Master Course in Parallel Computing
From Pen-and-Paper to Heterogeneous NUMA Cluster

Novica Nosović
University of Sarajevo
Faculty of Electrical Engineering
Dpt. of Computing and Informatics
e-mail: novica.nosovic@etf.unsa.ba

Workshop “Cooperation at Academic Informatics Education across Balkan Countries and Beyond: The Impact of Informatics to Society"
Jelsa (CRO), 2nd - 6th September 2019

Abstract

Serious course in parallel computer systems can hardly fit in 3 lecture hours per week, 14 weeks semester budget. Covering typical problems (what?) suitable for parallelization, solved with adequate algorithms (how?), coded with contemporary programming languages (tools?) and executed on accessible parallel machines (architectures?), requires more lecture hours and lab resources. Leading technical universities around the world do this in increased number of teaching hours and with extended access to variety of parallel hardware platforms. They also avoid "the textbook"-based teaching by offering extended list of conference papers in various parallel computers and programming technologies. Our approach is to maximize learning outcomes using (very) limited resources.

On Difficulty of Parallel Software

If (sequential programming = difficult)

parallel programming = difficult^2

Course Goals (working in groups)

- Write parallel program
- Achieve considerable speedup using “lessons learned”
- Use “low hanging fruit”
- Try to:
  - Be relevant, coherent and consistent
  - Find “piece of gold”
  - Write conference paper
  - Publish paper

Abstraction Layers

- Problem
- Algorithm
- Tools
- Architecture

Classroom Demo on Parallelization

- New concepts introduced: communication, synchronization, scheduling, load balancing, Amdhals law, message passing, task/data decomposition, map-reduce, hard/weak scaling, arithmetic intensity, high watermarking, sublinear/linear/superlinear speedup, multi/hyper-threading, shared/distributed memory, co-processors/accelerators, clustering,

The Class

Problem(s) in Measuring Speedup

64 byte cache line size

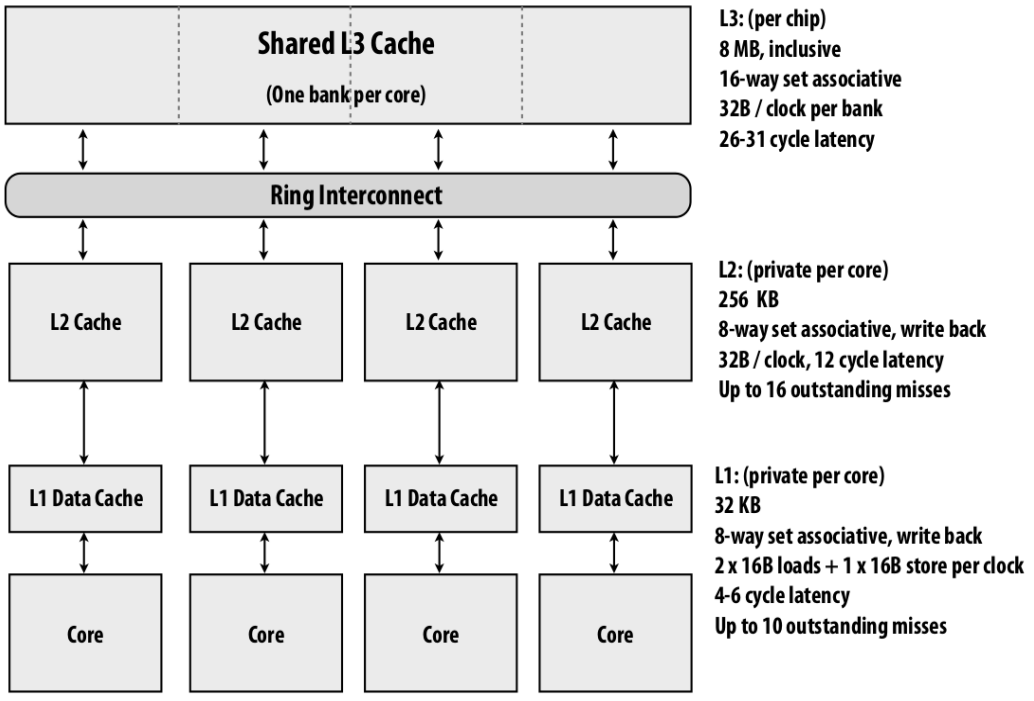

Parallella – credit card sized “supercomputer”

