Quantum Mechanical Aspects of Transit-Time Optical Stochastic Cooling

A. Charman* and J. S. Wurtele*

and thanks to M. S. Zolotorev, A. A. Zholents◇, and S. Heifets† for many useful discussions

*U.C. Berkeley Dept. of Physics
◇Argonne National Laboratory
†Stanford Linear Accelerator Center

Faster Stochastic Cooling

- To increase the cooling rate, and typically simultaneously decrease the equilibrium emittance achievable requires that incoherent effects be made small, necessitating:
  - lower particle densities
  - suppressed mixing between pickup and kicker
  - good mixing between cooling passes
  - amplifiers with
    - high gain
    - high total power
    - low noise
    - large bandwidth
    - high repetition rate
    - non-regenerative
- Existing stochastic cooling schemes based on microwave or RF technology are limited by the O(GHz) bandwidths available for high-gain amplifiers
  - typical RF stochastic cooling time-scales are minutes to hours or more
- In the case of muons, finite particle lifetimes require that any final cooling to boost luminosity be performed in at most a few lab-frame decay times, O(10 of microsecond) for
  - Relativistic lifetime O(1 millisecond) for muons at O(100 GeV)
- Stochastic cooling on the microsecond time-scale will require
  - reversible, adiabatic beam compression and stretching to reduce beam density during cooling
  - utilizing cooling signals at optical wavelengths, where solid state and/or parametric lasers are available, which can provide high gain over O(THz) bandwidths centered around O(1 micron) wavelengths

Challenges for Optical Stochastic Cooling (OSC)

- at such extremely fast frequencies, the pick-up signal cannot be manipulated electronically, but must be collected, manipulated, amplified, and directed to kicker through all-optical means
- to reduce longitudinal emittance, transverse optical fields must be made to effect longitudinal momentum kicks
- fields must be coherently amplified to high power levels
- between pickup and kicker, particle positions must be controlled to within less than one optical wavelength

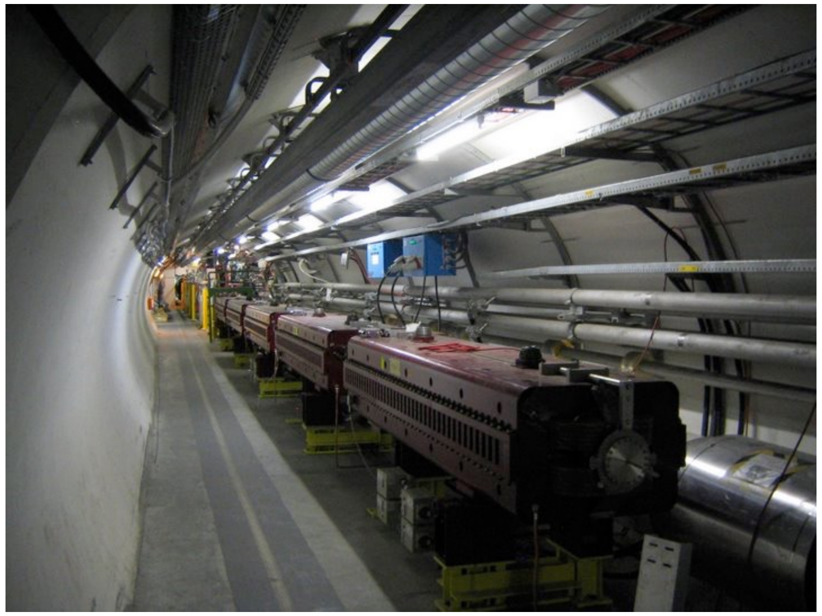

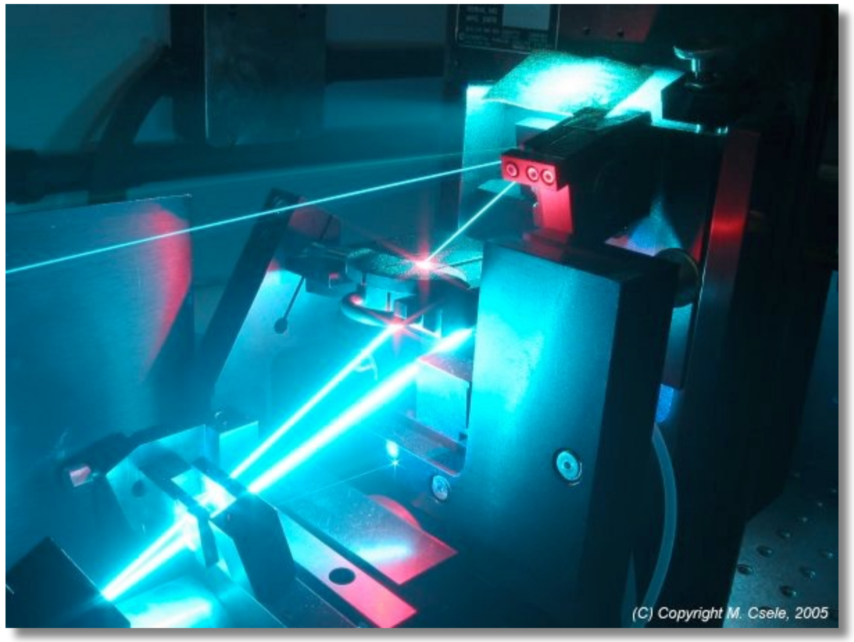

Transit-Time Optical Stochastic Cooling (OSC)

- Zolotorev, et al., have proposed a method of transit-time optical cooling, in which both the pick-up and kicker consist of large-field magnetic wigglers:
  - in the pickup wiggler particle quiver leads to spontaneous wiggler radiation
  - radiation is collected and greatly amplified in a low-noise, high bandwidth optical amplifier
  - While the light is being amplified, particles are directed into a bypass lattice, whose beam optics produce a time-of-flight delay proportional to the longitudinal and/or transverse deviation from a desired reference orbit
  - in the kicker wiggler, quivering particles interact resonantly with the electric field of the amplified radiation, exchanging energy depending on the relative phase
  - If time-of-flight delay is proportional in part to longitudinal momentum deviation, proper phasing can lead on average to a restoring force reducing momentum spread
  - If time-of-flight is also proportional in part to a transverse betatron coordinate, then dispersion beamline at kicker can also lead to cooling of transverse phase space

