Mining event log patterns in HPC systems

Ana Gainaru
joint work with Franck Cappello and Bill Kramer

HPC Resilience Summit 2010: Workshop on Resilience for Exascale HPC

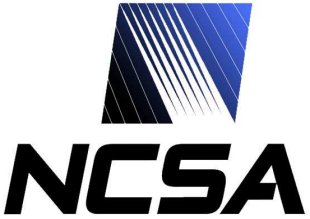

Table of contents

- Introduction
- Related work
- HELO (Hierarchical Event Log Organizer)
  - Offline clustering
  - Online classification
- Log files
- Results
- Conclusions

Introduction

- Find the representation of message types that exist in a log file
- Why?
  - Changes in the normal behavior of a message type could indicate a problem
  - Group of related messages - a better indicator of problems than individual messages
    - Anomalies are indicated by incomplete message sequences
  - Other open source tools perform poorly

Introduction

[2008-07-08 02:32:47][c1-0c1s5n0] 157 CMC Errors

Header Message

- Event: Header + Message
- Message
  - Constants - describe the message type
  - Variables – identify manipulated objects or states for the program
- Group template: *d+ CMC Errors*

Introduction

- HELO - Offline classification and online clustering

- Group wildcards: three types
  - d+ represents numeric tokens,
  - * represents any other single token
  - n+ represents all columns of tokens that have a value for some of the messages and don't exist for others.
- Example
  - *machine check interrupt (bit=0x1d): L2 dcache unit read parity error*
  - *machine check interrupt (bit=0x10): L2 DCU read error*
  - *machine check interrupt (bit=d+): L2 * * * n+*

Related work

- Supervised clustering
- Unsupervised clustering
  - Group messages based on the similarity between their descriptions
    - Pattern matching
    - Apriori
    - K-mean
    - Latent Semantic Indexing
- Advantages HELO

Other tools

- Loghound and SLCT
  - Limitations
    - High dimensional without having a fixed number of attributes
    - Not able to discover clusters irresponsive to how frequent the pattern instances appear in the input log file.
- IPLoM
  - Pattern matching algorithm
    - Searches for bijections between tokens from different messages
  - Limitations:
    - Syntactic depth of the mining process

Other tools

- StrAp
  - Offline and online
  - Numerical input data
  - Modifications made:
    - Unstructured text messages as input
    - Different lengths for messages
- MTE
  - Extracts two template sets:
    - Constants and variables
  - Limitation
    - Variable construction
    - *ciod: Error loading ./userfunc sqrt: invalid*

