Invasive Parallel Computing - An Introduction

J. Teich

Abstract:
Technology roadmaps foresee already today 1000 and more processors being integrated in a single MPSoC in the year 2020. For such systems, the control of many concurrent applications can obviously not be organized in a fully centralized way any more as it is done in today's multi-core processor systems. Also, feature variations are expected to become a severe problem threatening not only performance but also correctness of computations. One way shown how be able to cope with an expected increase of run-time uncertainties is to exploit flexibility of as well the code to be executed as the reconfigurability of the underlying hardware resources. The only major question is at what price this can and should be done, to what degree, and in the control of whom such adaptations shall take place.

In this introductory talk, we present a novel paradigm for an application-driven, decentralized as well as resource-aware organization of concurrent applications on future large scale MPSoCs called "Invasive Computing". The main idea of "Invasive Computing" relies on the vision that application-developers are typically conscious of the temporal computational demands of their programs and that these should be able to spread their load at run-time on processors, communication and memory resources themsolves in phases called invasion. Resource-awareness, on the other hand, means that such decisions whether to invade or retreat from resources should be done in reflection with the state of the underlying resources such as temperature-, reliability-, aging-, or fault-monitor information. As an example, spawning more and more threads to an already overloaded MPSoC CPU architecture might lead to less performance than choosing an alternative sequential or approximate computation.

In July 2011, the German Research Foundation (DFG) has estalished its transregional collaborative research center TR89 on "Invasive Computing" with Erlangen (FAU), Karlsruhe (KIT) und Munich (TUM) as the three participating research universities with the goal to investigate the paradigm of "Invasive Computing" intensively with respect to the development of new programming and language concepts, architectures of invasible resources, simulation and compiler support as well as application and demonstrator development.

Here, we give an overview of the basic principles of "Invasive Computing". In particular, we present a first language definition and implementation for a set of new and not yet existing parallel programming constructs on top of the language X10 as developed by IBM. Also, we will show how invasive programs may be efficiently simulated so to have a testbed for a) invasive application developers, b) resource-aware programming, and c) design space exploration of architectural tradeoffs such as numbers and types of processors, memory organization, etc. It will be finally outlined how and to what degree we may expect invasive computing to improve fault-resilience, scalability, efficiency and resource utilization by the analysis of invasive speedup and efficiency numbers.