Workloads prediction methods for proactive resource scaling in Kubernetes

Vitalii Omelchenko

Oleksandr Rolik

National Technical University of Ukraine "Igor Sikorsky Kyiv Polytechnic Institute", Kyiv, Ukraine

Introduction

- **Maximize quality of service and minimize amount of allocated computational resources.**
- Proactive scaling is more effective than reactive in typical workloads.
- The core part of a proactive solution is the forecasting technique.
- Cluster can include hundreds of different components with unique workload patterns.

Requirements to prediction methods

- versatility - ability to work with different time series without significant manual tuning or data preprocessing
- accuracy
- ability to work with complex seasonalities
- finite time/resources fitting

Prophet

- Prophet is a time series forecasting library developed at Facebook. The main goal of the development was to create a simple, transparent, and understandable model-generation algorithm that would allow for quick and reliable predictions.
- This algorithm is based on an additive regression model with several components.

$$y(t) = g(t) + s(t) + h(t) + e(t)$$

where g(t) - trend component, s(t) - seasonal component, h(t) - anomaly component and e(t) - an error function. In addition to the additive regression model, Prophet also uses the Fourier transform.
- Among the advantages of this model are the ability to work with any time series, the ability to work efficiently with large data sets and missing data, and flexibility in customization.

TBATS

- Acronym for the main components of the model: trigonometric seasonality, Box-Cox transformation, ARIMA, trend, and seasonality components.
- TBATS is designed to forecast complex time series with several seasonalities of different lengths. In TBATS, the original time series is subjected to the Box-Cox transformation, the primary purpose of which is to make the variance of the data stable, which is a crucial assumption for many statistical models, especially for linear regression and time series models.
- After that, the transformed time series is modeled as a linear combination of an exponentially smoothed trend, seasonal components, and ARMA components. Seasonal components are modeled by trigonometric functions using a Fourier series. TBATS is capable of adjusting some parameters independently using AIC .
- This method was chosen for this work because of its versatility and ability to adapt to time series of varying complexity, which means there is no need to adapt the model to the load of each existing component in the system.

NeuralProphet

- NeuralProphet is a time series prediction library developed on top of the PyTorch library. This library develops Facebook's Prophet model but uses neural networks. The main difference of this library is the ability to use the power of deep learning to predict time series with different trends and seasonality.
- The basis of this library is an autoregressive neural network, which combines classical autoregressive time series prediction methods with modern approaches based on artificial intelligence.
- This library consists of several components: trend, seasonality, regression of future variables, autoregression of historical variables, and regression of lagged variables. The following equation can describe their dependence:

Dataset

- workload pattern was generated with locust, data stored and obtained with prometheus
- value-time pairs
- complex seasonality - daily and weekly
- only cpu
- no preprocessing
- frequency - 1 hour
- anomalies are added by artificially distorting the data

