Abstract

The `xps` program dynamically displays the Unix processes as a tree or forest in an X Window, the roots on the left and the leaf processes (those with no children) on the right. The status of each process running, sleeping, stopped, etc., can be indicated by different colors. Different users can appear as different colors too.

Process selection can be made per user, all users, or through a regular-expression pattern.

In contrast to the terminal-based `pstree` or tree-widget based programs, the tree display uses diagonal lines, and effort is made to effectively use the full 2-dimensional area of the screen by balancing levels and centering the children of a node between their parent. A goal of the program is to give an idea of what's going on graphically as things may be constantly changing. Therefore the display algorithm tries to keep processes close to their parents to reduce the amount of scrolling to see localized process creation and destruction. Some effort is also given to make sure that the tree layout doesn't get wildly reorganized when there are small or localized changes. This makes it easier for the eye to pick up and recognize the changes over a potentially large display area.

We describe here criteria for tree animations such as this one and how the `xps` layout algorithm works.

There are some other miscellaneous features of `xps`. One can select viewing the processes by a single user, a regular expression for users, by all users, and perhaps show kernel processes. One can click on a process to get more information (via `ps` or a user-specified program) about that process, send a signal, or set the process priority, assuming you have the permission to do so.

Since programs of this ilk can consume a bit of CPU on their own, some effort has been made to turn off the update process when the program is iconified or not visible for some other reason such as being obscured by another window. Some attention has been paid to make algorithm display fairly fast in most situations, although it has to be admitted that this comes sometimes at the expense of a nicer layout.

The program is available from:
`ftp://netwinder.org/users/r/rocky/xps.tar.gz`
home page:
`http://www.netwinder.org/~rocky/xps-home`

What is xps? Why xps?

There are a number of front ends or GUIs that perform various underlying commands, and sometimes one wants to see what's going on behind the scenes.

Some examples include

- learning more about what is going on in configuration and installation processes; `make` or GNU `configure`, for example
- tracking down zombie creation
- understanding the processes for a user or contained in a process group; these generally cluster into subtrees
- helping distinguish a process via its lineage. For example one may have many `bash` sessions running. Some may be spawned from `sshd`, some from `telnet`, some via `emacs` or an `xinit` session.

To glean what's going on, many systems administrators use `ps` or `top` or a `top` variant, like `gtop`. The `ps` program can list the process id and its parent process id. However it is sometimes useful to display the parent-child relation graphically.

Considerations for doing a tree layout for xps: tree animations

`xps` uses a custom tree-layout algorithm for positions nodes and draws lines between them using low-level Xlib calls. Probably each time a new graphics library came out I considered ditching the custom tree-layout routine with a generic tree package. However to date, I've not found an acceptable one.

Tree widgets are often used for browsing file systems or menu systems. Generally these objects are fairly static. There is generally a way to collapse or expand a branch in the tree; this makes sense when information doesn't change all that much. But in the context of `xps`, sometimes the new or changed information is precisely what you want to see. So instead, in `xps`, symbolic or filtering rules are used to focus the display, rather than by zooming in or out from a user-selected tree branch. In particular, one can select the area of interest by a regular expression which is matched against the process-owner name.

In trees with menus or files, the display is pretty much static, and therefore the time users spent to customize the display may be cost effective; it may not be as important to spend time initially making connections between tree objects immediately visually distinct. Convention or experience seems to indicate that it is acceptable to use a Manhattan-metric line connection style, that is, where tree lines either go horizontally or vertically. See Figure 3 for a tree layout using Manhattan-metric lines.

The tree-layout requirements of `xps` are perhaps a bit more difficult to satisfy. In fact `xps` might be better thought of as an animation rather than a static charting program, and I am not aware of much literature on doing animations with trees or forests.

For an analogy, consider the differences between viewing a still-life painting and an animated cartoon. The cells of an animated cartoon do not have to be as perfectly rendered as a painting. Instead it is more important in an animation to make the *series* of cells relate to convey a story or action.

Below we gives a simple example of the kinds of problems faced.

In a world where things don't change, like a figure in a book, most people will find it most it pleasing for graphs with one or two children to look as in the left-hand part of Figure 1 rather than the left-hand part of Figure 2.

