



Prof. Stratos Idreos  
CS165: Data systems. Midterm1: 10/11/2017  
**Duration 2 and 1/2 hours. Return this sheet when done.**  
**Name and HUD:**

**Guidelines:** You may use any printed notes, papers or textbooks you want. No electronic devices are allowed. You need to score 100% to get full points. There are also 25% bonus points. For each part of the midterm give a step by step description and explanation of how you came up with your solution and how the solution works. In addition, explain in detail why you think it is the best choice, what other solutions exist and why you did not opt for them. Give as much information as possible so we can see your thought process. It is OK if your answer is not all the way complete and perfect as long as you clearly articulate the main tradeoffs and design points. Even in cases where you cannot reach an exact solution but you do have a feeling of the kinds of tradeoffs that exist, you should describe those as clearly as possible as this is a big part of the design process and will give you a lot of points (50%). In addition, if after you finish describing a solution you realize that this was actually not the optimal solution, please describe as much as possible why this is the case and what your new intuition says. Such answers will get the full credit if the design considerations that are described in the end are correct and complete. If you make assumptions for issues not covered by the set-up we provide, then you should be clear about what those assumptions are. There are purposely small “gaps” in the design below in some of the questions: you are expected to fill them in and describe the overall design and tradeoffs. For example: “Question X asks for...however, we also need to know the design for Z to solve this one because... I will assume that Z is done as follows... which is a good solution when ....because... And will assume Z is done as follows in all remaining cases because....This leaves me with the following tradeoffs...”

**Setting:** Assume a database with a single table R. The table has K attributes  $\{A_1, A_2, \dots, A_k\}$  and N tuples. Base data is stored one column-at-a-time as fixed-width and dense array. Each array  $A_i$  has a fixed width of  $A_i.width$  bytes. No index exists by default. Your hardware has a single CPU with 4 cores and a 4 layer memory hierarchy with the last level being persistent memory. Each memory level  $M_i$  has size  $M_i.size$  bytes and block size  $M_i.block$  bytes (the minimum transfer granularity). The following are true:

- $M_4.size > M_3.size > M_2.size > M_1.size$ ,  $M_4.block = 4 \times M_3.block$ ,  $M_3.block = 2 \times M_2.block$ ,  $M_2.block = M_1.block$
- Total data size  $< M_4.size$ , Total data size  $> M_3.size$ , The size of each attribute column  $> M_2.size$   
The CPU can read from any memory level directly. When data is not found in  $M_1$ , it looks in  $M_2$ , then in  $M_3$  and finally in  $M_4$ . When a block is not found in memory  $M_i$  we have one  $M_i$  miss. When data is not found in  $M_i$ , it is subsequently copied there. To make room for a new block in  $M_i$ , an existing block is evicted from  $M_i$  (with LRU) and copied in  $M_{i+1}$  (unless  $M_i=M_3$  in which case data exists already in  $M_{i+1}=M_4$  so there is no need to copy it).
- The CPU can write directly only to  $M_4$ .
- The TLB can hold up to 256  $M_4$  memory locations.

**Part 1: 65%** -> You are expected to use only one CPU core for Part 1.

- Create a query plan for the following query: select  $A_3+A_4$  from R where  $A_1 < 10$  or  $A_2 > 20$ . Your plan should assume that no indexes are available, use late tuple reconstruction, column-at-a-time processing and bit vectors to hold intermediate results. Define the set of DSL operators you will need (their input/output) and write the plan. (5%)
- Give the pseudocode for each DSL operator in the plan. (10%)
- Redo the above two steps using an index on column  $A_1$ . The index in this case is a fully sorted copy of column  $A_1$  (stored as a fixed-width, dense array). If you reuse any operators just say so; do not redefine them. (10%)
- Construct the costs of a) each operator of the two plans and b) the complete plans. As a metric use  $M_2$  misses. Describe these costs both using a textual description (2-3 sentences max per operator) and coming up with an equation in big O notation. (15%)
- Which of the two plans above is best, when and why? (10%)
- Using Parts 4 and 5 devise a solution that automatically picks the best plan in any scenario. (15%)

**Part 2: 60%** -> You are expected to use all 4 CPU cores for Part 2.

Assume we have Q queries arriving at the same time. Queries are of the same form as in Part 1. Your goal is to maximize throughput, i.e., the number of queries we can process per second.

- Using column-at-a-time processing as in Part 1, describe in pseudocode a shared scan operator, shared fetch operator, and shared addition operator that minimizes  $M_2$  misses without thrashing TLB. (15%)
- Do the same for the index in the form of a sorted column from Part 1.3). (10%)
- Construct again the total costs as in Part 1.4). (10%)
- Given your shared processing, would you change your answer in Part 1.5? (5%)
- Assume you can do any kind of design and you also need to respect specific latency requirements for each query as batches of queries arrive with an arbitrary arrival rate: Is your design above still the best one? If yes, explain why. If not, provide a new one (design, plan, costs, pseudocode). (20%)