Experiment 1. Three-week-long data, two seasonalities, no data distortion

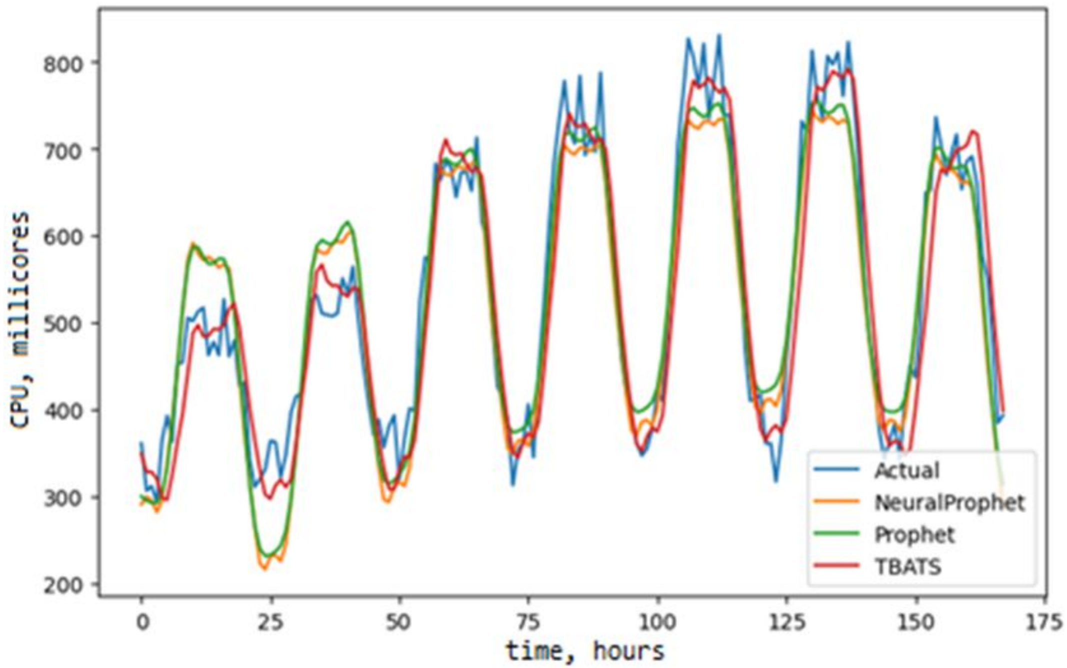

Figure X: Prediction of daily and weekly fluctuations without data distortion based on a three-week dataset

Experiment 1. Three-week-long data, two seasonalities, no data distortion

Model	RMSE	MAPE
TBATS	32.464	0.081
Prophet	46.16	0.0856
NeuralProphet	53.705	0.0891

Table X. Prediction of daily and weekly fluctuations without data distortion based on a three-week dataset

Experiment 2. Two-week-long data, two seasonalities, no data distortion

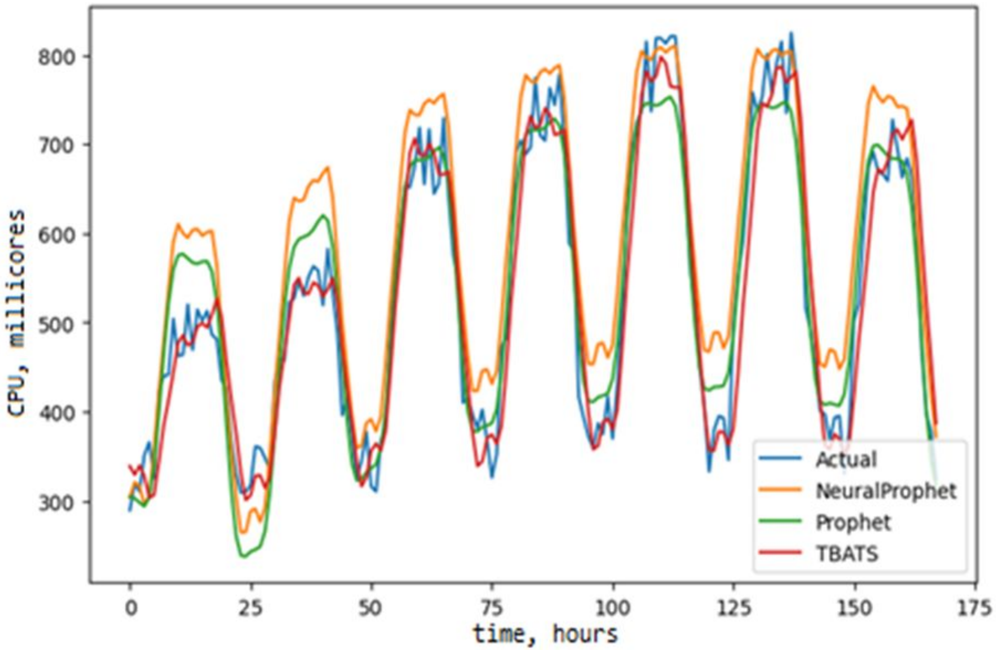

Figure X: [illegible] eek dataset

Experiment 2. Two-week-long data, two seasonalities, no data distortion

Model	RMSE	MAPE
TBATS	32.943	0.084
Prophet	51.236	0.089
NeuralProphet	71.378	0.130

Table X: Prediction of daily and weekly fluctuations without data distortion based on a two-week dataset

