

Challenges and Experiences in Implementing a Specifications Grading System in an Upper-Division Undergraduate Computer Networks Course

Dr. Mahima Agumbe Suresh, San Jose State University

Mahima Agumbe Suresh is an Assistant Professor at San Jose State University. She received her Ph.D. from the Department of Computer Science and Engineering at Texas A&M University in December 2015, advised by Dr. Radu Stoleru. After her Ph.D., she was a postdoctoral researcher at Xerox Research Labs, India, where she worked on crime analytics and process mining. Her research interests include edge computing, machine learning, modeling and system design for cyber-physical systems and the Internet of Things. She has published in several peer-reviewed conferences and journals and has been a program committee member at several conferences.

Challenges and Experiences in Implementing a Specifications Grading System in an Upper Division Undergraduate Computer Networks Course

Abstract

Computer Networks is an important course in most undergraduate curricula in computing disciplines. The course learning objectives cover a broad range of topics and skills. The students are expected to acquire knowledge about the basic functionality of the layered architecture of the Internet, while also demonstrating a deeper level of understanding of each layer. A higher level of mastery can either indicate a deeper understanding of each layer or an in-depth understanding of one layer. Traditional point-based grading systems fail to capture mastery over concepts and skills in the course. Students also find it difficult to ascertain where they stand in the class. Specifications grading formalizes the skills and understanding of concepts that students have to demonstrate in order to achieve a particular final letter grade in the course. The system also supports revising and resubmitting student work, thereby inculcating the interest to learn rather than achieve a minimum number of points. This paper presents the author's experience in developing and implementing a specifications grading system for an upper-division undergraduate computer networks course over two semesters. First, the paper presents the design of a strong specifications grading system by mapping the learning outcomes map to assessment targets. Further, it will elaborate on how the assessments can be bundled to form specifications for the final letter grade. It will also elaborate on the teaching and grading load to handle revisions. This system was implemented over two semesters - Fall 2021 and Spring 2022. The author's learnings from each attempt and the changes made in the second iteration will also be presented. Overall, the grading system was accepted positively across two semesters, without increasing the grading load on the instructor.

Introduction

Computer Networks is a core undergraduate course in most curricula in the computing disciplines. ACM curricula recommendations for computing list computer networks as an element of computing knowledge^{1,2}. It is typically taught as an introductory course to the field of computer networks in the upper division. The course has been taught for decades and has grown to incorporate the advancements in the field, as evidenced by the textbooks published in the field^{3,4,5}. The textbooks and the experts in the field have an informal consensus on the course learning outcomes.

Grading is an essential tool to assess students' achievement of course learning objectives. Point-based systems have been known to have shortcomings in misplaced student motivation^{6,7}. Specifically, students have more motivation to accumulate points rather than meet the course learning outcomes. Specifications grading⁸ is a new grading tool that makes it possible to systematically shift the focus toward student learning and developing a growth mindset. Past examples of successful implementations in STEM and other fields motivate this work^{9,10,11,12,13}.

This paper presents the author's experience in designing and implementing specifications grading for an upper-division computer networks course, CMPE 148 at the San Jose State University. This course is the first introduction to computer networks for students in computer and software engineering majors. The paper will cover the design of learning outcomes, mapping them to assessment targets, and bundling assessments into the final letter grade. The paper will also cover the details of how to design assessments to accurately capture student learning while maintaining a manageable teaching load. This system was implemented in Fall 2021 and Spring 2022. Each iteration highlighted different strengths and shortcomings in the system. The paper will finally present student feedback on the grading system.

Specifications Grading Features and Their Implementation

Linda Nelson chucks out the key features of specifications grading in her book that popularized the system⁸. With guidance from the book and based on previous implementations of specifications grading, specifically in STEM fields^{10,9}, the author incorporated the following key tenets in the grading system.

1. Pass/fail grades in place of points with clear specifications
2. Revision and retakes for most assessments
3. Bundles to determine final letter grade - higher grades imply deeper or broader understanding of concepts
4. Linking grades to learning outcomes

As mentioned before, Computer Networks is a well-established course with a clear course structures and syllabi. The grading system implementation had a course syllabus with the traditional point-based system with weighted averages for letter grades. So, this process describes a redesign of the grading system, rather than a fresh design from scratch.

