Locating and Addressing Performance Issues

Diomidis Spinellis
www.spinellis.gr

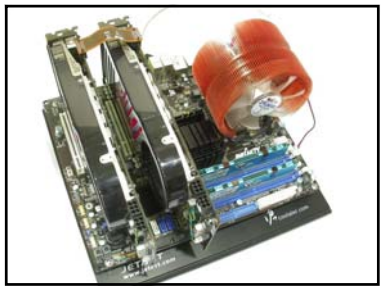

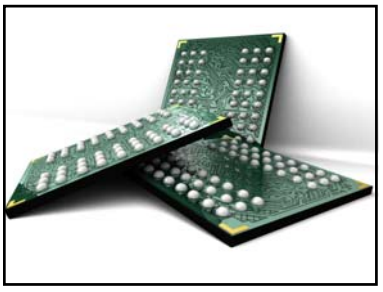

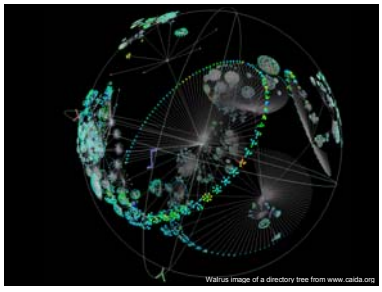

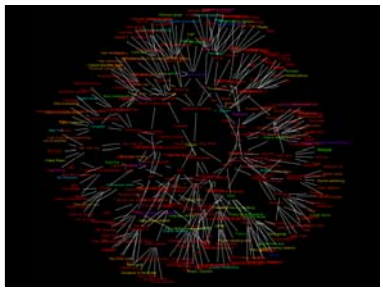

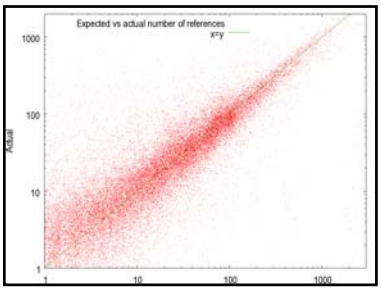
Expected vs actual number of references
x=y
Actual
1000
100
10
1
1
10
100
1000


RAIN BIRD
ESP MODULAR


PLAYSTATION 3

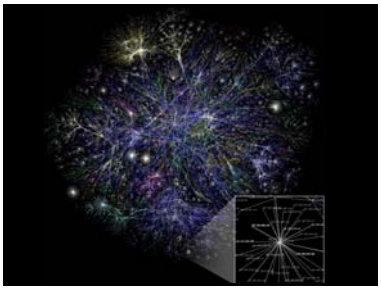

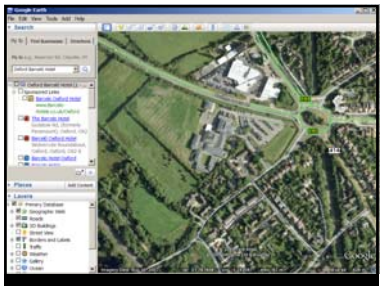

10

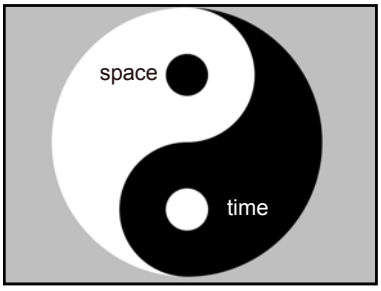
space
time

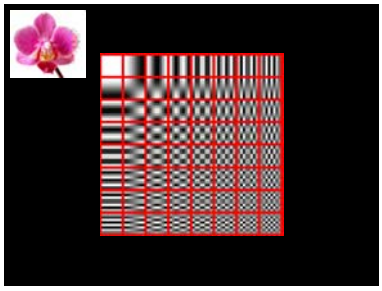

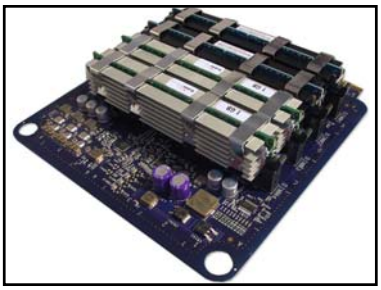

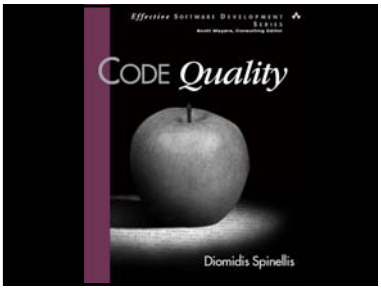
Effective Software Development Series
Code Quality
Diomidis Spinellis

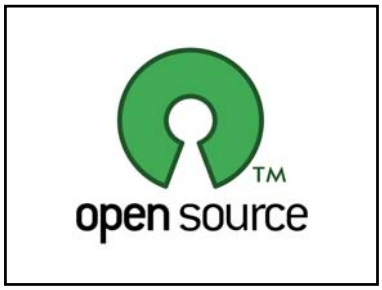
open source

[1]
Know your workload

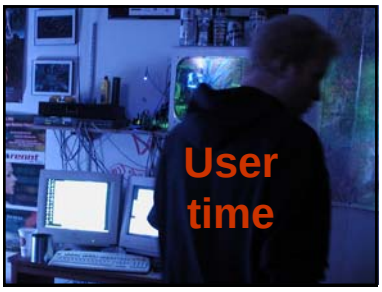

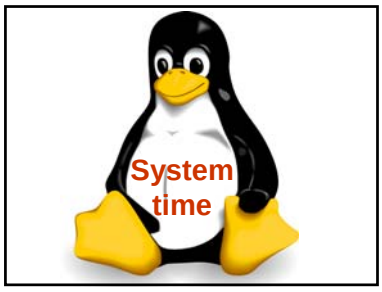

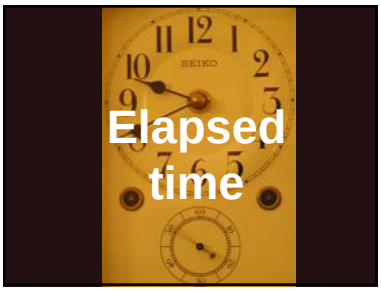


```
$ time cmd

real   1585m18.021s
user   1546m26.338s
sys      97m18.364s
```

```
$ /usr/bin/time cmd

92786.33 user
 5838.36 system
26:25:18 elapsed
```

