Log Analytics in HPC: A Data-driven Reinforcement Learning Framework

Zhengping Luo† , Tao Hou† , Tung Thanh Nguyen‡ , Hui Zeng‡ and Zhuo Lu†

† University of South Florida, Tampa, FL, USA, {zhengpingluo@mail., taohou@mail., zhuolu@}usf.edu

‡ Intelligent Automation Inc., Rockville, MD, USA, {tnguyen@, hzeng@}i-a-i.com

UNIVERSITY of SOUTH FLORIDA

High Performance Computing

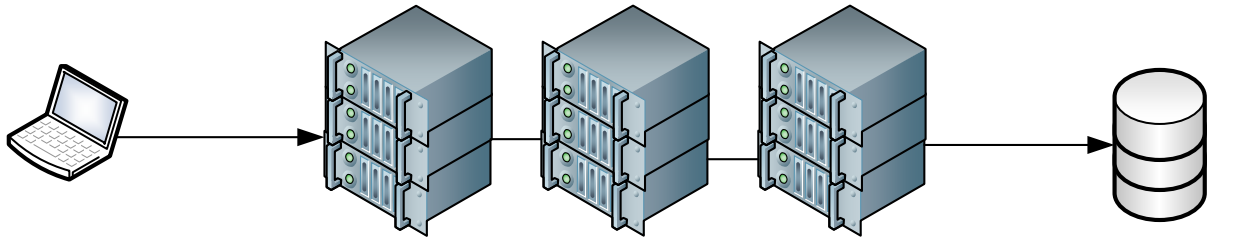

General log processing framework

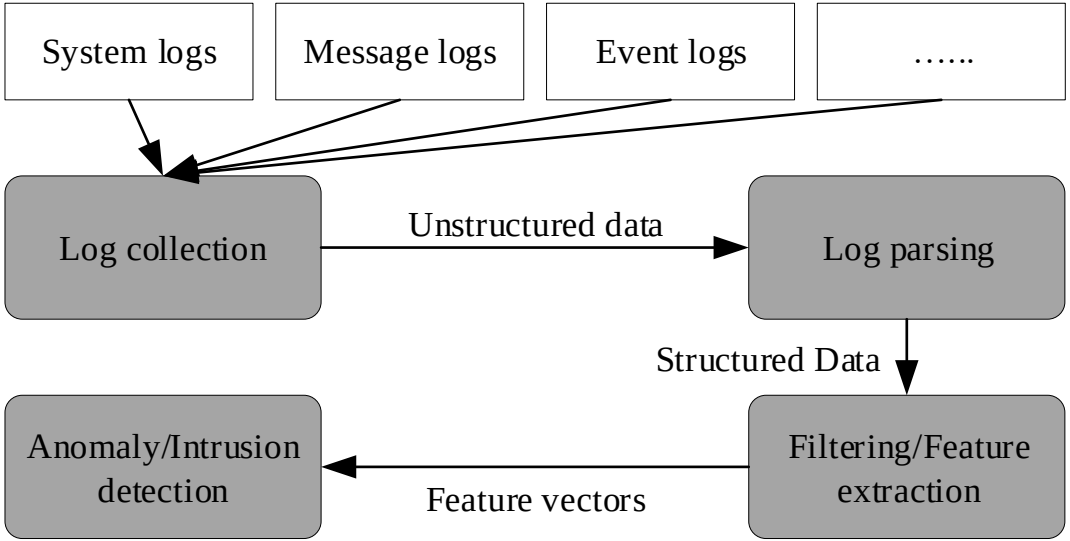

Message Passing Interface

- A communication protocol for parallel computing.
- Dominant model used in HPC (Sur et al. 2017).
- Popular implementation includes MPICH, Open MPI, etc.

Example MPI log file

```
MPI_Init entering at walltime 11201704.951285540, cputime 0.007714346 seconds in thread 0.
int argc=3
string argv[3]=["../../../skampi", "-i", "count1isend_recv.ski"]
MPI_Init returning at walltime 11201704.951286681, cputime 0.007715944 seconds in thread 0.
MPI_Comm_rank entering at walltime 11201704.951827833, cputime 0.007973521 seconds in thread 0.
MPI_Comm comm=2 (MPI_COMM_WORLD)
int rank=0
MPI_Comm_rank returning at walltime 11201704.951828922, cputime 0.007974872 seconds in thread 0.
MPI_Comm_size entering at walltime 11201704.951855949, cputime 0.008001787 seconds in thread 0.
MPI_Comm comm=2 (MPI_COMM_WORLD)
int size=2
MPI_Comm_size returning at walltime 11201704.951856793, cputime 0.008002713 seconds in thread 0.
MPI_Comm_dup entering at walltime 11201704.951865849, cputime 0.008011692 seconds in thread 0.
MPI_Comm oldcomm=2 (MPI_COMM_WORLD)
MPI_Comm newcomm=4 (user-defined-comm)
MPI_Comm_dup returning at walltime 11201704.951985484, cputime 0.008131085 seconds in thread 0.
MPI_Comm_get_attr entering at walltime 11201704.952071611, cputime 0.008217294 seconds in thread 0.
MPI_Comm comm=2 (MPI_COMM_WORLD)
int keyval=2 (MPI_IO)
int flag=1
```