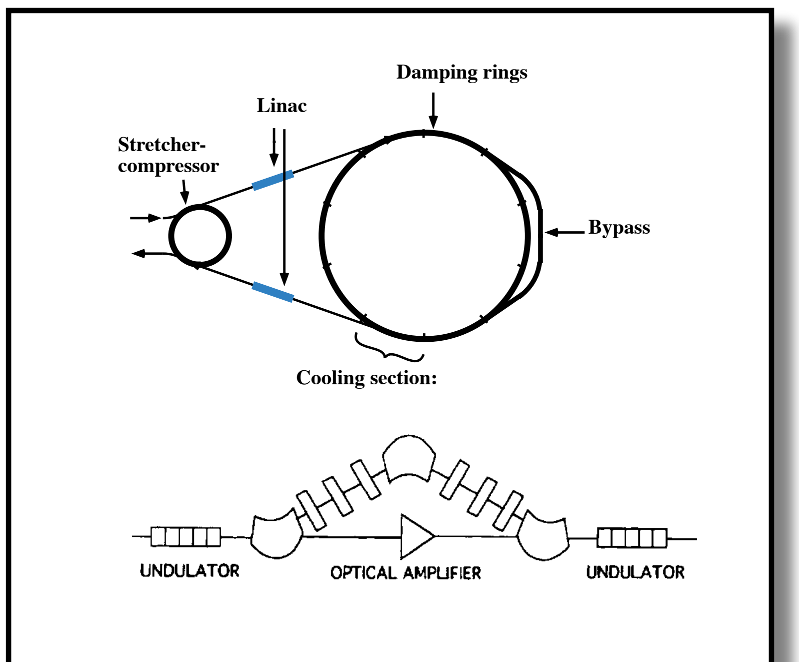

- before cooling, beam is energetically chirped in LINAC and passed through highly-dispersive ring to **reversibly** increase bunch length and decrease momentum spread to manageable levels; beam is re-compressed following cooling
- bypass and kicker beam optics must be carefully engineered, adjustable, and controllable through active feedback
- optical amplifier(s) must be stable, high-power, highly linear, variable-gain, low-noise (therefore actively cooled), and non-regenerative
- mixing (effectively random shifting of longitudinal particle positions on scale of radiation wavelength) must be thorough between kicker and next pass though pickup (easy), but negligible between pickup and kicker (difficult)

Spontaneous Wiggler Radiation

- Central wavelength of the wiggler radiation is downshifted from the undulator period by relativistic effects – for a planar wiggler:

$$\lambda_0 \approx \frac{\lambda_u}{2\gamma_0^2}\left(1+\frac{a_u^2}{2}\right)$$

- homogeneous "coherent" bandwidth of wiggler radiation scales inversely with number of wiggles:

$$\Delta\omega \approx \frac{1}{2N_u}\omega_0$$

- coherent contribution corresponds to a nearly diffraction-limited beam with angular spread $\Delta\theta \approx \frac{1}{\sqrt{N_u}\gamma_0}$ and spot size such that $\Delta r \Delta\theta \sim \frac{\lambda_0}{4\pi}$

- Assuming light fields in kicker are coherent over transverse extent of beam, and amplifier bandwidth is matched to coherent bandwidth of spontaneous radiation, the effective sample size will scale like $N_s \sim \rho_b \sigma_{b\perp}^2 L_s$, where $L_s = N_u\lambda_0$ is the sample length (= coherence length)

- average number of photons emitted per particle into coherent bandwidth is approximately

$$\alpha = \frac{e^2}{\hbar c} \approx \frac{1}{137}$$

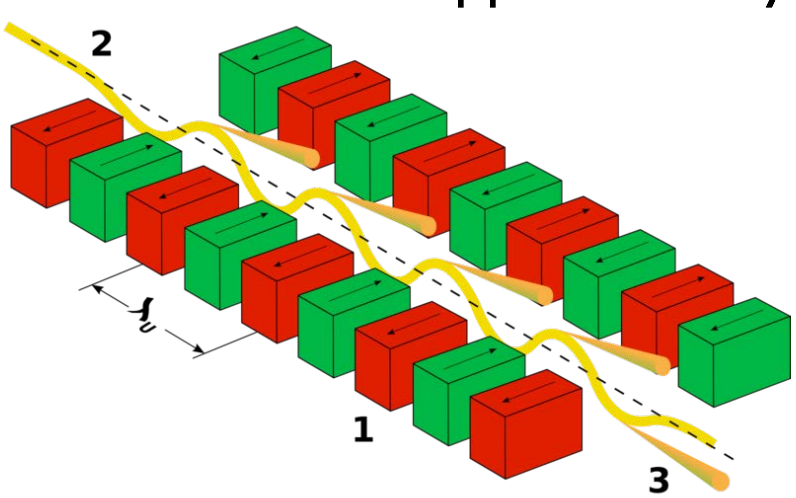

Quantum Cooling Catastrophe?

