

Problem Set 8

Crossing Abstractions

Out: 24 April

Due: Monday, 1 May (or later, see below for details)

Due Date

After we grade Exam 2, you will receive an email assessing your current performance in the class. It will include a recommendation for what you should do for Problem Set 8 and (possibly) the Final Exam. The recommendation might be one of these:

- **Excellent performance** — you have done enough to convince me you have mastery of the course objectives and will receive at least an A- in CS216. If you achieve at least Green Star performance on PS8 you will receive at least an A in CS216. You do not need to take the final. You may turn in PS8 any day until Monday, May 8 at 5pm. If you do choose to take the final, your performance on it could convince me you deserve higher than an A-, or that I was mistaken in my earlier judgement and you deserve a lower grade.
- **Good performance** — you have done well enough to convince me you deserve at least a B- in CS216. If you achieve at least Green Star performance on PS8 you will receive at least an B in CS216. You may choose to take the final, and if you perform well on the final you could still earn an A in the course. If you choose to take the final, you must turn in PS8 on Monday, May 1. If you choose not to take the final, you may turn in PS8 any day until Monday, May 8 at 5pm. If you do choose to take the final, as long as you achieve at least Green Star performance on PS8 you will not receive lower than a B in the class. (That is, there is no risk that taking the final will lower your grade).
- **Unconvincing performance** — you have not done enough yet to convince me you deserve a B- in CS216. You should work alone on PS8, and turn in PS8 on Monday, May 1 and take the final. You can still receive a high grade in CS216, but will need to demonstrate good understanding of the course material on PS8 and the final.

Collaboration Policy - Read *Very Carefully*

If you receive the excellent or good performance assessment, you may work alone or with one, two, or three other people of your choice on PS8. If you receive the unconvincing performance assessment, you should work alone on PS8.

You may consult any outside resources including books, papers, web sites and people you wish. You are also encouraged to discuss these problems with students in the class. You must acknowledge any people and outside resources you work with on your assignment. If you discuss the assignment with people other than your partner, you may not take any written materials out of your discussion. It is fine to bounce ideas off other people, but the answers you turn in must be your own.

You are **strongly encouraged** to take advantage of the staffed lab hours posted on the CS216 web site.

Purpose

- Demonstrate what you have learned in CS216
- Produce an interesting and useful program that involves crossing abstract layers and using several of the tools and techniques you learned in CS216.

Assignment

For this problem set, your goal is to create (or modify) a program that satisfies these goals:

1. It involves at least 2 program representation levels (most likely from Python, C, JVM, x86 Assembly) for sensible reasons
2. It involves at least 1 interesting algorithm and data structure
3. It is useful
4. It is creative
5. It is finished by the due date

To receive a green star on this assignment you must accomplish at least three of the goals. Achieving four of the goals is worth a gold star. Achieving all five may be worth a double gold star, especially if you find a good reason to incorporate the lowest representation levels (either JVM or x86 Assembly).

To satisfy the first goal, you should find a way to combine multiple programming levels to accomplish something that could not be easily accomplished using a single level alone. Some possibilities include:

- Creating a Python extension and program that uses it (this would involve using Python, C, and possibly x86 Assembly). You would need to justify extending Python by explaining why it is hard to do what you need with the provided language, or why better performance is needed for your application. (Google "python extension" to find lots of examples and documentation on creating Python extensions.)
- Create a C program that incorporates some assembly code for a good reason, such as being able to observe system properties not accessible directly or providing high performance for a fragment of performance critical code.

To satisfy the second goal, your program should involve an interesting algorithm and data structure. Ideally, it would be something you invent to solve a unique problem, but more likely it will involve implementing an established data structure and algorithm. To satisfy goal #2, your report should include an analysis of the asymptotic running time and memory use of the relevant part of your program.

To satisfy the third and fourth goals you will need to do something that actually has value (at least to yourselves, but ideally to many potential users) and that has not already been done by other people. If you are at a loss for ideas, you can satisfy goal 3 (but not goal 4) by improving the performance or functionality of SpamBayes, a spam filter implemented in Python.

To satisfy the fifth goal, you need to turn in a report and email a zip file containing a working implementation of your program by the due date. If you do not satisfy the fifth goal (finished), you should still turn in something by the due date. If you do not have a complete implementation, you should turn in a partial implementation, design description, and explanation of what is left to do.

If you are unsure if your idea for this problem set is acceptable, send me email describing it before Thursday, April 27. (You can, of course, ask questions or refine your idea after that, but it may be too late to get general feedback on your idea after Thursday.)

Report

On or before the due date (see above), you should turn in a paper report (either in class, or at my office, Olsson 236A). The report should contain these sections:

1. **Description (no more than 1 page)** — describe the program you created or the modifications you did to an existing program. If you aim to satisfy goal 3, you should also explain why it is useful and interesting. If you aim to satisfy goal 4, you should explain why it is original and how it differs from previous work on the same goal or why your goal is new and unique.
2. **Programming Tools (no more than 1 page)** — Identify the programming tools you used (or would plan to use to complete the program). Explain why you choose the tools you did. Discuss the advantages and disadvantages for the choices you made and consider alternatives.
3. **Data Representation and Algorithm Analysis** — If you aim to satisfy the second goal, describe the data representation and/or algorithmic problem you had to solve and describe your solution. Include an analysis of the asymptotic running time and memory use of your solution.
4. **Status** — Describe how much of your program is complete, and what is left to do.

In addition to submitting the report on paper, send a zip file of all your code and a README file explaining clearly how to run it, to cs216-staff@cs.virginia.edu. If you have test cases (you probably should!), include them in your zip file also. If your program is deployed as a web application, send the URL for the application in your email.

Mantras

Have Fun! — If you can't think of any program that it would be fun to work on, you might want to consider if computing is the right field for you.

Be Creative! — You can do anything you want so long as it is not destructive or unethical.

Be Realistic! — You only have 1 or 2 weeks for this. Try to think of something that you can make a good start on in a limited amount of time, but could continue to improve and expand later.

Try and pick something to work on that you find interesting and important enough that you will want to keep working on it over the summer, and that will have an impact if you are successful.