

# Visualization of Performance Anomalies with Kieker

## Bachelor's Thesis

Sören Henning

September 8, 2016



1. Introduction
2. Foundations
3. Approach
4. Evaluation
5. Conclusion and Future Work

- ▶ Performance is a significant quality characteristic
  - ▶ e.g., Amazon: 100 ms delay → 1% decrease in sales (Huang 2011)
  - ▶ e.g., Google: 500 ms delay → 20% drop in traffic (Huang 2011)

- ▶ Performance is a significant quality characteristic
  - ▶ e.g., Amazon: 100 ms delay → 1% decrease in sales (Huang 2011)
  - ▶ e.g., Google: 500 ms delay → 20% drop in traffic (Huang 2011)
- ▶ Detect performance issues to react on them
  - ▶ As soon as possible
  - ▶ Use monitoring (e.g., measure execution times)
  - ▶ Investigate these measurements for anomalies

- ▶ Performance is a significant quality characteristic
  - ▶ e.g., Amazon: 100 ms delay → 1% decrease in sales (Huang 2011)
  - ▶ e.g., Google: 500 ms delay → 20% drop in traffic (Huang 2011)
- ▶ Detect performance issues to react on them
  - ▶ As soon as possible
  - ▶ Use monitoring (e.g., measure execution times)
  - ▶ Investigate these measurements for anomalies
- ▶ Visualization can support interpretation of anomalies

## ΘPAD and ΘPADx

- ▶ Provide anomaly detection
- ▶ Part of Kieker
- ▶ Only R algorithms
- ▶ Problematic anomaly score
- ▶ No visualization
- ▶ More on this later

Development of an approach to detect performance anomalies with Kieker and visualize them

- ▶ G1: Migrate  $\Theta$ PAD to a TeeTime configuration



- ▶ G1: Migrate  $\Theta$ PAD to a TeeTime configuration
- ▶ G2: Implement multiple forecast algorithms

- ▶ G1: Migrate  $\Theta$ PAD to a TeeTime configuration
- ▶ G2: Implement multiple forecast algorithms
- ▶ G3: Provide a visualization of measured time series and detected anomalies

- ▶ G1: Migrate  $\Theta$ PAD to a TeeTime configuration
- ▶ G2: Implement multiple forecast algorithms
- ▶ G3: Provide a visualization of measured time series and detected anomalies
- ▶ G4: Evaluate the Implementation
  - ▶ G4.1. Feasibility evaluation
  - ▶ G4.2. Scalability evaluation

1. Introduction
2. Foundations
3. Approach
4. Evaluation
5. Conclusion and Future Work

- ▶ Performance Metrics (Koziolek 2008)
  - ▶ Time behavior and resource efficiency
- ▶ Time Series (Mitsa 2010)
  - ▶ Sequence of measurements at regular temporal intervals
- ▶ Anomaly Detection (Chandola, Banerjee, and Kumar 2009)
  - ▶ Anomaly: Abnormal data patterns
  - ▶ Detection: Compare measured values with reference model




Figure: Based on Bielefeld (2012) and Frotscher (2013)