In any one pass through the pickup, each particle only radiates $O(10^{-2})$ photons, so that cooling signal from each particle is very weak and presumably subject to
quantum mechanical fluctuations

Naive quantum mechanical considerations raise fundamental doubts about whether individual particles radiate too weakly to be cooled

or whether pick-up signal even contains the phase information needed for transit-time cooling, and whether this information can be reliably amplified and extracted:

- Do individual particles radiate in random "quantum jumps", on average once in every $O(\alpha^{-1})$ passes through pickup?
  - in naive picture (akin to M. Sand's treatment of synchrotron radiation damping), where particles emit photons in discrete Poissonian quantum jumps, it would seem particles rarely experience a large kick from cooling self-radiation but almost always feel large heating cross-radiation, leading to drastically slower cooling, over-correction or unstable feedback
- Even when individual particles radiate, is the phase even well-defined?
  - Heisenberg Uncertainty principle says: $\Delta\phi\Delta N \geq \frac{1}{2}$
  - Since photon number is small, phase is poorly determined: $\Delta\phi \sim \frac{1}{\Delta N} \sim \frac{1}{\sqrt{N_{\text{ph}}}} \gg 1$
- What about spontaneous emission or thermal noise in laser amplifiers?
  - amplifier noise expected to add additional, randomly-phased photons,
  - but this noise acts more or less like extra particles in sample
  - additive versus multiplicative

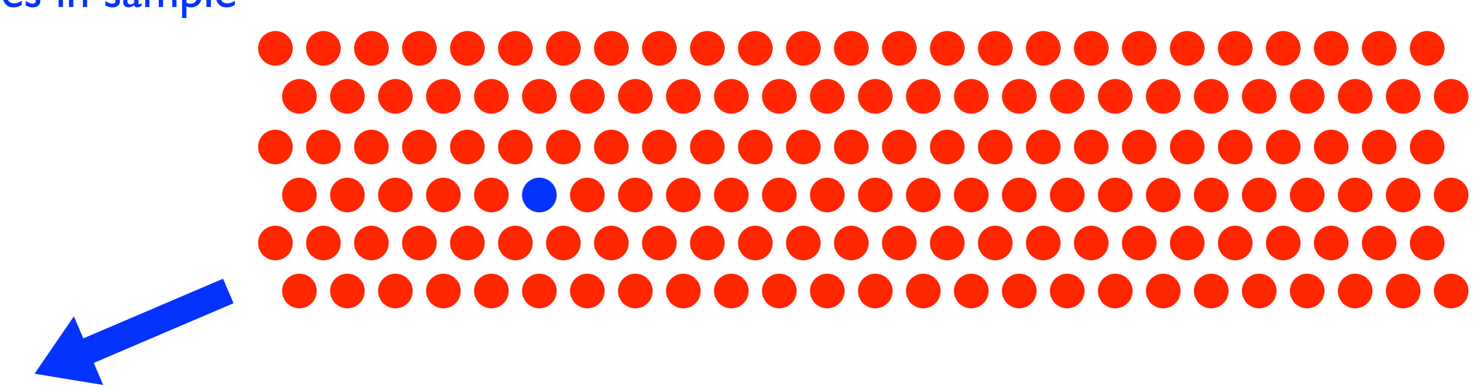

Naive QM Calculations Suggest Trouble

Neglecting end effects, diffraction and other details, energy kick per particle pass for a relativistic beam is roughly

$$\Delta\gamma_j \approx \frac{q a_u N_u \lambda_u}{2mc^2\beta_0\gamma_0}\sqrt{G}\Big[E_j\sin(\phi_j) + \sum_{k\neq j} E_k \sin(\phi_j - \phi_k) + \eta\Big]$$

Cooling rate longitudinal emittance can be approximated as

$$\tau_c^{-1} \approx -\tfrac{1}{2} f_c\Big[\frac{\langle\Delta\gamma_j^2\rangle + \langle 2\delta\gamma_j\Delta\gamma_j\rangle}{\langle\delta\gamma_j^2\rangle}\Big]$$

averages are over particle distribution function and any stochasticity in radiation emission and amplification

- Do individual particles radiate in random "quantum jumps", on average once in every $O(\alpha^{-1})$ passes through pickup?
  - in naive picture (akin to M. Sand's treatment of synchrotron radiation damping), where particles emit photons in discrete Poissonian quantum jumps, it would seem particles rarely experience a large kick from cooling self-radiation but almost always feel large heating cross-radiation, leading to drastically slower cooling, over-correction or unstable feedback
- With classical emission:
  - radiation amplitudes and phases are deterministic, given particle trajectory
  - with optimal gain and time-of-flight delays: $\tau_c^{-1} \approx f_c \frac{1}{1+\frac{N_s+N_n}{2\pi^2}}$
- With Poissonian photon emission but no additional phase uncertainty?
  - radiation intensity is random, with Poissonian emission probability $p \sim O(\alpha)$ per particle par pass,
  - cooling rate is suppressed: $\tau_c^{-1} \approx p f_c \frac{1}{1+\frac{N_s+N_n}{2\pi^2}}$
  - $O(\alpha^{-1})$ slower!
- With additional Heisenberg phase uncertainty?
  - radiation phase is subject to random fluctuations with large variance
  - cooling signal is corrupted, and rate is suppressed: $\tau_c^{-1} \approx e^{-N_{\rm ph}^{-1}} p f_c \frac{1}{1+\frac{N_s+N_n}{2\pi^2}}$
  - a factor of $O\big(e^{\alpha^{-1}}\big) \sim 10^{60}$ slower still!

Particle Dynamics in Wigglers are Classical

- A fully quantum-mechanical (or worse, QED) of particle dynamics would be very difficult K-J Kim's talk), but fortunately is not actually necessary
  - require classical statistical and dynamical behavior
  - must examine spin, longitudinal, and transverse degrees-of-freedom, including recoil effects
  - must examine both averages and fluctuations/uncertainties
- Careful examination reveals a plethora of constraints:
  - beam particles are non-degenerate: $\rho_b \ll \min\left[\left(\frac{\epsilon_\perp}{\lambda_c}\right)^3 \frac{\gamma}{\sigma_\perp^3}, \frac{1}{\lambda_c^3}\left(\frac{\delta\gamma}{\gamma}\right)^3\right]$
  - pair creation and other exotic QED effects are negligible: $\frac{\delta\gamma}{\gamma} \ll 1$
  - certain spin effects are negligible: $\frac{a_u}{\gamma} \ll 1$
  - a variety of other quantum spin, longitudinal, transverse effects can be ignored:

$$\frac{\hbar\omega_0}{mc^2} \ll \min\left[\delta\gamma, \gamma, a_u, a_u\gamma, \frac{a_u^2}{\gamma}, \frac{\gamma^3}{a_u}, N_u, \frac{\gamma}{\sqrt{N_u}}, \frac{\gamma}{N_u}, \frac{\gamma}{N_u^2 a_u}, \frac{\gamma^3}{N_u a_u}, \frac{\gamma}{\alpha N_u^2 a_u^2}, \frac{\delta p_\perp}{mc}, \frac{\gamma\delta p_\perp}{mc}\right]$$

