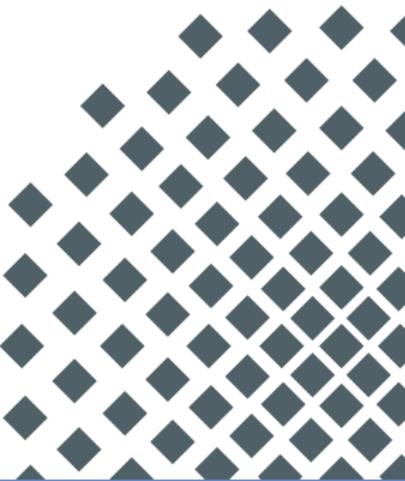


Investigating the Use of Bayesian Networks in the Hora Approach for Component-based Online Failure Prediction

Teerat Pitakrat, André van Hoorn

University of Stuttgart
Institute of Software Technology (ISTE)
Reliable Software Systems (RSS) Group
Stuttgart, Germany

Nov 27, 2014 @ SOSP 2014, Stuttgart





Motivation: Failure Management

[LOG IN or JOIN](#) **The Big Story:** Elon Musk reveals plans for high-speed Hyperloop

THE VERGE FEATURES • REVIEWS • PRODUCTS • SECTIONS • SHOWS • VIDEO • ABOUT • TIP US • FORUMS •

PREVIOUS STORY [Amazon refreshes Kindle Fire Silk browser interface, marking first major update...](#)

NEXT STORY [WikiLeaks source Bradley Manning apologizes: 'I believed I was going to help...'](#)

[APPS & SOFTWARE](#) [MICROSOFT](#) [WEB & SOCIAL](#)

59 COMMENTS

Microsoft recovering from seven-hour long Outlook.com outage

By Tom Warren on August 14, 2013 05:22 pm Email @tomwarren

DON'T MISS STORIES [FOLLOW THE VERGE](#) Like 149k Follow 281K followers



THE LATEST HEADLINES

 Dave Eggers' new novel asks what would happen if Google was truly evil

 China zoo faces uproar after 'African lion' revealed to be a dog

 From Val Kilmer to Mark Twain in 28 photos

 Kinect app that turns any surface into a 'touchscreen' debuts for \$149



Motivation: Failure Management

[LOG IN or JOIN](#) **The Big Story:** Elon Musk reveals plans for high-speed Hyperloop

THE VERGE FEATURES • REVIEWS • PRODUCTS • SECTIONS • SHOWS • VIDEO ABOUT TIP US FORUMS •

PREVIOUS STORY [Amazon refreshes Kindle Fire Silk browser interface, marking first major update...](#)

NEXT STORY [WikiLeaks source Bradley Manning apologizes: 'I believed I was going to help...'](#)

[APPS & SOFTWARE](#) [MICROSOFT](#) [WEB & SOCIAL](#)

59 COMMENTS

Microsoft recovering from seven-hour long Outlook.com outage

By Tom Warren on August 14, 2014 at 09:22 am tomwarren

AWS Server Issues Take Down Instagram, Vine, Airbnb And IFTTT

CHRIS VELAZCO

posted yesterday

49 Comments

Those of you looking to spend the rest of today watching people do other things on Instagram or Vine probably just had a rough time trying to do it. Both services went offline for over an hour, most likely because of issues with [Amazon Web Services](#). Instagram was the first of the two to [publicly acknowledge](#) its issues on Twitter, and Vine followed suit [half an hour later](#). The deluge of tweets that accompanied the services' initial hiccups first started at around 4 p.m. Eastern time, and only increased in intensity as users found they couldn't share pictures of their food or their meticulously crafted video

THE LATEST HEADLINES

Dave Eggers' new novel asks what would happen if Google was truly evil

China zoo faces uproar after 'African lion' revealed to be a dog

From Val Kilmer to Mark Twain in 28 photos

Kinect app that turns any surface into a 'touchscreen' debuts for \$149



Motivation: Failure Management

[LOG IN or JOIN](#) **The Big Story:** Elon Musk reveals plans for high-speed Hyperloop

THE VERGE FEATURES • REVIEWS • PRODUCTS • SECTIONS • SHOWS • VIDEO ABOUT TIP US FORUMS •

PREVIOUS STORY → Amazon refreshes Kindle Fire Silk browser interface, marking first major update...

NEXT STORY ← WikiLeaks source Bradley Manning apologizes: 'I believed I was going to help...'

[APPS & SOFTWARE](#) [MICROSOFT](#) [WEB & SOCIAL](#)

Microsoft recovering long Outlook.com outage

By Tom Warren on August 14, 2014 at 6:22 pm [@tomwarren](#)

AWS Server Issues Take Down Instagram, Airbnb And IFTTT

CHRIS VELAZCO

posted yesterday 49 C

Those of you looking to spend the rest of the day watching people do other things on Vine probably just had a rough time trying to do it. Both services went offline for over a hour, most likely because of issues with Amazon Services.

Instagram was the first of the two to [acknowledge](#) its issues on Twitter, and Vine followed suit [half an hour later](#).

The deluge of tweets that accompanied services' initial hiccups first started at 1 p.m. Eastern time, and only increased in intensity as users found they couldn't share pictures of their food or their meticulously crafted cocktails.

The Register®

Data Centre Columnists Cloud Forums Software Networks Security Policy Business Jobs Hardware Science Mobile Broadband Hosting

NETWORKS

Google goes dark for 2 minutes, kills 40% of world's net traffic

Systemwide outage knocks every service offline

By Neil McAllister, 17th August 2013 336 followers

146 You can all relax now. The near-unprecedented outage that seemingly affected all of Google's services for a brief time on Friday is over.

The event began at approximately 4:37pm Pacific Time and lasted between one and five minutes, according to the [Google Apps Dashboard](#). All of the Google Apps services reported being back online by 4:48pm.

The incident apparently blacked out every service Mountain View has to offer simultaneously, from Google Search to Gmail, YouTube, Google Drive, and beyond.

Big deal, right? Everyone has technical difficulties every once in a while. It goes with the territory.

T. Pitakrat, A. van Hoorn
Investigating the Use of Bayesian Networks in Hora for C-B OFF

Nov. 27, 2014 @ SOSP 2014

2 / 99



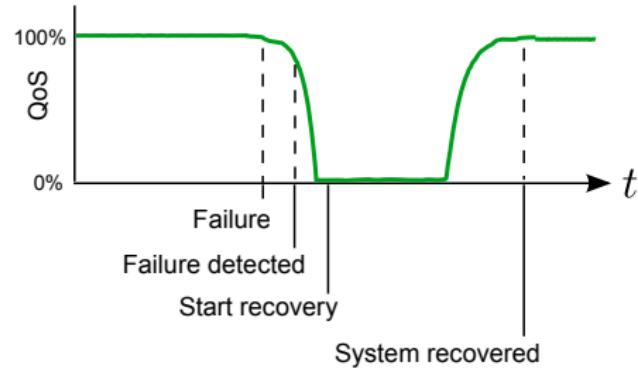
“A **service failure**, often abbreviated here to **failure**, is an event that occurs when the delivered service deviates from correct service.”

— Avizienis et al. [2004]



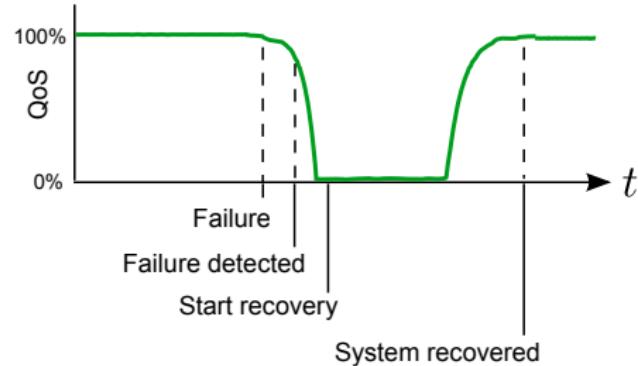
Motivation: Failure Management

Reactive

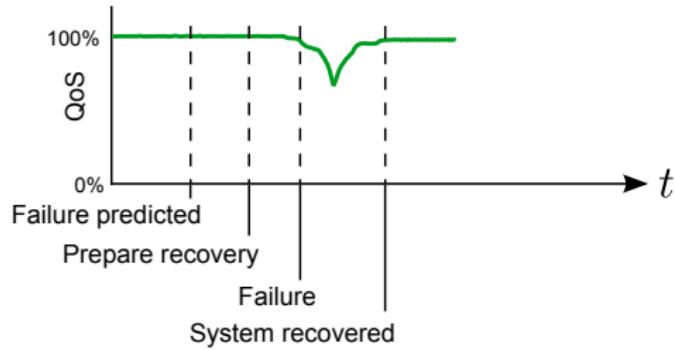


Motivation: Failure Management

Reactive



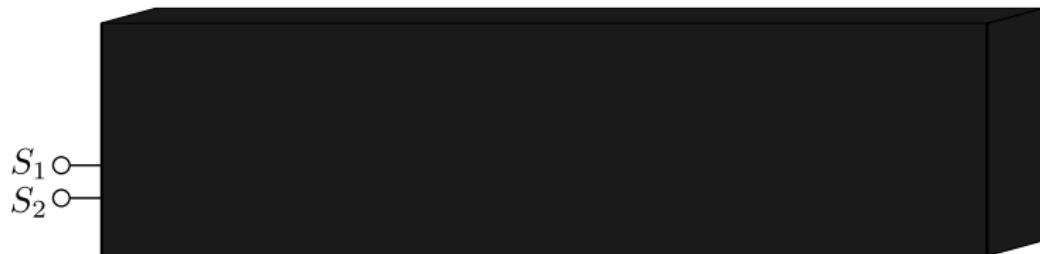
Proactive





- 1 Motivation: Failure Management
- 2 [Recap] Hora: Online Failure Prediction for CB Systems
- 3 Hora: Framework and Implementation
- 4 Evaluation
- 5 Conclusion

[Recap] Hora: Online Failure Prediction for CB Systems

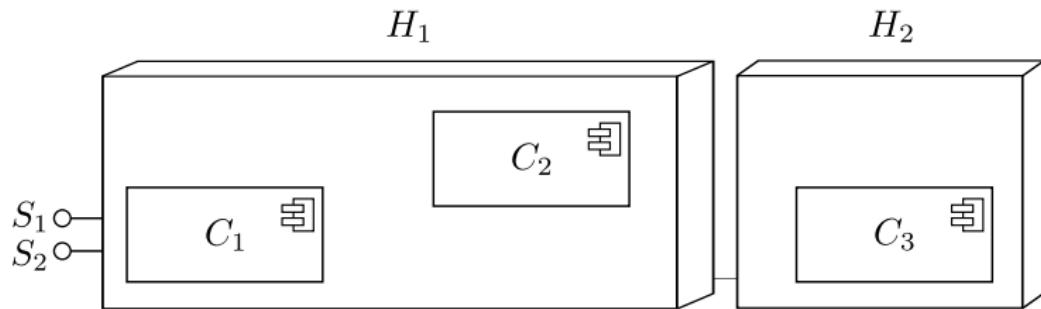


Amin et al. [2012], Liang et al. [2007]

