

**University of Pennsylvania**  
**Department of Electrical and System Engineering**  
**System-on-a-Chip Architecture**

ESE532, Fall 2019

Analysis Milestone

Wednesday, October 23

---

**Due:** Friday, Nov. 1, 5:00PM

**Group:** Forming your team, assessing the requirements and parallelism, and developing an initial functional implementation of the computation are group tasks for this milestone. You should discuss and refine your understanding of the task as a project team. You may share pseudocode and diagrams and collaborate on performance models within your team. You may not share or collaborate with anyone outside of your team.

For this milestone writeup is a group task.

1. Report the partners you have agreed to work with for the project duration.
2. Specifically considering processing the input stream at 1 Gb/s:
  - (a) Assuming hardware gathers up 64b words from the input stream, how many 1.2 GHz cycles do you have to process each 64b of data and maintain full throughput?
  - (b) Assuming hardware gathers up 64b words from the input stream, how many 200 MHz cycles do you have to process each 64b of data and maintain full throughput?
3. Working from the high-level description of the computations involved, assess the computational and memory requirements for each of the coarse-grained operations (Content-Defined Chunking, SHA-256, chunk matching for deduplication, LZW encoding).
  - (a) Summarize the computation of each coarse-grained operation in pseudocode. (write in your own words; credit sources)
  - (b) What memory is needed to support each task? (Identify what functions the memory serves and the size for each of the memory components.)
  - (c) What computational work is required per byte of input (or per chunk, where appropriate)? (How many adds, compares, multiplies, etc.)
  - (d) What memory operations are required per byte of input (or per chunk, where appropriate)?
  - (e) Based on this analysis and a simple model of processor execution (define as needed), what throughput (in Gb/s) can a single ARM processor achieve on this task? [This is a resource bound question.]

4. Identify and characterize parallelism available.
  - (a) What parallelism exists among the operations in the coarse-grained task flow?
  - (b) Within each operation, characterize opportunities for or inhibitions to thread-level data parallelism (i.e., what can be processed independently and what must be processed in a sequence).
  - (c) Within an operation thread, what opportunity is there for data-level parallelism?
  - (d) Within an operation thread, what opportunities exist for pipelined computations? What II is achievable? What dependencies determine the II? What is the depth of the pipelines?
  - (e) What is the latency bound for the entire deduplication/compression flow for a single chunk working with a maximum chunk size of 8KB.

Questions 3 and 4 should give you enough information to begin reasoning about how you can achieve the throughput goals (1Gb/s, 10Gb/s). We are deliberately asking you to consider these computations from the high-level description rather than a particular piece of sequential code so that you can reason about the fundamental requirements and opportunities apart from the specific artifacts of a particular, sequential implementation. As you begin to develop and benchmark your solutions, you should refer back to your analysis here. Is your solution performing worse than your analysis would predict? Is this because your solution is inefficient or encounters some unanticipated bottleneck? or is this because your original analysis missed something?

5. Develop placeholder encoder solution in C.
  - (a) Run on single ARM Cortex A53 processor.
  - (b) Work with sample input files we've provided and the provided Decoder. (see project handout.)
  - (c) Focus on decomposing into components and how those components communicate with each other.
    - What data structures, data format, and data sequence do you send between components?
  - (d) To scope this down, do **not** focus on full functionality of the components. Instead look for minimal, placeholder functionality. These should be simple—as simple as possible to handle the data. As a result, they will likely not be particularly good (e.g. poor compression) or fast (e.g. slow algorithms).
    - For example, you might use addition modulo  $2^{64}$  as a placeholder hash for SHA for initial integration as a placeholder to be replaced with the proper SHA-256 later. (using a simplistic hash like this will require that you check for false matches)
    - Hint: how can you minimize the initial functionality of each of the other components
    - The idea is that we will refine these with proper functionality for the next milestone. This helps enable you to always have something that works and allows you to work on components independently.
  - (e) Include a description in your report of:
    - your interfaces between components
    - your placeholder functions (what they do, how they are simplified)
    - status of the code – what overall functionality does it provide
  - (f) Turn in a tar file with your minimal functionality code to the designated assignment component in canvas.