SDC 19 SNIA INDIA

May 23-24, 2019
Bangalore, India

STORAGE DEVELOPER CONFERENCE

Recommendation Engine For Data Eviction Policy Selection

Divyank Shukla
Msys Technologies Pvt Ltd

Table of Contents

SDC 19
SNIA INDIA

Problem Description

- Ability to Advice clients :
  - “If we migrated[evicted] this extent our client will experience fall in IO latency by x% points”
  - “If we added one more VM, to the Controller, the IO latency for existing client will reduce by x% points”
  - “If we promoted this extent our IO latency for 80% of customers will improve by y%”

IO Controller Outline / Framework

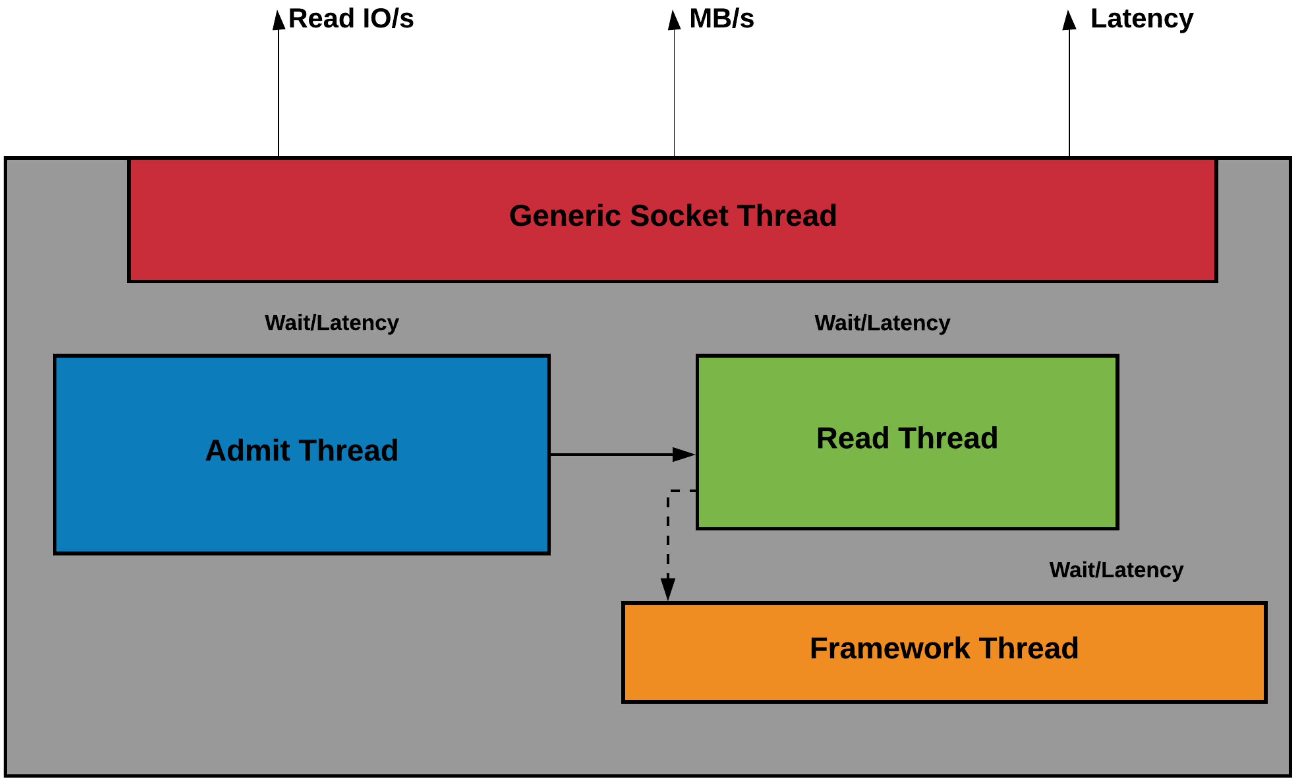

Framework and Metrics

- Threads for admitting, processing and communicating responses for Requests.
- Requests Split to Operations
- Operations (Ops) may get Queued
- Ops contend Locks, experience Wait times
- Ops's Responses communicated to clients.
- Request Latency, Wait times, Request Rates etc

Exploring The Problem Space

- Is it a N-class Classification Problem ?
- Is it a Regression Problem ?
- Is it a Queue Modelling Problem ?
- Can we model it as a Poisson Process ?
- Is it All of the Above ?

SDC 19 SNIA INDIA

Approaches

- Linear Regression
- Naïve Bayes Classification
- Multi-Resource Queue Modelling
- ARIMA (Time Based Forecasting)
- Predictive State Modelling
- Neural Networks
- Neural Networks with Internal States

Basic Statistics

- Are the variables dependent ?
- Cost function response to delta changes in dependent variables ?
- Is it linear, non-linear, parametric or not ?
- Does the response vary from time A to time B ?
- Bootstrapping, Markov Chain Monte-Carlo, Layers and Features learning

