Subject: Re: Map implementation
Posted by Novo on Wed, 10 Apr 2019 15:09:47 GMT

cbpporter wrote on Wed, 10 April 2019 05:39I also talked with a colleague and his hashmap version is supposedly 3x+ faster than stl, so it looks like it is very possible to greatly outperform it.

IMHO, it is impossible to create one ideal hash table which will greatly outperform STL in all possible scenarios.
Let's take a look at two scenarios.

1. One million hash tables containing one hundred records.
2. One hash table containing one hundred million records.

You will need two completely different implementations in these cases.

Two more scenarios.

1. Add data once and search for data most of the time.
2. Add/remove data most of the time and search for it occasionally.

Again, you will need two completely different implementations.