Related Approaches vs. Hora Approach



[Recap] Hora: Online Failure Prediction for CB Systems

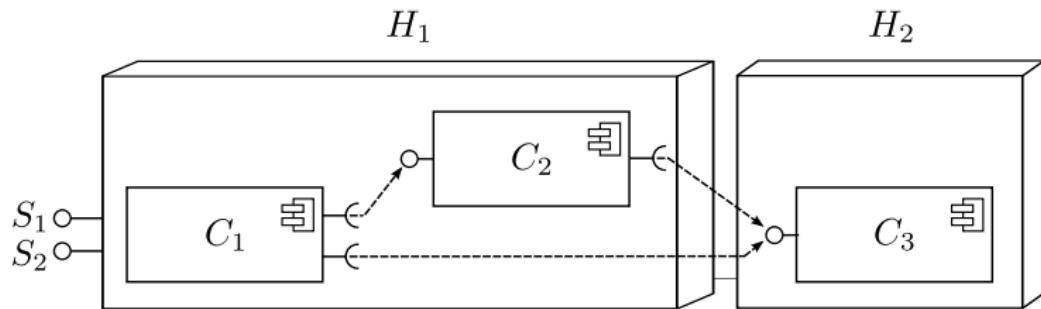


Bielefeld [2012]

Related Approaches vs. Hora Approach



[Recap] Hora: Online Failure Prediction for CB Systems

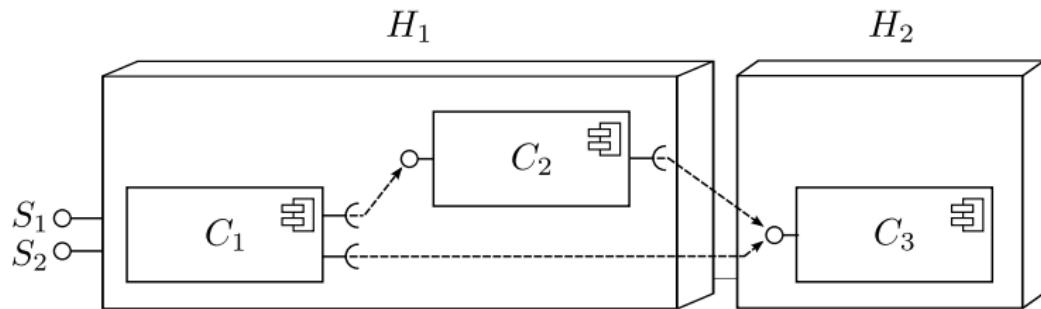


Pitakrat [2013], Pitakrat et al. [2014b]

Related Approaches vs. Hora Approach



[Recap] Hora: Online Failure Prediction for CB Systems



Component-level Prediction Models



Component Dependency

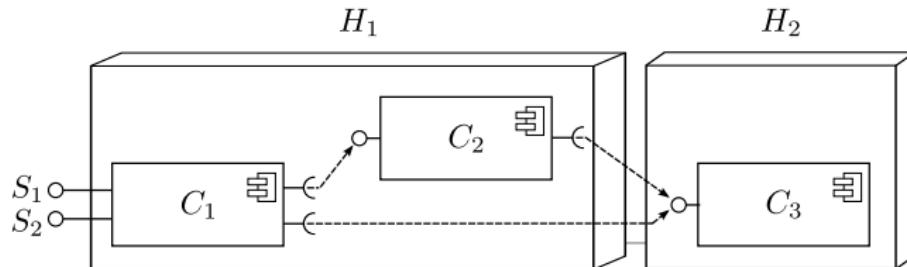


System-level
Prediction Model

Pitakrat [2013], Pitakrat et al. [2014b]

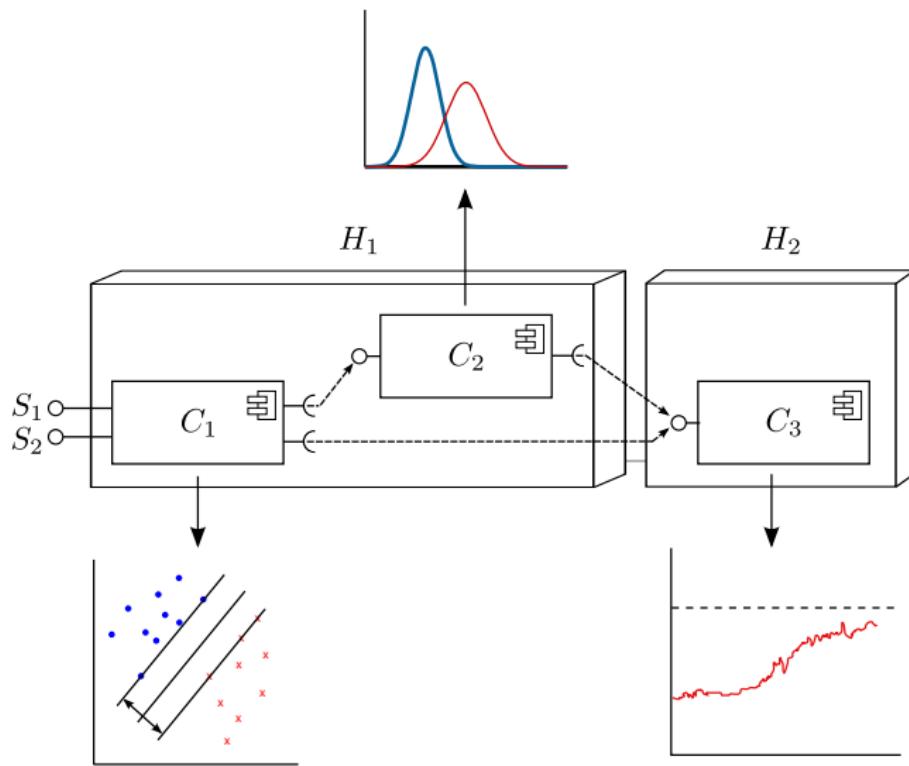


[Recap] Hora: Online Failure Prediction for CB Systems

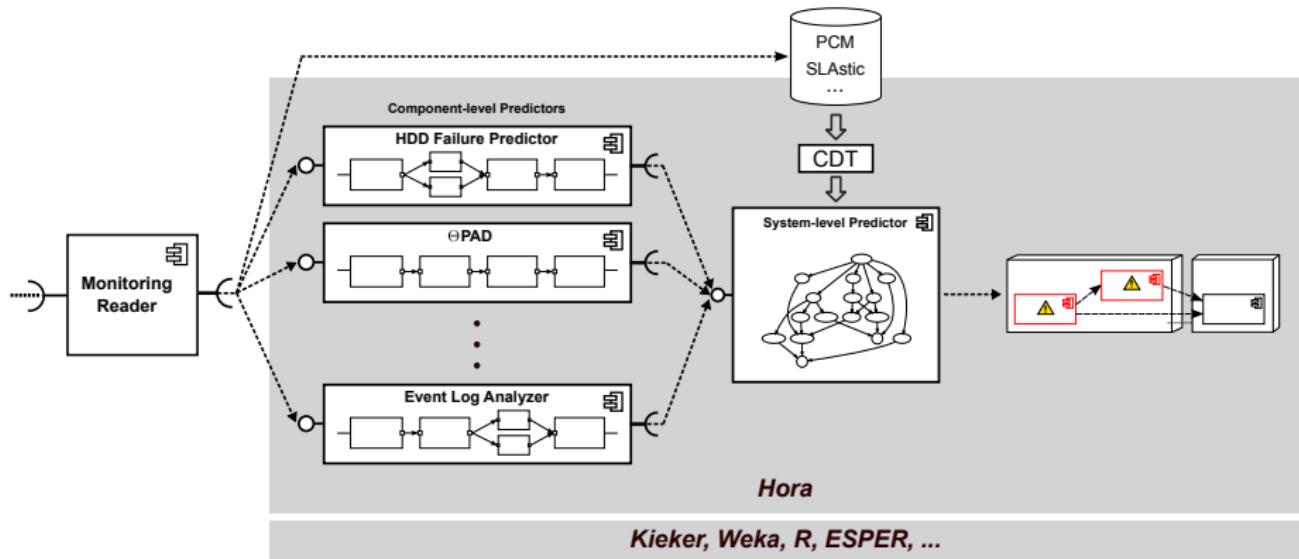




[Recap] Hora: Online Failure Prediction for CB Systems

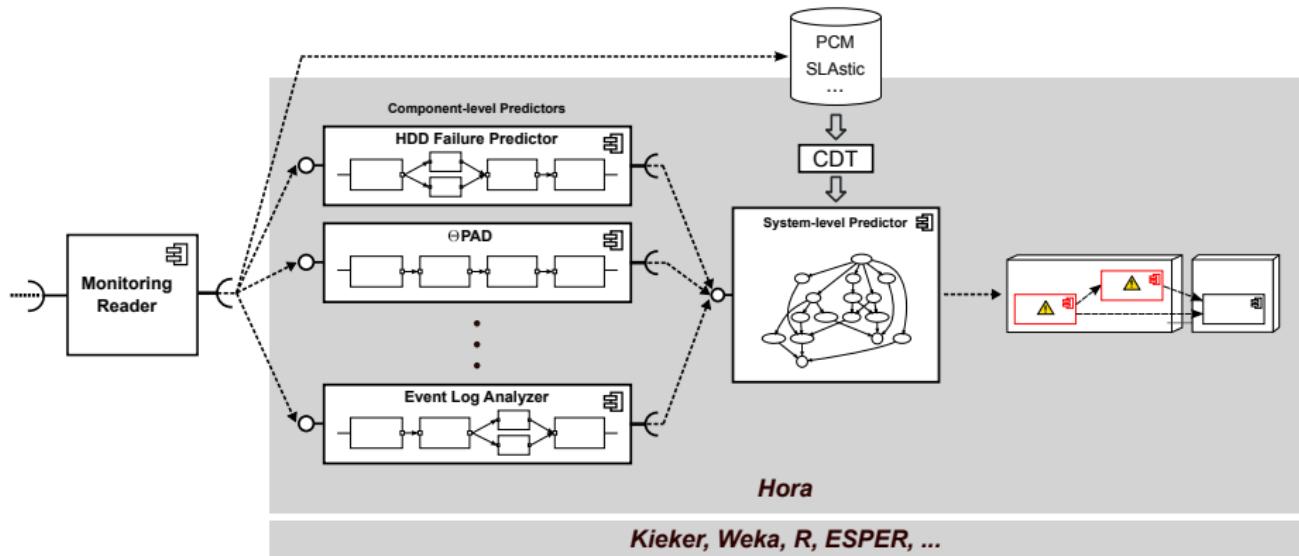


[Recap] Hora: Online Failure Prediction for CB Systems



Becker et al. [2009], Bielefeld [2012], Pitakrat et al. [2013; 2014a], van Hoorn [2014]

[Recap] Hora: Online Failure Prediction for CB Systems



Questions:

- ① What is a suitable model for System-level Prediction Model (SPM)?
- ② How to transform architectural models to CDT and to SPM?
- ③ How does Hora improve online failure prediction?

