

Easy use of Distributed TensorFlow Training on supercomputing facilities.

Deep Learning is a powerful tool for science, industry and other sectors that benefits from large datasets and computing capacity during models' design and training phases. TensorFlow (TF) Google's Machine Learning API is one of the tools most widely used for developing and training such deep learning models. There is a wide range of possibilities to configure a deep learning model however find the optimal model architecture can be a highly demanding computing task. Moreover, when involved datasets are very large the computing requirements increase and training processes can take a lot of time and hinder the design cycle. One of the most powerful capabilities of TF is its distributed computing capabilities, allowing portions of the automatic generated graph to be calculated on different computing nodes, and speeding up the training process. Deployment of distributed TF is not a straightforward task and it presents several issues, mainly related with its use under the control of local resources management systems and the usage of the right resources. In order to allow CESGA users to adapt their own TF codes to take advantage of TF and Finis Terrae II distributed computing capabilities, a complete python Toolkit has been developed. This Toolkit deals with several tasks that are not relevant in the models design, but necessary for exploiting the distributed capabilities, hiding the underlying complexity to final users. Additionally, an example of a successful industrial case, based on the Fortissimo 2 project experiment "Cyber-Physical Laser Metal Deposition (CyPLAM)", that uses this Toolkit, is presented. Thanks to the TF distributed capabilities, the computing capability of Finis Terrae and the use of the developed Toolkit, the time needed for training the largest model of this industrial case has been decreased from 8 hours (non- distributed TF) to less than 20 minutes.




Figure 1: Training time (left axis) and Speed Up (right axis) vs number of tasks for Distributed TensorFlow training for a CyPLAM model.

Primary authors: Dr FERRO COSTAS, Gonzalo (CESGA); Mrs COTELO QUEIJO, Carmen (CESGA); Dr GÓMEZ TATO, Andrés (CESGA)

Track Classification: R&D for computing services, networking, and data-driven science at the Iberian level.