

Nonlinear Single-Particle Dynamics in High Energy Accelerators

Part 8: Case Study – The ILC Damping Wiggler

Nonlinear Single-Particle Dynamics in High Energy Accelerators

This course consists of eight lectures:

1. Introduction – some examples of nonlinear dynamics
2. Basic mathematical tools and concepts
3. Representations of dynamical maps
4. Integrators I
5. Integrators II
6. Canonical Perturbation Theory
7. Normal form analysis
8. A case study

We have developed tools for:

- calculating particle motion in complex magnetic fields, including 2D fields (multipoles) and 3D fields (for example, insertion devices);
- analysing particle motion in beamlines, to evaluate the impact of nonlinear effects on a beam moving through the accelerator.

There are a range of analytical and numerical tools available for the above tasks. Which ones you use will depend on the system you are dealing with.

To illustrate the application of some of the methods developed in the previous lectures, we shall now consider a case study: the wiggler in the ILC damping rings.

As well as reviewing and applying material from previous lectures, we shall discuss two techniques we have not so far covered in any detail:

- use of a differential algebra (DA) code to construct a map in Taylor form, suitable for tracking;
- use of frequency map analysis to visualise effects such as tune shift with amplitude and resonances in storage rings.


We shall proceed as follows:

1. Introduction: the ILC damping rings and damping wiggler.
2. Fitting the wiggler field.
3. Constructing the dynamical map for the wiggler.
4. Analysis of the impact of the wiggler on the dynamics.

Ideally, there should be an additional section: optimisation of wiggler and storage ring optics design to minimise impact of nonlinear effects. But we have not fully achieved that step, yet!

Introduction: the ILC damping rings

The International Linear Collider (ILC) is a proposed facility for studies in high energy physics. The goal is to produce a luminosity of $2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, with centre of mass collision energy 500 GeV (initially).



To achieve the high luminosity, damping rings are used to provide highly stable, very low emittance beams.

In many respects, the damping rings are similar to synchrotron storage rings used (for example) in third generation light sources. The main difference is that after injection, the beam is stored just long enough to reach equilibrium (200 ms) before being extracted and accelerated to the interaction point.

Circumference	6476 m
Beam energy	5 GeV
Average current	400 mA
Natural emittance	0.6 nm
Transverse damping time	21 ms

Recall that the damping time τ in a synchrotron storage ring is given by:

$$\tau = 2 \frac{E}{U} T. \quad (1)$$

where E is the beam energy, U the energy loss per turn, and T is the revolution period.

The energy of the beam is a compromise between costs and emittance (both increase with energy) and collective effects (which become more severe at lower energy). The size of the ring is driven by the need to inject and store a train of more than 5000 bunches. And the damping time is determined by the injected and extracted beam emittances, and the store time:

$$\varepsilon(t) = \varepsilon(0)e^{-\frac{t}{2\tau}} + \varepsilon_{equ} \left(1 - e^{-\frac{t}{2\tau}}\right). \quad (2)$$

To achieve the required damping times for the specified circumference and beam energy, a large energy loss per turn of 10 MeV is needed. Roughly 90% of the energy loss per turn will be provided by a damping wiggler, of total length 215 m, and peak field 1.6 T.

There are many technical challenges associated with the wiggler. The power requirements make a normal-conducting electromagnetic wiggler unattractive. A permanent magnet wiggler is a possibility, but the magnetic material would be expensive, and there may be concerns with radiation hardness or activation.

Presently, the baseline is for a superferric wiggler, which can provide good field quality with wide aperture.

Dynamic octupole in an “ideal” wiggler

Even a wiggler with infinitely wide poles can have strong nonlinear effects. To understand these, let us perform a simple analysis of the dynamics in the field of an idealised wiggler.

Recall that, in Cartesian coordinates, the field of a wiggler may be represented by:

$$B_x = - \sum_{m,n} B_{mn} \frac{mk_x}{k_y} \sin mk_x \sinh k_y y \cos nk_s s, \quad (3)$$

$$B_y = \sum_{m,n} B_{mn} \cos mk_x \cosh k_y y \cos nk_s s, \quad (4)$$

$$B_s = - \sum_{m,n} B_{mn} \frac{nk_s}{k_y} \cos mk_x \sinh k_y y \sin nk_s s, \quad (5)$$

where:

$$k_y^2 = m^2 k_x^2 + n^2 k_z^2. \quad (6)$$

For an idealised wiggler, we can retain only the fundamental mode, $m = 0$, $n = 1$.

The idealised wiggler field is then:

$$B_x = 0, \quad (7)$$

$$B_y = B_{01} \cosh k_s y \cos k_s s, \quad (8)$$

$$B_s = -B_{01} \sinh k_s y \sin n k_s s. \quad (9)$$

Since we are making a relatively crude estimate of the dynamics, we shall work with non-canonical variables, and write the equations of motion as:

$$x'' = -\frac{B_y}{B\rho} + y' \frac{B_s}{B\rho}, \quad (10)$$

$$y'' = \frac{B_x}{B\rho} - x' \frac{B_s}{B\rho}. \quad (11)$$

If we assume that $|y'| \ll 1$, then the equation for x (10) becomes:

$$x'' \approx -\frac{B_y}{B\rho} = -\frac{B_{01}}{B\rho} \cosh k_s y \cos k_s s. \quad (12)$$

The solution is:

$$x \approx x_0 + x'_0 s + a \cosh k_s y \cos k_s s, \quad (13)$$

where a , the amplitude of the “wiggles”, is given by:

$$a = \frac{B_{01}}{k_s^2 B\rho}. \quad (14)$$

Note that for $y = 0$ the motion in x is linear.

