Performance Engineering for Graduate Students: A View from Amsterdam

Ana-Lucia Varbanescu
University of Twente
Enschede, The Netherlands
a.l.varbanescu@utwente.nl

Stephen Nicholas Swatman
University of Amsterdam
Amsterdam, The Netherlands
s.n.swatman@uva.nl

Anuj Pathania
University of Amsterdam
Amsterdam, The Netherlands
a.pathania@uva.nl

ABSTRACT

HPC relies on experts to design, implement, and tune (computational science) applications that can efficiently use current (super)computing systems. As such, we strongly believe we must educate our students to ensure their ability to drive these activities, together with the domain experts. To this end, in 2017, we have designed a performance engineering course that, inspired by several conference-like tutorials, covers the principles and practice of performance engineering: benchmarking, performance modeling, and performance improvement. In this paper, we describe the goals, learning objectives, and structure of the course, share students feedback and evaluation data, and discuss the lessons learned. After teaching the course seven times, our results show that the course is tough (as expected) but very well received, with high-scores and several students continuing on the path of performance engineering during and after their master studies.

CCS CONCEPTS

• **Social and professional topics** → **Model curricula**; **Computing education programs**; • **Software and its engineering** → *Software performance*.

ACM Reference Format:
Ana-Lucia Varbanescu, Stephen Nicholas Swatman, and Anuj Pathania. 2023. Performance Engineering for Graduate Students: A View from Amsterdam. In *Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis (SC-W 2023), November 12–17, 2023, Denver, CO, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3624062.3624102

1 INTRODUCTION

As high-performance computing (HPC) focuses on the design of applications (and systems) that maximize performance and efficiency. HPC is often defined in the context of big science, large-scale applications, but has been slowly expanding—sometimes with different names, such as HPDA (high-performance data analytics) or HTC (high-throughput computing)—towards medium-size applications and simulation, and even non-scientific applications like, for example, traditional data science or machine learning applications.

Maximizing performance was very often the task of library developers and/or expert developers, who focused specifically on a couple of codes and, with inside knowledge of the implementation and problem domain, managed to steadily improve these applications for the specific target machines—often supercomputers with somewhat custom architectures. However, the rapid development of multi- and many-core processors and accelerators, and their creative combinations in large-scale clusters and supercomputers, has significantly increased the demand for HPC, to the point where it has far exceeded the ability of experts to keep up. To cope with this massive demand, better tools and more systematic approaches are needed to provide highly-optimized, very efficient versions of HPC applications for the broad range of computing systems. We call the discipline that focuses on this systematic approach **Performance Engineering**. In this paper, we discuss the design and implementation of a graduate (i.e., MSc) level course on performance engineering. This is the first and only such course in The Netherlands to this date.

Our approach to performance engineering covers application and system characterization, performance modeling, and optimizations. We target multi-node heterogeneous platforms combining CPUs and GPUs (as accelerators), and use existing methods and tools—often used by experts—to demonstrate how the performance engineering process can be effectively conducted regardless of the specifics of the application and/or machine. To this end, the course covers both theoretical and practical aspects of the process: the lectures describe the concepts and methods, while the practical assignments enable students to actually test these methods and tools, and experience their strengths and limitations. The course further includes a project where, using an application of their choice, the students demonstrate all the stages of the performance engineering process, and ultimately provide a *better* version of the chosen application. In this paper, we describe in detail how these three components are combined.

We have taught the performance engineering course for seven years (in-person and, during the COVID-19 pandemic, online), and we have collected valuable feedback from 41 students who participated in the course. The feedback—from which we extract the data we use in this paper as evaluation data—indicates that students highly appreciate the course and its topics, even though they find the course difficult in terms of workload. The feedback, combined with the assessment results, have also provided us valuable insights into how to further improve the course; we also shared these lessons learned in this paper.

In summary, the main contribution of this work is to describe the first (and only) performance engineering course for HPC in The Netherlands, from design to implementation and evaluation, including lessons learned and future work suggestions. To this end, the remainder of this paper is structured as follows. In Section 2 we

SC-W 2023, November 12–17, 2023, Denver, CO, USA

ACM ISBN 979-8-4007-0785-8/23/11.
https://doi.org/10.1145/3624062.3624102

introduce the main terminology we use, and discuss related courses/tutorials that have inspired us in our design. Next, Sections 3 and 4 introduce the design and implementation of the course. In Section 5 we present the analysis of the students results (and feedback). Finally, we discuss several lessons learned in Section 6, and conclude the paper in Section 7.

2 BACKGROUND AND RELATED WORK

In this section we provide a short overview of the technologies and definitions we use in the course, as well as a brief discussion on related work.

2.1 Computing Systems

The course covers heterogeneous, potentially multi-node systems built from CPUs and GPUs. *CPUs* are representative multi-cores, with multiple caches and shared main memory. *GPUs* are representative many-cores, and are used as accelerators for the main CPU—i.e., the GPU is the accelerator *device* to the CPU *host*. *Multiple nodes*—potentially heterogeneous both intra- and inter-node—are representative for scaled-out systems, like clusters and supercomputers.

2.2 Programming Languages

Programming HPC has not yet reached consensus in terms of a standardized (set of) programming model(s) for its supercomputing systems. Therefore, based on popularity and accessibility, our course uses C/C++ as main programming language, combined with OpenMP (for shared-memory and accelerator programming), CUDA and OpenACC (for NVIDIA GPUs), and MPI for distributed computing. While our approach to performance engineering, and the methods we propose to support it, are agnostic to the programming models and languages being used, we find that most open-source tools are built to work with this mix of models. We further rely on assembly for several of the detailed, low-level analyses we teach in the course.

