



TRACKING DOWN MEMORY BOTTLENECKS WITH THE INTEGRATED ROOFLINE MODEL IN INTEL® ADVISOR


Cedric Andreolli, Technical Consulting Engineer, Intel

ISO3DFD

A memory bounded application ?

How does a stencil work ?

- For computing $P_{t+1}(x,y,z)$, we need to use all neighbors in the 3 dimensions of $P_t(x,y,z)$
- The stencil looks like a 3D cross in Iso3DFD



Iso3DFD – A wave propagation kernel

```
For (int iz=0; iz<n3; iz++)
For (int iy=0; iy<n2; iy++)
For (int ix=0; ix<n1; ix++) {
    int offset = iz*dimn1n2 + iy*n1 + ix;
    float value = 0.0;
    value += ptr_prev[offset]*coeff[0];
    for(int ir=1; ir<= 8 ; ir++) {
        value += coeff[ir] * (ptr_prev[offset + ir] + ptr_prev[offset - ir]);
        value += coeff[ir] * (ptr_prev[offset + ir*n1] + ptr_prev[offset - ir*n1]);
        value += coeff[ir] * (ptr_prev[offset + ir*dimn1n2] + ptr_prev[offset -
ir*dimn1n2]);
    }
    ptr_next[offset] = 2.0f* ptr_prev[offset] - ptr_next[offset] + value*ptr_vel[offset];
}
```

- 3D Finite Difference
- Acoustic isotropic, pressure only scheme,
- 16th order in space 2nd order in time
- No boundary conditions
- OpenMP (no MPI)

DEV01 – INTEL[®] VTUNE HOTSPOTS

Is everything as expected ?

Running the Integrated Roofline Model on Dev00

Dev00 is an unoptimized implementation of iso3DFD

- It's time for a demo
- Intel® Vtune™ Hotspots

time: 196.61 sec
throughput: 7.03 MPoints/s
flops: 0.43 GFlops

Running the Integrated Roofline Model on Dev00

Dev00 is an unoptimized implementation of iso3DFD

time: 196.61 sec
throughput: 7.03 MPoints/s
flops: 0.43 GFlops

time: 10.37 sec
throughput: 133.33
MPoints/s
flops: 8.13 GFlops

DEV02 – FIRST STEP WITH INTEGRATED ROOFLINE MODEL

Is everything as expected ?

Running the Integrated Roofline Model on Dev00

Dev01 only implements the threading

- It's time for a demo
- Intel® Advisor's new summary
- Column configurator
- Integrated Roofline Model
- Intel® VTune™ Memory Access

time: 10.37 sec
throughput: 133.33 MPoints/s
flops: 8.13 GFlops

Creation of dev02

Based on Intel® Advisor analysis, we need to reverse the loops

- Reversing the loops generated unit strides
- The performance is improving

time: 10.37 sec
throughput: 133.33 MPoints/s
flops: 8.13 GFlops

2.3x speedup

time: 13.43 sec
throughput: 439.59 MPoints/s
flops: 26.82 GFlops

DEV03 – VECTORIZATION

Following recommendations

Enabling vectorization

Dev02 has still many room for improvement


- It's time for a demo
- Intel Advisor's new summary
- Intel[®] Advisor's recommendations
- Roofline Compare
- Dependency analysis

time: 13.43 sec
throughput: 439.59 MPoints/s
flops: 26.82 GFlops

Dev03 – Vectorization done

Dev03 is enabling vectorization

- We only need to use a simple `#pragma omp simd` here




DEV04 – CACHE BLOCKING

What does the Roofline say ?

Enabling Cache Blocking

Dev03 has still many room for improvement

- It's time for a demo
- Integrated Roofline
- Check the position of the dot to their respective roofline
- Roofline Compare



Dev04 – Cache Blocking done

Just specifying export KMP_AFFINITY=scatter

time: 3.89 sec
throughput: 1518.47 MPoints/s
flops: 92.63 GFlops


1.1x speedup

time: 3.50 sec
throughput: 1686.58 MPoints/s
flops: 102.88 GFlops

CONCLUSION

Conclusion

- **13x speedup but the initial application was suboptimal**
- **Intel® Advisor helped for:**
 - **Solving memory accesses problems**
 - **Improving vectorization**
 - **Detecting L3 bottleneck**
- **Next step: Recommendation engine based on Integrated Roofline Model**



Legal Disclaimer & Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS". NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Copyright © 2015, Intel Corporation. All rights reserved. Intel, Pentium, Xeon, Xeon Phi, Core, VTune, Cilk, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

QUESTIONS & ANSWERS



**TECH.
DECODED**

