

The 1999—2000 USACO Report

by **Rob Kolstad**

<kolstad@usenix.org>

Head Coach, USACO

For several years, USENIX has sponsored the USA Computing Olympiad, an organization that fosters pre-college computing by offering contests whose winners ultimately travel to exotic foreign locales to represent the USA in international competitions.

I am the head coach and work with a fabulous set of people that includes Director Don Piele at the University of Wisconsin, Parkside; Greg Galperin from MIT (now on grad school leave at a startup); Hal Burch (Carnegie Mellon, interning at a startup); Russ Cox (Harvard undergrad); and Brian Dean (MIT grad student). Together we create programming tasks, grade them, evaluate performance, offer feedback and so on.

The 1999—2000 Contests

The USACO offers programming contests over the Internet throughout the school year. These contests are "open entry" (any pre-college student in the world can enter, as long as they have Internet access to see the problems). This year we continued to offer two divisions (one much less challenging than the other) into which students self-classified themselves. International participants often account for more than half of the total entries! Here's a quick table that summarizes participation for the 1999—2000 USACO season:

	Fall 1999		Winter 2000		Spring 2000		U.S. Open
	<i>easy</i>	<i>hard</i>	<i>easy</i>	<i>hard</i>	<i>easy</i>	<i>hard</i>	<i>hard</i>
Total	38	149	46	159	29	150	373
U.S.	9	64	22	54	16	57	221
	(24%)	(43%)	(48%)	(34%)	(55%)	(38%)	(59%)
Non-U.S.	29	85	24	105	13	93	152
	(76%)	(57%)	(52%)	(66%)	(45%)	(62%)	(41%)
Countries	15	31	16	32	9	32	35

The participation from the USA in the first contests was of great concern to me. I knew I was recruiting one to two dozen new programmers for each contest, but the total number of USA contestants was not increasing. I finally went to investigate to see which contestants were returning for future contests and to see if any pattern was recognizable.

The pattern was instantly and obviously clear. Those who score very, very low in a contest (e.g., scoring 80 out of 1,000 points) were quite likely to refrain from returning. Since this amounted to 10—25 contestants each time, the recruiting effort

(magazine articles, newsgroups) was required to yield a couple dozen new entries just to stay even!

Like a good marketer, I polled some members of the departing group and quickly found that they felt they need better training resources in order to compete with the world-class stars in the upper division.

In January, Hal Burch and I set out to create the "USACO Training Pages." Using material from the previous camp, problems from a myriad of sources, and Russ Cox as a problem analyst, we have created roughly 200 hours of training material for our competitors. If you know pre-college programmers who would like to improve their contest programming ability, please send them to <http://ace.delos.com/usacogate>.

The system includes sequencing, task grading, and a host of instructional materials on various algorithms. These pages are currently targeted at the student with a year or so of programming.

We're working to extend them "down" to the novice level this year.

The Training Pages have been discreetly announced through the USA since January. Over 300 pre-college students have tried at least some of the pages and the grader has evaluated 8,393 solution attempts submitted in both C/C++ and PASCAL. So far, it appears that the training pages are a super success and are creating students that, if this year's training camp is any indicator, are the equivalent of two training camps ahead of the game!

Regrettably, I cannot attribute the improved U.S. Open attendance results solely to the training pages. We improved our public relations program for announcing the U.S. Open so much that determining the reason for the dramatic increase is not possible.

The 1999—2000 Winners

Ranking the winners across four separate contests, some of which were more challenging than others (e.g., the U.S. Open) is a daunting task.

Consistency is an important factor; sometimes contestants "get lucky" once. Furthermore, the contests span the school year and the skill levels of the contestants change through the year.

The 1999—2000 season saw a spectacular performance by Reid Barton, a home-schooled junior from Boston. He won two of the contests (one of them outright). John Danaher (one of the many talented students at the Thomas Jefferson High School for Science and Technology located in Virginia) placed fifth, first, second, and fifth for a spectacular final year of eligibility. Sophomore Vladimir Novakovski, also a TJ student, earned a first and two seconds for the best performance in his age group. Veteran Percy Liang, a New Mexico student also in his final year, earned a first, second, and third.

The U.S. Open yielded some surprising results. Senior Jack Lindamood, a contender throughout the year, was the only U.S. student to achieve a perfect score and thus earned the U.S. Open championship crown. In his spare time, Jack programs DSP chips for Texas Instruments. Only nine points behind, Thuc (sounds like "took") Vu captured second place. A junior, Thuc has been in the U.S. only eight months and is still mastering the English language!