First, let us discuss the implementation of **pass/fail grades in place of points with clear specifications** and **revision and retakes for most assessments**. This is best discussed for each type of assessment. The course used four types of assessment: quizzes, homeworks, exams, and project.

Quizzes are meant to assess the student's understanding of the subject matter. Quizzes were already part of the traditional grading system as a way to ensure that students kept up with the pace of the course. However, it may not be in the best interest of students and might have been contributing to equity gaps in the course. For example, students who worked full-time or were caregivers had informally expressed that they would have liked more flexibility in taking quizzes.

With learning management systems (LMS) such as Canvas offering automated grading, quizzes can be administered outside the classroom. The course used multiple choice question banks and formula questions in the LMS and created quizzes that can be *repeated any number of times* until the student *passes*. If students answer 80% of the questions correctly, they pass the quiz. 80% of the questions tested students' understanding of the core concepts and 20% of the questions were deeper and more advanced. The course consisted of seven modules - Networking Overview, Application Layer, Transport Layer, Network Layer Data Plane, Network Layer Control Plane, Link and Physical Layer, and Other Topics (includes Network Security). There are three categories of quizzes as follows:

1. Required Quizzes (RQ) – A syllabus quiz and a prerequisite knowledge quiz.
2. Core Quizzes (CQ) – Quizzes that test core concepts in Computer Networks. Students who pass this course will be expected to have the knowledge tested in these quizzes
3. Advanced Quizzes (AQ) – Quizzes that test advanced concepts in Computer Networks. Students who pass these quizzes will have achieved advanced knowledge and skills.

Core quizzes are released upon completion of the modules they cover. In a 16-week semester, they were released approximately every three weeks (weeks 3, 6, 9, 12, and 15). Advanced quizzes were released at around week 14.

Homework assignments consisted of both theory and networking tools components. Theory components tested the understanding of the subject matter by applying concepts taught in class to solve numerical problems, write proofs, and reflect on ideas. The networking tools components were meant to provide the students with hands-on experience and familiarity with networking tools. Homework assignments were graded as “high pass”, “low pass”, “revision needed”, or “fail”. The homeworks consist of questions labeled “low pass” (LP) and “high pass” (HP). Students must pass all the LP questions to receive a “low pass” in the homework. Students must pass all the “low pass” questions and a fraction of the HP questions as specified in advance to receive a “high pass” in the homework. If a student receives a “low pass” or “revision needed” grade, the student may revise and resubmit their homework assignment. At most two resubmissions for each homework, within a stipulated deadline, were allowed. The LP questions test core knowledge in networks, whereas HP questions are either more in-depth or on advanced topics. For example, in the homework on Network Address Translation (NAT), the LP question asked students to figure out how to fill a NAT table given a scenario. The HP question asked students to perform a lab activity using the Wireshark tool to capture NAT packets. Homeworks are released upon completion of the modules they cover. In a 16-week semester, they were released approximately every four weeks (weeks 3, 7, 11, and 15).

Examinations are conducted twice in the semester. Real-world situations often do not have opportunities to revise one's work and have to be achieved in one shot after preparation, e.g., interviews. In this course, one midterm examination was conducted, with an option to retake the exam once. One final examination was conducted with no retake opportunities. Study guides were provided for each exam and a class lecture session was dedicated to reviewing materials prior to the exams. The questions were designed such that 70% of the questions test core networking knowledge that anyone who passes the course should be able to answer, 15% of the questions require proficiency in core topics, 10% of the questions require critical thinking, and the

last 5% of the questions show excellence.

Project was a way for students to demonstrate deeper understanding of concepts. A project was not necessary to pass the course. Students picked topics to study and summarize a specific topic in Computer Networks through a class report. Projects are graded as “high pass”, “low pass”, or “not applicable”. The final report cannot be revised. However, several project milestones are included to provide feedback on their work. This mimics the workplace setting where employers have code reviews and other avenues to receive feedback on their work before a final product release. In a 16-week semester, the students were required to submit a project proposal at around week 4, followed by project check-in milestones at weeks 8, 11, 14, and 16. The final project deliverable was due at the end of the finals weeks.

