



Deployment of Query Plans on Multicores

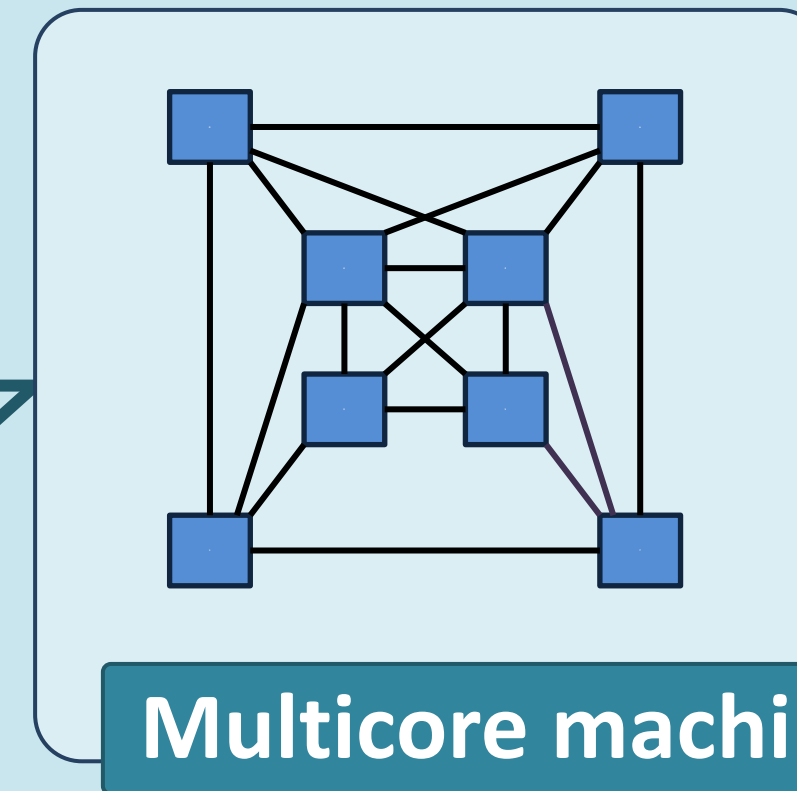
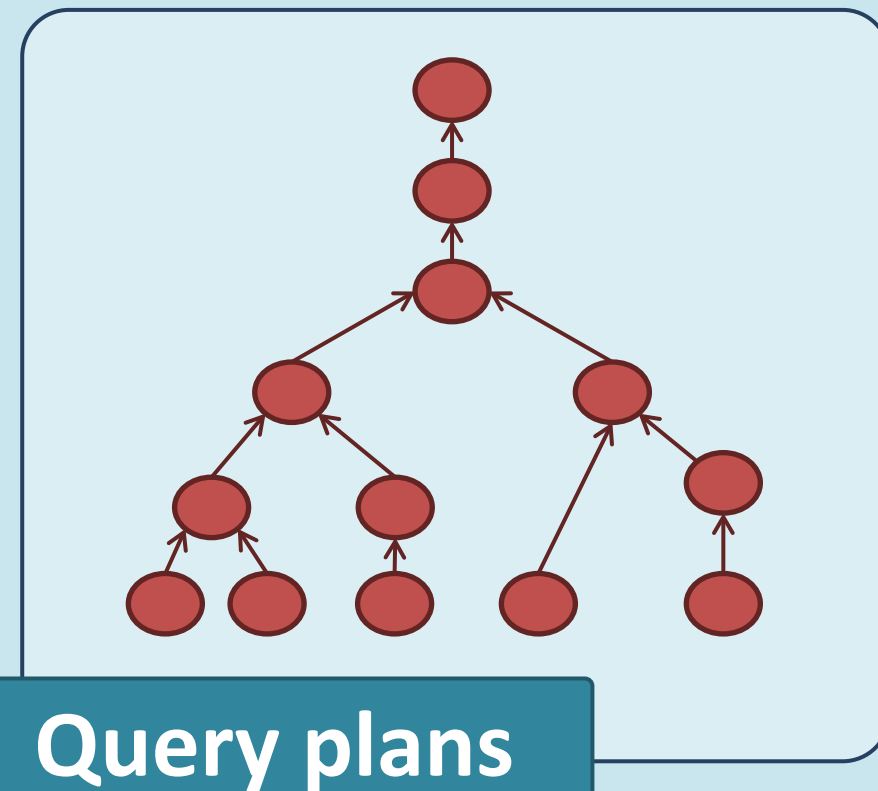


