



High-Performance Computing – Performance Engineering Case Study –

Christian Plessl & Michael Lass

High-Performance IT Systems Group
Paderborn University, Germany

version 1.1.0 2018-02-02



Paderborn
Center for
Parallel
Computing

- Vectorization
- Roofline model
- Case study: Performance optimization for n-body solver

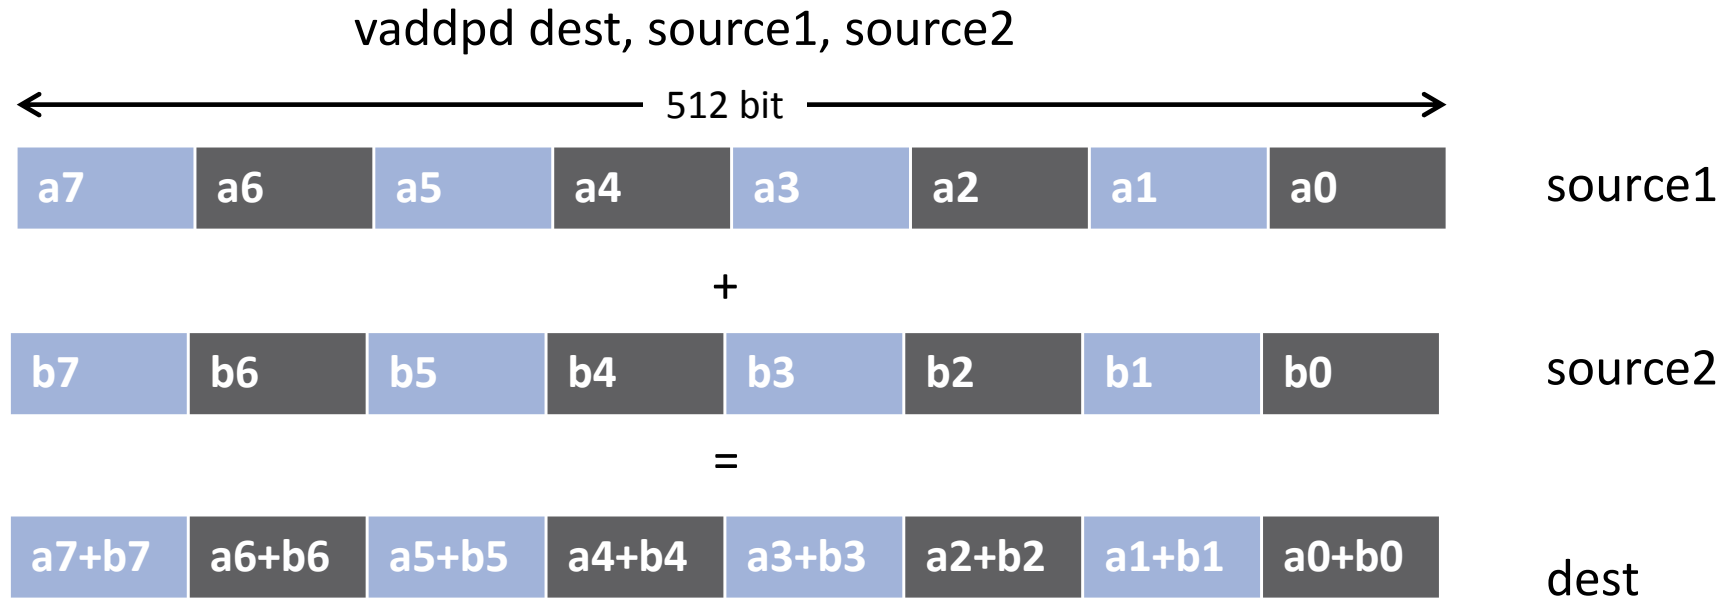
Vectorization

Vectorization

- Recent CPUs architectures have increasingly powerful SIMD, e.g. Intel:
 - width of SIMD registers
 - SSE (128bit): 2 double precision (DP) or 4 single precision (SP)
 - AVX (256bit): 4 DP / 8 SP
 - AVX512 (512bit): 8 DP / 16 SP
 - SIMD operations have become more versatile and efficient (see examples on next slides)
- Applications that do not use vectorization can only exploit a small fraction of the peak performance of modern CPUs
- We use the terms vector and SIMD instructions as synonyms; there are differences but they do not matter for this discussion

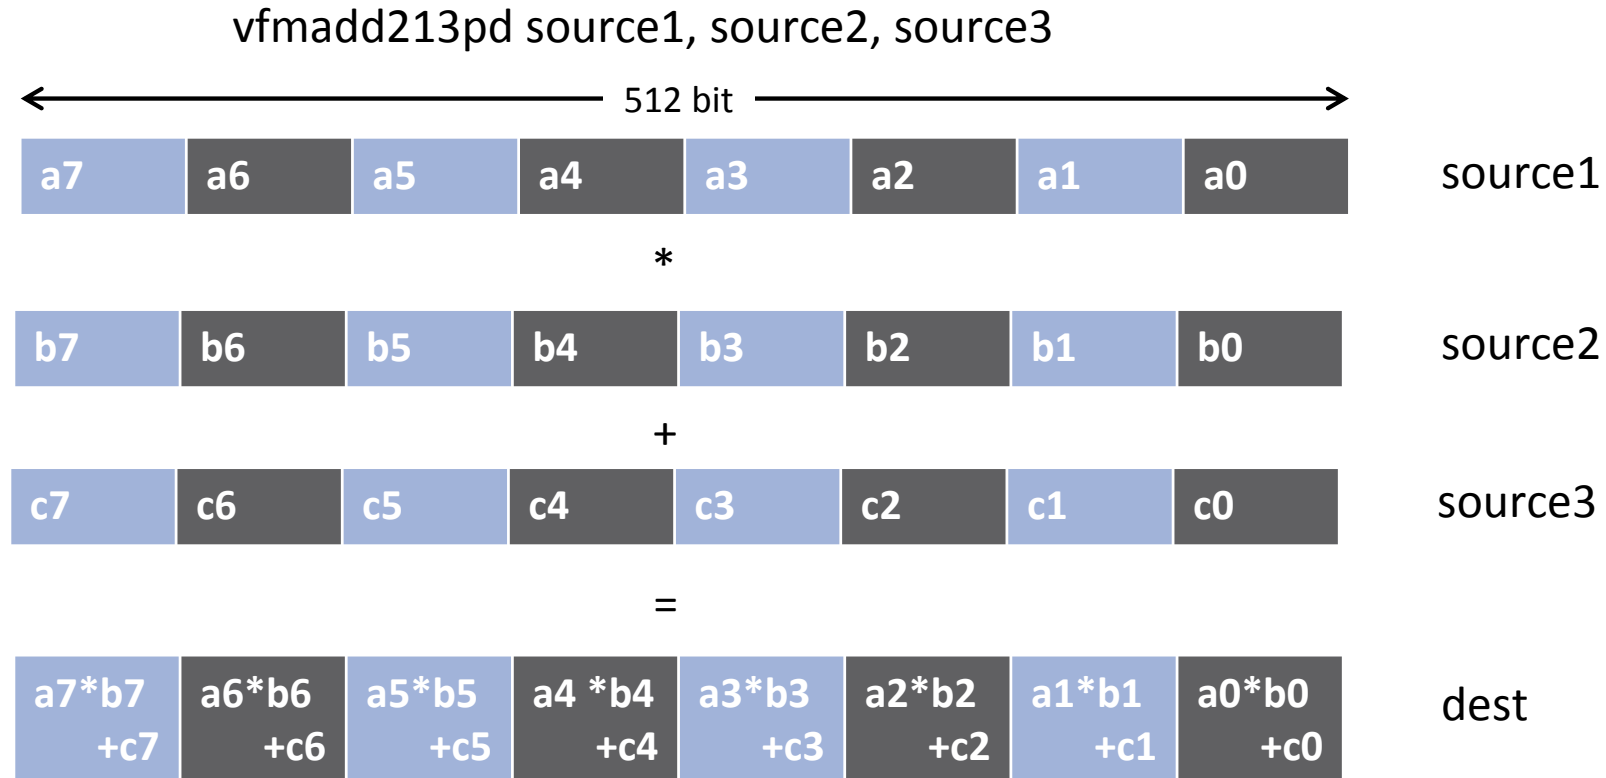


Examples for SIMD instructions



vector addition

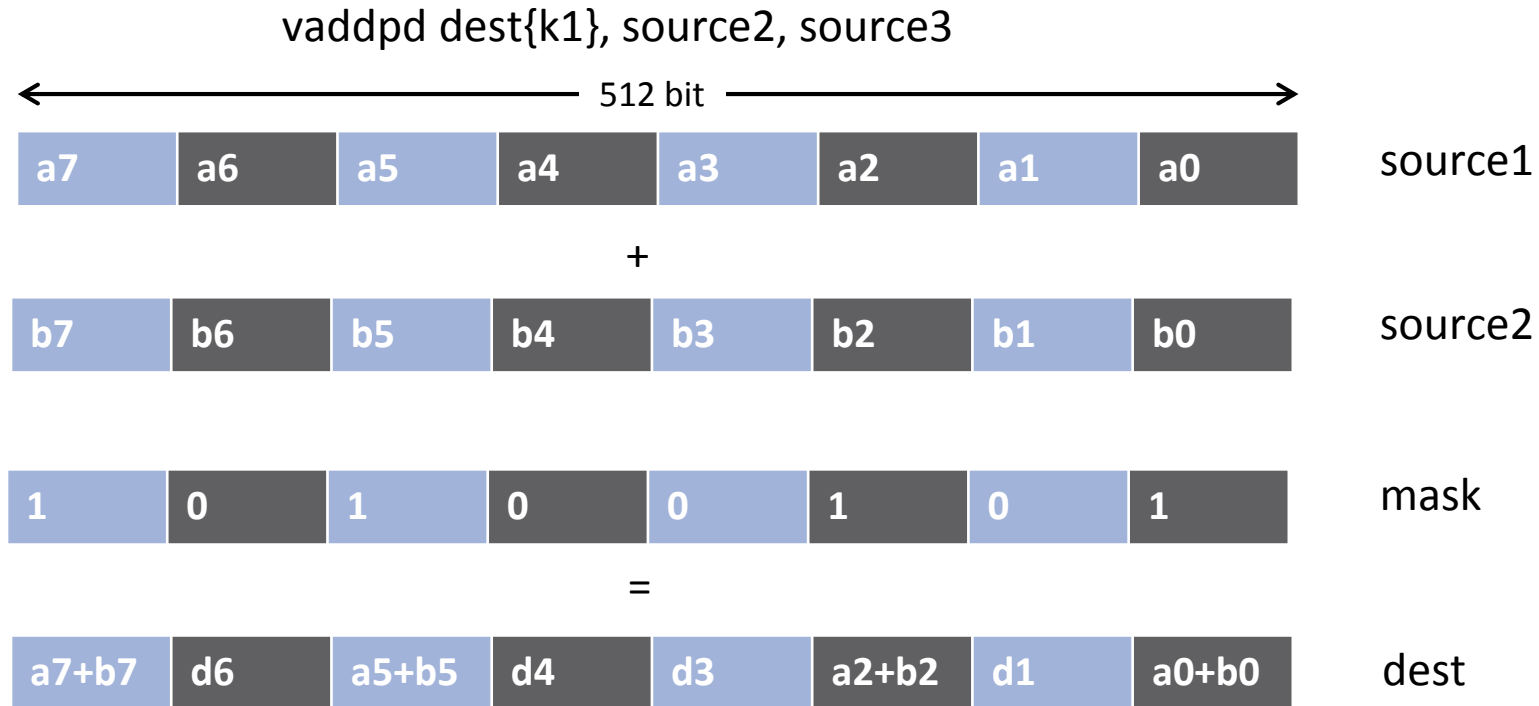
Examples for SIMD instructions (2)



fused multiply-add (FMA)

Combines two floating-point operations into a single instruction (multiplication and addition).
This instruction achieves the peak floating-point performance

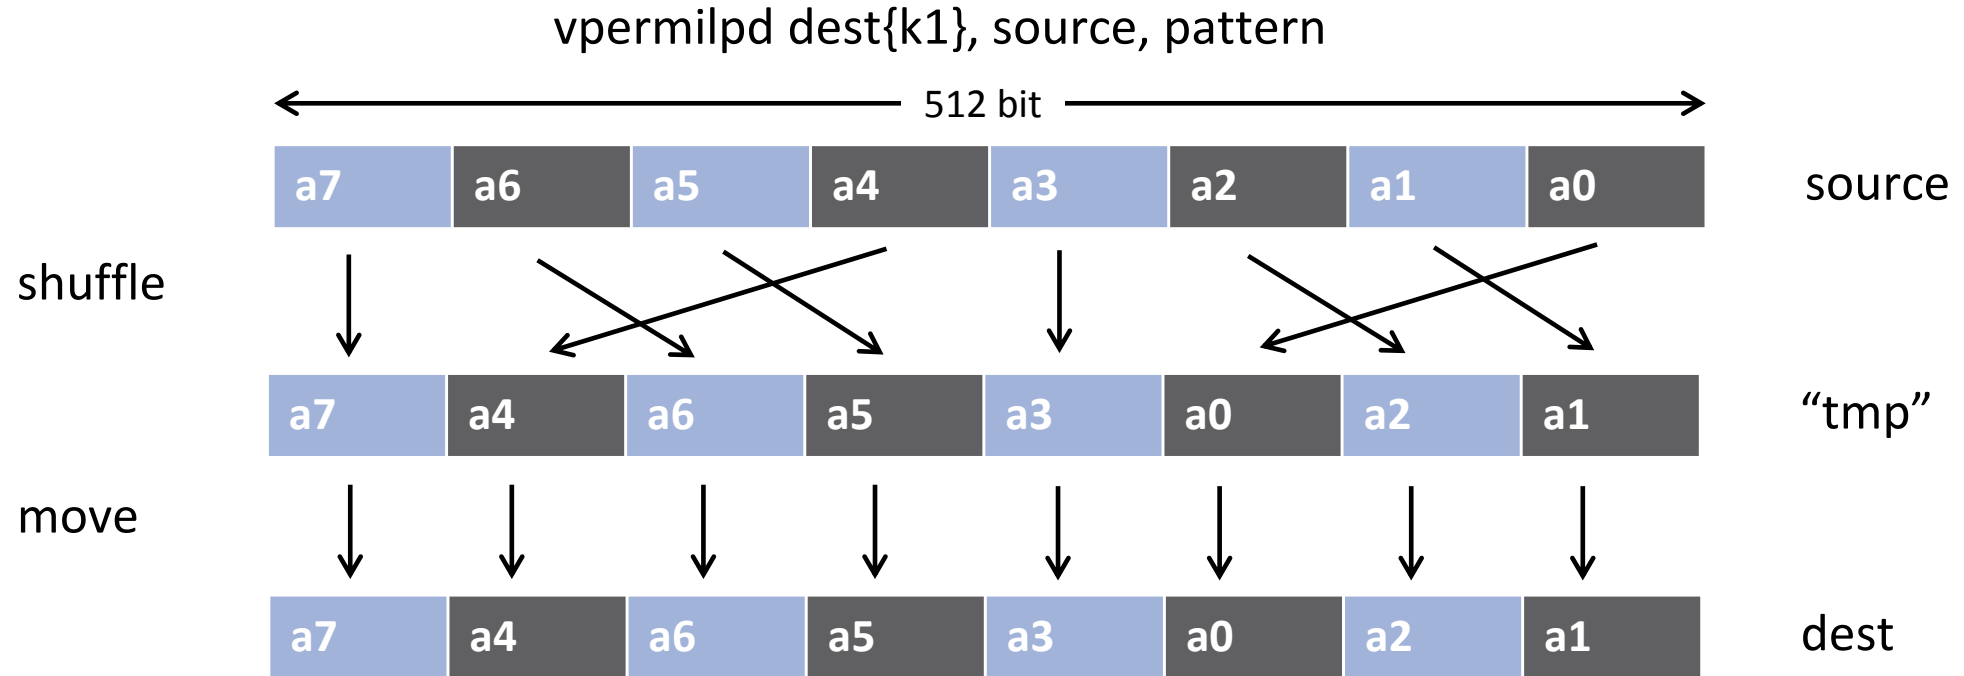
Examples for SIMD instructions (3)



vector addition with masks

Can selectively write results to vector depending on mask-bit. Essential for efficient vectorization of codes with conditional execution.

Examples for SIMD instructions (4)

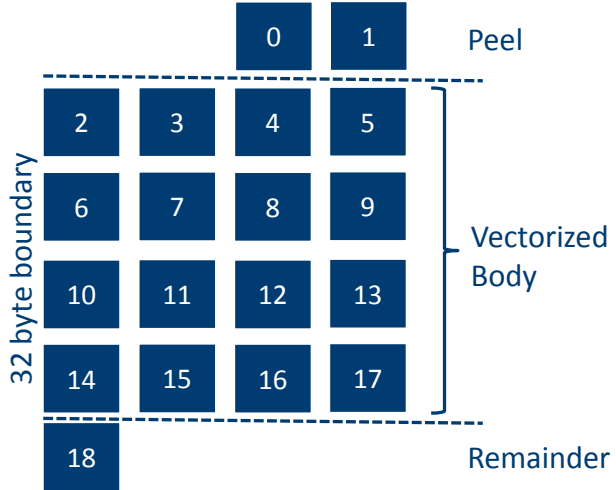


vector permutation

Loop Vectorization

- **Vectorization is typically applied to loops**
 - a couple of independent loop iterations are executed concurrently using SIMD instructions
- **Challenges**
 - number of loop iterations may be not evenly divisible by vector length
 - data to be processed (typ. arrays) may not be properly aligned for efficient transfer to SIMD registers
- **Anatomy of a vectorized loop**
 - **peel** (optional): used for unaligned iterations of loop, uses scalar instructions or slower SIMD instructions
 - **body**: fully vectorized, uses complete vector length
 - **remainder** (optional): process remaining iterations

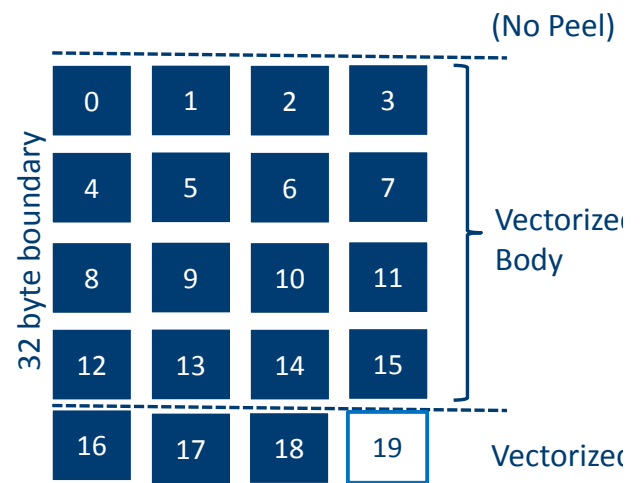
Loop Vectorization (2)



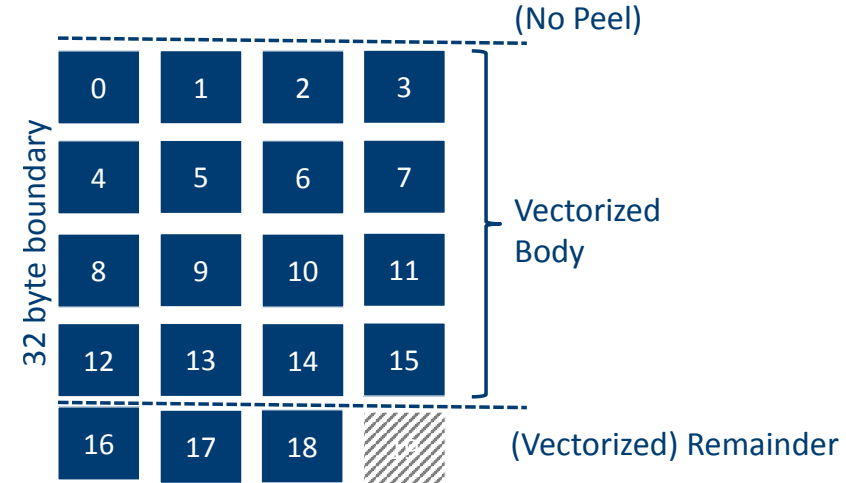
vectorized loop with peel and remainder (unaligned)

```
// xAVX
// 256 bits wide regs
// holds 4 x 64bit vals

void Func(double *pA)
{
    for (int i=0; i<19;i++)
        pA[i] = ...;
}
```



vectorized loop with remainder (aligned), remainder executed with scalar instructions



vectorized loop with remainder (aligned), remainder executed with SIMD instructions and masks

How to Generate Vectorized Code

- 1. Use optimized libraries that have vectorization built-in
 - examples: Intel Math Kernel Library (MKL), AMD Core Math Library (ACML), IBM Engineering and Scientific Subroutine Library (ESSL), etc.
- 2. Auto-vectorizing compiler
 - recent compilers support automatic vectorization of code
 - compilers must make conservative assumptions to avoid breaking code
 - rewrite performance-critical code sections to make them vectorizable or to convince the compiler that vectorization is safe
- 3. Code annotations to mark sections that are safe for vectorization
 - compiler-specific directives #pragmas (e.g. ICC has `#pragma ivdep`)
 - OpenMP offers a portable directive `#pragma omp simd`
- 4. Intrinsic or assembly code

Auto Vectorization

- Most recent compilers (Intel ICC, GCC, LLVM) support automatic vectorization of regular loops
- Intel ICC
 - enable vectorization with option `-vec` (automatically enabled with `-O3`)
 - `-vec-report` and `-opt-report` generate detailed diagnostics explaining what code parts have been successfully vectorized and which parts could not be vectorized and what prevented the compiler from vectorization
 - `-xHost` instructs compiler to use instructions for local CPU architecture
- GCC
 - `-O -ftree-vectorize` enables autovectorization (automatically enabled with `-O3`)
 - `-fopt-info-vec`, `-fopt-info-vec-missed` generates diagnostics about successful and failed vectorization of loops
 - `-march=native` instructs compiler to use instructions for local CPU architecture

Obstacles to Vectorization: Non-Contiguous Memory Access

```
for (i=0; i<=MAX; i++) {  
    c[i]=a[i]+b[i];  
}
```

code that is ideally suited for vectorization

loading data efficiently into SIMD registers requires that data is stored contiguously in main memory

```
// arrays accessed with stride 2  
for (int i=0; i<SIZE; i+=2)  
    b[i] += a[i] * x[i];
```

```
// inner loop accesses a with stride SIZE  
for (int j=0; j<SIZE; j++) {  
    for (int i=0; i<SIZE; i++) {  
        b[i] += a[i][j] * x[j];  
    }  
}
```

```
// indirect addressing of x using index array  
for (int i=0; i<SIZE; i+=2) {  
    b[i] += a[i] * x[index[i]];  
}
```

typical problems preventing vectorization due to non-contiguous memory access

Obstacles to Vectorization: Data Dependencies

```
A[0]=0;
for (j=1; j<MAX; j++) {
    A[j]=A[j-1]+1;
}
```

```
// equivalent to
A[1]=A[0]+1;
A[2]=A[1]+1;
A[3]=A[2]+1;
A[4]=A[3]+1;
```

Read-after-write (flow) dependency

Cannot execute several iterations concurrently, because values $A[j-1]$ required in iteration j is known only after iteration $j-1$ has finished

```
for (j=1; j<MAX; j++) {
    A[j-1]=A[j]+1;
}
```

```
// equivalent to
A[0]=A[1]+1;
A[1]=A[2]+1;
A[2]=A[3]+1;
A[3]=A[4]+1;
```

Write-after-read dependency

Not safe for general parallelization but safe for vectorization! We know that no iteration with a higher value of j can be executed before iteration with lower value of j

```
for (j=1; j<MAX; j++) {
    A[j-1]=A[j]+1;
    B[j]=A[j]*2;
}
```

```
// equivalent to
A[0]=A[1]+1;
A[1]=A[2]+1;
A[2]=A[3]+1;
A[3]=A[4]+1;
```

Write-after-read dependency

Not safe for vectorization because some $A[j]$ may be overwritten by the first SIMD instruction before they are read by the second SIMD instruction

Obstacles to Vectorization: Pointer Aliasing

```
void vadd(float *a, float *b, float *c, int n)
{
    for(int i=0; i<n; i++) {
        c[i] = a[i] + b[i];
    }
}
```

vector add: with potential aliasing, i.e. arrays a, b, and c might partially overlap



provide meta information to compiler to enable vectorization

- a, b and c are non-overlapping → use C99 restrict keyword
- a, b, n are read-only → use const modifier

```
void vadd(const restrict float *a,
          const restrict float *b,
          const restrict float *c,
          const int n)
{
    for(int i=0; i<n; i++) {
        c[i] = a[i] + b[i];
    }
}
```

Where Autovectorizers Tend to Fail

- Most frequent reason: data dependencies
- Further common reasons
 - potential pointer/array aliasing
 - unsuitable alignment
 - function calls in loop block
 - loop not countable (loop bound is not a runtime constant)
 - mixed data types
 - non-unit stride between elements
 - loop body too complex (register pressure)
 - profitability models deems vectorization as inefficient
- Many additional, but less likely reasons

OpenMP SIMD Constructs

The OpenMP simd Construct

- Conceptually, vectorization is conceptually similar to loop parallelization with OpenMP work sharing or reductions
- OpenMP 4.0 introduced the simd directive
 - explicit vectorization of a loop test
 - cuts loop into chunks that fit a SIMD vector register
 - by default no parallelization, but multi-threaded worksharing variant (simd for”) is available too

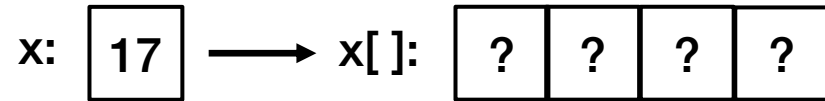
```
#pragma omp simd [clauses]  
for-loop
```

- Clauses are similar to for-directive

The simd Data Sharing Clauses

- **private (var-list)**

- expand scalar variables to uninitialized vectors



- **lastprivate(var-list)**

- **reduction(op : var-list)**

- expand scalar variable to vector
- compute reduction of elements of vector by applying operation op

```
double sum_all (double *a, double *b, int n)
{
    double tmp, sum;
    sum = 0.0;
    #pragma omp simd \
        private(tmp) reduction(+:sum)
    for (int i = 0; i<n; i++) {
        tmp = a[i] + b[i];
        sum += tmp;
    }
    return sum;
}
```

example: perform addition
reduction on arrays a and b

Further Clause for simd Construct

- `safelen (length)`
 - maximum number of loop iterations that can be processed concurrently without breaking a dependence
- `simdlen (length)`
 - preferred number of loop iterations to be executed concurrently
 - must be less or equal to `safelen` with present
- `linear (list[:linear-step])`
 - the value of a variable is a linear function of the iteration number, i.e.
$$x_i = x_0 + i * \text{linear-step}$$
- `aligned (list[:alignment])`
 - the memory location of each element in the list is aligned to the number of bytes specified in the optional alignment argument
 - if alignment argument is missing, default alignment is assumed

Further Clause for simd Construct (2)

- collapse (n)
 - perform loop fusion, i.e. specifies how many loops are associated with the construct
 - if more than one loop is associated with simd clause, all associated loops are collapsed into one larger iteration space that is then executed with SIMD instructions

```
void work( float*b, int n, int m ){
    int i;
    #pragma omp simd safelen(16)
    for (i = m; i<n; i++) {
        b[i] = b[i-m] - 1.0f;
    }
}
```

Loop simd Construct

- **Combine work sharing with vectorization**
 - distribute iterations of loop across the threads in a team
 - execute each loop chunk with SIMD instructions
- `#pragma omp for simd [clauses]`
for-loops

```
float sprod(float *a, float *b, int n)
{
    float sum = 0.0;
    #pragma omp for simd reduction(+:sum)
    for (int i = 0; i<n; i++) {
        sum += a[i] * b[i];
    }
    return sum;
}
```

example: scalar (dot) product

Function Vectorization

- Only simple functions are automatically vectorized
 - hence, function calls in OpenMP SIMD loops frequently prevent vectorization
- OpenMP allows to declare functions that can be safely called from a SIMD-parallel loop
- `#pragma omp declare simd [clauses] function-definition-or-declaration`
- Additional clause `uniform` to specify function arguments that are constant for all SIMD lanes

```
#pragma omp declare simd
float min(float a, float b) {
    return a < b ? a : b;
}

#pragma omp declare simd
float distsq(float x, float y) {
    return (x - y) * (x - y);
}

void example() {
    #pragma omp parallel for simd
    for (i=0; i<N; i++) {
        d[i] = min(distsq(a[i], b[i]), c[i]);
    }
}
```

Summary Vectorization

- Exploiting vectorization is key for performance on modern CPUs
- The cheapest way to vectorization is using numerical libraries and autovectorization
- OpenMP simd constructs provide a portable way to add explicit vectorization
- Vectorization is a deep topic, we have only scratched the surface here

Roofline Model

Roofline Model

- Application performance on CPUs is limited by
 - memory bandwidth, or (memory bound)
 - arithmetic performance (compute bound)
- Roofline model
 - analytical model of fundamental performance limits
 - describes an upper bound on the achievable performance based on the **operational intensity (also arithmetic intensity)**, which is the number of operations performed per data read from DRAM (FLOPS/byte)
 - Williams, Waterman and Patterson: “*Roofline: An Insightful Visual Performance Model for Multicore Architectures*”, Communications of the ACM, 4(52) 2009, <http://dx.doi.org/10.1145/1498765.1498785>
- Basic idea

$$GFLOP / s = \min \left\{ \begin{array}{l} \text{Peak Computational Performance} \\ \text{Memory Bandwidth} * \text{Operational Intensity} \end{array} \right.$$

Roofline Model (2)

$$GFLOP / s = \min \left\{ \begin{array}{l} \text{Peak Computational Performance} \\ \text{Memory Bandwidth} * \text{Operational Intensity} \end{array} \right.$$

Computational Peak performance

Peak FLOP = 2 x 2.6 x 6 x 16 =
499 single precision FLOP/s

- 2 sockets
- 2.6 GHz
- 6 core
- 16 single precision operation per SIMD instruction (8 MUL + 8 ADD with AVX)

Memory bandwidth

Peak Mem BW = 2 x 1.6 x 8 x 4 = 102.4 GB/s

- 2 sockets
- 1.6 GHz memory frequency
- 8 bytes per channel
- 4 memory channels per CPU

2-socket Oculus Node

Roofline Model (2)

Intel Sandy Bridge Microarchitecture:

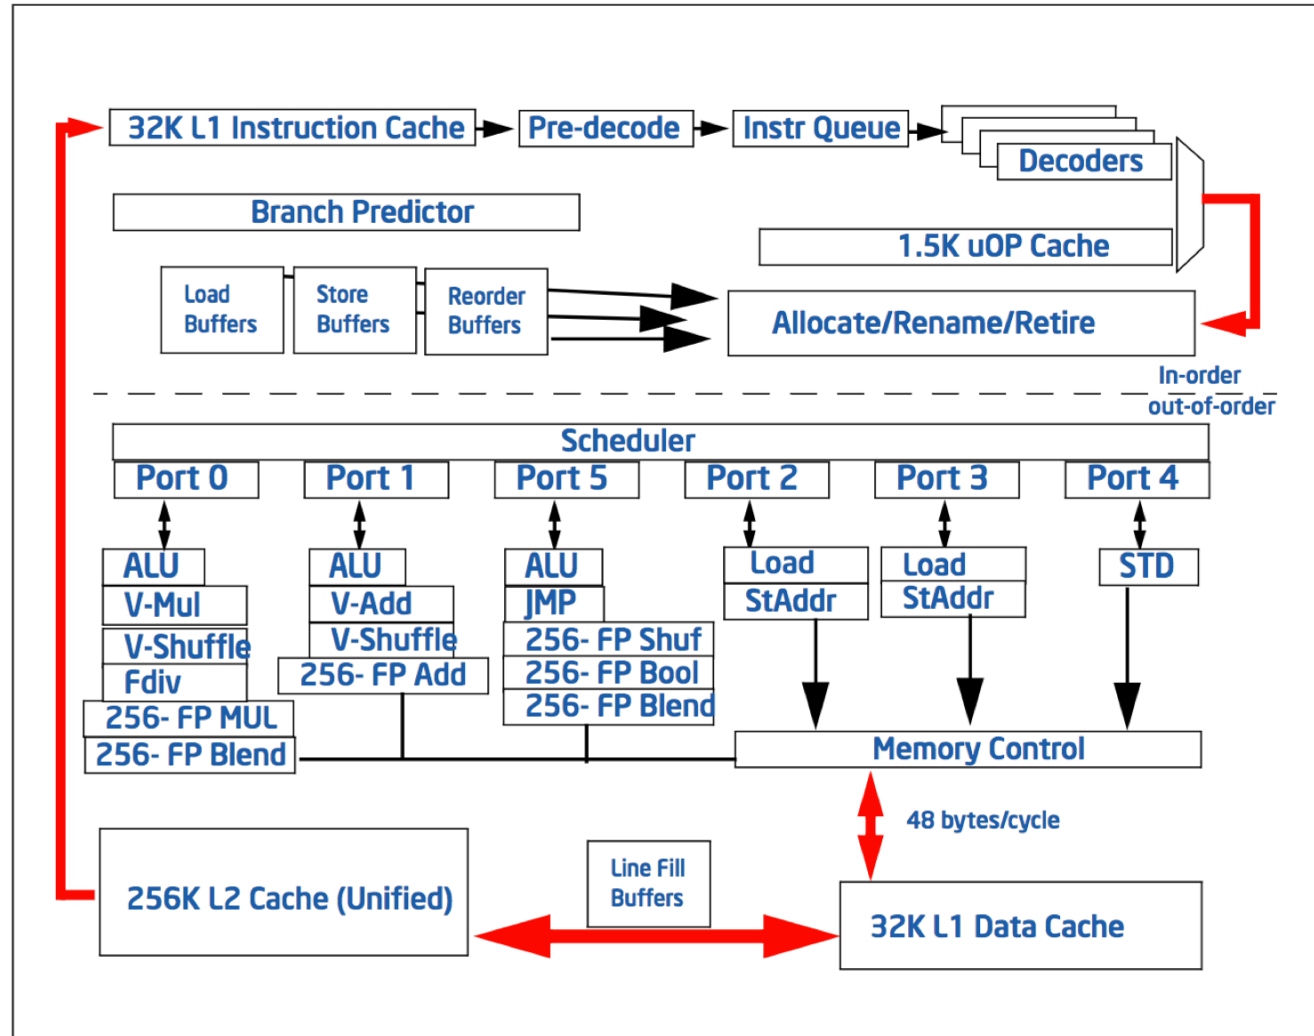
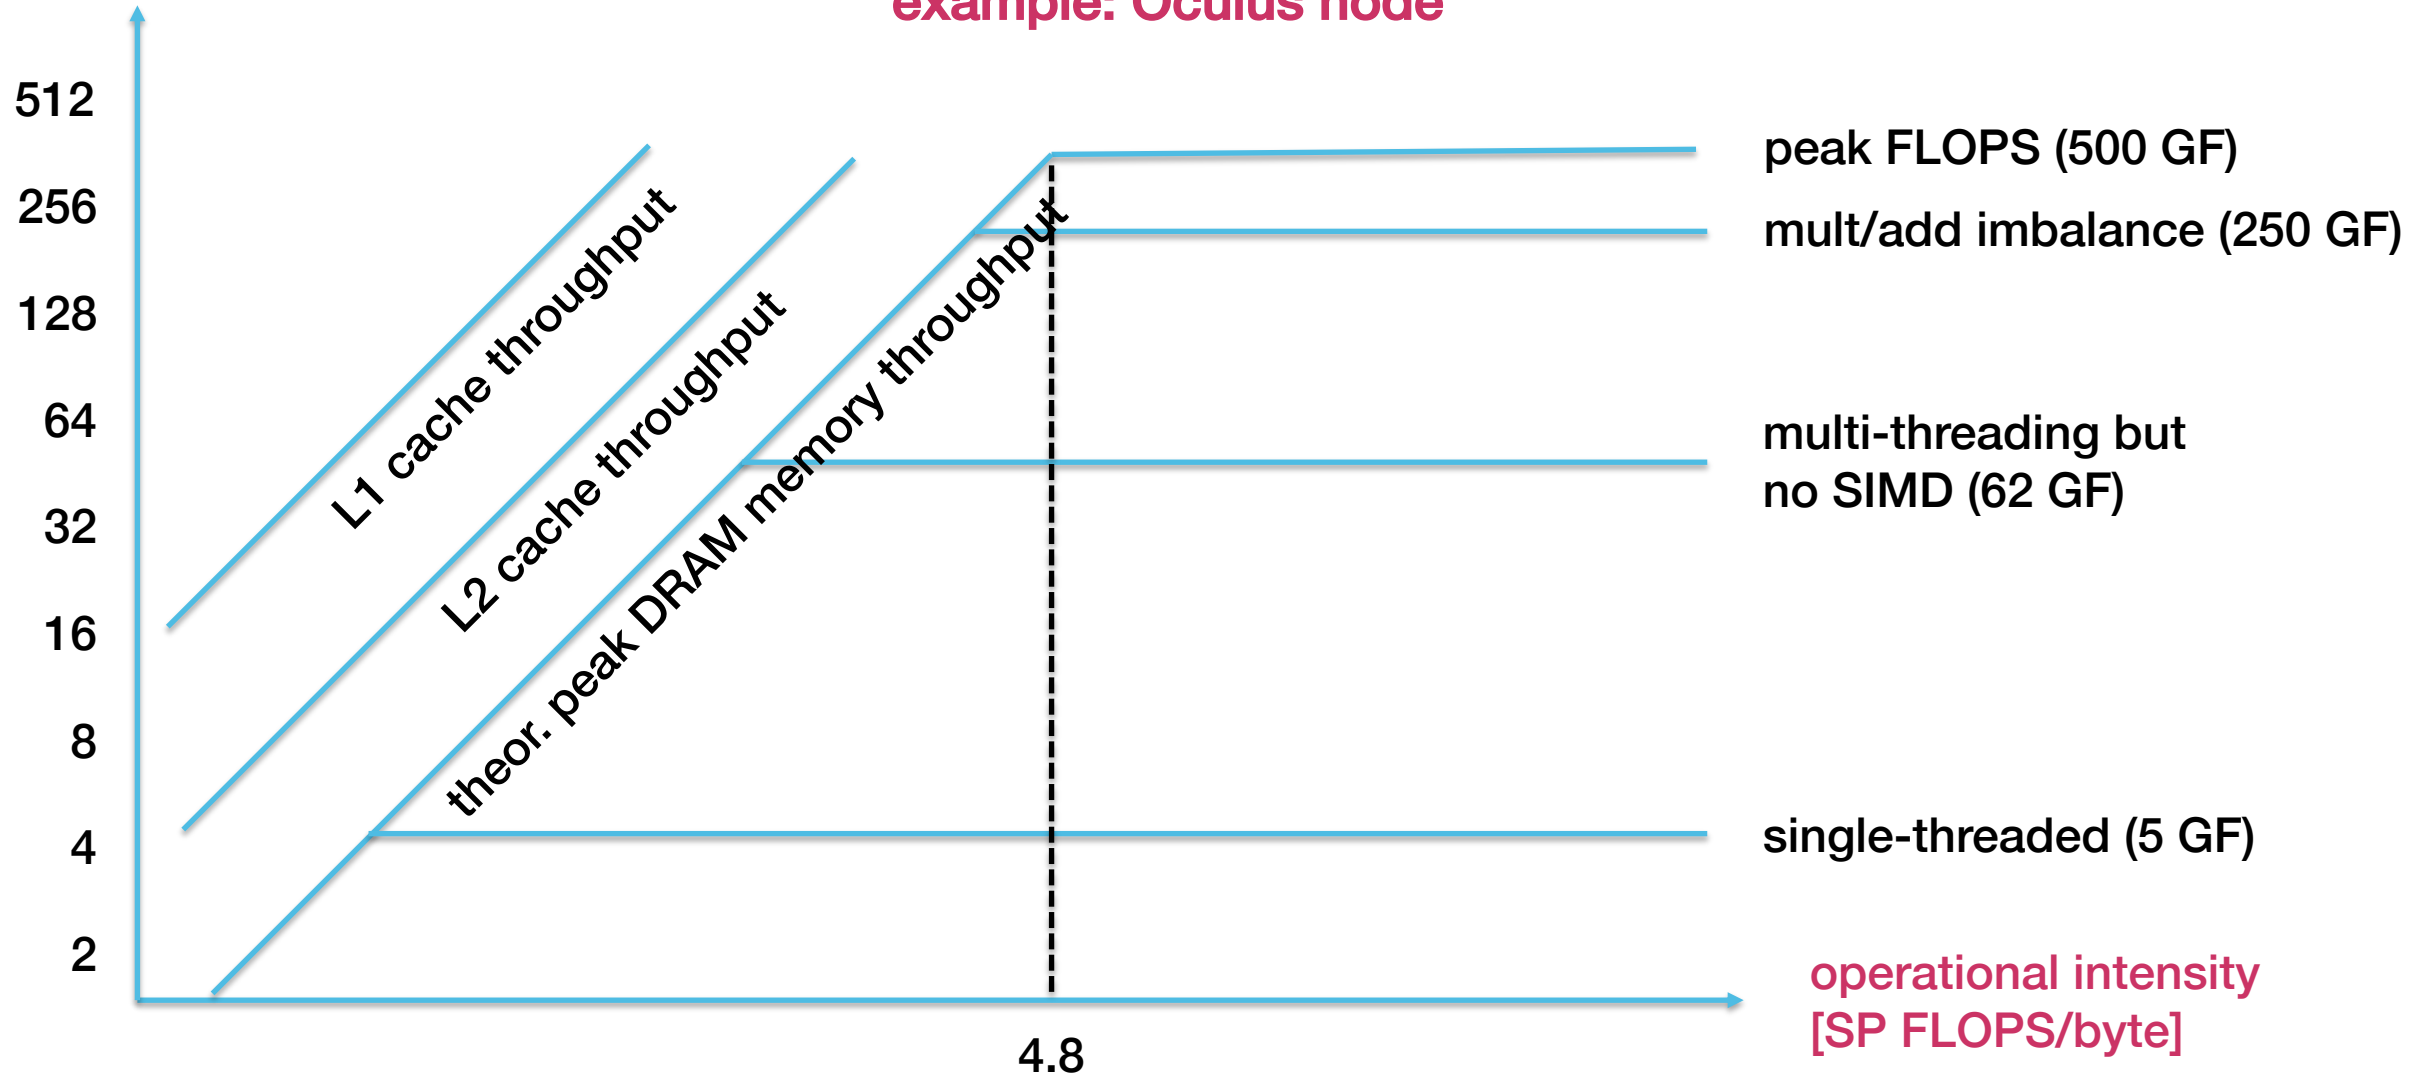


Figure 2-5. Intel Microarchitecture Code Name Sandy Bridge Pipeline Functionality

Roofline Model (3)

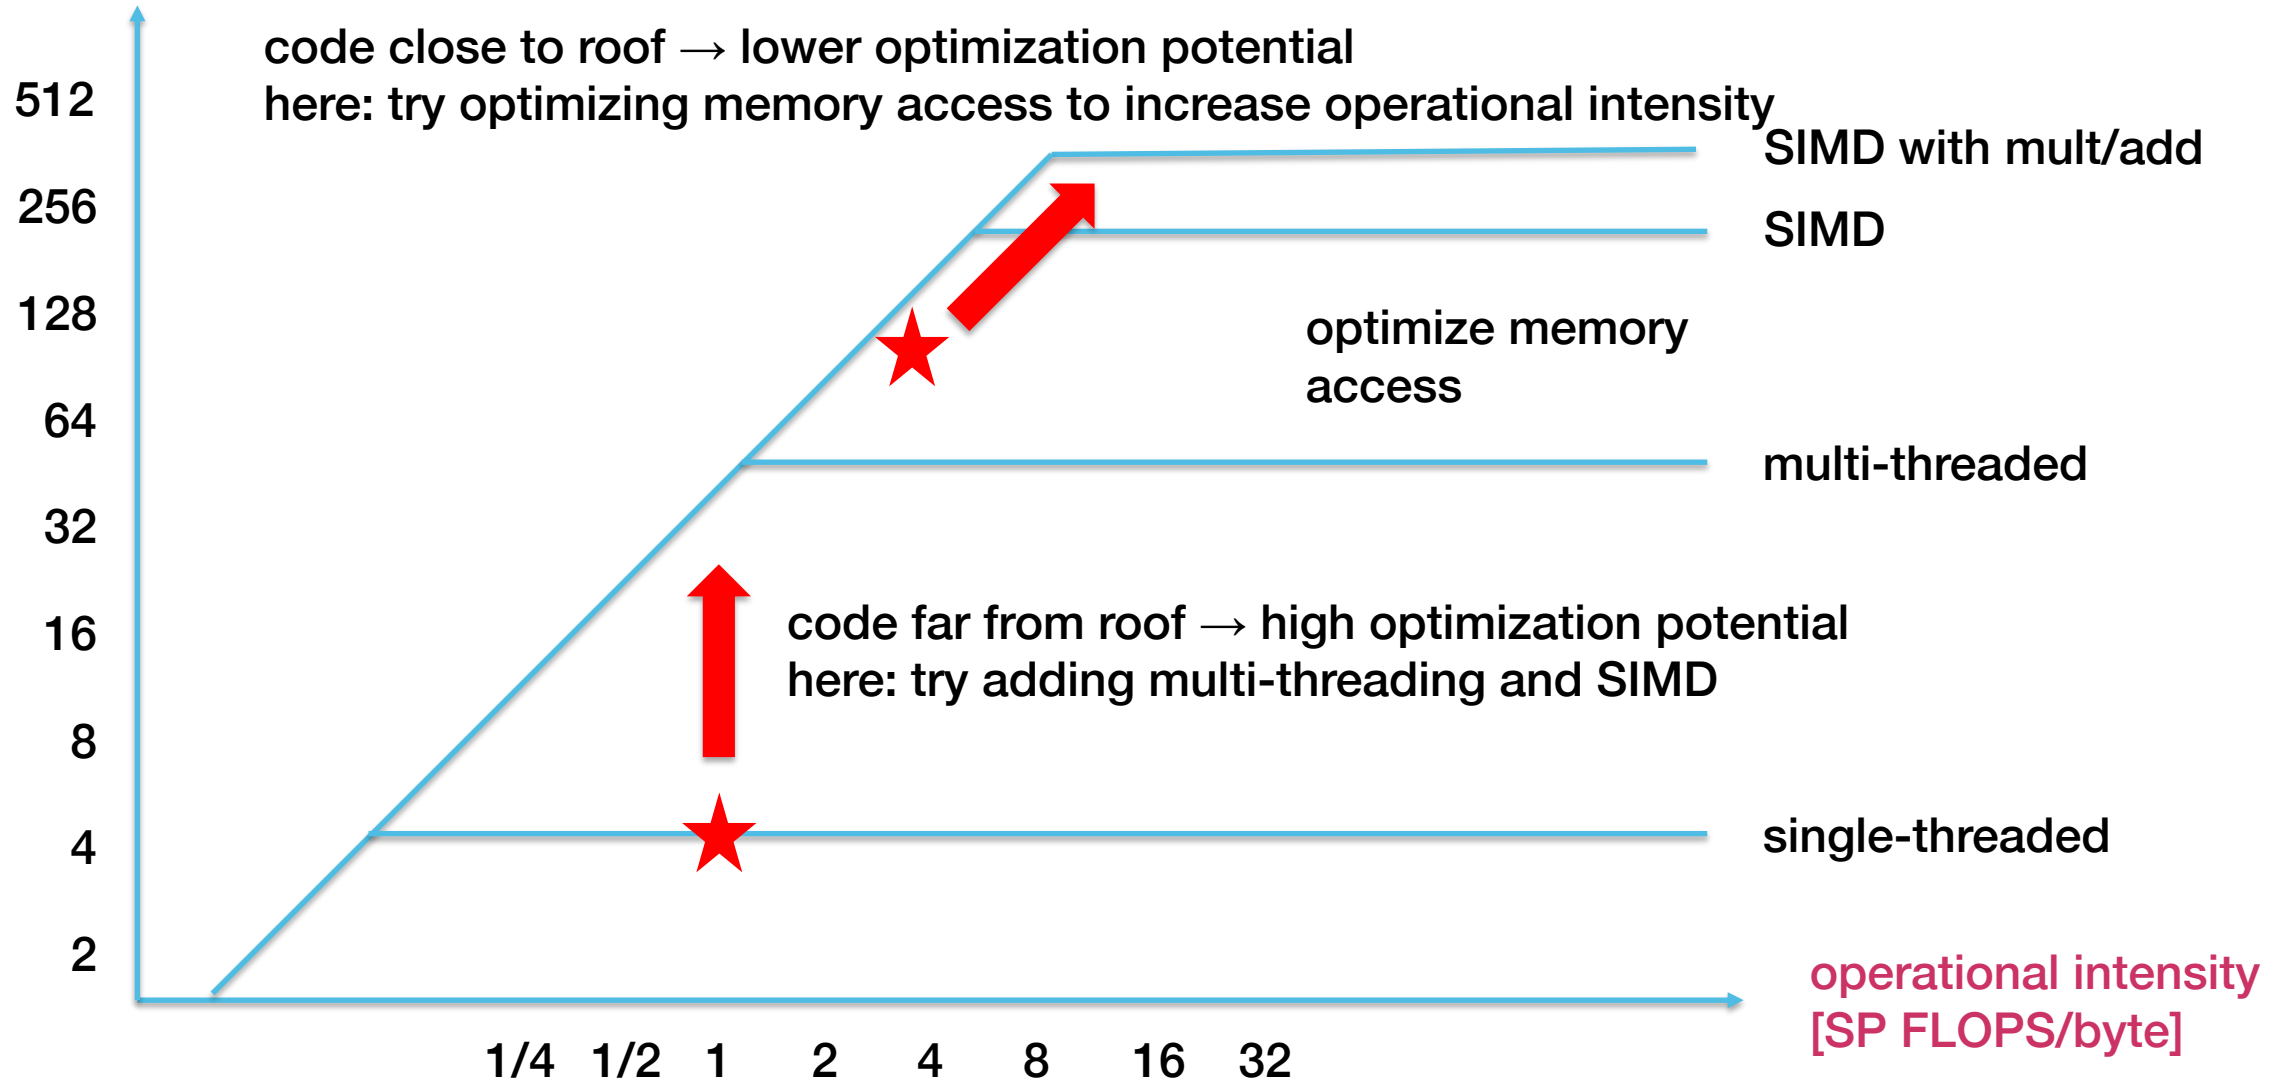
performance
limit [SP GFLOP/s]

example: Oculus node



How to Use Roofline Model for Optimization?

performance
limit [SP GFLOP/s]



Example 1: STREAM Triad

- Current HPC machines require about 5-10 FLOPS / byte to reach peak arithmetic capabilities
 - hence, we need about 40-80 operations per double-precision value read from DRAM to reach peak compute capability
 - very hard to achieve for real codes
 - the compute to memory balance is unlikely to improve in future (technological and economical tradeoffs in computer system design)
- Example: STREAM Triad
 - part of STREAM memory system benchmark
 - 2 FLOPs per iteration
 - transfers 24 bytes per iteration (read x[i], read y[i], write z[i])
 - $OI = 2 / 24 = 0.083$ FLOP/byte → clearly memory bound

```
#pragma omp parallel for
for(int i=0; i<N, i++) {
    z[i] = x[i] + alpha*y[i];
}
```

Example 2: 5-point Stencil

- **5-point constant coefficient stencil**
 - applies same operation to all points in 2D array
 - 5 FLOP per evaluation
 - 6 memory transfers (5 read, 1 write) per point
- **Naïve implementation**
 - do not consider any caching, read all data from DRAM
 - $AI = 5 / (6 * 8) = 0.104$ FLOP/byte → memory bound
- **Consider Caching**
 - if the cache can store two rows of the array all data access except for 1 read and 1 write per point can be
 - $AI = 5 / (2 * 8) = 0.3125$ FLOP/byte → still memory bound

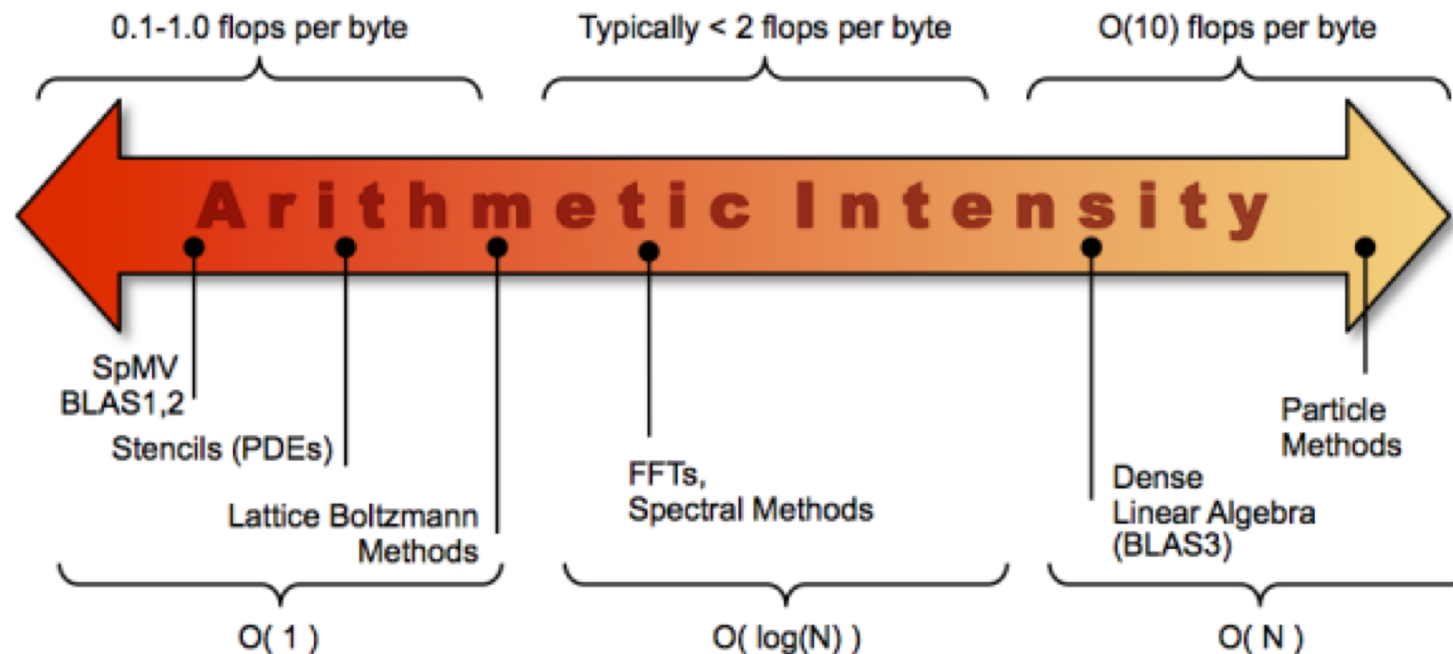
```
#pragma omp parallel for
for (int t=0; t<tmax; t++) {
    for(int y=1; y<ydim+1; y++) {
        for(int x=1; x<xdim+1; x++) {

            int idx = x + y*ystride;
            new[idx] = - 3.0*old[idx]
                    + old[idx-1]
                    + old[idx+1]
                    + old[idx-ystride]
                    + old[idx+ystride];

        }
    }
}
```

Operational Intensities

- Determining operational intensities for non-trivial applications is difficult
 - performance analysis tools can determine OI from profiling/performance counter data, e.g. Intel vTune Amplifier



Case study:
**Performance optimization of n-body
solver**

Performance Optimization: n-body

- Optimize the n-body application used earlier in this course

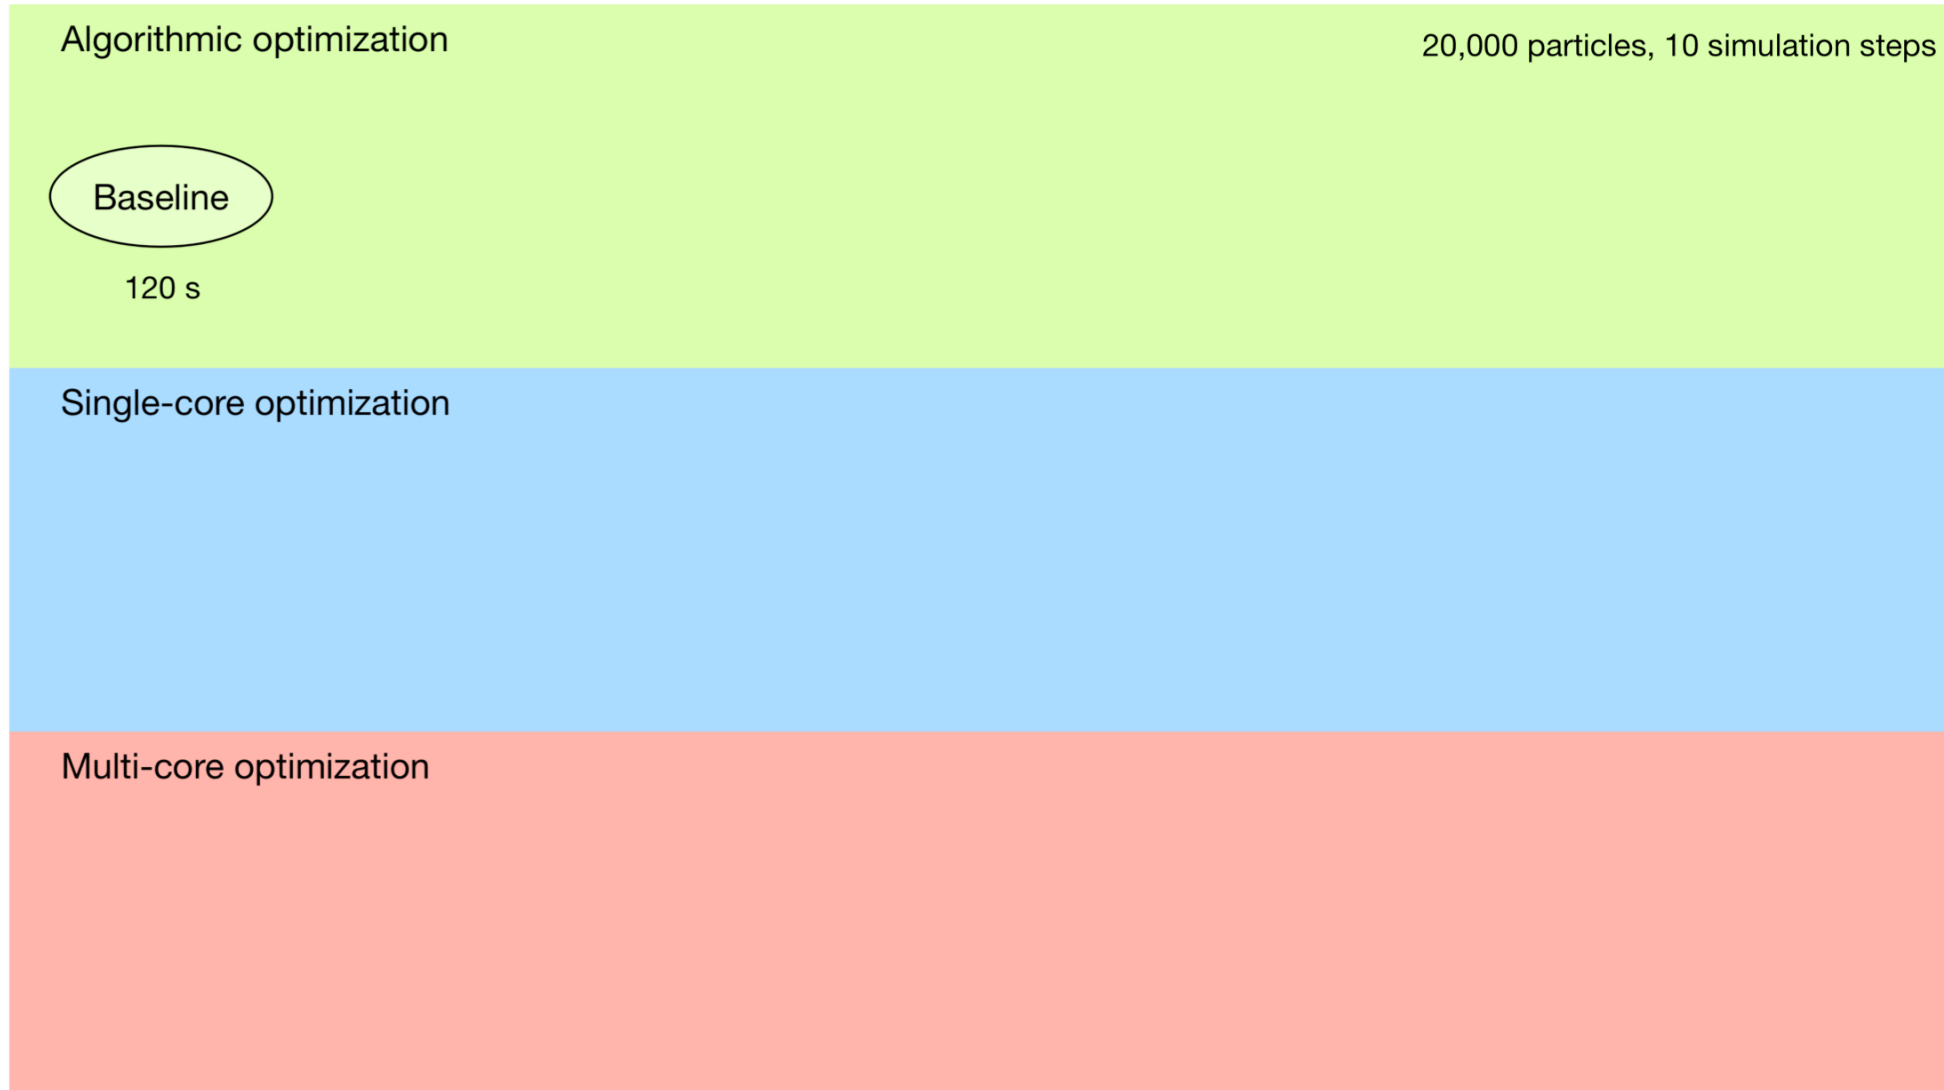
- Main application loop:

```
for (step = 1; step <= n_steps; step++) {  
    t = step*delta_t;  
    for (part = 0; part < n; part++)  
        Compute_force(part, forces, curr, n);  
    for (part = 0; part < n; part++)  
        Update_part(part, forces, curr, n, delta_t);  
}
```

- Compute_force(...):

```
for (k = 0; k < n; k++) {  
    if (k != part) {  
        f_part_k[X] = curr[part].s[X] - curr[k].s[X];  
        f_part_k[Y] = curr[part].s[Y] - curr[k].s[Y];  
        len = sqrt(f_part_k[X]*f_part_k[X] + f_part_k[Y]*f_part_k[Y]);  
        len_3 = len*len*len;  
        mg = -G*curr[part].m*curr[k].m;  
        fact = mg/len_3;  
        f_part_k[X] *= fact;  
        f_part_k[Y] *= fact;  
        forces[part][X] += f_part_k[X];  
        forces[part][Y] += f_part_k[Y];  
    }  
}
```

Performance Optimization: n-body



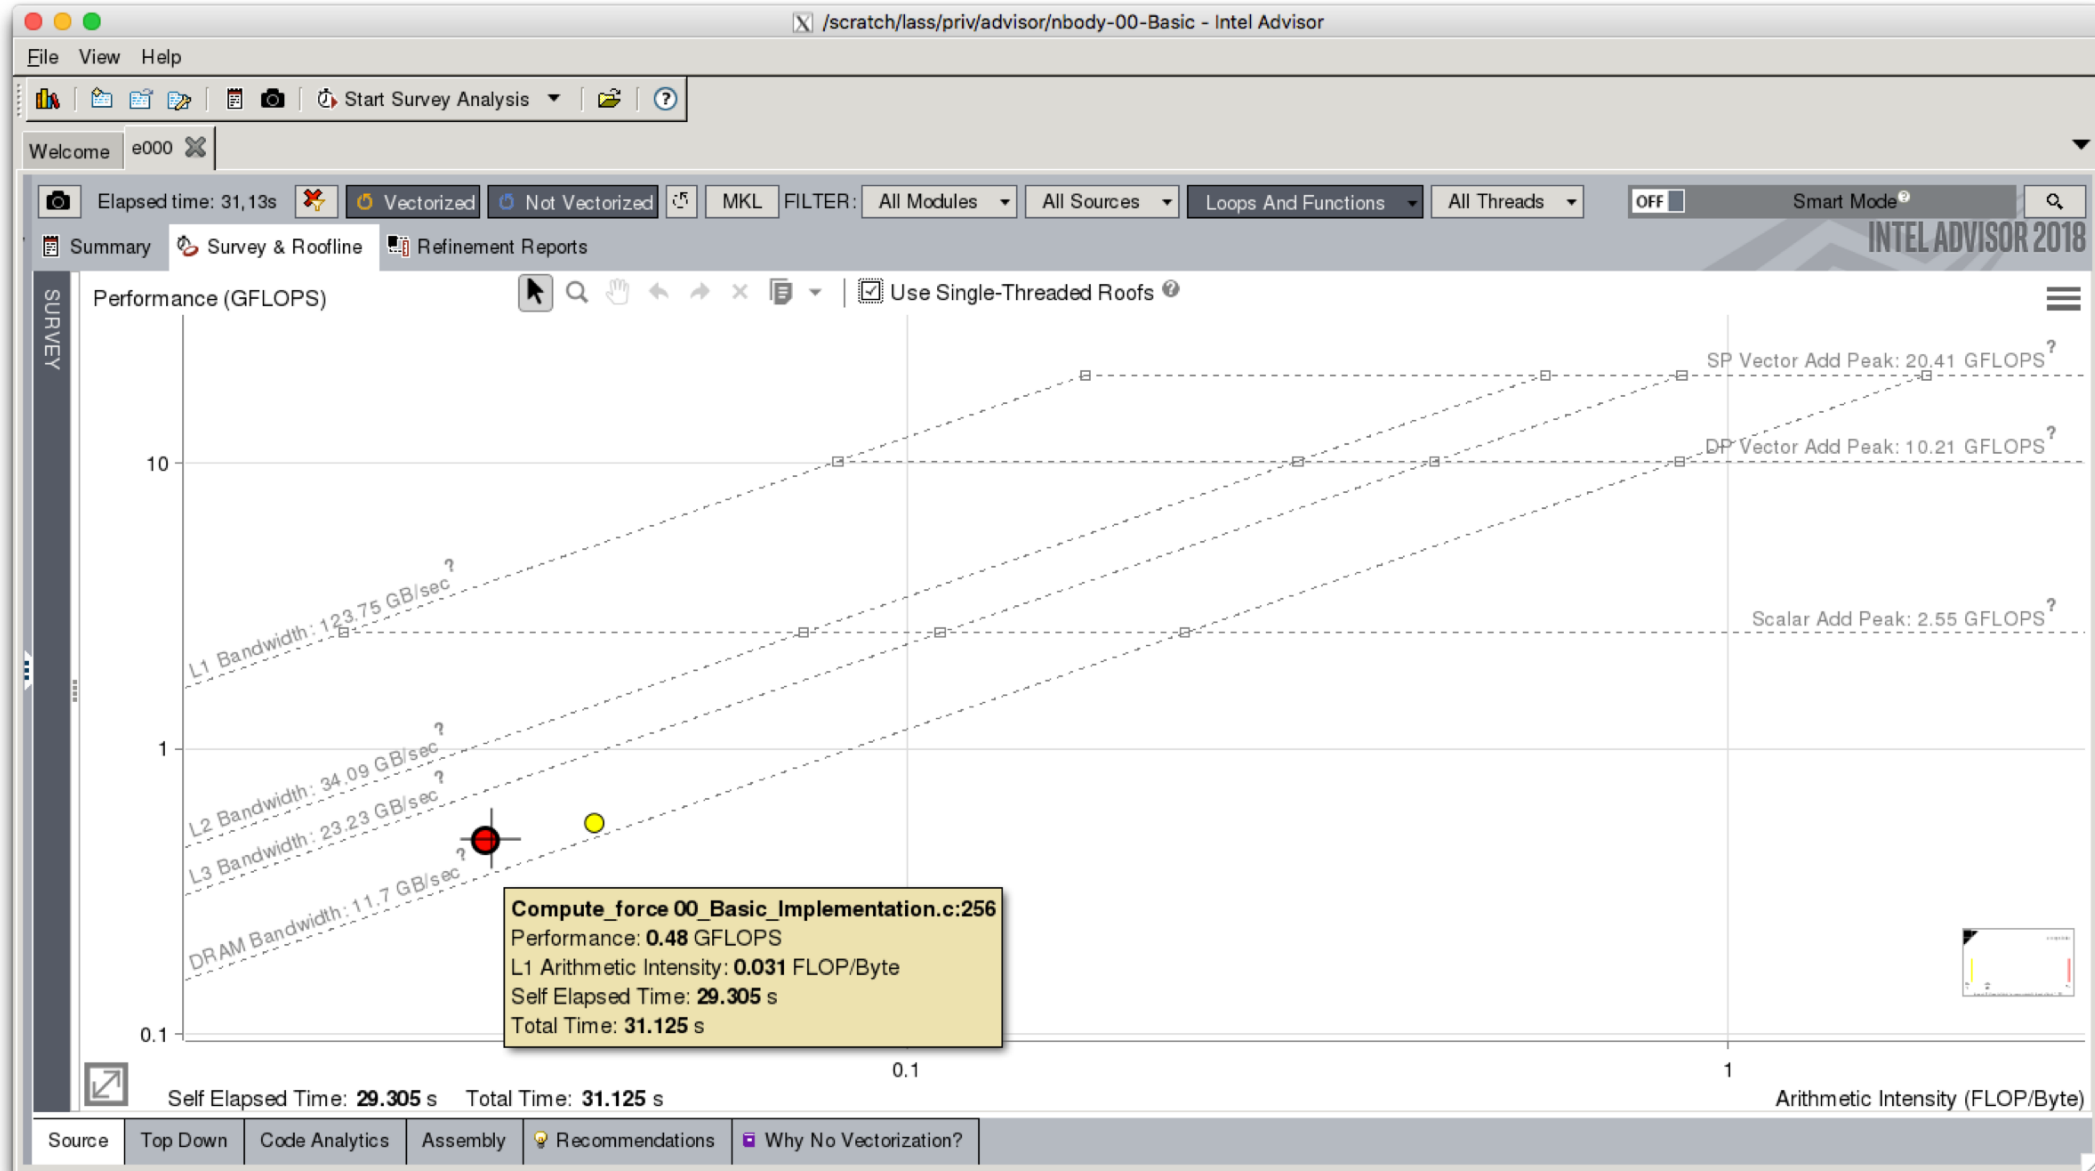
Performance Optimization: n-body

- Intel Advisor: Tool for prototyping and optimization of
 - Vectorization
 - Memory access patterns
 - Multi threading
- Since 2017, Advisor plots a roofline model
- Starting Advisor on our systems:

```
module load ps_xe_2018  
advixe-gui
```

Demo: Intel Advisor & n-body code

Performance Optimization: n-body

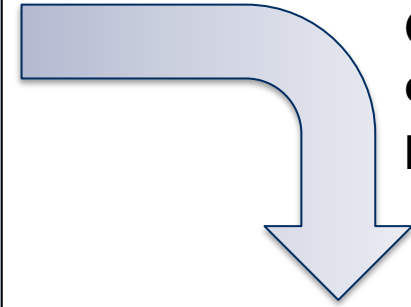


Baseline: non-optimized code

Performance Optimization: n-body

```
for (k = 0; k < n; k++) {  
    if (k != part) {  
        f_part_k[X] = curr[part].s[X] - curr[k].s[X];  
        f_part_k[Y] = curr[part].s[Y] - curr[k].s[Y];  
        len = sqrt(f_part_k[X]*f_part_k[X] + f_part_k[Y]*f_part_k[Y]);  
        len_3 = len*len*len;  
        mg = -G*curr[part].m*curr[k].m;  
        fact = mg/len_3;  
        f_part_k[X] *= fact;  
        f_part_k[Y] *= fact;  
        forces[part][X] += f_part_k[X];  
        forces[part][Y] += f_part_k[Y];  
    }  
}
```

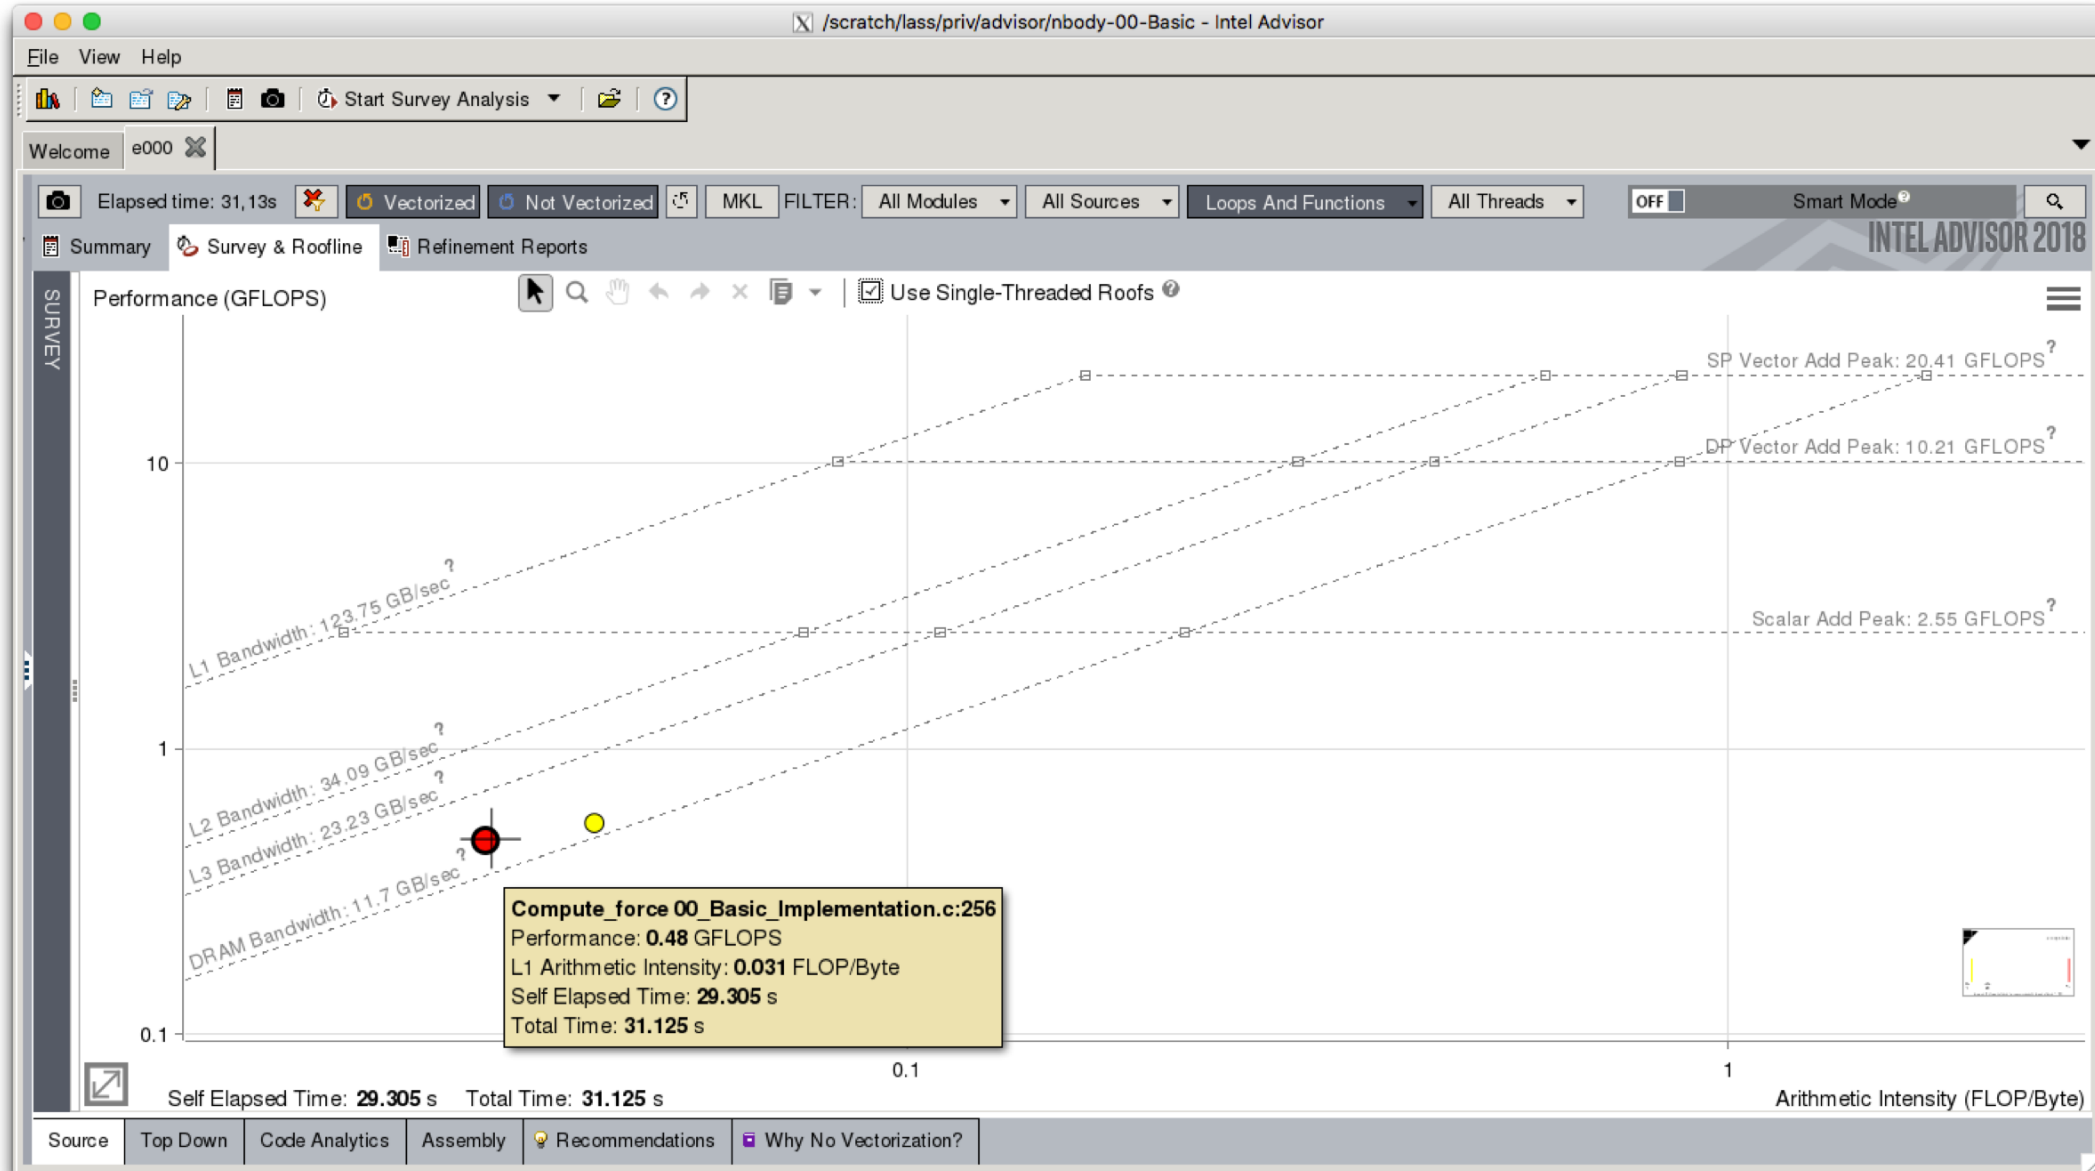
In which direction do we
move in the roofline model?



Compute forces only
once for each pair of
particles

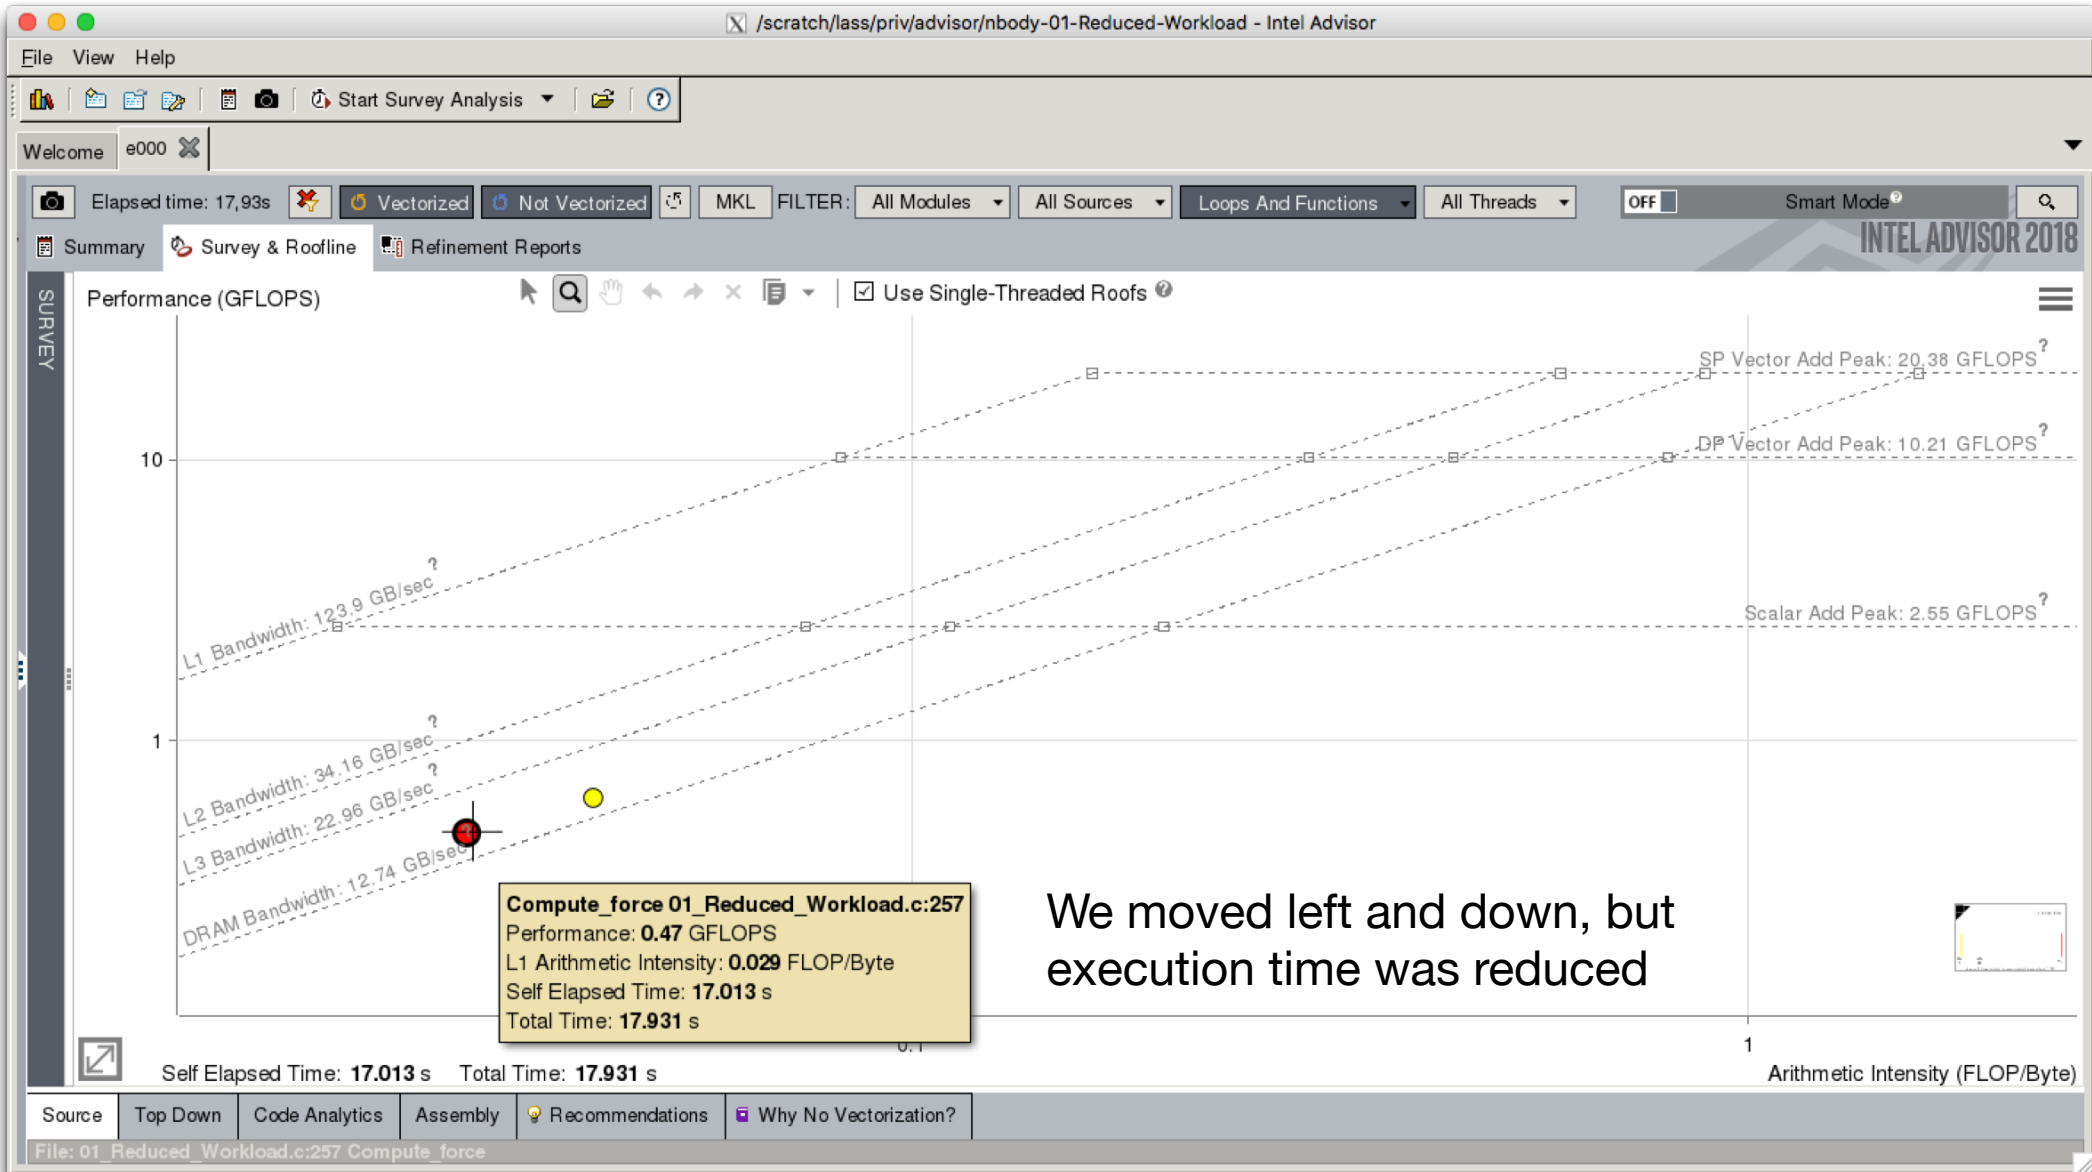
```
for (k = part+1; k < n; k++) {  
    f_part_k[X] = curr[part].s[X] - curr[k].s[X];  
    f_part_k[Y] = curr[part].s[Y] - curr[k].s[Y];  
    len = sqrt(f_part_k[X]*f_part_k[X] + f_part_k[Y]*f_part_k[Y]);  
    len_3 = len*len*len;  
    mg = -G*curr[part].m*curr[k].m;  
    fact = mg/len_3;  
    f_part_k[X] *= fact;  
    f_part_k[Y] *= fact;  
    forces[part][X] += f_part_k[X];  
    forces[part][Y] += f_part_k[Y];  
    forces[k][X] -= f_part_k[X];  
    forces[k][Y] -= f_part_k[Y];  
}
```

Performance Optimization: n-body



Baseline: non-optimized code

Performance Optimization: n-body

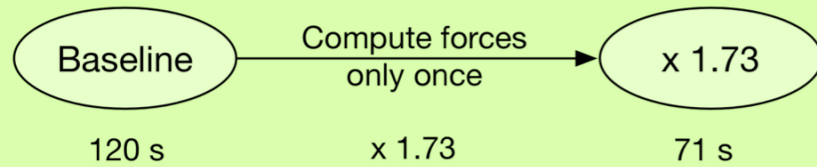


Algorithmically optimized code

Performance Optimization: n-body

Algorithmic optimization

20,000 particles, 10 simulation steps



Single-core optimization

Multi-core optimization

Performance Optimization: n-body

- Enable compiler optimizations and auto vectorization:

```
icc -O3 -xHOST -qopt-report -qopt-report-phase=vec -o nbody nbody.c
```

- Vectorization report (excerpt):

```
Begin optimization report for: Compute_force(int, vect_t *, struct particle_s *, int)
```

```
Report from: Vector optimizations [vec]
```

```
LOOP BEGIN at nbody.c(257,4)
```

```
remark #15300: LOOP WAS VECTORIZED
```

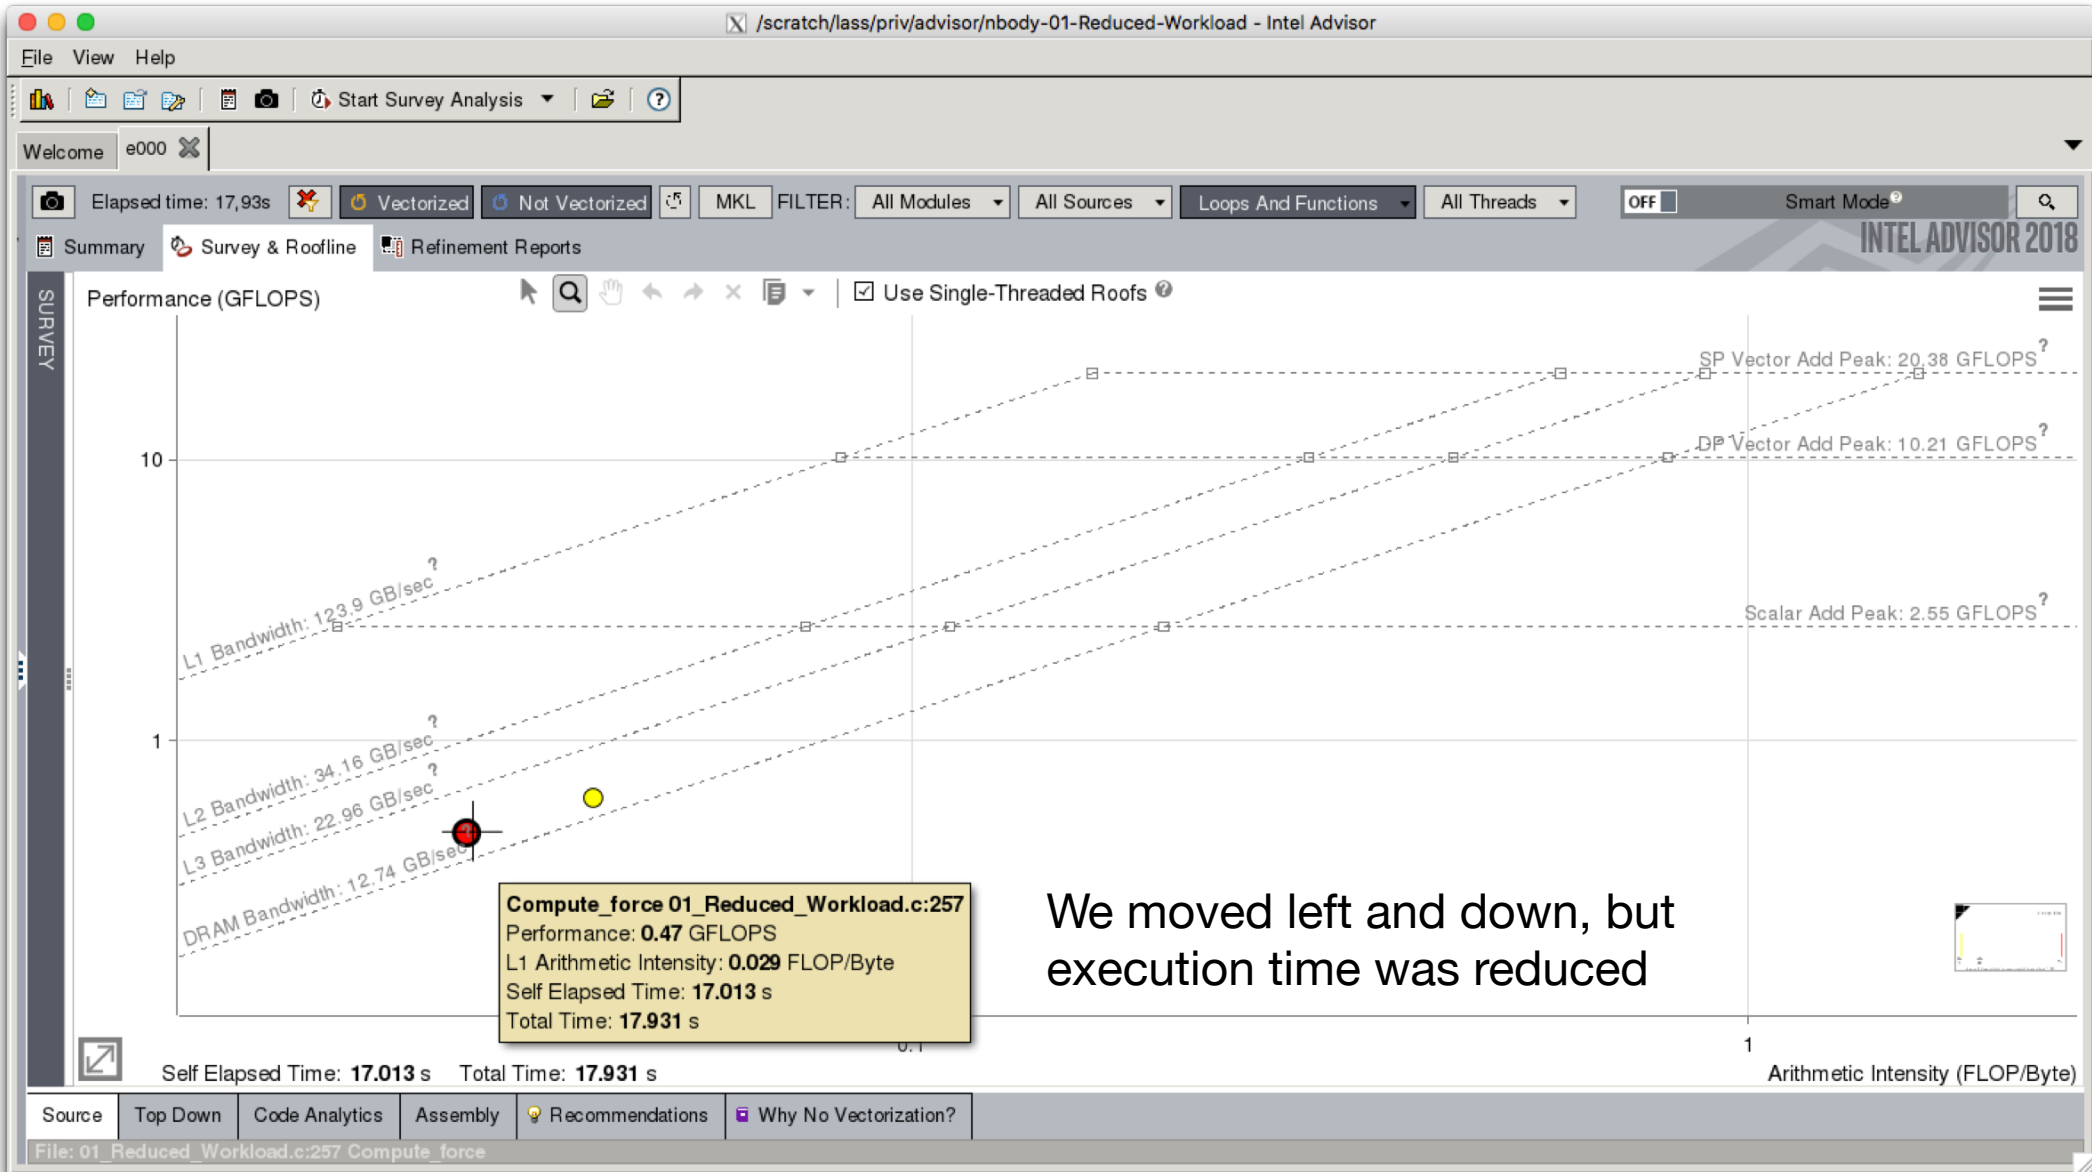
```
LOOP END
```

```
LOOP BEGIN at nbody.c(257,4)
```

```
<Remainder loop for vectorization>
```

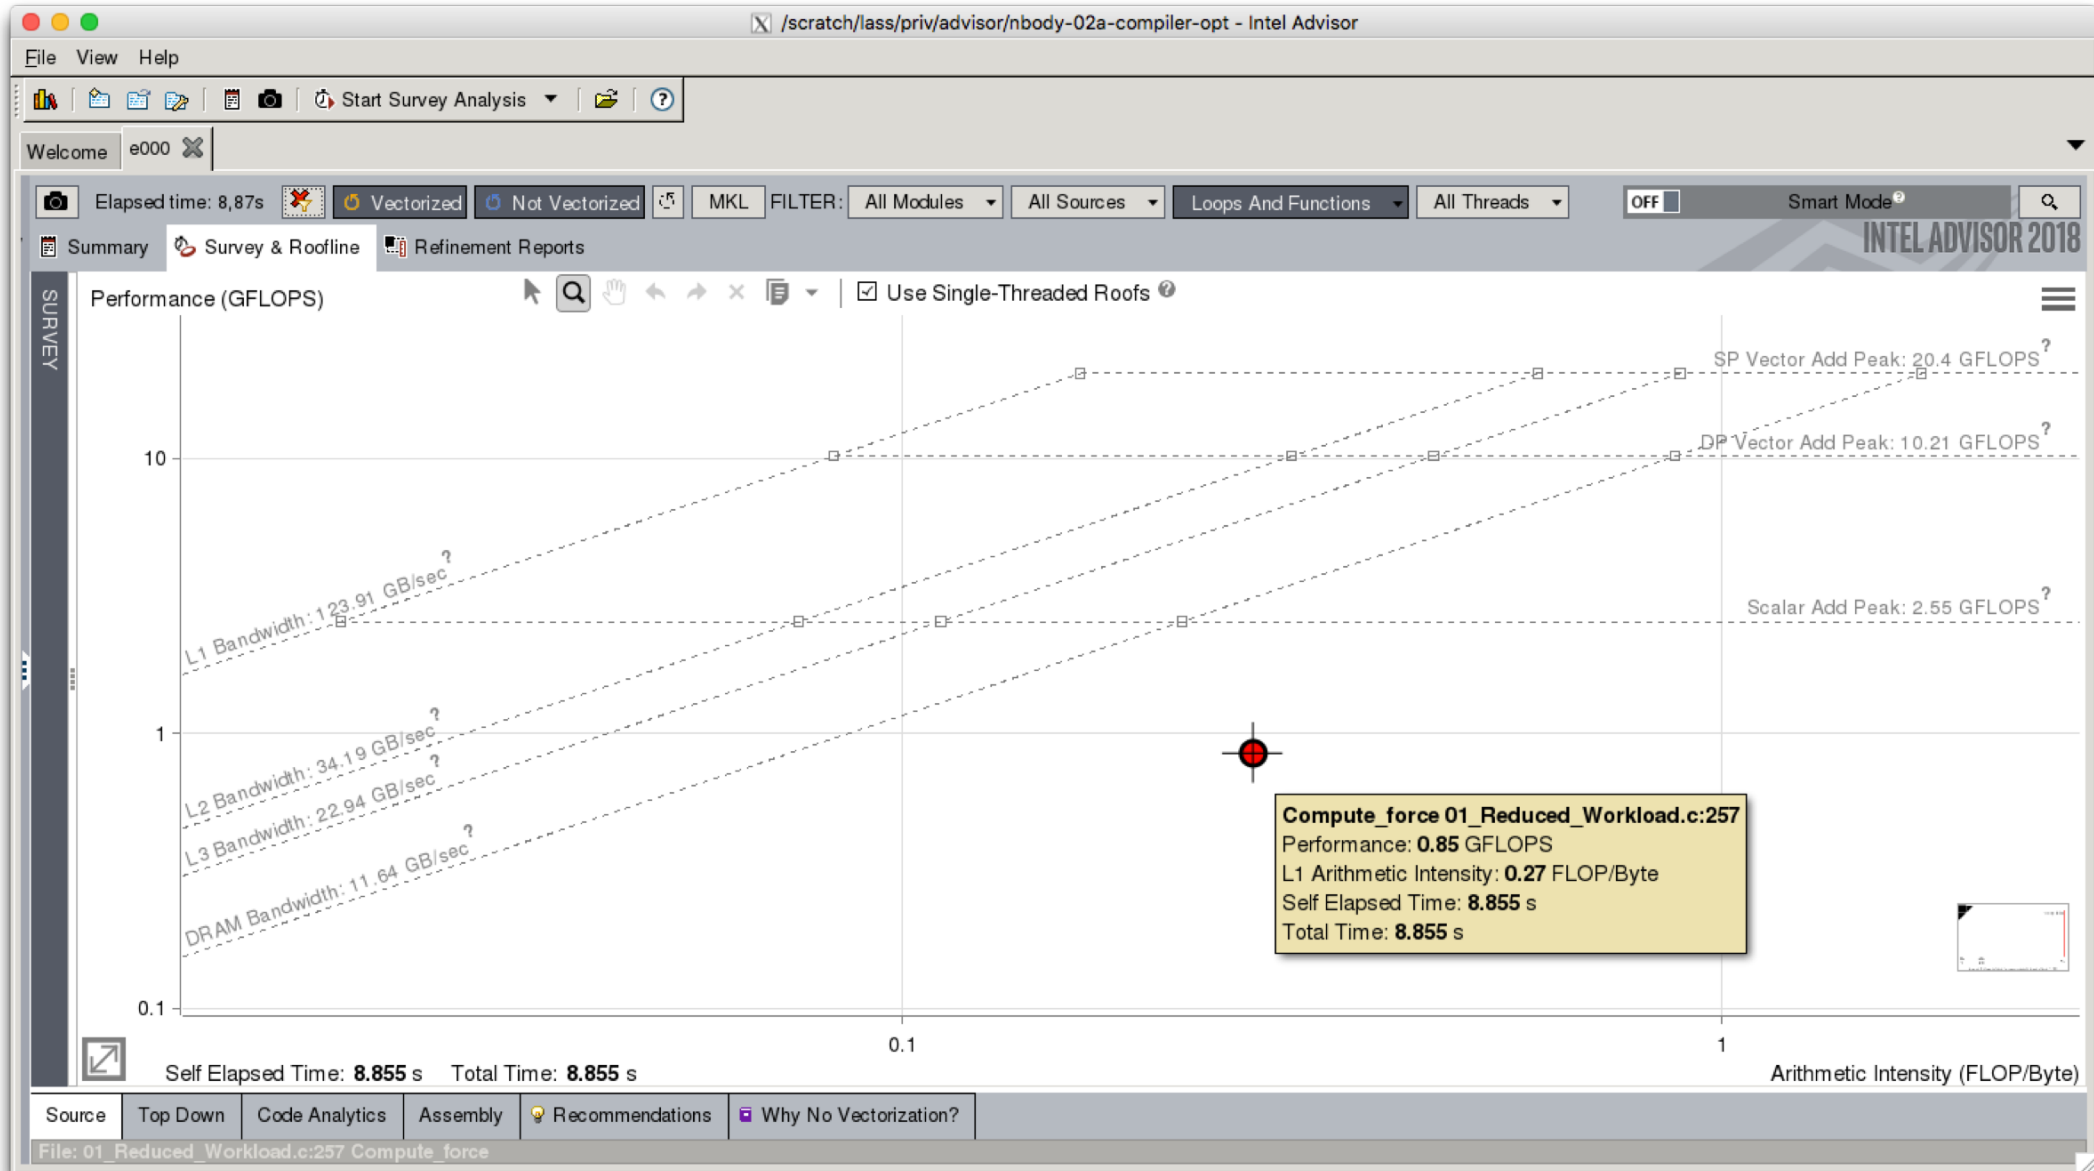
```
LOOP END
```


Performance Optimization: n-body



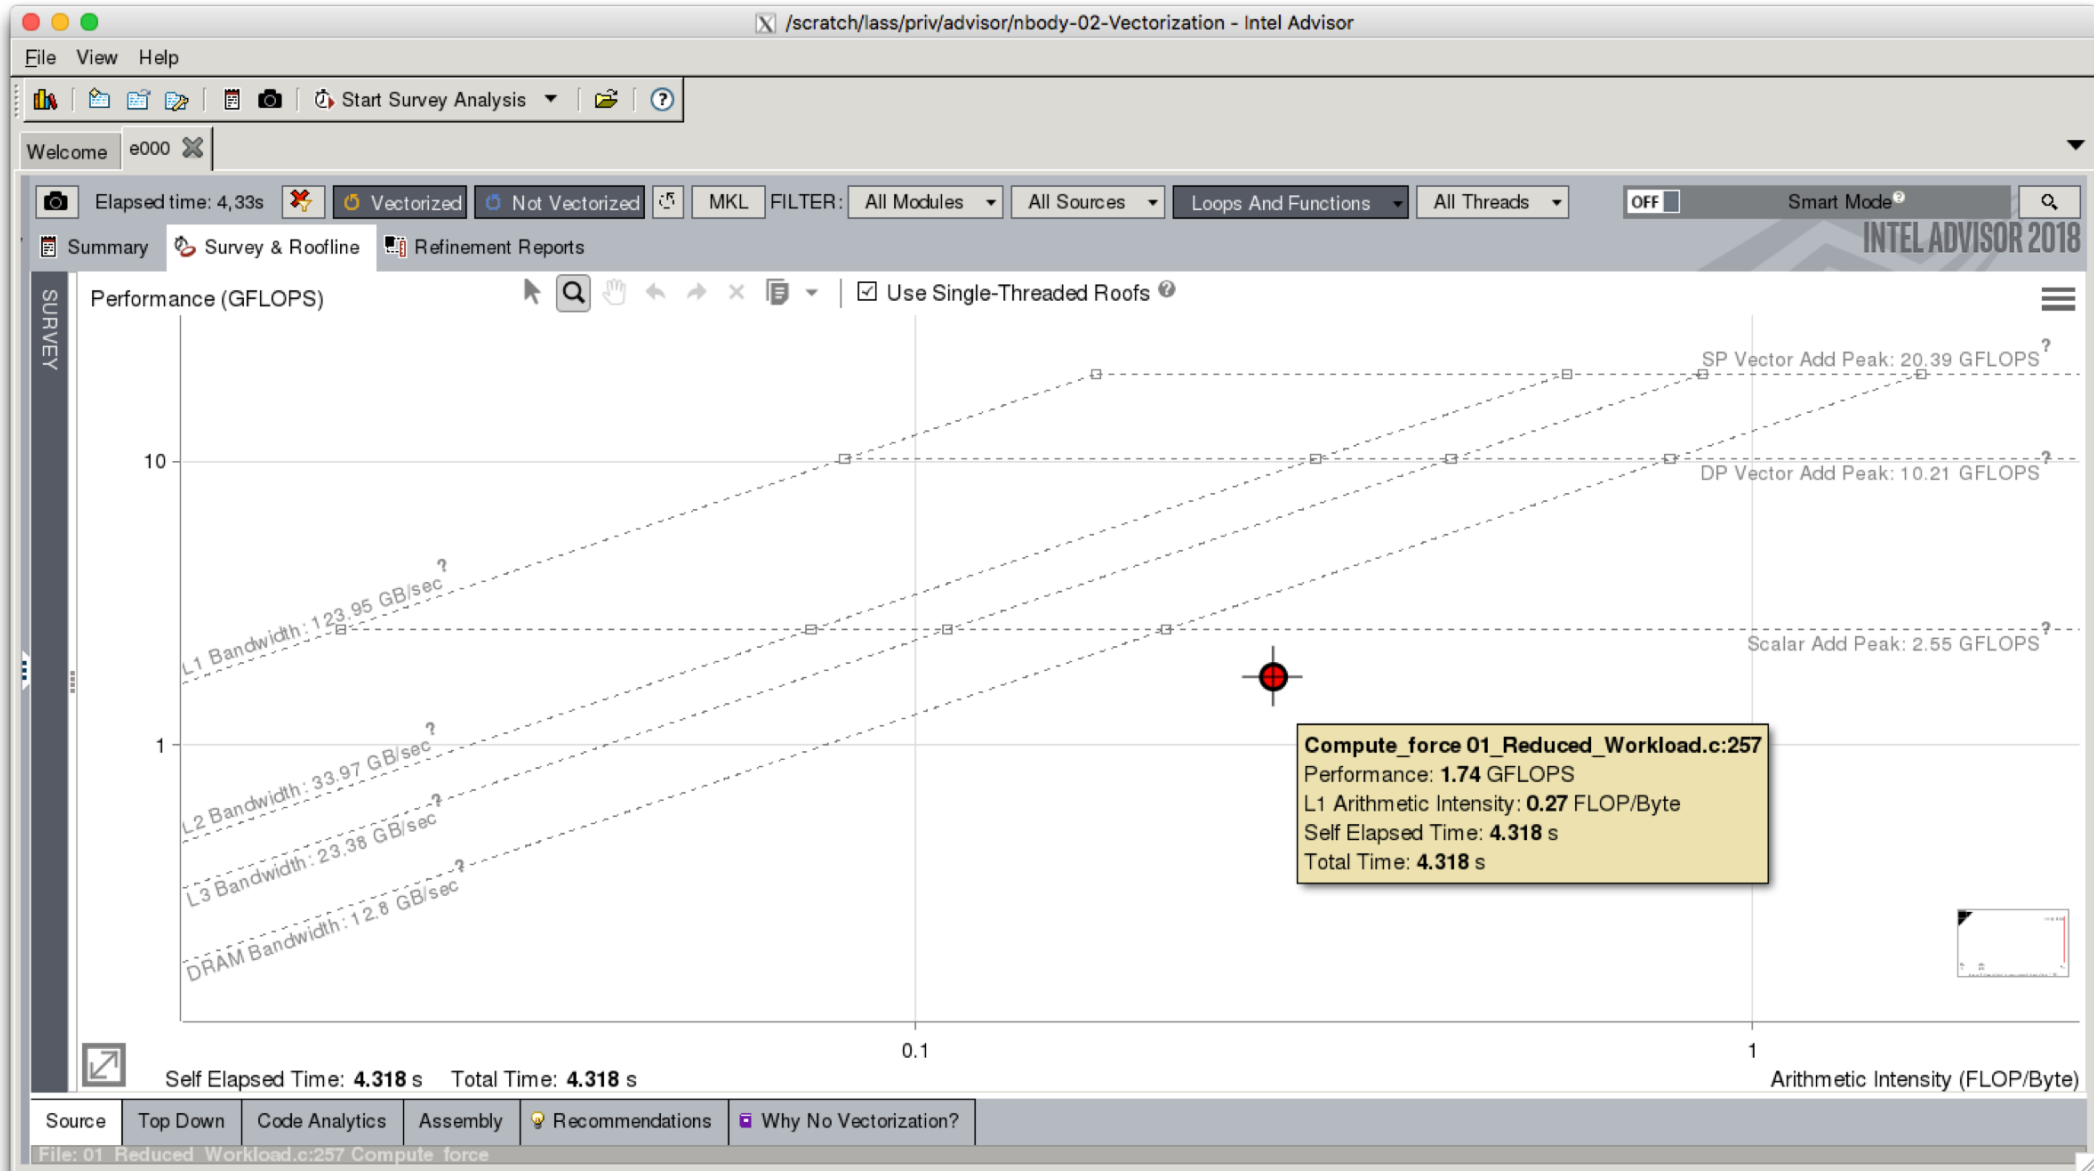
Algorithmically optimized code

Performance Optimization: n-body



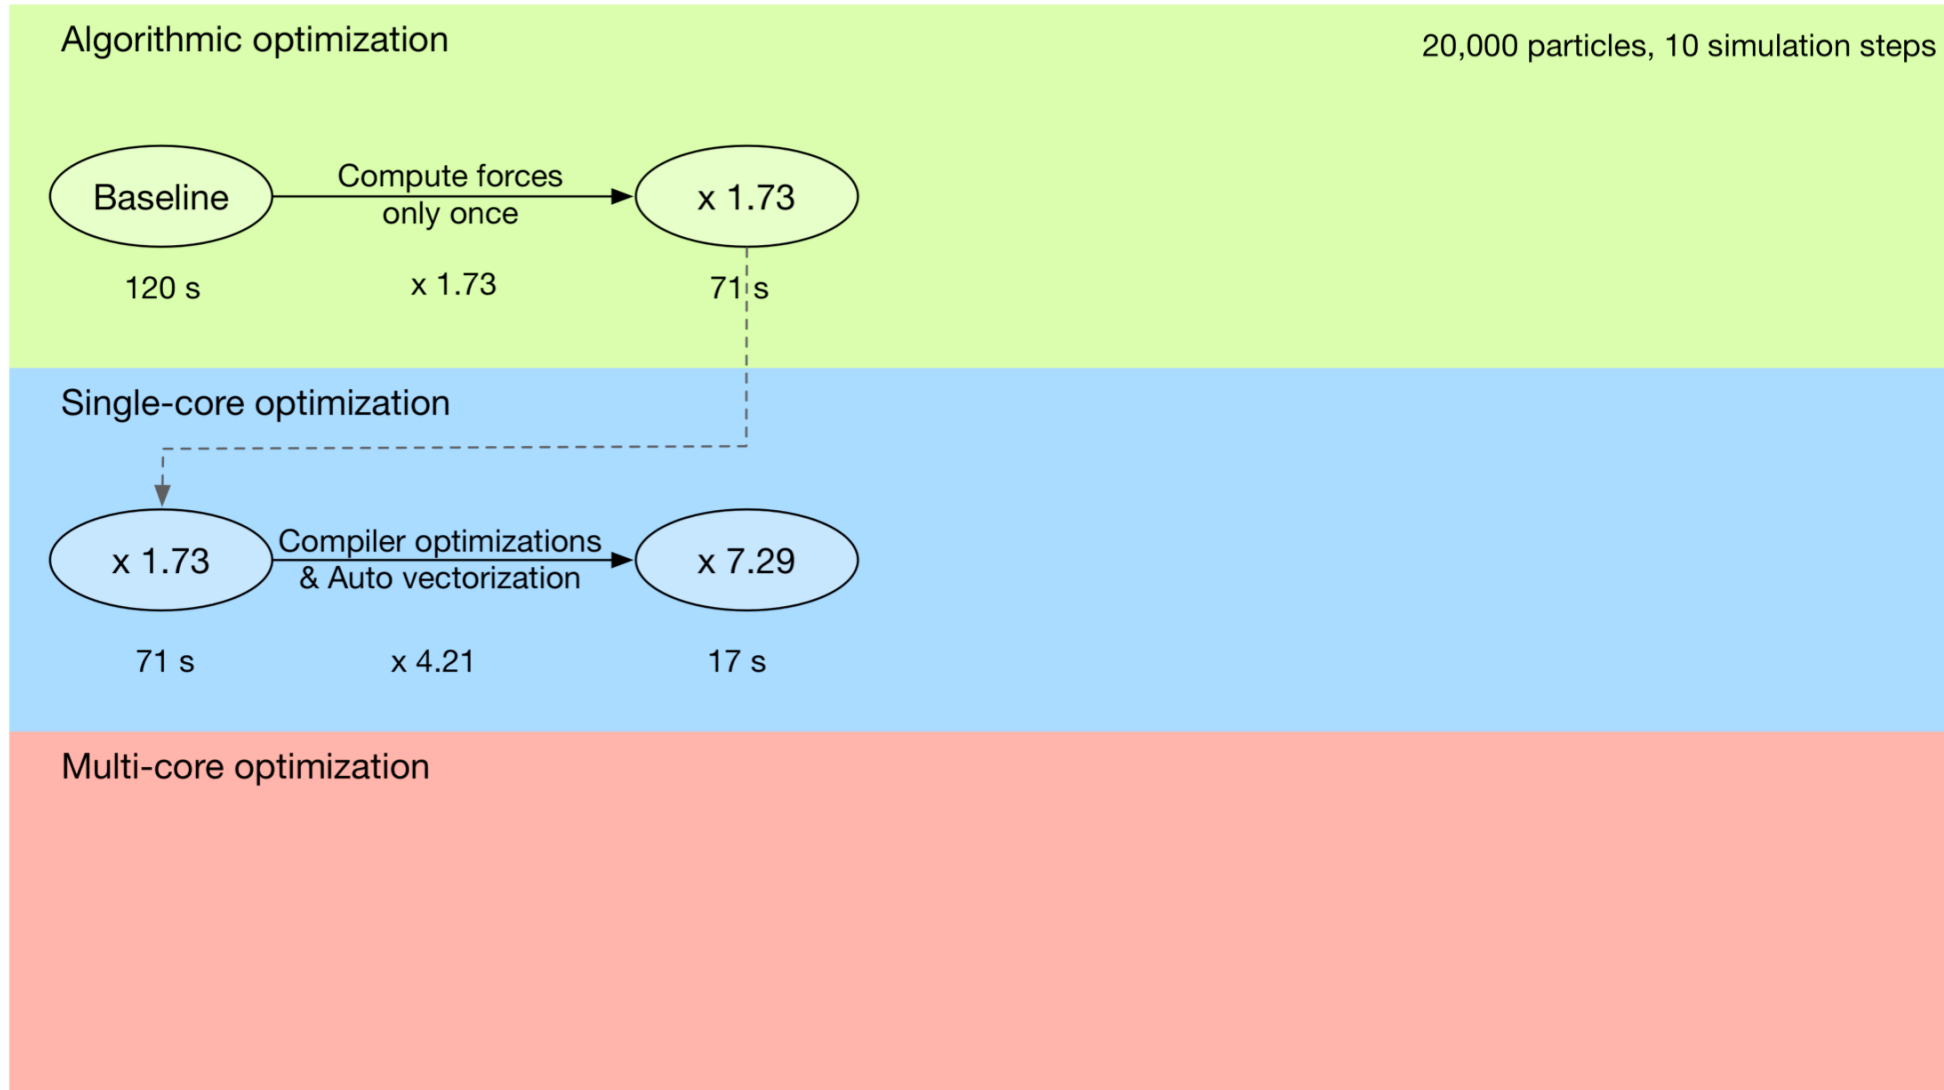
Compiler optimizations only

Performance Optimization: n-body



Compiler optimizations and auto vectorization

Performance Optimization: n-body




Demo: Intel Advisor Recommendations


Performance Optimization: n-body



- Advisor gives hints on how to improve performance:

The screenshot shows a navigation bar with tabs: Source, Top Down, Code Analytics, Assembly, Recommendations (with a lightbulb icon), and Why No Vectorization? (with a purple icon). Below the tabs, there is a link for "All Advisor-detectable issues: C++ | Fortran". The main content area displays a "Recommendation: Confirm inefficient memory access patterns" with a subtext: "There is no confirmation inefficient memory access patterns are present. To confirm: Run a Memory Access Patterns analysis." To the right, a grey box contains an "Issue: Possible inefficient memory access patterns present" with a detailed explanation: "Inefficient memory access patterns may result in significant vector code execution slowdown or block automatic vectorization by the compiler. Improve performance by investigating." At the bottom of this box is a blue lightbulb icon and a link: "Confirm inefficient memory ...".

Performance Optimization: n-body

Site Location	Loop-Carried Dependencies	Strides Distribution	Access Pattern	Max. Site Footprint	Site N
 [loop in Compute_force at 01_Reduced_Workload.c:...	No information available	33% / 67% / 0%	Mixed strides	361KB	loop_s

Memory Access Patterns Report	Dependencies Report	 Recommendations
-------------------------------	---------------------	--

ID		Stride	Type	Source	Nested Function	Variable references
▼ P1		10; 20	Constant stride	01_Reduced_Workload.c:259		block 0x2aaac112a010 allocated at 01_Reduced_Workload.c:48
<pre>257 for (k = part+1; k < n; k++) { 258 /* Compute force on part due to k */ 259 f_part_k[X] = curr[part].s[X] - curr[k].s[X]; 260 f_part_k[Y] = curr[part].s[Y] - curr[k].s[Y]; 261 len = sqrt(f_part_k[X]*f_part_k[X] + f_part_k[Y]*f_part_k[Y]);</pre>						

Performance Optimization: n-body

- Unit strides: Elements in memory accessed in consecutive order
- Constant strides: Fixed distance between accessed elements
- Irregular strides: Random access

- Why can non-unit strides be a problem?
 - Not all data transferred from memory is used
 - Elements need to be gathered into vectors
 - Memory prefetching may not provide required data

```
#define SIZE 1000
int i, j;
double *a = malloc(SIZE * SIZE * sizeof(double))
for (j=0; j < SIZE; j++) {
    for (i=0; i < SIZE; i++) {
        do_something(a[i*SIZE + j]);
    }
}
```


Performance Optimization: n-body

- In our application we use an array of structs

```
typedef double vect_t[DIM];
struct particle_s {
    double m; /* Mass */
    vect_t s; /* Position */
    vect_t v; /* Velocity */
};
[...]
```

```
curr = malloc(n*sizeof(struct particle_s));
```

- Distance between X coordinates of two consecutive particles in memory: 40 bytes

- Solution: Use struct of arrays instead

```
typedef double vect_t[DIM];
struct particles {
    double *m; /* Masses */
    vect_t *s; /* Positions */
    vect_t *v; /* Velocities */
};
[...]
```

```
curr = malloc(sizeof(struct particles));
curr->m = malloc(n * sizeof(double));
curr->s = malloc(n * sizeof(vect_t));
curr->v = malloc(n * sizeof(vect_t));
```

Performance Optimization: n-body

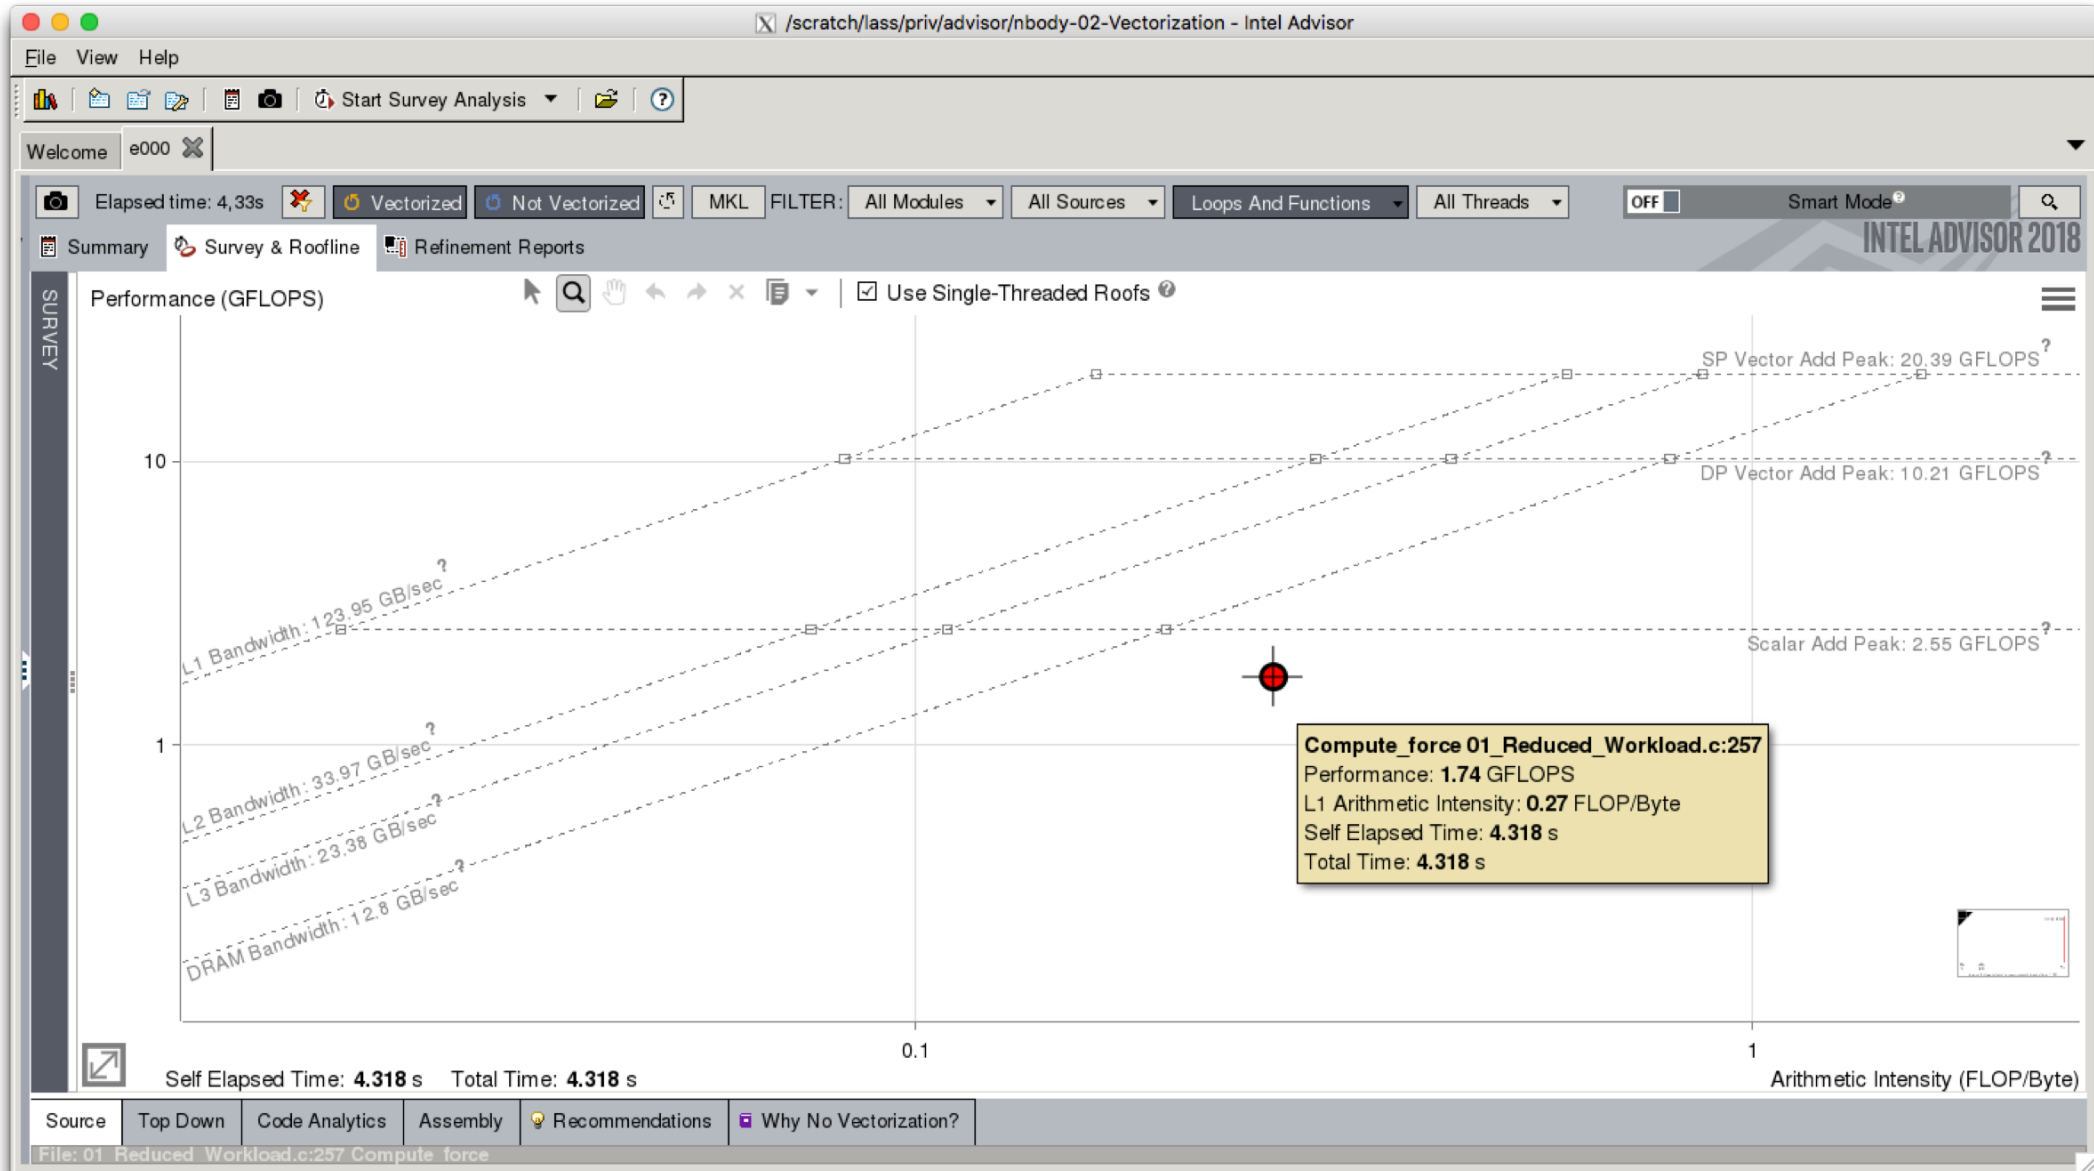
- Changing underlying data structure requires to change each access

```
f_part_k[X] = curr[part].s[X] - curr[k].s[X];
```

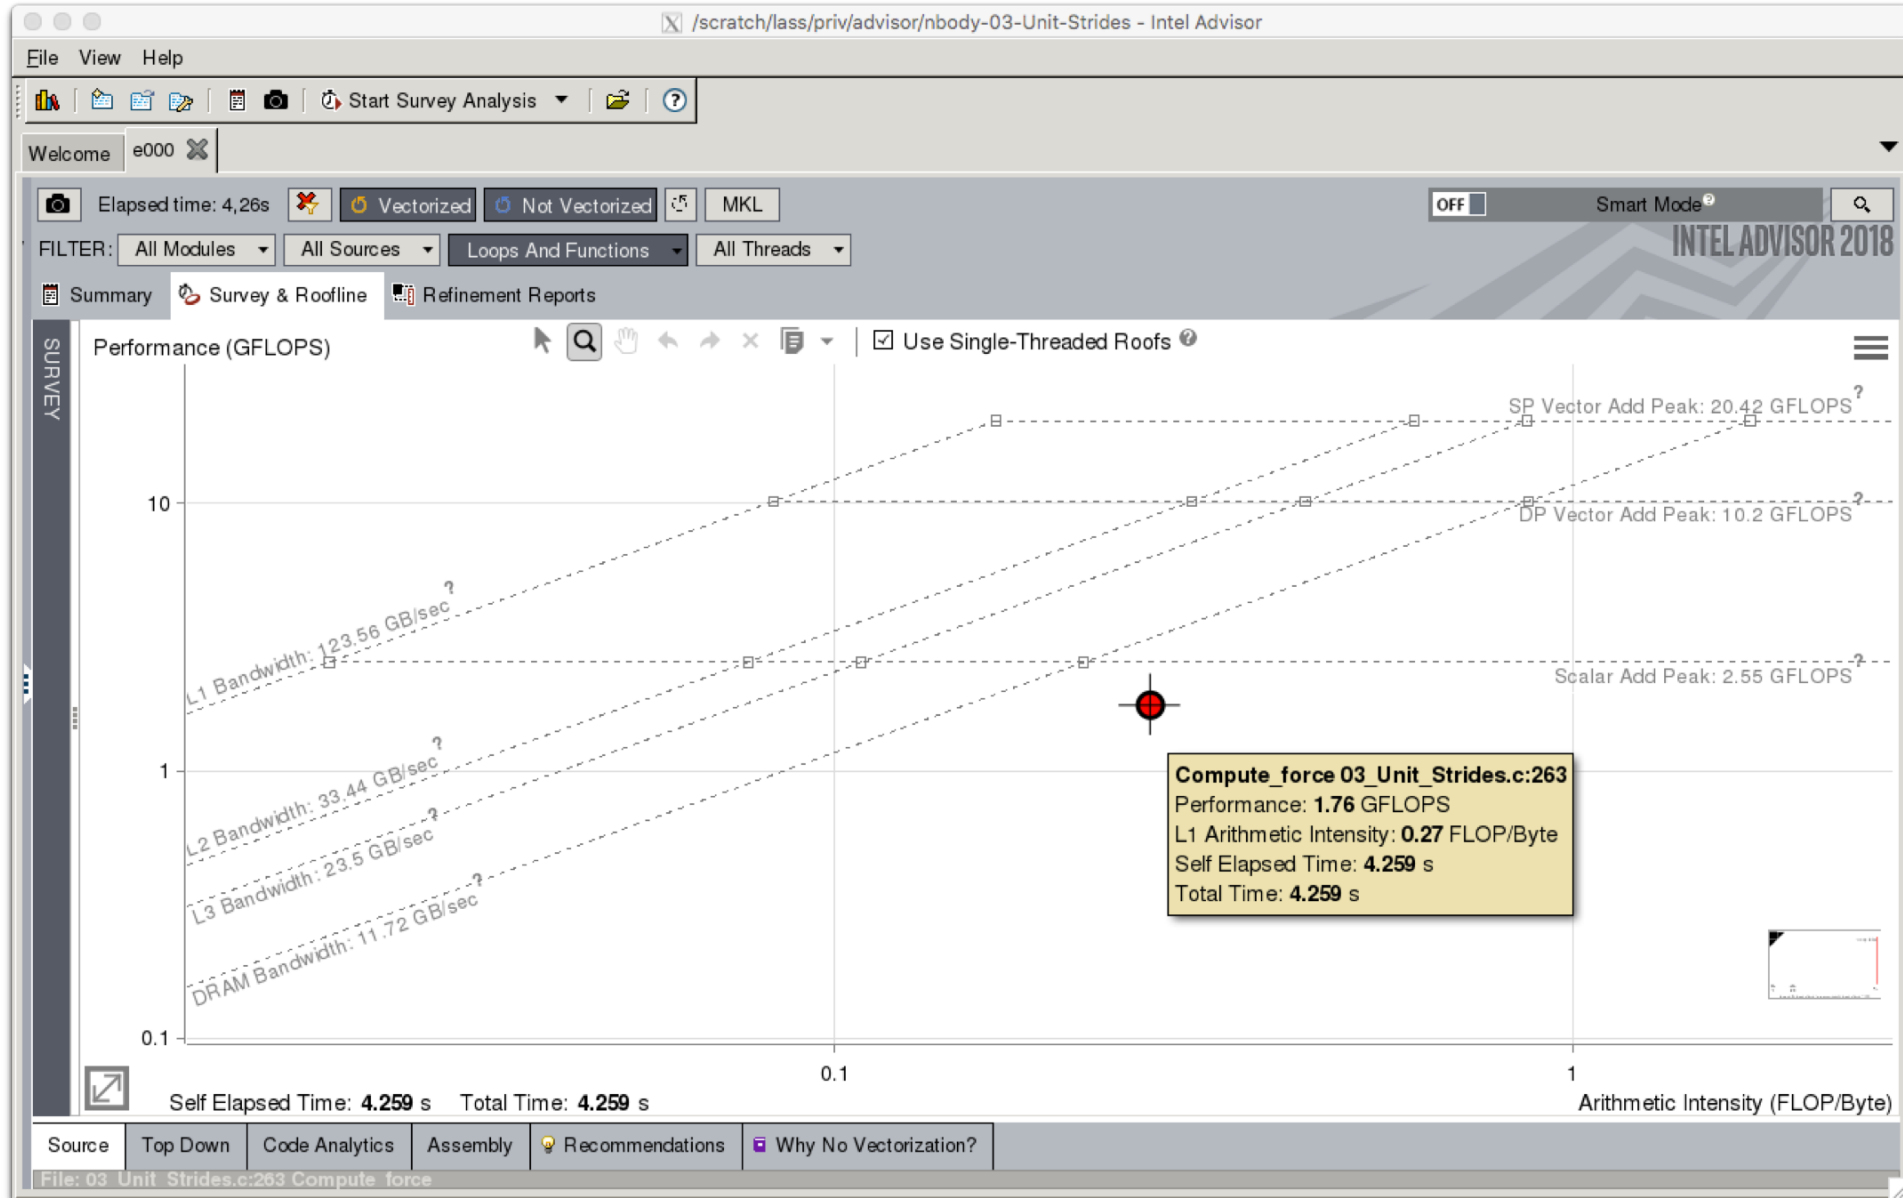


```
f_part_k[X] = curr->s[part][X] - curr->s[k][X];
```


Performance Optimization: n-body

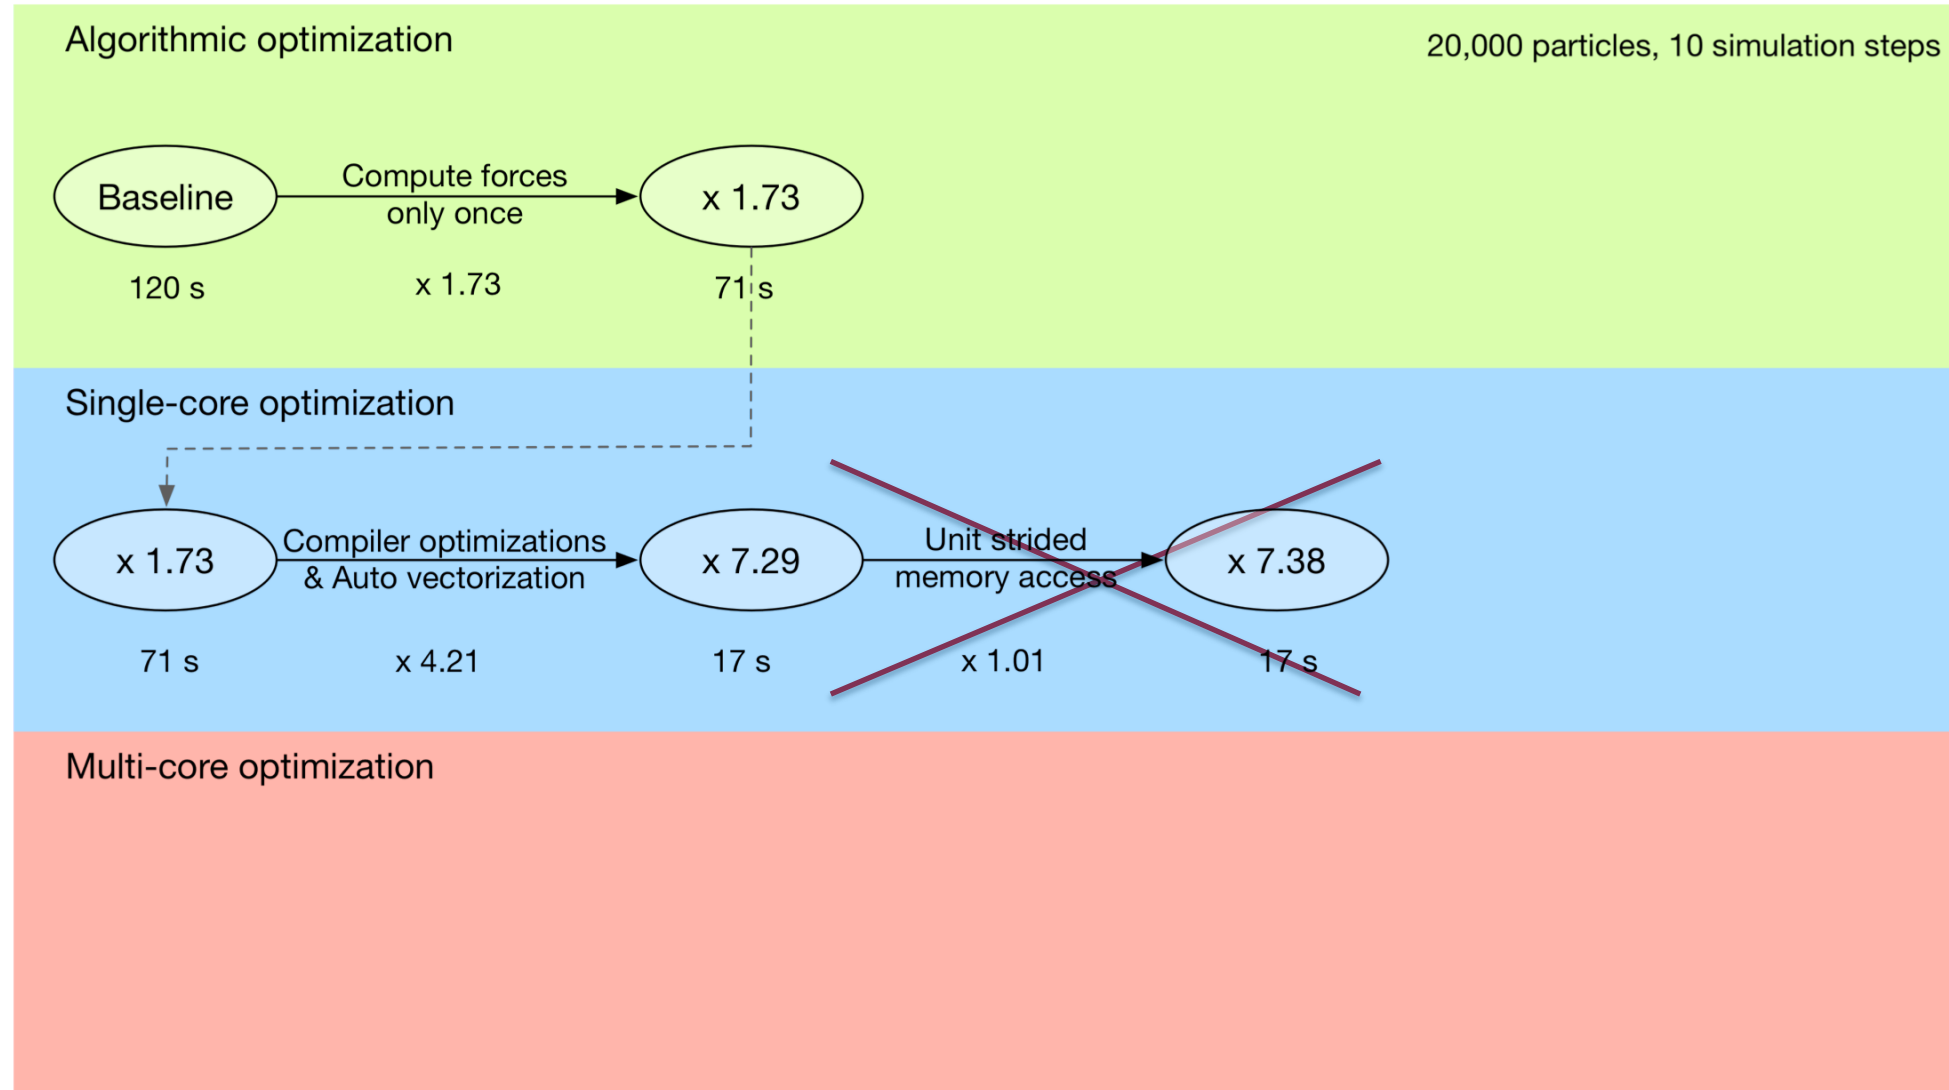


Performance Optimization: n-body



Optimized memory access patterns

Performance Optimization: n-body



Performance Optimization: n-body

- Enable multi-threading using OpenMP

```
for (step = 1; step <= n_steps; step++) {  
    t = step*delta_t;  
    memset(forces, 0, n*sizeof(vect_t));  
    for (part = 0; part < n-1; part++)  
        Compute_force(part, forces, curr, n);  
    for (part = 0; part < n; part++)  
        Update_part(part, forces, curr, n, delta_t);  
}
```

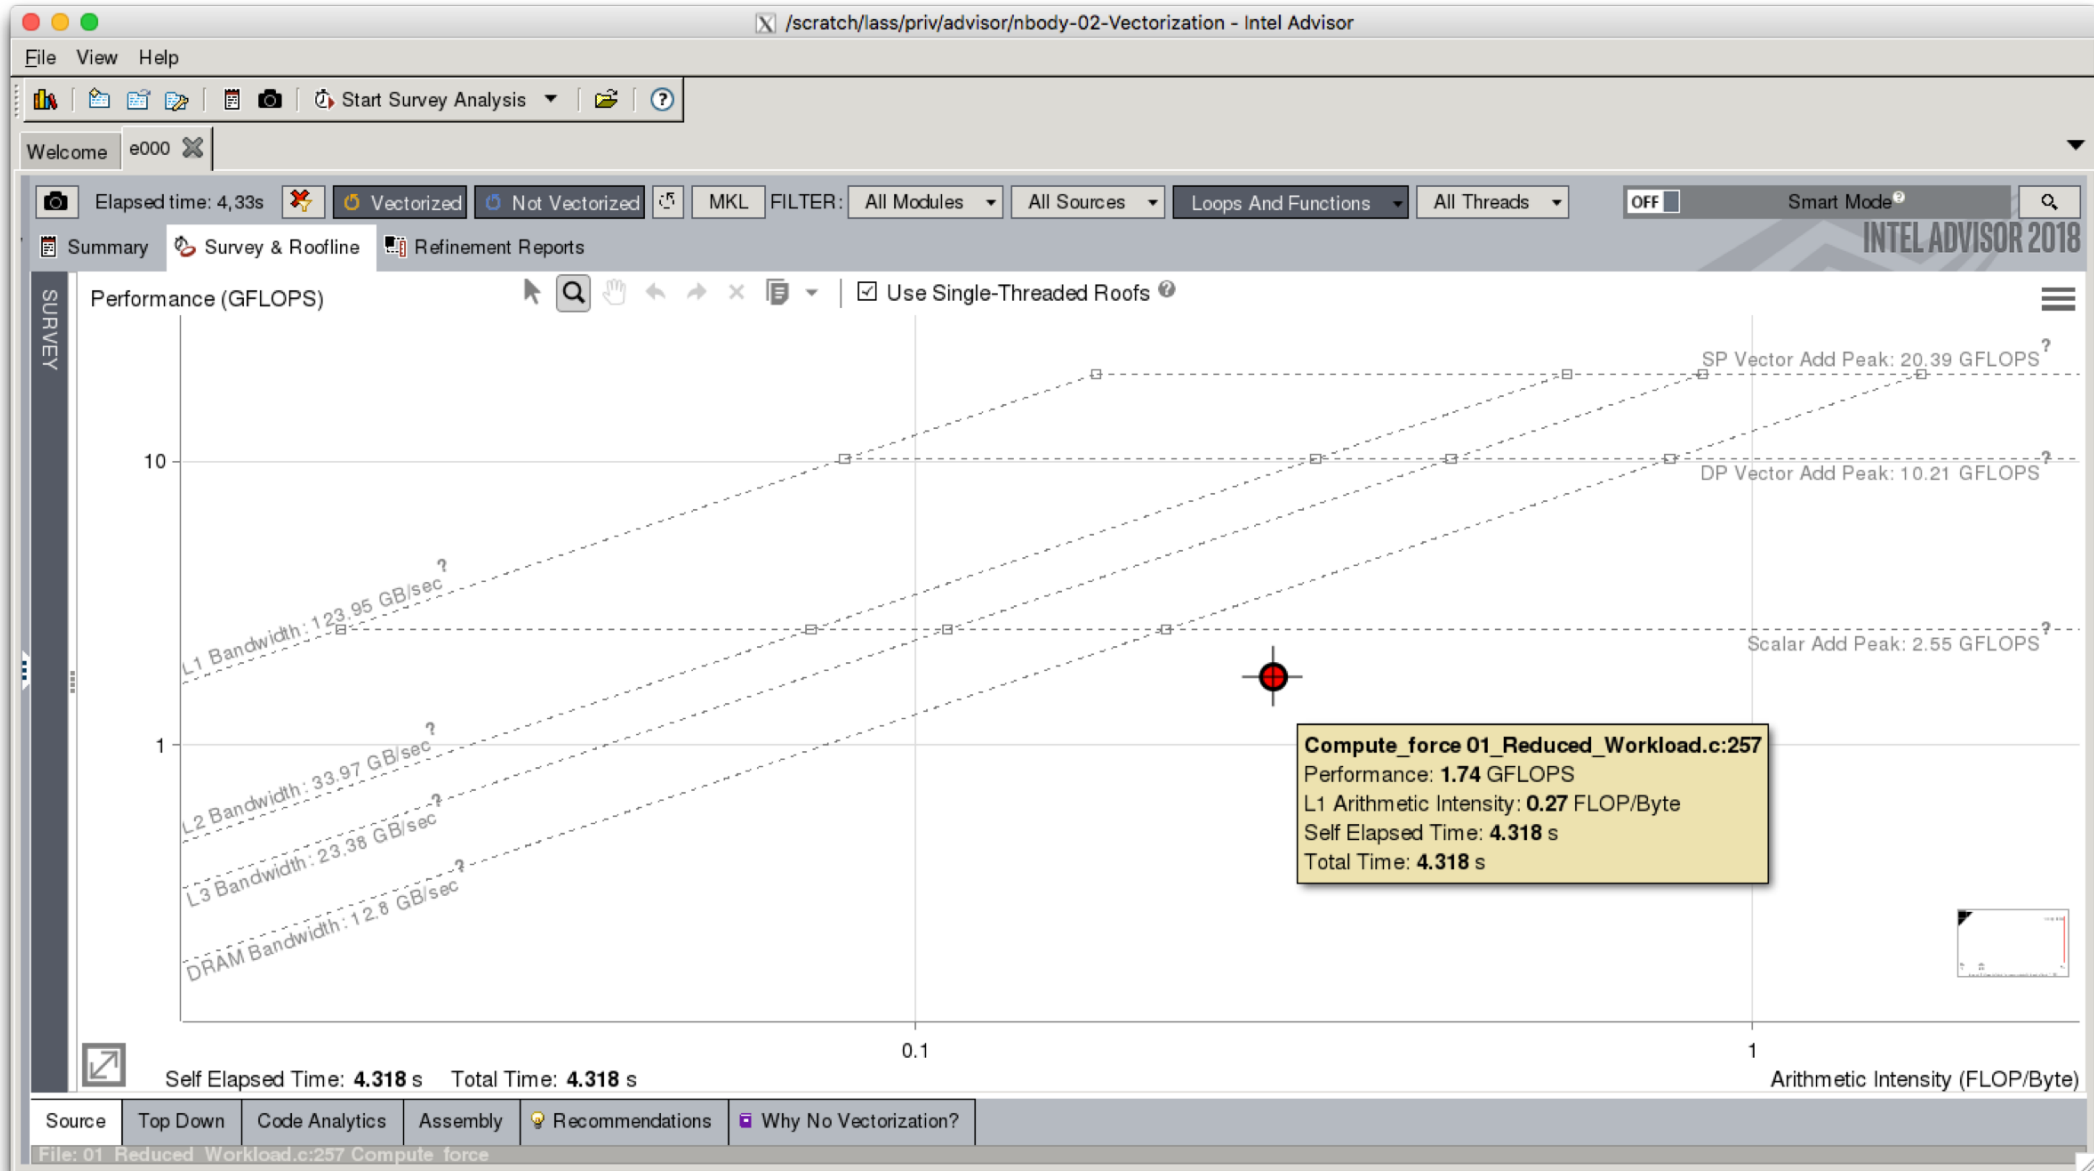


```
# pragma omp parallel num_threads(thread_count) default(none) \  
shared(curr, forces, thread_count, delta_t, n, n_steps, output_freq) \  
private(step, part, t)  
    for (step = 1; step <= n_steps; step++) {  
        t = step*delta_t;  
# pragma omp single  
        memset(forces, 0, n*sizeof(vect_t));  
# pragma omp for  
        for (part = 0; part < n-1; part++)  
            Compute_force(part, forces, curr, n);  
# pragma omp for  
        for (part = 0; part < n; part++)  
            Update_part(part, forces, curr, n, delta_t);  
    }
```

```
for (k = part+1; k < n; k++) {  
    [...]  
    forces[part][X] += f_part_k[X];  
    forces[part][Y] += f_part_k[Y];  
    forces[k][X] -= f_part_k[X];  
    forces[k][Y] -= f_part_k[Y];  
}
```

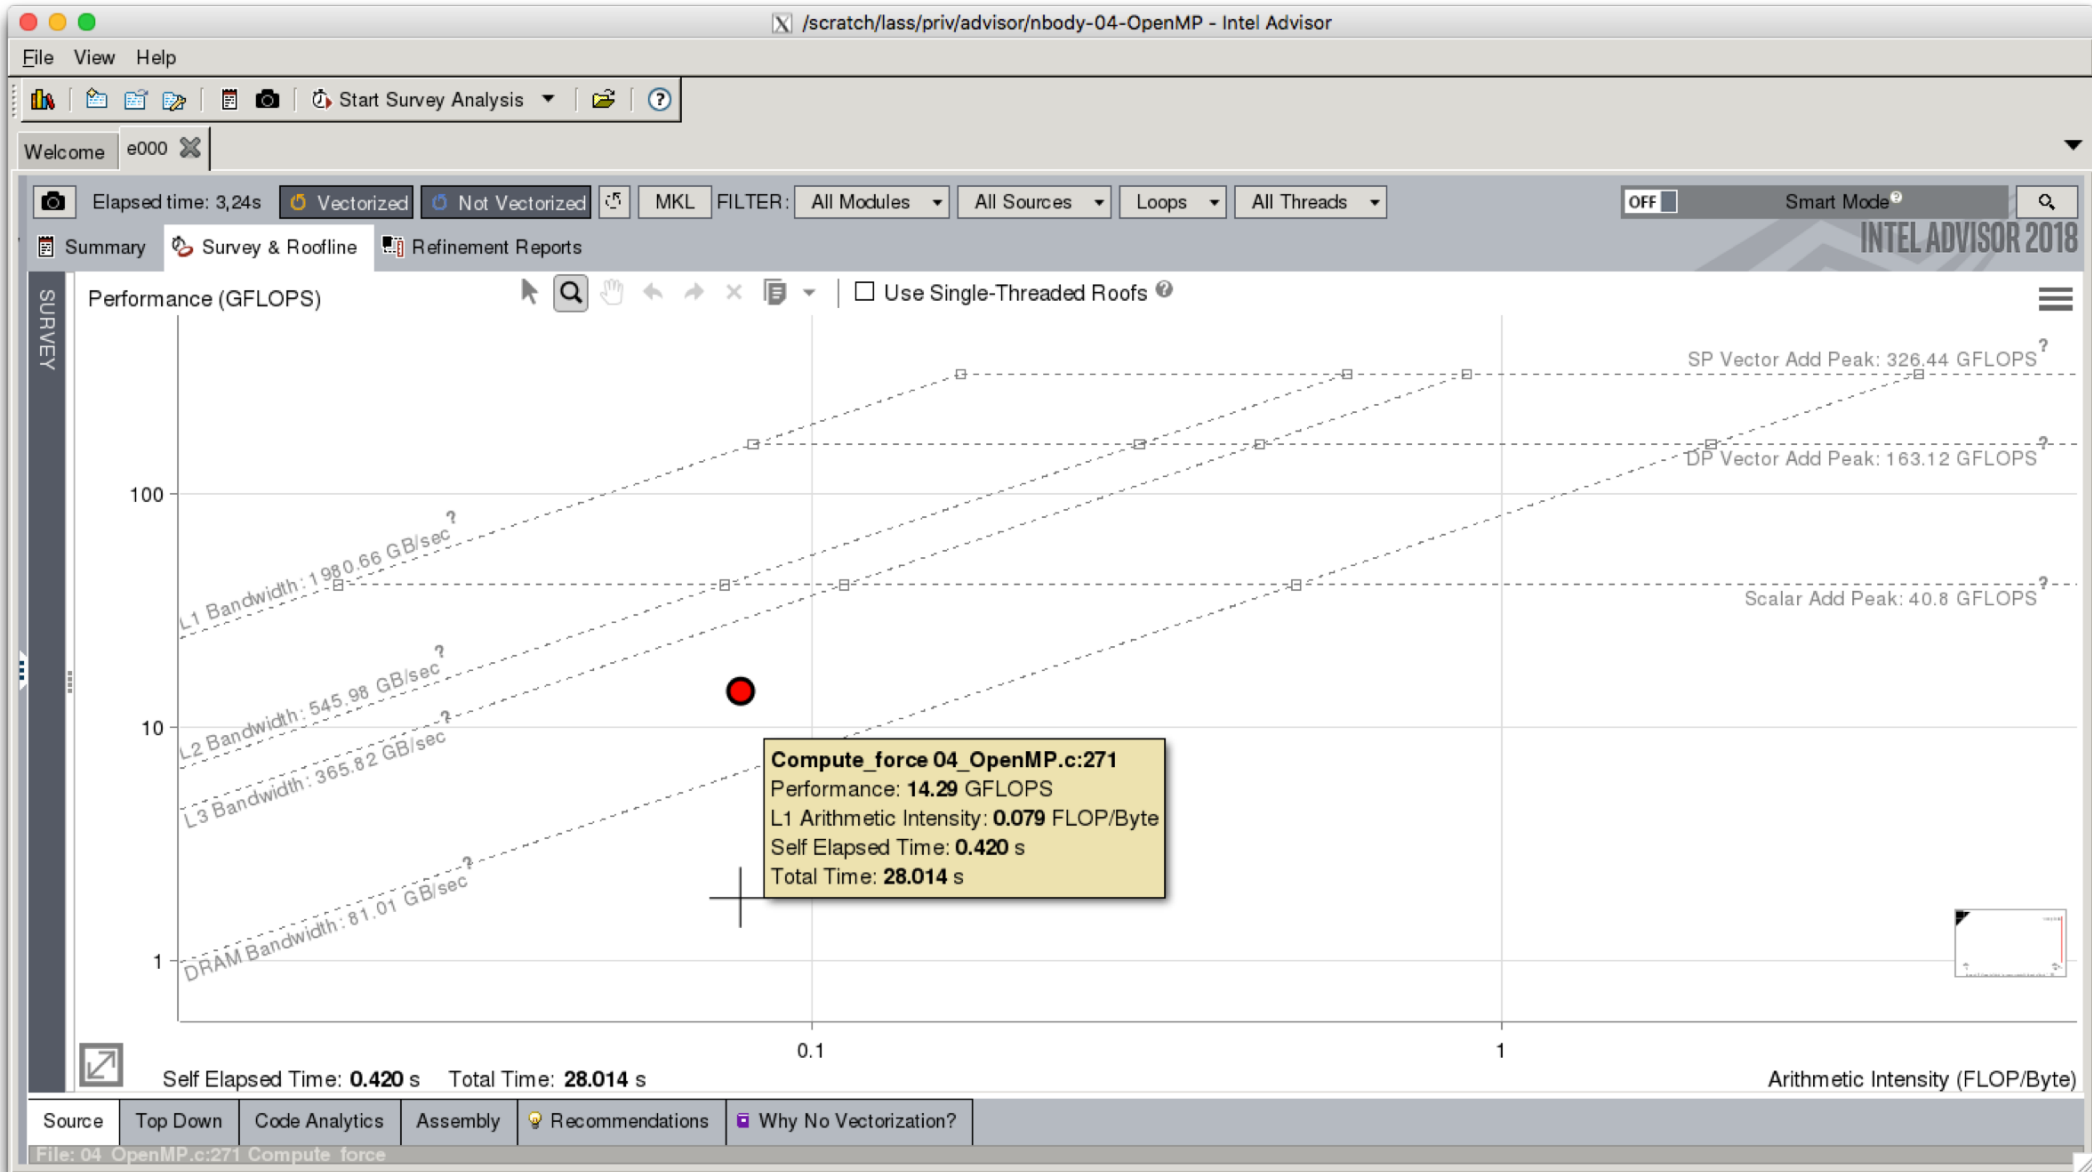
```
for (k = part+1; k < n; k++) {  
    [...]  
# pragma omp atomic  
    forces[part][X] += f_part_k[X];  
# pragma omp atomic  
    forces[part][Y] += f_part_k[Y];  
# pragma omp atomic  
    forces[k][X] -= f_part_k[X];  
# pragma omp atomic  
    forces[k][Y] -= f_part_k[Y];  
}
```

Performance Optimization: n-body



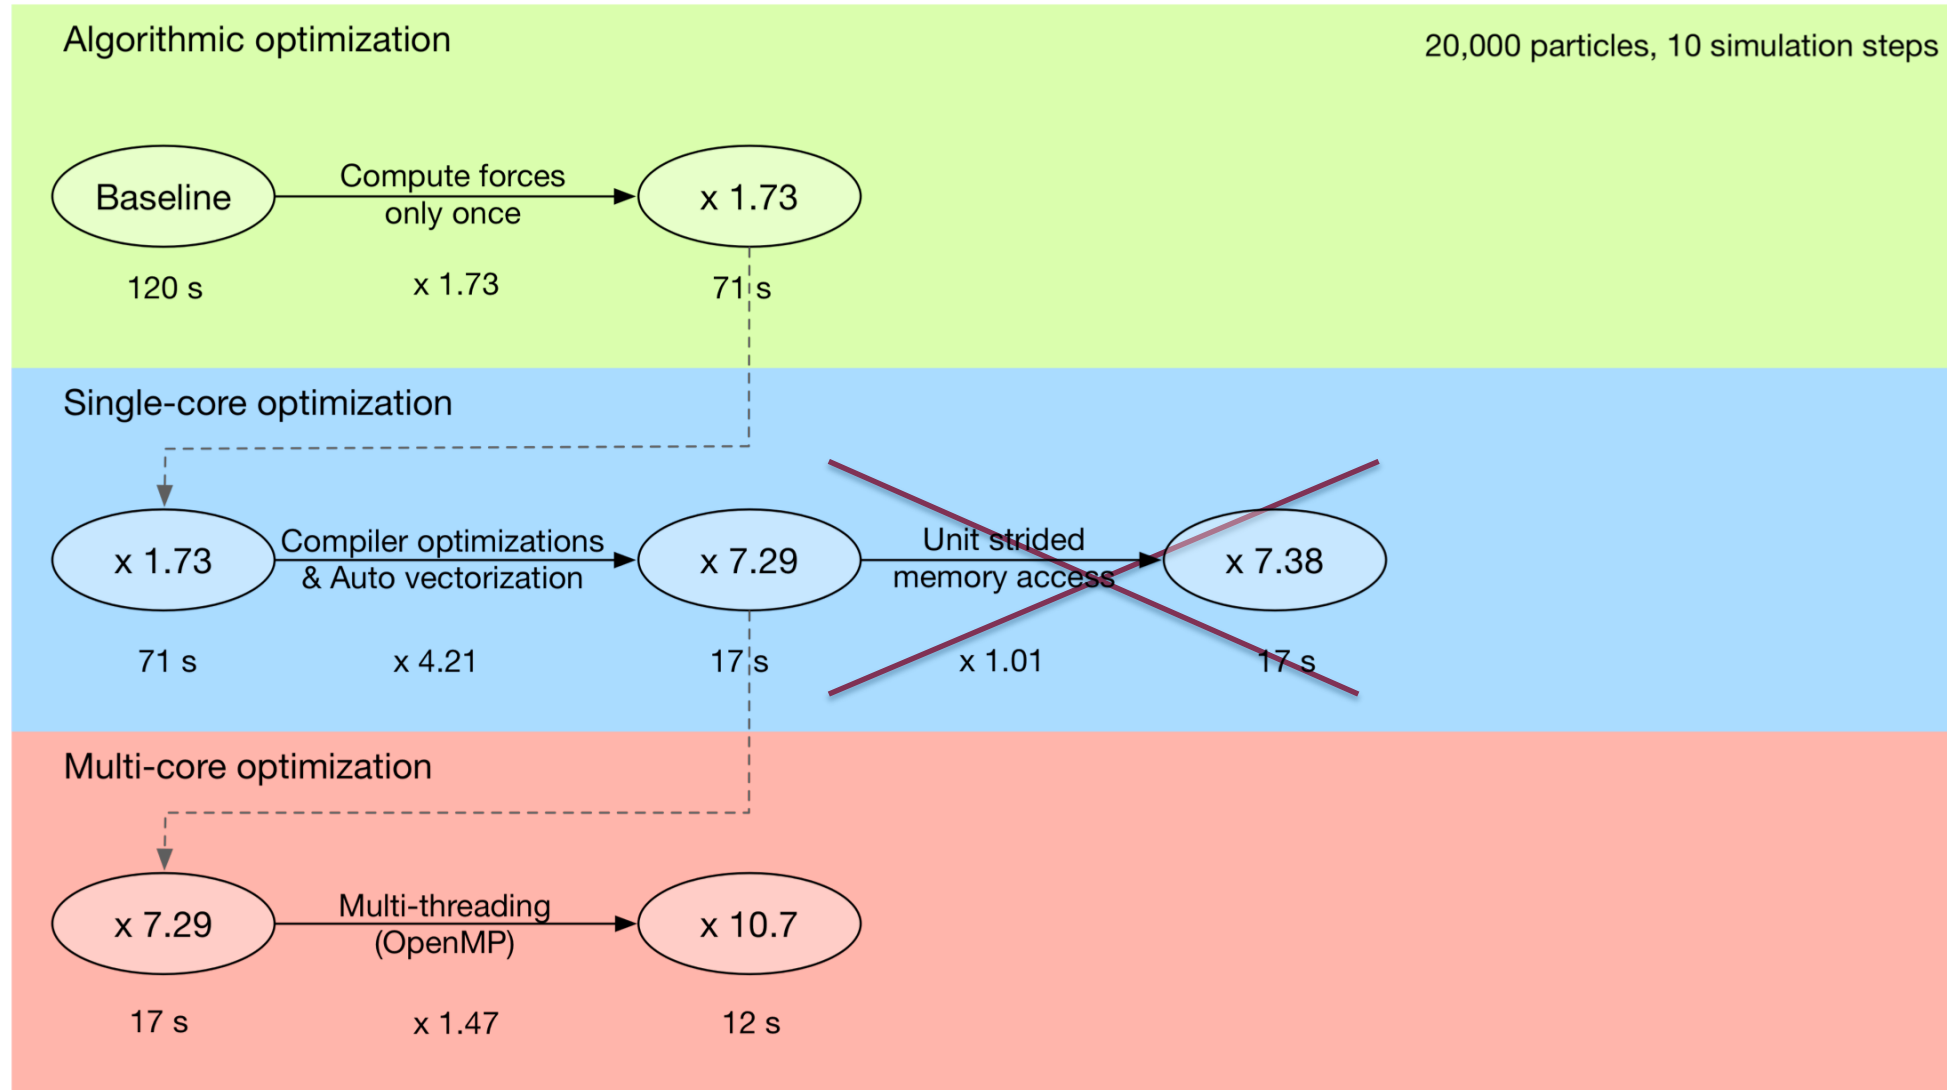
Compiler optimizations and auto vectorization

Performance Optimization: n-body



Parallelized with OpenMP

Performance Optimization: n-body



Performance Optimization: n-body

- Intel VTune: Profiling tool with lots of different analysis types
 - Basic / Advanced Hotspot Analysis
 - Concurrency / Locks & Waits
 - General Exploration: Efficient use of microarchitecture
 - Memory bandwidth
 - IO access
 - ...
- Our parallelization did not work out, so do a *Concurrency* analysis

- Starting VTune on our systems:

```
module load ps_xe_2018  
amplxe-gui
```

- Requesting a node with VTune support:

```
ccsalloc --res=rset=1:vtune=1:place=scatter:excl
```

Demo: Intel VTune Concurrency Analysis

Performance Optimization: n-body

Top Hotspots

This section lists the most active functions in your application. Optimizing these hotspot functions typically results in improving overall application performance.

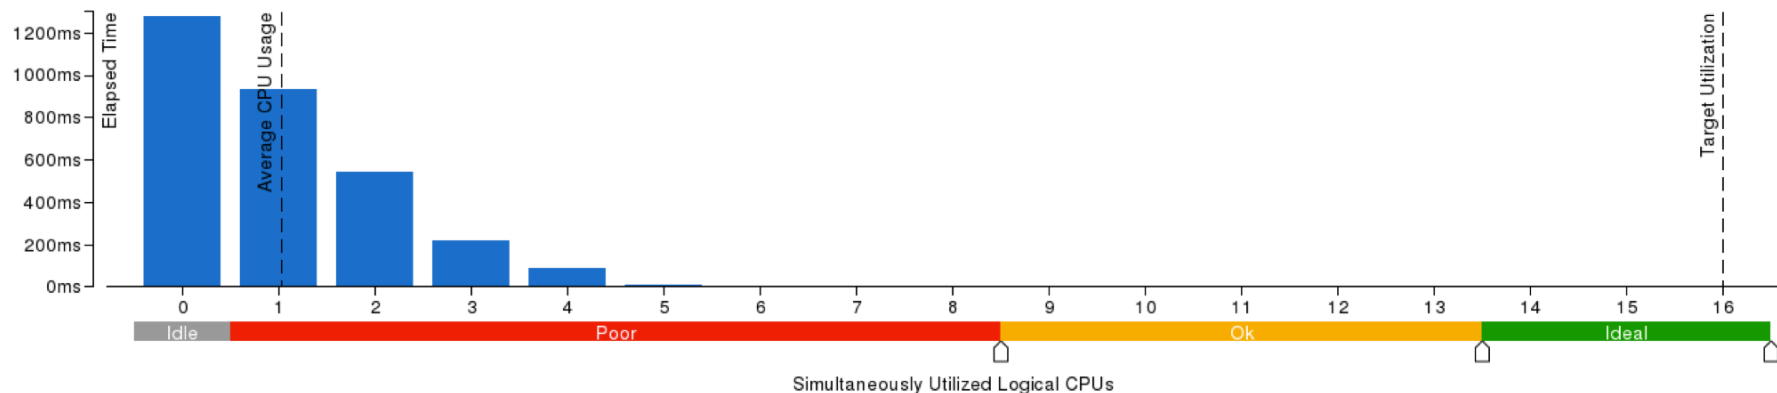
Function	Module	CPU Time [?]
__kmpc_barrier	libiomp5.so	17.811s ▾
__kmpc_atomic_float8_sub	libiomp5.so	14.548s ▾
__kmpc_atomic_float8_add	libiomp5.so	9.447s ▾
Compute_force	04_OpenMP	3.151s
[Import thunk __kmpc_atomic_float8_sub]	04_OpenMP	0.350s
[Others]		0.290s

Elapsed Time[?]: 3.083s 📄

↳ CPU Time [?] :	45.597s
↳ Effective Time [?] :	3.151s
↳ Spin Time [?] :	17.921s ▾
↳ Imbalance or Serial Spinning [?] :	17.491s ▾
↳ Lock Contention [?] :	0s
↳ Other [?] :	0.430s
↳ Overhead Time [?] :	24.525s ▾
↳ Creation [?] :	0s
↳ Scheduling [?] :	0s
↳ Reduction [?] :	0s
↳ Atomics [?] :	24.525s ▾
↳ Other [?] :	0s
↳ Wait Time [?] :	2.295s
Total Thread Count:	16
Paused Time [?] :	0s

CPU Usage Histogram 📄

This histogram displays a percentage of the wall time the specific number of CPUs were running simultaneously. Spin and Overhead time adds to the Idle CPU usage value.

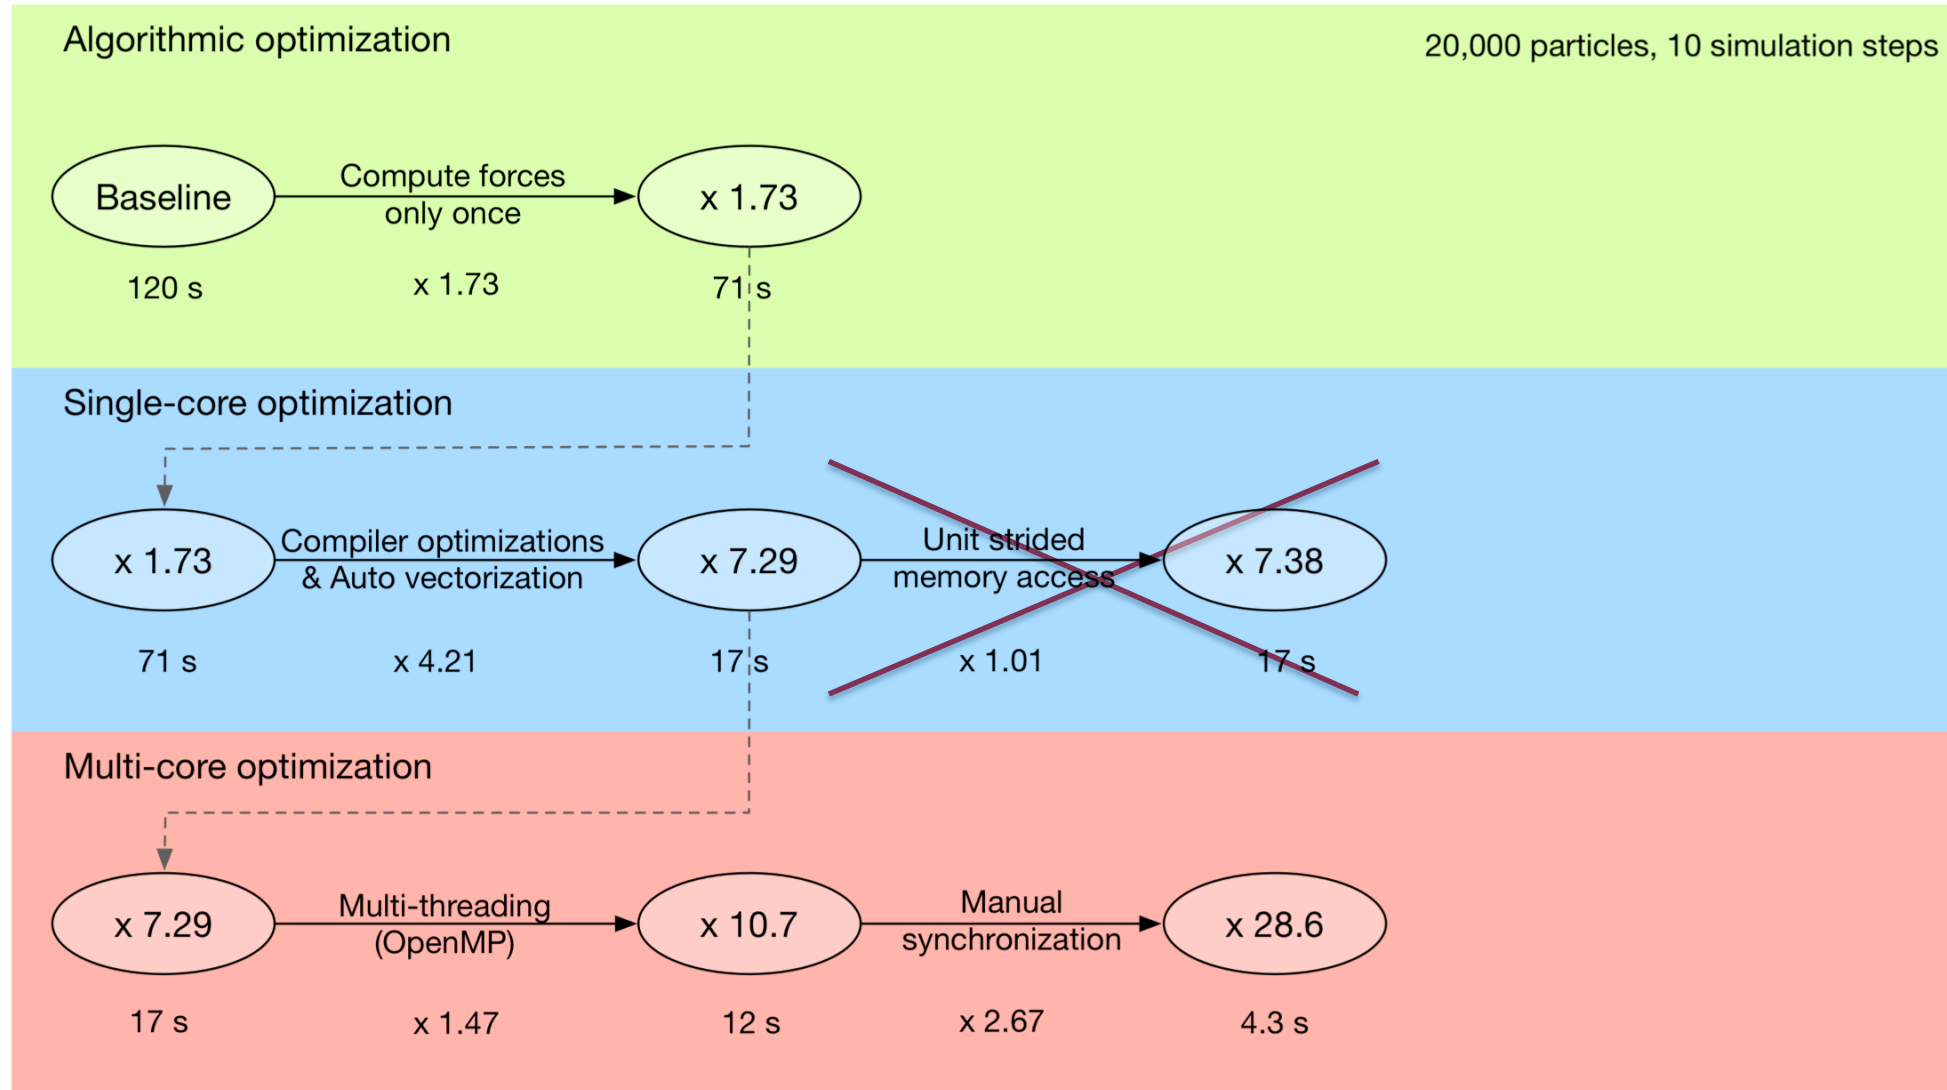


Performance Optimization: n-body

- Atomic operations are costly
- Solution:
 - Let each thread store its own (partial) forces
 - Aggregate results after all threads have finished

```
# pragma omp parallel num_threads(thread_count) default(none) \
shared(curr,forces,thread_count,delta_t,n,n_steps, \
      output_freq,loc_forces) \
private(step, part, t)
{
  int my_rank = omp_get_thread_num();
  int thread;
  for (step = 1; step <= n_steps; step++) {
    t = step*delta_t;
#   pragma omp for
    for (part = 0; part < thread_count*n; part++)
      loc_forces[part][X] = loc_forces[part][Y] = 0.0;
#   pragma omp for
    for (part = 0; part < n-1; part++)
      Compute_force(part, loc_forces + my_rank*n, curr, n);
#   pragma omp for
    for (part = 0; part < n; part++) {
      forces[part][X] = forces[part][Y] = 0.0;
      for (thread = 0; thread < thread_count; thread++) {
        forces[part][X] += loc_forces[thread*n + part][X];
        forces[part][Y] += loc_forces[thread*n + part][Y];
      }
    }
#   pragma omp for
    for (part = 0; part < n; part++)
      Update_part(part, forces, curr, n, delta_t);
  }
}
```

Performance Optimization: n-body



Demo: Intel VTune Concurrency Analysis

Performance Optimization: n-body

Top Hotspots

This section lists the most active functions in your application. Optimizing these hotspot functions typically results in improving overall application performance.

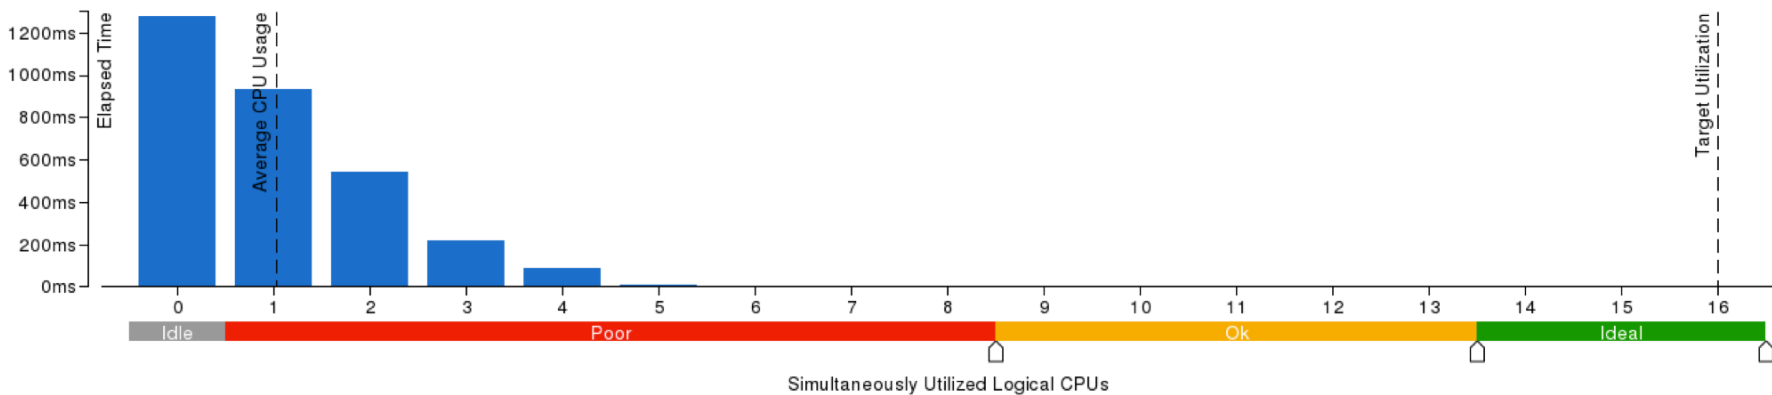
Function	Module	CPU Time
__kmpc_barrier	libiomp5.so	17.811s
__kmpc_atomic_float8_sub	libiomp5.so	14.548s
__kmpc_atomic_float8_add	libiomp5.so	9.447s
Compute_force	04_OpenMP	3.151s
[Import thunk __kmpc_atomic_float8_sub]	04_OpenMP	0.350s
[Others]		0.290s

Elapsed Time: 3.083s

CPU Time	45.597s
Effective Time	3.151s
Spin Time	17.921s
Imbalance or Serial Spinning	17.491s
Lock Contention	0s
Other	0.430s
Overhead Time	24.525s
Creation	0s
Scheduling	0s
Reduction	0s
Atomics	24.525s
Other	0s
Wait Time	2.295s
Total Thread Count	16
Paused Time	0s

CPU Usage Histogram

This histogram displays a percentage of the wall time the specific number of CPUs were running simultaneously. Spin and Overhead time adds to the Idle CPU usage value.



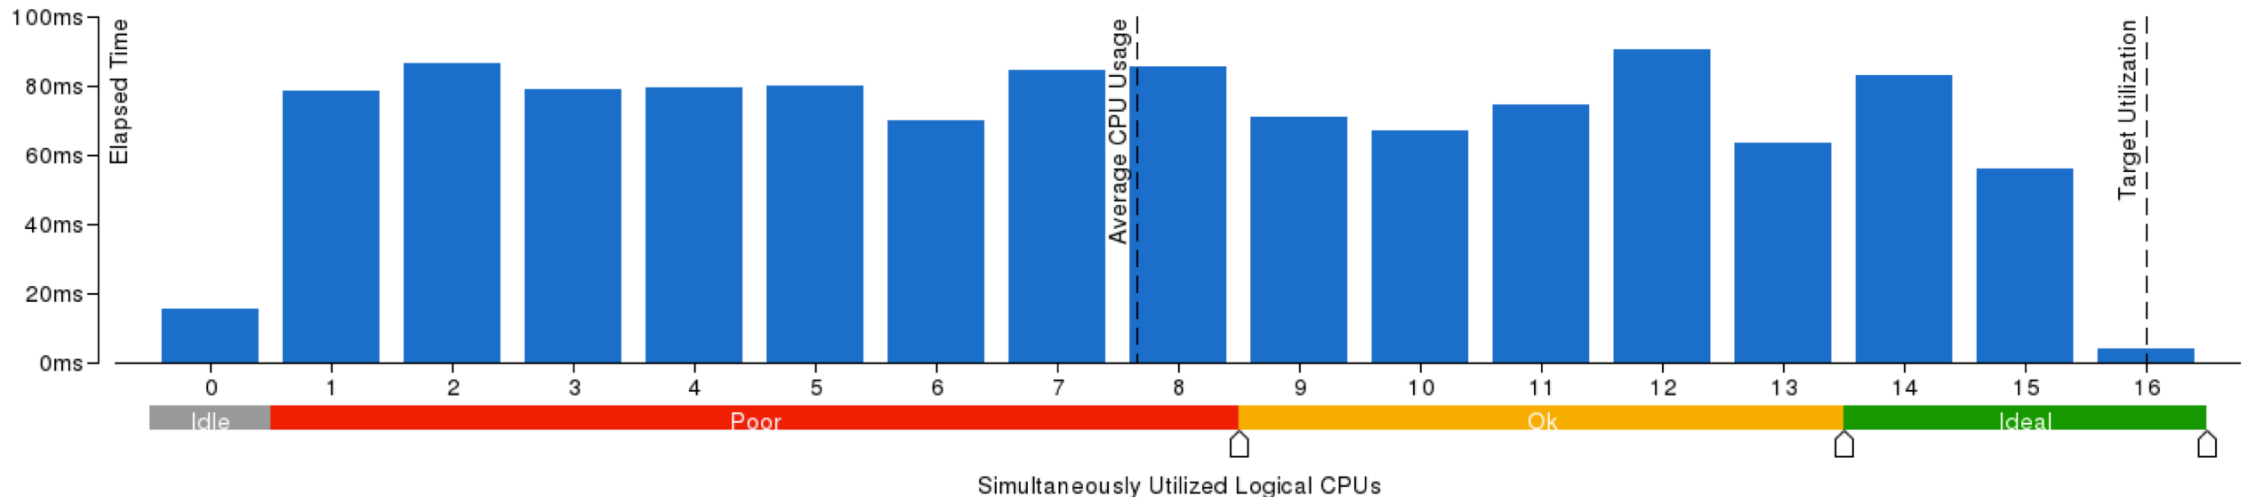
Performance Optimization: n-body

Elapsed Time [?]: 1.172s

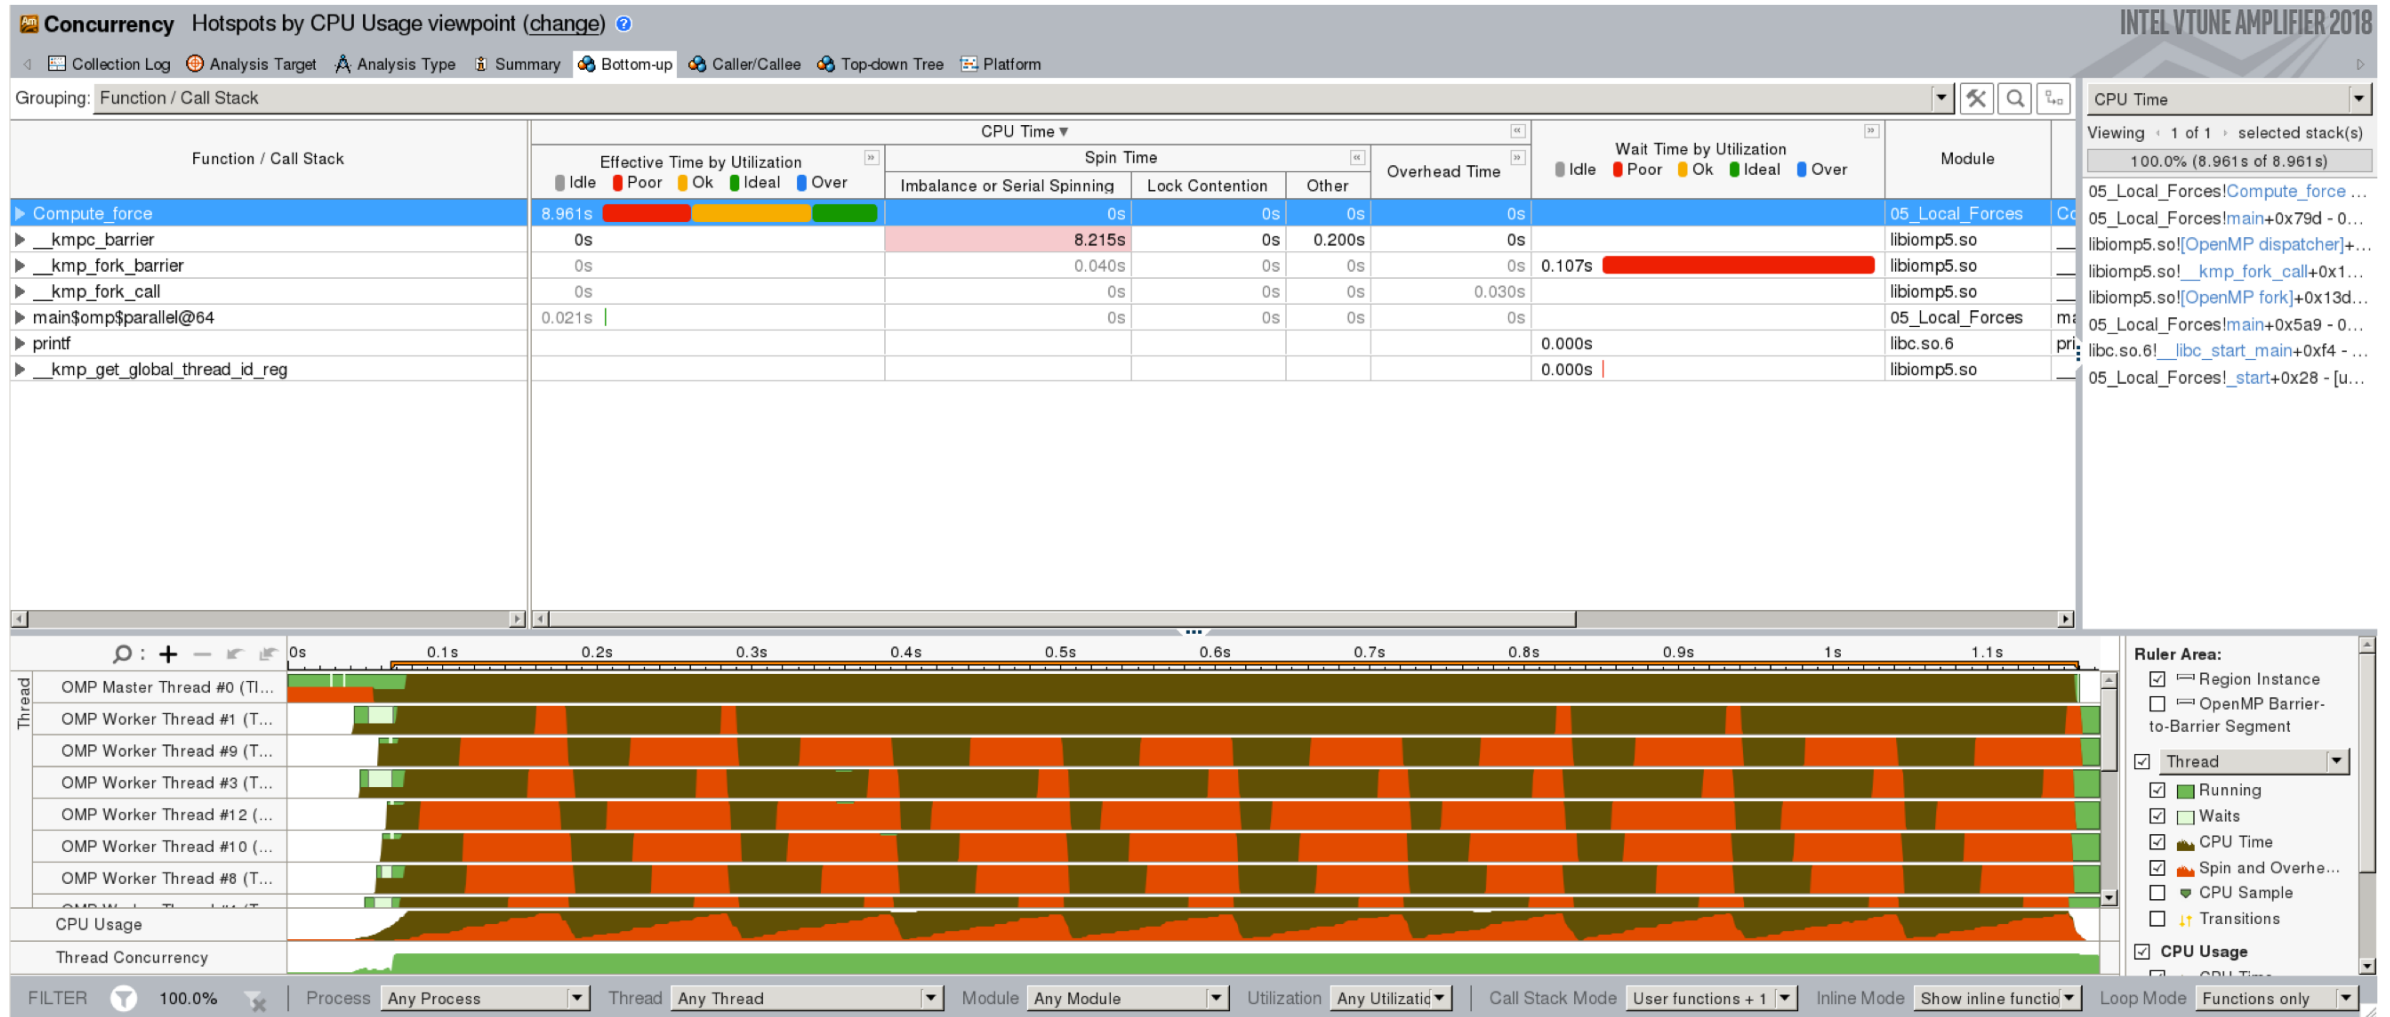
⊖ CPU Time [?] :	17.467s
⊙ Effective Time [?] :	8.982s
⊖ Spin Time [?] :	8.455s ▾
Imbalance or Serial Spinning [?] :	8.255s ▾
Lock Contention [?] :	0s
Other [?] :	0.200s
⊙ Overhead Time [?] :	0.030s
⊙ Wait Time [?] :	0.108s
Total Thread Count:	16
Paused Time [?] :	0s

⊖ CPU Usage Histogram

This histogram displays a percentage of the wall time the specific number of CPUs were running simultaneously. Spin and Overhead time adds to the Idle CPU usage value.



Performance Optimization: n-body



Performance Optimization: n-body

- **Work imbalance between the threads:**
 - First particle: all forces have to be calculated
 - Last particle: no forces have to be calculated
 - First thread has the most work, last thread the least
- **Solution: Cyclic assignments from particles to threads**

```
# pragma omp for schedule(static,1)  
for (part = 0; part < n-1; part++)  
    Compute_force(part, loc_forces + my_rank*n, curr, n);
```

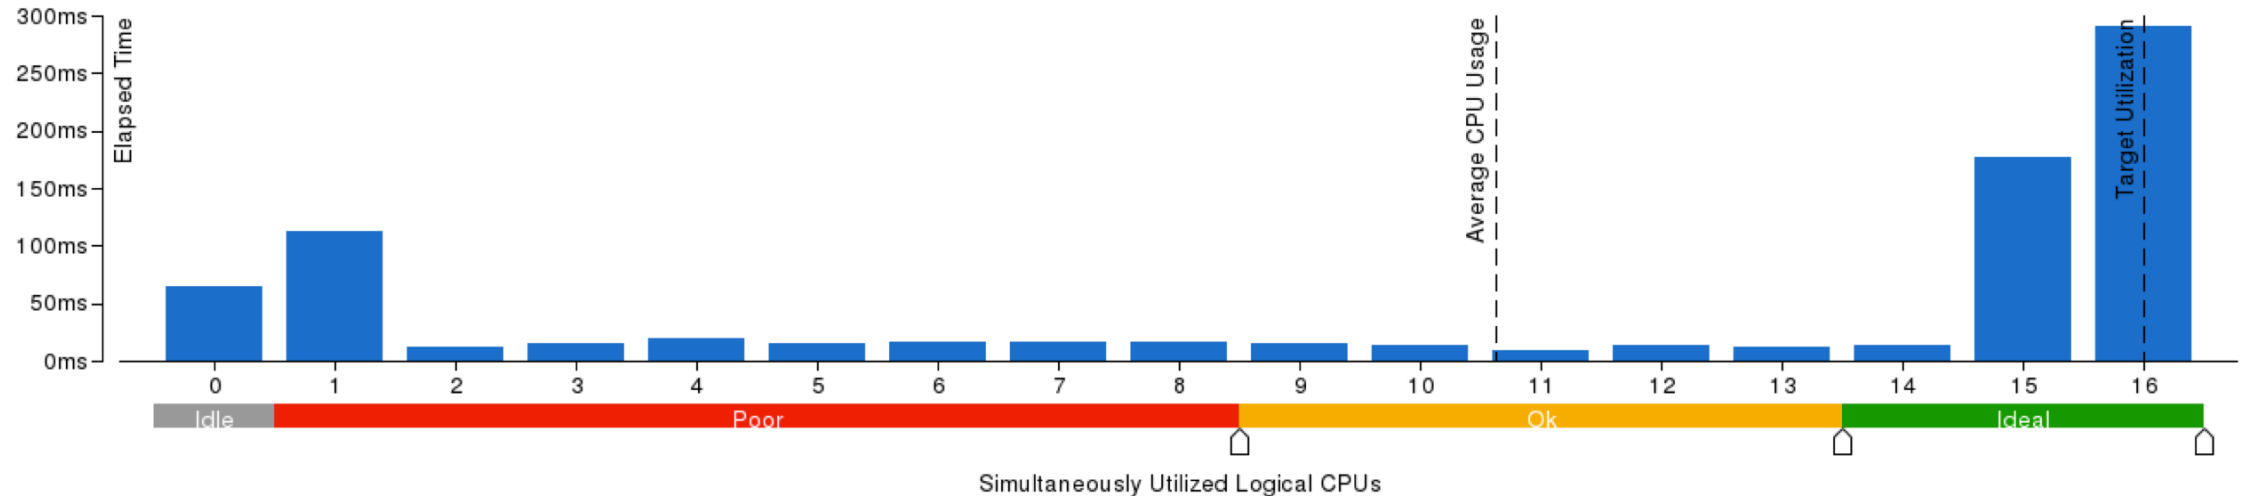
Performance Optimization: n-body

Elapsed Time [?]: 0.842s

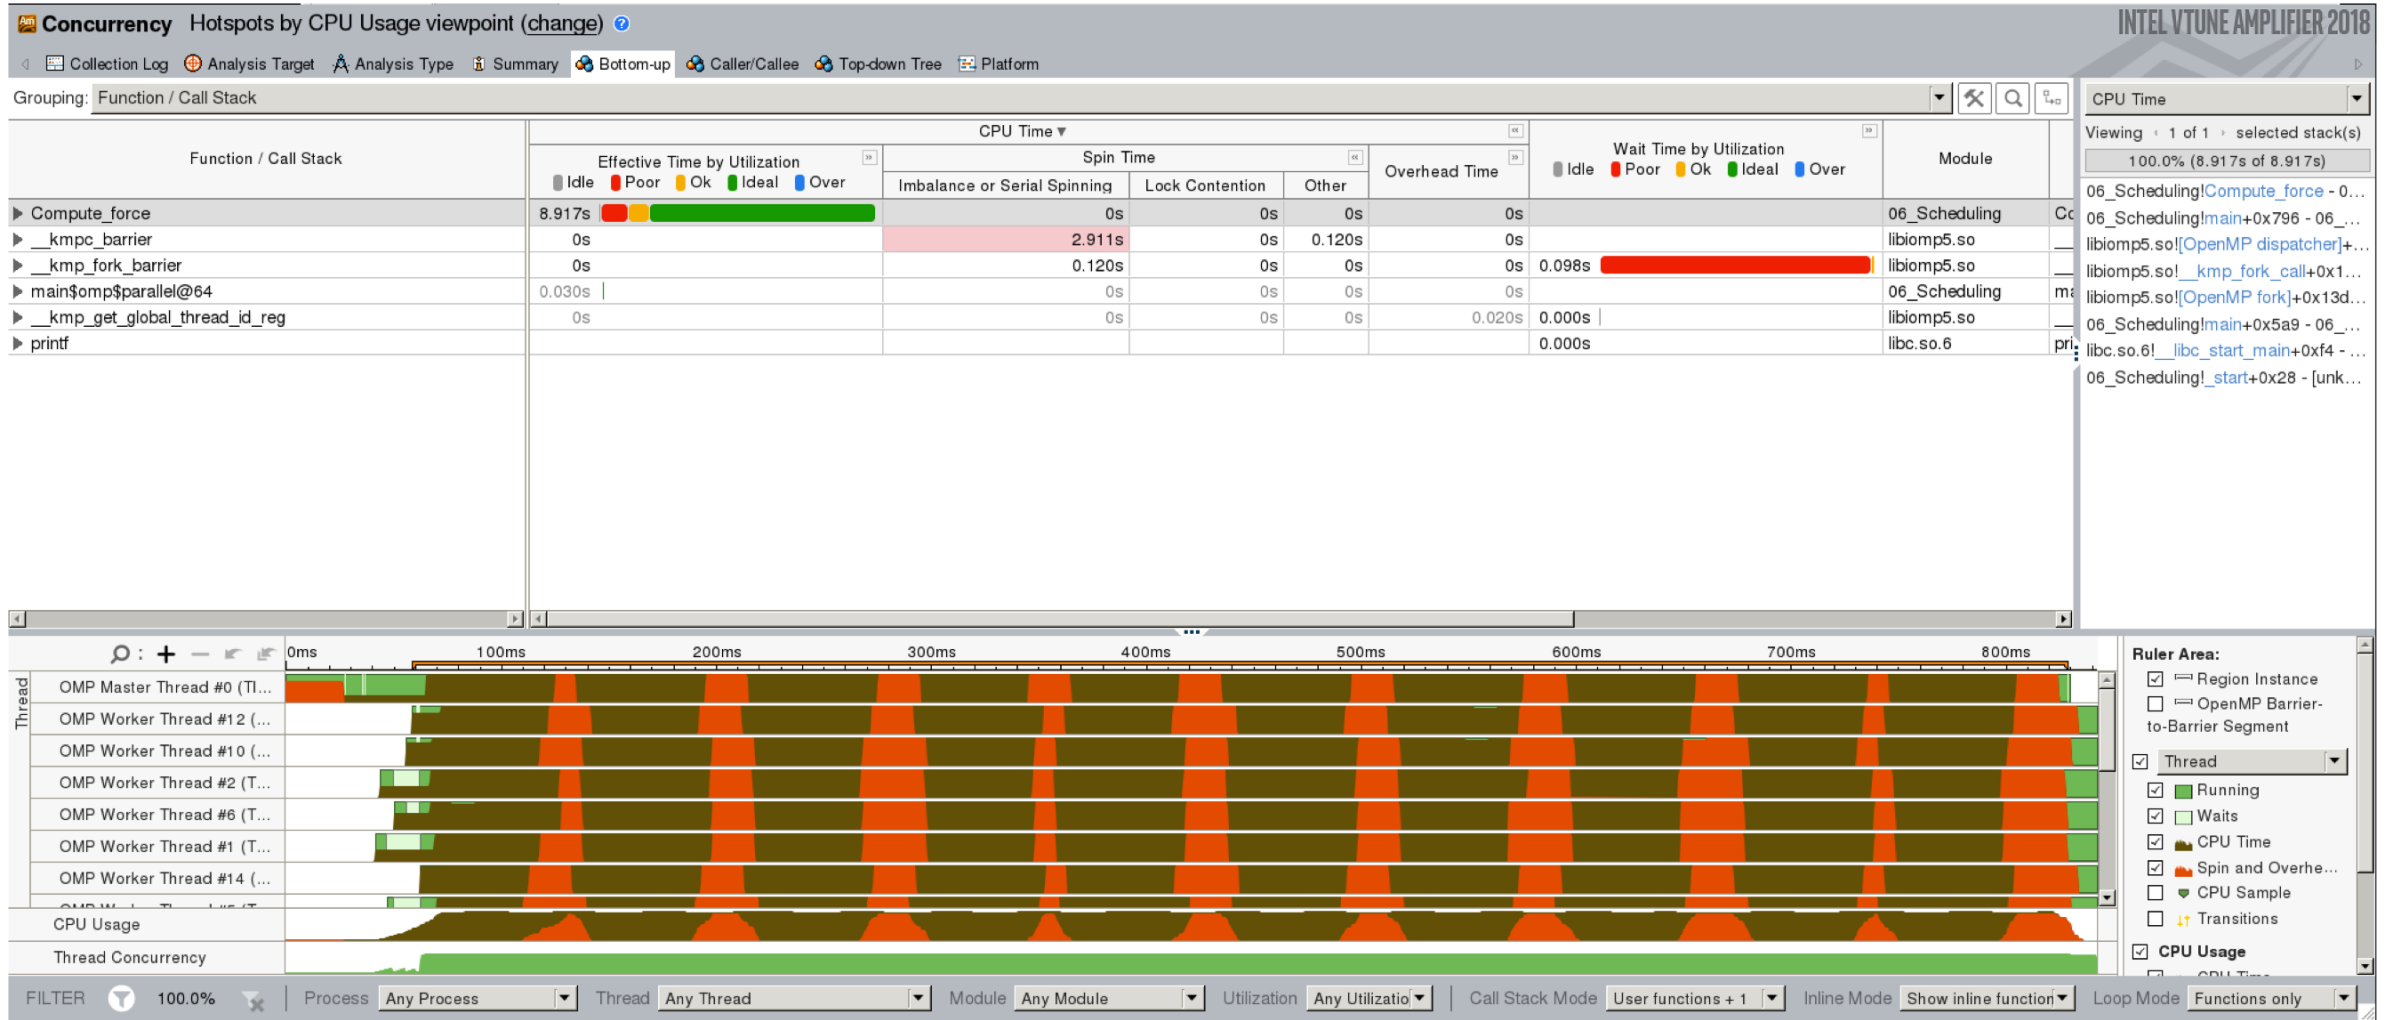
CPU Time [?] :	12.118s
Effective Time [?] :	8.947s
Spin Time [?] :	3.151s
Imbalance or Serial Spinning [?] :	3.031s
Lock Contention [?] :	0s
Other [?] :	0.120s
Overhead Time [?] :	0.020s
Wait Time [?] :	0.098s
Total Thread Count:	16
Paused Time [?] :	0s

CPU Usage Histogram

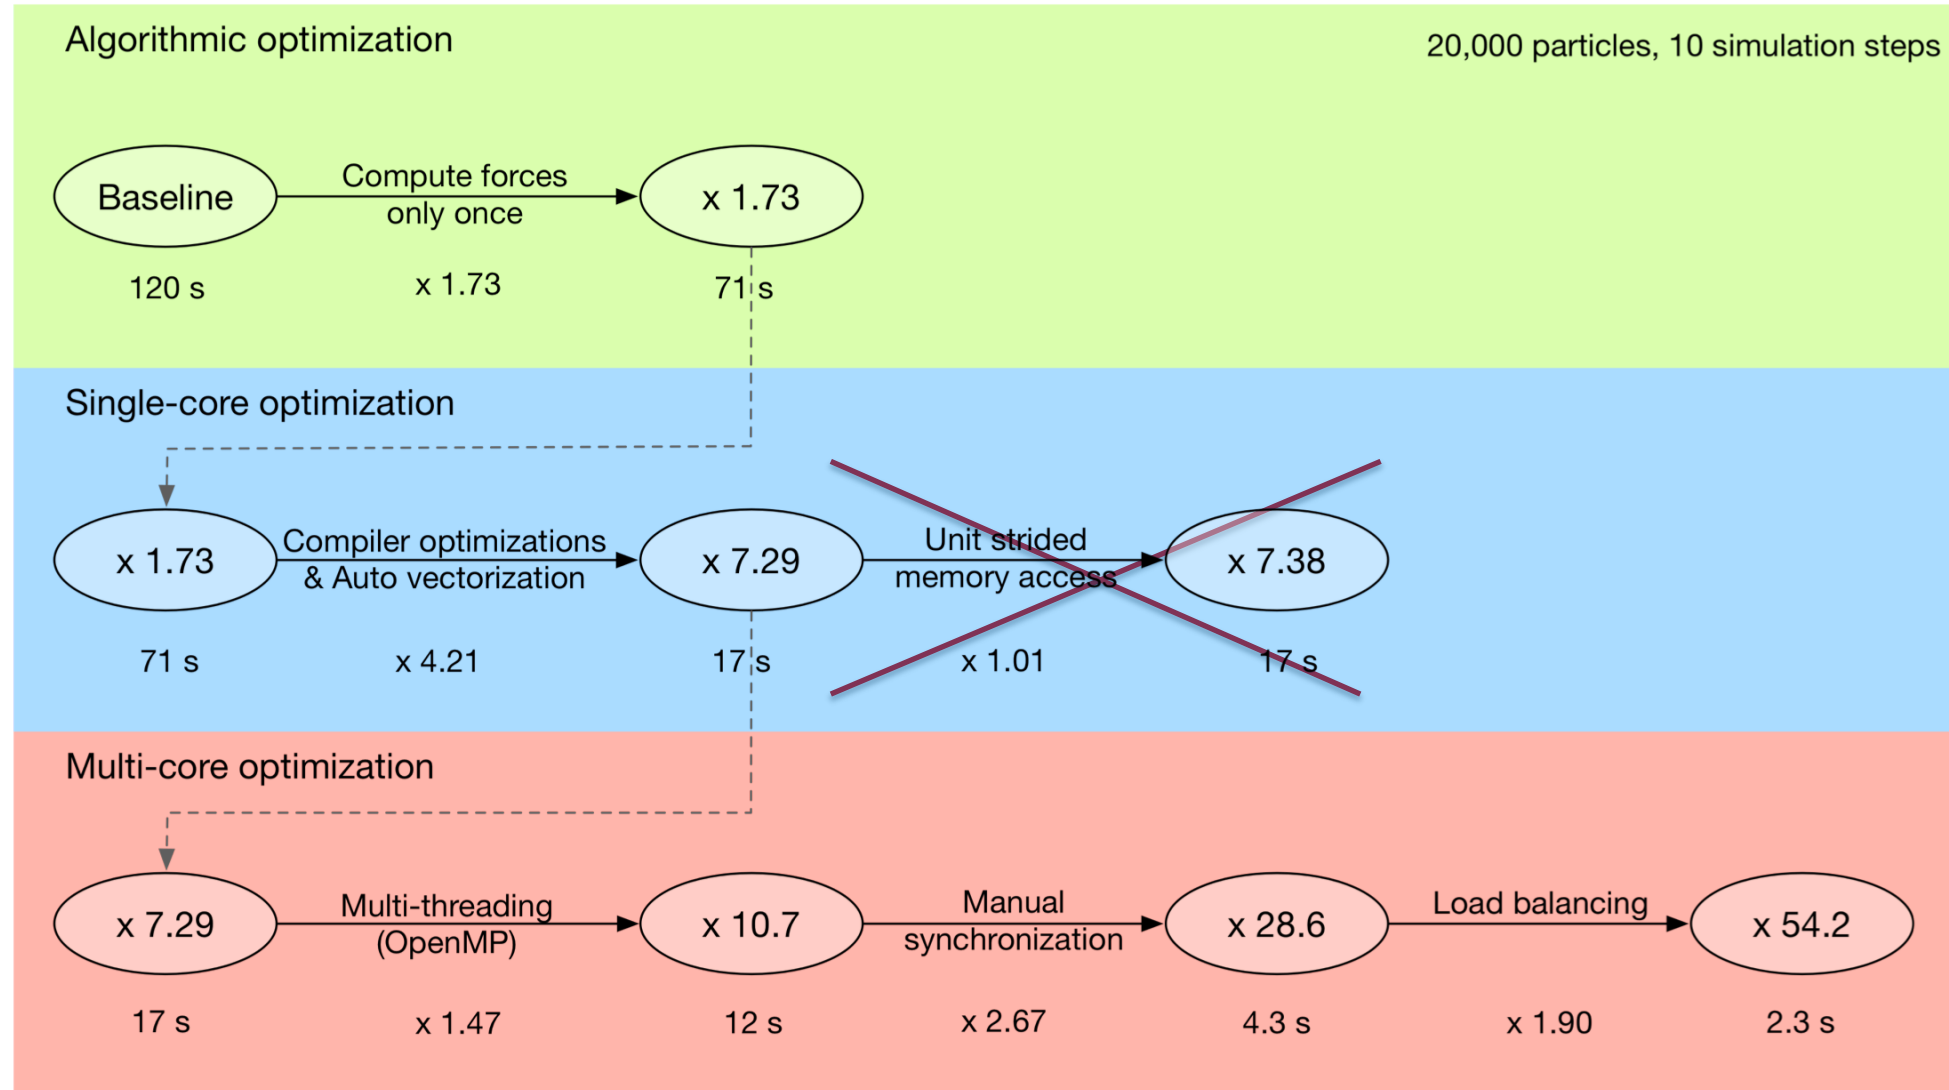
This histogram displays a percentage of the wall time the specific number of CPUs were running simultaneously. Spin and Overhead time adds to the Idle CPU usage value.



Performance Optimization: n-body

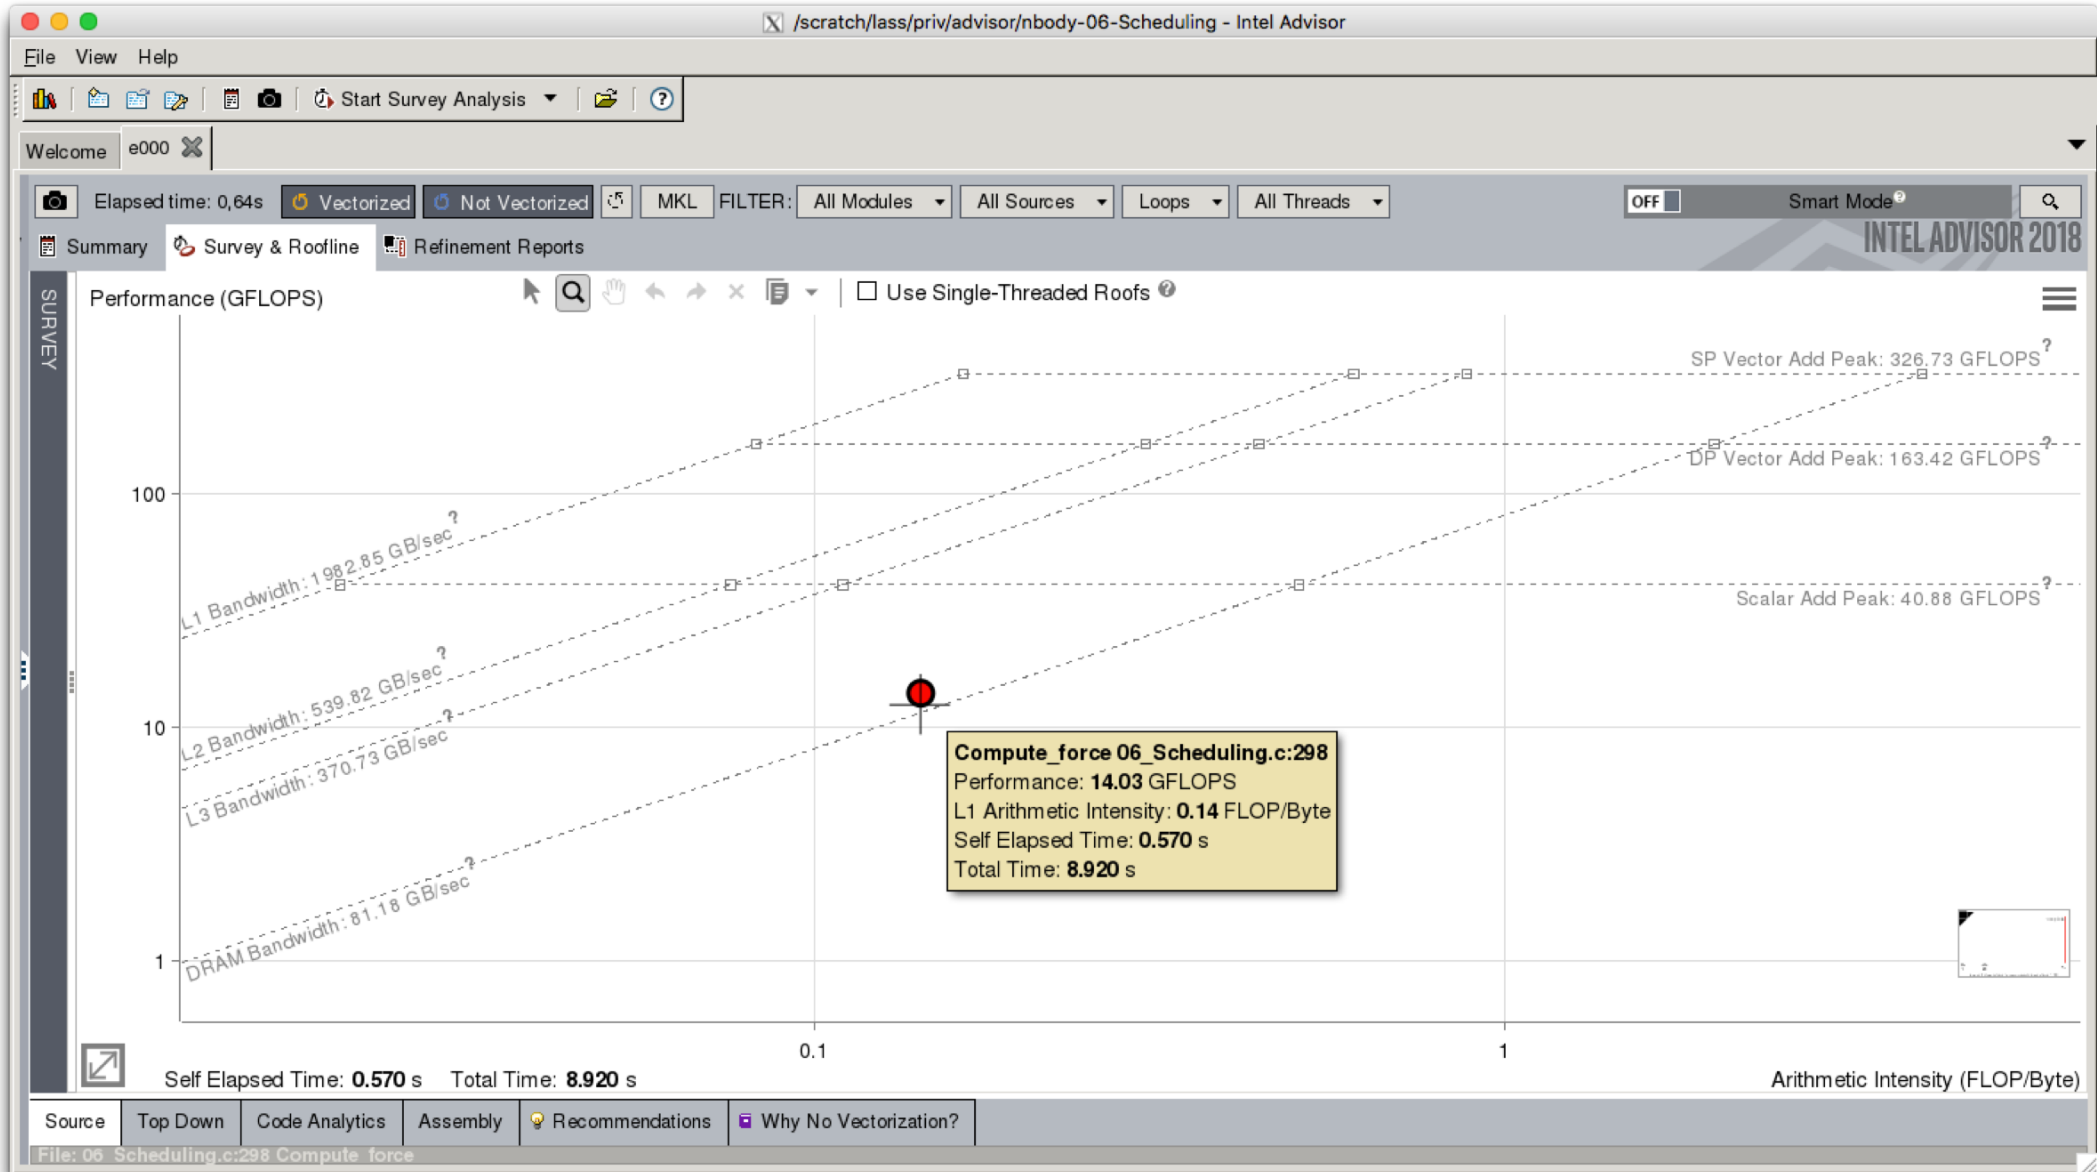


Performance Optimization: n-body



Demo: Back to Intel Advisor

Performance Optimization: n-body



OpenMP: Load balanced version






Performance Optimization: n-body



- Somewhere on our way we lost vectorization

```
for (k = part+1; k < n; k++) {  
  [loop in Compute_force at 06_Scheduling.c:298]  
  Scalar loop. Not vectorized: vector dependence prevents vectorization  
  No loop transformations applied
```

- Advisor can do a memory dependency analysis

Performance Optimization: n-body

Problems and Messages						
ID		Type	Site Name	Sources	Modules	State
P1		Parallel site information	loop_site_1	06_Scheduling.c	06_Scheduling	✓ Not a problem
P3		Read after write dependency	loop_site_1	06_Scheduling.c	06_Scheduling	🚫 New
P4		Read after write dependency	loop_site_1	06_Scheduling.c	06_Scheduling	🚫 New
P5		One task in parallel site	loop_site_1	06_Scheduling.c	06_Scheduling	🚫 New

Read after write dependency: Code Locations							
ID	Instruction Address	Description	Source	Function	Variable references	Module	State
▼ X4	0x402128	Read	 06_Scheduling.c:314	Compute_force		06_Scheduling	🚫 New
<pre>312 313 /* Add force into total forces */ 314 forces[part][X] += f_part_k[X]; 315 forces[part][Y] += f_part_k[Y]; 316 forces[k][X] -= f_part_k[X];</pre>							
▼ X5	0x402135	Write	 06_Scheduling.c:314	Compute_force		06_Scheduling	🚫 New
<pre>312 313 /* Add force into total forces */ 314 forces[part][X] += f_part_k[X]; 315 forces[part][Y] += f_part_k[Y]; 316 forces[k][X] -= f_part_k[X];</pre>							

Performance Optimization: n-body

- Solution: #pragma omp simd reduction(...)

```
double forces_accu_X = 0.0;
double forces_accu_Y = 0.0;

#pragma omp simd reduction(+:forces_accu_X,forces_accu_Y)
for (k = part+1; k < n; k++) {

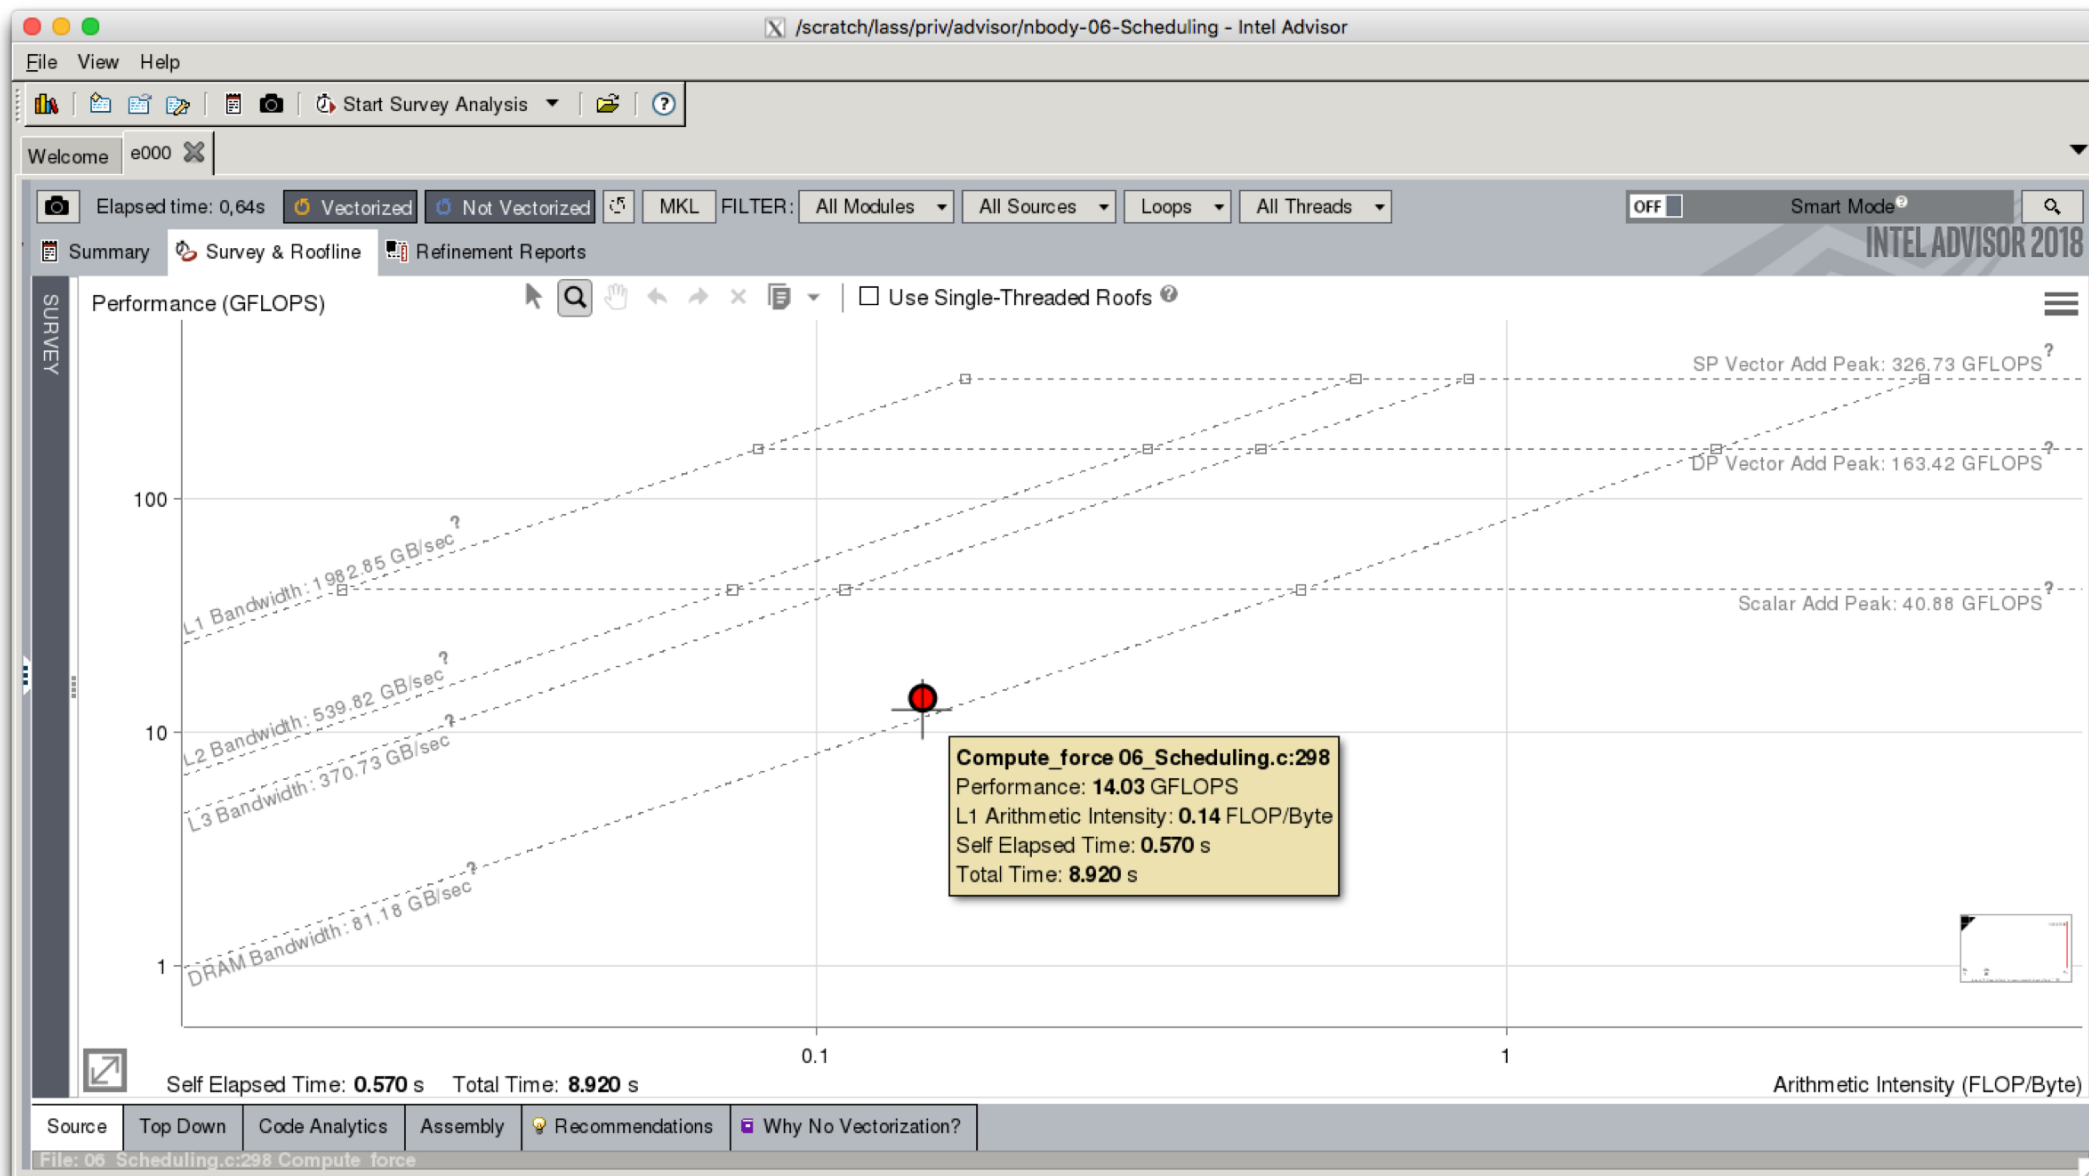
    /* Compute force between part and k */
    [...]

    /* Add force into total forces */
    forces_accu_X += f_part_k[X];
    forces_accu_Y += f_part_k[Y];
    forces[k][X] -= f_part_k[X];
    forces[k][Y] -= f_part_k[Y];

}

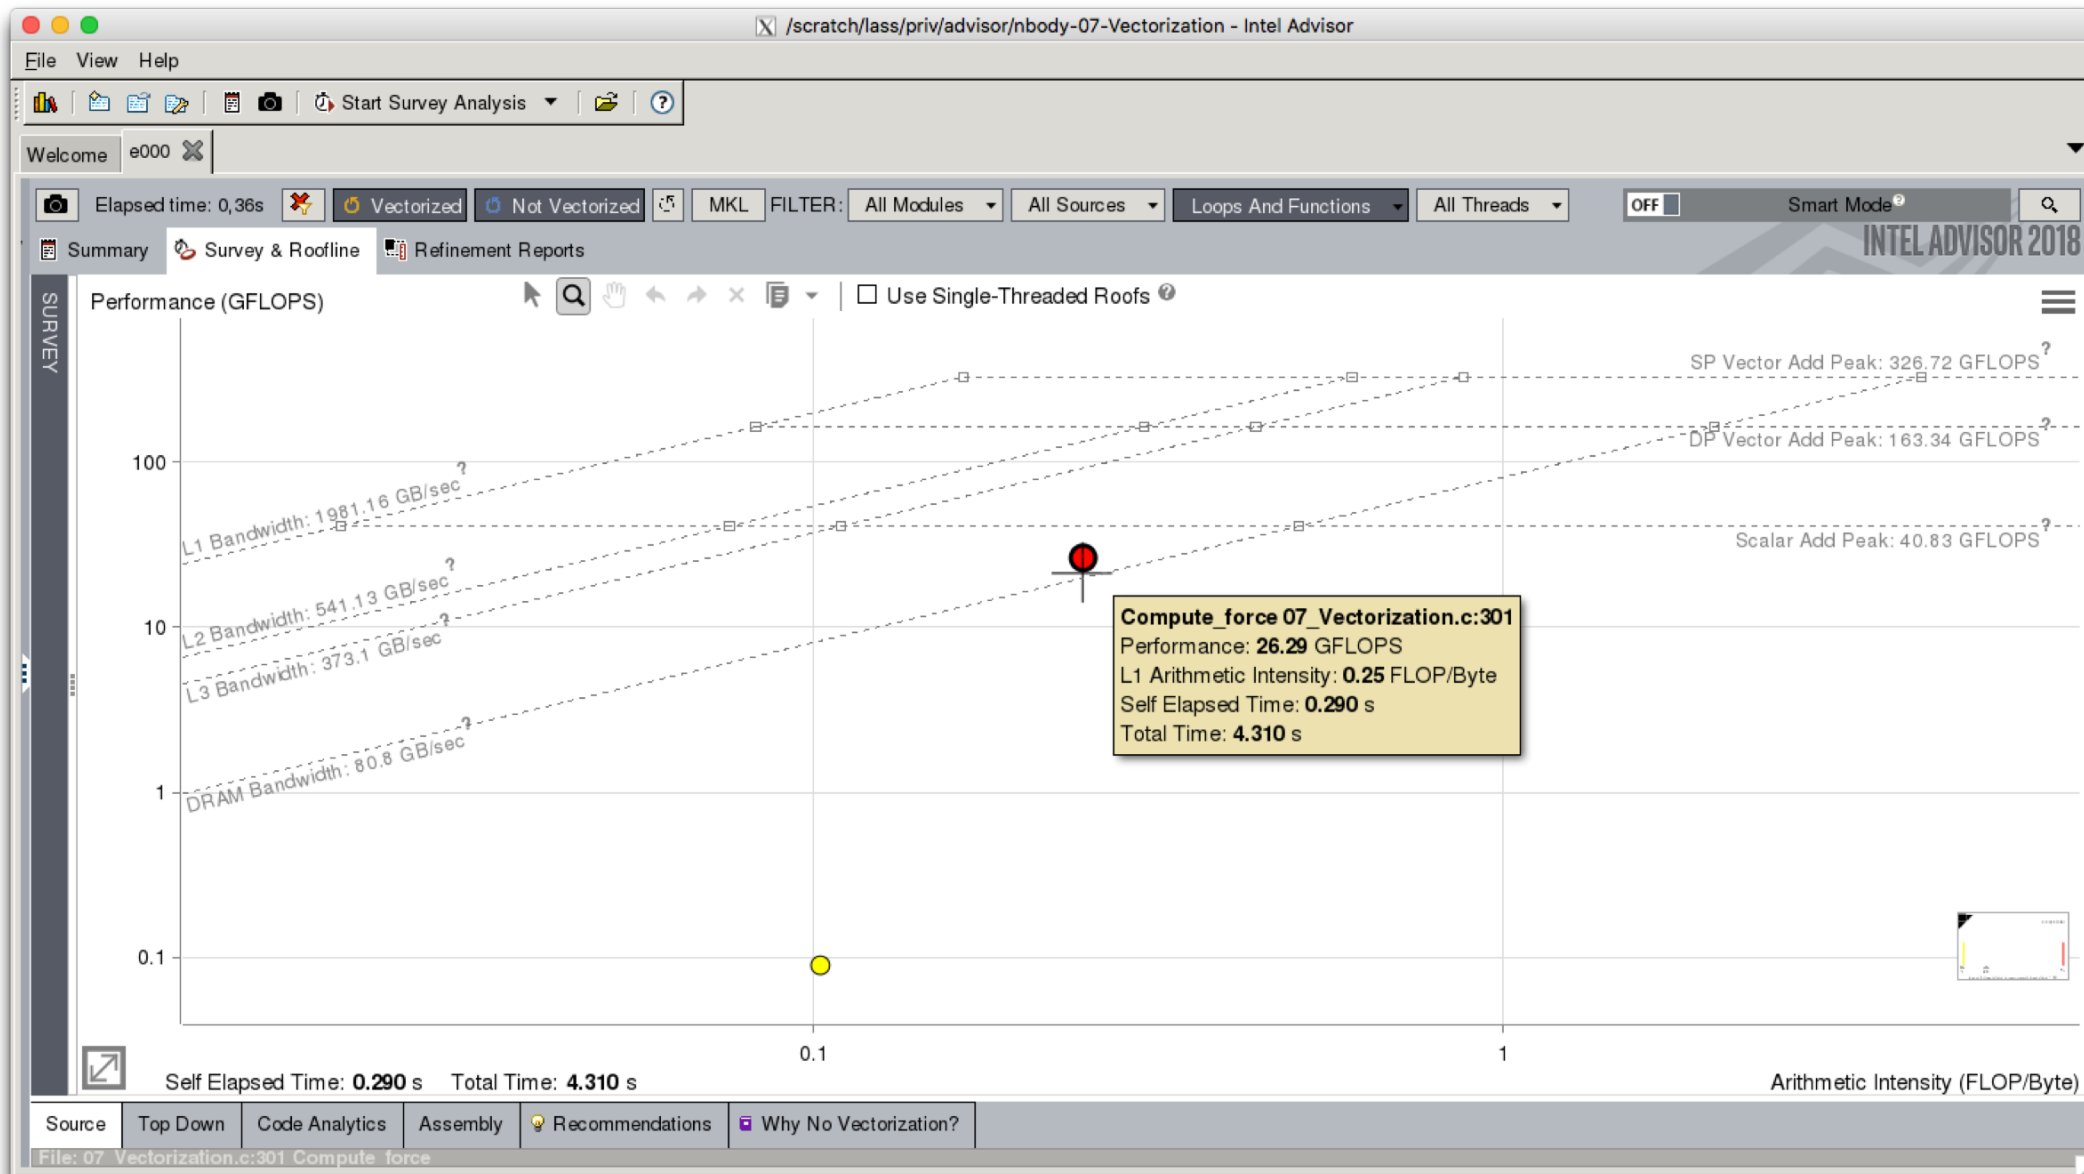
forces[part][X] += forces_accu_X;
forces[part][Y] += forces_accu_Y;
```

Performance Optimization: n-body



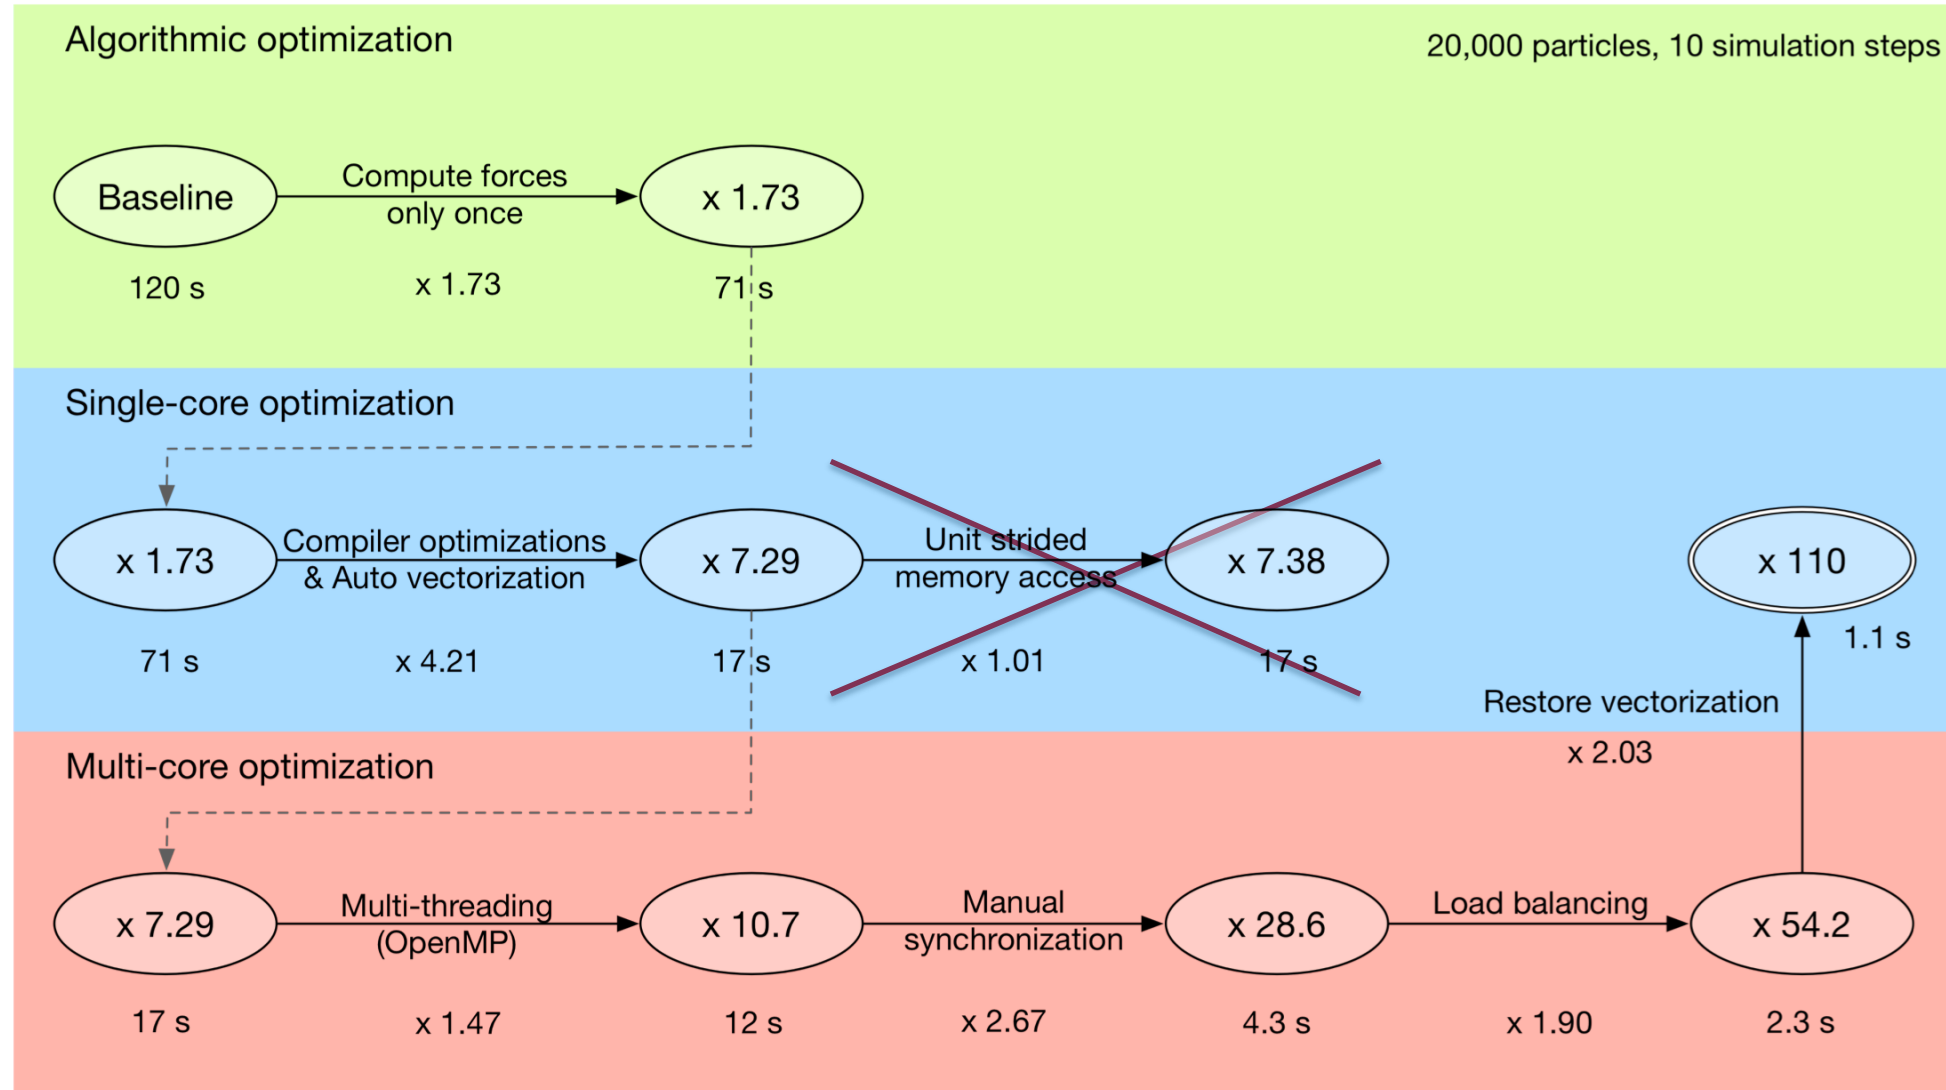
OpenMP: Load balanced version

Performance Optimization: n-body



OpenMP: Load balanced and vectorized version

Performance Optimization: n-body



Performance Optimization: n-body

- Still, we only achieve 26 GFLOP/s where 163 GFLOP/s should be possible
- Are we memory or compute limited?
- Do a *General Exploration* in VTune

Demo: Intel VTune General Exploration

Performance Optimization: n-body

⌵	Elapsed Time [?] : 8.359s	
	Clockticks:	270,202,660,000
	Instructions Retired:	160,763,720,000
	CPI Rate [?] :	1.681 ▴
	MUX Reliability [?] :	1.000
⌵	Filled Pipeline Slots:	
⌵	Unfilled Pipeline Slots (Stalls):	
⌵	Back-End Bound [?] :	79.5% ▴ of Pipeline Slots
⌵	Memory Latency:	
⌵	Memory Replacements:	
⌵	Memory Reissues:	
	Divider [?] :	63.3% ▴ of Clockticks
	Flags Merge Stalls [?] :	0.0% of Clockticks
	Slow LEA Stalls [?] :	0.0% of Clockticks
⌵	Front-End Bound [?] :	1.4% of Pipeline Slots
	Total Thread Count:	16
	Paused Time [?] :	0s

We spend > 60% of the time in our division!

Performance Optimization: n-body

- We achieved a 110x speedup by
 - Optimizing the algorithm
 - Using compiler optimizations
 - Using OpenMP for parallel execution
 - Reducing the synchronization effort
 - Balancing the load between threads
 - Making use of vectorization
- Tools can help in this process
 - We used VTune to find bottlenecks and load imbalance
 - We used Advisor to enable vectorization and optimize memory access patterns
- Optimization is an iterative process
 - You may have to revisit things you have optimized earlier

Profiling of MPI applications

Profiling of MPI applications

- Questions when optimizing an MPI application:
 - How much time is spent for communication?
 - How much data is transferred between the nodes?
 - Is the load balanced between all ranks?
- MPI implementations provide statistics and benchmarks to assist
- We focus on Intel MPI here

Profiling of MPI applications

- To collect basic statistics, export environment variable I_MPI_STATS

```
module load ps_xe_2018  
export I_MPI_STATS=all
```

```
# local execution
```

```
mpirun -n 16 ./mpi_nbody_red 1024 10 0.001 10 g
```

```
# execution on OCuLUS
```

```
ccsalloc -l -c 64 impi -- ./mpi_nbody_red 524288 10 0.001 10 g
```

Profiling of MPI applications

- stats.ipm
- IPM: Integrated Performance Monitoring

```
# [total] <avg> min max
# entries 64 1 1 1
# wallclock 21905.1 342.267 341.467 344.917
# user 21436.7 334.948 332.975 335.58
# system 23.0293 0.359832 0.180638 0.881514
# mpi 1231.47 19.2418 7.94775 33.3769
# %comm 5.62186 2.31862 9.74161
# gflop/sec NA NA NA NA
# gbytes 0 0 0 0
#
# [time] [calls] <%mpi> <%wall>
# MPI_Sendrecv_replace 753.383 40960 61.18 3.44
# MPI_Init 470.983 64 38.25 2.15
# MPI_Scatterv 6.15997 128 0.50 0.03
# MPI_Bcast 0.94546 384 0.08 0.00
# MPI_Type_commit 0.000687122 128 0.00 0.00
# MPI_Finalize 0.000604153 64 0.00 0.00
# MPI_Type_free 0.000304937 128 0.00 0.00
# MPI_Type_contiguous 0.000223637 64 0.00 0.00
# MPI_Comm_size 6.91414e-05 64 0.00 0.00
# MPI_Type_vector 6.31809e-05 64 0.00 0.00
# MPI_Comm_rank 6.24657e-05 64 0.00 0.00
# MPI_Type_create_resized 5.38826e-05 64 0.00 0.00
# MPI_Type_get_extent 1.40667e-05 64 0.00 0.00
# MPI_Wtime 1.26362e-05 128 0.00 0.00
# MPI_TOTAL 1231.47 42368 100.00 5.62
```

Profiling of MPI applications

- stats.txt
- detailed per-rank statistics

```
~~~~ Process 0 of 64 on node node01-002 lifetime = 344917304.99

Data Transfers
Src   Dst   Amount (MB)   Transfers
-----
000 -->000   0.000000e+00    0
000 -->001   4.000028e+00   348
000 -->002   0.000000e+00    0
000 -->003   0.000000e+00    0
000 -->004   0.000000e+00    0
[...]
000 -->063   1.600000e+02   640

[...]

Communication Activity
Operation   Volume (MB)   Calls   Min time   Avr time   Max time   Total time
-----
P2P
Csend       4.000028e+00   348     0.95      8.58      366.93    2985.95
CSendRecv   0.000000e+00    0       0.00      0.00      0.00      0.00
Send        0.000000e+00    0       0.00      0.00      0.00      0.00
SendRecv    1.600000e+02   640     445.13    24101.42  74103.12  15424910.31
Bsend       0.000000e+00    0       0.00      0.00      0.00      0.00
Rsend       0.000000e+00    0       0.00      0.00      0.00      0.00

[...]
```


Profiling of MPI applications

- Intel Trace Analyzer: Tool to visualize and analyze MPI traces
- Enable collection of traces by adding `-trace` to `mpirun`:

```
module load ps_xe_2018

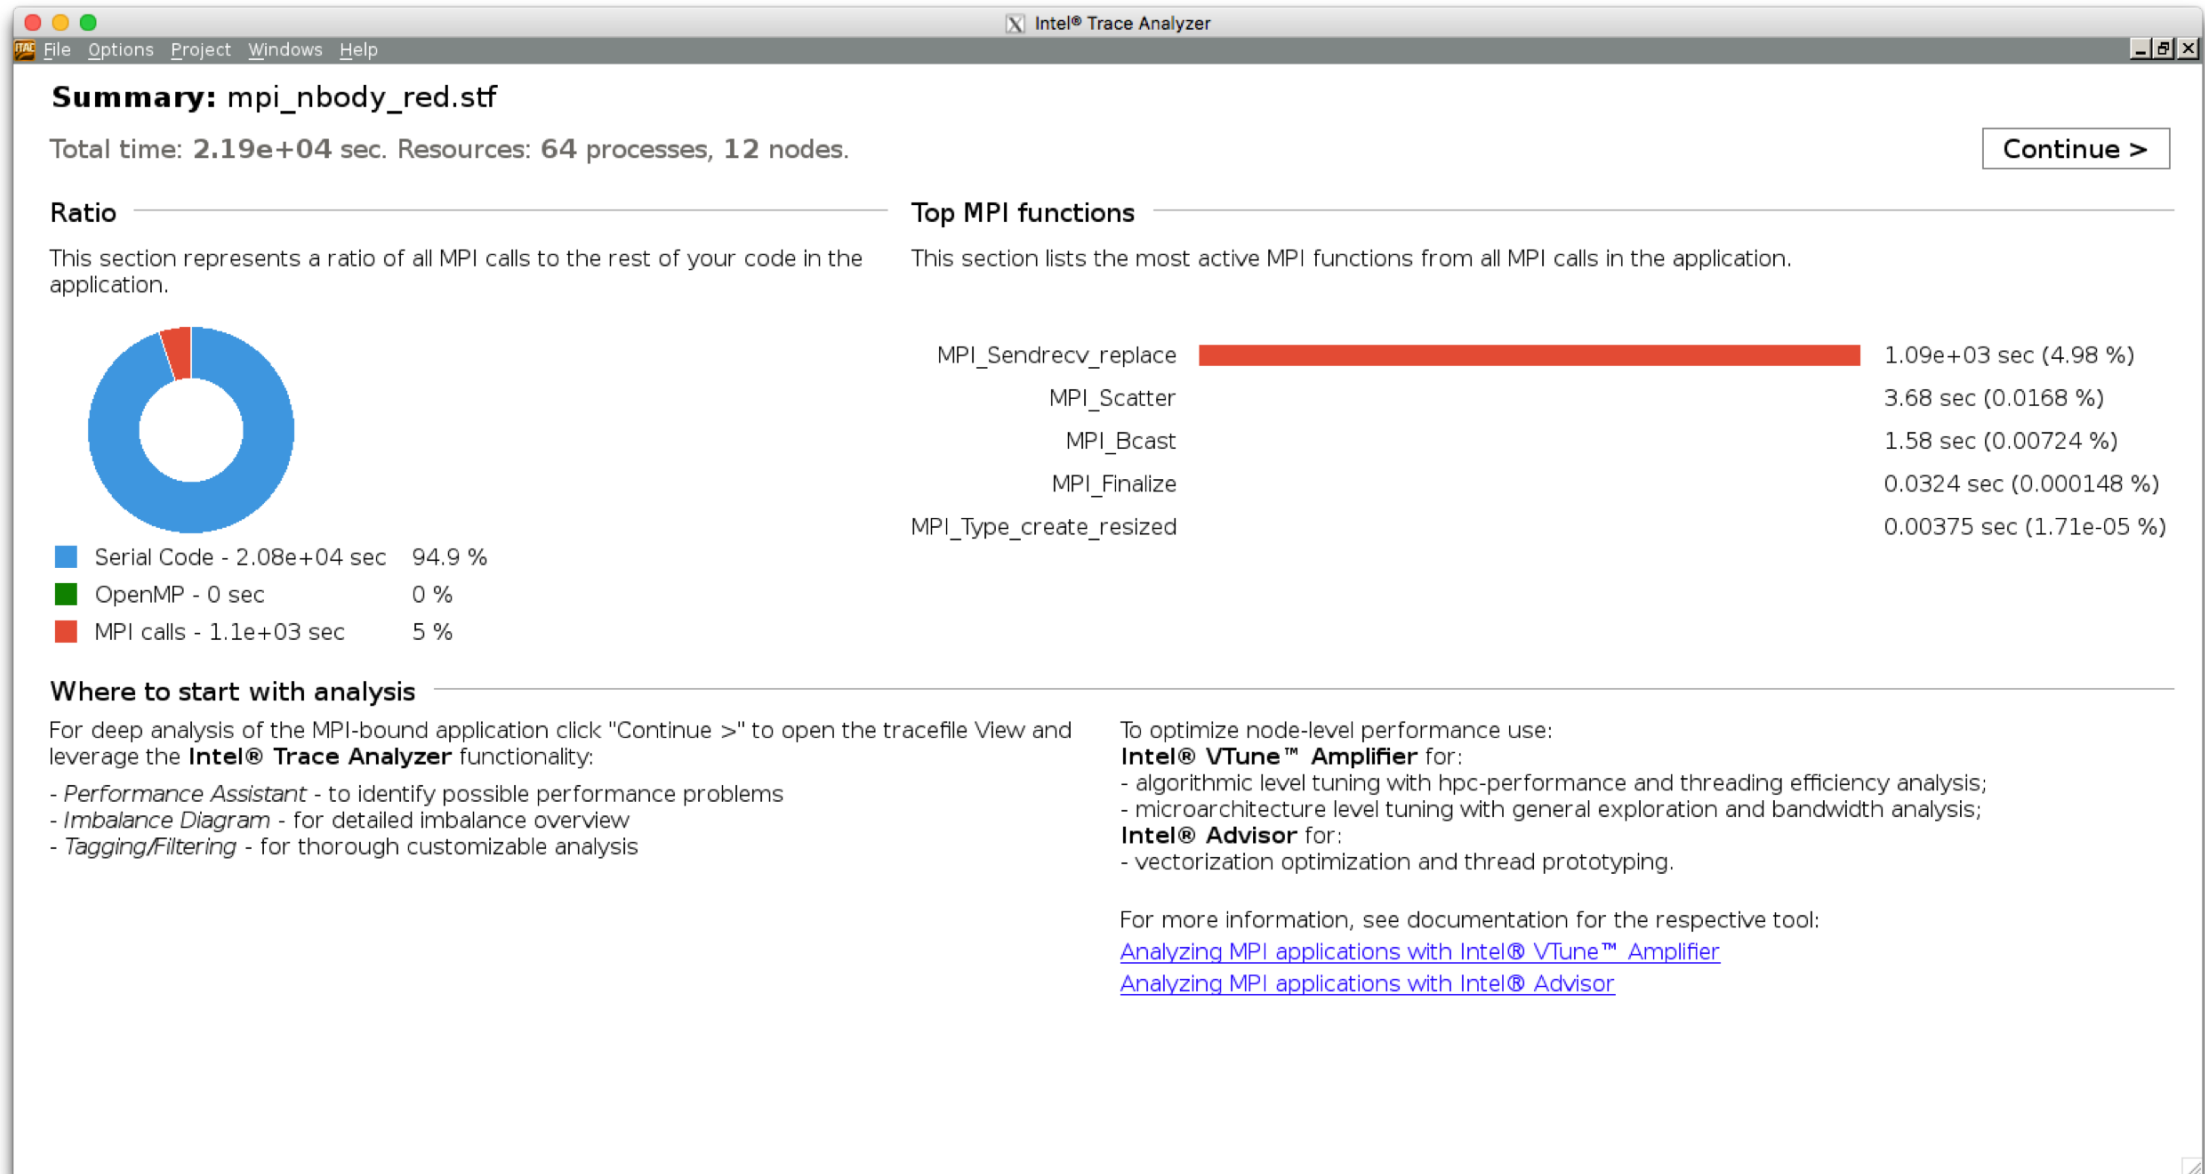
# local execution
mpirun -trace -n 16 ./mpi_nbody_red 1024 10 0.001 10 g

# execution on OCuLUS
ccsalloc -l -c 64 impi -trace -- ./mpi_nbody_red 524288 10 0.001 10 g

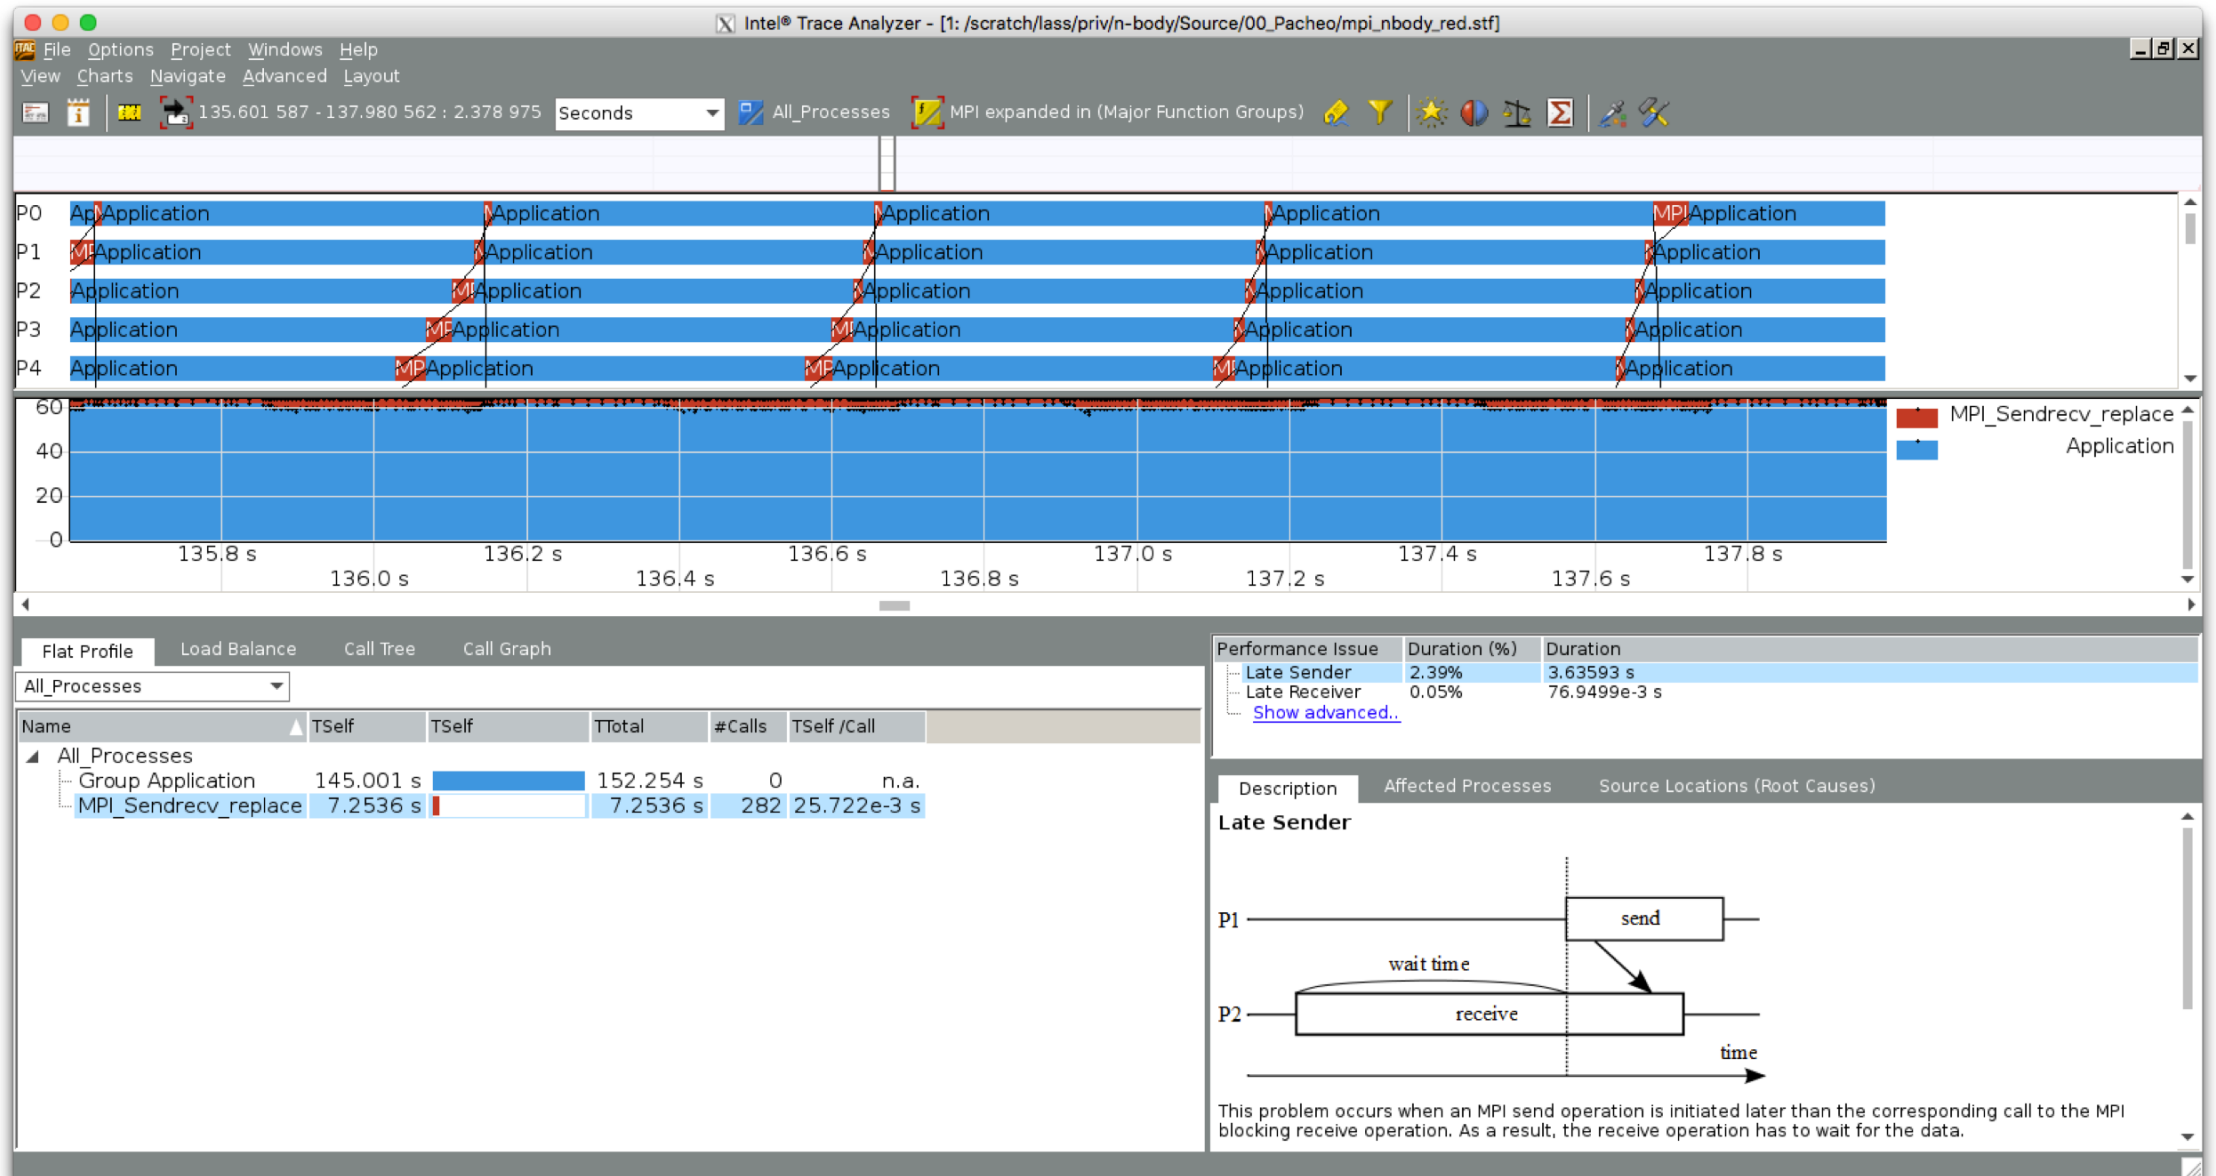
# start trace analyzer
traceanalyzer mpi_nbody_red.stf
```

Demo: Intel Trace Analyzer

Profiling of MPI applications



Profiling of MPI applications



Acknowledgements

- This lecture is based materials from these sources
 - Tutorial: Performance Tuning of Scientific Codes with the Roofline Model, SuperComputing 2017
 - Tutorial: Advanced OpenMP, Supercomputing 2017
 - Tutorial: Application Optimization and Vectorization, Intel 2017
 - Intel 64 and IA-32 Architectures Optimization Reference Manual
 - Intel Developer Zone documentation on vectorization

Change Log

- **1.1.0 (2018-02-02)**
 - added slides about MPI profiling
- **1.0.0 (2018-01-30)**
 - initial version of slides