

On Computational Models for Flash Memory Devices

Motivation


Flash memories

Between RAM memories and hard disks:

- are fast becoming the dominant storage on mobile computing
- have already replaced traditional hard disks on some devices

| Characteristic | RAM | Flash | Hard disk |
|--------------------|-----------|--------|------------|
| Volatile | Yes | No | No |
| Shock resistant | Yes | Yes | No |
| Physical size | Small | Small | Large |
| Storage capacity | Small | Large | Very large |
| Energy consumption | - | Medium | High |
| Price | Very high | Medium | Very cheap |


Hard Disks (HDD) vs Flash Disks (SSD)



Read/write blocks of data of size B randomly and sequentially

- HDDs, throughput provided by reads and writes is the same
- SSDs, results show that reads and writes are also done in blocks, **but** block size for reading is smaller than the block size for writing

Unit-Cost Model



Memory of size M , infinite flash disk, block transfers of consecutive data

Reads in blocks of size B_r , writes in blocks of size B_w


Cost of an algorithm: #items transferred. For r reads and w writes,

$$\text{cost} = r \cdot B_r + w \cdot B_w$$

Many existent algorithms in the I/O-model are easy to adapt in this model

Translation Layer (TL)

The algorithm uses the same block size (B_r) for reads and writes, similar to the I/O-model




The translation layer:

- groups B_r -sized blocks and writes B_w -sized blocks oblivious to algorithm
- accommodates the model (read block size B_r , write block size B_w)

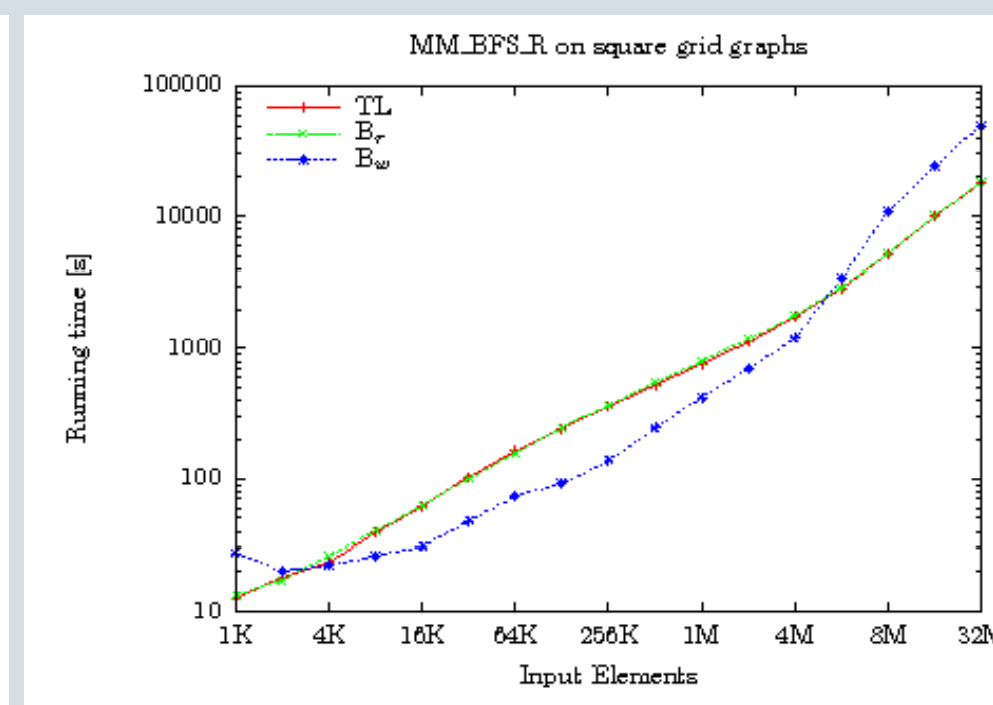
Experimental Results on Solid-State Disks

External Memory Sorting



Only sequential I/Os

External memory BFS



Mixed sequential and random I/Os

Naïve DFS (textbook algorithm)

| #vertices | TL | B_r | B_w |
|------------|----------|----------|---------|
| 262,144 | 0.218039 | 0.219124 | 1.10229 |
| 1,048,576 | 0.925051 | 0.958422 | 1.73248 |
| 4,194,304 | 3879.14 | 4641.21 | >20000 |
| 16,777,216 | 54465.7 | >1 day | >1 day |

Mostly random I/Os

Discussion

- Smallest running times are achieved in all cases when using the translation layer!
- As predicted by the unit-cost model, buffering B_r -sized read blocks into B_w -sized write blocks improves the performance

Conclusion

- The unit-cost model is validated by experimental results

References

Deepak Ajwani, Andreas Beckmann, Riko Jacob, Ulrich Meyer and Gabriel Moruz. *On Computation Models for Flash Memory Devices*. SEA 2009.