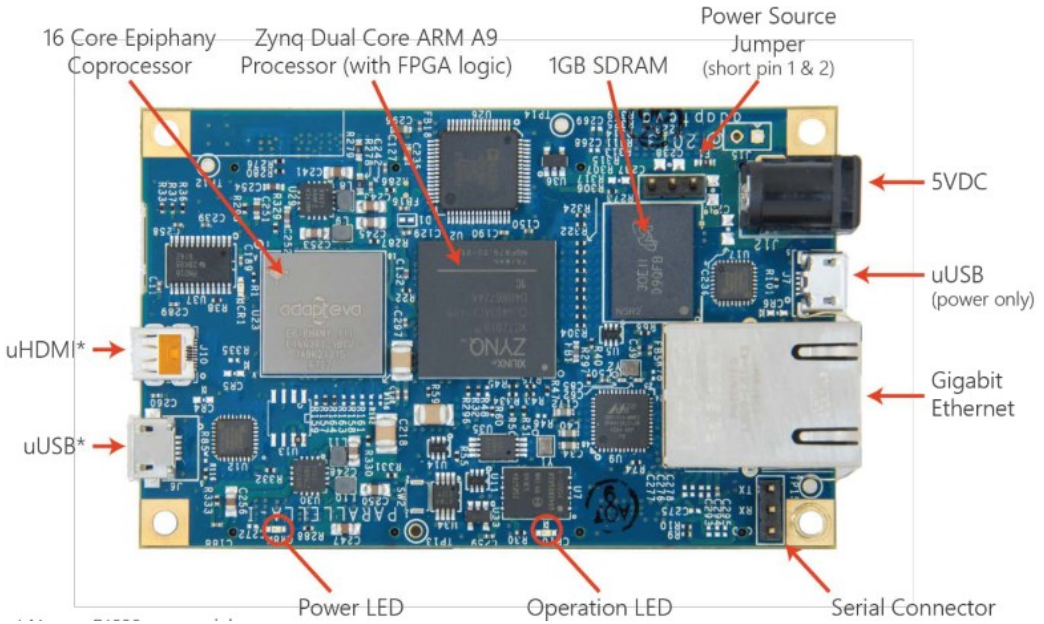

16 Core Epiphany Coprocessor

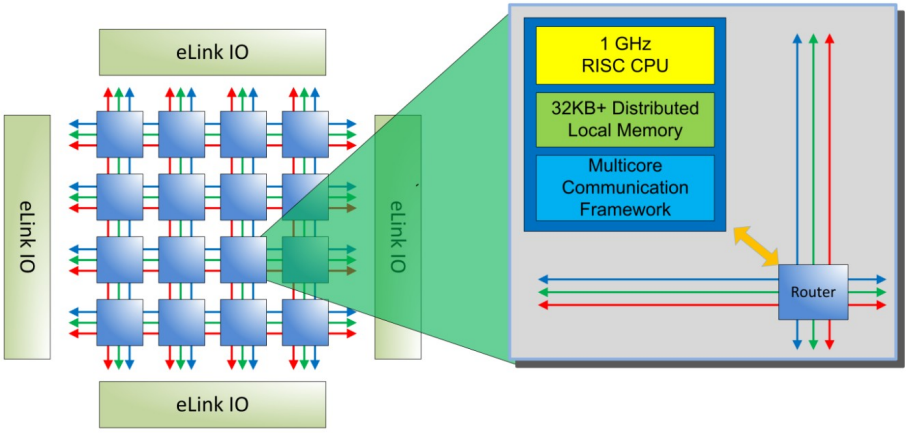