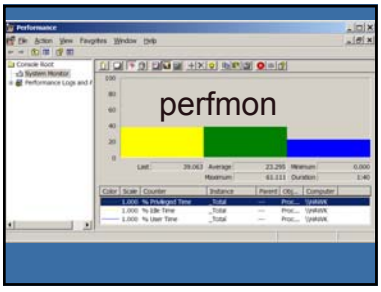

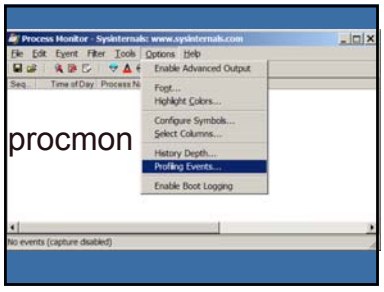

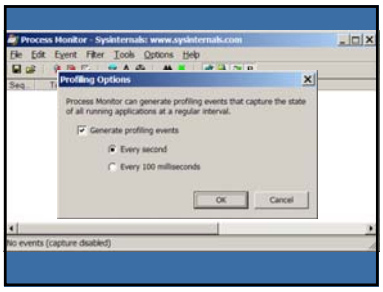
Process Monitor - Sysinternals: www.sysinternals.com
File Edit Event Filter Tools Options Help
Profiling Options
Process Monitor can generate profiling events that capture the state of all running applications at a regular interval.
Generate profiling events
Every second
Every 100 milliseconds
OK
Cancel
No events (capture disabled)

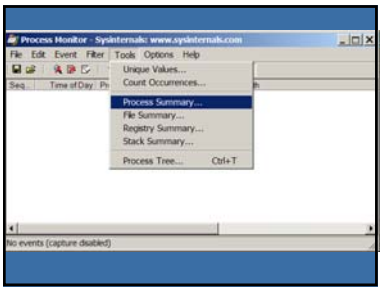
Process Monitor - Sysinternals: www.sysinternals.com
File Edit Event Filter Tools Options Help
Unique Values...
Count Occurrences...
Process Summary...
File Summary...
Registry Summary...
Stack Summary...
Process Tree... Ctrl+T
No events (capture disabled)

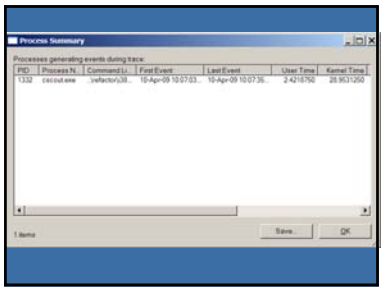
Process Summary
Processes generating events during trace:
Save...
OK

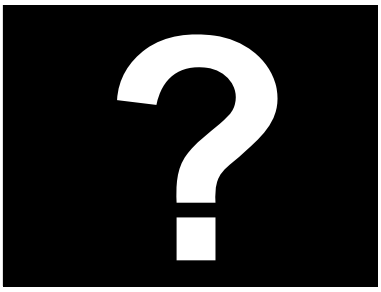
?


3


Elapsed time


System time

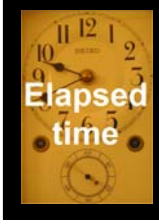

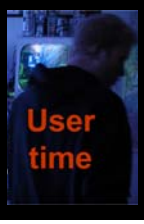

[2]

Look at your code

```
$ /usr/bin/time sed -n\
"s/\(.\)\(.\)\(.\)\3\2\1/(\1\2\3-\3\2\1)/p"\
/usr/share/dict/words

[...]
(col-loc)ation [...]
g(ram-mar) [...]
sh(red-der) [...]
s(nif-fin)g [...]

203.59 real 194.27 user 1.63 sys
```

CFLAGS+=-pg

```
$ make clean; make
[...]

$ /usr/bin/time sed -n\
"s/\(.\)\(.\)\(.\)\3\2\1/(\1\2\3-\3\2\1)/p"\
/usr/share/dict/words
[...]

$ ls -l sed.gmon
-rw-r--r--  1 dds dds 74358 Apr 15 sort.gmon

$ gprof sed
```

```
%   cumulative   self              self     total
 time   seconds   seconds    calls  ms/call  ms/call  name
 62.8     122.71   122.71 25539316     0.00     0.00  sstep
 14.5     151.09    28.39                             .mcount
  7.6     165.90    14.80  2171863     0.01     0.04  sslow
  5.1     175.77     9.87  1321712     0.01     0.04  sfast
  4.2     183.93     8.16  1086032     0.01     0.01  sbackref
  1.9     187.73     3.80   235881     0.02     0.68  smatcher
[...]
  0.0     195.40     0.00        1     0.00     0.03  vfprintf
```

```
 called/total       parents
%time    self descendents  called+self    name
                           called/total       children

                                          <spontaneous>
100.0    0.00      167.02                 main
         0.61      166.41       1/1           process
         0.00        0.00       1/2           fclose
         0.00        0.00       1/1           compile
         0.00        0.00       1/1           add_compunit
         0.00        0.00       1/1           add_file
         0.00        0.00       2/2           getopt
         0.00        0.00       1/1           setlocale
         0.00        0.00       1/1           cfclose
         0.00        0.00       1/1           exit
```

```
 called/total       parents
%time    self descendents  called+self    name
                           called/total       children

         3.80      157.44  235881/235881      regexec
 96.5    3.80      157.44  235881         smatcher
        14.80       76.11 2171863/2171863     sslow
         9.87       46.59 1321712/1321712     sfast
         8.16        0.00 1086032/1086032     sbackref
         0.47        1.43  218749/218760      malloc
         0.00        0.00     201/204         free
```

```
 called/total       parents
%time    self descendents  called+self    name
                           called/total       children

        46.59        0.00 9607679/25539316     sfast
        76.11        0.00 15841637/25539316    sslow
 73.5  122.71        0.00 25539316             sstep
```

