

A Framework for System Event Classification and Prediction by Means of Machine Learning

Teerat Pitakrat, Jonas Grunert, Oliver Kabierschke,
Fabian Keller and André van Hoorn

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

Dec 10, 2014 @ VALUETOOLS 2014, Bratislava



[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 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](#) [8+](#) [Like](#) [149K](#) [Follow](#) [261K 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](#)

59
COMMENTS

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

Search articles & products

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 from seven-hour long Outlook.com outage

By Tom Warren on August 14, 2013 04:22 PM Email @tomwarren

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

CHRIS VELAZCO

posted yesterday 49 Comments



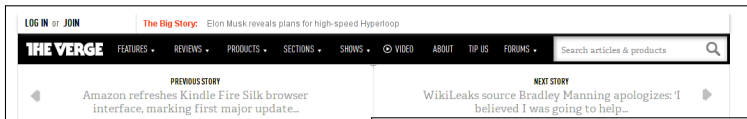
Those of you looking to spend the rest of today watching people do other things on Instagram or Vine probably just had a rough time trying to do it. Both services went offline for over an hour, most likely because of issues with **Amazon Web Services**.

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



Microsoft recovering long Outlook.com out

By Tom Warren on August 14, 2013 09:22 CDT | Email | @tomwarren

AWS Server Issues Take Down Instagram, Airbnb And IFTTT



CHRIS VELAZCO

posted yesterday



Those of you looking to spend the rest of the day watching people do other things on Instagram and Vine probably just had a rough time of it. Both services went offline for over a hour because of issues with Amazon's AWS cloud services.

Instagram was the first of the two to acknowledge its issues on Twitter, and followed suit half an hour later.

The deluge of tweets that accompanied the services' initial hiccups first started at 4:30 p.m. Eastern time, and only increased in intensity as users found they couldn't share pictures of their food or their meticulously crafted

49 C



146

RELATED STORIES

Google Apps goes titsup for millions - users REJOICE on Twitter

Google Drive goes titsup for MILLIONS of

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.

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 Search articles & products

PREVIOUS: Amazon refreshes Kindle interface, marking

APPS & SOFTWARE MICROSOFT WEB & SOCIAL

Microsoft's long Outlook

By Tom Warren on August 14, 2014 04:22 PM

AWS Server Issues Take Airbnb And IFTTT

CHRIS VELAZCO

posted yesterday

The deluge of tweets that accompany 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 crate

minutes, kills

service offline

precedented outage that seemingly brief time on Friday is over.

7pm Pacific Time and lasted long to the [Google Apps Dashboard](#), and being back online by 4:48pm.

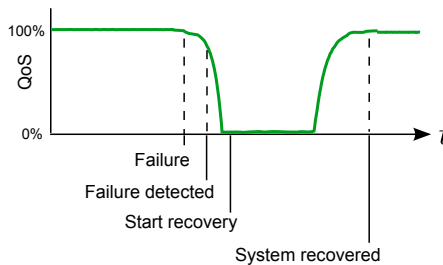
every service Mountain View has to reach to Gmail, YouTube, Google

Twitter Drive, and beyond.

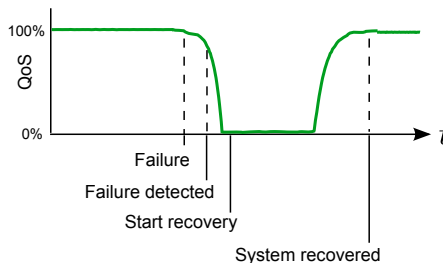
Google Drive goes titsup for MILLIONS of

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

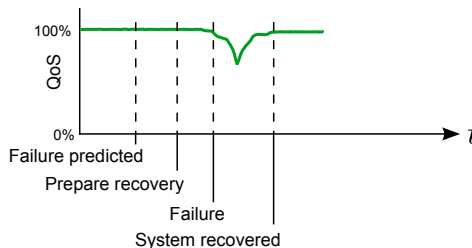
Reactive



Reactive



Proactive



- Log files can be used for
 - understanding system's behavior
 - diagnosing problems
 - detecting and predicting failures

- Log files can be used for
 - understanding system's behavior
 - diagnosing problems
 - detecting and predicting failures
- Example