Figure 1:

However, now consider the case where things are constantly changing. Suppose we want to chart what is going on when we have a process that forks a process, `child1`, and then forks a number of other processes sequentially. For a moment there will be `child2`; it dies, and then suppose a moment later a new processes, `child3`, is spawned. Suppose that now runs for a short period of time and dies; so we are back to the single `child1`; then a new process `child4` starts up and so on.

In this scenario, if you render the one or two child processes as in Figure 1, you will see a bit of flicker. This is distracting and does not assist comprehending what's really going on: that a new process is being started and finished while nothing is happening to `child1`.

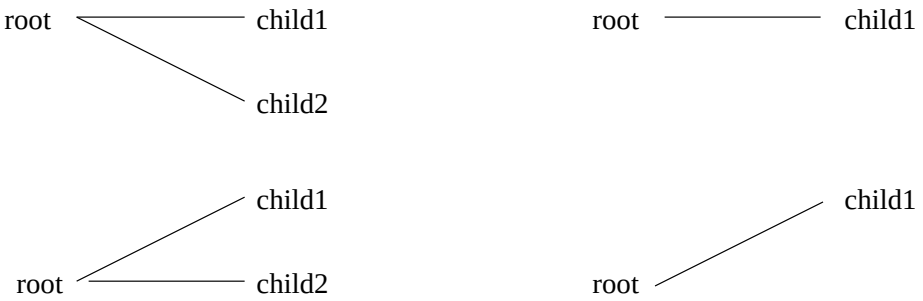

Figure 2:

If instead the graphs were rendered as in Figure 2, the display is more pleasing to the eye *in animation*; it is less spastic and more comprehensible even though the left-hand or right-hand side in isolation may not be the most esthetic rendering of the graph.

In sum, hysteresis of layout can help visualization.

Now consider the differences between a Manhattan-metric layout of the variety one often sees with a tree widget, and one that uses diagonal lines in an artificial but not uncommon situation as is often seen in `xps`.

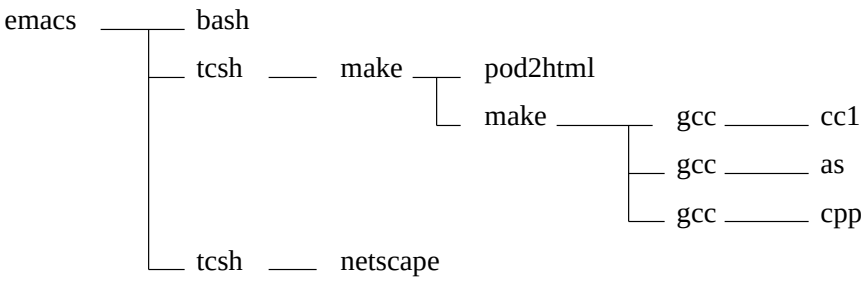

Figure 3:

Notice in Figure 3, how the root of the tree, `emacs`, is very far removed from its bottom-most node. This is not a problem with the Manhattan-metric connection style per se. However almost all most tree widget programs do this kind of layout; to set the position of a node, the layout algorithm can be extremely simple since it doesn't have to consider the positions of the children.

In addition to the lopsidedness, there's a bit of space that is wasted between `make` and `netscape` that doesn't appear in Figure 4. Furthermore, the blank area to the right of `bash` and `pod2html` is reclaimed reducing the overall dimensions of the graph. However, the more-compact layout is at not at the expense of readability.

Finally, compare Figure 3 and Figure 4 to see how the use of diagonal lines helps readability.

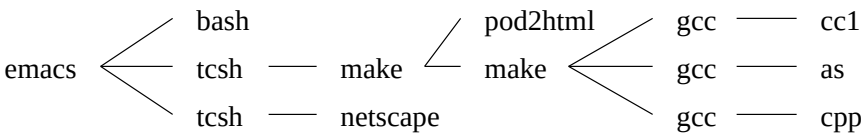

Figure 4:

Experience has suggested these display criteria for a tree-animation program such as `xps`:

- the layout needs to be fast — it is performed every second
- the layout should be compact

- the layout should have hysteresis — reduce flickering which is annoying to watch and makes finding important changes hard to spot
- the layout should be "pretty" on the display so users can understand the relationships quickly

We now go into some of these needs in more detail.