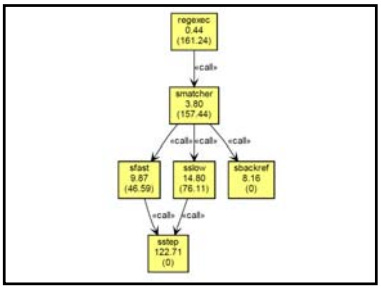

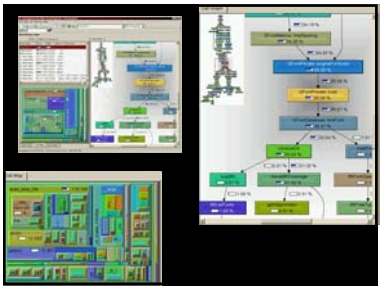

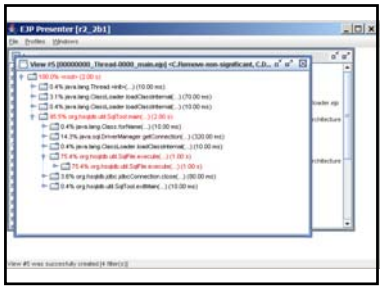

interactions

```
$ ltrace -c sed -n \
"s/\(.\)\(.\)\(.\)\3\2\1/(\1\2\3-\3\2\1)/p" words
[...]

% time     seconds  usecs/call     calls function
------ ----------- ----------- --------- --------
 33.12   47.109398         199    235882 regexec
 29.35   41.753647         177    235882 fgetln
 19.42   27.621235         117    235882 ungetc
 18.08   25.719493         108    237490 memmove
  0.02    0.035511         176       201 fputc
```

```
$ ltrace paste expr.c paste.c >/dev/null

fgets("/*\t$NetBSD: expr.c,v 1.5 1997/07"..., 2049,
0x08049260) = 0xbffff418

strchr("/*\t$NetBSD: expr.c,v 1.5 1997/07"..., '\n') =
"\n"

printf("%s", "/*\t$NetBSD: expr.c,v 1.5 1997/07"...) = 62

fgets("/*\t$NetBSD: paste.c,v 1.4 1997/1"..., 2049,
0x080493e8) = 0xbffff418

strchr("/*\t$NetBSD: paste.c,v 1.4 1997/1"..., '\n') =
"\n"
```

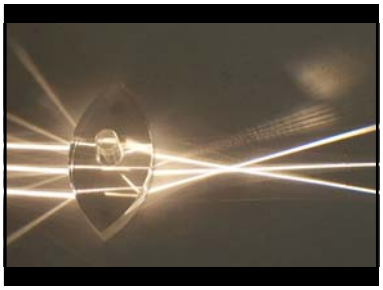

CFLAGS+=-pg -g -fprofile-arcs \
-ftest-coverage -O0

```
$ make clean; make
[...]

$ ./myprog
$ ./myprog

$ ls -l myprog.*
-rw-r--r--  1 dds   820 16:41 myprog.gcno
-rw-r--r--  1 dds   200 16:42 myprog.gcda


$ gcov myprog
```

```
        -:    0:Graph:myprog.gcno
        -:    0:Data:myprog.gcda
        -:    0:Runs:2
        -:    0:Programs:1

        -:    1:main(int argc, char *argv[])
        2:    2:{
        -:    3:    volatile int i, j, k;
        2:    4:    volatile int count = 0;
        -:    5:
      202:    6:    for (i = 0; i < 100; i++)
    20200:    7:        for (j = 0; j < 100; j++)
  2020000:    8:            for (k = 0; k < 100; k++)
  2000000:    9:                count++;
        -:   10:
        2:   11:}
```

[3]
Use better algorithms

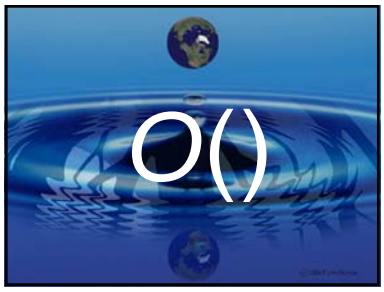

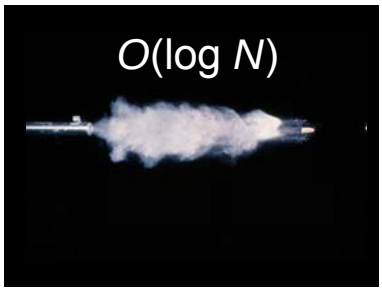

1,000 10
1,000,000 20
1,000,000,000 30

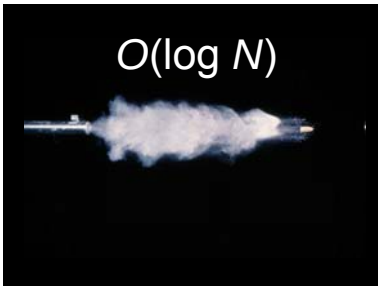

right mind

```
#include <list>
#include <utility>
#include <limits>

#include "person.h"

int matchQuality(const Person &a, const Person &b, int numPersons);

typedef std::list<Person> PersonContainer;
typedef PersonContainer::const_iterator PersonIterator;
typedef std::pair <PersonIterator, PersonIterator> PersonPair;

PersonPair
bestMatch(PersonContainer &pc)
{
    PersonPair bestMatch(pc.end(), pc.end());
    int q, bestMatchQuality = std::numeric_limits<int>::min();

    for (PersonIterator i = pc.begin(); i != pc.end(); i++)
        for (PersonIterator j = pc.begin(); j != pc.end(); j++)
            if ((q = matchQuality(*i, *j, pc.size())) >
                bestMatchQuality) {
                bestMatchQuality = q;
                bestMatch = PersonPair(i, j);
            }
    return bestMatch;
}
```

catastrophic

lookup

CREATE INDEX …

reverse

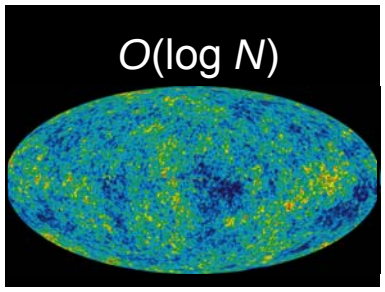

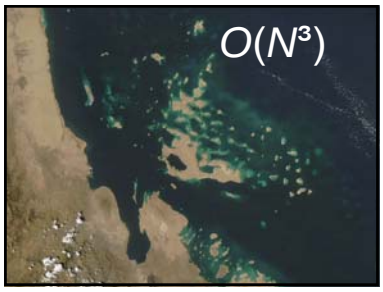

[4]
Know your computer

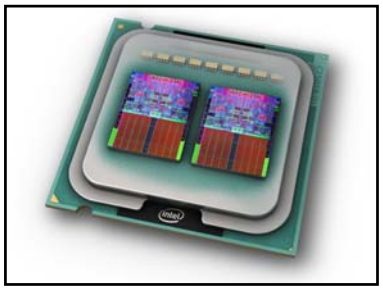

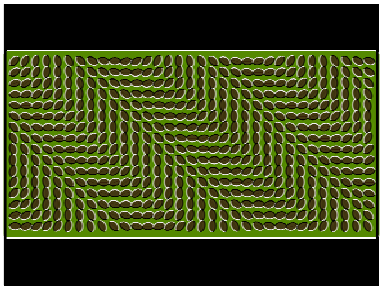

performance counters

oprofile
perfmon2

CPU_CLK_UNHALTED INST_RETIRED_ANY_P L2_RQSTS LLC_MISSES LLC_REFS
LOAD_BLOCK STORE_BLOCK MISALIGN_MEM_REF SEGMENT_REG_LOADS SSE_PRE_EXEC
DTLB_MISSES MEMORY_DISAMBIGUATION PAGE_WALKS FLOPS FP_ASSIST MUL DIV
CYCLES_DIV_BUSY IDLE_DURING_DIV DELAYED_BYPASS L2_ADS L2_DBUS_BUSY_RD
L2_LINES_IN L2_M_LINES_IN L2_LINES_OUT L2_M_LINES_OUT L2_IFETCH L2_LD
L2_ST L2_LOCK L2_REJECT_BUSQ L2_NO_REQ EIST_TRANS_ALL THERMAL_TRIP
L1D_CACHE_LD L1D_CACHE_ST L1D_CACHE_LOCK L1D_CACHE_LOCK_DURATION
L1D_ALL_REF L1D_ALL_CACHE_REF L1D_REPL L1D_M_REPL L1D_M_EVICT
L1D_PEND_MISS L1D_SPLIT SSE_PREF_MISS LOAD_HIT_PRE L1D_PREFETCH
BUS_REQ_OUTSTANDING BUS_BNR_DRV BUS_DRDY_CLOCKS BUS_LOCK_CLOCKS
BUS_DATA_RCV BUS_TRAN_BRD BUS_TRAN_RFO BUS_TRAN_WB BUS_TRAN_IFETCH
BUS_TRAN_INVAL BUS_TRAN_PWR BUS_TRANS_P BUS_TRANS_IO BUS_TRANS_DEF
BUS_TRAN_BURST BUS_TRAN_MEM BUS_TRAN_ANY EXT_SNOOP CMP_SNOOP BUS_HIT_DRV
BUS_HITM_DRV BUSQ_EMPTY SNOOP_STALL_DRV BUS_IO_WAIT L1I_READS L1I_MISSES
ITLB INST_QUEUE_FULL IFU_MEM_STALL ILD_STALL BR_INST_EXEC BR_MISSP_EXEC
BR_BAC_MISSP_EXEC BR_CND_EXEC BR_CND_MISSP_EXEC BR_IND_EXEC
BR_IND_MISSP_EXEC BR_RET_EXEC BR_RET_MISSP_EXEC BR_RET_BAC_MISSP_EXEC
BR_CALL_EXEC BR_CALL_MISSP_EXEC BR_IND_CALL_EXEC BR_TKN_BUBBLE_1
BR_TKN_BUBBLE_2 RS_UOPS_DISPATCHED RS_UOPS_DISPATCHED_NONE MACRO_INSTS
ESP SIMD_UOPS_EXEC SIMD_SAT_UOP_EXEC SIMD_UOP_TYPE_EXEC INST_RETIRED
X87_OPS_RETIRED UOPS_RETIRED MACHINE_NUKES_SMC BR_INST_RETIRED
BR_MISS_PRED_RETIRED CYCLES_INT_MASKED SIMD_INST_RETIRED HW_INT_RCV
ITLB_MISS_RETIRED SIMD_COMP_INST_RETIRED MEM_LOAD_RETIRED FP_MMX_TRANS
MMX_ASSIST SIMD_INSTR_RET SIMD_SAT_INSTR_RET RAT_STALLS SEG_RENAME_STALLS
SEG_RENAMES RESOURCE_STALLS BR_INST_DECODED BR_BOGUS BACLEARS
PREF_RQSTS_UP PREF_RQSTS_DN

ratios

CPU_CLK_HALTED /
INST_RETIRED

CYCLES_L1I_MEM_STALLED /
CPU_CLK_HALTED

L1D_REPL / INST_RETIRED

L2_LINES_IN / INST_RETIRED

[5]
Avoid I/O

boring

```
strace
truss
```

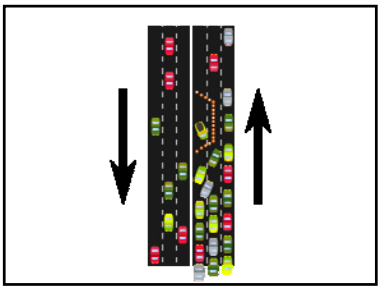

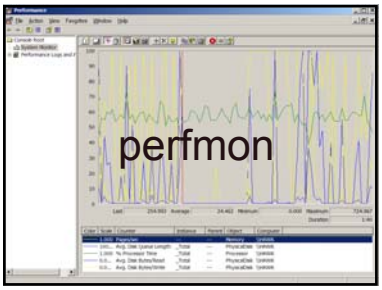


```
iostat
netstat
nfsstat
vmstat
```

example

```
$ /usr/bin/time cp /usr/share/dict/words wordcopy
    5.68 real         0.00 user         0.32 sys
```

```
$ iostat 1

 ad0
  KB/t tps  MB/s
  0.00   0  0.00
 32.80  15  0.47
 27.12  24  0.63
 37.31  26  0.94
 73.60  10  0.71
 35.24  33  1.14
 25.14  21  0.51
  7.00   4  0.03
  0.00   0  0.00
```

```
$ netstat 1

 input        (Total)           output
    packets  errs      bytes    packets  errs      bytes colls
          1     0         60          1     0        250     0
        219     0     237744        204     0     230648   113
        417     0     515196        418     0     513722   324
        383     0     467208        402     0     496650   292
        368     0     451418        381     0     470212   259
        425     0     519588        430     0     515714   301
        400     0     488434        400     0     496816   287
          9     0       6106         15     0      11886     7
          1     0         60          1     0        138     0
```

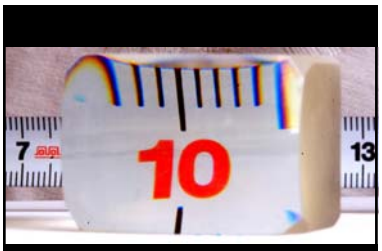


```
199.22.61.2 - - [16/Apr/2009:20:56:25 +0300] "GET /
HTTP/1.0" 200 6701
"http://www.google.ca/search?hl=fr&as_qdr=all&num=100&q=um
lgraph%3C&meta=" "Mozilla/4.0 (compatible; MSIE 6.0;
Windows NT 5.1; SV1; .NET CLR 1.1.4322; .NET CLR
2.0.50727; .NET CLR 3.0.04506.30)"

gb8.hydro.qc.ca - - [16/Apr/2009:20:56:25 +0300] "GET /
HTTP/1.0" 200 6701
"http://www.google.ca/search?hl=fr&as_qdr=all&num=100&q=um
lgraph%3C&meta=" "Mozilla/4.0 (compatible; MSIE 6.0;
Windows NT 5.1; SV1; .NET CLR 1.1.4322; .NET CLR
2.0.50727; .NET CLR 3.0.04506.30)"
```

```
$ /usr/bin/time logresolve <httpd-access.log >/dev/null
     1230.55 real         0.04 user         0.03 sys
```

```
$ netstat 1

 input        (Total)           output
    packets  errs      bytes    packets  errs      bytes colls
          7     0        486          8     0        108     0
         14     0        229         11     0        383     0
          3     0        336          3     0        324     0
          3     0        216          4     0        301     0
          3     0        667          3     0        216     0
          6     0         98          2     0        301     0
```

```
$ tcpdump port domain

16:15:33.283221 istlab.dmst.aueb.gr.1024 >
gns1.nominum.com.domain:9529 [1au] PTR?
105.199.133.198.in-addr.arpa. (57)

16:15:33.433305 gns1.nominum.com.domain >
istlab.dmst.aueb.gr.1024:9529*- 1/2/0 (122)
(DF) [tos 0x80]
```

```
| xargs -P 40
```

caching

locality of reference

[6]
Avoid OS interactions

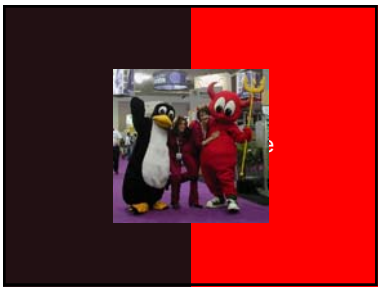

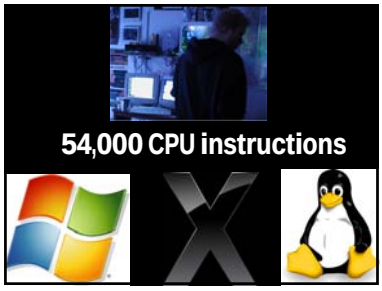


```
struct __ucontext {
    struct __sigset {
        __uint32_t __bits[4];
    } uc_sigmask;
    struct __mcontext {
        __register_t  mc_onstack, mc_rdi, mc_rsi, mc_rdx, mc_rcx, mc_r8,
                      mc_r9, mc_rax, mc_rbx, mc_rbp, mc_r10, mc_r11, mc_r12,
                      mc_r13, mc_r14, mc_r15, mc_trapno, mc_addr, mc_flags,
                      mc_err, mc_rip, mc_cs, mc_rflags, mc_rsp, mc_ss;
        long    mc_len;
        long    mc_fpformat;
        long    mc_ownedfp;
        long    mc_fpstate[64];
        long    mc_spare[8];
    } uc_mcontext;
    struct __ucontext *uc_link;
    struct {
        char    *ss_sp;
        __size_t ss_size;
        int     ss_flags;
    } uc_stack;
    int         uc_flags;
    int         __spare__[4];
} ucontext_t;
```

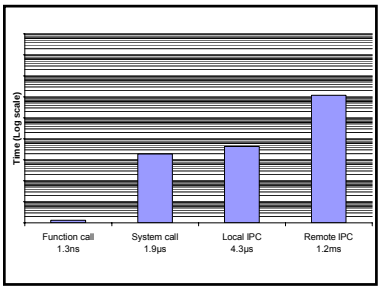


```
strace -c
```

```
$ strace -c ls -lR >/dev/null

% time     seconds usecs/call     calls    errors syscall
------ ----------- ----------- --------- --------- ----------
 25.69    0.002745           0    166243           lstat64
 22.02    0.002353           0    166018    166018 getxattr
 21.83    0.002333           0    166243    166243 lgetxattr
 14.93    0.001595           0    166250         7 stat64
  8.98    0.000960           0     44052           getdents64
  4.11    0.000439           0     22026        39 open
  1.38    0.000147           0     43939           fstat64
  0.95    0.000102           0     21993           close
[...]
```

