Dynamical Systems, Representation of Particle Beams

Alex Chao, CERN Accelerator School 2018

In the book “One Two Three…. Infinity”, George Gamow asked a question:

“How high can you count?”

Imagine an audience reply:

“One!”

“Two!”

“Three!” was the most advanced answer.

“Any higher?”

After a long pause, the audience:

“We give up. Any number larger is confusing. Let’s call it Infinity!”

There were no numbers of significance between three and infinity.

Who was the audience?

A primitive people? Maybe.

But accelerator physicists are no different.

The two big areas in accelerator beam dynamics are

- single particle dynamics
- collective effects

The single particle dynamics is one-particle model. The number of particles is "One!"

Collective effects are treated assuming the beam is a continuum. The number of particles is "Infinity!"

We basically have nothing in between, while the actual beam has, e.g. 10^{12} particles. We have therefore the dilemma, how do we represent and analyze the actual dynamical system, and then trust the results? Is 10^{12} closer to 1 or closer to ∞?

Single- particle models ?	<======	Actual beam 10^{12} particles	======>	Continuum models ?

In the following, I shall mention a few jargons and then make a few comments.
--- not in a particular order or not even for a particular purpose.

Single-particle models ---- 1

Two-particle models ---- 2

Continuum models ---- ∞

Multiparticle models ---- 10^6 (computer)

Dynamic aperture

Nearly integrable systems

Microbunches

Single-particle models – (just a recap of what you already learned)

The beam is represented by “point particles”.

These are particles with no internal structure, no size, but yet they have mass and charge, even spin!

Three reasons why this concept can only be flawed:

- It is inconceivable that something has a mass and a spin but no size.
- It violates uncertainty principle of quantum mechanics.
- Even worse, it leads to fierce divergences when the particle also has a charge – the idea of “point charges” is even worse than the idea of “point particles”.

But we accelerator physicists sweep these problems under the rug (as we will do too).

Swallowing the concepts of point particles and point charges, then, the beam is represented by a collection (10^{12} of them) of single-particles. Since they do not interact among themselves, these models describe the beam faithfully simply by repeating the analysis 10^{12} times.

The single-particle models have been very successful, yielding deep knowledge as elucidated by the other lectures.

Where does the success come from?

The backbone of the success comes from the Liouville theorem.

To appreciate the intricacy of the single-particle models, one must first appreciate the phase space. The Liouville theorem states that all those intricate subtle dynamical effects in phase space, predicted by the single-particle models, are rigorously preserved, and lasting a surprisingly long time, i.e. forever. All predictions by single-particle models, including the most intricate effects down to the finest details, last forever.

Seemingly obvious, this fact keeps returning to surprise us, e.g. echoes and free electron lasers.

The combination of phase space + Liouville theorem is the basis of a very large number of accelerator applications:

Courant-Snyder dynamics
emittance preservation
RF gymnastics
phase space displacement acceleration
KAM theorem
6D phase space gymnastics
emittance exchanges
echoes
free electron lasers
various harmonic generation techniques
steady state microbunching techniques
etc.

This is an awesome list. We have indeed come a very long way with single-particle models.

Single-particle models, continued --- a note on history

In the 1960s and 70s, accelerator nonlinear dynamics was done by Hamiltonian dynamics and the canonical perturbation theories.

- Single-resonance analysis was great success (because it is integrabe)
- But they break down with multiple resonances and they diverge by small denominators
- Most advanced tool was COSY 5th order by brute force integrating the EOM
- Cumbersome to proceed, seemed stuck

In the 1980s, two breakthroughs revolutionized the landscape:

- Lie algebra by Alex Dragt
- TPSA by Martin Berz
- Blossomed at the SSC driven by the SSC need at SSC Central Design Group

Today, the canonical perturbation theories are the past.
The Lie framework + the TPSA technique become the powerful industry standard for single-particle dynamics.