Mbps Plot Vs Normalized IOps

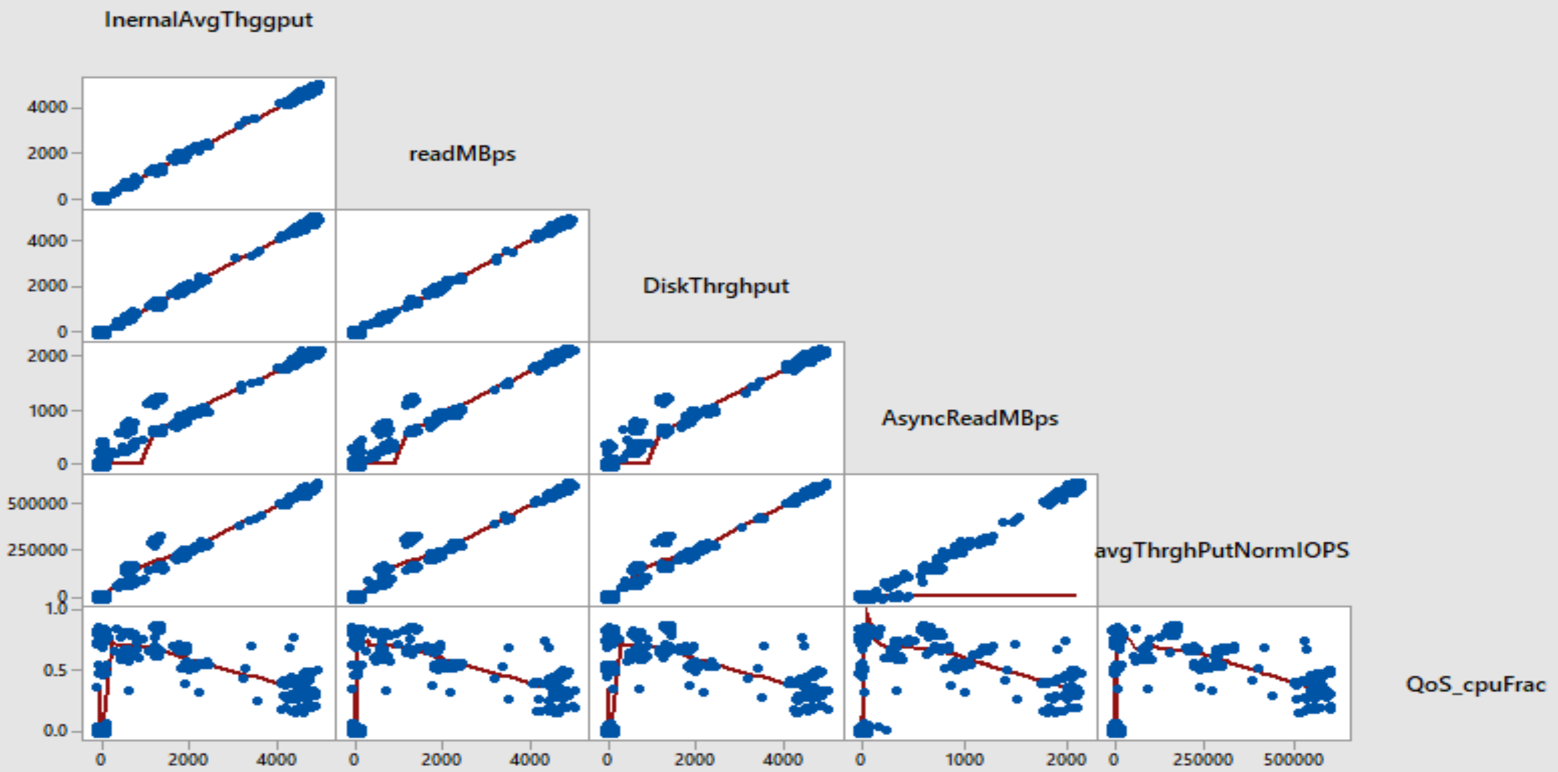

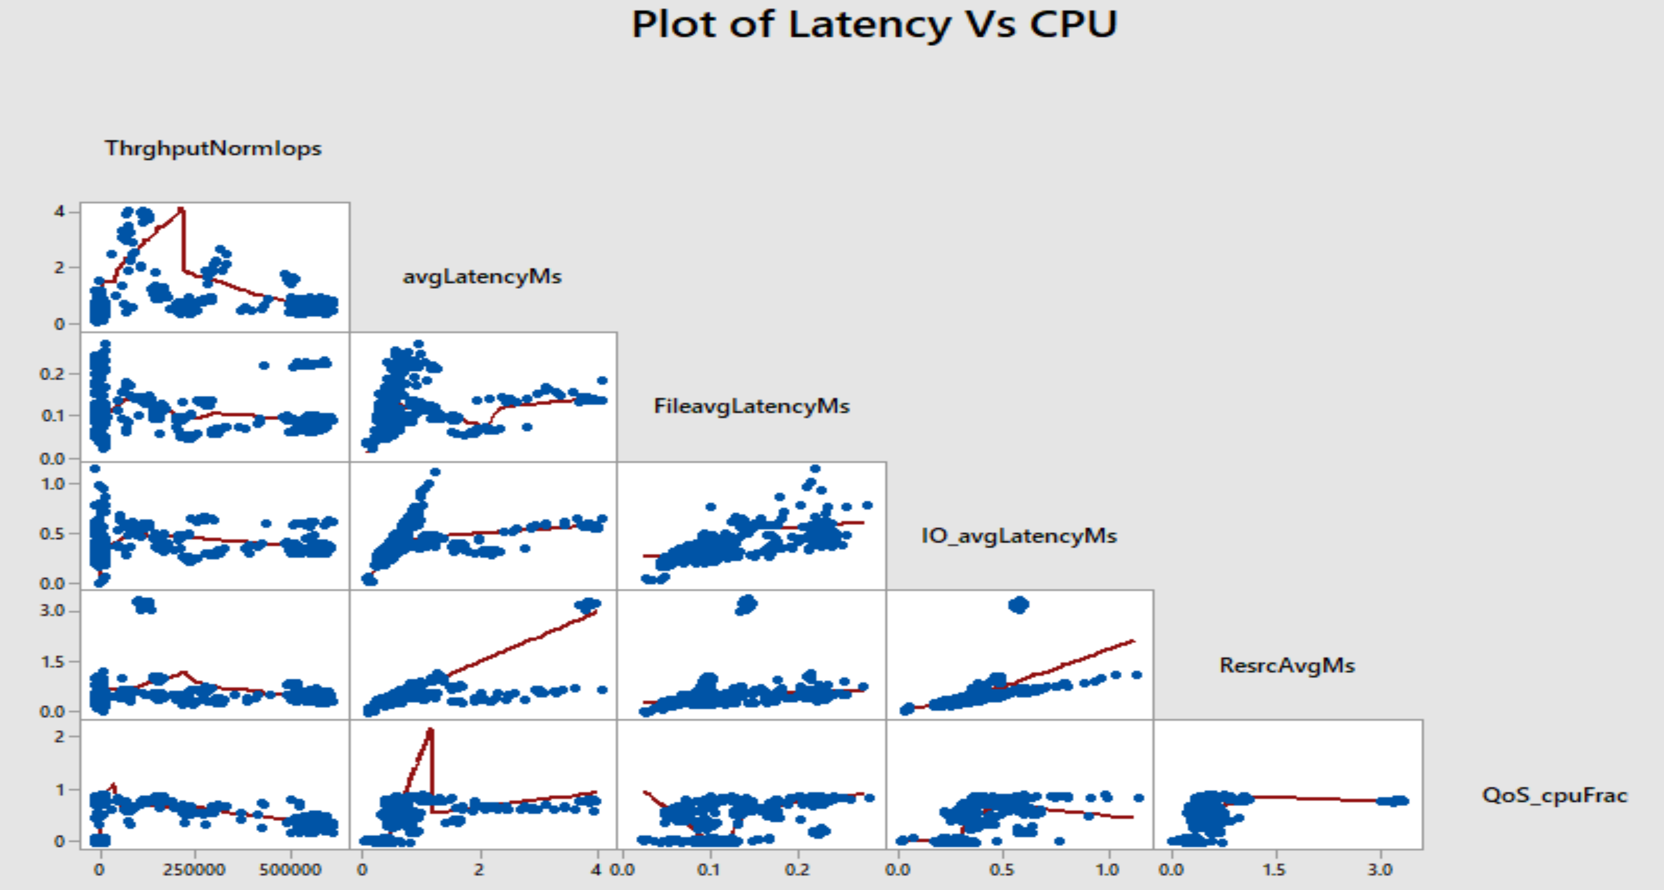

Plot of Latency Vs CPU
ThrghputNormIops
avgLatencyMs
FileavgLatencyMs
IO_avgLatencyMs
ResrcAvgMs
QoS_cpuFrac
4
2
0
0.2
0.1
0.0
1.0
0.5
0.0
3.0
1.5
0.0
2
1
0
0
250000
500000
0
2
4
0.0
0.1
0.2
0.0
0.5
1.0
0.0
1.5
3.0

System's Wait Matrix for Normalized Read

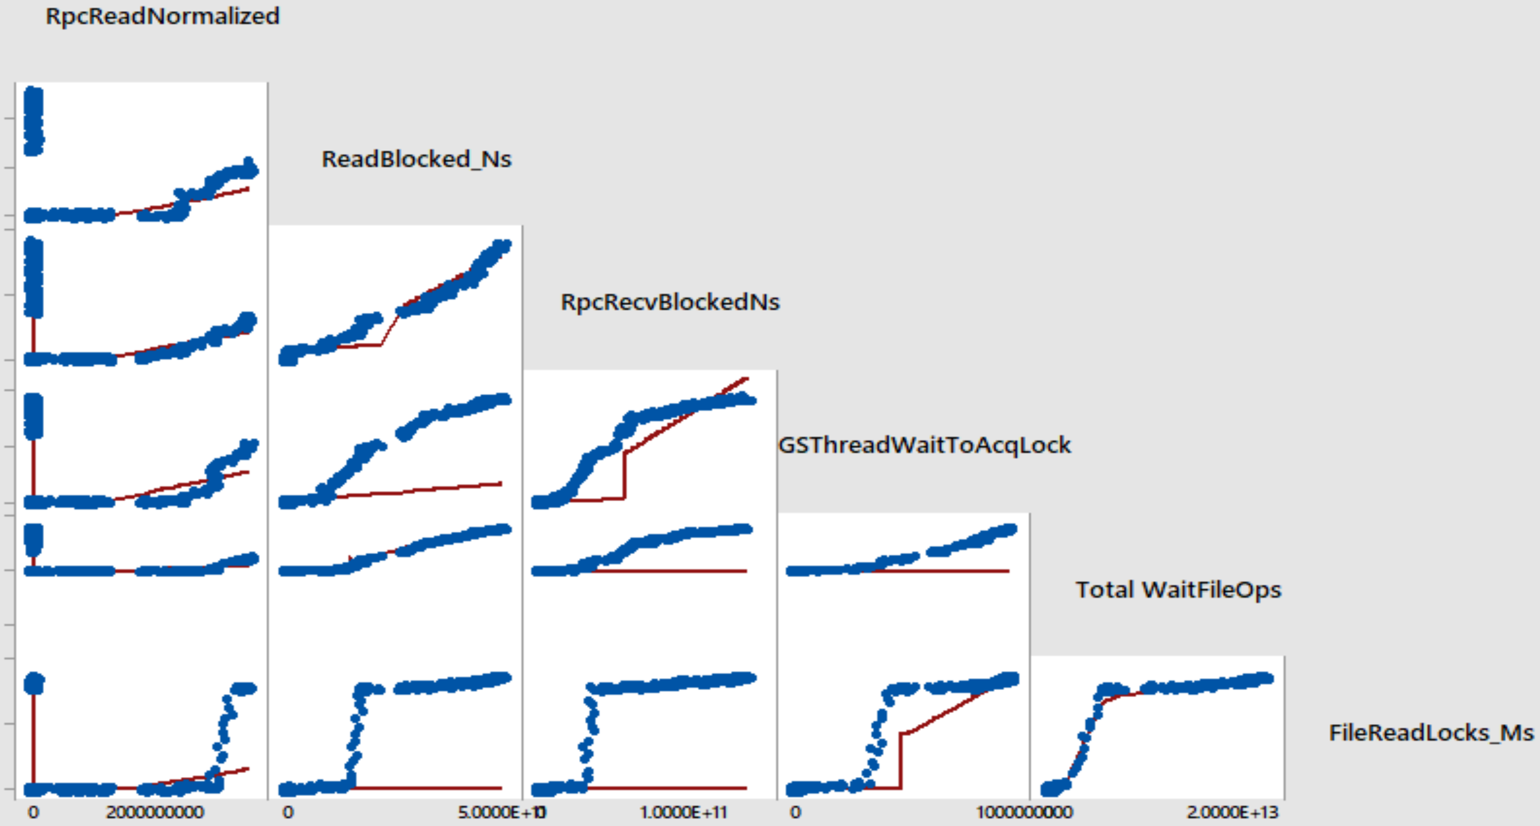

Data

A1 | timestamp

	A	B	C	D	E	F
1	timestamp	avgThroughputIOps	avgThroughputNormalizedIOps	avgLatencyMs	avgRequestsQueued	avgThroughputIOps
2	2019-04-18T11:01:33-07:00					
3	2019-04-18T11:02:34	0.0658182	0.0658182	0.735073	4.83812E-05	0.0658182
4	2019-04-18T11:03:34-07:00					
5	2019-04-18T11:04:35-07:00					
6	2019-04-18T11:05:36	0.0659136	0.0659136	0.712327	4.69521E-05	0.0659136
7	2019-04-18T11:06:37-07:00					
8	2019-04-18T11:07:37	0.0658619	0.0658619	1.08836	7.16811E-05	0.0658619
9	2019-04-18T11:08:38-07:00					
10	2019-04-18T11:09:39	0.0329149	0.0329149	0.926226	3.04866E-05	0.0329149
11	2019-04-18T11:10:40	0.0658619	0.0658619	0.574063	3.78089E-05	0.0658619

