload-intensity from monitoring data:

3. Categorization based on workload-intensity levels

Motivation

Foundations

Approach

Evaluation

Related Work

Conclusions
and future
work

1. Monitoring

2. Computation of workload-intensity from monitoring data:



→ next slides

3. Categorization based on workload-intensity levels

1. Monitoring

2. Computation of workload-intensity from monitoring data:

3. Categorization based on workload-intensity levels



- 1 The *pwi* range is divided into intervals (e.g., 15) of equal length
- 2 Bins are extended to minimum size (e.g., 100 observations)

- Key element of our approach: Four alternative workload-intensity metrics, denoted pwi (Platform Workload Intensity):

| Metric | Time metric | Execution environment | Operation weighting |
|---------|-----------------|-----------------------|---------------------|
| pwi_1 | Response times | Non-distributed | No weighting |
| pwi_2 | Execution times | Non-distributed | No weighting |
| pwi_3 | Execution times | Distributed | No weighting |
| pwi_4 | Execution times | Distributed | Learned |

Motivation

Foundations

Approach

Evaluation

Related Work

Conclusions
and future
work

Average number of concurrent traces during the time period between the start (call action) and the end of an operation execution.

Average number of concurrent traces during the time period between the start (call action) and the end of an operation execution.

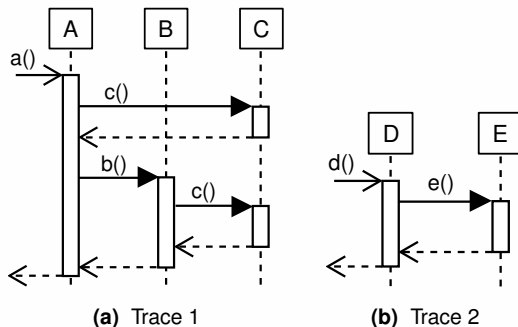




Figure: Example traces: UML Sequence Diagrams


Average number of concurrent traces during the time period between the start (call action) and the end of an operation execution.



Average number of concurrent traces during the time period between the start (call action) and the end of an operation execution.



Average number of concurrent traces during the time period between the start (call action) and the end of an operation execution.




An operation execution's pwi_2 is the average number of concurrent traces during its execution time period.

- Difference to pwi_1 : Execution time period instead of response time period
- No competition for resources during waiting for sub-calls

An operation execution's pwi_2 is the average number of concurrent traces during its execution time period.


- Difference to pwi_1 : Execution time period instead of response time period
- No competition for resources during waiting for sub-calls




An operation execution's *pwi*₃ is the average number of concurrent **active executions within the same execution environment** during its execution time period.

- *pwi*₃ extends *pwi*₂ for distributed systems.
- Assumption: Execution contexts have own hardware platform
- Hypothesis: Little competition for resources with executions in other execution environments.

An operation execution's pwi_3 is the average number of concurrent **active executions within the same execution environment** during its execution time period.



An operation execution's pwi_3 is the average number of concurrent **active executions within the same execution environment** during its execution time period.



*pwi*₄ extends *pwi*₃ by using the weight $w_{o,p} \in W$ for considering concurrent executions of p for evaluating o .

- *pwi*₁-*pwi*₃ equally consider different (local) operations
- Resource competition leads to high weights.

Computation of weight matrix W

- W is determined via machine learning from historical monitoring data
- Learning goal: maximum standard deviation reduction
- High computational costs if many operations are instrumented
- Convention: $w_{o,p}$ is 0, if o and p are not in the same execution environment
- Heuristic: Correlation matrix provides good starting values

Motivation

Foundations

Approach

Evaluation

Related Work

Conclusions
and future
work

Software system with 2 operations:

- *Wait*: Non-busy waiting for 300 ms.
- *Work*: CPU-intensive number crunching.

Software system with 2 operations:

- *Wait*: Non-busy waiting for 300 ms.
- *Work*: CPU-intensive number crunching.

Experiment setting:

- 120,000 random execution of *wait* and *work*
- 1-24 parallel executions

Motivation

Foundations

Approach

Evaluation

Related Work

Conclusions
and future
work

Software system with 2 operations:

- *Wait*: Non-busy waiting for 300 ms.
- *Work*: CPU-intensive number crunching.

Experiment setting:

- 120,000 random execution of *wait* and *work*
- 1-24 parallel executions

Results:

Weight matrix:

| | <i>work</i> | <i>wait</i> |
|-------------|-------------|-------------|
| <i>work</i> | 2.01 | -0.05 |
| <i>wait</i> | 1.03 | 0.05 |

Software system with 2 operations:

- *Wait*: Non-busy waiting for 300 ms.
- *Work*: CPU-intensive number crunching.