INFO: Reading file X

INFO: Reading complete

INFO: Executing Routine A

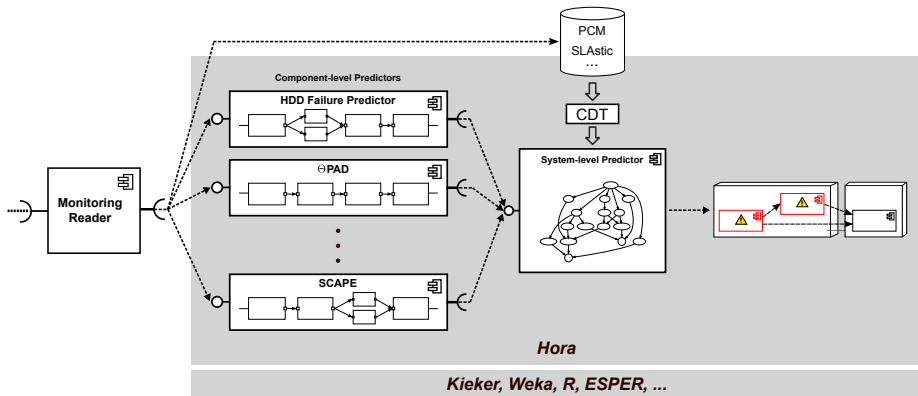
INFO: Reading file Y

FATAL: Critical Temperature in Segment Z

- Goals
 - Automatic classification of similar events
 - Automatic prediction of future events

- Goals
 - Automatic classification of similar events
 - Automatic prediction of future events
- Challenges
 - Log files are huge
 - Some information is redundant
 - Correlated events may not be close to each other

- Goals
 - Automatic classification of similar events
 - Automatic prediction of future events
- Challenges
 - Log files are huge
 - Some information is redundant
 - Correlated events may not be close to each other
- Approach: SCAPE framework
 - System event Classification And PrEdiction
 - Supports an extensible set of machine learning algorithms
 - Part of Hora approach for online failure prediction



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

1 Motivation: Failure Management

2 SCAPE Approach

3 Evaluation

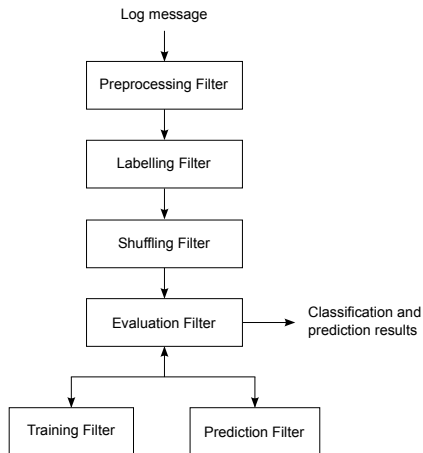
4 Conclusion

- Processing steps
 - 1 Event Preprocessing
 - 2 Event Classification
 - 3 Event Prediction

- Processing steps
 - 1 Event Preprocessing
 - 2 Event Classification
 - 3 Event Prediction
- Builds on
 - Kieker [van Hoorn et al. 2012]
 - Weka [Hall et al. 2009]

- Processing steps
 - 1 Event Preprocessing
 - 2 Event Classification
 - 3 Event Prediction
- Builds on
 - Kieker [van Hoorn et al. 2012]
 - Weka [Hall et al. 2009]
- Currently supports
 - Blue Gene/L log format
 - Weka's machine learning algorithms

- Processing steps
 - 1 Event Preprocessing
 - 2 Event Classification
 - 3 Event Prediction
- Builds on
 - Kieker [van Hoorn et al. 2012]
 - Weka [Hall et al. 2009]
- Currently supports
 - Blue Gene/L log format
 - Weka's machine learning algorithms



- Normalization [Liang et al. 2007]

- Filtering

- Normalization [Liang et al. 2007]
 - 1 Removing punctuation, e.g., . ; : ? ! = - [] | < > +
 - 2 Removing definite and indefinite articles, e.g., *a, an, the*
 - 3 Removing weak words, e.g., *be, is are, of, at, such, after, from*
 - 4 Replacing all numbers by the word `NUMBER`
 - 5 Replacing all hex addresses with N digits by the word `NDigitHex_Addr`
 - 6 Replacing domain specific identifiers by corresponding words such as `REGISTER` or `DIRECTORY`
 - 7 Replacing all dates by `DATE`
- Filtering

- Normalization [Liang et al. 2007]