SDC 19 SNIA INDIA

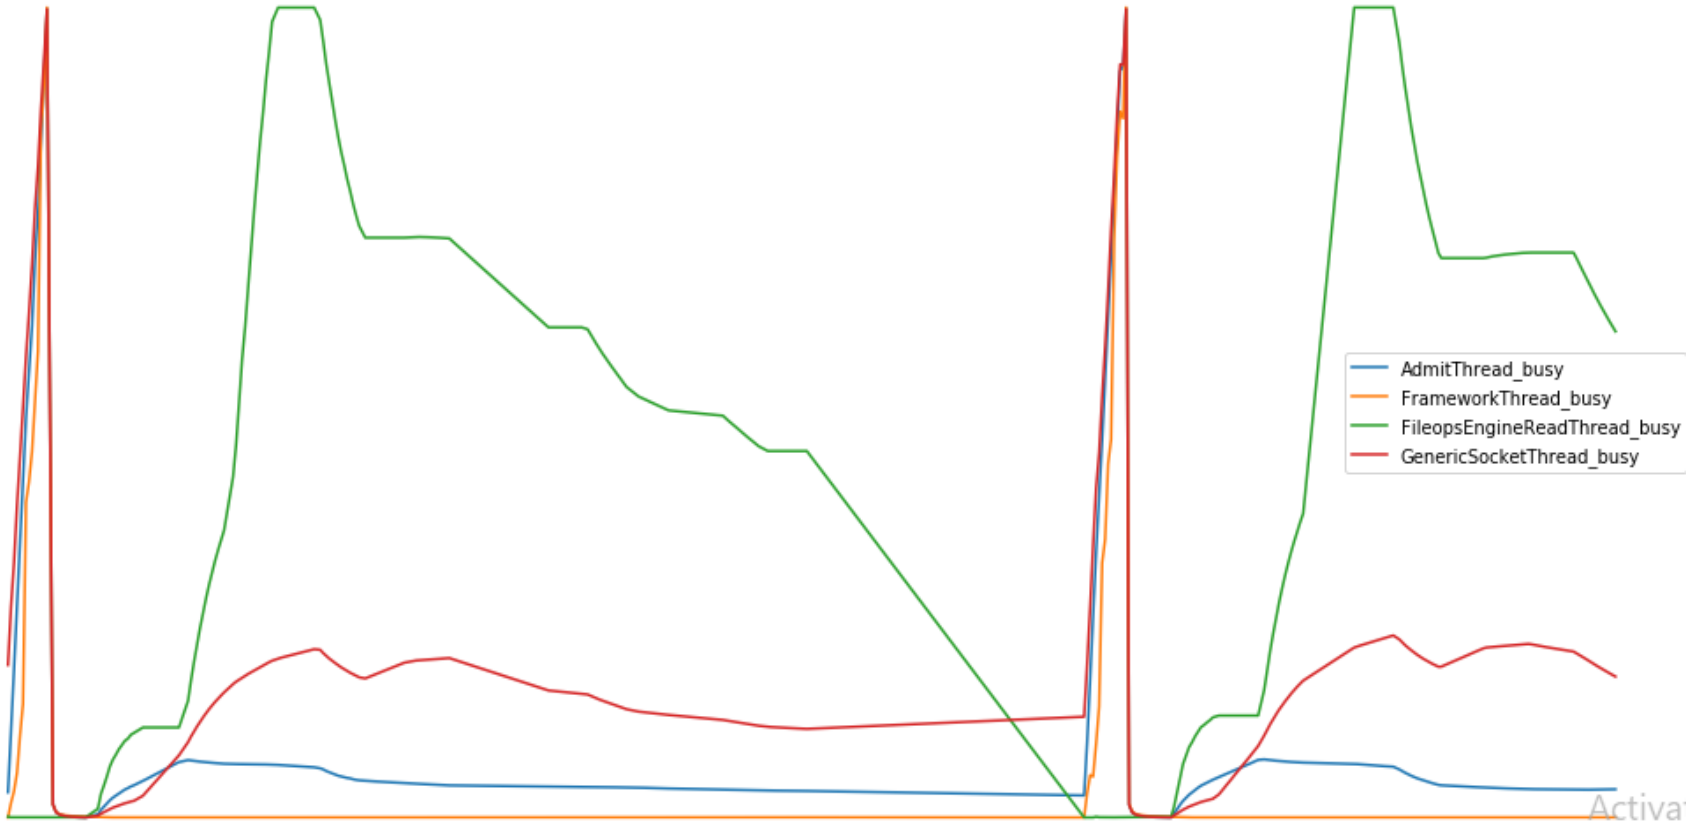
AdmitThread_busy
FrameworkThread_busy
FileopsEngineReadThread_busy
GenericSocketThread_busy

