Spring 2024 Cop 4516 Individual Final Contest Summary (by Arup Guha)

The first part of the summary will be a problem by problem analysis of how students did in the contest and some observations about the problems. The second part will be an analysis of the teams and how they did. Here is the first, by problem:

Problem A: Slow Typing
Solved by: 55 out of 60
First Solve: 3 minutes (Max Bagatini Alves)

This was definitely the easiest problem in the set. Luckily it ended up first alphabetically, so students discovered the problem quickly. The biggest challenges in this problem were precision and rounding. Doubles should be used and the output format must be exactly to 2 decimal places. At the 30 minute mark 35 student had successfully solve the problem. Over the course of the contest more students solved the problem. In fact, there were two students who solved the problem in the last hour increasing the total number of solves to 55.

Problem B: The Grand Social Gathering
Solved by: 37 out of 60
First Solve: 23 minutes (Ethan Voth)

There was a gap after Ethan solved the problem, but about 10 minutes later, a steady stream of correct solutions to this problem started coming in. Over the course of the whole contest, more and more students solved this problem. 14 students solved it in the first hour, 12 solved it in the second hour and 11 solved it in the last hour, so a fairly equal split in terms of how long students took to complete this problem.

Problem C: Sum of Divisors
Solved by: 12 out of 60
First Solve: 50 minutes (Juan Nicora)

The first correct submission to this problem came from Juan Nicora, who perhaps had an advantage using Python, which has a built in fast modular exponentiation function which other languages don't have. (Java does, but only if you create BigInteger objects.) In addition, Juan happened to know that for a prime number p, a^{-1} mod p can be computed as a^{p-2} mod p, as a consequence of Fermat's Theorem. 18 minutes later, Shian Morrison solved the problem in the intended manner, calculating each sum in a loop without using any built-in language features in Java. Seven students solved the problem in the second hour while 4 more solved it in the last hour. Over those 12 correct submissions, most of the possible techniques on the problem were used. Some used the same technique as the judge solution, manually computing each sum in the sum of products, avoiding mod inverse. Some students used code for mod inverse and the short formula and others did what Juan did and utilized Fermat's Theorem to find the modular inverse without modular inverse code.

Problem D: The Island of Valcu
Solved by: 0 out of 60
First Solve: N/A

This turned out to be the most challenging problem in the set and stymied all competitors. Several figured out that a multi-source BFS from all the volcanoes to the rest of the squares was required, but no one could figure out what to do with that information. In fact, the judges were concerned with the difficulty of this problem, but decided to keep it in the contest because it was a nice problem and a good deal of work had already been put into it.

Problem E: Company Tug of War
Solved by: 1 of 60
First Solve: 122 minutes (Ethan Voth)

Ethan had submitted on this problem much earlier but gotten a run time error. He eventually got off of the problem but when he returned to it realized that he hadn't modified the call stack depth, which is necessary for Python solutions of this problem. The most common error of the rest of the submissions was attempting a greedy algorithm when determining how to split up the two teams. A strongly related problem to this one: finding a subset of values from a list that add to a specified target is NP-Complete. That means that there's no known polynomial solution to solve the problem. This is why brute force is required. It's likely that some minor enhancements can be made to the brute force, but given the strict bounds, we can see that the straight brute force runs fast enough, so there's no need to further complicate the code.

Contest Summary

With the first problem in the set coincidentally being the easiest, most students discovered Slow Typing quickly. By the end of the first hour, 47 students had solved the problem, of which 35 had solved it by the 30 minute mark. A vast majority of these students moved on to the second problem in the set, Grand Social Gathering. This problem requires the use of a TreeMap to link names to frequencies and also requires finding the next key in alphabetical order in the map. By the end of the first hour, 14 students had solved The Grand Social Gathering. The only other problem solved within the first hour was Sum of Divisors. But the student who solved that hadn't solved The Grand Social Gathering. Thus, at the end of the first hour, 14 students had solved two problems and another 34 students were at 1 problem. (There was one student who solved The Grand Social Gathering but not Slow Typing at the end of the first hour.) The difference in penalty points at this point was relatively small. At top of the scoreboard Ethan Voth had 28 penalty points, followed by Jason Helman with 39 penalty points and Gabriel Rodriguez with 55 penalty points. Thus, at the one hour mark, the contest was still wide open.

Early in the second hour, Shian Morrison sprinted to the lead. At the 51 minute mark, he solved The Grand Social Gathering, but then, only 17 minutes later, he submitted the second correct solution to Sum of Divisors becoming the first student with 3 problems solved. As the halfway point of the contest approached at an hour and a half, the number of students with at least one correct solution increased to 53 out of 60 students in the competition. Of those 53 students, 20 had at least two correct solutions. For 27 minutes, Shian was the only student to have three problems solved. But then, at the 1:35 mark, Sathvik Prabhudeva solved Sum of Divisors to leapfrog into first place, and then only one minute later at the 1:35 mark, Gabriel Rodriguez solved it as well, jumping ahead of both Shian and Sathvik. Thus, at the 1:36 mark, first, second and third place all had 3 problems with 151, 164 and 188 penalty points, respectively. At this same point in time, two more students solved their first problem, pushing the number of students with at least 1 solve up to 55 out of 60. At the 1:47 mark, Ethan Voth, who had a very strong start to the contest solved Sum of Divisors, jumping into first place with three problems and 135 penalty points. By the end of the second hour, when the scoreboard froze, 6 students had solved 3 problems, 20 students had solved 2 problems, and 29 students had solved 1 problem.

Two minutes into the freeze, Ethan Voth became the first student to solve the Company Tug of War problem. He had submitted an incorrect solution which showed he knew the correct algorithm and must have left the problem to work on Sum of Divisors before coming back to it to find his bug. It turns out that it wasn't really a bug in the traditional sense. Ethan uses Python which has a very low default call stack depth. Ethan's fix was simply to make his call stack bigger than 10,000, the maximum number of possible recursive calls. A minute after Ethan got to 4 problems, Mike Rojas joined the set of students with 3 solves by getting Sum of Divisors, putting him in fifth place, the last place to earn a trophy. At the 2:15 mark another student solved Slow Typing (recognizing that he had to output to a 2 decimal places), bringing the total number of students with at least one solve to 56. At the 2:21 mark, Joseph Torres solved Sum of Divisors to get to his third problem, but didn't crack into the top 5 with the solve. Other than this submission, most of the correct submissions in the first half of the last hour got students to 2 solved problems. At the 2:36 mark Carson Breitbart solved Sum of Divisors to jump into 9^{th} place and be the 9^{th} person to have 3 problems solved. Then only one minute later Nicolas Garcia solved the same problem to become the 10^{th} person with 3 solves. But none of these submissions changed the top five, which, with 20

minutes remaining were: Ethan Voth (4, 297), Gabriel Rodriguez (3, 151), Sathvik Prabhudeva (3, 164), Shian Morrison (3, 188), and Mike Rojas (3, 220). With 19 minutes left in the contest a student solved Slow Typing bringing the total number of students with at least one problem to 57. At the 2:45 and 2:48 marks, Julio Lopez and Juan Nicora both solved The Grand Social Gathering. Both had already solved Sum of Divisors, so both moved to 3 problems solved. Neither cracked the top 5 due to penalty points. Interestingly enough, with those two solves, all students who solved Sum of Divisors also solved Slow Typing and The Grand Social Gathering. At the 2:57 mark, three minutes before the end of the contest, Hayden Sandler solved The Grand Social Gathering on his eighth submission to get to two problems, and this turned out to be the last correct solve for the contest.

Because there was only one student with four problems and several with three right after the two hour mark, no changes occurred to the top five in the last 57 minutes of the contest, so the trophy winners were exactly the set of students who were in the top 5 at the 2:03 mark:

5th Place: Mike Rojas (3 problems, 220 penalty)
4th Place: Shian Morrison (3 problems, 188 penalty)
3rd Place: Sathvik Prabhudeva (3 problems, 164 penalty)
2nd Place: Gabriel Rodriguez (3 problems, 151 penalty)
1st Place: Ethan Voth (4 problems, 297 penalty)

As is often the case with contests, particularly those with an increasing gradient of difficulty, having a quick start helps competitors greatly with their overall placement because (a) they accrue fewer penalty points on those problems and (b) they have more time to work on the future questions. Each of the students in the top five were able to solve 2 questions in under an hour and follow that up with another solve later in the contest. Eight other students solved 2 questions in under an hour but were not able to follow that up with another solve later in the contest.

All in all, it was a good contest. There was a nice gradient of difficulty and with time many students were able to eventually figure out how to solve several of the problems. It would have been nice if a few more students got to four problems, so perhaps the judges underestimated the difficulty of the last two problems. Technically, the contest went very well. There were no judging hitches and the data for each problem was solid – there were no known submissions that were accepted using invalid techniques.