

Some Performance Experiments for Simple Data Structures

*Francisco J. Ballesteros
Enrique Soriano
RoSAC-2011-1*

Laboratorio de Sistemas — Universidad Rey Juan Carlos
April 5 2011, Madrid, Spain.
<http://lsub.org>

ABSTRACT

Actual program performance is non-intuitive. Stroustrup, the author of C++, included measurements in a talk for ordered insertion in C++ vectors and lists. We measured the same issue using C in Plan 9, Mac OS X, and Linux, and C++ in Mac OS X. This document describes the results.

1. Which is faster, a list or a vector?

In a recent talk by Stroustrup in Madrid, it was pointed out that it is not intuitive if a list or a vector is a better data structure for ordered insertion for a given number of elements. In principle, according to bibliography, the list should win for container sizes greater than two (or a close number). However, a slide presented results from an experiment, which we reproduce in figure 1.

Perhaps surprisingly, in the experiment shown in the talk, the vector remains a better data structure for this purpose until 40.000 elements have been inserted (or a close number). This was for small element sizes.

However, things are even better (worse?). We reproduced the experiment in Plan 9 from Bell Labs (in C), Mac OS X (in C and C++) and Linux (in C). And we obtained some contradictory results from the resulting measures!

It seems that effects like hardware caches, operating system memory management, standard library implementations for the language used, etc. can in fact dominate what happens in the end to the performance of the program.

That is, in practice, results from complexity theory seems to be totally neglected. In our opinion, what happens is that software is so complex, and there are so many layers of software underneath the application code, that it is not even clear which data structures are better; at least for the simple case we describe here.




Figure 1 Ordered insertion times for C++ vector and list (Stroustrup, Madrid 2011 talk).

2. Reproducing the experiment

2.1. C program on Plan 9

We tried to reproduce the experiment, using Plan 9 from Bell Labs as the operating system running on a quadruple processor AMD64 Phenom 2.2GHz with 4GiB of memory installed. The machine has 64 bits per word, but the operating system and compilers installed keep it running in 32 bits (which is considered a machine word in what follows).

A C program was used to perform insertions on one of three different data structures. The program is reproduced in appendix A. The three data structures are: A regular array of elements (`Array`); A linked list of elements (`List`); and an array of pointers to elements (`Ptrs`). See figure 2.

Each element is represented as an integer plus some optional space, so that we could reproduce the experiment for different element sizes (all of them multiples of the machine word).

```

28  struct El
29  {
30      int n;          /* element value */
31      int dummy[];
32  };

```

The array grows as elements are added, growing for a customizable number of new elements each time.




Figure 2 Data structures used for the experiment: array (`Array`), list (`Node*`), and array of pointers (`Ptrs`). Filled boxes are elements.

```

34  struct Array
35  {
36      int nels;      /* number of elements used */
37      int naels;    /* number of elements allocated */
38      El *els;      /* array of elements */
39  };

```

Each linked list node is as expected:

```

41  struct Node
42  {
43      Node *next;   /* element in list */
44      int n;        /* element value */
45      int dummy[];
46  };

```

The array of pointers is similar to the array above, but refers to external elements instead of containing them:

```

48  struct Ptrs
49  {
50      int nels;      /* number of elements used */
51      int naels;    /* number of elements allocated */
52      El **els;     /* array of elements */
53  };

```

2.2. C++ program on Mac OS X

We used Mac OS X running on a 2.4GHz Core 2 Duo T7700 with 2GiB of memory installed. The C++ compiler was g++ version 4.2.1, and the libraries were libstdc++ 7.9.0 and libSystem 125.2.1.

The data structures used to compare lists and arrays are the STL implementations for `list<int>` and `vector<int>`. The source code for the program is included in appendix B.

2.3. C program on Mac OS X

We used the same Mac OS X machine than in the C++ set-up. The compiler was gcc version 4.2.1, and the standard library was libSystem 125.2.1. The C program is a port of the used in the Plan 9 experiment, using the same data structures.

2.4. C program on Linux

For the experiments on Linux, we used another machine, an Intel Pentium 4 CPU 2.40GHz with 512 MiB of RAM. The compiler was gcc version 4.4.1, and the standard C library was glibc 2.10.1. The C program is a port of the one used in the Plan 9 experiment, using the same data structures.

2.5. Experimental set up

Each experiment consisted on measuring the insertion of a given number of elements into one of the three data structures, with a fixed element size (and fixed increment size for C arrays). The elements inserted where integers (plus some optional space if required) taken in ascending order, in descending order, or in (pseudo-)randomized order. In the last case, the sequence of randomized integers was the same for all experiments, to make it fair.

Because these data structures are not isolated from the rest of the application when used in practice, 64 bytes of dynamic memory are allocated between each insertion in all the experiments. This memory is never released.

Measures of time are taken using *nsec(2)* in Plan 9 and *clock(3)* in C++ and C on Mac OS X and Linux. They include the insertion in the data structure and the allocation of memory for elements (allocation of elements in the case of the linked list and the array of pointers, and reallocation for the arrays). They do not include loops, extra allocations used by the program, and (pseudo-)random numbers generation.

3. Effect of increment in arrays

In the C implementation, the value for the increment in growing arrays may be important. This section tries to measure that effect. We inserted 10000 elements in randomized order into the array, for an element size of 4 bytes (1 integer in our experiment): once growing the array 1 element at a time, then growing it 16 elements at a time, and finally growing it 128 elements at a time. Figure 3 shows the time taken for the three experiments in nanoseconds, for Plan 9, Mac OS X, and Linux respectively. The relevant portion of code is as shown in this excerpt from appendix A:

```
94     if((a->naels%incr) == 0){
95         a->naels += incr;
96         a->els = realloc(a->els, a->naels*elsz);
97         if(a->els == nil)
98             return -1;
99     }
```

Figure 4 shows the results of the same experiment, using an element size of 64 bytes instead of 4 bytes, again, for Plan 9, Mac OS X, and Linux respectively.

For 4-byte elements, the graphs do not show significant difference regarding time (due to the scale). However, there can be seen important differences if the growing delta for the array is 128. For example, the Plan 9 program, for arrays up to 400 elements, is 83% faster with *incr=1* than with *incr=128* (mean of 25 independent executions incrementing the array size by 16 elements each time, from 16 to 400 elements).

For the same array sizes, the Mac OS X C program is 115% faster with *incr=1* than with *incr=128*. This is definitely non-intuitive! It seems that it is better to grow the array each time than it is to grow it from time to time. This is quite surprising, since `realloc` is called in each insertion when *incr=1*. Intuitively, one could expect to observe that the program performs better (at least, equally) with large increments. This is not the case.

Figure 4 shows the times for 64-byte elements. With *incr=128*, the Plan 9 program runs quickly out of memory (flat-dotted line in the graph above of figure 4). For increments of 1 and 16, the program can run for a longer number of elements. In Plan 9, *incr=1* performs worse than with *incr=16*, but it is still reasonable for 64-byte elements. In Mac OS X, *incr=1* results better, but comparable. In Linux (running on an older machine), results are comparable (but again, larger values of *incr* do not lead to better results).

For the next experiments, we use an increment of 1, growing the array each time.

3.1. Memory usage

Although figures are not shown here, the amount of memory used in these experiments over Plan 9 is quite different depending on the increment for array growth. In particular, with *incr=1* (growing the array one element at a time) the program consumes a lot less memory in the resulting process image (at the end of the experiment) than it consumes growing the array 128 elements at a time. An increment of 16 causes 14 times more memory to be consumed with respect to the increment of 1. An increment of 128 causes 141 more memory consumption. Also, using 64-byte elements and an increment of 128 makes the program run out of (virtual) memory in our Plan 9 system. Thus, the effect in memory footprint is not to be underestimated.

In what follows we consider only execution time, and not memory consumption.

4. Forward insertion experiment

Figure 5 shows the effect of forward insertion of 4-byte elements (1, 2, 3, etc.) in the data structures, using C in Plan 9. Inserting 4-byte elements in `Array` takes much less time than inserting on the other two data structures in the long run. For few elements (see the bottom graph) using `Ptrs` is worse than using `List`. However, for a number of elements between 1000 and 2000 elements `Ptrs` becomes better than `List`.

Figure 6 shows the times for 64-byte elements using the Plan 9 C program. For 64-byte elements things change. Instead of being faster, `Array` becomes slower, and `Ptrs` is not affected as much as the other two data structures. Also, there is a huge jump in execution time after inserting in the array about 3000 elements, which did not happen with 4-byte elements (probably would happen with a higher number of elements, not measured). Also, only for 64-byte elements, `List` is better than `Array` for less than 700 elements in the data structure (approx). No crossing point has been found for 4-byte elements: one is either better or worse than the other.

Inserting 4-byte elements using C++ in Mac OS X leads to the results shown in figure 7. Compare with figure 5 (the same experiment using C in Plan 9). Results are the opposite!

Times for inserting 4-byte and 64-byte elements using C in Mac OS X are shown in figures 8 and 9 respectively. Times for Linux are depicted in figures 10 and 11. For all these experiments, `Array` results better than `Ptrs` and `List`, in this order.

Results for C and C++ are the opposite. Moreover, results of 64-byte elements differ from the C program over Plan 9 and the C program over Linux and Mac OS X.

So, which data structure should we use?

5. Backward insertion

Figure 12 shows the effect of backward insertion of 4-byte elements (descending order) in the data structures using C in Plan 9. This time the list wins on the long run, as expected (in backward insertion, elements are always inserted in the head of the list). The same experiment, using 64-byte elements, leads to results shown in figure 13. Results are the equivalent, only that the vector gets worse due to the increase in element size.

Using C++ for 4-byte elements, we obtain the results shown in figure 14. Figures 15 and 16 show the results of inserting 4-byte and 64-byte elements using C in Mac OS X. Results of inserting 4-byte and 64-byte elements using C in Linux are depicted in figures 17 and 18 respectively.

For collections up to 400 4-byte elements, `Array` and `List` are comparable. For larger collections and larger elements, `List` wins, as expected.

6. Randomized insertion

We come to the experiment that motivated this work. This could be compared to the one shown by Stroustrup (but shouldn't).

We inserted 4-byte elements in randomized order into the data structures, using C on Plan 9. `Array` is better in the long run (but note the memory effects described above). For fewer elements (i.e., about 1500 or less) `List` becomes better. On the other hand, `Ptrs` seems to compete well with the other two ones. See figure 19.

In Plan 9, using 64-byte elements instead, the results are those shown in figure 20. Instead of being faster, `Array` becomes slower. The increase in element size makes the array take longer. For large collections, `Ptrs` is a good candidate in this case (better than the list in the long run).

Compare now figure 19 with results using C++, shown in figure 21. Surprisingly, our result is the opposite once more. Also, considering the number of elements, the result is also the opposite of the result shown by Stroustrup in his talk. Figure 26 shows our results for the scale used in the Stroustrup's graph.

Times for inserting 4-byte and 64-byte elements using C in Mac OS X results in the graphs depicted in figures 22 and 23. Like in Plan 9, for large collections, `Array` wins for 4-byte elements, and `Ptrs` wins for 64-byte elements.

The results of inserting a large number of 4-byte and 64-byte elements using C in Linux are depicted in figures 25 and 26 respectively. For huge collections, the array wins in this set up.

7. Summary

There is too much complexity. The cache hierarchies in the hardware, the operating system used, the C library and the standard library for the language used; all of them conspire to introduce effects that may even invert the results that you could expect. Clearly, in the end, the results obtained may be justified by different physical (that is, practical) effects and theory would be in accordance with the experiments if we consider such effects. However, it seems that we should use the simplest data structures that

simplify our programs, and do not pay attention to the data structures used before measuring our program in our particular compiler, system, and hardware platform.






Figure 3 Time (ns) for the C program inserting 4-byte elements into `Array` as a function of the number of elements for growing increments of 1, 16, and 128: Plan 9 (top), Mac OS X (middle), and Linux (bottom).



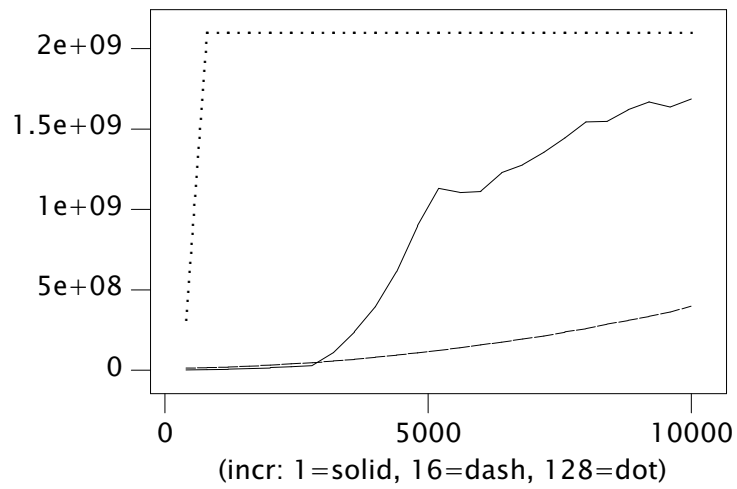


Figure 4 Time (ns) for the C program inserting 64-byte elements into `Array` as a function of the number of elements for growing increments of 1, 16, and 128: Plan 9 (top), Mac OS X (middle), and Linux (bottom).






Figure 5 Time (ns) for inserting 4-byte elements in ascending order as a function of the number of elements using C in Plan 9; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 6 Time (ns) for inserting 64-byte elements in ascending order as a function of the number of elements using C in Plan 9; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 7 Time (ns) for inserting 4-byte elements in ascending order as a function of the number of elements; for C++ STL vector (solid lines) and list (dashed lines), running on Mac OS X.






Figure 8 Time (ns) for inserting 4-byte elements in ascending order as a function of the number of elements using C in Mac OS X; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 9 Time (ns) for inserting 64-byte elements in ascending order as a function of the number of elements using C in Mac OS X; for Array (solid line), List (dashed line), and Ptrs (dotted line).





Figure 10 Time (ns) for inserting 4-byte elements in ascending order as a function of the number of elements using C in Linux; for Array (solid line), List (dashed line), and Ptrs (dotted line).





Figure 11 Time (ns) for inserting 64-byte elements in ascending order as a function of the number of elements using C in Linux; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 12 Time (ns) for inserting 4-byte elements in descending order as a function of the number of elements using C in Plan 9; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 13 Time (ns) for inserting 64-byte elements in descending order as a function of the number of elements using C in Plan 9; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 14 Time (ns) for inserting 4-byte elements in descending order as a function of the number of elements using C++ in Mac OS X; for C++ STL vector and list.






Figure 15 Time (ns) for inserting 4-byte elements in descending order as a function of the number of elements using C in Mac OS X; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 16 Time (ns) for inserting 64-byte elements in descending order as a function of the number of elements using C in Mac OS X; for Array (solid line), List (dashed line), and Ptrs (dotted line).





Figure 17 Time (ns) for inserting 4-byte elements in descending order as a function of the number of elements using C in Linux; for Array (solid line), List (dashed line), and Ptrs (dotted line).





Figure 18 Time (ns) for inserting 64-byte elements in descending order as a function of the number of elements using C in Linux; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 19 Time (ns) for inserting 4-byte elements in random order as a function of the number of elements using C in Plan 9; for `Array` (solid line), `List` (dashed line), and `Ptrs` (dotted line).






Figure 20 Time (ns) for inserting 64-byte elements in random order as a function of the number of elements using C in Plan 9; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 21 Time (ns) for inserting 4-byte elements in random order as a function of the number of elements using C++ in Mac OS X; for C++ STL vector and list.






Figure 22 Time (ns) for inserting 4-byte elements in random order as a function of the number of elements using C in Mac OS X; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 23 Time (ns) for inserting 64-byte elements in random order as a function of the number of elements using C in Mac OS X; for Array (solid line), List (dashed line), and Ptrs (dotted line).





Figure 24 Time (ns) for inserting 4-byte elements in random order as a function of the number of elements using C in Linux; for Array (solid line), List (dashed line), and Ptrs (dotted line).




Figure 25 Time (ns) for inserting 64-byte elements in random order as a function of the number of elements using C in Linux; for Array (solid line), List (dashed line), and Ptrs (dotted line).






Figure 26 Top: Stroustrup's results. Middle: Time (ns) for inserting 4-byte elements in random order as a function of the number of elements using C++ in Mac OS X; for C++ STL vector (solid) and list (dashed). Bottom: The same experiment performed in Linux (older machine).

Appendix A: Plan 9 C source code

listarry.c

```
1  #include <u.h>
2  #include <libc.h>

4  /*
5   *   Measure insertion into ordered sequences
6   */

8  enum
9  {
10     Incr = 16,
11     Num = 1000,
12     I2LN = 16,

14     Fwd = 0,
15     Bck = 1,
16     Rnd = -1,

18     Tarry = 0,
19     Tlist,
20     Tptrs,
21 };

23 typedef struct Array Array;
24 typedef struct Ptrs Ptrs;
25 typedef struct Node Node;
26 typedef struct El El;

28 struct El
29 {
30     int n;          /* element value */
31     int dummy[];
32 };

34 struct Array
35 {
36     int nels; /* number of elements used */
37     int naels; /* number of elements allocated */
38     El *els; /* array of elements */
39 };

41 struct Node
42 {
43     Node *next; /* element in list */
44     int n; /* element value */
45     int dummy[];
46 };

48 struct Ptrs
49 {
50     int nels; /* number of elements used */
51     int naels; /* number of elements allocated */
52     El **els; /* array of elements */
53 };
```



```

55 #pragma varargck type "A" Array*
56 #pragma varargck type "L" Node**
57 #pragma varargck type "P" Ptrs*

59 static int incr = Incr; /* in array realloc */
60 static int elsz = sizeof(int); /* number of bytes in element */
61 static int otherallocs; /* do mallocs to pollute space */

63 static void
64 usage(void)
65 {
66     fprintf(2, "usage: %s [-malpfbvrw] [-e nwords] [-i incr] [-n num]\n", argv0);
67     exits("usage");
68 }

70 static void*
71 anew(void)
72 {
73     return mallocz(sizeof(Array), 1);
74 }

76 static void*
77 lnew(void)
78 {
79     return mallocz(sizeof(Node**), 1);
80 }

82 static void*
83 pnew(void)
84 {
85     return mallocz(sizeof(Ptrs), 1);
86 }

88 static int
89 ains(void *x, int el)
90 {
91     int i;
92     Array *a = x;

94     if((a->naels%incr) == 0){
95         a->naels += incr;
96         a->els = realloc(a->els, a->naels*elsz);
97         if(a->els == nil)
98             return -1;
99     }
100    for(i = 0; i < a->nels && a->els[i].n < el; i++)
101        ;
102    if(i < a->nels)
103        memmove(&a->els[i+1], &a->els[i], elsz*(a->nels-i));
104    a->els[i].n = el;
105    a->nels++;
106    return 0;
107 }

109 static int
110 lins(void *x, int el)
111 {
112     Node **l, *n;

```

```

114     l = x;
115     for(; (n = *l) != nil && n->n < el; l = &n->next)
116         ;
117     n = malloc(sizeof(Node*+elsz);
118     if(n == nil)
119         return -1;
120     n->n = el;
121     n->next = *l;
122     *l = n;
123     return 0;
124 }

126 static int
127 pins(void *x, int el)
128 {
129     int i;
130     Ptrs *a = x;

132     if((a->naels%incr) == 0){
133         a->naels += incr;
134         a->els = realloc(a->els, a->naels*sizeof(El*));
135         if(a->els == nil)
136             return -1;
137     }
138     for(i = 0; i < a->nels && a->els[i]->n < el; i++)
139         ;
140     if(i < a->nels)
141         memmove(&a->els[i+1], &a->els[i], sizeof(El*)*(a->nels-i));
142     a->els[i] = malloc(elsz);
143     if(a->els[i] == nil)
144         return -1;
145     a->els[i]->n = el;
146     a->nels++;
147     return 0;
148 }

150 static int
151 Afmt(Fmt *f)
152 {
153     Array *a;
154     int i;

156     a = va_arg(f->args, Array*);
157     if(a == nil)
158         return fprintf(f, "<nilarray>");
159     fprintf(f, "[");
160     for(i = 0; i < a->nels; i++){
161         if(i > 0){
162             fprintf(f, ", ");
163             if(i < a->nels-1 && (i%I2LN) == 0)
164                 fprintf(f, "\n");
165         }
166         fprintf(f, "%d", a->els[i].n);
167     }
168     return fprintf(f, "]");
169 }

```

```

171 static int
172 Lfmt(Fmt *f)
173 {
174     Node **x, *l;
175     int i;

177     x = va_arg(f->args, Node**);
178     l = *x;
179     fmtprint(f, "(");
180     for(i = 0; l != nil; l = l->next){
181         if(i++ > 0){
182             fmtprint(f, ", ");
183             if((i%I2LN) == 0)
184                 fmtprint(f, "\n");
185         }
186         fmtprint(f, "%d", l->n);
187     }
188     return fmtprint(f, ")");
189 }

191 static int
192 Pfmt(Fmt *f)
193 {
194     Ptrs *a;
195     int i;

197     a = va_arg(f->args, Ptrs*);
198     fmtprint(f, "[");
199     for(i = 0; i < a->nels; i++){
200         if(i > 0){
201             fmtprint(f, ", ");
202             if(i < a->nels-1 && (i%I2LN) == 0)
203                 fmtprint(f, "\n");
204         }
205         fmtprint(f, "%d", a->els[i]->n);
206     }
207     return fmtprint(f, "]");
208 }

210 static vlong
211 test(int (*ins)(void*, int), void *a, int n, int dir)
212 {
213     vlong t0, t1, tot;
214     int i, r;
215     char *s;

```

```

217     tot = 0LL;
218     for(i = 0; i < n; i++){
219         r = nrand(n);
220         t0 = nsec();
221         switch(dir){
222             case Fwd:
223                 r = ins(a, i);
224                 break;
225             case Bck:
226                 r = ins(a, n-i);
227                 break;
228             default:
229                 r = ins(a, r);
230         }
231         if(r < 0){
232             print("no memory");
233             exits(nil);
234         }
235         t1 = nsec();
236         tot += (t1 - t0);
237         if(otherallocs){
238             s = malloc(64);
239             USED(s);
240         }
241     }
242     return tot;
243 }

```

```

245 static void
246 afree(void *x)
247 {
248     Arry *a;
249
250     a = x;
251     free(a->els);
252     free(a);
253 }

```

```

255 static void
256 lfree(void *x)
257 {
258     Node **a, *l, *n;
259
260     a = x;
261     l = *a;
262     while(l != nil){
263         n = l;
264         l = l->next;
265         free(n);
266     }
267     free(a);
268 }

```

```

270 static void
271 pfree(void *x)
272 {
273     Ptrs *a;
274     int i;

```

```

276     a = x;
277     for(i = 0; i < a->nels; i++)
278         free(a->els[i]);
279     free(a->els);
280     free(a);
281 }

283 static void
284 adump(void *x)
285 {
286     print("%A\n", x);
287 }

289 static void
290 ldump(void *x)
291 {
292     print("%L\n", x);
293 }

295 static void
296 pdump(void *x)
297 {
298     print("%P\n", x);
299 }

302 void
303 main(int argc, char*argv[])
304 {
305     void *x;
306     int dir, adt, n, vflag, wflag;
307     vlong d;
308     struct{
309         void>(*new)(void);
310         int(*ins)(void*,int);
311         void(*free)(void*);
312         void(*dump)(void*);
313     } fns[] = {
314         [Tarry]    {anew, ains, afree, adump},
315         [Tlist]   {lnew, lins, lfree, ldump},
316         [Tptrs]   {pnew, pins, pfree, pdump},
317     };

```

```

319     dir = Fwd;
320     adt = Tarry;
321     n = Num;
322     d = 0;
323     vflag = wflag = 0;
324     ARGBEGIN{
325     case 'a':
326         break;
327     case 'l':
328         adt = Tlist;
329         break;
330     case 'p':
331         adt = Tptrs;
332         break;
333     case 'f':
334         break;
335     case 'b':
336         dir = Bck;
337         break;
338     case 'r':
339         dir = Rnd;
340         break;
341     case 'i':
342         incr = atoi(EARGF(usage()));
343         if(incr < 1)
344             sysfatal("incr < 1");
345         break;
346     case 'n':
347         n = atoi(EARGF(usage()));
348         if(n < 1)
349             sysfatal("n < 1");
350         break;
351     case 'e':
352         elsz = atoi(EARGF(usage()));
353         if(elsz < sizeof(int))
354             sysfatal("elsz < 1");
355         if(elsz%sizeof(int))
356             sysfatal("elsz is not a sizeof(int) multiple");
357         break;
358     case 'm':
359         otherallocs = 1;
360         break;
361     case 'v':
362         vflag = 1;
363         break;
364     case 'w':
365         wflag = 1;
366         break;
367     default:
368         usage();
369     }ARGEND;
370     if(argc != 0)
371         usage();
372     fmtinstall('A', Afmt);
373     fmtinstall('L', Lfmt);
374     fmtinstall('P', Pfmt);
375     srand(13);
376     x = fns[adt].new();
377     d = test(fns[adt].ins, x, n, dir);

```

```
378     if(vflag)
379         fns[adt].dump(x);
380     print("%lld\n", d);
381     if(wflag){
382         print("waiting...\n");
383         sleep(3600 * 1000);
384     }
385     if(0)     fns[adt].free(x);
386     exits(nil);
387 }
```

—

Appendix B: C++ Source code

listarry.cpp

```
1  #include <iostream>
2  #include <vector>
3  #include <list>

5  enum
6  {
7      Num = 1000,

9      Fwd = 0,
10     Bck = 1,
11     Rnd = -1,

13     Tarry = 0,
14     Tlist,
15 };

18 using namespace std;

20 const long long CLK2MS = CLOCKS_PER_SEC / 1000000;
21 const int N = 500;

23 /*
24  * Definition of ARGBEGIN, ARGEND, ... omitted.
25  */

27 void usage()
28 {
29     cerr << "usage: " << argv0 << " [-alfbr] [-n num]" << endl;
30     exit(1);
31 }

33 int
34 main(int argc, char *argv[])
35 {
36     vector<int> array;
37     vector<int>::iterator vi;
38     list<int> lst;
39     list<int>::iterator li;
40     int i, r;
41     clock_t t0, t1;
42     long long tot;
43     void *s;
44     int adt, dir, n, otherallocs, vflag;
```



```

46     dir = Fwd;
47     adt = Tarry;
48     n = Num;
49     otherallocs = vflag = 0;
50     ARGBEGIN{
51     case 'a':
52         break;
53     case 'l':
54         adt = Tlist;
55         break;
56     case 'f':
57         break;
58     case 'b':
59         dir = Bck;
60         break;
61     case 'r':
62         dir = Rnd;
63         break;

65     case 'n':
66         n = atoi(EARGF(usage()));
67         if(n < 1){
68             cerr << "n < 1" << endl;
69             return 1;
70         }
71         break;
72     case 'm':
73         otherallocs = 1;
74         break;
75     case 'v':
76         vflag = 1;
77         break;
78     }ARGEND;
79     if(argc != 0)
80         usage();

82     tot = 0;
83     srand(13);
84     for(i = 0; i < n; i++){
85         switch(dir){
86         case Fwd:
87             r = i;
88             break;
89         case Bck:
90             r = n-i;
91             break;
92         default:
93             r = rand() % n;
94         }

```

```

96         t0 = clock();
97         /* No dispatching here */
98         if(adt == Tarry){
99             for(vi = arry.begin(); vi != arry.end(); vi++)
100                 if(*vi >= r){
101                     arry.insert(vi, r);
102                     break;
103                 }
104             if(vi == arry.end())
105                 arry.push_back(r);

107         }else{
108             for(li = lst.begin(); li != lst.end(); li++)
109                 if(*li >= r){
110                     lst.insert(li, r);
111                     break;
112                 }
113             if(li == lst.end())
114                 lst.push_back(r);
115         }
116         t1 = clock();
117         tot += t1 - t0;
118         if(otherallocs)
119             s = malloc(64);
120     }

122     if(vflag){
123         cout << "{";
124         if(adt == Tarry)
125             for(vi = arry.begin(); vi != arry.end(); vi++)
126                 cout << " " << *vi;
127         else
128             for(li = lst.begin(); li != lst.end(); li++)
129                 cout << " " << *li;
130         cout << "}" << endl;
131     }
132     cout << tot * CLK2MS << "000" << endl;
133     return 0;
134 }

```

—