The Exascale Era and What to Expect in 2016+

Andreas S. Kronfeld
Fermilab
1 March 2011
Workshop on the Hadronic Light-by-light Contribution
to the Muon Anomaly (g–2)
Institute for Nuclear Theory, University of Washington

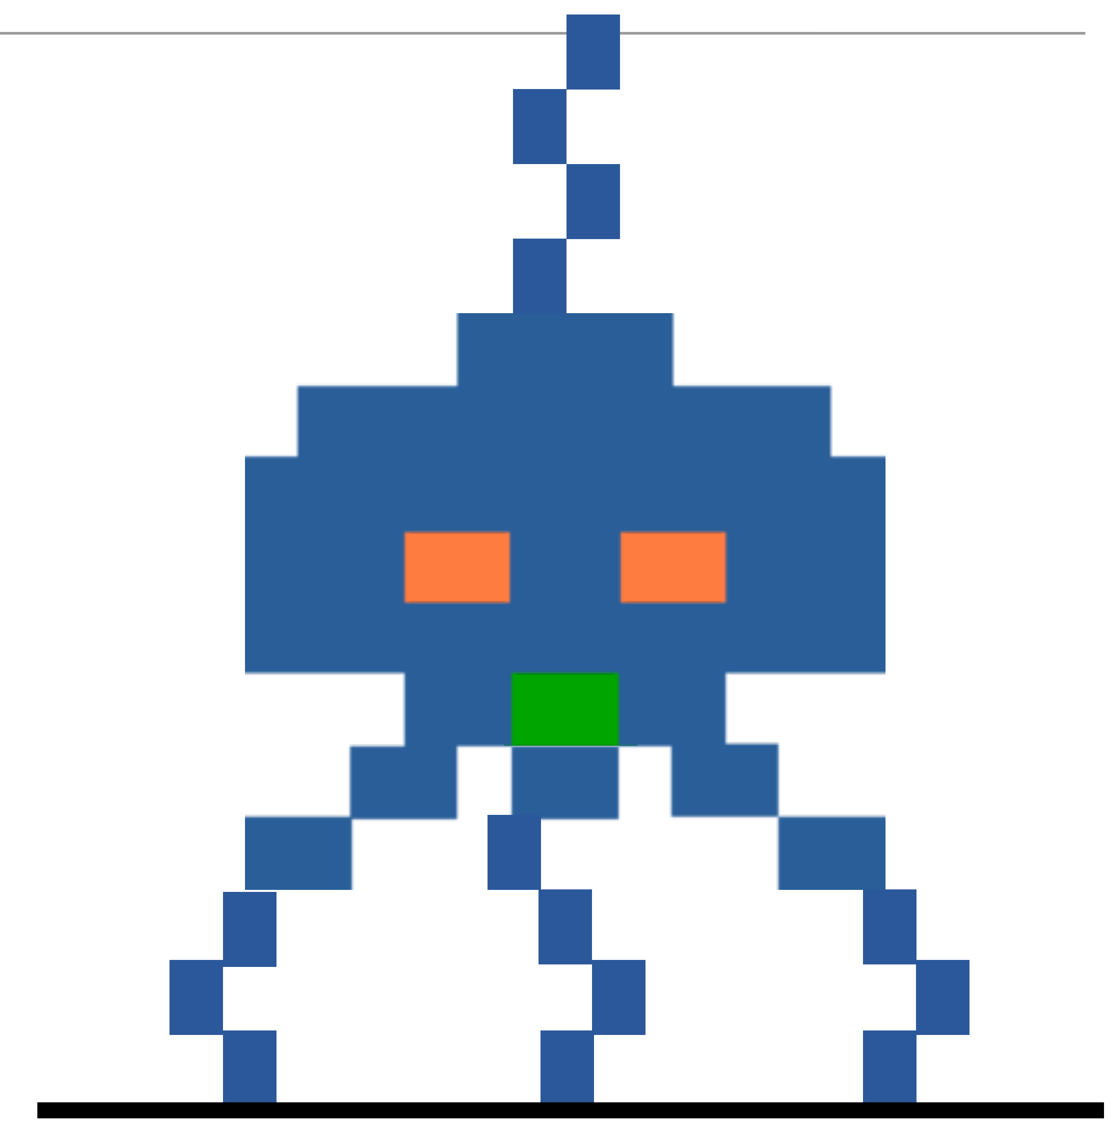

Fostering Dialogue

- We'd like to ask you to give a talk at the workshop….
- I'd hoped this would be a chance to sit back and learn for a change….
- What we have in mind is a talk on the future computing needs for g–2 and the anticipated resources.
- But I would have to talk to lots of people first. I don't really know what I'm talking about.
- That's right, but you're good at that.

Outline

- Computing needs for lattice QCD:
  - general outline and scaling laws;
  - achievements with past & present resources;
  - g–2 special needs.
- Remarks on available resources: USQCD-centric.
- Forecasts for computing and for g–2:
  - note that I brought sunglasses but no umbrella to Seattle.

Lattice Gauge Theory in a Nutshell

Lattice Gauge Theory

K. Wilson, *PRD* **10** (1974) 2445

- Invented to understand asymptotic freedom without the need for gauge-fixing and ghosts [Wilson, hep-lat/0412043].

- Gauge symmetry on a spacetime lattice:

    - mathematically rigorous definition of QCD functional integrals;

$$\langle \bullet \rangle = \frac{1}{Z} \int \mathcal{D}U \, \mathcal{D}\psi \mathcal{D}\bar{\psi} \exp\left(-S\right) [\bullet]$$

    - enables theoretical tools of statistical mechanics in quantum field theory and provides a basis for constructive field theory.

- Lowest-order strong coupling expansion demonstrates confinement.

Numerical Lattice QCD

- Nowadays "lattice QCD" usually implies a numerical technique.
- Integrate the functional integral on a $N_S^3 \times N_4$ lattice (spacing a) numerically:

$$\langle \bullet \rangle = \frac{1}{Z} \int \mathcal{D}U \; \mathcal{D}\psi \, \mathcal{D}\bar{\psi} \; \exp(-S) [\bullet]$$
$$= \frac{1}{Z} \int \mathcal{D}U \det(\not{D} + m) \exp(-S) [\bullet']$$

- Finite lattice: can evaluate integrals on a computer; dimension ~ 10^8, using *importance sampling*.
- Healthy research field to devise MC algorithms.

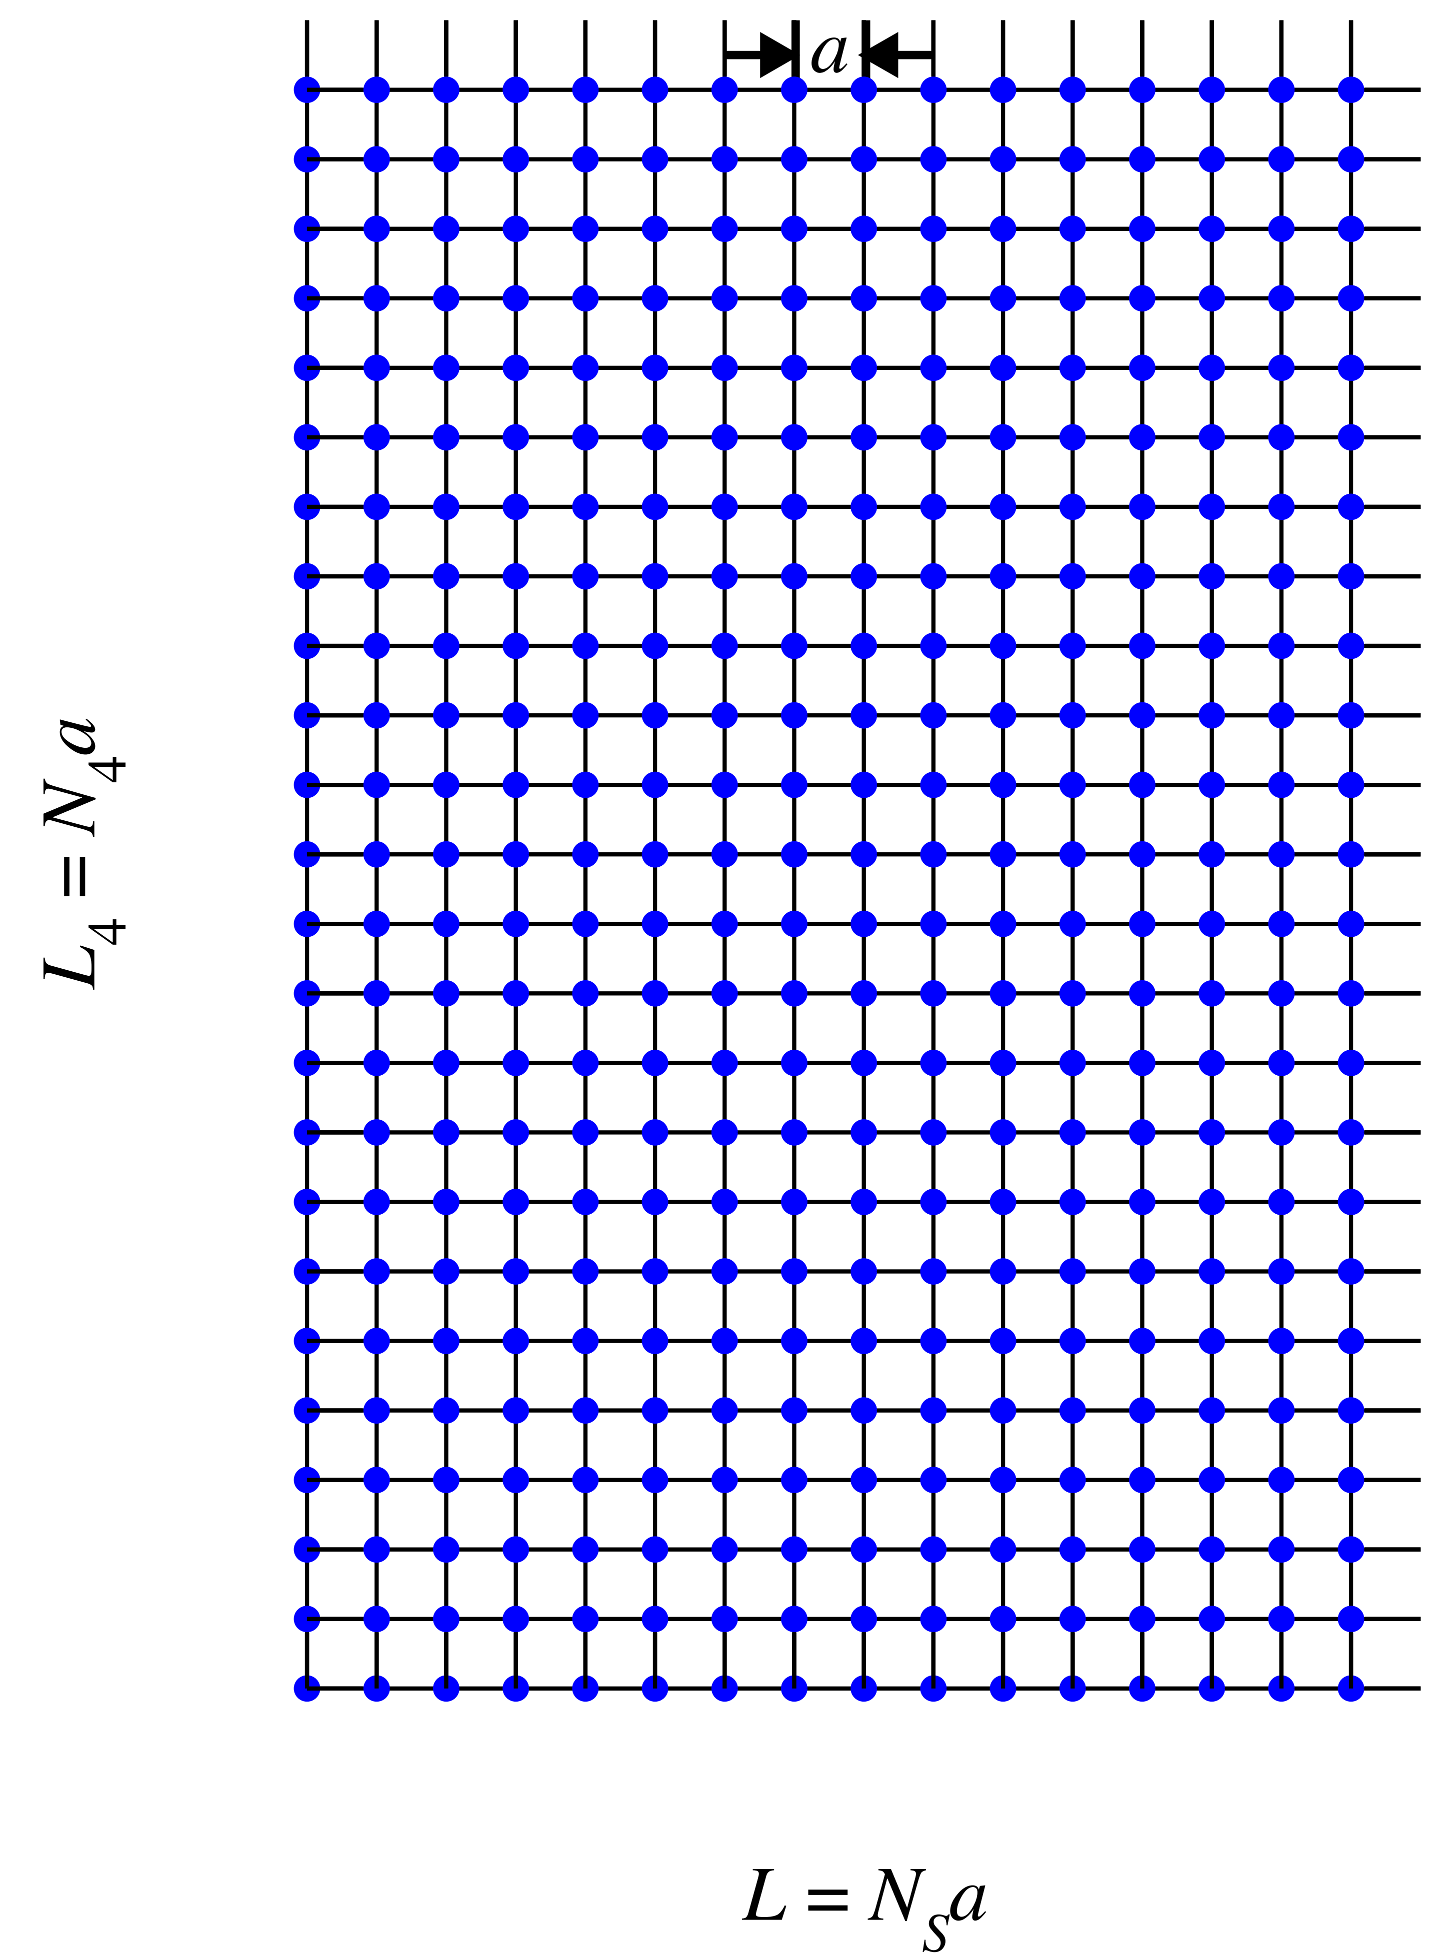

Numerical Lattice QCD

- Nowadays "lattice QCD" usually implies a numerical technique.

- Integrate the functional integral on a $N_S^3 \times N_4$ lattice (spacing a) numerically:

$$\langle \bullet \rangle = \frac{1}{Z} \int \mathcal{D}U \, \mathcal{D}\psi \, \mathcal{D}\bar{\psi} \, \exp(-S) \, [\bullet]$$

hand

$$= \frac{1}{Z} \int \mathcal{D}U \, \det(\not{D} + m) \exp(-S) \, [\bullet']$$

- Finite lattice: can evaluate integrals on a computer; dimension ~ 10^8, using *importance sampling*.

- Healthy research field to devise MC algorithms.