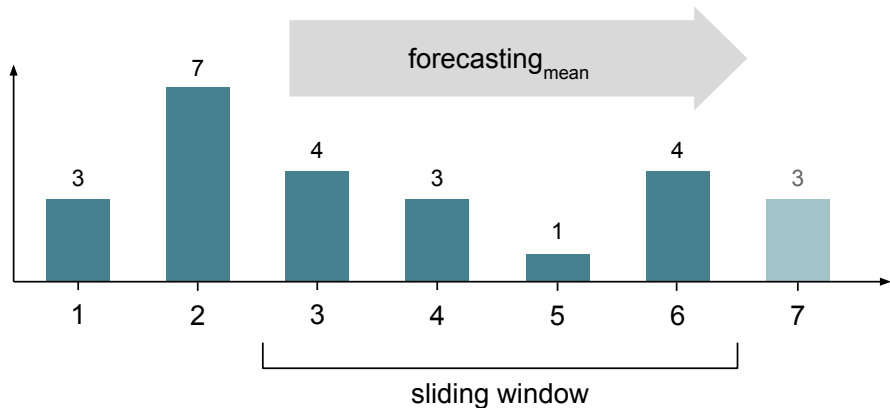
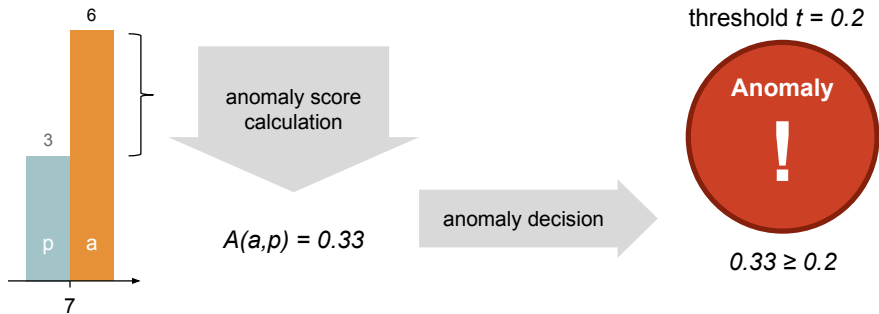


Figure: Based on Bielefeld (2012) and Frotscher (2013)







- ▶ Microservices Architectural Pattern (Wolff 2015)
- ▶ Kieker Monitoring Framework (Hoorn et al. 2009)
- ▶ TeeTime Pipe and Filter Framework (Wulf, Ehmke, and Hasselbring 2014)
- ▶ And further technologies (see next slides)





TeeTime ≡

1. Introduction
2. Foundations
3. Approach
4. Evaluation
5. Conclusion and Future Work



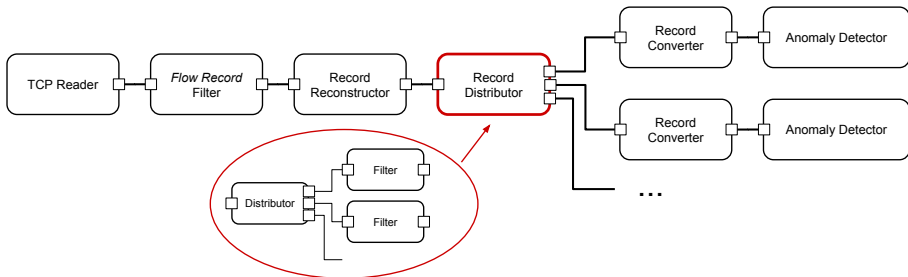







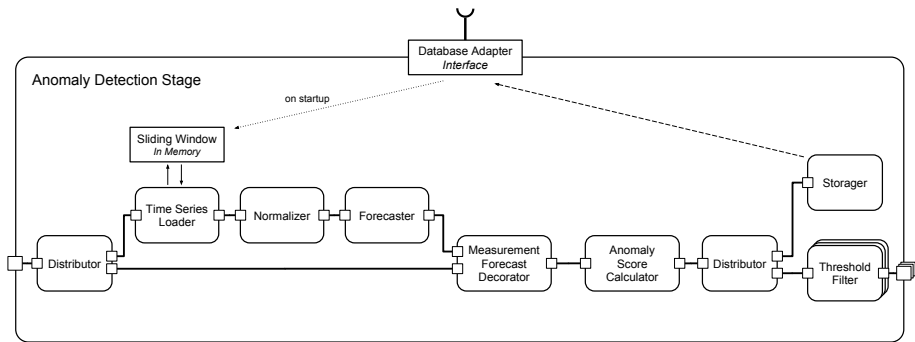
### Filter:

- ▶ Operation signature
- ▶ Class signature
- ▶ Host name
- ▶ ...