MPI commands

Sending	MPI_Send, MPI_ Rsend, MPI_Isend
Receiving	MPI_Recv, MPI_Irecv
Collective	MPI_Allgatherv, MPI_Allreduce, MPI_Alltoallv MPI_Barrier, MPI_Bcast, MPI_Gather MPI_Gatherv, MPI_Reduce, MPI_Testall MPI_Waitall, MPI_Waitany
Other	MPI_Iprobe, MPI_Test, MPI_Barrier

The MPI commands that were commonly used [DeMasi et al. 2013].

Malicious detection

- Reinforcement learning

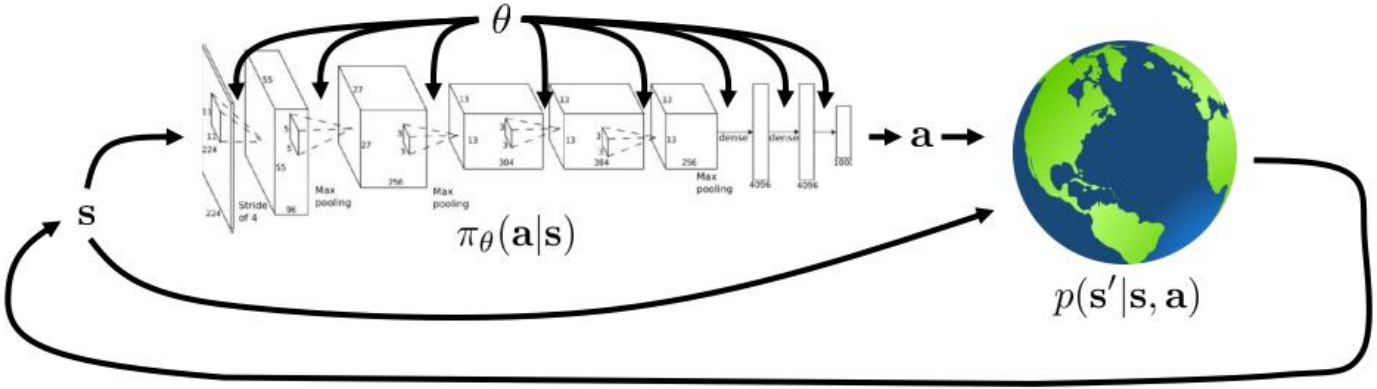

$$\underbrace{p_\theta(\mathbf{s}_1, \mathbf{a}_1, \ldots, \mathbf{s}_T, \mathbf{a}_T)}_{p_\theta(\tau)} = p(\mathbf{s}_1) \prod_{t=1}^{T} \pi_\theta(\mathbf{a}_t|\mathbf{s}_t) p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$$

$$\theta^\star = \arg\max_\theta E_{\tau \sim p_\theta(\tau)} \left[\sum_t r(\mathbf{s}_t, \mathbf{a}_t) \right]$$