Now we write the equation for y (11):

$$y'' \approx -x' \frac{B_s}{B\rho} = -ak_s \frac{B_{01}}{2B\rho} \sinh 2k_s y \sin^2 k_s s. \quad (15)$$

We can estimate the total vertical deflection through one period of the wiggler by:

$$\Delta y' = \int_0^{\lambda_w} y'' ds \approx -\frac{\pi}{2k_s^2} \left(\frac{B_{01}}{B\rho} \right)^2 \sinh 2k_s y, \quad (16)$$

where $\lambda_w = 2\pi/k_s$ is the wiggler period. Note that we assume $x_0 = x'_0 = 0$.

Expanding the sinh function:

$$\Delta y' \approx -\frac{\pi}{k_s} \left(\frac{B_{01}}{B\rho} \right)^2 \left[y + \frac{2}{3} k_s^2 y^3 + \dots \right]. \quad (17)$$

We see that the vertical deflection is given by an infinite series in odd powers of y . The linear term indicates that there is vertical focusing.

The third order term is what we would expect from an octupole. This is sometimes referred to as a “dynamic” octupole, since its origin is not an octupole field (which would vary as the third power of the distance from the centre of the magnet), but the result of the sinusoidal horizontal trajectory through a field that varies as the *fourth* power of y .

The strength of the dynamic octupole is given by:

$$[k_3\ell]_{\text{dynamic}} = -4\pi k_s \left(\frac{B_{01}}{B\rho} \right)^2. \quad (18)$$

Note that dynamic multipoles behave rather differently from regular multipoles: their effects in the horizontal and vertical planes are generally different, and they have an unusual scaling with particle energy.

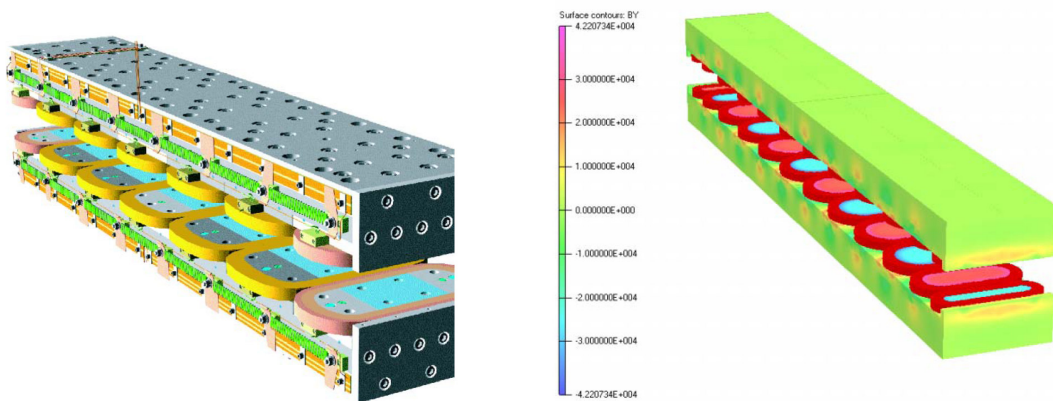
Higher order modes in a wiggler field can lead to significant modification of the dynamics, compared to a simple analysis based on an idealised field.

This motivates a more detailed analysis, using detailed field map produced from a magnetic modelling code.

To proceed further, we need a specific wiggler design...

The ILC damping wigglers

The present baseline design specifies wigglers based on the CESRc superferric wigglers (below, left). These have seven poles with peak field 2.1 T, period 0.4 m. The modified design for ILC (below, right) will have twelve poles, peak field 1.6 T, and period 0.4 m.



J. Urban, G. Dugan, M. Palmer, J. Crittenden.


Before the present baseline was selected, there were extensive studies of the impact of nonlinear dynamics in the wiggler, and the potential impact on the performance of the damping rings.

One of the main concerns was that the wiggler fields would cause the trajectories of particles to become unstable at large betatron or synchrotron amplitude. In other words, the wigglers might limit the dynamic aperture of the rings.

Since the injected positron beam in particular has a very large emittance, and the *average* injected beam power during operation would be 225 kW, injection losses resulting from dynamic aperture limitations could have a serious impact.

The ILC damping wigglers

The “configuration studies” were based on a permanent magnet wiggler, which was the best available model at the time.



TESLA TDR (2001)

The superferric wiggler (since adopted as the baseline) appears to have a better field quality than the permanent magnet wiggler, though a thorough, detailed study has yet to be carried out.


The first step is to obtain a numerical field map from a modelling code such as OPERA. In principle, one can integrate numerically the equations of motion for particles moving through this field. However, when the total length of wiggler is over 200 m, and we are interested in hundreds of thousands of particle-turns, this is not a practical approach.

