

EXPERIENCE OF USING MULTIMONITOR WORKSTATIONS UNDER XFREE86 4.X IN VEPP-5 CONTROL ROOM

D.Yu.Bolkhovityanov, R.G.Gromov, I.L.Pivovarov, A.A.Starostenko
The Budker Institute of Nuclear Physics, Novosibirsk, Russia

Abstract

Modern PC workstations often provide more CPU power than required for most control applications. On the other hand, the screen space is always in short supply. One possible solution is to use more PCs, but in fact we need only more screens, not more keyboards, mice, etc. PC architecture allows using more than one videocard, and the X Window protocol is aware that there can be more than one screen. Until the release of XFree86 version 4 there was no freely available server capable of driving multiple “heads.” We have been using multiheaded workstations under XFree86 in the VEPP-5 control room since early 2000 (currently 4 4-headed PCs plus several dual-headed). The “Xinerama” mode (one-large-screen) is better suited for accelerator control system than “several separate screens.” When moving to this configuration we’ve encountered a number of, mostly human-related problems. Some of these required modifications to the X server. Additionally, the “style” of performing control has slightly changed.

1 NEED IN MORE SCREEN SPACE

Historically automation at BINP is based on CAMAC. Home-made Odrenok [1] machines were used as both crate controllers and as the main computational power. The information was displayed via CAMAC-based display controllers, which gave a 256×256 color pixel picture. That allowed sufficient display space for most tasks.

On the new VEPP-5 facility the computation and high-level control was moved from crate controllers to Intel-based workstations. So, the aging CAMAC display hardware wasn’t an option.

Modern video cards and monitors have resolutions large enough to simply put the contents of all 256×256 displays on them.

This approach was taken by the VEPP-4 team, which exploits a large number of legacy programs using CAMAC display controllers. They made an emulation library, which redirects graphic output of such programs to X11 windows, but there was no reason for VEPP-5 to go this way.

2 POSSIBILITY

The PCI bus allows multiple videocards in one computer. One card is treated as the primary (the one on which the boot

screen appears), and others are inactive until a multihead-aware system is loaded. The AGP slot looks like just one more PCI slot.

From the very beginning, X theoretically allowed the use of several screens on one host. These screens are referred to as *hostname:N.0*, *hostname:N.1*, etc., where *N* after the colon is a display number (typically 0) and 0, 1, etc., after the dot is a screen number.

However, in practice, XFree86 up to version 3.x inclusive didn’t support multihead. That capability appeared in a long-awaited version 4.0, released in early 2000.

3 TRADITIONAL MULTIHEAD VS XINERAMA

The traditional X multihead presents each screen separately. Consequently, when a window is created, it is placed on one of these screens, and cannot span screens, or be moved from one screen to another (see Fig.1a).

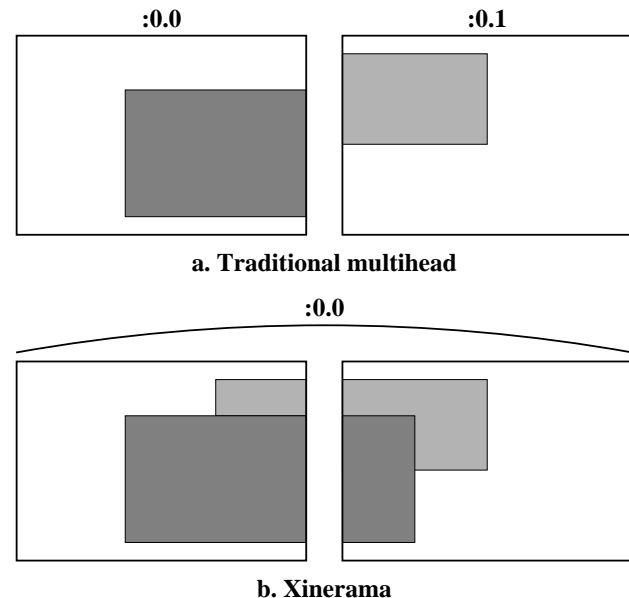


Figure 1: Traditional and Xinerama multihead

On the other hand, Xinerama makes multiple physical screens behave as a single screen, transparently to the clients (see Fig.1b) [2]. So, the windows can be freely moved between screens.

Since the situation on the screen of a control computer isn't static (there's often a need to group windows in different ways, to move more important windows to a "more visible" screen), Xinerama is much better suited for use in a control room than a traditional multihead.