- Caveats:
  - "fundamental criterion" $\frac{N_u\lambda_c}{\gamma} \ll \lambda_0$ is not typically the most stringent in a short wiggler
  - these conditions are strictly neither logically necessary nor sufficient for classical behavior
  - decoherence and statistical mechanics come to the rescue...
- Furthermore, particle motion is determined by external fields only:
  - radiation reaction is negligible
  - Vlasov self-fields in beam are marginally negligible

Fundamental Theorems of Quantum Optics

Neglecting end effects, diffraction and other details, energy kick per particle pass for a relativistic beam is roughly

$$\Delta\gamma_j \approx \frac{q a_u N_u \lambda_u}{2mc^2 \beta_0 \gamma_0} \sqrt{G} \left[E_j \sin(\phi_j) + \sum_{k \neq j} E_k \sin(\phi_j - \phi_k) + \eta \right]$$

Cooling rate longitudinal emittance can be approximated as

$$\tau_c^{-1} \approx -\tfrac{1}{2} f_c \left[\frac{\langle \Delta\gamma_j^2 \rangle + \langle 2\delta\gamma_j \Delta\gamma_j \rangle}{\langle \delta\gamma_j^2 \rangle} \right]$$

averages are over particle distribution function and any stochasticity in radiation emission and amplification

- Do individual particles radiate in random "quantum jumps", on average once in every $O(\alpha^{-1})$ passes through pickup?
  - in naive picture (akin to M. Sand's treatment of synchrotron radiation damping), where particles emit photons in discrete Poissonian quantum jumps, it would seem particles rarely experience a large kick from cooling self-radiation but almost always feel large heating cross-radiation, leading to drastically slower cooling, over-correction or unstable feedback
- With classical emission:
  - radiation amplitudes and phases are deterministic, given particle trajectory
  - with optimal gain and time-of-flight delays: $\tau_c^{-1} \approx f_c \frac{1}{1 + \frac{N_s + N_n}{2\pi^2}}$
- With Poissonian photon emission but no additional phase uncertainty?
  - radiation intensity is random, with Poissonian emission probability $p \sim O(\alpha)$ per particle par pass,
  - cooling rate is suppressed: $\tau_c^{-1} \approx p f_c \frac{1}{1 + \frac{N_s + N_n}{2\pi^2}}$
  - $O(\alpha^{-1})$ slower!
- With additional Heisenberg phase uncertainty?
  - radiation phase is subject to random fluctuations with large variance
  - cooling signal is corrupted, and rate is suppressed: $\tau_c^{-1} \approx e^{-N_{\rm ph}^{-1}} p f_c \frac{1}{1 + \frac{N_s + N_n}{2\pi^2}}$
  - a factor of $O\left(e^{\alpha^{-1}}\right) \sim 10^{60}$ slower still!

Quantum Mechanical Amplifiers

Early ideas by C. Townes and other pioneers refined and clarified by C. Caves and others research gravity wave detectors

- Axioms of QM alone are sufficient to characterize the best-case action of linear amplifiers and lower limits on amplifier noise
  - noise typically can be traced to spontaneous emission in the gain medium, but detailed modeling is not necessary
  - input (pre-amplification) and output (post) amplification modes are linearly related in Heisenberg picture
  - in order that time evolution remain unitary, the amplifier must also add extra noise to each amplified mode
  - at input, the minimal extra noise is equivalent to another set of vacuum fluctuations, or "half a photon per mode"
- Why does quantum mechanics "extract its due twice"?
  - original input noise is amplified along with the signal because of linearity: a linear amplifier cannot know what part of the input state will later be regarded as noise or as signal
  - additional noise, equivalent to another half-photon, or independent set of vacuum fluctuations, is added as a consequence of the uncertainty principle: a phase-insensitive linear amplifier amplifies both quadrature components of a field mode, which do not commute, to arbitrarily levels that can be classically measured without further back-action, so at least one additional DOF must be involved
- Effects on QM states of the fields
  - most easily calculate using quasi-distribution functions or their Fourier transforms (quasi-characteristic functions)
  - coherent states (displaced vacuum states) are transformed into displaced thermal states
  - classical (positive-*P*) states remain classical, or become even more so
  - non-classical (negative-*P*) tend to become more classical, through convolution with added noise
  - note high-gain amplifier does not appreciably change SNR, but only magnifies fields to levels where both phase components can be measured classically