Instead, we can fit a set of “modes” to the wiggler field. This provides a semi-analytical description of the field that can be used for constructing a dynamical map. The dynamical map forms the basis for the tracking studies.

Fitting a field map

The procedure of fitting a field map to numerical field data has already been described in Lecture 5. Essentially, the mode coefficients may be obtained from a 2D discrete Fourier transform of the radial field component on a cylinder inscribed within the wiggler.

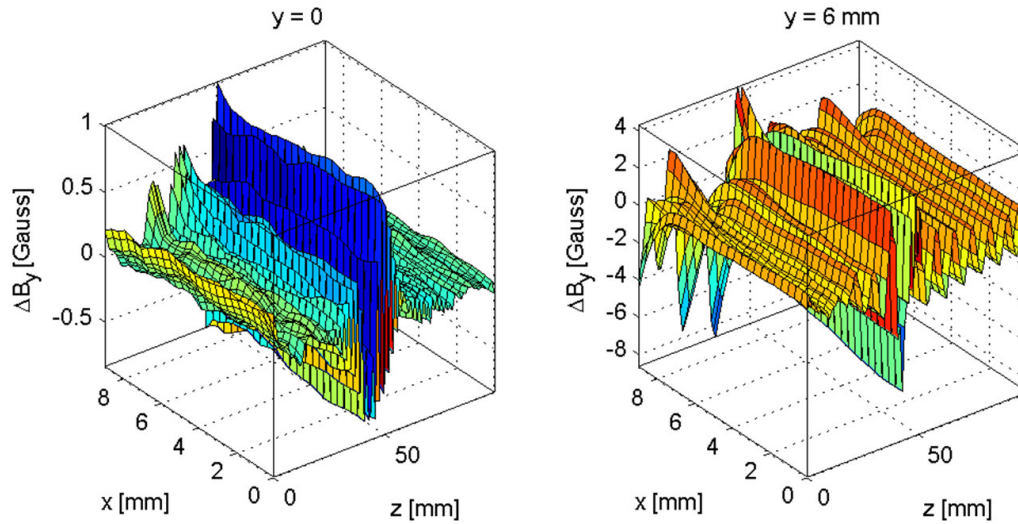
$$\begin{aligned}
 B_\rho &= \int dk_s \sum_m \tilde{B}_m(k_s) I'_m(k_s \rho) \sin m\phi \cos k_s s, \\
 B_\phi &= \int dk_s \sum_m \tilde{B}_m(k_s) \frac{m}{k_s \rho} I_m(k_s \rho) \cos m\phi \cos k_s s, \\
 B_s &= - \int dk_s \sum_m \tilde{B}_m(k_s) I_m(k_s \rho) \sin m\phi \sin k_s s.
 \end{aligned}$$




The behaviour of the modified Bessel function $I_m(k_s \rho)$ causes residuals of the fit to reduce exponentially towards the axis of the cylinder.

Fitting a field map

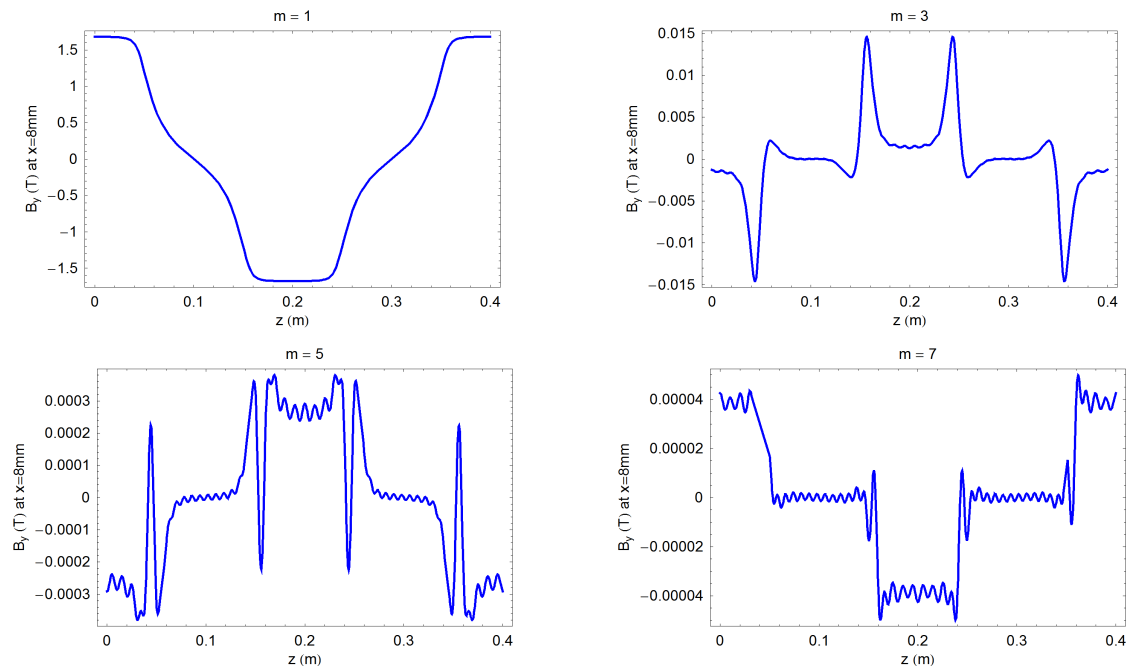
For the TESLA wiggler design, we obtain a good-quality fit by fitting on a cylinder of radius 9 mm, using 18 azimuthal and 100 longitudinal modes.



Fitting a field map



The coefficients of the mode decomposition in a cylindrical basis provide the multipole components as a function of distance along the wiggler axis.



Constructing the dynamical map

The field modes allow us to reconstruct the field, or the vector potential, at any point within the wiggler.

Since the field is three-dimensional, it is appropriate to use the Wu-Forest-Robin integrator, again described in Lecture 5. The second-order version of the integrator may be written:

$$\mathcal{M}(\Delta\sigma) \approx e^{-\frac{\Delta\sigma}{4} \cdot H_1} e^{-\frac{\Delta\sigma}{2} \cdot H_3} e^{-\frac{\Delta\sigma}{4} \cdot H_1} e^{-\Delta\sigma \cdot H_2} e^{-\frac{\Delta\sigma}{4} \cdot H_1} e^{-\frac{\Delta\sigma}{2} \cdot H_3} e^{-\frac{\Delta\sigma}{4} \cdot H_1}$$

where:

$$H_1 = -\left(\frac{1}{\beta_0} + \delta\right) + \frac{1}{2\beta_0^2\gamma_0^2} \left(\frac{1}{\beta_0} + \delta\right)^{-1} + \frac{\delta}{\beta_0} + \frac{p_x^2}{2\left(\frac{1}{\beta_0} + \delta\right)} + p_s,$$

$$H_3 = -a_s, \quad e^{-\Delta\sigma \cdot H_2} = e^{I_y} e^{-\Delta\sigma \cdot \tilde{H}_2} e^{-I_y},$$

and:

$$I_y = \int_0^y a_y(x, y', s) dy', \quad \tilde{H}_2 = \frac{p_y^2}{2\left(\frac{1}{\beta_0} + \delta\right)}.$$

Each factor in the map represents a finite set of arithmetical operations on the dynamical variables: in other words, the map can be written algebraically in closed form

Given the vector potential as a function of the coordinates, we can write a computer code to apply the map, and hence integrate the equations of motion for a particle moving through the field.

If we were limited to doing this numerically, there would not be much advantage over simply integrating through the original numerical field data.

However, by using a differential algebra code, we can construct a Taylor series to represent the result of integrating through one complete period of the wiggler...

Constructing the dynamical map: DA codes

A differential algebra code allows the user to define a DA type of variable. This can be used in the same way as any other variable, except it represents a Taylor series to some order, in a set of “basis” variables DA(1), DA(2), and so on.

For example, suppose that x and y are of type DA. The lines of code:

```
x := DA(1);  
y := cos(x);
```

would result in y having the “value”:

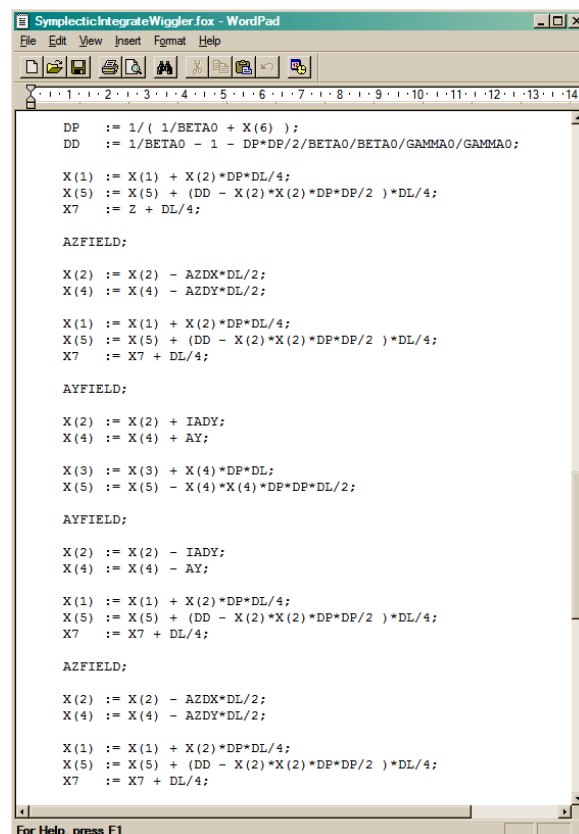
$$1 - \frac{1}{2}DA(1)^2 + \frac{1}{24}DA(1)^4 - \dots$$

I use the DA code COSY (developed by Martin Berz, http://www.bt.pa.msu.edu/index_cosy.htm) for constructing dynamical maps from sets of field mode coefficients.

A code “snippet”, showing the actual implementation of the Wu-Forest-Robin integrator, is given on the next slide, with a section of the output (consisting of the coefficients in each of the Taylor series for each dynamical variable) on the slide after.

