

Seems Like Old Times

Michael D. O'Dell Editor-in-Chief

Greetings!

This issue of *Computing Systems* is very special: it is an all-Research edition. (By “Research,” I mean the research group at AT&T Bell Labs where UNIX was born.) Once upon a time, the *Bell System Technical Journal* and its special “UNIX editions” were an immensely valuable conduit between the people within Research and the outside world. While we certainly won’t claim the same level of importance (if for no other reason than times have changed and more information is more widely available), we are very pleased to have this collection of papers, all by authors who should be quite familiar to the community.

Brian Kernighan and Christopher Van Wyk are names known by any serious *troff* user. Their paper, “Page Makeup by Postprocessing Text Formatter Output,” presents their approach to the quite difficult problem of page layout. Of particular interest is that their paper was not typeset by the usual USENIX personnel; rather, it was submitted as camera-ready copy, output directly by the program which the paper describes. The only thing we did was supply them with page numbers and answer some stylistic layout questions. I believe the appropriate citation is: *Quod erat demonstrandum*.

Rob Pike’s paper, “A Concurrent Window System,” proposes an alternative to the event-based programming model presented by most popular window systems. Since, from one point of view, events provide a way to implement a limited form of internal multiprogramming, he proposes an environment where the usual event-processing work is encapsulated into a group of interacting concurrent processes. This approach yields a window system

which is modular, efficient, and above all, simple and comprehensible. In these days of standardized complexity, this step back to reconsider the underlying issues is both refreshing and revealing.

Rounding out this issue, we have two papers on computer viruses: Tom Duff's paper, "Experiences with Viruses on UNIX Systems," and Douglas McIlroy's paper, "Virology 101." Both discuss what is clearly a rather sensitive topic: the innards of working viruses and how they are fundamentally part of the computing landscape. While these papers include the expected admonitions like "Don't try this at home, kids," it is likely that someone will take issue with their publication because "it might lead to someone trying them." My response to this position is the following analogy.

Killing another human is trivially simple. One needn't use a handgun; there are adequate poisons readily available down at the garden center, a brick or a crowbar to the skull works well (if a little messy), and one can always use an automobile as a lethal weapon, particularly after an excess of imbibing. So, I argue that discussing the techniques available with which to do-in someone does not make one an accessory to murder. Further, a paucity of information about how to commit murder isn't a credible prophylactic, either. Such things are too easy to discover. The reason why a normal person doesn't go around murdering other people is not because he hasn't the means or the knowledge; rather, it is because he subscribes to the belief that the act of murder is morally repugnant and **WRONG**. Likewise, the reason a computer scientist does not commit a virus is not because it is forbidden knowledge, but because he believes the act is morally repugnant and **WRONG**.

There is a good case to be made for very carefully controlled experiments to evaluate the level of system vulnerability and the efficacy of potential countermeasures, but the nature of this work parallels biological virus research. Biological virology experiments are carried out in special laboratories whose sole purpose is the containment of potentially virulent organisms. If software experiments need to be done (and they do seem necessary at the moment), we must develop the software equivalent of a P4 Containment Laboratory. Neither biological virus nor software virus experimentation can be done safely in a basement laboratory.

Messrs. Duff and McIlroy will attest to this, first-hand. (It seems there might be interesting research in that direction.)

That's the papers in this issue. And now for my usual plea: please take five minutes to consider what you and your associates are doing, and then submit a manuscript. It doesn't hurt much, the availability of your work will benefit the UNIX community and seeing your name in print is so gratifying.

Until next time...