Lattice QCD at the Exascale

The 2nd R-CCS International Symposium
K to Fugaku

RIKEN Center for Computational Science

February 17-18, 2020

N.H. Christ

Columbia University

Outline

- Overview of Lattice QCD
- Algorithms and architectures
- Recent calculations
  - Muon *g*-2
  - $K_L \rightarrow \mu^+ \mu^- / \pi^0 \rightarrow e^+ e^-$
  - K_L-K_S mass difference
- Exascale plans

The RBC & UKQCD collaborations

BNL and BNL/RBRC

Yasumichi Aoki (KEK)
Taku Izubuchi
Yong-Chull Jang
Chulwoo Jung
Meifeng Lin
Aaron Meyer
Hiroshi Ohki
Shigemi Ohta (KEK)
Amarjit Soni

UC Boulder

Oliver Witzel

CERN

Mattia Bruno

Columbia University

Ryan Abbott
Norman Christ
Duo Guo
Christopher Kelly
Bob Mawhinney
Masaaki Tomii
Jiqun Tu
Bigeng Wang
Tianle Wang
Yidi Zhao

University of Connecticut

Tom Blum
Dan Hoying (BNL)
Luchang Jin (RBRC)
Cheng Tu

Edinburgh University

Peter Boyle
Luigi Del Debbio
Felix Erben
Vera Gülpers
Tadeusz Janowski
Julia Kettle
Michael Marshall
Fionn Ó hÓgáin
Antonin Portelli
Tobias Tsang
Andrew Yong
Azusa Yamaguchi

University of Liverpool

Nicolas Garron

MIT

David Murphy

Peking University

Xu Feng

University of Regensburg

Christoph Lehner (BNL)

University of Southampton

Nils Asmussen
Jonathan Flynn
Ryan Hill
Andreas Jüttner
James Richings
Chris Sachrajda

Stony Brook University

Jun-Sik Yoo
Sergey Syritsyn (RBRC)

Particle Physics – Overview

- Search for new phenomena at the highest energies.
- Goal of the Large Hadron Collider at CERN.

- High precision results at lower energies may have greater reach: Belle II at KEK.
- To compare with Standard Model need lattice QCD predictions: HPC plays a critical role