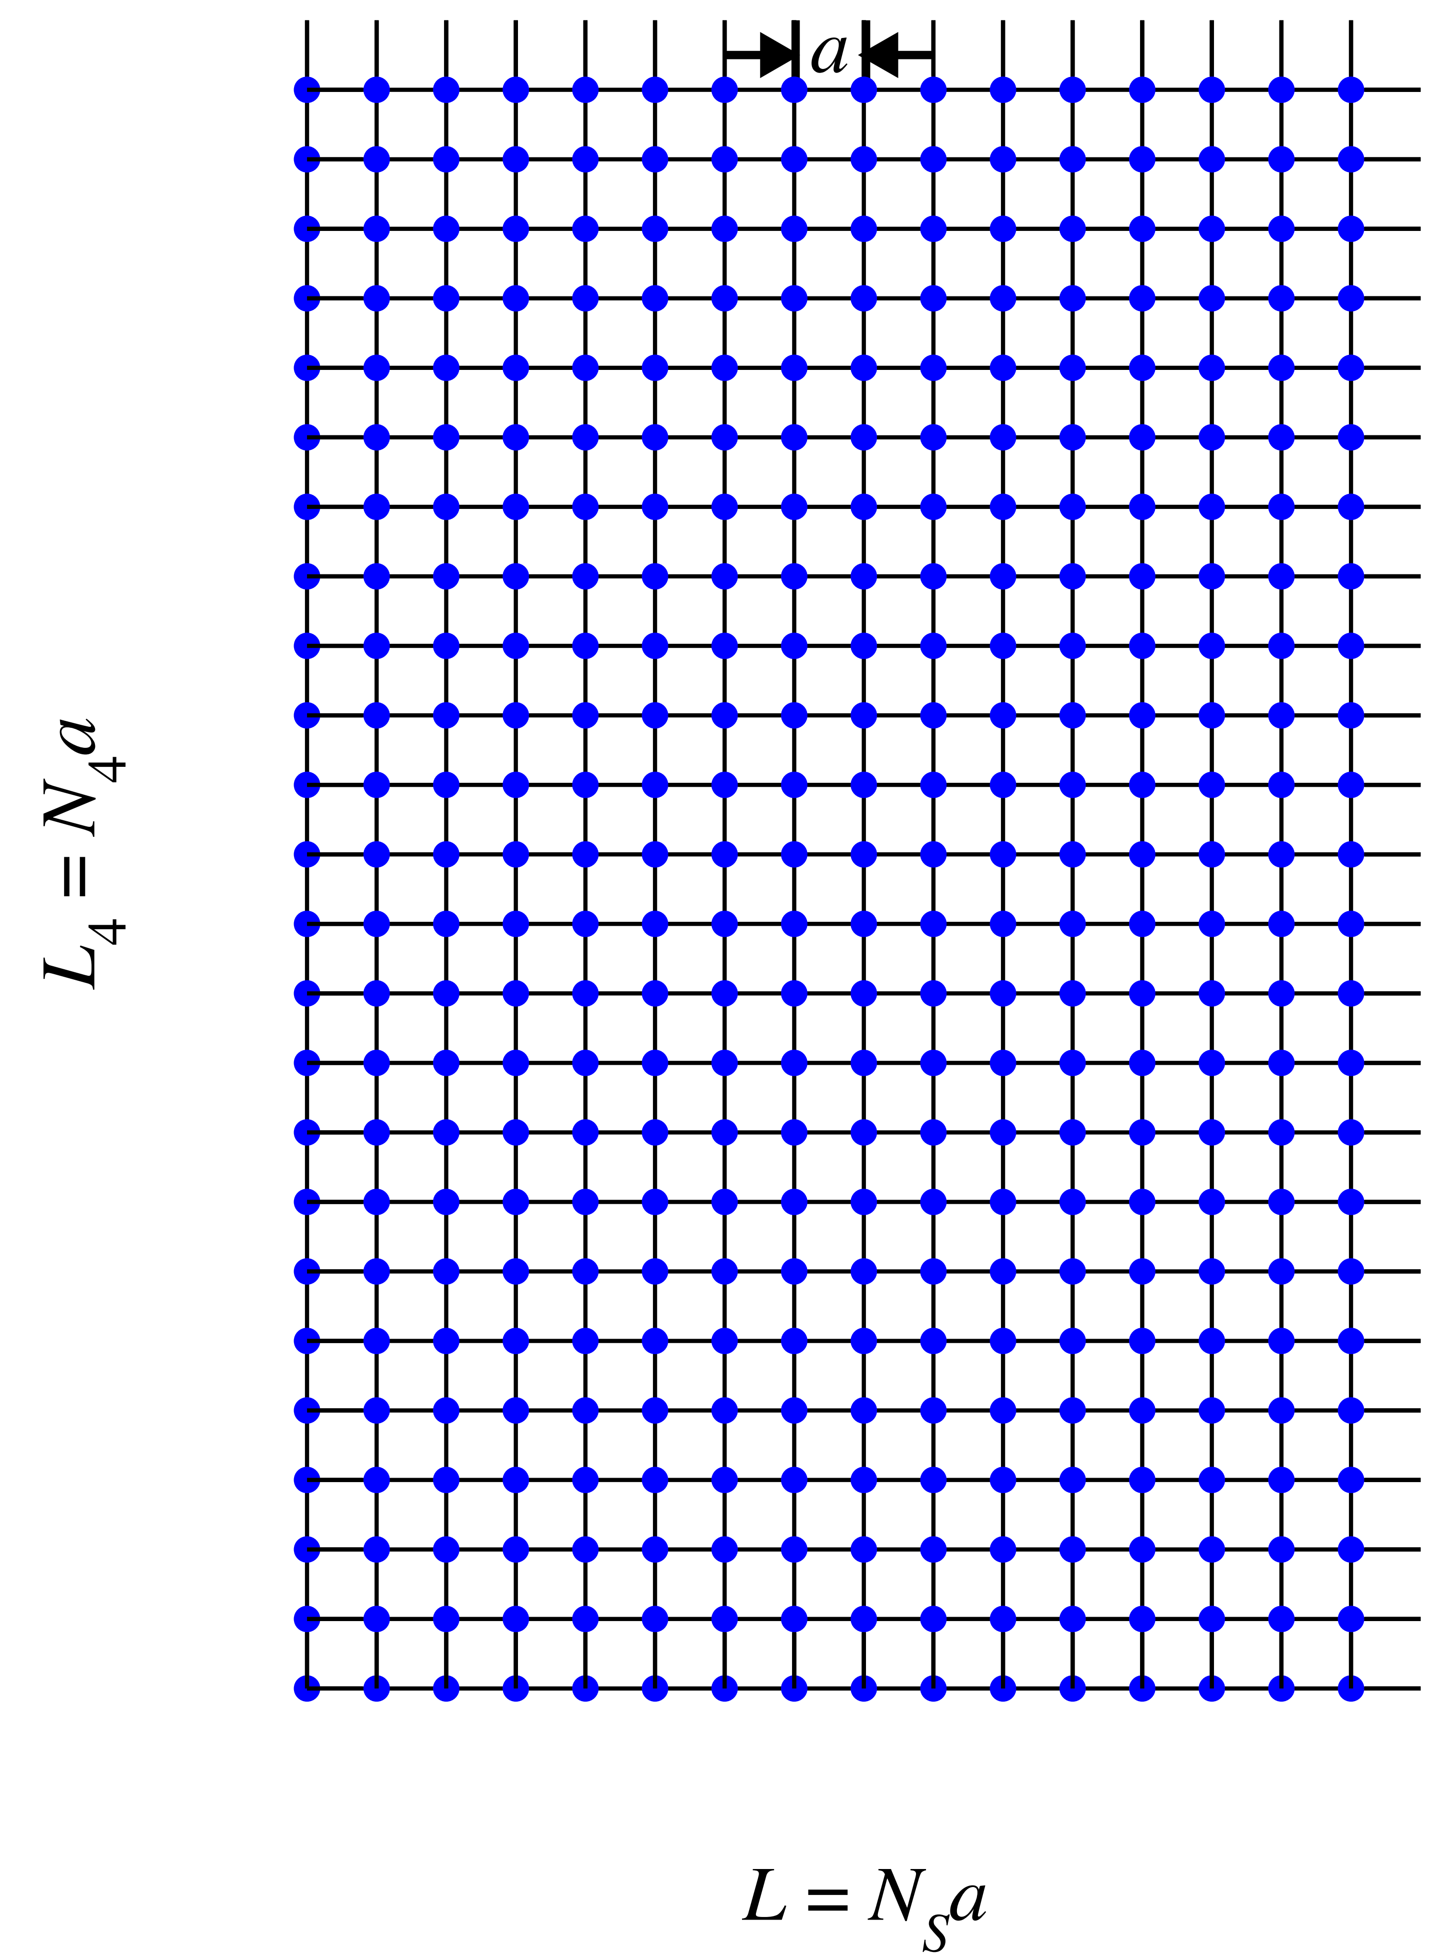

Numerical Lattice QCD

- Nowadays "lattice QCD" usually implies a numerical technique.

- Integrate the functional integral on a $N_S^3 \times N_4$ lattice (spacing a) numerically:

$$\langle \bullet \rangle = \frac{1}{Z} \int \underbrace{\mathcal{D}U}_{\text{MC}} \; \underbrace{\mathcal{D}\psi \, \mathcal{D}\bar{\psi}}_{\text{hand}} \exp(-S) \, [\bullet]$$

$$= \frac{1}{Z} \int \mathcal{D}U \det(\not{D} + m) \exp(-S) \, [\bullet']$$

- Finite lattice: can evaluate integrals on a computer; dimension ~ 10^8, using *importance sampling*.

- Healthy research field to devise MC algorithms.

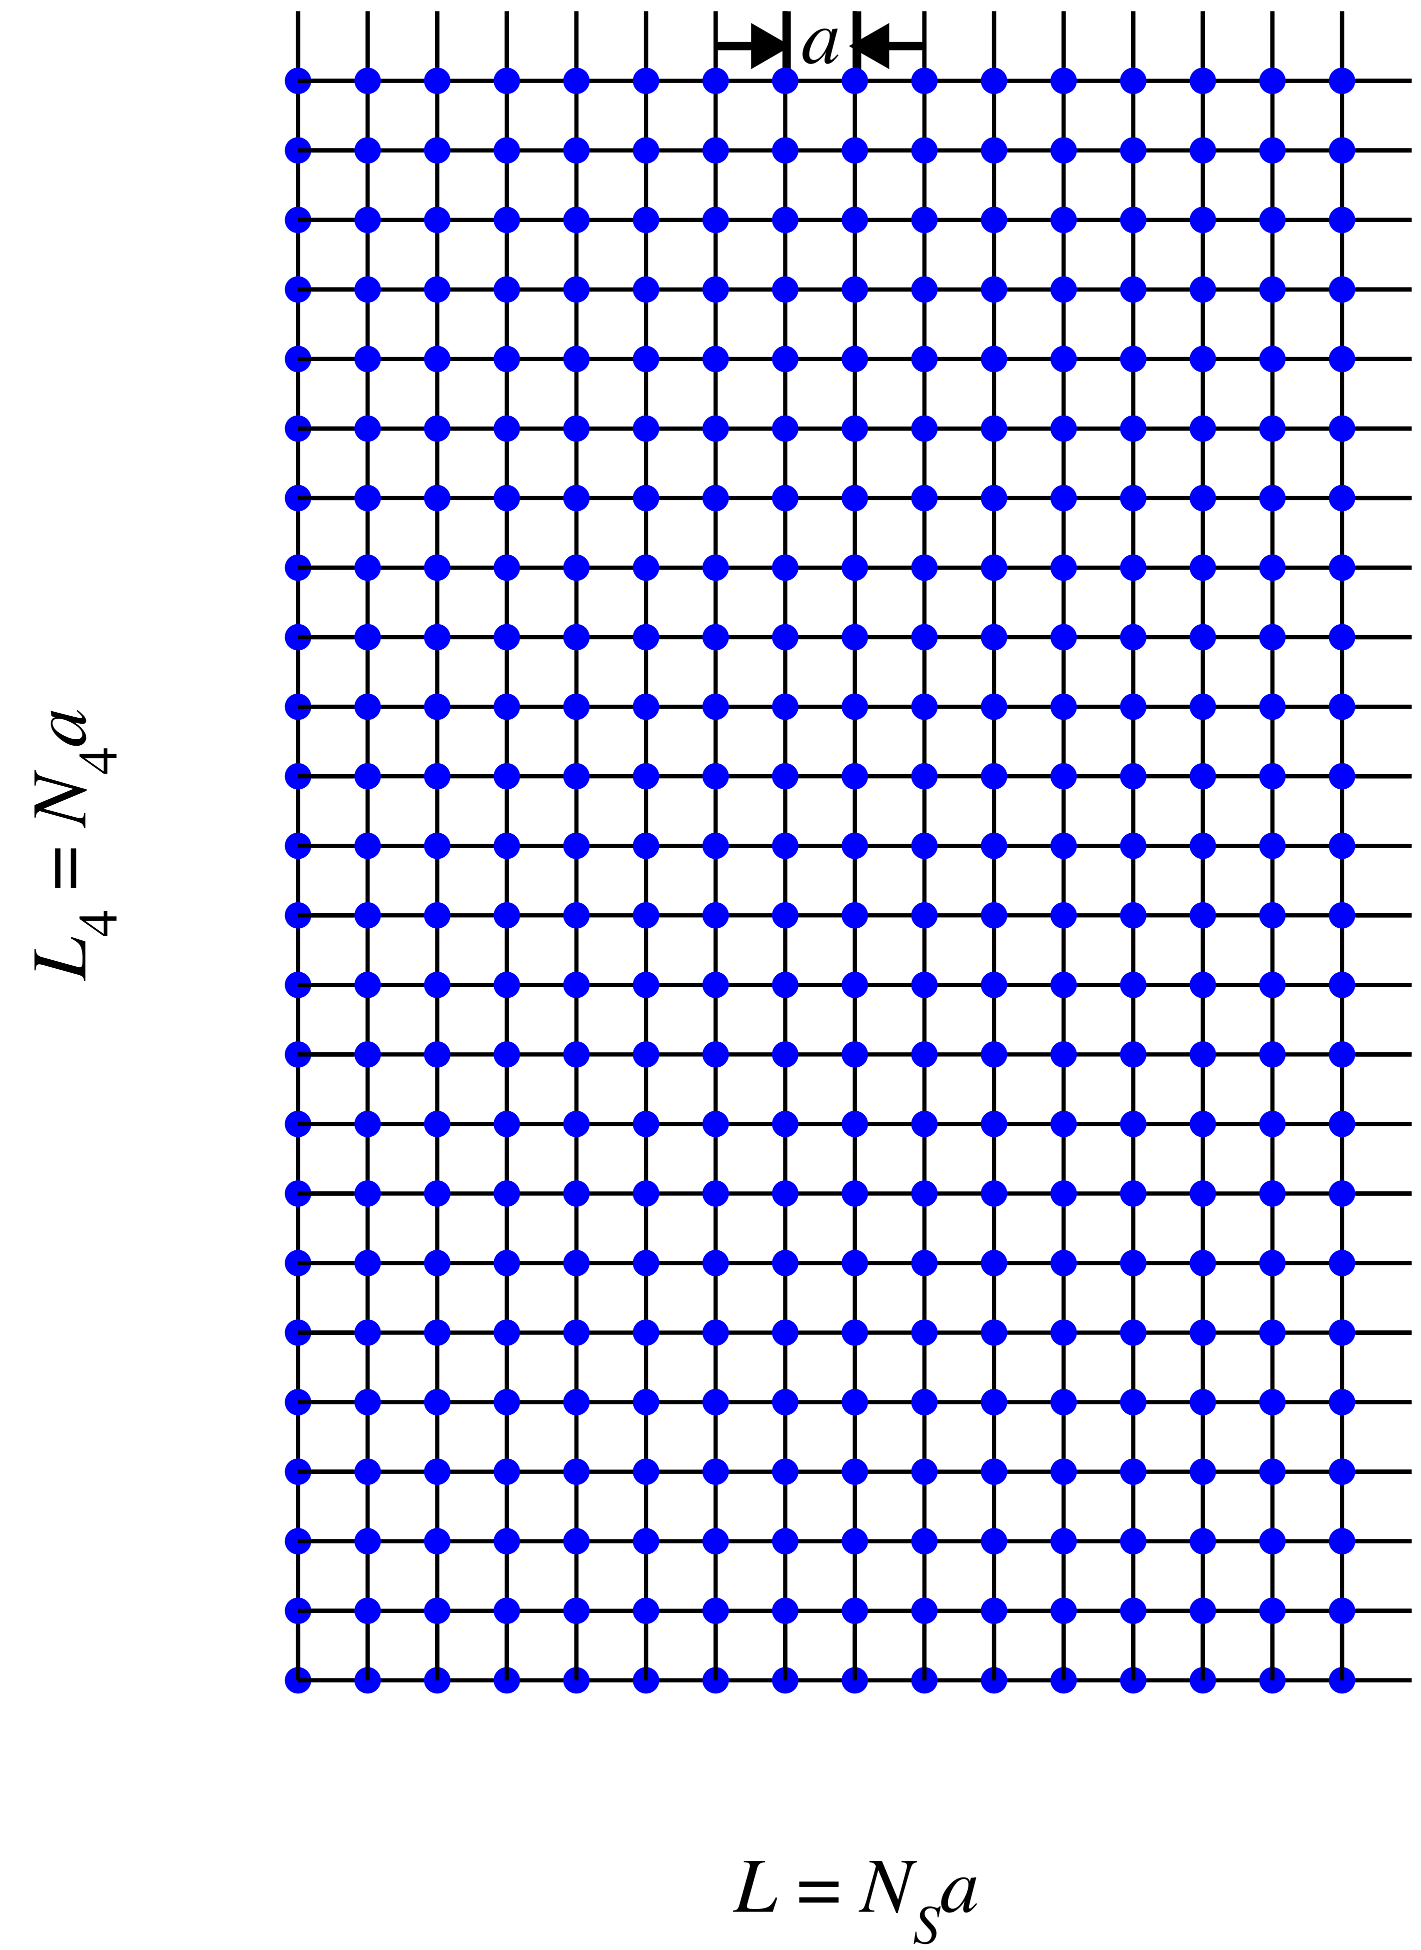

Some algorithmic issues

e.g., ASK, hep-lat/0205021

- lattice $N_S^3 \times N_4$, spacing a
- memory $\propto N_S^3 N_4 = L_S^3 L_4 / a^4$
- $\tau_g \propto a^{-(4+z)}$, $z = 1$ or 2.
- $\tau_q \propto (m_q a)^{-p}$, $p = 1$ or 2.
- Imaginary time:
  - static quantities
  - or Euclidean Green functions

- (notation)
- dimension of spacetime = 4
- critical slowing down
- especially dire with sea quarks
- thermodynamics: $T = (N_4 a)^{-1}$

$$\langle \bullet \rangle = \frac{1}{Z} \int \mathcal{D}U \, \mathcal{D}\psi \mathcal{D}\bar{\psi} \exp(-S) \, [\bullet]$$
$$= \mathrm{Tr}\{\bullet \, e^{-\hat{H}/T}\} / \mathrm{Tr}\{e^{-\hat{H}/T}\}$$

- Some compromises:
  - finite human lifetime ⇒ Wick rotate to Euclidean time: $x^4 = ix^0$;
  - finite memory ⇒ finite space volume & finite time extent; nonzero lattice spacing;
  - finite CPU power ⇒ light quarks heavier than up and down; nonzero lattice spacing.
- The first introduces no error, but can be an obstacle (*e.g.*, fragmentation functions).
- Finite volume unimportant for stable hadrons (like external photons and muons)…
  - … but strong effects for resonances (like rho) and massless internal particles (photons).