Note that $x(1)$, $x(2)$ etc. are DA types, representing the dynamical variables.

AYFIELD and AZFIELD are procedures that “evaluate” the vector potential (and the appropriate integrals) at the current integration step. Of course, the “evaluation” is in terms of the DA variables, so the vector potential is also represented as a Taylor series.



```
DP := 1 / ( 1/BETA0 + X(6) );
DD := 1/BETA0 - 1 - DP*DP/2/BETA0/BETA0/GAMMA0/GAMMA0;

X(1) := X(1) + X(2)*DP*DL/4;
X(5) := X(5) + (DD - X(2)*X(2)*DP*DP/2 ) *DL/4;
X(7) := X(7) + DL/4;

AZFIELD;

X(2) := X(2) - AZDX*DL/2;
X(4) := X(4) - AZDY*DL/2;

X(1) := X(1) + X(2)*DP*DL/4;
X(5) := X(5) + (DD - X(2)*X(2)*DP*DP/2 ) *DL/4;
X(7) := X(7) + DL/4;

AYFIELD;

X(2) := X(2) + IADY;
X(4) := X(4) + AY;

X(3) := X(3) + X(4)*DP*DL;
X(5) := X(5) - X(4)*X(4)*DP*DP*DL/2;

AYFIELD;

X(2) := X(2) - IADY;
X(4) := X(4) - AY;

X(1) := X(1) + X(2)*DP*DL/4;
X(5) := X(5) + (DD - X(2)*X(2)*DP*DP/2 ) *DL/4;
X(7) := X(7) + DL/4;

AZFIELD;

X(2) := X(2) - AZDX*DL/2;
X(4) := X(4) - AZDY*DL/2;

X(1) := X(1) + X(2)*DP*DL/4;
X(5) := X(5) + (DD - X(2)*X(2)*DP*DP/2 ) *DL/4;
X(7) := X(7) + DL/4;
```

```

TESLACartnM28N80R8_DynamicalMapSymp.dat - WordPad
File Edit View Insert Format Help
-----
I COEFFICIENT ORDER EXPONENTS
1 1.000239793277584 1 1 0 0 0 0
2 0.4000309742401165 1 0 1 0 0 0
3 -.9867121431839730E-07 1 0 0 0 0 1
4 0.6783757565489493E-05 2 2 0 0 0 0
5 0.4526818014348874E-01 2 1 1 0 0 0
6 0.4526625824014737E-02 2 0 2 0 0 0
7 0.1220542881615254E-03 2 0 0 2 0 0
8 -.8551798174309719E-01 2 0 0 1 1 0
9 -.4799476576527517E-03 2 1 0 0 0 1
10 -.4001115817640771 2 0 1 0 0 1
11 -.1258216037294967E-01 2 0 0 2 0 0
12 0.2961588480913085E-06 2 0 0 0 0 2
-----
I COEFFICIENT ORDER EXPONENTS
1 -.4031800934778396E-13 0 0 0 0 0 0
2 0.1199017293337175E-02 1 1 0 0 0 0
3 1.000239793277579 1 0 1 0 0 0
4 -.4933770693550698E-06 1 0 0 0 0 1
5 -.1017707265693920E-03 2 2 0 0 0 0
6 -.4070656232032187E-04 2 1 1 0 0 0
7 -.2263680386123829E-01 2 0 2 0 0 0
8 -.9115148927280254E-04 2 0 0 2 0 0
9 -.1620086389709910E-03 2 0 0 1 1 0
10 -.1200813133452189E-02 2 1 0 0 0 1
11 -.4799253247135486E-03 2 0 1 0 0 1
12 0.2261261873337336E-01 2 0 0 2 0 0
13 0.9874990670609626E-06 2 0 0 0 0 2
-----
I COEFFICIENT ORDER EXPONENTS
1 0.9993439537479655 1 0 0 1 0 0
2 0.3999100315053072 1 0 0 0 1 0
3 -.1356841446148197E-04 2 1 0 1 0 0
4 -.8554441427080117E-01 2 0 1 1 0 0
5 -.4526035649920165E-01 2 1 0 0 1 0
6 -.9051686985569584E-02 2 0 1 0 1 0
7 0.1312286980151998E-02 2 0 0 1 0 1
8 -.3997114342602809 2 0 0 0 1 0
-----
I COEFFICIENT ORDER EXPONENTS
-----
For Help, press F1
    
```

The drawback with this approach is that the Taylor series is simply truncated at some specified order. That means we lose symplecticity.


Of course, it would be possible to construct, from the truncated Taylor map, a symplectic map in implicit form, as a mixed-variable generating function (Lecture 3). However, applying such maps tends not to be very efficient in practice.

Alternatively, we can hope that by generating a Taylor map to sufficiently high order, that the symplectic error will be small. But in that case, there may be little benefit in using a symplectic integration routine to construct the map, over a more conventional integration routine, e.g. Runge-Kutta.

It is interesting to compare the maps produced by the symplectic integrator in COSY, with a high-order Runge-Kutta routine built-in to COSY. The plot on the next slide compares the coefficients of corresponding terms in the maps for the transverse variables, for 5th order maps.