- 1 Removing punctuation, e.g., . ; : ? ! = - [] | < > +
- 2 Removing definite and indefinite articles, e.g., *a, an, the*
- 3 Removing weak words, e.g., *be, is are, of, at, such, after, from*
- 4 Replacing all numbers by the word `NUMBER`
- 5 Replacing all hex addresses with N digits by the word `NDigitHex_Addr`
- 6 Replacing domain specific identifiers by corresponding words such as `REGISTER` or `DIRECTORY`
- 7 Replacing all dates by `DATE`

- Filtering

- Adaptive Semantic Filter (ASF) [Liang et al. 2007]
 - Removes highly correlated events (uses Phi correlation coefficient)
- Duplicate Removal Filter (DRF)
 - Removes similar events

```
4 torus receiver x+ input pipe error(s) (dcr 0x02ec) detected
1 torus receiver x- input pipe error(s) (dcr 0x02ed) detected
191790399 L3 EDRAM error(s) (dcr 0x0157) detected
2 L3 EDRAM error(s) (dcr 0x0157) detected
Error receiving packet, expecting type 57
3 torus receiver y+ input pipe error(s) (dcr 0x02ee) detected
3 torus receiver z- input pipe error(s) (dcr 0x02f1) detected
```

Before normalization

```
4 torus receiver x+ input pipe error(s) (dcr 0x02ec) detected
1 torus receiver x- input pipe error(s) (dcr 0x02ed) detected
191790399 L3 EDRAM error(s) (dcr 0x0157) detected
2 L3 EDRAM error(s) (dcr 0x0157) detected
Error receiving packet, expecting type 57
3 torus receiver y+ input pipe error(s) (dcr 0x02ee) detected
3 torus receiver z- input pipe error(s) (dcr 0x02f1) detected
```

Before normalization

```
number torus receiver x input pipe error detected
number torus receiver x input pipe error detected
number register edram error detected
number register edram error detected
error receiving packet expecting type number
number torus receiver y input pipe error detected
number torus receiver z input pipe error detected
```

After normalization



```
number torus receiver x input pipe error detected
number torus receiver x input pipe error detected
number register edram error detected
number register edram error detected
error receiving packet expecting type number
number torus receiver y input pipe error detected
number torus receiver z input pipe error detected
```

Before filtering



```
number torus receiver x input pipe error detected
number torus receiver x input pipe error detected
number register edram error detected
number register edram error detected
error receiving packet expecting type number
number torus receiver y input pipe error detected
number torus receiver z input pipe error detected
```

Before filtering

```
number torus receiver x input pipe error detected
number register edram error detected
error receiving packet expecting type number
number torus receiver z input pipe error detected
```

