



Dynamic Adaptivity in Support of Extreme Scale

Storage Performance Isolation: an Investigation of Contemporary I/O Schedulers

Sarala Arunagiri and Patricia Teller
The University of Texas at El Paso
and
Seetharami Seelam*
IBM T J Watson Research Center

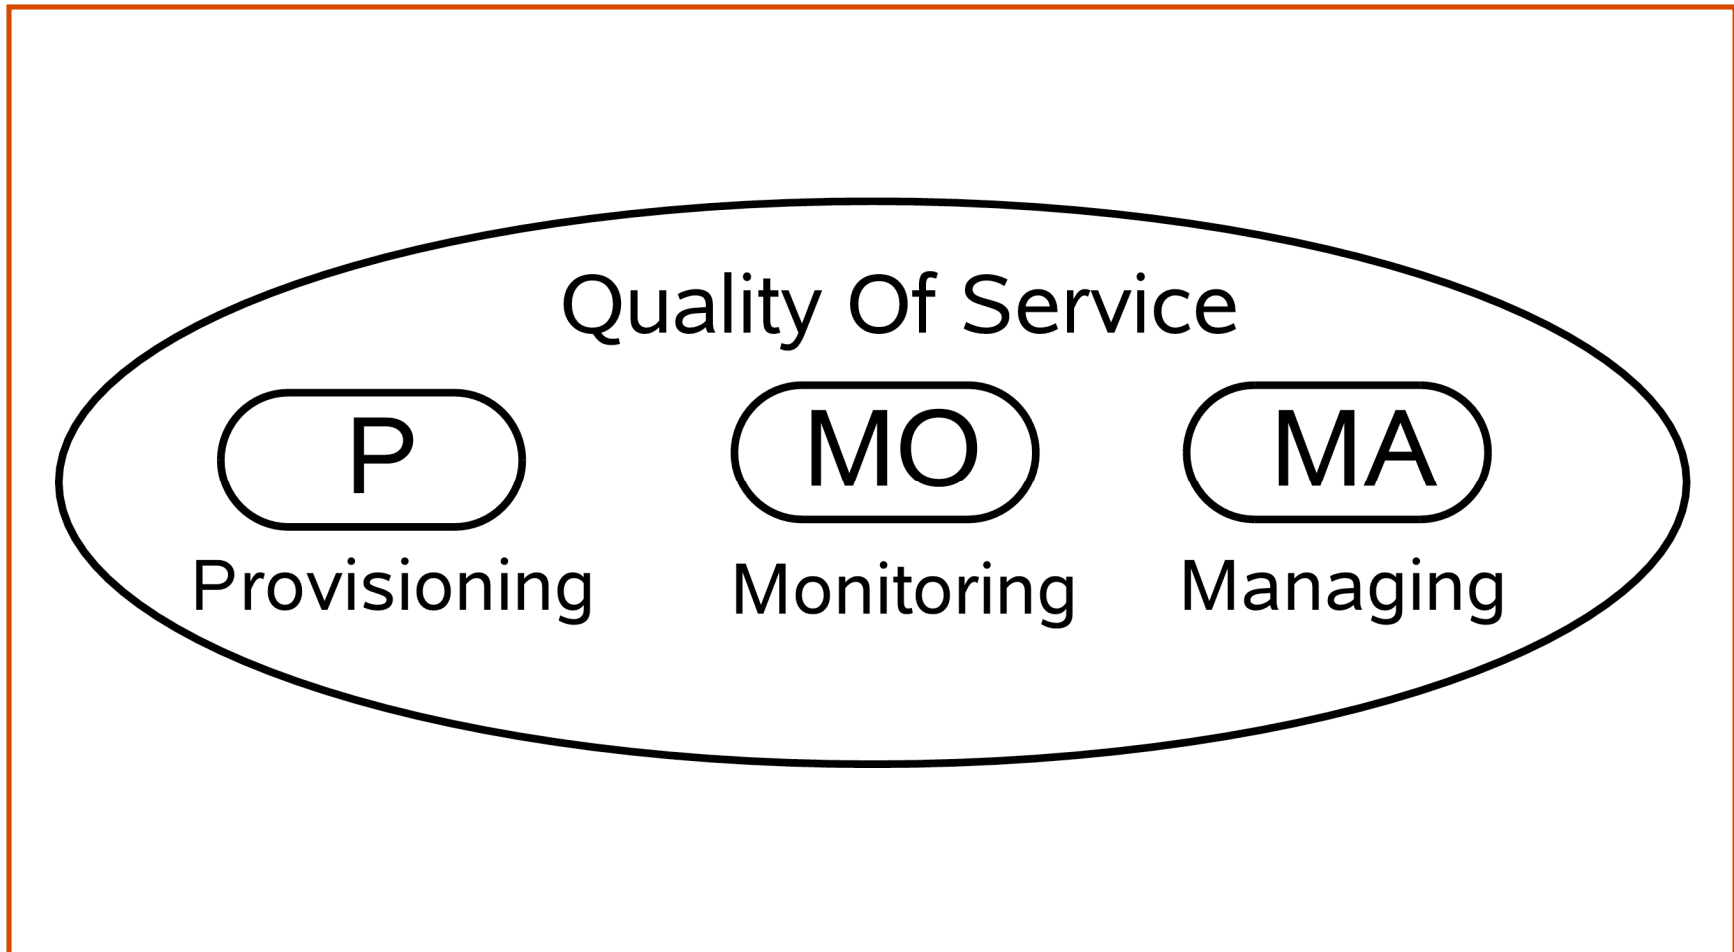
02/14/2007

FAST 2007

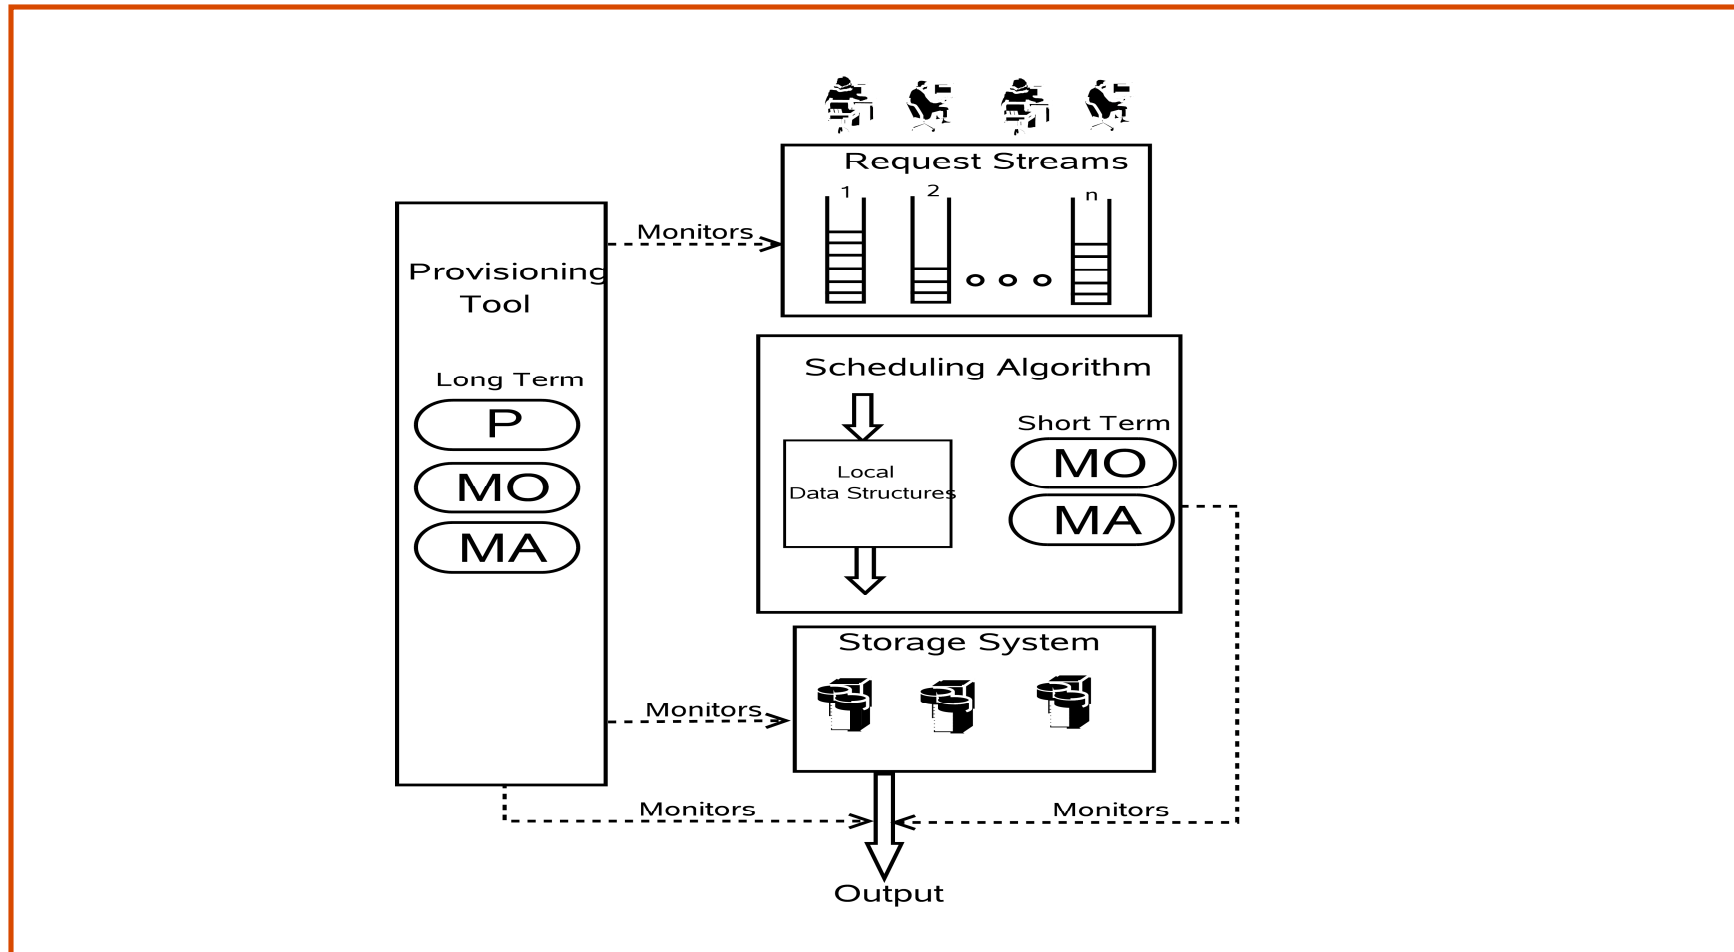
*This work was done while the author was affiliated with UTEP

