
Subject: shmmap questions

Posted by [Russell Ryan](#) on Mon, 10 Dec 2012 06:14:13 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi gang,

I'm working on a Markov Chain Monte Carlo code to model some images, specifically using the parallel tempering algorithm. Since this algorithm is very computationally expensive, I'm using the IDL_IDLBridge to multithread the process (basically running several chains in parallel). But to compute the likelihoods I need a large array of model luminosities, which if I simply passed to each bridge I'd quickly run out of memory. So now pass the data using shared memory, and shmmap. Everything seems to work when the array of model is smallish, but the code gives the message:

% SHMMAP: Existing file too short for desired mapping.

if I make the model array large. I can't really say how large "large" is, but I wondered if this wasn't related to the post on D. Fanning's site: http://www.idlcoyote.com/code_tips/bridge.html . I'm doing this on Mac OSX, and I noticed that the shmmax variable in the kernel was set to 128 Mb. I raised that level to 1 Gb using the command: "sudo sysctl -w kern.sysv.shmmax=1073741824". I still get the error, and I'm not sure where to go next. Any ideas?

Any help is greatly appreciated!

Russell