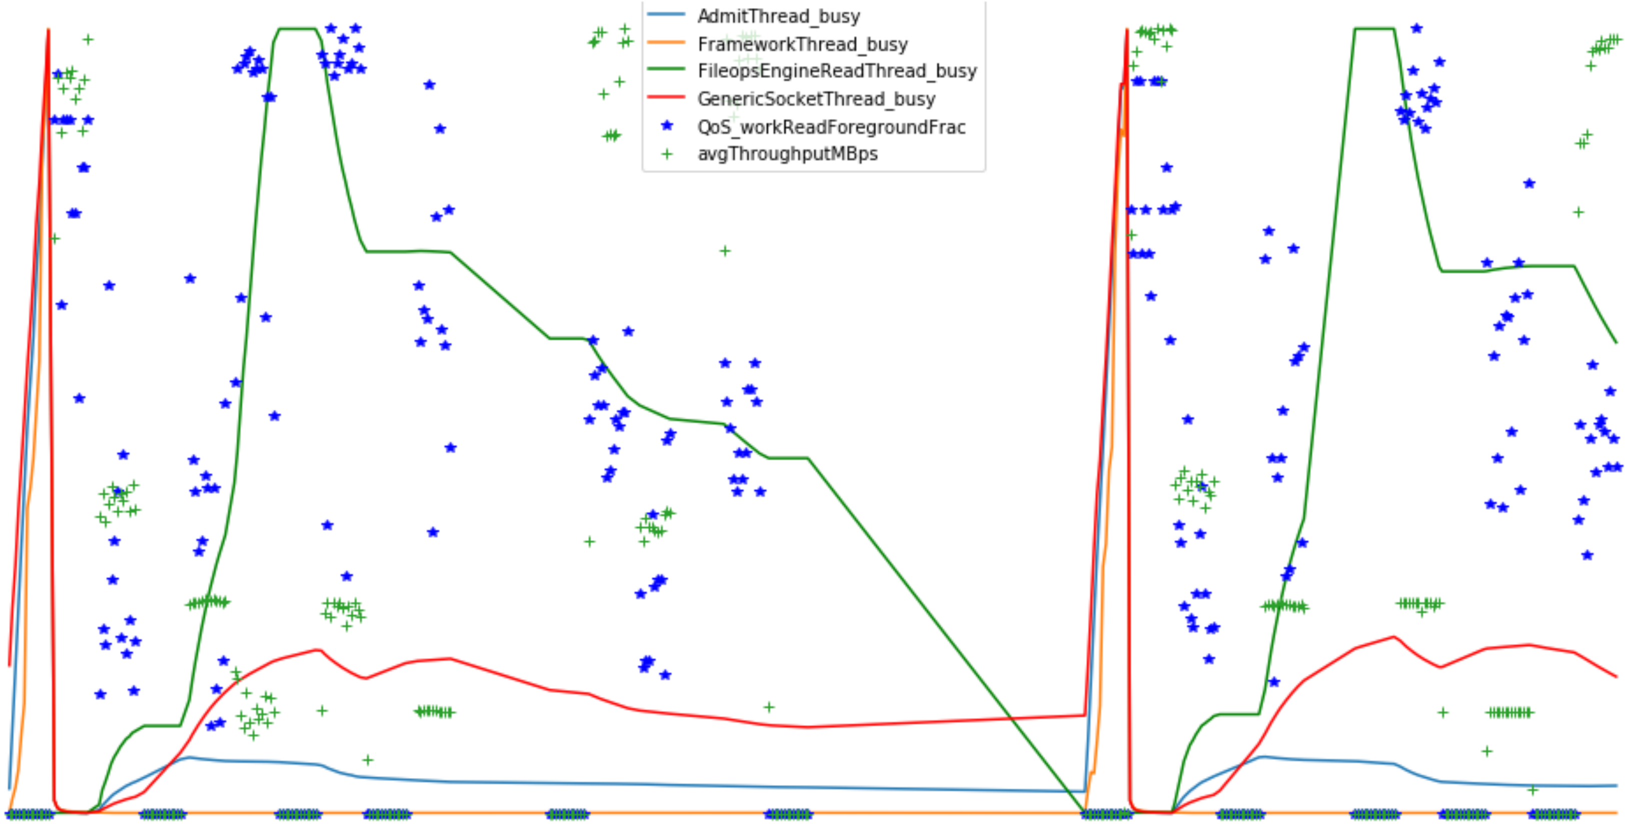
AdmitThread_busy
FrameworkThread_busy
FileopsEngineReadThread_busy
GenericSocketThread_busy
QoS_workReadForegroundFrac
avgThroughputMBps