1

Functional Components of QoS



Metrics Used



The Question

Given limited

- Amount of over provisioning
- Knowledge of I/O stream characteristics

Question: Do the scheduling algorithms
provide Absolute Performance Isolation?
Most schedulers/tools/frameworks
Demonstrate performance isolation using
workloads satisfying implicit assumptions



Dynamic Adaptivity in Support of Extreme Scale

Investigation

<i>Schedulers/tools</i>	<i>Absolute Performance Isolation</i>
Cello disk scheduling framework	Yes
CFQ-CRR(P)	Yes
YFQ	No
Facade virtual store controller	No
Interposed 2-Level scheduler	No
Triage workload controller, both file and block access	No
SLEDS-storage controller, block level storage service	Policy can be set to achieve the goal. Heuristic presented does not provide Absolute Performance Isolation.
Chameleon-storage resource arbitrator	Policy can be set to achieve the goal.



Dynamic Adaptivity in Support of Extreme Scale

More Questions Targeted at Storage Service Providers

- In practice, from a storage service provider's point of view, is Absolute Performance Isolation a desirable property?
 - Are the current schedulers satisfactory in this context?
 - Is there a cost-benefit in trying to design new algorithms to tackle this problem?
- ... and many more questions



Dynamic Adaptivity in Support of Extreme Scale

Current Work

**Design of a Scheduler that addresses
the issue of
Absolute Performance Isolation**

For more information

<http://research.utep.edu/daises>

sarunagiri@utep.edu

seelam@us.ibm.com

pteller@utep.edu