



UvA-DARE (Digital Academic Repository)

High performance N-body simulation on computational grids

Groen, D.J.

Publication date
2010

[Link to publication](#)

Citation for published version (APA):

Groen, D. J. (2010). *High performance N-body simulation on computational grids*.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.



Contents

1	Introduction	1
1.1	Simulating N -body systems	2
1.1.1	Direct-method integration methods	2
1.1.2	Alternative N -body integration methods	4
1.2	High-performance N -body simulations	4
1.2.1	Parallelization and speedup	5
1.2.2	Communication in parallel simulations	7
1.2.3	N -body simulations using special-purpose hardware	7
1.3	Distributed infrastructures and computational grids	8
1.4	Thesis overview	9
1.4.1	Chapter 2	10
1.4.2	Chapter 3	10
1.4.3	Chapter 4	10
1.4.4	Chapter 5	11
1.4.5	Chapter 6	12
2	Distributed Direct N-body Simulations on a Global Grid of GRAPEs	13
2.1	Introduction	13
2.2	Experiment setup	15
2.3	Results of grid calculations	17
2.3.1	Timing results of N -body calculations	17
2.3.2	Profiling of the N -body simulations	18
2.4	Modelling the performance of the grid	19
2.4.1	Single PC	20
2.4.2	Grid of PCs with copy algorithm	22
2.4.3	Grid of PCs with ring algorithm	23

2.4.4	Single PC with GRAPE	24
2.4.5	Grid of PCs with GRAPE and copy algorithm	26
2.4.6	Grid of PCs with GRAPE and ring algorithm	26
2.5	Results of the performance model	27
2.5.1	Future prospects	27
2.6	Discussion and conclusions	30
3	Living Simulations: Galaxy Merger using Direct and Tree N-body Integration	33
3.1	Introduction	33
3.2	Living application	34
3.2.1	Rationale	34
3.2.2	How the living application works	34
3.2.2.1	Security considerations	36
3.2.3	Living simulation	36
3.3	Simulating galaxy mergers as a living simulation	37
3.3.1	Motivation	37
3.3.2	Implementation	38
3.3.3	Experiment setup	39
3.3.4	Results	41
3.4	Conclusion	43
4	Simulating the Universe on an Intercontinental Grid of Supercomputers	45
4.1	Introduction	45
4.2	The intercontinental grid	46
4.2.1	Demand on the network	48
4.3	The simulation environment	49
4.4	Simulating the universe	50
4.5	Concluding remarks	51
5	High Performance Gravitational N-body Simulations on a Planet-wide Distributed Supercomputer	57
5.1	Introduction	57
5.2	Overview of SUSHI	58
5.2.1	Parallelization across supercomputers	58
5.2.1.1	Communication scheme	59
5.2.1.2	Domain decomposition	59
5.2.1.3	Implementation of communication routines	60
5.3	Performance model	60
5.3.1	Single supercomputer	61
5.3.2	Multiple supercomputers	62
5.3.3	Scalability across sites	63
5.4	Experiments	64
5.4.1	DAS-3 experiment setup	64

5.4.2	DAS-3 results	65
5.4.3	Gravitational Billion Body Project experiment setup	68
5.4.3.1	Network configuration	69
5.4.4	GBBP results	70
5.5	Scalability of N -body simulations across supercomputers	77
5.5.1	Speedup and efficiency predictions for TreePM simulations	77
5.5.2	Speedup and efficiency predictions for tree and direct-method simulations	79
5.5.2.1	Performance model for the tree algorithm	81
5.5.2.2	Modelling of block time steps	81
5.5.2.3	Predictions	81
5.5.3	Bandwidth analysis for cosmological simulations	83
5.6	Conclusion	84
6	A Light-Weight Communication Library for Distributed Computing	87
6.1	Introduction	87
6.2	Related work	88
6.3	Architecture of MPWide	88
6.3.1	Design	88
6.3.1.1	Data transport in the wide area network	89
6.3.1.2	Functionality and programming interface	89
6.3.2	Forwarder	91
6.3.3	Implementation	92
6.4	Benchmarking MPWide	92
6.4.1	Results	93
6.4.1.1	Local tests	93
6.4.1.2	National tests	94
6.4.1.3	International tests	95
6.5	Testing performance in a production environment	95
6.5.1	Test experiments	96
6.5.1.1	Results on DAS-3 Dutch grid	97
6.5.1.2	Results on Amsterdam and Helsinki supercomputers	98
6.5.1.3	Results on Amsterdam and Tokyo supercomputers	98
6.5.2	Production	99
6.6	Conclusions and future work	99
7	Conclusions	103
7.1	Thesis summary	103
7.2	Conclusions and Recommendations	105
7.2.1	Simulations on a planet-wide distributed supercomputer	105
7.2.2	Recommendations for planet-wide N -body simulations	107
7.2.3	Creating a planet-wide distributed supercomputer	107
7.2.3.1	Configuration of a global supercomputer	108
7.2.3.2	Scheduling and running on a global supercomputer	109

7.2.4 Future perspectives and applications	110
Nederlandse Samenvatting	111
Acknowledgements	115
List of Publications	119
Bibliography	123