HELO algorithm - Offline

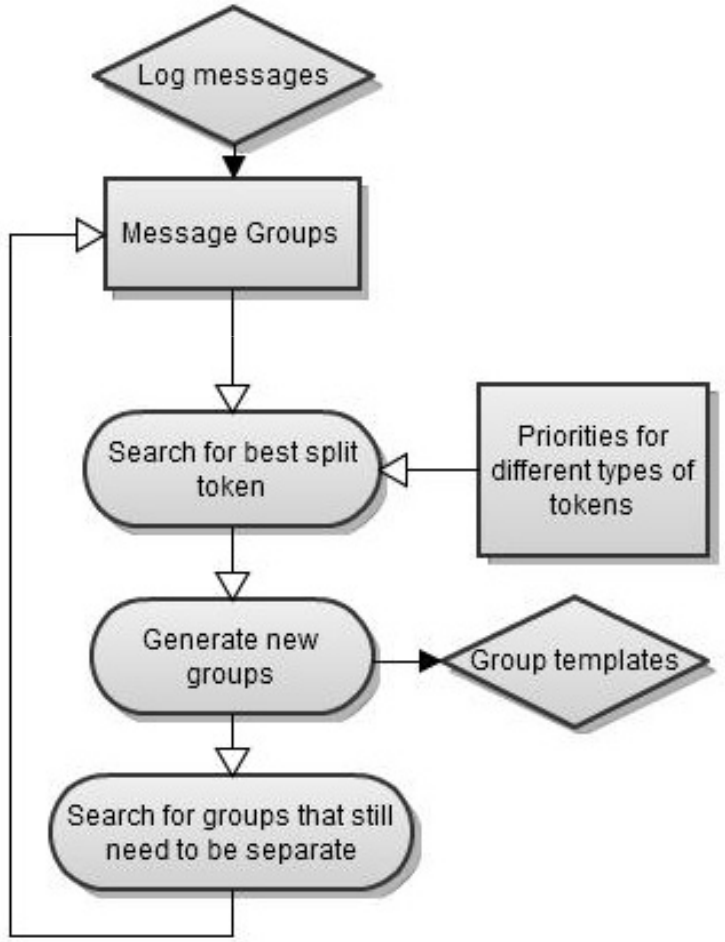

- Cluster goodness
  - Percentage of constant words
  - Over the average message length.
  - Default value: 40%

Splitting process

- Three type of words:
  - Numeric values – least priority
  - Hybrid tokens – extract the English words
  - English words – are left the way they are
- The column with:
  - The least number of distinct words, the most number of English words

Added 8 subnets and 409600 addresses to DB
address parity check..0
address parity check..1
Added 10 subnets and 589500 addresses to DB
data TLB error interrupt

Group reorganization

- If the splitting process splits constants
- Similarity between group templates 80%

- Example:
  - node card * check: missing u11 node
  - node card * check: missing u01 node

