



Convolutional LSTM models to estimate network traffic

NOTED

Joanna Waczyńska

Edoardo Martelli

Sofia Vallecorsa

Tony Cass

Edward Karavakis

May 19, 2021

1 Introduction & Motivation

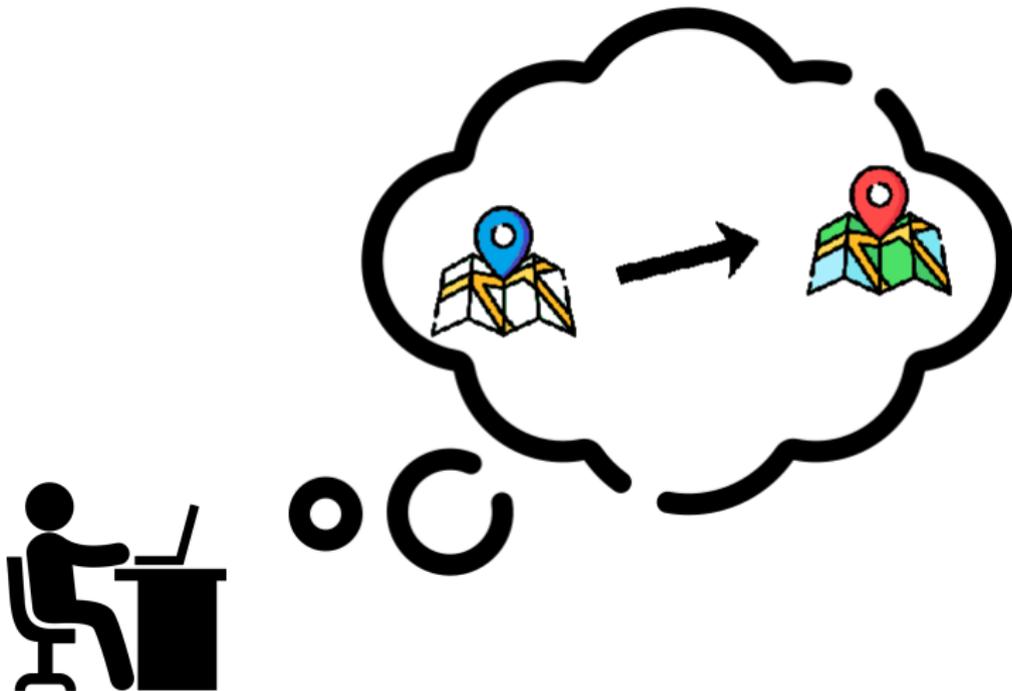
2 Traffic forecasting

- Data and models
- Results

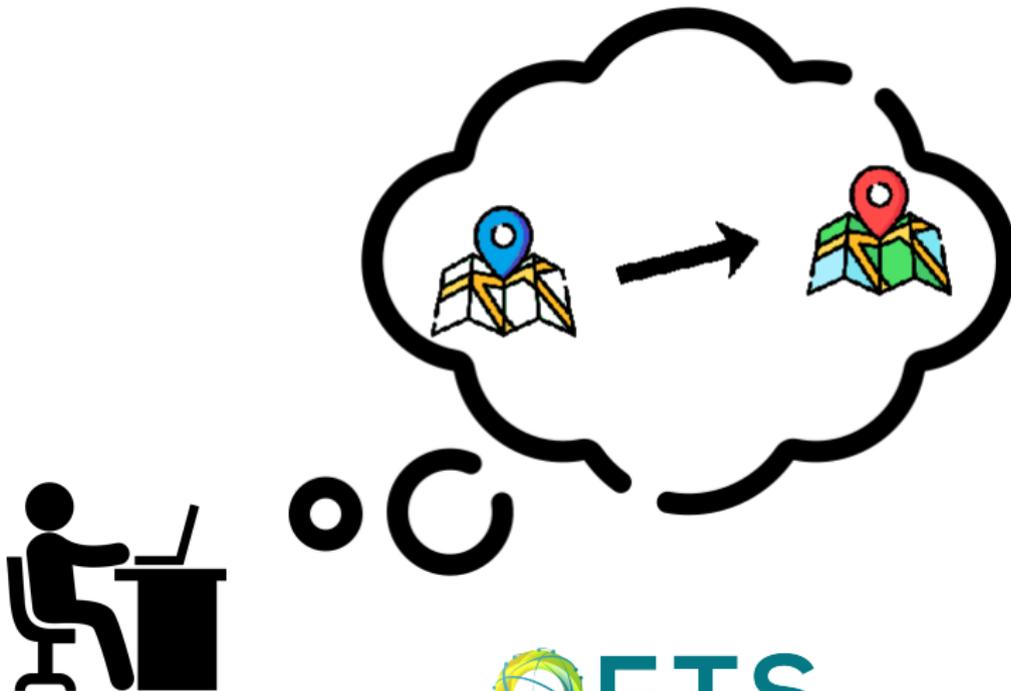
3 Conclusion

Section 1

Introduction & Motivation

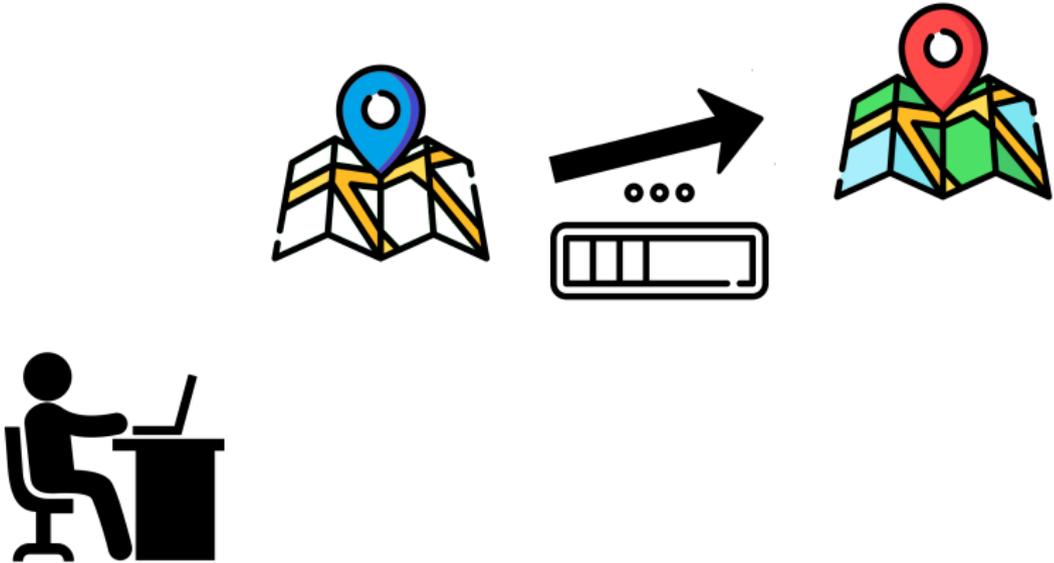


Icons made by Freepik from www.flaticon.com



 **FTS**
File Transfer Service

Icons made by Freepik from www.flaticon.com



Icons made by Freepik from www.flaticon.com



Icons made by Freepik from www.flaticon.com

LHCOPN

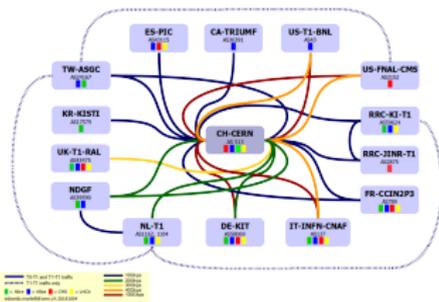


Figure: LHCOPN (Large Hadron Collider Optical Private Network) topology

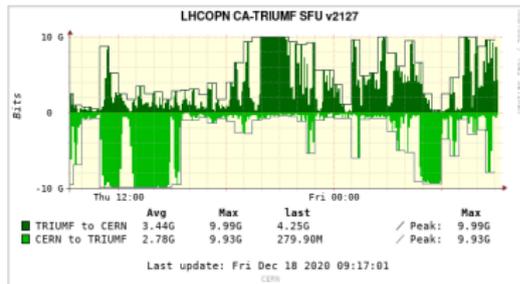
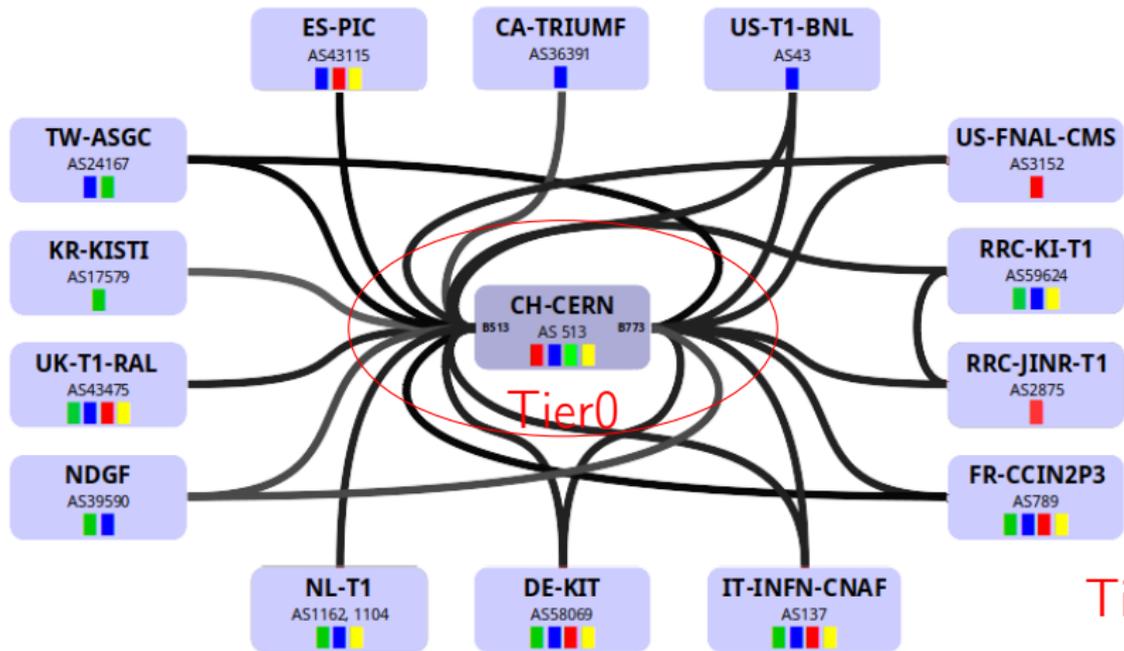
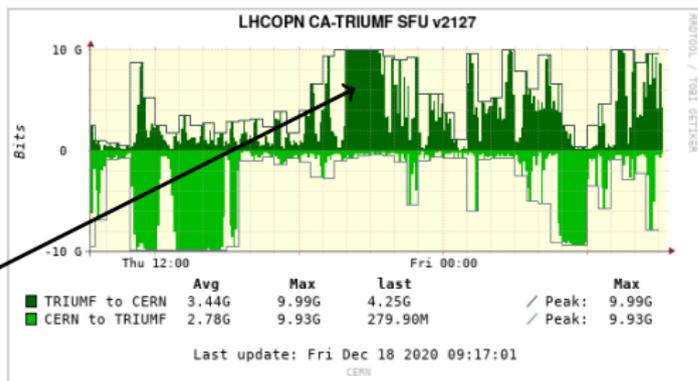


Figure: Focusing only on **one link** – network traffic observed on the LHCOPN path between CERN and TRIUMF. Link saturation occurs in both directions.



