Subject: Re: New containers - naming
Posted by mirek on Sun, 03 Feb 2013 20:17:58 GMT
View Forum Message <> Reply to Message

dolik.rce wrote on Sun, 03 February 2013 14:39Hi Mirek,

That great, now Il be even more addicted to NTL...

Just to make sure, InVector stands for "insert vector"?

Well, either that or that it supports "in" positions. Of course, better ideas are welcome

Quote:
IMHO it would be great if the name of the new containers contained the information how it differs from the regular ones.
I think the biggest advantage of this will be the range searches, so what about something like RangeIndex, RangeVector and RangeArray?

Well, my original line of thinking was that the main difference is that keys are ordered... (also note std::unordered_map etc...). But OrderedIndex is perhaps too long, so "Order".

Quote:
BTW: Will you tell us how it works, or is it left as an exercise for the reader? I suspect you told me once about this idea in past over a beer, but I want spoil it just yet for other curious programmers here

The basic storage is Vector<Vector<T>>, the size of inner vectors is kept between some thresholds, otherwise they are split / merged. For such small amount of data, Vector was always faster at inserting/removing than any node based structures.

The key to provide relatively fast index access is sort of numeric binary tree, which is easy to maintain on inserts/removes (removes are not implemented yet) and still quite fast for operator[]. Plus, there is a per thread cache to speed up simple scans. And a good thing is that Find[Upper/Lower]Bound methods can compute the index of an element very cheaply.

Mirek