4 XINERAMA PROBLEMS

4.1 Technical problems

When joining screens, Xinerama leaves only depths, which are common to all screens. So, it is impossible to join a 16-bit screen with a 24-bit one. Additionally, the 24+8 "overlay" feature of Matrox cards is lost, since the 2nd head doesn't support it.

However, the main inconvenience is that since all screens look like a single one to all clients, the window managers happily place windows between screens, maximize them on all screens, etc.

We use FVWM [4] in the VEPP-5 control room, so we invested some time in its initial xineramification, which was completed by the FVWM team (now Xinerama support in FVWM is probably the most complete and configurable among all WMs). Currently, most WMs are Xinerama-aware, but some toolkits still aren't.¹

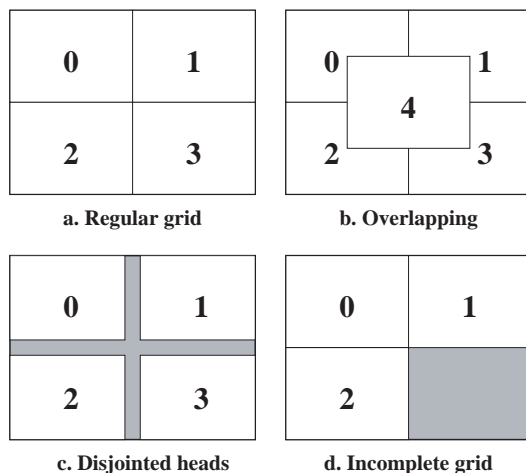


Figure 2: Possible Xinerama layouts

There is one more exotic problem. Technically, Xinerama makes a single large desktop with a size of a bounding rectangle of all screens, and screens themselves function as viewports to the desktop. So, Xinerama allows the placement of heads in many different ways (see Fig.2).

- Heads can form a regular grid – that's the most common case.
- They can overlap (a so-called "clone/zoom" mode). This is used very rarely, since the position of a "zoom" screen is fixed, and can't be moved (e.g. following the pointer).

¹A frequent case: a window with Yes and No buttons is centered, so that [Yes] goes to one monitor and [No] to the other one.

- Can be disjointed.²
- Or the grid can be incomplete.

In the two latter cases, there are "black holes" on the desktop, which aren't visible on any monitor, which can't be reached with the mouse. The consequences are worst in the last case, because complete windows can disappear in the black hole.

Additionally, there are still some problems with the software, which either requires a direct access to a framebuffer, or uses a fullscreen mode (various video capturing and movie playing programs).³ However, thanks to XVideo extension, these problems became very rare.

4.2 Human Problems

Some of our software developers are greedy: when they see so much display space, they say: "Hey, let's move this and this to another screen, and my program will just fill this screen." The common rule is "some programs tend to grow to occupy all screen space." So, the appetites of some people need reduction.

Another problem is that mouse pointers often gets lost on a large screen space. Finally, a patch for X server was developed [3], which allows 1) doublesizing the pointer and/or 2) changing the default colors from black&white to something more visible. We use red doublesized pointers, which provides good visibility.

5 HOW MANY HEADS TO USE

The most common multihead layout in the world is two heads: side by side horizontally, or one above another (if 2nd head is used rarely).

Three heads are hard to use: the layout will either be as Fig2d, which is inconvenient, or lined up. In the latter case it takes too much time to move the pointer between the first and the last screens. When we *had* to use three heads, we put monitors in the shape of an "r," but the X layout was "three heads vertically." That setup was extremely confusing for operators.

Four heads give the best balance between the "as much screen space as possible" principle and convenience of use. When used in a 2×2 grid, as shown in Fig.2a, there are no black holes, and the distance between heads is small.

6 HARDWARE

6.1 Criteria for selecting video cards

First, hardware should be multihead-capable (e.g., 3Dfx cards are known not to work in multihead mode under XFree86 at all). Second, it must have good support in XFree86 and be very stable. Third, it should produce an

²That's a pathological case; more often screens of different sizes are used (e.g. 1024×768 and 800×600), which has the same effect.

³Up to XFree86 4.1 Xawtv behaved very funny: the window frame could be on one screen, and undecorated video picture – on another.

excellent picture, and have good 2D performance (3D isn't important). Fourth, video hardware should occupy as few PCI slots as possible.

6.2 Solution we use

There were 4 main manufacturers: ATi, Matrox, nVidia and S3 (the latter is almost dead now). We chose Matrox because it satisfied all our criteria, and we already had very positive experience with their products.

