

## 💵 INDIANA UNIVERSITY

Universität Bielefeld

# Conjugate gradient solvers on Intel<sup>®</sup> Xeon Phi<sup>™</sup> and NVIDIA<sup>®</sup> GPUs

September 10, 2014 | Patrick Steinbrecher |

Olaf Kaczmarek Swagato Mukherjee Christian Schmidt Mathias Wagner

## Outline

- Lattice QCD
- Implementation
- Comparison





## Outline



- Lattice QCD
- Implementation
- Comparison



## Lattice QCD measurements

Chiral Condensate

$$\left\langle \overline{\psi}\psi\right\rangle = \frac{\partial \ln \mathcal{Z}}{\partial m_q} = \frac{n_f}{4} \left\langle \operatorname{Tr} M^{-1} \right\rangle$$

## Lattice QCD measurements

Chiral Condensate

 $\langle \overline{\psi}\psi\rangle = \frac{\partial \ln \mathcal{Z}}{\partial m_a} = \frac{n_f}{4} \left\langle \operatorname{Tr} M^{1} \right\rangle$  Solve: Mx = y

$$x = M^{-1}y$$

## Lattice QCD measurements

Chiral Condensate

 $\langle \overline{\psi}\psi \rangle = \frac{\partial \ln \mathcal{Z}}{\partial m_q} = \frac{n_f}{4} \left\langle \operatorname{Tr} M^{-1} \right\rangle$  Solve: Mx = y $x = M^{-1}y$ 

• derivatives of  $\ln Z$  w.r.t. the chemical potential  $\mu$ 

$$\frac{\partial^{2} \ln \mathcal{Z}}{\partial \mu^{2}} = \left\langle \frac{n_{f}}{4} \frac{\partial^{2} \left( \ln \det M \right)}{\partial \mu^{2}} \right\rangle + \left\langle \left( \frac{n_{f}}{4} \frac{\partial \left( \ln \det M \right)}{\partial \mu} \right)^{2} \right\rangle$$
$$\operatorname{Tr} \left( M^{-1} \frac{\partial^{2} M}{\partial \mu^{2}} \right) - \operatorname{Tr} \left( M^{-1} \frac{\partial M}{\partial \mu} M^{-1} \frac{\partial M}{\partial \mu} \right) \qquad \operatorname{Tr} \left( M^{-1} \frac{\partial M}{\partial \mu} \right)$$

(simplest case)

condition:  
$$\lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_{ki}^* \eta_{kj} = \delta_{ij}$$

$$\operatorname{Tr} M^{-1} = \lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_k^{\dagger} M^{-1} \eta_k$$

(simplest case)

$$\operatorname{Tr} M^{-1} = \lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_k^{\dagger} M^{-1} \eta_k$$

• evaluate traces using N random noise vectors  $\eta$  error ~ N<sup>-1/2</sup>

condition:  
$$\lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_{ki}^* \eta_{kj} = \delta_{ij}$$

(simplest case)

$$\lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_{ki}^* \eta_{kj} = \delta_{ij}$$

12.4

$$\operatorname{Tr} M^{-1} = \lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_k^{\dagger} M^{-1} \eta_k$$

. .

• evaluate traces using N random noise vectors  $\eta$  error ~  $N^{-1/2}$ • up to 1500  $\eta$  per gauge field configuration

(simplest case)

$$\operatorname{Tr} M^{-1} = \lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_k^{\dagger} M^{-1} \eta_k$$

- evaluate traces using N random noise vectors  $\eta$  error ~  $N^{-1/2}$ • up to 1500  $\eta$  per gauge field configuration
- up to 20000 configurations for each temperature
  - implies  $\sim 10^7$  inversions / temperature

condition:  $\lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_{ki}^* \eta_{kj} = \delta_{ij}$ 

(simplest case)

99% of runtime

condition:  $\lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_{ki}^* \eta_{kj} = \delta_{ij}$ 

$$\operatorname{Tr} M^{-1} = \lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_k^{\dagger} M^{-1} \eta_k$$

