

Homework #2 (CMSC 396H, Spring 2016)

Due February 10, 9:00AM

1 Overview

The goal of this assignment is to gain experience with critically reading, analyzing, and writing about technical papers. We'll do this by all reading and discussing the same paper:

“Vanish: Increasing Data Privacy with Self-Destructing Data”, R. Geambasu, T. Kohno, A.A. Levy, H.M. Levy, In *USENIX Security*, 2009.

In class on February 10, we will have an interactive discussion about this paper, wherein we will discuss and debate the merits of the paper.

Writing Task 1: Paper review. For future papers, you will be writing a concise synopsis of and insight into the work, but for this assignment, you will be doing a slightly more thorough review of the paper, to get into the habit of thinking critically about the strengths and weaknesses of a paper. To this end, you will be writing your review of the “Vanish” paper in the template of what is common for *peer-reviewed* conferences and journals:

- **Paper Summary** (roughly one paragraph in length): A neutral description of the paper. Some common things to include in the summary: what problem does the paper seek to solve, how does it try to do so, what are some of the techniques it uses to evaluate or build the solution, and what are some of the main take-away results.
- **Paper Strengths** (itemized list of about one sentence each; details come later): What you think the contributions are; what you think the paper does right/well, or what you found interesting. Was it a good problem, a good solution, an evaluation rooted in a realistic setting, was the paper well-written, etc.
- **Paper Weaknesses** (itemized list of about one sentence each; details come later): What do you think the paper did *not* do well.
- **Detailed Comments** (at a minimum, one short paragraph for each of the paper weaknesses): This is the part of the review that provides the rationale for behind each of the items you listed as weaknesses of the paper: Why did you disagree with the problem, the solution, the results, how it wrote about related work, etc. You can also add extra detail about why you *liked* aspects of the paper. Essentially, this is where a review presents its case for why the paper should be accepted for publication (or not). What could the authors do to improve the paper to better meet your expectations (e.g., run another experiment, prove a property about some aspect of the system, etc.)?

There are a couple differences between this review and those in the peer-review process; primarily, when evaluating whether or not to accept a paper into a conference or journal, each review would also include a numeric score, ranging from “definitely reject” to “definitely accept”. Also, in the peer-review process, reviews would be delivered to the paper’s authors: I will not be sending the authors your reviews!

2 Paper Strengths and Weaknesses: What makes research good?

There are many ways you can reason about the strengths and weaknesses of a paper. In fact, the target can move over time: the demand for larger datasets or more true-to-life simulations often grows as a particular field matures.

That said, there are a few aspects of a paper that are consistent, reasonable measures for the paper's quality. These include: is the paper novel? Does it compare to related work thoroughly and accurately? Are there any technical flaws? Are there potential flaws to the way an experiment was run or analyzed that lead you to question the paper's conclusions?

This paper maps out various ways you may consider evaluating a paper, and describes what the steps before, during, and after a paper review are (when reviewing for a conference or journal):

“Writing reviews for systems conferences”, T. Roscoe, March 2007.

Cultivating your own view of “good research”. Read Roscoe's paper, and let it help inform how you approach your review of the “Vanish” paper. But as you grow as researchers, you will not only learn how your research community values and evaluates a paper's quality—you will also cultivate your own perspectives and beliefs as to what makes a paper good (or not).

3 Submitting

You may simply post your write-ups as *private* Piazza posts to me.

These are due by 9AM the morning of our next class (Wednesday February 10), so that I can have time to read through all of them and discuss them during class.