  - node card * check: missing * node

Online classification

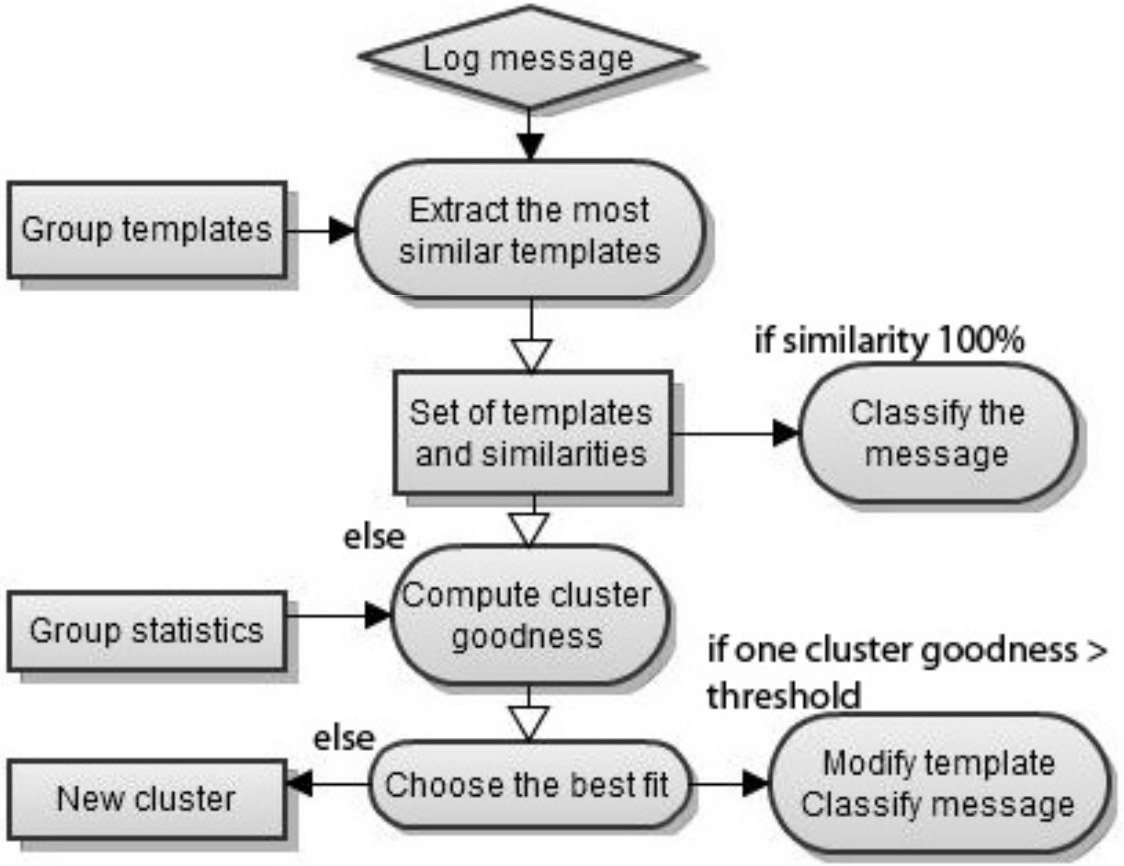

Log files

System	Messages	Time	Log type
BlueGene/L	4,747,963	6 months	event and login logs
Mercury	>10 million	3 months	event logs
PNNL	4,750	4 years	event logs
Cray XT4	3,170,514	3 months	event, syslog, console
LANL	433,490	9 years	cluster node outages

Table 1. Log data statistics.

- Extracted groups from each log file manually to compute the performance
- All logs have a description and different characteristics

Log files

- LANL has a friendly format
- Cray has a large amount of event patterns
- Mercury has a large amount of total messages, a few hundred thousand events per day
- PNNL has a large number of groups but having a small amount of messages
- BlueGene, Mercury and Cray put a lot of semantic problems

Definitions

- Information retrieval measures:
  - True positives
  - False positives
  - False negatives
  - Precision - measure of exactness
  - Recall – measure of completeness
  - F-measure - evenly weights precision and recall into a single value

Experiments

- Offline/online
- Offline: two cases
  - Measure the corrected found groups
  - Measure the corrected classified messages
- Online
  - Determine the percentage of corrected classified events