- evaluate traces using N random noise vectors  $\eta$  error ~  $N^{-1/2}$ • up to 1500  $\eta$  per gauge field configuration
- up to 20000 configurations for each temperature
  - implies  $\sim 10^7$  inversions / temperature

(simplest case)

99% of runtime

condition:  $\lim_{N\to\infty}\frac{1}{N}\sum_{ki}^{\prime\prime}\eta_{ki}^{*}\eta_{kj}=\delta_{ij}$ 

$$\operatorname{Tr} M^{-1} = \lim_{N \to \infty} \frac{1}{N} \sum_{k=1}^{N} \eta_k^{\dagger} M^{-1} \eta_k$$

- evaluate traces using N random noise vectors  $\eta$ error  $\sim N^{-1/2}$ up to 1500  $\eta$  per gauge field configuration
- up to 20000 configurations for each temperature
  - implies  $\sim 10^7$  inversions / temperature



## **HISQ Dslash** w = D v

$$w_{n} = \sum_{\mu=0}^{4} \left[ \left( U_{n,\mu} v_{n+\mu} - U_{n-\mu,\mu}^{\dagger} v_{n-\mu} \right) + \left( N_{n,\mu} v_{n+3\mu} - N_{n-3\mu,\mu}^{\dagger} v_{n-3\mu} \right) \right]$$















September 10, 2014

Patrick Steinbrecher

## **Multiple right-hand sides**

Memory

## **Multiple right-hand sides**

const. gauge fields

Memory





## DS(







September 10, 2014









## Multiple right-hand sides (rhs)



## Multiple right-hand sides (rhs)

- gauge compression gives no significant perf. increase
- multiple rhs improves only Dslash operator
  - Linear Algebra still scales linearly



## Implementation

GPU:	since 2009	MIC:	since 2014 $\hookrightarrow$ more to explore more to improve
<u>code</u>	CUDA	<u>code</u>	C++,OpenMP Intrinsics
<u>cards</u>	M2075 GTX <sup>™</sup> 580 GTX <sup>™</sup> Titan K20 K40 ····	<u>cards</u>	Xeon Phi <sup>™</sup> 5110P Xeon Phi <sup>™</sup> 7120P

#### used for

#### used for

 $\hookrightarrow$  will be used on larger clusters

- A. Bazavov et al., arXiv:1404.4043 [hep-lat]
- A. Bazavov et al., arXiv:1404.6511 [hep-lat]
- A. Bazavov et al., Phys. Rev. Lett. 111, 082301 (2013)
- A. Bazavov et al., Phys. Rev. Lett. 109, 192302 (2012)

## Intel<sup>®</sup> Xeon Phi<sup>™</sup>

- x86 based many-core processor
  - runs a Linux µOS
  - in-order execution
  - up to 61 cores at 1.238 GHz
  - core boost for 7120 series
  - theoretical bandwidth 352  $\mathrm{GB/s}$



## Intel<sup>®</sup> Xeon Phi<sup>™</sup>

- x86 based many-core processor
  - runs a Linux µOS
  - in-order execution
  - up to 61 cores at 1.238 GHz
  - core boost for 7120 series
  - theoretical bandwidth 352  $\rm GB/s$
- parallelism through
  - large SIMD registers



#### Mask Registers:



Eight 16-bit k registers per core.

#### Vector Registers:



Thirty-two 512-bit zmm registers per core.
# Intel<sup>®</sup> Xeon Phi<sup>™</sup>

- x86 based many-core processor
  - runs a Linux µOS
  - in-order execution
  - up to 61 cores at 1.238 GHz
  - core boost for 7120 series
  - theoretical bandwidth 352  $\mathrm{GB/s}$
- parallelism through
  - large SIMD registers
  - 4 hardware threads per core



#### Mask Registers:



Eight 16-bit k registers per core.

#### Vector Registers:



Thirty-two 512-bit zmm registers per core.