Some Jargon

- QCD observables (quark integrals by hand):

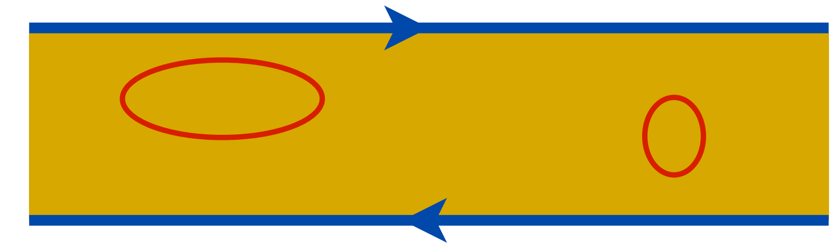

$$\langle \bullet \rangle = \frac{1}{Z} \int \mathcal{D}U \prod_{f=1}^{n_f} \det(\not{D} + m_f) \exp\left(-S_{\text{gauge}}\right) [\bullet']$$

sea   valence: $(\not{D} + m)^{-1}$

- *Quenched* means replace $\det$ with 1. (Obsolete.)

- *Un*quenched means not to do that.

- *Partially* quenched (usually) doesn't mean "n_f too small", but $m_{\text{val}} \neq m_{\text{sea}}$, or $\not{D}_{\text{val}} \neq \not{D}_{\text{sea}}$ ("mixed action").

Sea Quarks

- Staggered quarks, with rooted determinant, $\mathrm{O}(a^2)$.
- Wilson quarks, $\mathrm{O}(a)$:
  - twisted mass term—auto $\mathrm{O}(a)$ improvement $\Rightarrow \mathrm{O}(a^2)$;
  - tree or nonperturbatively $\mathrm{O}(a)$ improved $\Rightarrow \mathrm{O}(a^2)$.
- Ginsparg-Wilson (domain wall or overlap), $\mathrm{O}(a^2)$:
  - $\not{D}\gamma_5 + \gamma_5\not{D} = 2a\not{D}^2$ implemented w/ $\mathrm{sign}(\not{D}_\mathrm{W})$.

Sea Quarks

- Staggered quarks, with rooted determinant, $\mathrm{O}(a^2)$.
- Wilson quarks, $\mathrm{O}(a)$:
  - twisted mass term—auto $\mathrm{O}(a)$ improvement $\Rightarrow \mathrm{O}(a^2)$;
  - tree or nonperturbatively $\mathrm{O}(a)$ improved $\Rightarrow \mathrm{O}(a^2)$.
- Ginsparg-Wilson (domain wall or overlap), $\mathrm{O}(a^2)$:
  - $\not{D}\gamma_5 + \gamma_5\not{D} = 2a\not{D}^2$ implemented w/ $\mathrm{sign}(\not{D}_\mathrm{W})$.

fast

↓

clean

2+1 Sea Quarks!

HPQCD, MILC, Fermilab Lattice, hep-lat/0304004

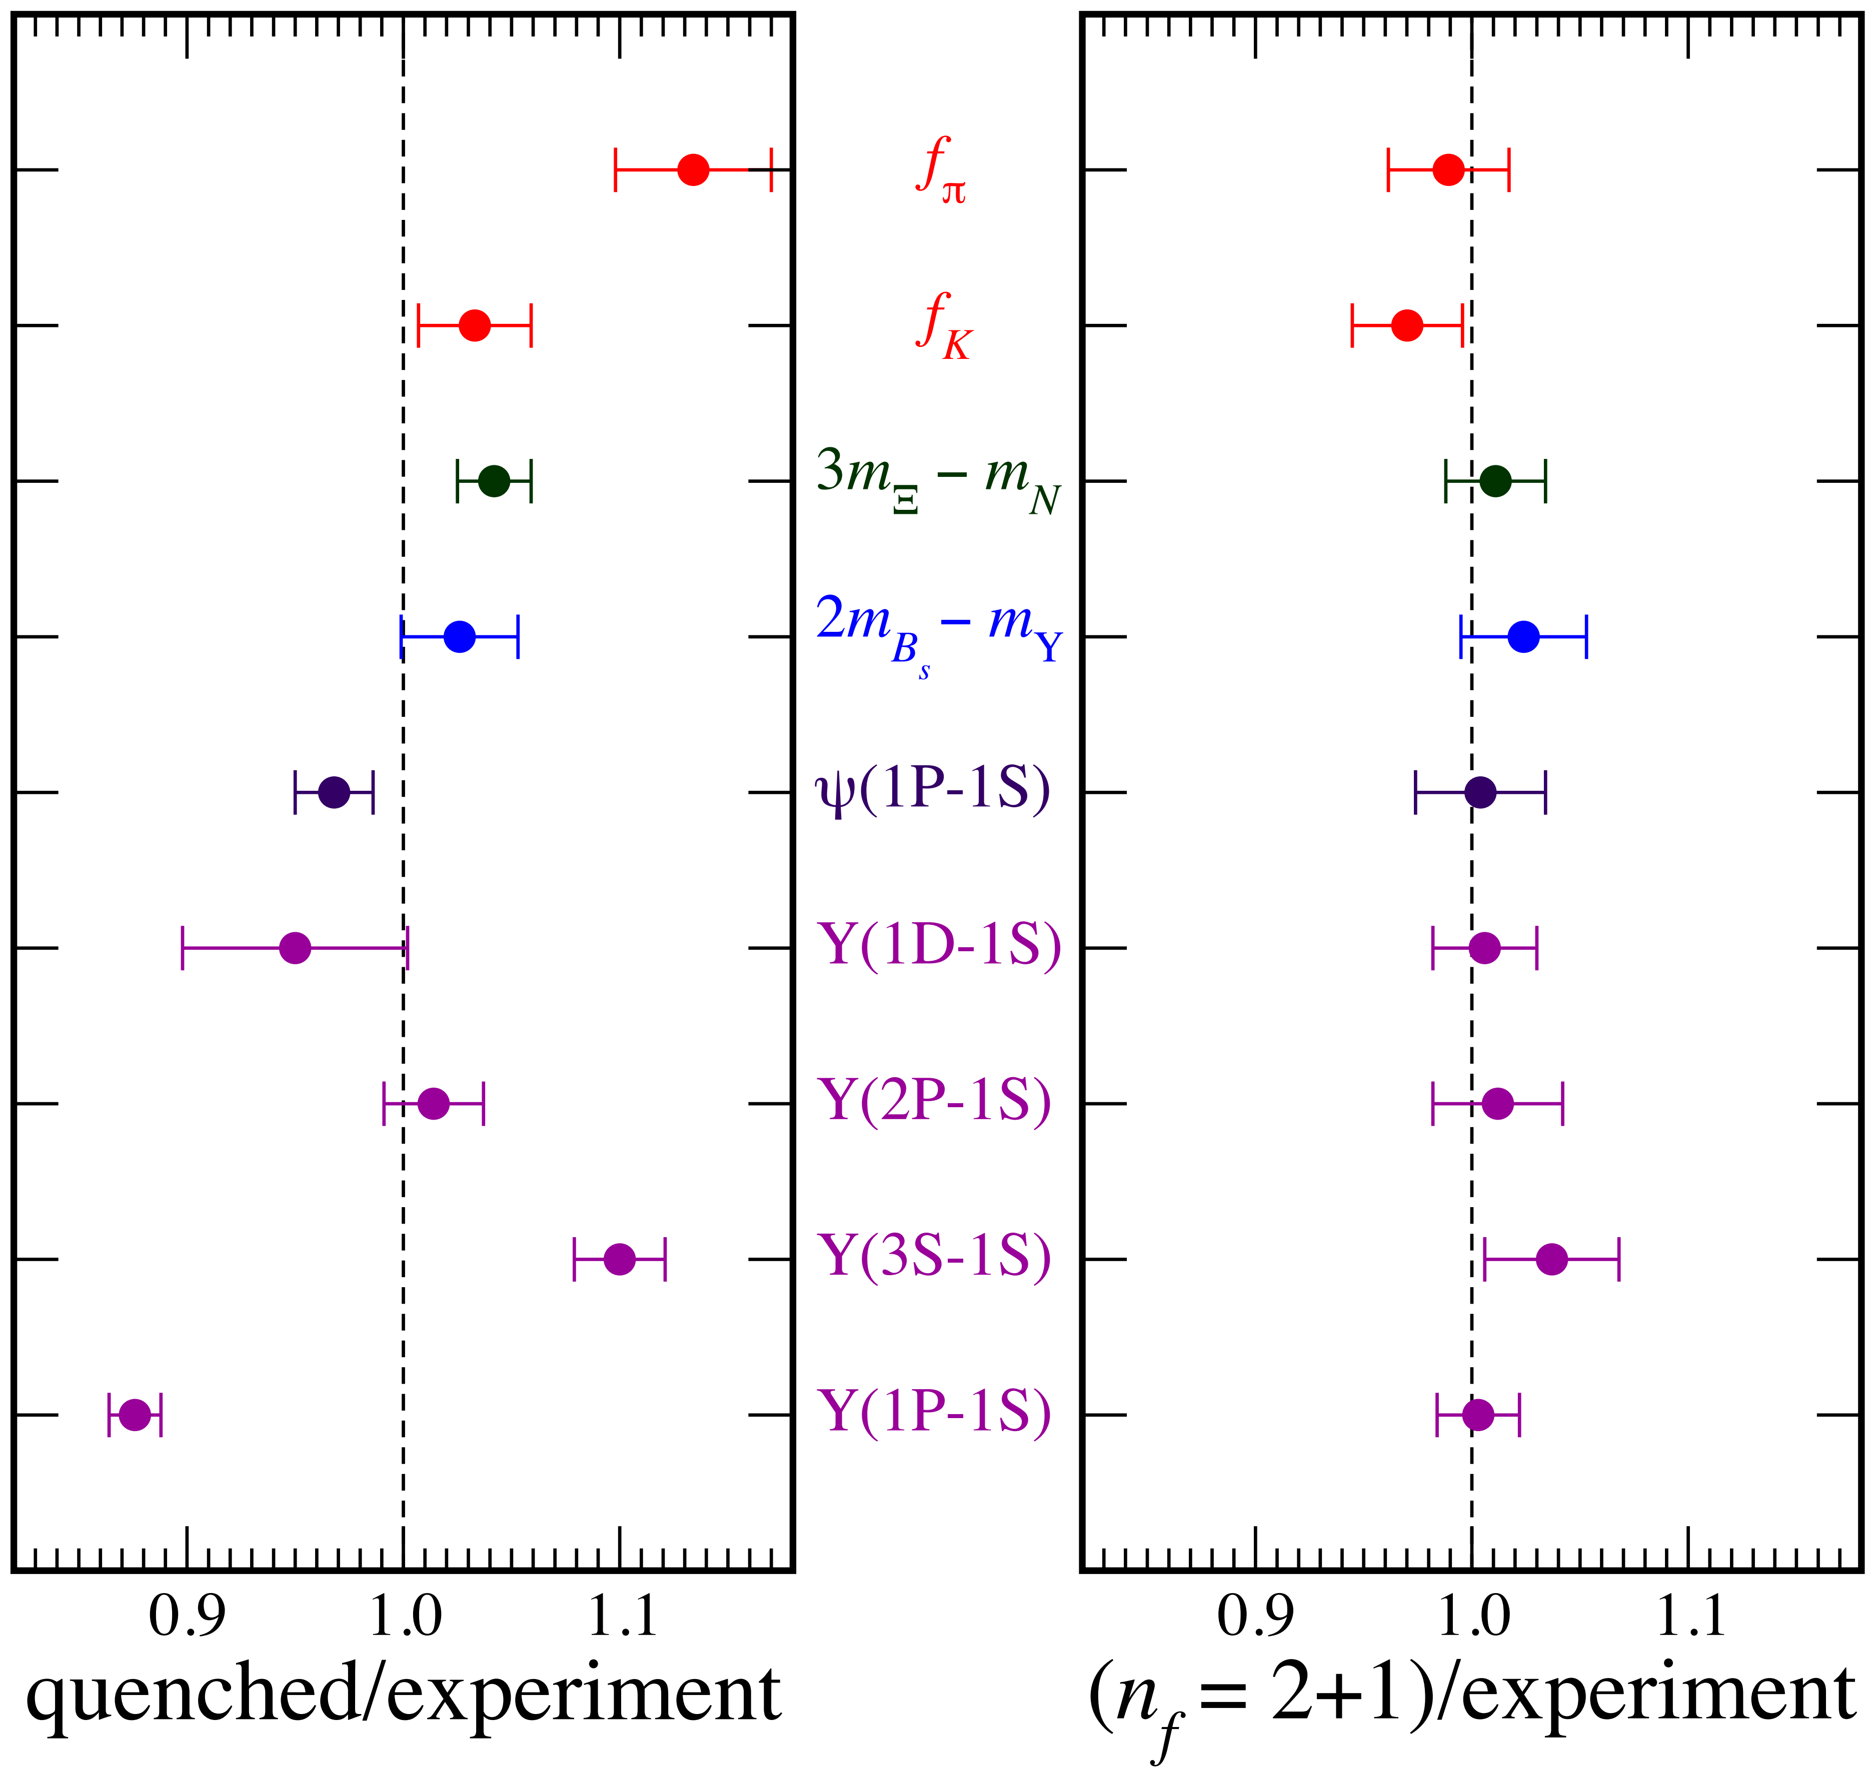

- $a = 0.12$ & 0.09 fm;
- O(a^2) improved: asqtad;
- FAT7 smearing;
- $2m_l < m_q < m_s$;
- π, K, Υ(2S-1S) input.
- Updates with smaller a, and smaller m_q.

