


# Modeling Variations in Load Intensity over Time


Third International Workshop on Large-Scale Testing (LT 2014), co-located with ICPE'14

Jóakim v. Kistowski, Nikolas Herbst, Samuel Kounev | March 22nd, 2014


DESCARTES RESEARCH GROUP




## ■ Page Requests for the German Wikipedia




## ■ Page Requests for the German Wikipedia



## ■ Page Requests for the German Wikipedia



# Motivation



Additive decomposition into seasonal part, trend, and remainder. Created using *BFAST* [1].

- **Problem:** No means to effectively capture and reproduce the varying load intensity of real-world cloud systems
- **Idea:** Close this gap by creating a meta-model for load intensity variations
- **Benefits:** Enable more precise communication and creation of realistic load scenarios for benchmarking
- **Actions:** Creation of the Descartes Load Intensity Model (DLIM) and High-Level DLIM

- **Problem:** No means to effectively capture and reproduce the varying load intensity of real-world cloud systems
- **Idea:** Close this gap by creating a meta-model for load intensity variations
- **Benefits:** Enable more precise communication and creation of realistic load scenarios for benchmarking
- **Actions:** Creation of the Descartes Load Intensity Model (DLIM) and High-Level DLIM

- **Problem:** No means to effectively capture and reproduce the varying load intensity of real-world cloud systems
- **Idea:** Close this gap by creating a meta-model for load intensity variations
- **Benefits:** Enable more precise communication and creation of realistic load scenarios for benchmarking
- **Actions:** Creation of the Descartes Load Intensity Model (DLIM) and High-Level DLIM



- **Problem:** No means to effectively capture and reproduce the varying load intensity of real-world cloud systems
- **Idea:** Close this gap by creating a meta-model for load intensity variations
- **Benefits:** Enable more precise communication and creation of realistic load scenarios for benchmarking
- **Actions:** Creation of the Descartes Load Intensity Model (DLIM) and High-Level DLIM

## User Behavior Models (Markov Chains)

- van Hoorn et al. (2008) [2]: probabilistic and intensity-varying workloads
- Roy et al. (2013) [3]: workload volatility of a VoD system

## Workload Models


- Barford et al. (1998) [4]: file popularity and distribution (Web)
- Casale et al. (2012) [5]: bursts
- Beich et al. (2010) [6]: data popularity and user separation (Cloud)

## Statistical Model

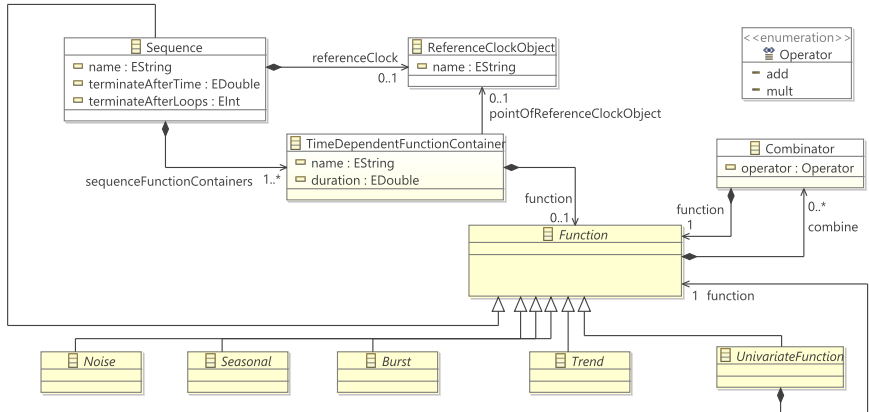
- Feitelson (2002) [7]: workload representativity through statistical characteristics

# Descartes Load Intensity Meta-Model

- Describe arrival rate variations over time
- Provides structure for piece-wise mathematical functions
- Independent of work/request type




# Descartes Load Intensity Meta-Model (DLIM)



# DLIM Example Instance

- Created using LIMBO <sup>1</sup>
- Contains *Seasonal* part, *Trends*, and *Burst*

- ◆ Sequence DLIM\_example
  - ▲ ◆ Combinator MULT
    - ▲ ◆ Sequence trends
      - ▲ ◆ Time Dependent Function Container trend1
        - ◆ Sin Trend 1.0
      - ▲ ◆ Time Dependent Function Container trend2
        - ◆ Sin Trend 1.7
      - ▲ ◆ Time Dependent Function Container trend2
        - ◆ Sin Trend 1.5
    - ▲ ◆ Combinator ADD
      - ▲ ◆ Sequence burstContainer
        - ◆ Time Dependent Function Container burstOffset
          - ▲ ◆ Time Dependent Function Container burst
            - ◆ Exponential Increase And Decline 8.0
    - ▲ ◆ Time Dependent Function Container day
      - ◆ Sin 4.0