2.3 Performance Engineering

For the purpose of our course, we adapted the definition of performance engineering from the discipline of software engineering [12] as follows: *Software Performance Engineering (SPE) is a systematic, quantitative approach to the cost-effective development of software systems to meet performance requirements. SPE is a software-oriented approach that focuses on architecture, design, and implementation choices. SPE provides the information needed to build software that meets performance requirements within a given budget (defined in terms of time, cost, or efficiency).*

Further inspired by [5, 12], we define the process of performance engineering as a seven-stage, iterative process, as follows:

Stage 1 Collect and analyse (user) performance requirements.
Stage 2 Understand current performance.
Stage 3 Assess feasibility of the requirements.
Stage 4 Assess suitable approaches to meet the requirements (including algorithm and/or system (co-)design).
Stage 5 Apply tuning and optimization.
Stage 6 Assess progress and iterate back to steps 3–5.
Stage 7 Analyse and document the process and the final result.

The lectures cover all stages, while the practical aspects of the course specifically target stages 2–6, allowing students to learn how to use (and further become proficient using) different tools for each of these stages.

2.4 Related Work

We have started our course from the concepts of "performance engineering" found in the software engineering discipline [12]. We combined these aspects with methods and tools presented in scientific articles such as the Roofline model [17], the ECM model [11], microbenchmarking [18] and benchmarking [9, 10]. We further incorporate the technical and practical aspects described in tutorials presented in the relevant HPC venues, such as single-core performance engineering (at SC and ICPE), performance modeling using the Roofline model for CPUs and GPUs (at SC, NVIDIA TechDays), performance modeling for distributed applications (at SC, ISC), or the polyhedral model (at HiPEAC). As such, we have designed a course that provides a unique, novel combination of theoretical concepts and methodological aspects for performance engineering.

While ours is the first and only course on performance engineering in The Netherlands, we have also drawn inspiration from courses like Computer Systems from CMU[1], Performance Engineering from MIT[2], Performance Modeling from TUDelft[3] and the University of Edinburgh[4], and Queuing Theory from MIT[5]. These courses provided inspiration in terms of matching theoretical topics with tools and assignments. However, to the best of our knowledge, no other academic course currently taught provides the same range and mix of topics as our course.

3 COURSE DESIGN

In this section we describe the goals of our course, the learning objectives, and the structure we propose. We provide further details on the actual implementation in Section 4.

3.1 Goals and Learning Objectives

The course focuses on the modern aspects of performance engineering in the context of parallel applications and distributed heterogeneous (super)computing systems. To this end, the course introduces performance metrics and measurement techniques, performance analysis, benchmarking and microbenchmarking, performance models (analytical and statistical), and performance prediction. The course further demonstrates how to apply these techniques for several applications running on multi-node computing systems, featuring combinations of multi-core CPUs and many-core GPU accelerators. The ultimate goal of the course is to bundle these techniques together and provide students the opportunity to *create their own performance engineering toolbox, which they can successfully use to deploy a systematic approach for performance engineering on any application.*

[1] http://15418.courses.cs.cmu.edu/spring2016/lectures
[2] https://ocw.mit.edu/courses/6-172-performance-engineering-of-software-systems-fall-2018/
[3] https://repository.tudelft.nl/islandora/object/uuid%3Aea4a3f85-597c-479c-b045-15da8dc4850b
[4] https://www.inf.ed.ac.uk/teaching/courses/pm/
[5] https://ocw.mit.edu/courses/15-072j-queues-theory-and-applications-spring-2006/

The course is designed to meet eight learning objectives [15]. Specifically, at the end of the course, students will be able to:

Objective 1 Quantify (using the appropriate tools and methods) the performance of an application running on a computing system using the appropriate metric;
Objective 2 Demonstrate and compare several performance modeling methods, and assess their usefulness for practical problems;
Objective 3 Classify and use several performance prediction methods, and compare their applicability in practice;
Objective 4 Design an empirical performance analysis process for any application, interpret its results, and recommend solutions for performance improvement;
Objective 5 Design and use a suitable model for accurate performance prediction for a given application;
Objective 6 Apply and assess different (existing) optimization techniques to parallel and distributed codes;
Objective 7 Design and develop a complete performance engineering process, apply it successfully on any given application, and assess its outcome in terms of performance gain;
Objective 8 Use different performance engineering tools (e.g., profilers, microbenchmarks and benchmarks, performance counters libraries, etc.).

3.2 Prerequisites

Performance engineering is a multi-disciplinary topic, combining computer systems and computer architecture, algorithm design, parallel and distributed algorithms, and parallel programming. For students to succeed in following this course, we formulated the following prerequisites:

Prerequisite 1 Computer organization and architecture basics, including data representation, CPU architecture and functionality, assembly language literacy, memory hierarchy and caching, multi- and many-core architectures;
Prerequisite 2 Computer systems fundamentals, including shared-memory systems, distributed systems, and heterogeneous systems design and implementation;
Prerequisite 3 Parallel algorithms design and development, and basic skills in C/C++ as support language.
Prerequisite 4 Parallel and distributed programming basics, including OpenMP, CUDA, OpenCL, and MPI.
Prerequisite 5 Basic statistics and data analysis methods, including regression techniques.

