Why is the Verlet algorithm so useful for molecular dynamics simulations?

Hans C. Andersen, Stanford University

May 8, 2016

The velocity Verlet algorithm

The usual Cartesian Hamiltonian

$$H(\mathbf{r},\mathbf{p}) = \mathbf{p}^2/2m + U(\mathbf{r})$$

Hamilton's equations in terms of positions, velocities, and accelerations

$$\dot{\mathbf{r}}(t) = \mathbf{v}(t)$$
$$\dot{\mathbf{v}}(t) = \mathbf{a}(t) = b\left(\mathbf{r}(t)\right) \equiv -\nabla_{\mathbf{r}} U(\mathbf{r}(t))/m$$

$$\dot{\mathbf{r}}(t) = \mathbf{v}(t)$$
$$\dot{\mathbf{v}}(t) = \mathbf{a}(t) = b\left(\mathbf{r}(t)\right) \equiv -\nabla_{\mathbf{r}} U(\mathbf{r}(t))/m$$

Verlet algorithm (the velocity Verlet version):

Given $\mathbf{r}(t)$, $\mathbf{v}(t)$, $\mathbf{a}(t)$, calculate

Error

$$\mathbf{r}(t+h) = \mathbf{r}(t) + \mathbf{v}(t)h + (1/2)\mathbf{a}(t)h^2 \qquad O(h^3)$$
$$\mathbf{a}(t+h) = b\left(\mathbf{r}(t+h)\right) \qquad 0$$
$$\mathbf{v}(t+h) = \mathbf{v}(t) + (1/2)\left[\mathbf{a}(t) + \mathbf{a}(t+h)\right]h \qquad O(h^3)$$

Now we have $\mathbf{r}(t+h)$, $\mathbf{v}(t+h)$, $\mathbf{a}(t+h)$.

Then iterate.

Algorithmic error of the Verlet algorithm

Local error. The local error is $O(h^3)$ for small h. The new positions and momenta after one time step have errors that are $O(h^3)$ for small h.

Global error. The global error is $O(h^2)$ for small enough h. If the trajectory has a total duration of T, the final positions and momenta have an error of $O(h^2)$ for small enough h.

If $U(\mathbf{r})$ has continuous second partial derivatives,

these error estimates are correct.

If the potential energy function has significant discontinuities in the second derivative,

the accuracy of the Verlet algorithm is degraded significantly.

- The Verlet algorithm is not exact for any nonzero time step, no matter how small. It makes errors on each time step. The errors in later steps amplify errors made in earlier steps because of the chaotic nature of the dynamics.

- The usual statement of global error is correct, but misleading for very long trajectories and time steps typically used.

- For the typical long runs that are routine these days and for typical system sizes used, there is little chance that the trajectory calculated using the algorithm is anywhere near the actual trajectory at long times.

- Should we be concerned about this?

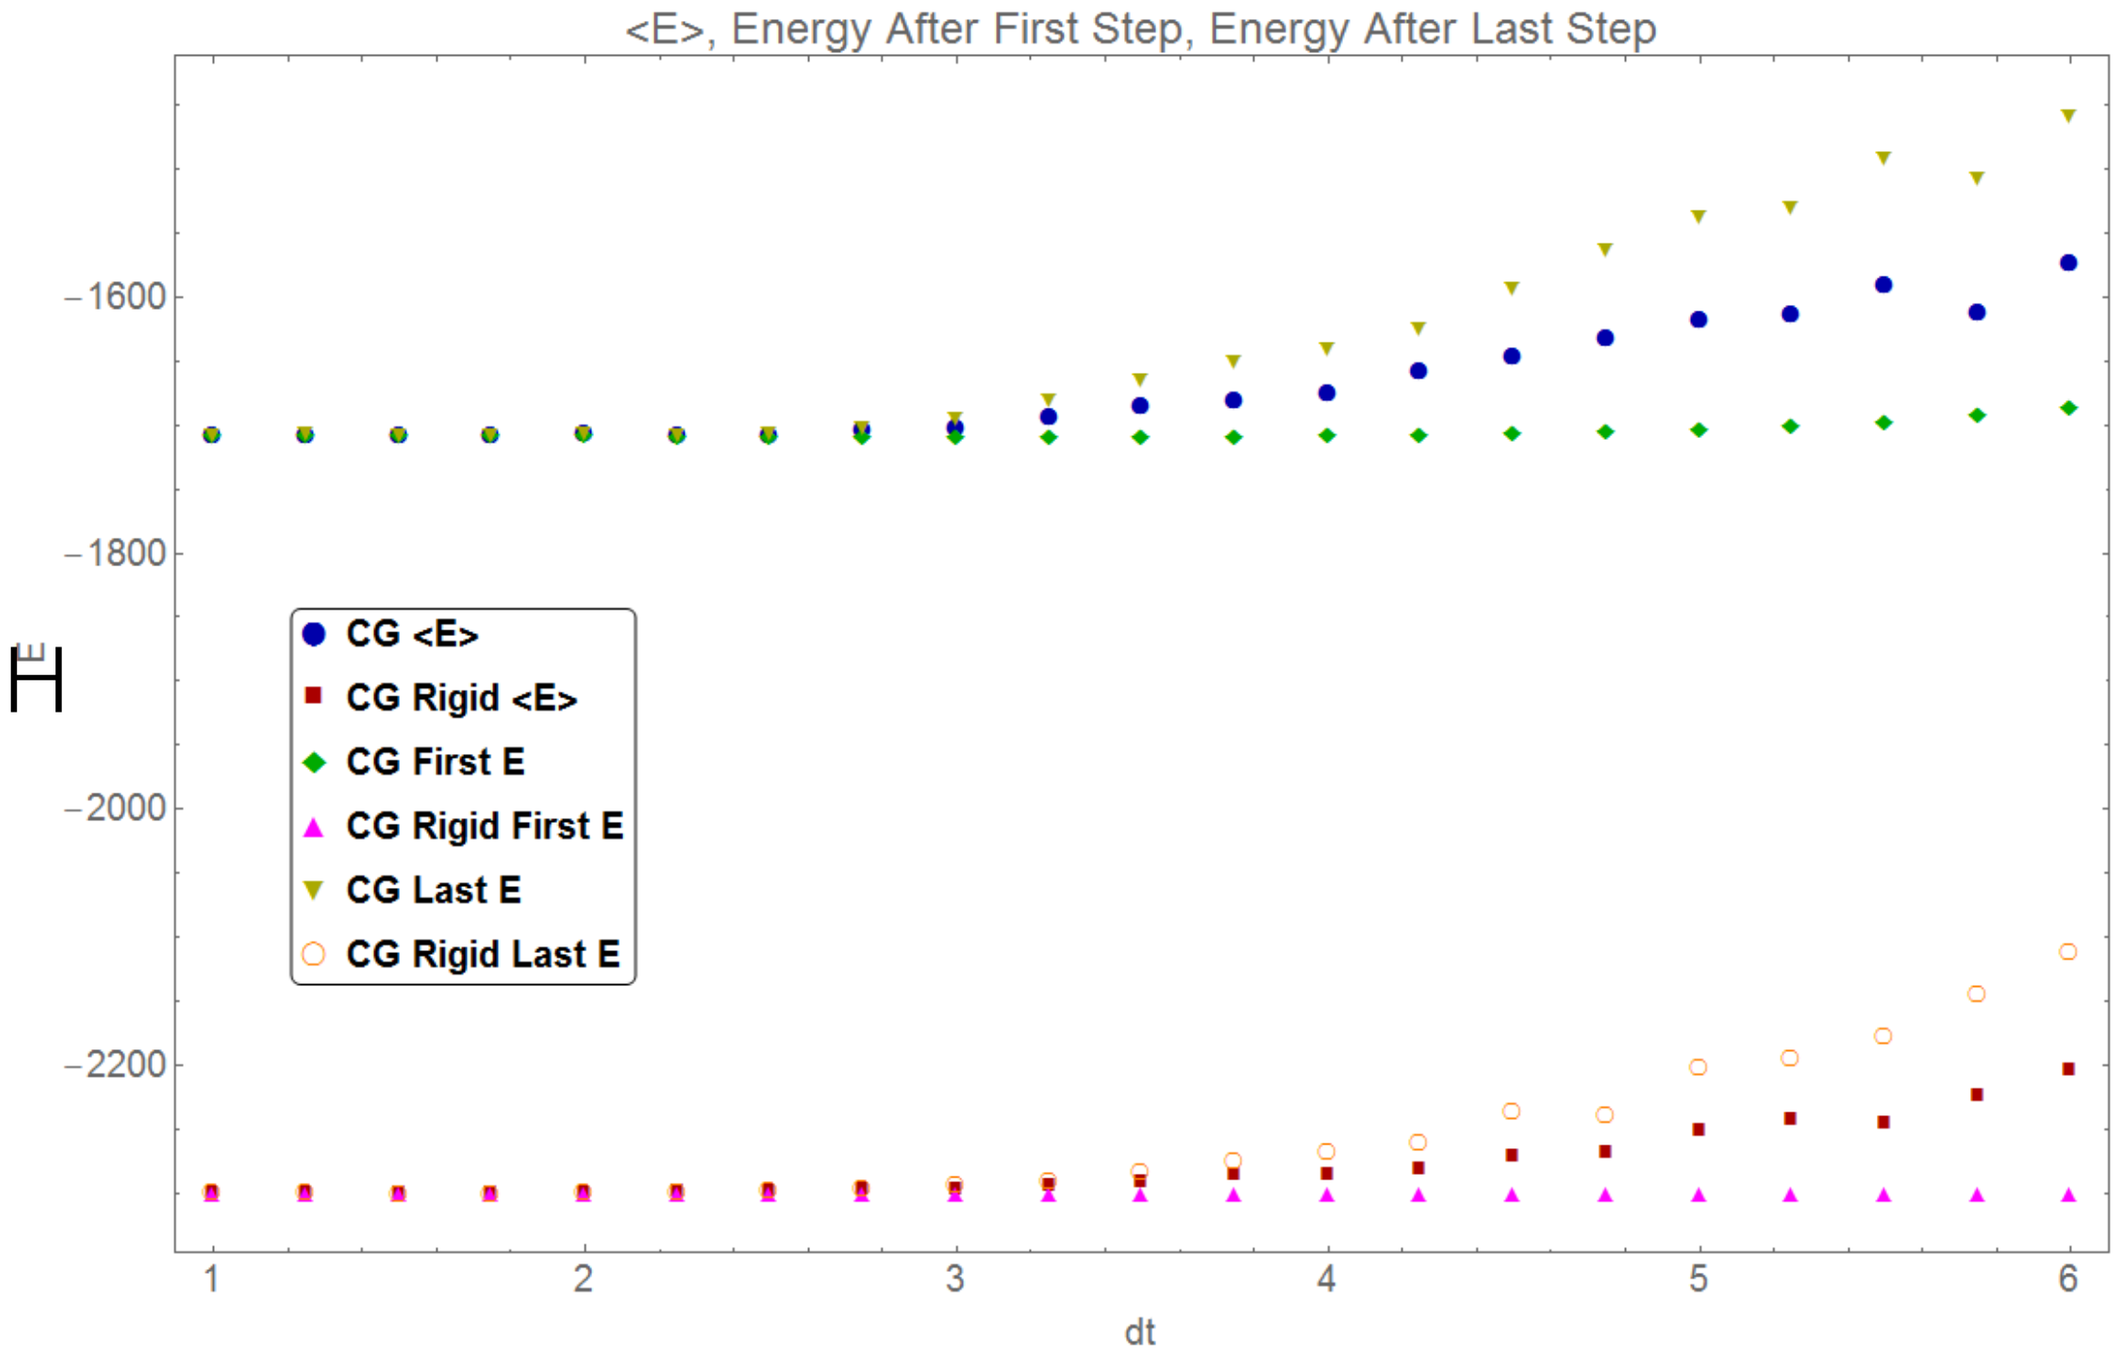
<E>, Energy After First Step, Energy After Last Step
−1600
−1800
−2000
−2200
H
CG <E>
CG Rigid <E>
CG First E
CG Rigid First E
CG Last E
CG Rigid Last E
1
2
3
4
5
6
dt

time step h

rms H fluctuation

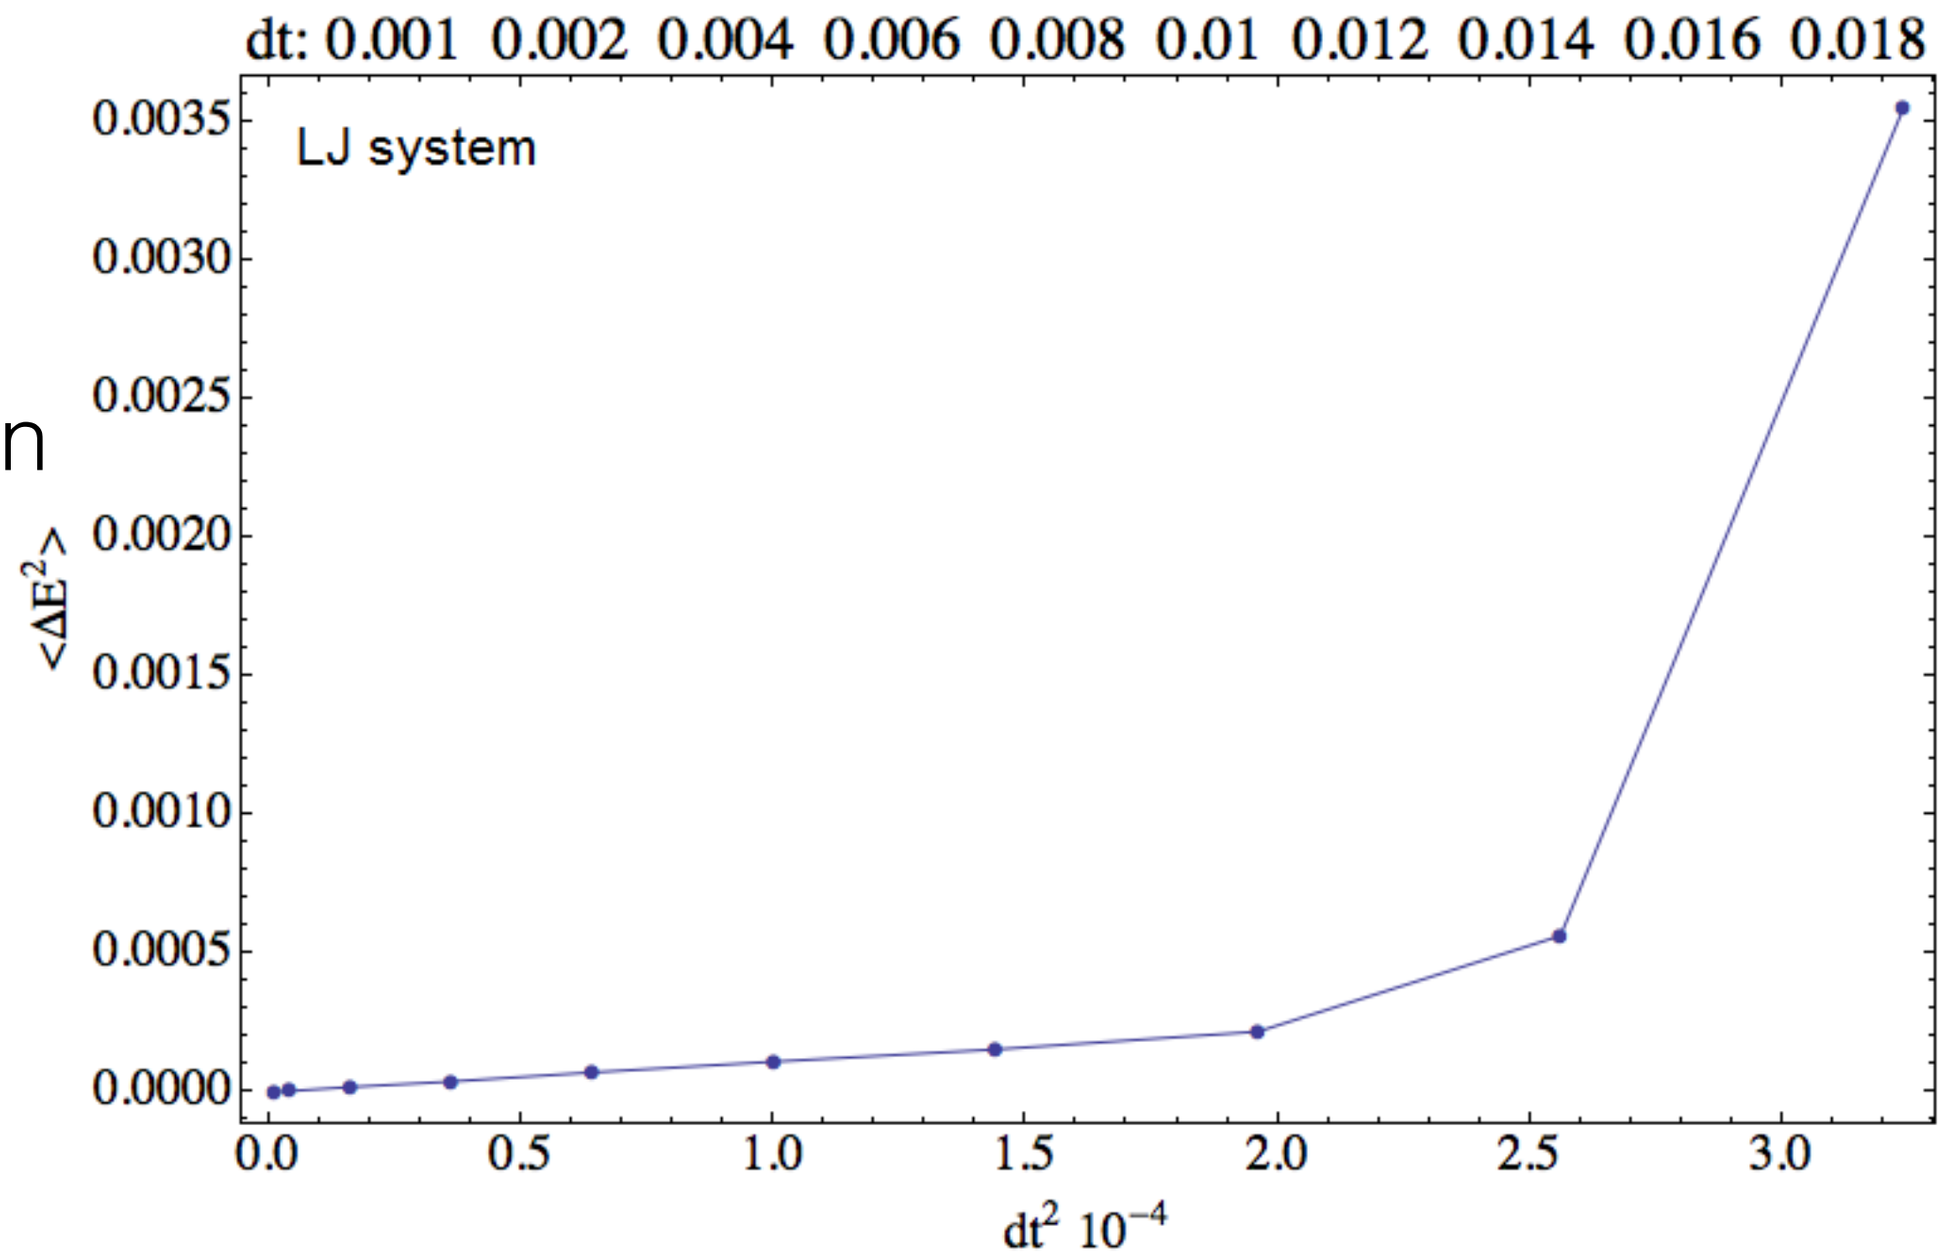

(time step h)^2

rms H fluctuation

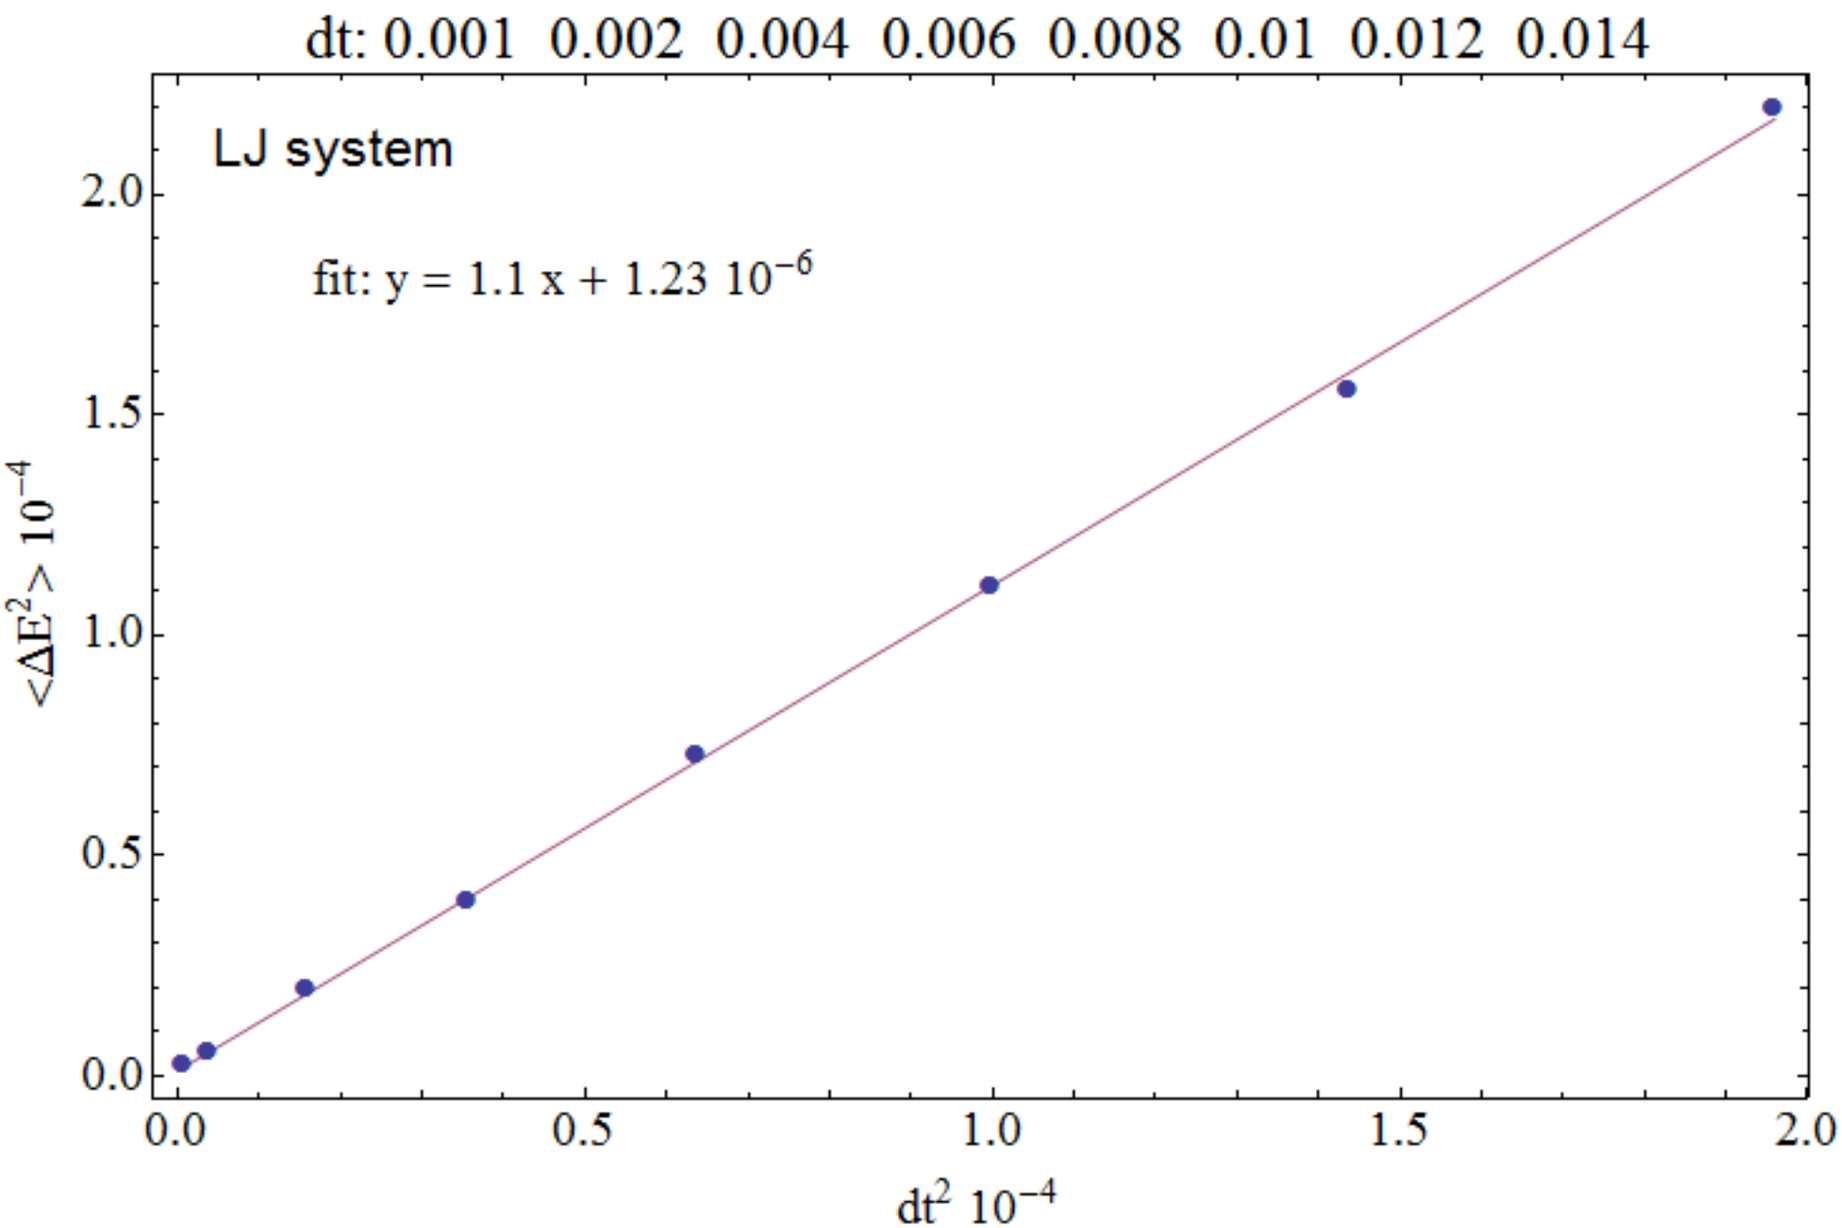

(time step h)^2

rms H fluctuation

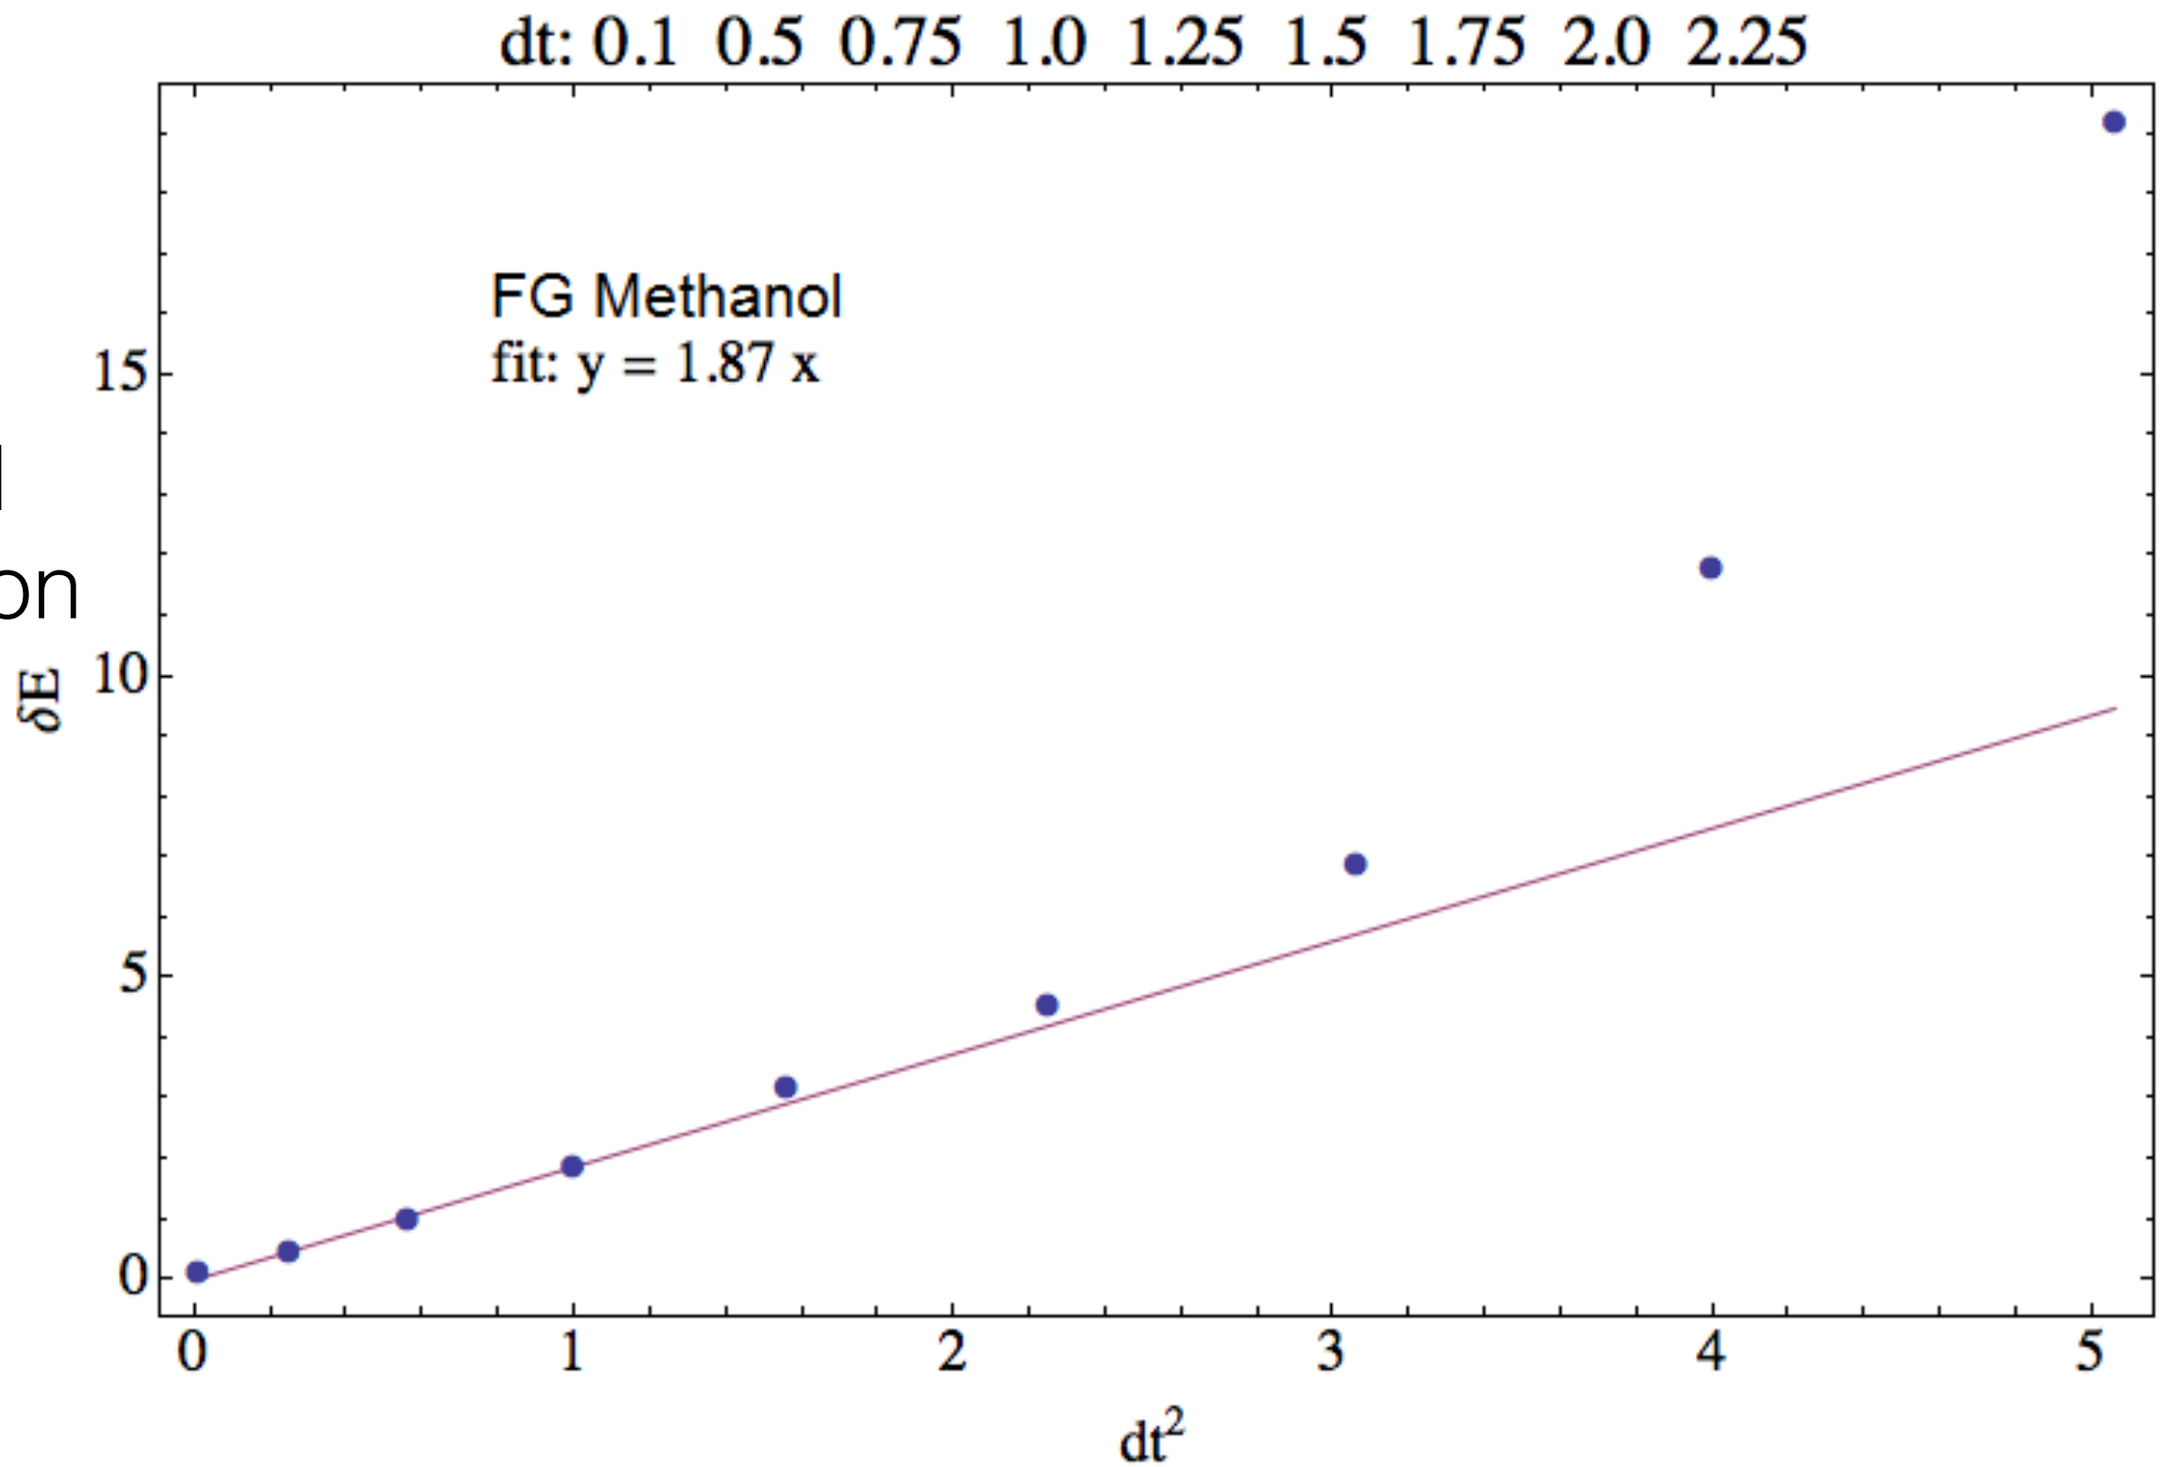

(time step h)^2

Strategies for sleeping at night

- Ignorance
- Denial and rationalization
- Learn about the shadow Hamiltonian

The Verlet algorithm and the shadow Hamiltonian

Assertion. The Verlet algorithm is symplectic. Under certain circumstances the following statement is correct.

The discrete phase points along a trajectory calculated using

- a specific initial state,
- a nonzero time step h,
- the Verlet algorithm,
- a Hamiltonian $H(\mathbf{r}, \mathbf{p})$

lie on the exact continuous trajectory that

- has the same initial state, and
- is generated by a "shadow" Hamiltonian, H_S, where

$$H_S = H + \Delta H \quad \text{and} \quad \Delta H = O\left(h^2\right)$$

In what sense is this statement correct?

- *If the potential energy function in H is analytic,* an asymptotic expansion for H_S exists in powers of h.
- *If the potential energy function in H does not have two continuous partial derivatives,* the statement is not correct.
- Depending on the number of continuous derivatives that the potential has and the magnitude of the discontinuities,

the statement is correct in the sense that accurate *approximations* for such a shadow Hamiltonian can be constructed to various orders in h.

In practice, to get the full accuracy of the Verlet algorithm it is important to use a potential energy function that has two continuous derivatives (or at worst, only very small discontinuities in the second derivative). If this is done, then the statement is correct..

This explains some of our observations

- the rms energy fluctuation for a long run is linear in h^2 for very long trajectories (a debugging tool that my students have used for over thirty years).

- energy is not necessarily conserved in freezing of a LJ supercooled liquid.

- our experience in doing molecular dynamics simulations of melting in 2 dimensions.

The existence of the shadow Hamiltonian has implications for

I. Understanding of the accuracy of dynamical simulation results

II. Debugging of programs that use the Verlet algorithm

III. Choosing the time step for the Verlet algorithm for equilibrium simulations

IV. Extrapolating to the limit of zero time step

I. Understanding the accuracy of dynamical simulation results

When one uses the Verlet algorithm with a Hamiltonian whose potential energy function has **two continuous partial derivatives** the result is a trajectory that:

- is not an accurate trajectory for the Hamiltonian of interest,
- is the **exact** trajectory for a Hamiltonian that is **very close** to the Hamiltonian of interest.

Is the thing we are calculating from a trajectory sensitive to small changes in the Hamiltonian?

1. **The end point of the trajectory** is probably something that is very sensitive.

Don't expect to get a good answer for the end point or for intermediate points.

2. **The average of some phase function over a very long trajectory** (e.g. a radial distribution function) is equal to a microcanonical equilibrium average of that phase function. In general, equilibrium properties are not sensitive to small changes in the Hamiltonian except near critical points.

Expect to get good results for the the average of a phase function over a very long trajectory. This includes both static quantities and time correlation functions.

3. Other properties? **Specific analysis is required.**

II. Debugging of programs that use the Verlet algorithm

Consider the rms fluctuation of the actual Hamiltonian of interest from its average along a very long trajectory calculated with the Verlet algorithm.

The result is equal to the rms fluctuation of the actual Hamiltonian of interest from its average in a microcanonical ensemble for the shadow Hamiltonian. The latter is $O(h^2)$.

Hence a plot of the rms fluctuation of the actual Hamiltonian from its average as a function of h^2 should be a straight line through the origin, for small h.

rms H fluctuation

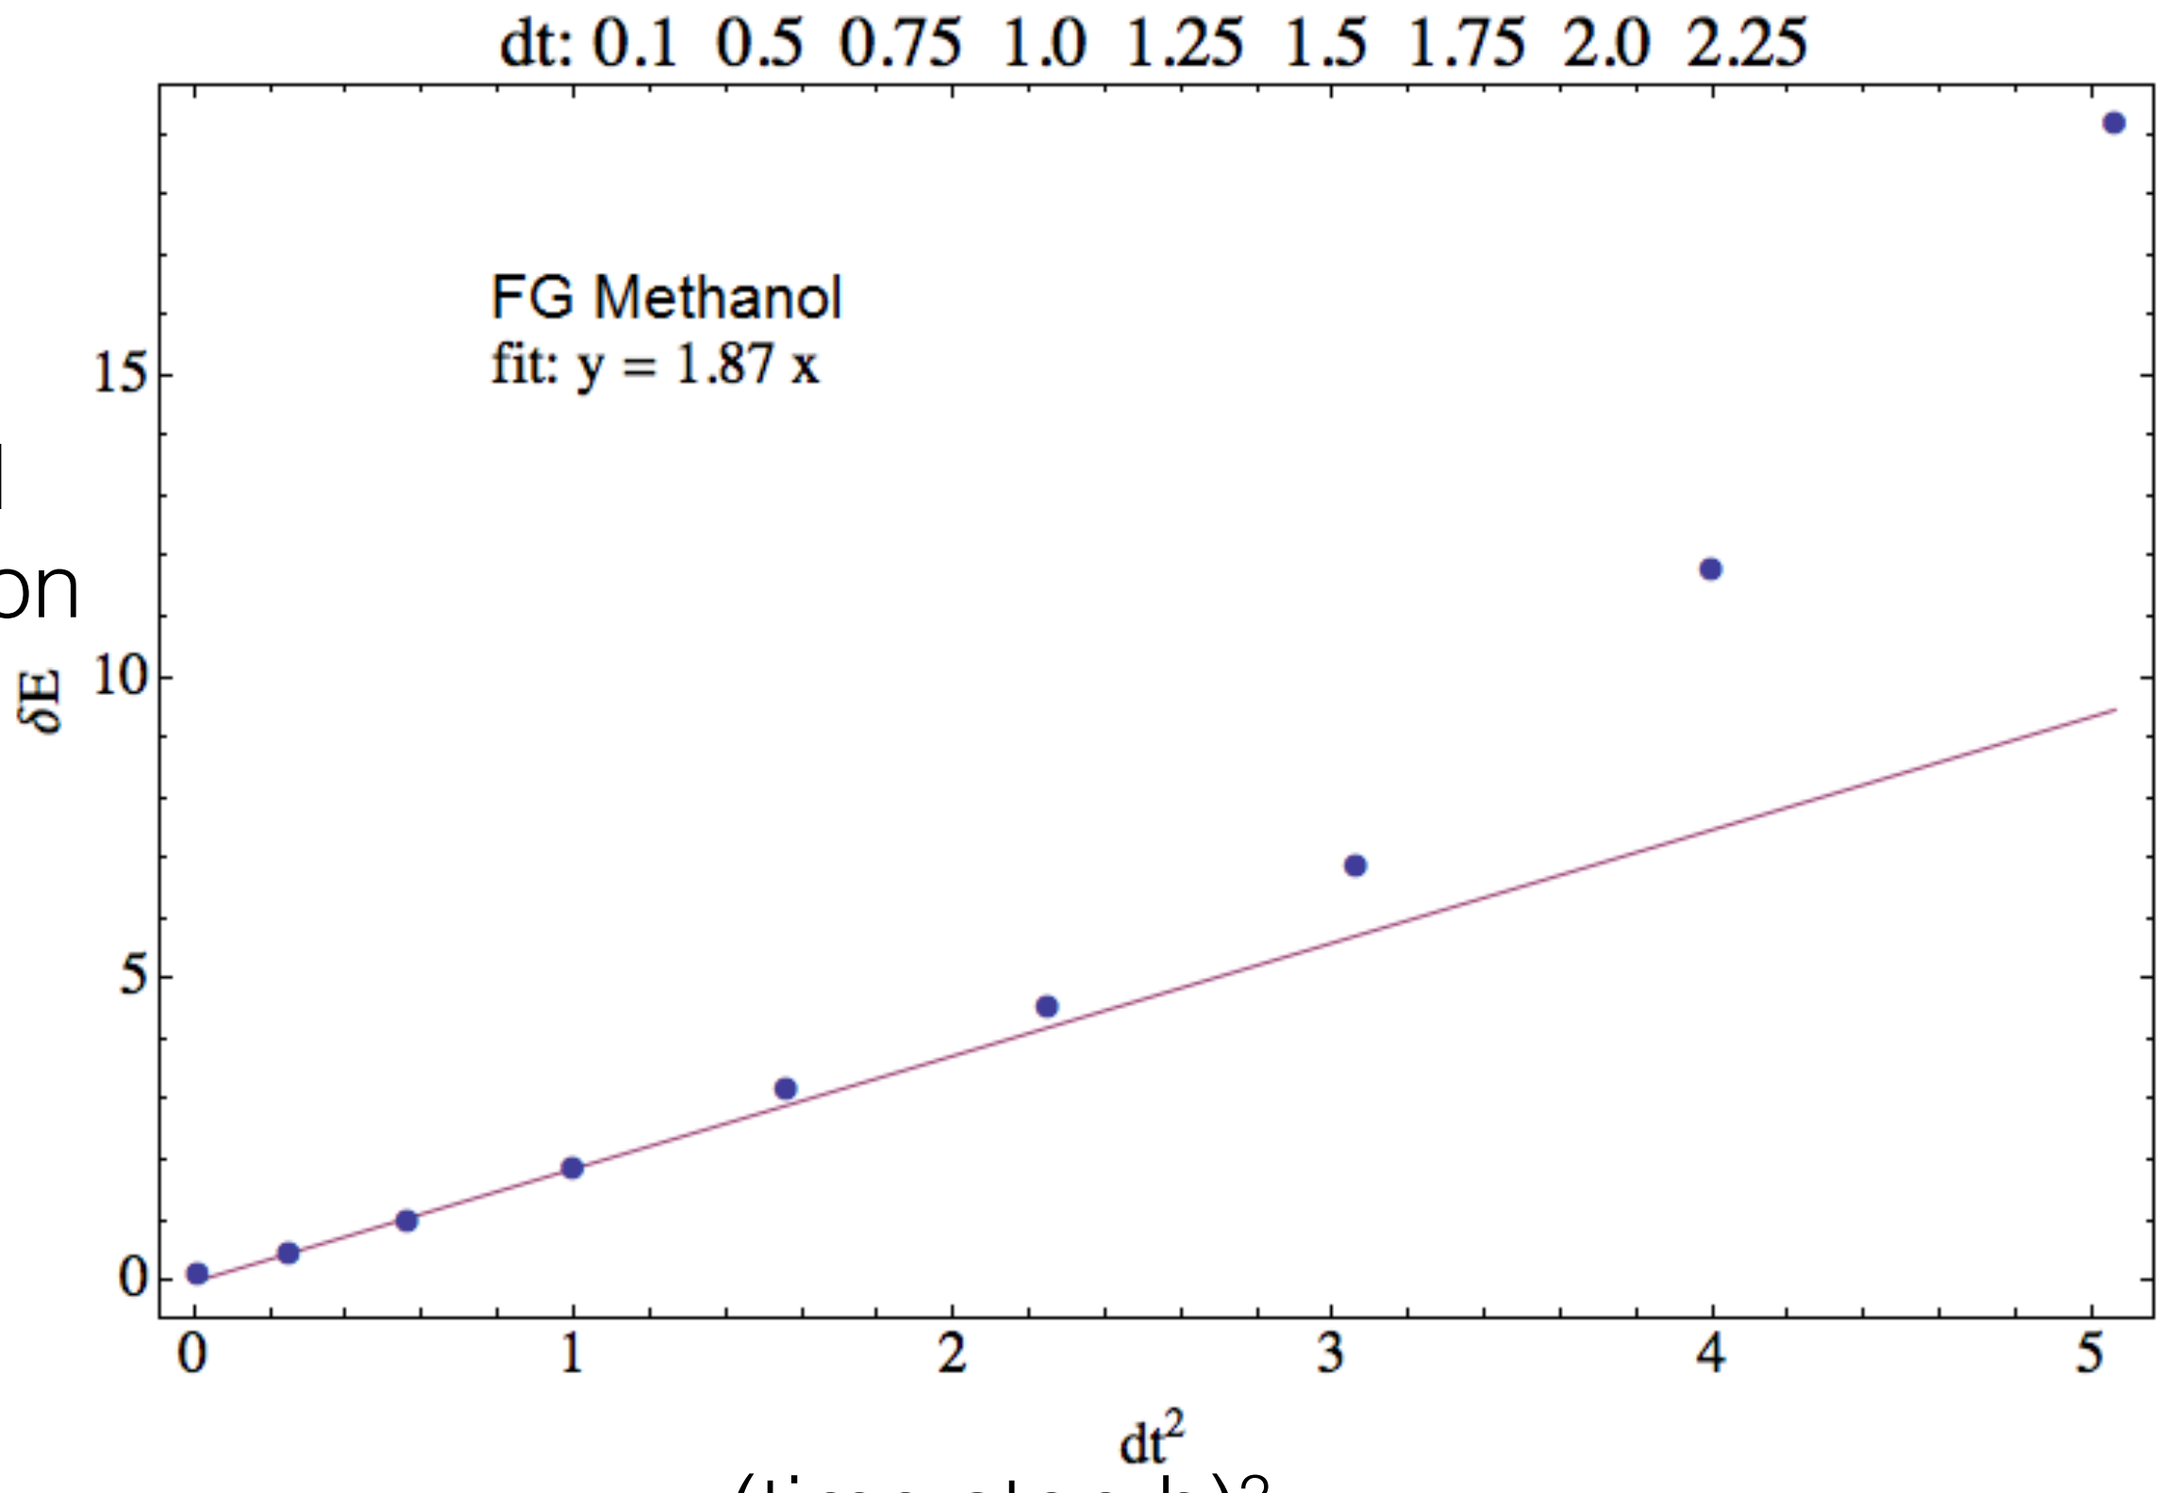

(time step h)^2

A convincing straight line through the origin is a strong confirmation that:

- the forces used to calculate dynamics are the negative gradients of the potential function in the Hamiltonian used to calculate the energies
- the potential function has two continuous derivatives

III. Choosing the time step for the Verlet algorithm for equilibrium simulations

Choose a time step that is in the linear regime of the previous curves for a trajectory that is of the length that you will be performing. This will give you the stability that you need in the trajectory and avoid making a huge error.

Depending on how much accuracy you need and your computational resources, pick somethng at the middle or low end of the linear regime, rather than the high end.

IV. Extrapolating to the limit of zero time step

For any property of a very long trajectory that is to be calculated,

if two or more simulations are performed using different time steps, all in the linear regime, and the results are plotted vs. h^2,

a linear extrapolation of the result to zero time step can be performed to estimate the value of the property for an exact trajectory as determined by the Hamiltonian of interest.

Why the Verlet algorithm is so useful

- Although the trajectories are not accurate for the Hamiltonian of interest, they are exact (except for numerical precision error) for the shadow Hamiltonian. This can be used to decide how accurate the results can be expected to be for the Hamiltonian of interest.

- Plotting the rms energy fluctuation as a function of h^2 for a long time trajectory provides

  a very useful way of spotting programming errors and fixing them;

  a rational basis for choosing the time step to use.

- Plotting phase function averages calculated from long trajectories with two or more different time steps vs. h^2 and linearly extrapolating to $h^2 = 0$ provides a valid way of estimating the value of the property for an exact trajectory of the same length starting from the same initial state.

My thanks to:

- Bill Swope
- Tariq Andrea
- Marcia Grabow
- Ken Bagchi
- Aram Davtyan
- Avisek Das
- and many others

END of talk

Slides not used
save for reference

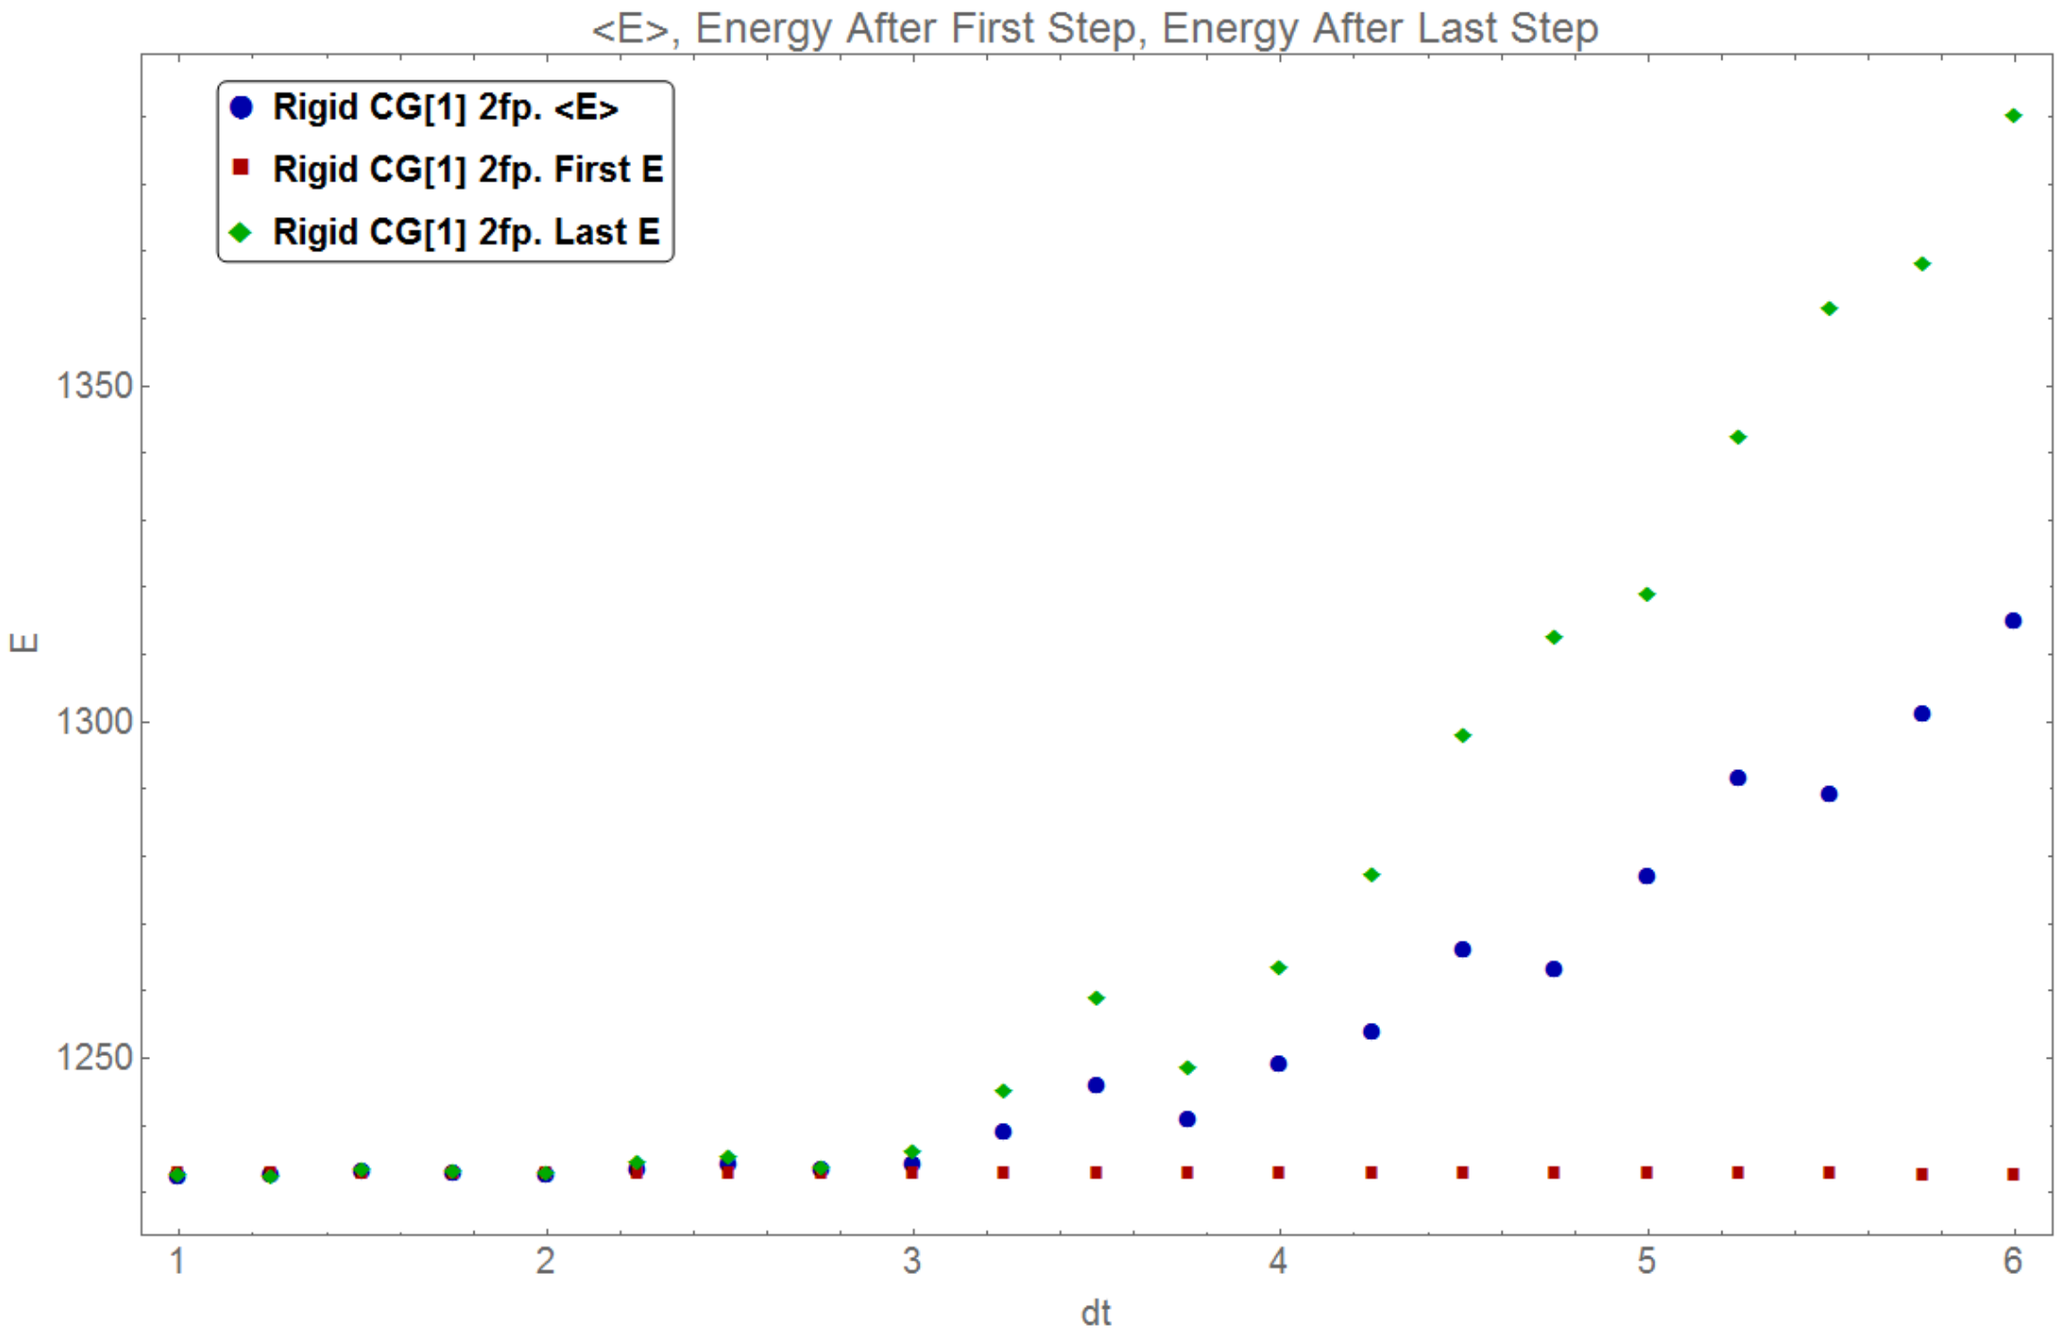
<E>, Energy After First Step, Energy After Last Step
Rigid CG[1] 2fp. <E>
Rigid CG[1] 2fp. First E
Rigid CG[1] 2fp. Last E
E
1350
1300
1250
1
2
3
4
5
6
dt