Experiment setting:

- 120,000 random execution of *wait* and *work*
- 1-24 parallel executions


Results:

Weight matrix:


| | <i>work</i> | <i>wait</i> |
|-------------|-------------|-------------|
| <i>work</i> | 2.01 | -0.05 |
| <i>wait</i> | 1.03 | 0.05 |

Standard dev. reduction (%):

| | pwi_4 |
|-------------|--------------|
| <i>work</i> | 72.5 ± 2 |
| <i>wait</i> | 18.8 ± 9 |



(a) wait



(b) work

- 1 Motivation
- 2 Foundations
- 3 Workload-intensity-sensitive timing behavior analysis
- 4 Empirical evaluation**
- 5 Related work
- 6 Conclusions and future work

Motivation

Foundations

Approach

Evaluation

Related Work

Conclusions
and future
work

Evaluation Metric

Reduction of standard deviation (in percent) in relation to the original dataset for each operation and in total weighted by the number of observations per operation.

- Evaluation and simulation techniques can benefit from “reduction” of standard deviation, e.g.,
 - in terms of requiring less observations,
 - providing tighter confidence intervals,
 - requiring less or shorter simulation runs [?].

Evaluation method:

- Results for $pwi_1 - 3$ can directly be computed
- Evaluation of pwi_4 requires two separate data sets for training, and one for cross-validation
- Operations with less than 600 observations are accounted 0% reduction

Motivation

Foundations

Approach



Evaluation

Related Work

Conclusions
and future
work


Setting

- 5-node distributed variant¹ of the iBATIS JPetStore
- 34 software operation instrumented
- Probabilistic, multi-user workload using Markov4JMeter
- Real workload intensity curve, scaled to max. 80% capacity utilization



¹Instrumented sources available at <http://sourceforge.net/projects/kieker>

Results



- Standard deviation is reduced in average from 35% for pwi_1 up to 56% for pwi_4 .
- Log-transforming the pwi values, before defining bins additionally improves standard deviation reduction by 29% in average.
- For pwi_4 , this results in a standard deviation reduction of 65%.
- For some operations, there is no benefit.


Setting

- Telecommunication signaling system of Nokia Siemens Networks
- 8 instrumented operations on two clustered nodes


Nokia Siemens
Networks



- Test workload using the companies own workload simulator
- Less than 15% of CPU utilization peak




Results




- pwi_4 performs best in the comparison.
- For all pwi metrics, standard deviation reduction additionally increases by more than 30% if the logarithm of the pwi values are used for defining timing behavior classes.
- Traces do not cross execution environments $\Rightarrow pwi_2 = pwi_3$.

Setting

- Customer portal for ordering photo prints and other photo products of CeWe Color AG, Europe's largest digital photo service provider.
- Large number of monitoring points: 161
- Low utilization: CPU utilization (averaged) stays below 15%
- Real workload - Kieker monitoring framework used in production environment:



Results



- pwi_4 performs best in the comparison of the four alternative methods (26.46%, 29.15% for log.).
- Single execution environment monitored $\Rightarrow pwi_2 = pwi_3$.
- 0% benefit was accounted for several operations with too few observations.

- 1 Motivation
- 2 Foundations
- 3 Workload-intensity-sensitive timing behavior analysis
- 4 Empirical evaluation
- 5 Related work**
- 6 Conclusions and future work

Motivation

Foundations

Approach

Evaluation

Related Work

Conclusions
and future
work

- ? : Requests are grouped by request complexity.
- ? : Workload intensity changes related to the day time are used in network data analysis.
- ? : Requests are grouped according to resource usage.
- ? : Control-flow (Caller context).
- ? : Control-flow (Stack content).
- ? : Control-flow (Trace context).

Motivation

Foundations

Approach

Evaluation

Related Work

Conclusions
and future
work

- 1 Motivation
- 2 Foundations
- 3 Workload-intensity-sensitive timing behavior analysis
- 4 Empirical evaluation
- 5 Related work
- 6 Conclusions and future work**

Motivation

Foundations

Approach

Evaluation

Related Work

**Conclusions
and future
work**

Approach summary

- Goal: “Reduce” variance for statistical measurement analysis
- Workload-intensity metrics pwi_1 - pwi_4
- Categorization based on workload-intensity
- No additional monitoring requirements

Empirical evaluation results

- Applicability in real, distributed, enterprise software systems
- Observation: A significant part of the variance in timing behavior could be controlled by considering workload intensity.
- pwi_4 (operation specific weights) performed best.
- No big difference between pwi_1 (response times), and pwi_2 (execution times) in the case studies.

- Application in the context of anomaly detection.
- Comparison of the standard deviation reduction with the *pwi* workload-intensity metrics with that resulting from other timing behavior influences, such as parameter values, request types, and control flow context, in standard deviation reduction.
- Comparison of the *pwi* workload-intensity metrics with other workload intensity metrics, such as CPU utilization, load average, and arrival rate.