After filtering



Count	Label	Message
1	KERNBIT	KERNEL FATAL ddr: redundant bit steering failed, sequencer timeout
1	KERNEXT	KERNEL FATAL external input interrupt (unit=0x03 bit=0x01): tree header with no target waiting
1	KERNTLBE	KERNEL FATAL instruction TLB error interrupt
1	MONILL	MONITOR FAILURE monitor caught java.lang.IllegalStateException: while executing CONTROL Operation
2	LINKBLL	LINKCARD FATAL MidplaneSwitchController::clearPort() bl_clear_port failed: R63-M0-L0-U19-A
2	MONNULL	MONITOR FAILURE While inserting monitor info into DB caught java.lang.NullPointerException
3	KERNFLOAT	KERNEL FATAL floating point unavailable interrupt
3	KERNRTSA	KERNEL FATAL rts assertion failed: personality->version == BGLPERSONALITY_VERSION in void start() at start.cc:131
3	MMCS	MMCS FATAL L3 major internal error
5	KERNPROG	KERNEL FATAL program interrupt
10	APPTORUS	APP FATAL external input interrupt (unit=0x02 bit=0x00): uncorrectable torus error
10	MASNORM	BGLMASTER FAILURE mmcs_server exited normally with exit code 13
12	MONPOW	MONITOR FAILURE monitor caught java.lang.UnsupportedOperationException: power module U69 not present and is stopping
14	KERNNOETH	KERNEL FATAL no ethernet link
14	LINKPAP	LINKCARD FATAL MidplaneSwitchController::parityAlignment() pap failed: R22-M0-L0-U22-D, status=00000000 00000000
16	KERNCON	KERNEL FATAL MailboxMonitor::serviceMailboxes() lib_ido_error: -1033 BGLERR_IDO_PKT_TIMEOUT
18	KERNPAN	KERNEL FATAL kernel panic
24	LINKDISC	LINKCARD FATAL MidplaneSwitchController::sendTrain() port disconnected: R07-M1-L1-U19-E
37	MASABNORM	BGLMASTER FAILURE mmcs_server exited abnormally due to signal: Aborted
94	KERNSERV	KERNEL FATAL Power Good signal deactivated: R73-M1-N5. A service action may be required.
144	APPALLOD	APP FATAL cioid: Error creating node map from file /p/gb2/draeger/benchmark/datt16k_062205/map16k_bipartiz
166	LINKIAP	LINKCARD FATAL MidplaneSwitchController::receiveTrain() iap failed: R72-M1-L1-U18-A, status=beaaabff ec000000
192	KERNPOW	KERNEL FATAL Power deactivated: R05-M0-N4
209	KERN SOCK	KERNEL FATAL MailboxMonitor::serviceMailboxes() lib_ido_error: -1019 socket closed
320	APPCHILD	APP FATAL cioid: Error creating node map from file /p/gb2/cabot/miranda/newmaps/8k_128x64x1_8x4x4.map
342	KERNMC	KERNEL FATAL machine check interrupt
512	APPBUSY	APP FATAL cioid: Error creating node map from file /p/gb2/pakin1/sweep3d-5x5x400-10mk-3mmi-1024pes-sweep/sweep.map
720	KERNMNT	KERNEL FATAL Error: unable to mount filesystem
816	APPOUT	APP FATAL cioid: LOGIN chdir(/p/gb1/stella/RAPTOR/2183) failed: Input/output error
1503	KERNMICRO	KERNEL FATAL Microloader Assertion
1991	APPTO	APP FATAL cioid: Error reading message prefix on CioStream socket to 172.16.96.116:41739, Connection timed out
2048	APPUNAV	APP FATAL cioid: Error creating node map from file /home/auselton/bgl/64mps.sequential.mapfile
2370	APPRES	APP FATAL cioid: Error reading message prefix after LOAD_MESSAGE on CioStream socket to 172.16.96.116:52783
3983	KERNRTSP	KERNEL FATAL rts panic! - stopping execution
5983	APPREAD	APP FATAL cioid: failed to read message prefix on control stream CioStream socket to 172.16.96.116:33399
6145	KERNREC	KERNEL FATAL Error receiving packet on tree network, expecting type 57 instead of type 3
23338	KERNTERM	KERNEL FATAL rts: kernel terminated for reason 1004rts: bad message header
31531	KERNMNTF	KERNEL FATAL Lustre mount FAILED : bglio11 : block_id : location
49651	APPSEV	APP FATAL cioid: Error reading message prefix after LOGIN_MESSAGE on CioStream socket
63491	KERNSTOR	KERNEL FATAL data storage interrupt
152734	KERNDTLB	KERNEL FATAL data TLB error interrupt
4399503	-	KERNEL INFO instruction cache parity error corrected

```
4 torus receiver x+ input pipe error(s) (dcr 0x02ec) detected
1 torus receiver x- input pipe error(s) (dcr 0x02ed) detected
191790399 L3 EDRAM error(s) (dcr 0x0157) detected
2 L3 EDRAM error(s) (dcr 0x0157) detected
Error receiving packet, expecting type 57
3 torus receiver y+ input pipe error(s) (dcr 0x02ee) detected
3 torus receiver z- input pipe error(s) (dcr 0x02f1) detected
```

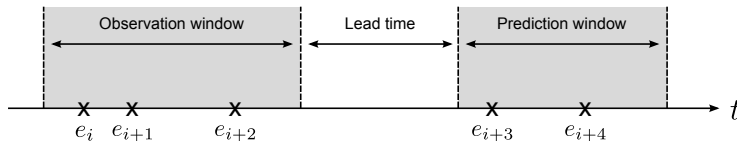
Before classification

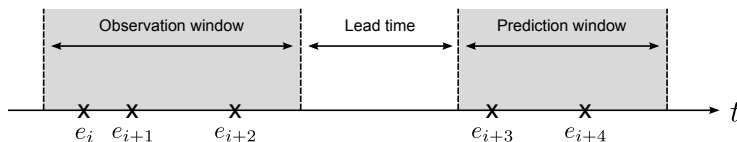
```
4 torus receiver x+ input pipe error(s) (dcr 0x02ec) detected
1 torus receiver x- input pipe error(s) (dcr 0x02ed) detected
191790399 L3 EDRAM error(s) (dcr 0x0157) detected
2 L3 EDRAM error(s) (dcr 0x0157) detected
Error receiving packet, expecting type 57
3 torus receiver y+ input pipe error(s) (dcr 0x02ee) detected
3 torus receiver z- input pipe error(s) (dcr 0x02f1) detected
```