# Intel<sup>®</sup> Xeon Phi<sup>™</sup>

- x86 based many-core processor
  - runs a Linux µOS
  - in-order execution
  - up to 61 cores at 1.238 GHz
  - core boost for 7120 series
  - theoretical bandwidth 352  $\mathrm{GB/s}$
- parallelism through
  - large SIMD registers
  - 4 hardware threads per core
- peak fp32/fp64
  - 2.42/1.21 TFlop/s



#### Mask Registers:

16 bits 0 15

Eight 16-bit k registers per core.

#### Vector Registers:



Thirty-two 512-bit zmm registers per core.

# Intel<sup>®</sup> Xeon Phi<sup>™</sup>

- x86 based many-core processor
  - runs a Linux µOS
  - in-order execution
  - up to 61 cores at 1.238  $\rm GHz$
  - core boost for 7120 series
  - theoretical bandwidth 352  $\rm GB/s$
- parallelism through
  - large SIMD registers
  - 4 hardware threads per core
- peak fp32/fp64
  - 2.42/1.21 TFlop/s
- offload / native mode



#### Mask Registers:

16 bits 0 15

Eight 16-bit k registers per core.

#### Vector Registers:



Thirty-two 512-bit zmm registers per core.

## **Knights Corner architecture**



### Problem:

## $18 \, \mathrm{mod} \, 16 \; = \; 2$







mix color components of matrices



### Solution:

### mix color components of matrices



easiest single-precision approach "16-fold site fusion"

Patrick Steinbrecher

٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠
٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠
٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠
•	•			•	•	•	•		•				•	•	•	•	•	•	•						•	•	•		•		٠
•	•		•	•	•	•	•						•		•	•	•								•		•				
-					-	-	-			-							-	-							-				-	-	-
						•						•										•	•							•	•
•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
•	•	•	٠	٠	•	٠	•	•	•	•	•	•	٠	٠	٠	٠	•	•	•	•	•	•	•	٠	٠	•	٠	•	٠	•	•
٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠
٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠
٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠
٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠	٠
•	•			•	•	•	•								•	•	•										•				
																														-	-
•	•	•	•	•	٠	•	•	•	•	•	•	•	•	•	•	•	٠	•	•	•	•	•	•	•	٠	•	•	•		•	•

Site fusion		
<ul> <li>even</li> </ul>	• • • • • • • • • • • • • • • • • • •	• • • • • • • • • • • • • • • •
• odd	· · · · · · · · · · · · · · · · · · ·	
	· · · · · · · · · · · · · · · · · · ·	••••

- even-odd preconditioning
  - need to fuse sites of the same parity

Site fusion		•••••••••••••••••••••••••••••••••••••••
• even		•
• odd	· · · · · · · · · · · · · · · · · · ·	•
	· · · · · · · · · · · · · · · · · · ·	•

- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse 16 sites for single-precision



- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse 16 sites for single-precision



- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse 16 sites for single-precision



- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse 16 sites for single-precision



- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse 16 sites for single-precision



- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse 16 sites for single-precision
  - $32^3 \times 8 \longrightarrow 2 \times 32^2 \times 8$  (less indexing)
  - constraints on lattice size
  - high register pressure

.

### naive 16-fold site fusion



### naive 16-fold site fusion















very natural implementation — use SIMD registers like scalars

September 10, 2014

Patrick Steinbrecher

### 16-fold

- easy to implement
- very good performance for 1 right-hand side
- high register pressure
  - 30 registers for a fused matrix-vector product
- perfect arithmetic vectorization

### 16-fold

- easy to implement
- very good performance for 1 right-hand side
- high register pressure
  - 30 registers for a fused matrix-vector product
- perfect arithmetic vectorization

only 32 registers/core

### 16-fold

- easy to implement
- very good performance for 1 right-hand side
- high register pressure
  - 30 registers for a fused matrix-vector product
- perfect arithmetic vectorization

only 32 registers/core

### 8-fold

- requires align/blend/mask instructions
  - some "padding" Flops
- very good performance for 1 right-hand side
- reduced register pressure
  - 16 registers for a fused matrix-vector product