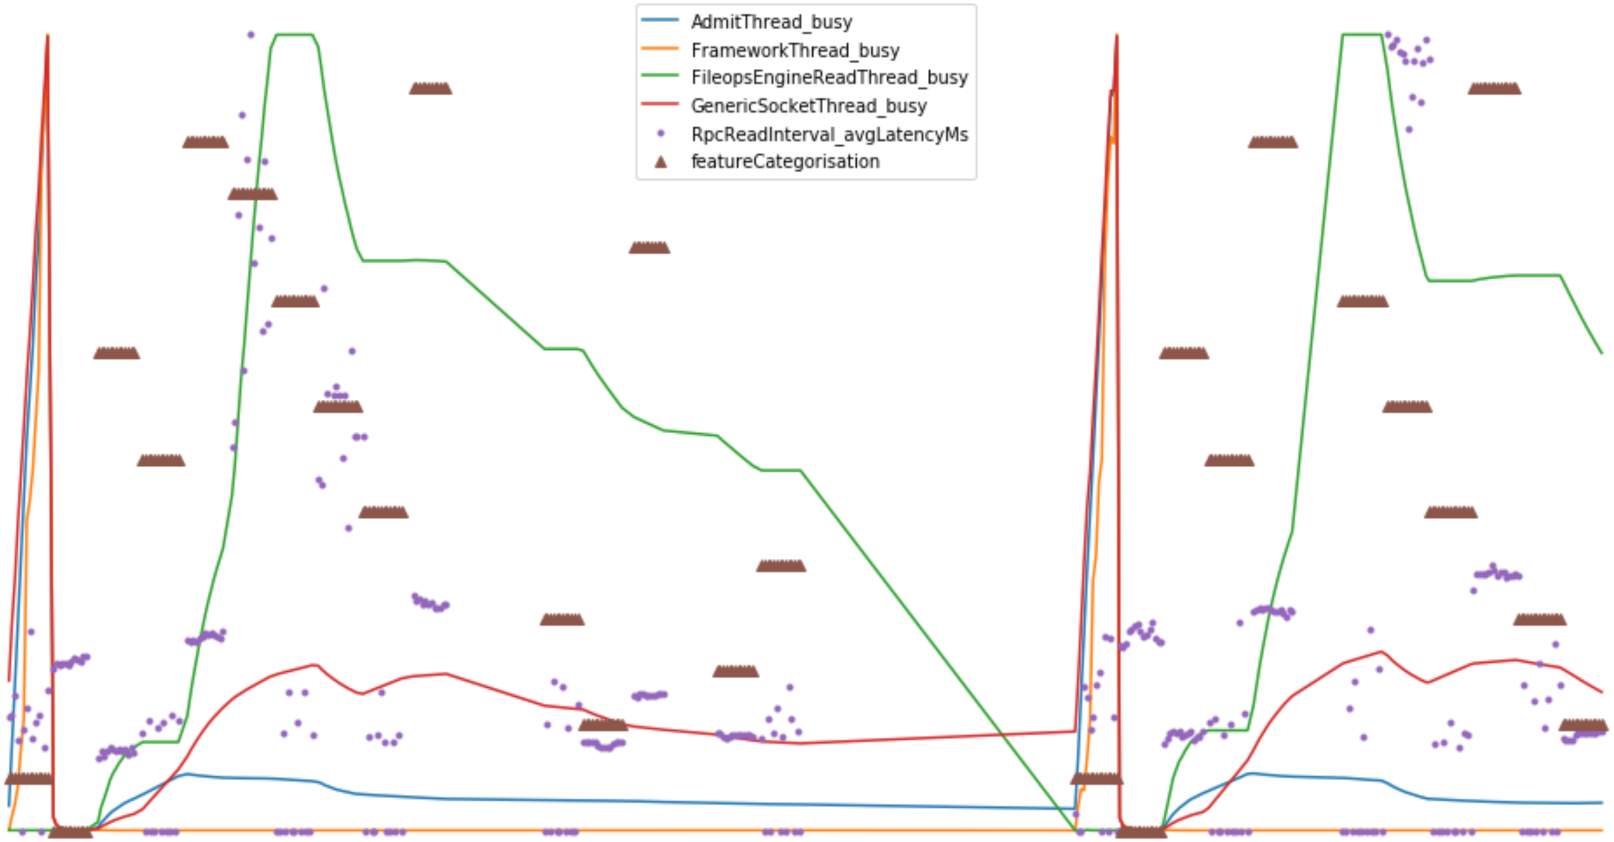

SDC 19 SNIA INDIA

Inference / Observations

- Projections are continuous in low dimensional space and similar histories get clustered
- Parameter sharing among similar histories
- K-State Markov Model (HMM, Memory less model)
- Predictive State Representative Models
- Time Varying Models
- StateFul Neural Networks Models (RNN ?)

Context Sensitive Models

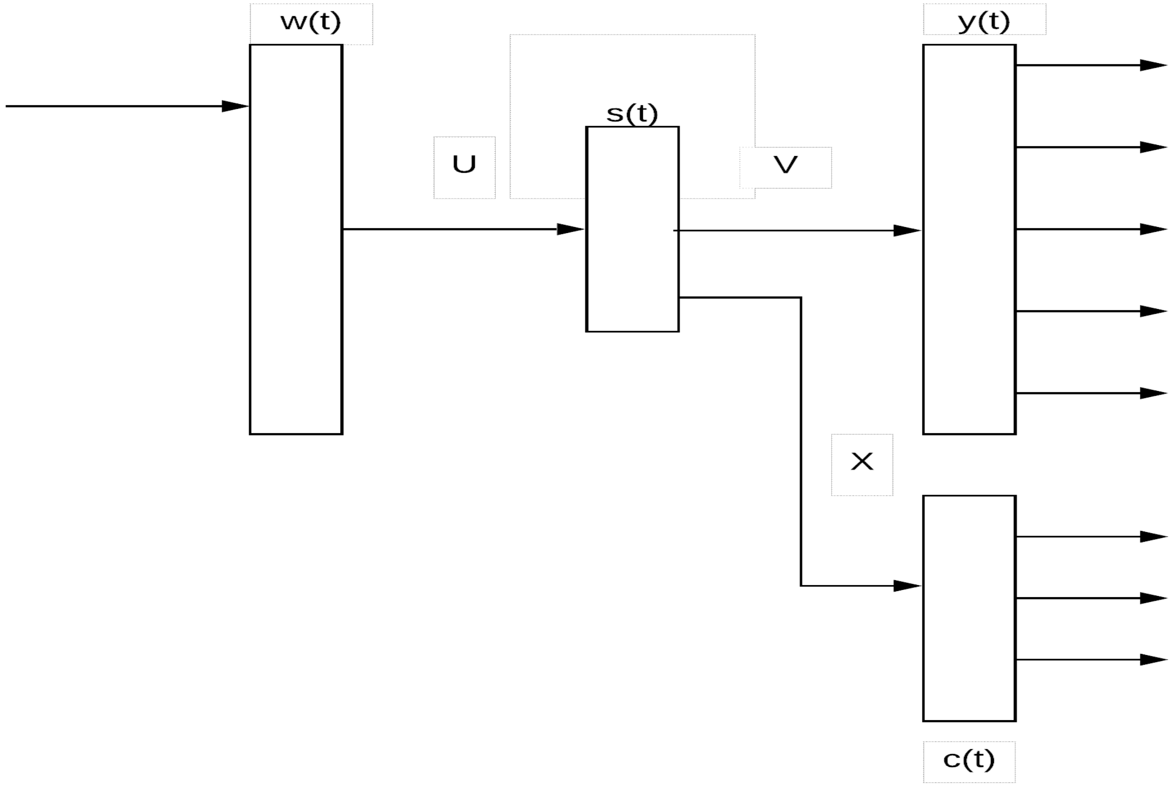

Time Sensitive Models

- ARIMA
- NN
- RNNs
- RNNs with LSTM (Vanilla LSTM)
- RNN Bidirectional LSTM
- RNN ESN LSTM

ARIMA

- Pros
  - Dependent on relationship with past observations
  - Uses differences of Raw Observations
- Cons
  - Lag or the size of the Moving Average window be known
  - The number of times, the Raw Observations, are differenced
  - The number of Raw Observations

SDC 19
SNIA INDIA

Neural Networks

- Pros
  - Works slightly better when it comes to Classification Problems
  - Can be deployed to predict almost any thing
- Cons
  - No insights for the reason why it did or did not work
  - Treats samples as just observations
  - Has no mechanism to predict for Time series data

Context Sensitive Time Variant Models

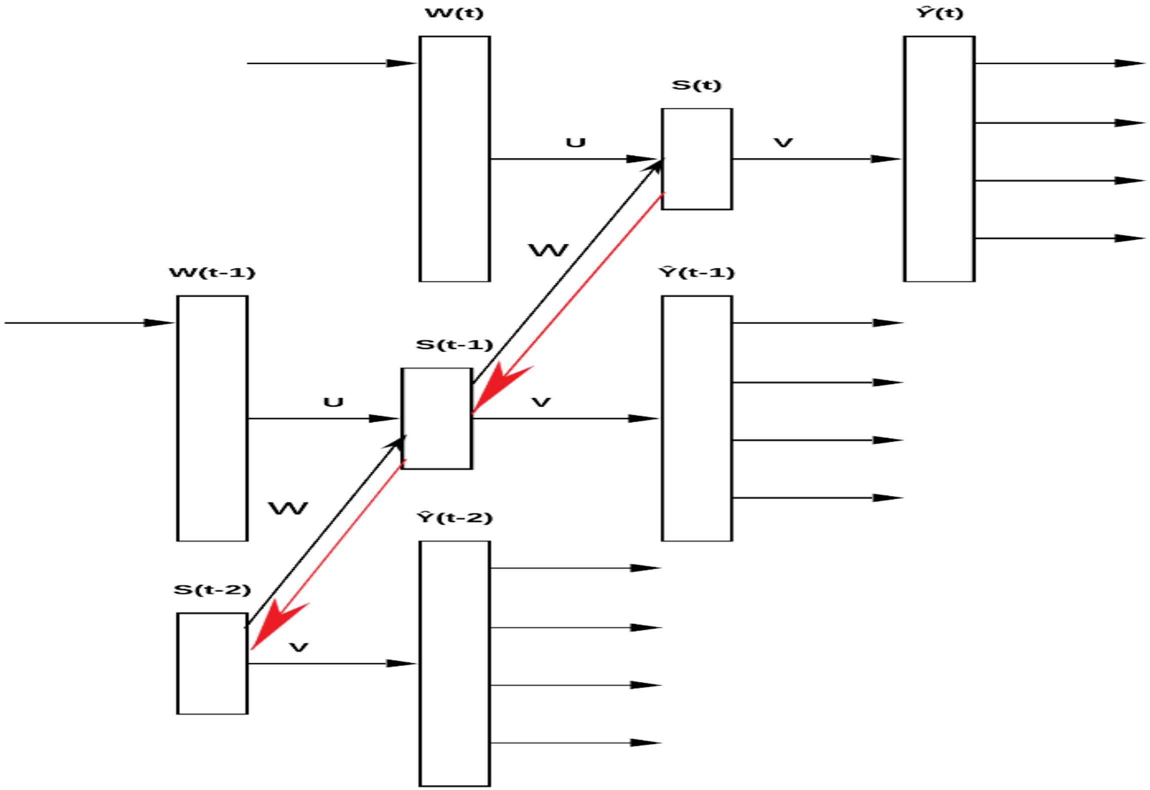

RNN

- Pros
  - Works better when data has time variations
  - Memorizes sequences well
  - Has dynamic state with context-dependent computation (Vs HMM)
- Cons
  - Are not necessarily Inductive
  - Problem of Vanishing and Exploding gradients
  - Results vary abruptly based on Models/Layers chosen

RNNs with LSTM

- Pros
  - Embeds integrators for memory storage in network
  - Remembers sequence to predict the outcomes

- Cons
  - Internal States must be selected to use memory, not abuse.
  - Works very bad when all the data is correlated
  - Poor results when times series responses are chaotic

RNN Stacked LSTMs

- Pros
  - Create new representations at high levels of abstractions
  - Increasing depths trade off with fewer neurons that trains faster
- Cons
  - Involve huge computational resources
  - Not a silver bullet

SDC 19
SNIA INDIA

Results

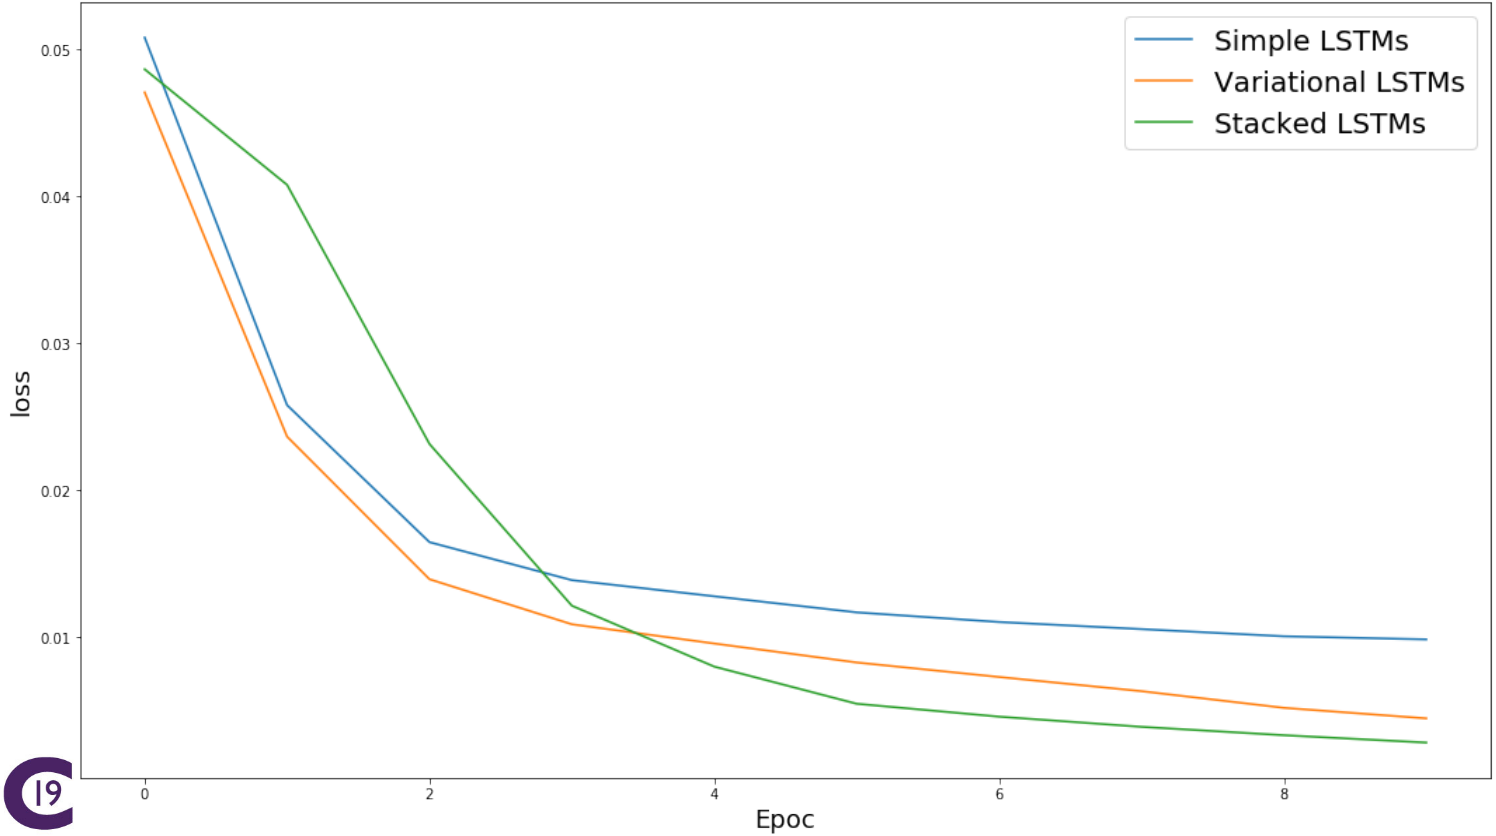
Simple LSTMs
Variational LSTMs
Stacked LSTMs
loss
0.05
0.04
0.03
0.02
0.01
0
2
4
6
8
Epoc