ControlFlow vs. DataFlow

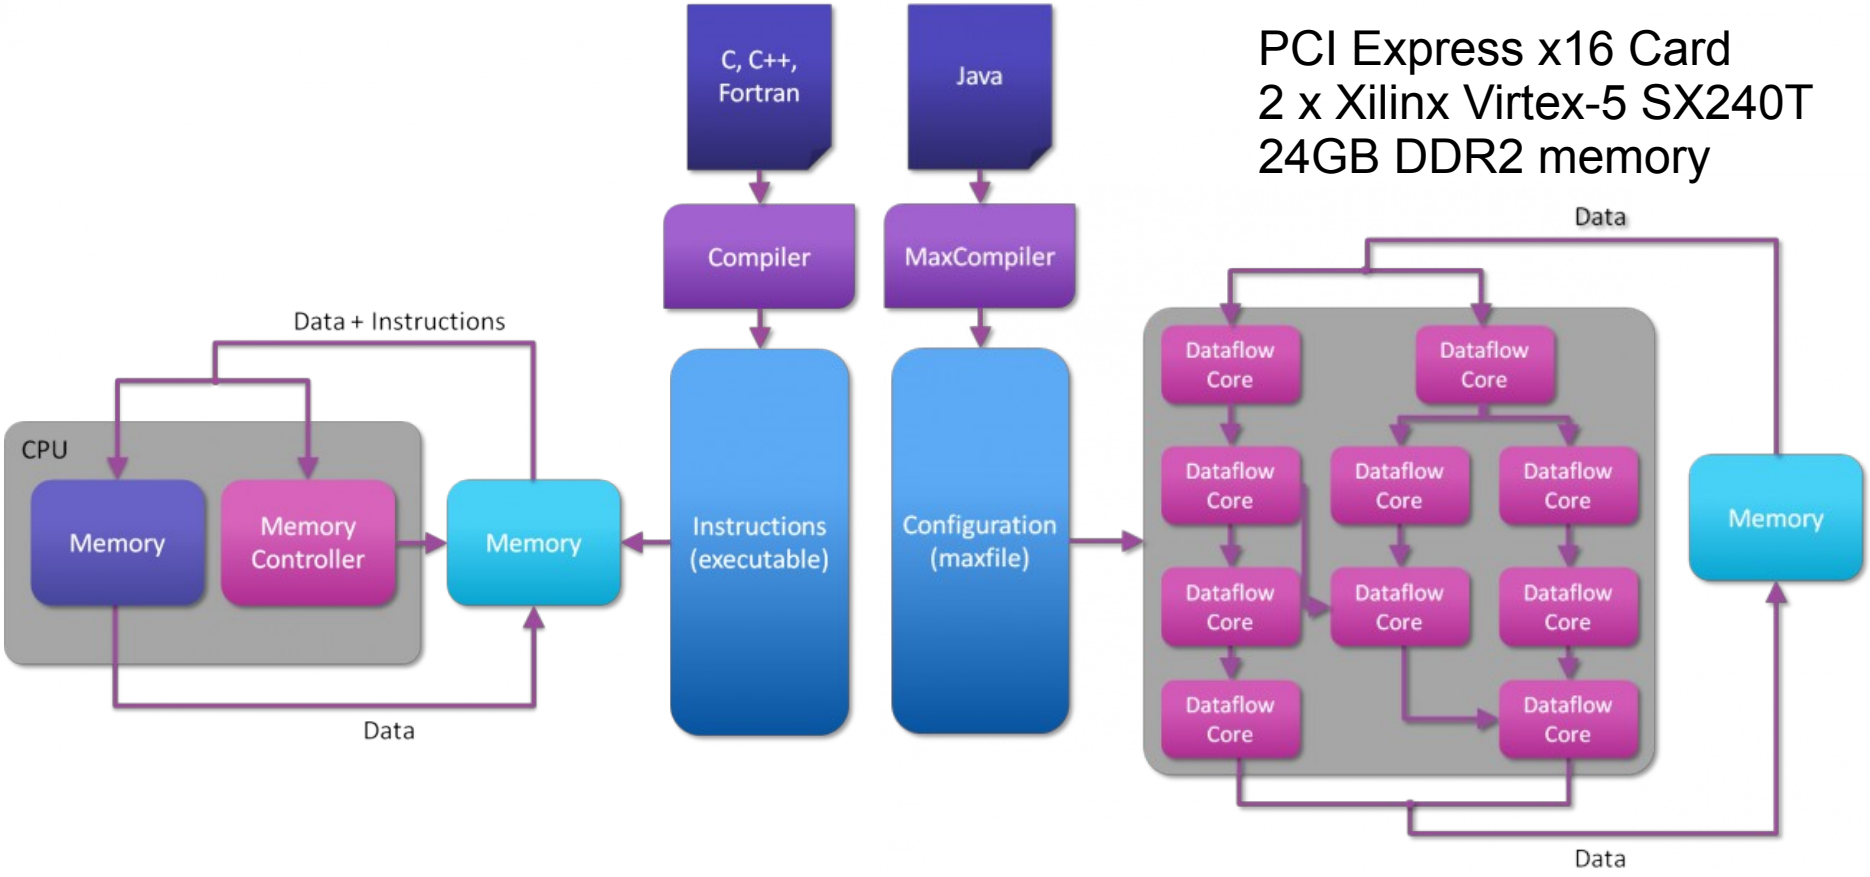

Conventional Heterogeneous Multicore System Architecture