The performance engineering course is embedded in the computer science master program which is realized as a joint degree program between the University of Amsterdam and the Vrije Universiteit Amsterdam. Students in the program have knowledge of all the aforementioned prerequisites to a certain extent. The course is mandatory for students in the "Parallel Computing Systems" concentration, and is available as an elective course to all other students. Although we provide quick refreshers for all the topics discussed above, we observe that students who have prior knowledge in these topics (in the sense of having actively taken courses, solved assignments, or developed small projects in these topics) perform better, because they can build their own performance engineering toolbox and expertise on top of assimilated knowledge.

Overall, the most successful students in the course are computer science students, but we have had several successful graduates from computational science (where modeling and numerical methods are focused much more than systems and performance aspects) and exchange students with diverse backgrounds. These successes indicate that the course is self-contained and comprehensive.

3.3 Course Structure

To reach the course goal and achieve its learning objectives, we have structured the course along three types of activities: lectures, assignments, and project.

Lectures. We designed the theoretical part of the course to cover the principle and methods of performance engineering. To this end, we teach *how* to correctly measure and communicate performance data, design microbenchmarks and benchmarking suites, design and validate different performance models, and design and optimize parallel applications for modern computing systems. Lectures combine fundamental systems and applications knowledge with modern, state-of-the-art methods from literature.

Assignments. The course relies on assignments to encourage students to link the theoretical aspects presented in the course with existing processing and tools that facilitate the application of these theoretical aspects in practice. The assignments formulate a general problem, provide a starting point in the form of code templates, and require students to follow specific performance engineering steps to analyze and/or improve the code. All assignments rely heavily on empirical research—from experimental design to data analysis—but use simple applications to limit the coding effort required for completion. All assignments also require detailed documentation to allow students to learn not only how to execute performance engineering tasks, but also how to document them for non-expert stakeholders (e.g., domain experts, application owners, system designers).

Project. While assignments are self-contained and designed to focus on specific aspects of the performance engineering process, the project enables students to experience the deployment of performance engineering for a larger, real case-study application. This enables them to better understand the limitations and challenges of the provided methods and tools, and the manual effort still required when actually aiming to systematically improve application performance, aiming at efficient HPC applications. For the project, the students must achieve four important milestones:

Milestone 1 Define an application of interest and formulate a performance problem to be solved.
Milestone 2 Formulate a plan to deploy performance engineering methods to solve the proposed performance problem. The required elements of the plan are: benchmarking, performance analysis, requirements analysis, performance modeling and prediction, performance optimization, reflection.
Milestone 3 Document the performance engineering process.
Milestone 4 Present their intermediate and final results to an audience of their peers.

3.4 Assessment

The course uses continuous assessment, combining an exam, in-class quizzes, assignments, and a project which includes a report and two presentations.

The exam assesses the theoretical knowledge of the students, targeting mostly the knowledge they have acquired during the lectures. The exam is included to ensure that students do understand the fundamentals of the tools they employ, which in turn demonstrates they can think critically about the output of the tools they use, as well as they can port their knowledge to new systems, tools, and applications. We find that such an exam is a real discriminating factor between students who only master the usage of tools and those who truly understand the methodological aspects of performance engineering. We use in-class quizzes to stimulate students to acquire such knowledge. The exam is, by design, individual—compared with the rest of the course components, which can be solved individually or, ideally, in teams of 2–4 students.

Assignments are provided to encourage students to practice with all tools we believe to be important for a performance engineer's arsenal. We find such assignments work when they are focused and targeted at specific elements in the performance engineering process. The challenging aspect of these assignments is to find he right balance between workload and learning benefits.

Finally, we assess the completion of the chosen project, including the presentation and report the students provide. We assess the project based on the feasibility of the proposed plan to address the identified performance problem, the suitability of the employed methods and tools, and the communication aspect (including the reports and the presentation). We do not factor in the success (or lack thereof) in achieving a specific performance goal, but rather focus on the analysis and reflection aspects—i.e., how and why certain steps succeed or fail.

Overall, the final students grades are calculated as a weighted average of these three grades. We chose to give the largest weight to the project, as many of our students lack experience with working on real applications and systems, while being mostly exposed to simple, "in-lab" assignments that are dedicated to show how methods and tools *do* work. In the case of the projects, students mostly learn how and why many such methods and tools are limited and/or fail in a real-life scenario, and how they can use their theoretical knowledge to compensate for these limitations. We further choose equal weights for the assignments and the theoretical exam, but allow for slack for the students to slightly compensate one with the other, if/when needed (more details in Section 4).

4 COURSE IMPLEMENTATION

In this section we present the actual implementation of the proposed course. Specifically, we (briefly) discuss the alignment of learning objectives to the specific lectures and assignments, and the actual topics we cover with each of them. The course has been running in a block of 8 weeks, with three weekly scheduled activities: lectures, labs (dedicated to the practical assignments), and seminars (dedicated to the projects). The final week is dedicated to the exam and project final presentation.

Table 1: Overview of topics covered in the course, as well as the stages of the performance engineering process (Section 2.3) and the learning objectives (Section 3.1) which motivate each topic.

	Stages							Learning Obj.							
Topic	1	2	3	4	5	6	7	1	2	3	4	5	6	7	8
Basics of performance	✓							✓							
Code tuning and optimization				✓									✓	✓	
Roofline model and extensions		✓						✓	✓	✓				✓	
Analytical modeling		✓	✓					✓	✓	✓	✓				
(Micro)benchmarking	✓		✓								✓		✓	✓	✓
Data-driven and stat. modeling	✓		✓					✓	✓	✓			✓	✓	
Simulation and simulators		✓		✓		✓					✓		✓	✓	✓
Perf. counters and patterns		✓			✓	✓							✓	✓	✓
Scale-out to distributed systems		✓	✓	✓	✓	✓	✓				✓	✓	✓	✓	✓
Queuing theory			✓	✓					✓	✓					
Polyhedral model					✓								✓	✓	✓

4.1 Lectures and Topics

The course covers the topics listed in Table 1, aligned with both the performance engineering process (Section2.3) and the learning objectives (Section 3.1).

