IOI 2010 Tasks: General Information

There will be four tasks on each competition day.

The run-time CPU limit is 10 seconds and the run-time memory limit is 256MB for every test run, unless a larger limit is specified in the task statement.

The compile-time CPU limit is 10 seconds and the compile-time memory limit is 256MB.

Submitted programs must not do any input or output, must not use the network, must not do any system calls, must not attempt to gain points other than by solving the task as specified, and must not attempt to interfere with the operation of the grading system. Failure to follow this rule may result in a wrong submission or run-time error.

The use of standard mathematics, string and algorithm libraries is permitted.

Unless otherwise stated, the word *number* means *non-negative integer*.

Unless otherwise stated, numbers will not exceed 1,000,000,000. [It is possible to use `int` (C/C++) or `longint` (Pascal) to manipulate such a number or the sum of two such numbers.]

The phrase *between N and M* shall be interpreted to mean *all numbers not less than N and not greater than M, where N≤M*.

Contestants may submit at most once per task per minute. More frequent submissions will result in *submission failed*.

Contestants may submit at most 1MB per task. Note that the `submit` command submits a zip file containing *all* files in the current folder with the name `*.c` `*.cpp` `*.pas` or `*.txt`, whichever matches the file named on the `submit` command. This zip file must not exceed 1MB in size.

For each submission, the score for a certain number of test runs will be available to the contestant. The contestant may release all remaining test scores using the *release tokens* of the web grader interface. Contestants will have two release tokens per task, which each regenerate 30 minutes after use.

The final score for each task will be the *maximum* of the following: *rounded to the nearest integer*:

- The sum of the scores of the subtasks for any *released* submission
- The sum of the scores of the subtasks for the last submission, whether the last submission is released or not

The first subtask of every task is designed to be solvable by all contestants.

Each subtask has either a fixed score, or a score determined by a formula given in the subtask.

The total score for most tasks is 100 points. For some tasks, it is possible (but not likely, in the

opinion of the Scientific Committee) that a contestant may score up to 110 points.

Detailed test results will include one of the following:

- OK
- OK *subtask*
- OK score *score*
- Wrong result
- Invalid library call
- Too many library calls
- CPU Time limit exceeded
- Forbidden *system call*
- Run-time error
- *a RunC failure message*

Additional RunC output may also be shown in the detailed test results.