The product nVidia doesn't provide specifications for their cards, so the XFree86 driver is very lacking. Instead nVidia provides a binary-only driver. ATi cards are not-so-good, and there are myriads of subversions, which affects stability of the driver.

We use MillenniumIIs as PCI cards, but any PCI card with 4M or more memory⁴ would do (Millenium, G100, G200).

6.3 Multiheaded videocards

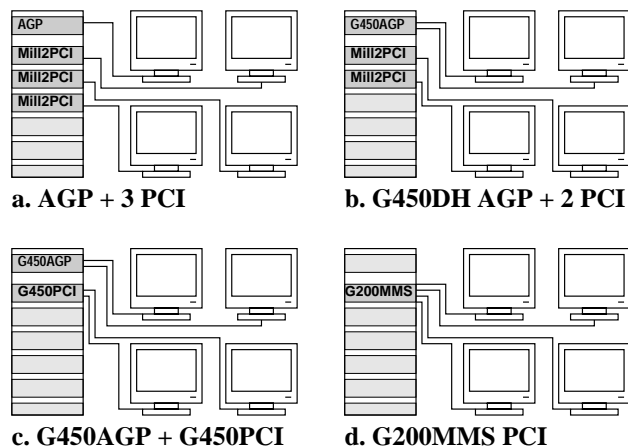


Figure 3: Hardware options for 4 heads

One more advantage of Matrox cards is that since 1999 they have two heads on one card (G400DH, G450, G550). So, to have 4 heads, an AGP G450 plus two PCI MillenniumIIs were enough (see Fig. 3b), and only two PCI slots are used (which are always in deficit in control machines).

Matrox produces a PCI version of G450, which allows the use of only one PCI slot in addition to the AGP slot (see Fig. 3c). Unfortunately, when running as a non-primary card, the G450 (either AGP or PCI) requires an additional driver module, which is available as binary-only from Matrox (so-called HAL module [5]). However, we hope the native XFree86 support will soon become better, thus making HAL redundant.

Currently two of our 4-headed PCs are equipped with G450AGP+2×MilleniumII, and two are G450AGP+G450PCI.

Theoretically, there exists even a better choice – G200MMS, which supports 4 heads on one card (see Fig. 3d). However, it exists in a PCI version only, so if

we need 4 heads total, it occupies the same one PCI slot as the G400AGP+G450PCI. Additionally, G200MMS is almost impossible to find in Russia.

6.4 Motherboards

We chose the ASUS P3B-F (Intel 440BX chipset), which has 1 AGP slot, 6 PCI slots, and one ISA slot (one position is shared). So, you get either 6 PCI slots, and 0 ISA slots, or 5 PCI slots, and 1 ISA slot. The main requirement was a presence of an ISA slot, since we still use old ISA hardware.

7 FUTURE ENHANCEMENTS

One feature our operators wish to have is the ability to control programs on adjacent computers with their mouse. A program *x2x* [6] exists which does exactly this, but it doesn't work with Xinerama. So, we plan to “xineramify” *x2x*.

Currently, the world is moving toward the use of TFT monitors, as they are safer for people. However, most TFTs have a limited viewing angle, which is inappropriate in multiheaded system, and those TFTs which are okay (like SGI 1600SW) are too expensive. So, currently we use 17" CRT displays, but plan to replace them when affordable TFTs appear.

Finally, consistent and ergonomic placement of windows on a 4-monitor desktop is a time-consuming task. So, we are planning to implement some sort of automation for this. Currently, we are experimenting with the X resource database (the *WINDOWNAME.geometry* resource).

8 REFERENCES

- [1] G.Piskunov, “CAMAC-embedded 24-bit computer”, *Au-tometriya*, N4 (1986) pp. 32-38 (in Russian).
- [2] Dennis Baker, “Using the Xinerama Extensions to MultiHead XFree86 V. 4.0+”, November 2, 2000. <http://www.linuxdoc.org/HOWTO/Xinerama-HOWTO.html>
- [3] D.Yu. Bolkhovityanov, “Visible Cursor Patch for XFree86 4.0.3”, <http://www.inp.nsk.su/~bolkhov/files/bigcursor/>
- [4] Official FVWM Homepage, <http://www.fvwm.org/>
- [5] Matrox Graphics, Inc., “Matrox beta drivers for XFree86”, <ftp://ftp.matrox.com/pub/mga/archive/linux/2001/beta.133.143/>
- [6] David Chaiken, “X to X connection”, <http://gatekeeper.dec.com/pub/DEC/SRC/x2x/>

⁴1152×864 @ 32bpp ≈ 4M RAM for framebuffer.