- TPSA provides the most efficient maps.
- Assuming convergence (ignore chaos), Lie algebra extracts from TPSA to obtain analysis of
  one-turn effective Hamiltonian, normal forms.
- TPSA and Lie algebra are most powerful if and only if used together (arguably).

But this has been 30 years! Perhaps we should not call them a “modern approach”. Note the Hamiltonian canonical perturbative theories had lived only 20 years!

But have we replaced the old Hamiltonian dynamics completely? Are we done? The answer is no!

The very basic structure has remained, as evidenced by the stubborn Courant-Snyder language of the beta-functions.

The formalism based on the beta-functions, dispersion functions, $\mathcal{H}$-function, etc. is a remnant of the now extinct perturbative Hamiltonian dynamics and is inconsistent with the TPSA. The correct use of TPSA necessarily abandons the use of these special functions, which occupy half of every accelerator physics textbook!

An alternative framework developed in 1979 was called SLIM. It advocated no use of these special functions and is consistent with TPSA to the linear order. Use of SLIM and TPSA shall shorten the standard accelerator physics textbooks by half!

Two-particle models

Setting aside quantum mechanics, single-particle models are a tremendous success.

Biggest weakness occurs when particles interact electromagnetically with each other or with the vacuum chamber environment. The field a particles sees is then no longer prescribed by the external fields alone. Pursuing multi-particles along the line of single-particle models immediately become impossible as the number of particles is increased.

Some progress can be made if we have only two particles in the beam. The two particles interact with the environment or with each other. We thus arrive at the two-particle models.

When the two point-particles interact with each other, the analysis developed in single-particle models can still be applied but it becomes cumbersome. So far, we have only results in simplified models, nowhere near the sophistication of the single-particle models ---- no phase space manipulation, no KAM (at least not yet).

As oversimplified as they are, however, these two-particle models yielded important analytical insights towards the unraveling of collective effects. We shall leave out this discussion below.

Instead, let us mention another insightful consequence of the two-particle models, i.e. here we found two approaches to describe the beam dynamics.

1. We can consider the motion of $x_1(t)$ and $x_2(t)$ as two individual point particles evolving in time.
2. Or we can describe it by a superposition of two “modes”, a + mode in which the two particles move together ($x_1(t) + x_2(t)$), and a – mode with the two particles move oppositely to each other ($x_1(t) - x_2(t)$).

These are two representations of the beam dynamics, the “particle representation” and the “mode representation”.

The two representations are completely equivalent. They necessarily yield exactly the same final results.

The particle representation is also called a “time domain” approach. The mode representation is also called a “frequency domain” approach. Again, these two approaches necessarily give identical final results.

It is called time domain because in single particle representation, we focus on the time evolution of the two particles $x_1(t)$ and $x_2(t)$.

It is called frequency domain in mode representation because we focus on the evolution of the two eigen-modes.

The two representations are preferred by different people. Simulation programs might prefer time domain. Beam stability analysis might bias toward the frequency domain. Neither should claim advantage over the other.

We see that single-particle models are exclusively using time domain. As we shall see later, continuum analysis uses exclusively the frequency domain. The two prime areas of accelerator physics use two completely opposite beam representations. The two-particle models serve as an intermediate of these two pictures. They use both representations.

Continuum models

The opposite extreme to single-particle models is the continuum models. The beam is now represented by a continuum of distribution in 6D phase space (note: not in 3D real space!). All discreteness of particles are smoothed out. Beam evolution is determined by Vlasov equation.

Without “particles”, the analysis of beam motion is now described as a superposition of “modes”.

	Single-particle dynamics	**Collective effects**
Number of particles	1	∞
Beam representation	Collection of point particles	Continuum in phase space distribution
Dynamics approach	Time domain	Frequency domain
Analyses	x(t)	modes

Two-particle models
are intermediate.

Binary models

So we have tools developed for single-particle, two-particle, and continuum models. We let go of the attempt to represent 10^{12} particles.

But we still miss the important collision events. These are supplemented by the binary models.

Binary models are different from two-particle models.

Two-particle models: the whole beam has only two macro-particles, interacting with each other.

Binary models: the 10^{12} particles in the beam interact in two-by-two pairs.

In the continuum models, a drastic approximation is being made. In the Vlasov approach, we ignore collisions among point particles, i.e. we include the averaged smooth collective fields of the continuum but we ignore interaction between individual particles.

This approximation might work for the collective effects due to wakefields. It is however a drastic over-simplification for space charge effects recognizing that Coulomb fields diverges between individual point particles and therefore has to be extremely singular and granular, while the sum of fields from a continuum beam is smooth.

Can we really replace 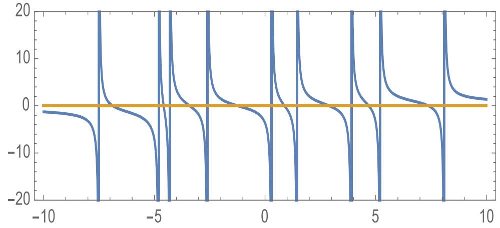

by 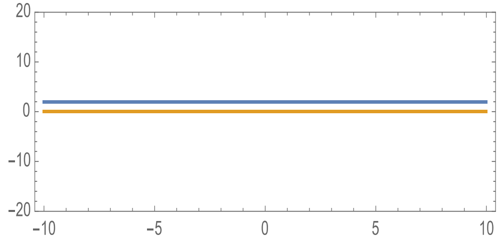

??

The way we try to deal with the problem, partially, is to consider a smooth collective effect, but then supplement it by a few binary models such as intrabeam and Touschek modifications.

One might ask: Is that sufficient?

Beam-beam models

The situation of smoothing out a granular force by a smooth force gets worse with the beam-beam effect.

Which picture should be used to model the on-coming beam? Have we missed something in making this approximation?

Is the beam-beam limit determined by the smooth kick (as we have been doing all these years)? Or part of it is actually due to the granularity of the on-coming beam distribution? Have we left out some important noise-like diffusion effect?

For single beams, we still have the supplemental Tousheck and intrabeam binary models. We do not have a corresponding beam-beam binary models.

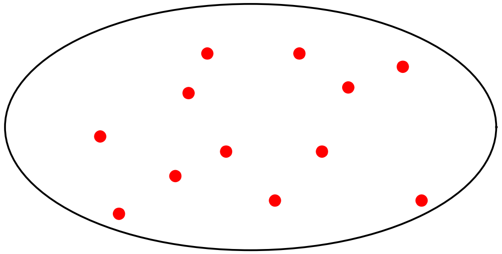

the actual beam

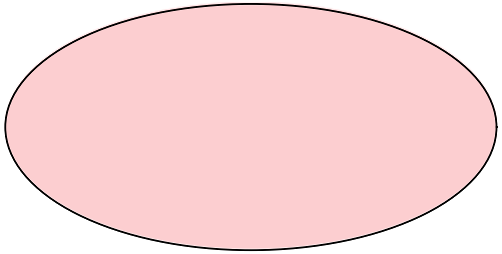

the model we use

10^6-particle models

Back to the single-particle models. Obvious next step is to add more particles in the beam representation. However, when there are 3 or more particles, the analysis becomes cumbersome even for simplified models. We resort to computer simulations.

One then study collective effects by multi-particle simulations. As in the two-particle models, we have two ways to proceed:

- represent the beam by a collection of single point-particles who interact with wake fields
- divide the otherwise continuum phase space into grids, each grid represented by a "macro-particle".

Both are time domain approaches (but they are not quite the same concepts). One simulates the motion of a particle, the other describes evolution of an element in phase space.

Choosing a finite number of particles in a time-domain computer simulation is then equivalent to truncation to the highest mode number in a frequency-domain analytical calculation.