Note that the green lines divide terms of different order; within each order, the coefficients are arranged in ascending order (the horizontal axis is simply an index). Because we plot the coefficients on a logarithmic scale, we take the absolute value; and negative coefficients appear on a downward sloping curve.

There are some discrepancies between the maps for a few coefficients: these can be explained by the fact that the symplectic integrator makes the paraxial approximation in the Hamiltonian, whereas the Runge-Kutta integrator does not.



We can also compare the “symplectic error” in each map, by calculating the Jacobian J ; for a symplectic map, the Jacobian will satisfy:

$$J^T \cdot S \cdot J = S, \tag{19}$$

where S is the usual antisymmetric matrix.


We can quantify the symplectic error in a map by calculating:

$$\Delta = J^T \cdot S \cdot J - S, \tag{20}$$


and finding the value of the terms of different order in each component of Δ , for unit values of the dynamical variables.

If we do this for the maps constructed using the symplectic and Runge-Kutta integrators, we can show the results on plots as follows... (note that each point corresponds to a term of the specified order in one component of Δ).

Runge-Kutta integrator: symplectic error



Wu-Forest-Robin integrator: symplectic error



Note that we only plot the error in Δ up to order $N - 1$ for a map of order N , for the following reason.

Suppose that we take a symplectic map in Taylor form, and introduce an error in the coefficient of a term of order $N + 1$. In the Jacobian, this will appear at order N ; then, when constructing the product $J^T \cdot S \cdot J$, the error appears at order N and higher (assuming that the Jacobian contains terms of zeroth order and higher).

If we truncate a symplectic map at order N , then we effectively introduce errors into all coefficients in the map at order $N + 1$ and higher. So, by the above argument, we expect Δ to contain non-zero terms of order N and higher.

In practice, non-zero terms appear in Δ at order lower than N even with a symplectic integration routine, because of limits on numerical precision.

However, we see that the symplectic error is smaller for the symplectic integration, than it is for the Runge-Kutta integration.

Of course, since the exact solution to the equations of motion must be symplectic, improving the accuracy of the Runge-Kutta integration should reduce the symplectic error.

The real distinction between the two integration routines, is that the Wu-Forest-Robin integrator should give us a map in which the symplectic error up to the order of the map is the result only of limited numerical precision, even if the map itself is not accurate (e.g. because the integration step size is too large).

Characterising the dynamics

Once we have selected the maps to use for various elements in the lattice, we are ready to perform analysis of the “global” dynamics; i.e. we can investigate how particles behave over many turns through the ring.

There are a number of different methods to characterise the dynamics of a storage ring, including evaluation of the tune shifts and resonance strengths, using canonical perturbation theory (Lecture 6) or normal form analysis (Lecture 7).


The expressions we derived in canonical perturbation theory are most readily applied when the dominant nonlinearities come from multipoles, which may or may not be the case in a wiggler-dominated storage ring.

Normal form analysis is most appropriate when using specialised codes for constructing and manipulating maps in the form of Lie transformations.

Instead of perturbation theory or normal form analysis, we apply frequency map analysis. In principle, the steps are quite straightforward:

1. Track a particle for some (large) number of turns through the lattice, with some given initial coordinates in phase space, and record its phase space coordinates each time it passes the starting point in the ring.
2. Perform a discrete Fourier transform on the tracking data to determine the betatron tunes.
3. Repeat for a range of different initial coordinates in phase space.

The idea behind frequency map analysis is that plotting the tunes obtained from tracking data will indicate the tune shifts (by the area covered on the tune diagram) and the resonances that are strongly driven (by the clustering of points around particular resonance lines)...



A further refinement is to indicate the variation in tune over many periods of the ring by different colours. For example, if tracking over 1000 turns, one can take the difference between the tunes over the first 500 and the final 500 turns.

Generally, smooth or regular motion is characterised by a very slow change in tune over the trajectory of the particle.

A very rapid rate of change of tune can indicate irregular or chaotic motion. We often find that this occurs where several resonance lines cross (recall the Chirikov criterion).

The trick with frequency map analysis is to determine the tunes with high precision from a relatively small sample of tracking data.

For N turns, a simple discrete Fourier transform will yield the tunes with precision $1/N$. This is usually insufficient to derive meaningful data from tracking data collected in a realistic length of time.

An improved precision can be achieved by searching (numerically) for the frequency ω that maximises the function:

$$\sum_{n=1}^N x_n \exp\left(i\omega \frac{n}{N}\right). \quad (21)$$

This technique is equivalent to interpolating between points close to a peak in a discrete Fourier transform.

There are various techniques that can be used to improve the precision further. For example, a Hanning filter can be applied to the data before performing the Fourier transform; and then interpolation can be used to identify the “true” position of the peak in the spectrum, even if it does not happen to lie exactly on a sample point in tune space.

With care, it can be possible to determine the tunes with precision of order $1/N^4$ from N turns of tracking data. The original technique, developed by Laskar, is known as “numerical analysis of the fundamental frequencies” (NAFF)...

