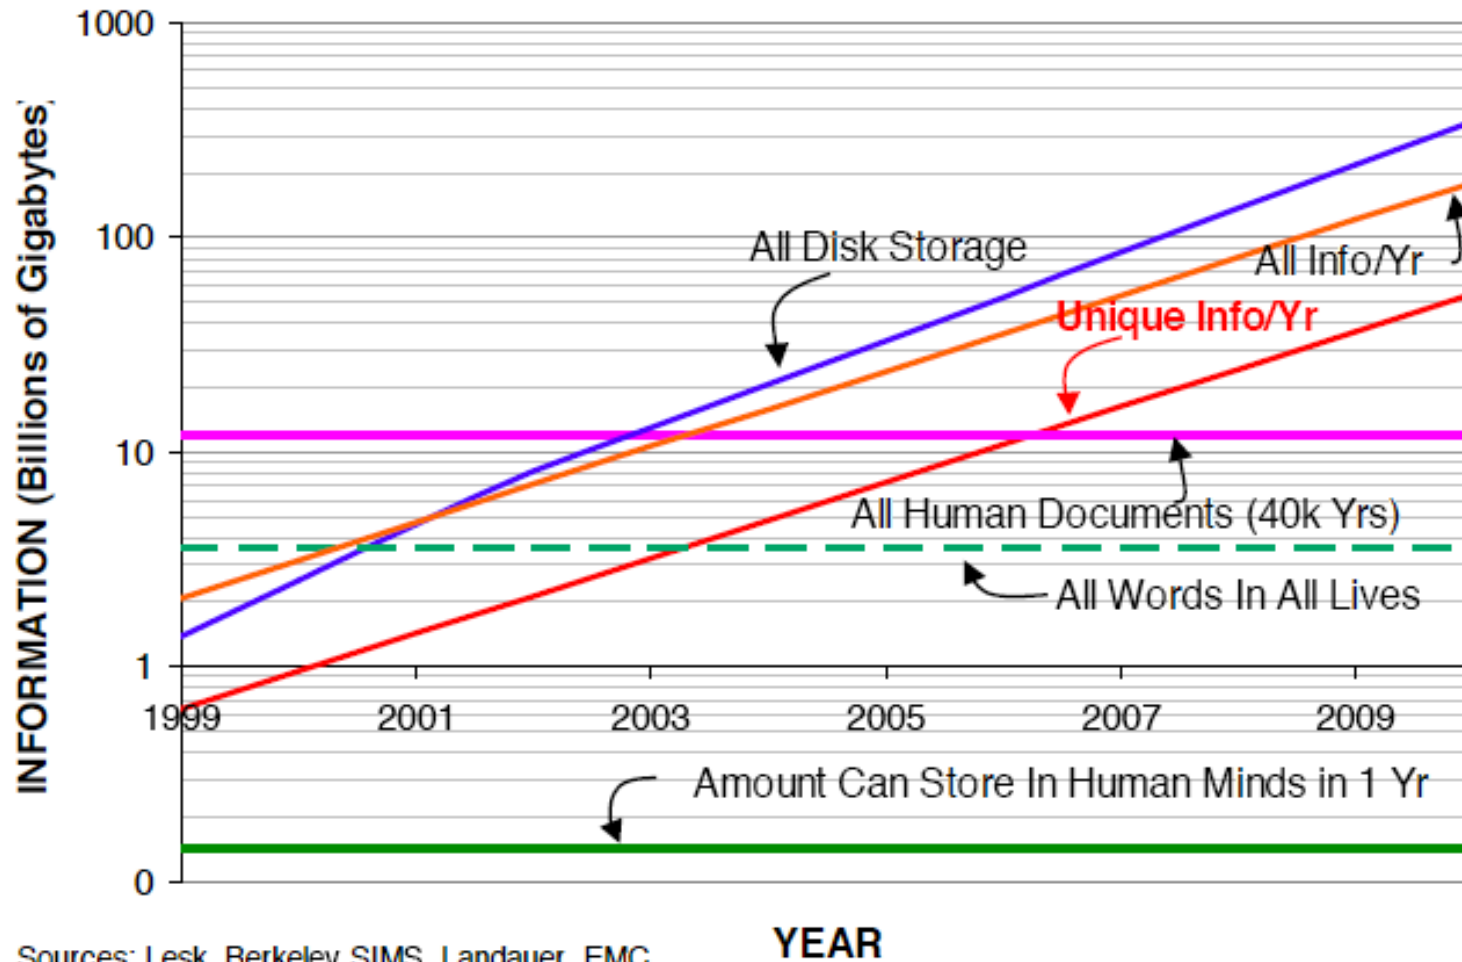


High-Performance Visualization

Thomas Fogal

Growth of data



VisIt / IceT

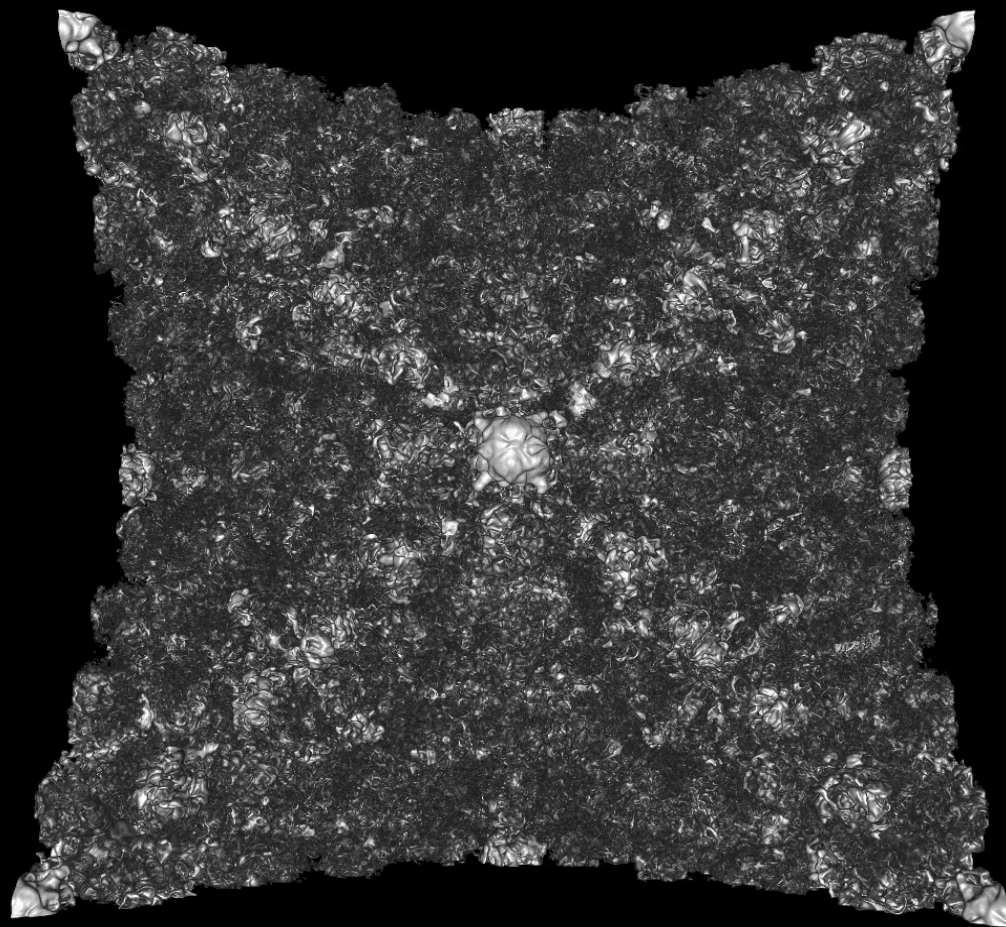
VisIt: An End-User Tool For Visualizing and Analyzing Very Large Data



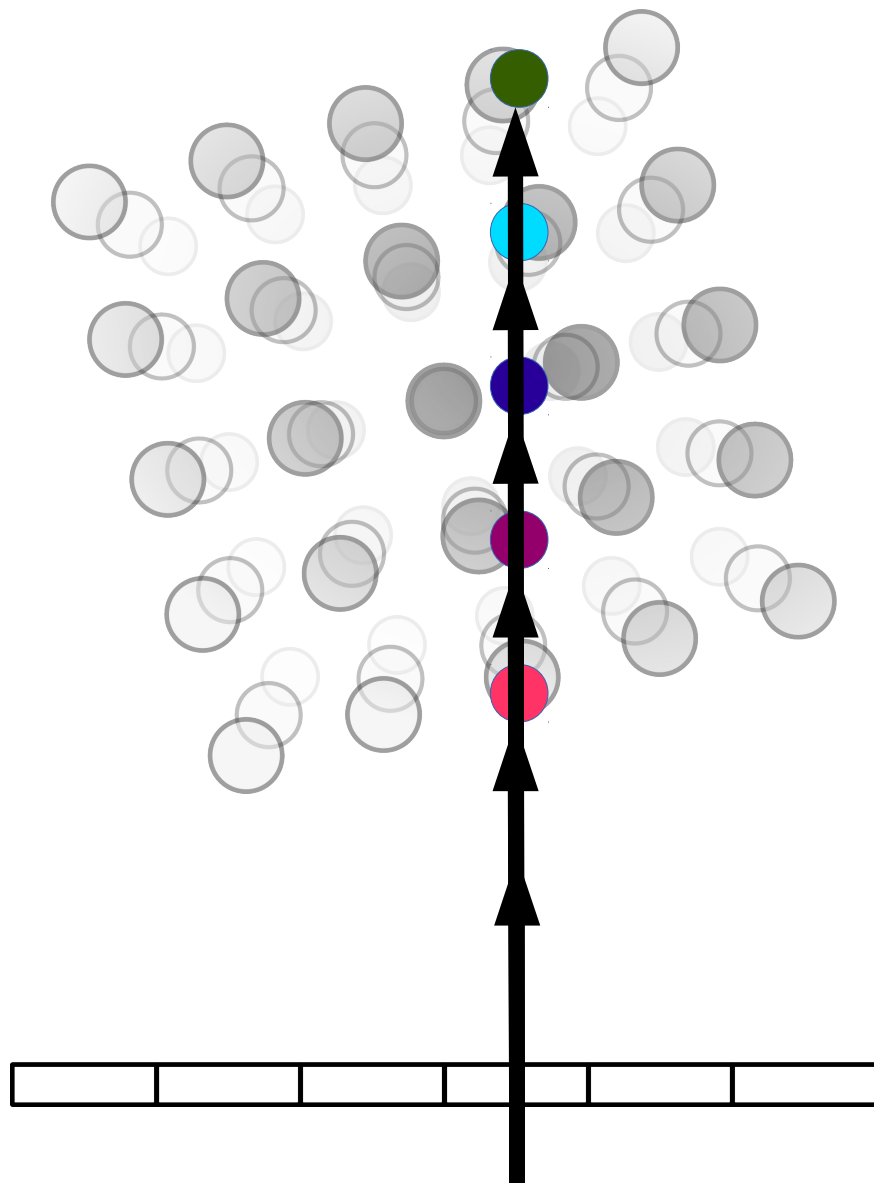
Hank Childs^{*,§}, Eric Brugger[†], Brad Whitlock[†], Jeremy Meredith[‡], Sean Ahern[‡], Kathleen Bonnell[†], Mark Miller[†], Gunther H. Weber^{*}, Cyrus Harrison[†], David Pugmire[‡], Thomas Fogal[¶], Christoph Garth[§], Allen Sanderson[¶], E. Wes Bethel^{*}, Marc Durant[∇], David Camp^{*,§}, Jean M. Favre^{||}, Oliver Rübel^{*}, Paul Navrátil[△], Matthew Wheeler^α, Paul Selby^α, and Fabien Vivodtzev[±]

^{*} Lawrence Berkeley National Laboratory, [†] Lawrence Livermore National Laboratory, [‡] Oak Ridge National Laboratory, [§] University of California at Davis, [¶] University of Utah, [∇] Tech-X Corporation, ^{||} Swiss National Supercomputing Center, [△] Texas Advanced Computing Center, ^α Atomic Weapons Establishment, [±] French Atomic Energy Commission, CEA/CESTA

Abstract. VisIt is a popular open source tool for visualizing and analyzing data. It owes its success to its foci of increasing data understanding, large data support, and providing a robust and usable product, as well as its underlying design that fits today's supercomputing landscape. In this short paper, we describe the VisIt project and its accomplishments.