Before classification

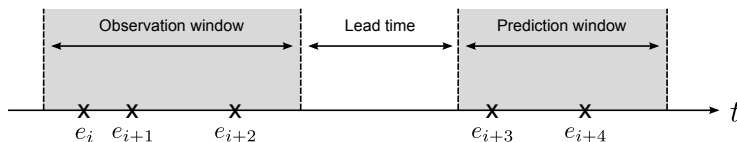
```
- 4 torus receiver x+ input pipe error(s) (dcr 0x02ec) detected
- 1 torus receiver x- input pipe error(s) (dcr 0x02ed) detected
- 191790399 L3 EDRAM error(s) (dcr 0x0157) detected
- 2 L3 EDRAM error(s) (dcr 0x0157) detected
KERNREC Error receiving packet, expecting type 57
- 3 torus receiver y+ input pipe error(s) (dcr 0x02ee) detected
- 3 torus receiver z- input pipe error(s) (dcr 0x02f1) detected
```

After classification

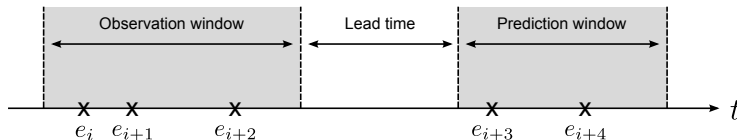




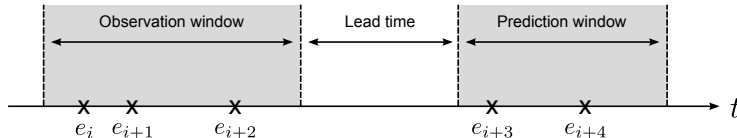
- 4 torus receiver x+ input pipe error(s) (dcr 0x02ec) detected
 - 1 torus receiver x- input pipe error(s) (dcr 0x02ed) detected
 - 191790399 L3 EDRAM error(s) (dcr 0x0157) detected
 - 2 L3 EDRAM error(s) (dcr 0x0157) detected
- KERNREC Error receiving packet, expecting type 57
- 3 torus receiver y+ input pipe error(s) (dcr 0x02ee) detected
 - 3 torus receiver z- input pipe error(s) (dcr 0x02f1) detected



Observation window	[<ul style="list-style-type: none"> - 4 torus receiver x+ input pipe error(s) (dcr 0x02ec) detected - 1 torus receiver x- input pipe error(s) (dcr 0x02ed) detected - 191790399 L3 EDRAM error(s) (dcr 0x0157) detected - 2 L3 EDRAM error(s) (dcr 0x0157) detected]
Prediction window	[<p>KERNREC Error receiving packet, expecting type 57</p> <ul style="list-style-type: none"> - 3 torus receiver y+ input pipe error(s) (dcr 0x02ee) detected - 3 torus receiver z- input pipe error(s) (dcr 0x02f1) detected]



Observation window	[<ul style="list-style-type: none"> - 4 torus receiver x+ input pipe error(s) (dcr 0x02ec) detected - 1 torus receiver x- input pipe error(s) (dcr 0x02ed) detected - 191790399 L3 EDRAM error(s) (dcr 0x0157) detected - 2 L3 EDRAM error(s) (dcr 0x0157) detected
Prediction window	[<p>KERNREC Error receiving packet, expecting type 57</p> <ul style="list-style-type: none"> - 3 torus receiver y+ input pipe error(s) (dcr 0x02ee) detected - 3 torus receiver z- input pipe error(s) (dcr 0x02f1) detected



Observation window	[<ul style="list-style-type: none"> - 4 torus receiver x+ input pipe error(s) (dcr 0x02ec) detected - 1 torus receiver x- input pipe error(s) (dcr 0x02ed) detected - 191790399 L3 EDRAM error(s) (dcr 0x0157) detected - 2 L3 EDRAM error(s) (dcr 0x0157) detected
Prediction window	[<p>KERNREC Error receiving packet, expecting type 57</p> <ul style="list-style-type: none"> - 3 torus receiver y+ input pipe error(s) (dcr 0x02ee) detected - 3 torus receiver z- input pipe error(s) (dcr 0x02f1) detected

