CL^{2}QCD

Lattice QCD based on OpenCL

Christopher Pinke
with M. Bach, O. Philipsen & A. Sciarra

GOETHE UNIVERSITÄT FRANKFURT AM MAIN

Institute for Theoretical Physics
Goethe-University Frankfurt

HIC for FAIR Helmholtz International Center

GPU Computing in High Energy Physics
University of Pisa, September 10-12, 2014

Strong Interactions: Quantum Chromodynamics (QCD)

$$\mathcal{S}_{\text{QCD}} = \int \mathrm{d}^4x \; \bar{\psi}(x)\; D[A_\mu]\; \psi(x) \; - \frac{1}{4}\mathcal{F}^a{}_{\mu\nu}[A_\mu](x)\mathcal{F}^{a,\mu\nu}[A_\mu](x)$$

Fermion part: Quarks ψ — Gauge part: Gluons A_μ

Confinement

Quarks carry 3 colours transmitted by Gluons. Confinement of quarks and gluons into colourless objects.

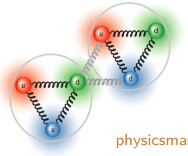

physicsmasterclasses.org

Running coupling

Politzer, Wilczek, Gross (NP 2004)

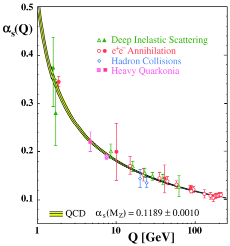

Bethke (2007)

QCD non-perturbative/strong at low energies

Studies from first principles → Lattice QCD

Exploring the QCD phase diagram

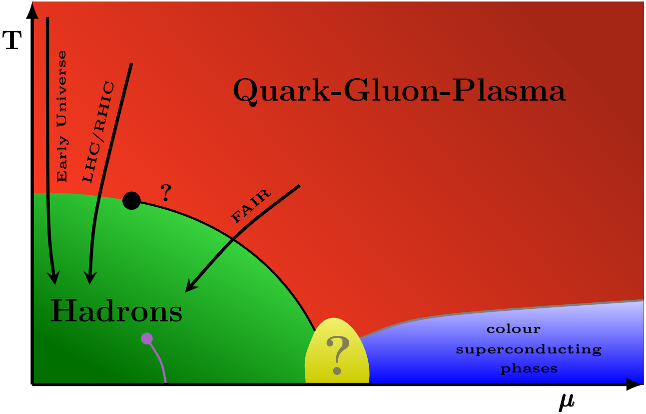

Critical Endpoint (CEP): Does it exist?

Exploring the QCD phase diagram

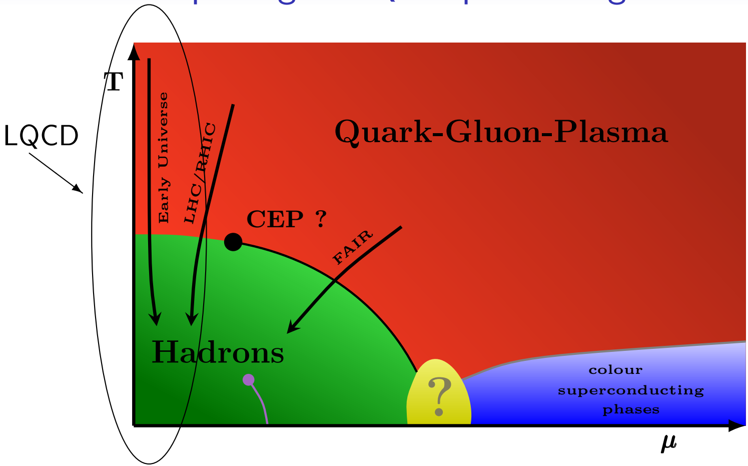

Sign-Problem of LQCD:

$\det(D+\mu)^* = \det(D-\mu^*) \Rightarrow$ Importance sampling ill-defined

Exploring the QCD phase diagram

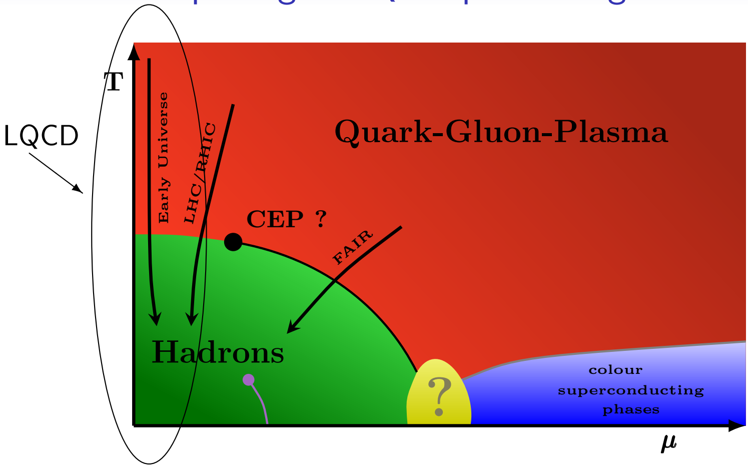

Sign-Problem of LQCD:

Solutions: Taylor expansion, Reweighting, *Imaginary μ*

Lattice QCD (I)

Discretize spacetime:

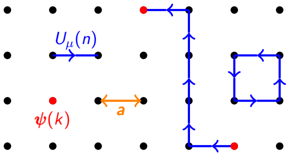

$$\mathcal{S}_{\mathrm{LQCD}} = \mathcal{S}_{\mathrm{QCD}} + a\mathcal{S}_1 + a^2\mathcal{S}_2 + \dots$$

Continuum limit: $a \to 0$

Fermions on the lattice

Different discretizations:

- Kogut-Susskind/Staggered
  Remnant chiral symmetry, cheap, Rooting
- Wilson
  Theoretically sound, no chiral symmetry
- Ginsparg-Wilson
  Chiral symmetry, very expensive

Consistency with experiment

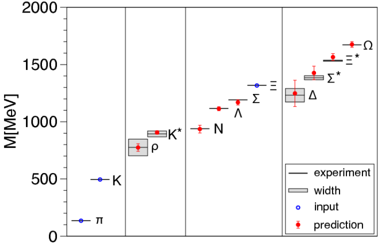

BMW Collaboration, 2008

Lattice QCD (II)

Apply Importance Sampling: Hybrid Monte Carlo (HMC)

Duane et.al (1987)

$$\langle \mathcal{O} \rangle = \frac{1}{Z} \int \mathcal{D}[U]\, \mathcal{O}[U]\, \mathrm{e}^{-S_{\mathrm{eff}}[U]} \approx \frac{1}{N} \sum_m \mathcal{O}[U_m]$$

- Propability $\rho \sim \mathrm{e}^{-S_{\mathrm{eff}}[U]}$, $S_{\mathrm{eff}}[U] \sim S_{\mathrm{gauge}} + \ln \det D$
- Fermion determinant: $\det D \sim \phi^\dagger D^{-1}\, \phi$
- Sign-Problem of Lattice QCD:

$$\det(D+\mu)^* = \det(D-\mu^*)$$

$\Rightarrow$ Importance sampling ill-defined at finite, real μ

Lattice QCD (III)

```
r=b-D*x
p=r
rsold=r'*r
for i < # steps:
  Dp=D*p
  alpha=rsold/(p'*Dp)
  x=x+alpha*p
  r=r-alpha*Dp
  rsnew=r'*r
  if sqrt(rsnew)<eps
    break
  beta = rsnew/rsold
  p=r+beta*p
  rsold=rsnew
```

Inversion of the Fermion Matrix

$$D\,\psi = \phi \Leftrightarrow \psi = D^{-1}\,\phi$$

- Dominant part of typical LQCD simulation (for small m_{quark})
- Use Krylov-subspace based solver (iterative)

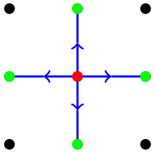

Most expensive ingredient: Derivative part $\not{D}$:

Wilson: read/write: 2880 Bytes

$\leftrightarrow$ perform 1632 FLOP

$\rightarrow$ Low numerical density ($\sim$ 0.5 FLOP/Byte)

$\Rightarrow$ LQCD always memory bandwidth limited

Lattice QCD at Finite Temperature

Lattice Volume: $V = N_\sigma^3 * N_\tau$ — N_σ: spatial extent, N_τ: temporal extent

Finite Temperature on the lattice:

$$T = \frac{1}{a(\beta) N_\tau} \quad \text{and} \quad N_\tau \ll N_\sigma$$

$\left(\text{Compare } T = 0 : N_\tau \gg N_\sigma \right)$

Simulation strategy

- Find critical temperature
  $\Rightarrow$ Need multiple β values (Temperature scan)
- Check volume scaling (Phase transitions $\sim$ non-analyticities)
  $\Rightarrow$ Need multiple spatial volumes (N_σ) (Finite Size Analysis)

$\Longrightarrow$

- Small/Moderate problem sizes
- Studies inherently parallel (trivially)

Graphics Processing Units (GPUs)

LQCD always memory bandwidth limited ⇒ GPUs well-suited

Performance

	Chip	Peak SP {GFLOPS}	Peak DP {GFLOPS}	Peak BW {GB/s}
AMD Radeon HD 5870	Cypress	2720	544	154
AMD Radeon HD 7970	Tahiti	3789	947	264
AMD FirePro S10000	Tahiti	2×3410	2×850	2×240
NVIDIA GeForce GTX 680	Kepler	3090	258	192
NVIDIA Tesla K40	Kepler	4290	1430	288
AMD Opteron 6172	Magny-Cours	202	101	43
Intel Xeon E5-2690	Sandy Bridge EP	371	186	51

LQCD & GPUs

Virtually all applications based on vendor-specific NVIDIA's CUDA
See the **QUDA** library: https://github.com/lattice/quda

GPU clusters at Frankfurt

	LOEWE-CSC	SANAM
GPU nodes	786	304
GPUs/node	1 × AMD 5870	2 × AMD S10000
CPUs/node	2 × Opteron 6172	2 × Xeon E5-2650

Future: LOEWE-CSC-Update & new LQCD-Cluster at GSI

Alternative to CUDA: OpenCL

https://www.khronos.org/opencl/

Open standard for heterogeneous computing platforms
⇒ GPUs and CPUs can be used together within same framework

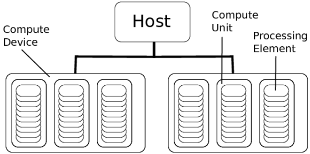

CL^{2}QCD (I)

http://code.compeng.uni-frankfurt.de/projects/clhmc

- New LQCD code based on OpenCL
  Bach et. al., Comput.Phys.Commun. 184 (2013)
- Successfully applied in physics studies on GPU clusters Loewe-CSC & Sanam

Features

- First OpenCL application for Wilson fermions focusing on **Twisted Mass Wilson fermions** Frezzotti & Rossi 2003
- **Staggered fermions** in standard formulation Kogut & Susskind 1975
- Improved gauge actions
- Standard inversion and integration algorithms
- ILDG-compatible IO
- RANLUX Pseudo-Random Number Generator (PRNG)
  Lüscher 1994

CL^{2}QCD (II)

Executables

- **HMC:** Generation of gauge field configurations for $N_f = 2$ (Twisted Mass) Wilson type fermions using HMC algorithm;
- **RHMC:** Generation of gauge field configurations for staggered type fermions using Rational HMC algorithm Clark & Kennedy (2007)
- **SU3HEATBATH:** Gen. of gauge field conf. for $SU(3)$ Pure Gauge Theory using heatbath algorithm e.g. Cabibbo & Marinari (1982)
- **INVERTER/GAUGEOBSERVABLES:** Measurements of fermionic/gauge observables on given gauge field conf.