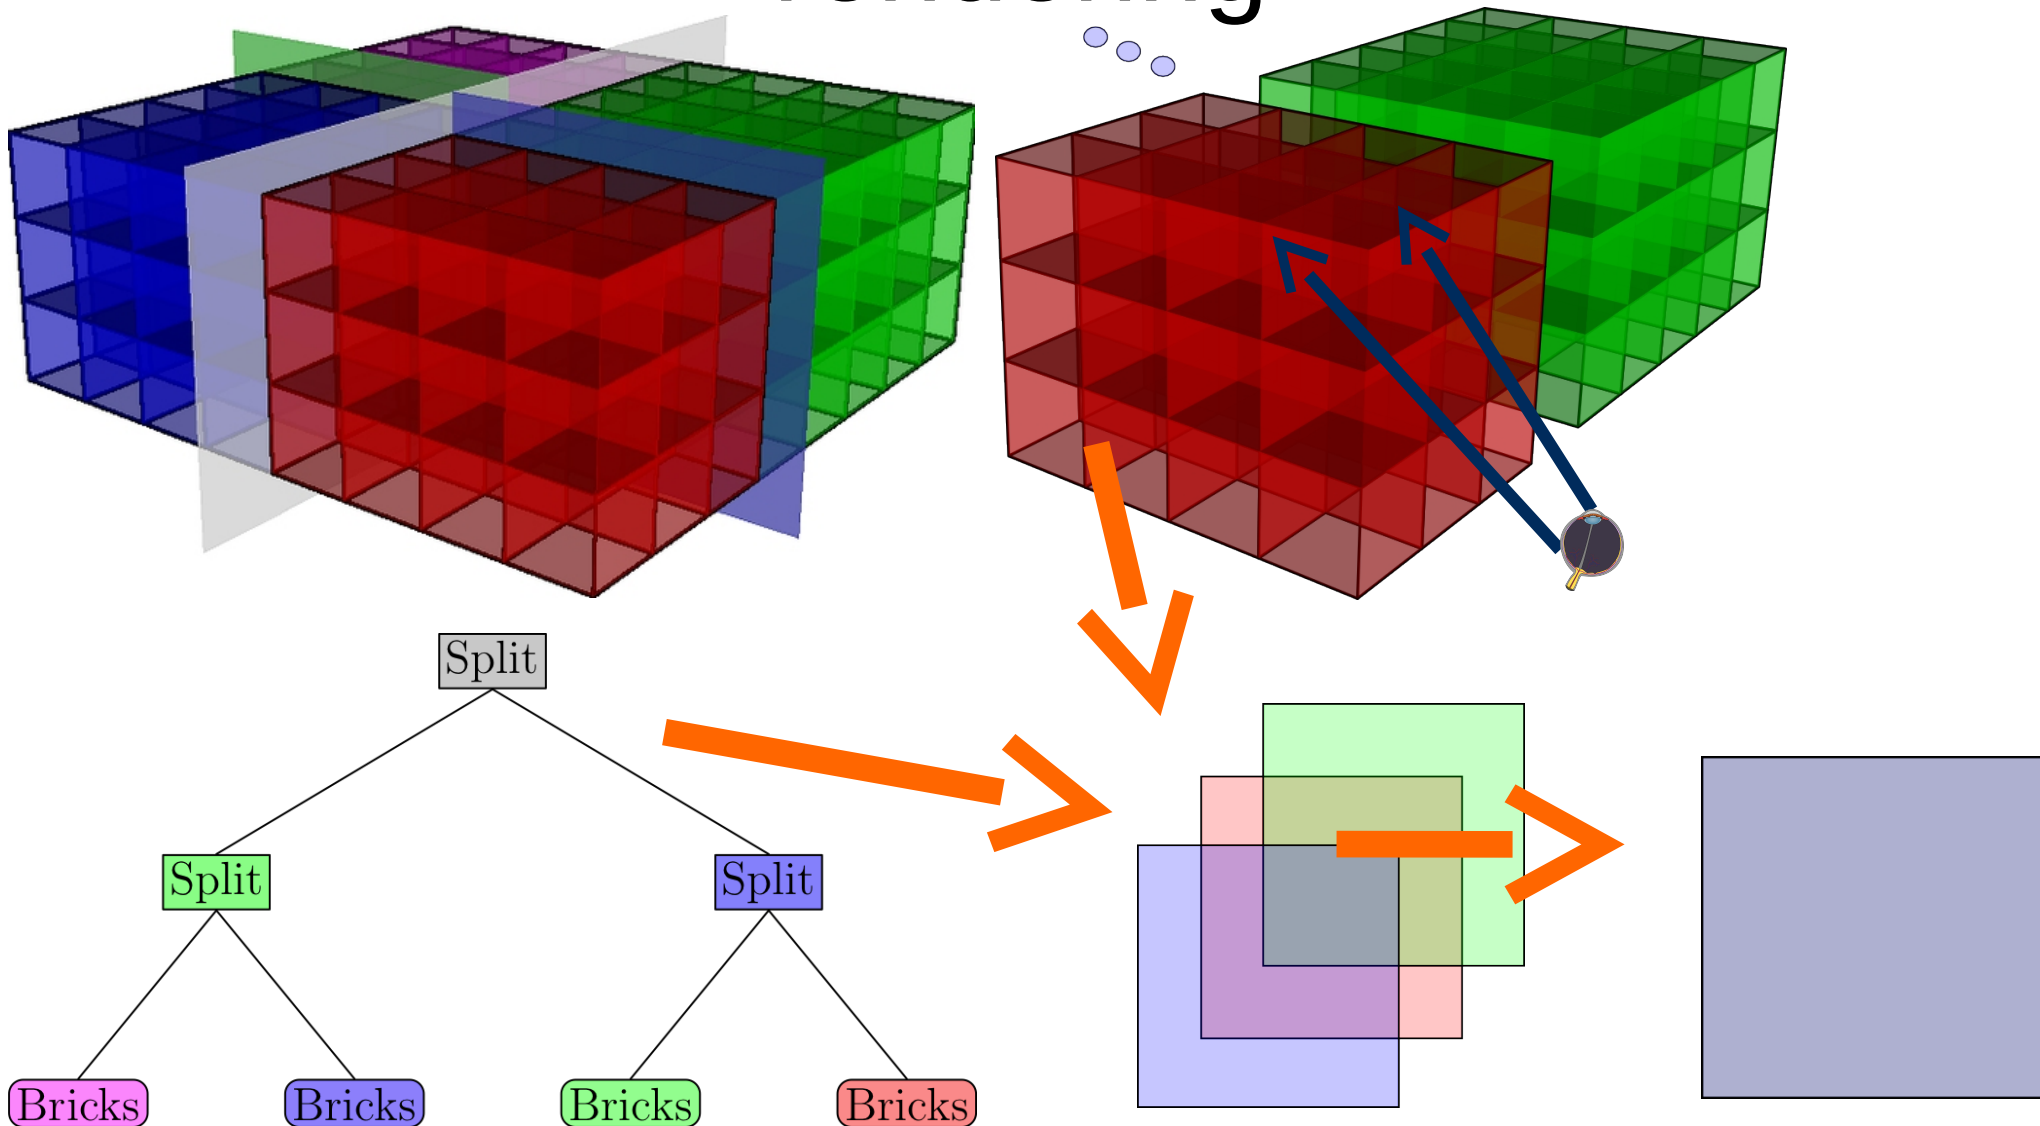


VolRen Background



● + ● + ● + ● + ●

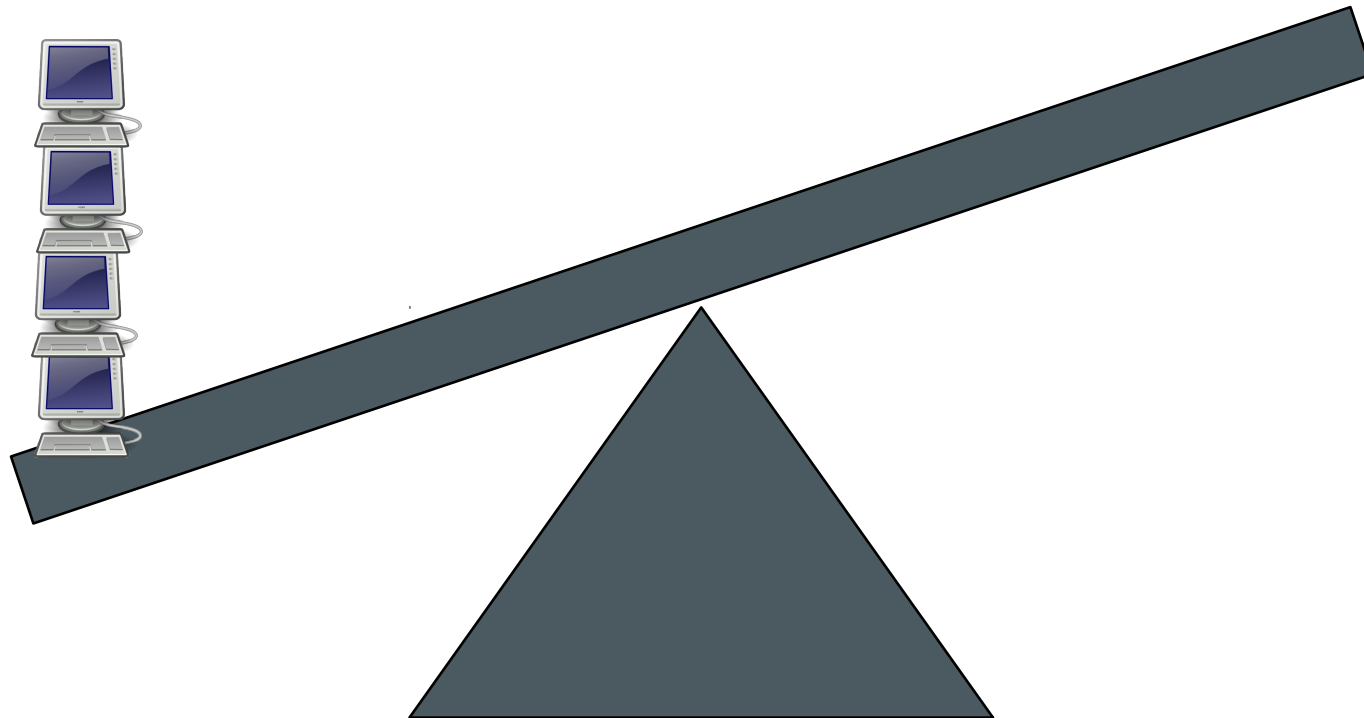
Traditional large-scale volume rendering



