WSO 2006

3º Workshop de Sistemas Operacionais

Campo Grande – MS

17/julho/2006

ESCOLA POLITÉCNICA DA UNIVERSIDADE DE SÃO PAULO

New Visual Characterization Graphs for Memory System Analysis and Evaluation

Edson T. Midorikawa
edson.midorikawa@poli.usp.br
Hugo Henrique Cassettari
hugohc@terra.com.br
LAHPC - PCS
Escola Politécnica da USP

Topics

- Introduction
- Previous work
- New Proposed Graphs
- Case study: performance of adaptive page replacement algorithms
- Conclusion and Future works

Introduction

- Memory systems (still) are one of the most important components of modern computer systems.
- They work based on the *Principle of Locality of References.*
- Workload characterization is essential for memory management research.

Previous Work

- WSO 2004
  - **LRU-WAR** (*LRU with Working Area Restriction*)
- WSO 2005
  - **3P** (*Three Pointers*)
- Performance comparison among algorithms:
  - FIFO, LRU, Clock, …
  - FBR, 2Q, …
  - SEQ, EELRU, LIRS, ARC, CAR, CART

A simple question

Why algorithm **A** has
better performance
than algorithm **B**
for program **P**
under condition **C?**

Previous Work

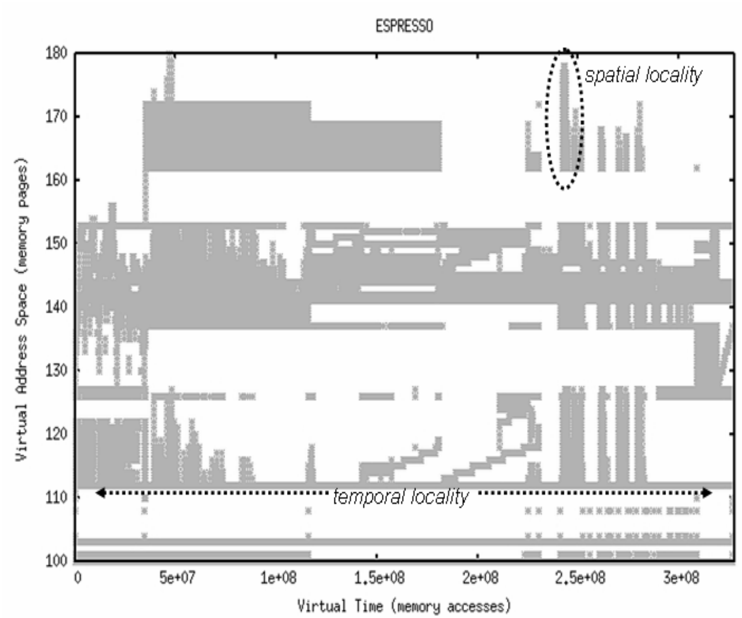

Memory access map

Locality surface

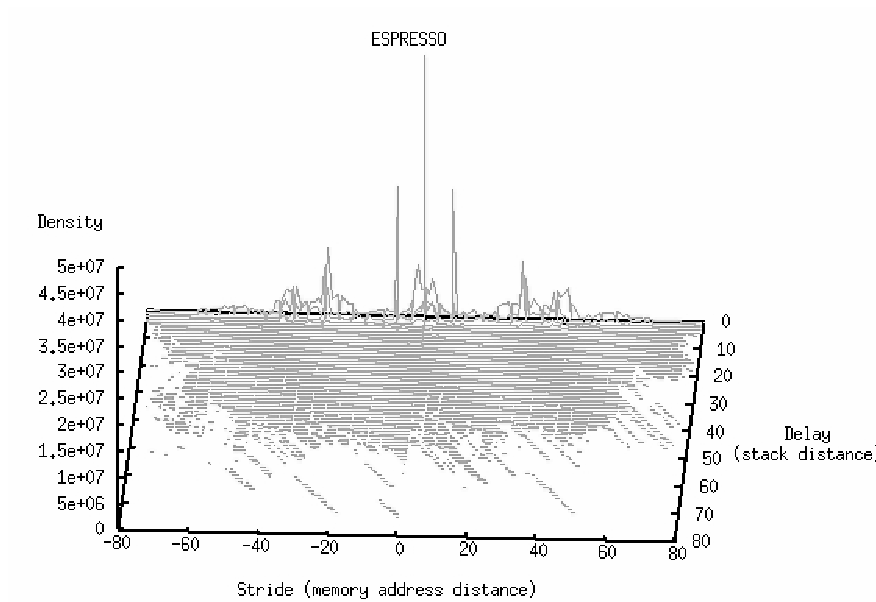

Objective

- Introduce new visual resources to study memory access behavior, workload characterization, and performance prediction.
