

Editor's Message: Ogres, Villagers, & the Art of Writing

I read and evaluate about a hundred papers every year, mostly in the area of computer systems. This experience has convinced me that, with a little care, authors of systems papers could be far more effective in communicating their ideas (of course, this would also increase the likelihood of publication). I would like to share some ideas on the art of writing effective systems papers.

An effective systems paper tells an age-old story that I call "Ogres and Villagers." Consider the troubles of a village that is attacked by ogres descending from the hillside every night. The villagers build fences to defend themselves but the fences are inadequate and the villagers suffer the consequences. You, the writer of the paper, have to come up with an ogre-proof fence. You should figure out what is needed to do the job, you have to build it, and then you should test it to see how well it works. The ideal outcome is that your fence keeps the ogres away and the villagers live happily ever after.

The purpose of your paper is to tell this story clearly, convincingly, and thoroughly. First, make sure that we (the readers of this story) know who the villagers are and who the ogres are. Don't assume that we know an ogre when we see one. In the related work section, tell us what attempts the villagers have already made to defend themselves. Convince us that these solutions don't work. If possible, analyze the reasons why other solutions have failed. Maybe the fences were too flimsy. Maybe the perimeter of the village is sandy allowing the ogres to burrow underneath. The more thoroughly you set the stage, the more convincing your solution.

Second, tell us about your solution. What makes it ogre-proof? What materials did you use? How much did it cost? How long did it take to build? Did you need to use experts, or can it be built by anyone? Did you use pre-fabricated materials from elsewhere? Why or why not? Did your solution borrow elements from prior attempts? Describe the solution in sufficient detail to enable other villagers to build their own fences.

Third, tell us how well your solution worked. Did it stop the ogres? Did it at least delay them for a while, giving the villagers time to hide? Or did it fail? And if so, why? This is also the chance for you to tell us the best way to measure the "ogre proof-ness" of a fence. Should it be measured by dollars spent per ogre stopped? Or by the mean ogre delay time? Or by the increase in the villagers happiness per dollar spent?

Make sure that your conclusions are statistically valid. Showing that the defence worked adequately for one night is not as convincing as showing that it worked every night over a period of two years. Also, be careful not to overstate your claims. Do not tell us that your fence protects against trolls and giants if you only tested it against ogres.

It may have been too expensive to build the fence of your dreams, so you may have simulated it. If so, make sure you convince us that your simulation is accurate. Does it correctly model an ogre's behaviour? Does it correctly model the physics of a fence? If possible, compare the prediction made by simulation with the actual behaviour of a fence in the field and show that the prediction is accurate.

A paper is poorly written if you fail to tell this story convincingly. If you jump right into the solution without setting the background, we may not even begin to understand what the story was about. If you write only about your own solution without describing other attempts, we may suspect that ogre-proof fences already exist, and that you are merely re-inventing the wheel. If the metrics you used for evaluation are obviously the wrong ones - measuring the mean thickness of the fence, for example, rather than its effectiveness - we may find the paper unconvincing. The evaluation in many systems papers leaves much to be desired: describing a single hour's observation of a mockup fence tested against villagers dressed up as ogres! Such poorly-written papers frustrate their reviewers who retaliate with low evaluation scores and summary rejection. In my experience, the common complaint that a reviewer did not understand a paper is usually because the authors failed to describe the solution sufficiently.

In contrast, well-written papers are a pleasure to read. The characters shine through, the need for a solution is obvious, the fence is succinctly described, and the paper leads the readers through a detailed evaluation of the performance of the fence. These are papers that are impossible to reject.

I have focused on systems papers, which is the work I know best. However, I suspect that the Ogres and Villagers story might be applicable to other areas as well. I look forward to hearing from readers about whether or not that is the case.

S. Keshav
CCR Editor