Compositing scalability

Rendering

Compositing

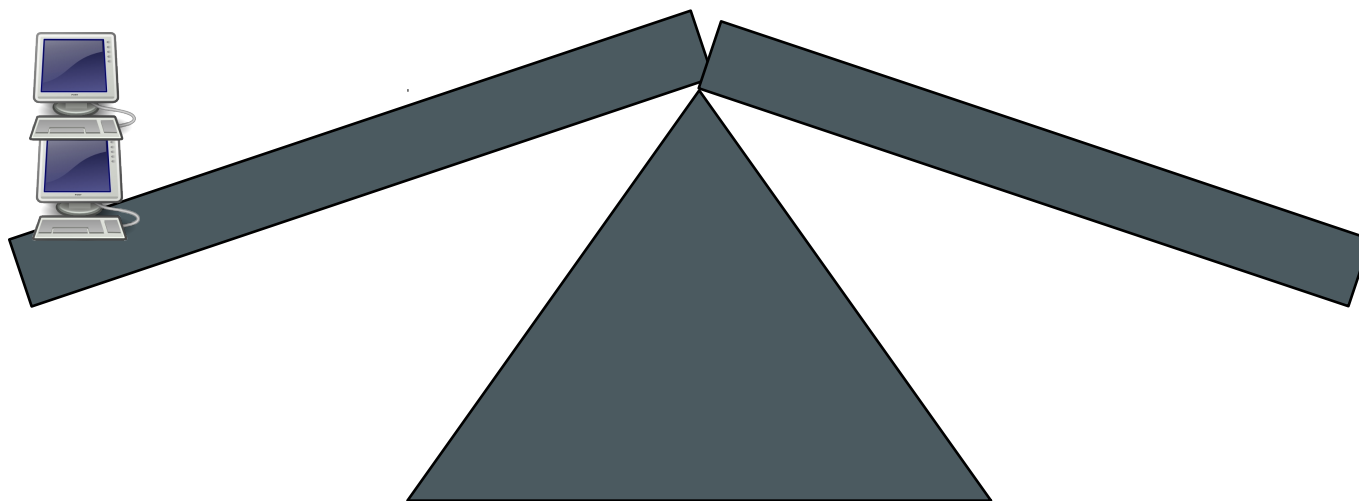




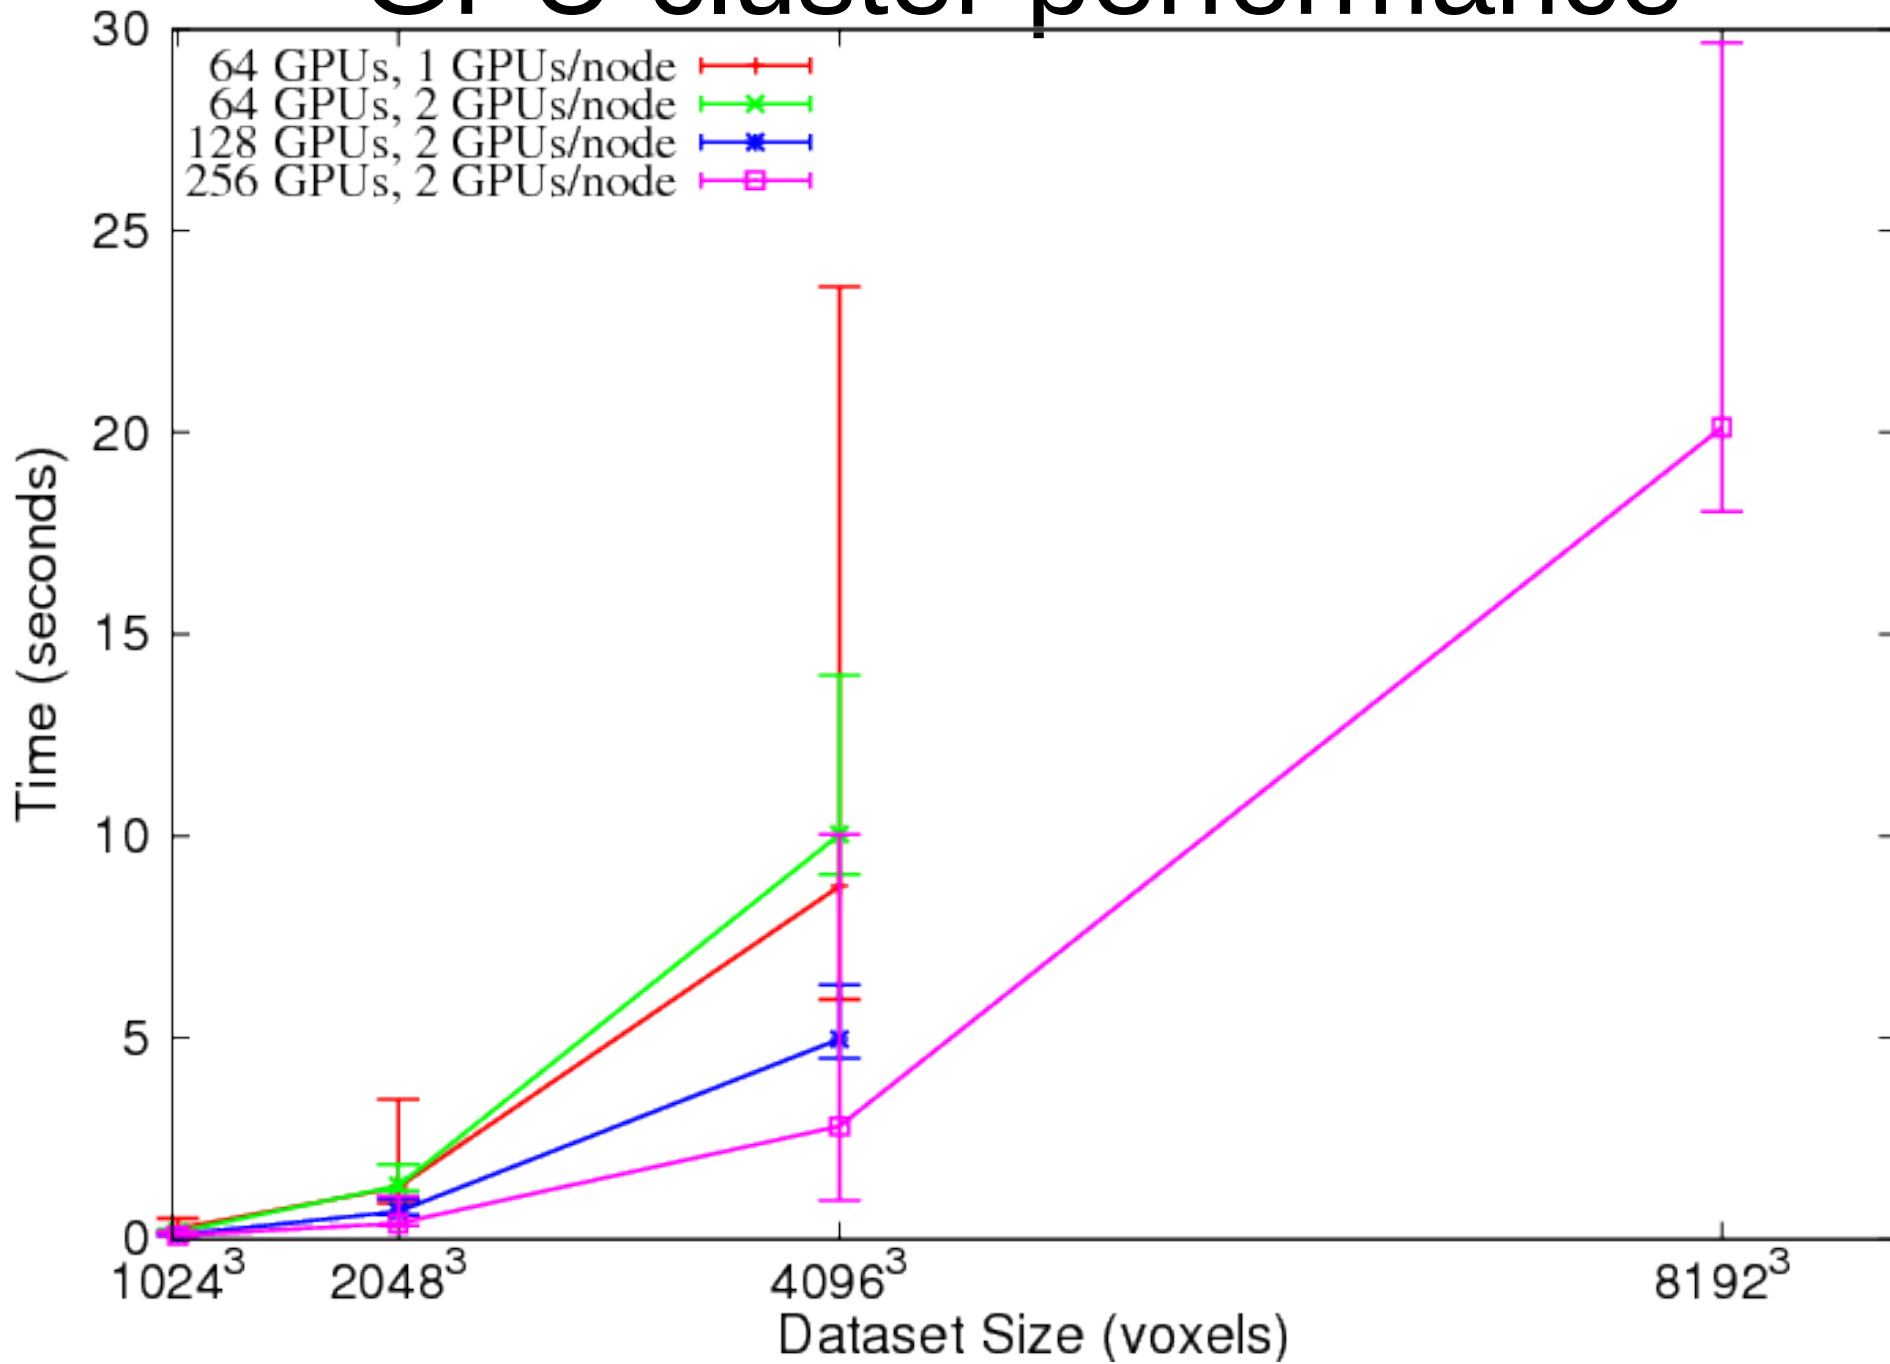
Compositing scalability

Rendering

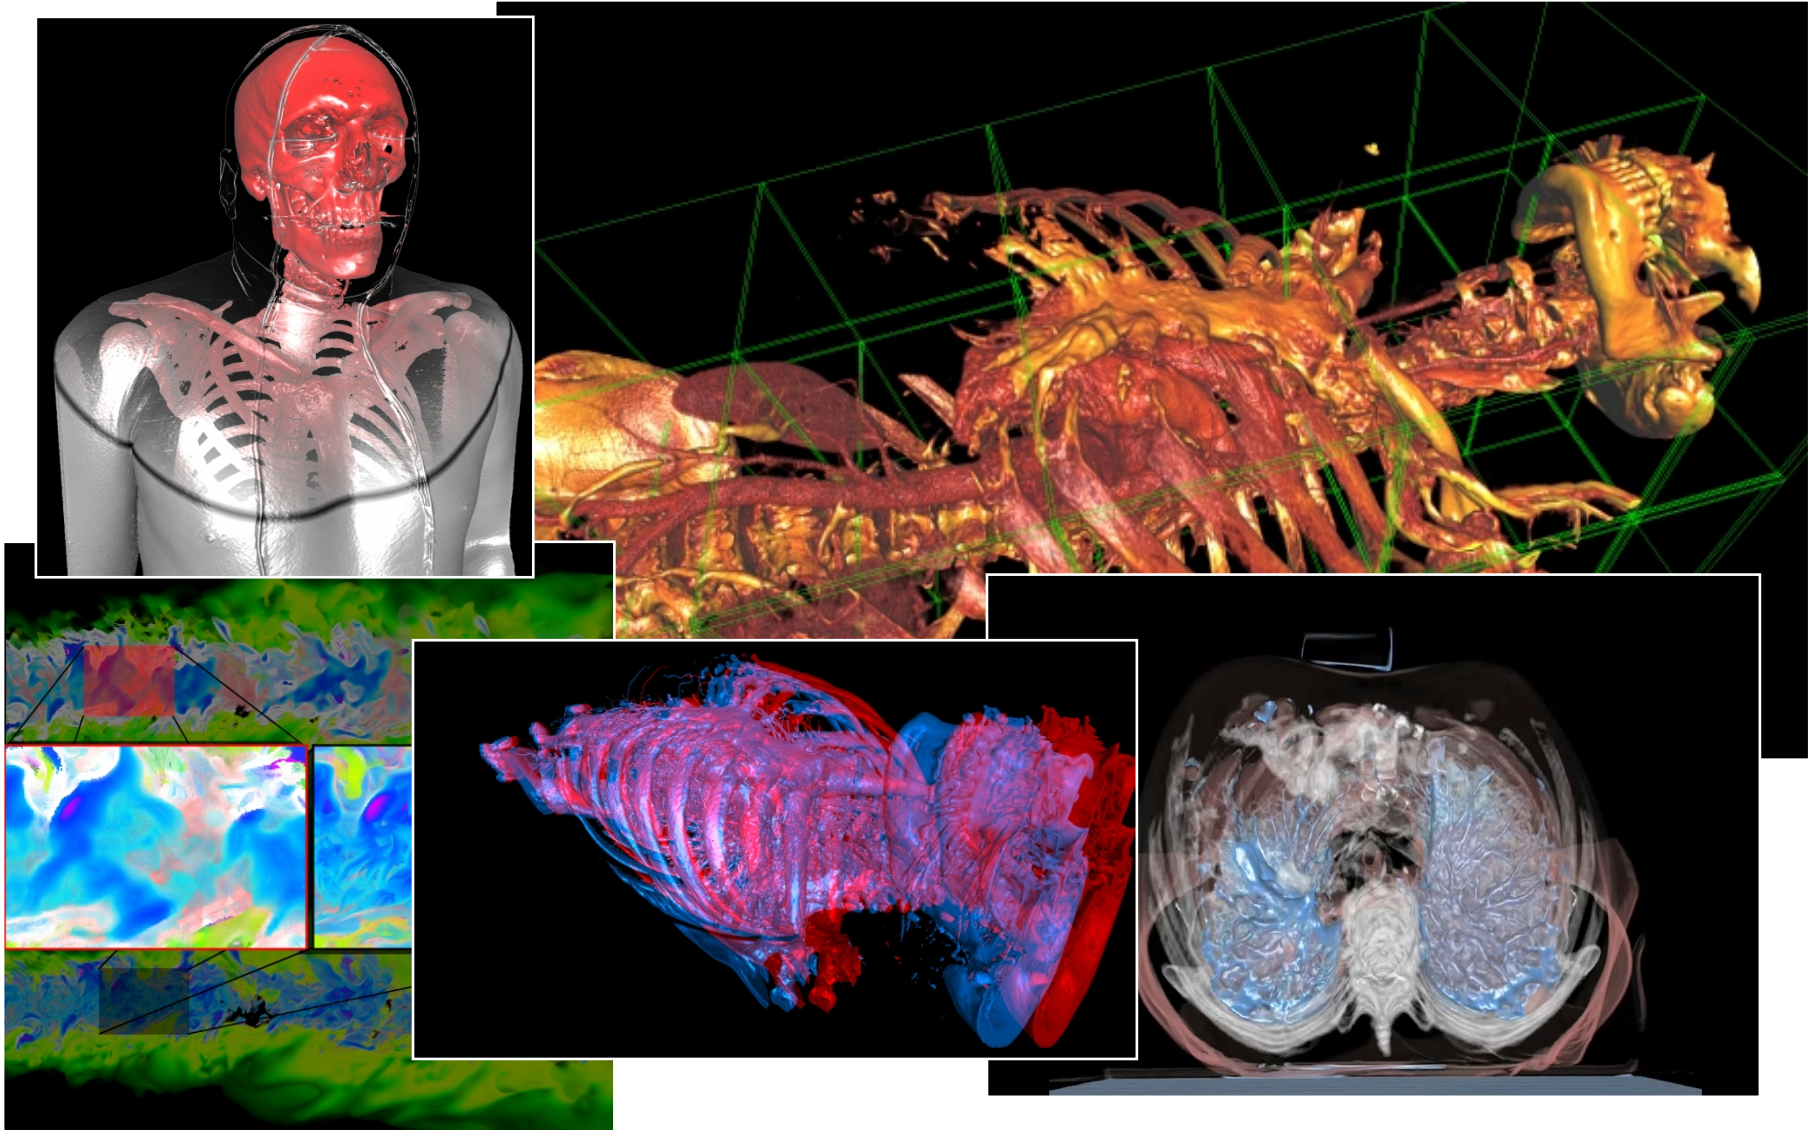
Compositing



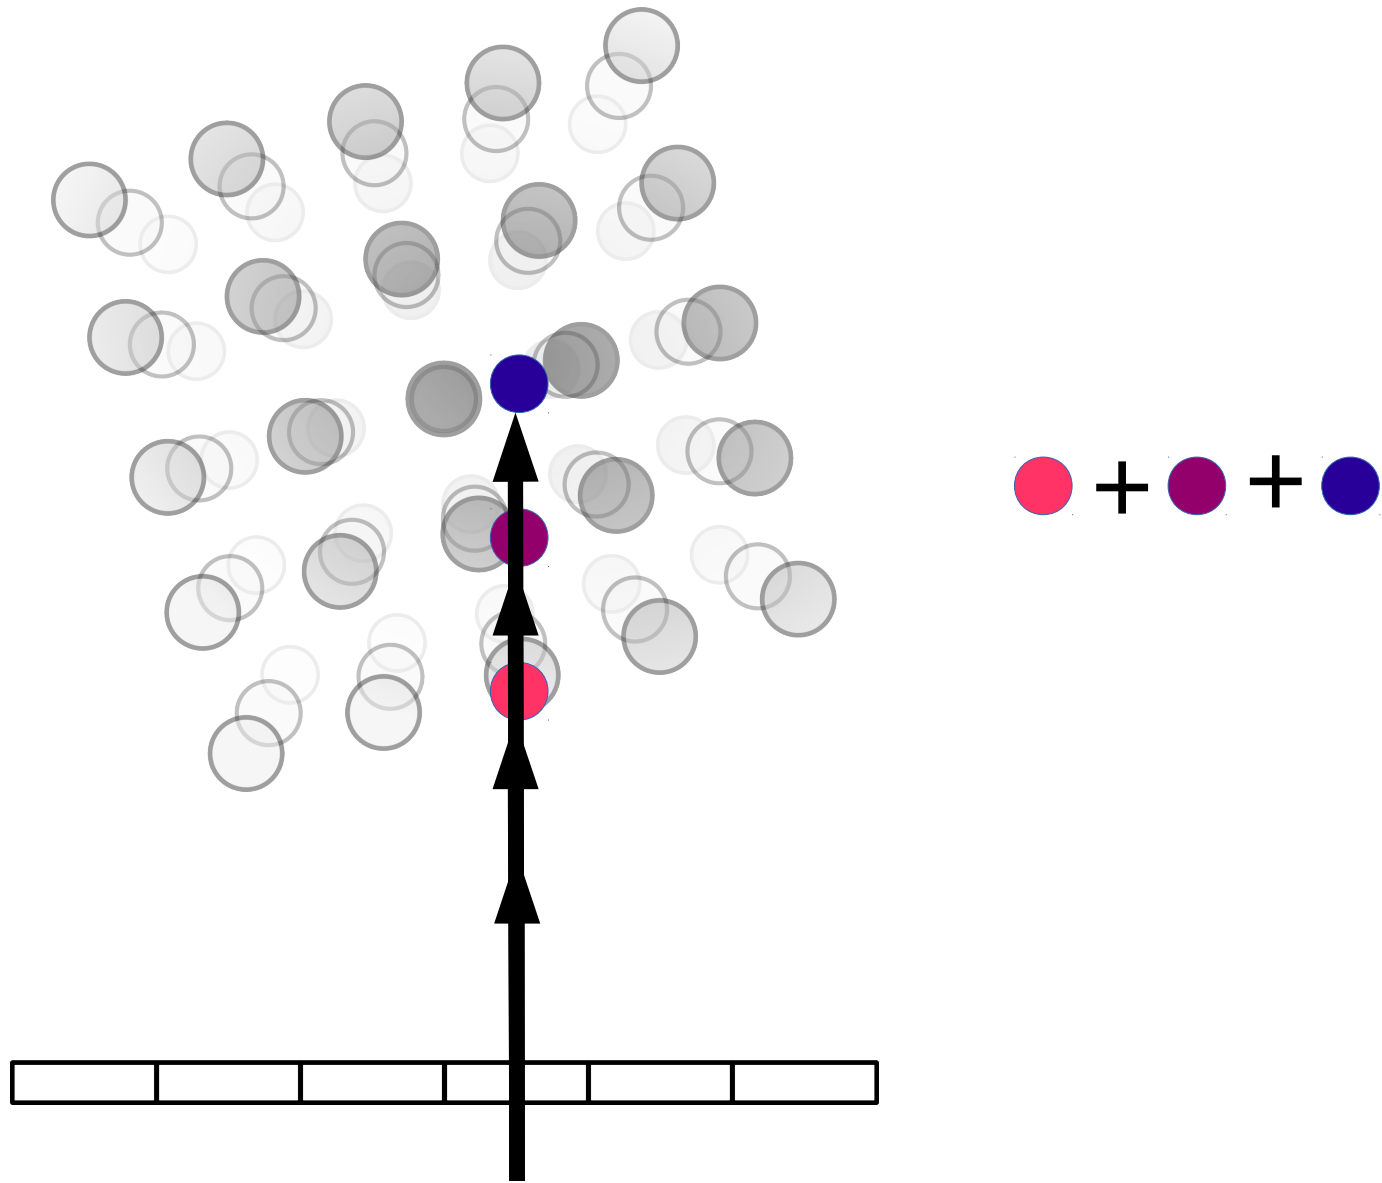
GPU cluster performance



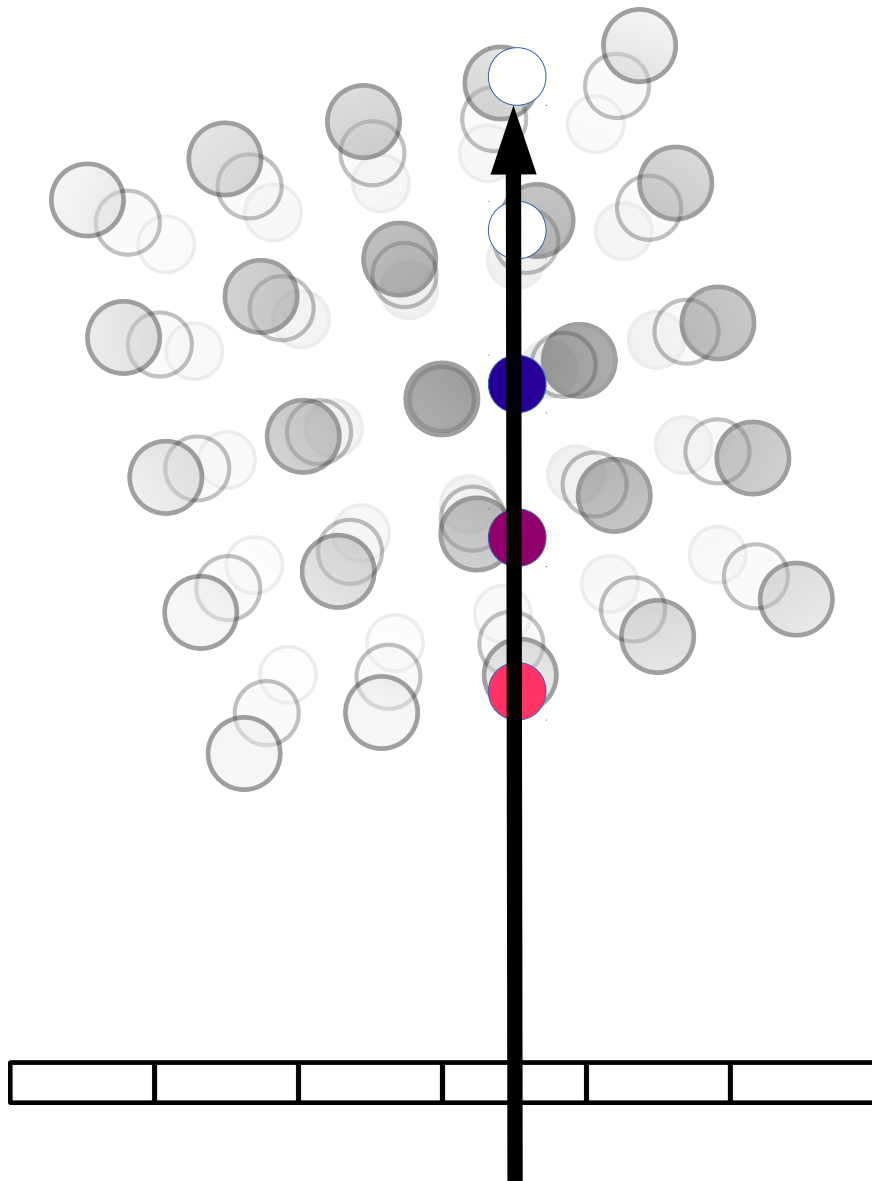
Single node volume visualization



Ray saturation

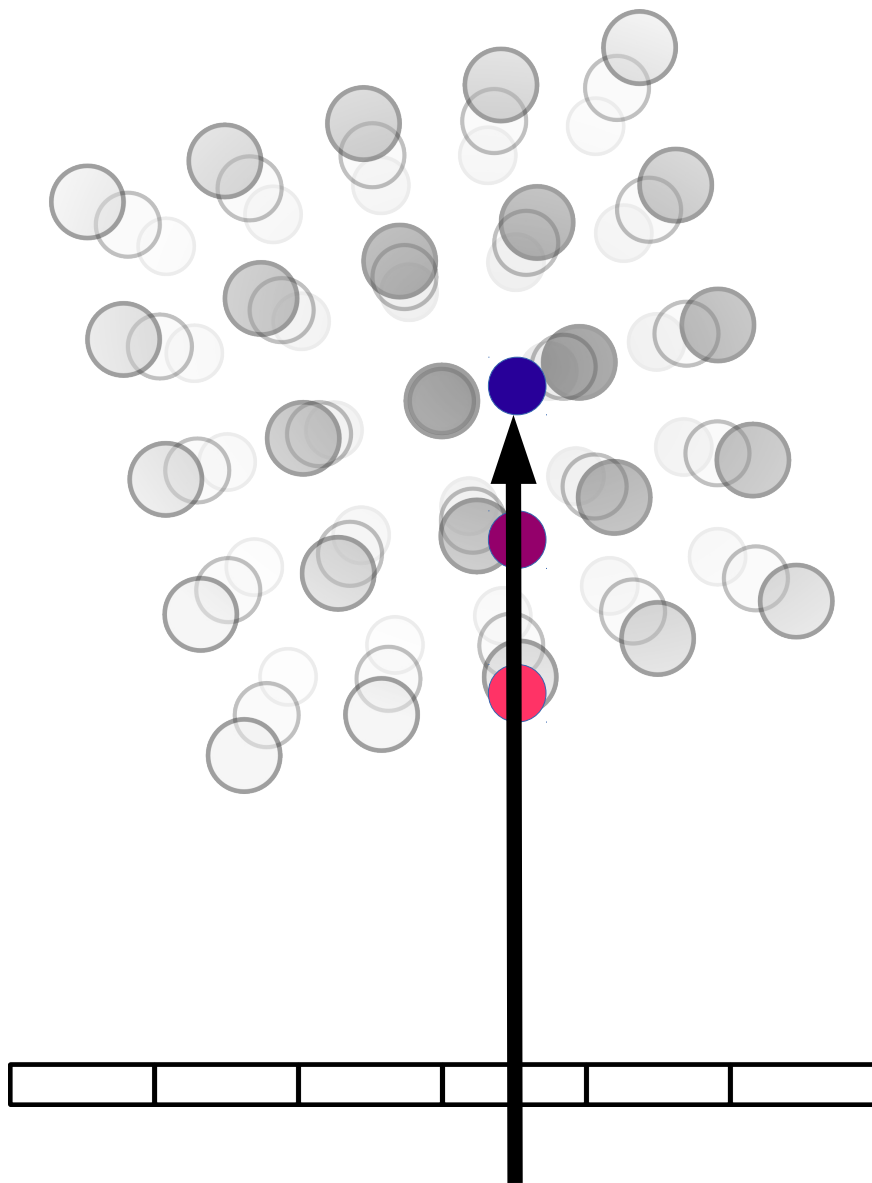


Full Alpha



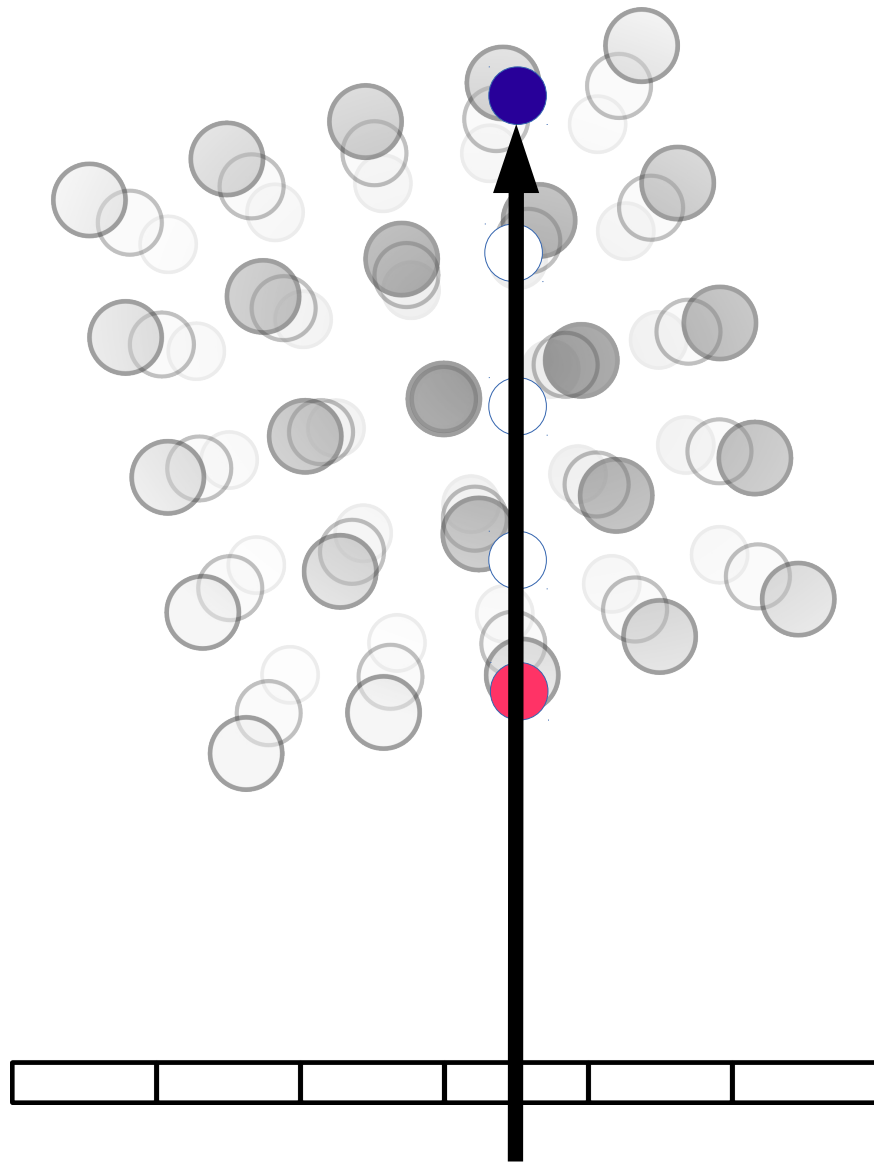
$$\text{red circle} + \text{purple circle} + \text{blue circle} + \text{white circle} + \text{white circle}$$

Early Ray Termination



$$\text{●} + \text{●} + \text{●} = \text{VRI}$$

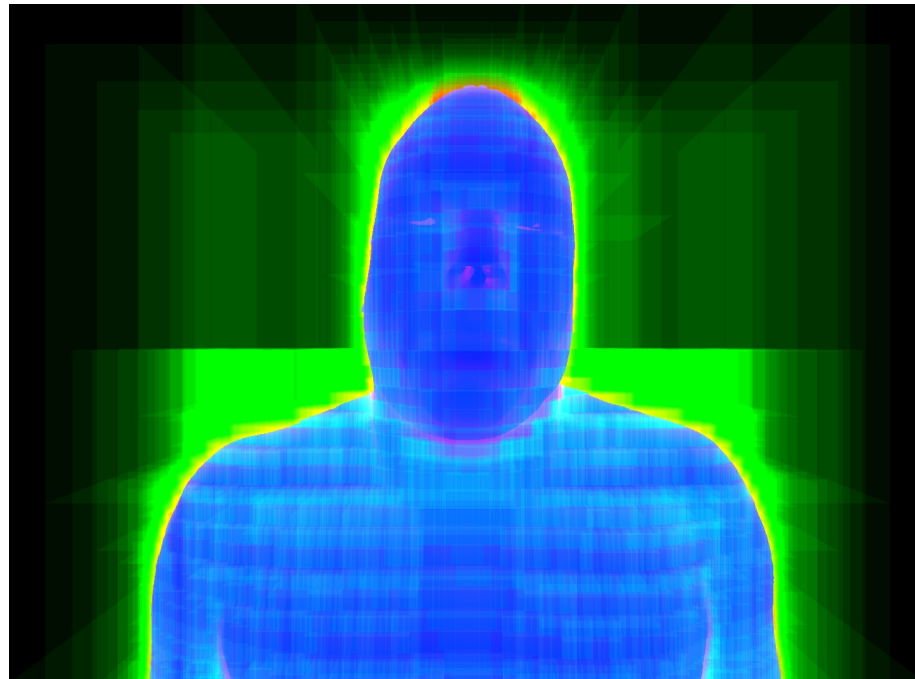
Empty Space Leaping



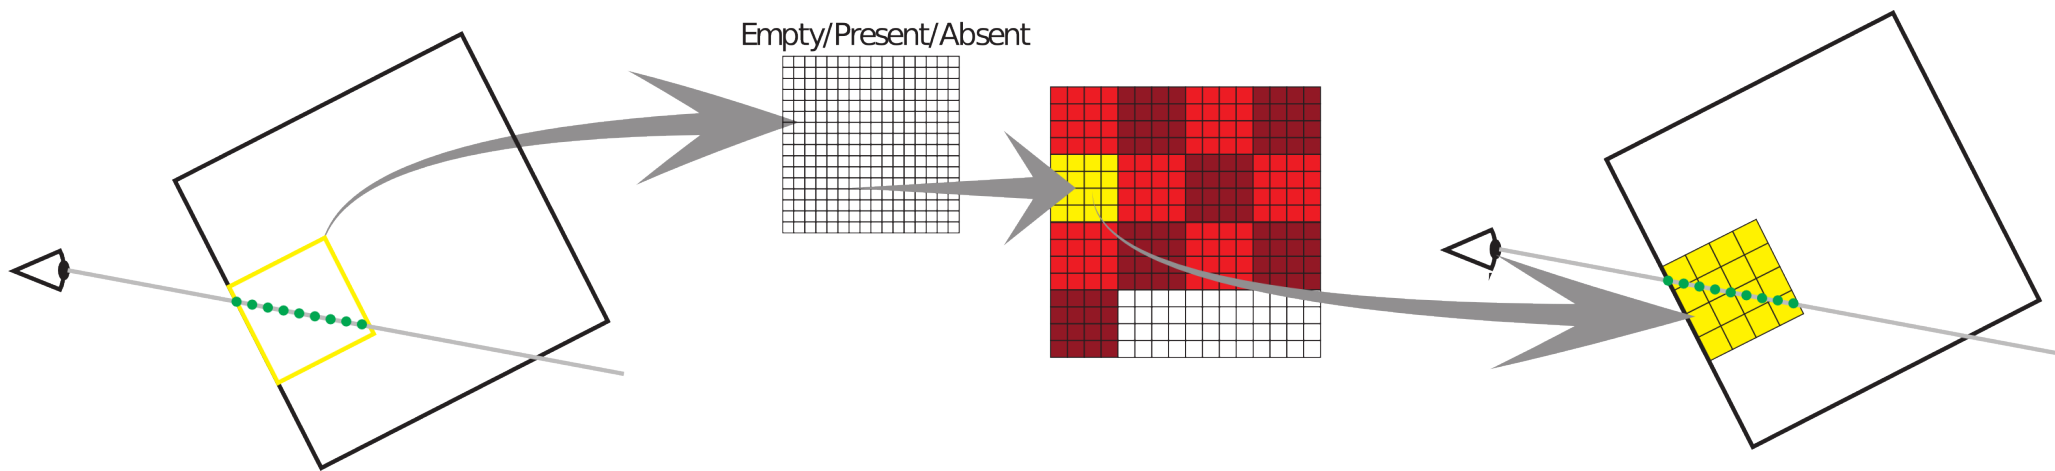
● + ● = VRI

What's important for performance

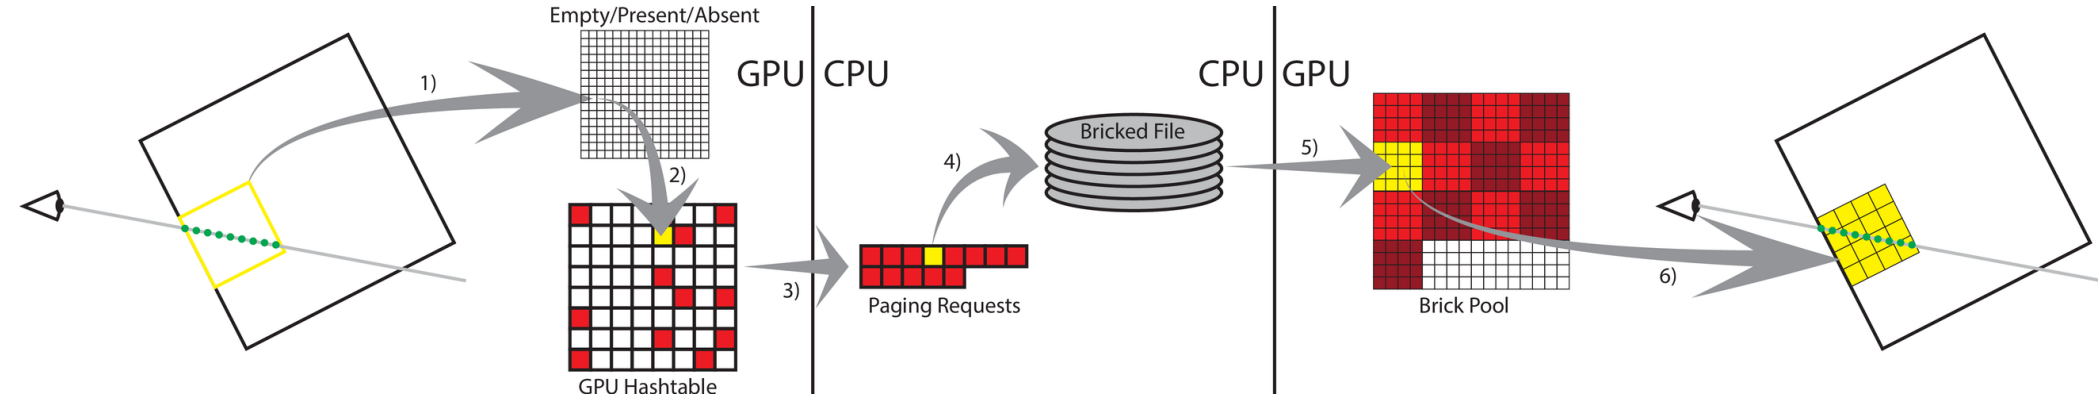
- Identifying densely-sampled regions
- Transition to coarse sampling quickly
- Communicate data needed to IO



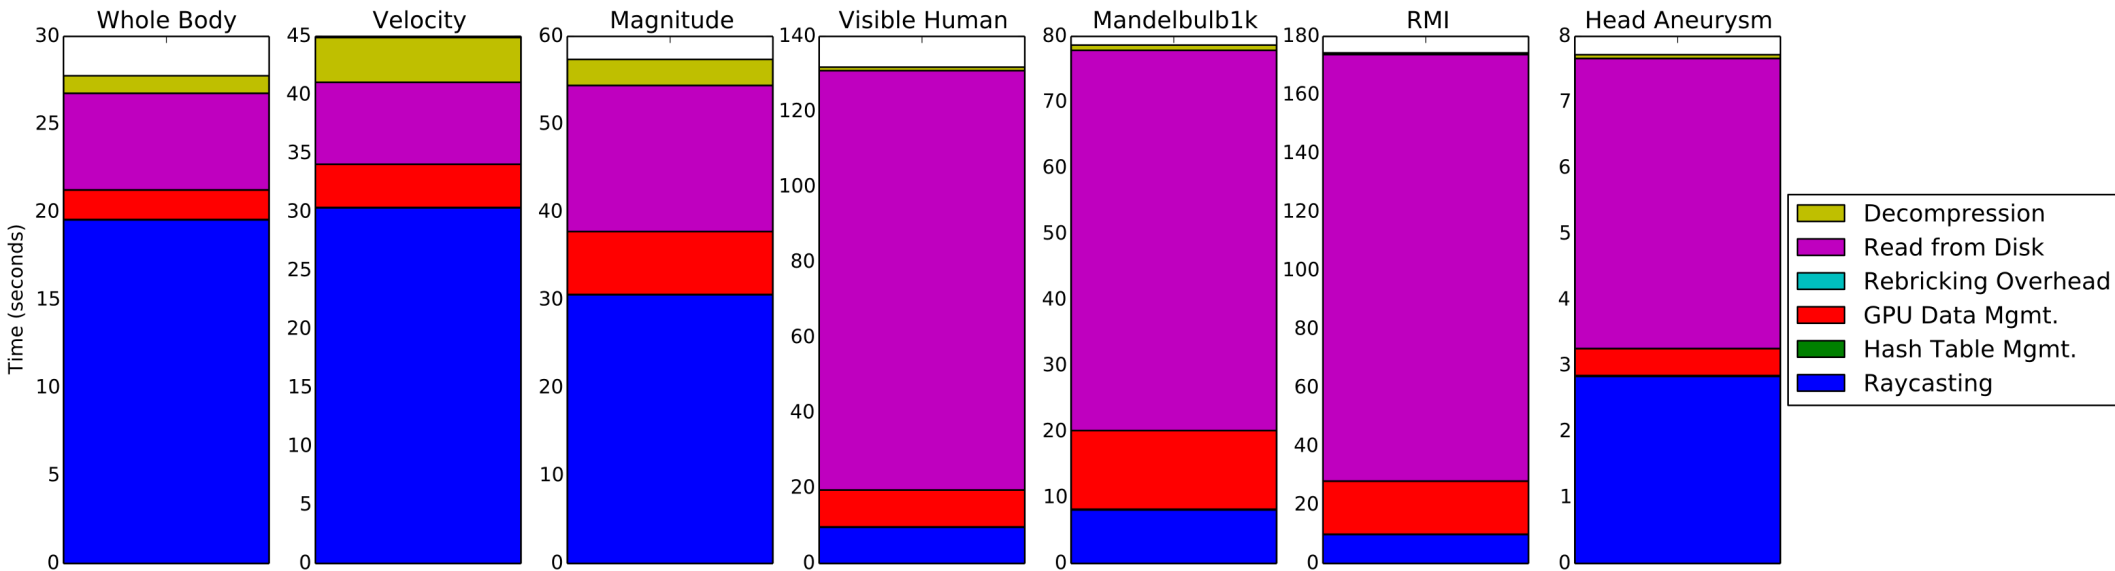
Ray-Guided Rendering



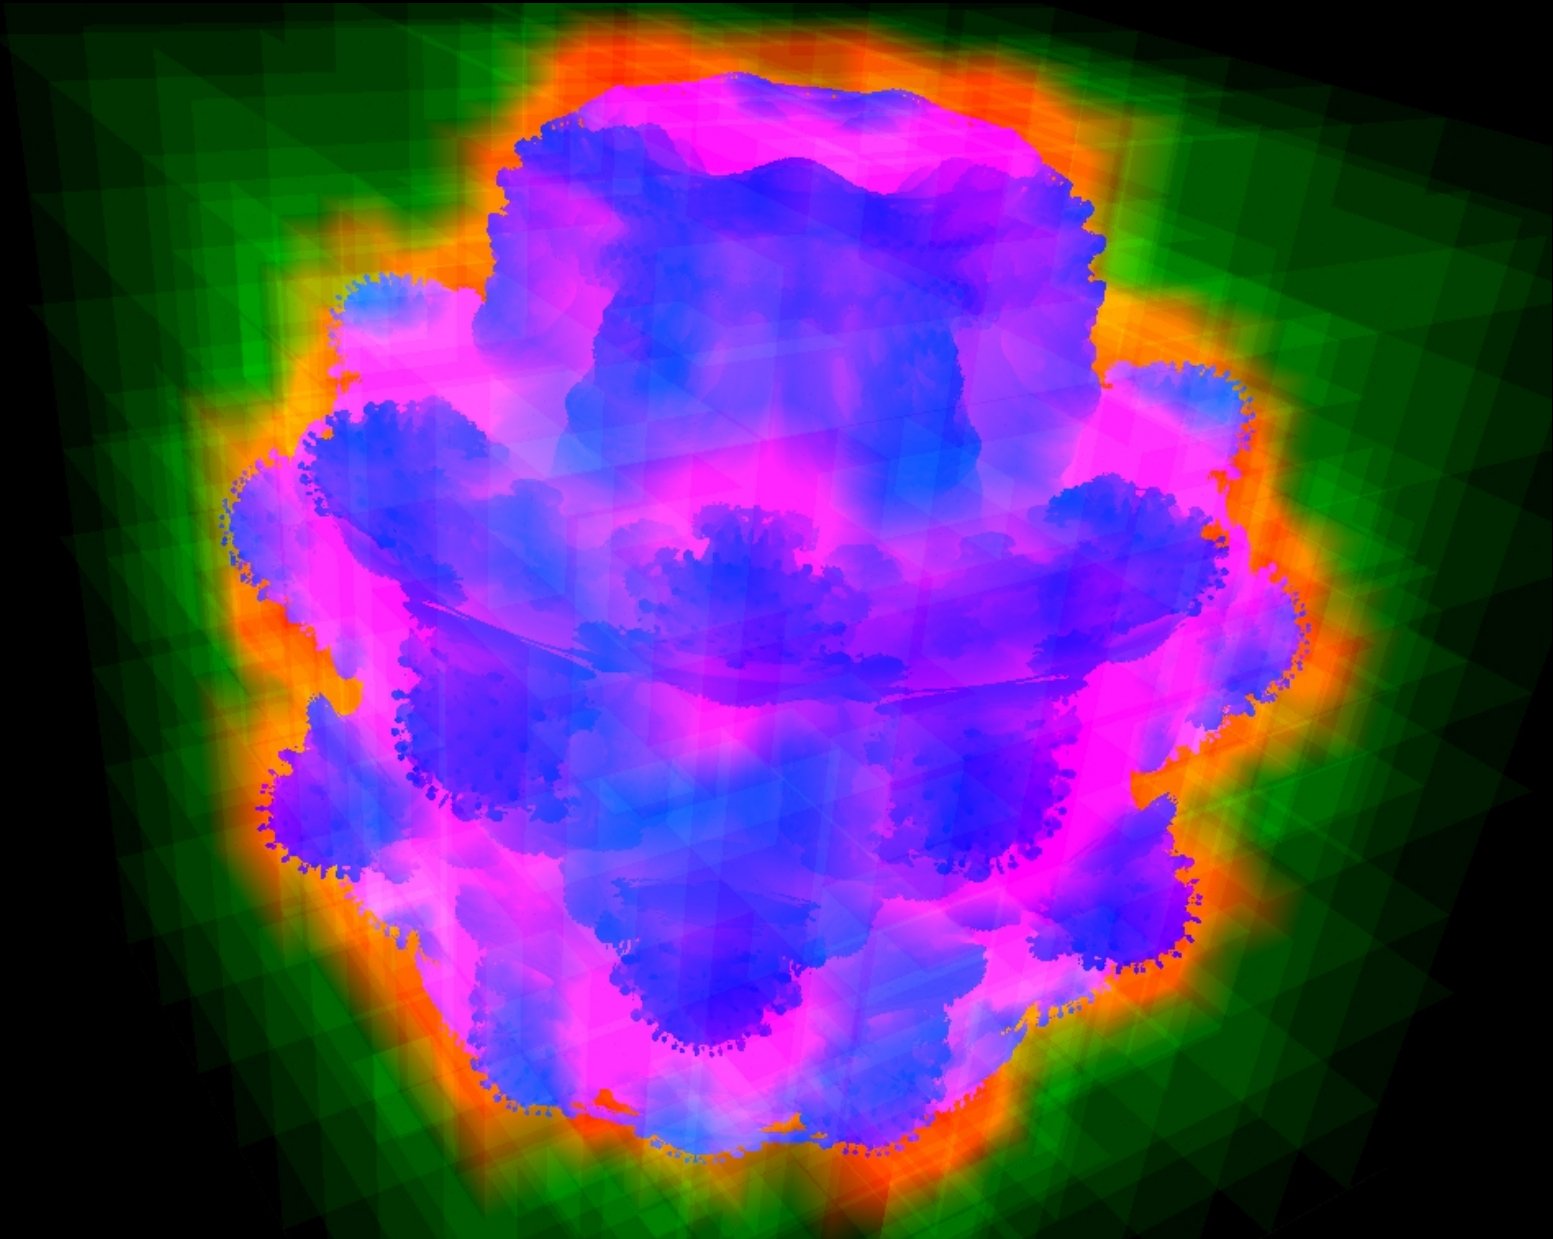
Ray-Guided Rendering



Where does the time go?

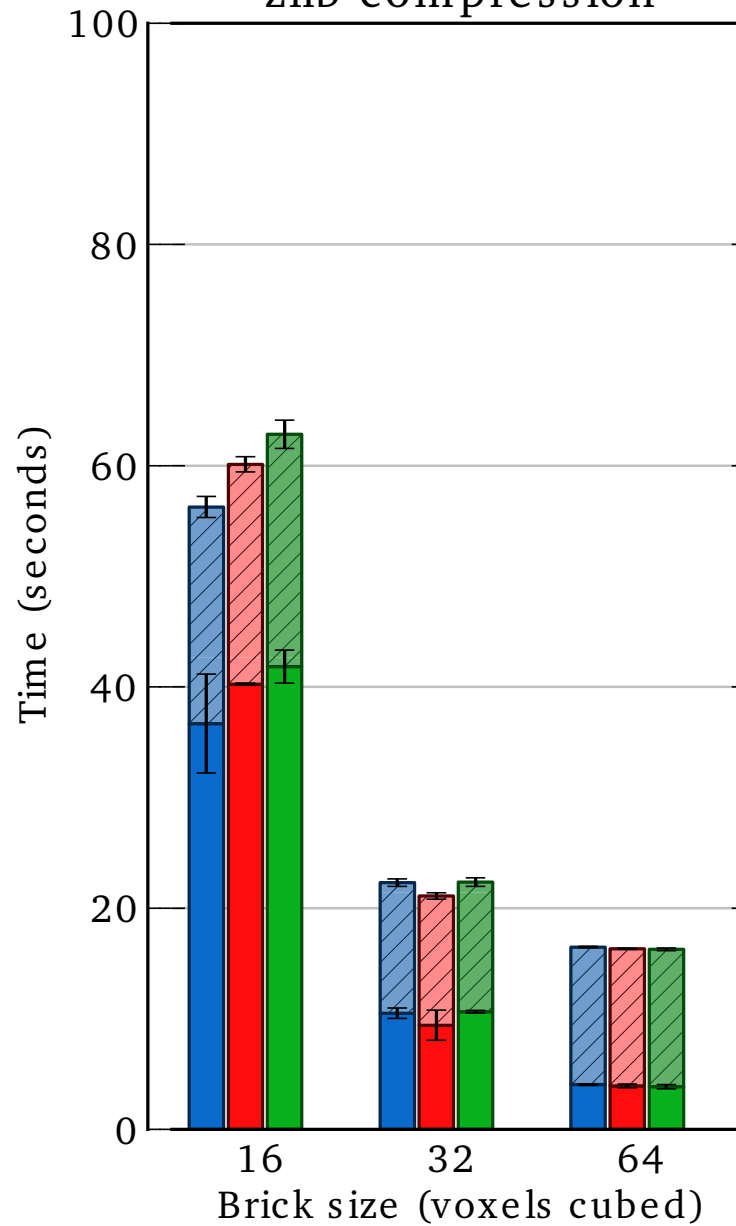


Brick Size



Brick Size: IO

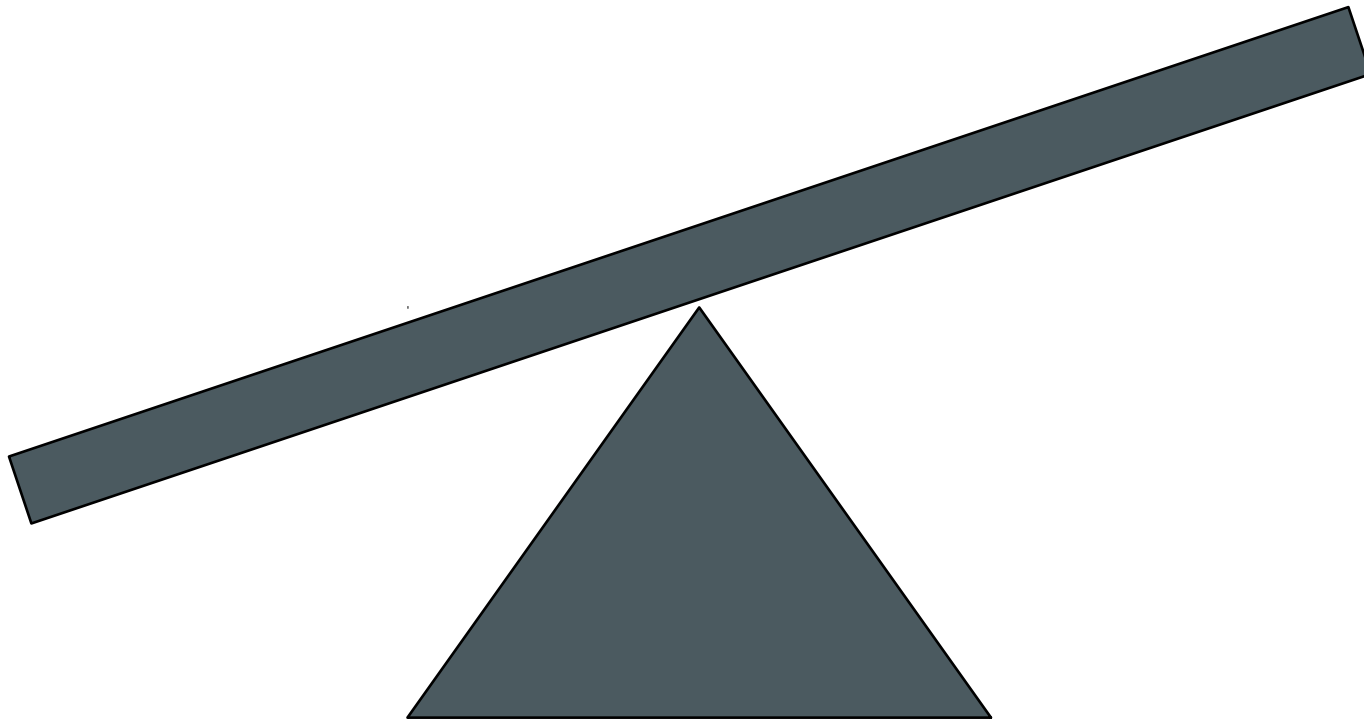
zlib compression



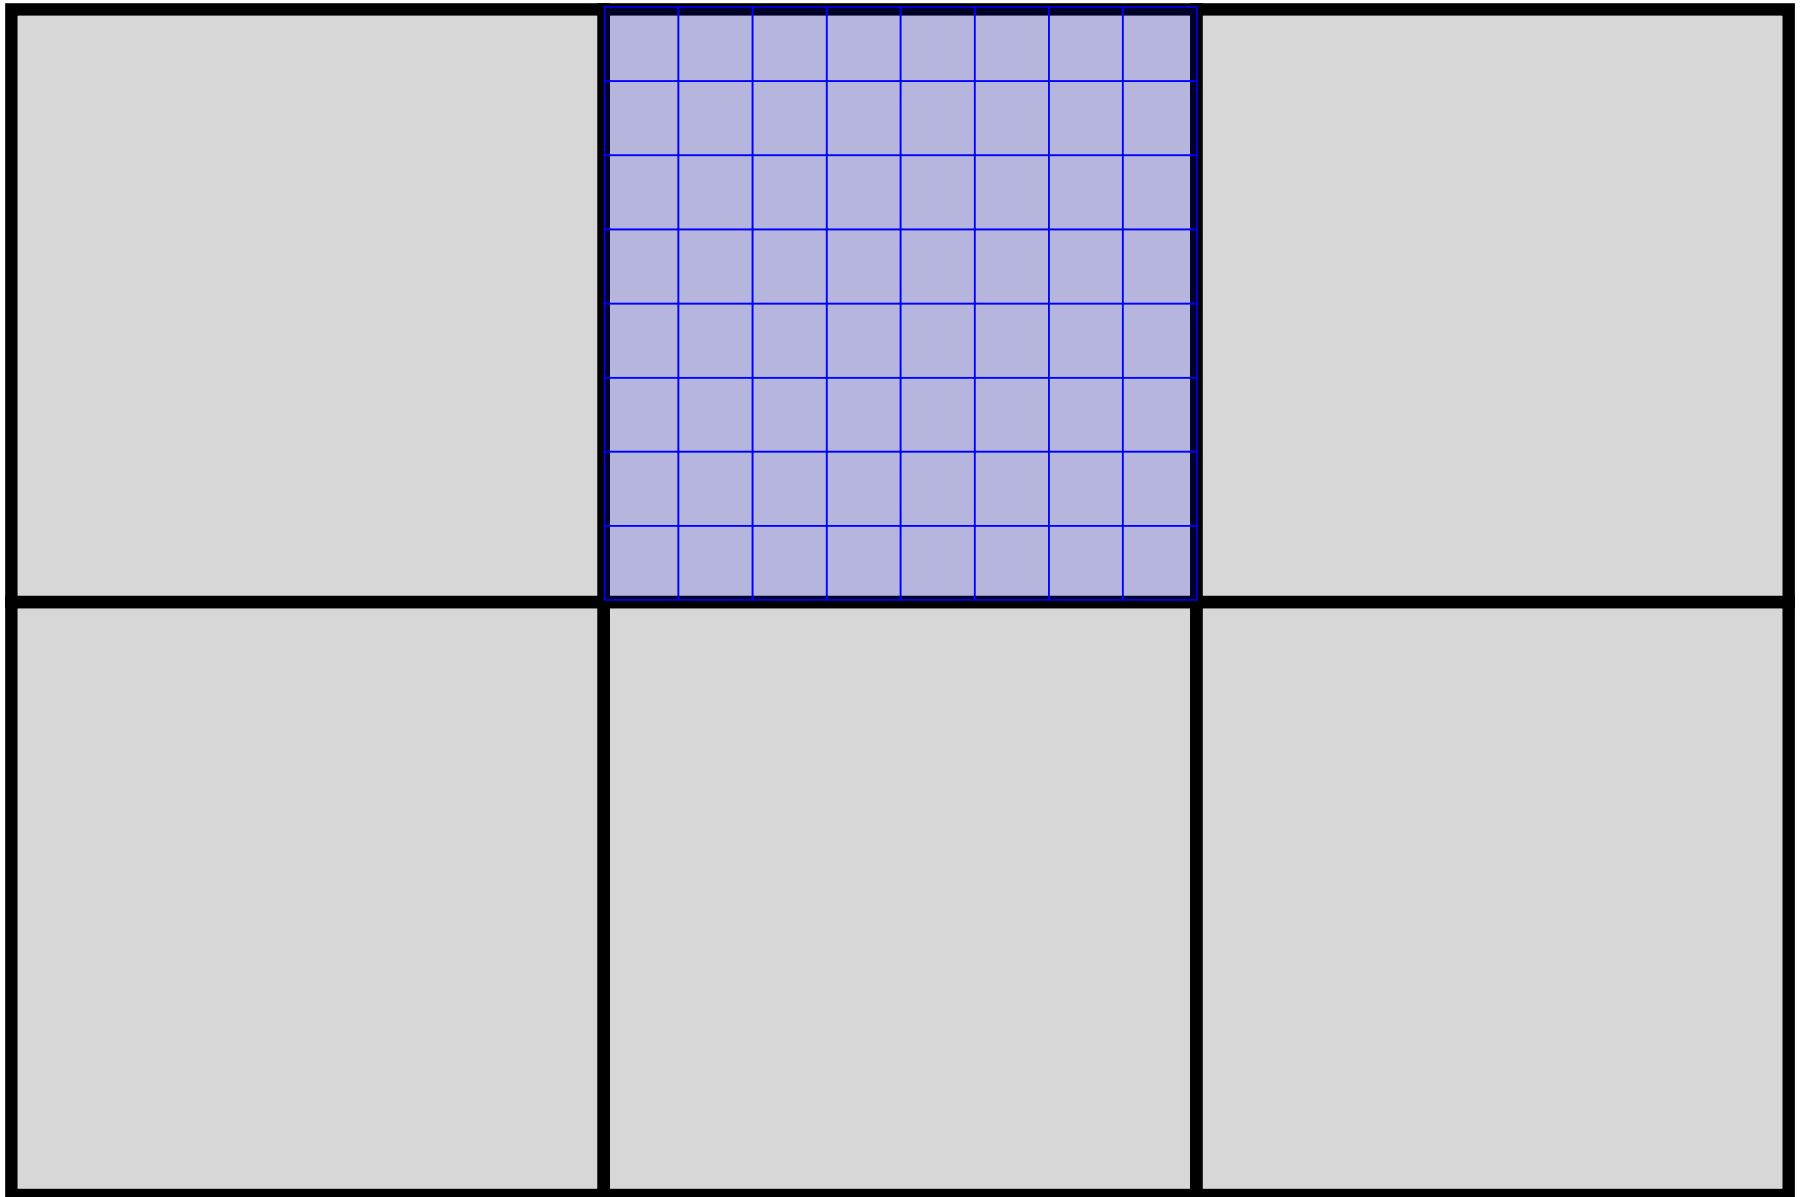
Brick size balancing

Render
performance

IO
performance



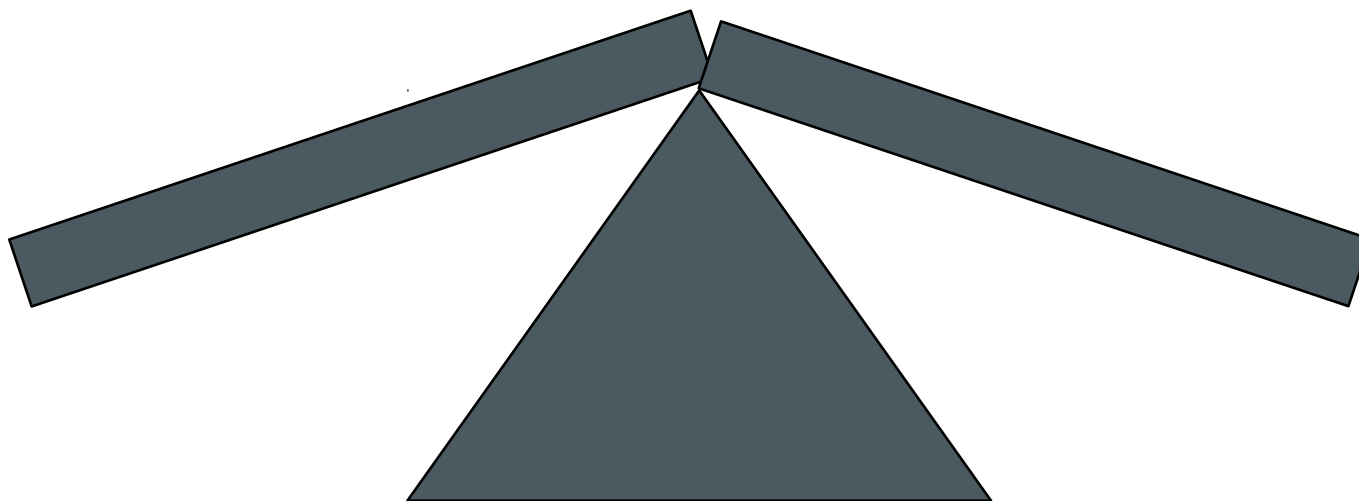
Dynamic Bricking



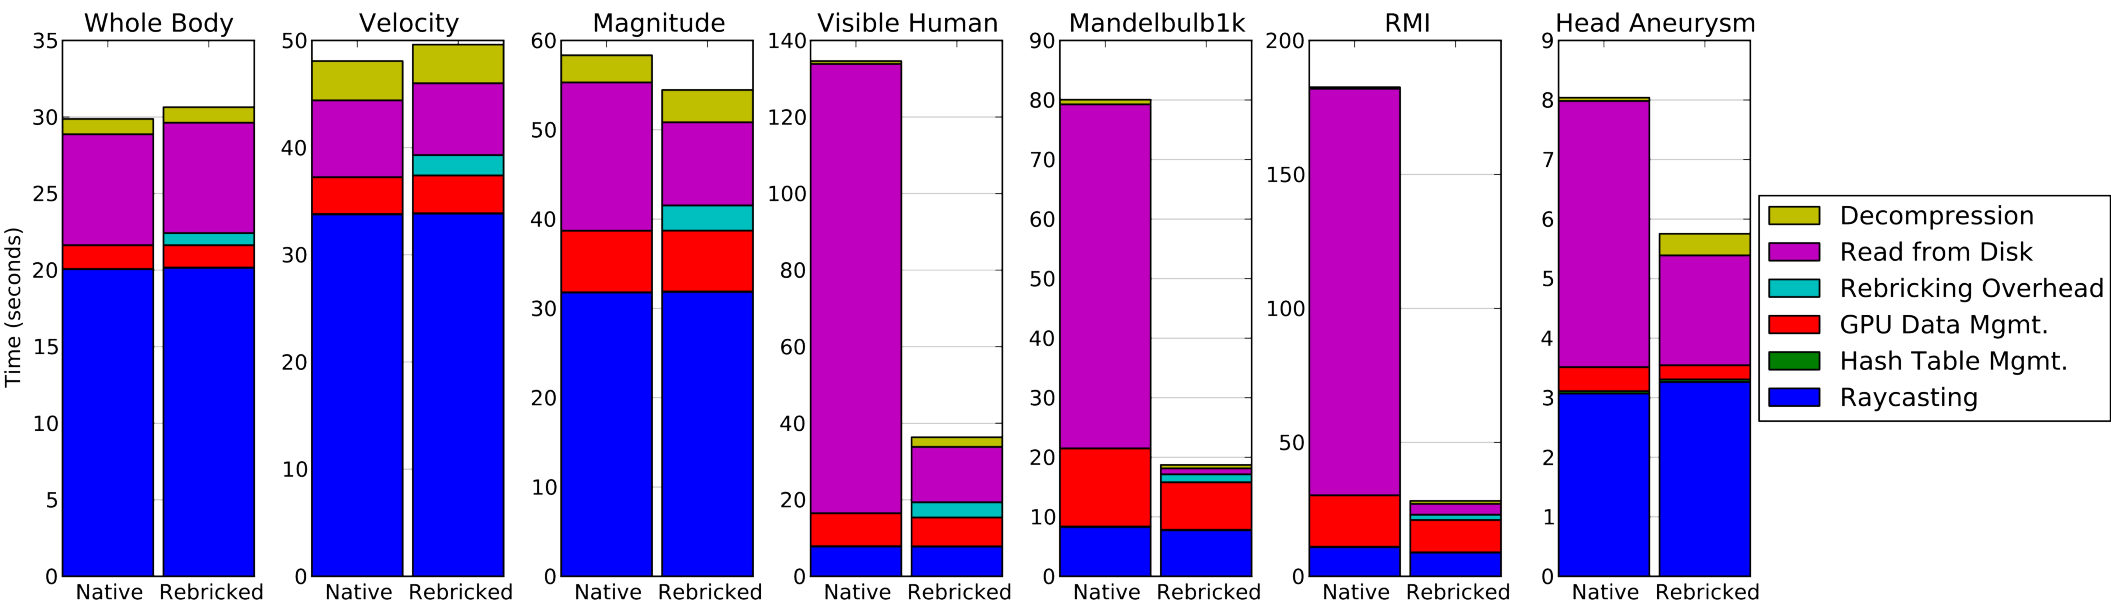
Dueling subsystems

Rendering

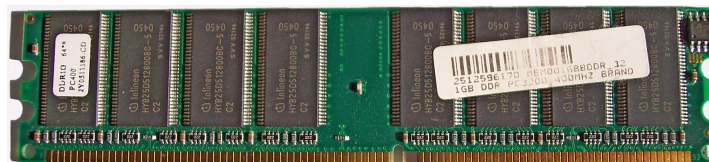
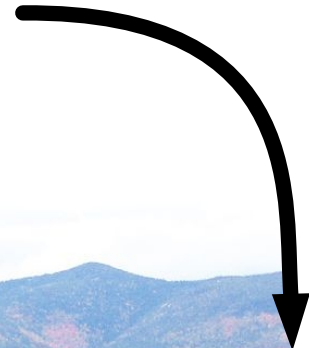
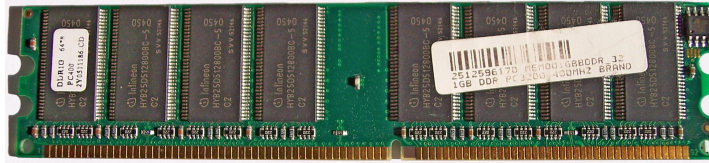
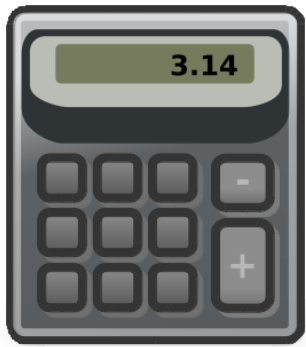
IO

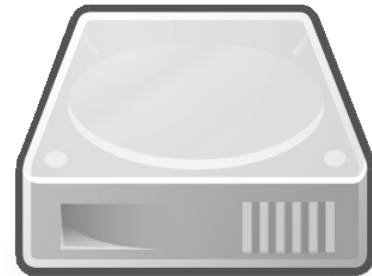
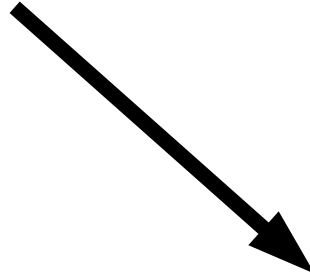
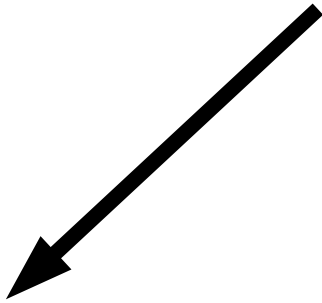
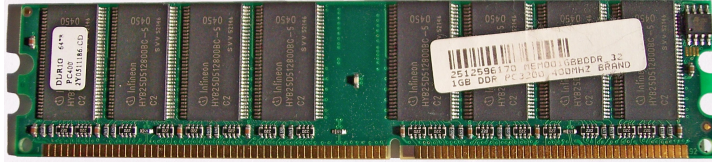
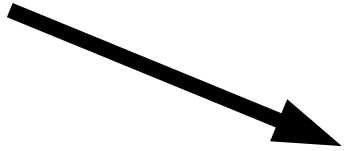


Dynamic bricking reduces I/O



In situ visualization



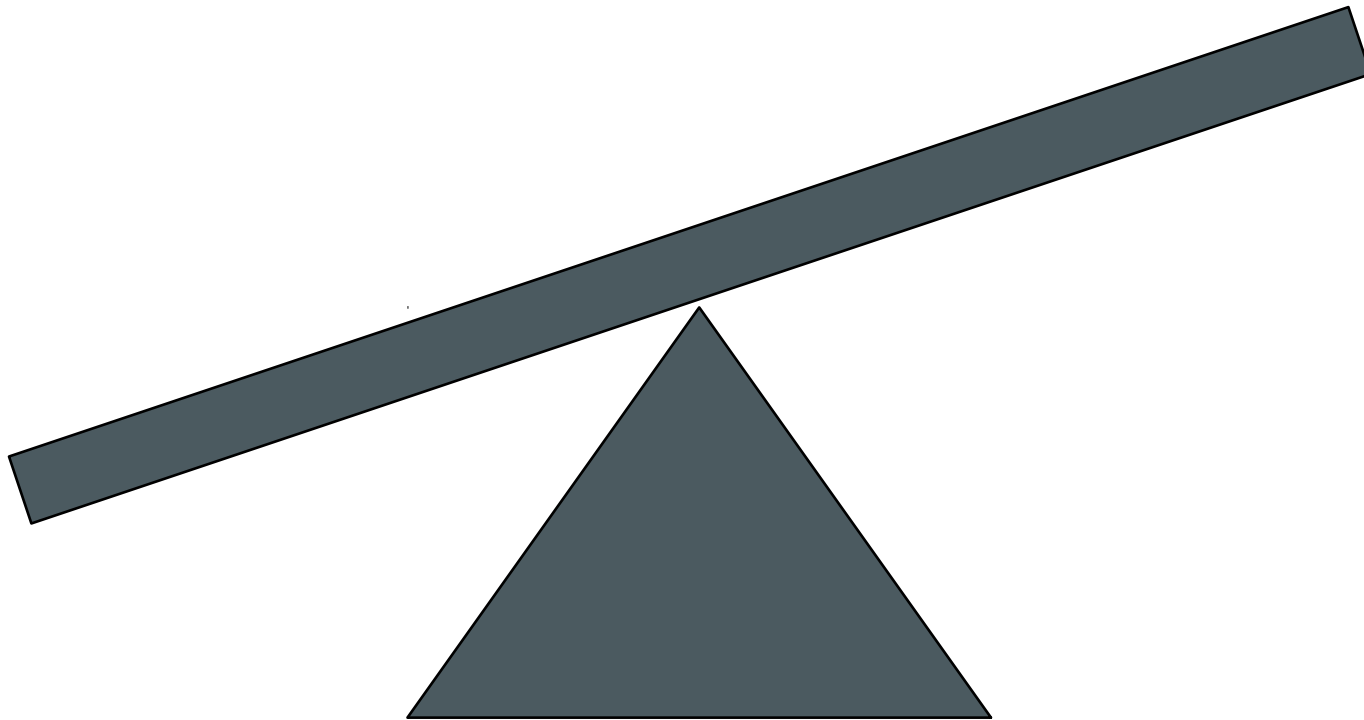


Story time

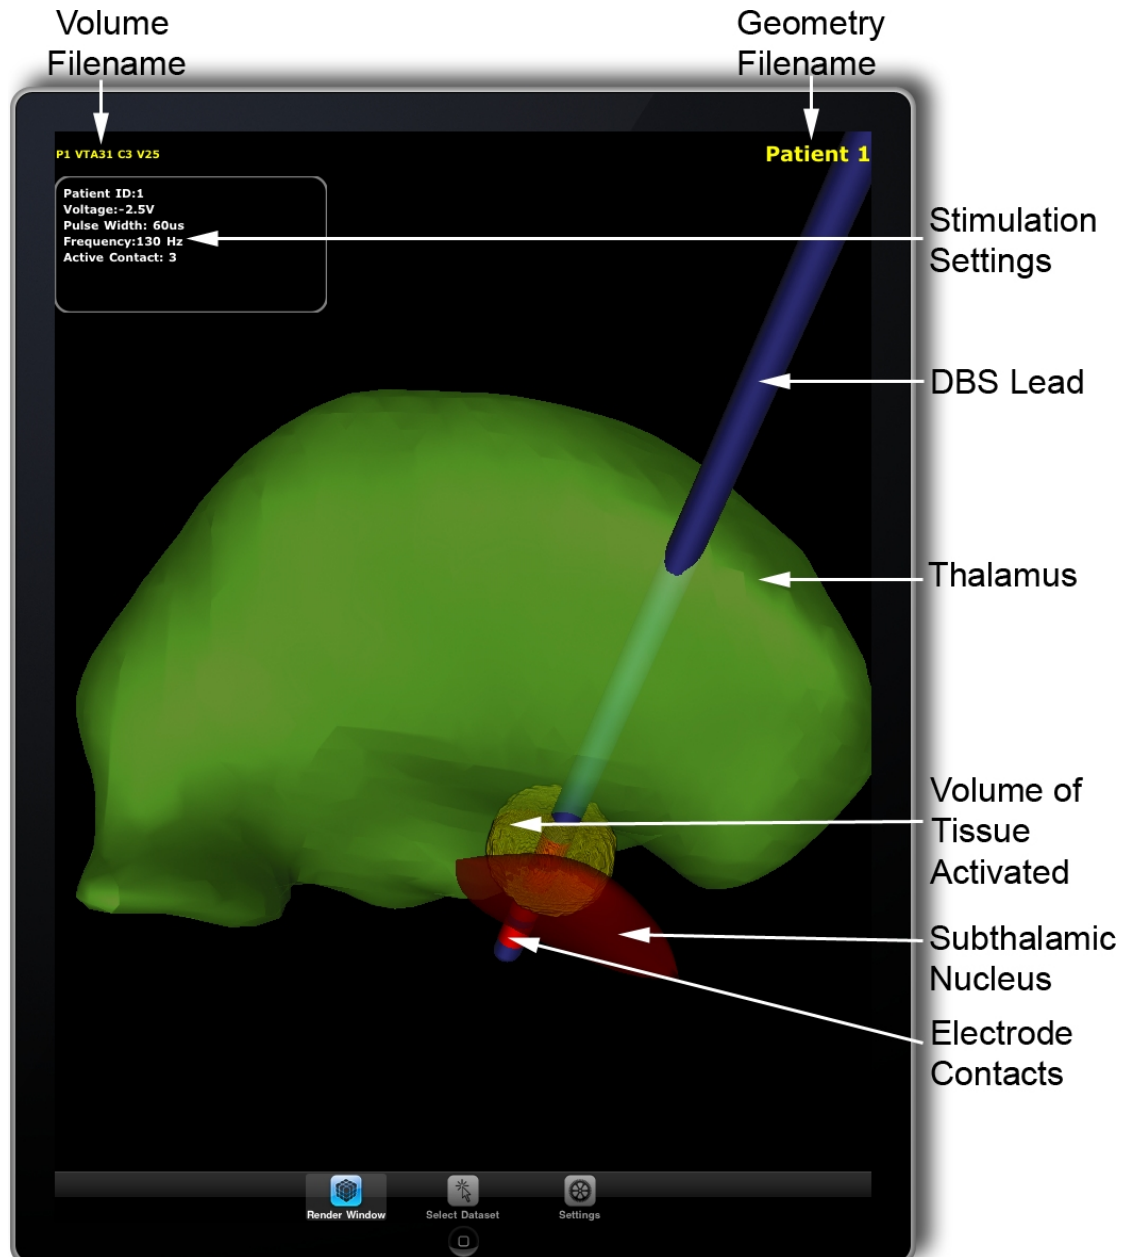
Tradeoffs

Working in
situ vis

Seeing the
light of day



Simplicity



Where to insert visualization

```
gdb -q ./a.out \  
-ex "break fopen" \  
-ex "break main" \  
-ex "run" \  
-ex "condition 1 strstr(filename, \".csv\")" \  
-ex "disable 2" \  
-ex "cont" \  
-ex "up"
```

Where to insert visualization

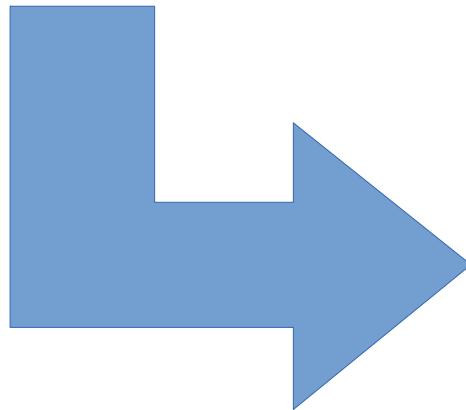
```
gdb -q ./a.out \  
-ex "break fopen" \  
-ex "break main" \  
-ex "run" \  
-ex "condition 1 strstr(filename, \".csv\")" \  
-ex "disable 2" \  
-ex "cont" \  
-ex "up"
```

Where to insert visualization

```
gdb -q ./a.out \  
-ex "break fopen" \  
-ex "break main" \  
-ex "run" \  
-ex "condition 1 strstr(filename, \".csv\")" \  
-ex "disable 2" \  
-ex "cont" \  
-ex "up"
```

Where to insert visualization

```
gdb -q ./a.out \  
-ex "break fopen" \  
-ex "break main" \  
-ex "run" \  
-ex "condition 1 strstr(filename, \".csv\")" \  
-ex "disable 2" \  
-ex "cont" \  
-ex "up"
```



Filename.c:1204

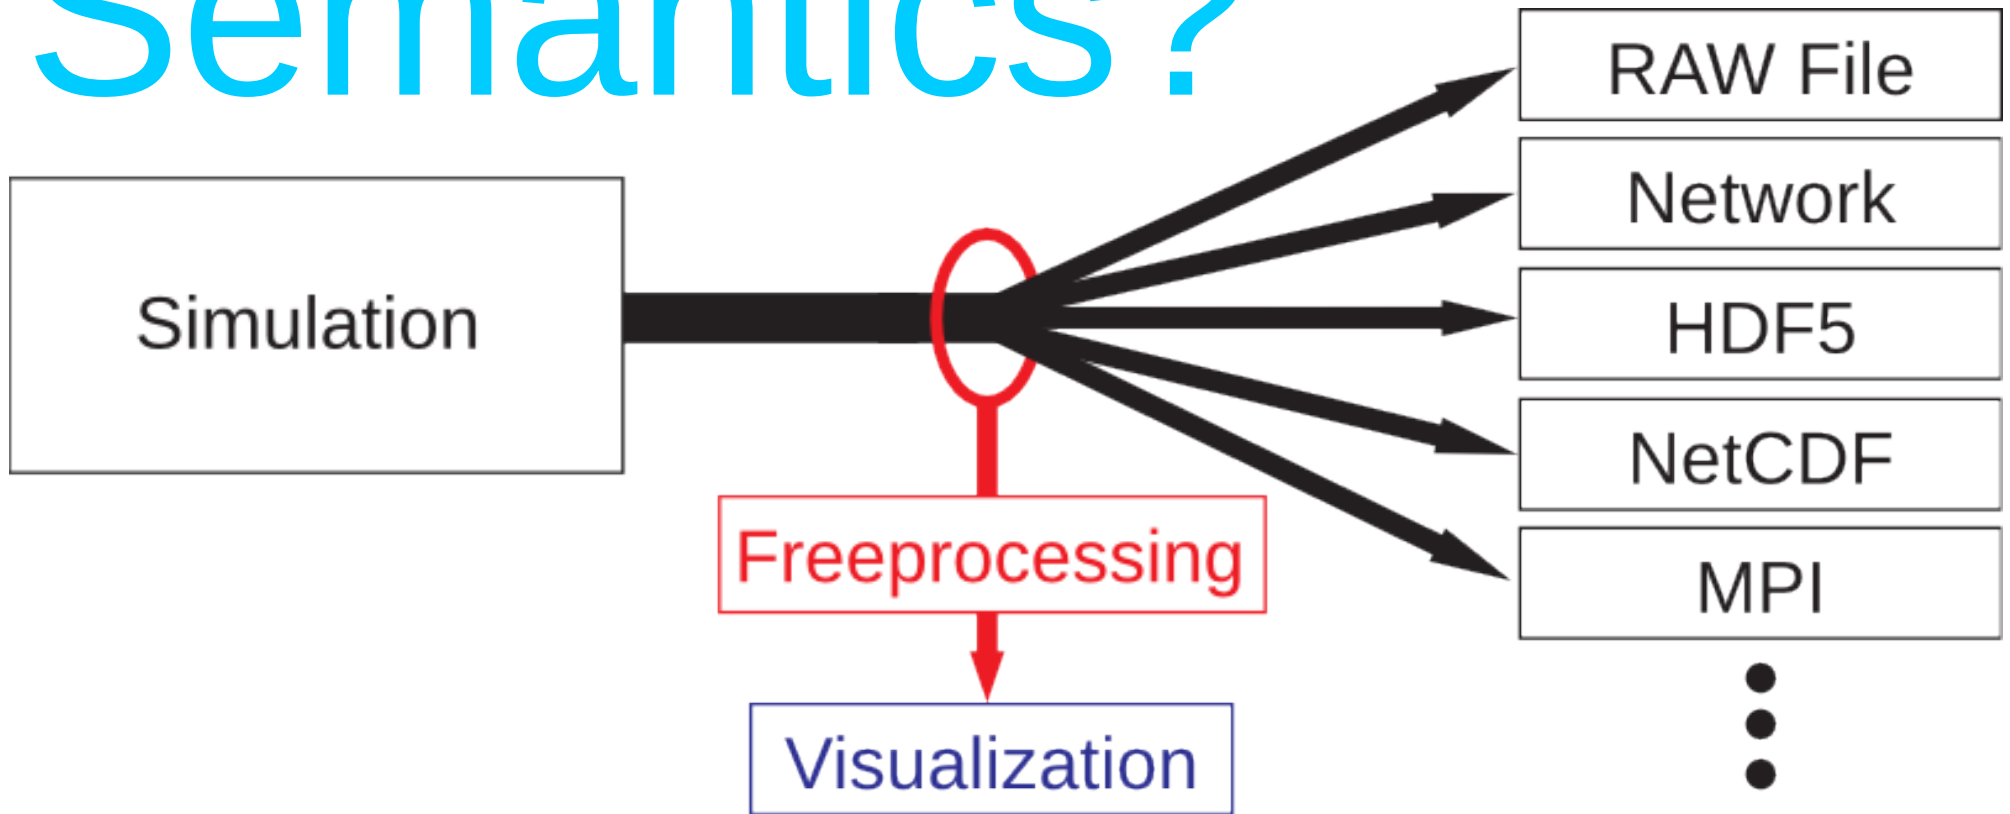
SCIENTIFIC PROGRESS GOES "BOINK"



A Calvin and Hobbes Collection by Bill Watterson

Interception

Semantics?



"Inside every large problem is a small problem struggling to get out."

– Tony Hoare

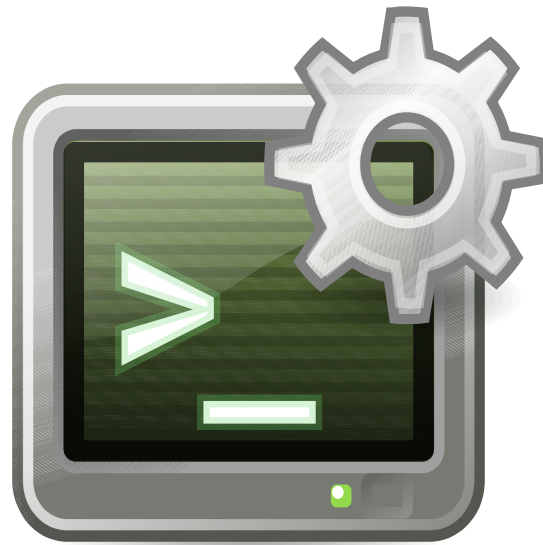
Data *model* + Program

{

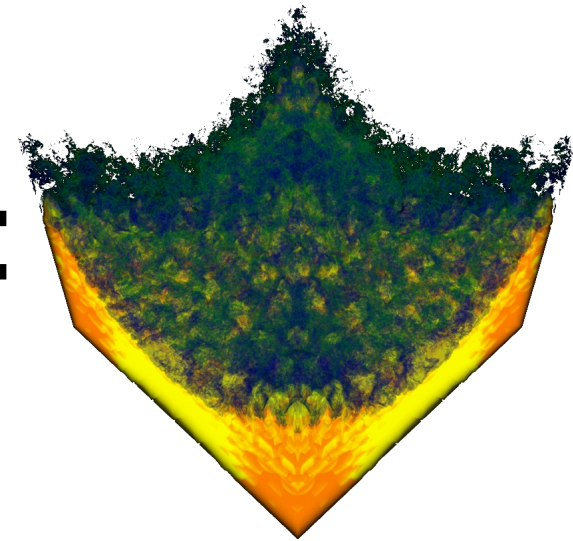
3D,
isotropic

}

+



=



ImageVis3D

```
for (uint32_t z=i0verlap; z < bricksize.z-i0verlap; z++)
for (uint32_t y=i0verlap; y < bricksize.y-i0verlap; y++)
for (uint32_t x=i0verlap; x < bricksize.x-i0verlap; x++)
{
    DOUBLEVECTOR3 vGradient = ComputeGradient(
        pTempBrickData, normalizationFactor, iCompcount,
        bricksize, x,y,z);
    if(vGradient.length() > fMaxGradMagnitude)
        fMaxGradMagnitude = vGradient.length();
}
```

LAMMPS

```
for(i=0; i<subNbx+3; i++)
  for(j=0; j<subNby+3; j++)
    for(k=0; k<subNbz+3; k++) {
      u_lb[i][j][k][0]=0.0;
      u_lb[i][j][k][1]=0.0;
      u_lb[i][j][k][2]=0.0;
      density_lb[i][j][k] = densityinit;
    }
```

libgfortran

```
for (y = 0; y < ycount; y++) {  
    bbase_y = &bbase[y*bystride];  
    s = (GFC_INTEGER_16) 0;  
    for (n = 0; n < count; n++)  
        s += abase[n*axstride] *  
            bbase_y[n*bxstride];  
    dest[y*rxstride] = s;  
}
```

Enzo

```
for (int k = k1; k <= k2; k++) {
    for (int j = j1; j <= j2; j++) {
        for (int i = i1; i <= i2+1; i++) {
            int idx3d = (k*jdim+j)*idim + i;
            float vdiff1 = 0.0f, wdiff1 = 0.0f;
            if (lj1)
                vdiff1 = (vslice[idx3d-dimx] +
vslice[idx3d-1-dimx])
                    - (vslice[idx3d+dimx] +
vslice[idx3d-1+dimx]);
```

PsiPhi

```
DO k = 1, nK
  DO j = 1, nJ
    DO i = 1, nI
      velFluctDisp = velFluctDisp +
        (VelFluct(i,j,k) - meanDb1)**2
    END DO
  END DO
END DO
```


Enzo

```
for (int k = k1; k <= k2; k++) {
    for (int j = j1; j <= j2; j++) {
        for (int i = i1; i <= i2+1; i++) {
            int idx3d = (k*jdim+j)*idim + i;
            float vdiff1 = 0.0f, wdiff1 = 0.0f;
            if (lj1)
                vdiff1 = (vslice[idx3d-dimx] +
vslice[idx3d-1-dimx])
                    - (vslice[idx3d+dimx] +
vslice[idx3d-1+dimx]);
```

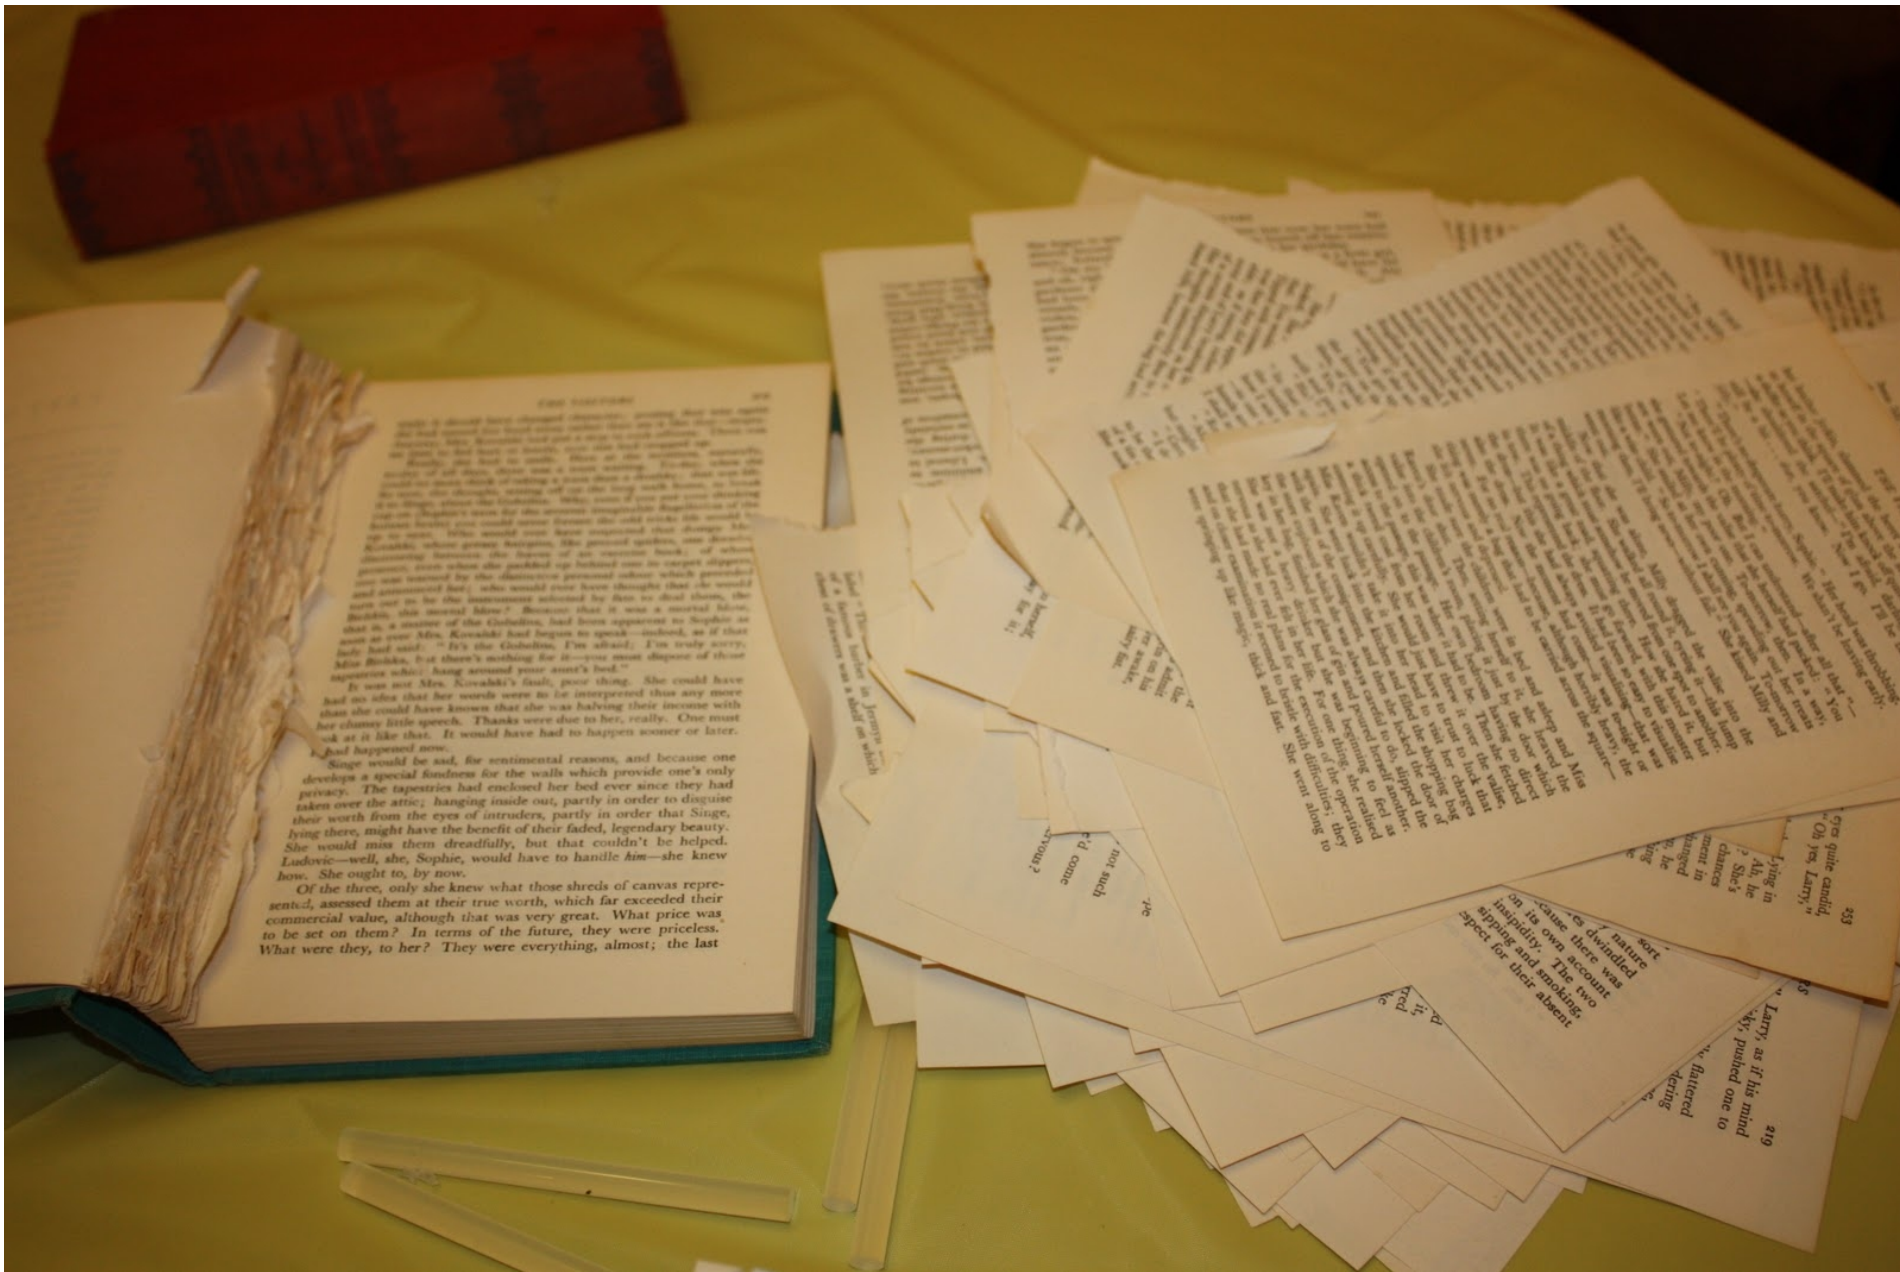
Two components to identify

- Where?
- How?
 - Finding that 'where'
 - Inserting the code

Where?

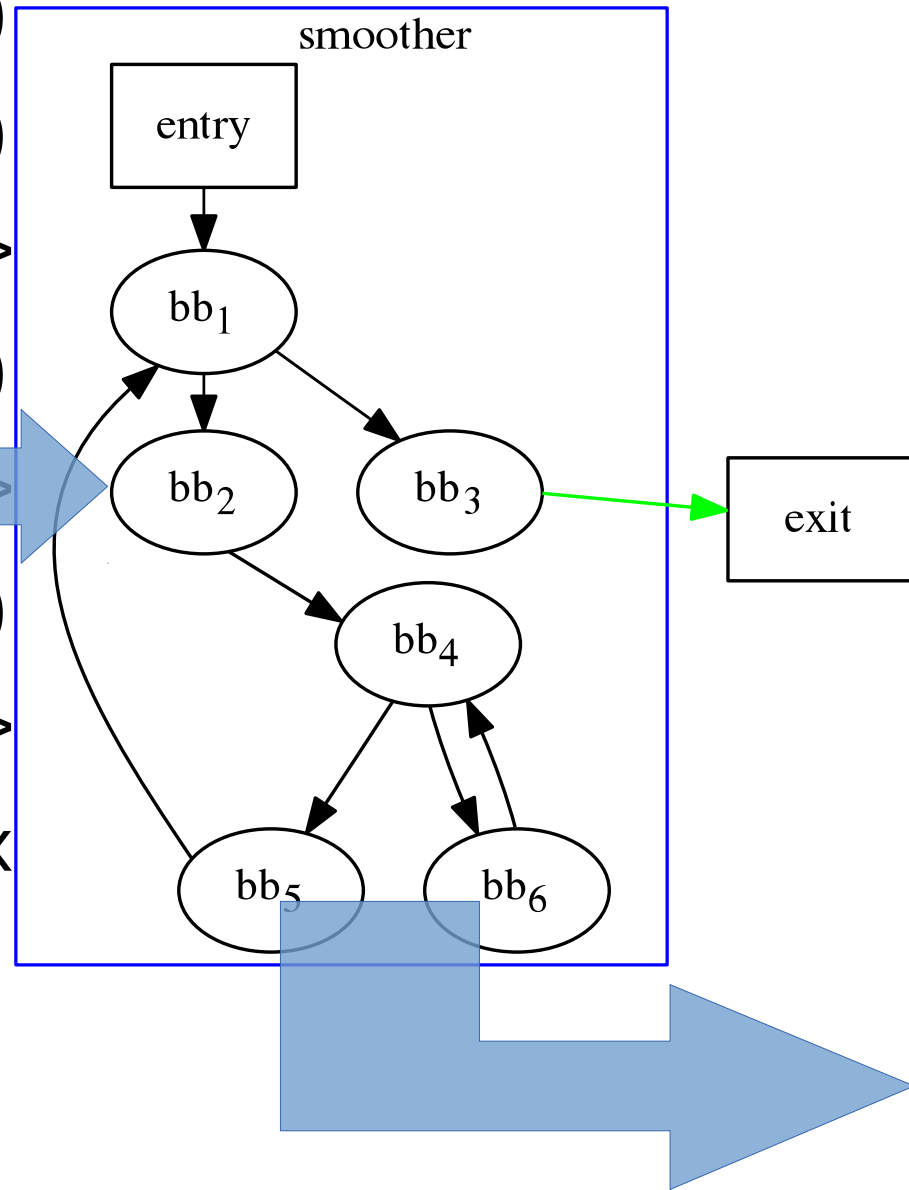
```
for(size_t y=0; y < dims[1]; ++y) {  
    const size_t row=y*dims[0];  
    for(size_t x=0; x < dims[0]; ++x)  
        array[row+x] = ...  
}  
vis(array, dims[0], dims[1], ...)
```

Page fault



How: dynamic analysis

```
mov    %rax, -0x6b8(%rbp)
movq   $0x0, -0x6f0(%rbp)
jmpq   400bba <s3+0x14e>
movq   $0x0, -0x6e8(%rbp)
jmpq   400b97 <s3+0x12b>
movq   $0x0, -0x6e0(%rbp)
jmp    400b78 <s3+0x10c>
mov    -0x6f8(%rbp), %rax
add    $0x8, %rax
mov    (%rax), %rax
```



Decompilation

```
for(size_t y=0; y < dims[1]; ++y) {  
    const size_t row=y*dims[0];  
    for(size_t x=0; x < dims[0]; ++x)  
        array[row+x] = ...  
}
```

Summary

A dark gray equilateral triangle pointing upwards, containing the text 'GPU' in white.

GPU

A dark gray equilateral triangle pointing upwards, containing the text 'Sampling' in white.

Sampling

A dark gray equilateral triangle pointing upwards, containing the text 'Simplicity' in white.

Simplicity

?



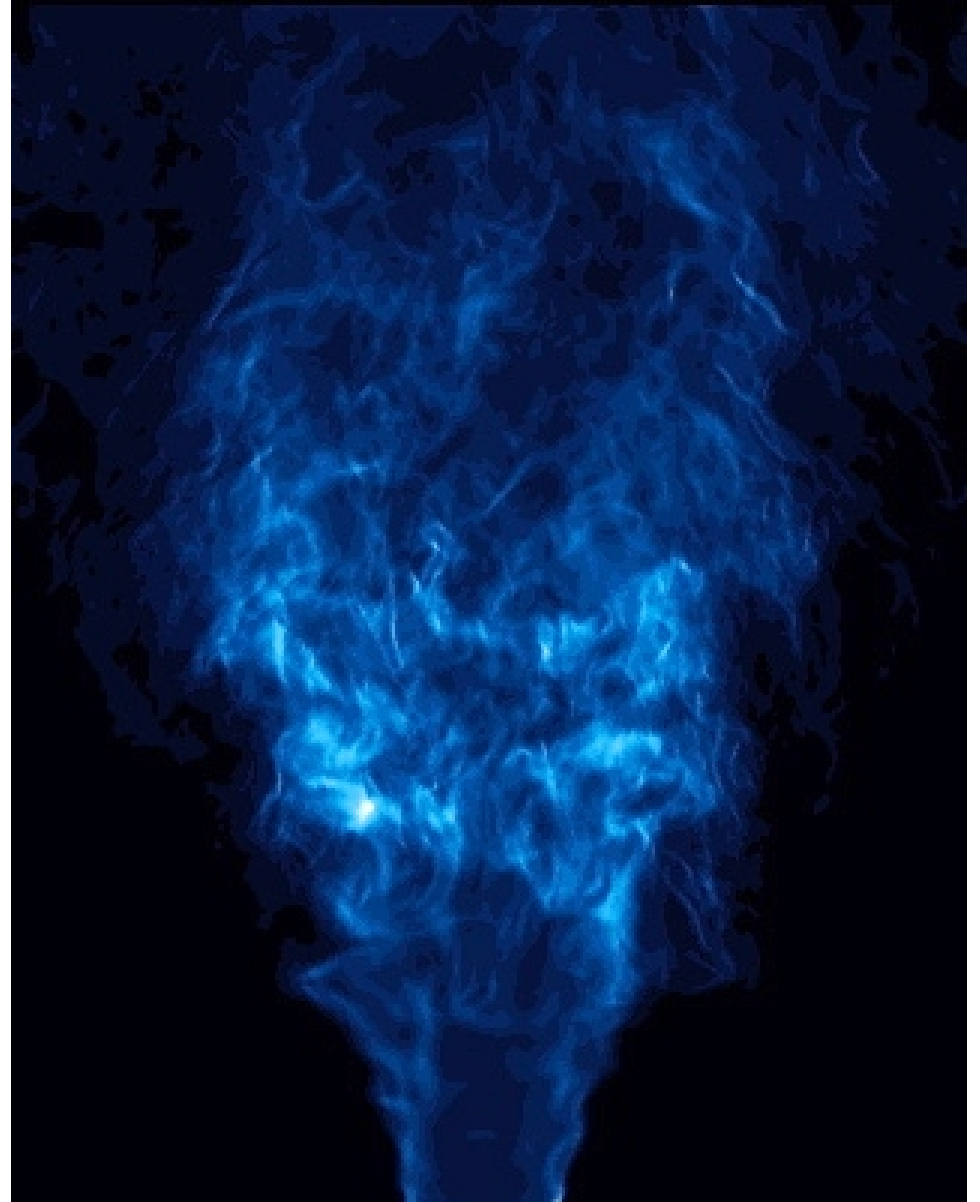
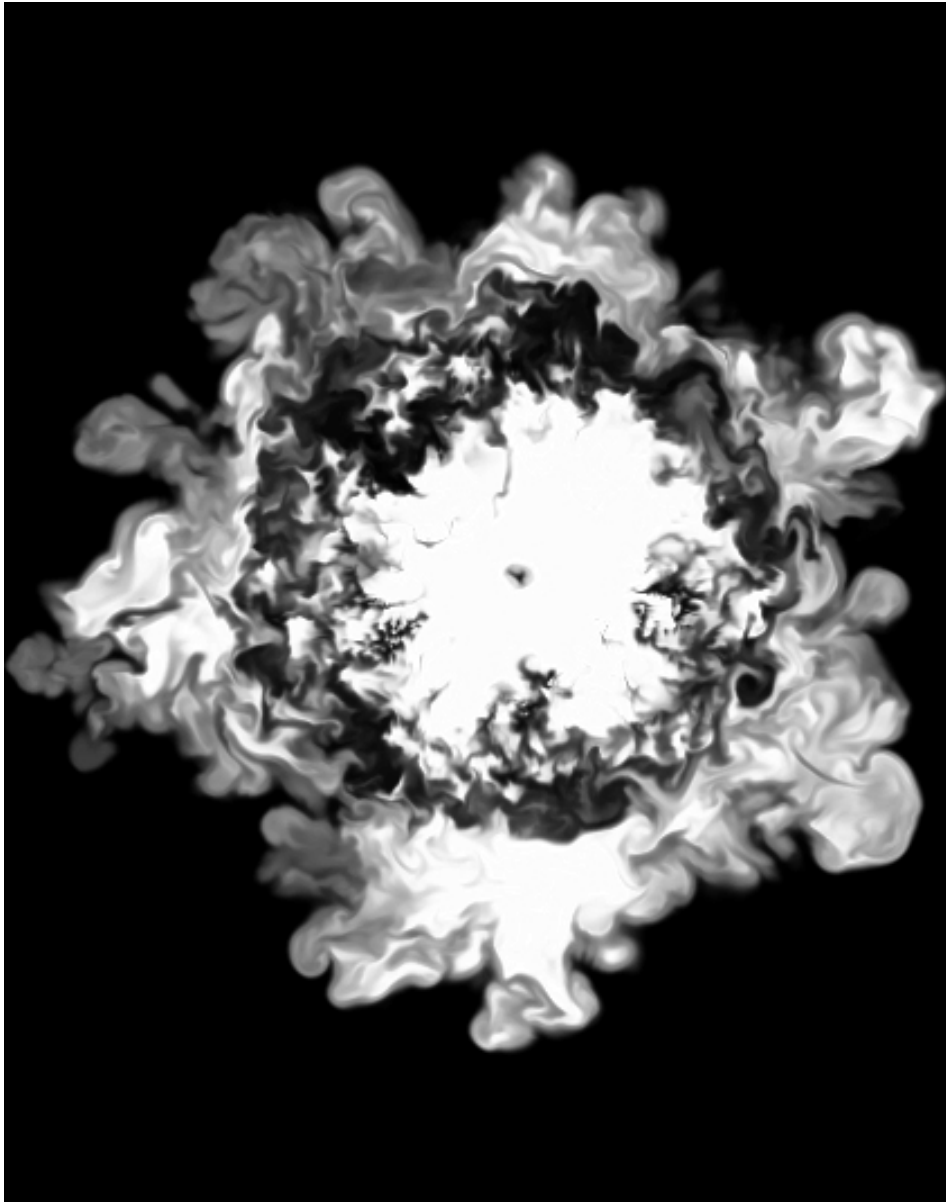
Thanks: Flickr User:jurvertson

Binary Instrumentation

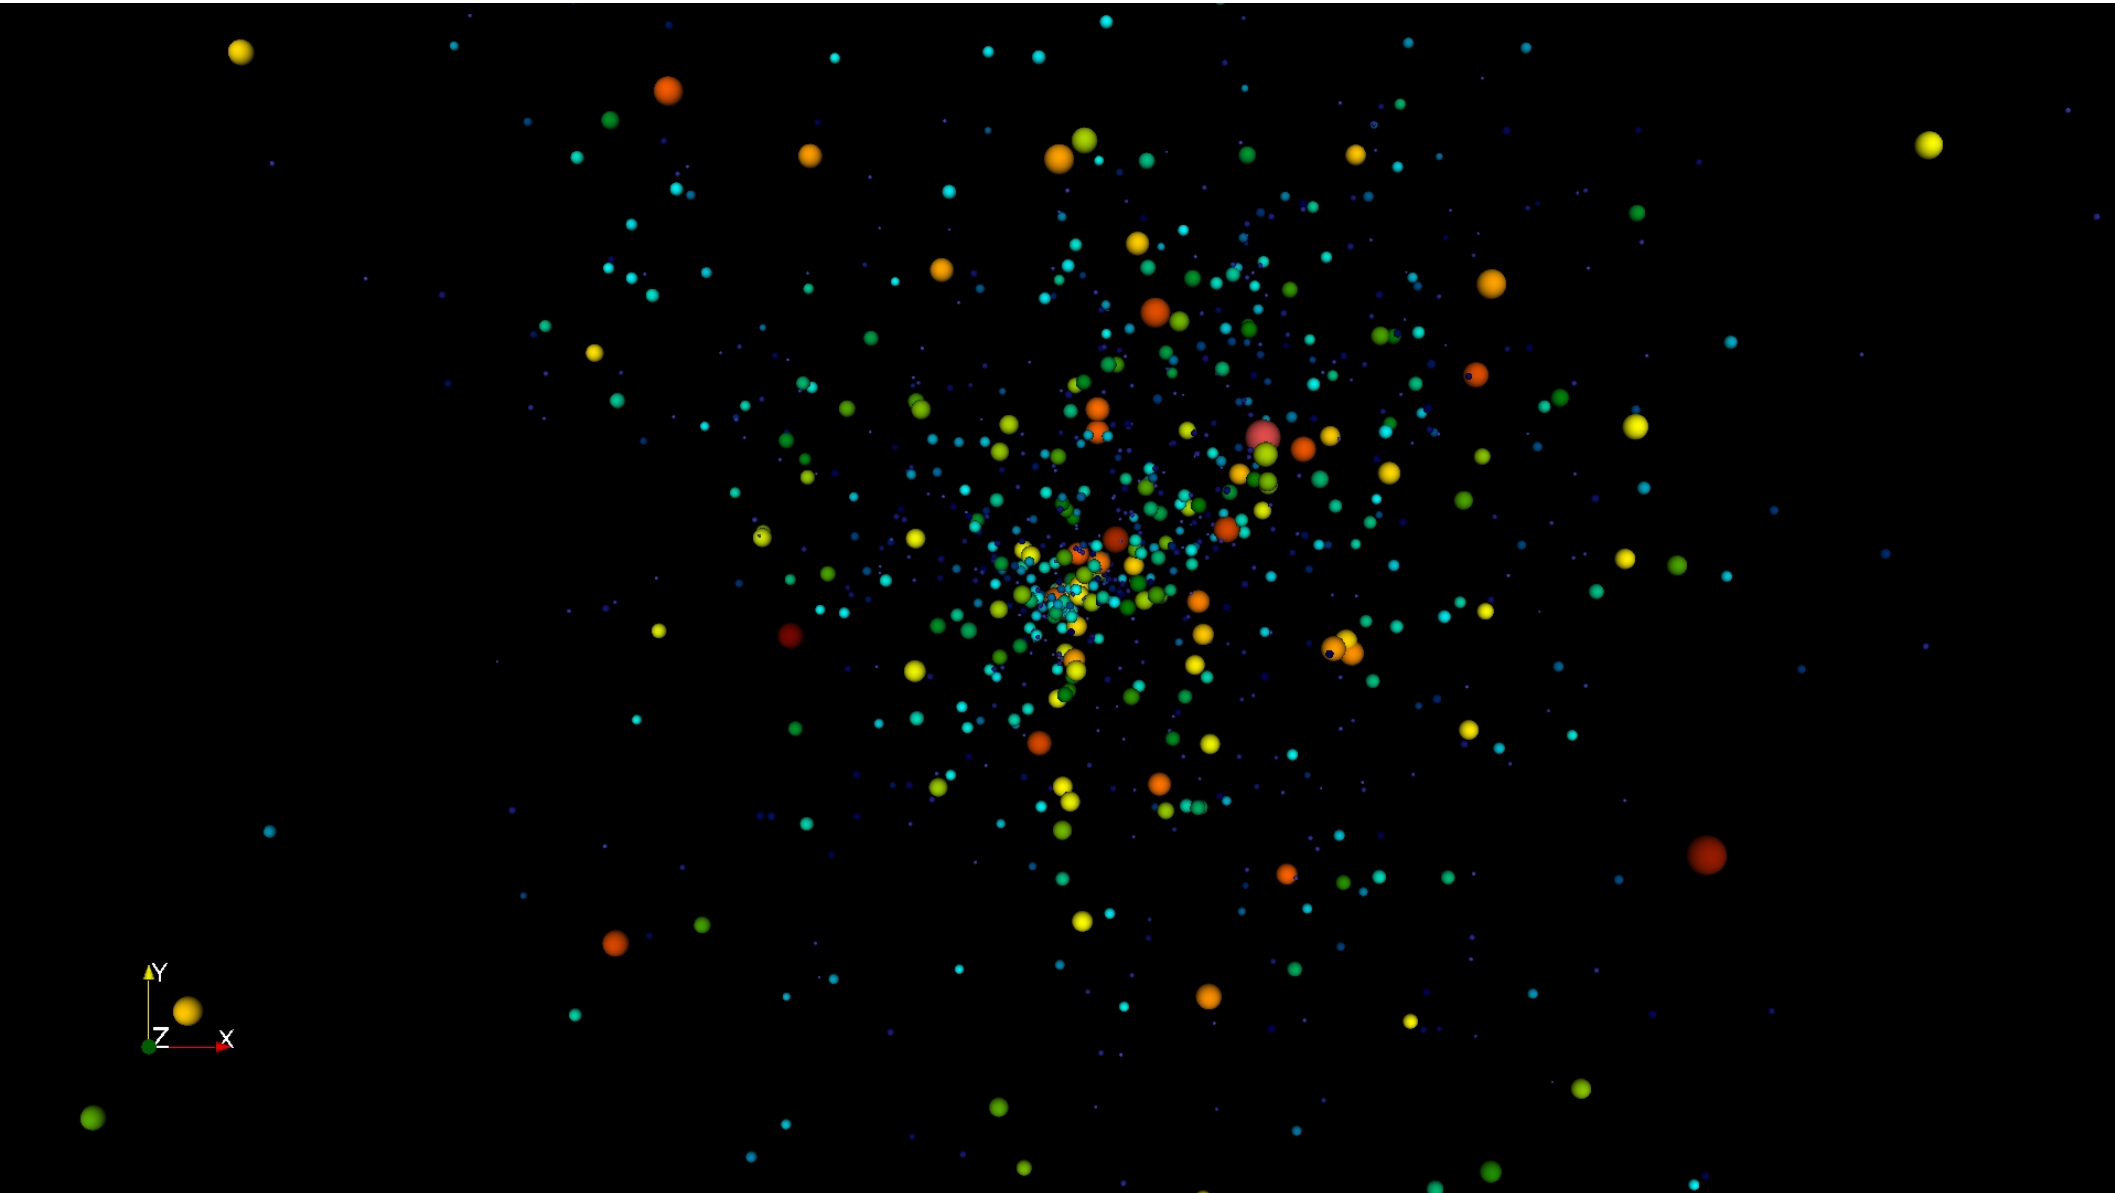
File: libsitu.so				ASCII Offset: 0x000007D0 / 0x0000C7DE (%04)												
000007A0	44	32	00	00	00	00	00	00	00	00	00	00	00	00	00	D2.....
000007B0	00	5F	5F	67	6D	6F	6E	5F	73	74	61	72	74	5F	5F	.__gmon_start__
000007C0	5F	69	6E	69	74	00	5F	66	69	6E	69	00	5F	49	54	__init_fini_ITM
000007D0	5F	64	65	72	65	67	69	73	74	65	72	54	4D	43	6C	__deregisterTMClon
000007E0	6E	65	54	61	62	6C	65	00	5F	49	54	4D	5F	72	65	neTable.ITM_reg
000007F0	69	73	74	65	72	54	4D	43	6C	6F	6E	65	54	61	62	isterTMCloneTabl
00000800	65	00	5F	5F	63	78	61	5F	66	69	6E	61	6C	69	7A	e.__cxa_finalize
00000810	00	5F	4A	76	5F	52	65	67	69	73	74	65	72	43	6C	.__Jv_RegisterCla
00000820	73	73	65	73	00	67	65	74	70	69	64	00	73	74	64	sses.getpid.stdo
00000830	75	74	00	66	69	6C	65	6E	6F	00	69	73	61	74	74	ut_FILENO.isatty
00000840	00	5F	5F	61	73	73	65	72	74	5F	66	61	69	6C	00	.__assert_fail.s
00000850	79	6D	62	5F	64	62	67	00	76	70	72	69	6E	74	66	ymb_dbg.vprintf.
00000860	70	75	74	73	00	73	74	72	6E	63	6D	70	00	73	79	puts.strncmp.sym
00000870	62	5F	70	61	72	73	65	5F	6F	70	74	69	6F	6E	73	b_parse_options.
00000880	73	74	72	64	75	70	00	73	74	72	63	68	72	00	73	strdup.strchr.st
00000890	72	6C	65	6E	00	66	72	65	65	00	67	65	74	65	6E	rlen.free.getenv
000008A0	00	66	6E	6D	61	74	63	68	00	66	6F	70	65	6E	00	.__fnmatch.fopen.f
000008B0	65	72	72	6F	72	00	66	65	6F	66	00	66	73	63	61	error.feof.fscan
000008C0	66	00	5F	5F	65	72	72	6E	6F	5F	6C	6F	63	61	74	f.__errno_locati
000008D0	6F	6E	00	64	6C	65	72	72	6F	72	00	64	6C	6F	70	on.dlerror.dlope
000008E0	6E	00	64	6C	73	79	6D	00	73	74	72	6E	63	61	73	n.dlsym.strncase

^G Help ^C Exit (No Save) ^T goTo Offset ^X Exit and Save ^W Search

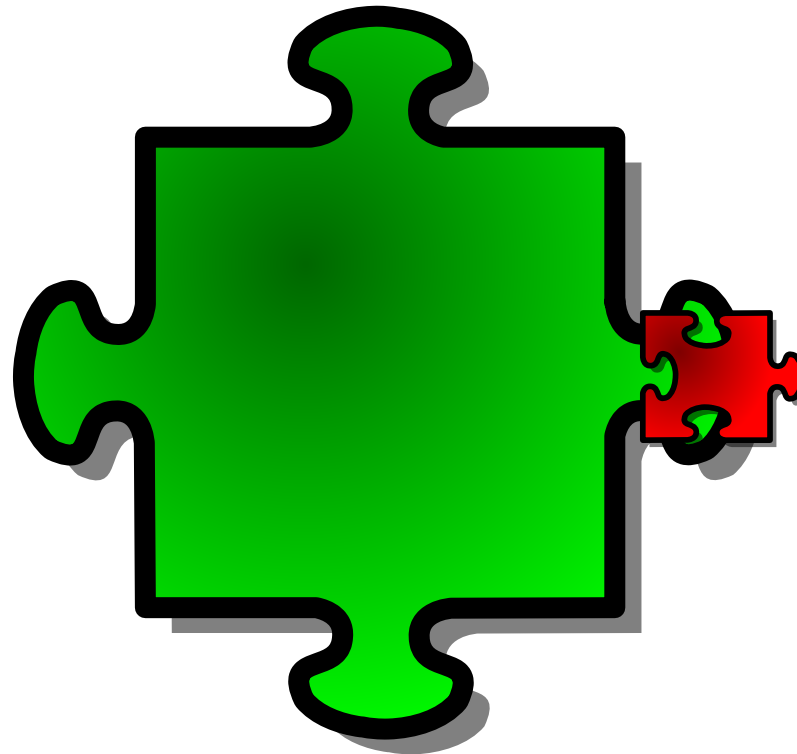
Traditional use cases are viable



What *about* the data?



Small connectors encourage ad hoc computations



Network-based *in situ*

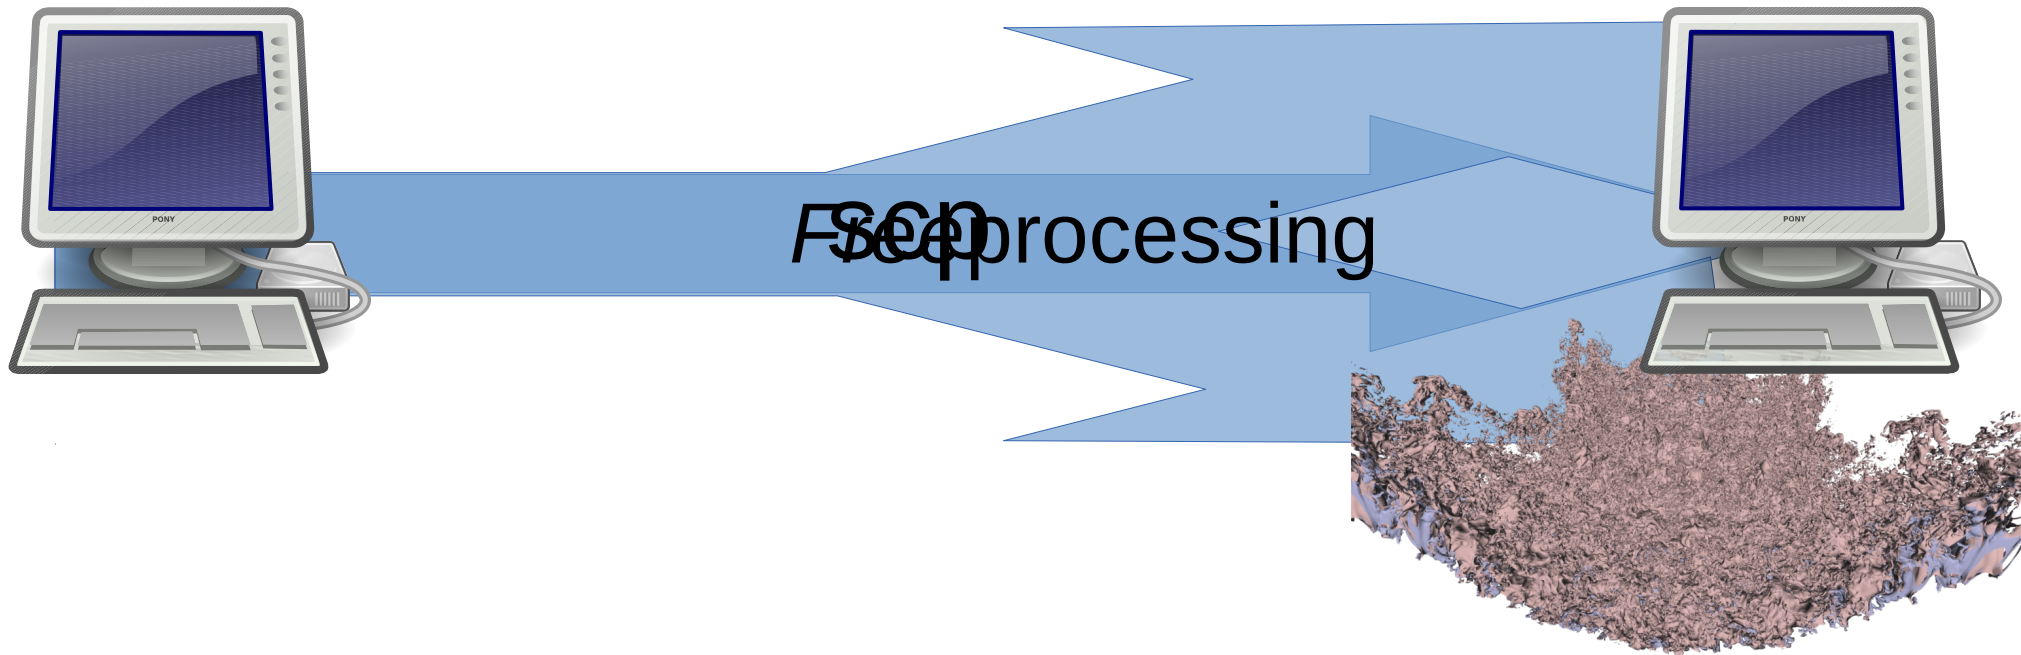


Image Credits

- Wikipedia User:Mysid, User:Lobsterbake, User:Someone35
- Flickr user jurvertson
- Tango project
- Silo web site
- My HPC students