In treating multiparticle collective effects, time domain applications use “wakefields”. Frequency domain uses “impedances”. Wakefields are more used in simulations. Impedances are more used in analyses:

Simulations	Analyses
Time domain	Frequency domain
Wakefields	Impedances

These multi-particle models typically go up to 10^6 particles.

10^6 calculations per step for collective single-particle models.

10^6 (10^6 -1)/2 calculations per step for intrabeam binary models.

Should we try to simulate 10^{12} particles?

Assuming computer power is available, question is what for? Why bother?

- Single-particle models allow exploration of detailed phase space without collective effects
- Two-particle models give the qualitative understanding of collective effects
- 10^6 computer simulations give accurate information on lower order modes and instability thresholds
- What is there to learn from a 10^{12} particle simulation?

Dynamic aperture

This is one old problem intrinsically difficult because it is not integrable.

All soluble cases ASSUME a priori integrability, e.g.

single isolated resonance

convergence in power series perturbative expansions

All "predictions" of dynamic aperture are based on assumed integrability.

The problem of convergence is intrinsic. All perturbation theories must assume convergence. It can not be avoided simply by switching tools from canonical perturbation theories to TPSA, Lie algebra, or TPSA+LieAlgebra. The intrinsic problem is still there as much now as ever before!

So, does TPSA converge? How fast do they converge (if any)? To what order can we truncate TPSA? There might never be an answer to this question. Our approach can only be: let us assume the convergence and push to the limit to see how/when it breaks down!

For those cases, simulations are the only dependable tool --- Or are they?

For a workable accelerator, we need a sizable region of stability in phase space to place the beam. KAM theorem, however, does not give any sizable stability region for the beam at all. Mathematicians have been astounded to know we require such a large region of stability. And yet storage rings do work!

The catch lies in the fact that the mathematicians, the KAM, addresses stability for infinite number of turns. But we ask merely for stability of 10^{10} turns. These are completely different issues!

Note that the earth has evolved around the sun only for a few 10^9 turns. KAM would set grave doubt on the stability of the earth.

Extensive efforts were made for example in the dynamic aperture studies for the Superconducting Super Collider in the 1990s. One of the results was shown in a "survival plot". The particle nearest the "dynamic aperture" was found to stay in the storage ring for one million turns happily but got lost in the very last 30 turns.