<E>, Energy After First Step, Energy After Last Step

Rigid CG[1] 4fp. <E>
Rigid CG[1] 4fp. First E
Rigid CG[1] 4fp. Last E

E

5000
4950
4900

1
2
3
4
5
6

dt

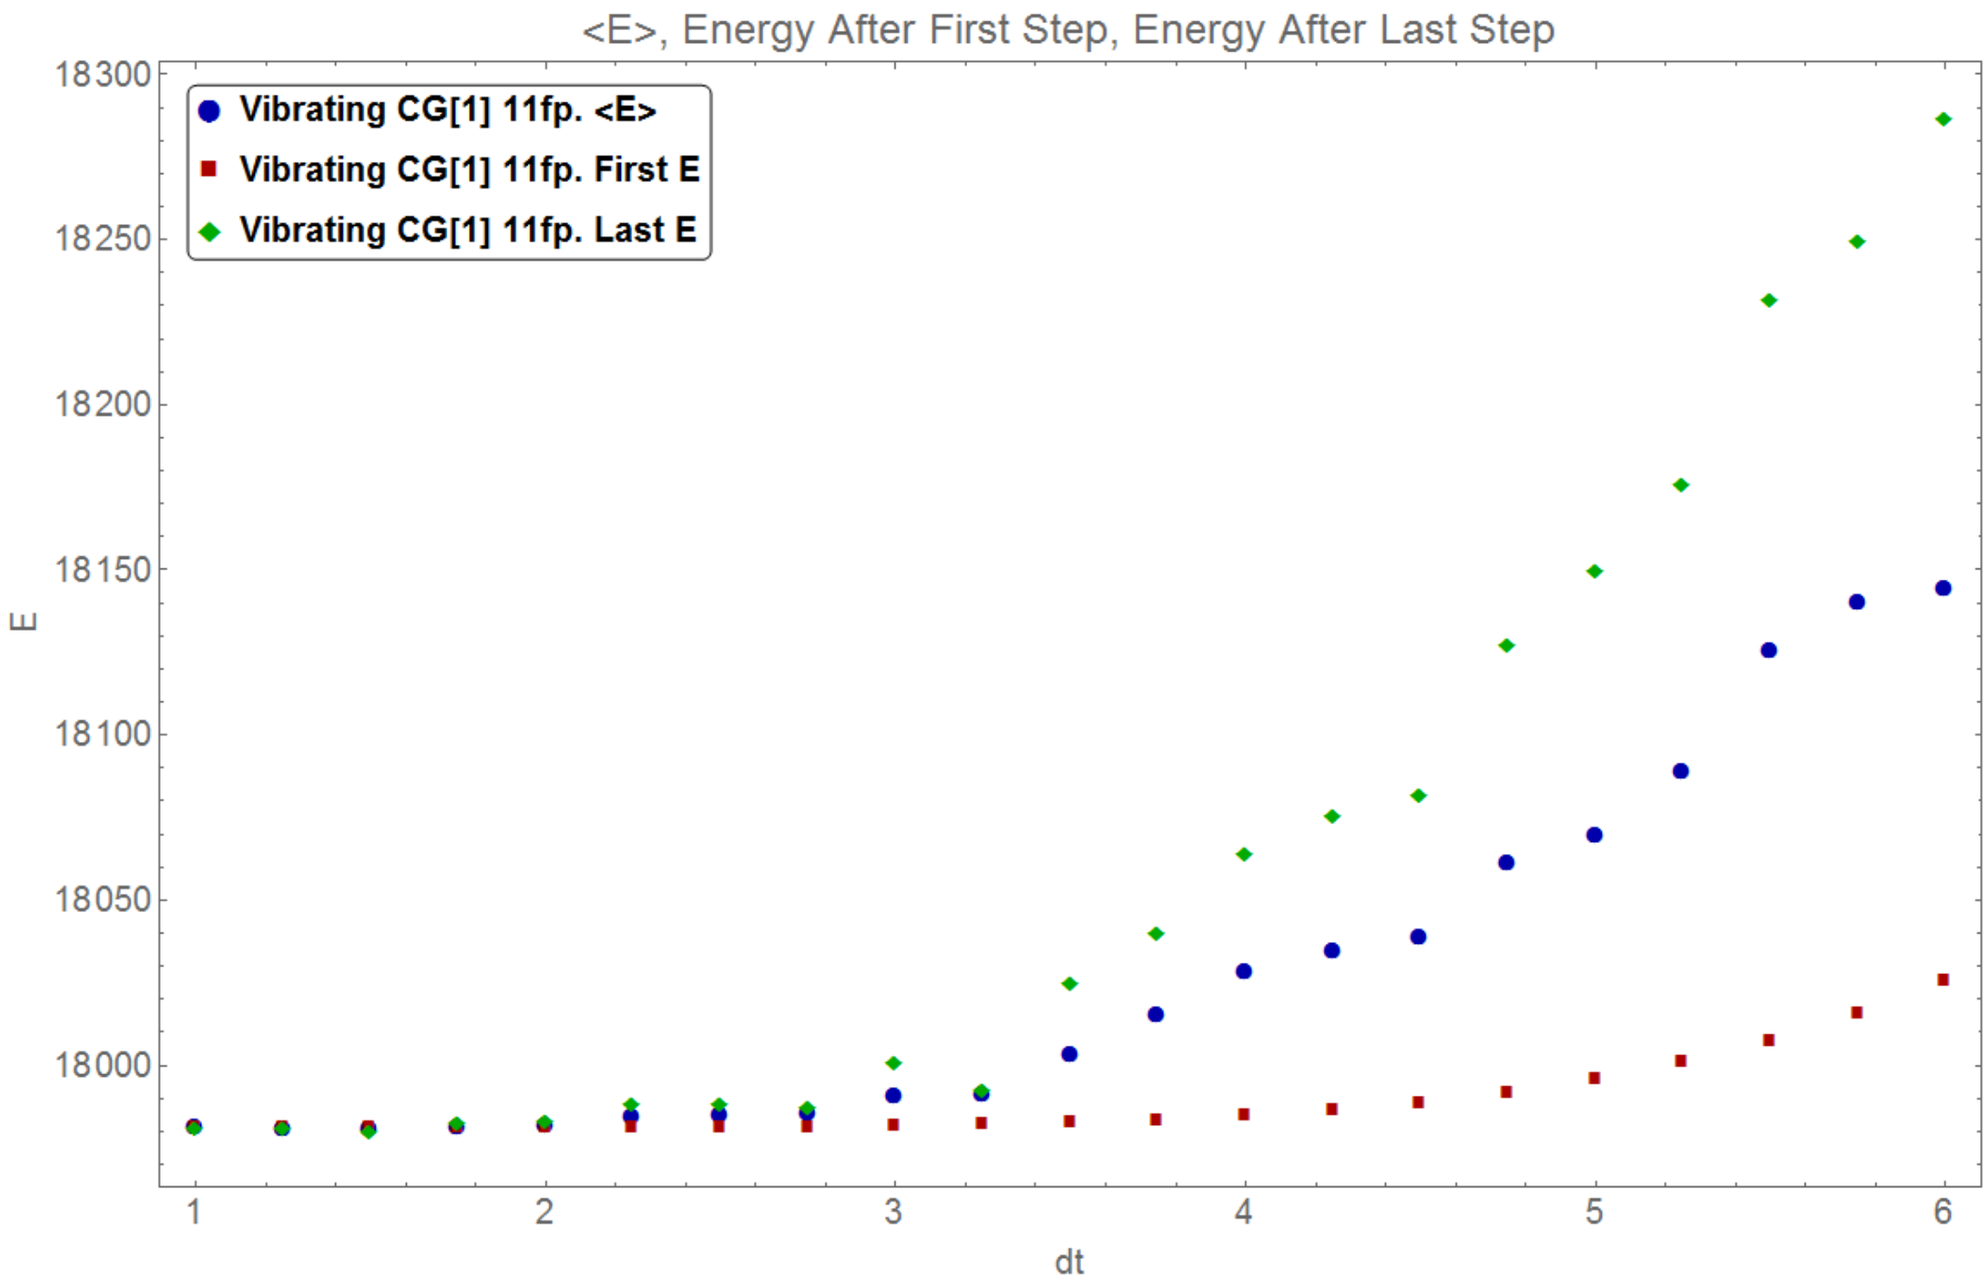
<E>, Energy After First Step, Energy After Last Step
Vibrating CG[1] 11fp. <E>
Vibrating CG[1] 11fp. First E
Vibrating CG[1] 11fp. Last E
18300
18250
18200
18150
18100
18050
18000
E
1
2
3
4
5
6
dt

<E>, Energy After First Step, Energy After Last Step

Rigid CG[1] 11fp. <E>
Rigid CG[1] 11fp. First E
Rigid CG[1] 11fp. Last E

17550
17500
17450
17400
E

1
2
3
4
5
6
dt

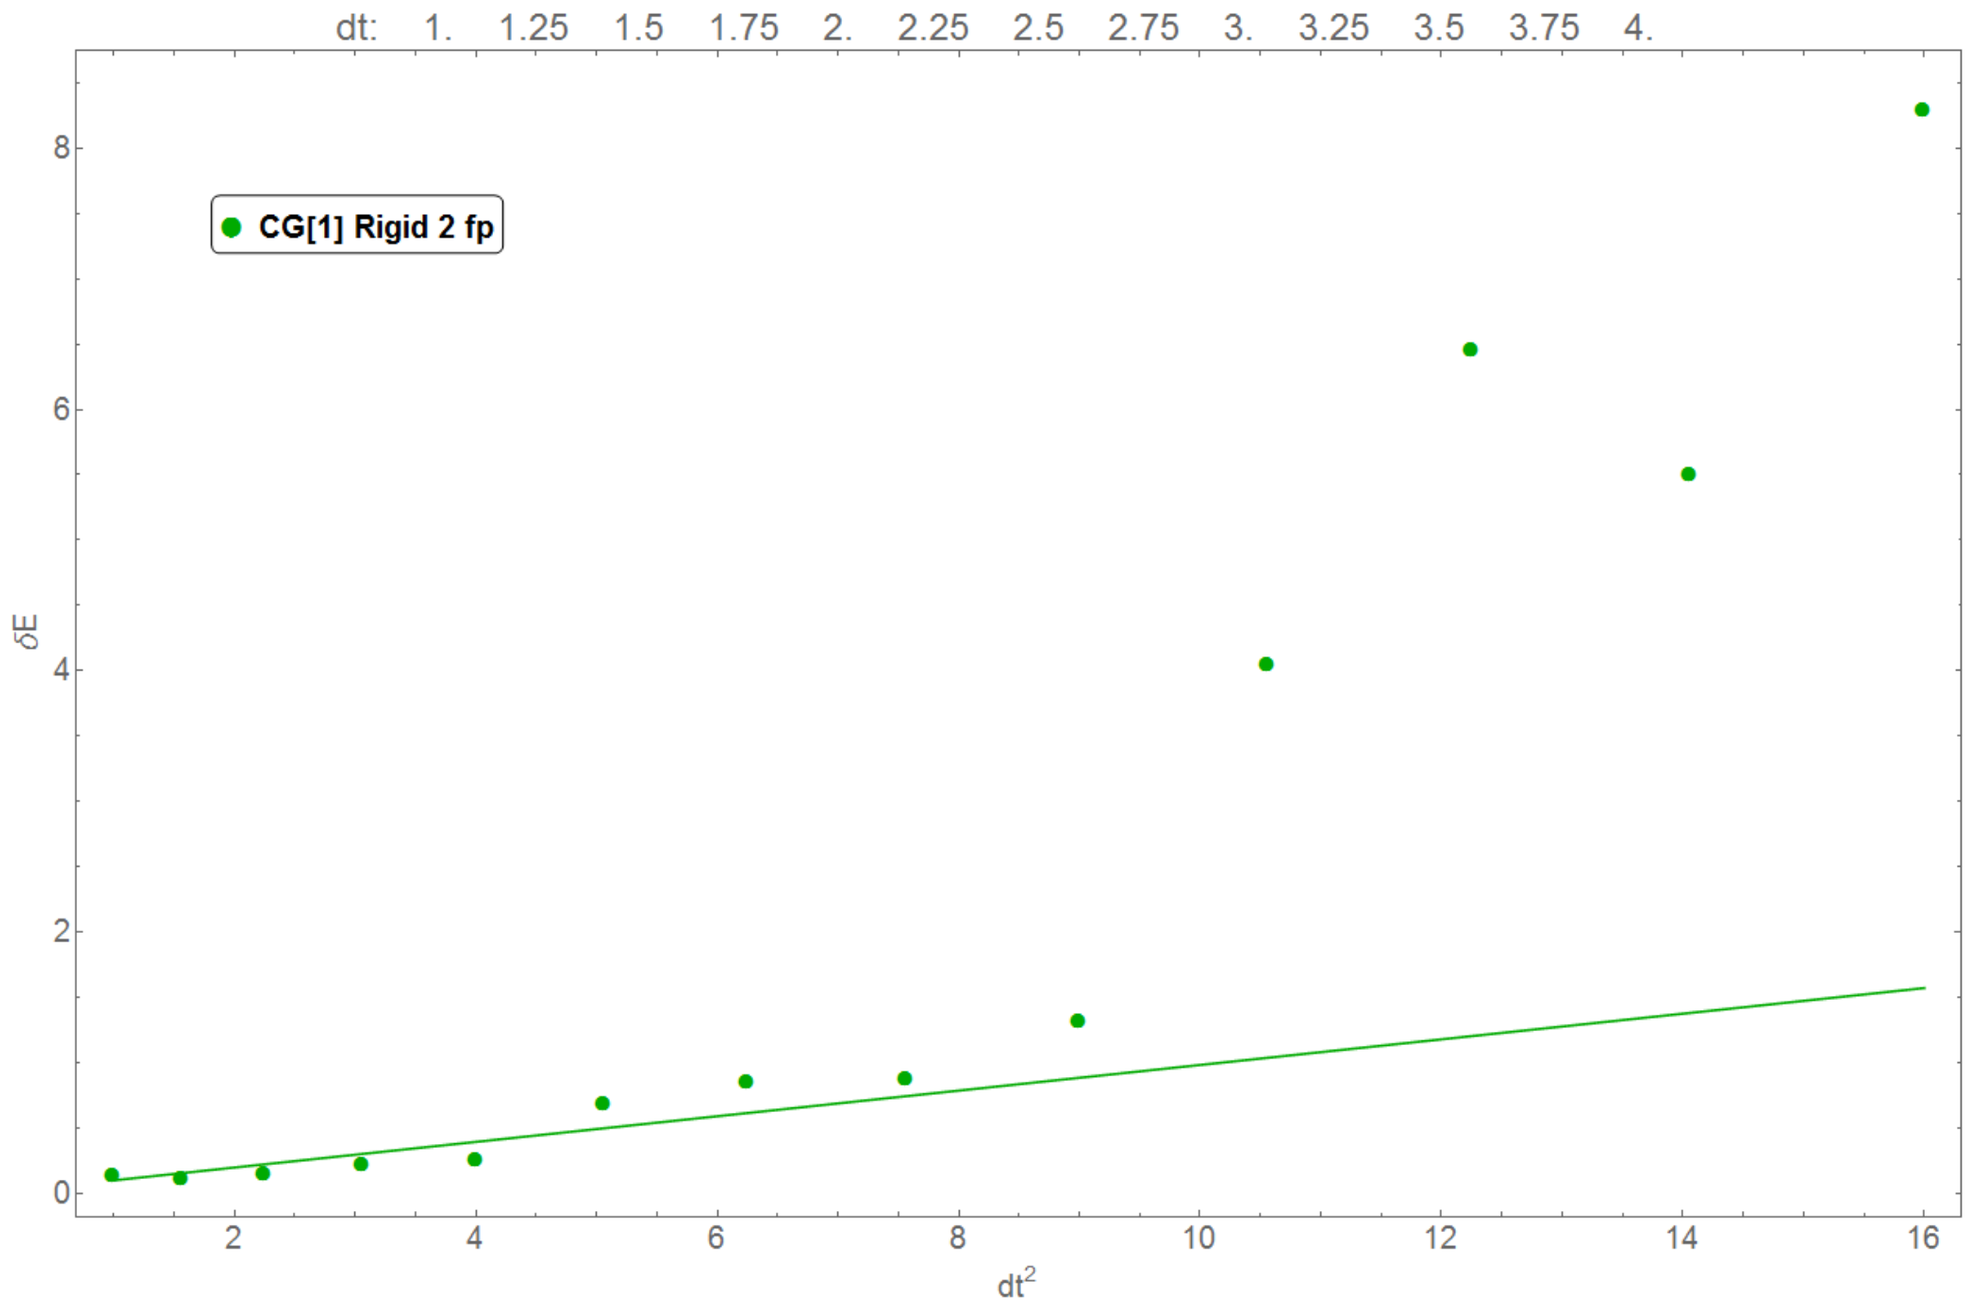
dt: 1. 1.25 1.5 1.75 2. 2.25 2.5 2.75 3. 3.25 3.5 3.75 4.
CG[1] Rigid 2 fp
δE
8
6
4
2
0
2
4
6
8
10
12
14
16
dt²