This course uses **tokens** to provide students flexibility in completing the assessments. Students start the semester with 2 tokens in their wallet. Additional tokens were earned by attending office hours or posting meaningful questions or answering others’ questions on the class discussion forum.

Tokens could be exchanged for the following purposes:

- One token in exchange for an opportunity to revise and resubmit one homework assignment. At most two resubmissions are allowed for each homework assignment
- Two tokens in exchange for an opportunity to retake the midterm exam on the specified date
- One token in exchange for a one-day deadline extension for homework assignments or project milestones. There will be a limit to the number of one-day extensions, which will be mentioned in the corresponding assessment. Unless otherwise mentioned, deadlines will be strictly enforced.

For homework assignments, if the student fails to submit their assignment by the posted deadline, their submission will receive a grade of “revision needed”. If they fail to submit the assignment by the revision deadline, the submission will receive a grade of “fail”.

For project milestones, if the student fails to submit their assignment by the posted deadline, their submission will receive a grade of “incomplete”. However, only the final submission before the final deadline will be considered in determining the letter grade. Students are encouraged to use the project milestone deadlines to receive feedback on their reports and/or demo/videos.

Next, let’s discuss how these assignments are “bundled” into a letter grade. The final letter grade was determined using the specifications below. Note that all components of a grade must be met to achieve it. The table below shows the specifications for Spring 2022, incorporating feedback from Fall 2021. Students must meet all the requirements for a particular grade to earn it.

Grade	Quizzes	Homework	Exam Average	Individual Project
A plus	Pass RQ + 5 CQ + 1 AQ	4 High Pass*	>= 95%	3 HP
A	Pass RQ + 5 CQ + 1 AQ	4 High Pass	>= 90%	2 HP + 1 LP
A minus	Pass RQ + 5 CQ + 1 AQ	3 High Pass, 1 Low Pass	>= 90%	1 HP + 2 LP
B plus	Pass RQ + 5 CQ + 1 AQ	2 High Pass, 2 Low Pass	>= 85%	2 HP
B	Pass RQ + 5 CQ + 1 AQ	2 High Pass, 2 Low Pass	>= 80%	1 HP + 1 LP
B minus	Pass RQ + 5 CQ + 1 AQ	1 High Pass, 3 Low Pass	>= 80%	1 HP
C plus	Pass RQ + 5 CQ	4 Low Pass	>= 75%	Not required
C	Pass RQ + 5 CQ	4 Low Pass	>= 70%	Not required
C minus	Pass RQ + 5 CQ	3 Low Pass	>= 70%	Not required
D plus	Pass RQ + 4 CQ	3 Low Pass	>= 65%	Not required
D	Pass RQ + 4 CQ	2 Low Pass	>= 60%	Not required
D minus	Pass RQ + 4 CQ	2 Low Pass	>= 60%	Not required
F	Did not meet specifications for D minus			

* - Student passed all high pass questions in all 4 homeworks.

Finally, let's discuss how we **link grades to learning outcomes**. The first step to designing a grading system that captures students' mastery over course learning outcomes is to design measurable learning outcomes that are at the right granularity. Course learning outcomes are often 4-8 measurable statements that capture what students are capable of when they complete the course. Learning outcomes for specifications grading break down these measurable statements into finer granularity. Although computer networks is a vast field, the learning objectives can be broken down as follows. "Core" outcomes are the outcomes that students who pass the course demonstrate having gained mastery over, and "advanced" outcomes indicate that students have achieved higher grades. "Advanced" outcomes are marked with a ★ symbol. With each outcome, the corresponding assessments are also listed.

