

Recitation 7: VM and Malloclab (pt.1)

Friday, October 10th

Introduction

The goal of this activity is to gain a better understanding of TLB benefits, using the lens of locality of access and size of working set. Each group make an analysis of the implications that the scenario has on the TLB or page activity.

Scenario

Large Data Computing: We are running a large-data computation task, processing data on the magnitude of terabytes. Suppose we have a reasonably good, regular access pattern to data, as well as a reasonable page size (eg. 4KB).

Discuss the features of the workload. We simply present them for the purpose of this recitation, but make sure to briefly discuss in your group about why these features are true!

Workload Features: Good spatial locality from accessing, but huge working set across the lifetime of the program. This is all running on one process so there are minimal context switches.

Questions to Answer

Using the categories of locality of access and size of working set, try to see how the scenario does with respect to each category. Think about:

1. Is there good locality in the scenario? What is the total size of the working set (large or small)?
2. What do the observations in part 1 indicate about TLB behavior? What implications do they have?
3. What changes can we make to improve the benefits we gain from TLB or avoid the pitfalls of TLB? (eg. cache design, page size, ect...)

It may be helpful to organize your thought in the table provided:

Factor	Does the TLB Help?	Why or Why Not?	What are some improvements we can make?
Locality of Access			
Size of Working Set			