

2. User Request (Direct-Mapped Cache Description)

Please read the attached file describing a fully associative CPU cache. Now, I want a reasonably short description of a direct-mapped cache, including the concepts of a CPU address being decoded into a tag, index, and offset. In addition, please use the concept of a comparator circuit that compares a tag from the CPU address with a tag from a cache line. Please note that a fully associative cache requires one comparator for each cache line, but a direct-mapped cache requires only a single comparator. (To Google Gemini, Oct. 20, 2025)

A **direct-mapped CPU cache** is the simplest cache organization where a **memory block can be placed in only one specific cache line** (or “slot”), determined by a portion of the memory address. This mapping is not flexible, meaning a unique **index** (derived from the address) points directly to the single cache line where the block must reside.

Address Decoding and Mapping

When the CPU requests an address, the address is logically divided into three parts:

1. **Block Offset:** Identifies the specific data unit (e.g., a byte) within a data **block** (or “cache line”).
2. **Index (or Set Index):** Used to select the specific cache line (or set) where the memory block might be stored. For a direct-mapped cache, the index directly selects the *one and only one* line a block can occupy.
3. **Tag:** The remaining, most significant bits of the address. This uniquely identifies the memory block *within* the main memory, distinguishing it from other blocks that map to the same **Index**.

Hit Check and Comparators

When the CPU requests data, a **cache hit** occurs only if the following conditions are met for the single cache line pointed to by the **Index**:

1. The **Valid Bit** (V) of the cache line is set to 1 ($V = 1$).
2. The **Tag** extracted from the CPU’s requested address **matches** the Tag stored in that specific cache line.

The check for the Tag match is performed by a **comparator circuit**.

- **Direct-Mapped Cache:** Because each memory block can map to only **one** specific cache line, a direct-mapped cache requires only a **single comparator** circuit. This single comparator compares the Tag from the CPU address with the Tag read from the one indexed cache line.

- **Fully Associative Cache Contrast:** In contrast, a fully associative cache must search *every* cache line simultaneously for a match, requiring **one comparator for every single cache line** in the entire cache. This is why a direct-mapped design is simpler and faster in terms of lookup logic, though it is more prone to conflict misses.