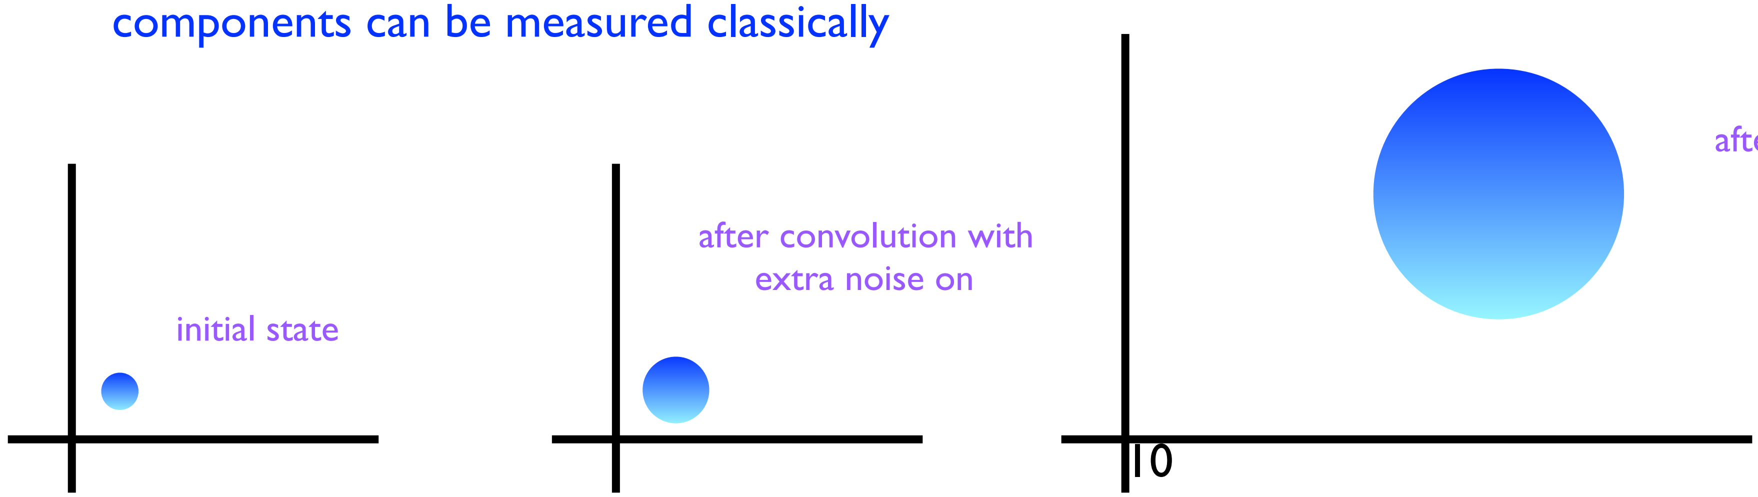

Can the Quantum Noise Limit Be Achieved?

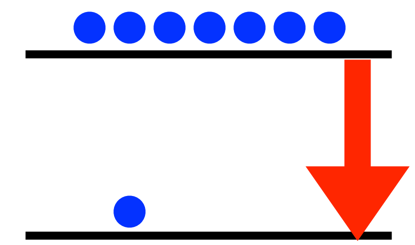

- Coupling to various Heat Baths?
  - the initial radiation state into which particles radiate is thermal (blackbody) rather than vacuum, but for IR or optical frequencies at reasonable operating temperatures, these are indistinguishable
  - approaching the quantum noise limit for OSC far easier than in the RF or microwave regimes
  - even if some thermal noise is added to the the minimal quantum noise, it is also additive, and not expected to be catastrophic
- Finite Pump Strength
  - finite pump strength ⇒ finite population inversion ⇒ less-than-ideal gain and "excess" spontaneous emission noise
  - depends on details of amplifier, but a simple two-state inverted population model suggests that the effective number of noise photons per mode at input is approximately

$$\bar{N}_{\rm n} = \tfrac{1}{2} + (P - \tfrac{1}{2})|1 - \tfrac{1}{G}|$$

$G \geq 1$ is power gain

$P = \frac{n_{\rm exc}}{n_{\rm exc} - n_0} \geq 1$ is population inversion factor

note $\bar{N}_{\rm n} \to 1^+$ as $P \to 1^+$ and $G \to \infty$

  - Back-of-the-envelope estimate for multi-stage Ti:Sapphire Laser suitable for OSC might correspond to $G \sim O\left(10^4\right) \gg 1$ and $P \sim O\left(\frac{7}{4}\right)$ corresponding to a noise level only 25% higher than the ideal lower bound
- Multiplicative Noise?
  - Quantum mechanics neither requires nor prevents an actual amplifier from exhibiting some multiplicative noise in addition to the required additive noise, for example due to pump fluctuations
  - But it will not be of the drastic all-or-nothing character that proved so deleterious to cooling in the Poissonian emission model....

Some Formulae

- Charman and JW *Quantum Mechanical Treatment of Transit-Time Optical Stochastic Cooling of Muons* http://arxiv.org/abs/0905.0485
- General discussion of many of the issues presented today

Energy gain in kicker

$$mc^2\Delta\gamma_j \approx q \int_{t_j+\Delta t_j}^{t_j+\Delta t_j+T_j} dt\, \boldsymbol{v}_{\perp j}(t)\cdot \boldsymbol{E}(\boldsymbol{x}_j(t),t),$$

First order statistics

$$\langle \Delta\gamma_j \,|\, \gamma_j, \gamma_0\rangle \approx \frac{q a_u N_u \lambda_u \sqrt{G} E_0}{2mc^2\beta_0\gamma_0}\Theta\left(1-\frac{|\phi_{jj}|}{2\pi N_u}\right)\left[1-\frac{|\phi_{jj}|}{2\pi N_u}\right]\sin(\phi_{jj})$$

Second order statistics

$$\begin{aligned}\langle (\Delta\gamma_j)^2 \,|\, \gamma_j, \gamma_0\rangle \approx \left(\frac{q a_u N_u \lambda_u \sqrt{G} E_0}{2mc^2\beta_0\gamma_0}\right)^2 \Bigg[&\Theta\left(1-\frac{|\phi_{jj}|}{2\pi N_u}\right)\left(1-\frac{|\phi_{jj}|}{2\pi N_u}\right)^2 \sin^2(\phi_{jj}) \\ &+ \frac{N_b}{L_b}\frac{N_u\lambda_0}{3}\left(1-\frac{3}{8\pi^2 N_u^2}\right)\Theta\left(1-\frac{|\phi_{jj}|}{k_0 L_b}\right)\left(1-\frac{|\phi_{jj}|}{k_0 L_b}\right) \\ &+ \left\{\frac{1}{N_{\mathrm{ph}}}\frac{1}{\pi^2}\frac{N_u^2}{b_L(b_L^2-4N_u^2)}\sin^2(\pi b_L) + \right. \\ &\left. + \frac{1}{N_{\mathrm{ph}}}\frac{1}{8\pi}\left\{2\,\mathrm{Si}\,(2\pi b_L) - \mathrm{Si}\,(2\pi[b_L-2N_u]) - \mathrm{Si}\,(2\pi[b_L+2N_u])\right\}\right\}\Bigg]\end{aligned}$$

$$b_L = \frac{1}{2}N_u\frac{\Delta\omega_I}{\omega_0}$$

Discussion

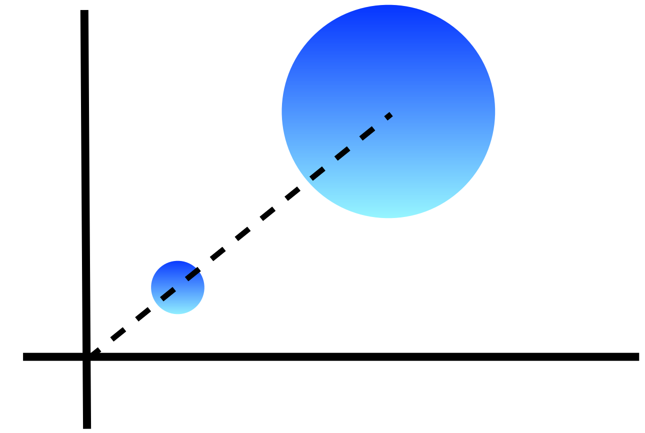

- Why "Naive" Quantum Treatments Fail
  - particles do radiate into number states, but rather Glauber coherent states, or mixtures thereof
  - photon number is not sharply defined, but phase information is partly available, and uncertainties are minimal
  - particles do not radiate whole photons stochastically in a series of quantum jumps
  - particles really do radiate "a fraction of a photon" on every pass, available for amplification and feedback
  - pickup field is on average exactly that expected classically, with additional "vacuum fluctuations"
  - in kicker, amplified radiation is essentially classical, and particles respond linearly to fields rather than via discrete photon absorption
  - a linear amplifier does not multiply photons, but acts unitarily on the full QM state of radiation field, effectively scaling amplitude after adding some extra noise at input
  - nothings emits, absorbs, prepares, measures, projects into, multiplies, or otherwise singles out Fock states
  - environmental decoherence only expected to make radiation appear more classical
- Why does Sand's Treatment Work For Synchrotron Radiation Damping (SRD)?
  - OSC involves negligible radiation reaction but strong feedback which is linear in the fields– actual state of the radiation field is important
  - SRD involves strong radiation reaction effects, but no feedback– after emission, radiation DOF can be "traced out" to obtain a Focker-Planck equation for electron DOF
  - As long as 1st and 2nd moments are reasonably accurate, Poissonian "quantum jump" model for SRD mimics correct random walk statistics for electrons while not faithfully representing the actual emission process or radiation phase space
- The Bottom Line:
  - quantum mechanical effects do not destroy cooling, but just introduce additional (additive) noise approximately equivalent (prior to amplification) to $O(\alpha^{-1})$ additional particles in sample, setting an ultimate upper limit on cooling rates and efficacy of further beam dilution

What We Have Learned

"… the wiggler is capable of reducing the wavefunction [of an electron] by causing it to emit a photon. The electron's position can therefore be deduced from subsequent measurements."

S. Benson, and J. M. J. Madey, "Shot and Quantum Noise in Free Electron Lasers," 1985

This is basically **incorrect**. The wiggler magnet itself does not reduce the particle's wavefunction—only a subsequent photon-counting measurement will do that, which never happens in the stochastic cooling situation. And besides, with many electrons in a sample length, observation of a photon can tell us very little about the position or momentum of any one electron. If subsequent measurements on the electron or radiation reveal information about the electron's position, then its wavefunction ``collapses"

"We find that a classical description of the input fields and of the amplification process is completely valid provided we take correctly into account the response of the amplifier to the input zero-point fields. This result is valid for inputs of arbitrarily small power."

J. P. Gordon, W. H. Louisell, and L. R. Walker, "Quantum Fluctuations and Noise in Parametric Processes," 1963

This is basically **correct**, apart perhaps from a possible factor of two, arising, as we have seen, from the extra set of amplified vacuum fluctuations which must necessarily be added to the field to ensure that Heisenberg's uncertainty relations remain valid for the non-commuting components of the field, in addition to the original amplified vacuum fluctuations, which must be amplified along with the signal, because the amplifier cannot know what part of the input field is to be regarded as signal and what part is considered noise. But besides that, radiation is typically not emitted into number states, and not amplified in the number state basis.

"No phenomenon is a phenomenon until it is an observed phenomenon." –John Wheeler

Ultra-fast to Moderately Fast Cooling

- Feared Quantum Cooling Catastrophe was a Red Herring!
  - quantum mechanical effects do not prevent, destroy, or radically slow cooling, but just introduce some additional (additive) noise approximately equivalent (prior to amplification) to $O(\alpha^{-1})$ additional particles in each particle's sample
  - ultra-fast OSC may face many technical challenges, but no problem in principle with quantum fluctuations or quantum mechanical amplifiers
  - OSC works with with weak pickup signals just as predicted classically, and can be calculated essentially classically, with some extra noise
- Examples
  - Ultra-fast OSC for final cooling in a muon collider is probably beyond current technology, but may be achievable in coming decades

    cooling a $\gamma \sim O(10^3)$ beam with $N_b \sim O(10^9)$ on a time-scale $\tau_c \sim O(10^{-7}\text{s})$ in order to achieve an $O(10^{-3})$ relative decrease in longitudinal emittance would require stretching to a size where $N_s \sim O(10^2)$ and about $O(10)$ lasers amplifiers, each producing $O(10^2\ \text{W})$ of average power at $O(10^2\ \text{Hz})$ repetition rate

  - For ultra-high luminosity proton/anti-proton beams, OSC could be used to cool moderately faster than RF stochastic cooling to lower diffusive losses at walls:

    cooling of an $O(10^3\ \text{GeV})$ anti-proton beam with $N_b \sim O(10^{10})$, $N_s \sim O(10^5)$ on a time-scale $tau_c \sim O(10^3\text{s})$ would require an amplifier producing about $O(2\ \text{W})$

  - Although synchrotron radiation damping is typically efficient, OSC could also be used for fast electron cooling

    cooling within an $O(10^{-1}\ \text{Gev})$ electron storage ring with $N_b \sim O(10^9)$, with $N_s \sim O(10^5)$ on a time-scale $\tau_c \sim O(10^{-1}\ \text{s})$ could work with an amplifier delivering only about $O(10^{-3}\ \text{W})$ of power

Unfinished Business?

