

Pulsar Searching with Many-Cores

Alessio Sclocco, Rob V. van Nieuwpoort

Published in: Facing the Multicore-Challenge III, September 19-21, 2012, Stuttgart, Germany. Lecture Notes in Computer Science Volume 7686, 2013, pp 123-124.

Abstract

Pulsars are rapidly rotating neutron stars whose signal is received on Earth periodically. They are relatively newly discovered astronomical objects (the first was discovered only in 1967) and elusive ones: so far only two thousand of them are known. Their properties, especially their big mass and precise period, can be used to probe space and gravitation. This makes them important not only for astronomers, but for physicists and other scientists as well: the discovery of the first binary pulsar by Hulse and Taylor in 1973 [1] has been so important for proving general relativity and other aspects of gravitation that they won the 1993 Nobel prize for physics. Thus, being able to discover new pulsars is an important goal of current radio astronomy.

The process of finding a new pulsar is, however, difficult and time consuming: it involves a brute-force search over hundreds of thousands of parameter combinations. Moreover, the amount of data that needs to be searched is huge: the input of a typical observation using a radio telescope like LOFAR [2] is in the order of hundreds of terabytes. During the process petabytes of intermediate results are produced and analyzed. Thus, searching for pulsars clearly is a big data problem.

The challenges are not only limited to the amount of data that needs to be processed: the signal received from a pulsar is usually quite faint and can be completely covered by radio frequency interference and, even when there is no artificial noise, a long integration time may be necessary to properly detect its profile. There is, furthermore, the inter-stellar medium to take into account: typical effects of the interaction between the emission and the medium are dispersion, scintillation and scattering. And this is not the only interaction: if the pulsar is orbiting with a companion of some kind, the interaction between them modifies the signal and makes it even more difficult to detect from Earth.

We propose to reduce the time necessary to search for new pulsars using many-core accelerators, e.g. modern GPUs. We design and implement the whole searching pipeline, using OpenCL to build the three most important computational kernels of this application: dedispersion, folding and the signal to noise computation. An overview of our pulsar searching pipeline is presented in Figure 1. Using our prototypical tool we are able to find pulsars in test data sets generated with Duncan Lorimer's SIGPROC. To overcome one of the main problems of using accelerators, i.e. the expensive memory transfers between device and host memory, we keep all the intermediate results in device memory. Moreover we use auto-tuning, a technique that we know to be effective with many-core architectures [3], to identify which parameters are the best suited for each specific device used to run our prototype.