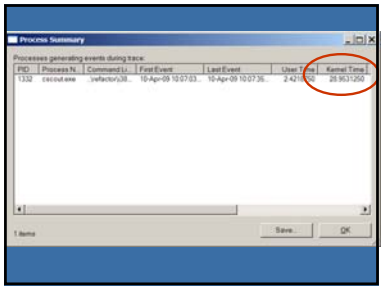

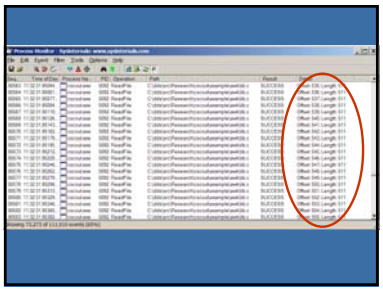

```
$ strace cscout awk.cs
[...]
_llseek(6, -8190, [535], SEEK_CUR)      = 0
read(6, "ng or publicity pertaining\nto di"..., 8191) = 8191
_llseek(6, -8190, [536], SEEK_CUR)      = 0
read(6, "g or publicity pertaining\nto dis"..., 8191) = 8191
_llseek(6, -8190, [537], SEEK_CUR)      = 0
read(6, " or publicity pertaining\nto dist"..., 8191) = 8191
_llseek(6, -8190, [538], SEEK_CUR)      = 0
read(6, "or publicity pertaining\nto distr"..., 8191) = 8191
_llseek(6, -8190, [539], SEEK_CUR)      = 0
read(6, "r publicity pertaining\nto distri"..., 8191) = 8191
_llseek(6, -8190, [540], SEEK_CUR)      = 0
read(6, " publicity pertaining\nto distrib"..., 8191) = 8191
_llseek(6, -8190, [541], SEEK_CUR)      = 0
read(6, "publicity pertaining\nto distribu"..., 8191) = 8191
_llseek(6, -8190, [542], SEEK_CUR)      = 0
read(6, "ublicity pertaining\nto distribut"..., 8191) = 8191
_llseek(6, -8190, [543], SEEK_CUR)      = 0
read(6, "blicity pertaining\nto distributi"..., 8191) = 8191
```

```
ifstream in(fname.c_str(), ios::binary);

for (;;) {
        off_t x = in.tellg();
        if ((val = in.get()) == EOF)
                break;
        // [...]
}

in.close();
```

```
class fifstream {
private:
        ifstream i;             // The underlying ifstream
        bool dirty;             // True if we can't use mypos
        long mypos;             // Cached position
public:
        fifstream(const char *s, ios_base::openmode mode) :
            i(s, mode), dirty(true) {
                assert(mode & ios::binary);
        }
        ifstream::pos_type tellg() {
                if (dirty) {
                        dirty = false;
                        mypos = (unsigned long)i.tellg();
                }
                return (ifstream::pos_type)mypos;
        }
        ifstream::int_type get() {
                ifstream::int_type r = i.get();
                if (r == EOF)
                        dirty = true;
                else
                        mypos++;
                return r;
        }
        fifstream &putback(char c) {
                i.putback(c);
                mypos--;
```

[7]

Look behind your back

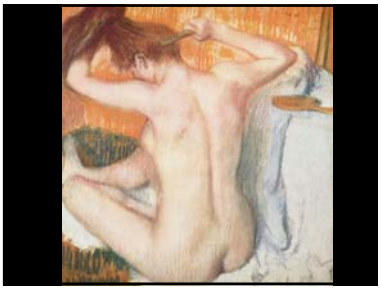

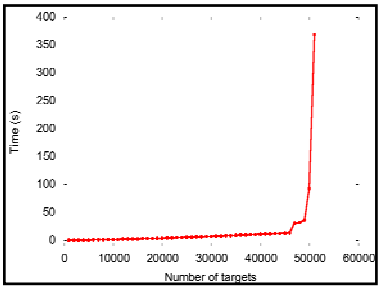

trash

VM

locality of reference (again)

unacceptable

3

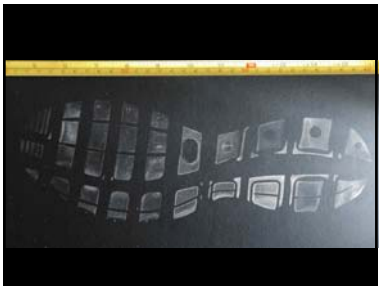

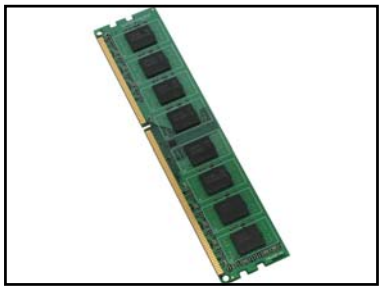

locality

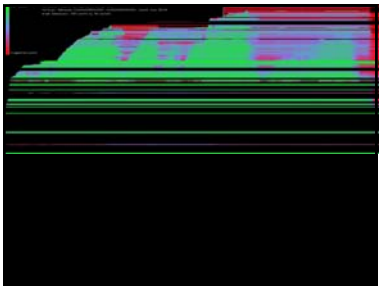

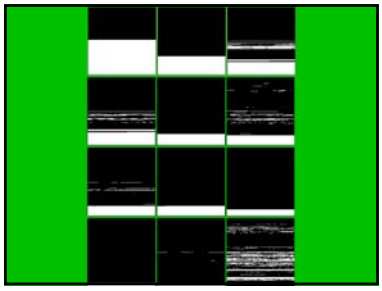

interrupts

```
$ cat /proc/interrupts
           CPU0       CPU1
  0:      14487          0   IO-APIC-edge      timer
  1:          2          0   IO-APIC-edge      i8042
  4:    1292152          0   IO-APIC-edge      serial
  6:          5          0   IO-APIC-edge      floppy
  7:          0          0   IO-APIC-edge      parport0
  8:          1          0   IO-APIC-edge      rtc0
  9:          0          0   IO-APIC-fasteoi   acpi
 12:          4          0   IO-APIC-edge      i8042
 16:   12124066          0   IO-APIC-fasteoi   uhci_hcd:usb1
 20:    9637759          0   IO-APIC-fasteoi   ata_piix
220:   12784761          0   PCI-MSI-edge      eth0
NMI:          0          0   Non-maskable interrupts
LOC:  133422880   33597793   Local timer interrupts
RES:     744742      21805   Rescheduling interrupts
CAL:     287367        539   function call interrupts
TLB:      85455     324100   TLB shootdowns
TRM:          0          0   Thermal event interrupts
SPU:          0          0   Spurious interrupts
```

