

---

Subject: Re: Map implementation

Posted by [mirek](#) on Sun, 14 Apr 2019 06:03:45 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Novo wrote on Sat, 13 April 2019 21:07mirek wrote on Wed, 10 April 2019 11:37

Quote:

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.

That might be true, however I do not see a way how to improve Index for either (I see some accumulated knowledge how to improve it for both, but that is another story).

IMHO, this is impossible because in the first case your main concern is the memory usage. Tiny overhead multiplied by million is a huge problem. And in the second case performance is the main issue.

If the performance is the issue, then the memory is the issue too. The game starts at L1 cache size, which can correspond to hundreds of records.

---