Investigated parameters:

- Size of observation window
- Lead time
- Size of prediction window
- Sensitivity

- 1 Motivation: Failure Management
- 2 SCAPE Approach
- 3 Evaluation**
- 4 Conclusion

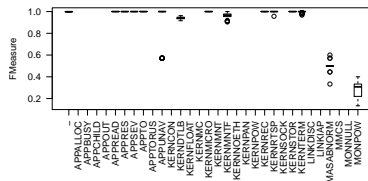
- Research questions
 - RQ1: How do different machine learning algorithms perform for system event classification and prediction?
 - RQ2: What is the impact of event preprocessing on the size of the dataset and on the event classification?

- Research questions
 - RQ1: How do different machine learning algorithms perform for system event classification and prediction?
 - RQ2: What is the impact of event preprocessing on the size of the dataset and on the event classification?
- Blue Gene/L supercomputer [Oliner and Stearley 2007]
 - 131,072 processors and 32,768 GB of RAM
 - 4,747,963 event messages collected over 215 days
 - 10-fold cross-validation

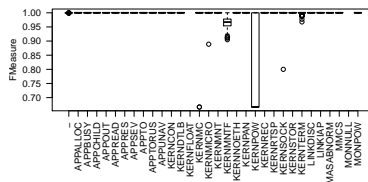
[illegible]

Evaluation

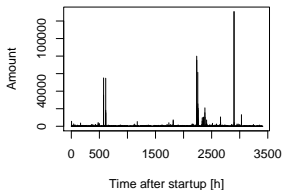
Count	Label	Message
1	KERNBIT	KERNEL FATAL ddr: redundant bit steering failed, sequencer timeout
1	KERNEXT	KERNEL FATAL external input interrupt (unit=0x03 bit=0x01): tree header with no target waiting
1	KERNTLBE	KERNEL FATAL instruction TLB error interrupt
1	MONILL	MONITOR FAILURE monitor caught java.lang.IllegalStateException: while executing CONTROL Operation
2	LINKBLL	LINKCARD FATAL MidplaneSwitchController::clearPort() bl_clear_port failed: R63-M0-L0-U19-A
2	MONNULL	MONITOR FAILURE While inserting monitor info into DB caught java.lang.NullPointerException
3	KERNFLOAT	KERNEL FATAL floating point unavailable interrupt
3	KERNRTSA	KERNEL FATAL rts assertion failed: personality->version == BGLPERSONALITY_VERSION in void start() at start.cc:131
3	MMCS	MMCS FATAL L3 major internal error
5	KERNPROG	KERNEL FATAL program interrupt
10	APPTORUS	APP FATAL external input interrupt (unit=0x02 bit=0x00): uncorrectable torus error
10	MASNORM	BGLMASTER FAILURE mmcs_server exited normally with exit code 13
12	MONPOW	MONITOR FAILURE monitor caught java.lang.UnsupportedOperationException: power module U69 not present and is stopping
14	KERNNOETH	KERNEL FATAL no ethernet link
14	LINKPAP	LINKCARD FATAL MidplaneSwitchController::parityAlignment() pap failed: R22-M0-L0-U22-D, status=00000000 00000000
16	KERNCON	KERNEL FATAL MailboxMonitor::serviceMailboxes() lib_ido_error: -1033 BGLERR_IDO_PKT_TIMEOUT
18	KERNPAN	KERNEL FATAL kernel panic
24	LINKDISC	LINKCARD FATAL MidplaneSwitchController::sendTrain() port disconnected: R07-M1-L1-U19-E
37	MASABNORM	BGLMASTER FAILURE mmcs_server exited abnormally due to signal: Aborted
94	KERNSEV	KERNEL FATAL Power Good signal deactivated: R73-M1-N5. A service action may be required.
144	APPALLOD	APP FATAL cioid: Error creating node map from file /p/gb2/draeger/benchmark/dat16k_062205/map16k_bipartiz
166	LINKIAP	LINKCARD FATAL MidplaneSwitchController::receiveTrain() iap failed: R72-M1-L1-U18-A, status=beaaabff ec000000
192	KERNPOW	KERNEL FATAL Power deactivated: R05-M0-N4
209	KERN SOCK	KERNEL FATAL MailboxMonitor::serviceMailboxes() lib_ido_error: -1019 socket closed
320	APPCHILD	APP FATAL cioid: Error creating node map from file /p/gb2/cabot/miranda/newmaps/8k_128x64x1_8x4x4.map
342	KERNMC	KERNEL FATAL machine check interrupt
512	APPBUSY	APP FATAL cioid: Error creating node map from file /p/gb2/pakin1/sweep3d-5x5x400-10mk-3mmi-1024pes-sweep/sweep.map
720	KERNMNT	KERNEL FATAL Error: unable to mount filesystem
816	APPOUT	APP FATAL cioid: LOGIN chdir(/p/gb1/stella/RAPTOR/2183) failed: Input/output error
1503	KERNMICRO	KERNEL FATAL Microloader Assertion
1991	APPTO	APP FATAL cioid: Error reading message prefix on CioStream socket to 172.16.96.116:41739, Connection timed out
2048	APPUNAV	APP FATAL cioid: Error creating node map from file /home/auselton/bgl/64mps.sequential.mapfile
2370	APPRES	APP FATAL cioid: Error reading message prefix after LOAD_MESSAGE on CioStream socket to 172.16.96.116:52783
3983	KERNRTSP	KERNEL FATAL rts panic! - stopping execution
5983	APPREAD	APP FATAL cioid: failed to read message prefix on control stream CioStream socket to 172.16.96.116:33399
6145	KERNREC	KERNEL FATAL Error receiving packet on tree network, expecting type 57 instead of type 3
23338	KERNTERM	KERNEL FATAL rts: kernel terminated for reason 1004rts: bad message header
31531	KERNMNTF	KERNEL FATAL Lustre mount FAILED : bglio11 : block_id : location
49651	APPSEV	APP FATAL cioid: Error reading message prefix after LOGIN_MESSAGE on CioStream socket
63491	KERNSTOR	KERNEL FATAL data storage interrupt
152734	KERNDTLB	KERNEL FATAL data TLB error interrupt
4399503	-	KERNEL INFO instruction cache parity error corrected