Standard Model

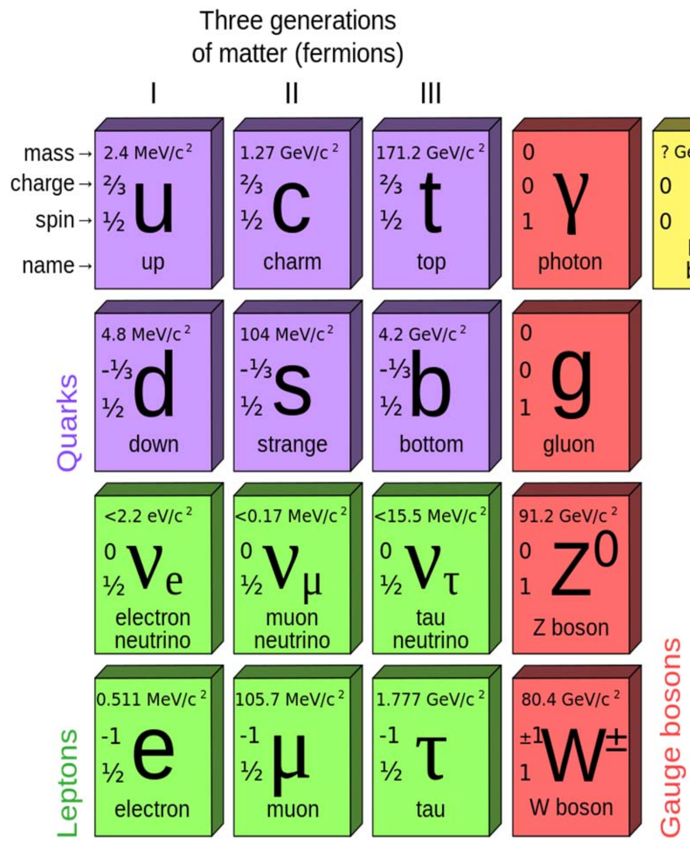

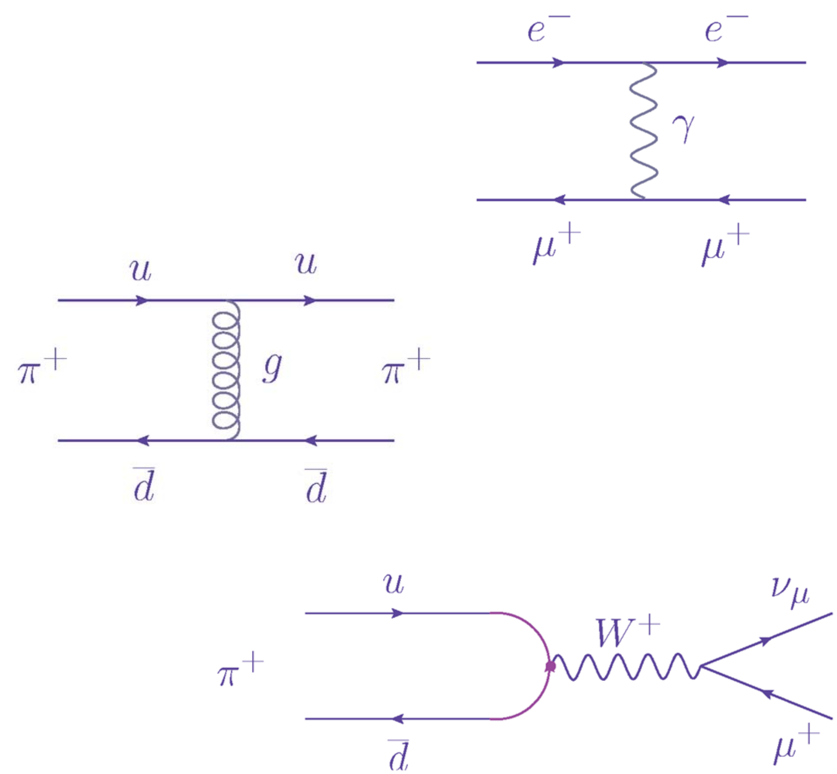

- γ (photon), *g* (gluon) and W^+ (weak) exchange

Standard Model

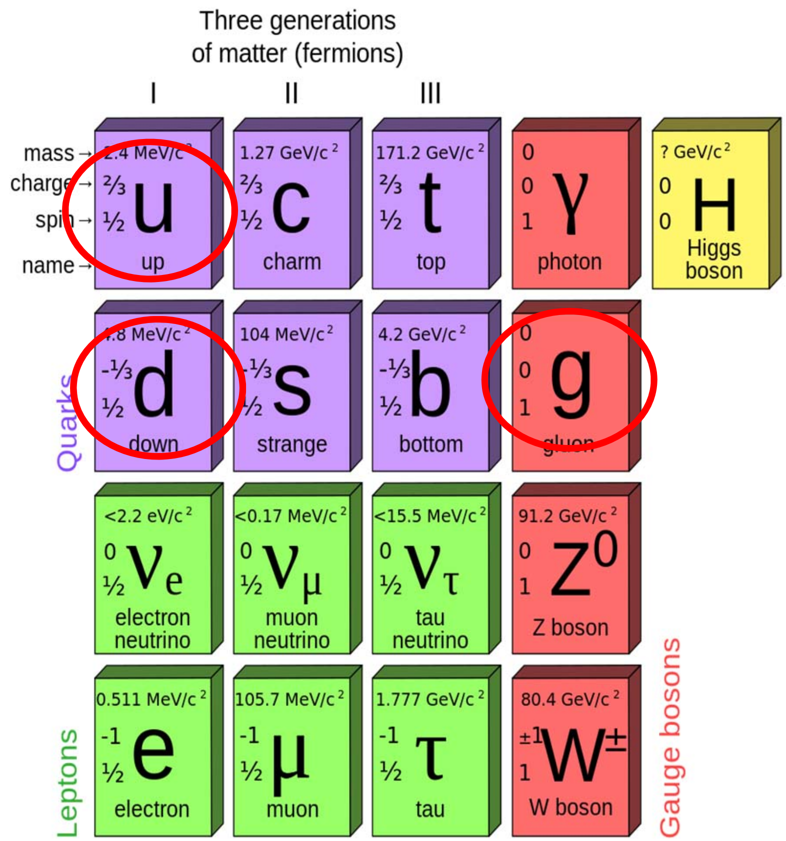

- Proton and neutron in atomic nucleus composed for ***u*** and ***d*** quarks.
- These are bound by the exchange of gluons ***g.***

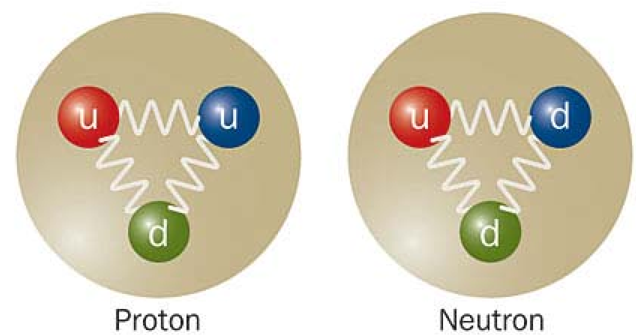

- ***u*** and ***d*** quarks nearly massless and highly relativistic
- Copious virtual pair creation → many body problem

Lattice QCD

Lattice QCD

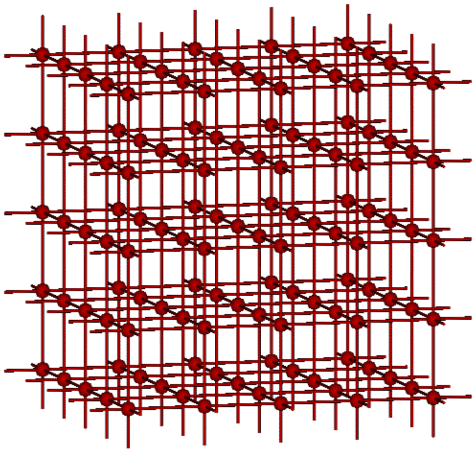

- Introduce a space-time lattice
- Evaluate the Euclidean Feynman path integral
  - Study $e^{-H_{QCD}\,t}$
  - Precise non-perturbative formulation

$$\sum_n \langle n | e^{-H(T-t)} \mathcal{O} e^{-Ht} | n \rangle = \int d[U_\mu(n)] e^{-\mathcal{A}[U]} \det(D+m) \mathcal{O}[U](t)$$

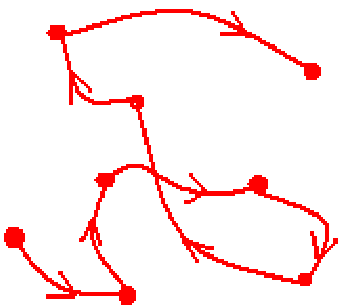

- Use Monte Carlo importance sampling with hybrid, MD/Langevin evolution
- ~40 samples give sub-percent errors