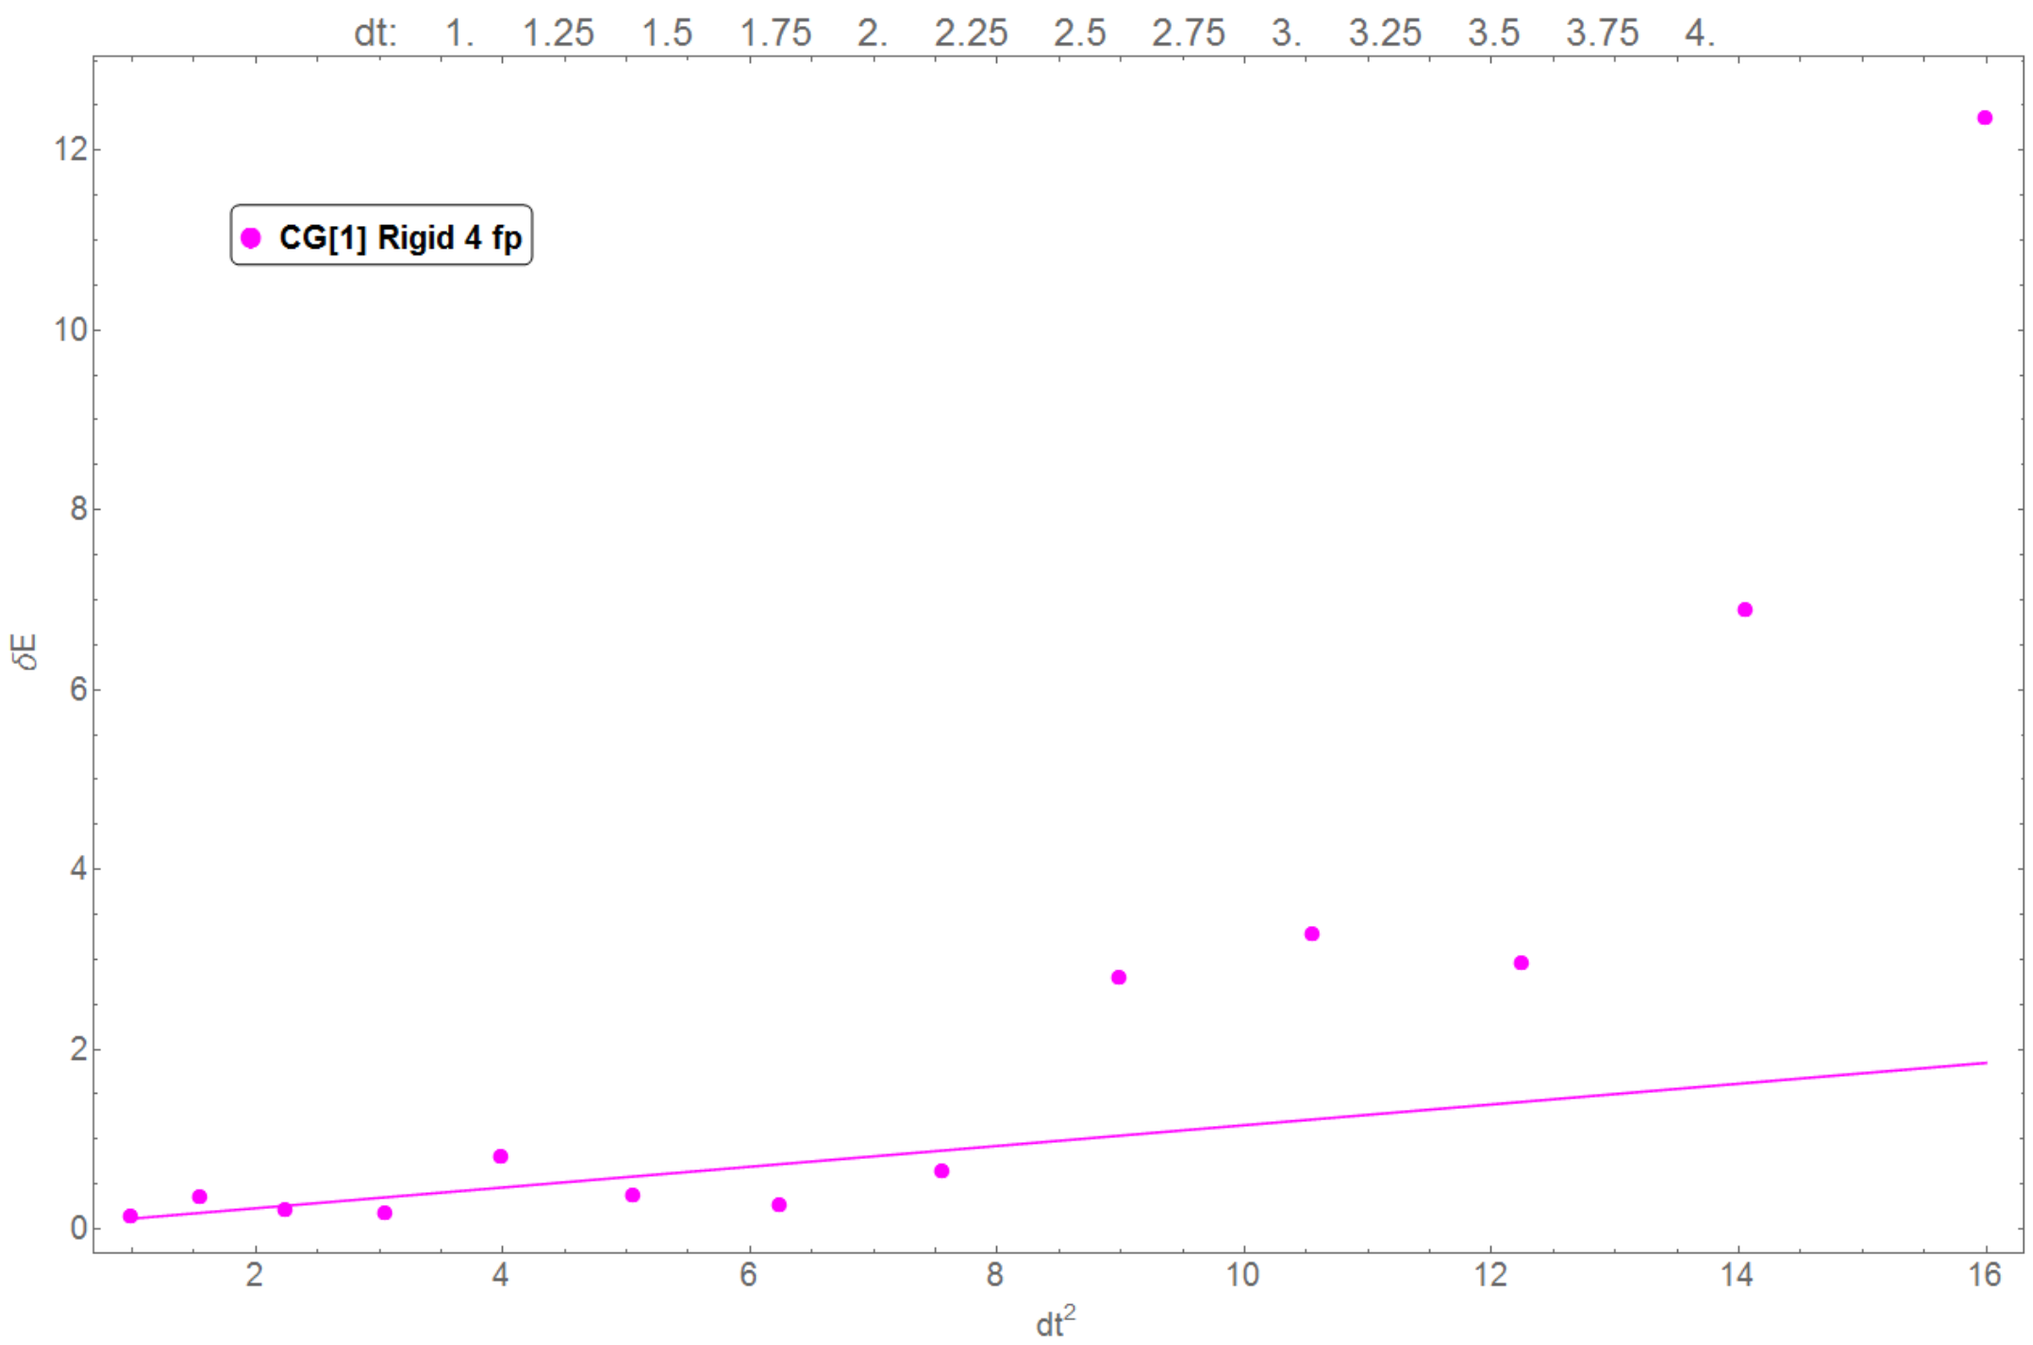
dt: 1. 1.25 1.5 1.75 2. 2.25 2.5 2.75 3. 3.25 3.5 3.75 4.
CG[1] Rigid 4 fp
δE
dt²
0
2
4
6
8
10
12
14
16