KiekPAD

Anomaly Detection

demo-method ▾









- ▶ Usage of Arne Johanson's CanvasPlot (Johanson 2016)

1. Introduction
2. Foundations
3. Approach
- 4. Evaluation**
5. Conclusion and Future Work

# Feasibility Evaluation

## Scenarios



### Evaluation



# Feasibility Evaluation

Scenario: Seasonal with Anomaly

Evaluation



- ARIMAForecaster
- ExponentialWeightedForecaster
- LinearWeightedForecaster
- LogarithmicWeightedForecaster
- MeanForecaster
- RegressionForecaster

- ▶ Take time for record processing in analysis

- ▶ Take time for record processing in analysis
- ▶ Evaluate: Execution time  $\leq$  measurement frequency ?



- ▶ Take time for record processing in analysis
- ▶ Evaluate: Execution time  $\leq$  measurement frequency ?
- ▶ For all parameter combinations:

**Measurement frequencies** 2 ms, 5 ms, 10 ms, 50 ms, 100 ms, 150 ms, 200 ms

**Sliding window** 10,000 ms, 50,000 ms, 100,000 ms, 150,000 ms, 200,000 ms, 400,000 ms

**Normalization interval** 10 ms, 20 ms, 100 ms, 200 ms, 500 ms, 1000 ms, 2000 ms

**Forecast algorithm** ARIMAForecaster, RegressionForecaster

**Normalization algorithm** MeanAggregator

## Some examples:

freq.	sld. window	norm. intvl.	forecaster	∅ exec. time
5	10,000	200	Regression	1.64
5	400,000	20	Regression	4.35
50	50,000	200	ARIMA	69.98
100	100,000	500	ARIMA	78.24
150	200,000	2,000	ARIMA	187.21

...

all values in ms

1. Introduction
2. Foundations
3. Approach
4. Evaluation
5. Conclusion and Future Work

- ▶ Further development of the  $\Theta$ PAD Approach
  - ▶ Proving infrastructure via Docker containers
  - ▶ Immediately record processing

- ▶ Further development of the  $\Theta$ PAD Approach
  - ▶ Proving infrastructure via Docker containers
  - ▶ Immediately record processing
- ▶ Native Java algorithms for anomaly detection

- ▶ Further development of the  $\Theta$ PAD Approach
  - ▶ Proving infrastructure via Docker containers
  - ▶ Immediately record processing
- ▶ Native Java algorithms for anomaly detection
- ▶ Providing an interactive, real time visualization

- ▶ Further development of the  $\Theta$ PAD Approach
  - ▶ Proving infrastructure via Docker containers
  - ▶ Immediately record processing
- ▶ Native Java algorithms for anomaly detection
- ▶ Providing an interactive, real time visualization
- ▶ All implementations available as open source:  
[github.com/SoerenHenning](https://github.com/SoerenHenning)

- ▶ Handling of fast incoming measurements
  - ▶ Aggregate before analysis ( $\Theta$ PAD)
  - ▶ Cache time series operations



- ▶ Handling of fast incoming measurements
  - ▶ Aggregate before analysis ( $\Theta$ PAD)
  - ▶ Cache time series operations
- ▶ Parallelized and distributed analysis
  - ▶ Is or will be supported by TeeTime

- ▶ Handling of fast incoming measurements
  - ▶ Aggregate before analysis ( $\Theta$ PAD)
  - ▶ Cache time series operations
- ▶ Parallelized and distributed analysis
  - ▶ Is or will be supported by TeeTime
- ▶ Take advantage of Cassandra's features for data storage

- ▶ Handling of fast incoming measurements
  - ▶ Aggregate before analysis ( $\Theta$ PAD)
  - ▶ Cache time series operations
- ▶ Parallelized and distributed analysis
  - ▶ Is or will be supported by TeeTime
- ▶ Take advantage of Cassandra's features for data storage
- ▶ Configuration via Rest/GUI

- Bielefeld, Tillmann Carlos (2012). “Online Performance Anomaly Detection for Large-Scale Software Systems”. *Diploma Thesis*. Germany: Kiel University.
- Chandola, Varun, Arindam Banerjee, and Vipin Kumar (2009). “Anomaly detection: A survey”. In: *ACM computing surveys (CSUR)* 41.3, p. 15.
- Frotscher, Tom (2013). “Architecture-Based Multivariate Anomaly Detection for Software Systems”. *Master’s Thesis*. Germany: Kiel University.
- Hoorn, André van et al. (2009). *Continuous Monitoring of Software Services: Design and Application of the Kieker Framework*. Tech. rep. TR-0921. Department of Computer Science, Kiel University, Germany.