Fast layout:

In designing any monitoring program, such as `xps`, a cardinal rule should be: don't trash the system you are trying to monitor.

Because layout needs to be fast and work for a large number of nodes, positioning decisions should be pretty much linear. Currently one pass from root to the leaves is made. I've considered making a backward pass as well—from leaves to roots.

The program has a number of hacks to avoid layout recalculation when it is not needed. It checks to see if processes have changed; if not, nothing is changed. However this is tricky since `xps` shows the status of each process (e.g., whether it is running or sleeping or stopped), and if this alone changes some redisplay of the node color (not position) is needed. If the display becomes iconified or fully obscured, no layout is computed.

Pretty layout:

Experience has shown that it is preferable to draw straight lines for parent–child connections rather than diagonal lines since diagonal lines take longer to draw and can appear jagged unless antialiasing is used. (Antialiasing also generally takes more time to compute and draw.) On the other hand, when needed, I find diagonal lines are easier to follow than Manhattan-metric lines.

Similarly, I've found that lining things up horizontally and vertically helps.

Layout heuristics:

The key to a program like this is its tree-forest display algorithm. In the previous section we described desirable qualities; in this section we'll give a rough idea of the actual heuristics used. In the next section we go over how this is implemented and the time complexity of the layout algorithm.

Trees tend to be narrow at the root and get bushy as one moves towards the leaves. Therefore the approach used in `xps` is put nodes into fixed levels (which run horizontally) as we move from the roots to the leaves (left to right here).

The algorithm we use pretty much makes one pass to keep things fast. We keep track of the maximum number of nodes at a given level as we move from roots

to leaves. As the breadth increases, the gap between levels decreases up to a minimum threshold. When things get too tight, the overall dimensions of the graph increases; in display, the tree breadth is shown in the vertical direction.

Figure 5 gives an example of how levels are redivided as we move deeper in a tree.

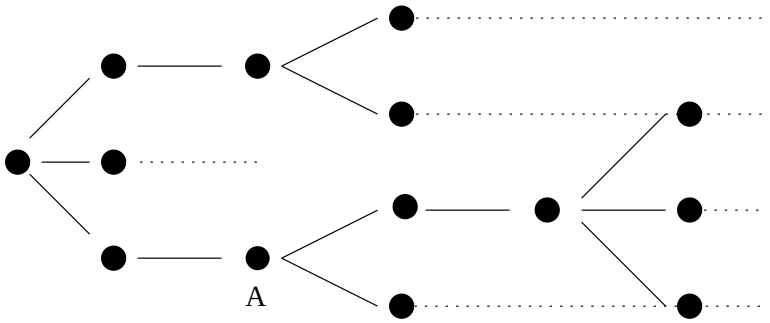

Figure 5:

Putting things in levels increases the likelihood that parent-child nodes will line up and that they will stay lined up over time. Actually, although it might be useful to have `xps` take into account old positions of node as discussed in a previous section, right now it doesn't; it is just an artifact of the way the layout occurs that this tends to happen.

After creating the levels for a given depth, the dotted lines in Figure 5, `xps` next positions the node for that depth. To this end, the program keeps track of its level position or *virtual rank* in that depth, and compares it to the virtual rank of its parent. The virtual rank differs from the number of nodes placed so far (or *rank*), when we've decided to leave a level slot empty so as to position a child closer to its parent. In Figure 5, the node marked "A" has a virtual rank of 3, while its rank is 2. This is because we've moved the node down to its level slot 3 so it will line up with its parent.

Note that since the number of levels only increases, and when this happens inter-level space between level generally decreases, if the virtual rank of child is less than or equal to the virtual rank of its parent, it can be positioned no further down than its parent.

When setting the position of a node, `xps` tries to position the virtual rank of the middle child close to or less than the virtual rank of the parent, but not so much that the overall breadth is increased.