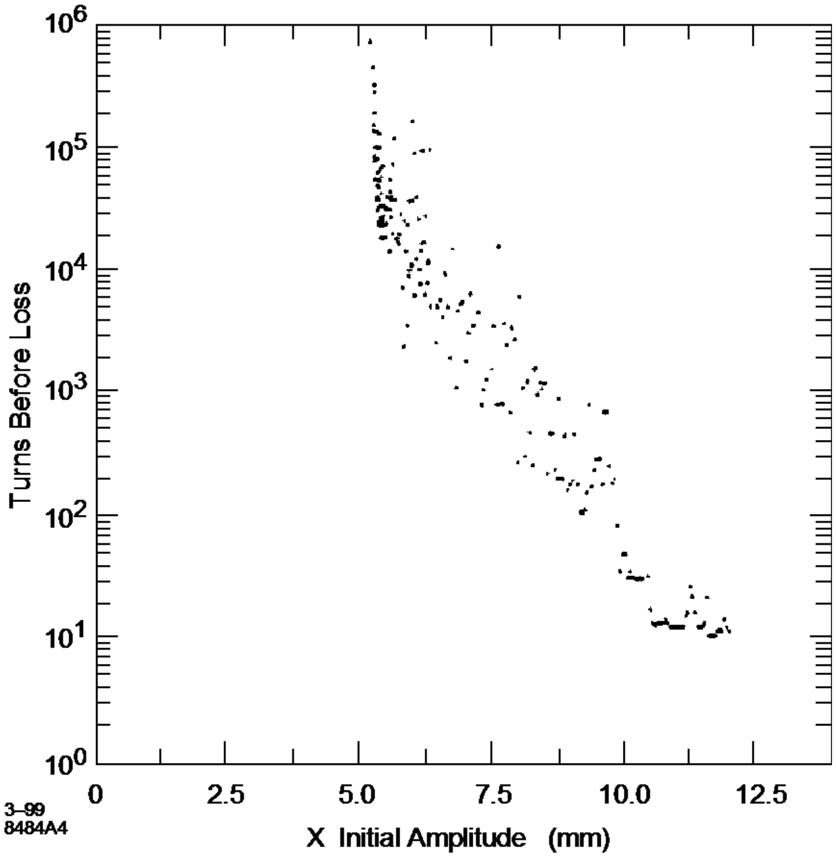

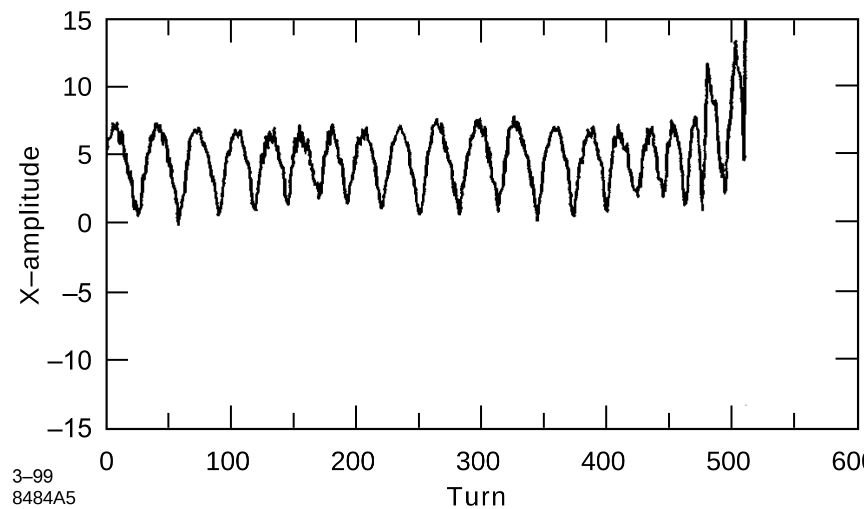

What do we learn from this?

- particle loss mechanism is very subtle
- dynamic aperture prediction is difficult
- there is no "early alarm" of instability.
  (e.g. you can't predict 10^6 stability from 10^3 tracking.)

Nearly integrable systems

Accelerator design necessarily starts initially with an integrable system as its "design baseline". So far, this initial integrable design baseline has been chosen almost exclusively to be the linear, uncoupled Courant-Snyder system.

Our job has been to maximize the stability region when various perturbations are added. The initial system must be chosen to be robust against all envisioned perturbations. For example, choice of working point must avoid lower-order resonances, etc. All nonlinearities are considered perturbations $\sim \varepsilon$, and chaos occurs as soon as integrability is broken, endangering the dynamic aperture.

Is there a way to find another initial integrable system that is more robust than the linear uncoupled Courant-Snyder system? This is being tested at Fermilab by the IOTA project. The idea considers the initial system to be nonlinear, in fact very nonlinear with a large "detuning". Such a system is believed to be able to tolerate larger perturbations.

* Conventional accelerators:

  Both detuning (stability mechanism) and chaos (instability mechanism) originate from nonlinear perturbations. ➔ Detuning ~ Chaos ~ ε

* IOTA:

  Detuning originates from zero-th order design ➔ Detuning >> Chaos

If proven correct, it says what we have been doing in the past 70 years is unwise!