- Huang, Cheng (2011). *Public DNS System and Global Traffic Management*. Accessed: 2016-08-25. URL: <http://research.microsoft.com/en-us/um/people/chengh/slides/pubdns11.pptx.pdf>.
- Johanson, Arne (2016). *CanvasPlot*. Accessed: 2016-08-25. URL: <https://a-johanson.github.io/canvas-plot>.
- Koziolek, Heiko (2008). “Dependability Metrics”. In: ed. by Irene Eusgeld, Felix C. Freiling, and Ralf Reussner. Berlin, Heidelberg: Springer-Verlag. Chap. Introduction to Performance Metrics, pp. 199–203. ISBN: 3-540-68946-X, 978-3-540-68946-1.
- Mitsa, T. (2010). *Temporal Data Mining*. Chapman & Hall/CRC Data Mining and Knowledge Discovery Series. CRC Press. ISBN: 9781420089776.



Wolff, E. (2015). *Microservices: Grundlagen flexibler Softwarearchitekturen*. Dpunkt.Verlag GmbH. ISBN: 9783864903137.

Wulf, Christian, Nils Christian Ehmke, and Wilhelm Hasselbring (2014). "Toward a Generic and Concurrency-Aware Pipes & Filters Framework". In: *Symposium on Software Performance 2014: Joint Descartes/Kieker/Palladio Days*.

# Feasibility Evaluation

## Scenario: Constant with Anomaly

### Feasibility Evaluation





- ARIMAForecaster
- LinearWeightedForecaster
- MeanForecaster
- ExponentialWeightedForecaster
- LogarithmicWeightedForecaster
- RegressionForecaster

# Feasibility Evaluation

## Scenario: Constant with Anomaly - Detail

Feasibility Evaluation





- ARIMAForecaster
- ExponentialWeightedForecaster
- LinearWeightedForecaster
- LogarithmicWeightedForecaster
- MeanForecaster
- RegressionForecaster



# Feasibility Evaluation

Scenario: Linearly Increasing with Anomaly

Feasibility Evaluation





ARIMAForecaster      LinearWeightedForecaster      MeanForecaster  
ExponentialWeightedForecaster      LogarithmicWeightedForecaster      RegressionForecaster

# Feasibility Evaluation

## Scenario: Linearly Increasing with Anomaly - Detail

### Feasibility Evaluation





- ARIMAForecaster
- ExponentialWeightedForecaster
- LinearWeightedForecaster
- LogarithmicWeightedForecaster
- MeanForecaster
- RegressionForecaster

# Feasibility Evaluation

Scenario: Seasonal with Anomaly

Feasibility Evaluation





— ARIMAForecaster      — LinearWeightedForecaster      — MeanForecaster  
— ExponentialWeightedForecaster      — LogarithmicWeightedForecaster      — RegressionForecaster

# Feasibility Evaluation

## Scenario: Seasonal with Anomaly - Detail

### Feasibility Evaluation





- ARIMAForecaster
- LinearWeightedForecaster
- MeanForecaster
- ExponentialWeightedForecaster
- LogarithmicWeightedForecaster
- RegressionForecaster

# Feasibility Evaluation

## Scenario: Exponential Increasing

### Feasibility Evaluation





- ARIMAForecaster
- LinearWeightedForecaster
- MeanForecaster
- ExponentialWeightedForecaster
- LogarithmicWeightedForecaster
- RegressionForecaster

# Feasibility Evaluation

## Scenario: Exponential Increasing - Detail

### Feasibility Evaluation




- ARIMAForecaster
- LinearWeightedForecaster
- MeanForecaster
- ExponentialWeightedForecaster
- LogarithmicWeightedForecaster
- RegressionForecaster

KiekPAD

Anomaly Detection

demo-method ▾



KiekPAD

Anomaly Detection

demo-method ▾

Predictions

Anomaly Scores

Thresholds

-0.3

0.3

Refresh Interval

500