Challenges

- Numerical control of the path integral
  - 96^3 x 192 lattice → 5 x 10^9 dimensional integral
  - Long autocorrelation times → misleading errors
- Extend Schrodinger quantum mechanics into Euclidean time
  - Unphysical time evolution
  - Time-like processes challenging

- Fermion sign problem
  - $\mu_{\text{baryon}} \neq 0$ → path integral weight not positive
  - Poor signal to noise ratio in baryon correlators

$$\left\langle \left| q(x)\bar{q}(y) \right|^2 \right\rangle \sim e^{-|x-y| m_\pi} \qquad \left\langle (q(x))^{3B} (\bar{q}(y))^{3B} \right\rangle \sim e^{-|x-y| B m_N}$$

Historical overview

- 1974 Gauge-invariant lattice theory formulated. [Ken Wilson]
- 1980 String tension computed with physical dependence on the gauge coupling. [Mike Creutz]
- 1990 Fermion loops routinely included.
- 1992 Fermion doubling problem solved. [David B. Kaplan]
- 2002 Routine use of chiral quarks.
- 2014 Routine calculation with physical *u*, *d* and *s* quark masses.

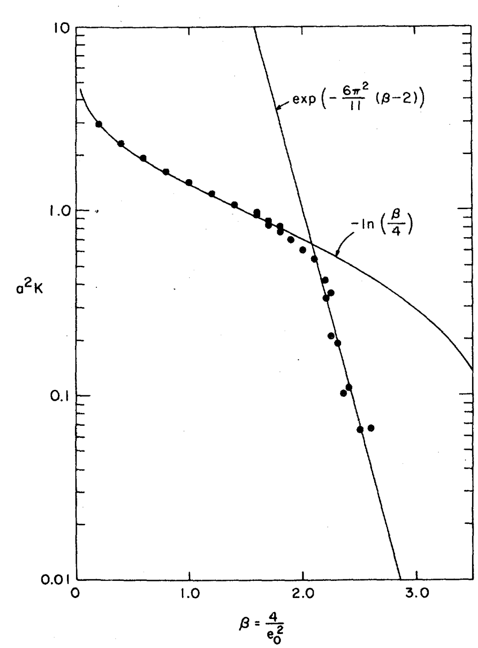

Lattice QCD – 2020

- Physical quark masses (ChPT not needed)
- Chiral quarks (doubling problem solved)
- Large physical volumes: (6 -10 fm)^3
- Small lattice spacing: 1/*a* = 2.774 GeV
- Some quantities with ~0.1% precision
- Now an essential tool for many standard model tests.

Algorithms and Architectures

Architecture

- Regular lattice problem:
  easy for parallel computer.
- Difficulty scales as ~ L^6.
- Weak scaling not sufficient
- Large demand on network bandwidth:
  - Local 4D volume grows with node flops
  - Off-node bandwidth grows as (node flops)$^{3/4}$
  - Mira → Summit

    0.2 TF/s /20 GB/s → 42 TF/s /25 GB/s
- Serious algorithm challenge

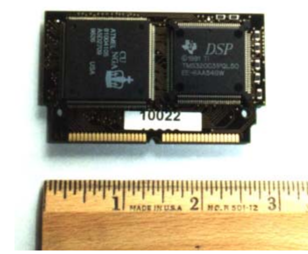

QCDSP – 600 Gflops

- RIKEN/Columbia/BNL
- Gordon Bell Prize 1998

Numerical methods

- Use low eigenmodes to solve $\not{D}\, G_n = h_n$ for multiple right-hand sides (deflation).
- 96^3 x 192 volume requires 5K eigenvectors (160 TB).
- Compress using local coherence: 30x (Clark, Jung & Lehner, arXiv:1710.06884)
- Use All-Mode-Averaging technique (Blum, Izubuchi & Shintani, arXiv:1208.4349 [hep-lat])
  - Loosen CG stopping condition $10^{-8} \rightarrow 10^{-4}$.
  - Obtain accurate result for 8 out of 128 time slices
  - $O(10^{-8})$ accurate result = $\langle G_{10^{-4}} \rangle_{124} + \langle G_{10^{-8}} - G_{10^{-4}} \rangle_8$
  - Achieve 5-20 x speed-up.

Numerical methods - cont

- Split-Grids (Boyle and Jung)
  - Deflate multiple right-hand sides on full machine
  - Perform separate CG inversions on many subpartions
  - 4x speed-up on Cori at NERSC (Cray XC40/KNL)
- Multi-splitting preconditioned conjugate gradient (Tu, Guo, and Mawhinney, arXiv:180408593)
  - Precondition with Dirac operator without off-node comms.
  - Use many local inversions, exploit large node performance.
  - Use tensor cores on Summit and obtain 1.5x speedup
  - Must include halo sites in product $D^{\dagger}_{pcond}\ D_{pcond}$.