Experiment 3. One-week-long, two seasonalities, no data distortion

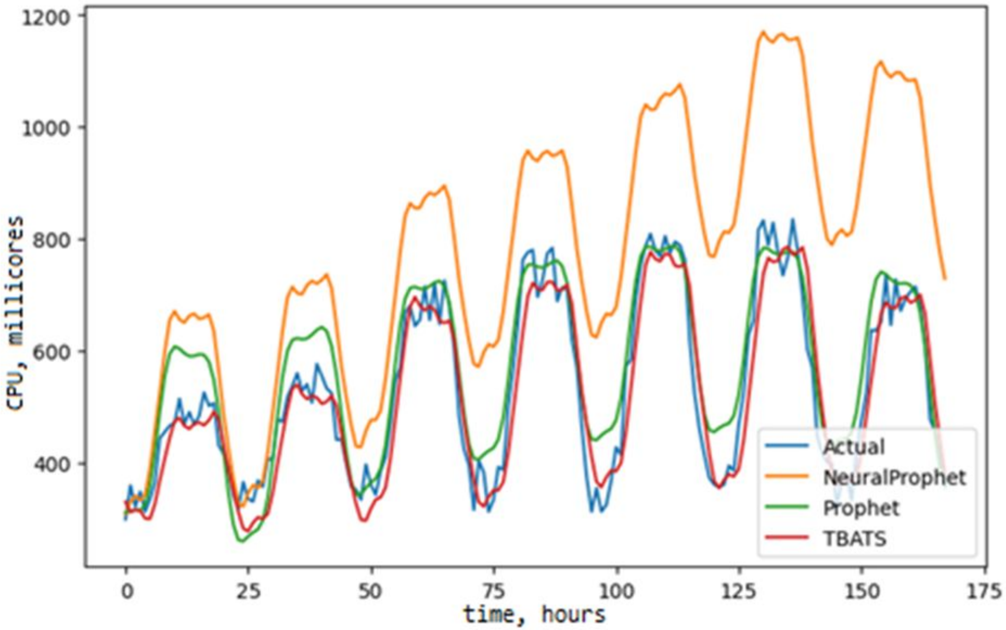

Figure X: Prediction of daily and weekly fluctuations without data distortion based on a one-week dataset

Experiment 3. One-week-long, two seasonalities, no data distortion

Model	RMSE	MAPE
TBATS	35.006	0.085
Prophet	64.427	0.117
NeuralProphet	281.056	0.493

Table X: Prediction of daily and weekly fluctuations without data distortion based on a one-week dataset

