Well, also being absolutely not an expert of GC algorithms, I always think that to allocate a lot of memory just because there is enough of it it's not a great practice.
First thing, you must look to what does OS in respect of free memory.... If your OS tells you that you've got 1 GB free ram, even when it's already swapping out another GB on disk, what this behaviour does is slowing down your system.

IMHO a truly efficient GC should be hardware implemented, or at least have a strong hardware support. Doing it software you'll face everytimes with some sort of problem, as latency, memory inefficiency or both.
And it should also collect freed memory asap.
Cbpporter spoke about some sort of 'threshold' for collecting garbage.... but what should be such a threshold ? On oldest OS that was simple, no swapping mechanics, so you could decide, 30% of physical ram, ok. On modern OS, you can't. That depends on too many factors, on how many processes are running, and so on. You could implement a maybe efficient GC only at OS level, not at the application level, I guess.

No doubt that manual memory allocation is more efficient than GC, even if it can be a bot slower on the short time.
And a good framework can help to keep things simple.

Id rather extend C++ (or make some more modern language, without GC) to include some helpful features, than switch to less efficient languages.

Ciao

Max