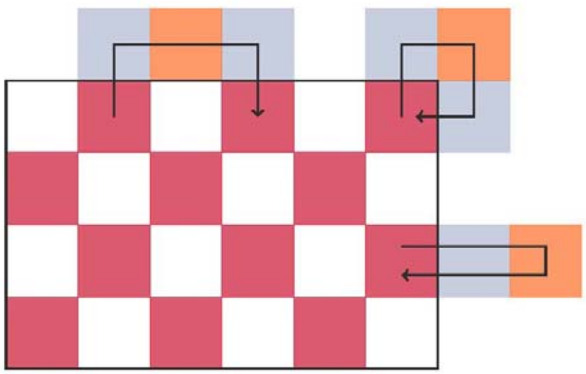

New Directions

- Use machine learning to infer the results for complex quantities from those of simple ones. (Boram Yoon, LANL)

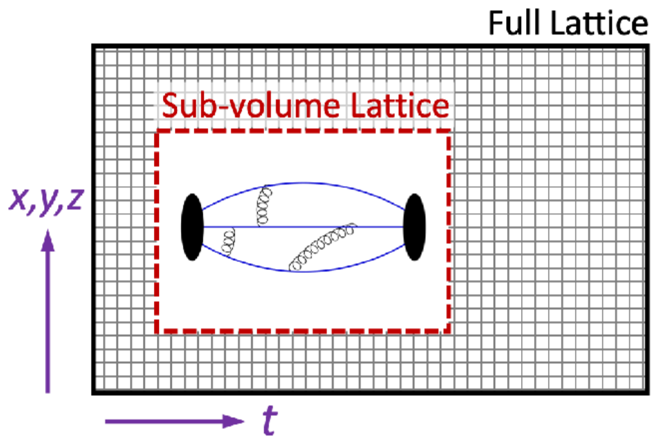

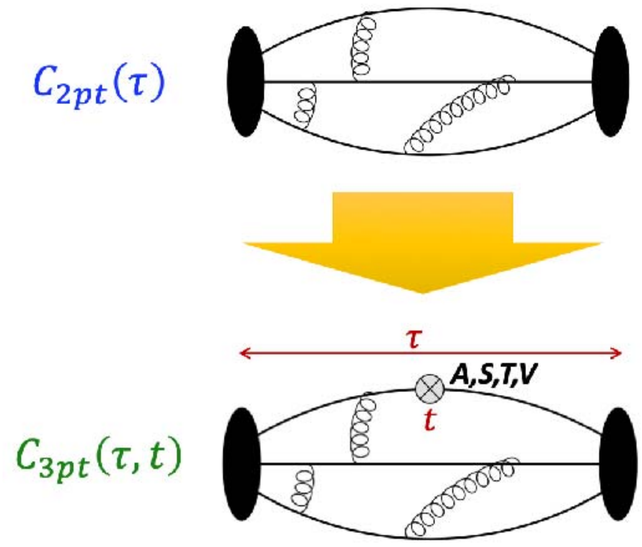

Correct small errors with a few exact results.

- Use machine learning to generate Markov samples. (Albergo, Kanwar and Shanahan, arXiv:1904.12072)

New Directions

- Use machine learning to simplify local MSPCG approximation (Jiqun Tu: Columbia → NVIDIA)

$$M_{L_s}^{-1} \simeq T^\dagger M_{L_s'}^{-1} T + \mu \mathbb{I}$$

- Learn optimal 1536 parameters in *T*

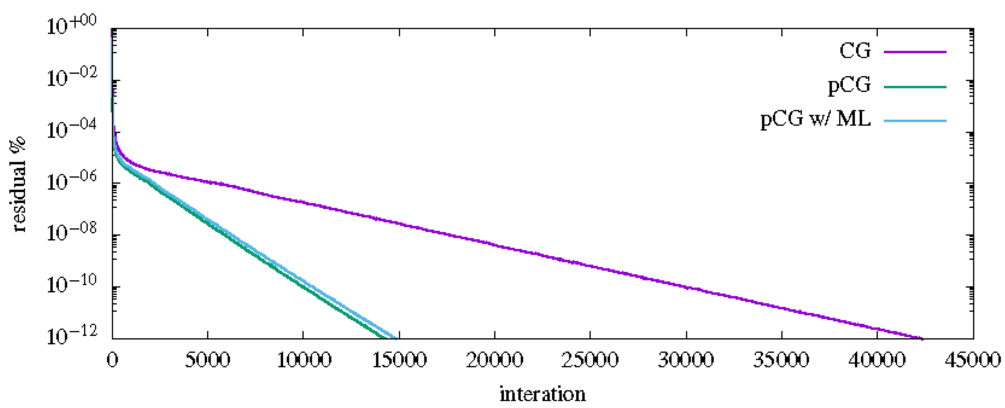

- Use Quantum Computing to generate real-time evolution?

Muon anomalous magnetic moment

g - 2 for the muon

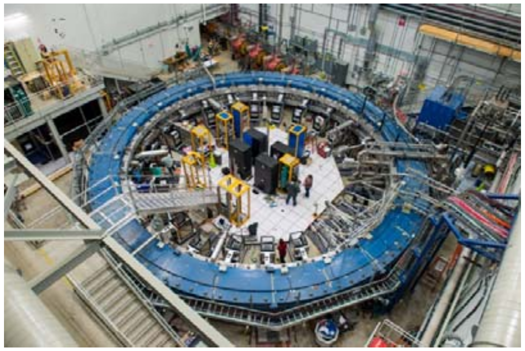

- Anomalous moment: $a_\mu = (g_\mu - 2)/2$
- BNL E821 expt:

  $a_\mu = 11699208.9 \pm 6.3 \times 10^{-10}$
- 3σ difference between the standard model prediction and experiment:

  FNAL E989

  $a_\mu^{EXP} - a_\mu^{SM} = 27.4(2.7)(2.6)(6.3) \times 10^{-10}$
- Effects of quark and gluons enter at order α_{EM}^2:

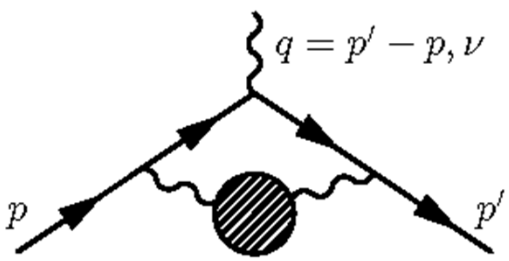

HVP
$692.5(2.7) \times 10^{-10}$

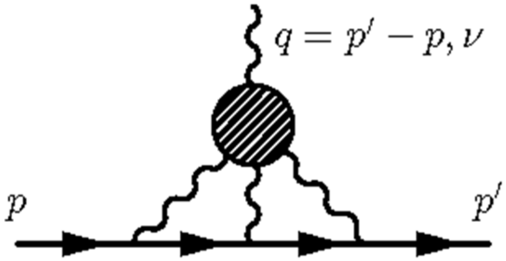

HLbL
$10.5(2.6) \times 10^{-10}$

Brookhaven, Columbia, Connecticut, Nagoya, RIKEN

Thomas Blum (Connecticut)

Norman Christ (Columbia)

Masashi Hayakawa (RIKEN & Nagoya)

Taku Izubuchi (BNL & RIKEN BNL)

Luchang Jin (Columbia → Connecticut)

Chulwoo Jung (Brookhaven)

Christoph Lehner (Regensburg & BNL)

Use lattice QCD to calculate HLbL

- Compute connected and leading disconnected parts. $\vec{\mu} = \frac{1}{2}\int d^3r \left(\vec{r} \times \vec{j}(\vec{r})\right)$

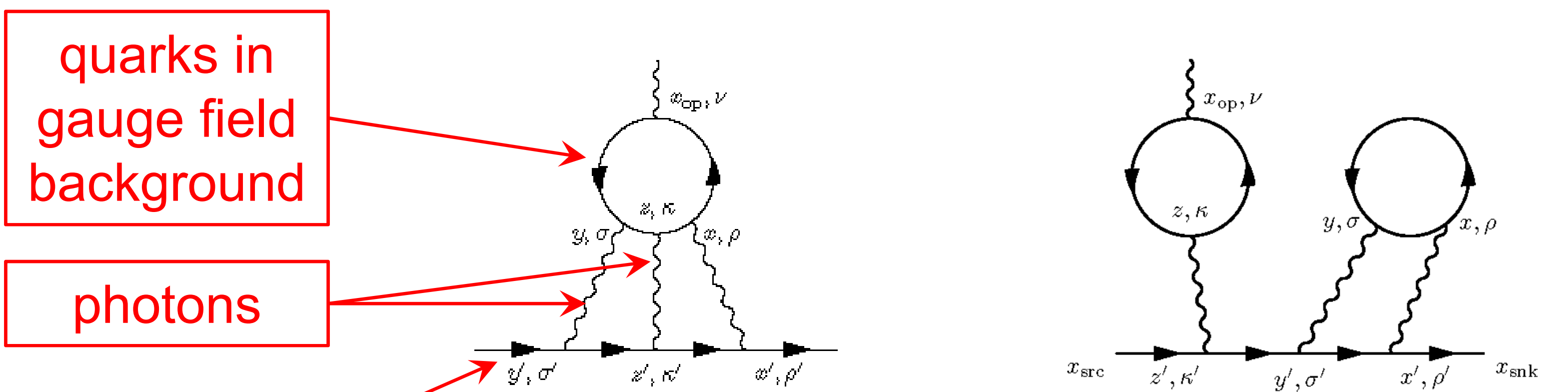

- Important challenge for lattice QCD
  - Treat E&M through an expansion in $\alpha_{\rm EM}$
  - Massless photon introduces new problems
- Sum stochastically over x and y

Use lattice QCD to calculate HLbL

- Use Lattice QCD for quark loop and gluons
- Evaluate photons and muon parts analytically

