Systems, Networking, and Cybersecurity
Spring, 2017

Distributed: January 9, 2017 (11:59PM)
Due: January 23, 2017 (11:59PM)

Honor Code. This examination is conducted under the University's Graduate Honor System Code. Students are encouraged to draw from other papers than those listed in the exam to the extent that this strengthens their arguments. However, the answers submitted must represent the sole and complete work of the student submitting the answers. Material substantially derived from other works, whether published in print or found on the web, must be explicitly and fully cited. Note that your grade will be more strongly influenced by arguments you make rather than arguments you quote or cite.

Written answers. The answers to the questions on this exam must be submitted no later than the due date listed above. Answers must be submitted in a single PDF document emailed to the exam coordinator (Changhee Jung, chjung@vt.edu).

Oral Exam. The written exam will be followed by an oral exam, where the student is expected to defend his/her solutions. Unless specifically requested, the student is not expected to make a formal presentation. In the oral exams, faculty may ask questions about any paper in the reading list to assess the student's understanding of the subject. Oral exams will be scheduled individually for each student.

Assessment. After the oral examination, the examining faculty will determine the student's score for the examination process. The score is between 0 – 3 points, depending on the student's performance on both the written and oral components. These points may be applied toward the total score of 6 points necessary to qualify for the Ph.D. The assessment criteria, as defined by GPC, are as follows:

- 3: Excellent performance, beyond that normally expected or required for a PhD student.
- 2: Performance appropriate for PhD-level work. Prime factors for assessment include being able to distinguish good work from poor work, and explain why; being able to synthesize the body of work into an assessment of the state-of-the-art on a problem (as indicated by the collection of papers); being able to identify open problems and suggest future work.
- 1: While the student adequately understands the content of the work, the student is deficient in one or more of the factors listed for assessment under score value of 2. A score of 1 is the minimum necessary for an MS-level pass.
- 0: Student's performance is such that the committee considers the student unable to do PhD-level work in Computer Science.

Systems, Networking, and Cybersecurity
Spring, 2017

Questions on the paper "Validity of the single-processor approach to achieving large scale computing capabilities", AFIPS'67

Give a brief overview of the core ideas of this paper. Describe whether you view the conclusions of the paper as optimistic or pessimistic for the future of parallel and distributed computing. Describe examples of the impact of this work following its publication. Describe the meaning of this work in the context of today's systems. Are the conclusions still valid? Cite specific other papers that support your arguments as well as papers that built on this work and how they impacted the field.

Systems, Networking, and Cybersecurity
Spring, 2017

Questions on the paper “A bridging model for parallel computation", Communications of the ACM, vol. 33, issue 8

Give a brief overview of the core ideas of this paper. Describe whether you view the conclusions of the paper as optimistic or pessimistic for the future of parallel and distributed computing. Describe examples of the impact of this work following its publication. Describe the meaning of this work in the context of today’s systems. Are the conclusions still valid? Cite specific other papers that support your arguments as well as papers that built on this work and how they impacted the field.

Systems, Networking, and Cybersecurity
Spring, 2017

Questions on the paper "Power-aware speedup", IPDPS'07

1. Give a brief overview of the core ideas of this paper. Describe whether you view the conclusions of the paper as optimistic or pessimistic for the future of parallel and distributed computing. Describe examples of the impact of this work following its publication. Describe the meaning of this work in the context of today's systems. Are the conclusions still valid? Cite specific other papers that support your arguments as well as papers that built on this work and how they impacted the field.

2. Compare and contrast the Power-aware Speedup model with the first two papers by Amdahl and Valiant. How are they related and what are the key differences between these models? Cite examples and any other references to support your claims.

Systems, Networking, and Cybersecurity
Spring, 2017

Questions on the paper "Early Detection of Configuration Errors to Reduce Failure Damage"

1. The Table 1 and 2 show the severity and diagnosis time of latent versus non-latent errors among the customers' configuration issues of COMP-A. How can we determine the severity level and diagnosis time? Provide a reasonable mythology.

2. Do you agree with the methodology, findings, and implications of the study in Section 2 that analyzed the root causes and characteristics of LC errors? Criticize their methodology. Would you do anything differently? Anything else you might want to know through this open-source application study?

3. PCheck supports running the generated checkers periodically in a separate thread to capture the TOCTTOU errors. What is TOCTTOU errors and why does it matter in the context of PCheck?

4. Discuss soundness and completeness of the PCheck, i.e., sources of false negatives and false positives. Discuss potential solutions to address them.

Systems, Networking, and Cybersecurity
Spring, 2017

Questions on the paper "Shuffler: Fast and Deployable Continuous Code Re-Randomization"

1. Discuss why existing address space randomization does not work to prevent ROP attacks, i.e., why do we need Shuffler?

2. Criticize the threat model of the paper. Describe if each of four assumptions in Section 2.1 is reasonable.

3. Explain why and how Suffer encrypt return addresses.

4. Criticize their security analysis in Section 6. Are you satisfied with their methodology and evaluation? Why do you think so? Do you think differently? If so, what should they have been done differently?

Systems, Networking, and Cybersecurity
Spring, 2017

Questions on the paper "Atlas: leveraging locks for non-volatile memory consistency", OOPSLA'14

1. Give a brief overview of the core ideas of this paper. Describe whether you view the conclusions of the paper as optimistic or pessimistic for the current (e.g., FRAM) and future (e.g., MRAM) of persistent memory.

2. Criticize Atlas' failure recovery model and explain how Atlas can re-execute the failure-atomic section (FASE) that encounters an application crash due to power failure. In particular, you should explain how Atlas deals with the inputs (memory and register) to the FASE for recovery.

3. What is the main reason for building the dynamic happens-before graph (HBG)? Why do we even need it? If you do not want to build the happens-before graph, e.g., due to either the lack of the idle core to run the helper thread or the runtime overhead of managing the graph, how would you change the original semantic of the FASE to still maintain the crash consistency?

Systems, Networking, and Cybersecurity
Spring, 2017

Questions on the paper "Daino: a high-level framework for parallel and efficient AMR on GPUs", SC'16

1. Give a brief overview of the core ideas of this paper and criticize the work; you need to provide at least 3 positive and 3 negative points of the work.

2. Explain why the proposed framework imposes the restriction on pointer aliasing? Similarly, explain why the framework cannot efficiently transform the applications that use irregular data access patterns at the block level?

3. The proposed performance model achieves only 78%~%85 of the theoretical AMR speedup for the benchmarks evaluated. What are the main reasons for the imprecision of the performance model? What would you do to improve the accuracy of the performance model?