Results – Offline – Case 1

Performance for corrected clustered templates

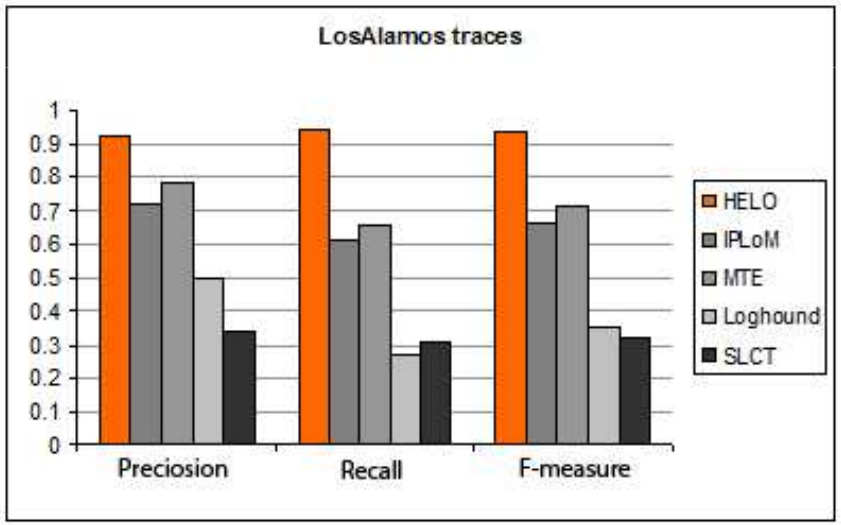

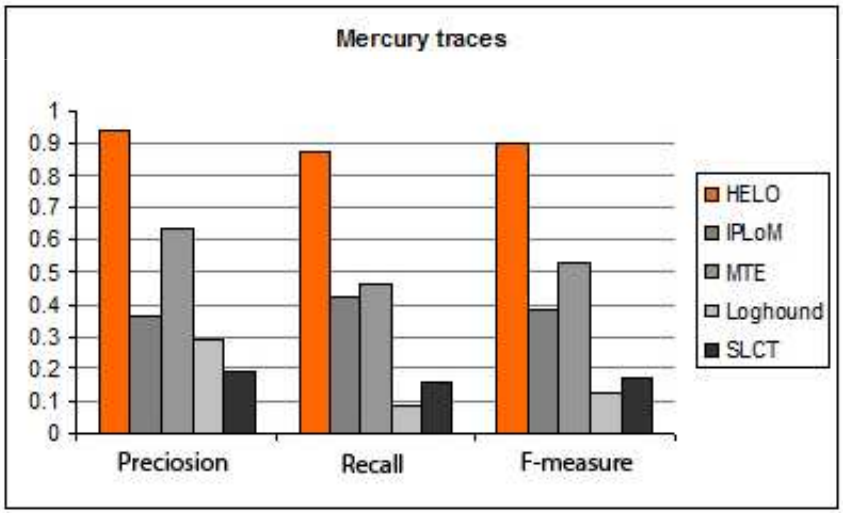

Results – Offline – Case 1

- Semantic problems
  - *fpr1 = 0x10055620000003e1004562008000815*
  - *lr = 0x00205034 xer = 0x00000002*
- Message length
  - *Corrective Measures SDE / DS2100 (upper) need to be replaced*
  - *Corrective Measures Upper DS2100 in need of Replacement*
- Message frequency