- Present a case study with the proposed visual characterization graphs.

New Proposed Graphs

- We introduce 6 new visual characterization graphs:
  - Memory access surface
  - IRG graph
  - IRR graph
  - IRR surface
  - IRR histogram
  - IRR cumulative graph

Case Study

- We present a case study of *performance evaluation of page replacement algorithms* for modern virtual memory systems
  - *Traditional*: FIFO, LRU, FBR, 2Q
  - *Adaptive*: ARC, EELRU, LIRS, LRU-WAR
  - *Online*: Clock, CAR, CART, 3P

- Applications (trace files):
  - *Compress*: file compression utility.
  - *Espresso*: a circuit simulator.
  - *Grobner*: a formula-rewrite program.
  - *Sprite*: from the Sprite network file system.

Case Study

- Simulation results

Trace	Average Percentage Increase in Page Faults from the Optimum Case											
	FIFO	LRU	FBR	2Q	ARC	EELRU	LIRS	LRU-WAR	CLOCK	CAR	CART	3P
Compress	191.97%	121.07%	98.24%	94.79%	120.81%	72.63%	122.41%	79.46%	145.99%	127.91%	103.95%	94.66%
Espresso	174.85%	67.73%	62.88%	80.27%	68.19%	52.24%	66.84%	66.88%	114.96%	114.14%	102.99%	80.60%
Gnobner	233.96%	139.12%	115.66%	116.53%	139.17%	72.91%	113.86%	125.40%	139.20%	139.34%	118.21%	108.06%
Sprite	33.29%	17.70%	17.90%	25.21%	22.93%	19.04%	40.72%	21.36%	18.67%	23.68%	26.27%	25.13%