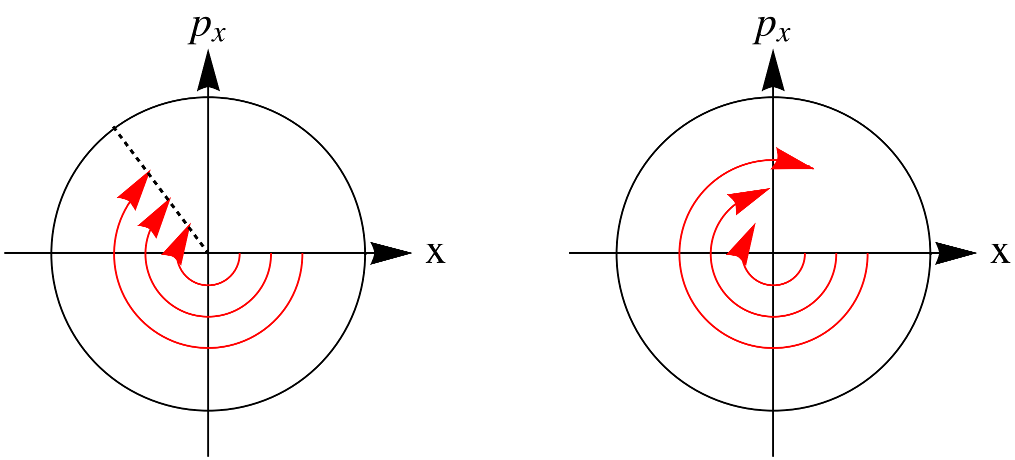

No detuning
All particles move in unisom
susceptible to resonances

Large detuning
Shearing motion
desensitizes resonances

Note that it is not that such systems are still integrable after adding perturbations. It is just that the region of instability can be much reduced.

Microbunches

One particular area of concern is microbunches. Microbunches are a new development, particularly due to the advent of the free electron lasers. Microbunch beam dynamics will be a focus of accelerator physics for many years to come.

Survival of microbunches is a vivid manifestation of the Liouville theorem. Phase space conservation works all the way down to sub-nanometer levels!

But microbunching is challenging beam dynamics! The challenge is due to the DYNAMIC RANGE of the physics involved.

- the electron bunch length ~ 1mm
- the microstructures within the electron bunch ~ 0.1nm
- Dynamic range of 10^7

There is no way to simulate the beam dynamics as we did for single-particle models. The dynamic range is too big.

To address the physics of microbunches, we invented yet another technique, i.e. we only simulate the time evolution of one particular frequency component $\lambda_0 = \lambda_u (1 + K^2)/ 2 \gamma^2$, of the beam distribution and EM field distribution.

In particular, this one-frequency model of FEL physics does not contain information at any wavelength < λ_0. Results on the behavior of beam distribution becomes questionable when the beam bunch is shorter than λ_0. Currently there are no simulation codes to deal with such cases.

So we now have a very peculiar beam dynamics model that is a mixture of time-domain and frequency-domain. Such a mixture has always been dangerous and error-prone and is mostly forbidden in the past.

The GENESIS code, a mixture code, for example, for FEL simulations must be used with extreme care.

Incidentally, frequency filtering is a key ingredient of FEL physics not only in simulations. The analysis of FEL physics does the same, as evidenced by the famous cubic FEL equation or the SASE mechanism.

Summary 1 The present landscape

The present landscape might look like this:

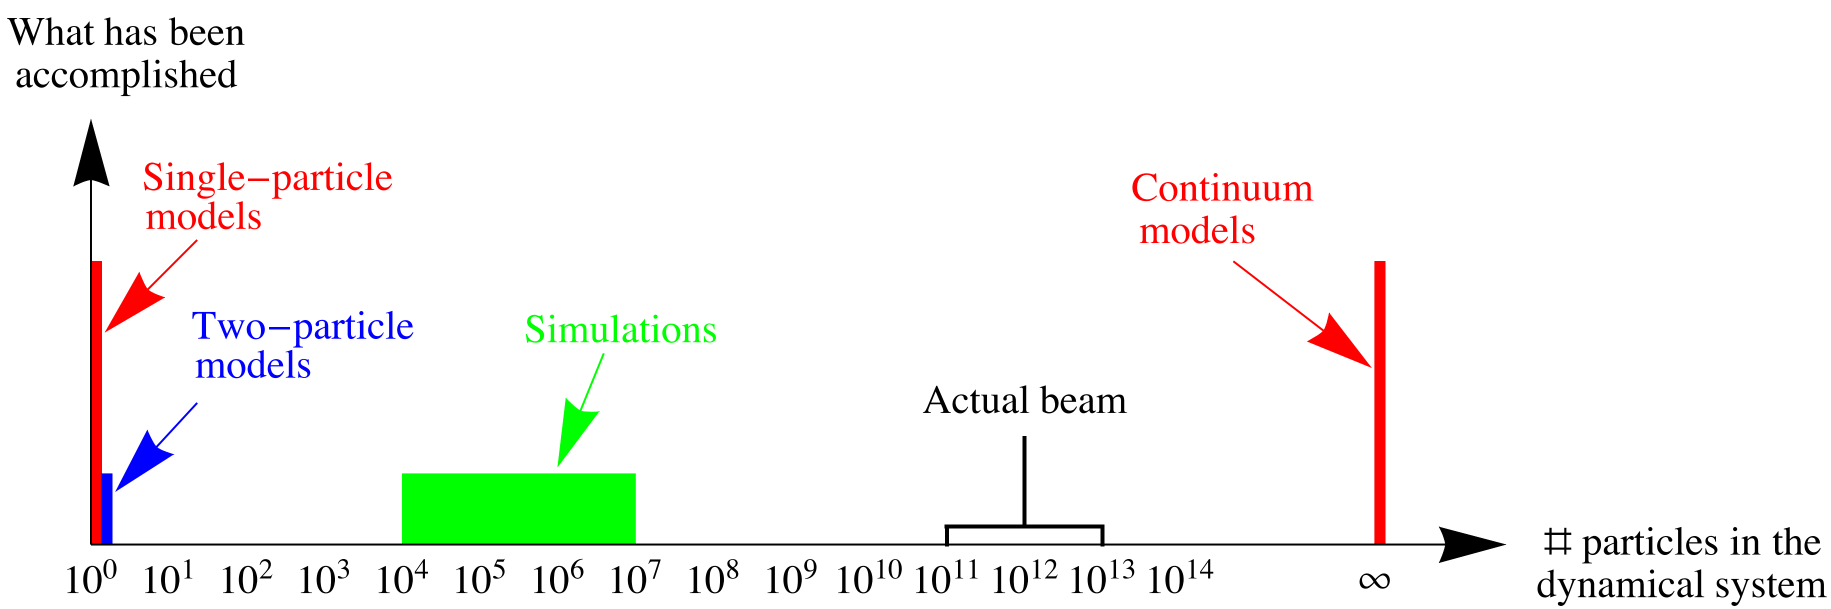

Our approach so far has been to represent the actual beam as

- A collection of non-interacting point particles ---- 1
- Two interacting particles + binary models ---- 2
- One million interacting particles + simulation ---- 10^6 (or 3)
- A continuum of phase space distribution ---- ∞

SUMMARY 2 Possible future

TPSA + Lie algebra =>

~~Canonical perturbation theories~~

SLIM =>

~~Courant–Snyder functions $\alpha(s), \beta(s), \gamma(s), \eta(s), \mathcal{H}(s)$~~

Advent of microbunches
$\Rightarrow$ new (dangerous) territory Mixture of frequency and time domains

Mixing time– and frequency– domains

IOTA
$\Rightarrow$ Change the integrable baseline from Courant-Snyder systems to heavily detuned systems

Nonlinear detuned systems as baseline design