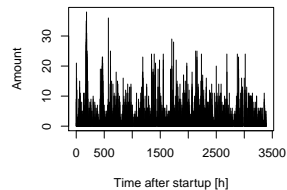
Naive Bayes with normalized log



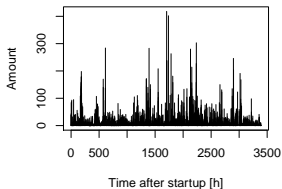
C4.5 with normalized log



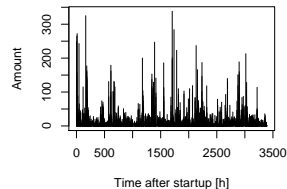
Original log



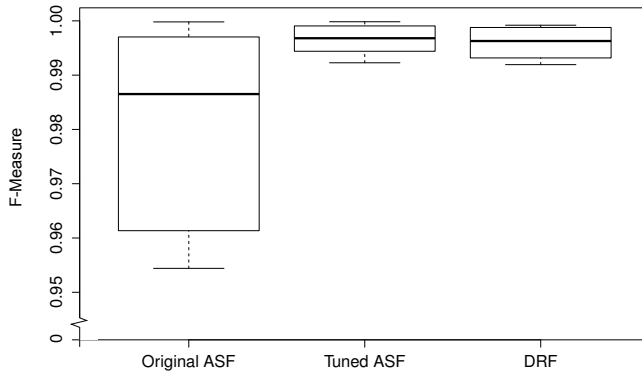
Original ASF

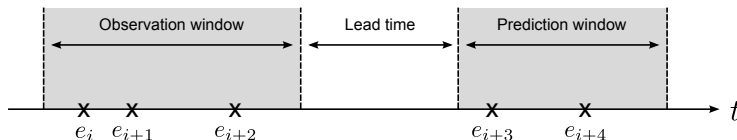


Tuned ASF



DRF





Investigated parameters:

- Size of observation window
- Lead time
- Size of prediction window
- Sensitivity

Algorithm	Lead time (sec)						
	0	60	120	300	600	1200	2800
NaiveBayes	0.663	0.589	0.547	0.517	0.506	0.511	0.506
C4.5	0.877	0.672	0.634	0.627	0.624	0.640	0.625