Examine multi-dimensional effects in beam and radiation dynamics more carefully...

- Use full betatron particle orbits for reasonable lattice parameters
- Allow for transverse and other optical effects so far neglected:
  - dispersion
    - since cooling depends critically on phasing, any dispersion in the amplifier or optical transport system will be important
    - average time-of-flight delay within passband can be easily compensated for, but any group velocity dispersion or nonlinear phase modulation will be deleterious
  - diffraction/transverse effects
    - certain paraxial effects may be very important
    - diffraction over the interaction region
    - Guoy phase
  - near-field effects
    - collection optics and amplifier may not be in far field of pickup wiggler
    - certain paraxial effects may be very important

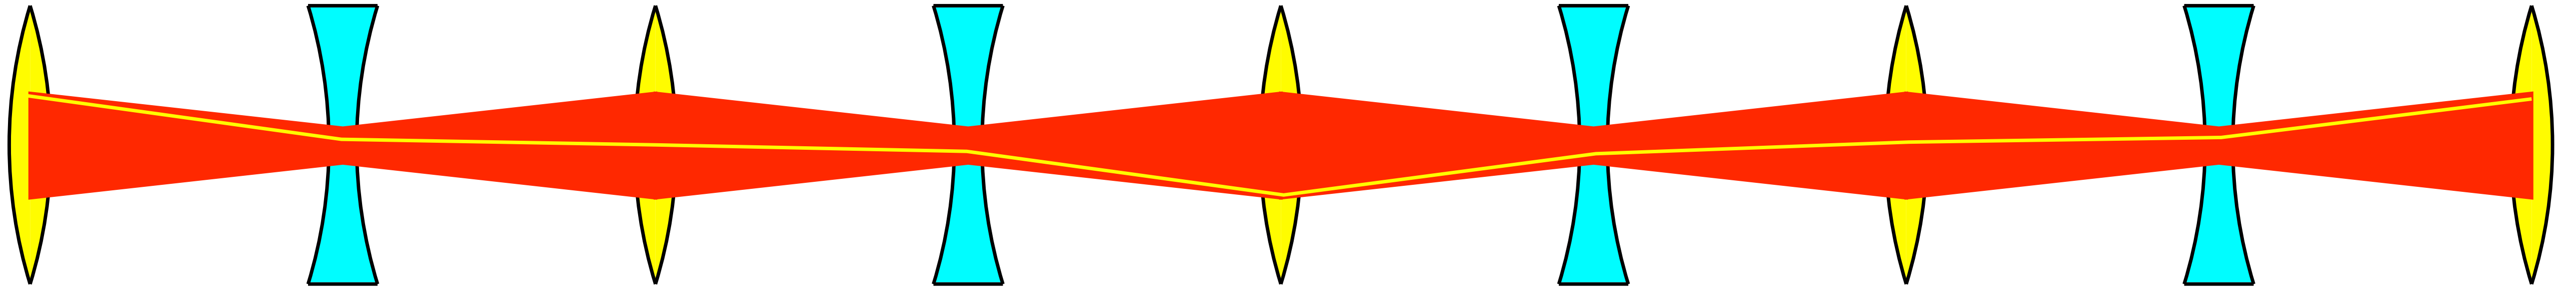

Background and References

- Address questions to either Andy (acharman@berkeley.edu) Charman and myself (wurtele@berkeley.edu). We present a different approach to that of S. Heifets and M. S. Zolotorev. *Quantum theory of optical stochastic cooling*. Physical Review E **65**:016507 (2001). Answers are essentially the same.

- Details first appeared in Andy's 2007 UCB Ph.D. Thesis:

  "Random Aspects of Beam Physics and Laser-Plasma Interactions":
  http://raman.physics.berkeley.edu/papers/Charman_Thesis.pdf

- Preprint Charman and JW:
  Quantum Mechanical Treatment of Transit-Time Optical Stochastic Cooling of Muons
  http://arxiv.org/abs/0905.0485

- Perhaps IOTA-related: Charman, Penn, and JW:
  A Hilbert-space formulation of and variational principle for spontaneous wiggler radiation
  http://arxiv.org/abs/physics/0501018. *A Hilbert−Space Variational Principle for Spontaneous Wiggler and Synchrotron Radiation* http://scitation.aip.org/content/aip/proceeding/aipcp/10.1063/1.2409159

Extra

Stochastic Cooling

- Cooling through coupling and feedback to additional degrees-of-freedom
  - dynamics in reduced single-particle phase space of beam particles is Non-Liouvillian
  - with suitable choice of signal, manipulation, amplification, and feedback, longitudinal and/or transverse emittance can be reduced
- Cooling occurs in spite of, or perhaps because of noise in beam
  - cooling forces accumulate coherently as drift
  - heating/noise accumulates incoherently, as diffusion
- Stochastic cooling is non-evaporative
  - cooling is not achieved at expense of removing particles in tail
  - beam brightness can be increases as emittance is reduced
- Generic Stochastic Cooling Pass:
  - interaction of charge particle beam in pickup generates an EM cooling signal
  - beam is prepared, and signal is amplified, manipulated, and fed back on beam in kicker
  - beam particles might be scrambled a it in mixer to ensure noise accumulates diffusively