<sup>1</sup>LIMBO: [www.descartes-research.net/tools](http://www.descartes-research.net/tools)


- DLIM Benefits:
  - Powerful
  - Expressive
  - Deriving arrival rates or request time-stamps is easy
  
- DLIM Drawbacks:
  - Instances can become complex
  - Large trees may be unintuitive
  
- Solution: **High-Level DLIM**
  - Few parameters for load intensity variation description
  - Easier to grasp
  - Strictly structured into one *Seasonal*, *Trend*, recurring *Burst*, and *Noise* parts

# HLDLIM Seasonal and Trend parts


## ■ HLDLIM *Seasonal* part:



## ■ HLDLIM *Trend* part:



## ■ HLDLIM *Burst* part:



## ■ HLDLIM *Noise* part:

### Uniform Distribution

- Minimum Noise Rate
- Maximum Noise Rate



# HLDLIM Example Instance

## ■ Seasonal part:


Period:	24
Peaks per Seasonal:	1
Base Arrival Rate:	4
First Peak Arrival Rate:	12

## ■ Burst part:

First Burst Offset:	46
Burst Peak Arrival	
Rate:	8
Burst Width:	4


## ■ Trend part:

- Number of Seasonal Periods within one Trend: 1
- Trend List: 16, 20, 14




# DLIM / HLDLIM Differences

## DLIM:




## HLDLIM:



# Preliminary Evaluation

- Trace: German Wikipedia project-counts
- Fitted using best effort approach

	median error	mean error
<b>DLIM</b>	8.95%	16.08%
<b>HLDLIM</b>	9.64%	16.25%



- Two Meta-Models for load intensity variation description
  - **DLIM**: Powerful and expressive
  - **HLDLIM**: Abstract and concise
- Available as part of LIMBO <sup>2</sup>
- Enables creation of custom load intensity variations for open workload based benchmarking
- Capable of accurately capturing real-world load intensity variations



<sup>2</sup>LIMBO: [www.descartes-research.net/tools](http://www.descartes-research.net/tools)

# Thank you for your Interest!

- Which parameters are missing in HLDLIM, which are superfluous?
- How would you go about using DLIM for run-time resource managements? Which additional use cases do you envision for DLIM?

- LIMBO: [www.descartes-research.net/tools](http://www.descartes-research.net/tools)



# References



J. Verbesselt, R. Hyndman, G. Newnham, and D. Culvenor, "Detecting trend and seasonal changes in satellite image time series," *Remote Sensing of Environment*, vol. 114, no. 1, pp. 106 – 115, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S003442570900265X>



A. van Hoorn, M. Rohr, and W. Hasselbring, "Generating probabilistic and intensity-varying workload for web-based software systems," in *Proceedings of the SPEC international workshop on Performance Evaluation: Metrics, Models and Benchmarks*, ser. SIPEW '08. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 124–143. [Online]. Available: [http://dx.doi.org/10.1007/978-3-540-69814-2\\_9](http://dx.doi.org/10.1007/978-3-540-69814-2_9)



S. Roy, T. Begin, and P. Goncalves, "A complete framework for modelling and generating workload volatility of a vod system," in *Wireless Communications and Mobile Computing Conference (IWCMC), 2013 9th International*, 2013, pp. 1168–1174.



P. Barford and M. Crovella, "Generating representative web workloads for network and server performance evaluation," in *Proceedings of the 1998 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*, ser. SIGMETRICS '98/PERFORMANCE '98. New York, NY, USA: ACM, 1998, pp. 151–160. [Online]. Available: <http://doi.acm.org/10.1145/277851.277897>



G. Casale, A. Kalbasi, D. Krishnamurthy, and J. Rolia, "Burn: Enabling workload burstiness in customized service benchmarks," *IEEE Transactions on Software Engineering*, vol. 38, no. 4, pp. 778–793, 2012.



A. Beitch, B. Liu, T. Yung, R. Griffith, A. Fox, and D. A. Patterson, "Rain: A workload generation toolkit for cloud computing applications," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2010-14, Feb 2010. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-14.html>



D. Feitelson, "Workload modeling for performance evaluation," in *Performance Evaluation of Complex Systems: Techniques and Tools*, ser. Lecture Notes in Computer Science, M. Calzarossa and S. Tucci, Eds. Springer Berlin Heidelberg, 2002, vol. 2459, pp. 114–141. [Online]. Available: [http://dx.doi.org/10.1007/3-540-45798-4\\_6](http://dx.doi.org/10.1007/3-540-45798-4_6)