

Technical Report of Baltic Olympiad Informatics 2014

Motiejus JAKŠTYS

Technical Committee, National Olympiad Informatics, Lithuania
e-mail: motiejus@jakstys.lt

Abstract. In this article we present the technical and organizational details of Baltic Olympiad Informatics (BOI¹) 2014 held in Palanga, Lithuania.

BOI has been held since 1995. During that years, we have accumulated a lot of experience and polished the process. The article presents some of the organizational and technical issues and how we dealt with them. In conclusions, we highlight some recommendations for preparation for the next years. Other countries organizing national or international Olympiads might benefit from the recommendations and experiences.

Keywords: informatics olympiads, linux, desktop preparation.

1. Introduction

Baltic Olympiad Informatics (BOI) is an annual informatics competition by high school students started in 1995 by Lithuania, Latvia and Estonia. Later, more countries joined BOI: Finland, Sweden, Norway, Denmark, Germany and Poland. The goal of the Olympiad is to bring gifted students to international environment as a preparation for IOI.

The key part of the Olympiad is preparing high-quality tasks, tests and checkers. The evaluation system must be robust, the evaluation criteria must be clear. These responsibilities are called "Scientific part", and have a Scientific Committee dedicated to them.

Besides Scientific committee, there are two supporting committees: Organizational and Technical. Organizational part is something that is generic for every Olympiad, be it for informatics, physics, mathematics or of the Old Prussian Literature. Responsibilities

¹ BOI is also used as abbreviation for the Balkan Olympiad in Informatics. Within this article, BOI refers only to the Baltic Olympiad.

include finding the accommodation venue, organizing arrivals and departures, meals and leisure activities.

Technical committee is responsible for evaluation servers, student workstations, public internet, private network, student workstation network configuration, and technical support of the Organizational Committee.

In this article we will briefly remind the history of Baltic Olympiad Informatics (section 2). After historical introduction, we will present work of Organizational Committee (section 3) and of Technical Committee (section 4) which was carried out in BOI2014. The article describes some challenges we faced in BOI2014, and how each of the challenge was handled. The cases will be followed by related work, concluded by suggestions and recommendations for the following events.

2. Short History

Baltic Olympiad Informatics increased from three countries in 1995 to nine countries in 2014.

Since 1996, each participating country arrives with 6 students and 2 deputy leaders. On rare occasions, a country might be permitted to arrive with more than 6 students; a situation where 7 students are needed might happen when the performance between 6th and 7th in the national contest is very similar.

Up to this day, the style of BOI is kept similar to IOI. We will further elaborate on the similarities and differences in organization between the two. Detailed history of Baltic Olympiad Informatics are covered in (Bulotaitė *et al.*, 1997, Poranen *et al.*, 2009).

3. Organizing BOI2014

The very first work organizing the Olympiad is carried out by Organizational Committee. The dates have to be agreed, venues picked, countries invited any many other things. We will cover them in this section.

3.1. Schedule

Compared to IOI, BOI schedule is more compressed and fully fits to four-five days.

First day: arrivals, opening ceremony. At night, first set of tasks is selected and translated.

Second day: competition day 1. At night, second set of tasks is selected and translated.

Third day: competition day 2.

Fourth day: closing ceremony, party.

Fifth day: departures.

Short duration brings in many advantages. First, the organizers, which are usually university students or full-time employees, need to take less days off from their primary activities. Secondly, it is cheaper for the organizers: less nights to pay for accommodation means more money can be spent on other things like prizes and leisure activities.

Why can BOI allow short duration, but not IOI? That is mainly possible because of close distance: all delegations are able to arrive in a short time (usually, all delegations arrive in the allocated half-day period). What is more, there are no major time zone differences which would require long adaptations.

3.2. *The Venue*

In BOI2014, participating countries take turns when hosting the Olympiad. BOI2013 was hosted in Rostock, Germany. BOI2014 was Lithuania's turn.

During National finals, Lithuania is hosting each Olympiad in a different location every year. Olympiads are usually hosted at schools in different towns through out the country. Given a similar number of participants and team leaders to the national contest (50 contestants and 20 organizers in National Olympiad compared to 60 contestants 40 team leaders and organizers in BOI), it was decided to use the same style for BOI2014: host the Olympiad in an interesting venue (e.g. health resort, school in an interesting town), rather than a technically convenient one (e.g. University campus).

After evaluating possible venue options, a health resort in a sea resort Palanga was chosen. The health resort could comfortably accommodate the staff, students and guests. For leisure activities, besides being close to the sea, surroundings of Palanga are rich with notable natural points of interest and historical landmarks.

Organizing the Olympiad in a health resort gave advantages for organizational matters (accommodation, meal, activities), but raised technical challenges. Computer and power networks had to be created from scratch. Computers needed to be rented before the Olympiad and fully set-up with tables, chairs and cabling.

In the following section, we will discuss the technical setup in the health resort.

4. Technical Infrastructure

Like mentioned in section 3.2, the technical infrastructure needed to be created from scratch. At the venue, the following were readily available:

- Uplink internet connection (DSL).
- Power sockets.
- Large hall.

Technical and Organizational Committees needed to rent tables, chairs and computers, set them all up, create and connect to a power grid.

4.1. Physical Infrastructure

All student workstations were in a large hall. See Fig. 1 for schematics of the layout. Red lines and circles show power grids. Black crossed circles are WiFi access points. White squares are desks with computers.

Since all the infrastructure needed to be created from scratch, we tried to make it as simple as possible. For example, to avoid cutting and pulling many Ethernet cables to set up wired networking, we opted for wireless. This could have been done because the workstations were laptop computers, and thus all had ability for wireless networking.

During planning stages, there were concerns about interference of many devices using wireless network in such a small area. To avoid that, we created 3 non-overlapping channels (1, 6, 11) and configured two access points per channel. The workstations were

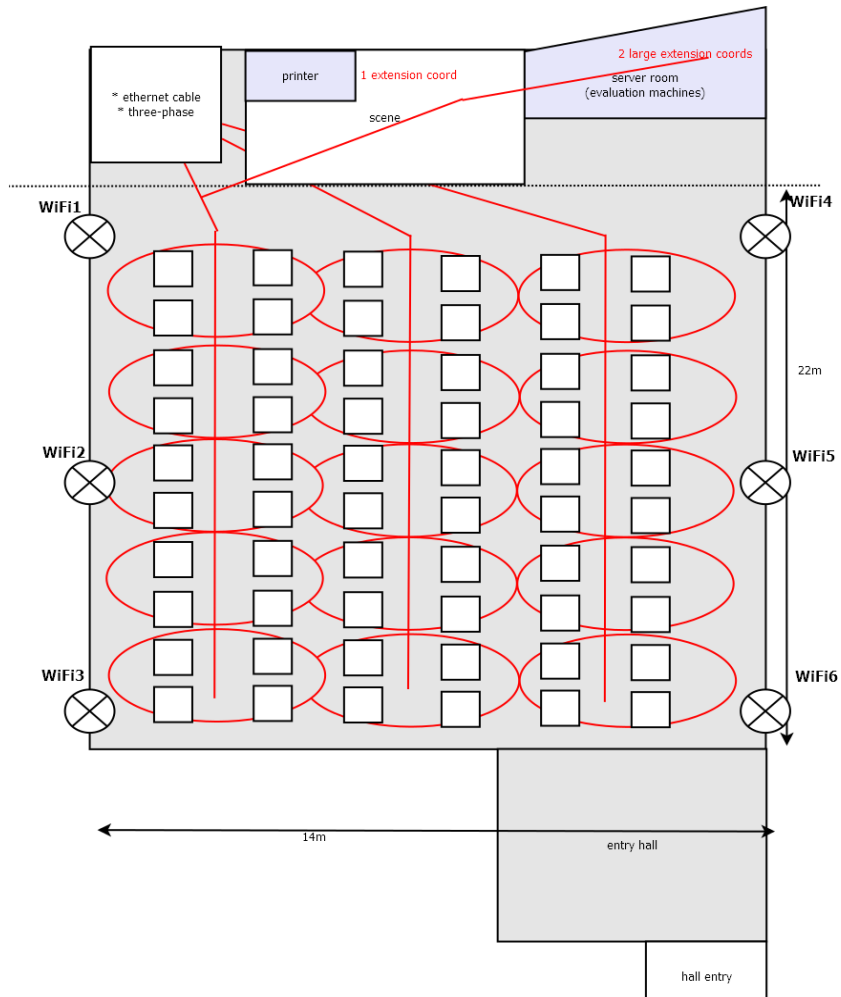


Fig. 1. Hall schematics.

configured to pick all access points at random (the same priority in `wpa_supplicant` configuration file). Networking was used only for accessing the Contest Management System's web interface, thus low traffic, and worked perfectly.

According to Power Supply Unit specifications, a laptop can draw up to 200 W of power. 60 laptops can draw up to 12 kW. Including network equipment, contest management system and evaluation servers, this figure goes up to 15 kW. A single commodity power socket can provide up to 8 kW. To accommodate our needs, we were provided with a three-phase commodity power network, each phase capable delivering 8 kW each, giving 24 kW of total capacity.

4.2. Student Workstations

Student workstations are providing students a comfortable environment to develop the contest tasks. Workstation software has to be recent and comfortable. It also should restrict contestants from cheating, e.g. sharing solutions with each other.

4.2.1. Software

Until 2001, DOS programming environment was used in the IOI contest. That changed in 2005 when it was replaced to Linux (Jyrki Nummenmaa, 2005). Like in IOI, in BOI we permit and support only Linux operating system. As a task to Technical Committee, an Operating System image with necessary configuration and software needs to be created.

When thinking what to base the Linux distribution for BOI2014 on, we had the following requirements in mind:

- Should have all the necessary packages included (compilers, IDEs, documentation, ...).
- Contestants should be able to test the environment easily, so it needs to be easy to run it outside of the Olympiad. Ideally, a downloadable ISO file which can be booted from a CD or USB memory stick.
- The distribution might be reused for other contest (or) in the following years. Since the requirements in the following contests might change, we want to make the distribution configuration and creation as automatic as possible.

Because of the following requirements, we chose Debian-Live. Debian-Live is a framework for creating Debian-based bootable system images. It works as follows: given some configuration (e.g. list of packages to be installed, users, passwords and other settings), the Debian-Live machinery creates a bootable ISO image. The bootable ISO image is a live image which can be booted and is perfectly suitable for playing around.

However, the system on the live image does not have a permanent file system. This is a problem during the contest, because, in case of accidental reboot, the contents of user's home directory will be lost, including all the produced solution files. To avoid that, the live ISO image is installed to a hard drive, which permits permanent storage, and the hard-drive image is created for usage during the contest. Hard-drive based installation from the generated ISO file is very similar to the live ISO image, just a few extra hard-drive specific tweaks are necessary.

After some preparation by Technical Committee, now ISO generation process is single-click². Normally, installing the ISO to a (virtual) hard drive automatically is a bit more involved because, by default, the Debian Installer asks questions. Luckily, these questions can be pre-answered in a *preseed* file, and the installation to the (virtual) hard-drive can be one-click too³.

Like mentioned before, a few tweaks are still needed to bring use the installed hard-drive image in the Olympiad. First, security: there are many things we want to restrict users to do:

- Connect to unauthorized networks.
- Change network configuration.
- Start and stop services.
- Use external media.

To support the restrictions above, we created scripts which modify the resulting hard-drive image to accommodate these requirements.

4.2.2. Booting the Machines

We had two major challenges for imaging and booting the machines: we had no wired connectivity to transfer the images, and we were not permitted to alter the on-disk Windows installation. After the contest, the laptops must be sent back in the condition they were received.

Since we cannot alter Windows (and change the Master Boot Record, MBR), it brings up a challenge of booting the Linux image. Booting from WiFi network is not an option, because it will overload the wireless network, and even to bootstrap PXE is not an option (since motherboards do not support netboot over wireless). Therefore we devised a way to store the operating system and user files on the hard drive without modifying the existing Windows installation. Here's how: the full disk image was placed on C:\ drive of the Windows partition. The MBR stays untouched. That way the partition table could stay intact, and it is easy to clean up a single file afterwards.

Placing a hard-drive image on NTFS drive still requires some work to boot it. Motherboard must find stage0 boot loader in the MBR, which should eventually start booting the hard-drive image on the NTFS partition.

Here is the sequence we devised to reach the partition on disk:

- Computer boots kernel and initrd from the USB stick.
- Initrd has a script which finds C:\.
- C:\ is mounted on /mnt/looproot. /mnt/looproot/liox.raw is the full Linux system image.
- Root partition on /mnt/looproot/liox.raw is bound to /dev/loop0.
- initrd binds /dev/mapper/loop0p1 as a root partition and passes the command to /init of the root partition. The booting continues as normal as if root filesystem were on /dev/mapper/loop0p1 with the kernel command-line argument `root=/dev/mapper/loop0p1`.

² or rather, a `make iso` invocation

³ `make vm`

Here is the snippet of the script from `initrd` which executes the actions above:

```
mkdir -p /mnt/looproot
mount -t ntfs-3g ${LOOPROOT} /mnt/looproot
losetup /dev/loop0 /mnt/looproot/${LOOPSRC}
kpartx -a /dev/loop0
```

`${LOOPROOT}` and `${LOOPSRC}` can be passed via cmdline. For example:

```
initrd initrd.img
kernel vmlinuz looproot=/dev/sda2 \
loopsrc=liox.raw \
root=/dev/mapper/loop0p1
```

It is still necessary to bootstrap the kernel and `initrd` so they can do the work described here. Small USB sticks were chosen for this: when a (re)boot is necessary, the operators plug in a USB stick. After kernel and `initrd` are loaded, the stick can be taken out and the operator can proceed to the next machine. For 60 laptops on-site, with this approach, it takes around five minutes to boot all machines for three trained operators.

4.2.3. *Uploading the Image*

Copying the disk image to the machines is different every year. In other words, we did not find the best solution yet which would prove its value for the long-term.

In BOI2014 we were walking around with external hard drives and USB sticks having the copying of the image scripted. It is laborious and it takes time, but is very simple. Because its simplicity and reliability, this copying solution has been used for many years. Over the years organizing the National Olympiad, the Technical committee tried different network-based approaches, but, if the approaches work in the labs, they usually fail in real-life due to misbehaving overloaded networks.

4.2.4. *Networking*

Main server was running the following services:

- Contest Management System – all services except Workers.
- HTTP proxy – `nginx`.
- DHCP and DNS services – `dnsmasq`.

DHCP was serving the IP addresses to student workstations, which were querying for IP addresses over WiFi.

Since the workstations were laptops and physical network infrastructure did not exist, we opted for using WiFi for networking. We used 6 access points in the room. Each AP had a different ESSID, but were connected to the same bridged L2 network.

Students' network access to other public networks (which can also be WiFi Access Points from the cell phones in their pockets) was restricted by pre-configuring the network and not granting the permissions to change network settings. This is done in `wpa_supplicant.conf`.

This network setup is simple: there was only one DHCP server running in the network, and all the IP addresses were dynamically configured. Dynamic configuration

(DHCP) brings a nice side-effect that all the network configuration is in a single place (in our case, `dnsmasq.conf`). What is more, having the same L2 network does not require any switch configuration.

Having everything in the same L2 network has security drawbacks: students are blocked from accessing each other only in their workstations. If a student manages to get superuser privileges on the machine, it would be only a small step away to communicate with other machines on the network, and, for example, to share and discuss the problems and solutions.

Another disadvantage is contest server reliability. If the server goes down due to hardware failure, it will take non-trivial amount of time to recover. Taking the hard drives out, plugging them to another machine, reconnecting the network would all take time and students' ability to interact with the contest management system would be impaired. This risk could be mitigated by having a standby server ready to take over in case the first one goes down.

4.3. *Internet Access*

Providing internet access for Scientific Committee and Public proved to be quite a challenge on the facilities. In the next sections, we describe our experiences in setting it up.

4.3.1. *First Attempt*

About one month before the contest, we arrived at the planned contest location to inspect the facilities. Among other things, we tested internet connectivity. The uplink was a DSL line with the following properties: small latency (up to 20ms to known servers) with 10Mb/2Mb of bandwidth. Since internet should have been primarily used for accessing public internet by staff and students, we agreed these will be sufficient.

After arrival to the contest location to do the preparation and preliminary network setup, internet connection quality significantly degraded. Degradation manifested by IP packet loss of up to 80%. Internet connection is critical for the Scientific Committee, since the tasks and tests are shared via a third-party internet-accessible service GitHub. The packet loss was so significant that initial cloning of the task repositories was failing. This needed to be fixed.

While investigating, we could not relate the problem with anything from our side. After testing various options to reduce the packet loss, all in vain (like putting all our infrastructure under a single a NAT host), we decided to look for the cause of the problem in the facilities. Late in the afternoon, we discovered an Ethernet hub which we could not identify the purpose of. Disconnecting the hub fixed the packet losses and removed jitter, and did not seem to break anything: we received no complaints from the staff, nor internet access was broken any on-premise machines we had access to.

Later at night, a staff member contacted us with a complaint that credit card reader does not work. The night before the Olympiad, the facilities were hosting a wedding party, and they were unable to service the customers' credit cards for the extra drinks. We turned the hub back on, and the staff reported the card problem fixed. Further throughout

the Olympiad, we did not dare to disconnect the hub any more and had to resort to alternative means of internet connectivity.

4.3.2. *Second Attempt*

We connected a 3G USB dongle to a netbook, and connected that netbook to the main server. With this setup, all "internal" network access (contestants' machines and WiFi of the Scientific Committee) went through that dongle. As long as nobody were downloading large files, latency and throughput were acceptable.

The 3G dongle remained in use throughout the duration of the Olympiad.

5. Related Work

(Poranen *et al.*, 2009) is an accurate article describing the history and scientific part of Baltic Olympiad Informatics. The article was a joint effort of a number of hosting and participating BOI countries. Besides details about BOI tasks, history and organization, each country gives a brief overview of their national Olympiad.

(Dagienė and Skūpienė, 2007) describes experiences of National Olympiad of Lithuania. Like mentioned in section 3.2, challenges in organizing Baltic Olympiad Informatics and National Olympiad Informatics are very similar. This article highlights the reasons for creating the Olympiad, its history, challenges and organization.

(Imajo, 2011) describes a case study running 60 laptop computers under a single wireless station. The article describes Japanese National Olympiad Experience using a single Apple AirPort Extreme wireless. According to the article, it was a success. At BOI2014 we used six commodity routers instead of one: "This time we relied on AirPort Extreme wireless routers because they can officially manage 50 wireless connections at the same time according to Apple Inc., although ordinary wireless routers can manage about 20 at the most."

(Blackham, 2013) describes the technical experiences of running IOI. Due to the number of participant workstations, setup and characteristics in IOI are quite different. Interestingly, (Blackham, 2013) pushes the images using UDP multicast. BOI and IOI are quite different in scale and requirements for safety:

- Even though the Olympiad was hosted on "on an extremely reliable power grid", the there was backup power supply prepared nevertheless. At BOI, we were also on an extremely reliable power grid, however, did not invest in renting and setting up the generators.
- From the contest system side, all of the critical pieces of infrastructure had replicas or hot backups: the database for Contest Management System had a was replicated and in hot-standby mode. Core servers and network equipment were backed up by UPSes to help survive the power outage. At BOI, the main server had a software RAID-1 array with a prepared machine to put the disks to in case of breakage.
- Each participant's workstation was connected to their own L2 network. Since this requires support from the networking gear, we went for a simpler approach – each participant got their own L3 network instead of L2. Being simpler, the approach

is less safe: if two participants are able to gain superuser privileges on the workstation, they would be able to configure the network so they could see each other. However, exploiting this fact requires both gaining the elevated rights and coordination, which we deem an acceptable risk.

6. Further Recommendations

Hosting the event in a sea resort proved to be a very good choice. Contestants lived only a few hundred meters away from the sea, which provided very good opportunities for relaxation close by. Good location had the price of running the contest in a health resort without technical infrastructure (only electricity granted), which, in the end, was fully sorted out. The laptops and desks were rented, and, thanks to sufficient manpower and time, assembled on time. As long electricity and basic internet facilities are granted, it is sensible to prioritize accommodation, leisure and attraction facilities rather than technical infrastructure.

For the next BOI, we would recommend to keep researching for an efficient way to distribute the distribution images over the network. It might be worth looking at udcast proposed by (Blackham, 2013).

It should go without saying that upon arrival to the contest location, the systems should be preconfigured as much as possible. However, it is important to leave flexibility for ad-hoc changes on every stage. Things that are prepared one day or the night before the contest should be simple rather than perfect. It is acceptable for things to be dirty, as long as they are documented, can be properly revised after the contest.

References

- Blackham, B. (2013). A behind the scenes look at IOI 2013. *Technical Report*, University of Queensland. <http://tiny.cc/ioi2013>
- Bulotaitė, J., Diks, K., Opmanis, M., Prank, R. (1997). *Baltic Olympiads in Informatics*.
- Dagienė, V., Skūpienė, J. (2007). Contests in programming: quarter century of Lithuanian experience. *Olympiads in Informatics*, 1, 37–49.
- Imajo, K. (2011). Contest environment using wireless networks: a case study from Japan *Olympiads in Informatics*, 5, 26–31.
- Jyrki Nummenmaa, C.I. (2005). *Linux Programming Environment on CD*. <http://www.ioinformatics.org/newsletters/html/ioinews5.htm>.
- Poranen, T., Dagienė, V., Eldhuset, A., Hyyro, H., Kubica, M., Laaksonen, A., Opmanis, M., Pohl, W., Skūpienė, J., Soderhjelm, P., Truu, A. (2009). Baltic olympiads in informatics: challenges for training together. *Olympiads in Informatics*, 3, 112–131.



M. Jakštys is a software engineer at Amazon.com, Inc. He is a chairman of the Technical Committee in National Olympiad Informatics (Lithuania) since 2014 and was a chairman of Technical Committee in Baltic Olympiad Informatics 2014 in Palanga. Motiejus was a Deputy Leader of the Lithuanian team in IOI 2014.