Futures and Scope

- RNN ESN LSTMs
- Pros
  - Work well for chaotic time series
  - Reservoirs show non-linearity w.r.t inputs
  - Outputs show linear regression
- Cons
  - Training is a challenge

SDC 19
SNIA INDIA

Scope and Futures

- CSPs can use this Recommendation Engine to make decisions regarding Eviction Policy to improve SLAs.
- Traditional methods of modelling unable to catch up with the complexities of the system
- RNN, LSTMs, Stack LSTM, ESN provide us with huge scope to tune the model
- Models Vs Complex or Stacked Models share similar relationship as series and Fourier series

Queue Analysis

- Dispatching Discipline
  - Priority Based, FCFS
- Distribution of Arrivals
  - Poisson
- Distribution of Service Times
  - Depends on Request Size
  - Rate of Requests
  - The state of the system.

Queue Analysis

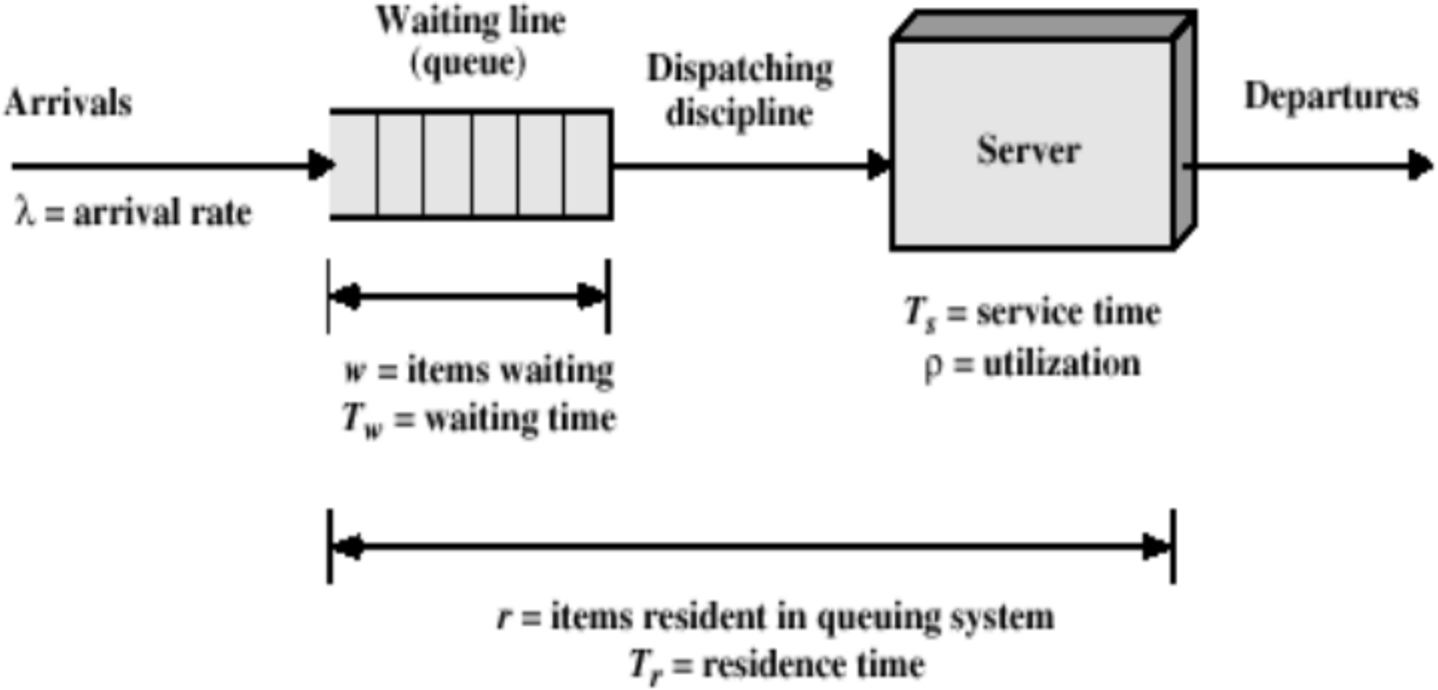

Law's At Play

- $L = \lambda * \dot{W}$ Little's Law
- Service Time/Response Time = 1 – Utilization
  - $(R_T - S_T)/R_T$ = Utilization
- $W(t+1) = W(t) + \sum_{z=0}^{T} s(t-z-1)\, e_h(t-z)^T$
- Central Limit Theorem

Bibliography and Credits

- [Reservoir Computing Approaches to Recurrent Neural Network Training]
- [Learning Non-Linear Models] (https://www.cs.cmu.edu/~rsalakhu/papers/dynamic.pdf)
- Credits for Graphs and Data Cleaning – Rahul Tripathi
- https://drive.google.com/open?id=1I98SeI_LV4FOkY10PEfV_YLwQV-5WDgf

SDC 19 SNIA INDIA

Q&A

- “Introduction of Numbers as Co-ordinates is an Act of Violence”

  - Hermann Weyl

- “All models are wrong some are more useful”

SDC 19 SNIA INDIA