



FP7 Coordination and Support Action to fund 50 technology transfer projects (TTP) in computing systems. This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement n° 609491.

SemBoost: accelerating semantic intelligence applications on massively parallel architectures

Paolo Burgio, Micaela Verucchi, Marko Bertogna
HiPeRT Lab, Università di Modena e Reggio Emilia, Italy

TTP Problem

Tools for Semantic Analysis process structured text


- E.g., web pages
- To extract semantic information
- Almost in real-time (few milliseconds)



Still cannot exploit massively parallel architectures



GP (General Purpose) GPUs



Next-generation Many-core accelerators

TTP Solution

SemBoost project


Application → Search in a network
 Using GP-GPUs
 Data structure → Network of concepts

CompactArr (CA) is an ad-hoc data structure explicitly designed for host+accelerator systems

1. Compact
2. Cache-friendly
3. Easy to move across different memory spaces


Effectively achieving performance through parallelization requires:

1. Adequate application partitioning
2. **Ad-hoc data structure**



Traditional/"Naive" representations are:

- Pointer-based
- Sparse in memory
- Hard to move across different memory spaces



Comes from **LightKer** project @UNIMORE



TTP Impact

Cache-friendly

- Speedup also on CPUs
- **2-3X against pointer-based implementation**

1

CompactArr

2

Reduced memory occupation


- <2MB against 26MB
- **~15X reduction**

3

Enables adoption of GPU accelerators

- **20% speedup vs. pure CPU**

Single search time	List-based version	CA
AVG time (ns)	642	257
MAX time	3139	1164
MIN time	112	59



TTP Facts

Contact: Marko Bertogna
 E-mail: marko.bertogna@unimore.it
 TETRACOM contribution: 30000 €
 Duration: 4/1/2015-2/1/2016



UNIMORE
UNIVERSITÀ DEGLI STUDI DI
MODENA E REGGIO EMILIA