Figure 6 shows what goes on here. In the upper-left diagram we show the virtual ranks and before readjustment; to the right of that, after the readjustment at the new depth. In the botomm-left we see what happens when the new level is attached to a node further down; since centering the children would increase the overall breadth and dimensions of the tree we do not center around the middle child. The bottom-right graph show a heuristic not employed by `xps` but might be if a backward pass were made: readjusting the parent.

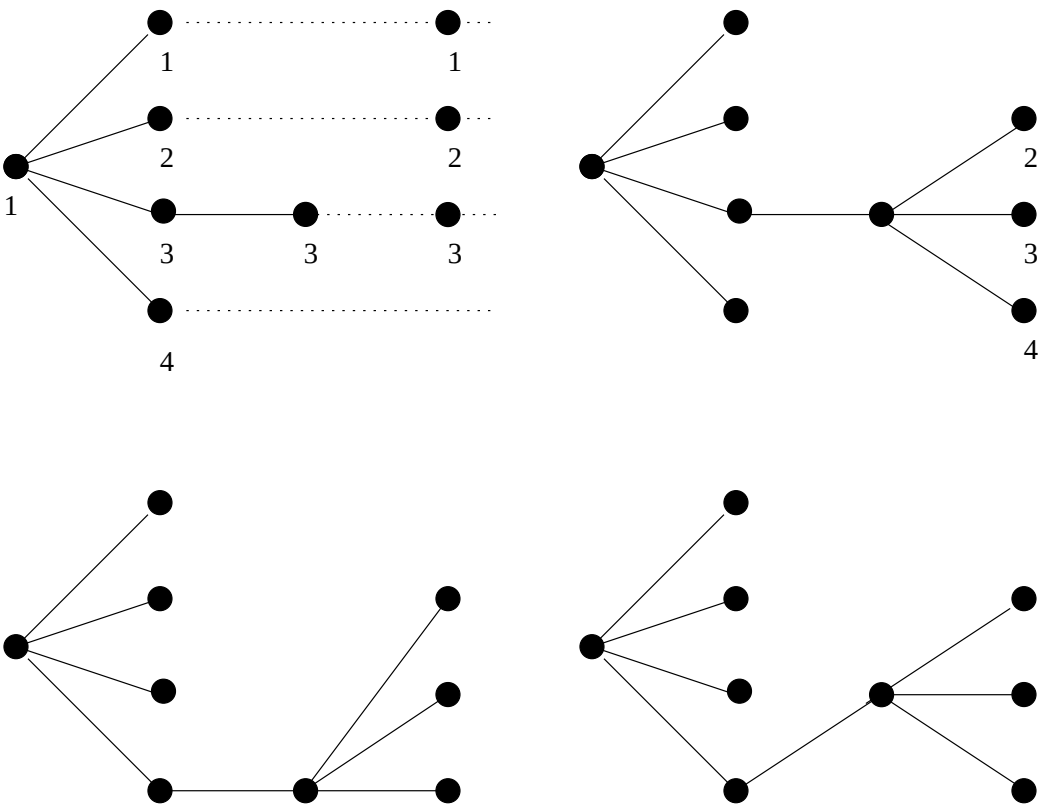

Figure 6:

There is one layout problem that should be mentioned and is shown in Figure 7.

We see here that in drawing lines from the end of a node, nodes with a short name can sometimes cross over a neighboring node name. Currently `xps` does not try to prevent this from happening. We have experimented with doing a layout as in Figure 8 without much success. Most of the time, though, the lines do not go through process labels.

Layout Algorithm implementation

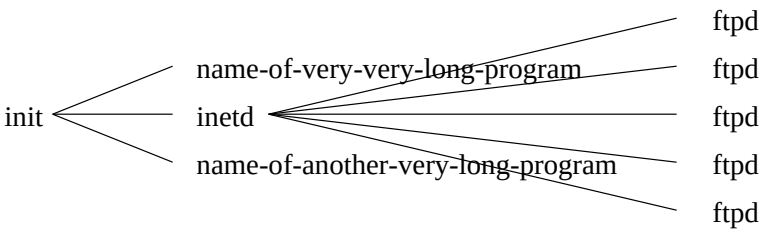

Figure 7:

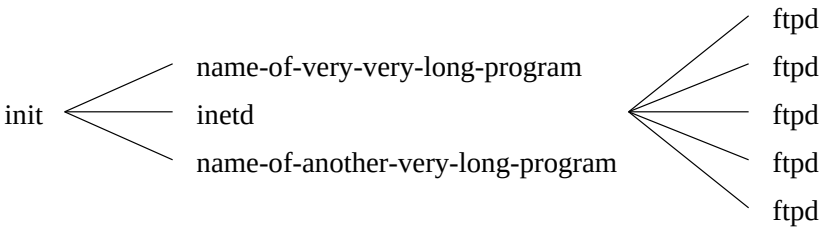

Figure 8:

The program needs to do a topological sort. A depth-first search is done to assign depth number.

The roots are then sorted by uid and pid. There is generally a small number of roots, often one. At depths deeper than the root, nodes are first arranged by attachment to a parent.

Within the children of a node, sorting is done by a bubble sort. Many people with computer science training cringe when they hear this, because they have been taught a bubble sort takes $O(n^2)$ time in the worst case, while there are many sorting algorithms that take $O(n\ log\ n)$ time in the worst or average case. However, a bubble sort has an advantage our most sorting algorithms: when used on *almost* sorted data, the running time is linear.

In the context of `xps`, the processes are often pretty much sorted by process id. The bubble-sort code we use does a check to stop early when the items are fully sorted. Also, it should be noted that the number of items to be sorted is often relatively small, since the sort is performed only on children of a node. A bubble sort generally has less overhead than most other sorting programs. Still, a check on the number of children should be made and a sort like quicksort could be used if the number of children exceeds some threshold like 40 (which is in my experience rare).

The sorting to arrange things into levels by parent is $O(npc^2)$ for n nodes, p non-leaf nodes, and where c is the maximum number of children a parent has.

This is a rough analysis and looks unpromising, but in practice is probably pretty good. I suppose a radix sort could be done on the levels. Then we'd have $O(n + pc^2)$.

Still someone might want to time various algorithms. The program is not linear. Faster tree-layout and arranging algorithms would be helpful in handling larger trees.

Other niceties of `xps`

`xps` does a sort by user id and by process id within that. This tends to group related processes together and the process id arranges things by age. It has been suggested that if it is desirable to sort precisely by age, that should be done instead of sorting by process id.

`xps` has the ability to filter out processes by user id or userid–regular expression. Each user is assigned a different color. Reducing the amount of display has a two-fold benefit: it not only unclutters the display but it also reduces the amount of time needed for display.

Finally, `xps` has the ability to point and kill. (The interface hasn't been developed to the point of arcade games, but still it might give the systems administrator a heightened feeling of power; in contrast to mass kill programs like `skill`, it can be satisfying to see the process die before your very eyes.)

A screenshot from the program is given in Figure 9.

Future?

Just about everything could be improved. The distribution contains an extensive list of things to do. Above were some suggestions for how tree layout might be improved.

The thing I would most like to see is `xps` ported to the KDE `qt` and the GNOME `gtk+` and `glib` libraries. Any volunteers?

Display is done by drawing on a canvas. This is primitive. Process names are not widget labels, just X strings drawn on a canvas. Figuring out the process under the mouse and showing which process is selected is a bit low-level and not tool-like. Currently, a horizontal and a vertical linear search is done to find the process id under the mouse. Binary search might be faster.

Alternatively, if a full-fledged widget for the nodes of a tree were used, then X (or an X-toolkit) would worry about when the widget is selected; presumably it uses an algorithm at least as good as binary search. Making the selected node look, well, *selected* would then be done by the toolkit.

It would be nice to benchmark the various toolkit approaches for efficiency.

One might experiment to compare speed and intuitiveness of using a canned tree widget versus a hacked layout customized for this application. Personally, I prefer the tree layout, but the code is not toolkit idiomatic.

Currently, the program contains coordinates of nodes before its position gets updated. However nothing else related to the history of the layout is saved, such as how long a node has been in the same position. It might be interesting to experiment with algorithms which make incremental improvements over time using say local heuristics. Analogous is the splay-tree algorithm which tends to make a tree balanced over time by making small localized changes as nodes are accessed.

