

Identifying and Harnessing Concurrency for Parallel and Distributed Network Simulation

zur Erlangung des akademischen Grades eines

Doktors der Ingenieurwissenschaften/
Doktors der Naturwissenschaften

von der Fakultät für Informatik
des Karlsruher Instituts für Technologie (KIT)

vorgelegte

Dissertation

von

Philipp Andelfinger

aus Weingarten

Tag der mündlichen Prüfung: 10. Februar 2016

Erster Gutachter: Prof. Dr. rer. nat. Hannes Hartenstein
Karlsruhe Institute of Technology (KIT)

Zweiter Gutachter: Dr. Kalyan Perumalla
Oak Ridge National Laboratory (ORNL)

Ich erkläre hiermit, dass ich die vorliegende Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel verwendet habe.
Karlsruhe, 08. Dezember 2015

Zusammenfassung

Da Netzwerke von Computern inhärent parallele Systeme darstellen, können Simulationen von Computernetzen häufig durch parallele und verteilte Ausführung auf mehreren Prozessoren substantiell beschleunigt werden. Simulationsmodelle von Computernetzen können sich jedoch grundlegend in ihrer Eignung für eine Parallelisierung unterscheiden. Obwohl sich bereits eine Vielzahl von vorausgehenden Arbeiten mit der effizienten parallelen Ausführung von Netzwerkmodellen befasst hat, besteht ein Mangel an Verfahren, die Entscheidungen bezüglich der Parallelisierung von Netzwerkmodellen und die Wahl geeigneter Hardware-Plattformen, Simulatorarchitekturen und Synchronisationsansätze unterstützen. Zusätzlich hat die breite Verfügbarkeit kostengünstiger Manycore-Hardware das Spektrum möglicher Realisierungen von Simulatoren über die Möglichkeiten traditioneller CPU-basierter Ansätze hinaus erweitert.

Diese Dissertation betrachtet die effiziente Ausführung von Netzwerksimulationen aus zwei Perspektiven: zunächst werden Evaluationsmethoden vorgeschlagen, die eine Abschätzung des Parallelisierungspotentials von Netzwerkmodellen erlauben. Im Anschluss werden Ansätze vorgestellt, die das identifizierte Parallelisierungspotential mittels moderner Manycore-Hardware effizient ausnutzen.

Identifizierung von Nebenläufigkeit: wir stellen einen analytischen Ansatz vor, der die durchschnittliche Anzahl an Recheneinheiten abschätzt, die von einem idealisierten parallelen oder verteilten Simulationslauf eines gegebenen Netzwerkmodells ausgelastet werden können. Die Abschätzung erfolgt auf Basis von Modellwissen und einfachen Netzwerkstatistiken, die in sequentiellen Simulationsläufen gewonnen werden. Da das vorgestellte Verfahren nicht auf automatisierten Methoden wie der “Critical Path Analysis” beruht, erlauben die Schätzungen ein Verständnis von Zusammenhängen zwischen den Kommunikationsmustern im simulierten Netzwerk und der resultierenden Nebenläufigkeit des Netzwerkmodells. Ein Verständnis dieser Zusammenhänge kann Modelloptimierungen und die Auswahl geeigneter Simulatorarchitekturen unterstützen. Das Verfahren basiert auf einer näherungsweise Bestimmung des Simulationsfortschritts unter dem Synchronisationsalgorithmus YAWNS. Wir legen den Zusammenhang zwischen Critical Path Analysis und YAWNS dar und beweisen die Gültigkeit unseres Ansatzes sowie existierender Arbeiten, die unsere Annahmen teilen. Eine Evaluation der Akkuratheit konkreter Schätzungen wird anhand einer Anwendung des Schätzverfahrens auf Implementierungen dreier Netzwerkmodelle in bekannten Netzwerksimulatoren durchgeführt.