[8]
Know your storage

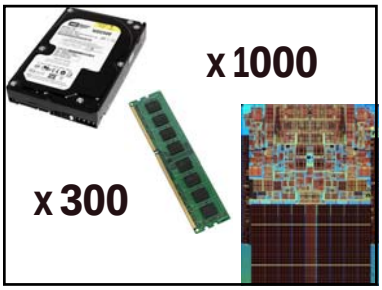

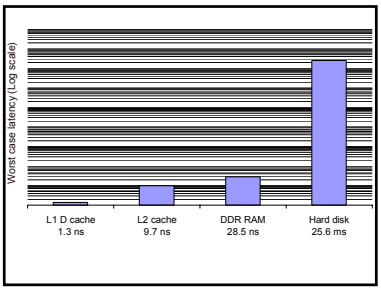

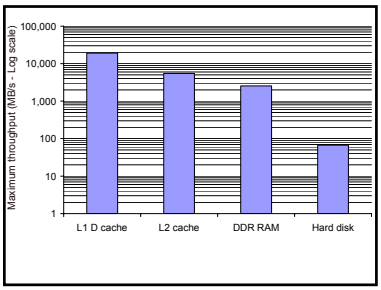

practice?

minimize

locality of reference (again²)

```
const int N = 518;

double a[N][N], r[N][N];

for (int i = 0; i < N; i++)
    for (int j = 0; j < N; j++) {
        r[i][j] = 0;
        for (int k = 0; k < N; k++)
            r[i][j] += a[i][k] * a[k][j];
    }
```

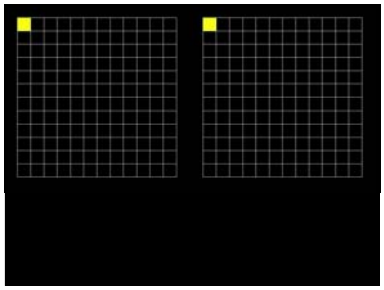

```
const int N = 518;

double a[N][N], r[N][N];
const int TILE = 7;


assert(N % TILE == 0);

for (int i = 0; i < N; i += TILE)
  for (int j = 0; j < N; j += TILE)
    for (int k = 0; k < N; k += TILE)
      for (int ii = i; ii < i + TILE; ii++)
        for (int jj = j; jj < j + TILE; jj++) {
          r[ii][jj] = 0;
          for (int kk = k; kk < k + TILE; kk++)
            r[ii][jj] += a[ii][kk] * a[kk][jj];
        }
```

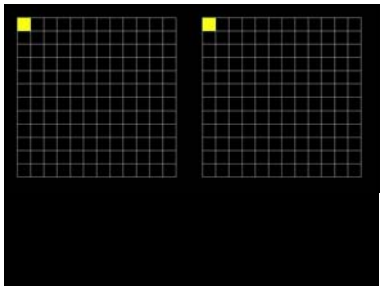

oprofile

```
# opreport image:timetile
CPU: Core 2, speed 2126.44 MHz (estimated)
Counted L1D_REPL events (Cache lines allocated in
the L1 data cache) with a unit mask of 0x0f (No
unit mask) count 2000
Counted L2_LINES_IN events (number of allocated
lines in L2) with a unit mask of 0x70 (multiple
flags) count 2000

    L1D_REPL:2000| L2_LINES_IN:2000|
  samples|      %|  samples|      %|
-----------------  -----------------
   169847 100.000      10525 100.000 timetile
```

```
# opreport image:timetile
CPU: Core 2, speed 2126.44 MHz (estimated)
Counted L1D_REPL events (Cache lines allocated in
the L1 data cache) with a unit mask of 0x0f (No
unit mask) count 2000
Counted L2_LINES_IN events (number of allocated
lines in L2) with a unit mask of 0x70 (multiple
flags) count 2000

    L1D_REPL:2000| L2_LINES_IN:2000|
  samples|      %|  samples|      %|
-----------------  -----------------
     4276 100.000        711 100.000 timetile
```

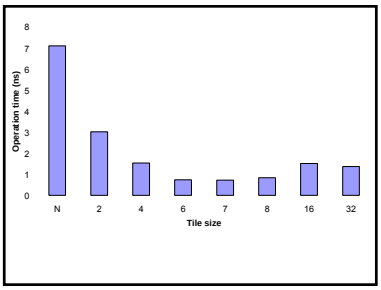

Don't try this at home!

The March of Progress

```
1978: C
    printf("%10.2f", x);

1988: C++
    cout << setw(10) << setprecision(2) << showpoint << x;

1996: Java
    java.text.NumberFormat formatter = java.text.NumberFormat.getNumberInstance();
    formatter.setMinimumFractionDigits(2);
    formatter.setMaximumFractionDigits(2);
    String s = formatter.format(x);
    for (int i = s.length(); i < 10; i++)
        System.out.print(' ');
    System.out.print(s);

2004: Java
    System.out.printf("%10.2f", x);

2008: Scala and Groovy
    printf("%10.2f", x)
```

[9]
Know your data

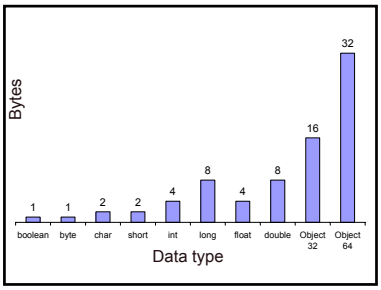

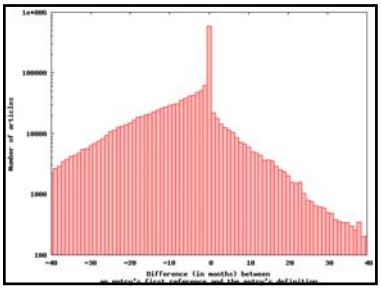


```
class RangeMap {
private:
        vector<bool> active;
public:
        static const int NMONTH = (2009 - 2001) * 12;
        RangeMap() : active(NMONTH, false) {}
};

class EntryDetails {
private:
    RangeMap defined, stub;
    vector <int> nrefs;
    string definer, referrer, firstRef;
    time_t firstDef;
    int numReferences, numContributors;
    int numRevisions, numReverts;
public:
    EntryDetails() : defined(), stub(),
nrefs(RangeMap::NMONTH, 0), firstDef(-1),
numReferences(0), numContributors(0), numRevisions(0),
numReverts(0) {}
};
```

```
class RangeMap {
    private BitSet active;
    public static final int NMONTH = (2009 - 2001) * 12;
    public RangeMap() {
        active = new BitSet(NMONTH);
    }
}

class EntryDetails {
    RangeMap defined, stub;
    ArrayList <Integer> nrefs;
    String definer, referrer, firstRef;
    Date firstDef;
    int numReferences, numContributors;
    int numRevisions, numReverts;
    public EntryDetails() {
        defined = new RangeMap();
        stub = new RangeMap();
        nrefs = new ArrayList<Integer>(RangeMap.NMONTH);
    }
}
```