Unit Tests

OpenCL hardware/platform independent

Computing architecture known at runtime only
$\Rightarrow$ Kernel compilation at runtime

- Varying computing platforms $\Rightarrow$ Need regression tests
- Follow Test Driven Development Beck (2002) and Clean Code Martin (2008) concepts
- LQCD fcts. local $\Rightarrow$ well testable

CL^{2}QCD - Code Structure (I)

- **Host program** of CL^{2}QCD set up in `C++`
  $\Rightarrow$ Independent program parts & extension capabilities
- Cross-platform compilation using CMake `http://www.cmake.org`
- Two main components:
  - `physics` package: High-level functionality
  - `hardware` package: Low-level functionality

  `meta` package: Control of program execution and I/O
- ***All parts*** of simulation code carried out using OpenCL kernels (double precision)
- Kernels in a certain way detached from host part (host can continue independently of kernel execution status)
  $\Rightarrow$ Clear separation into administrative part (host) and performance-critical calculations (kernels)
- **OpenCL kernels source files:**
  - OpenCL language based on `C99`
  - Compilation and execution handled within `hardware` package

CL^{2}QCD - Code Structure (II)

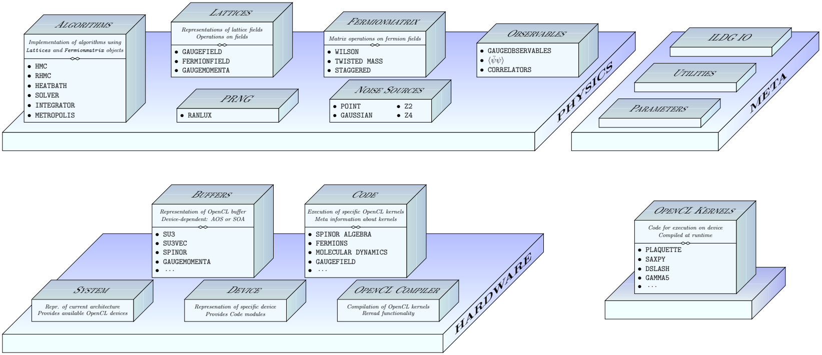

Memory access

Different memory access patterns:

Array of Structures (AoS) $\leftrightarrow$ *Structure of Arrays* (SoA)

Example: Fermion field ψ:

AoS	$\psi(0).e0$	$\psi(0).e1$	$\psi(0).e2$	$\cdots$	$\psi(0).e11$	$\psi(0).e12$	$\psi(1).e0$	$\psi(2).e0$	$\cdots$
SoA	$\psi(0).e0$	$\psi(1).e0$	$\psi(2).e0$	$\cdots$	$\psi(N).e0$	$\psi(0).e1$	$\psi(1).e1$	$\psi(2).e1$	$\cdots$

ψ: 12-comp. complex vector

GPU: SoA preferable for optimal memory access

$\not{D}_{\text{Wilson}}$ Performance

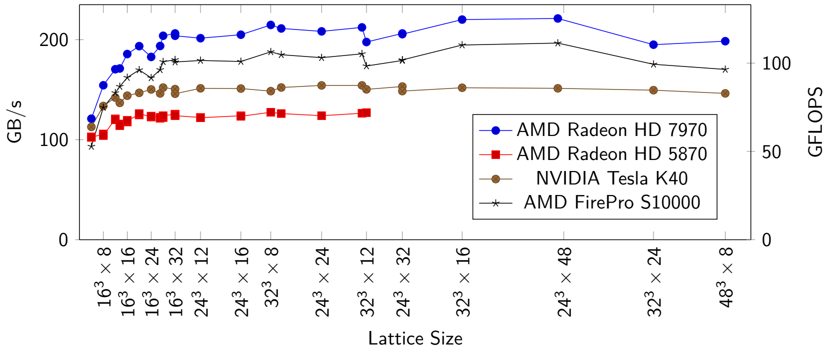

Peak Performances:

HD 7970:	264 GB/s,
HD 5870:	154 GB/s,
S10000:	240 GB/s,
K40:	288 GB/s

- AMD GPUs: $\sim$ 75% of peak BW
- NVIDIA: $\sim$ 60% (No spec. tuning)

HMC Performance

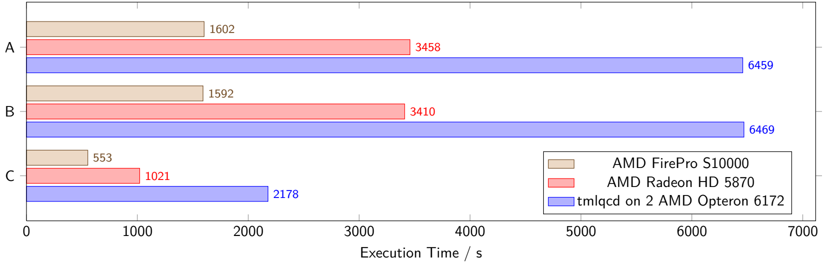

Speedup, compared to 2 LOEWE-CSC CPUs (= 1 node)

- AMD Radeon HD 5870 (LOEWE-CSC): ~ 2
- AMD FirePro S10000 (SANAM): ~ 4

Multi GPU

Scaling of the CG solver on SANAM (4 GPUs per node):

Strong scaling:

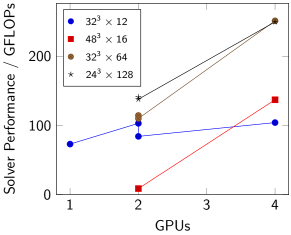

Weak scaling:

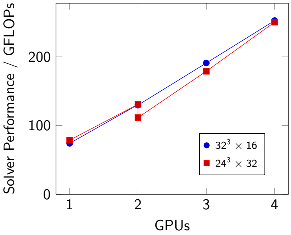

- Time direction splittable (within node)
- Nice scaling for large enough lattices
- Fin. T studies: Spatial splitting preferable

$D_{\text{staggered}}$ Performance

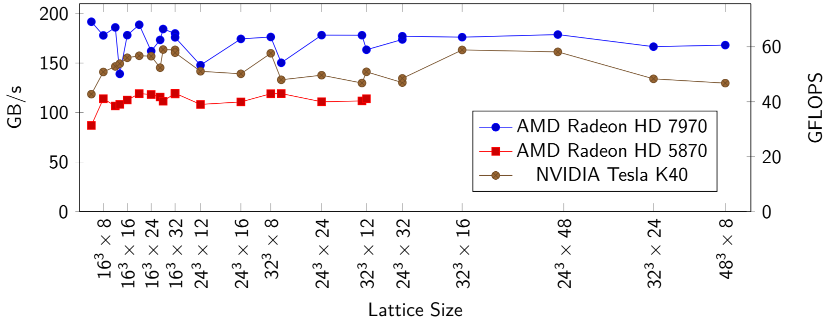

Peak Performances:
HD 7970: 264 GB/s,
HD 5870: 154 GB/s,
K40: 288 GB/s

- Staggered working set much smaller than Wilson
- No optimizations carried out (yet)
- AMD HD 5870: $\sim 75\%$ of peak BW
- Newer hardware requires more investigation

Summary

- New LQCD software CL^{2}QCD based on OpenCL
- Shows very good performance
- Available at:

`http://code.compeng.uni-frankfurt.de/projects/clhmc`

- Successfully applied in physics studies at finite temperature

Perspectives

- Implementation of new features as required by physics studies
- Tuning for specific hardware (esp. staggered code)