Um zusätzlich zu den Eigenschaften des untersuchten Netzwerkmodells auch die zur Durchführung verwendete Hardware und den Synchronisationsansatz zu berücksichtigen, stellen wir ein Werkzeug vor, das die Laufzeit paralleler und verteilter Simulationen auf Basis sequentieller Simulationen und Hardware-Messungen prädiziert. Das Werkzeug führt eine Simulation einer geplanten parallelen oder verteilten Simulation durch ("Simulation zweiter Ordnung"). Wir zeigen, dass im Falle von Netzwerkmodellen mit nicht-trivialen Berechnungen pro simulierter Nachricht eine angemessen akkurate Prädiktion erreicht wird.

Nutzung von Nebenläufigkeit: traditionelle Ansätze der parallelen und verteilten Simulation greifen zur Ausführung des Netzwerkmodells auf einen Verbund von CPUs zurück. In einigen existierenden Arbeiten ergaben sich für Modelle von Peer-to-Peer-Netzwerken durch Parallelisierung der Simulation nur geringe Laufzeitverbesserungen. Wir analysieren und vergleichen zwei Partitionierungsstrategien für Modelle von Netzwerken, die auf dem Protokoll Kademia basieren. Ein Beispiel für ein solches Netzwerk ist die BitTorrent DHT, eines der größten öffentlichen Peer-to-Peer-Netzwerke. Mittels einer Partitionierung der Simulation basierend auf der logischen Topologie des Netzwerkes erreichen wir eine Beschleunigung der Simulation um einen Faktor von 6,0 im Vergleich mit einer sequentiellen Ausführung sowie eine nahezu lineare Reduktion des Speicherbedarfs pro Rechenknoten.

Da die mittels einer CPU-basierten parallelen und verteilten Ausführung erreichte Beschleunigung einer Simulation die hierzu erforderlichen Hardware-Ressourcen nicht in jedem Falle rechtfertigen kann, untersuchen wir die Ausführung von Netzwerksimulationen auf Grafikprozessoren (Graphics Processing Units, GPUs). Heutige GPUs sind in der Lage, allgemeine Berechnungen auf hunderten oder tausenden paralleler Recheneinheiten durchzuführen. Eine im Arbeitsplatzrechner eines Forschers vorhandene kostengünstige GPU kann dazu dienen, die Wartezeit zwischen Änderungen an einem Netzwerkmodell und dem Erhalten von Simulationsergebnissen zu verringern.

Zunächst vergleichen und evaluieren wir Architekturen zur GPU-Beschleunigung rechenaufwändiger Schritte einer CPU-basierten detaillierten Simulation drahtloser Netzwerke. Obwohl eine einzelne simulierte Nachrichtenübertragung bereits Möglichkeiten zur parallelen Verarbeitung bietet, zeigen unsere Messungen, dass eine signifikante Beschleunigung der Simulation es erforderlich macht, mehrere Nachrichtenübertragungen aggregiert zu betrachten. Um die Korrektheit der Simulation zu gewährleisten, muss hierbei die Möglichkeit von Interaktionen zwischen mehreren Sendevorgängen berücksichtigt werden. Unsere Ergebnisse demonstrieren daher, dass bereits im betrachteten Fall einer GPU-Beschleunigung einzelner Schritte einer durch eine CPU verwalteten Simulation Synchronisationsmechanismen aus dem Feld der parallelen und verteilten Simulationen erforderlich sind.