Below is a list of the students chosen to attend the USACO training camp, an eight-day affair held just south of Milwaukee at the University of Wisconsin's Parkside branch.

- Reid Barton, Arlington, MA
- Jacob Burnim, Silver Spring, MD
- Kevin Caffrey, Oakton, VA
- Tomasz Czajka, Stalowa Wola, Poland
- Adam D'Angelo, Redding, CT
- John Danaher, Springfield, VA
- Richard Eager, Alexandria, VA
- Percy Liang, Phoenix, AZ
- Jack Lindamood, Dallas, TX
- Yuran Lu, Presque Isle, ME
- Vladimir Novakovski, Springfield, VA
- Gregory Price, Falls Church, VA
- Gary Sivek, Burke, VA
- Steven Sivek, Burke, VA
- Thuc Vu, Westminster, CA
- Tom Widland, Albuquerque, NM

The list contains a whopping seven students from TJHSST in Virginia, including identical twins Gary and Steven Sivek. We finally have a finalist from California, which is underrepresented in all our contests for no reason I can determine. Danaher, Liang, and Novakovski were our returning veterans. Four campers (Barton, Danaher, Liang, and Lu) earned a head start on this year's international competition at the Balkan Olympiad a couple of months ago. (See the July *login*: for an account of that competition.)

Tomasz Czajka, a Polish veteran who won several USACO contests, was invited as our international representative this year. He provided stellar competition to our contestants in both formal and informal competitions.

Training Camp

Training camp has two direct goals: improve the skills of our top programmers and identify our four best candidates for gold medals at the international events throughout the rest of this season.

Indirectly, the camp and its rewards provide goals for all US contestants throughout the competitive year.

Contestants arrived at camp on Tuesday, June 20, and departed early on June 28. Two five-hour contest days highlighted the experiences, with dozens of instructional modules, problem attacks, and less formal exercises in between. This year we incorporated several parallel sessions with one coach/four students in an effort to increase personalized instruction and decrease the "Well, Joe knows the answer so I don't need to think" phenomenon. I think our effort was quite successful, and we'll continue this trend.

We implemented Dynamic Programming Day, during which Russ Cox organized lectures, drills, lab problems, and even a set of eight different student presentations on dynamic programming. The presentations were stunning and the students nailed the DP problems on the first contest.

The "training notebook," begun in earnest last year, grew to almost 150 pages of tutorials, drills, and programs. We're proud of this resource and believe it helps students throughout their competitive careers.

This year's innovations included not only written solution analyses for the contest problems but also in-depth graphical analyses of the test data and its creation.

Informal events included the Nine Men's Morris competition in which campers wrote programs that played against each other in the traditional English (and other) drinking game. Polish contestant Tomasz Czajka's program was undefeated. A (Frisbee-style) Disc Golf Competition was won by Jacob Burnim, while the traditional USACO Quiz Show was won by Gary Sivek, who edged out his twin brother Steven 2700-2400 in the final round.

The first round of competition at the camp was an easier round with very high scores (mostly due to success on dynamic programming problems). For the second round, we cranked up the difficulty (see below).

After Russ Cox graded all the problems and the group of coaches reviewed those scores and the finishes throughout the year, four students and one alternate were chosen for the international USACO traveling team. They are:

- Reid Barton
- John Danaher
- Percy Liang
- Gregory Price
- Jacob Burnim (Alternate)



L. to R: Gregory Price, Percy Liang, John Danaher, Reid Barton

The next big contest is slated for late August in Cluj, Romania (in the state of Transylvania); the t-shirts will be fabulous. The international finals (IOI — International Olympiad on Informatics) will be held the last week of September in Beijing, China. The U.S. team carries with it high hopes for multiple gold medals.

The Future

It's just amazing how much time an endeavor such as this can take, especially once a publicly available set of Web pages is available. We plan to continue the set of individual contests on the Internet next year with the same style of divisions (or maybe even pseudo-divisions that let contestants evaluate themselves better in appropriate subgroups).

This next year will be a milestone for the competitions, because we are finally leaving behind the 16-bit DOS-based compilers that have been de rigueur since the IOI's inception. These dinosaurs allow contestants access to only 640KB of memory and generate terribly slow code. I have chaired (I think I was the chair, anyway) a committee that is in charge of updating the technology. The current plan is to switch to GNU C/C++ and Free Pascal on the day after the IOI ends in China.

The new compilers enable: automatic problem grading without the problems of rebooting DOS, file systems that might or might not become corrupted after each program execution, lack of multitasking, and lack of memory protection. This changes the landscape for available contest types and dramatically eases grading (i.e., IOI grading for 300 competitors should take closer to a fraction of an hour than the current 0.5 to 2 days).

We're looking to grow the level of our competitions and the number of competitors throughout the next year.

In order to foster computer programming as a more recognized sport, I'm looking at holding a set of dual meets (just like track, tennis, swimming, or football) through the school year. Each meet will pit one school's team against another's — even if the teams are proximal only via the Internet. It's an easy matter to conduct Web registration, offer a variety of scoring mechanisms (like swimming? tennis?) and team organizations (singles, doubles, etc.) to make some exciting competitions (probably more so for the contestants than the spectators).

The parallelism that enables all the competitions to be just the same except for the competitors will be our little high-leverage secret.

The problem-grading system used by the Web training pages has proven very reliable and amazingly efficient. Email responses are returned to the submitter often within just a couple of seconds.

Opportunities

A couple of items really need a large group focus in our effort. First of all, we need publicity. If you know students who might be interested in competing, please send them to <http://www.usaco.org> and <http://ace.delos.com/usacogate>. The USACO page shows some of the accomplishments of our team over the years and always contains up-to-date info on the latest contests and events.

Second, we need additional programming tasks. These are a very special niche of problem type that require months (half a year? more?) to learn how to create. They suffer from the "Goldilocks Syndrome": they must not be too hard; they must not be too easy; they must be "just right." If you would like to help contribute to our "problem budget" (which could easily exceed 100 problems for the upcoming year), please contact me at kolstad@ace.delos.com.

Acknowledgments and Thanks

The USENIX Association's sponsorship of USACO is invaluable. Our time is spent creating contests, software, and educational opportunities instead of dressing up and begging for money. We're working hard to improve our program.

Director Don Piele implements a huge number of tasks, including: the USACO Web page, administering the budget, acting as USENIX liaison, serving as camp operations manager, planning all physical plant items for camp (including transportation), and meals. He is like the genie in Aladdin's lamp: wish for something (a three-hole punch, for example) and it appears within hours.

The four coaches (Brian, Greg, Hal, Russ) work with me throughout the year on problem creation, solution creation, test data creation, grading, and answering student questions from the training pages.

All of our assistants donate their time (almost ten days at UW-Parkside just for training camp and its setup). It's a great group to work with and continues to create new and exciting ways to improve our camp.

Thanks!

Two Challenging Tasks

Here's an easier problem (three stars out of four) created this year by Hal Burch. It is cowified to keep with the spirit of attending camp in Wisconsin:

Cow Talk

Like the natives of Africa and their drums, the cows moo to each other over long-distance networks of relay-cows. Cow A moos to Cow B; Cow B passes that message to Cow C, and so on. Cows only moo to specified partners, even if some other cow happens to be close by; and when two cows are partners, it is always the case that either can moo to the other to relay messages.

The speed of sound being finite, some relay links take longer than others. In this communication area, we can calculate the relay time seconds by dividing the distance between the cows by 100. Let's call the set of cows-who-relay along with the list of specific, relaying cow-partners a cow-network. The cow-network is always set up so that it is possible for any cow to talk to any other.

Some cow-networks are better than others. A cow-network's merit is measured in terms of the "average communication time between cows." This number is the sum of the total communication times between all possible pairs of cows divided by the number of pairs of cows; thus a lower number is better. When two cows wish to communicate, they always talk using the relay-cows which give the quickest possible moo transmission.

Given a set of cows, their (x,y) grid locations, and their communication partners, determine the link between two (hitherto unconnected) partners which, when added, creates a cow-network with the smallest possible figure of merit.

This year's hardest problem (four stars) was "Kinergarten" (kine, you know, is an archaic plural for cow). This problem was couched in terms of a teacher taking the young calves on a field trip in which they each held on to a set of strings in order to avoid separating. After a good amount of verbiage, the problem devolves to:

Kinergarten [Burch & Galperin, 2000]

Given a set of nodes (no more than 150) and weighted arcs between some of them, count the total possible number of different minimal spanning trees. A given arc weight appears no more than 10 times. 640KB memory limit. Five seconds CPU time limit on a P2/300.

The problem, already quite challenging, is complicated by the fact that the total number doesn't fit in just a few bits, thus requiring the implementation of parts of a "bignum" package just to keep track of the final result.

Conclusion

The USACO is alive and well. Our camp participants thought the organization must have at least several dozen affiliates, given the Web pages, contests, and general "look and feel" of the public interface.

This year was our most successful for increasing participation and raising the bar. The upcoming challenges are dominated by continuing to gain affiliates, maintain our quality, and prepare for the 2003 IOI, which is to be held in the USA. Our biggest challenge for 2003 is finding sponsors for the IOI.