Discuss and illustrate the computer network architecture

- Discuss the 7-layer of the Open Systems Interconnection (OSI) model - Quiz 1, Exam
- Discuss each TCP/IP architecture layer functionality - Quiz 1, Exam
- Map a TCP/IP protocol to the layer it implements - Quiz 1, Exam
- ★ Describe cross-layer architectures - Project
- ★ Describe cellular architectures - Project

Use networking commands and tools

- Run and understand the output of basic network commands (e.g., ping, ifconfig, traceroute, nslookup, telnet) - Homework 2, 3, 4
- Able to perform a packet capture using a tool such as Wireshark - All homeworks
- ★ Analyze basic components of a packet capture - All homeworks
- ★ Advanced analysis of packets captured via tools - Homework 4

- ★ Run and understand advanced network commands such as ip, tcpdump - Project

Explaining the functionality of the application layer

- Discuss the general functionality of the application layer - Quiz 2, Homework 2, Exam
- Explain the most basic application layer protocols - HTTP and DNS - Quiz 2, Homework 2, Exam
- ★ Explain application layer standards - SMTP, P2P, CDN - Project
- Able to write basic socket programs - Homework 2
- ★ Able to write advanced socket programs - Homework 2 and Project

Explaining the functionality of the transport layer

- Discuss the general functionality of the transport layer including multiplexing and demultiplexing in TCP and UDP - Quiz 3, Homework 3, Exam
- Explain the TCP congestion control - Quiz 3, Homework 3, Exam
- Explain TCP connection establishment - Quiz 3, Homework 3, Exam
- ★ Compare and contrast TCP variants - Project
- Explain the segment structure and functionality of UDP - - Quiz 3, Homework 3, Exam

Explaining the functionality of the network layer

- Discuss the basic data plane functionality of a router - Quiz 4, Homework 4, Exam
- Discuss data plane addressing basics including IPv4, NAT - Quiz 4, Homework 4, Exam
- ★ Discuss data plane advanced topics including IPv6, SDN, NDN - Project
- Explain control plane routing algorithms including BGP, OSPF - Project, Advanced Quiz
- ★ Explain the specific details of control plane routing algorithms - Project
- ★ Explain control plane routing algorithms in special settings - Project
- ★ Describe the SDN control plane - Project
- ★ Discuss network management protocols and ICMP - Project, Advanced Quiz

Explaining the functionality of the link layer

- Discuss framing and the concept of Medium Access Control (MAC) protocols - Quiz 5, Advanced Quiz, Exam
- ★ Discuss MAC protocol details such as CSMA-CD, CSMA-CD - Quiz 5, Advanced Quiz, Exam
- Outline link Layer addressing and Address Resolution Protocol - Quiz 5, Advanced Quiz, Exam
- ★ Discuss switched LANs - Project

- ★ Describe the wireless link and WiFi - Project
- ★ Describe Ethernet - Project
- ★ Discuss wireless LANs - mobility, advanced features, low power protocols - Project

Other outcomes

- ★ Explain datacenter networking - Project
- ★ Explain physical layer theoretical foundations - Project
- ★ Understand the propagation loss model on the wireless link - Project
- Demonstrate basic mathematics skills including Boolean algebra and Fourier analysis - All assessments
- Communicate effectively with fellow engineers - Exam and project

Results of the Implementation

Evaluating the success of the grading system is fairly subjective. The paper will now report on the following metrics of evaluation:

1. Teaching Load
2. Load on Teaching Assistant/Grader
3. Student Perception
4. Grade Distribution

Teaching Load: In comparison with the traditional grading system, the specifications grading systems carry an overhead in two main areas - assessments and office hours. Managing the grading system in traditional tools such as Excel was challenging. Therefore, the author incorporated features in the learning management system and wrote Python scripts to assist in grading. Most learning management systems include grading schemes that can be customized. The author leveraged these features extensively.

Quizzes: Setting up the quizzes require careful planning and working with the learning management system to set up question banks and formula questions. Compared to the traditional grading with no reattempts, this system requires the instructor to create significantly more questions. However, once the quizzes are released, there is no need to grade them. Feedback can be incorporated into the learning management system. Formula questions can be hard to create if we have follow-up questions, e.g., asking students to figure out the network prefix using a subnet mask and then asking for the number of hosts in the subnet as a follow-up. For such questions, the author set them up as fill-in-multiple-blank questions in a question bank.

Homeworks: The bulk of the work in redesigning homeworks is to map questions to learning outcomes and determine if they are LP questions or HP questions. In addition, there is additional burden to handle revisions and re-submissions.