Schließlich stellen wir einen rein GPU-basierten Simulationsansatz vor, in welchem neben dem Netzwerkmodell auch die gesamte Simulationslogik auf einer GPU ausgeführt wird. Da auf eine Interaktion zwischen einer CPU des Host-Systems und der GPU weitestgehend verzichtet wird, eignet sich der rein GPU-basierte Ansatz auch im Falle von Netzwerkmodellen, deren Simulationsereignisse jeweils nur ge-

ringfügige Berechnungen erfordern. Im Gegensatz zu existierenden Arbeiten werden in unserem Ansatz die simulierten Knoten zu Gruppen zusammengefasst, welche jeweils gemeinsam betrachtet werden. Diese Aggregation erlaubt es, die Auslastung der Recheneinheiten der GPU gegenüber dem Verwaltungsaufwand der Simulation abzuwägen, indem der Grad an Aggregation basierend auf Messungen der Simulationsleistung dynamisch zur Laufzeit angepasst wird. Eine Leistungsbewertung unserer Implementierung des Ansatzes anhand eines Modells Kademia-basierter Netzwerke und des Benchmark-Modells PHOLD zeigt eine Beschleunigung der Simulationen um einen Faktor von bis zu 19,5 bzw. 27,5 im Vergleich mit einer sequentiellen CPU-basierten Ausführung, sowie eine Ereignisrate von bis zu $6,8 \times 10^6$ bzw. $39,3 \times 10^6$ Ereignissen pro Sekunde auf einer einzelnen GPU.

Abstract

Since computer networks are inherently parallel systems, simulations of computer networks can in many cases be accelerated substantially through parallel and distributed execution on a set of interconnected processors. Still, simulation models of computer networks vary significantly in their parallelization potentials. Although an enormous number of works consider the efficient parallel execution of specific network models, there is still a lack of guidelines that help in decisions on parallelization and in the selection of hardware platforms, simulator architectures and synchronization approaches that enable an efficient execution. Further, the advent of commodity many-core devices has broadened the range of possible simulator realizations beyond the possibilities of traditional CPU-based approaches.

This dissertation addresses the efficient execution of simulation models of computer networks from two perspectives: we propose evaluation methods to determine a model's parallelization potential and provide simulator realizations that efficiently exploit the identified potentials using modern many-core hardware.

Identifying Concurrency: we propose an analytical approach to estimate the concurrency of network models, i.e., the number of processors that can be occupied in an idealized parallel and distributed simulation run based only on model knowledge and simple network statistics from sequential simulation runs. By not relying on an automated “black-box” method such as critical path analysis, our estimation approach exposes the relationships between the communication patterns in a simulated network and the resulting concurrency of the simulation. Insights into these relationships may guide model optimizations and the selection of suitable simulator architectures. Our estimations approximate the progress of simulations performed using the well-known synchronization algorithm YAWNS. After clarifying the relationship between critical path analysis and YAWNS, we provide a proof that shows the validity of the general approach and that substantiates the results of previous works that share our assumptions. Empirical results on the example of three network models implemented in popular simulators demonstrate that the estimations are sufficiently accurate to evaluate the models' parallelization potential, while avoiding an automated “black-box” analysis of sequential simulation runs.

In order to take into account both the properties of the considered network model as well as the execution hardware and synchronization approach, we present a tool that predicts the runtime of parallel and distributed simulations on the basis of sequential simulation runs and hardware measurements. The tool performs a simulation of an envisioned parallel and distributed simulation (“second-order simulation”). We show

that reasonably accurate runtime predictions are achieved for distributed simulation runs in the case of network models where simulated messages require non-trivial amounts of computation.

Harnessing Concurrency: traditionally, parallel and distributed simulations rely on interconnected CPUs for execution of the simulation model. In some previous works, models of peer-to-peer networks have been reported to benefit only to a small degree from parallelization in CPU-based execution environments. We analyze and compare two partitioning strategies for models of Kademlia-based networks such as the BitTorrent DHT, one of the largest public peer-to-peer networks. When applying a partitioning strategy based on the logical topology of the network, we achieve a simulation speedup of 6.0 compared to a sequential execution and a near-linear reduction of the memory requirements per execution node.

Since high-performance CPU-based parallel and distributed simulations can consume enormous amounts of hardware resources that may not be justified by the achieved runtime reductions, we consider the acceleration of network simulations using graphics processing units (GPUs). GPUs have evolved to support general-purpose computations on hundreds or thousands of parallel processing elements. An inexpensive GPU in a researcher's workstation can be used to shorten the feedback loop between changes to the network model and the retrieval of the corresponding simulation results.

We compare and evaluate architectures for a GPU-based coprocessing of detailed wireless network simulations. Although each simulated transmission already provides opportunities for parallel processing, we show that significant runtime reductions require an aggregated consideration of multiple transmissions in parallel. Since, in order to maintain correctness, the potential for interactions between multiple transmissions must be considered, the results show that even a simple GPU-based coprocessing requires synchronization mechanisms from the field of parallel and distributed simulation.

Subsequently, we propose a fully GPU-based simulation approach that executes all simulation logic as well as the network model on a GPU. Due to avoiding most interaction between a host CPU and the GPU, the fully GPU-based approach is applicable to network models where individual events require only small amounts of computation. Contrary to existing works, our approach aggregates sets of simulated nodes that are considered jointly. The aggregation enables the exploitation of a tradeoff between the utilization of the GPU's processing elements and simulation overheads by dynamically adapting the degree of aggregation according to performance measurements at simulation runtime. We conduct a performance evaluation of our implementation on the example of a model of Kademlia-based networks and the well-known PHOLD benchmark model. Using a single commodity GPU, we achieve a speedup of up to 19.5 and event rates up to 6.8×10^6 for the model of Kademlia-based networks. A speedup of up to 27.5 and event rates of up to 39.3×10^6 events per second of wall-clock time are achieved in the case of the PHOLD model.

Contents

| | |
|--|-----------|
| Zusammenfassung | i |
| Abstract | v |
| List of Figures | xi |
| List of Tables | xv |
| 1 Introduction | 1 |
| 1.1 Contributions | 3 |
| 1.2 Thesis Outline | 5 |
| I Fundamentals | 7 |
| 2 Parallel and Distributed Network Simulation | 9 |
| 2.1 Synchronization of Simulated Time | 11 |
| 2.1.1 Conservative Algorithms | 12 |
| 2.1.2 Optimistic Algorithms | 13 |
| 2.2 Performance Evaluation | 13 |
| 2.2.1 Taxonomy | 14 |
| 2.2.2 Strategies | 16 |
| 2.3 Simulator Adaptation | 19 |
| 2.3.1 Lookahead | 19 |
| 2.3.2 Partitioning | 20 |
| 2.3.3 Synchronization | 21 |
| 3 Considered Network Models | 23 |
| 3.1 Peer-to-Peer Overlay Network | 24 |
| 3.2 TCP/IP in a Fixed Topology | 26 |
| 3.3 Wireless Ad-Hoc Communication | 26 |
| 3.4 PHOLD Benchmark Model | 28 |
| 3.5 Comparison | 29 |

| | | |
|-----------|--|-----------|
| II | Identifying Concurrency | 31 |
| 4 | Identifying Concurrency – Introduction | 33 |
| 5 | Analytical Concurrency Estimation Approach | 35 |
| 5.1 | Fundamental Algorithms | 37 |
| 5.2 | Methodology | 40 |
| 5.2.1 | Consideration of Fixed Lookahead | 40 |
| 5.2.2 | Relationship between Critical Path Analysis and Synchronization Algorithms | 40 |
| 5.2.3 | Analytical Concurrency Estimation Model | 42 |
| 5.2.4 | Limited Deviation between ACPA- and YAWNS-Based Concurrency | 43 |
| 5.3 | Network Model Analysis | 48 |
| 5.3.1 | Peer-to-Peer Overlay Network | 48 |
| 5.3.2 | TCP/IP in a Fixed Topology | 51 |
| 5.3.3 | Wireless Ad-Hoc Communication | 54 |
| 5.4 | Evaluation | 56 |
| 5.4.1 | Sensitivity Analysis | 56 |
| 5.4.2 | Validation of Estimations | 59 |
| 5.5 | Towards a Consideration of Variable Event Processing Times | 62 |
| 5.5.1 | Refined Concurrency Estimation Model | 62 |
| 5.5.2 | Impact of Variable Event Processing Times | 64 |
| 5.6 | Discussion | 65 |
| 5.7 | Conclusions | 66 |
| 6 | Second-Order Network Simulation | 67 |
| 6.1 | Methodology | 68 |
| 6.1.1 | Modeling Levels | 68 |
| 6.1.2 | Prediction Workflow | 68 |
| 6.1.3 | Model Components | 69 |
| 6.1.4 | Hardware Measurements | 70 |
| 6.1.5 | Second-Order Simulator Operation | 71 |
| 6.2 | Performance Predictions | 73 |
| 6.2.1 | Experiments | 74 |
| 6.2.2 | Evaluation | 74 |
| 6.3 | Discussion | 76 |
| 7 | Identifying Concurrency – Conclusions | 77 |

| | | |
|------------|---|------------|
| III | Harnessing Concurrency | 79 |
| 8 | Harnessing Concurrency – Introduction | 81 |
| 9 | CPU-Based Distributed Simulation of Kademlia-Based Networks | 83 |
| 9.1 | Related Work | 84 |
| 9.2 | Partitioning Schemes | 85 |
| 9.2.1 | ID-Based Partitioning | 86 |
| 9.2.2 | Location-Based Partitioning | 87 |
| 9.3 | Simulator Evaluation | 88 |
| 9.3.1 | Performance | 89 |
| 9.3.2 | Synchronization Efficiency | 91 |
| 9.4 | Conclusions | 96 |
| 10 | GPU-Based Parallel Simulation | 97 |
| 10.1 | General-Purpose Computation on Graphics Processing Units | 98 |
| 10.1.1 | The Graphics Pipeline | 99 |
| 10.1.2 | NVIDIA CUDA | 100 |
| 10.2 | Related Work | 101 |
| 10.2.1 | Hybrid CPU-GPU-Based Simulation | 101 |
| 10.2.2 | Fully GPU-Based Simulation | 103 |
| 10.3 | Hybrid CPU-GPU-Based Simulation of Wireless Networks | 107 |
| 10.3.1 | Proposed Simulator Architectures | 107 |
| 10.3.2 | Evaluation | 108 |
| 10.3.3 | Discussion | 110 |
| 10.4 | Fully GPU-Based Parallel Simulation of Kademlia-Based Networks | 110 |
| 10.4.1 | Proposed Simulation Approach | 111 |
| 10.4.2 | Evaluation | 115 |
| 10.4.3 | Discussion | 127 |
| 10.5 | Conclusions | 127 |
| 11 | Harnessing Concurrency – Conclusions | 131 |
| 12 | Conclusions and Outlook | 133 |
| | Bibliography | 139 |

List of Figures

| | | |
|-----|--|----|
| 2.1 | A simulation round using YAWNS: events in $\{t_{\min}, t_{\min} + 1, \dots, t_{\max}\}$ create no events with timestamps below t_{\max} and are thus safe to be executed in parallel. | 13 |
| 2.2 | Taxonomy of the main factors critical to PADS performance (L: Lookahead, P: Partitioning, C: Communication, S: Synchronization). . . . | 15 |
| 2.3 | Example of an event precedence graph for a simulation of a network of three nodes. An edge $e_1 \rightarrow e_2$ signifies the precedence relation “event e_1 must be executed before event e_2 ”. Dashed rectangles signify groups of events that can be executed in parallel. The left-hand graph shows the precedence structure on the node level, corresponding to the full concurrency in the network simulation itself. In the right-hand graph, nodes A and B are assigned to the same logical process, which is reflected by a reduction in the number of events that can be executed in parallel. | 15 |
| 3.1 | A single campus network of the topology defined in the context of the NMS program (Figure adapted from [PR11]). | 27 |
| 3.2 | Sequence of events during a transmission in the wireless network model. | 27 |
| 5.1 | Critical path analysis of a precedence graph with fixed event processing times. | 37 |
| 5.2 | Synchronization using YAWNS. Events with timestamps $\leq t_{\max}$ can be processed safely. | 39 |
| 5.3 | Relationships between concurrency results of ACPA, YAWNS and our estimation approach. | 41 |
| 5.4 | Full precedence graph. | 45 |
| 5.5 | Reduced precedence graph. | 45 |
| 5.6 | All cases of event visibility and node assignment when considering events for execution using ACPA on an arbitrary reduced precedence graph. Leaf nodes: number of events processed in the current execution. | 47 |
| 5.7 | Processing of a reduced precedence graph using ACPA. Dashed lines indicate the positions of the YAWNS lookahead windows. Boxes indicate events that are processed by ACPA in a single execution. In this example, in all but the first execution, ACPA processes three events per execution. | 48 |

| | | |
|------|--|----|
| 5.8 | Upper bound of $X_{\text{YAWNS}}/X_{\text{ACPA}}$ | 48 |
| 5.9 | Event patterns in Kademlia_A : lookups are composed of α overlapping sequences of RPCs. Node 1 performs a lookup with $\alpha = 1$ and $\alpha = 2$ | 50 |
| 5.10 | A campus network in the NMS network model (Figure adapted from [PR11]). | 51 |
| 5.11 | Event patterns in the NMS model: a single packet is transmitted from node 1 to 3 via node 2. Node 3 replies with an acknowledgement. | 53 |
| 5.12 | Concurrency of a single transmission in Wireless_A | 55 |
| 5.13 | Sensitivity analysis of Kademlia_A | 57 |
| 5.14 | Sensitivity analysis of the NMS model | 58 |
| 5.15 | Sensitivity analysis of Wireless_A , varying the number of nodes and the beacon rate. | 58 |
| 5.16 | Comparison of YAWNS (Y) with ACPA concurrency (C). | 60 |
| 5.17 | Comparison of analytical estimate (C_{est}) with ACPA concurrency (C). | 60 |
| 5.18 | Expected and observed distribution of the number of events per node in the <i>active</i> category in each lookahead window for a run of Kademlia_A with $n = 1000$, $\lambda_{\text{user}} = 300$, $\alpha = 8$, and 0% packet loss. | 61 |
| 5.19 | Expected and observed distribution of the number of events per node of the <i>hub</i> category in each lookahead window for a run of the NMS model with 16 campus networks, 1 node per LAN, and 1Gbps of hub bandwidth. | 62 |
| 5.20 | Comparison of our analytical estimate (C_{est}) with ACPA concurrency (C) of Wireless_A | 62 |
| 5.21 | Distribution of per-event processing time. | 65 |
| 6.1 | Levels of abstraction in modeling of simulations. | 69 |
| 6.2 | Data flow during performance prediction. | 69 |
| 6.3 | Finite state machine description of the LP behavior in SONSim. | 72 |
| 6.4 | Accuracy of runtime estimations of the Kademlia_B model. | 75 |
| 6.5 | Accuracy of runtime estimations of the PHOLD model. | 75 |
| 6.6 | Accuracy of runtime estimations of the PHOLD model with an artificial processing time of $100\mu\text{s}$ per event. | 75 |
| 9.1 | Example of ID-based partitioning of a simulated network into 4 logical processes. Each logical process contains peers with IDs sharing a common prefix. | 86 |
| 9.2 | Binary tree structure of the Kademlia routing table of a peer with ID prefix 101. Dashed lines denote edges leading to leaf nodes not sharing the peer's prefix. Each doubling of LPs leads to a cut that displaces the peers in a single leaf node of the routing table to a remote LP. | 87 |
| 9.3 | In the location-based partitioning scheme, peers are assigned to ranges of latitudes ("discs", left) or longitudes ("slices", middle), or to regions of small diameter and equal size (right) on the earth's surface. | 88 |
| 9.4 | Memory usage per LP for a network size of 1 million peers, varying the number of logical processes and the partitioning scheme. | 90 |

| | | |
|-------|---|-----|
| 9.5 | Simulation runtime for a network size of 1 million peers, varying the number of logical processes and the partitioning scheme. | 90 |
| 9.6 | An LP determines its EOT by considering the timestamps of the earliest possible incoming remote event and of the next locally scheduled event. If the earliest of these events will trigger the creation of a remote event, given the lookahead τ , the lowest possible timestamp of the new event is $EOT := \min(EIT, t_{i+1}) + \tau$ | 93 |
| 9.7 | Chronological sequence of activities performed by LP_A as an example for waiting times due to synchronization. LP_A waits for its EIT to advance (1.) before executing further events (2.). At $t_i + 250ms$, a remote event arrives from LP_B (3.). | 93 |
| 10.1 | Hybrid CPU-GPU-based simulation architectures | 108 |
| 10.2 | Speedup achieved by GPU-based parallelization of individual signal processing algorithms compared to sequential execution on a CPU. | 109 |
| 10.3 | Speedup of the proposed hybrid CPU-GPU-based architectures when compared to sequential execution on a CPU. | 109 |
| 10.4 | Fully GPU-based simulation. | 111 |
| 10.5 | During event execution, newly created events are appended to the target LP's FEL in an unsorted fashion. In a subsequent step, the new events are inserted into the FEL in non-decreasing timestamp order. | 113 |
| 10.6 | To resize LPs, FELs are aligned to the start of their respective memory area boundaries and subsequently relocated according to the new boundaries. | 115 |
| 10.7 | Event rate for Kademlia_C varying the memory access synchronization method of the GPU-based simulator variant, the amount of traffic in the simulated network, and the number of peers. | 120 |
| 10.8 | Event rate for Kademlia_C varying the simulator variant, the amount of traffic in the simulated network, and the number of peers. | 121 |
| 10.9 | Event rate of GPU-based simulation of the PHOLD model with $\lambda = 100$ | 122 |
| 10.10 | Event rate of GPU-based simulation of the PHOLD model with $\lambda = 1$ | 123 |
| 10.11 | Event rate of GPU-based simulation of the PHOLD model with $\lambda = 0.01$ | 124 |
| 10.12 | Event rate of GPU-based simulation of the PHOLD model when varying the number of GPU threads per block. | 125 |

List of Tables

| | | |
|------|---|-----|
| 3.1 | Comparison of the considered network models. | 29 |
| 5.1 | Symbols used in algorithms. | 42 |
| 5.2 | Ratio between refined and basic estimations for Kademlia_A with $\alpha = 8$ and 0% timeouts. | 65 |
| 5.3 | Ratio between refined and basic estimation results for the NMS network model. | 65 |
| 9.1 | Average distance between remote peers using location-based partitioning. | 89 |
| 9.2 | Percentage of messages to local peers [%] depending on the partitioning scheme. | 90 |
| 9.3 | Percentage of time spent in the different execution states during simulation runtime. | 91 |
| 9.4 | EOT and EIT quality when varying the number of LPs. | 95 |
| 10.1 | Time complexity of the simulation tasks in the GPU-based simulation approaches, disregarding potential serialization of operations due to parallel access to data structures by multiple threads. | 117 |
| 10.2 | Symbols used in the time complexity analysis. | 117 |
| 10.3 | Percentage of runtime spent on simulation steps for Kademlia_C with $d_{\max} = 10s$ | 122 |
| 10.4 | Event rates [10^6 events/s] using fixed-sized and adaptive LPs to execute Kademlia_C | 126 |
| 10.5 | Event rates [10^6 events/s] when varying the LP size and the maximum speedup achieved through the aggregation of nodes in simulations of the PHOLD with 0% remote traffic. | 126 |