The course has been taught with different numbers of lectures: we started from 7 lectures (one per week), but, given the growing list of topics, we eventually extended the number of lectures to 10; we give two lectures per week in the first three weeks, where more theoretical knowledge is needed. The lectures are provided in the performance engineering repository [16].

4.2 Labs and Assignments

The course features four practical assignments, which are briefly introduced in the following paragraphs. The actual assignments, including code, data, and templates for reporting (where needed), are available in the course repository [16]. For all assignments we allow students to use their own machines and we further provide access to the DAS-5 HPC cluster (featuring job isolation and dedicated hardware resources via a SLURM-based scheduler) [1].

Assignment 1: The Roofline Model. The goal of this assignment is to allow students to use one of the simplest performance modeling tools available to performance engineering: the Roofline model [17]. However, the model relies on a deep understanding of computer systems and their main performance indicators, the balance between compute- and memory-bound in systems and applications, and a characterization of the application itself. In this assignment, the students are provided a basic matrix multiplication code, and are asked to provide a Roofline model for the sequential code, then further optimize the code (based on the information from the Roofline model)—thus addressing the identified bottleneck(s)—and reapply the same modeling technique. This exercise aims to demonstrate that the model is able to capture different versions of the same code. Finally, the assignment also requires the students to implement and Roofline-model a parallel version of matrix multiplication. Again, the goal is to demonstrate how the model of both the system and the application change when parallelism is added.

To bootstrap this assignment, we provide sequential code for matrix multiplication, and suggest optimizations like loop reordering and loop tiling. We do not aim to achieve the best possible performing matrix multiplication, but rather to have different versions with different performance envelopes, to demonstrate the sensitivity of the model.

Finally, the assignment requires the students to practice their experimental design skills, as we request to provide models and evaluations of the model using different input datasets. We also suggest tools that can calculate and plot the model automatically, but want the students to reflect on the difference between modeling by hand and by tool. Please note that very little coding is required for this assignment in terms of adapting matrix multiplication (around 100 lines of code), but a lot of effort is required for systematic evaluation and model validation (for which we recommend and eventually provide scripts and templates).

Assignment 2: Analytical Modeling and Microbenchmarking. In this assignment, students need to provide an analytical model of their matrix multiplication versions - sequential and parallel, calibrate these models using microbenchmarking, and evaluate the models against measured performance data. We further add a second kernel to the mix, where we ask students to also model basic histogram calculation, aiming to add data-dependent behavior as additional modeling challenge. The students can reuse their code and/or provided code for both kernels.

The goal for this assignment is threefold: we want students (1) to observe and understand the levels of granularity in analytical models, and the additional calibration challenges that come with those, (2) to get familiar with microbenchmarking as a model calibration tool, and (3) expose students to models that, although potentially very inaccurate, can still provide important insight into the performance of a given application. The students learn by trial and error to find the right level of granularity (ranging from coarse, at function level, to very fine, at ASM instruction level) where the kernels can be accurately model, and understand the difference in the details provided by these different models.

This assignment also provides students the opportunity to use the tabulated performance data for different processors [3], or different microbenchmarking tools, like STREAM or uops. They can further get familiar with detailed performance profilers like perf, NVIDIA's nvprof/nsight, as well as with instruction scheduler simulators like IACA, OSACA [8], or LLVM-MCA.

Assignment 3: Statistical Modeling. In this assignment, students must demonstrate they can work around the limitations of analytical modeling by using machine-learning models. To this end, they are required to collect performance data from a relevant set of inputs, and model—using statistical methods—the expected performance. Finally, they must evaluate the prediction accuracy of the proposed model. To do so, they revisit matrix multiplication, and are provided with a new kernel: sparse matrix-vector multiplication (SpMV). We provide three versions of the kernel—based on the three classical storage models, CSR, CSC, and COO—and request them to model the sequential version, one parallel version of their choice, and compare the results against an analytical model.

With this assignments, we aim to showcase the challenges of defining and collecting training data, of feature engineering, and of empirical validation for such models. We further showcase the interpretability of the models by comparison, by exposing students to two extremes: the highly-explainable analytical model vs. the black-box statistical models. Depending on the types of data they choose to collect, this assignment can further expose students to tools for collecting performance counter data and/or other detailed performance indicators. However, these are separately introduced and demonstrated in assignment 4.

Assignment 4: Performance Counters and Performance Patterns. This final assignment goes into further detail into collecting detailed performance data for one of the provided kernels. We opted for SpMV, but the same exercise can be applied for matrix multiplication. Furthermore, to better illustrate the use of performance counters in the context of performance anomalies hypotheses, we introduce the concept of performance patterns (inspired by Treibig et al. [14]) and encourage students to understand the correlation of performance patterns and observed counters values. For the latter part, we ask students to develop a simple (synthetic) kernel to demonstrate some of this performance patterns, and show they can be identified and fixed using performance counters data.

One important aspect of this assignment is to expose students to clear performance patterns that often appear in HPC applications and algorithms, and teach them how they can identify them with empirical tools. Furthermore, in this assignment, students also get familiar with tools like Linux perf, PAPI [2], Likwid [13], Intel VTune, as well as NVIDIA Nsight Systems and Nsight Compute.

4.2.1 Coverage and Timeline. The current assignments cover all technical skills for performance analysis, modeling, and prediction. We note that we do not cover well the modeling and analysis of distributed systems (where we discuss and demonstrate tools like Vampir [6], Score-P [7], and the Scalasca approach [4]), but we found that we have insufficient time in this format to also cover these tools/approaches into an actual assignment.

We provide assignments somewhat sequentially: first assignment 1 (2 weeks deadline), followed by assignment 2 (2 weeks deadline, some days overlapping with assignment 1), and followed by both assignments 3 and 4, which we release at the same time, and the deadline is the end of the 8-week course period (effectively, students have 3 weeks for both assignments). This approach allows students to mix and match different parts of the assignments, often guided by what they need for their project.

We also provide 2–3 in-person lab sessions (with TAs support) to help students complete each of these assignments. Lab sessions focus on methods and tools demonstration, using simpler/different kernels than those from the assignments, but allowing students to observe how the tools are being used in practice.

4.3 Seminars and the Project

To support the development of the project, we kick-off the project in week 1, and provide the following timeline for students to plan their work:

Week 1 Brief introduction of the project goals and several examples at high-level (dedicated seminar). We focus on describing the goals of the project, and helping students understand what a performance challenge could be for a different application.

Week 2 Students work on providing a prototype of the sequential/reference version. We note that this version can be based on their own code or any open-source code they find useful as reference implementation.

Week 3 Students must define their evaluation strategy and define an experimental setup (dedicated seminar). We focus here on the right infrastructure, metrics, and measurement methods needed for the assessment of current and improved performance of the target application.

Week 4 Students should provide a first performance model of the code, and potentially apply the first optimizations. This leads to creating more prototypes starting from the reference code.

Week 5 Students have a skeleton of their report and provide a short (5-minute) talk to introduce their application, status and goals, and first model.

Weeks 6–7 Students provide more prototypes and a full performance engineering process, evaluating each prototype and motivating (through models) the following prototype(s). These weeks focus on performance analysis, modeling, and improvement.

Week 8 Students finalize project report and presentation, and reflect on whether and how they managed to meet the performance requirements.

For weeks 2, 4–6 we schedule invited talks from previous students and/or performance engineering experts that discuss their approach (and successes) in applying performance engineering. The remainder of the seminar time is allocated to consultation and discussion about the challenges of different projects.

4.4 Grading

The student grades are calculated as a weighted average using the grade for the exam, the assignments, and the project, and using in-class quizzes as bonus points. In accordance with the system ubiquitously adopted in higher education in The Netherlands, final grades lie between 1 and 10: the worst and best possible grades, respectively. In this system, a grade of 5.5 or higher is considered a passing grade. Grades are calculated using Equation 1, where G_P is the project grade, G_A is the assignments grade, G_E is the exam grade—a simple sum of the points achieved for correct answers—and S_Q is the score for in-class assignments.

$$G = \max(1, \min(10, 0.5 \times G_P + 0.3 \times G_A + 0.3 \times (G_E + S_Q/70))) \quad (1)$$

Note that our grading scheme clearly rewards the most important part of the course, the project, but it also provides some slack to enable students to excel in either the theoretical or the practical aspects of the course (i.e., exam or assignments, respectively), and rewards them for in-class participation through quizzes.

To calculate the project grade, we combine the grade for the project itself (i.e., the application of performance engineering methods and tools to the application at hand), G_P^p, the report, G_P^r, and the presentations (midterm and final, averaged), G_P^t, as follows:

$$G_P = 0.4 \times G_P^p + 0.3 \times G_P^r + 0.3 \times G_P^t \quad (2)$$

For the assignments, we use a points system that enables students to work in different groups. Each assignment has an allocated number of points (10p, 9p, 11p, 12p for assignments 1, 2, 3, and 4, respectively) and the final grade is calculated as the sum of the assignment points, G_A^i, $1 \leq i \leq 4$, divided by a factor that accounts for the number of students per team (we remind the reader the assignments and project are to be executed in teams of 1–4 students). Specifically, the calculation is described in Equation 3.

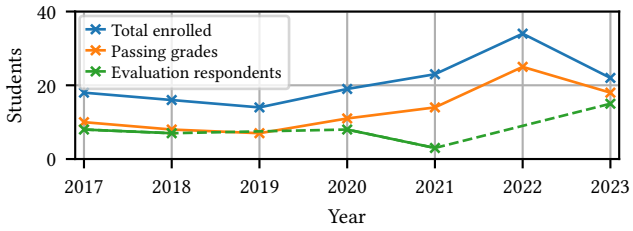

Figure 1: The number of students who enrolled in the course, and the number of course evaluation respondents. Note that the evaluation for the 2019 and 2022 courses are unavailable.

$$G_A = 10\frac{\sum_{1 \leq i \leq 4} G_A^i}{N} \quad \text{where} \quad N = \begin{cases} 32 & \text{for 1 student} \\ 36 & \text{for 2 students} \\ 40 & \text{for 3–4 students} \end{cases} \quad (3)$$

5 EVALUATION

In this section we reflect on the student performance and the feedback we have received for this course.

5.1 Students' Performance

The total number of students who have enrolled in the course over seven years is 146. Due to the high workload and advanced content, 15–50% drop out of the course; in total, 93 students have completed the course with a passing grade. In the following paragraphs, we present a few insights on the progress of the students in this course, based on their grades.

The average grade for the students passing the course is 8. This is higher than other courses due to the many different ways to achieve points. In fact, we find that students that finalize the course also get very good results, while those who would usually struggle fail and drop out early on. This is also a side-effect of the relative low impact of the final exam grade, which is often used as a discriminating factor between good and very good student performance, and rarely as a pass/fail discriminator.

The average assignments grade is around 8. Given the clear nature of the assignments, and the extensive support we provide in terms of tools, demonstration, and feedback, students provide very good solutions for the assignments. Traditionally, reports are of somewhat lower quality (often due to time pressure). To improve this aspect, we provided reporting templates. which have helped raising the average grade in the last couple of years.

The average exam grade for the students is around 7.5; this is due to the difficult nature of the exam, and the fact that we cover more theoretical aspects of performance engineering. These are also aspects that are not practiced extensively during the labs. We find that such exams are very good to discriminate individual talent/preparation within the teams which, otherwise, will receive the same grade. We observe that in-class quizzes clearly help with

good performance in the exam, as students often understand where and how to emphasize on specific topics. However, we do admit these quizzes may take a long time to create and grade, and more detailed feedback would help even further with the comprehension of the more theoretical aspects of the course.

For projects, the average grade is 8. This somewhat expected, given that we follow closely their progress, and we provide many feedback and consultation moments. The grades are not higher due to time pressure and, often, due to challenging priorities and project management within the team(s). However, all projects we have seen have been successful to a certain degree, with grades ranging from 6.5 (where parallelism was not correctly applied) to 10 (for work that we still use as example when we kick-off the project every year).

On a side note, we find it interesting to provide insight into the projects students have chosen to complete in this course. Recurring projects are, in the decreasing order of popularity: 2D stencil code optimization (due to an assignment from a previous course), game of life (also an assignment in a previous course), graph processing (due to one of the recurring invited lectures). However, we have also seen exotic applications, like content generation for games (RushHour), game optimization (EVE Online), solving Wordle, or FFT optimizations.

5.2 Students' Feedback

Overall, students grade the course between 8 and 9 (out of 10) every year, including during the pandemic. This is exceptionally high for courses in the Dutch academic system. They consistently grade the acquired knowledge, applicability, and practical relevance above 90%. Students appreciate the connection between the course, assignments, and project, and appreciate the ability to work on their own application. They further enjoy the final project presentations, where they not only get to brag about their own results, but also learn about new applications and new methods/tools to analyze and improve these applications.

On the downside, students are more critical of the high workload of the course: they claim to spend 20–50% more time than officially allocate for the course. To alleviate this problem, we created lab templates (to allow the reporting to be more efficient) and provided more flexible deadlines for the different assignments. We also further emphasized the importance of automation in their empirical analysis (from data collection to plotting), which seems to be the most time consuming aspect of the assignments and/or project.

Table 2: Aggregated student responses to evaluation questions. Data was taken in 2017, 2018, 2020, 2021, and 2023.

(a) Responses to questions on a scale from "Firmly Disagree" (corresponding to a numeric value 1) to "Firmly Agree" (corresponding to 5), where a higher score is considered better.

Statement	Firmly Disagree	Disagree	Neutral	Agree	Firmly Agree	*M*
"The course …"						
Taught me a lot	0	0	1	17	18	4.5
Was clearly structured	0	2	3	19	13	4.2
Was intellectually challenging	0	0	2	9	25	4.6
"I acquired, learned, or developed …"						
Factual knowledge	0	0	1	13	13	4.4
Fundamental principles	0	1	2	16	11	4.2
Current scientific theories	0	3	5	13	9	3.9
To apply subject matter	0	0	0	7	22	4.8
Professional skills	0	0	3	13	15	4.4
Technical skills	0	0	6	14	9	4.1
"… helped me understand the subject"						
Assignment 1	0	1	1	12	16	4.4
Assignment 2	0	0	1	11	16	4.5
Assignment 3	1	1	1	17	10	4.1
Assignment 4	0	1	1	12	13	4.4

*N*ote: In 2017-2018, assignments were evaluated with a single score; these scores have been duplicated across the four separate assignments in this table.

(b) Responses to questions on a scale from "Very Low" (numeric value 1) to "Very High" (numeric value 5), where a score between 3 and 4 is considered optimal.

Statement	Too Low	Low	Just Right	High	Too High	*M*
"The … of the course was"						
Workload	0	0	11	14	11	4.0
Level	0	1	16	13	6	3.7

6 LESSONS LEARNED

Before concluding this paper, we summarize our insights after these five years of teaching in a set of six lessons learned. They are:

Lesson 1 Performance Engineering is appealing when treated like a puzzle. We appeal to students' curiosity to understand why applications behave weirdly on different systems. To allow for this, one needs to prepare topics and assignments that showcase such behavior.

Lesson 2 Provide both methods and tools for each part of the course. Students appreciate the theory much better when they can link it to concrete examples. Furthermore, this mix enables them to build an intuition about how to navigate the performance engineering process and, in the long term, how to further improve it.

Lesson 3 Do not underestimate empirical analysis efforts. We often notice that students spend a lot of time in empirical analysis. This is often the case when experimental design is missing, and/or automation is not properly defined. We spend time and provide many examples on how this should be done, to allow them to build such automated setups correctly and efficiently.

Lesson 4 Projects stimulate creativity, and students should be allowed exploration time and space. As such, we provide no endline for our projects. Instead, we want them to try different things and report, after critical reflection, on their findings.

Lesson 5 Stimulate critical thinking to reporting on both positive and negative results. There are no negative results for our projects/assignments: we grade the process and the actual insights, and not the ultimate speed-up or high-accuracy models students

achieved. Rather, understanding why and how methods and tools work and fail is fundamental to this course.

Lesson 6 This is an intensive course for both teachers and students. Keeping the material up-to-date, adding new tools and infrastructure, and eliminating deprecated content are difficult when preparing the course. However, we are proud to offer something that students can rely on and immediately apply for their next performance engineering project in real-life.

7 CONCLUSION

As HPC and (super)computing systems increase in diversity, more effort is needed to keep improving the performance and efficiency of applications for newer systems. In turn, this means more specialists are needed to directly work with domain experts for HPC applications, but also to design and build better tools to facilitate this process. To create these specialists, we need our students to be aware, able, and willing to contribute to performance engineering. We strongly believe this can only happen if students are properly trained in the systematic nature of the performance engineering process.

To facilitate this education, we designed and implemented the first (and only) graduate-level course on performance engineering in The Netherlands. Our course combines lectures, assignments, and a project to combine theoretical, methodological, and practical aspects (and tools) to demonstrate and practice the full performance engineering process. After teaching the course for seven years, we can observe that such a course (1) is feasible for graduate-level students with basic knowledge of computer architecture and programming, (2) provides them with a deeper understanding and practical, applicable knowledge of efficient systems and application design, and (3) improves their ability to communicate with different stakeholders, including applications' end-users and system designers. Students' feedback indicates the course is challenging and interesting in its open approach to assignments and project. The course has been evaluated among the top courses in the program, and it has lead to 10–25% of the students in each year's cohort to be interested in taking on a performance engineering project for their MSc degree.

To further improve the course, we identified one main challenge to be tackled and three topics to be further developed. The main challenge facing this course is the continuous updating of the material: students appreciate it because it is up-to-date with the current tools and practices from the field, yet this often requires yearly efforts to test and update software, libraries, and tools. This can be extra-challenging when the course relies on shared machines, with very limited user rights. In terms of topics to be improved, we list the three critical ones: (1) supporting various vendors hardware, (2) including additional metrics—such as energy-efficiency—more prominently, and (3) expanding the distributed computing aspect of the course, potentially extending towards shared systems like cloud computing, the computing continuum, and the use of virtual machines or containers.

REFERENCES

[1] Henri Bal, Dick Epema, Cees de Laat, Rob van Nieuwpoort, John Romein, Frank Seinstra, Cees Snoek, and Harry Wijshoff. 2016. A Medium-Scale Distributed System for Computer Science Research: Infrastructure for the Long Term. *Computer* 49, 05 (5 2016), 54–63. https://doi.org/10.1109/MC.2016.127

[2] S. Browne, J. Dongarra, N. Garner, G. Ho, and P. Mucci. 2000. A Portable Programming Interface for Performance Evaluation on Modern Processors. *The International Journal of High Performance Computing Applications* 14, 3 (2000), 189–204. https://doi.org/10.1177/109434200001400303

[3] Agner Fog et al. 2011. Instruction tables: Lists of instruction latencies, throughputs and micro-operation breakdowns for Intel, AMD and VIA CPUs. *Copenhagen University College of Engineering* 93 (2011), 110.

[4] Markus Geimer, Felix Wolf, Brian J. N. Wylie, Erika Ábrahám, Daniel Becker, and Bernd Mohr. 2010. The Scalasca performance toolset architecture. *Concurrency and Computation: Practice and Experience* 22, 6 (2010), 702–719. https://doi.org/10.1002/cpe.1556

[5] Lizy Kurian John and Lieven Eeckhout (Eds.). 2006. *Performance Evaluation and Benchmarking (1st ed.).* CRC Press, Boca Raton, Florida, United States of America.

[6] Andreas Knüpfer, Holger Brunst, Jens Doleschal, Matthias Jurenz, Matthias Lieber, Holger Mickler, Matthias S. Müller, and Wolfgang E. Nagel. 2008. The Vampir Performance Analysis Tool-Set. In *Tools for High Performance Computing*. Springer Berlin Heidelberg, Berlin & Heidelberg, Germany, 139–155.

[7] Andreas Knüpfer, Christian Rössel, Dieter an Mey, Scott Biersdorff, Kai Diethelm, Dominic Eschweiler, Markus Geimer, Michael Gerndt, Daniel Lorenz, Allen Malony, Wolfgang E. Nagel, Yury Oleynik, Peter Philippen, Pavel Saviankou, Dirk Schmidl, Sameer Shende, Ronny Tschüter, Michael Wagner, Bert Wesarg, and Felix Wolf. 2012. Score-P: A Joint Performance Measurement Run-Time Infrastructure for Periscope,Scalasca, TAU, and Vampir. In *Tools for High Performance Computing 2011*. Springer Berlin Heidelberg, Berlin & Heidelberg, Germany, 79–91.

[8] Jan Laukemann, Julian Hammer, Johannes Hofmann, Georg Hager, and Gerhard Wellein. 2018. Automated Instruction Stream Throughput Prediction for Intel and AMD Microarchitectures. In *2018 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)* (Dallas, Texas, USA). IEEE, New York, New York, USA, 121–131. https://doi.org/10.1109/PMBS.2018.8641578

[9] John D. McCalpin. 1991-2007. *STREAM: Sustainable Memory Bandwidth in High Performance Computers.* Technical Report. University of Virginia, Charlottesville, Virginia. http://www.cs.virginia.edu/stream/ A continually updated technical report. http://www.cs.virginia.edu/stream/.

[10] John D. McCalpin. 1995. Memory Bandwidth and Machine Balance in Current High Performance Computers. *IEEE Computer Society Technical Committee on Computer Architecture (TCCA) Newsletter* 2, 19-25 (12 1995), 19–25.

[11] Johannes Seiferth, Christie Alappat, Matthias Korch, and Thomas Rauber. 2018. Applicability of the ECM Performance Model to Explicit ODE Methods on Current Multi-core Processors. In *High Performance Computing*, Rio Yokota, Michèle Weiland, David Keyes, and Carsten Trinitis (Eds.). Springer International Publishing, Cham, 163–183.

[12] Connie Smith. 2003. *Performance Evaluation – Stories and Perspectives - Chapter 16: Software performance engineering.* Austrian Computer Society, Vienna, Austria.

[13] Jan Treibig, Georg Hager, and Gerhard Wellein. 2010. LIKWID: A Lightweight Performance-Oriented Tool Suite for x86 Multicore Environments. In *2010 39th International Conference on Parallel Processing Workshops* (San Diego, California, United States of America). IEEE, New York, New York, United States of America, 207–216. https://doi.org/10.1109/ICPPW.2010.38

[14] Jan Treibig, Georg Hager, and Gerhard Wellein. 2012. Performance Patterns and Hardware Metrics on Modern Multicore Processors: Best Practices for Performance Engineering. In *Euro-Par Workshops*. Springer Berlin Heidelberg, Berlin & Heidelberg, Germany, 451–460. https://doi.org/10.1007/978-3-642-36949-0_50

[15] University of Amsterdam. 2022. *Course catalogue (2022-2023).* University of Amsterdam. https://studiegids.uva.nl/xmlpages/page/2022-2023-en/search-programme/programme/7283/251111

[16] Ana-Lucia Varbanescu and Stephen Nicholas Swatman. 2023. Performance Engineering Course repository. https://github.com/alvarbanescu/PerfEngCourse

[17] Samuel Williams, Andrew Waterman, and David Patterson. 2009. Roofline: An Insightful Visual Performance Model for Multicore Architectures. *Commun. ACM* 52, 4 (apr 2009), 65–76. https://doi.org/10.1145/1498765.1498785

[18] Henry Wong, Misel-Myrto Papadopoulou, Maryam Sadooghi-Alvandi, and Andreas Moshovos. 2010. Demystifying GPU microarchitecture through microbenchmarking. In *2010 IEEE International Symposium on Performance Analysis of Systems & Software (ISPASS)* (White Plains, New York, United States of America). IEEE, New York, New York, United States of America, 235–246. https://doi.org/10.1109/ISPASS.2010.5452013

A ARTIFACT INFORMATION

In accordance with the Supercomputing Transparency and Reproducibility Initiative, this appendix provides additional detail about the artifacts relevant for this manuscript.

A.1 Artifact Description

Our paper presents our experience with teaching performance engineering to graduate students. In order to allow maximal re-use of our materials and to provide transparency about the performance of the course, we have made all the lecture slides and assignments—both the documents and the framework code—available, as well as the quantitative data used in this paper and the software used to create figures and tables from that data. The remainder of this section details the different artifacts which we provide; a schematic overview of the artifacts is given in Figure2.

We provide three artifacts: the frameworks designed to provide support for students' assignments, and the scripts to be used to generate the figures and tables in this paper:

- SW-1 Assignment frameworks implemented in C which are used by students. Available in `assignments/code/`.
- SW-2 Python script to generate Figure1 with information about student counts. Uses DATA-1. Available as `scripts/make_plots.py`.
- SW-3 Python script to generate Table 2 with information about student evaluations about the quality of the course. Uses DATA-2. Available as `scripts/make_tables.py`.

Furthermore, we provide two (anonymized) data artifacts about the historical performance of the course:

- DATA-1 Structured historical data about the number of enrolled students, the number of students with a passing grade, and the number of evaluation respondents. Available as `data/students.csv`.
- DATA-2 Structured historical data about select questions answered by students in the evaluations of the course. Available as `data/metrics.csv`.

Finally, we provide two artifacts in the form of documents relating to the lectures and the assignments:

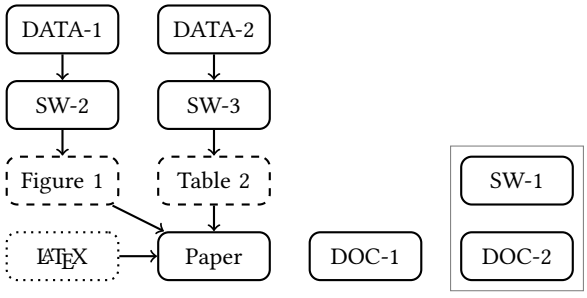

Figure 2: Dependency graph between the paper, its figures, and the associated artifacts. Artifacts with solid borders are provided as-is, artifacts with dashed borders can be deterministically reproduced using the provided artifacts, and artifacts with dotted borders can be provided upon request.

- DOC-1 Lecture slides as they were used in the 2022 edition of the course, available as PDF documents in `slides/`.
- DOC-2 Assignment documents as they were used in the 2022 edition of the course, available as PDF documents in `assignments/`.

A.2 Artifact Availability

All of the aformentioned documents, software, and data are available in our GitHub repository [16].

A.3 Artifact Evaluation

Our assignments require a heterogeneous system, using a CPU and NVIDIA GPU. We have used both Intel and AMD CPUs, although the tools we recommend are Intel-specific. We have use GPUs of compute capability between 3.0 and 7.2. We have not tried AMD GPUs, due to limited hardware access and available TA support.

The aforementioned software frameworks were not altered in any way for this paper: our code focuses solely on providing the seed the students need to use for their development. All the code can be compiled with off-the-shelf C/C++ compilers. All the provided code is developed by teachers and/or TAs, or uses open-source code available online (e.g., code for reading matrices in the matrix market format).