Predictions

(*i.e.*, calculation preceded measurement)

Fermilab Lattice, MILC, HPQCD, hep-ph/0408306, hep-lat/0411027, hep-lat/0506030

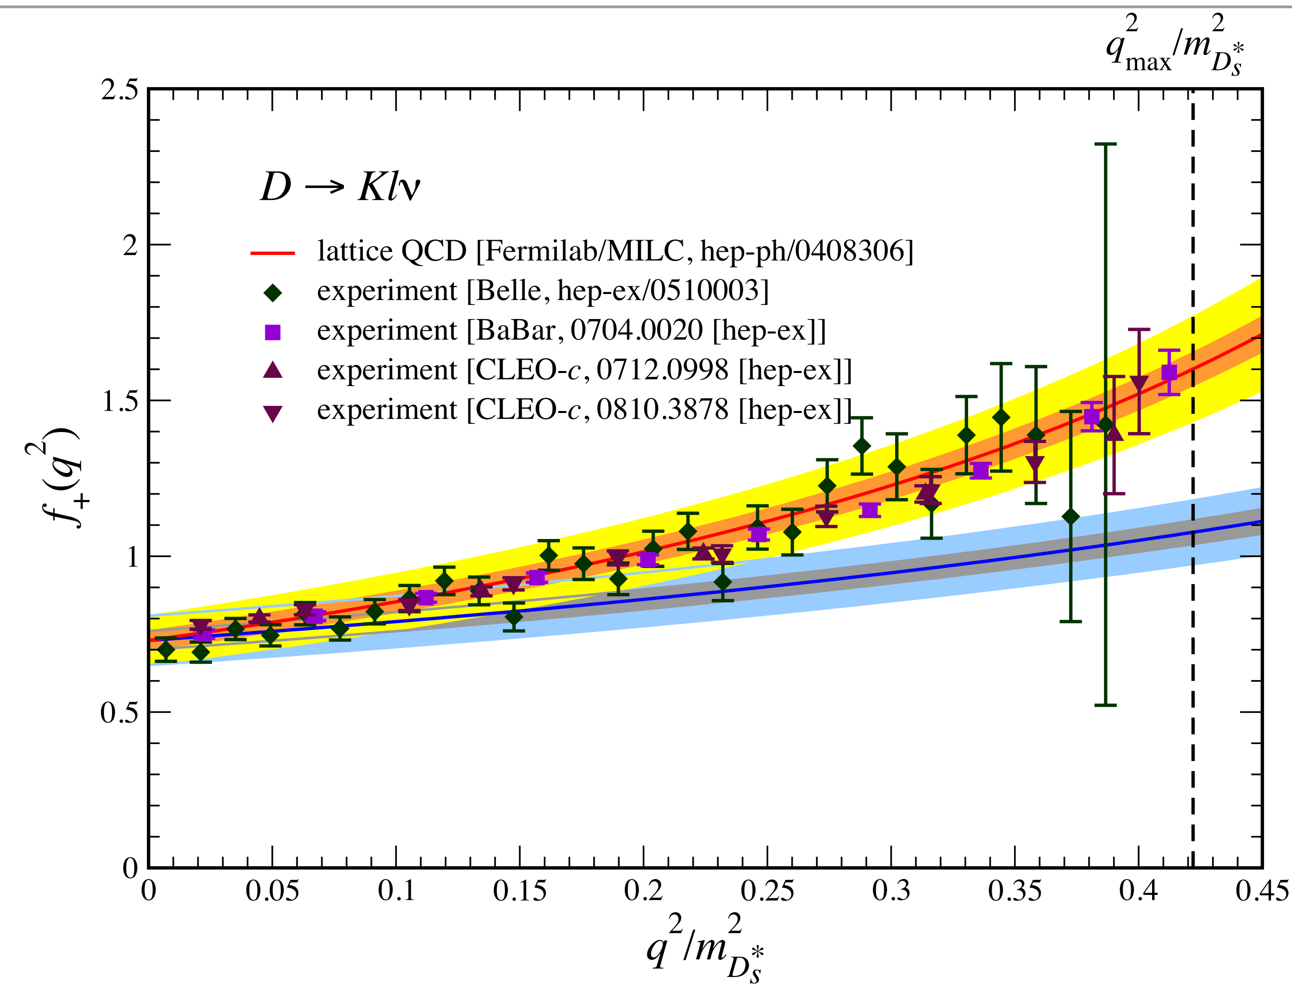

- Semileptonic form factor for $D \to Kl\nu$
- Mass of B_c meson
- Charmed-meson decay constants

2004

2005

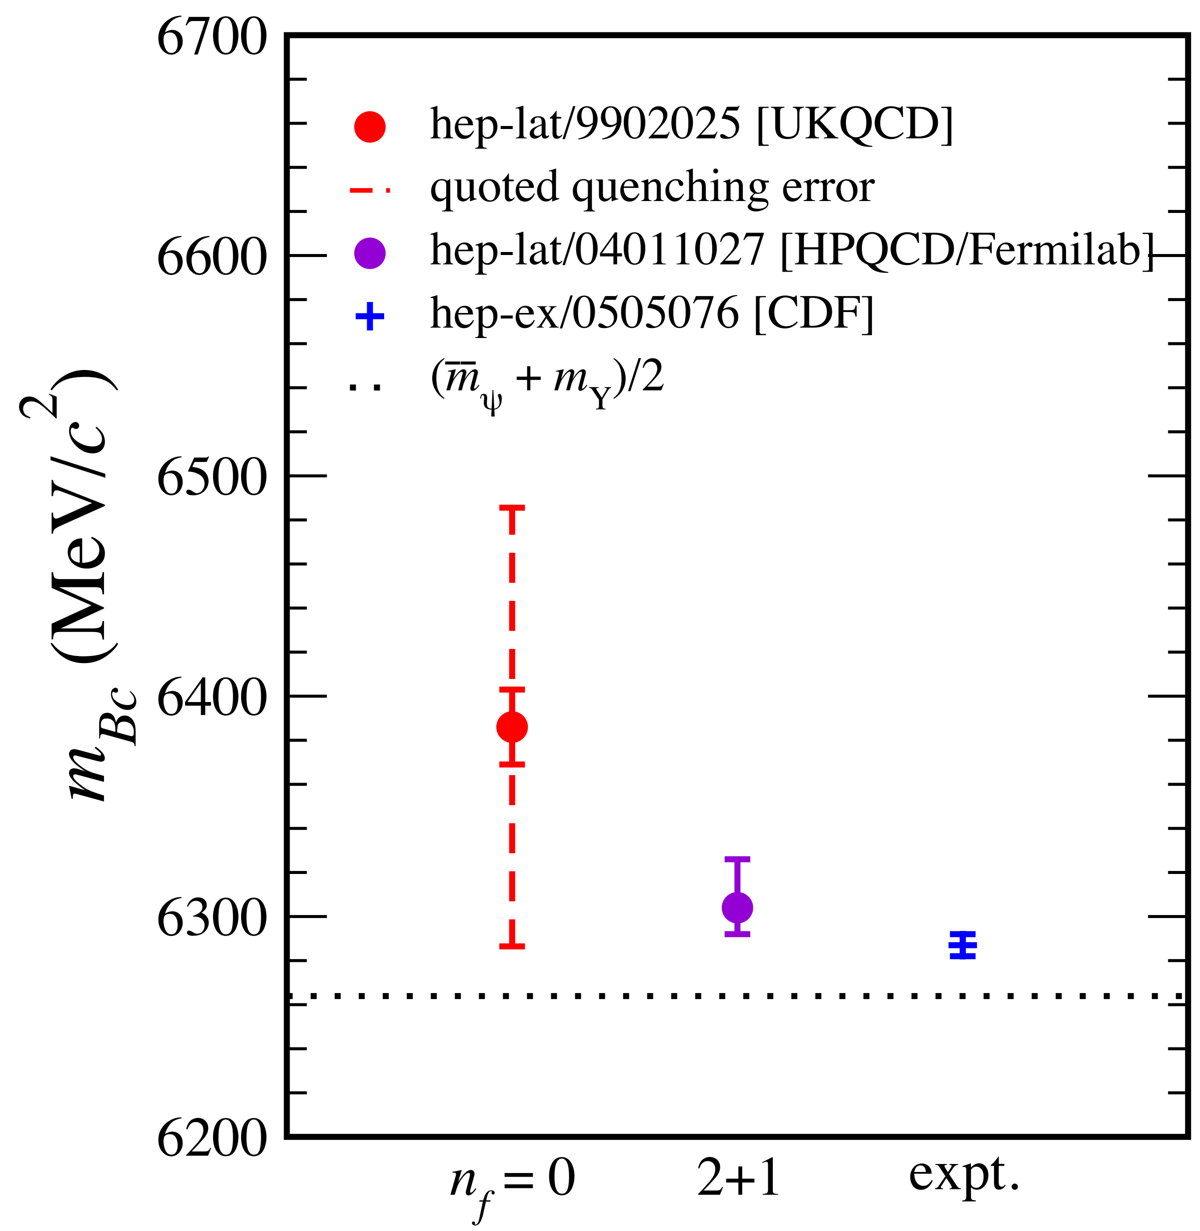

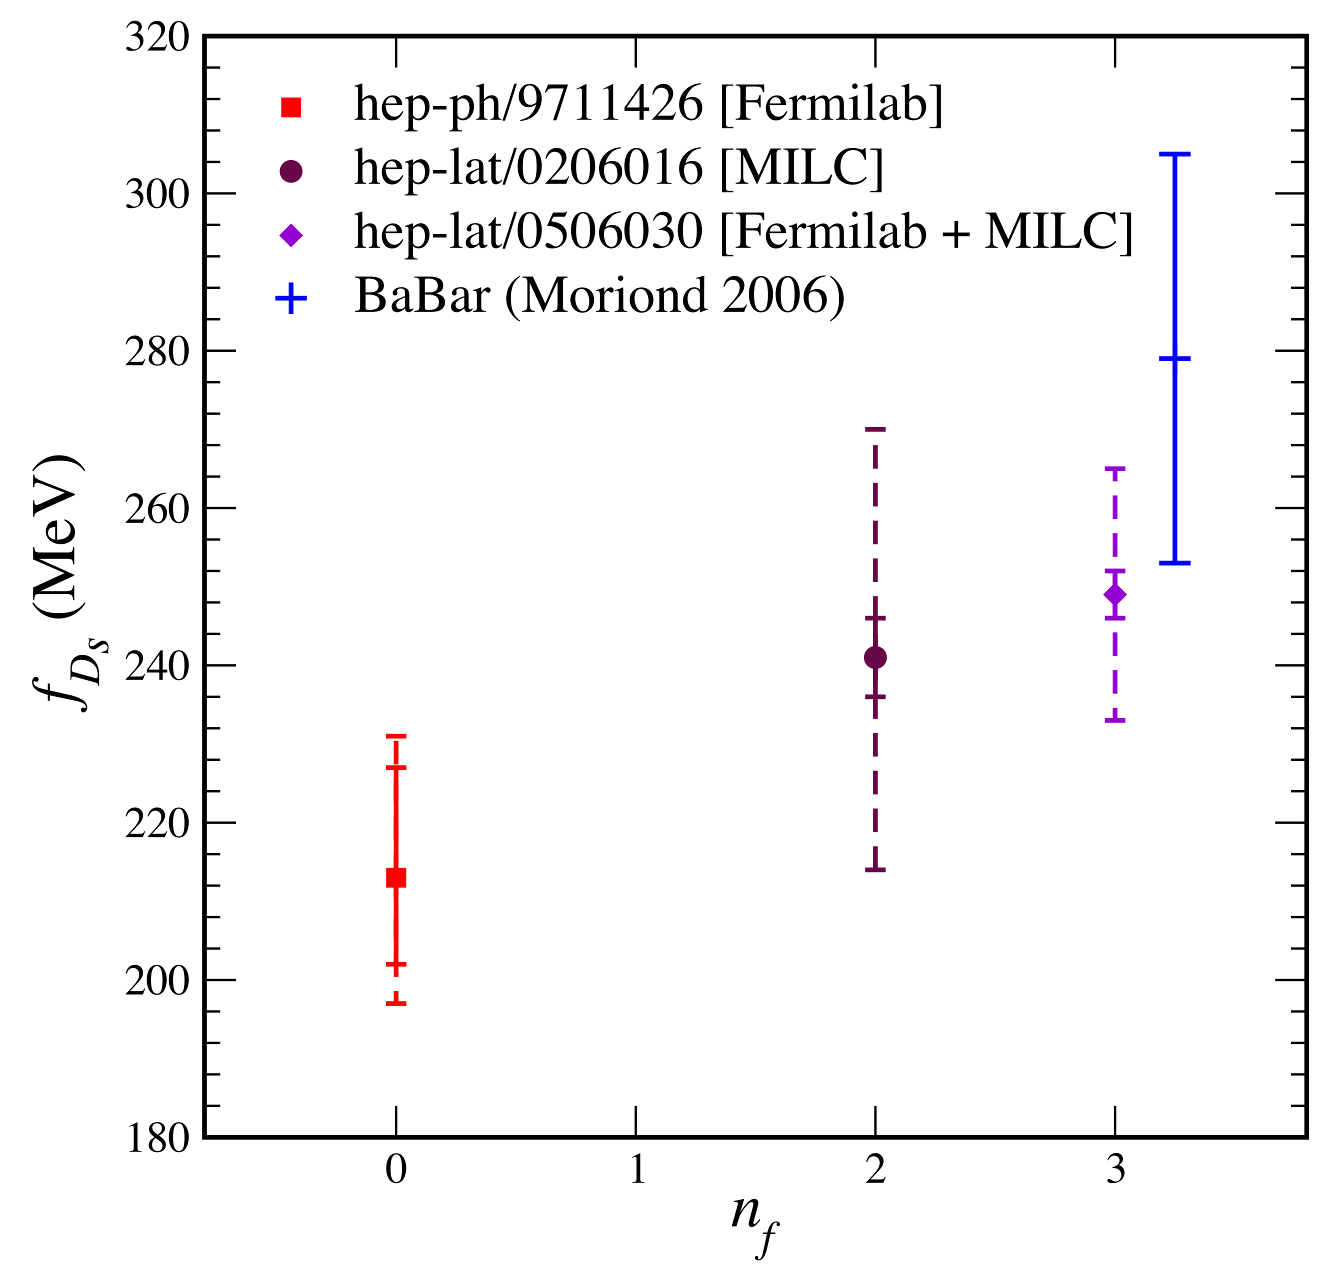

Hadron Spectrum 3

BMW Collaboration: *Science* **322** (2008) 1224

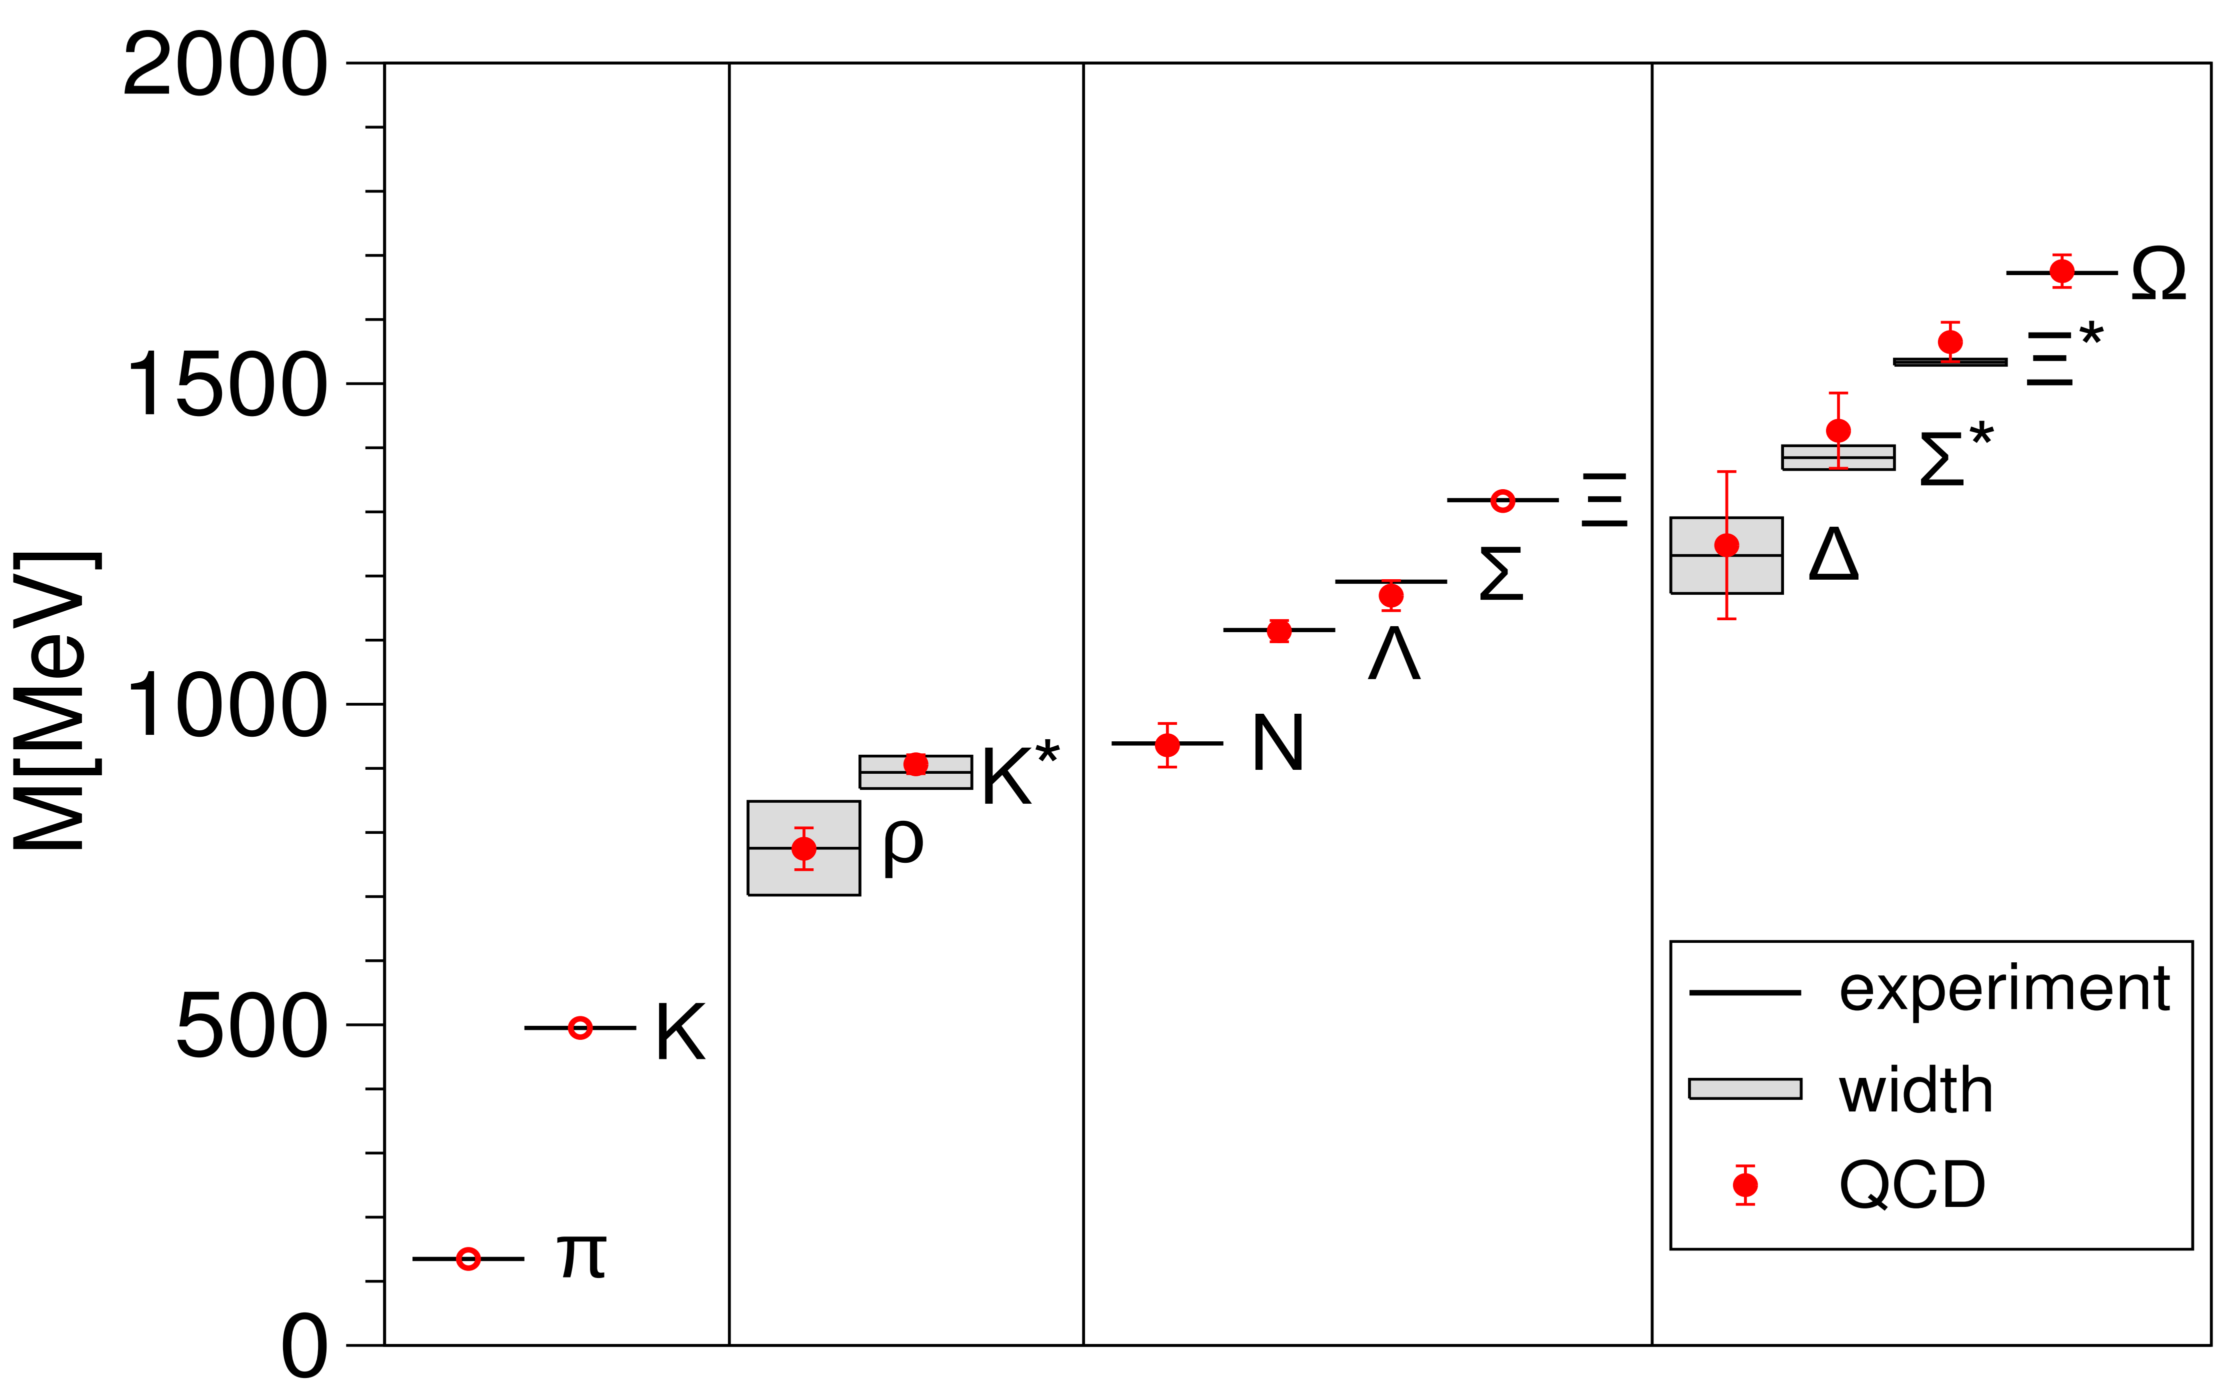

- $a = 0.125, 0.085, \& \ 0.065$ fm;
- tree O(a) Wilson;
- 6× stout smearing;
- $2m_l < m_q < 1.7m_s$;
- π, K, Ξ input.

QCD postdicts the low-lying hadron masses!

Lattice QCD for $g\text{–}2$

- With lattice QCD, one can compute FT$\langle V_\mu(x)V_\nu(0)\rangle$ or FT$\langle V_\mu(x)V_\nu(y)V_\rho(z)V_\sigma\rangle$ (from first principles) and convolute the result with QED Feynman diagrams.

- In addition to usual worries (continuum limit, physical pion cloud), need $q \sim m_\mu$, so might expect to need box-size a few times $\pi/m_\mu \sim 6$ fm.

- Structure in Green functions expected at two QCD scales: $m_\pi \approx 1.3m_\mu$ and $m_\mu \approx 7m_\varrho$; also need to match onto pQCD regime.

- The VP 2-pt function has 2 (1) form factors; the HL×L has 138 (43).

- In the end, need only two numbers, HVP (≈ 7000) to 0.2%, HL×L (≈ 100) to 5%, to match measurement of approved experiment Fermilab E989.

- Probably need cleverness, not just brute force.

Sea Quarks are Necessary for g–2

- Not just for processes sketched in the top figure (for both vacuum polarization and HL×L).
- All fermion lines/loops connected to initial or final state must be treated separately:
  - “disconnected diagrams”—
  - present because photon is flavor singlet;
  - really, really demanding.
- Any fully disconnected calculations?

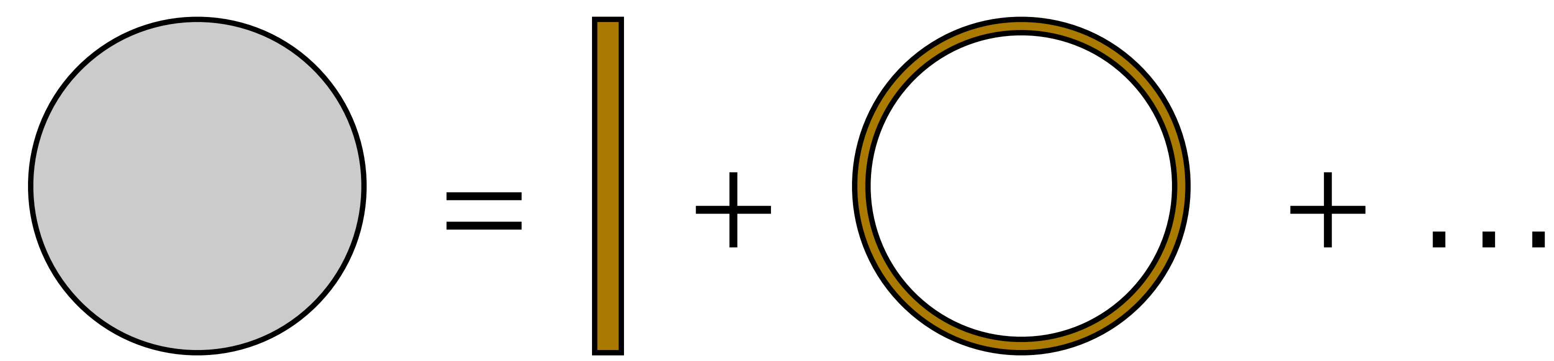

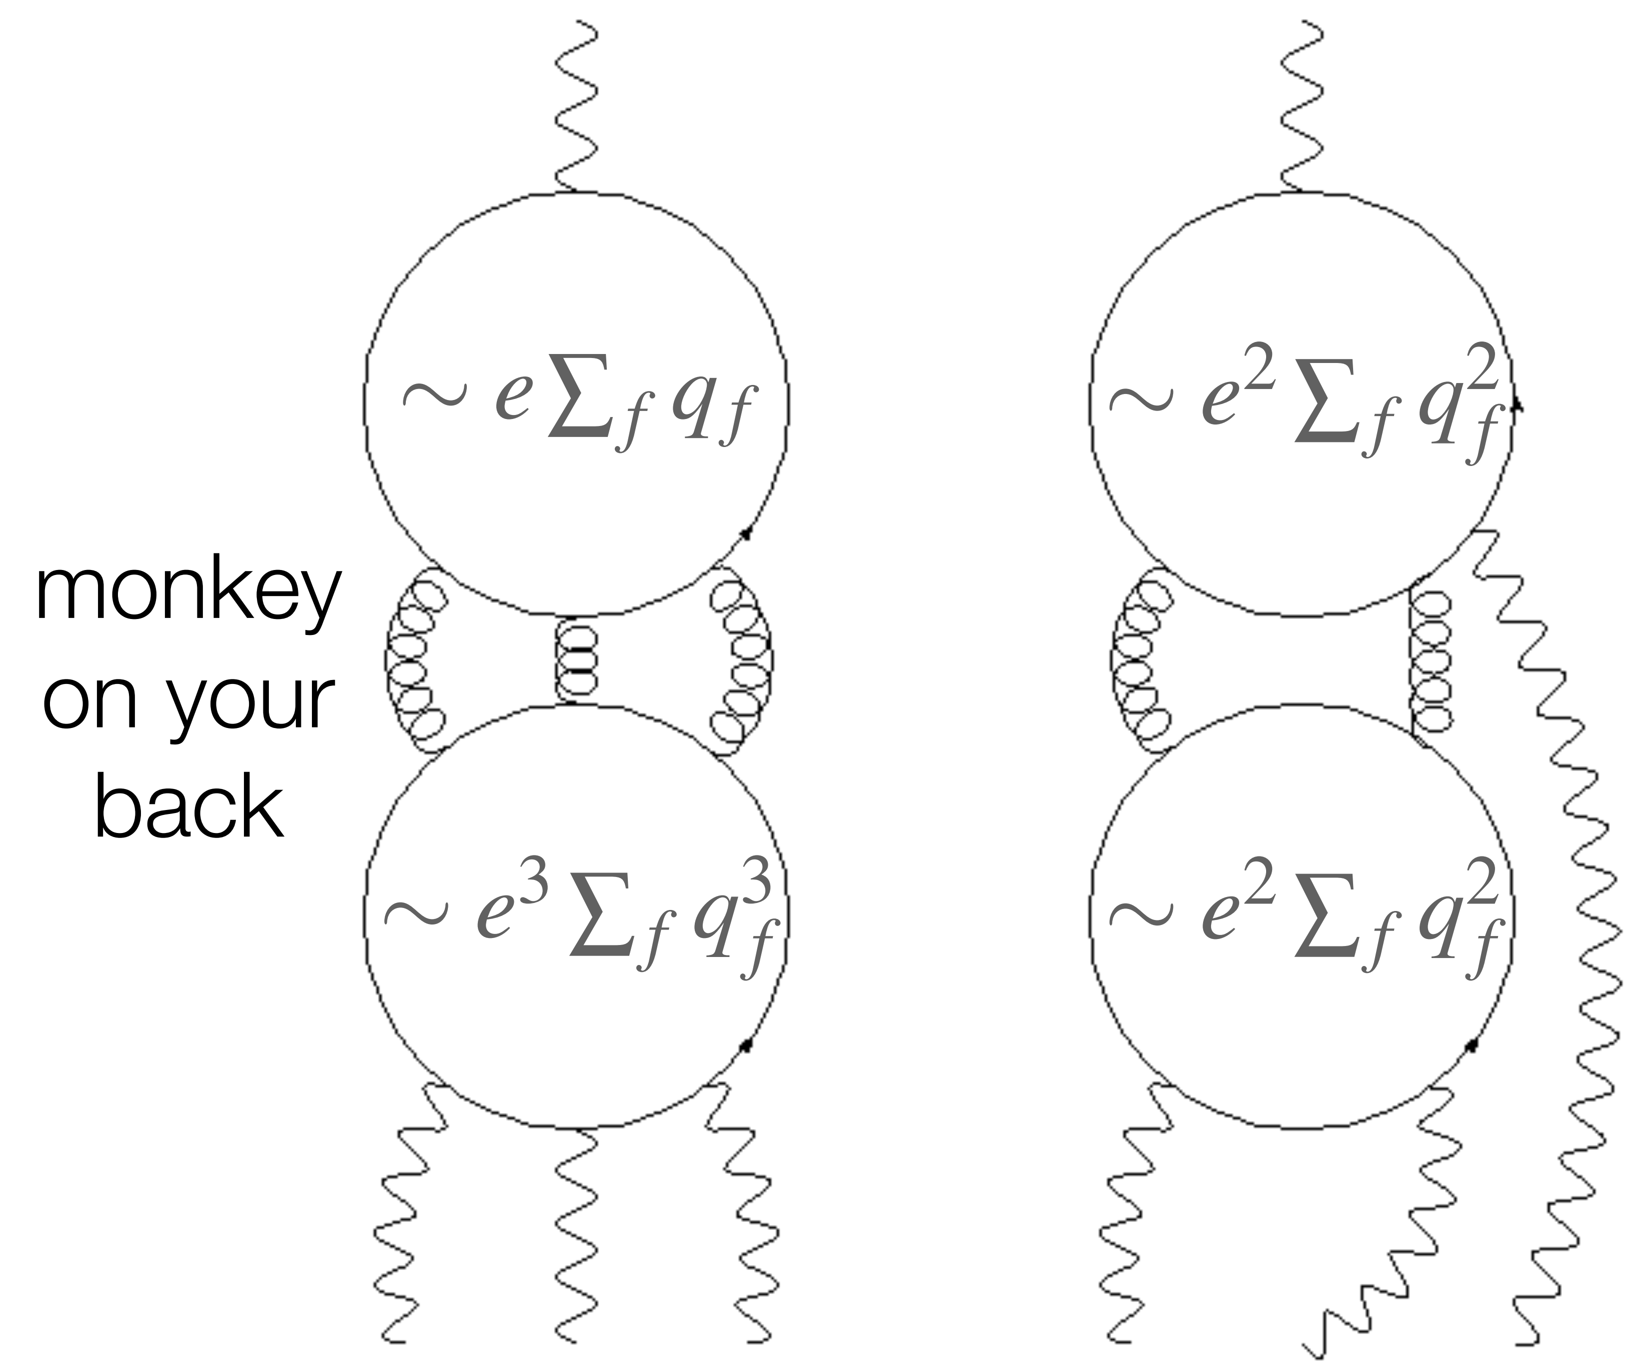

Error Budgets for Muon $(g-2)$

error ∝ perimeter; area ∝ weight in sum in quadrature

stats
syst

BNL E821 → FNAL E989

HL×L
HVP (lo)
HVP (ho)
EW
QED

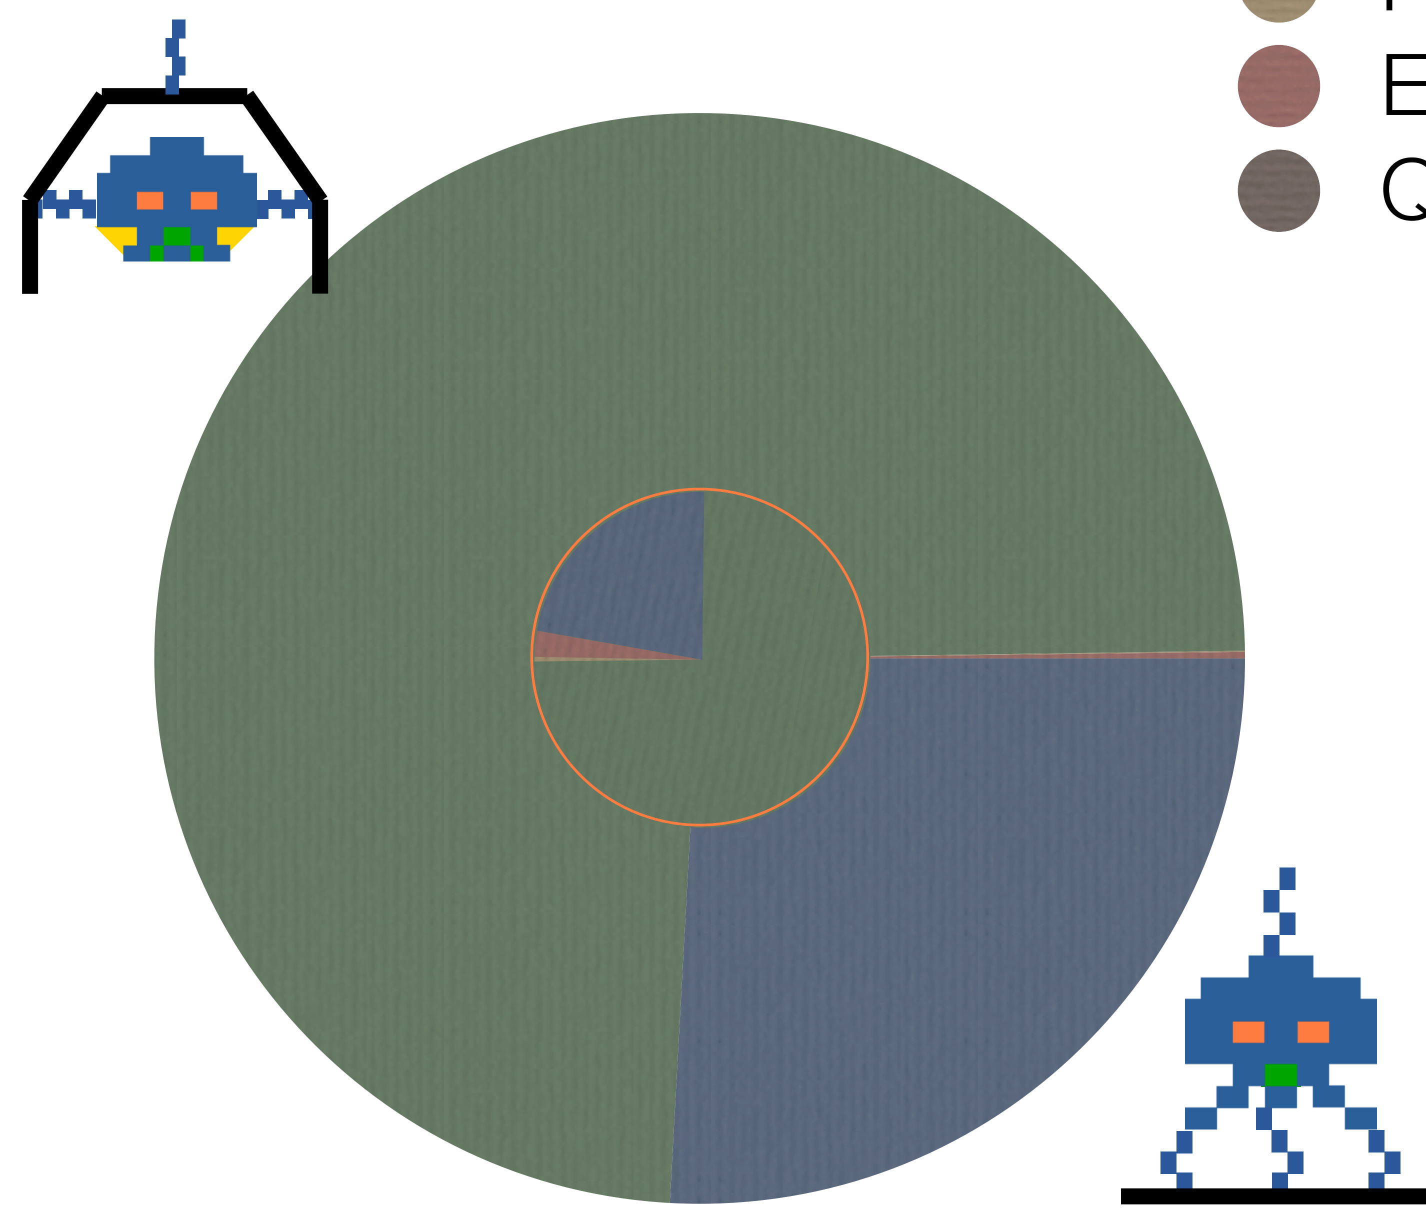

Standard Model Calculation

Hardware for Theorists

from the Terascale to the Exascale

What's in a Prefix?

- Teraflop, terabyte, *terascale*: 10^{12}
  - 10 teraflop/s, *etc.*
  - 100
- Petascale: 10^{15}
  - 10
  - 100
- Exascale (aka extreme scale): 10^{18}
- Zetta- and yotta- are next….

Timeline for Lattice QCD

- ~1998 delivered to lattice QCD/region
- ~2003
- ~2008
- ~2013
- ~2018
- ~2023
- ~2028 (*2021*)

peak horse-power of centralized computing facilities cross these thresholds 5–7 years earlier

Case Study: USQCD Collaboration

- Originated in 1999 with the formation of the "LQCD Executive Committee" by several wise men with the encouragement of the US DOE HEP, NP, and SciDAC offices:
  - SciDAC = Scientific Discovery through Advanced Computing.
- 2001 SciDAC grant supports software development and test clusters.
- 2003–04: ~ $6M for 5 Tflop s^{-1} QCDOC and small clusters (1 yr = 8000 hr ≈ 3×10^7 s).
- 2005-09: $9.2M for clusters at JLab and Fermilab—up to 23 Tflop s^{-1} yr—plus operations;
  - also INCITE grants at Argonne and Oak Ridge—up to 15 Tflop s^{-1} yr.
- 2010-14: $18.15M + $4.96M (ARRA) for clusters—up to ~500 Tflop s^{-1} yr;
  - plus access to Pflop s^{-1} computers at Argonne, Oak Ridge, NCSA.

Price-Performance (USQCD Experience)

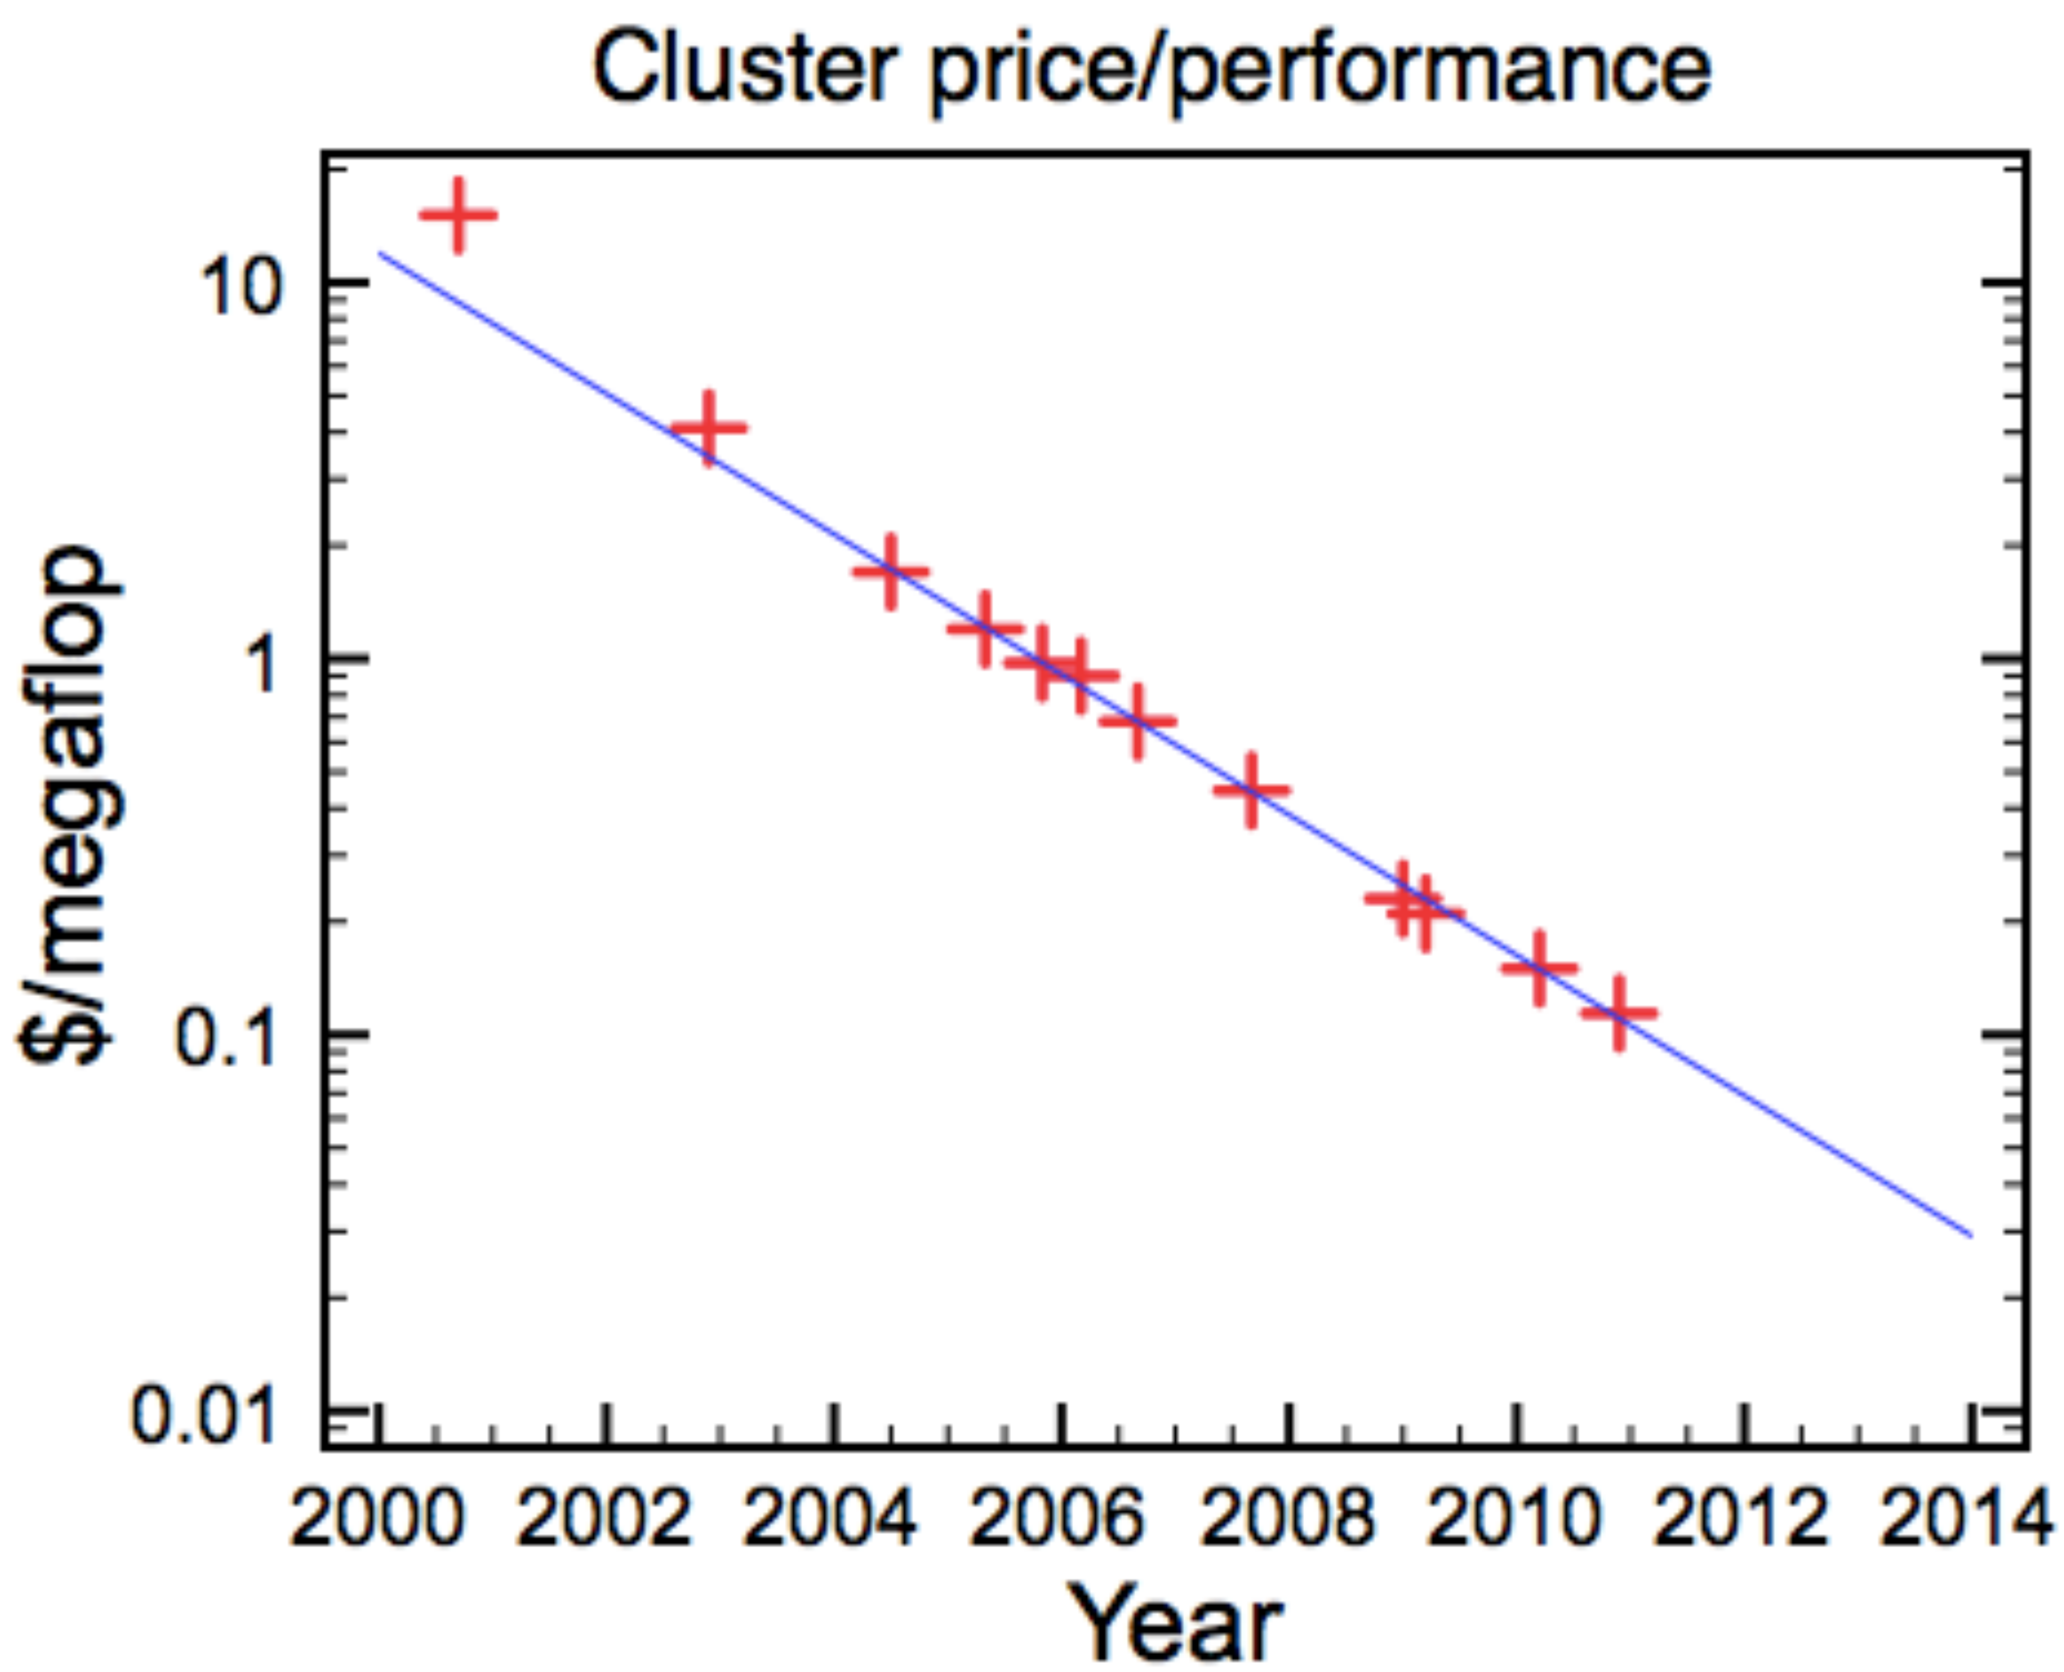

- In the past decade, cost of real-life lattice QCD computing has dropped by two orders of magnitude.
- Data are measured sustained flop/s on clusters at Fermilab and JLab.
- Fit yields \$1/(Mflop/s)$\times 2^{-t/19}$, where t = months since September 2005.
- While past performance does not guarantee future results, the trend survived a transition from faster clock-speeds to multiple-socket, multi-core CPUs @ ~2 Ghz.

Design Considerations

e.g., D.J. Holmgren, hep-lat/0410049

- Good design requires balance among:
  - floating point performance;
  - memory bandwidth: width and speed of the front-side bus;
  - min(I/O, network) bandwidth of the nodes;
  - bandwidth and *latency* of the network fabric.