Thanks

Derik Lieber wrote the original version of this program.

Many people have read and made helpful suggestions on this paper including Stuart Frankel, Ph.D., George MacDonald, and Mike Welles.

Related work

`pstree` by Werner Almesberger found on Linux distributions shows a static Manhattan-metric process tree relation using character-oriented graphics.

Another program named `pstree`, by Lars Christensen, does about the same thing and runs on other Unixes. `ftp://ftp.thp.Uni-Duisburg.DE/pub/source/`

I. Herman, G. Melançon, M. S. Marshall "Graph Visualisation and Navigation in Information Visualisation" can be consulted for a survey on graph visualization and navigation techniques, used in information visualization. For example it cites the classic paper on tree layout, E.M. Reingold and J.S. Tilford, "Tidier Drawing of Trees," IEEE Transactions on Software Engineering, SE-7(2), pp. 223-228, (1981). However a number of subsequent papers including those found in the survey given above question the tree heuristics used in this paper. For `xps`, I have not found this paper all that useful, although it is an interesting read.
`http://www.cwi.nl/InfoVisu/Survey/StarGraphVisuInInfoVis.html`

Finally George MacDonald's `treeps` program is similar. A home page for this is `http://www.slip.net/~gmd/tps/treeps.htm`

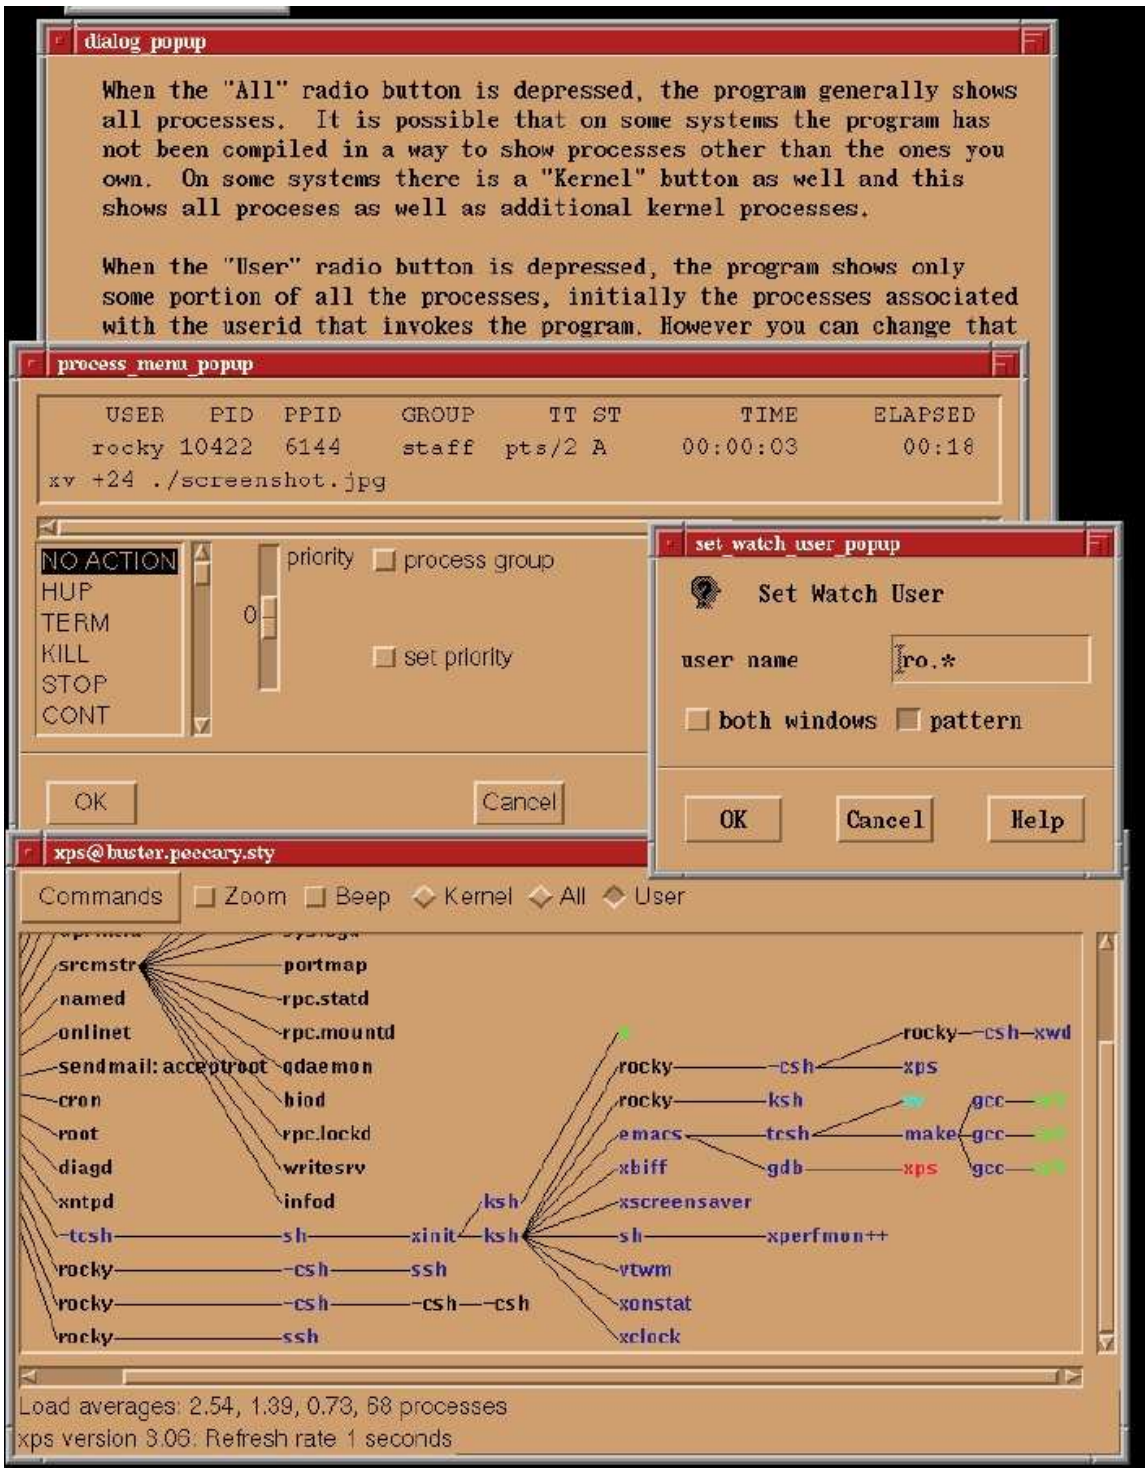
dialog_popup
When the "All" radio button is depressed, the program generally shows all processes. It is possible that on some systems the program has not been compiled in a way to show processes other than the ones you own. On some systems there is a "Kernel" button as well and this shows all proceses as well as additional kernel processes.
When the "User" radio button is depressed, the program shows only some portion of all the processes, initially the processes associated with the userid that invokes the program. However you can change that
process_menu_popup
USER PID PPID GROUP TT ST TIME ELAPSED
rocky 10422 6144 staff pts/2 A 00:00:03 00:18
xv +24 ./screenshot.jpg
NO ACTION
HUP
TERM
KILL
STOP
CONT
priority
process group
0
set priority
OK
Cancel
set_watch_user_popup
Set Watch User
user name
ro.*
both windows
pattern
OK
Cancel
Help
xps@buster.peecary.sty
Commands
Zoom
Beep
Kernel
All
User
srcmstr
named
onlinet
sendmail: acceptroot
cron
root
diagd
xntpd
portmap
rpc.statd
rpc.mountd
qdaemon
biod
rpc.lockd
writesrv
infod
tcsh
sh
xinit
ksh
rocky
csh
ssh
rocky
csh
csh
csh
rocky
ssh
ksh
rocky
csh
xps
rocky
ksh
emacs
tcsh
make
gcc
gcc
xbiff
gdb
xps
xscreensaver
sh
xperfmon++
vtwm
xonstat
xclock
Load averages: 2.54, 1.39, 0.73, 68 processes
xps version 3.06. Refresh rate 1 seconds

Figure 9: