
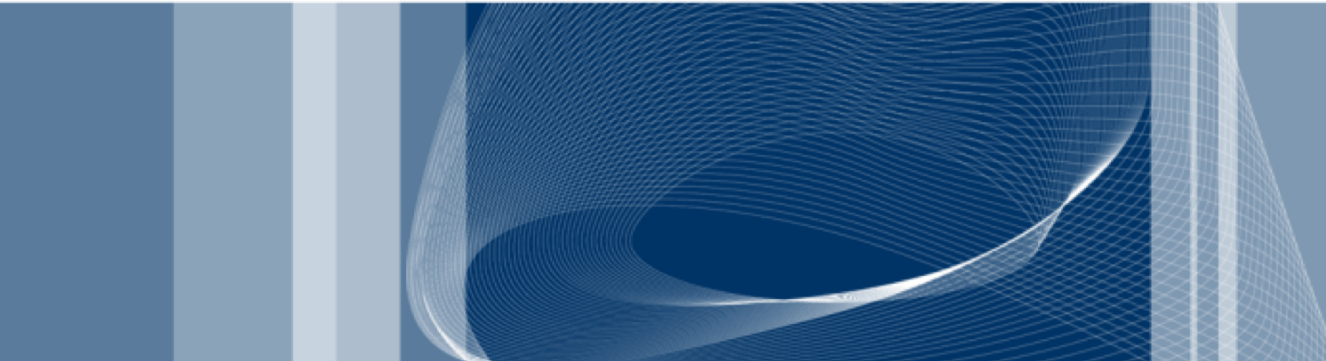


International Conference on Algorithms and Architectures for Parallel Processing – ICA3PP 2016

15 December 2016

 POLITECNICO DI MILANO



MODELING PERFORMANCE OF HADOOP APPLICATIONS: A JOURNEY FROM QUEUEING NETWORKS TO STOCHASTIC WELL FORMED NETS

Danilo Ardagna, Simona Bernardi, **Eugenio Gianni**,
Soroush Karimian Aliabadi, Diego Perez-Palacin,
José Ignacio Requeno



Outline



Background and motivations



Proposed models



Experimental results



Conclusions and future work



Outline



Background and motivations



Proposed models




Experimental results



Conclusions and future work

Background — Big Data



- The Apache Hadoop framework and its ecosystem are enabling technologies for the Big Data paradigm
- They provide automatic parallelization over huge datasets, as well as scalability and fault tolerance



Background – Cloud Offerings



Among the major Cloud providers, Amazon EC2 offers a very interesting pricing model



On demand instances are reliable VMs that users can lease by the hour according to their needs



Spot instances are unreliable VMs hosted on spare capacity with a very competitive, but fluctuating, price



Motivation



The Big Data market amounts to \$ 26.85 billion and is expected to grow steadily (Markets and Markets)



Big Data analyses will be increasingly supported by public Clouds, up to 40% by 2020 (IDC)

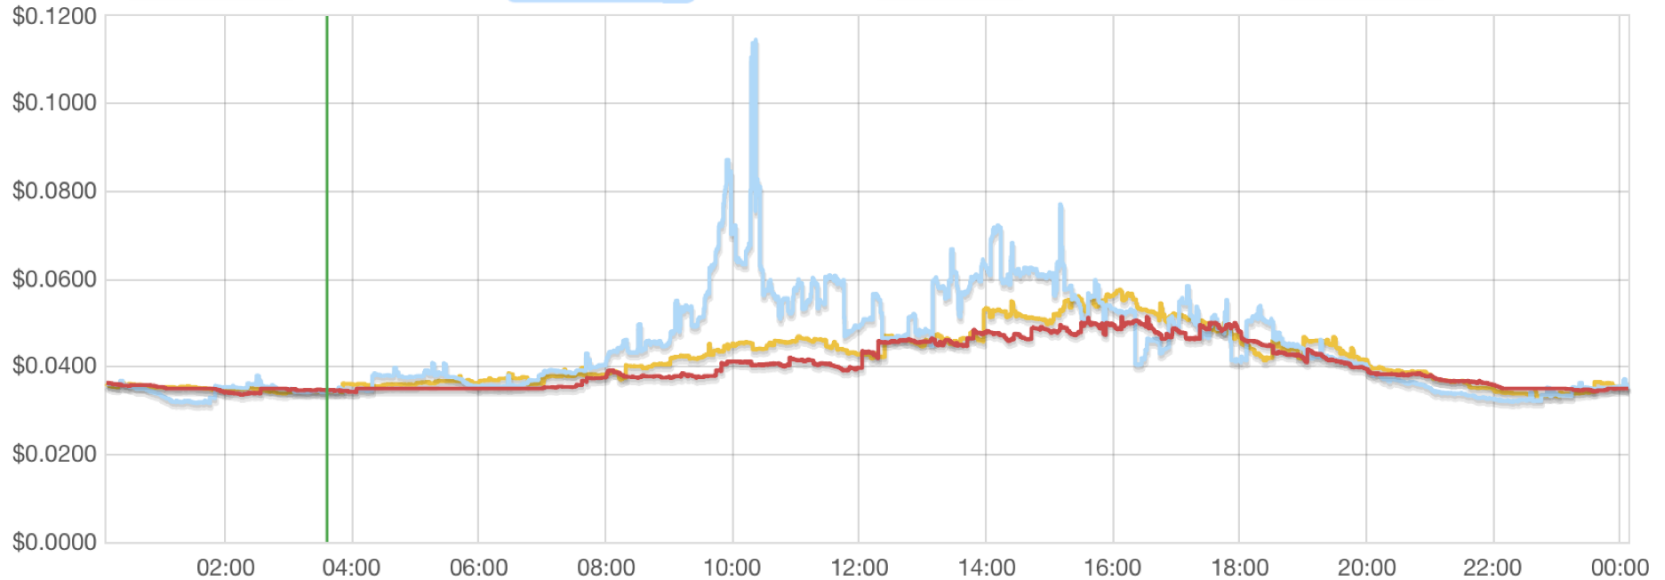


Hadoop 2 overhauled the resource management system to increase cluster utilization, calling for new modeling efforts

Savings Enabled by Spot VMs

Spot Instance Pricing History

Product: Linux/UNIX Instance type: m4.xlarge Date range: 1 day Availability Zone: All Zones



Availability Zone | Price

eu-west-1a	\$0.0343
eu-west-1b	\$0.0342
eu-west-1c	\$0.0346

Date

12/13/2016, 3:35:31 AM UTC+0100



On demand price: \$ 0.238 per hour



Outline



Background and motivations



Proposed models



Experimental results



Conclusions and future work



Models Overview



Different modeling paradigms enable trading off some accuracy for shorter simulation times



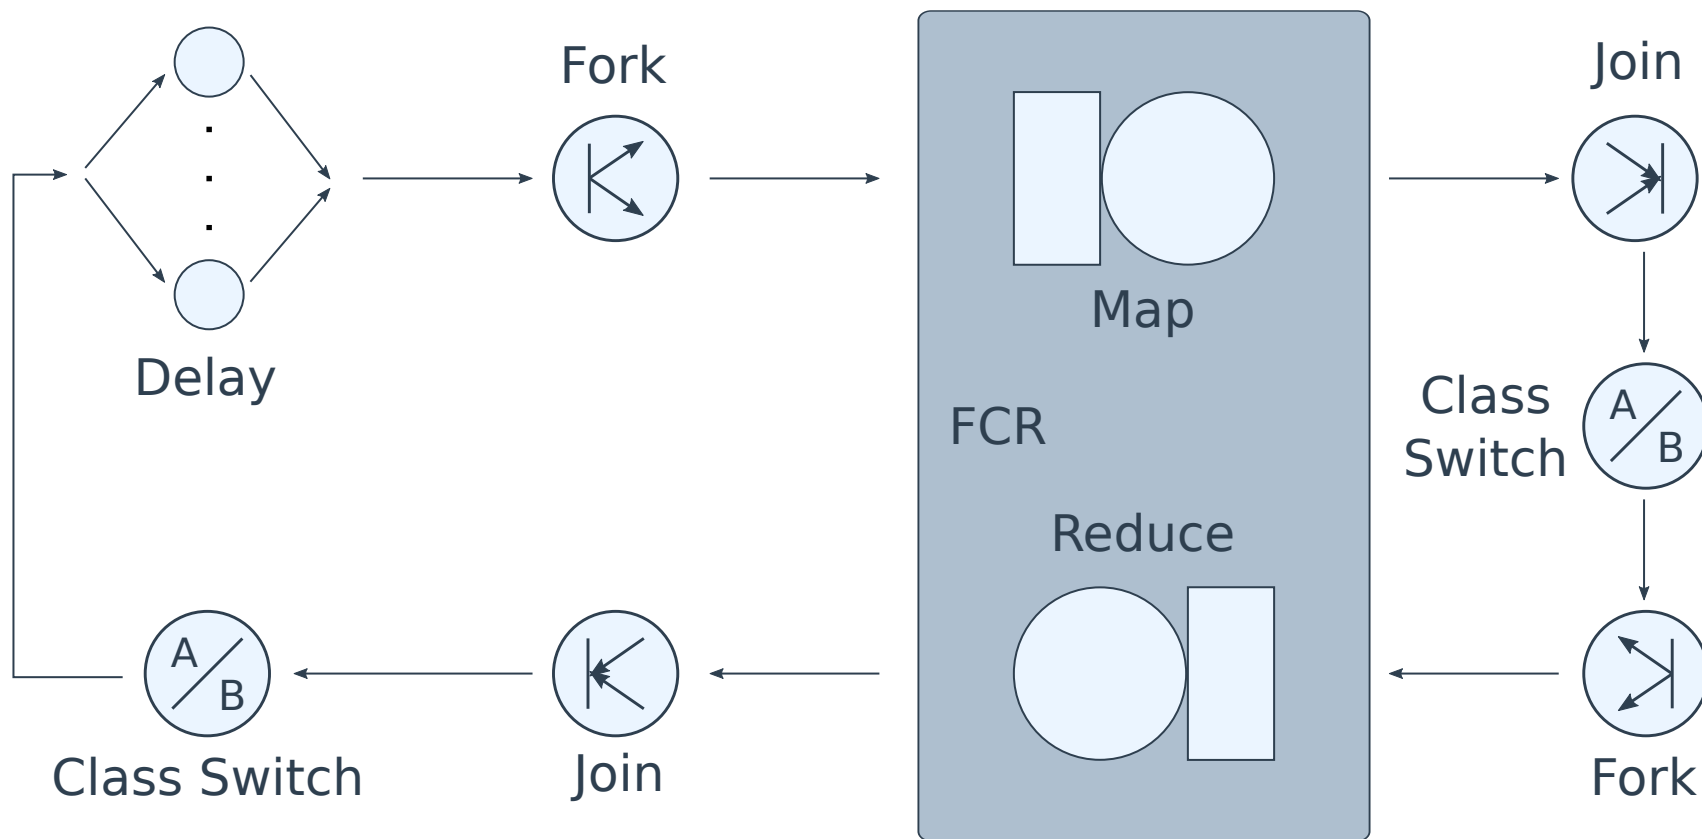
Queueing networks provide a first rough approximation, suitable when the simulation time is constrained



Stochastic well formed nets are powerful enough to accurately model the Hadoop 2 framework



Queueing Network





QN Pros and Cons



Cannot investigate more elaborate scenarios

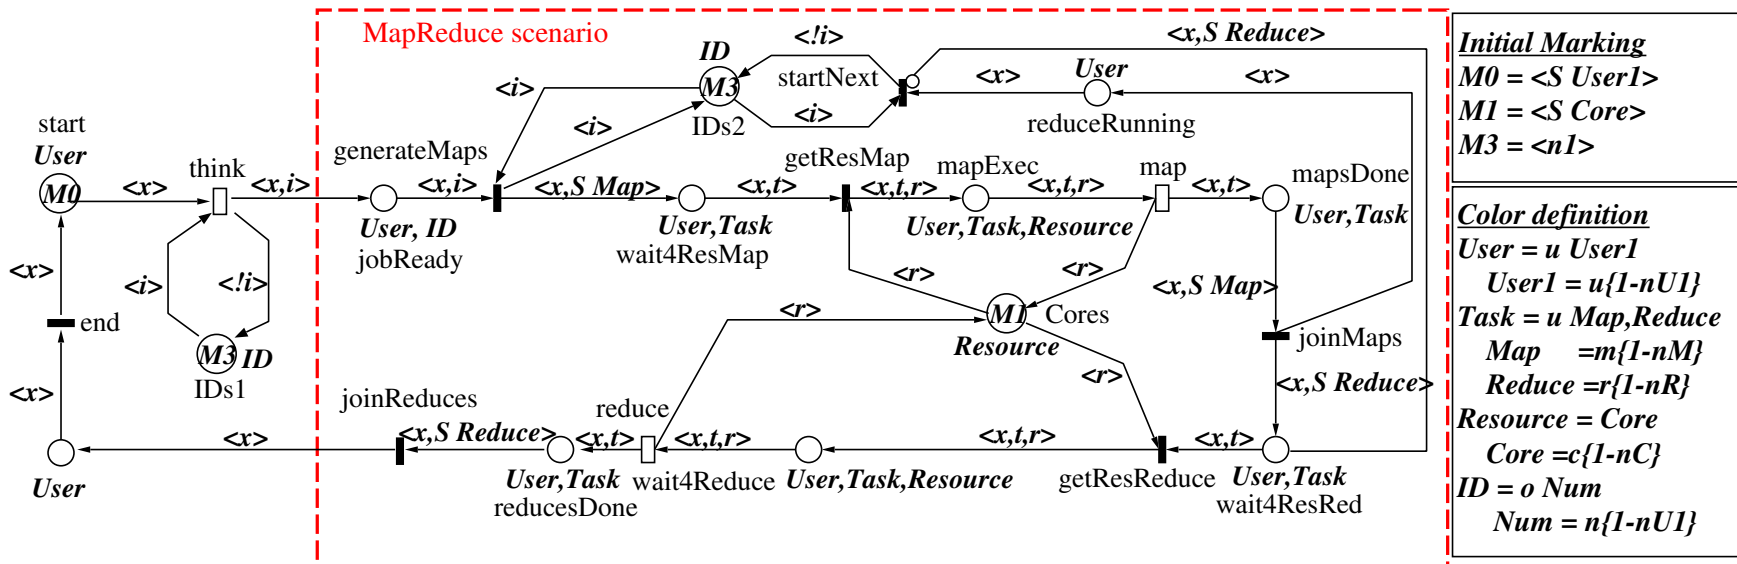


Cannot enforce the FIFO policy at job level as the Capacity Scheduler does



Faster simulations

Stochastic Well Formed Net





SWN Pros and Cons

- ✓ Completely describes the Capacity Scheduler FIFO policy
- ✓ It is possible to go in greater detail, adding complexity and investigating elaborate scenarios
- ✗ Comparably slow simulations



Outline



Background and motivations



Proposed models



Experimental results



Conclusions and future work



Experimental Setting



Five SQL queries run on Apache Hive and profiled to obtain task execution time distributions



TPC-DS benchmark dataset generator, with scale factor ranging from 250 GB to 1 TB



Stopping criteria set up with 10% accuracy and 95% confidence interval


Query	Users	Cores	Scale [GB]	n^M	n^R	T [ms]	τ_{QN} [ms]	ϑ_{QN} [%]	τ_{SWN} [ms]	ϑ_{SWN} [%]
R1	1	240	250	500	1	55410	50753.34	-8.40	50629.58	-8.63
R2	1	240	250	65	5	36881	27495.31	-25.45	37976.82	2.97
R3	1	240	250	750	1	76806	77260.03	0.60	83317.27	8.48
R4	1	240	250	524	384	92197	78573.96	-14.72	89426.51	-3.01
R1	1	60	500	287	300	378127	411940.93	8.94	330149.74	-12.69
R3	1	100	500	757	793	401827	524759.36	30.59	507758.68	26.36
R3	1	120	750	1148	1009	661214	759230.77	14.82	698276.75	5.61
R4	1	60	750	868	910	808490	844700.85	4.48	806366.51	-0.26
R3	1	80	1000	1560	1009	1019973	1053829.78	-1.00	1020294.84	0.03
R5	1	80	1000	64	68	39206	36598.32	-6.65	38796.47	-1.04
R1	3	20	250	144	151	1002160	1038951.05	3.67	909217.89	-9.27
R1	5	20	250	144	151	1736949	1215490.20	-30.02	1428894.40	-17.74
R2	3	20	250	4	4	95403	112050.45	17.45	99219.94	4.00
R2	5	20	250	4	4	145646	97619.46	-32.97	88683.10	3.09
R1	5	40	250	144	151	636694	660241.29	3.70	613577.53	-3.63
R2	3	40	250	4	4	86023	105785.41	22.97	119712.30	-17.81
R2	5	40	250	4	4	90674	103173.38	13.78	117582.82	29.68



Average error: 14.13% for QNs, 9.08% for SWNs

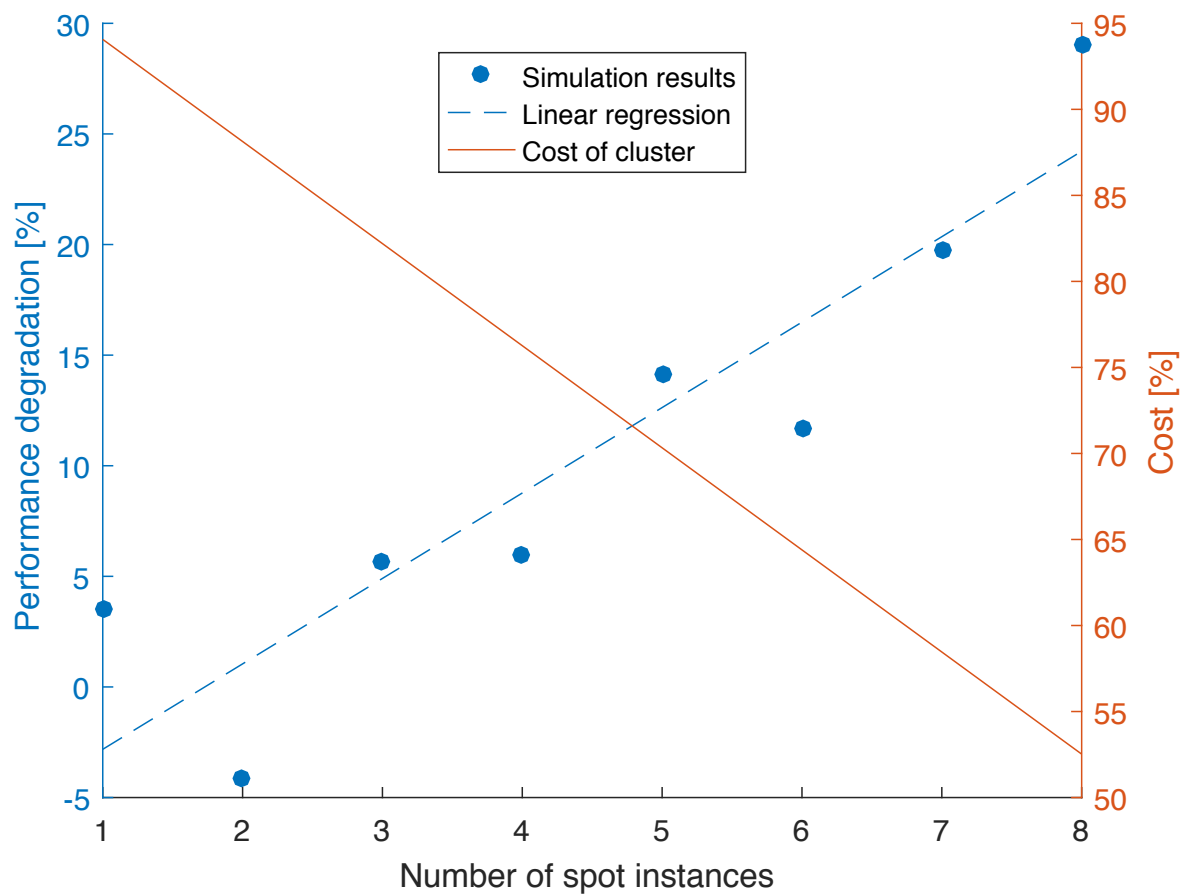


Degradation vs Recovery Time





Cost & Performance vs Spot VMs





Outline



Background and motivations



Proposed models



Experimental results



Conclusions and future work



Conclusions and Future Work



QNs and SWNs provide valuable models of the Hadoop 2 framework, enabling tradeoffs between accuracy and simulation time



Via SWNs, we studied the impact of cheap, but unreliable, computing resources



Extend the models to more complex applications, possibly based on other widely adopted frameworks, like Apache Spark or Storm



**Thanks for your
attention!**