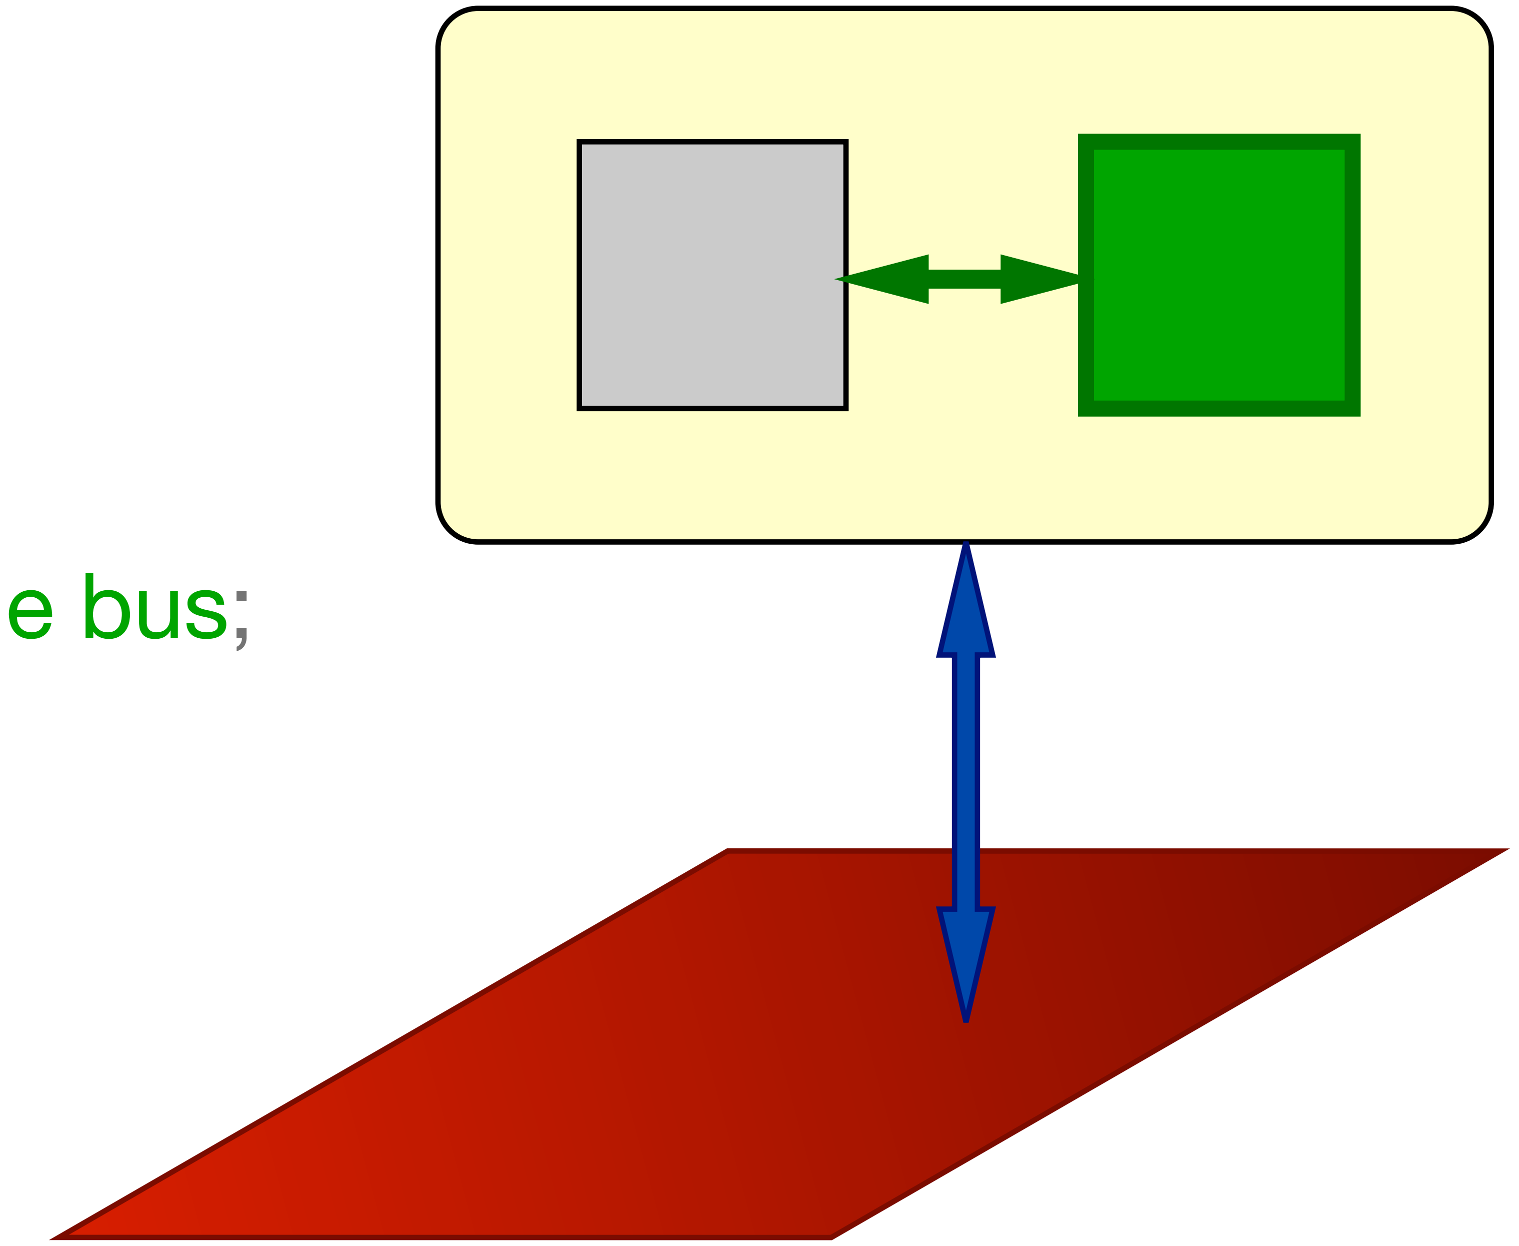

- Any of these can limit the capacity or, especially, capability of a supercomputer.
- Latency is an issue for lattice QCD, because the basic datum, an SU(3) matrix, is small.

Design Considerations

e.g., D.J. Holmgren, hep-lat/0410049

- Good design requires balance among:
  - multi-socket
  - floating point performance;
  - memory bandwidth: width and speed of the front-side bus;
  - min(I/O, network) bandwidth of the nodes;
  - bandwidth and *latency* of the network fabric.

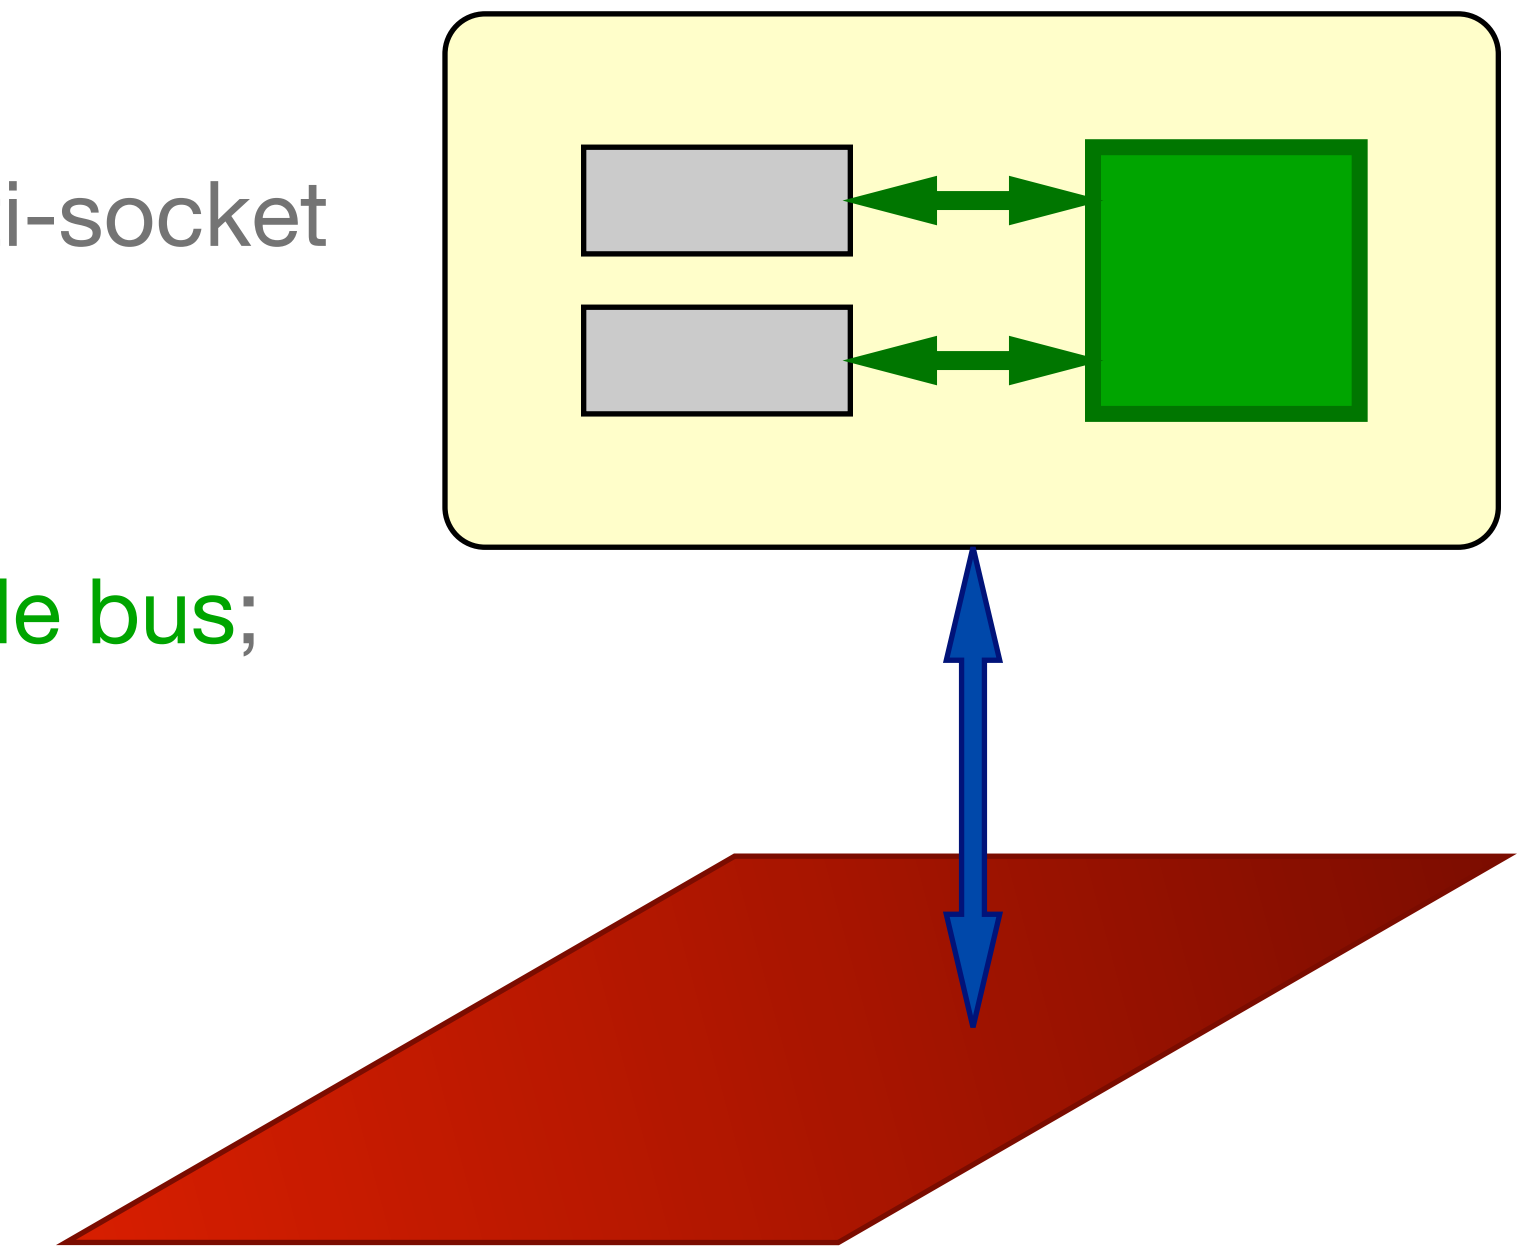

- Any of these can limit the capacity or, especially, capability of a supercomputer.
- Latency is an issue for lattice QCD, because the basic datum, an SU(3) matrix, is small.

Design Considerations

e.g., D.J. Holmgren, hep-lat/0410049

- Good design requires balance among:
  - floating point performance;
    - multi-socket
    - multi-core
  - memory bandwidth: width and speed of the front-side bus;
  - min(I/O, network) bandwidth of the nodes;
  - bandwidth and *latency* of the network fabric.

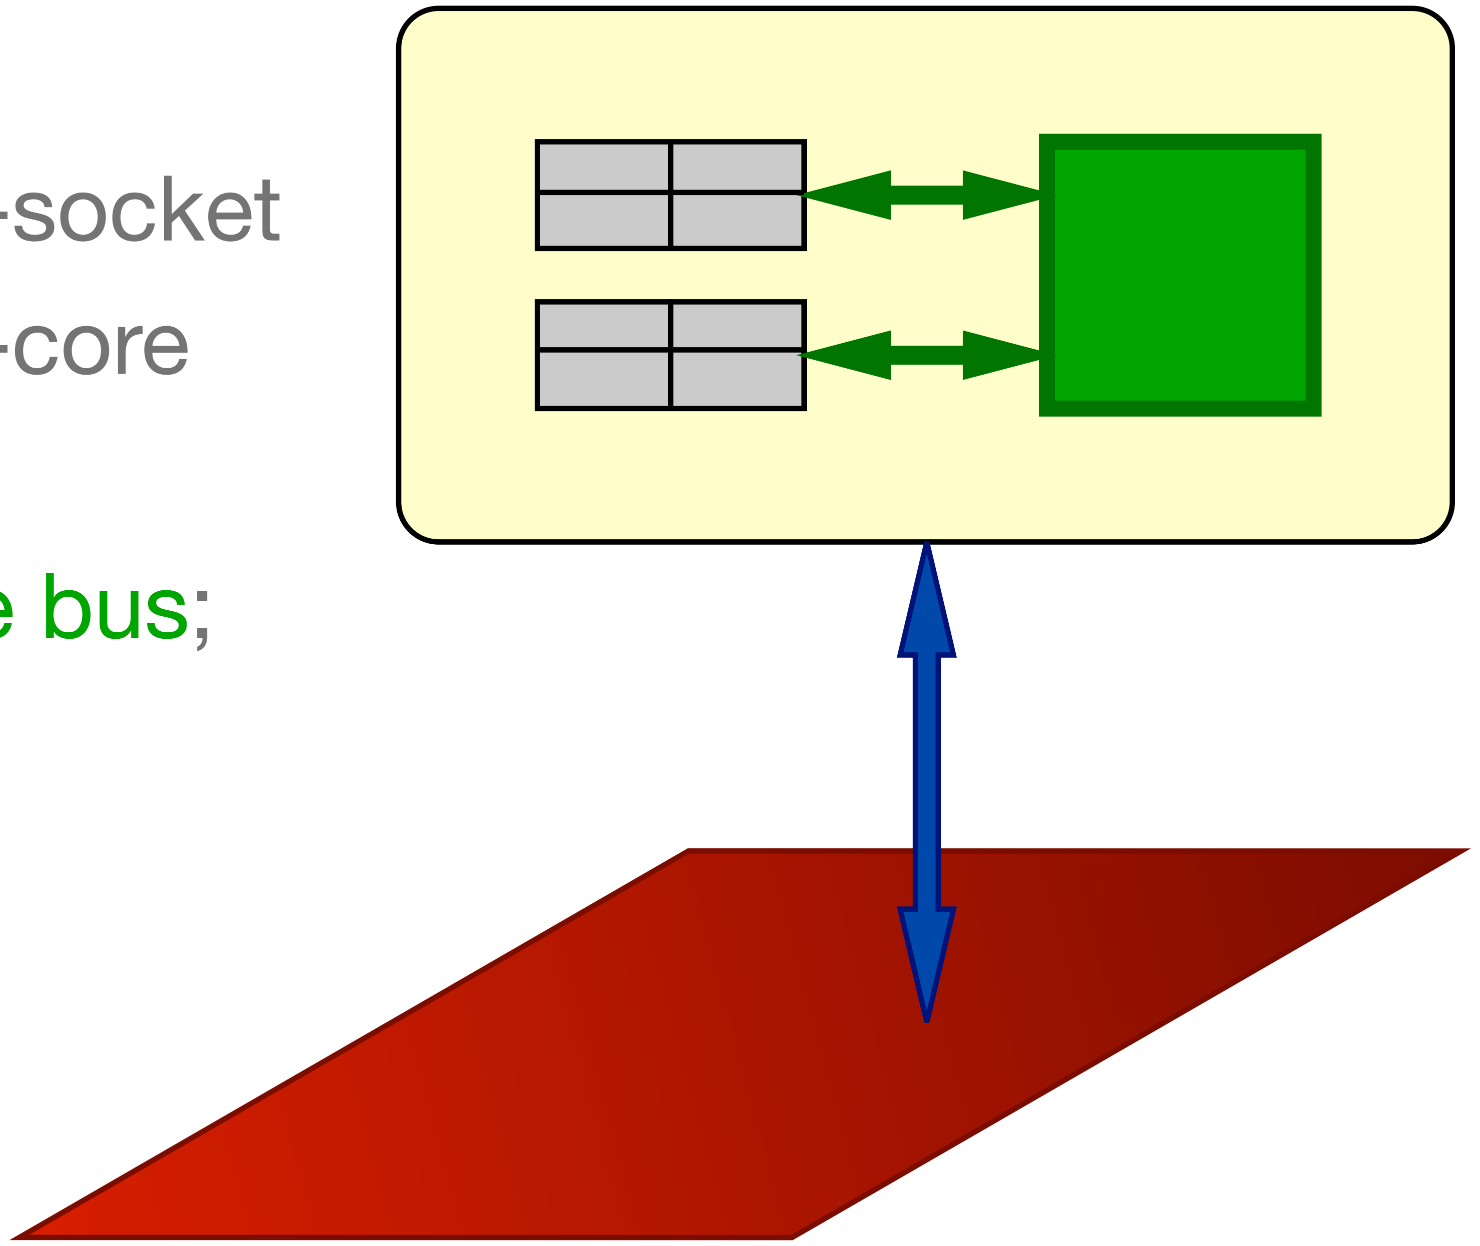

- Any of these can limit the capacity or, especially, capability of a supercomputer.
- Latency is an issue for lattice QCD, because the basic datum, an SU(3) matrix, is small.