Memory access surface

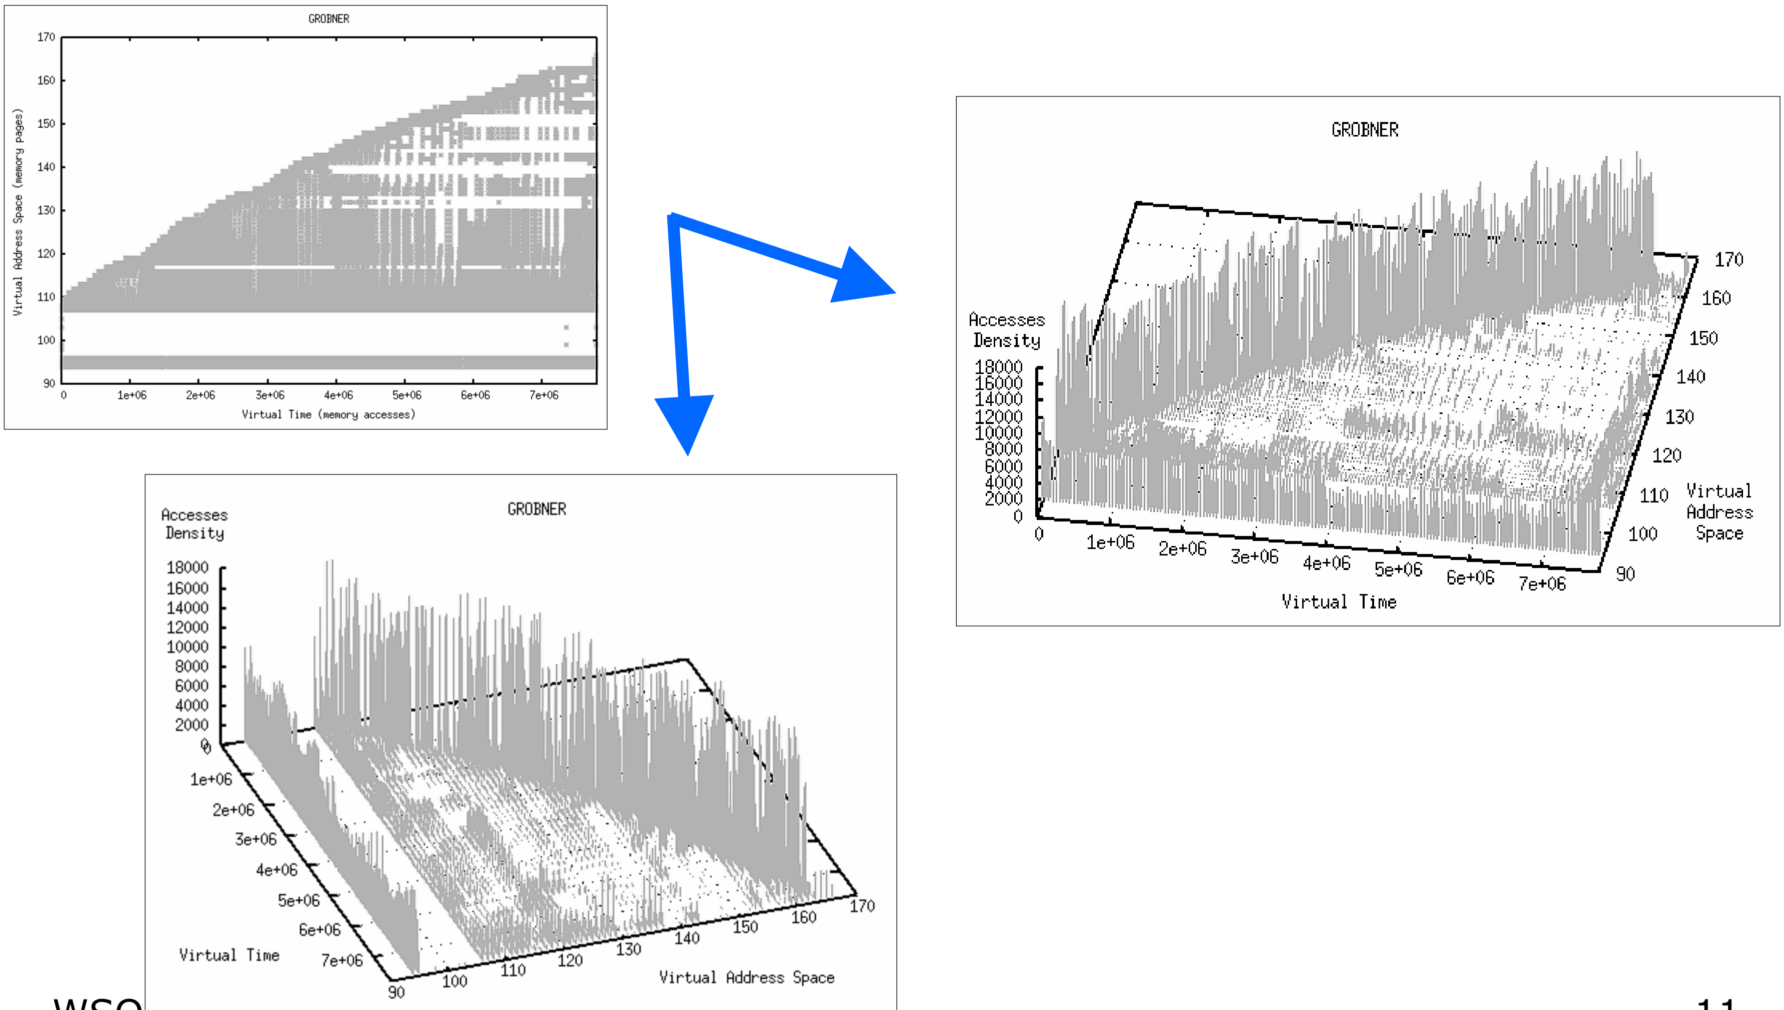

Case Study

- Grobner
  - It uses almost all pages in the virtual address space.
  - High access density:
    - Sequential access pattern.
    - Temporal reuse of low address pages.
  - Good results for page replacement algorithms that detect sequential patterns.

Definitions

- IRG (inter-reference gap)
  - Time distance between successive references of a memory page.
- IRR (inter-reference recency)
  - reuse distance or recency
  - number of other unique pages accessed between two consecutive references to the same memory page.

IRG Graph

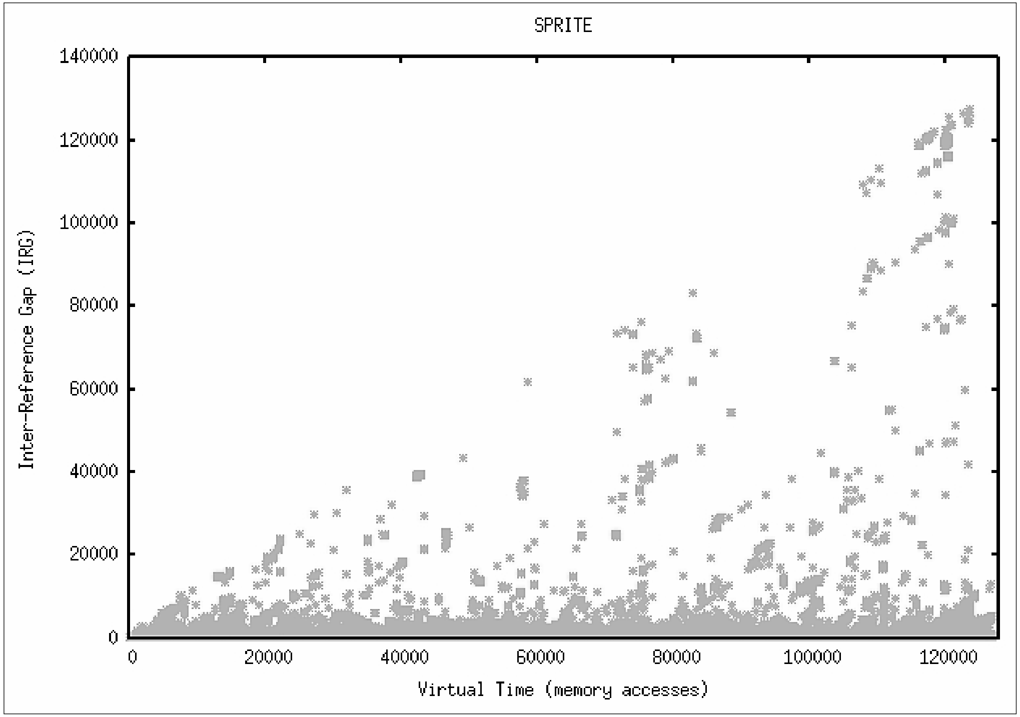
SPRITE
Inter-Reference Gap (IRG)
Virtual Time (memory accesses)
0
20000
40000
60000
80000
100000
120000
140000

IRR Graph

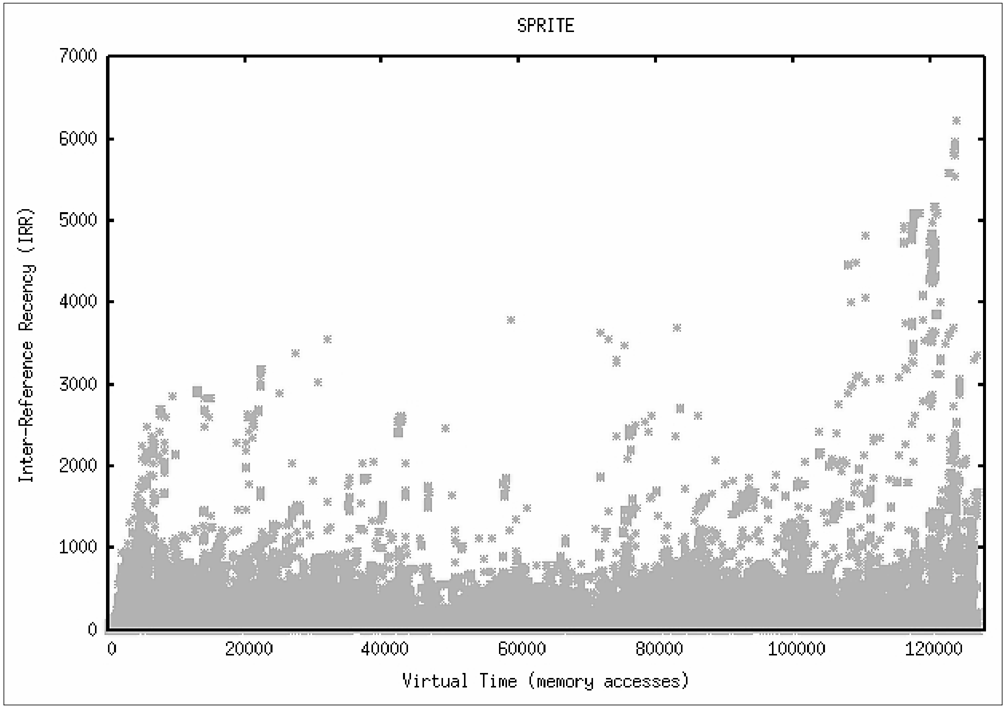
SPRITE
Inter-Reference Recency (IRR)
0
1000
2000
3000
4000
5000
6000
7000
0
20000
40000
60000
80000
100000
120000
Virtual Time (memory accesses)

IRR Surface

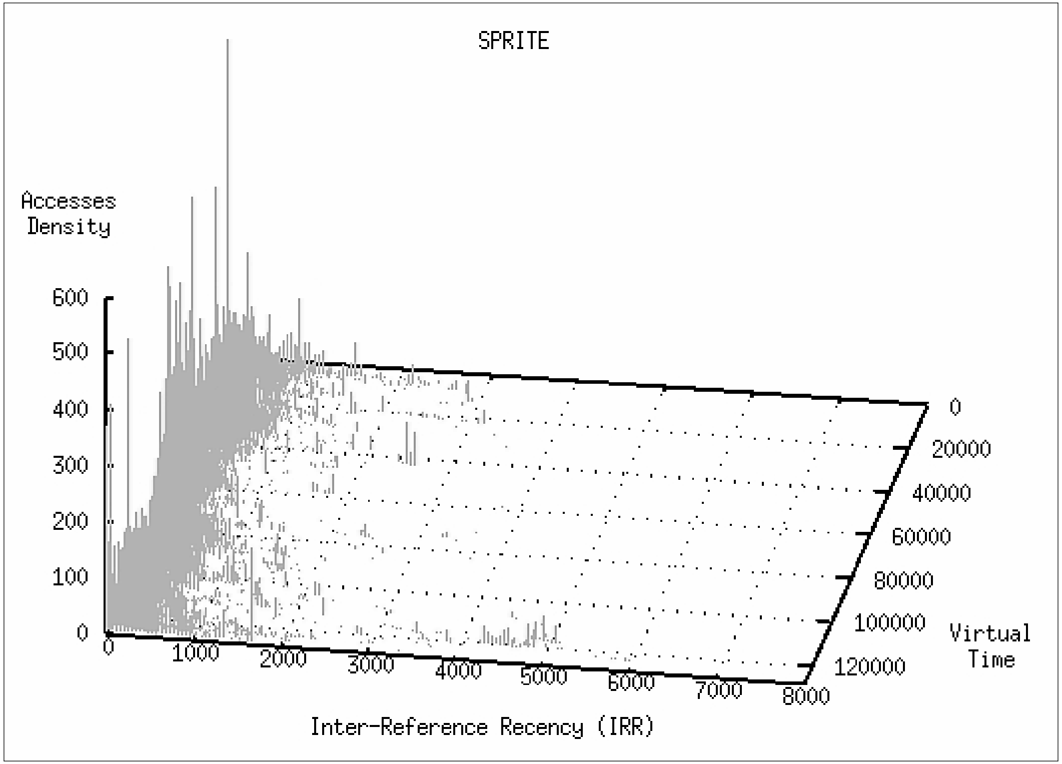
SPRITE
Accesses Density
600
500
400
300
200
100
0
0
1000
2000
3000
4000
5000
6000
7000
8000
Inter-Reference Recency (IRR)
0
20000
40000
60000
80000
100000
120000
Virtual Time