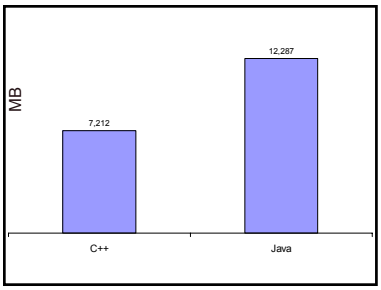

massive

C++

packing

```
struct {
        int     a;
        char    ca;
        int     b;
        char    cb;
};
```

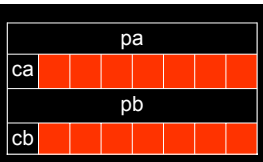


```
unsigned isActive : 1;
```

BitSet

& |= &= ~

[10]
Manage your memory

Valgrind

```
valgrind  --tool=massif \
httpd -f `pwd`/httpd.conf -X


ms_print massif.out.22917
```

n	time(i)	total(B)	useful-heap(B)	extra-heap(B)	stacks(B)
0	0	0	0	0	0
1	286,224	41,784	41,704	80	0
2	689,141	42,616	42,728	88	0
3	762,893	48,232	47,414	818	0
4	1,027,152	48,584	47,753	831	0
5	1,225,118	56,856	55,985	871	0
6	1,668,460	65,072	64,189	883	0
7	2,410,250	65,424	64,528	896	0
8	2,630,440	66,968	65,822	1,138	0
9	2,738,276	77,806	75,798	1,282	0
10	3,141,858	85,216	84,002	1,214	0
11	3,766,669	93,432	92,206	1,226	0
12	3,915,357	90,928	89,982	946	0
13	4,299,294	91,288	90,334	954	0
14	4,320,894	90,928	89,982	946	0
15	4,580,563	91,288	90,321	959	0
16	4,592,456	91,336	90,340	987	0
17	5,970,497	91,688	90,688	1,000	0
18	6,071,121	93,584	92,334	1,250	0
19	6,259,346	94,528	93,204	1,324	0
20	7,397,604	94,168	92,852	1,316	0
21	7,587,145	127,040	125,676	1,364	0
22	7,662,481	172,248	170,816	1,432	0
23	7,780,835	221,544	220,040	1,504	0
24	7,985,688	279,056	277,468	1,588	0
25	8,243,447	353,000	351,304	1,696	0

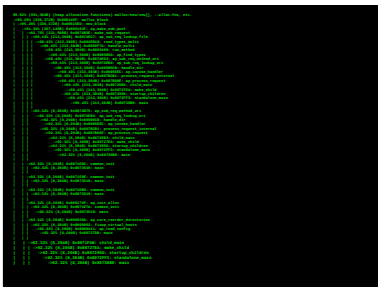

```
99.52% (351,304B) (heap allocation functions) malloc/new/new[],
--alloc-fns, etc.

->95.28% (336,372B) 0x806149F: malloc_block
| ->95.29% (336,372B) 0x80615ED: new_block
|   ->81.34% (287,140B) 0x806164F: ap_make_sub_pool
|   | ->62.76% (221,508B) 0x807A83E: make_sub_request
|   | | ->60.43% (213,304B) 0x807AEC7: ap_sub_req_lookup_file
|   | | [...]
|   | ->02.32% (8,204B) 0x8071FAB: child_main
|   | | ->02.32% (8,204B) 0x80727EA: make_child
|   | |   ->02.32% (8,204B) 0x8072B5D: startup_children
|   | |     ->02.32% (8,204B) 0x8072FF3: standalone_main
|   | |       ->02.32% (8,204B) 0x80738B9: main
|   | | [...]
->93.10% (10,812B) in 28 places, all below massif's threshold
(01.00%)
```

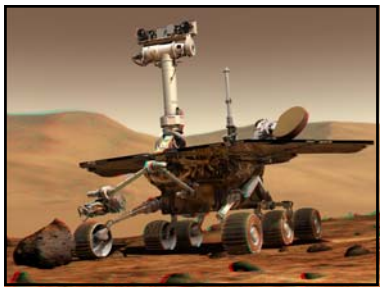

memcheck

0

But wait there’s more!

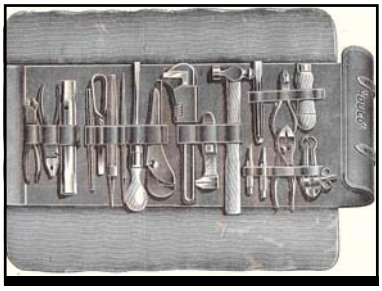

dtrace

```
$ dtrace -n 'syscall:::entry { trace(execname); }'

dtrace: description 'syscall:::entry ' matched 225 probes
CPU     ID                    FUNCTION:NAME
  1    124                      ioctl:entry   dtrace
  1    124                      ioctl:entry   dtrace
  1    260                  sysconfig:entry   dtrace
  1    260                  sysconfig:entry   dtrace
  1    194                  sigaction:entry   dtrace
  0    422                    open64:entry   gtar
  0     38                     write:entry   gtar
  0     36                      read:entry   gtar
  0    406                   fstat64:entry   gtar
  0     42                     close:entry   gtar
  0    402                    stat64:entry   gtar
  0    422                    open64:entry   gtar
  0     38                     write:entry   gtar
  0     36                      read:entry   gtar
  0     36                      read:entry   gzip
  0     36                      read:entry   gzip
  0     36                      read:entry   gzip
  0     36                      read:entry   gzip
  0     36                      read:entry   gzip
  0     38                     write:entry   gzip
```

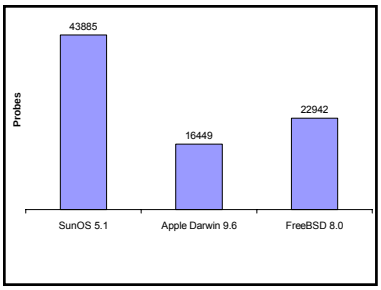

DTraceToolkit

http://www.spinellis.gr/