I have made some substantial changes to Core memory allocator and index, improving performance of some synthetic benchmarks.

Allocator now much better handles big blocks, which improves e.g. performance of adding ~20000 elements to Vector<int> 3 times. Also, memory pages of most categories can be now reused in another category. We have now 3 categories of blocks <1KB, <64KB and <32MB/220MB (32 bit cpu/64 bit cpu). MemoryTryRealloc is now properly implemented and used in library. mingw performance is improved with TLS workaround.

sizeof(Index) is now 40 (was ~90). Adding elements to Index is now faster.

Frankly, in retrospective it was all mostly a lot of work for really small gains as all low-hanging fruits were already picked years ago. But large blocks handling in allocator is quite nice improvement...