Apparatus Architecture

Apparatus is a complex tool. It's more of a framework with its own ecosystem then some simple plugin. In order to use it effectively and consciously you have to understand how it actually works. We're not talking about the very specifics of the implementation however but the main top-level architectural concepts. We will begin our acquaintance with a top level "singletonian" entity called *Machine*.

Machine

Machine is the main system of Apparatus. It's a manager for the things global and thereby a global singleton itself. It's actually a UObject but its lifespan is defined by its internal state, not the standard garbage collecting procedures. If Machine has some Mechanicals defined within it, or there are still some Subjects in the "scene" it will be retained and remain available. Only when it's no longer needed and actually empty (meaningless) it will be available for a disposal.

Within machine in particular and Apparatus as a whole, two "worlds" exist. Two levels of ECS data with their own unique features and optimizations.

The API documentation page for the UMachine class is of course available for your reference.

Low-Level Traits

Let's start with the lower layer first. The *Traits* subsystem was actually developed later, but it's now at the core of the functionality and provides the needed base functionality for the upper layer to work properly.

ECS was once developed in performance in mind. Packing and storing the data linearly in memory. What could be simpler? While it's actually not that easy to implement this for dynamically structured entities and requires some sophisticated bookkeeping, the whole notion is correct. The whole hardware layer of CPUs and RAM is really tailored towards this memory organization. Modern day CPUs have some large cache capacities which are utilized more efficiently when used with data pieces stored next to each other.

Unreal Engine's own memory model doesn't guarantee this level of linearity and using custom allocators is rather quirky or not viable at all. That's why we created the Traits subsystem.

Traits are primarily based on Structs. Those are managed exclusively by Apparatus and are stored in special buffers called Chunks. Right how it's supposed to be - one by one, sequentially, no gaps.

Traits are in turn assembled into collections (or it wouldn't be ECS after all). Those collections are called *Subjects* and are referenced through special *Handles*, not pointers. They are absolutely GC-independent and are disposed explicitly.

The whole design maximizes the performance of the Mechanics running on the Subjects, but it actually has some limitations comparing to the higher-level *Details*.

High-Level Details

Unlike Traits, Details are not Structs. They are "Unrealean" types of things - Objects (or UObjects to be more specific). This makes them really versatile in terms of utilizing existing Unreal Engine's functionality. Not only that but they also support hierarchical filtering and even multi-detail iterating.

Details are always stored in their respective *Subjectives* - this is a special type of container that is directly associated with an Unreal Actor or User Widget. Subjectives are not iterated directly however but through a special caching storage called *Belt*. It's a sparse type of storage and is used as an optimization mainly, containing only references to original details.

Please note, that all of the Subjectives are actually Subjects internally. They all have a Subject so Handle in them. This means you can add traits to them. You can interchangeably utilize the both worlds together when/if needed. It's all up to you.

Enchaining

One of the main technical goals of Apparatus is to effectively process some very large amounts of subjects and subjectives (our own term for ECS' entities) under a specific filter. With that in mind a special concept of *enchaining* was developed.

Enchaining is a process of collecting all of the currently available belts and chunks matching a certain filter and embodying them in a special type of array called *chain*. Chains are managed by the machine singleton and you won't be creating them manually, even when using a C++ workflow.

Instead you have to use a global (static) **Senchain methods**, passing them a filter of the desired selection. You can either enchain belts or chunks together. While they are being enchained and the chain actually exists the respective belts or chunk are also *locked*.

Locking

During iterating the chains and their respective chunks and belts we have to guarantee a certain level of immutability for them. We don't want to process the same subject twice, for example, as it may be tossed around chunks while being structurally modified inside the currently ongoing iterating. The belt/chunk locking functionality is designed specifically for that purpose.

When the chunk or belt becomes enchained, its internal locks counter is incremented essentially making it somewhat frozen to the evaluating mechanic that is going be working on that chain. You can worry-free use the chain cursor and let Apparatus handle all the immutability issues for you.

Filtering

Filtering is an essential part of the ECS paradigm implementation. It lets you select specific subjects and subjectives to work upon. Using the word *select* in this context is not by chance as the term could be very familiar to a database programmer, for example. Technically it's quite the same. You

define a WHERE clause with a set of conditions to meet. These can be both inclusive (positive) and exclusive (negative).

Apparatus uses all sorts of different optimization schemes and caches to make the filtering process as fast as possible. You shouldn't worry too much about that.

API documentation for filters.

Iterating

Once you have your Subjects spawned and set. Belts or Chunks enchained you're ready to iterate on them to deliver the necessary logic of the game or application. This is done through a very common concept of *Iterators* and *Cursors*.

Both Belt and Chunk have their own types of Iterators, but you would rarely use them directly. Instead you'll almost always use Chain Cursors. They are essentially Iterators with a naming chosen to eliminate some possible ambiguity. For now you should only use the default (implicit) Cursor as the threading is still a planned feature and you would rarely need to iterate a Belt (or a Chunk) with multiple different Cursors.

The API documentation for Begin and Advance methods is provided accordingly.

From:

http://turbanov.ru/wiki/ - Turbopedia

Permanent link:

http://turbanov.ru/wiki/en/toolworks/docs/apparatus/architecture?rev=1623075154

Last update: **2021/06/07 14:12**