# 16-fold

Site fusion

- easy to implement
- very good performance for 1 right-hand side
- high register pressure
  - 30 registers for a fused matrix-vector product
- perfect arithmetic vectorization

only 32 registers/core

### 8-fold

- requires align/blend/mask instructions
  - some "padding" Flops
- very good performance for 1 right-hand side
- reduced register pressure
  - 16 registers for a fused matrix-vector product



Patrick Steinbrecher

in-order cores are very sensitive to cache misses

- in-order cores are very sensitive to cache misses
  - 1 thread/core

- in-order cores are very sensitive to cache misses
  - 1 thread/core  $\rightarrow$  cache miss

- in-order cores are very sensitive to cache misses
  - 1 thread/core  $\rightarrow$  cache miss  $\rightarrow$  stall

- in-order cores are very sensitive to cache misses
  - 1 thread/core  $\rightarrow$  cache miss  $\rightarrow$  stall
  - need to use 4 threads per core

- in-order cores are very sensitive to cache misses
  - 1 thread/core  $\rightarrow$  cache miss  $\rightarrow$  stall
  - need to use 4 threads per core
- only L2 hardware prefetcher (no L1)
  - does a good job for a linear memory access

- in-order cores are very sensitive to cache misses
  - 1 thread/core  $\rightarrow$  cache miss  $\rightarrow$  stall
  - need to use 4 threads per core
- only L2 hardware prefetcher (no L1)
  - does a good job for a linear memory access
- access pattern of Dslash is too complicated for the compiler
  - insert software prefetches (SWP) by hand
  - hard part is to figure out the prefetch distance

- in-order cores are very sensitive to cache misses
  - 1 thread/core  $\rightarrow$  cache miss  $\rightarrow$  stall
  - need to use 4 threads per core
- only L2 hardware prefetcher (no L1)
  - does a good job for a linear memory access
- access pattern of Dslash is too complicated for the compiler
  - insert software prefetches (SWP) by hand
  - hard part is to figure out the prefetch distance



HISQ inverter runs  $2 \times$  faster with software prefetches

### Performance gains on Xeon Phi<sup>™</sup>



## **GPU** implementation

- coalesced memory layout  $\rightarrow$  SoA
- 1 threads calculates 1 lattice site
  - automatic tuning of block size

Instruction Cache																				
	Wa	np Sal	HALF			Wa	rp Sche	suier			We	rp Sol	Helder	Warp Scheduler						
•	spate		Dispe	ich 🛛	D	spate		Dispe	lch	Dispetch Dispetch						Dispatch Dispatch				
							Real	stor	File ()	16.52	6 x 3	2-bit								
+					+	+	÷.			÷	+		•	+	+		+			
	Core	6010	DP Unit	Care	Core	6010	OP Unit		sru	Core	Core	Core	OP Unit	Core	Core	Core	OP UK	1097	SFL	
	C	C	DP Unit	Care	Core	C 010	OP Unit			Core	Core	C	DP Unit	Core	Core	C	00 UK	1.047	ar.	
Care	C	C	DP Unit	Core	Core	C	OP Line		SFU	Core	Core	C	DP Unit	Core	Core	C	or us	1.0187	221	
C	C	C	DP Unit	C	C	C	07 UH		870	C	C	C	DP Unit	Core	C	C	DP UK	1.011	871	
C	C	Ce.+2	CP Unit	C	C	C	09 Det		870	C	C	C	CP Unit	C	C	C	DP to	LOST	871	
Care	Care	C	CP U-II	Care	Care	Care	DP Det	LOUT	seu	Care	C	C	DP Unit	Core	Care	C	-	LOUT	551	
Care	Care	<b>Care</b>	DP laws	Care	Care	Care	02 per	LOBT	sru	Care	Care	Garre	CP Laws	Care	Care	Garre	-	1.087	561	
Core	Care	Care -	CP Unit	Core	Care	Core	02 pet	1.081	sru	Care	Care	Gare	DP Unit	Core	Care	Care	09 U.S.	1.091	561	
Core	Care	C 8170	DP Unit	Core	Core	Core	07 044		sru	Core	Care	<b>Care</b>	DP Unit	Core	Core	<b>Care</b>	02.04	1.097	561	
Care	C	C	CP Unit	Core	Care	C 010	07 Unit	1.081	97U	Care	Care	6010	DP Unit	Core	Core	<b>Core</b>		1.041	sn.	
Care			TOP lines	Core.	Core.		02.044	1000	2711	Com.	Core.		CIP Halt	Com	6 mm			1000	50	
			COLUMN 1				02.044	1057	7711	Core.			DO UNIT	Core.				1000	-	
_			-					1007					Tel II-II					1000		
									-											
										-							_			
C.Int			-	C. Inte	-		-			Care	- are			- Com	C.See		_	Loan 1	-	
Care	Care	Cere	CP Unit	Care	Cere	Cere	DP Del	LOIET	880	Care	Care	Cere	OP Unit	Care	Care	Care	OP UK	LOST	871	
								108		SI NI	ce:10								redit	
							48 K	ana B Re	ad-O	niv D	ata C	acho							-	
	Tex	1	Tex			Tex		Te			Тех	1	Ter	ς.	Tex			Tex		
	Tex		Tex			Tex		Tes			Tex		Ter			Tex		Tes		