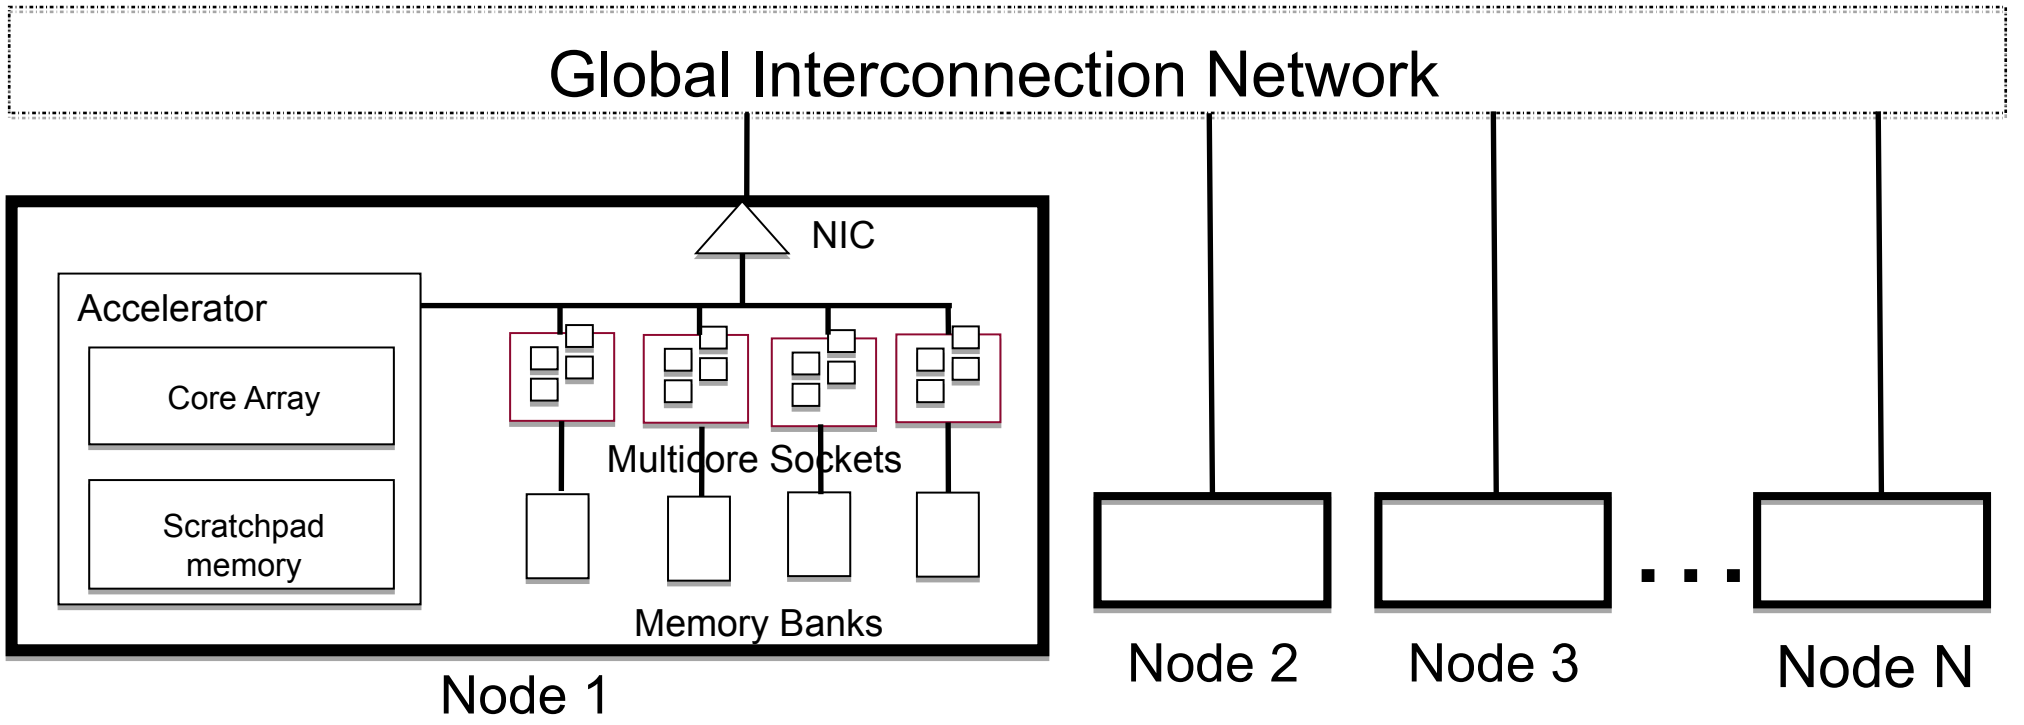

20 Dell PowerEdge SC1435`s with IB Switch

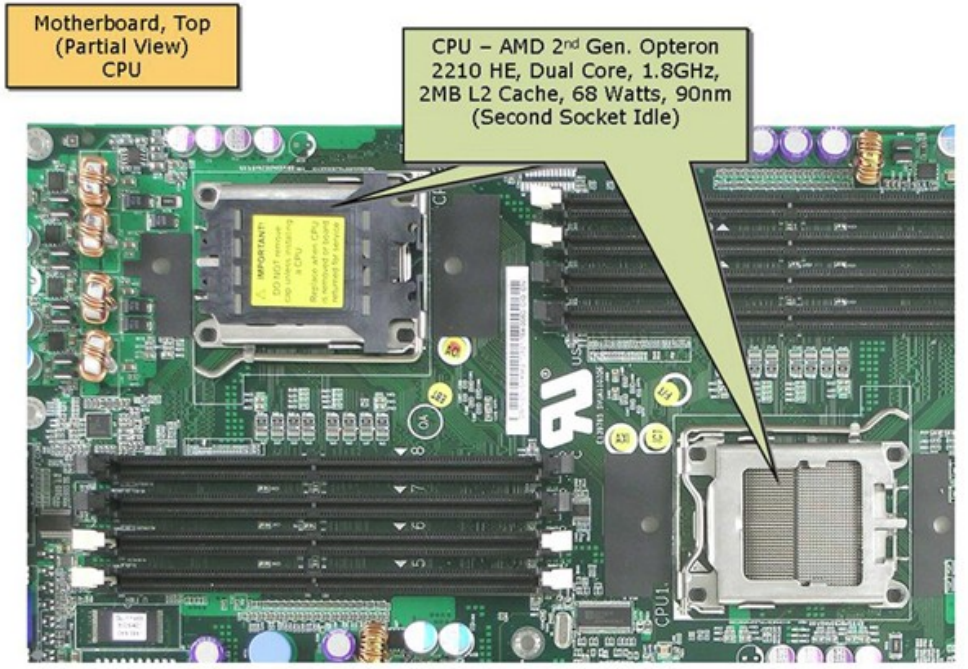

Source: IHS

NUMA Configuration with Broadcom Chipset

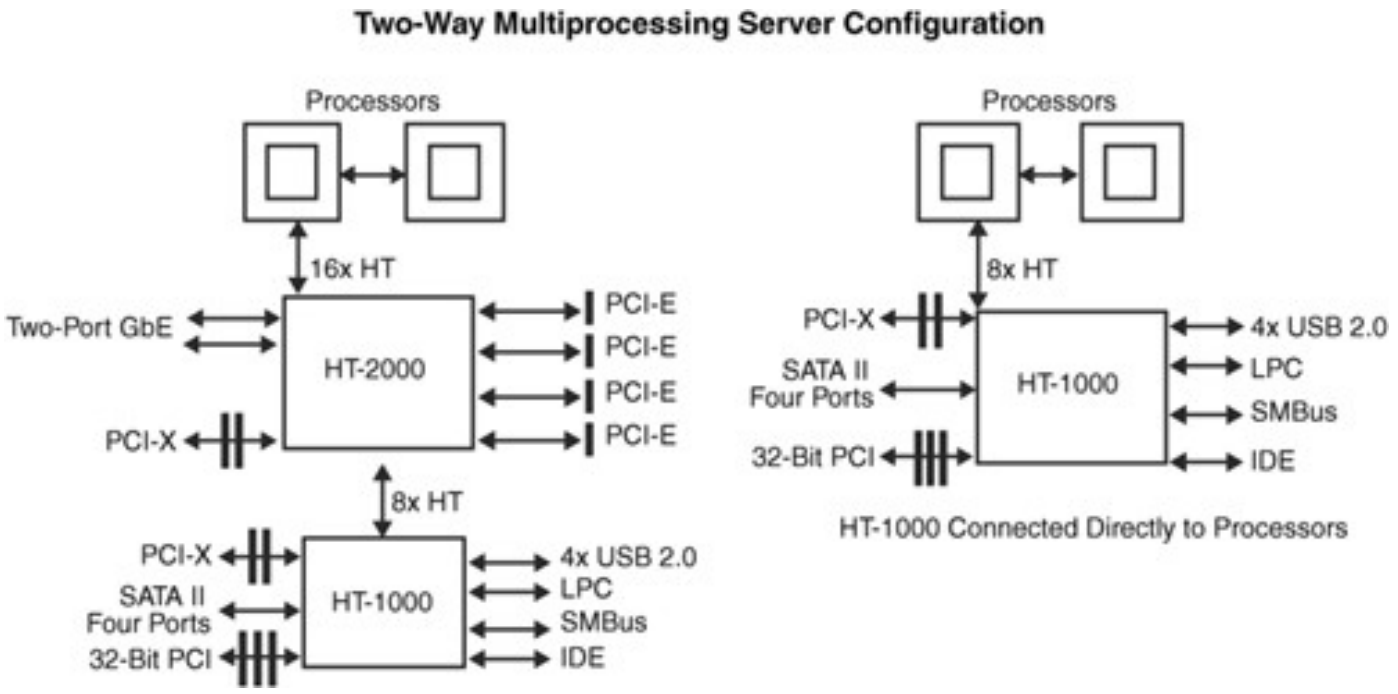

Load balancing is hard

My Expectations

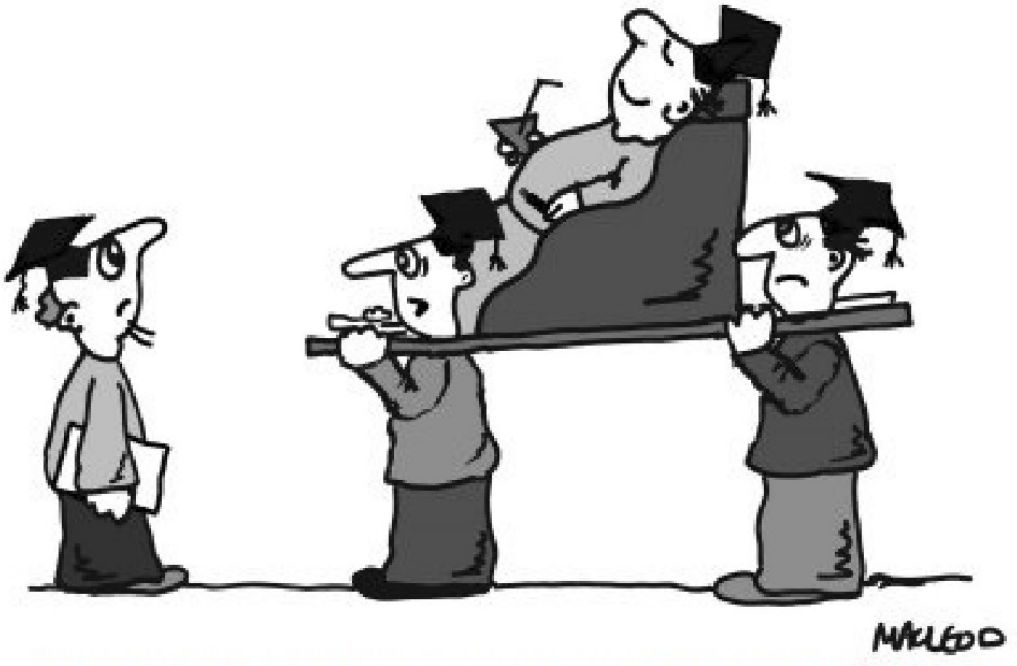

Results achieved

- Predictable results can not be considered “piece of gold”