

Complex Systems Made Easy

4,2 million investment: optimised programming to bring high performance computing and embedded systems together

Smartphones, MP3 players, cameras etc. all become smaller, but the technical expectations towards them constantly increase. At the other end of the spectrum, modern supercomputers have nearly reached exascale (one quintillion (10^{18}) processing units in a single system), but the performance hardly increases anymore. Up until 2006, faster processors were able to deliver better performance and thus meet the increasing expectations, but the effective speed of processor has not really increased since. Instead, manufacturers hope to meet the growing requirements by an increasing amount of processing units and specialised processors. “The potential of common processor architectures is mostly exhausted”, warns Professor Stefan Wesner, coordinator of a newly funded project to address this problem:

Programming Large Scale Heterogeneous Infrastructures

The international research project on “Programming Large Scale Heterogeneous Infrastructure” (acronym “POLCA”) funded by the European Commission under Framework Programme 7 with a budget of 4.2 million Euros aims at improving usability and energy efficiency of modern complex systems considerably. Classical approaches to the problem are no longer appropriate for modern hardware.

Modern embedded systems fight with the increasing performance demands that can only be met through parallelising the respective application. “Our goal consists in bringing high performance computing capabilities to embedded systems, such as they are used in smartphones, but in particular in specialised micro controllers, such as used for oil drilling, earthquake detection, radio analysis etc.” says Manuel Carro, Deputy Director of the IMDEA Software Institute.

This is a completely new field to the embedded systems domain, whereas it is one of the classical key topics addressed in High Performance Computing. On the other hand, modern supercomputers struggle with properly exploiting the growing number of specialised processors, which is typical for the domain of embedded systems. “Modern supercomputers are no longer homogeneous but integrate a lot of different specialised systems to deal best with dedicated problems,” says Colin Glass, head of the department for Scalable Computing & Optimization Methods at the High Performance Computing Centre, Stuttgart, “which makes efficient programming increasingly difficult – only a handful experts worldwide can do that now.”

Exploiting the specialisation of a system can also reduce the energy consumption of a program considerably. That is particularly relevant considering the growth of energy usage through the increasing number of processing cores and the resulting CO₂ production. It is therefore only natural to foster exchange experts between these two domains. The POLCA initiative brings experts from two companies and six research institutes all over Europe together to specifically address this problem. The project thereby unites key players from both the embedded system and the high performance computing domain.

POLCA makes complex parallel systems programmable

POLCA will realise a programming model and development tool-chain that will enable the programmer to fully exploit the specifics of the destination environment and to integrate multiple instances into a single system. In other words, the POLCA model will allow development of distributed applications over a wide range of processor types. This will prepare for the next generation of applications that have to deal with a large scale and scope of processors to compensate for the lack of performance.

“The processor market is dominated by the United States mostly because the entry point for specialised processors is hindered by the lack of appropriate programming models. But Europe produces some of the best specialised processors,” says Stefan Wesner. The POLCA model will enable software developers to exploit the capabilities of modern computer systems to their full extent without being constrained to a specific supplier. This strengthens Europe’s position in the highly competitive processor market considerably.

“Enabling programmers to write computer programs which run at high performance across multiple very different computer systems is a key challenge for the computer industry. Maxeler dataflow computers already provide major advantages but do require some extra programming effort.” States Oliver Pell, VP of High Performance Computing at Maxeler Technologies, “An efficient POLCA approach could minimize development time, thus reducing the barrier to entry for new customers to start utilizing such technology.”

Lotfi Guedria Department Deputy Manager at CETIC adds that “modern devices, such as FPGAs, integrate a variety of advanced programmable resources. This offers more flexibility to the programmer for a wider scope of application areas and allows for significant performance increase, as CETIC has shown in several cases. However, this requires additional hardware expertise from the developer.”

“Designing efficient hardware is in itself complex enough,” continues Gerard Rauwerda, CTO of Recore Systems, “yet that complexity is nothing compared to the difficulties programmers experience when they need to program those systems for optimal utilisation. Within POLCA we will develop the means to tackle those difficulties from both sides – with the programmer as the great beneficiary.”

The Consortium

The project consortium is constituted by the German Universities of Ulm and Stuttgart, as well as the Dutch University of Twente and the Spanish “Universidad Politécnica de Madrid”, along with the Belgium “Centre of Excellence in Information and Communication Technologies” (CETIC), the Spanish “IMDEA Software Institute”, the British computing platform manufacturer “MAXELER Technologies”, as well as the fabless semiconductor company “RECORE Systems” from the Netherlands.

For more information please check out <http://www.polca-project.eu> or contact Lutz Schubert [{lutz.schubert@uni-ulm.de}](mailto:lutz.schubert@uni-ulm.de)