Sergey Levine et al. 2017

Log analytics using Reinforcement learning

- ReLog framework

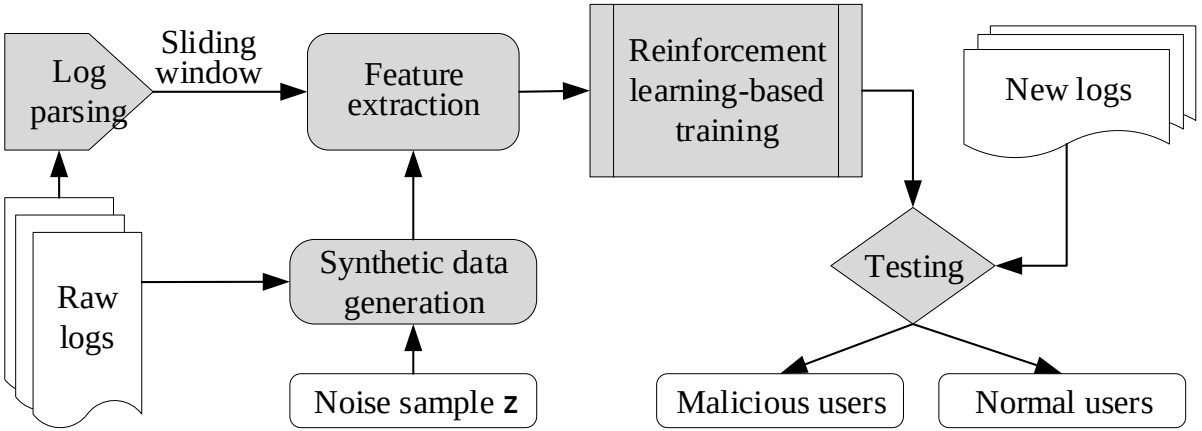

Training data generation

- Generative adversarial networks

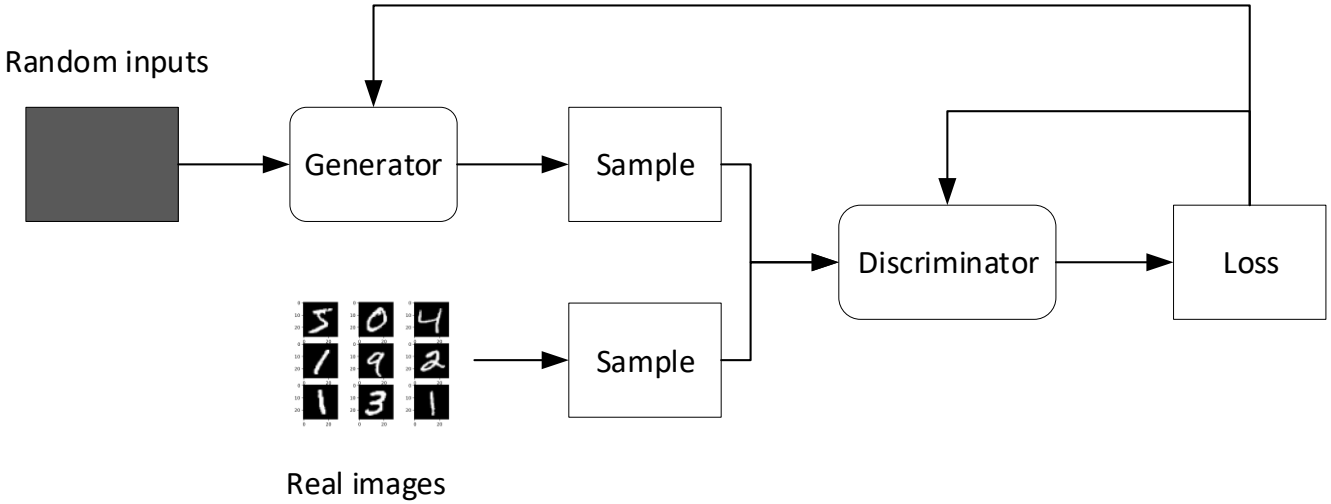

Occurring frequency of MPI commands

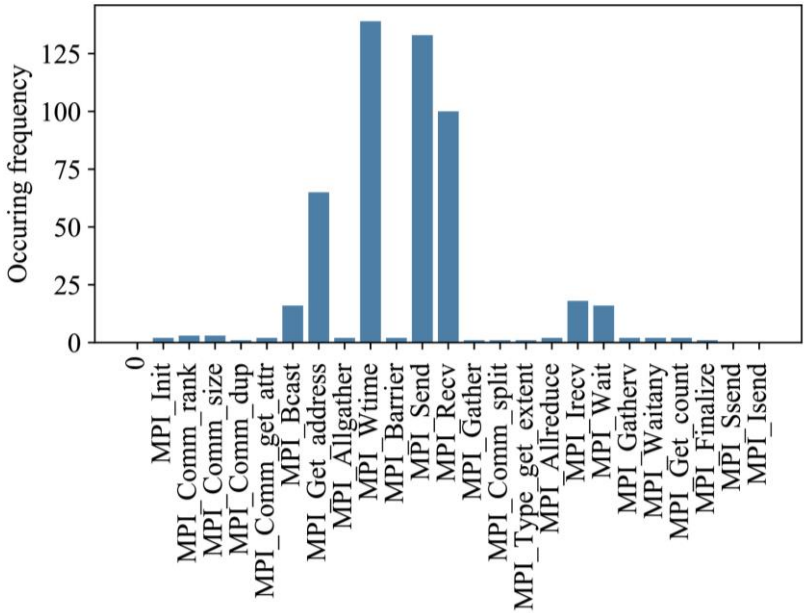

Training loss of ReLog

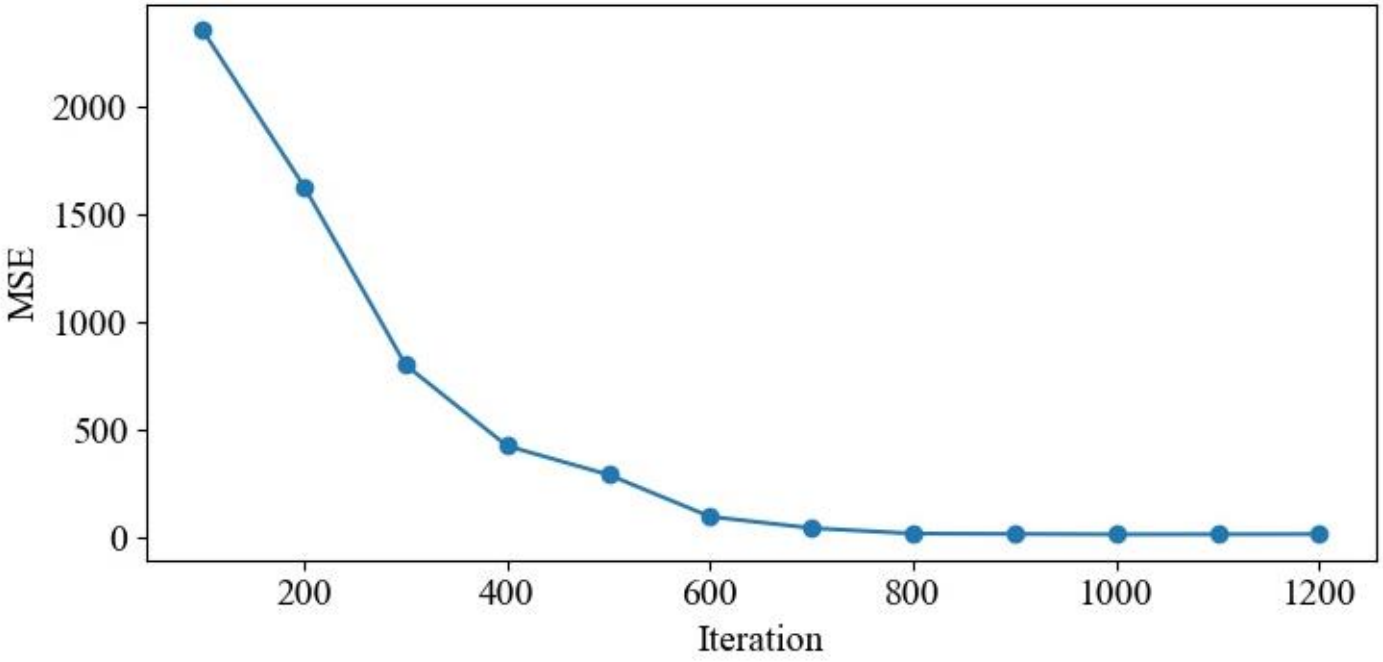

Detection performance

TABLE I
RELATIONSHIP BETWEEN SLIDING WINDOW SIZE AND DETECTION ACCURACY.

Window size	**100**	**120**	**140**	**160**	**180**	**200**	**220**
Detection accuracy	0.36	0.42	0.54	0.78	0.93	0.93	0.93

TABLE II
COMPARISON OF RELOG WITH OTHER EXISTING METHODS

Detection methods	**Time cost (seconds)**	**Detection accuracy**
DeepLog [5]	56	0.91
SVM [9]	13	0.86
ReLog	107	0.93

References

- DeMasi, Orianna, Taghrid Samak, and David H. Bailey. "Identifying HPC codes via performance logs and machine learning." In *Proceedings of the first workshop on Changing landscapes in HPC security*, pp. 23-30. 2013.
- https://sites.google.com/view/icml17deeprl
- Sur, Sayantan; Koop, Matthew J.; Panda, Dhabaleswar K. (4 August 2017). "MPI and communication---High-performance and scalable MPI over Infini Band with reduced memory usage". *High-performance and Scalable MPI over InfiniBand with Reduced Memory Usage: An In-depth Performance Analysis. ACM*. p. 105.

Thank you all for your time

UNIVERSITY of SOUTH FLORIDA