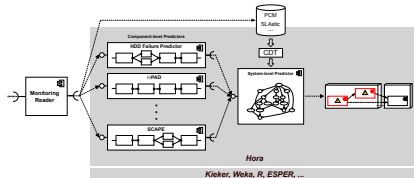
Algorithm	Prediction window (sec)						
	60	120	300	600	1200	2800	4800
NaiveBayes	0.491	0.493	0.485	0.506	0.511	0.532	0.553
C4.5	0.579	0.578	0.598	0.624	0.640	0.625	0.635

Algorithm	Number of past observations						
	1	2	3	4	6	8	16
NaiveBayes	0.603	0.517	0.506	0.500	0.501	0.501	0.503
C4.5	0.621	0.626	0.624	0.624	0.624	0.626	0.634

Algorithm	Sensitivity						
	1%	5%	10%	20%	40%	80%	100%
NaiveBayes	0.546	0.522	0.516	0.506	0.462	0.519	0.399
C4.5	0.523	0.572	0.609	0.624	0.691	0.234	-

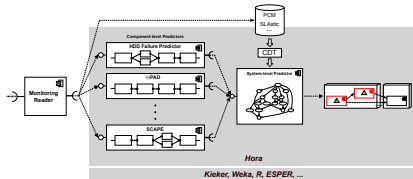
- 1 Motivation: Failure Management
- 2 SCAPE Approach
- 3 Evaluation
- 4 Conclusion**

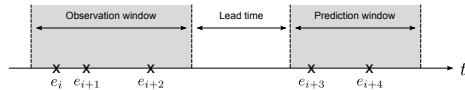
Conclusion



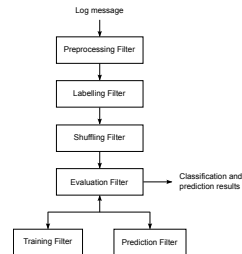
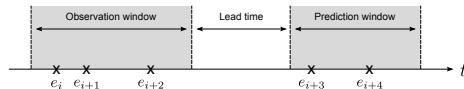
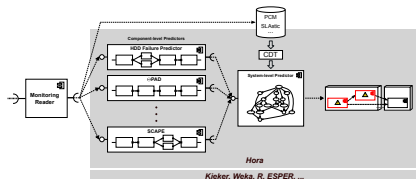


Conclusion

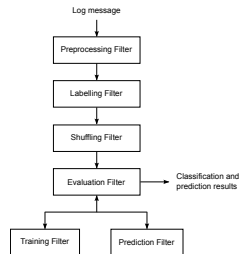
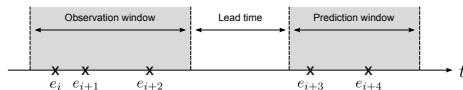
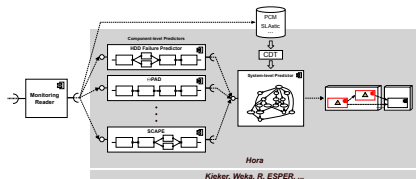
[illegible]

[illegible]

Conclusion

[illegible]

Conclusion

[illegible]

Supplementary material:
<http://www.iste.uni-stuttgart.de/rss/people/pitakrat/scape>

- Improve event prediction
- Extend evaluation settings
 - Evaluate with event log from other systems
- Integrate SCAPE into Hora framework
 - Combine with architectural model to infer the failure probability of other components

- S. Becker, H. Koziolok, 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.
- M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA data mining software: An update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.
- Y. Liang, Y. Zhang, H. Xiong, and R. K. Sahoo. An adaptive semantic filter for Blue Gene/L failure log analysis. In *Proc. Int'l Parallel and Distributed Processing Symp.*, pages 1–8, 2007.
- A. Oliner and J. Stearley. What supercomputers say: A study of five system logs. In *Proc. 37th Annual IEEE/IFIP Int'l Conf. on Dependable Systems and Networks*, pages 575–584, 2007.
- 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)*, 2014.
- A. van Hoorn. *Model-Driven Online Capacity Management for Component-Based Software Systems*. PhD thesis, Kiel, Germany, 2014. Dissertation, Faculty of Engineering, Kiel University.
- A. van Hoorn, J. Waller, and W. Hasselbring. Kieker: A framework for application performance monitoring and dynamic software analysis. In *Proc. 3rd ACM/SPEC Int'l Conf. on Performance Engineering*, pages 247–248. ACM, 2012.