Jana Giceva, Gustavo Alonso, Timothy Roscoe, Tim Harris
Department of Computer Science, ETH Zurich Oracle Labs

1. Problem Definition

Complex query plans

- Large number of operators
- Complex data-dependencies



Manycore systems

- Complex architectures
- Asymmetric interconnect topology

User requirements



Performance



Efficient resource utilization

User requirements for

- Good and stable performance
- Machine under-utilization avoidance

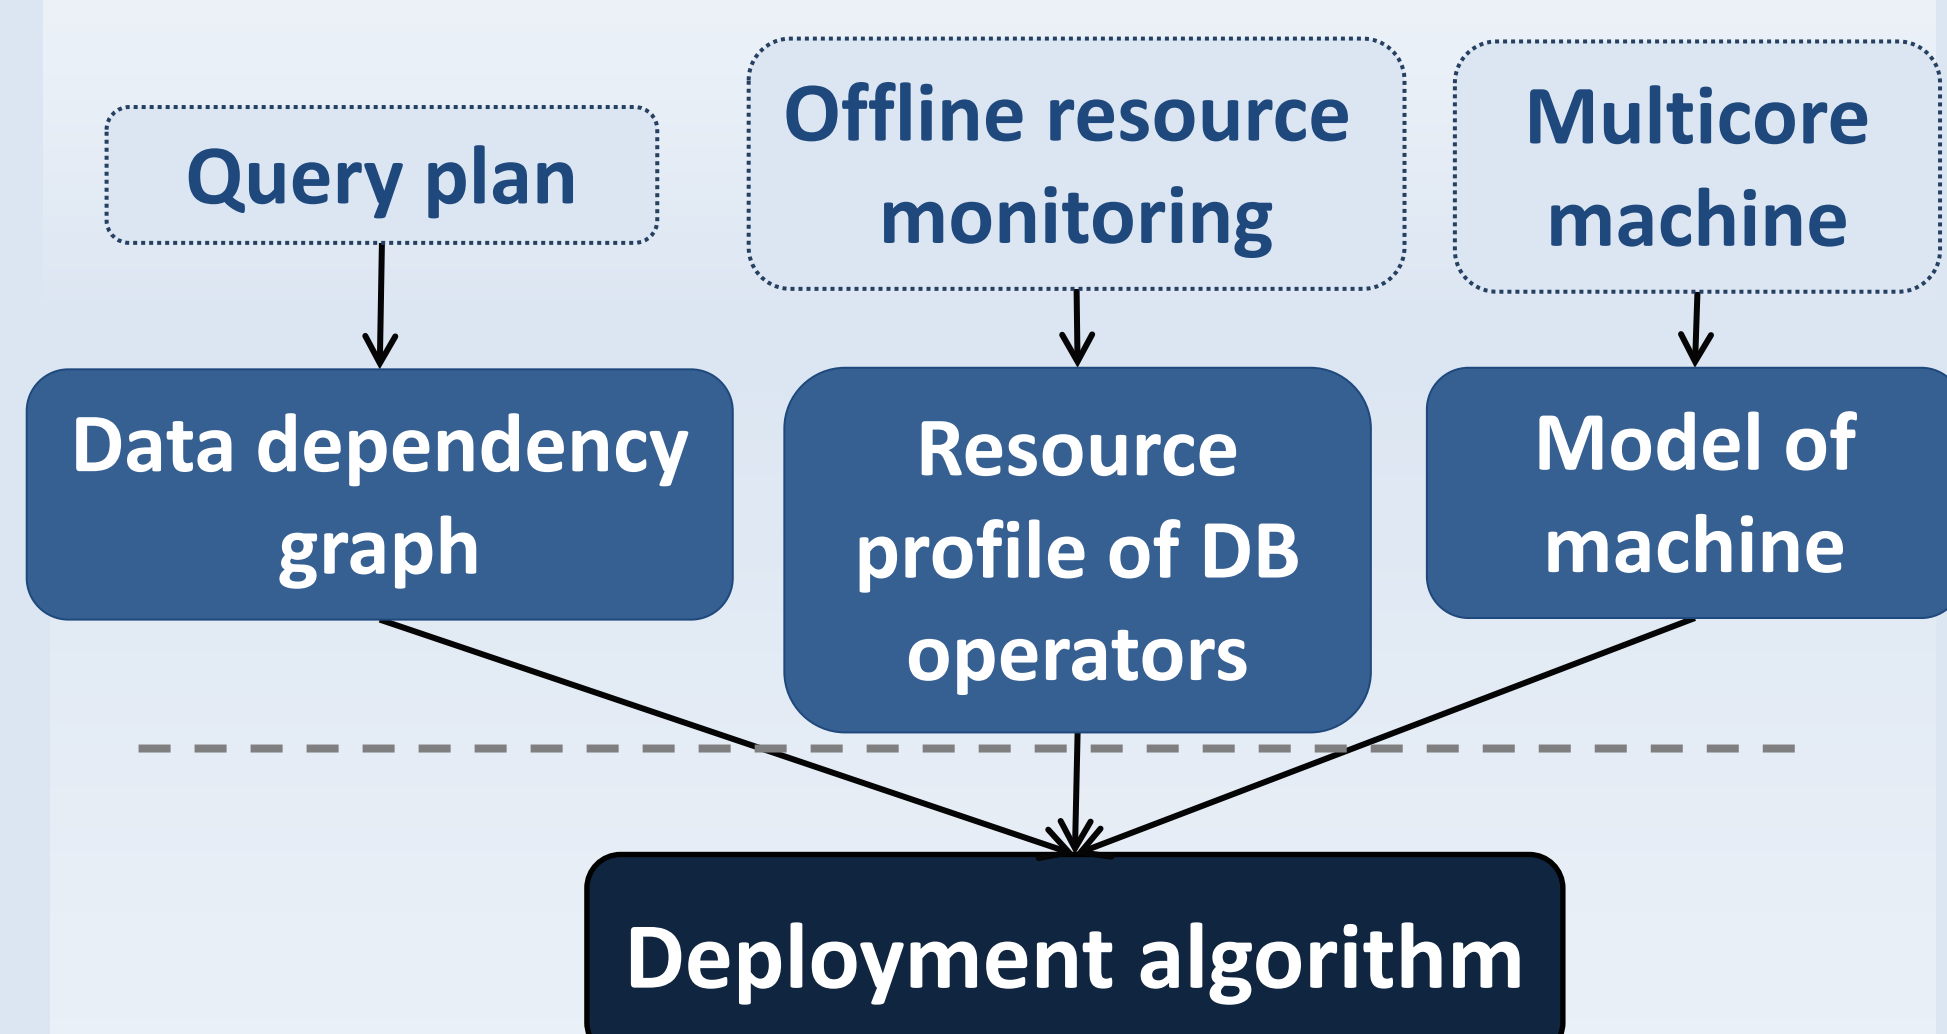
2. Approaches

Until now:

- Default operating system scheduling
- Operator-per-core deployment

New:

3. Our deployment algorithm



Minimize resource usage, provide NUMA-aware deployment & enhance data-locality.

GOAL: Deployment that minimizes resource footprint without affecting performance

3. Algorithm Phases

Input

Data-dependency graph

Algorithm phases

Collapse operator pipelines

Collapse into one scheduling unit operator pipelines that can be temporally ordered.

Resource profile of DB operators

Find minimum number of CPU cores

1st phase of bin packing of operators to find minimum number of CPU cores.

NUMA-node properties

Find minimum number of NUMA nodes

2nd phase of bin packing to determine minimum number of NUMA nodes.

Interconnect topology

Deployment mapping

Data-locality by finding densest k-item sub-graph on NUMA interconnect graph.

Model of multicore machine

Deployment decision

4. Evaluation

Setup:

- Query plan: SharedDB's TPC-W with 44 relational operators
- Dataset size: 20 GB
- Machine: AMD Magnycours

Results:

Approaches	# Cores	Throughput [WIPS]		Latency percentiles [ms]		
		Average	Stdev	50 th	90 th	99 th
Default OS	48	317	31	8	72	82
Operator-per-core	44	425	54	14	22	36
Deployment algorithm	6	428	32	15	23	36

Analysis:

1. Default Operating System scheduling

- ✗ Poorer performance
- ✗ Reduced stability and predictability
- ✗ Uses the whole machine

2. Operator-per-core deployment

- ✓ Good performance
- ✓ Stable and predictable
- ✗ Over-provisions the system

3. Our deployment algorithm

- ✓ Good performance
- ✓ Stable and predictable
- ✓ Minimizes resource requirement

Comparable performance with 1/7th of resources