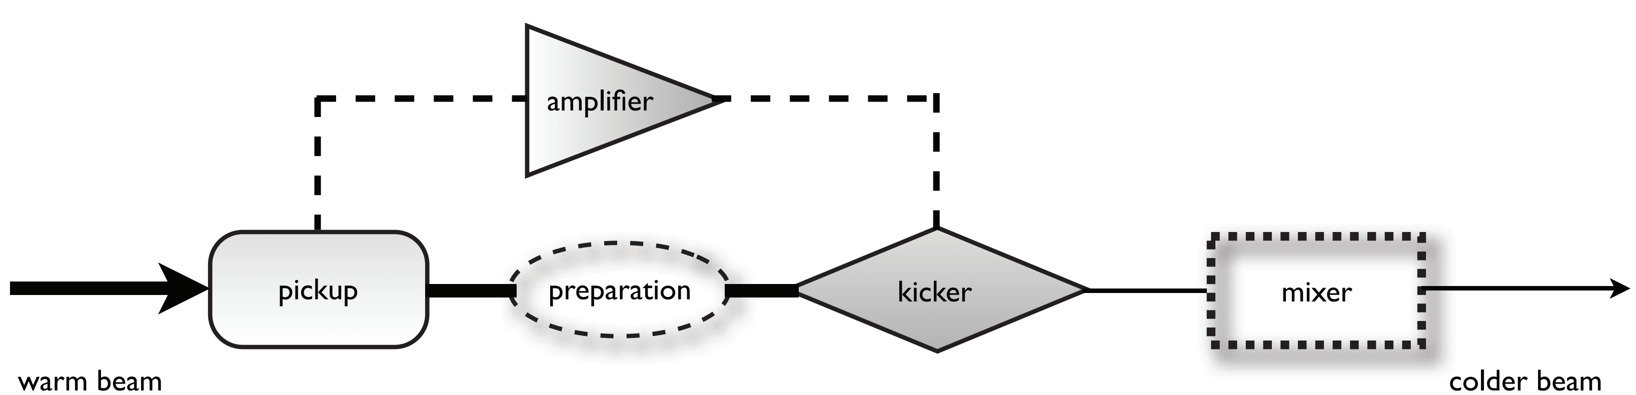

Stochastic Cooling Rates

For a beam perturbation linear in the kicker fields, the instantaneous cooling rate for a cooled degree-of-freedom (e.g., energy spread, or transverse momentum) will be:

frequency of passage through cooling section(s)

net power gain

adjusted as a function of time t to attempt to optimize cooling rate

cooling rate

$$\tau_c^{-1}(t) \approx f_c \left[B(t)\sqrt{G(t)} - \tfrac{1}{2} A(t) G(t) \right]$$

$$B(t) > 0$$

depends on the cooling scheme employed and the current particle distribution function, and represents the ***drift****, or* ***coherent cooling*** *effects arising from interaction of each particle with its own amplified feedback signal*

from self-fields

from other fields

$$A(t) = A_0 + A_1\,(N_s + N_n) + \ldots > 0$$

represents the ***diffusive****, or* ***incoherent heating*** *effects arising from interaction of each particle with its own and neighboring particles' signals, as well as amplifier or other sources of noise*

N_s is the effective ***sample size****, or average number of particles with whose signals a given particle interacts in the kicker*

N_n is a measure of extra ***amplifier noise*** *introduced into the signal, expressed as an equivalent number of extra particles in the sample*

Single-Particle Versus Beam Statistics

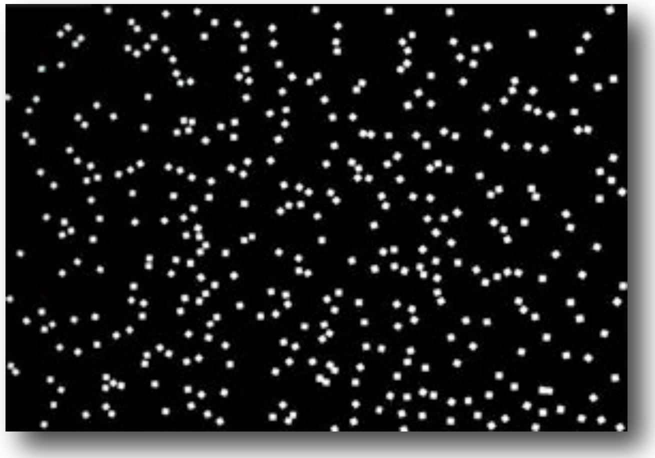

- In conventional (RF) Stochastic Cooling
  - many small kicks over many passes
  - leads naturally to a smoothed-out time-averaging amenable to a Fokker-Planck analysis
- In fast Optical Stochastic Cooling:
  - magnitude and even direction of any particular particle's energy kick on any one pass are subject to large uncertainty, and cannot be estimated reliably
  - beam is cooled with a relatively small number of relatively large kicks, so time-averaging or smoothing is less effective at improving predictability for individual particles
  - however, particles remain largely uncorrelated during cooling, so while prediction of individual particle behavior is unreliable, averaging over all particles in beam is expected to give $O(\frac{1}{\sqrt{N_b}})$ improvement in predictive accuracy
  - what we really care about are quantities like

$$\delta\gamma_{\mathrm{RMS}}(t) = \left[\frac{1}{N_b}\sum_{j=1}^{N_b}\left(\gamma_j(t) - \gamma_0\right)^2\right]^{\frac{1}{2}}$$

Bandwidth, Sample Size, and Cooling Rate

The typical **sample size** scales like:

transverse spot size of beam

particle density in beam

velocity of beam

Sample Size

$$N_s \sim \rho_b \min\left[\sigma_{b\perp}^2, S_c\right] v_b \Delta\omega^{-1}$$

transverse coherence area of kicker fields

limiting bandwidth of pickup/amplifier/kicker

Because the incoherent heating contribution grows faster with amplifier gain

than the coherent cooling term, at any time there is a locally optimal gain $G(t) = \left[\frac{B(t)}{A(t)}\right]^2$

maximizing the instantaneous cooling rate $\tau_c^{-1}(t) = \frac{1}{2} f_c \frac{B(t)^2}{A(t)}$

Typically, the locally optimal gain starts off relatively large, and then tends to decrease as the beam cools and approaches an asymptotic equilibrium distribution with finite emittance in which the heating and cooling contributions balance