

Grid Computing

2012/1/30

金光浩(11D37060 松岡研究室)

PATUS: A Code Generation and Auto-Tuning Framework For Parallel Stencil Computations

Shoaib Kamilyz, Cy Chany, Leonid Olikery,
John Shalfy, Samuel Williamsy

In Parallel & Distributed Processing Symposium (IPDPS), 2010
IEEE International

松岡研究室
金光浩 11D37060

Overview

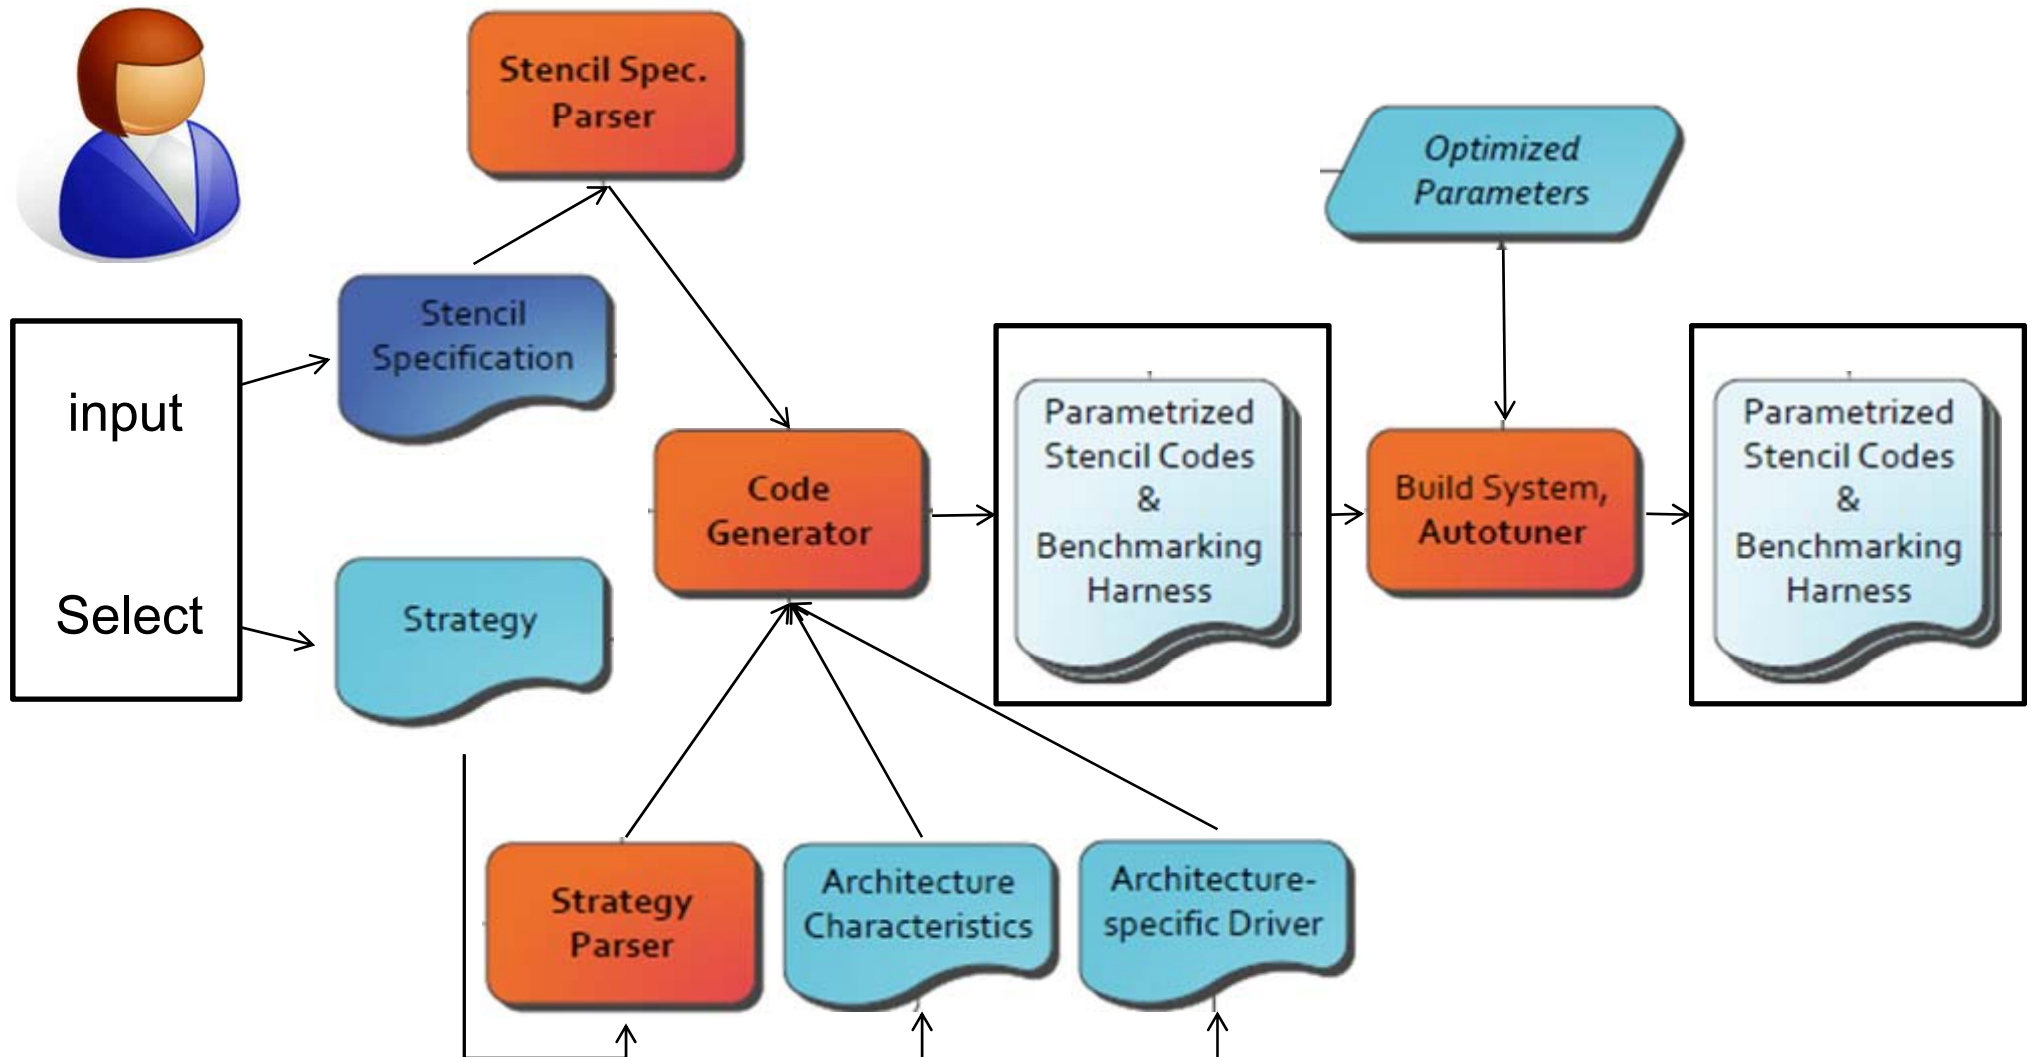
PATUS is a **code generation** and **auto-tuning framework** for **stencil computations** targeted at modern multi- and many-core processors, such as multi-core **CPUs** and **GPUs**.



Using a small domain specific language (**DSL**), the **user defines the stencil kernel** using a C-like syntax. The framework **generates the code** for a compute kernel from a specification of the stencil operation and a Strategy:

They **leverage** the auto-tuning methodology to find the optimal hardware **architecture-specific** and **Strategy-specific parameter configuration**.

Process



Examples-Input

Source path: [svn/ trunk/ examples/ edge.stc](#)

```
1 stencil edge
2 {
3     domainsize = (1 .. width, 1 .. height);
4     t_max = 1;
5
6     operation (float grid u)
7     {
8         u[x, y; t+1] =
9             -12 * u[x, y; t] +
10             2 * (u[x - 1, y; t] + u[x + 1, y; t] + u[x, y - 1; t] + u[x, y + 1; t]) +
11             u[x - 1, y - 1; t] + u[x + 1, y - 1; t] + u[x - 1, y + 1; t] + u[x + 1, y + 1; t];
12     }
13 }
```

Source path: [svn/ trunk/ examples/ wave-float.stc](#)

```
1 stencil wave
2 {
3     domainsize = (1 .. x_max, 1 .. y_max, 1 .. z_max);
4     t_max = 1;
5
6     operation (float grid u, float param dt_dx_sq)
7     {
8         u[x, y, z; t+1] = 2 * u[x, y, z; t] - u[x, y, z; t-1] +
9             dt_dx_sq * (
10                 -15/2 * u[x, y, z; t] +
11                 4/3 * (
12                     u[x+1, y, z; t] + u[x-1, y, z; t] +
13                     u[x, y+1, z; t] + u[x, y-1, z; t] +
14                     u[x, y, z+1; t] + u[x, y, z-1; t]
15                 )
16                 -1/12 * (
17                     u[x+2, y, z; t] + u[x-2, y, z; t] +
18                     u[x, y+2, z; t] + u[x, y-2, z; t] +
19                     u[x, y, z+2; t] + u[x, y, z-2; t]
20                 )
21             );
22     }
23 }
```

1.A rectangular domain

2. Define Initial condition by user

3. Working on boundary condition

Strategy

Source path: svn/

Directories	Filename
▼ svn	cacheblocked.stg
branches	
tags	gpu.stg
▼ trunk	simple.stg
.settings	
▼ arch	
CPU_OpenMP	
CPU_OpenMP_PAPI	
GPU_CUDA	
bin	
▶ doc	
examples	
lib	
runtime	
▶ src	
strategy	
▶ tools	
wiki	

Strategy-1

Source path: [svn/](#) [trunk/](#) [strategy/](#) simple.stg

```
1  /**
2  * Naive parallel strategy using coordinates:
3  * Iterate over the grid and parallelize the outer most loop.
4  */
5  strategy naive_parallel (domain u, auto int chunk)
6  {
7      // do all the timesteps
8      for t = 1 .. stencil.t_max
9      {
10         // apply the stencil to all the points in the domain
11         for point p in u(:, t) parallel schedule chunk
12             u[p; t+1] = stencil (u[p; t]);
13     }
14 }
```

The parameter `chunk` to the `schedule` keyword defines how many consecutive blocks one thread is given.

Strategy-2

Source path: [svn/ trunk/ strategy/ cacheblocked.stg](#)

```
1  /**
2  * Cache-blocking strategy.
3  */
4  strategy cacheblocked_domain (domain u, auto dim cb, auto int chunk)
5  {
6      // iterate over time steps
7      for t = 1 .. stencil.t_max
8      {
9          // iterate over subdomain
10         for subdomain v(cb) in u(:,; t) parallel schedule chunk
11         {
12             // calculate the stencil for each point in the subdomain
13             for point p in v(:,; t)
14                 v[p; t+1] = stencil (v[p; t]);
15         }
16     }
17 }
```

Define the block size[of each dimension] and access order[schedule].

blocks v of size cb over the root domain u

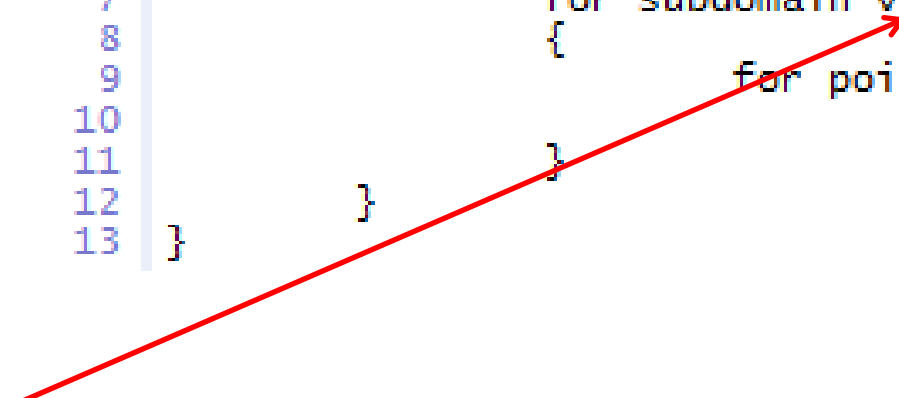
The cb will be interfaced with the auto tuner.

values for $cb = (c1; c2; ::; cd)$, where d is the dimensionality of the stencil.

Strategy-3

Source path: [svn/ trunk/ strategy/ gpu.stg](#)

```
1 strategy gpu_blocked (domain u, auto int cbx)
2 {
3     // iterate over time steps
4     for t = 1 .. stencil.t_max
5     {
6         // iterate over subdomain
7         for subdomain v(cbx, 1 ...) in u(:, t) parallel
8         {
9             for point pt in v(:, t)
10                v[pt; t+1] = stencil (v[pt; t]);
11        }
12    }
13 }
```



1-dimensional thread blocks and grids,
3-dimensional thread blocks and a 2-dimensional grid,
3-dimensional thread blocks and grids

GPU-Platform



patus

Patus is a code-generation and auto-tuning framework for parallel stencil computations.

[Project Home](#)[Downloads](#)[Wiki](#)[Issues](#)[Source](#)[Checkout](#) **Browse** [Changes](#)

Source path: `svn/`

Directories

- ▼ [svn](#)
 - [branches](#)
 - [tags](#)
 - ▼ [trunk](#)
 - [.settings](#)
 - ▼ [arch](#)
 - [CPU_OpenMP](#)
 - [CPU_OpenMP_PAPI](#)
 - [GPU_CUDA](#)
 - [bin](#)
 - ▶ [doc](#)
 - [examples](#)
 - [lib](#)
 - [runtime](#)

Filename

- [Makefile](#)
- [driver.cu](#)

Makefile

Source path: [svn/](#) [trunk/](#) [arch/](#) [GPU_CUDA/](#) Makefile

```
1 #
2 # Makefile for Patus stencil benchmark
3 #
4 # Note: $(PATUS_*) variables will be automatically replaced by the
5 # required runtime files by Patus.
6 #
7
8 CC = nvcc
9 NVCCFLAGS = -O3 -arch=sm_13 -I/home/christen/NVIDIA_GPU_Computing_SDK/C/common/inc
10 #NVCCFLAGS = -O0 -g -arch=sm_13 -I/home/christen/NVIDIA_GPU_Computing_SDK/C/common/inc
11
12 bench: kernel.cu driver.cu $(PATUS_RUNTIME_FILES)
13     $(CC) $(NVCCFLAGS) -o $@ $+
14
15 clean:
16     rm -rf *.o bench
```

Driver.cu

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <stdint.h>
4 #include <cuda.h>
5 #include <cutil.h>
6
7 typedef uint64_t gpu_ptr_t;
8
9 #pragma patus forward_decls
10
11 int main (int argc, char** argv)
12 {
13     int i;
14     cudaError_t res;
15
16     // prepare grids
17     #pragma patus declare_grids
18     #pragma patus allocate_grids
19
20     #pragma patus declare_GPU_grids
21     #pragma patus allocate_GPU_grids
22     #pragma patus copy_grids_to_GPU
23
24     #pragma patus initialize_grids
25     cudaThreadSynchronize ();
26     res = cudaGetLastError ();
27     if (res != cudaSuccess)
28     {
29         printf ("CUDA Error [Initialization]: %s.\n", cudaGetErrorString (res));
30         #pragma patus deallocate_grids
31         cudaThreadExit ();
32         return -1;
33     }
34 }
```

Exchange this part for **input**

or **predefined**

* Not pre-treatment of C language

Java interface

Source path: svn/

- ▼ svn
 - branches
 - tags
 - ▼ trunk
 - .settings
 - ▶ arch
 - bin
 - ▶ doc
 - examples
 - lib
 - runtime
 - ▶ src
 - strategy
 - ▶ tools
 - wiki

Filename

[TableLayout.jar](#)

[jgap.jar](#)

[log4j-1.2.16.jar](#)

Experiment

The Anelastic Wave Propagation code AWP-ODC
of the Southern California Earthquake Center (SCEC)

AMD Opteron “24 Cours” and an NVIDIA C2050 Fermi GPU

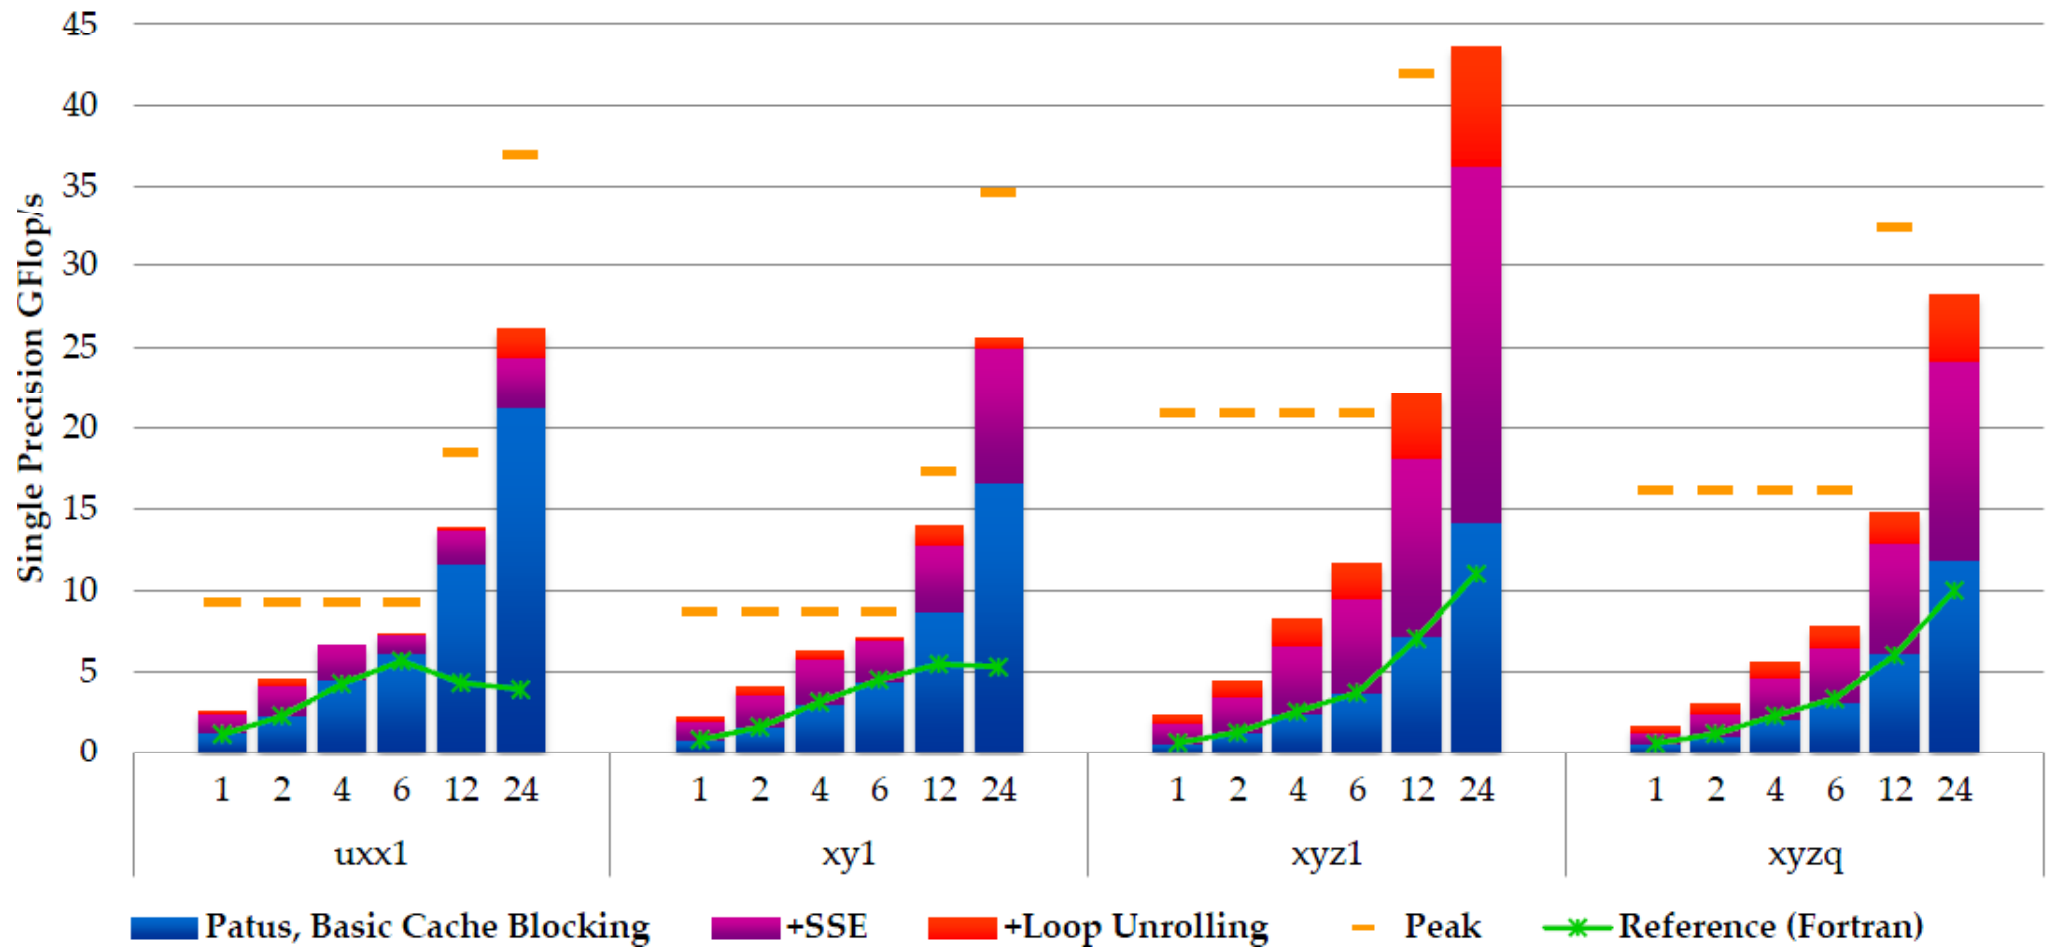
GNU gfortran/gcc 4.5.2 compilers
CUDA 4.0 and NVIDIA’s nvcc compiler

a 188×188×152 domain, single precision

$$u(i,j,k)=u(i,j,k)+(dth/d)*(\begin{aligned} &+c1*(xx(i,j,k)-xx(i-1,j,k))+ \\ &+c2*(xx(i+1,j,k)-xx(i-2,j,k))+ \\ &+c1*(xy(i,j,k)-xy(i,j-1,k))+ \\ &+c2*(xy(i,j+1,k)-xy(i,j-2,k))+ \\ &+c1*(xz(i,j,k)-xz(i,j,k-1))+ \\ &+c2*(xz(i,j,k+1)-xz(i,j,k-2))) \end{aligned}$$

CPU-result

Performance and Scaling of AWP-ODC Kernels
AMD Opteron "Magny Cours"



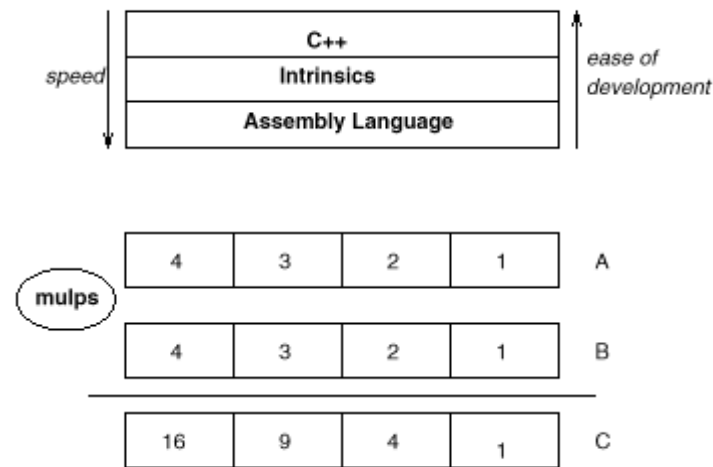
SSE

The Intel [Streaming SIMD Extensions](#) technology enhances the performance of floating-point operations. Intel C\C++, Microsoft's Macro Assembler support SSE. Pentium II and Pentium III Xeon SIMD instructions can greatly increase performance when exactly the same operations are to be performed on multiple data objects.

Intel's first [IA-32](#) SIMD effort was the [MMX](#) instruction set. [MMX](#) had two main problems: it re-used existing [floating point](#) registers making the CPU unable to work on both [floating point](#) and [SIMD data](#) at the same time, and it only worked on integers.

[SSE floating point instructions](#) operate on a new [independent](#) register set (the XMM registers), and it adds a few integer instructions that work on MMX registers.

```
#include <xmmintrin.h>
__m128 a, b, c;
a = _mm_set_ps(4, 3, 2, 1)
b = _mm_set_ps(4, 3, 2, 1)
c = _mm_set_ps(0, 0, 0, 0)
c = _mm_mul_ps(a, b);
```



Arithmetic Instructions

- addps, addss
- subps, subss
- mulps, mulss
- divps, divss
- sqrtps, sqrtss
- maxps, maxss
- minps, minss

Loop unrolling

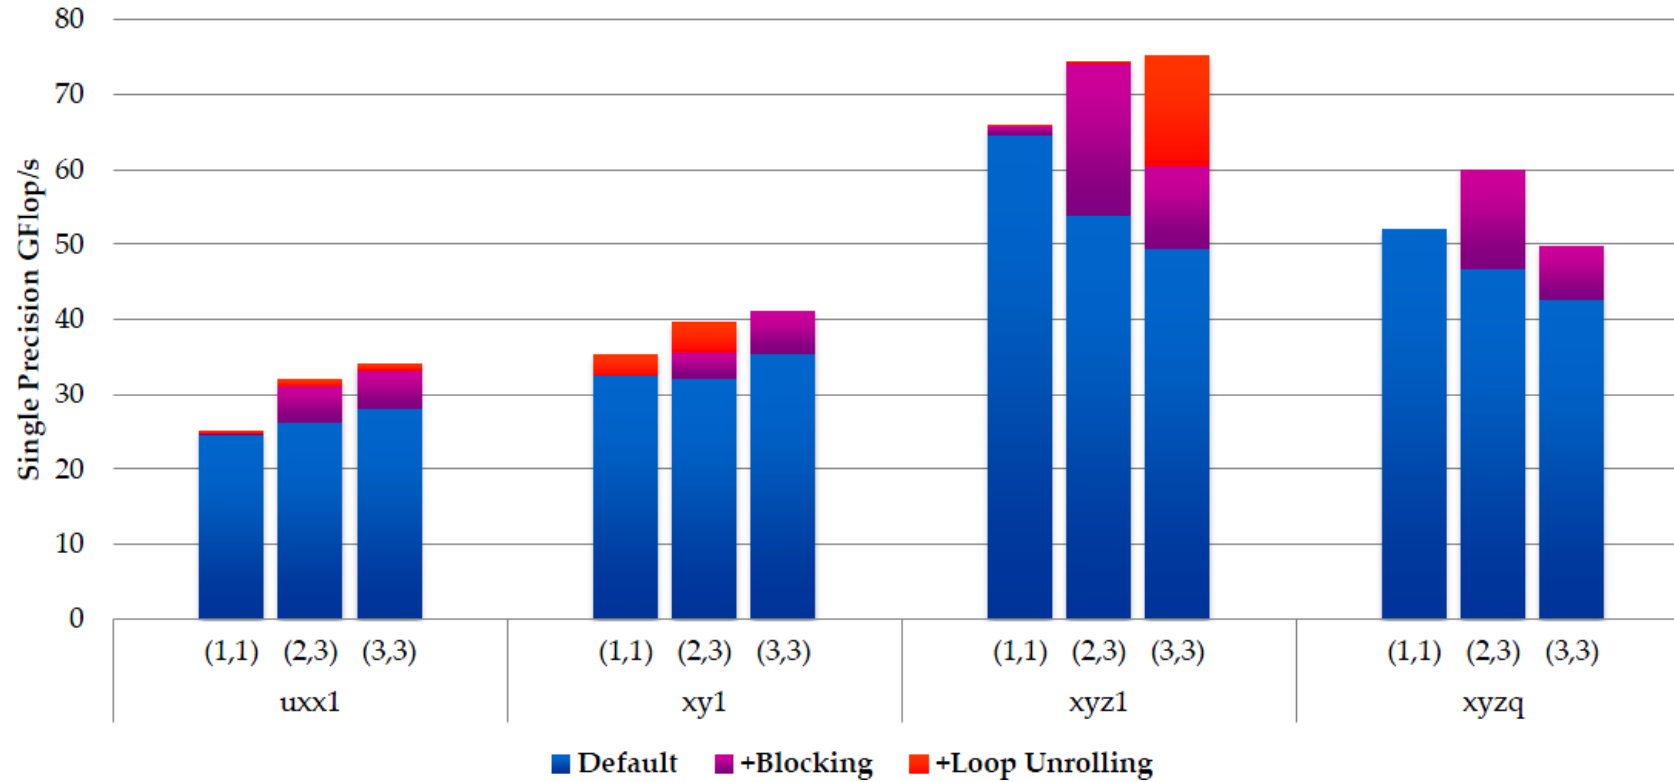
```
for (i = 0; i < 100; i++)  
    A[i] = A[i] + B[i] * C
```

You can unroll the loop, as we have below, giving you the same operations in fewer iterations with less loop overhead. You can imagine how this would help on any computer. Because the computations in one iteration do not depend on the computations in other iterations, calculations from different iterations can be executed together. On a superscalar processor, portions of these four statements may actually execute in parallel:

```
for (i = 0; i < 100; i += 4)  
{  
    A[i] = A[i] + B[i] * C;  
    A[i+1] = A[i+1] + B[i+1] * C;  
    A[i+2] = A[i+2] + B[i+2] * C;  
    A[i+3] = A[i+3] + B[i+3] * C  
}
```

GPU

Performance of AWP-ODC Kernels
NVIDIA Tesla C2050



1-dimensional thread blocks and grids,
3-dimensional thread blocks and a 2-dimensional grid,
3-dimensional thread blocks and grids

Conclusion

PATUS, a [code generation](#) and [auto-tuning](#) framework for general [stencil](#) computations.

It is for both programmers in need of an efficient implementation of a stencil kernel for a given hardware architecture, but who do not want to care about [hardware-specific](#) tuning, and for domain experts who want to experiment.

The framework still has limitations (restriction to shared memory architectures, no special boundary treatment, lacking support for temporal blocking schemes), which we intend to overcome in the future.

PATUS is open source software and licensed under the GNU Lesser GPL. It can be obtained from <http://code.google.com/p/patus/>.

question?

Auto-tuning

[Panorama](#) [8], [9] was a research compiler for tiling iterative stencil computations in order to minimize cache misses. [2008](#)

[Berkeley stencil auto-tuner](#) [10] seeks to substitute an annotated stencil computation in [Fortran95](#) automatically by an optimized version. [2007](#)

[The Pochoir stencil compiler](#) [11] applies the cache oblivious ideas initially formulated by Frigo and Strumpen [12] to stencil codes with ideally many time steps. [2008](#)

[Mint](#) [14] targets NVIDIA GPUs as hardware platforms and translates traditional, but annotated, C code to CUDA C and applies hardware-specific optimizations specifically tailored for stencil computations. [2007](#)