Applications Session

Ad-hoc Architectures for modern DBMS: a HW/SW Co-Design Approach

L. Pomante[1], P. Di Felice[2]
[1]*University of L'Aquila – DEWS, ITALY, luigi.pomante@univaq.it*
[2]*University of L'Aquila – DIEI, ITALY, paolino.difelice@univaq.it*

Abstract

In the last years, database technology has become fundamental in several application domains (e.g. *Geographic Information Systems*, *Bioinformatics Information Systems*, etc.). Very often, such domains involve the heavy use of application-specific data-types (e.g. spatial geometries, protein structures, etc.), and related operators, very different from the ones that classical DBMS (*Data Base Management Systems*) have coped within the past. These custom data types present new challenges to DBMS designers because they upset traditional implementation guidelines that focused mainly on the optimization of the I/O cost.

The consequence of this "revolution" is the need of a specific support in order to efficiently manage such data. In fact, several attempts have been made recently to add some kind of "accelerators" (exploiting GPU and FPGA). Actually, the idea of a dedicated *DB machine* is very old and periodically proposed in the past. However, the technological limitations of the related epochs have always "pushed-back" the idea of a *DBMS computer architecture*. Nowadays, with the availability of cost-effective FPGA and application-specific processors the idea of providing ad-hoc support to DBMS is gaining (again) a lot of consideration.

However, existing works lack of generality because they target acceleration of specific operators according only to the experience of the designer. Moreover, acceleration cannot be the only issue: with the diffusion of portable devices other factors (i.e. energy-consumption, reliability, cost, etc.) should be taken into account also for different important scenarios (e.g. portable DBMSs with standard data types and operator, resource-limited smartcard, etc…).

Trying to fill such a gap, this work presents a co-design methodology (based on an existing system-level embedded systems co-design methodology) for the definition of ad-hoc architectures for modern DBMS (e.g. ORDBMS). The main goals of the proposed methodology are:

- to analyze the DBMS specification to identify operators that could benefit from ad-hoc executors
- to explore the solutions design space
- to define the architecture that optimize the relevant design metrics

Future work will be devoted to the development of an effective tool-chain, and to the design flow validation by means of real-world case studies.