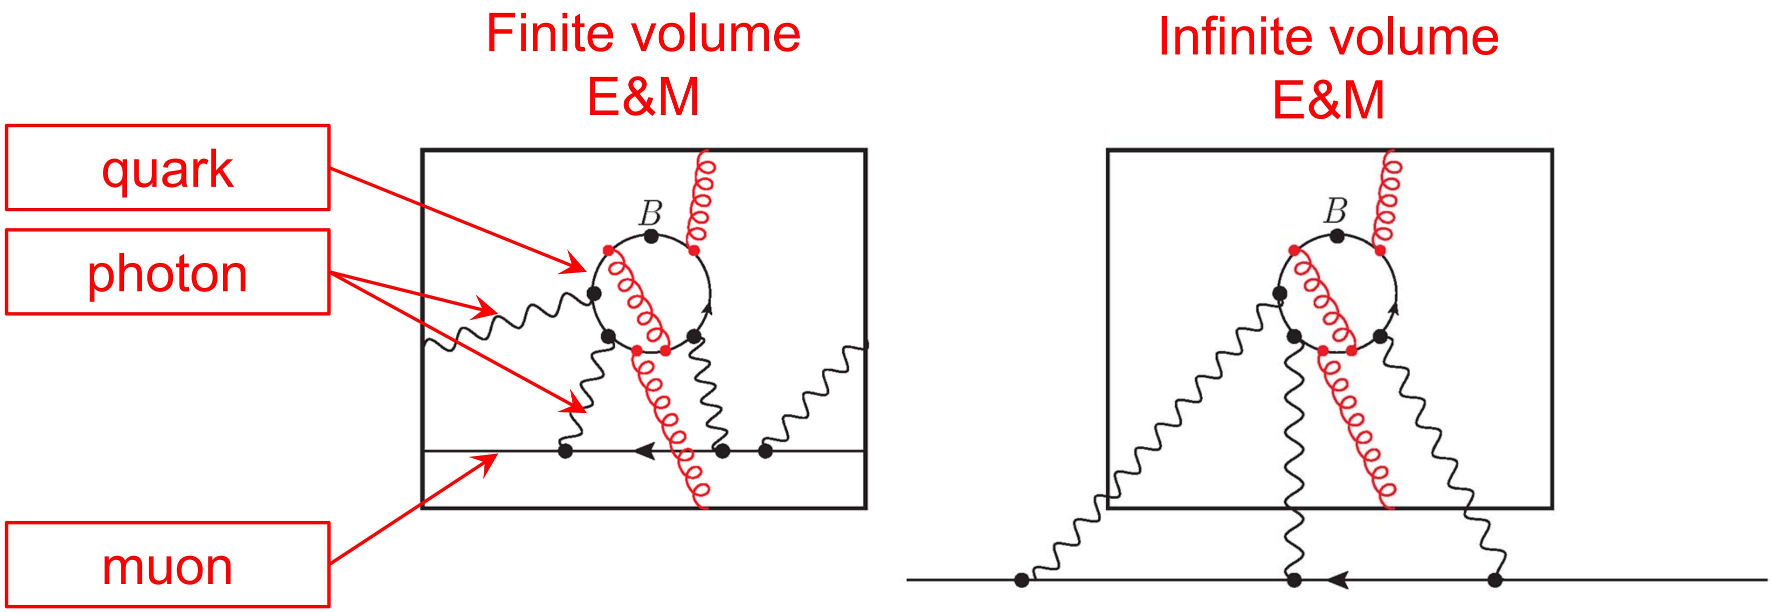

- With physical mass and continuum limit
- $\mu = 7.20(3.98)\text{stat}(1.65)\text{sys} \times 10^{-10}$ from QED_L

(arxiv:1911.081230)

$$K_L \to \mu^+ \mu^-$$

$$\pi^0 \to e^+ e^-$$

Physics of $K_L \to \mu^+ \mu^-$

- A second order weak, ``strangeness changing neutral current''

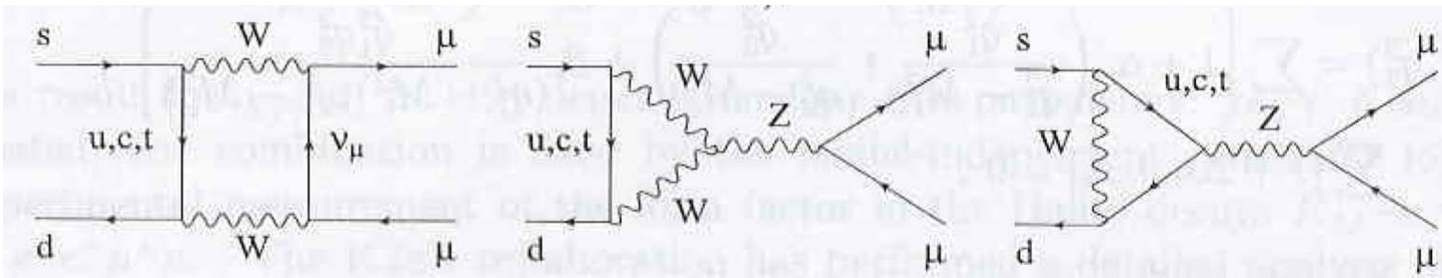

(Cirigliano, *et al.* , Rev. Mod. Phys., **84**, 2012)

- $K_L \to \mu^+ \mu^-$ decay rate is known:
  - BR($K_L \to \mu^+ \mu^-$) = (6.84 ± 0.11) x 10^{-9}
- Large ``background'' from two-photon process:

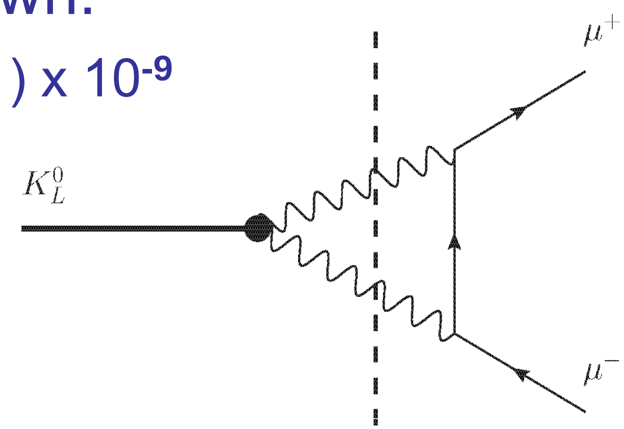

Calculation $\pi^0 \rightarrow e^+ e^-$

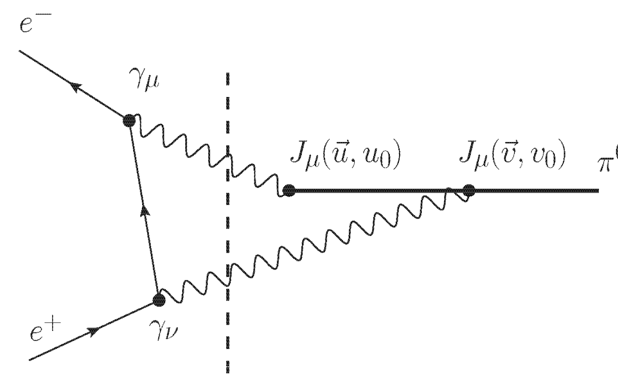

- Avoid Euclidean calculation
- Evaluate in Minkowski space
- Wick rotate internal time integral:

E&M (complex) QCD (real)

$$\mathcal{A}_{\pi^0 \to e^+ e^-} \to \int d^4 w \, \widetilde{L}(k_-, k_+, w)_{\mu\nu} \langle 0 | T\left\{ J_\mu\left(\frac{w}{2}\right) J_\nu\left(-\frac{w}{2}\right) \right\} | \pi^0(\vec{P} = 0) \rangle$$

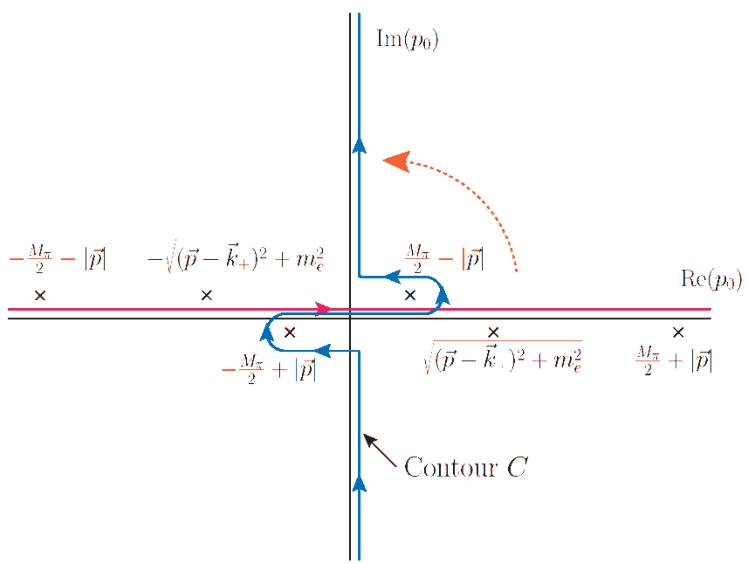

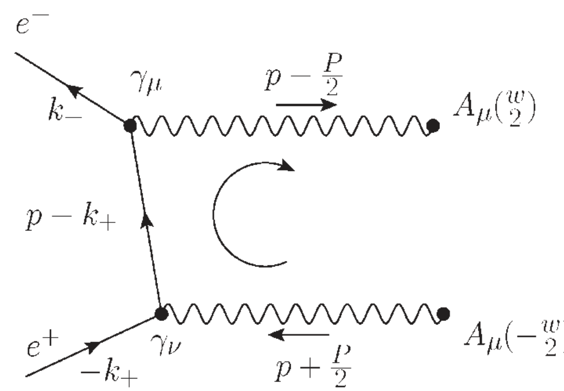

Lattice Results for $\pi^0 \to e^+ e^-$

(Yidi Zhao)

$$\mathcal{A}_{\pi^0 \to e^+ e^-} \to \int d^4w \, \widetilde{L}(k_-, k_+, w)_{\mu\nu} \langle 0 | T\left\{ J_\mu(\frac{w}{2}) J_\nu(-\frac{w}{2}) \right\} | \pi^0(\vec{P} = 0) \rangle$$

- Lattice result is complex:
  - Exponentially small FV corrections
  - Physical kinematics, $1/a \leq 1.73$ GeV :
    - Im(A) = 35.94(1.01)(1.09) [Expt: 35.07(37)]
    - Re(A) = 20.39(72)(70). [Expt: 21.51(2.02)]
- First step in predicting $K_L \to \mu^+ \mu^-$
  - Lattice QCD calculation needed to remove two-photon background.
  - Would allow ~10% test of standard model prediction for rare, 2nd order weak decay

$K_L - K_S$ mass difference

$K^0 - \overline{K}^0$ system

- $K^0 - \overline{K}^0$ are distinct anti-particles: ($\overline{s}\,d$) and ($\overline{d}\,s$) bound states
- These are mixed by the strangeness-violating weak couplings:

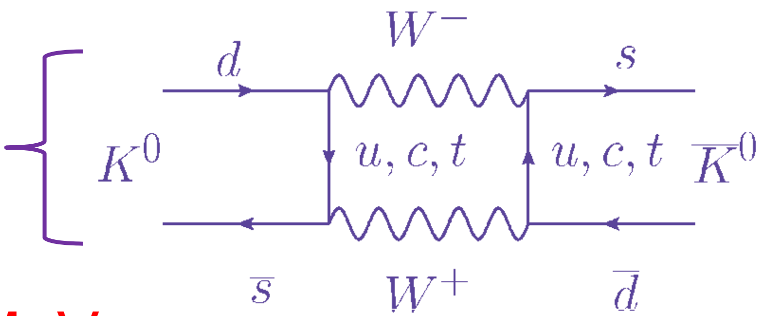

- $M_{K_L} - M_{K_S} = 3.483(6) \times 10^{-12}$ MeV
- Sensitive to 1000 TeV energy scale, 1000 x LHC energies
- Evidence for charm quark energy scale first found here.
- Effects of QCD can now be computed from first principles!