Results – Offline – Case 2

Performance for corrected clustered messages

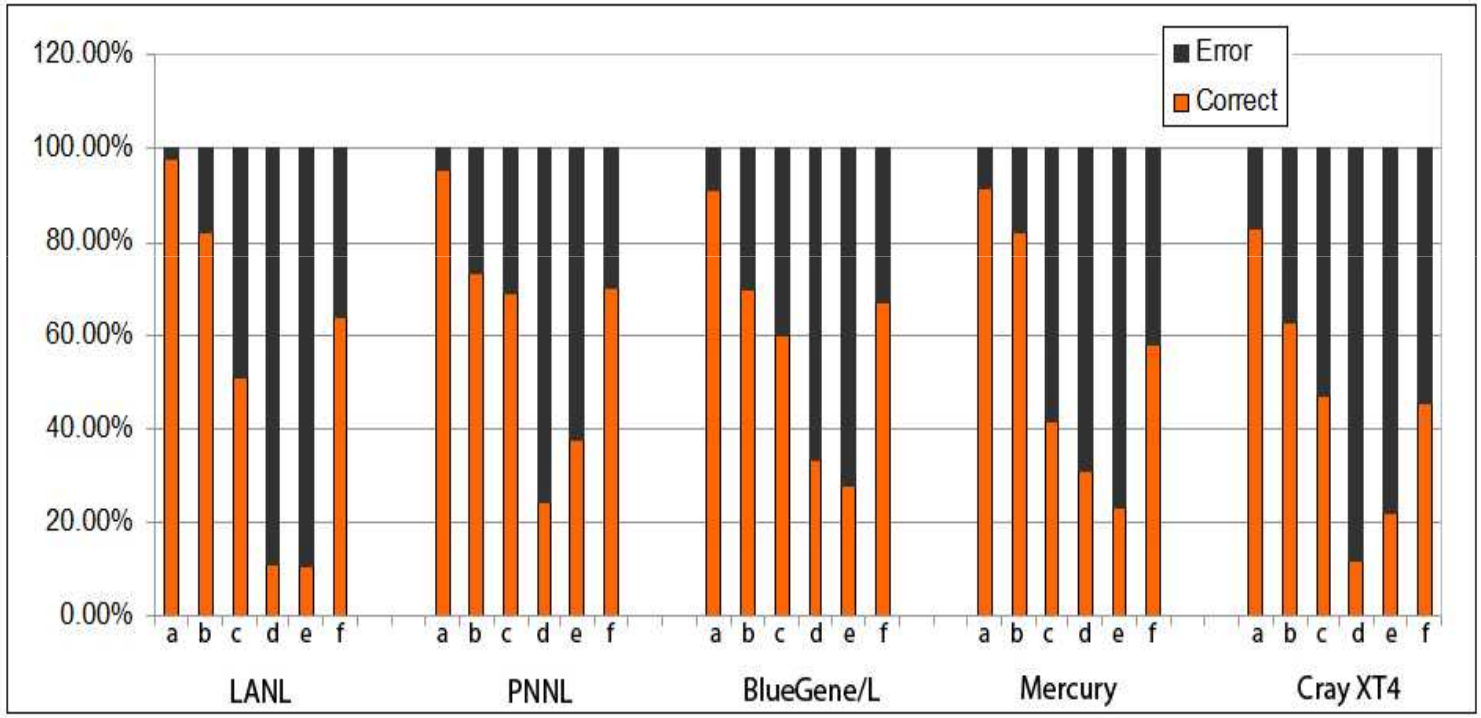

a)HELO b)StrAp c)IPLoM d)Loghound e)SLCT f)MTE

Online

- Compare HELO with StrAp
- Divide each log into 10 sets:
  - One for training
  - 9 for testing
- The output:
  - Array of group ids, one value for each message received for classification.

Online

Performance for corrected clustered incoming events

- For each training set

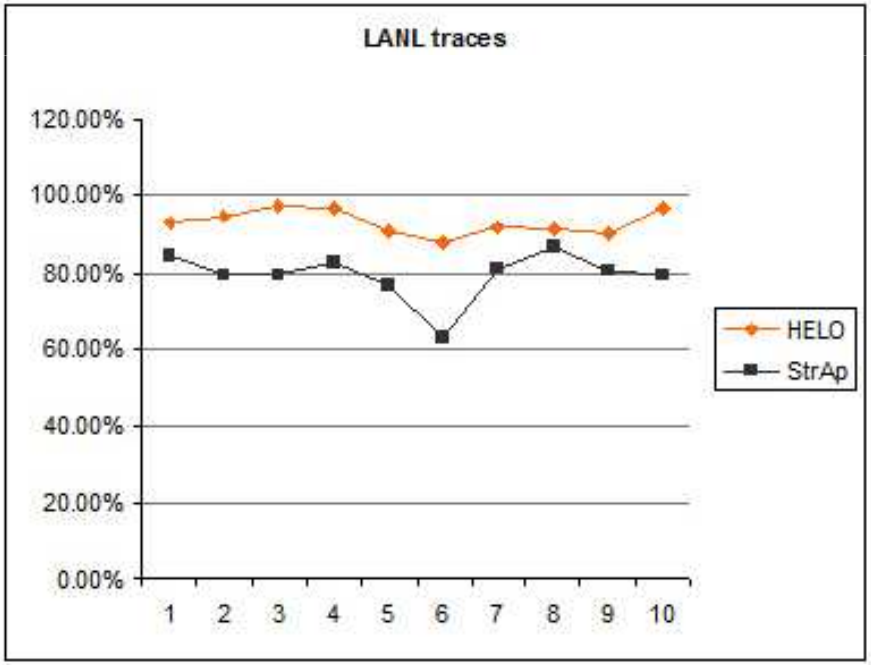

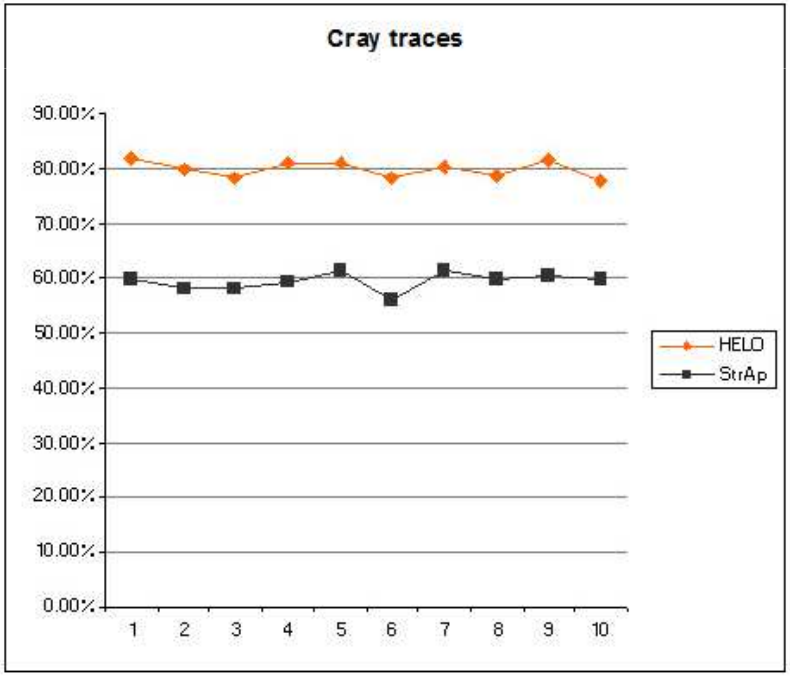

Online

- Different methodology for both tools
- Training set with semantic problems
  - The distance between the two tools will be higher
- Many cluster messages with different length
  - The distance between the tool is smaller

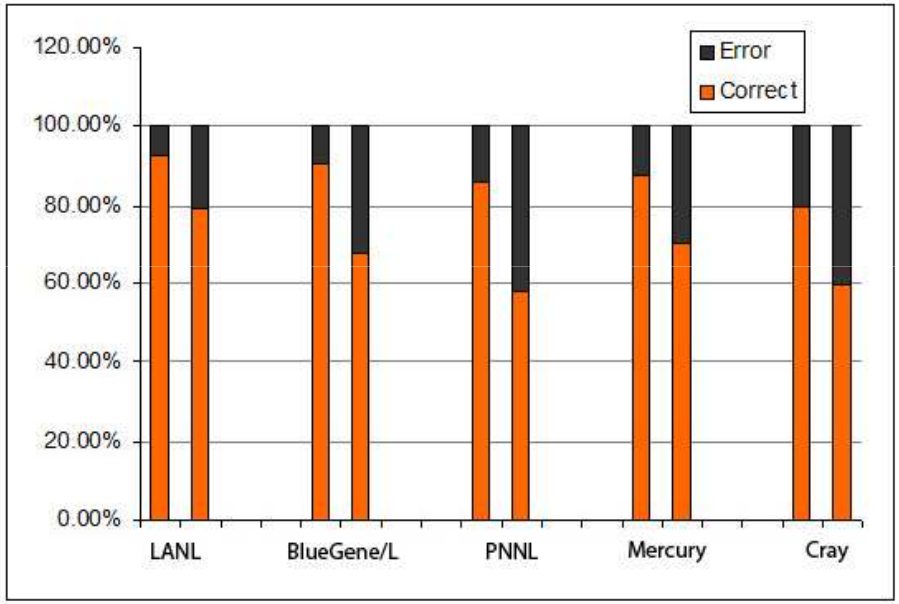

Mean value for all test cases

Conclusions

- Event analysis needs an automatic and efficient clustering approach
- HELO extracts group templates
  - Are used to describe events
  - Are user-friendly
- Comparison with 5 different tools for 5 different log files

Conclusions

- Other tools:
  - Do not scale well for the size and dimensionality of logs
  - Have limitations in the syntactic depth of the mining
  - Have problems with messages with different length
  - Are unable to adapt the templates to new messages
- HELO performance:
  - Average precision and recall of 0.9
  - Increase the correct number of groups by a factor of 1.5
  - Decrease the number of false positives and negatives by an average factor of 4.

Future work

- Correlations between templates
  - Message sequences – time or location
- Analyzing changes in the normal behavior of a message type
  - Precursor for faults
  - Influences on other message types

Q&A

- Thank you

Ana Gainaru
againaru@ncsa.illinois.edu