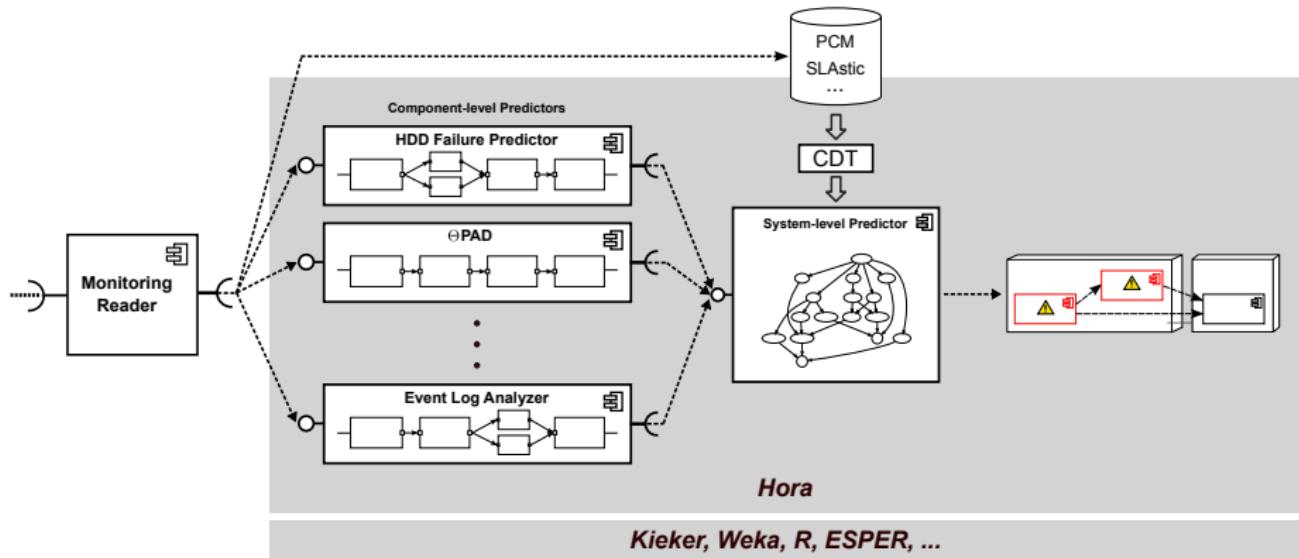
Becker et al. [2009], Bielefeld [2012], Pitakrat et al. [2013; 2014a], van Hoorn [2014]



Hora: Framework and Implementation

- 1 Motivation: Failure Management
- 2 [Recap] Hora: Online Failure Prediction for CB Systems
- 3 Hora: Framework and Implementation
- 4 Evaluation
- 5 Conclusion

Hora: Framework and Implementation



Questions:

- ① What is a suitable model for System-level Prediction Model (SPM)?
- ② How to transform architectural models to CDT and to SPM?
- ③ How does Hora improve online failure prediction?



Hora: Framework and Implementation

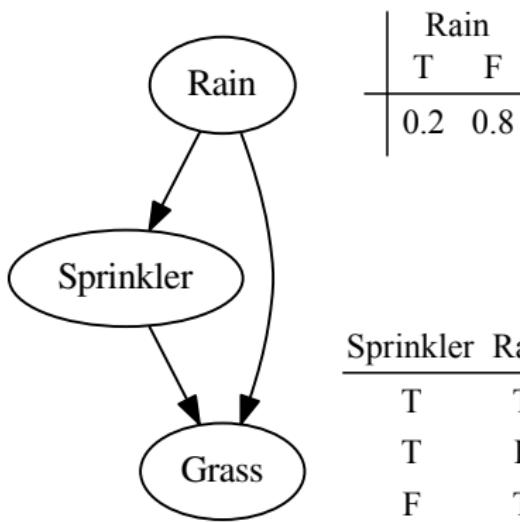
Bayesian network: probabilistic graphical model

Bayesian network library used in Hora: <https://github.com/kutschkem/Jayes>

Hora: Framework and Implementation

Bayesian network: probabilistic graphical model

Rain	Sprinkler	
	T	F
T	0.01	0.99
F	0.4	0.6



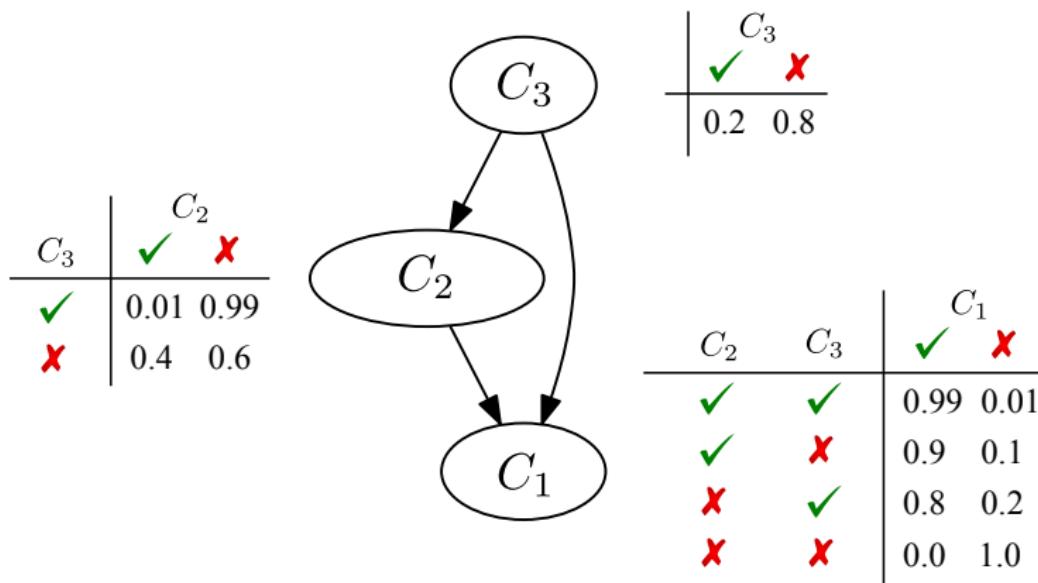
Rain	T	F
	0.2	0.8

Sprinkler	Rain	Grass wet	
		T	F
T	T	0.99	0.01
T	F	0.9	0.1
F	T	0.8	0.2
F	F	0.0	1.0

Bayesian network library used in Hora: <https://github.com/kutschkem/Jayes>

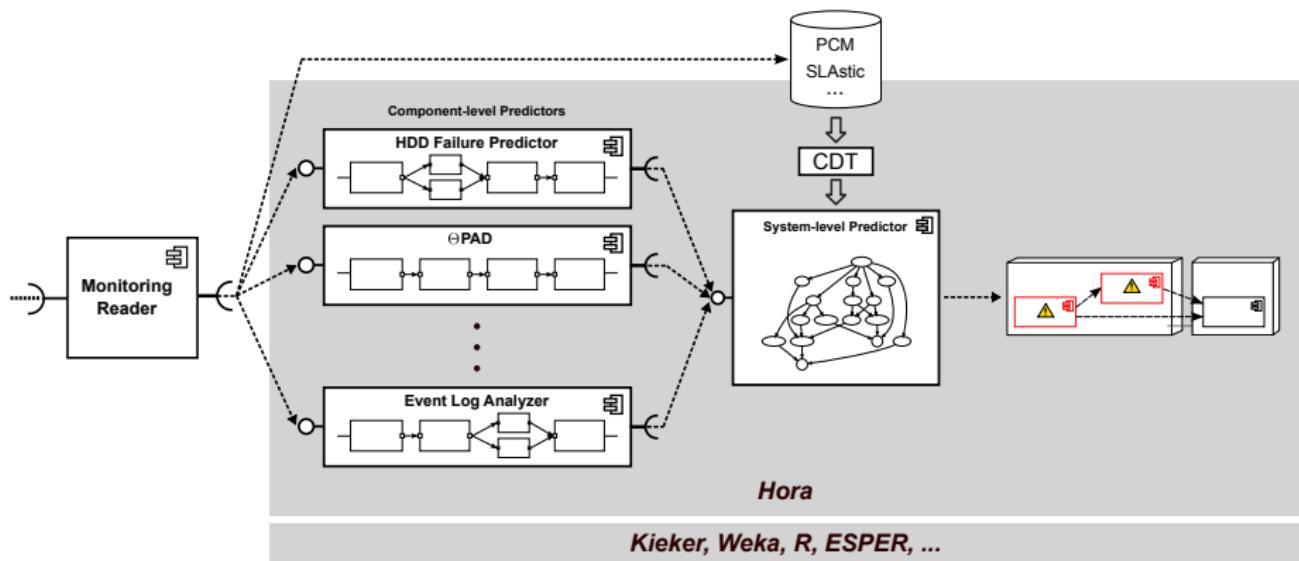
Hora: Framework and Implementation

Bayesian network: probabilistic graphical model



Bayesian network library used in Hora: <https://github.com/kutschkem/Jayes>

Hora: Framework and Implementation



Questions:

- ① What is a suitable model for System-level Prediction Model (SPM)?
- ② **How to transform architectural models to CDT and to SPM?**
- ③ How does Hora improve online failure prediction?