Lattice Version

- Evaluate standard, Euclidean, 2nd order $\overline{K}^0 - K^0$ amplitude:

$$\mathcal{A} = \langle 0|T\left(K^0(t_f)\frac{1}{2}\int_{t_a}^{t_b} dt_2 \int_{t_a}^{t_b} dt_1 H_W(t_2) H_W(t_1) K^{0\,\dagger}(t_i)\right)|0\rangle$$

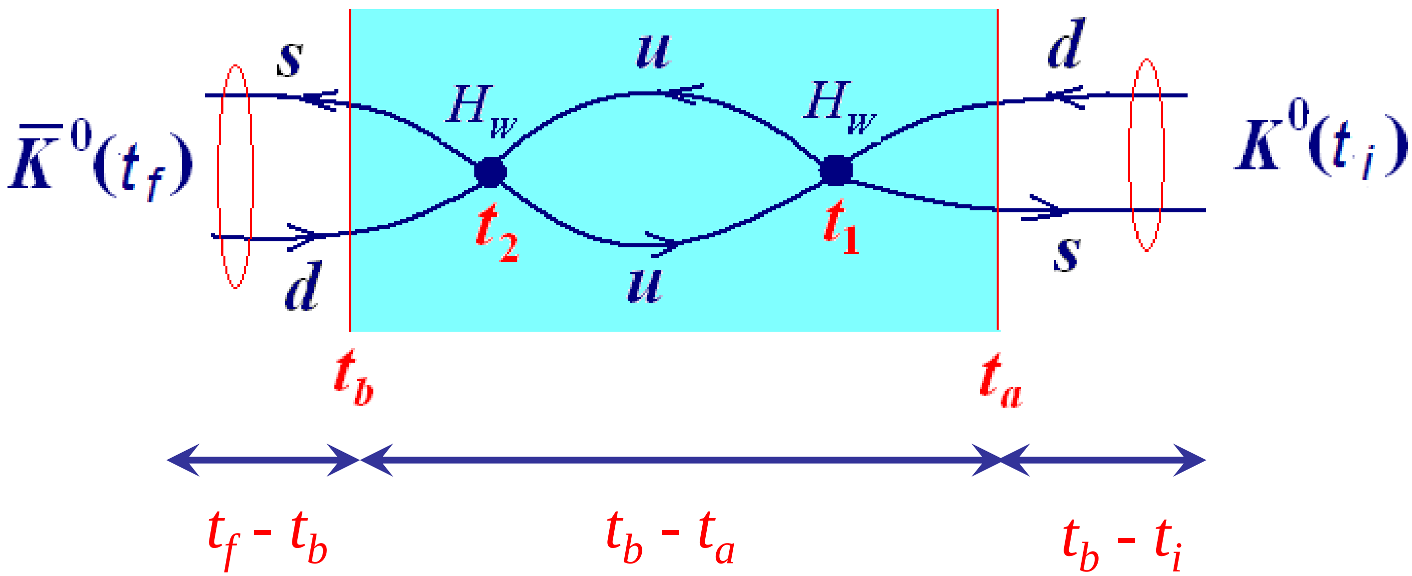

ΔM_K Preliminary Results

(Bigeng Wang)

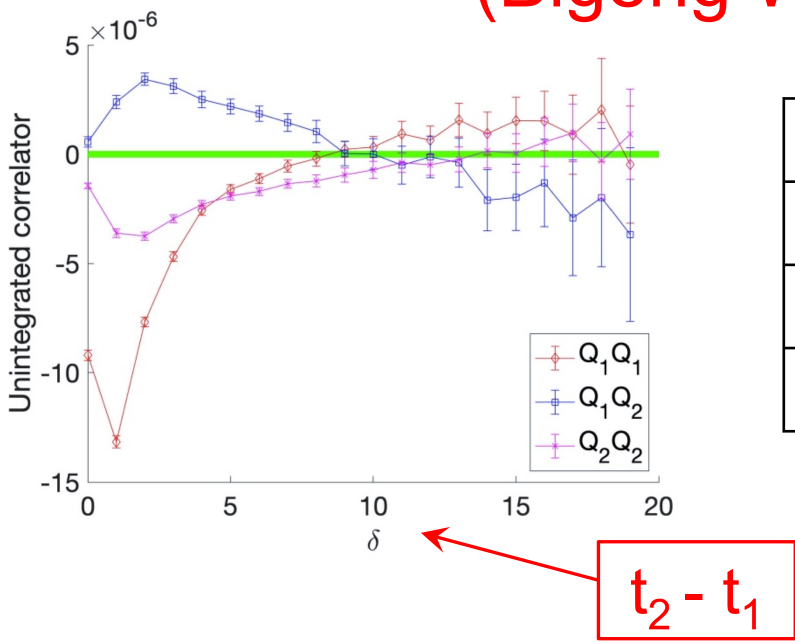

	ΔM_K x 10^{+12} MeV
ΔM_K	6.7(0.6)
Δ_{FV}	0.27(18)
Expt.	3.483(6)

- Physical light, strange and charm masses
- 64^3 x 128, $1/a$ = 2.36 GeV
- Integrate: $0 \leq \delta \leq 10$
- 152 configurations
- a^2 errors 20% ?

Many new lattice results

- Study two-nucleon problem (S. Aoki, *et al.* HAL QCD)
- Use Euclidean lattice data to constrain spectral functions in inclusive decays (S. Hashimoto, *et al.*)
- Results for *Kl*3 decay on largest lattice to date (K.-I. Ishikawa, *et al.*, PACS Collaboration)
- Calculate long-distance part (5%) of direct CP violation in $K \rightarrow \pi\pi$ decay: ε_K (J. Karpie):
- $K \rightarrow \pi\pi$ decay and direct CP violation (RBC/UKQCD)
  re (ε'/ε) = (1.38 ± 7) x 10^{-4} (lattice)
  (16.6 ± 0.23) x 10^{-4} (Expt)

Lattice QCD at the Exascale

- Important opportunities to discover new physics beyond the standard model.
- Much work still needed to compensate for the increasingly weak inter-node network.
- Target smaller lattice spacing
  - Increase accuracy of charm quark physics
  - Use open boundary conditions (Luscher & Schaefer)
  - Fourier accelerate: HMC → RMHMC (Girolami & Calderhead)
  - Studying QCD at β=100: