Advanced Computing Training Program
Final Report

Participant Name:	**Helena Pais**
Affiliation and Position:	**Researcher, Centro de Física da Universidade de Coimbra**
Period:	**January 14 – March 28 2019**

The general idea of my visit to TACC was to a learn, in what time permits, how to manage an HPC cluster. During my time there, I joined the HPC group. There were two points approached: from an user point of view, and from an administrator point of view.

From the user point of view, we optimized a code, that calculates the equation of state in the subsaturation density region, including heavy nuclear clusters, in β–equilibrium cold stellar matter in the Thomas-Fermi approximation within a relativistic mean-field approach. We used profiling tools like gprof, and vectorization. We changed the compiler, and did other small changes throughout the code. These brought us an overall improvement of about 10%. This was done in collaboration with Lars Koesterke.

From an administrator point of view, we built a virtual workload manager (SLURM) cluster, using Virtual Box and Ansible, an open-source change management software tool, which we also learned. Many features were discussed, like the accounting database set up, using MySQL/MariaDB, prologs and epilogs scripts. Besides that, and in order to share directories among the nodes, we installed a NFS server/client, using Ansible. We also made sure that all our nodes were synchronized, and for that purpose, we installed a NTP server/client as well. At the end, we were able to run jobs on it. This was done in collaboration with Cyrus Proctor and Nick Thorne. We were also interested in understanding how the ticketing system works at TACC. We went through the ticketing software used at TACC, Request Tracker (RT), and we also discussed about other open-source solutions, like OSTicket. This was done in collaboration with Virginia Trueheart. Finally, we installed in our cluster TACC Stats, a tool for collection and analysis of the system data. This kind of tool can be useful for the system administrators and managers as it collects data, from CPU to memory to filesystem usage, system load and process statistics. By analyzing this data, future and more efficient improvements can be made in new clusters. This tool was installed using Ansible, and a short report was written where all the necessary steps were described. These involved, among others, building and installing RPMs packages, and installing and creating a PostgreSQL database. This was done in collaboration with Todd Evans.

Overall, my stay was very pleasant and productive, as I was able to learn a few new things, namely a profiling tool (gprof), a change management tool (Ansible), installing and managing SLURM clusters, installing a NFS and NTP server/client, installing and creating databases (MySQL/MariaDB, PostgreSQL), building and installing RPMs packages, and a tool for collection and analysis of jobs data, TACC Stats. Moreover, I had the opportunity to see how the group works and interacts, and I am thankful to the people there, and the HPC group in particular, for their kind hospitality and helpfulness. I hope to have the opportunity to put my new skills in practice in a near future, either in Portugal or elsewhere. I think this program can be very useful to researchers that are interested in learning or developing new computational tools for their projects. However, I think there's still place for improvement, in what concerns the overall setup of the program: there should be a more solid dialog between both parts (Portuguese and American) for a program to be put in order. In particular, all the possible offers should be

discussed and advertised accordingly, so that both parts, the Portuguese researchers that want to participate, and the ones that are tutoring know what they should expect.