

## Ruby master - Misc #15402

### Shrinking excess retained memory of container types on promotion to uncollectible

12/11/2018 08:43 PM - methodmissing (Lourens Naudé)

<b>Status:</b>	Open
<b>Priority:</b>	Normal
<b>Assignee:</b>	
<b>Description</b>	
<p>I've been toying with the idea of the viability of attempting to reclaim over provisioned memory from buffer capacity of container objects like Array and String, effectively reducing the footprint of retained memory of such objects.</p> <p>GC at the moment covers these dominant paths:</p> <ul style="list-style-type: none"><li>• Collection of shallow memory: unreferenced object slot with values encoded on the object</li><li>• Collection of retained memory: unreferenced object slot with off ruby object heap pointer to String buffer, Array buffer etc. (heap.aux)</li><li>• Finalization hooks like reclaiming resources for Tempfile for example</li></ul> <p>I explored in <a href="https://github.com/ruby/ruby/pull/2037">https://github.com/ruby/ruby/pull/2037</a> (more details and data points on the PR) a forth one:</p> <ul style="list-style-type: none"><li>• Shrinking over provisioned buffer capacity of Array (also applies to String and likely others) on promotion to uncollectible (Also garbage collect excess retained space from types with a first class capacity and buffer on promotion to uncollectible)</li></ul> <p>Sharing here for feedback in case anyone has ideas for a more appropriate hook, or additional precondition for such a hook. Or if excess buffer capacity can even be considered first class garbage in a GC context.</p> <p>I chose Array as a proof of concept because the type already have this optimization through <code>ary_shrink_capa</code> through <code>ary_make_shared</code> and the threshold for not encoding members on the object is quite low at 3 elements. Plausible that many framework / boot specific long lived arrays are larger than that and because the growth factor on expansion is 2x, also likely a fair amount of over provisioned capacity.</p> <p>Results of the changeset:</p> <ul style="list-style-type: none"><li>• Benchmark <code>so_binary_trees</code> - 26% reduction in total memory usage</li><li>• Also a very noticeable 24% difference with <code>app_lc_fizzbuzz</code></li><li>• General few bytes reduction for almost all core benchmarks</li><li>• Mainline redmine after boot - 1.5% or 45kb reduction in Array retained memory size</li></ul> <p>Implementation caveats:</p> <ul style="list-style-type: none"><li>• Promotion to uncollectible may be a bad heuristic for shrinking buffer capacity.</li><li>• Needed to create a new <code>rb_ary_shrink_capa</code> function as <code>ary_shrink_capa</code> is private API and has several assertions (frozen and shared check) that hard fails during GC. That way shrinking responsibility and accounting remains the responsibility of <code>array.c</code> - the GC just calls it (same as with the <code>memsize</code> APIs)</li><li>• I tried running it during GC which worked fine for benchmarks like <code>so_binary_trees</code> but failed under GC stress and larger heaps because of <code>TRY_WITH_GC</code> via <code>objspace_xrealloc</code>, which can invoke GC</li><li>• A reasonable workaround for this was to use the postponed job API which is used by GC for object finalization, but that's one job per object space, not 1 per Array being shrunk, which may hit the 1000 item postponed job buffer for some heaps. It degrades gracefully though with fallback being the optimization simply not being applied to the excess objects in the set.</li><li>• Have no idea about the future of the postponed job API and if this is an appropriate use case</li><li>• <code>RVALUE_PAGE_OLD_UNCOLLECTIBLE_SET</code> only special cases Array at the moment - it's easy to support other types</li></ul> <p>Outliers to still evaluate:</p> <ul style="list-style-type: none"><li>• Fragmentation does not get significantly worse through reallocs for specific rare cases post GC (no data)</li><li>• The effect of <code>objspace_malloc_increase</code> called by <code>objspace_xrealloc</code> on GC frequency (I think not much given the small reduction on retained usage, but have no data to prove yet)</li><li>• How well the postponed job pattern scales to large heaps and how much of the job slots are consumed (no data)</li></ul> <p>Thoughts on exploring more types or is the pattern tainted / broken to begin with?</p>	