Case Study

- Sprite
  - IRG graph: most of accesses occur in virtual time intervals between 1 and 5,000 references.
  - IRR graph + IRR surface: most accesses occur in the first 1,000 positions in the LRU stack (15% of the footprint).
  - It presents pages with strong temporal locality being alternating with less accessed pages:
    - Not favorable to adaptive algorithms.

IRR Histograms

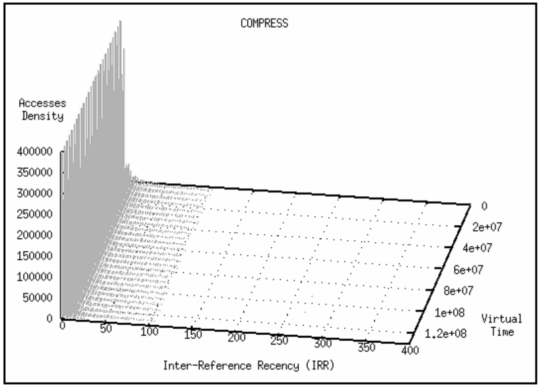

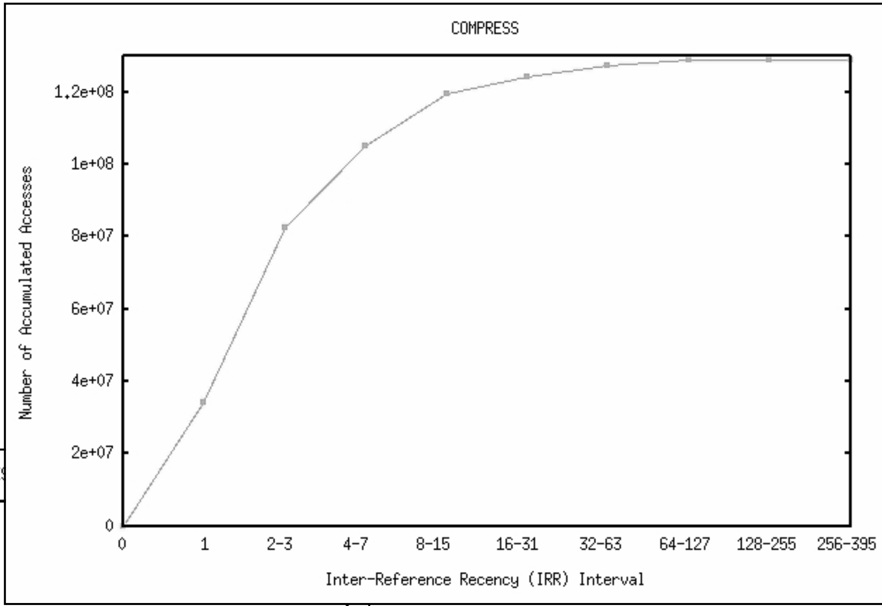

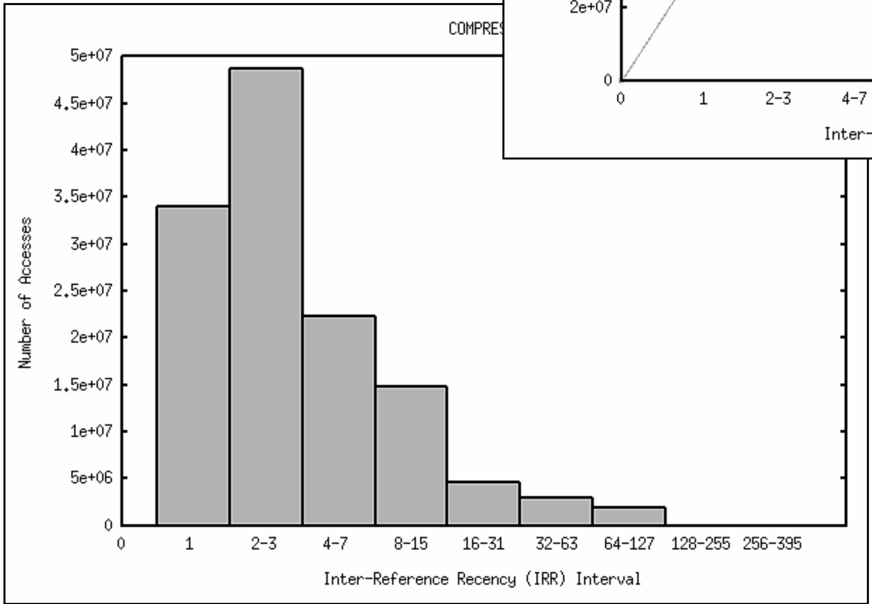

IRR Histogram

IRR Cumulative Histogram

Case Study

- Compress
  - A memory size equal to 8 pages is enough to maintain the page fault rate as low as 18,7%.
  - With 64 pages: 1.45%.

  - It presents high temporal locality;
  - It provides more accurate distinction among different processing phases by adaptive algorithms:
    - Good performance of adaptive algorithms.

Conclusion

- In this paper we introduced six new graphs for studying locality of references.
- Visual aspects:
  - Memory access patterns;
  - Temporal and spatial localities;
  - Real distribution of memory accesses;
  - Reuse frequency;
  - LRU stack position.
- The performance of page replacement algorithms was analyzed using these new graphs.

Future Work

- More case studies
  - Parallel applications

- New metric: **IRR-n** (number of distinct pages referenced among n+1 consecutive accesses to the same page).

- Integration and enhancement of the tools available in *Elephantools*.

Thank you.

Questions?

ESCOLA POLITÉCNICA DA UNIVERSIDADE DE SÃO PAULO

Extra slides

Locality Surface

- Technique to quantify temporal and spatial locality in programs
- Introduced to cache memory studies.
- Characterize the memory accesses taking into account the complete program:
  - Z-axis: number of occurrences of specific delay (stack distance) and stride (difference between memory addresses) values obtained among inter-referenced pages
  - X-axis:
  - Y-axis:

Case Study

- Simulation experiments

Trace	Number of Distinct Pages	Memory Sizes Considered in Simulations	Number of Simulations
Compress	396	10, 15, 20, ..., 385, 390, 395	78
Espresso	77	10, 11, 12, ..., 73, 74, 75	66
Grobner	67	10, 11, 12, ..., 63, 64, 65	56
Sprite	7075	100, 200, 300, ... 6800, 6900, 7000	70