Exams: The added load was to grade the exams in time to allow for a retake opportunity and design two midterm exams in place of one.

Projects: The workload increase was mostly in the added time to provide meaningful and actionable feedback in the project milestones.

Office hours were a lot busier than usual because of the token system. It was in the best interest of the students. It was a welcome change to see students attempt to understand concepts rather than haggle for extra points in their assessments.




Figure 1: Graph that captures the impact of Specifications Grading with the questions: “The grading system (Chart 1) (i) Increased my academic performance in this course; (ii) Increased my engagement in the class; (iii) Increased my confidence taking this course; (iv) Increased my curiosity for the subject matter (Chart 2) (i) Increased my ability to set goals and devise a plan to achieve them; (ii) Increased my ability to enjoy challenging assignments; (iii) Increased my ability to learn from my mistakes; (iv) Increased my ability to use feedback to improve my learning (Chart 3) (i) Increased my ability to overcome challenges and achieve my goals; (ii) Decreased the level of stress in the course; (iii) Helped me develop a growth mindset”

Load on Teaching Assistant/Grader: In comparison to the traditional grading system, the load on the grader was the same because of the automatically graded quizzes. The graders’ time was redirected to provide meaningful feedback to the students in the revisions and re-submissions.

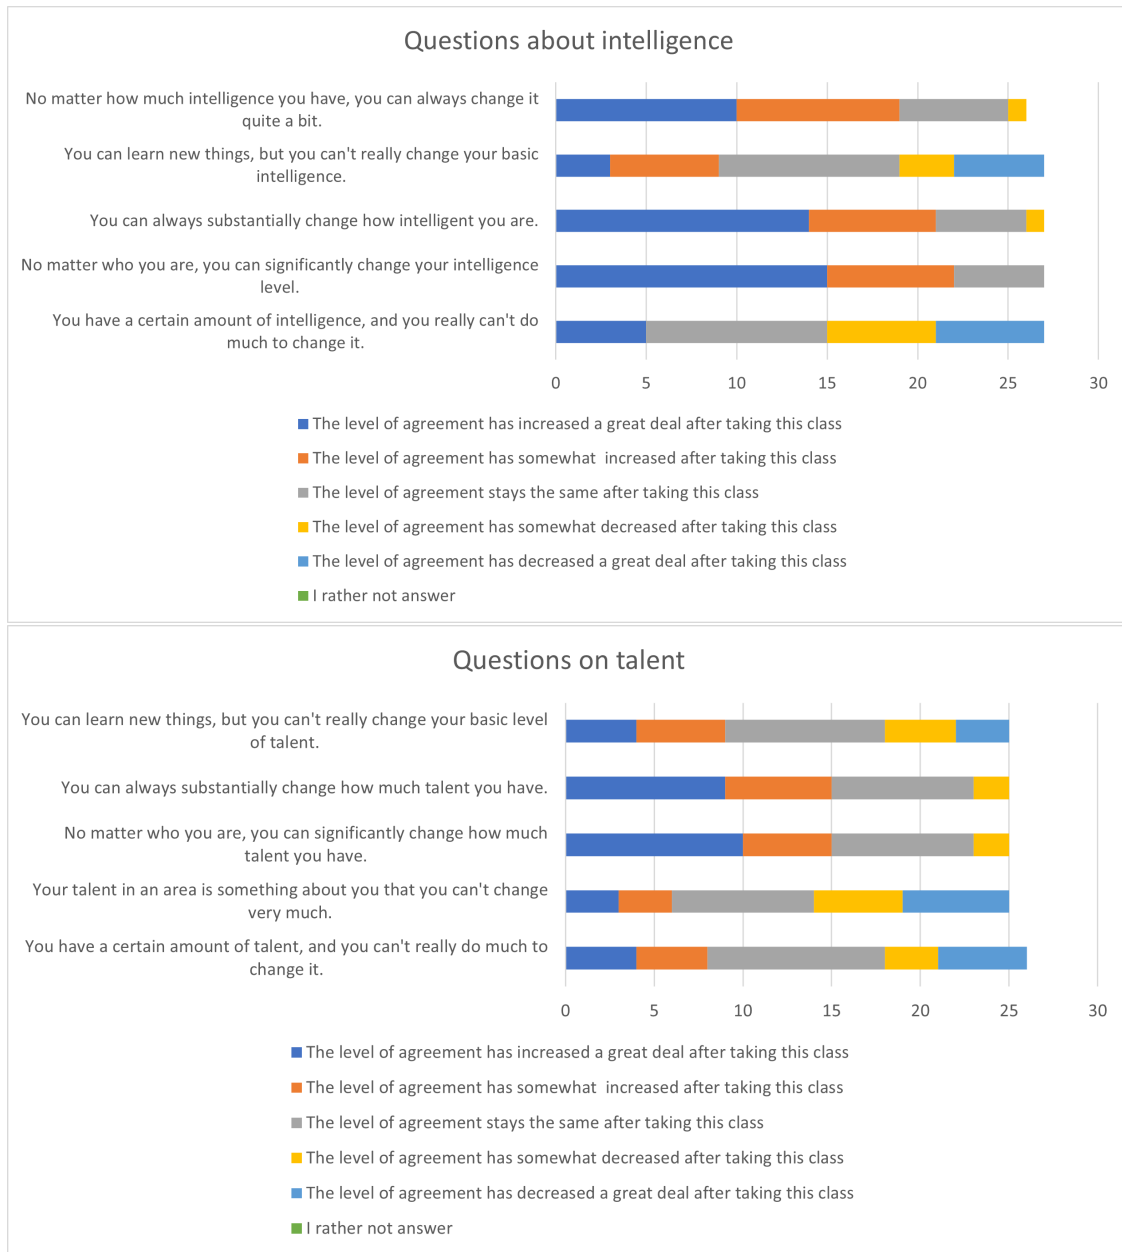


Figure 2: Graph that captures the impact of Specifications Grading with the questions on the “growth mindset”

Student Perception Overall, the students provided positive feedback at the end of the semester. The biggest win for me was that no student complained that their learning experience was negatively impacted by the system. Most students – even informally – shared with me that the system helped them focus on the learning.

To quote a student’s feedback: *“The specification grade setup was something new and something I was initially skeptical about, but I believe it was a good choice in the end. It promoted learning the material and gave everyone a chance to get the grade they desired.”* The quote is taken from a

survey that was sent in Spring 2022 for students to submit voluntary anonymous feedback on the grading system. The questions focused on the outcome of the implementation and the growth mindset^{9,8}. Students gave positive responses as depicted in Figure 1 and 2. Figure 1 indicates that most students strongly agreed that the grading system had a positive impact on their overall learning experience in the class. Figure 2 shows that students were developing a growth mindset at the end of the class.

Grade Distribution




Figure 3: Graphs that captures the grade distribution. If the grading system would take the best three categories among quizzes, homeworks, exams, and projects, the grade change would not have been significantly different. Each grade level includes + and - level grades. For example, B- is one grade level above a C+. Finally, the last graph show the category in which students received their lowest grade, which was the category that determined their final grade. Students whose grade would not change because of one category

Figure 3 shows the grade distribution for the course in the Fall 2021 and Spring 2022 semesters. In Fall 2021, student grades show inflation. One of the main reasons for this is that the HP questions in the homeworks did not account for the fact that students would get opportunities to revise and resubmit. They needed to be more challenging. In addition, no student received an A+ because they were constrained by their group project.

In Spring 2022, the projects were made individual and the HP questions were made more challenging. However, students complained that the system was too unforgiving. They recommended that the system take the best three out of four categories - quizzes, homeworks,

exams, and project - to determine the final grade. In terms of mastery, this would mean that it would suffice for students to demonstrate a lower level of mastery in advanced topics to achieve a higher grade. Therefore, chart 2 in figure 3 present an analysis of the same. If one category was eliminated, some students could receive 4 grade levels higher. For example, a student receiving B- could have received an A. In order to understand this better, the final chart in figure 3 presents the category where a student received the lowest grade. Students were mostly constrained by their grade in their projects, which is one of the only ways to show mastery over advanced topics in the course. Overall, this shows that for higher grades, the system is able to capture mastery over the learning outcomes listed before. However, for grade levels lower than C+, the best three out of four categories system can still be applied without diluting it.




Figure 4: Graphs that captures the equity gaps in the grading system over four categories: (i) Underrepresented minorities, (ii) First generation, (iii) economically challenged, (iv) Gender

Equity is an important issue to be considered while experimenting with new grading systems. Figure 4 depicts the equity gaps in terms of percentage of failing students and overall GPA. The equity gap in failure rate in my classes overall, including courses taught using the traditional grading system is -10.7%, and the GPA gap is 0.72. In the computer networks course with specifications grading, these numbers were -4.6% and 0.39. The course did not show a significant cause for concern in any of the categories.

Learnings from the Experience

In Spring 2022, the author experimented with a few changes to incorporate feedback from the Fall 2021 experience. The key issue with the implementation in Fall 2021 was that grading was slow. This was partly due to issues with finding grading assistance on time. The author ended up doing a major portion of the grading for the first half of the semester. The author also planned for five homeworks, which was challenging to fit into the schedule. Additionally, group projects did not fit well into specifications grading while bundling individual and group assessments. In Spring 2022, the number of homeworks was reduced to four instead of five and rehired the same grader to avoid any delays with onboarding. The project was changed to be individual and the specifications were made significantly more clear. In a different course, the author also experimented with including individual and group components in one grading bundle with

specifications grading. For the group assessments, the system included an individual contributions component, which would impact the final grade for a student.

Overall, the grading load was not higher than the traditional grading system. The author had fewer conflicts with students and little to no discussions about partial credit. It was a positive experience and the author would recommend this system to other Computer Networks course instructors.

References

- [1] Alison Clear and Allen S Parrish. Computing curricula 2020. *Computing Curricula*, 2020.
- [2] Association for Computing Machinery (ACM) Joint Task Force on Computing Curricula and IEEE Computer Society. *Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science*. Association for Computing Machinery, New York, NY, USA, 2013. ISBN 9781450323093.
- [3] James F Kurose and Keith W Ross. Computer networking: A top-down approach. edition. *Addison Wesley*, 2007.
- [4] AS Tanenbaum. Computer networks 2nd edition prentice-hall. *Inc., Englewood Cliffs, NJ*, 1988.
- [5] David J Wetherall and Andrew S Tanenbaum. *Computer networks*. Pearson Education, 2013.
- [6] Thomas R Guskey and Jane M Bailey. *Developing grading and reporting systems for student learning*. Corwin Press, 2001.
- [7] Mark Carl Rom. Grading more accurately. *Journal of Political Science Education*, 7(2):208–223, 2011.
- [8] Linda B Nilson. *Specifications grading: Restoring rigor, motivating students, and saving faculty time*. Stylus Publishing, LLC, 2015.
- [9] Kate J McKnelly, William J Howitz, Taylor A Thane, and Renée D Link. Specifications grading at scale: Improved letter grades and grading-related interactions in a course with over 1,000 students. 2022.
- [10] William J. Howitz, Kate J. McKnelly, and Renée D. Link. Developing and implementing a specifications grading system in an organic chemistry laboratory course. *Journal of Chemical Education*, 98(2):385–394, 2021. doi: 10.1021/acs.jchemed.0c00450.
- [11] Dennis Earl. Two years of specifications grading in philosophy. *Teaching Philosophy*, 45(1):23–64, 2022.
- [12] Ella Tuson and Tim Hickey. Mastery learning and specs grading in discrete math. In *Proceedings of the 27th ACM Conference on on Innovation and Technology in Computer Science Education Vol. 1, ITiCSE '22*, page 19–25, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392013. doi: 10.1145/3502718.3524766. URL <https://doi.org/10.1145/3502718.3524766>.
- [13] Brian P. Helmke. Specifications grading in an upper-level bme elective course. In *2019 ASEE Annual Conference & Exposition*, number 10.18260/1-2–33278, Tampa, Florida, June 2019. ASEE Conferences. <https://peer.asee.org/33278>.