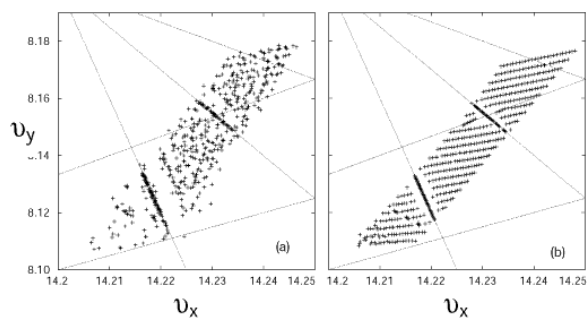
Some references:

- J. Laskar, “The chaotic behaviour of the solar system: a numerical estimate of the size of the chaotic zones,” *Icarus* 88 (1990) 266–291.
- J. Laskar, C. Froeschle, A. Celletti, “The measure of chaos by the numerical analysis of the fundamental frequencies,” *Physica D* 56 (1992) 253–269.
- J. Laskar, “Frequency analysis for multi-dimensional systems: global dynamics and diffusion,” *Physica D* 67 (1993), 257-281.
- H. S. Dumas, J. Laskar, “Global dynamics and long-time stability in Hamiltonian systems via numerical frequency analysis,” *Phys. Rev. Lett.* 70 (1993), 2975–2978.

Frequency map analysis can be a powerful tool for developing an insight into the dynamical behaviour of particles in an accelerator lattice. In particular, it can be useful for identifying resonances that may be limiting the dynamic aperture.

Unfortunately, it is not always easy to fix problems identified by FMA. For example, retuning the lattice to move the working point further away from a particular resonance may lead to an increase in the tune shifts with amplitude, with the result that there is a negligible improvement in the dynamic aperture.

Nevertheless, FMA is now routinely used as a design tool for new storage rings. It is even possible, with modern diagnostics, to collect high-precision turn-by-turn data, to allow the construction of experimental frequency maps, to compare with machine models.



The Advanced Light Source at LBNL was the first machine to construct an experimental frequency map.

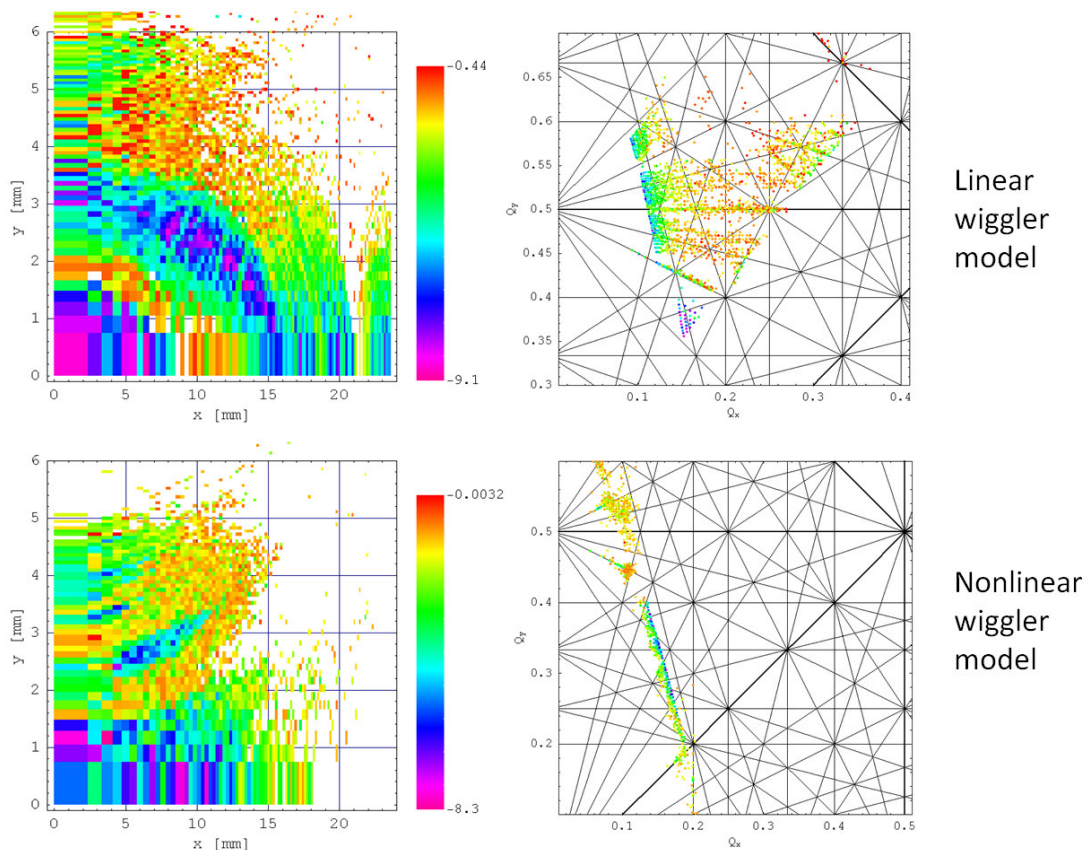
D. Robin, C. Steier, J. Laskar, L. Nadolski, "Global dynamics of the Advanced Light Source revealed through experimental frequency map analysis," *Phys. Rev. Lett.* 85 (2000), 558.

Throughout this course, we have ignored non-symplectic effects, such as radiation damping.

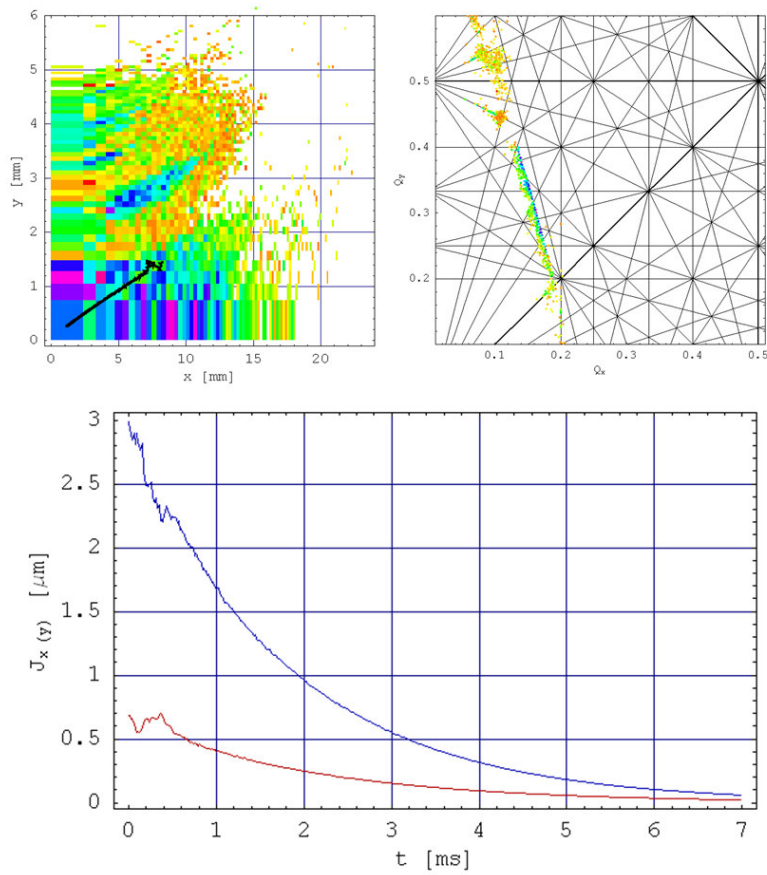
However, if damping is reasonably slow, then much of the analytical methods we have developed can still be applied.

In fact, under certain (rather special) conditions, damping can reveal the presence of nonlinear resonances in the beam distribution...

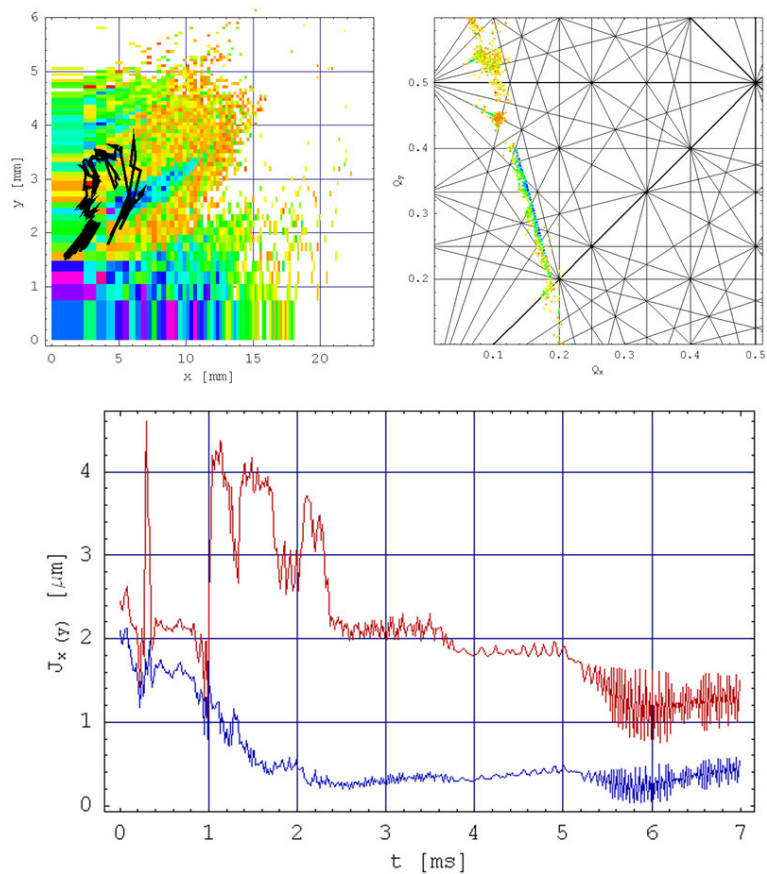
Nonlinear dynamics and damping: the NLC damping rings




Nonlinear dynamics and damping: the NLC damping rings



Nonlinear dynamics and damping: the NLC damping rings



Observation of a third-order resonance in the ALS



D. Robin, C. Steier, J. Safranek, W. Decking, "Enhanced performance of the ALS through periodicity restoration of the lattice," proceedings EPAC 2000.