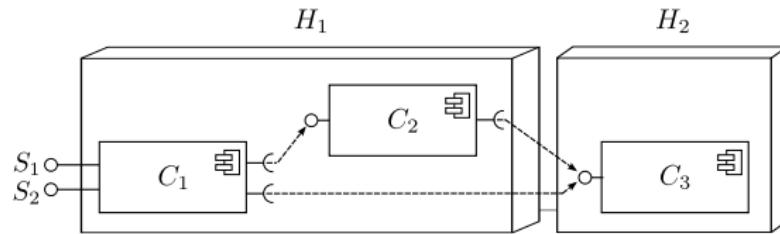
Component Dependency Table

Component Dependencies

Hora: Framework and Implementation



University of Stuttgart
Germany



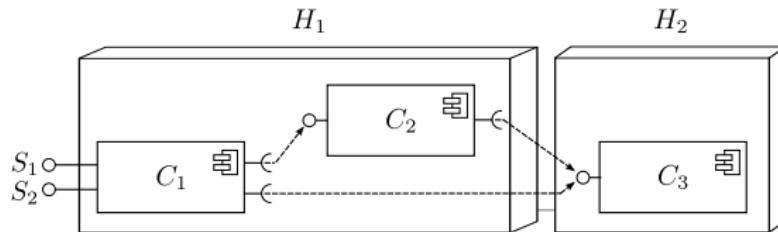
Component Dependency Table

Component Dependencies

Hora: Framework and Implementation



University of Stuttgart
Germany



	C_1	C_2	C_3	H_1	H_2
C_1		•	•	•	
C_2			•	•	
C_3					•
H_1					
H_2					

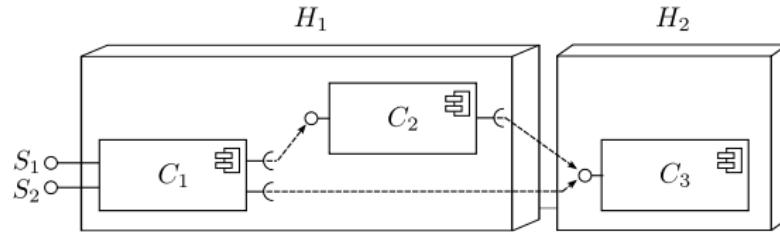
Component Dependency Table

Dependency Calling Probability

Hora: Framework and Implementation



University of Stuttgart
Germany

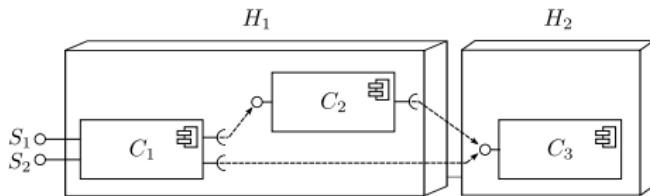


	C_1	C_2	C_3	H_1	H_2
C_1		0.5	0.5	1.0	
C_2			0.5	1.0	
C_3					1.0
H_1					
H_2					

System-level Prediction Model



Hora: Framework and Implementation

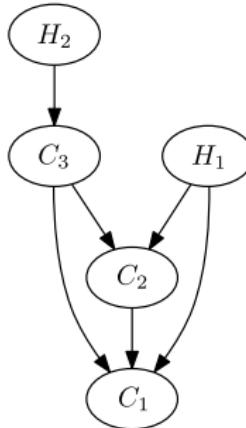
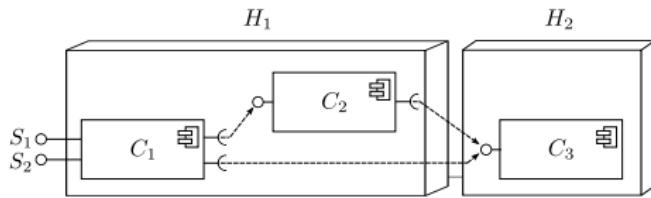


	C_1	C_2	C_3	H_1	H_2
C_1		0.5	0.5	1.0	
C_2			0.5	1.0	
C_3					1.0
H_1					
H_2					

System-level Prediction Model



Hora: Framework and Implementation

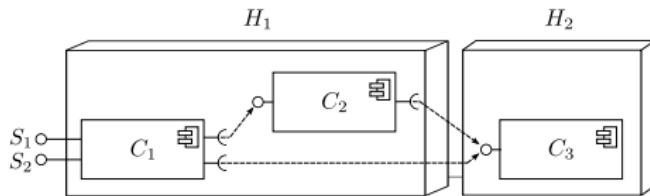


	C_1	C_2	C_3	H_1	H_2
C_1		0.5	0.5	1.0	
C_2			0.5	1.0	
C_3					1.0
H_1					
H_2					

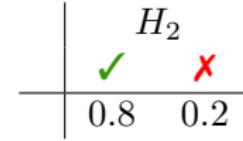
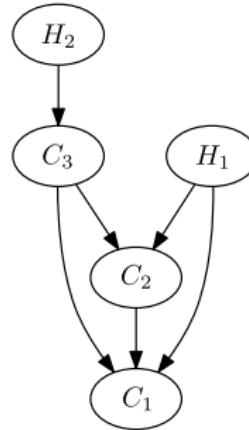
System-level Prediction Model



Hora: Framework and Implementation



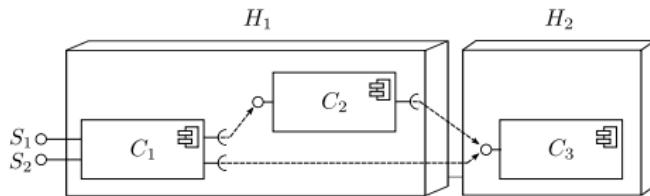
	C_1	C_2	C_3	H_1	H_2
C_1		0.5	0.5	1.0	
C_2			0.5	1.0	
C_3					1.0
H_1					
H_2					



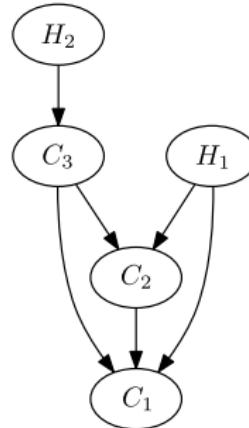
System-level Prediction Model



Hora: Framework and Implementation



	C_1	C_2	C_3	H_1	H_2
C_1		0.5	0.5	1.0	
C_2			0.5	1.0	
C_3					1.0
H_1					
H_2					

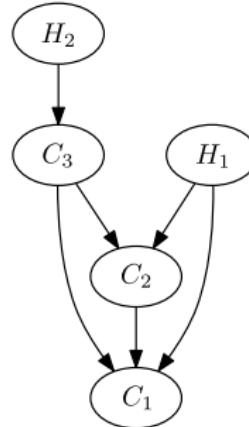
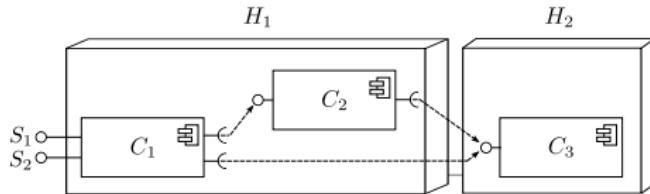


H_2	C_3	
	✓	✗
✓	0.9	0.1
✗	0.0	1.0

System-level Prediction Model



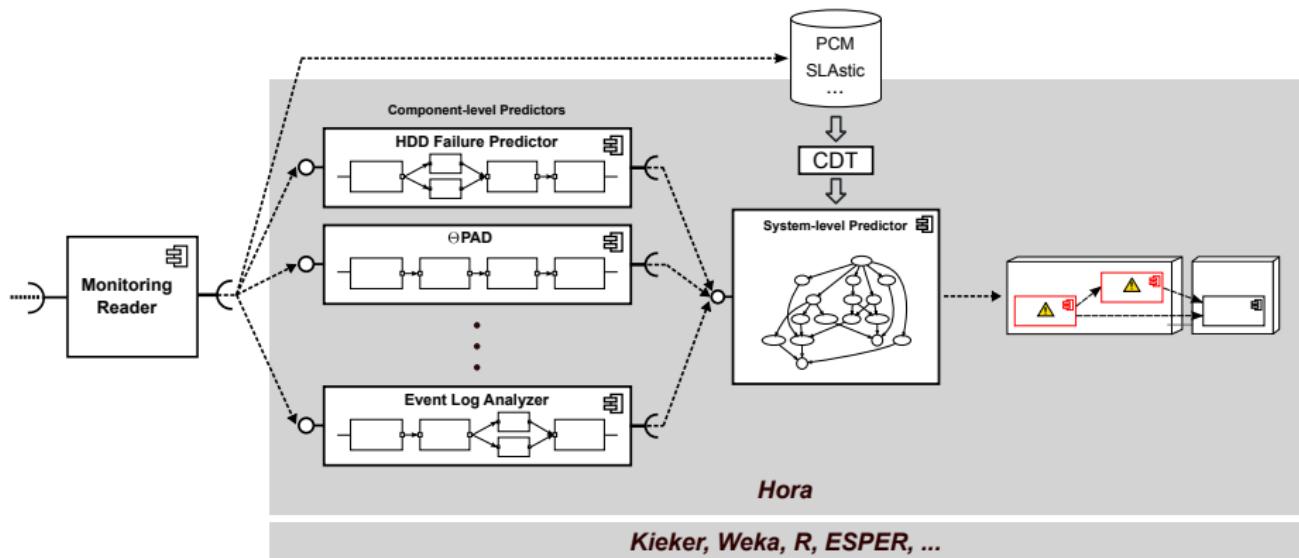
Hora: Framework and Implementation



	C_1	C_2	C_3	H_1	H_2
C_1		0.5	0.5	1.0	
C_2			0.5	1.0	
C_3					1.0
H_1					
H_2					

H_1	C_3	C_2	0.9	0.1
✓	✓	✓	0.5	0.5
✓	✗	✗	0.0	1.0
✗	✓	✗	0.0	1.0
✗	✗	✗	0.0	1.0

Hora: Framework and Implementation



Questions:

- ① What is a suitable model for System-level Prediction Model (SPM)?
- ② How to transform architectural models to CDT and to SPM?
- ③ **How does Hora improve online failure prediction?**



Evaluation

- 1 Motivation: Failure Management
- 2 [Recap] Hora: Online Failure Prediction for CB Systems
- 3 Hora: Framework and Implementation
- 4 Evaluation
- 5 Conclusion

Evaluation

- System under analysis



Evaluation

- System under analysis



- Fault injection



Evaluation

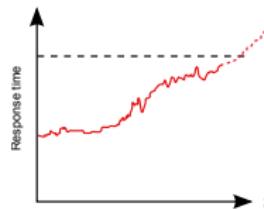
- System under analysis



- Fault injection

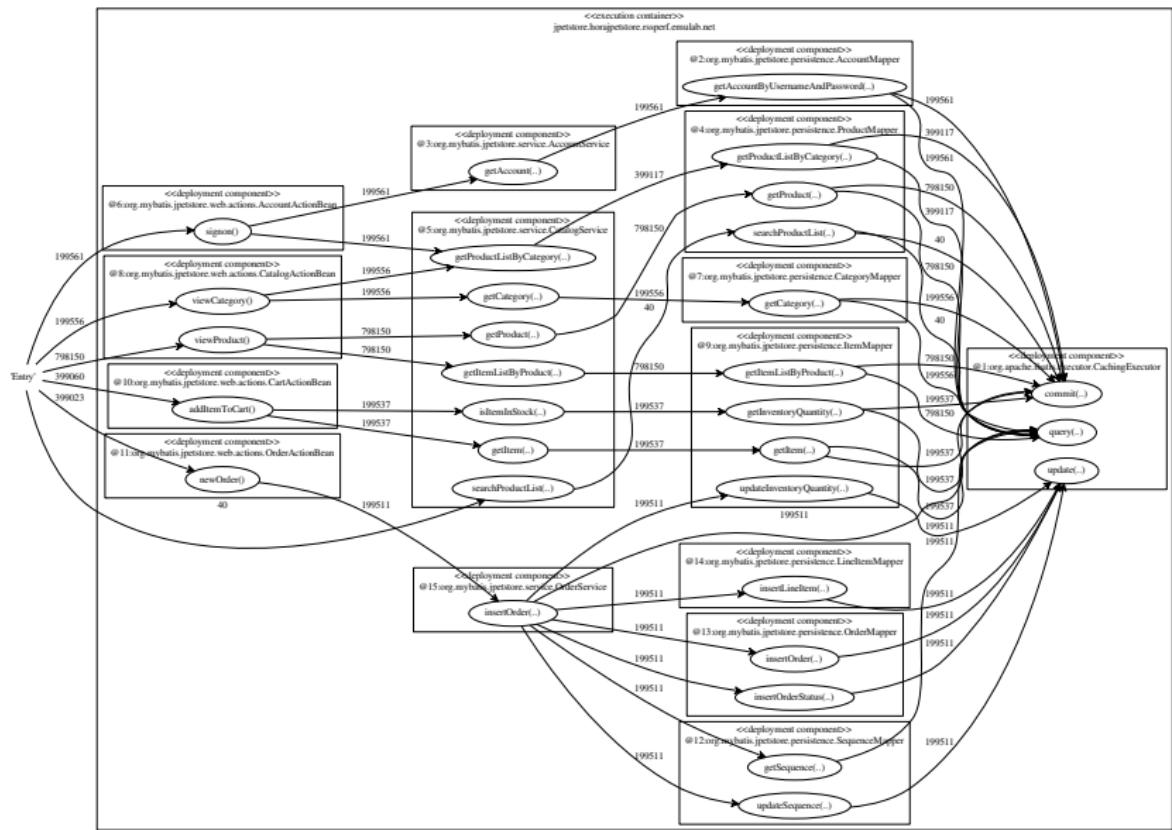


- Predicting response time violation using time series forecasting



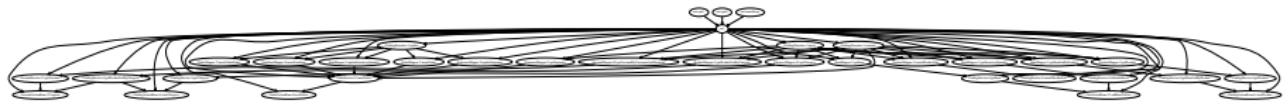
Operation Dependency Graph

Evaluation



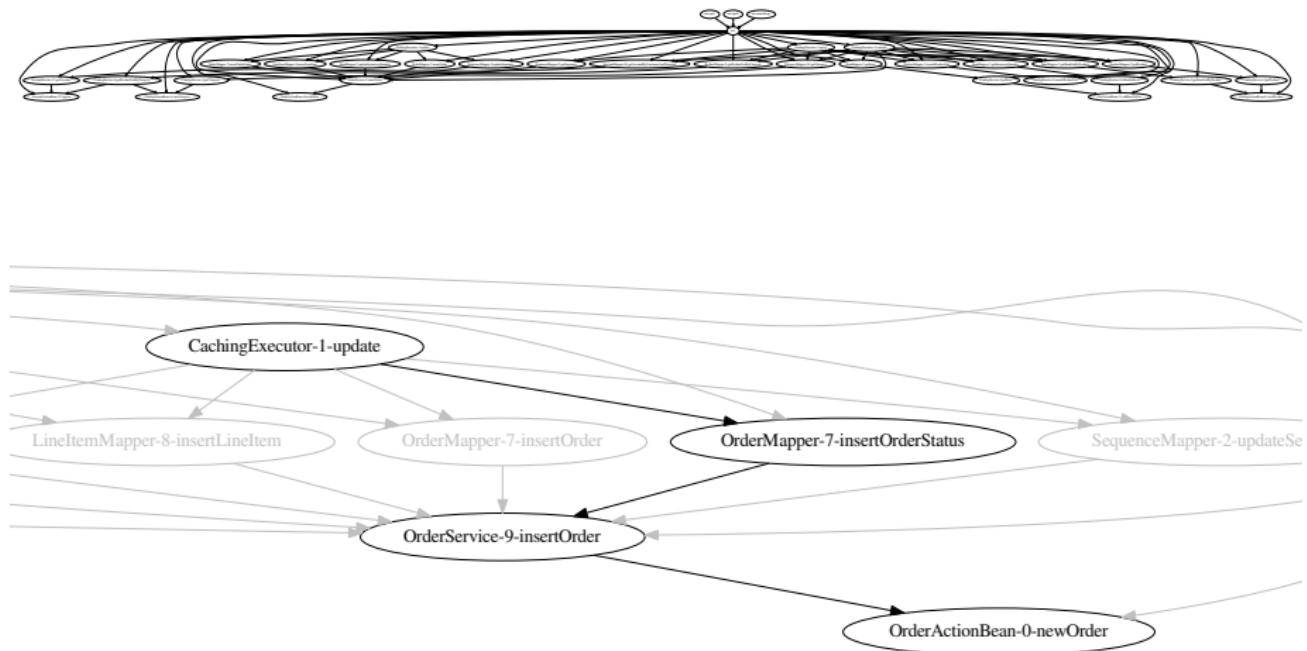


Evaluation





Evaluation



Component Failure Prediction

CachingExecutor-1-update

Evaluation

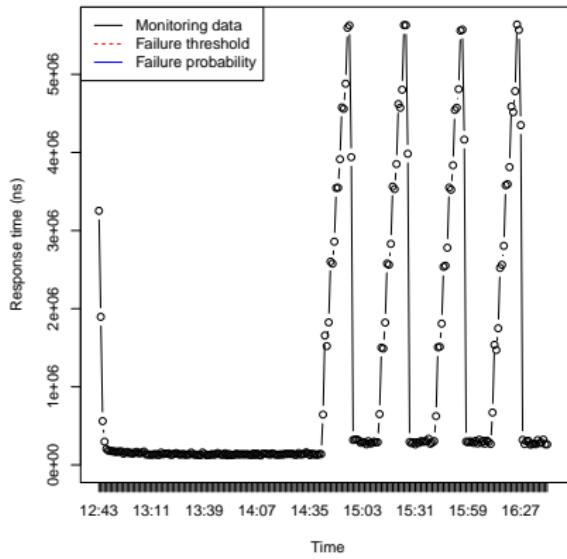


University of Stuttgart
Germany

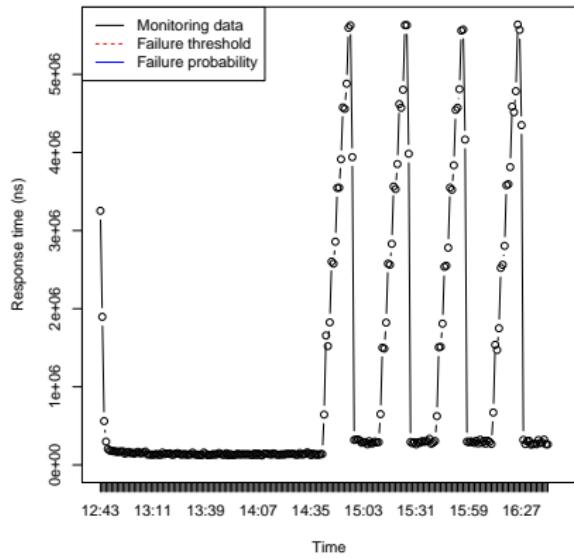
CPM

SPM

CachingExecutor-1-update



CachingExecutor-1-update



Component Failure Prediction

CachingExecutor-1-update

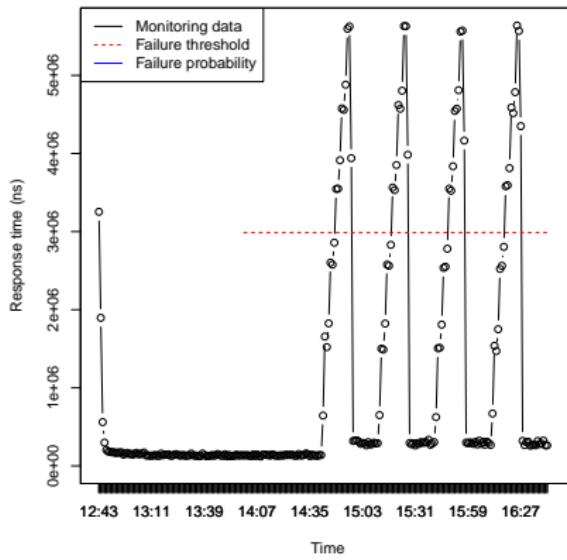
Evaluation



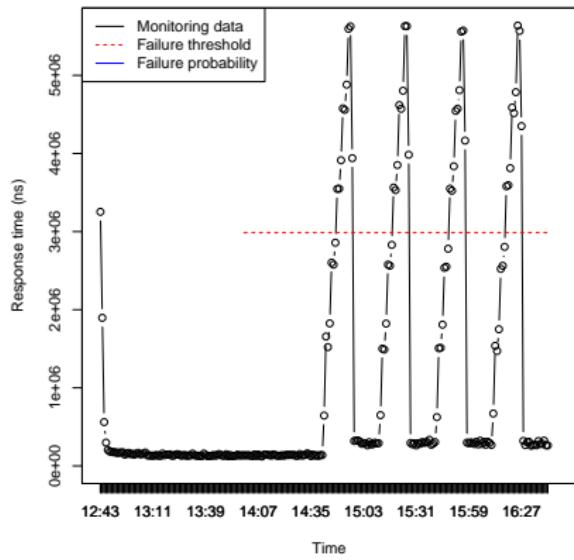
CPM

SPM

CachingExecutor-1-update



CachingExecutor-1-update



Component Failure Prediction

CachingExecutor-1-update

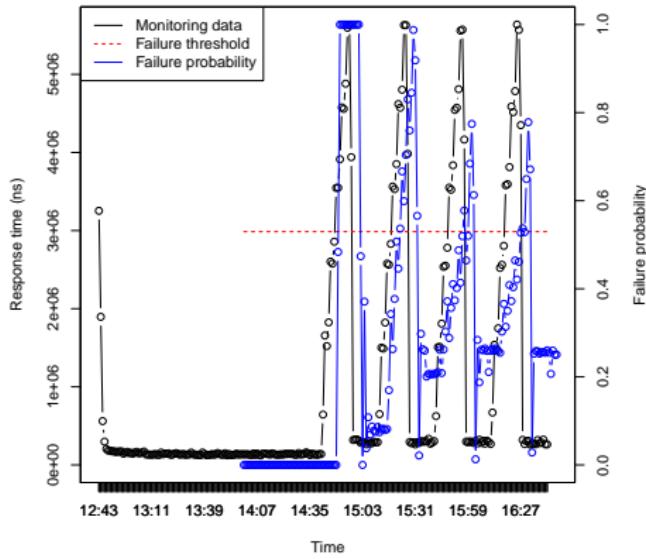
Evaluation



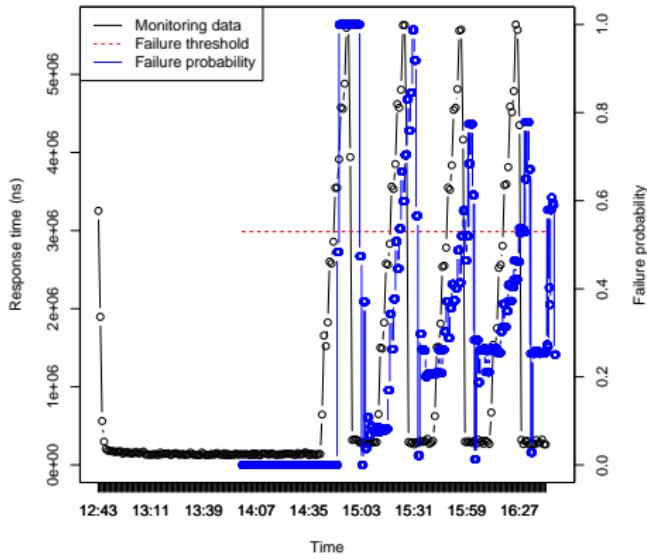
CPM

SPM

CachingExecutor-1-update



CachingExecutor-1-update



Component Failure Prediction

CachingExecutor-1-update

Evaluation

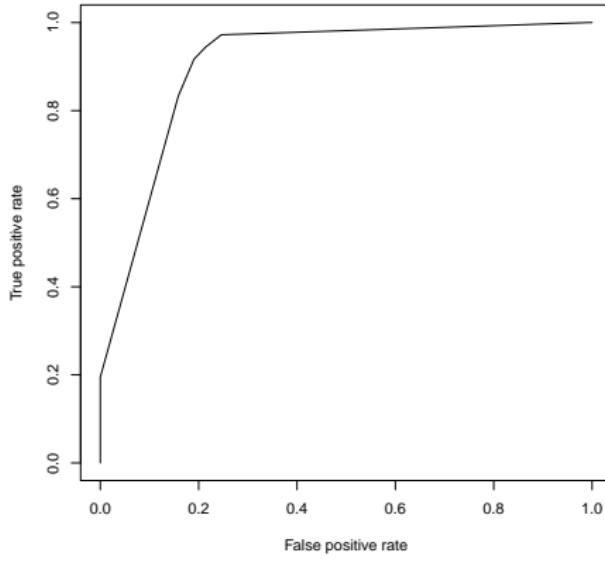


University of Stuttgart
Germany

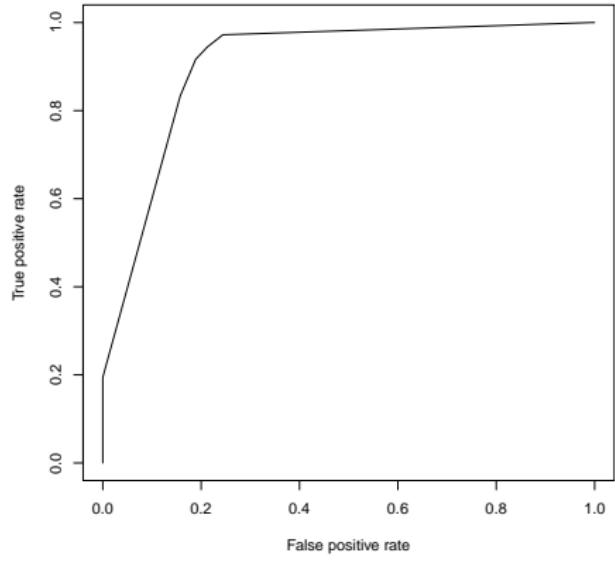
CPM

SPM

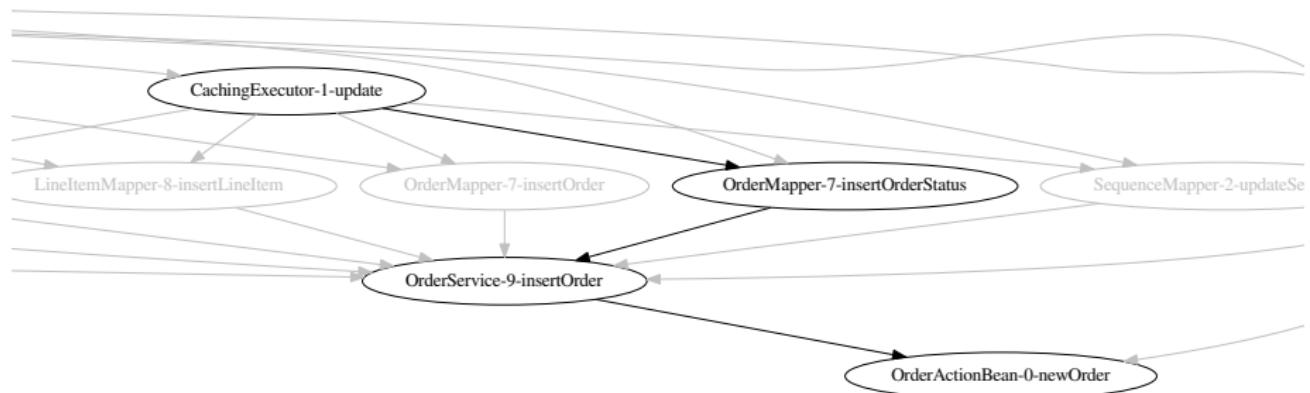
CachingExecutor-1-update ROC Curve



CachingExecutor-1-update ROC Curve



Evaluation



Component vs. System Failure Prediction

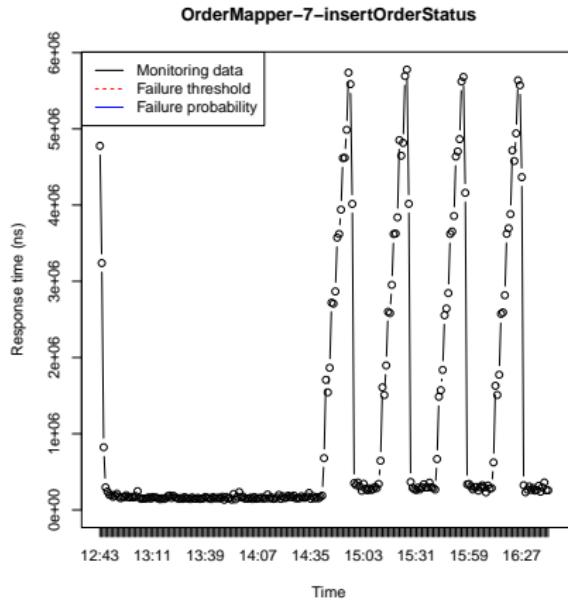


University of Stuttgart
Germany

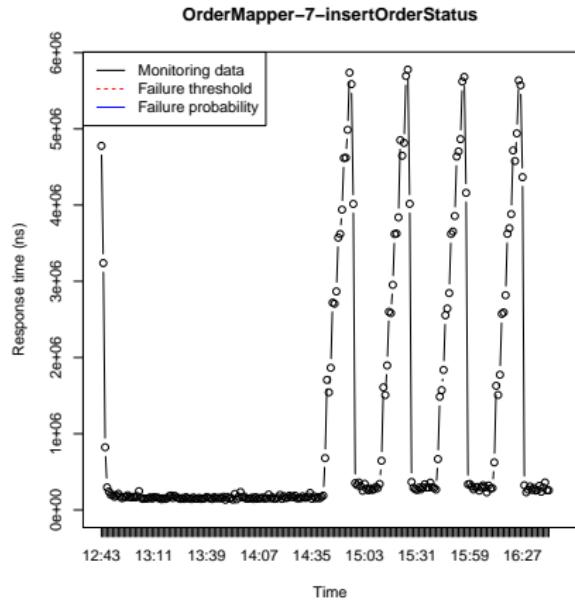
OrderMapper-7-insertOrderStatus

Evaluation

CPM



SPM



Component vs. System Failure Prediction

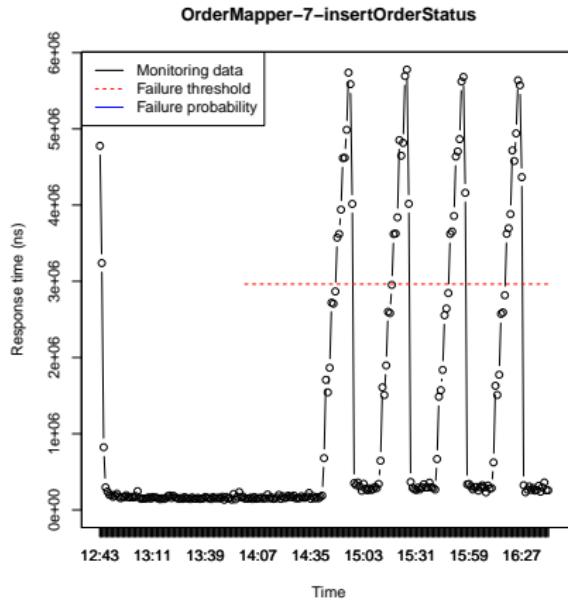


University of Stuttgart
Germany

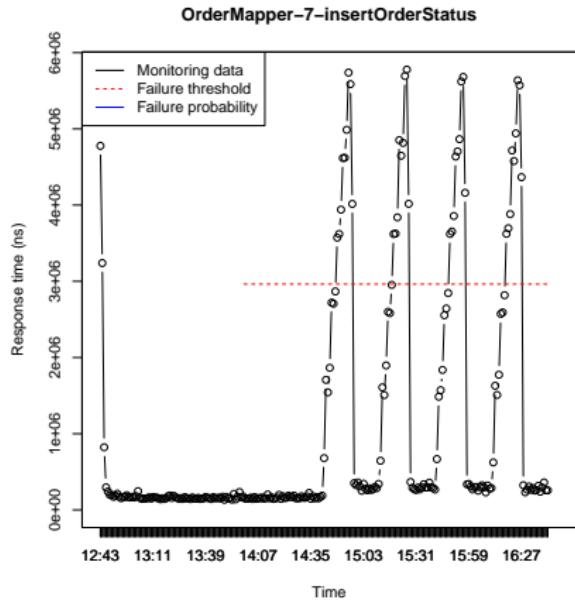
OrderMapper-7-insertOrderStatus

Evaluation

CPM



SPM



Component vs. System Failure Prediction

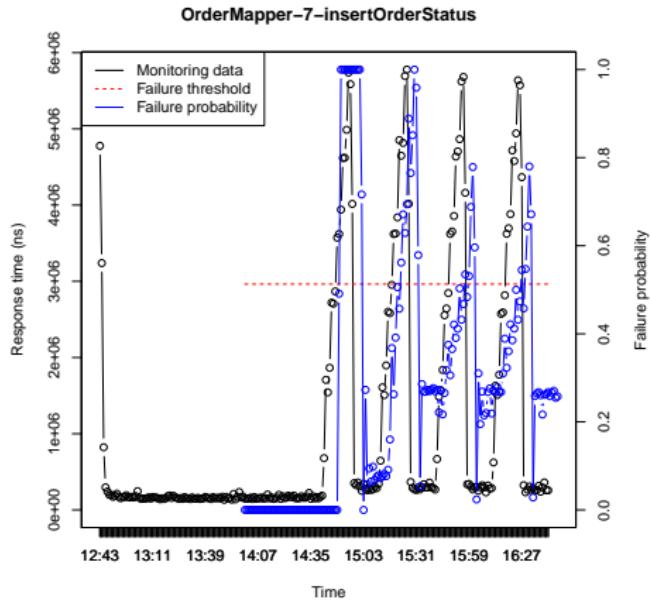


University of Stuttgart
Germany

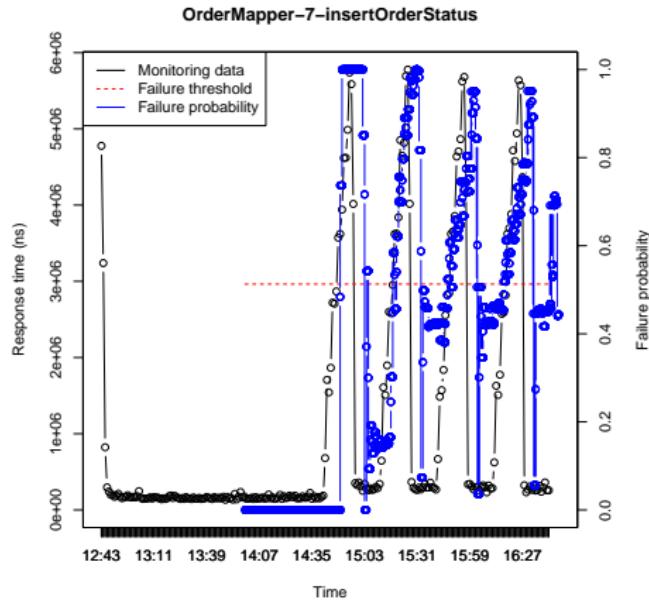
OrderMapper-7-insertOrderStatus

Evaluation

CPM



SPM



Component vs. System Failure Prediction



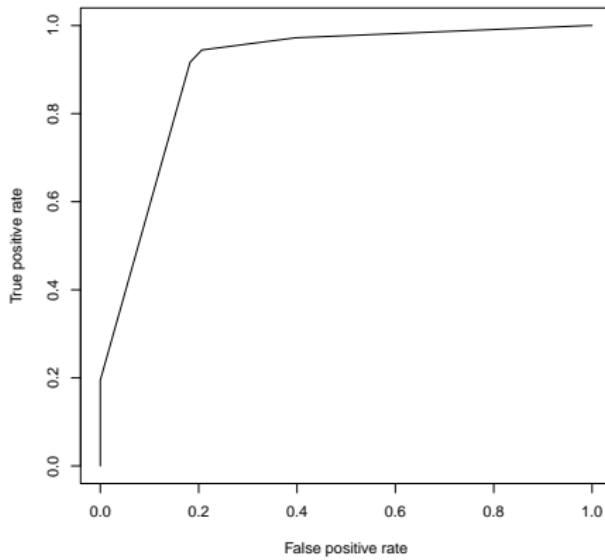
University of Stuttgart
Germany

OrderMapper-7-insertOrderStatus

Evaluation

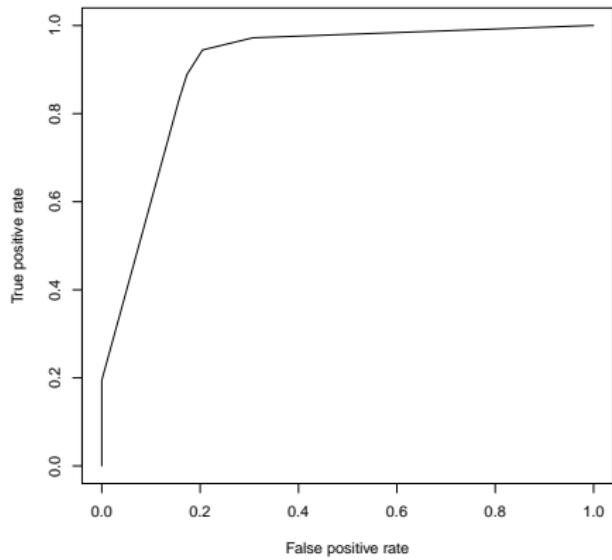
CPM

OrderMapper-7-insertOrderStatus ROC Curve

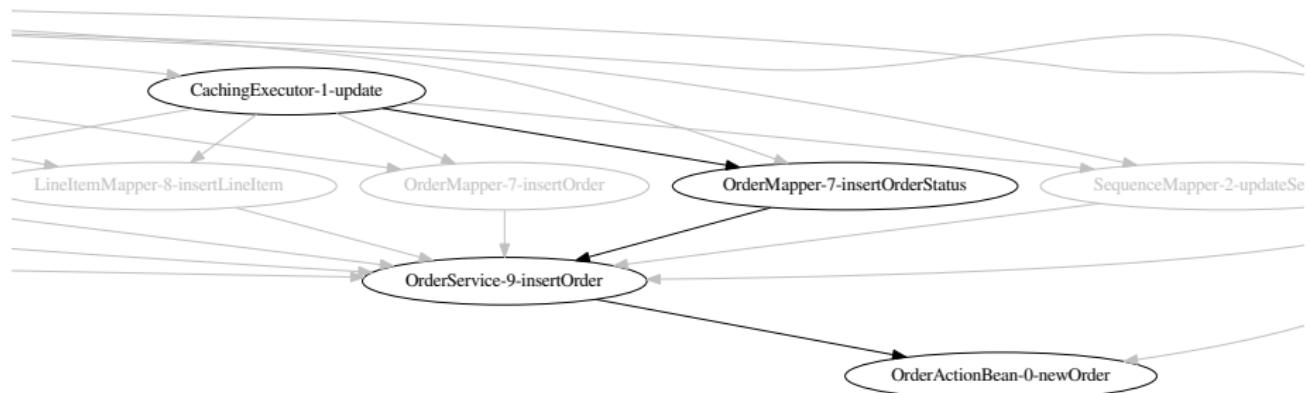


SPM

OrderMapper-7-insertOrderStatus ROC Curve



Evaluation



Component vs. System Failure Prediction

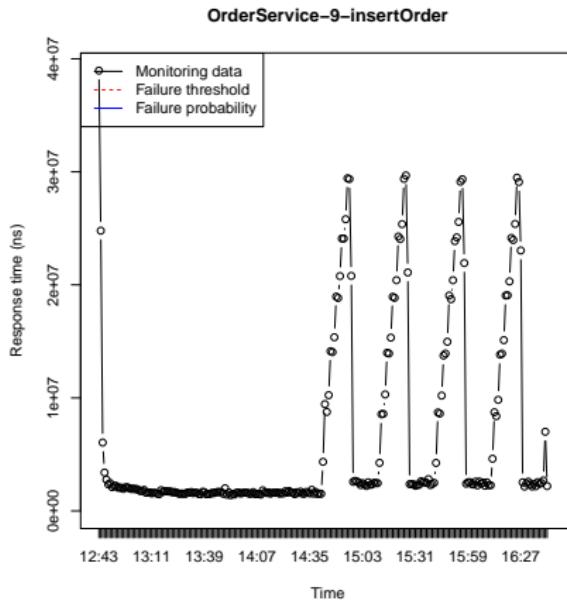


University of Stuttgart
Germany

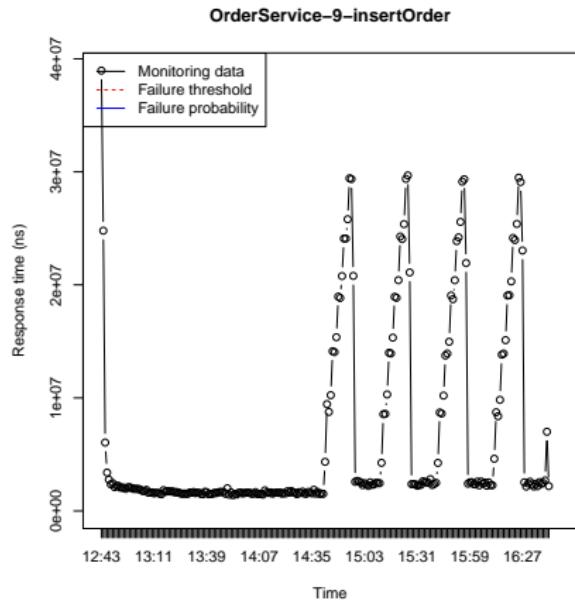
OrderService-9-insertOrder

Evaluation

CPM



SPM



Component vs. System Failure Prediction

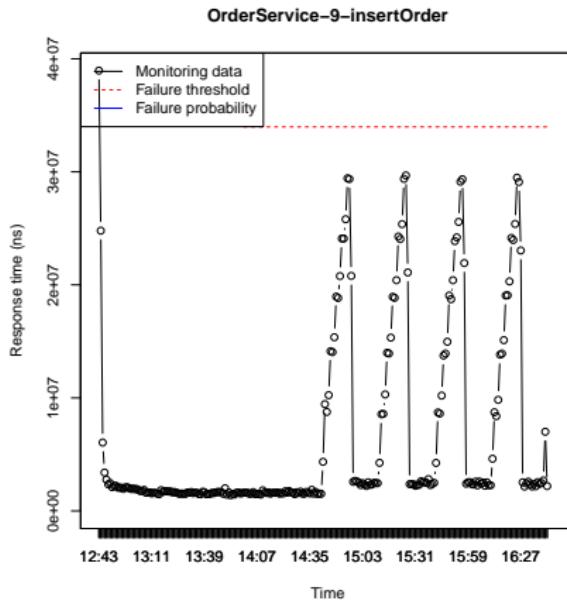


University of Stuttgart
Germany

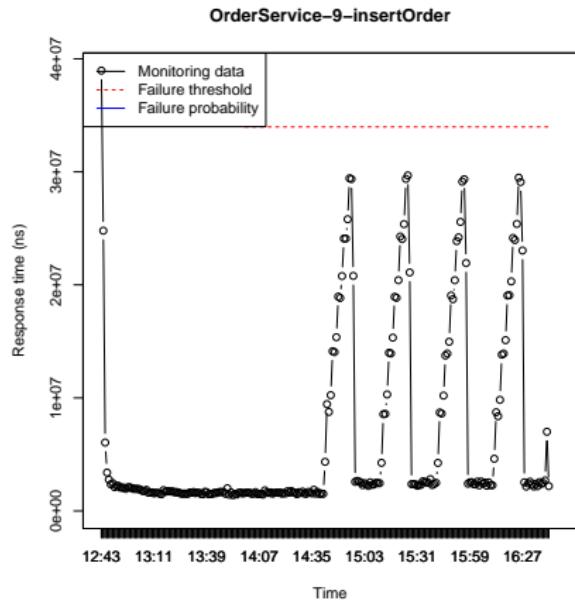
OrderService-9-insertOrder

Evaluation

CPM



SPM



Component vs. System Failure Prediction

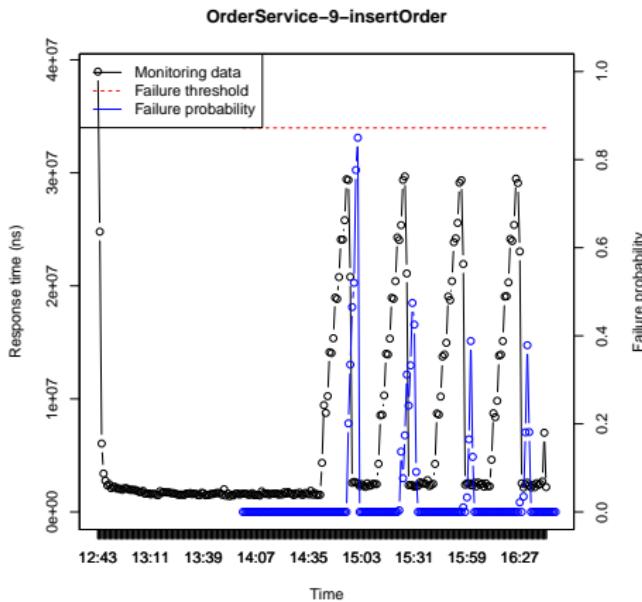


University of Stuttgart
Germany

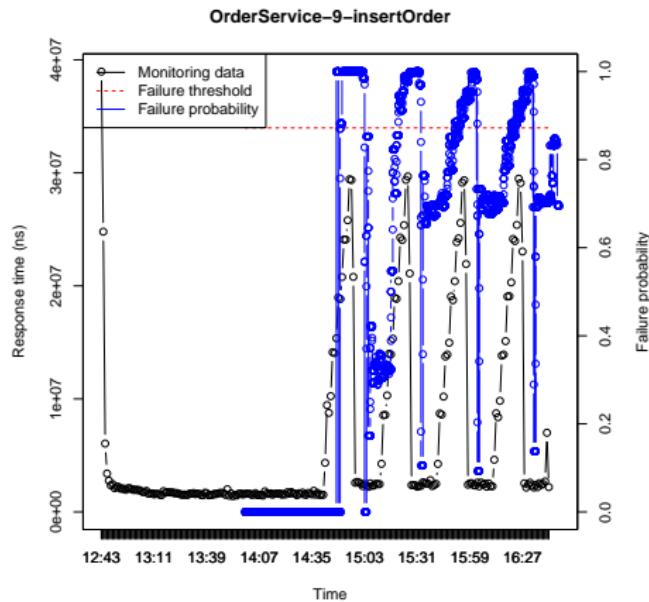
OrderService-9-insertOrder

Evaluation

CPM



SPM



Evaluation

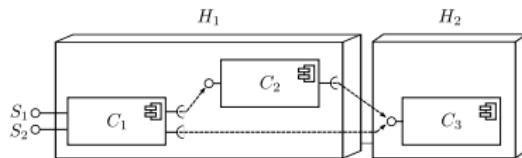
- CPMs perform quite well but can still be improved
- SPM shows potential to predict failure propagation using component dependency
- Although the failure propagates to other components, the predictions are not very good as the failure thresholds are not set properly



Conclusion

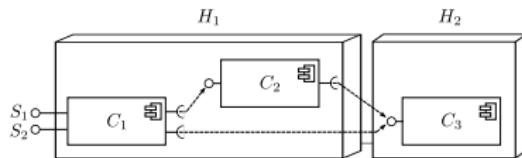
- 1 Motivation: Failure Management
- 2 [Recap] Hora: Online Failure Prediction for CB Systems
- 3 Hora: Framework and Implementation
- 4 Evaluation
- 5 Conclusion

Conclusion



Component-level Prediction Models
+
Component Dependency \Rightarrow System-level
Prediction Model

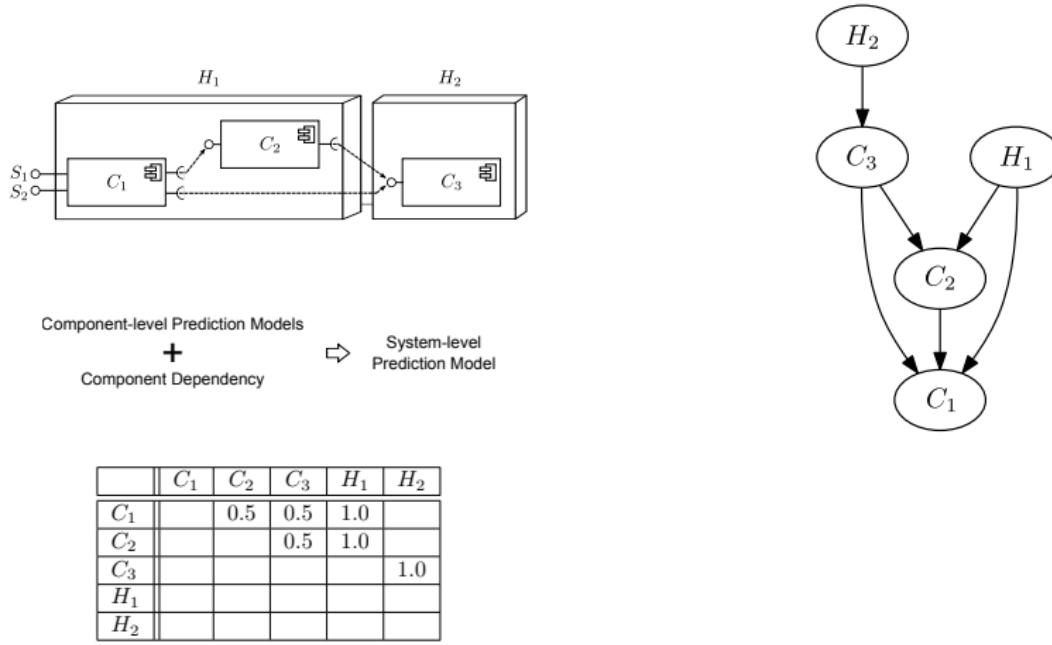
Conclusion



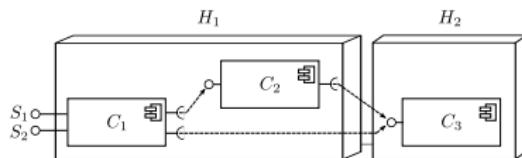
Component-level Prediction Models
+
Component Dependency \Rightarrow System-level
Prediction Model

	C_1	C_2	C_3	H_1	H_2
C_1		0.5	0.5	1.0	
C_2				0.5	1.0
C_3					1.0
H_1					
H_2					

Conclusion



Conclusion

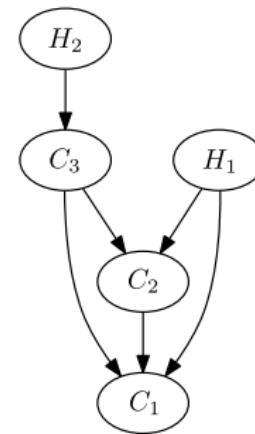


Component-level Prediction Models



Component Dependency

⇒ System-level
Prediction Model



	C_1	C_2	C_3	H_1	H_2
C_1		0.5	0.5	1.0	
C_2				0.5	1.0
C_3					1.0
H_1					
H_2					

		C_2	
H_1	C_3	✓	✗
✓	✓	0.9	0.1
✓	✗	0.5	0.5
✗	✓	0.0	1.0
✗	✗	0.0	1.0

Conclusion

- Improve prediction of CPMs
 - Apply different prediction techniques
- Calibrate the configuration parameters of Hora
 - Failure threshold
 - Lead time
 - ...
- Extend evaluation settings
 - Evaluate with distributed lab study
 - Evaluate with large-scale production systems
- Develop a failure injection framework/benchmark



- A. Amin, A. Colman, and L. Grunske. An approach to forecasting QoS attributes of web services based on ARIMA and GARCH models. In *Proceedings of the 19th IEEE International Conference on Web Services*, pages 74–81, June 2012.
- A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–33, 2004. ISSN 1545-5971. doi: 10.1109/TDSC.2004.2.
- S. Becker, H. Koziolek, and R. Reussner. The Palladio component model for model-driven performance prediction. *Journal of Systems and Software*, 82(1):3–22, 2009.
- T. C. Bielefeld. Online performance anomaly detection for large-scale software systems. Master's thesis, Mar. 2012. Diploma Thesis, Kiel University.
- Y. Liang, Y. Zhang, H. Xiong, and R. K. Sahoo. Failure prediction in ibm bluegene/l event logs. In *Proceedings of the 7th IEEE International Conference on Data Mining*, pages 583–588, 2007.
- T. Pitakrat. Hora: Online failure prediction framework for component-based software systems based on kieker and palladio. In *Proceedings of the Symposium on Software Performance: Joint Kieker/Palladio Days 2013*, CEUR Workshop Proceedings, 2013.
- T. Pitakrat, A. van Hoorn, and L. Grunske. A comparison of machine learning algorithms for proactive hard disk drive failure detection. In *Proceedings of the 4th International ACM Sigsoft Symposium on Architecting Critical Systems*, pages 1–10. ACM, 2013.
- T. Pitakrat, J. Grunert, O. Kabierschke, F. Keller, and A. van Hoorn. A framework for system event classification and prediction by means of machine learning. In *Proceedings of the 8th International Conference on Performance Evaluation Methodologies and Tools (ValueTools 2014)*, 2014a.
- T. Pitakrat, A. van Hoorn, and L. Grunske. Increasing dependability of component-based software systems by online failure prediction. In *Proceedings of the 10th European Dependable Computing Conference (EDCC)*, pages 66–69, May 2014b.
- A. van Hoorn. *Model-Driven Online Capacity Management for Component-Based Software Systems*. PhD thesis, Kiel, Germany, 2014. Dissertation, Faculty of Engineering, Kiel University.