7n: 396 nodes
dual-socket
quad-core
dble Infiniband


kaon: 600 nodes
dual-socket
dual-core
dble Infiniband


10q: 224 nodes
dual-socket
quad-core
dble Infiniband

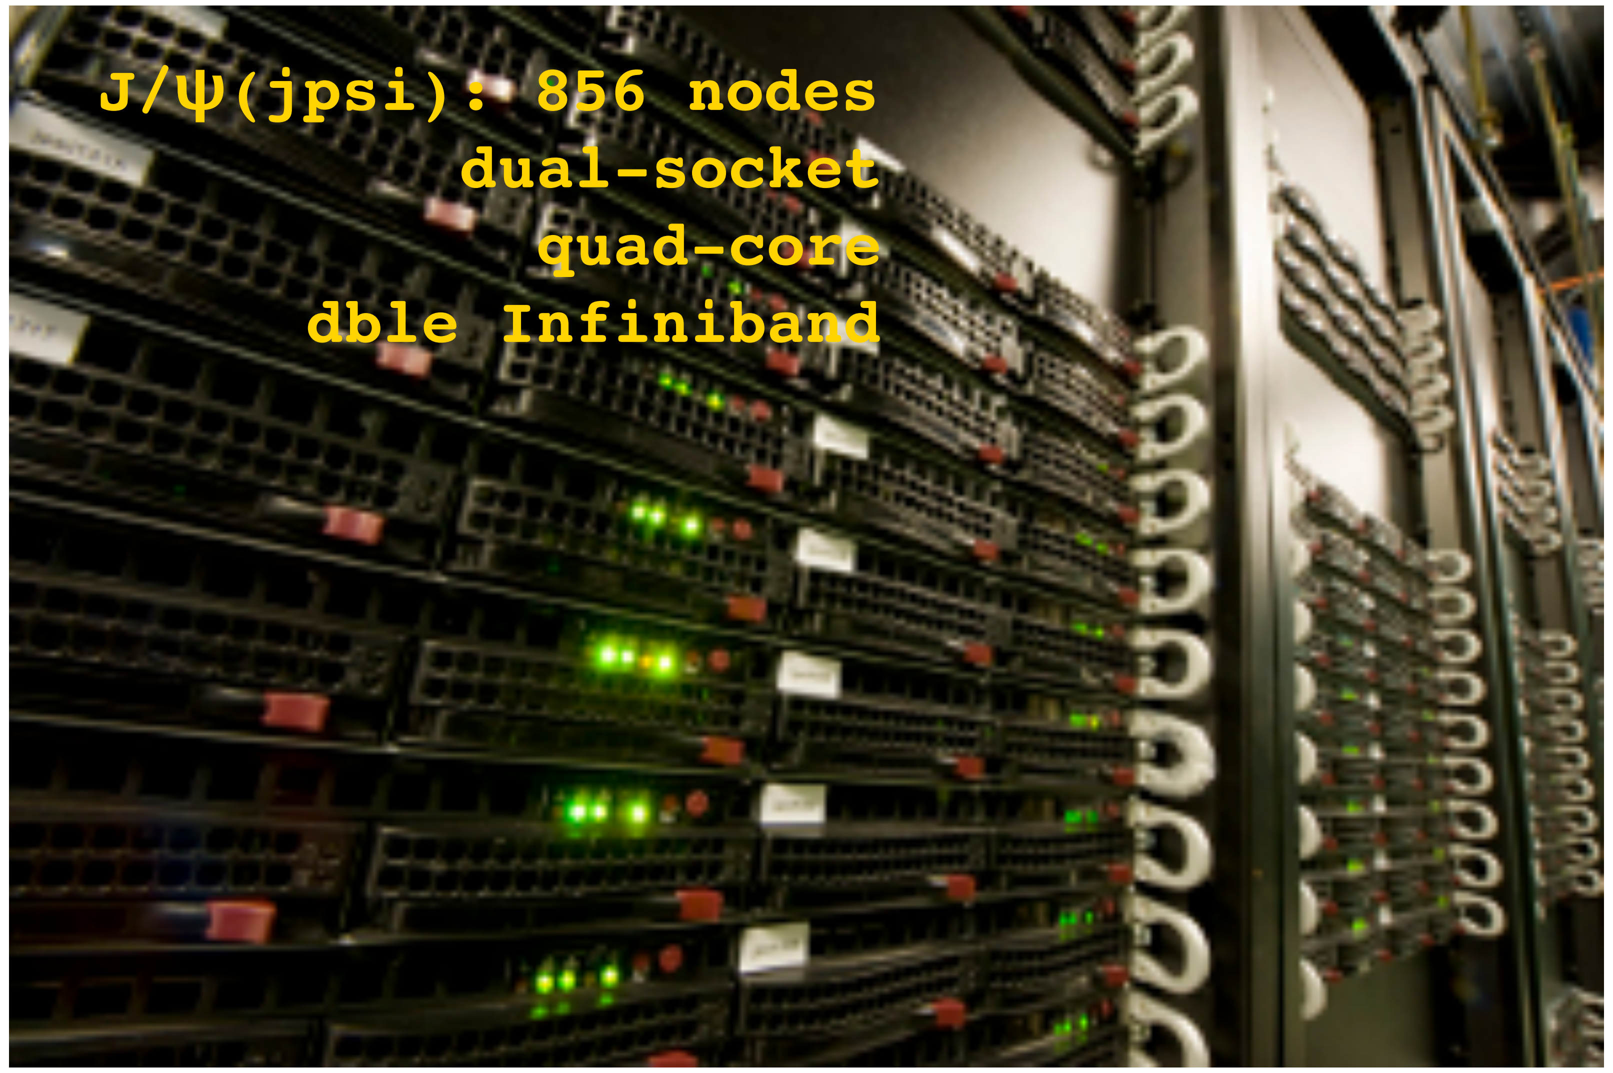
J/ψ(jpsi): 856 nodes
dual-socket
quad-core
dble Infiniband

Technological Convergence

- Lattice QCD has thrived on home-designed and -built computers (PACS-CS, CP-PACS, APE, APE-next, QCDOC, QCDSP, ACPMAPS, GF-11, …), commercial supercomputers at labs (esp. KEK), and general-purpose facilities (NCSA, NERSC, …, Jülich, Kobe, ...).
- Design features are converging:
  - multi-CPU, multi-core motherboards with Linux on every node;
  - fast, low-latency interconnection, which can cost as much as CPUs.
- Commercial markets in video games, search (*e.g.*, Google), industrial design drive cost.
- New technologies: graphics cards (aka GPU) & multi-threading (not 8 cores but 128–1024).
- Infectious ideas: cosmology clusters.

GPUs: Disruptive Technology?

- GPU has better price/performance via 100s of cores & some special-function support, but—
  - poor (for latQCD) memory bandwith;
  - harder to program, but tools and skills improving.
- Several code bases within USQCD software: QUDA, ….
- Job speed up, ÷6–15: job is the right unit, since few-to-several GPUs hang off 1 CPU.
- Dirac inverter sped up even more: opens the door for new jobs arrangements?
  - Will all-to-all propagators become more cost effective and greatly help g–2?


10g:
18+32 nodes
dual-socket
quad-core
dble Infiniband
32 have 4 1-Tflop/s GPUs

Ensembles and Software

Asqtad Data Mine

A. Bazavov et al., arXiv:0903.3598, and refs therein.

MILC + USQCD

(2002–2009)

$a = 0.18$ fm not shown
$a = 0.15$ fm

am_l / am_s	$m_\pi L$	Lattice	# Lats
$a = 0.12$ fm			
0.40/0.40	29.4	$20^3 \times 64$	332
0.20/0.20	19.6	$20^3 \times 64$	341
0.10/0.10	13.7	$20^3 \times 64$	339
0.05/0.05	9.7	$20^3 \times 64$	425
0.04/0.05	8.7	$20^3 \times 64$	351
0.03 / 0.05	7.6	$20^3 \times 64$	564
0.02 / 0.05	6.2	$20^3 \times 64$	1758E
0.01 / 0.05	4.5	$20^3 \times 64$	2023E
0.01 / 0.05	6.3	$28^3 \times 64$	241
0.007 / 0.05	3.8	$20^3 \times 64$	1852E
0.005 / 0.05	3.8	$24^3 \times 64$	1802E
0.03 / 0.03	7.6	$20^3 \times 64$	359
0.01 / 0.03	4.5	$20^3 \times 64$	346

am_l / am_s	$m_\pi L$	Lattice	# Lats
$a \approx 0.09$ fm			
0.0124 / 0.031	5.78	$28^3 \times 96$	1996C
0.0093 / 0.031	5.04	$28^3 \times 96$	1138C
0.0062 / 0.031	4.14	$28^3 \times 96$	1946C
0.00465 / 0.031	4.11	$32^3 \times 96$	540C
0.0031 / 0.031	4.21	$40^3 \times 96$	1012C
0.00155 / 0.031	4.80	$64^3 \times 96$	700R
0.0062 / 0.0186	4.09	$28^3 \times 96$	985C
0.0031 / 0.0186	4.22	$40^3 \times 96$	642N
0.0031 / 0.0031	4.20	$40^3 \times 96$	440R
$a \approx 0.06$ fm			
0.0072 / 0.018	6.33	$48^3 \times 144$	625
0.0054 / 0.018	5.48	$48^3 \times 144$	617C
0.0036 / 0.018	4.49	$48^3 \times 144$	771
0.0025 / 0.018	4.39	$56^3 \times 144$	800N
0.0018 / 0.018	4.27	$64^3 \times 144$	826C
0.0036 / 0.0108	5.96	$64^3 \times 144$	483N
$a \approx 0.045$ fm			
0.0028 / 0.0140	4.56	$64^3 \times 192$	861N

DWF Data Mine

P. Boyle [RBC+UKQCD], arXiv:0710.5880.

RBC+UKQCD using resources of
RBRC, UKQCD, USQCD

	a (fm)	volume	am_l/am_s	$m_\pi L$	lgfs
2006	0.114	24^3x64	0.03/0.04	9.0	~200
			0.02/0.04	7.6	~200
2007			0.01/0.04	5.7	~800
			0.005/0.04	4.5	~800
2008	0.081	32^3x64	0.008/0.03	5.5	~600
			0.006/0.03	4.8	~900
2009			0.004/0.03	4.0	~800

USQCD SciDAC-2 API for QCD

R. Brower & the USQCD Software Committee

SciDAC 1

SciDAC 2

Application Software
Chroma, CPS, MILC

- Level 4
  - **QCD Physics Toolbox** — shared code, visualization, performance tools
  - **Workflow** — & Data Analysis Tools
- Level 3
  - **QOP** — Dirac inverters, gauge force: **op**timized in assembler
  - **User Environment** — runtime, accounting, grid
- Level 2
  - **QDP** — (QCD **d**ata **p**arallel)
  - **QIO** — binaries || XML, ILDG
- Level 1
  - **QLA** — (QCD **l**inear **a**lgebra)
  - **QMP** — (QCD **m**essage **p**assing)
  - **QMC** — (QCD **m**ulti-**c**ore interface)

- With this framework, junior—and even senior—researchers have taken up lattice QCD.

- So could you: muon g–2 seems like a fruitful ground for collaboration.

Forecasts: Needs *vs.* Resources

Needs for g–2

- Let's assume that the monkey-on-your-back topology can be safely neglected (likely).
- Let's assume that the HVP to needed precision comes along with HL×L (not obvious).
- Let's focus on QCD+QED: easier to forecast one number than many form factors.
- BCHIYY find 100% error using 10^{-2} Tflop s^{-1} yr, and planning "reasonable" calculation with 10 Tflop s^{-1} yr. Target 10% (5%) needs—naïvely—a factor of 100 (400) more computing:
    - 1–5 Tflop s^{-1} yr needed.
- *Caveats*: with 100% error it is hard to foresee obstacles both surmountable and unsurmountable. Estimate is, thus, more likely to be over-pessimistic or over-optimistic than accurate.

Resources for g–2

- "Luminosity" formula: resource = f_{g-2} × budget × Moore's Law; f_{g-2} = fraction for g–2:
  - USQCD Moore's Law: $2^{t/1.6}$ Tflop s^{-1} (\$M)$^{-1}$; (now t = years since 2005/09)
  - USQCD budget experience: 2.9×$2^{t/10.5}$ \$M yr^{-1}; (omits Tea Party effects)
  - TB *et al.* are increasing f_{g-2} from 10^{-4} to 10^{-2}.
- Predict resource of 5 Tflop s^{-1} yr in 2016.
- Coincides with forecast of need.

Other g–2 Forecasts

- Hashimoto: all-to-all propagators are useful for π-γ*-γ* form factor and can be used to build up whole HL×L. Numerous form factors more daunting than the lattice calculation itself.

- Jansen: working on π-γ*-γ* now. Full HL×L is difficult to estimate (subject is new); a first guess for a project with fixed external-momentum sources is 10^7 core-hr (≈ 10 Pflop s^{-1} yr).

- Schierholz/Rakow (QCSF): 5% calculation of HL×L underway, with HVP a by-product; method sketched at Lattice 2008.

- Wittig: 10–20% accuracy for HL×L will require exascale resources, so skeptical that it will be done by 2016. Even HVP requires large volumes and near-physical quark masses.

Computing Outlook

- Davies: STFC providing a grant for clusters (Cambridge) and an IBM BlueGene/Q (Edinburgh) to be shared among lQCD (via UKQCD), astrophysics, ….
- Hashimoto: KEK receiving IBM BlueGene/Q 2011-12 with peak 1.2 Pflop s^{-1}, mostly lQCD. Kobe flagship will be 10 Pflop s^{-1}, with perhaps 10% for lattice QCD. Sustain of peak.
- Jansen: Jülich has a BlueGene/P at 1.2 Pflop s^{-1} (peak) and LRZ München will soon have 110,000-core Xeon-based cluster at 3 Pflop s^{-1} (peak). LQCD receives 20-25% of total; code sustains 20-25% of peak.
- Wittig: Supercomputer centers at Jülich, Stuttgart, München will try to keep pace. Other labs (e.g., GSI) and universities will have medium-sized computers dedicated to lQCD.

Observations

What Can Lattice QCD Do for You?

- Lattice QCD calculations of hadron masses, matrix elements, α_s, and quark masses are mature (arguable not tenured).
- Lattice QCD calculations of g–2 are no longer in their infancy, but still in the toddler stage.
- Several different ideas are being investigated (for both HVP and HL×L):
  - brute force; GPU acceleration of propagators;
  - momentum insertions; all-to-all propagators;
  - QED+QCD.
- Computing is well-supported: success breeds success.

What Can You Do for Lattice QCD?

- You know the structure of the four-point function better than we do.
- What is really needed to improve the accuracy of the HL×L amplitude?
  - What can be said model independently?
  - Better parametrizations? cf., form factors for semileptonic decays.
- Are there combinations of effective field theories, models, and targeted lattice-QCD calculations that provide a useful improvement, for the time being?

Thanks

- Don Holmgren
- Paul Mackenzie
- George Fleming
- Christine Davies
- Gerrit Schierholz
- Organizers
- Audience
- Shoji Hashimoto
- Karl Jansen
- Hartmut Wittig
- Paul Rakow
- Tom Blum

Thanks

- Don Holmgren
- Paul Mackenzie
- George Fleming
- Christine Davies
- Gerrit Schierholz
- Organizers
- Audience
- Shoji Hashimoto
- Karl Jansen
- Hartmut Wittig
- Paul Rakow
- Tom Blum