## **GPU** implementation

- coalesced memory layout  $\rightarrow$  SoA
- 1 threads calculates 1 lattice site
  - automatic tuning of block size
- memory access
  - exploit texture cache for gauge links
  - load vectors through the standard load path (L1)


- coalesced memory layout  $\rightarrow$  SoA
- 1 threads calculates 1 lattice site
  - automatic tuning of block size
- memory access
  - exploit texture cache for gauge links
  - load vectors through the standard load path (L1)
- multiple right-hand sides optimizations
  - register blocking



- multiple right-hand sides optimizations
  - texture cache blocking





- multiple right-hand sides optimizations
  - texture cache blocking





explore best possible combination of cache and register blocking

- use GPU boost on K40  $\longrightarrow$  875  $\rm MHz$ 
  - memory-bound problems stay well within thermal limits

### GPU talk by Mathias Wagner

GTC 2014  $\longrightarrow$  http://goo.gl/AglZBM

## Comparison



September 10, 2014

Patrick Steinbrecher

<sup>1</sup>Theoretical bandwidth, ECC off <sup>2</sup>Stream Triad, 4 threads/core, ECC on

Slide 19

## Comparison



September 10, 2014

Patrick Steinbrecher

<sup>1</sup>Theoretical bandwidth, ECC off <sup>2</sup>Stream Triad, 4 threads/core, ECC on

Slide 20



# Thank you for your attention!

## Things to improve / test on the Xeon Phi<sup>™</sup>

- divide lattice into L2 friendly blocks
  - let 4 threads work as a team in one block
  - synchronization at the end of each block



better cache locality less TLB pressure

- run Linear Algebra with less than 4 threads per core?
  - better Stream Triad bandwidth for 1 thread/core

## Intrinsics

### Definition

Intrinsics are inlined assembly-coded C functions, which can be optimized by the compiler. They do not require explicit instruction or register scheduling through the programmer as in asm.

## Intrinsics

### Definition

Intrinsics are inlined assembly-coded C functions, which can be optimized by the compiler. They do not require explicit instruction or register scheduling through the programmer as in asm.

- full control over SIMD instructions
- MIC Intrinsics are very similar to ISA on Intel<sup>®</sup> CPUs
  - mostly easier on the Xeon Phi<sup>™</sup>

## Intrinsics

#### Definition

Intrinsics are inlined assembly-coded C functions, which can be optimized by the compiler. They do not require explicit instruction or register scheduling through the programmer as in asm.

- full control over SIMD instructions
- MIC Intrinsics are very similar to ISA on Intel<sup>®</sup> CPUs
  - mostly easier on the Xeon Phi<sup>™</sup>

due to more and mask instructions