Experiment 4. Three-week-long, two seasonalities, 10% anomalies

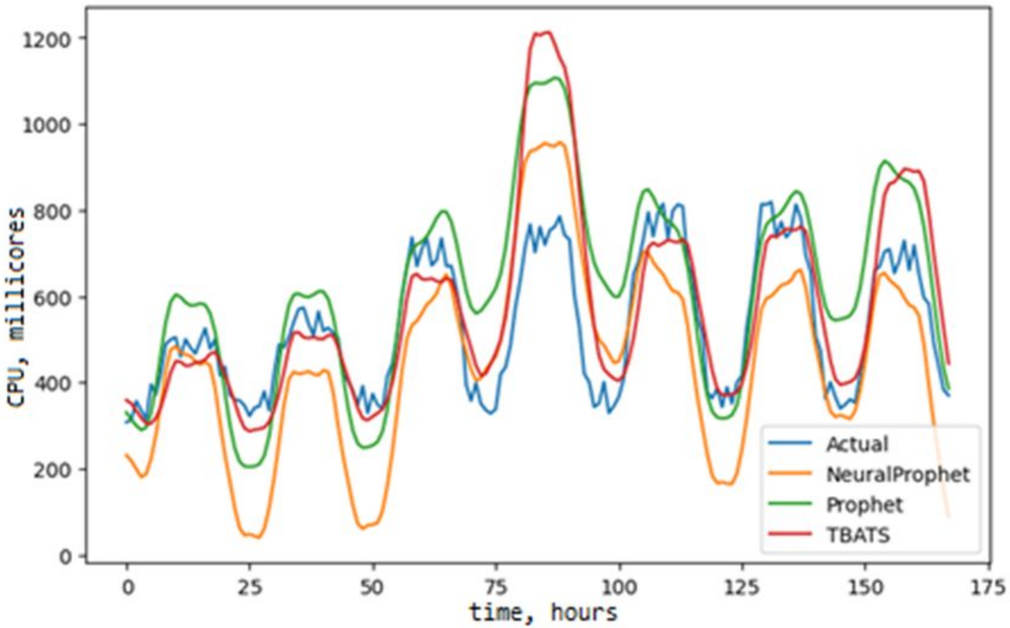

Figure X: Prediction of daily and weekly fluctuations with three input periods and 10% of days with anomalies

Experiment 4. Three-week-long, two seasonalities, 10% anomalies

Model	RMSE	MAPE
TBATS	135.949	0.174
Prophet	163.716	0.263
NeuralProphet	161.726	0.297

Table X: Prediction of daily and weekly fluctuations with three input periods and 10% of days with anomalies

Experiment 5. Three-week-long, two seasonalities, 20% anomalies

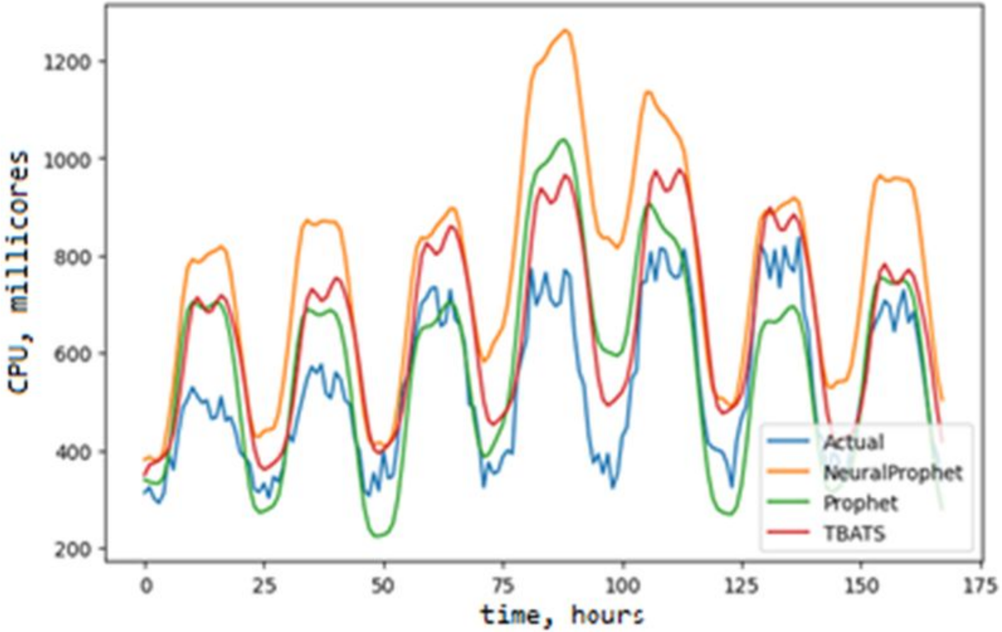

Figure X: Prediction of daily and weekly fluctuations with three input periods and 20% of days with anomalies

Experiment 5. Three-week-long, two seasonalities, 20% anomalies

Model	RMSE	MAPE
TBATS	134.313	0.243
Prophet	134.848	0.209
NeuralProphet	274.135	0.476

Table X: Prediction of daily and weekly fluctuations with three input periods and 20% of days with anomalies

Conclusions

- The results of the experiments show that, in general, all three selected models are able to predict complex load patterns with several seasonalities without preliminary data processing and with anomalies quite accurately.
- TBATS is more accurate than Prophet and NeuralProphet in all the experiments, but the difference in accuracy is no more than 16%.
- TBATS and Prophet predicted the load quite accurately with only one week of data, and the accuracy degradation is only 10%. NeuralProphet requires additional parameter settings or data pre-processing in some of the experiments.
- Anomalies significantly affect the accuracy, and pre-processing of historical data is necessary.