Tier1

The goal of the project



We would like to **optimise transfers** of LHC data **eliminating network saturation** along the chosen path.

Solution - how to avoid saturation

Automatically

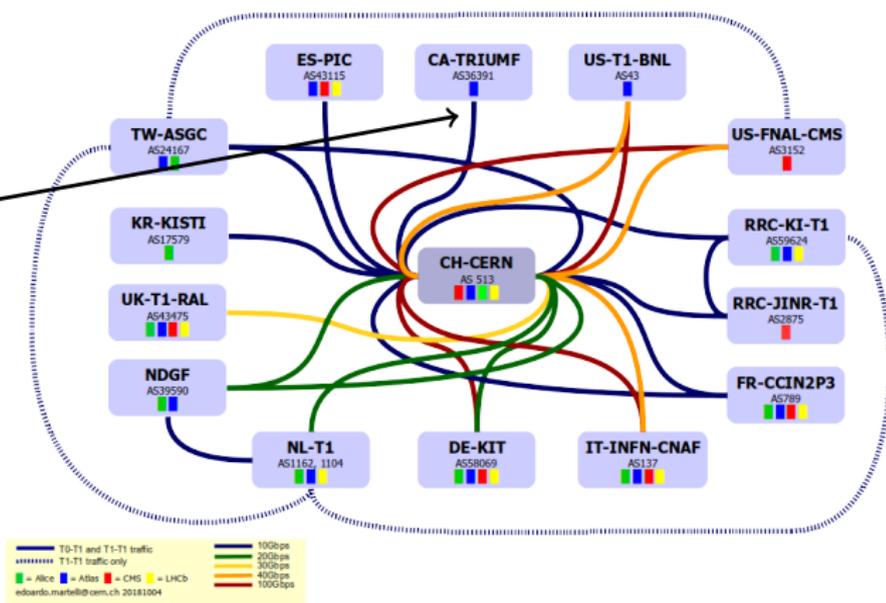
recognise when the link will be saturated for a long period of time, and **automatically modify** the configuration of network devices (SDNC) (Add extra path/link to balancing traffic).

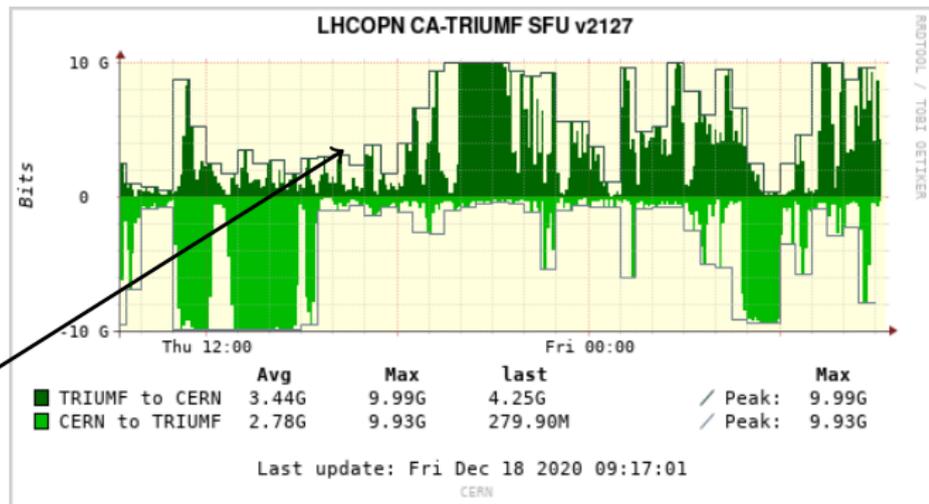
Ok, we know how, but when?

Let's assume that we have data representing aggregated information about all transfers between sites on long time in the network.

LHCOPN

X (input) -
 Aggregated
 information about
 all transfers from
 CA-Triumf to another
 sites (CH-CERN
 is the Tier0 and
 others in LHCOPN
 are the Tier1)





ROOTOOL / TOBI OETIKER

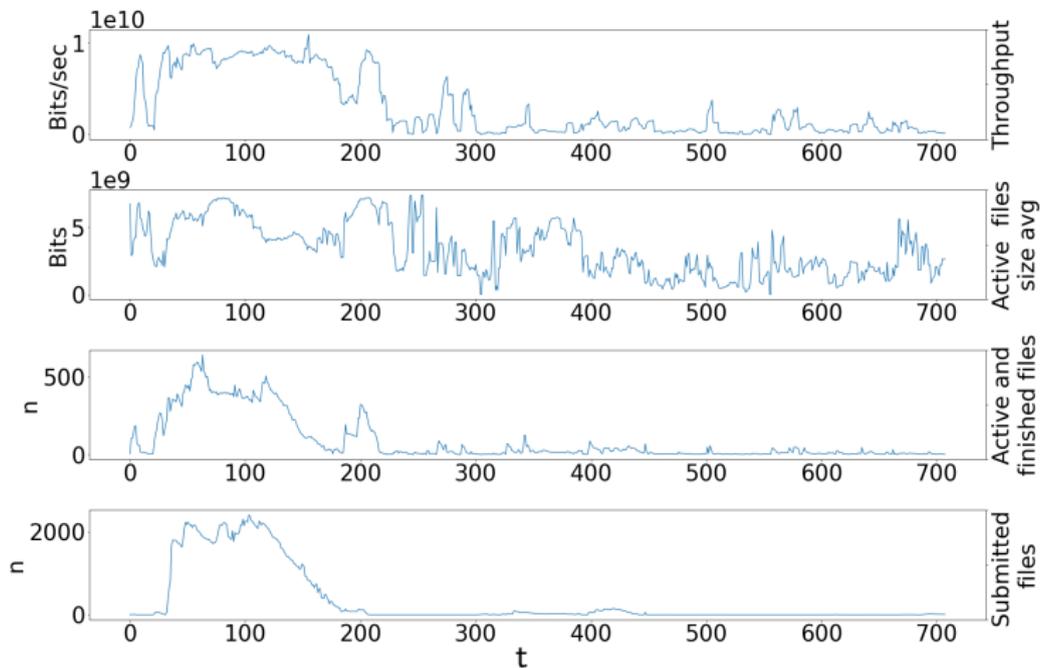
We would like to predict Y (*output*)
 - Traffic from Netstat.cern.ch

Section 2

Traffic forecasting

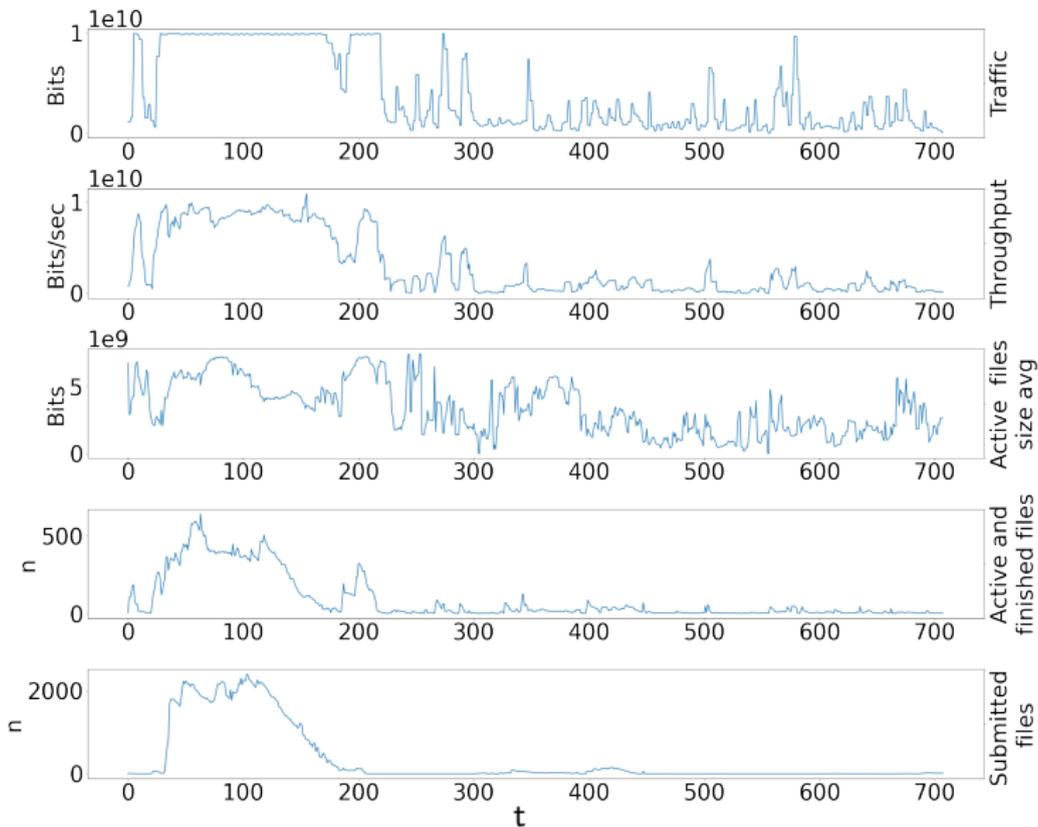
Subsection 1

Data and models

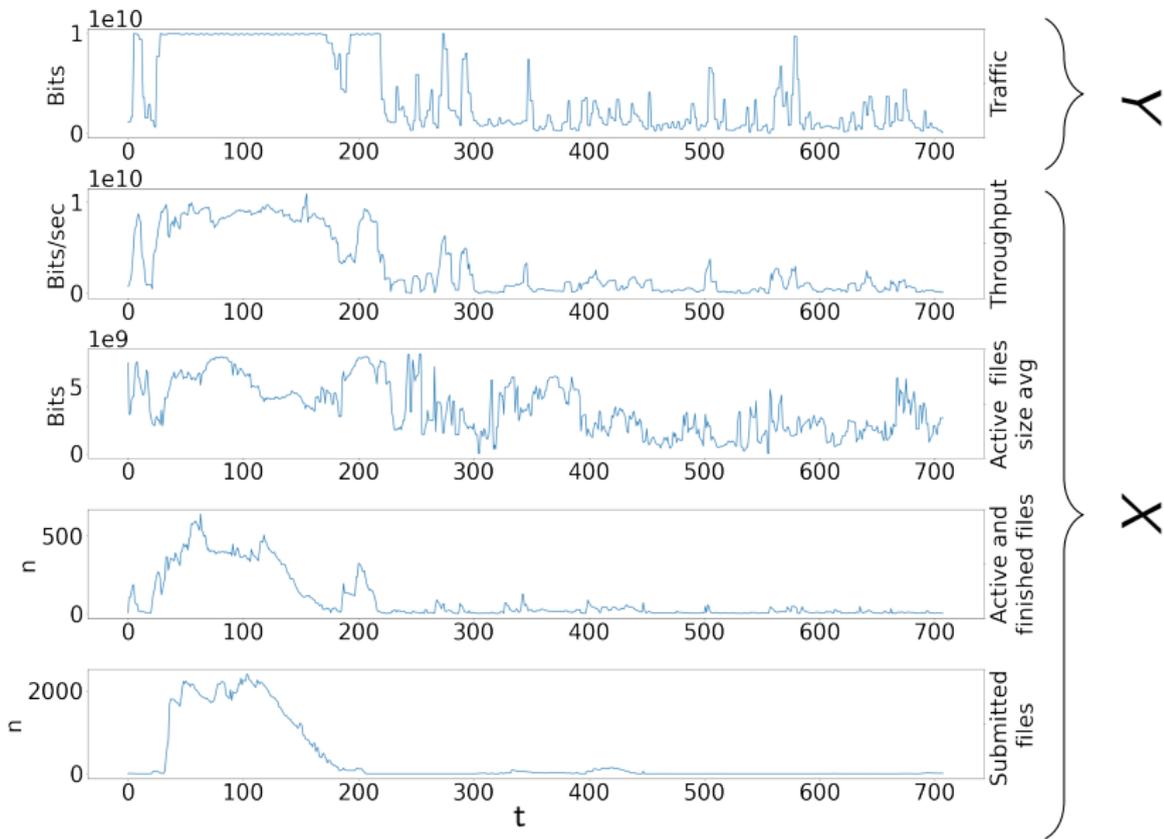


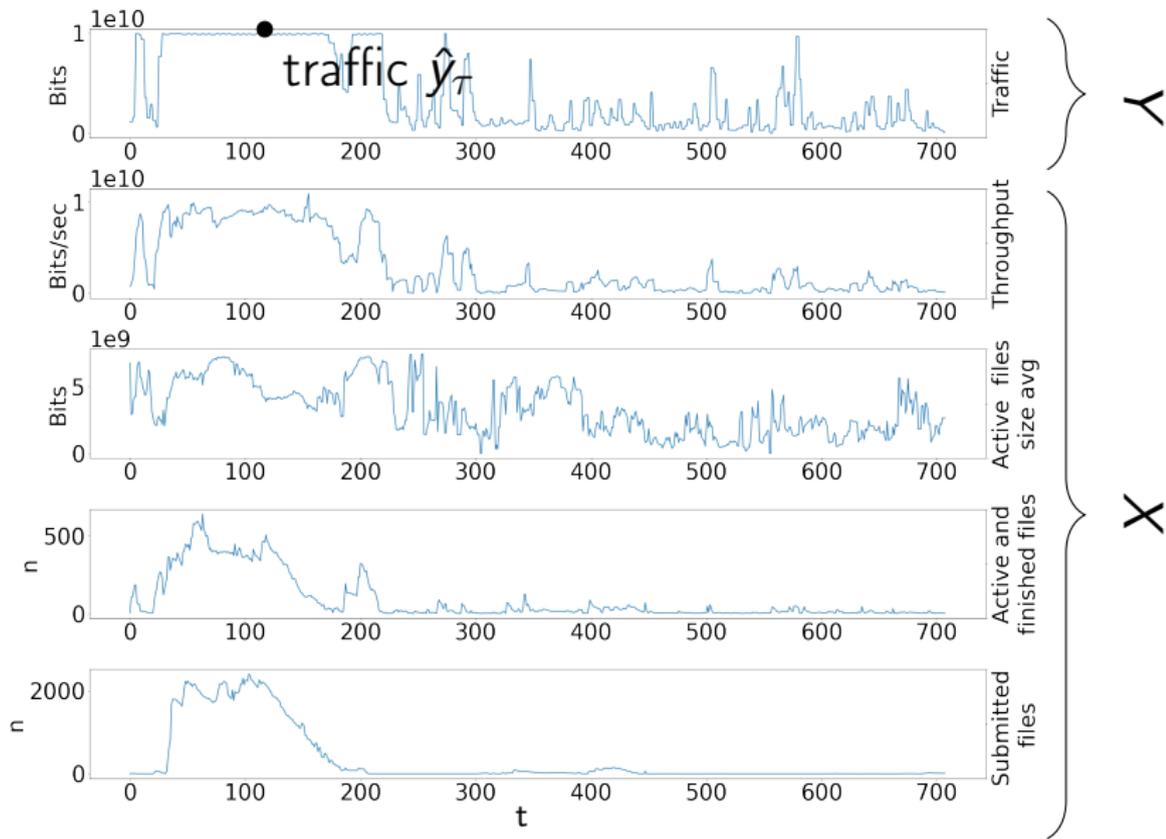
*The figure shows data used as the test data set during modeling

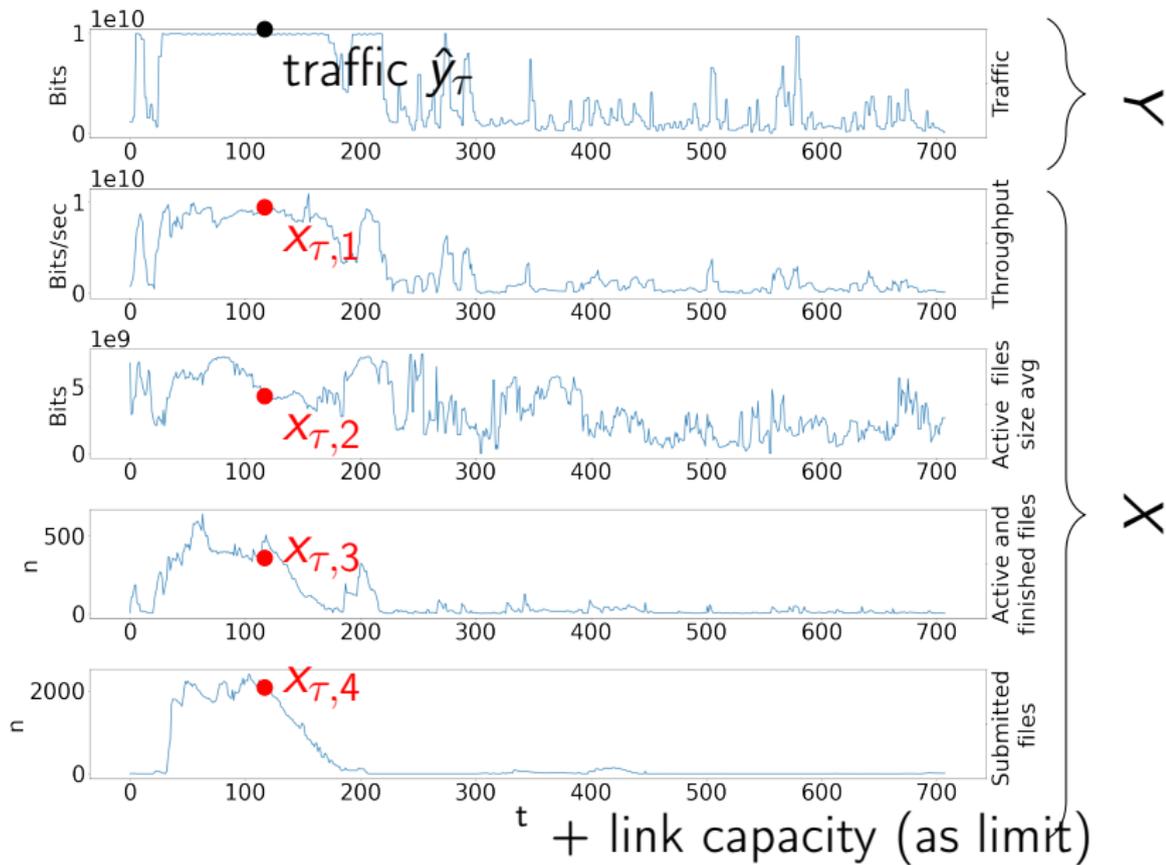
Y (real) – data from NetStat

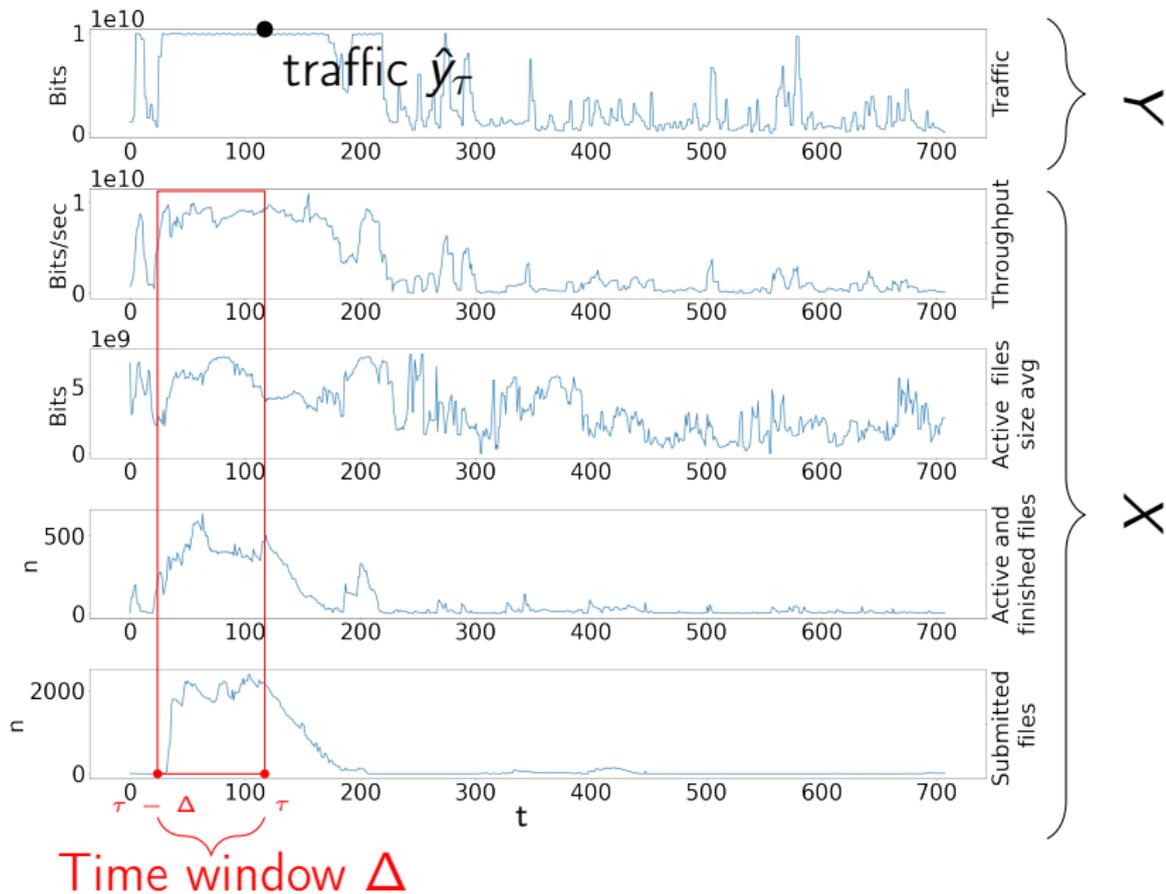


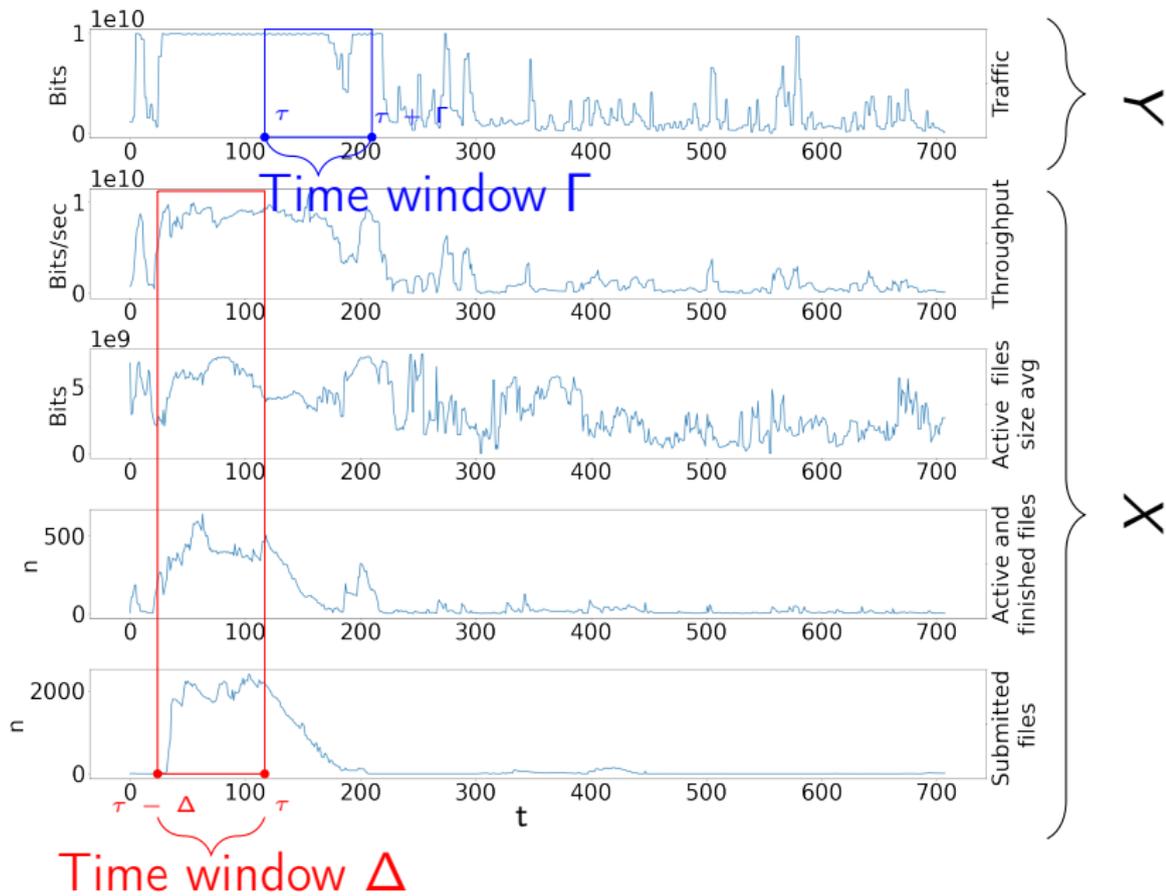
When should we add/
remove an extra link?











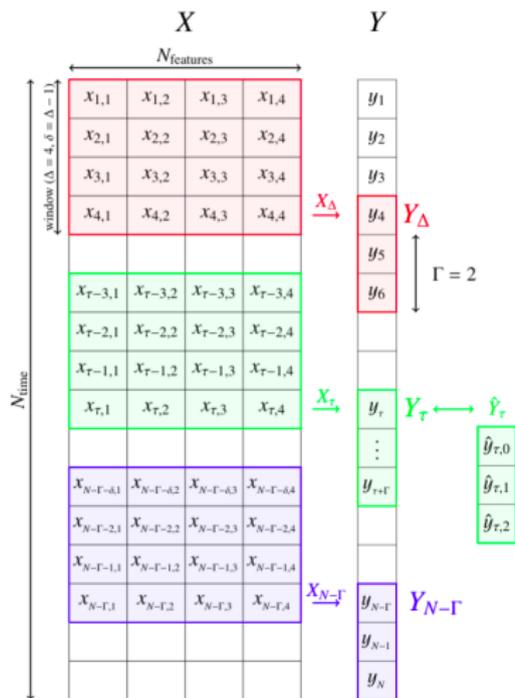
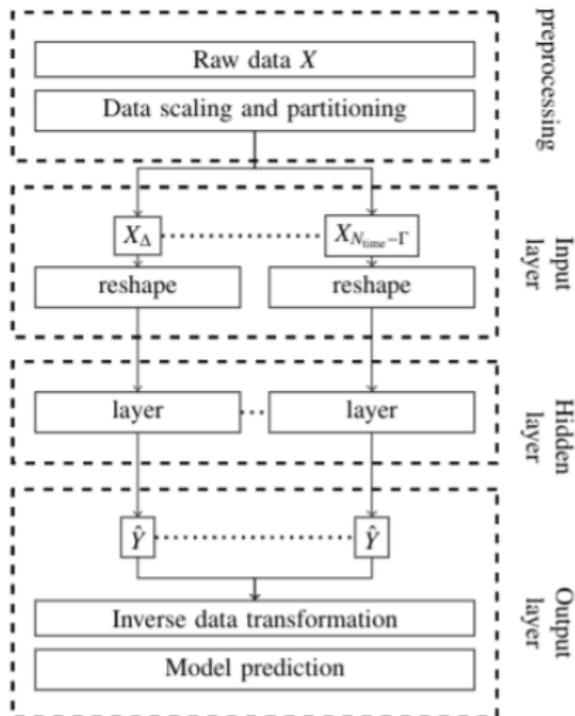


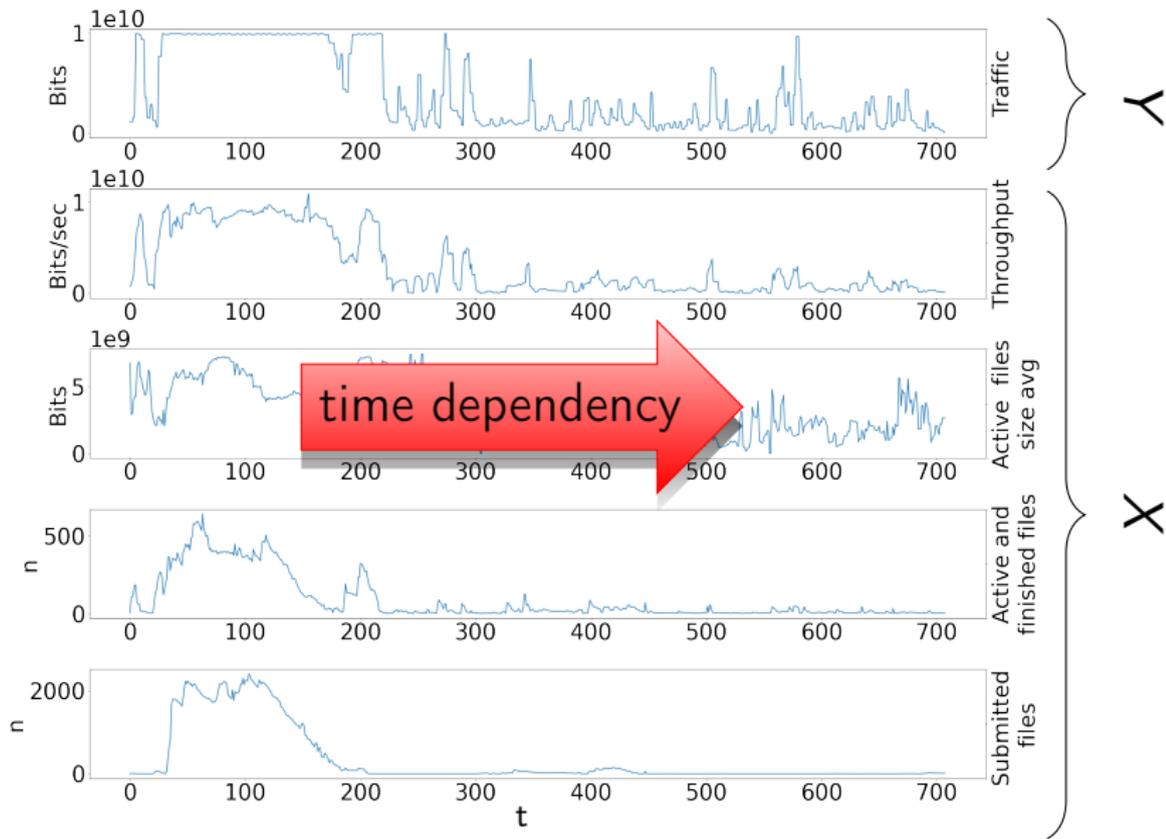
Figure: Schema input data set X , and dependent variables Y .

Example features: throughput, active files size avg, active and finished files, submitted files (queue size).

General model scheme



For some architectures, we consider additional pre-processing steps and calculate the throughput exponential moving average over the last 15 minutes, and estimated size values

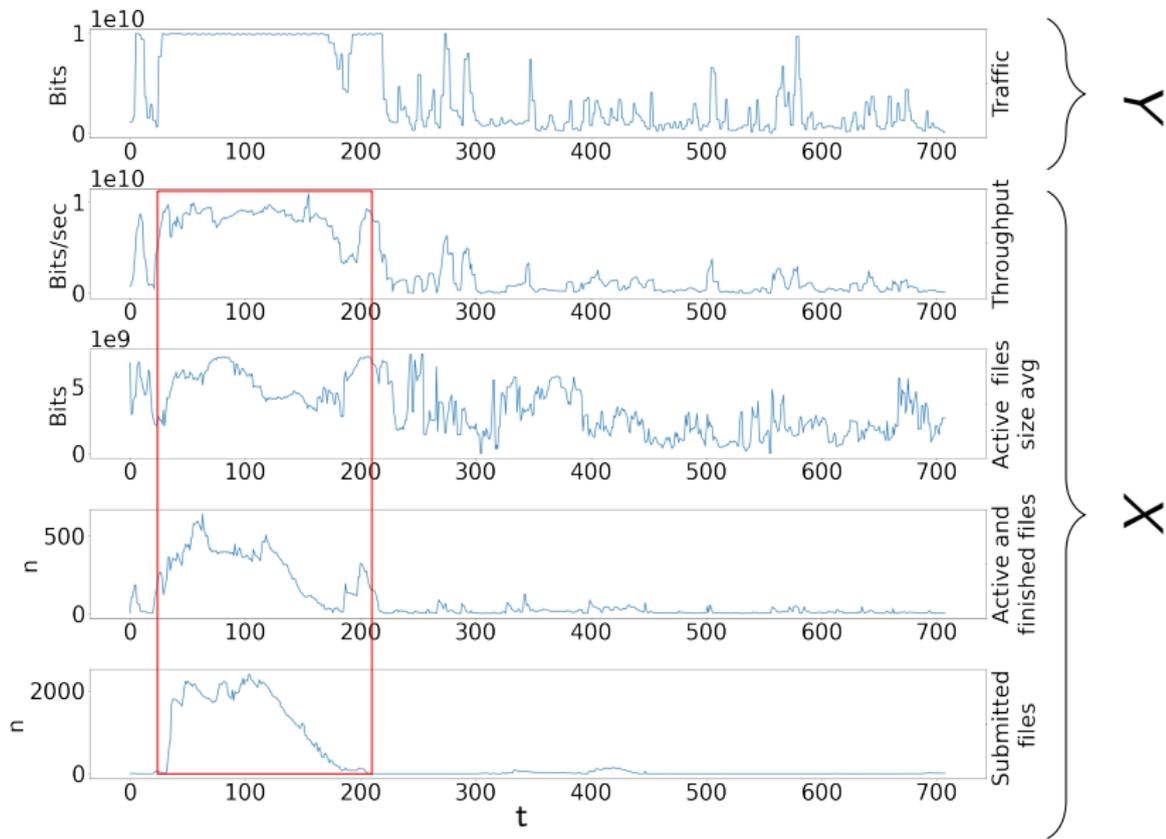


What is LSTM? Why LSTM?

- The Long Short-Term Memory network is an artificial recurrent neural network.
- LSTM has the ability to capture sequence **pattern** information **on time-series** analyses.



Illustration by Freepik Storyset



Why convolution? How does it work?

- convolutional layers are used to **extract correlations** across multiple features.

*In CNN and CNN-LSTM models we used kernel 3x3, in conv-LSTM we used kernel 1xN_features.

layer types

- convolution layer - Conv2d (model CNN)
- LSTM layer - LSTM (model LSTM)
- one convolution layer + one LSTM layer - Conv2d + LSTM (model CNN-LSTM)
- Convolutional LSTM layer - ConvLSTM2D (Conv-LSTM)

Model schemes

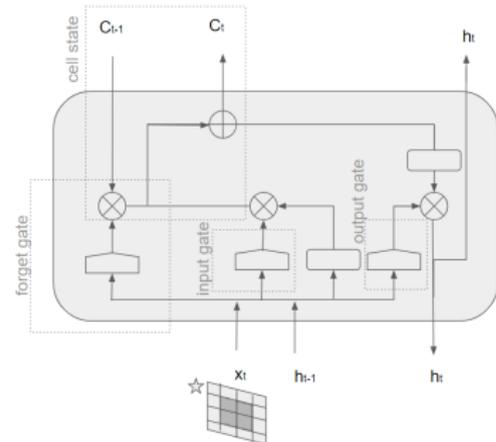


Figure: CNN-LSTM

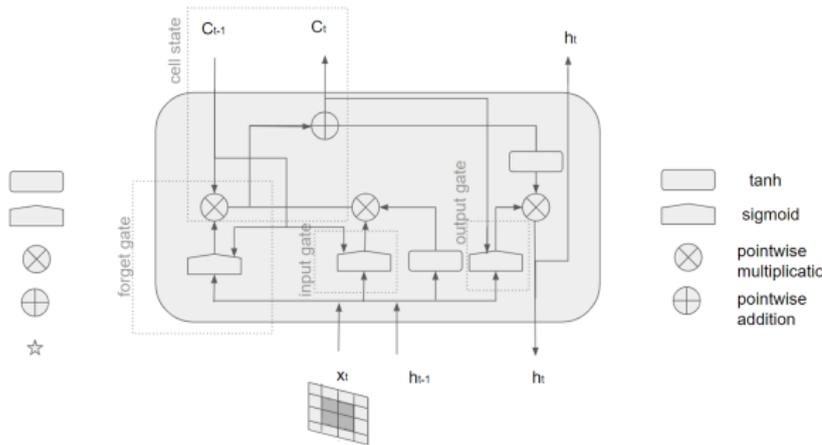
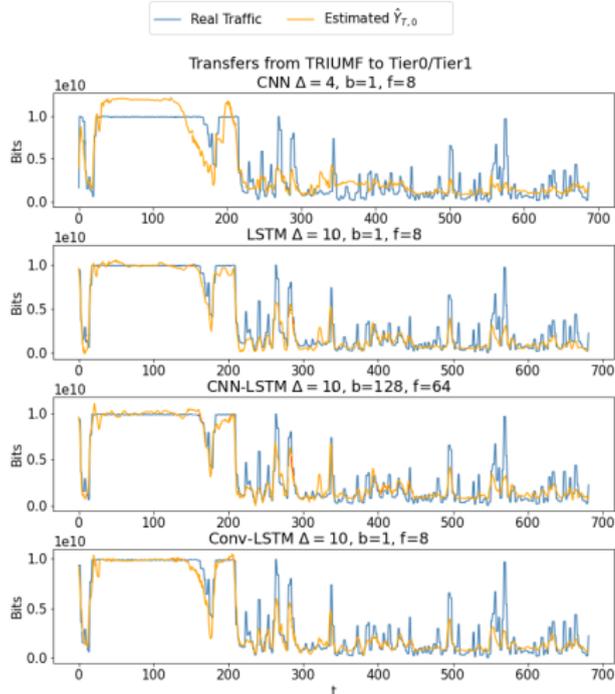


Figure: Conv-LSTM

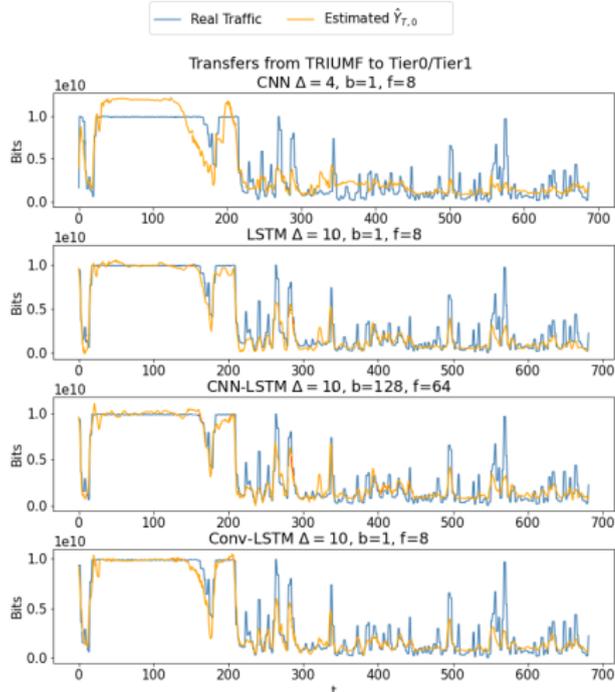
Subsection 2

Results



Prediction of instantaneous traffic

Figure: Effects of applying models with the best configuration on the test dataset. 1. We trained our model to predict traffic based on information about transfers from the FTS (from TRIUMF-SFU to Tier0/Tier1). Forecasting is based on aggregated information about transfers from last 2Δ minutes; b (batches); f (filters).

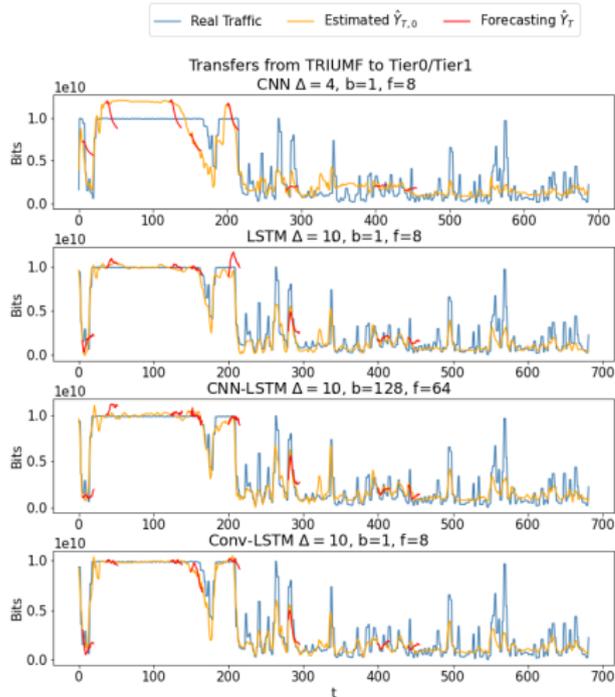


Prediction of the instantaneous traffic

Table: Model comparison. Error = mean square error for Γ window.

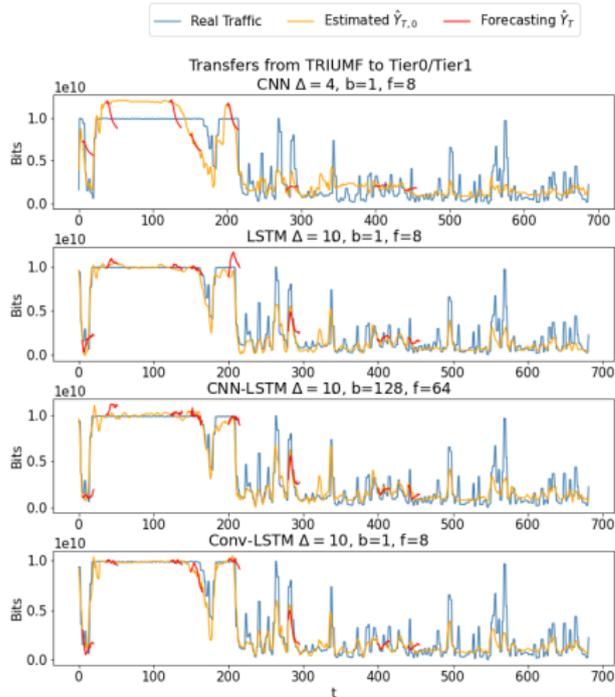
Δ	Model	Batch - Filters \ Units	Error ₀
4	CNN	1 - 8	0.206
10	LSTM	128 - 64	0.025
	CNN-LSTM	128 - 64	0.021
	CONV-LSTM	1 - 8	0.036

Figure: Effects of applying models with the best configuration on the test dataset. 1. We trained our model to predict traffic based on information about transfers from the FTS (from TRIUMF-SFU to Tier0/Tier1). Forecasting is based on aggregated information about transfers from last 2Δ minutes; b (batches); f (filters).



Prediction of future traffic

Figure: Effects of applying models with the best configuration on the test dataset. 1. We trained our model to predict traffic based on information about transfers from FTS (from TRIUMF-SFU to Tier0/Tier1). Forecasting is based on aggregated information about transfers from last 2Δ minutes; We predict the next Γ samples. b (batches); f (filters).



Prediction of future traffic

Table: Model comparison. Error = mean square error for Γ window.

Δ	Model	Batch - Filters \ Units	Error $_{\Gamma}$
4	CNN	1 - 8	0.206
10	LSTM	128 - 64	0.185
	CNN-LSTM	128 - 64	0.188
	CONV-LSTM	1 - 8	0.125

Figure: Effects of applying models with the best configuration on the test dataset. 1. We trained our model to predict traffic based on information about transfers from FTS (from TRIUMF-SFU to Tier0/Tier1). Forecasting is based on aggregated information about transfers from last 2Δ minutes; We predict the next Γ samples. b (batches); f(filters).

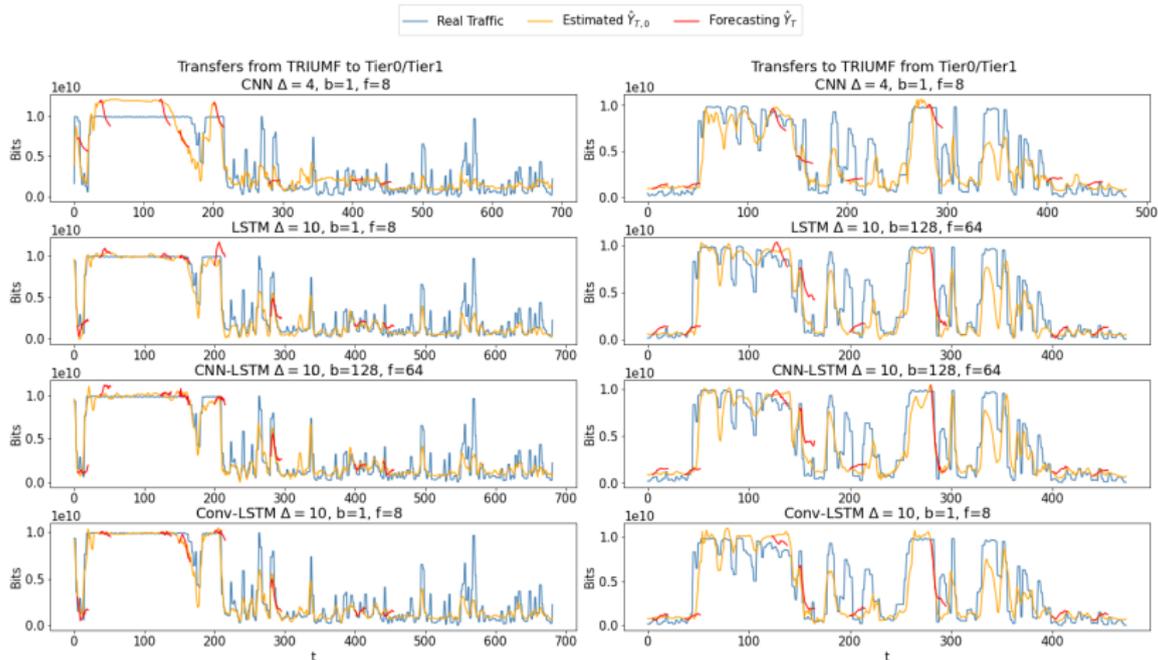


Figure: Effects of applying models with the best configuration on the test dataset. 1. We trained our model to predict traffic based on information about transfers from FTS (from TRIUMF-SFU to Tier0/Tier1). 2. Pictures present results on two data sets: from TRIUMF-SFU to Tier0/Tier1 and from Tier0/Tier1 to TRIUMF-SFU. Forecasting is based on aggregated information about transfers from last 2minutes; b (batches); f(filters). Forecasting \hat{Y}_T is predicted Γ future values for chosen t . Here $\Gamma = 15$ (30 min).

Table: Comparison of model parameters on the test data set representing transfers from TRIUMF to Tier0/Tier1. $\Gamma = 15$ (30 minutes). S is the standard deviation over 10 training repetitions.

Δ	Model	Batch - Filters \ Units	$Error_{\Psi}$	$S(Error_{\Psi})$	$Error_{\Psi,0}$	$S(Error_{\Psi,0})$
4	CNN	1 - 8	0.206	0.007	0.206	0.009
	LSTM	128 - 64	0.224	0.008	0.042	0.005
	CNN-LSTM	128 - 64	0.233	0.015	0.060	0.007
	CONV-LSTM	1 - 8	0.159	0.012	0.048	0.007
10	CNN	1 - 8	0.223	0.095	0.223	0.010
	LSTM	128 - 64	0.185	0.012	0.025	0.006
	CNN-LSTM	128 - 64	0.188	0.011	0.021	0.006
	CONV-LSTM	1 - 8	0.125	0.008	0.036	0.008

- $Error_{\Psi,0}$ – $MSE_{\Psi,0}(\Delta, \Gamma)$ for input window time Δ and output window time Γ . We calculate the mean square error only for estimation traffic $y_{\tau,0}$ on time τ , where $\tau \in \Psi$
- $Error_{\Psi}$ – $MSE_{\Psi}(\Delta, \Gamma)$ We calculate the mean square error for all estimation traffic $y_{\tau,\gamma}$ on time τ , where $\tau \in \Psi$, and $\gamma \in [0, \dots, \Gamma]$
- Ψ period when observed link (TRIUMF -> CERN) was overloaded.

Section 3

Conclusion

Conclusion

- Our work shows that **CNN-LSTM** and **Conv-LSTM** architectures can indeed enable us to detect network saturation

Conclusion

- Our work shows that **CNN-LSTM and Conv-LSTM architectures can indeed enable us to detect network saturation**
- They provide **great forecasting** accuracy even over long time periods (**up to 30 minutes**) based on short history (time window).

Conclusion

- Our work shows that **CNN-LSTM and Conv-LSTM architectures can indeed enable us to detect network saturation**
- They provide **great forecasting** accuracy even over long time periods (**up to 30 minutes**) based on short history (time window).
- We consider CNN-LSTM as the best prediction for instantaneous traffic prediction

Conclusion

- Our work shows that **CNN-LSTM and Conv-LSTM architectures can indeed enable us to detect network saturation**
- They provide **great forecasting** accuracy even over long time periods (**up to 30 minutes**) based on short history (time window).
- We consider CNN-LSTM as the best prediction for instantaneous traffic prediction
- Conv-LSTM as the most suitable model to predict the end of saturation.

Conclusion

- Our work shows that **CNN-LSTM and Conv-LSTM architectures can indeed enable us to detect network saturation**
- They provide **great forecasting** accuracy even over long time periods (**up to 30 minutes**) based on short history (time window).
- We consider CNN-LSTM as the best prediction for instantaneous traffic prediction
- Conv-LSTM as the most suitable model to predict the end of saturation.
- **FTS analysis helps us optimise transfers and delivers better performance for users! :-)**

Thank you for your attention!

Contact:

Joanna Waczyńska
joanna.waczynska@gmail.com
joanna.waczynska@cern.ch
IT-CS-CE

Edoardo Martelli
edoardo.martelli@cern.ch
IT-CS-CE

Sofia Vallecorsa
Sofia.Vallecorsa@cern.ch
IT-DI-OPL

Edward Karavakis
edward.karavakis@cern.ch
IT-ST-PDS

Tony Cass
tony.cass@cern.ch
IT-CS



home.cern

supplement

MSE – Mean Square Error

We consider the one-step MSE and calculate an average MSE for the Γ -steps forecasting during \mathcal{N} period when N samples were observed:

$$\text{MSE} = \text{MSE}_{\mathcal{N}}(\Delta, \Gamma) = \sum_{i=0}^{\Gamma} \text{MSE}_i, \quad (1)$$

where $\text{MSE}_i = \frac{1}{N-\Gamma-\Delta} \sum_{t=\Delta}^{N-\Gamma} (y_{t,i} - \hat{y}_{t,i})^2$, for $i \in \{0, \dots, \Gamma\}$.

$\text{MSE}_{\Psi,0}(\Delta, \Gamma)$ means MSE calculating during Ψ period for hyperparameters Δ and Γ . 0 index means result is calculate only for $\hat{y}_{\tau,0}$, where $\tau \in \{\Delta, \dots, N - \Gamma\}$.

How have we chosen hyperparameters?

- Time window Γ
- Time window Δ

Forecasting

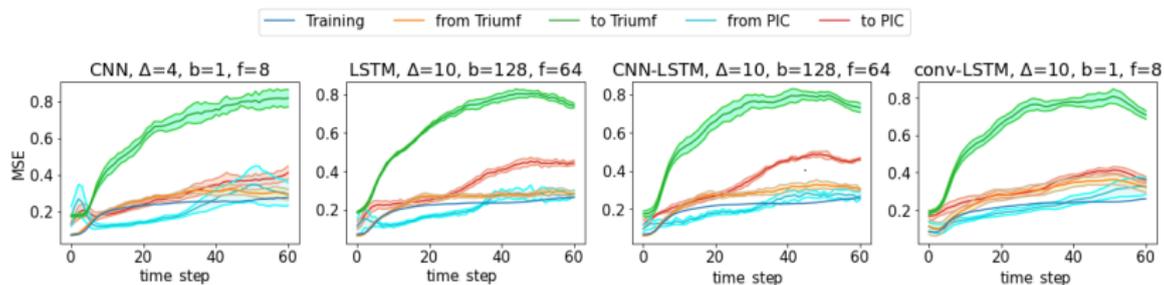


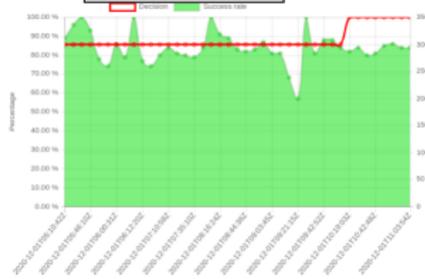
Figure: Average MSE and its variance with respect to the forecasting steps (here: $\Gamma = 60$).

FTS details - how transfer report format look like

Details for davs://eosatlas.cern.ch -- davs://webdav-at1.pic.es

DESTINATION

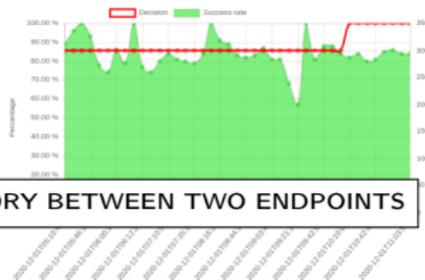
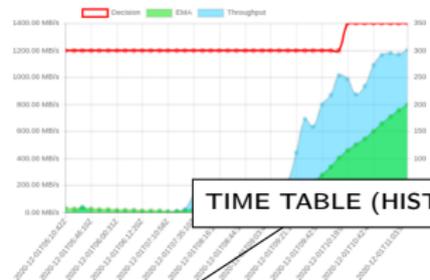
SOURCE



First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (Last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	806.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	104	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	606.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	88.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.40 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed





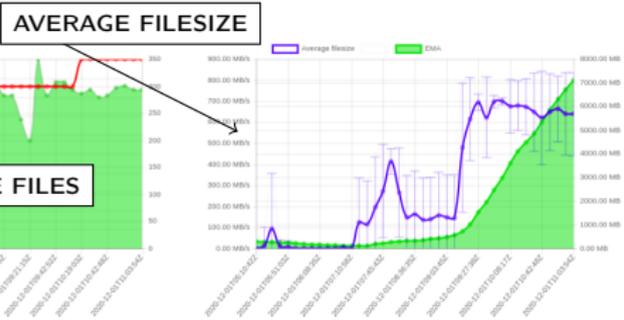
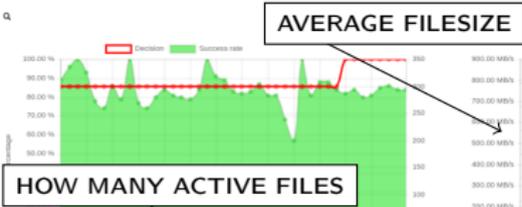
TIME TABLE (HISTORY BETWEEN TWO ENDPOINTS)



First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (Last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	806.38 MiB/s	0	Range fixed
2020-12-01T11:03:55Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T11:04:59Z	350	104	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T11:04:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T11:04:48Z	350	186	2367	81.00%	1.07 GiB/s	606.82 MiB/s	0	Range fixed
2020-12-01T11:05:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T11:07:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T11:09:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T11:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.40 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed





First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (Last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	806.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	194	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	606.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	609.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.40 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed





HOW MANY SUBMITTED FILES

Success rate (Last 1min)

First Previous 1 Next Last

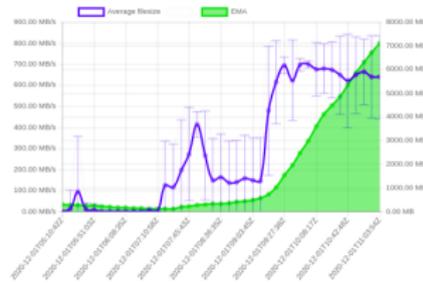
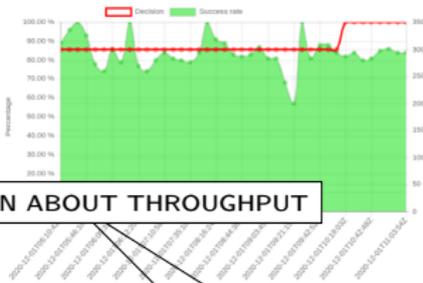
Timestamp	Decision	Running	Queue	Success rate (Last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	806.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	104	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	606.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	609.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.40 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed





First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (Last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T10:30:54Z	350	183	2137	84.00%	1.18 GiB/s	806.38 MiB/s	0	Range fixed
2020-12-01T10:31:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:34:59Z	350	104	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	606.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	88.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	583.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	484.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	878.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	448.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

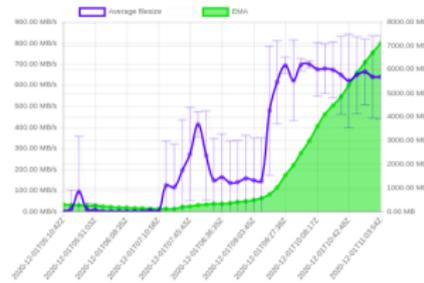


INFORMATION ABOUT THROUGHPUT

First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (Last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	806.38 MiB/s	0	Range fixed
2020-12-01T10:59:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:54:59Z	350	194	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	606.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	80.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	404.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.40 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed



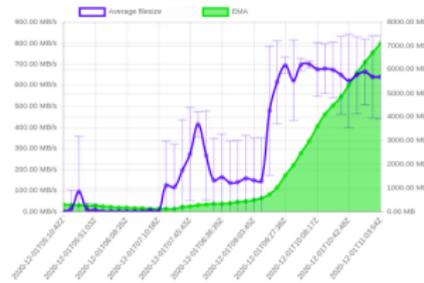


First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (Last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T10:35:54Z	350	183	2137	84.00%	1.18 GiB/s	806.38 MiB/s	0	Range fixed
2020-12-01T10:35:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:34:59Z	350	104	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	606.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	88.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	484.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	878.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	192	2884	88.00%	799.26 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	199	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	448.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

NEW BULK



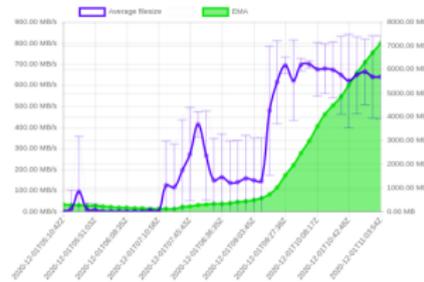


First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (Last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T10:35:54Z	350	183	2137	84.00%	1.18 GiB/s	806.38 MiB/s	0	Range fixed
2020-12-01T10:35:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T10:34:59Z	350	104	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T10:49:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T10:42:48Z	350	186	2367	81.00%	1.07 GiB/s	606.82 MiB/s	0	Range fixed
2020-12-01T10:35:21Z	350	183	2443	88.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T10:27:34Z	350	127	2556	84.00%	871.47 MiB/s	503.33 MiB/s	0	Range fixed
2020-12-01T10:19:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T10:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	484.31 MiB/s	0	Range fixed
2020-12-01T09:55:39Z	300	118	2798	88.00%	878.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T09:42:52Z	300	192	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T09:35:03Z	300	199	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T09:27:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T09:21:15Z	300	194	2929	57.00%	448.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T09:15:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T09:09:51Z	300	17	26	81.00%	133.48 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T09:03:45Z	300	26	24	81.00%	110.26 MiB/s	55.61 MiB/s	0	Range fixed

NEW BULK





First Previous 1 Next Last

Timestamp	Decision	Running	Queue	Success rate (Last 1min)	Throughput	EMA	Diff	Explanation
2020-12-01T11:03:54Z	350	183	2137	84.00%	1.18 GiB/s	806.38 MiB/s	0	Range fixed
2020-12-01T11:05:35Z	350	183	2174	84.00%	1.14 GiB/s	755.59 MiB/s	0	Range fixed
2020-12-01T11:04:59Z	350	194	2250	86.00%	1.15 GiB/s	709.42 MiB/s	0	Range fixed
2020-12-01T11:04:15Z	350	185	2312	85.00%	1.14 GiB/s	657.39 MiB/s	0	Range fixed
2020-12-01T11:02:48Z	350	186	2367	81.00%	1.07 GiB/s	606.82 MiB/s	0	Range fixed
2020-12-01T11:05:21Z	350	183	2443	88.00%	933.49 MiB/s	546.35 MiB/s	0	Range fixed
2020-12-01T11:07:34Z	350	127	2556	84.00%	871.47 MiB/s	583.33 MiB/s	0	Range fixed
2020-12-01T11:09:03Z	350	125	2611	82.00%	985.44 MiB/s	462.42 MiB/s	0	Range fixed
2020-12-01T11:08:17Z	300	133	2694	84.00%	1014.78 MiB/s	484.31 MiB/s	0	Range fixed
2020-12-01T11:05:39Z	300	118	2798	88.00%	870.12 MiB/s	336.48 MiB/s	0	Range fixed
2020-12-01T11:04:52Z	300	133	2884	88.00%	799.20 MiB/s	277.19 MiB/s	0	Range fixed
2020-12-01T11:05:03Z	300	132	2909	81.00%	631.83 MiB/s	219.19 MiB/s	0	Range fixed
2020-12-01T11:07:38Z	300	192	2947	100.00%	689.37 MiB/s	173.34 MiB/s	0	Range fixed
2020-12-01T11:09:21:15Z	300	194	2929	57.00%	440.40 MiB/s	116.00 MiB/s	0	Range fixed
2020-12-01T11:05:34Z	300	195	2896	68.00%	228.99 MiB/s	79.96 MiB/s	0	Range fixed
2020-12-01T11:09:09:51Z	300	17	26	81.00%	133.40 MiB/s	63.40 MiB/s	0	Range fixed
2020-12-01T11:03:45Z	300	26	24	81.00%	138.36 MiB/s	55.61 MiB/s	0	Range fixed

LINK WASN'T EMPTY



home.cern