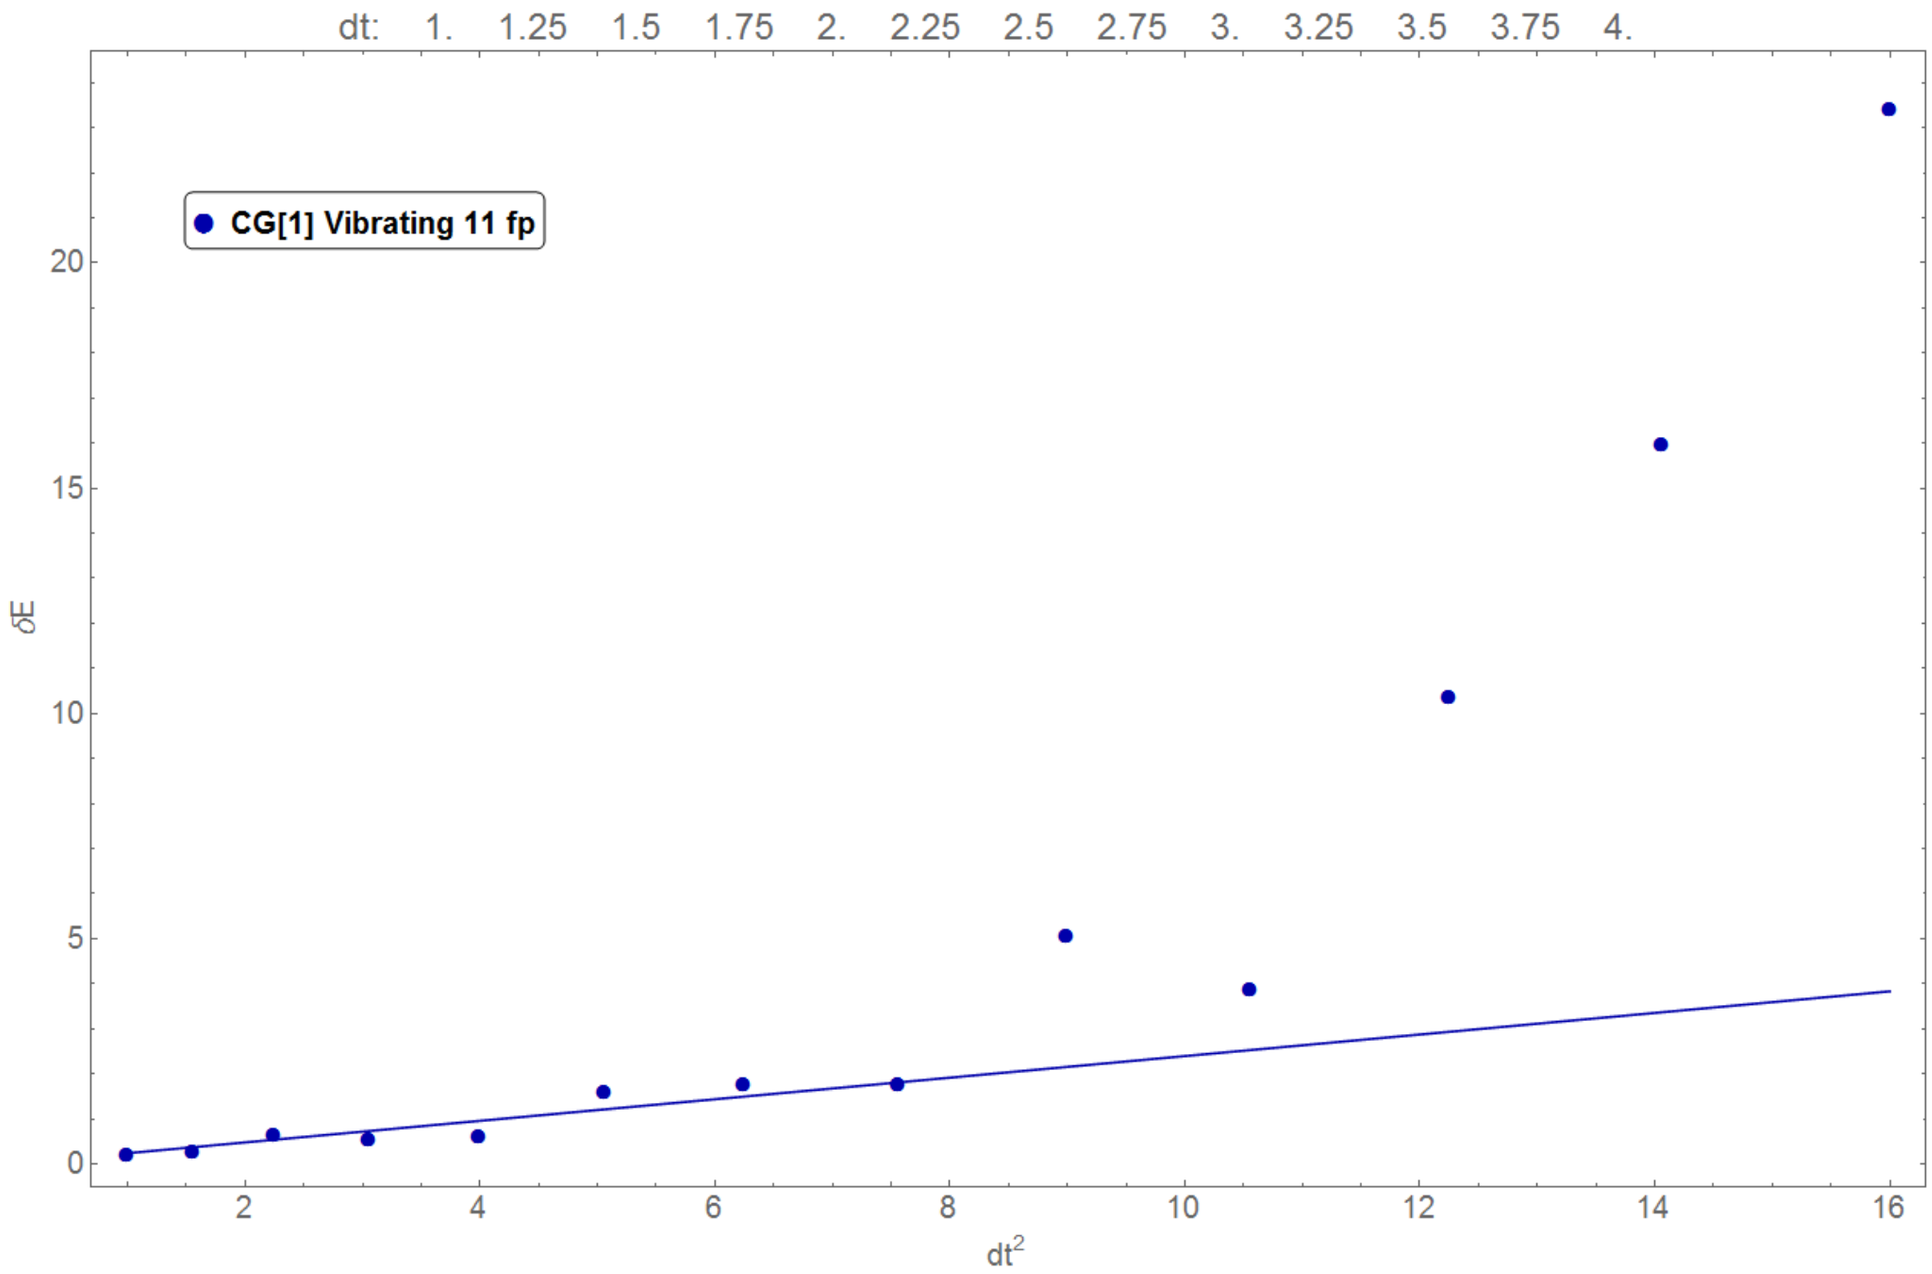
dt: 1. 1.25 1.5 1.75 2. 2.25 2.5 2.75 3. 3.25 3.5 3.75 4.
CG[1] Vibrating 11 fp
δE
20
15
10
5
0
2
4
6
8
10
12
14
16
dt²

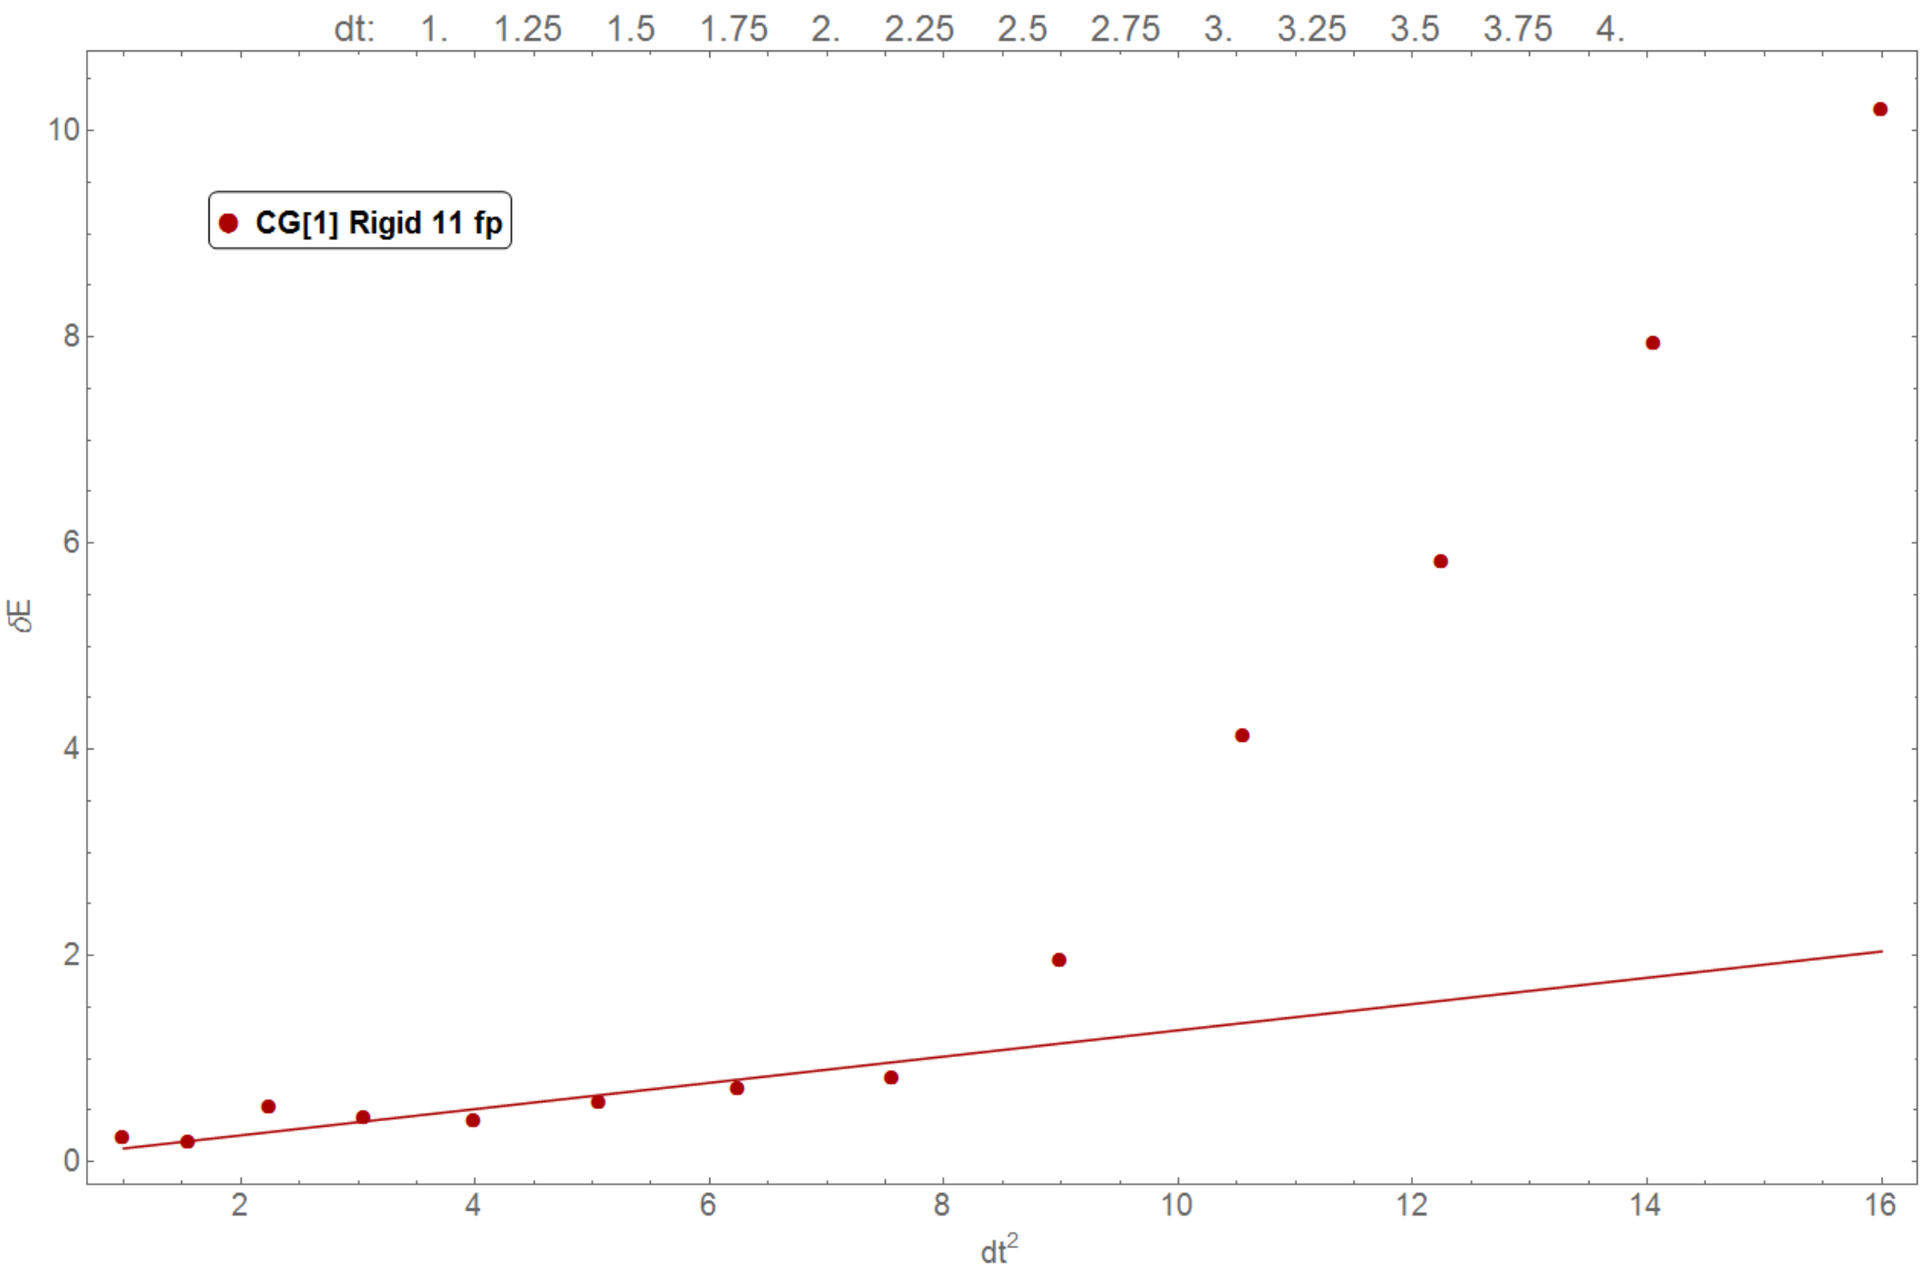
dt: 1. 1.25 1.5 1.75 2. 2.25 2.5 2.75 3. 3.25 3.5 3.75 4.
CG[1] Rigid 11 fp
δE
dt²
0
2
4
6
8
10
2
4
6
8
10
12
14
16

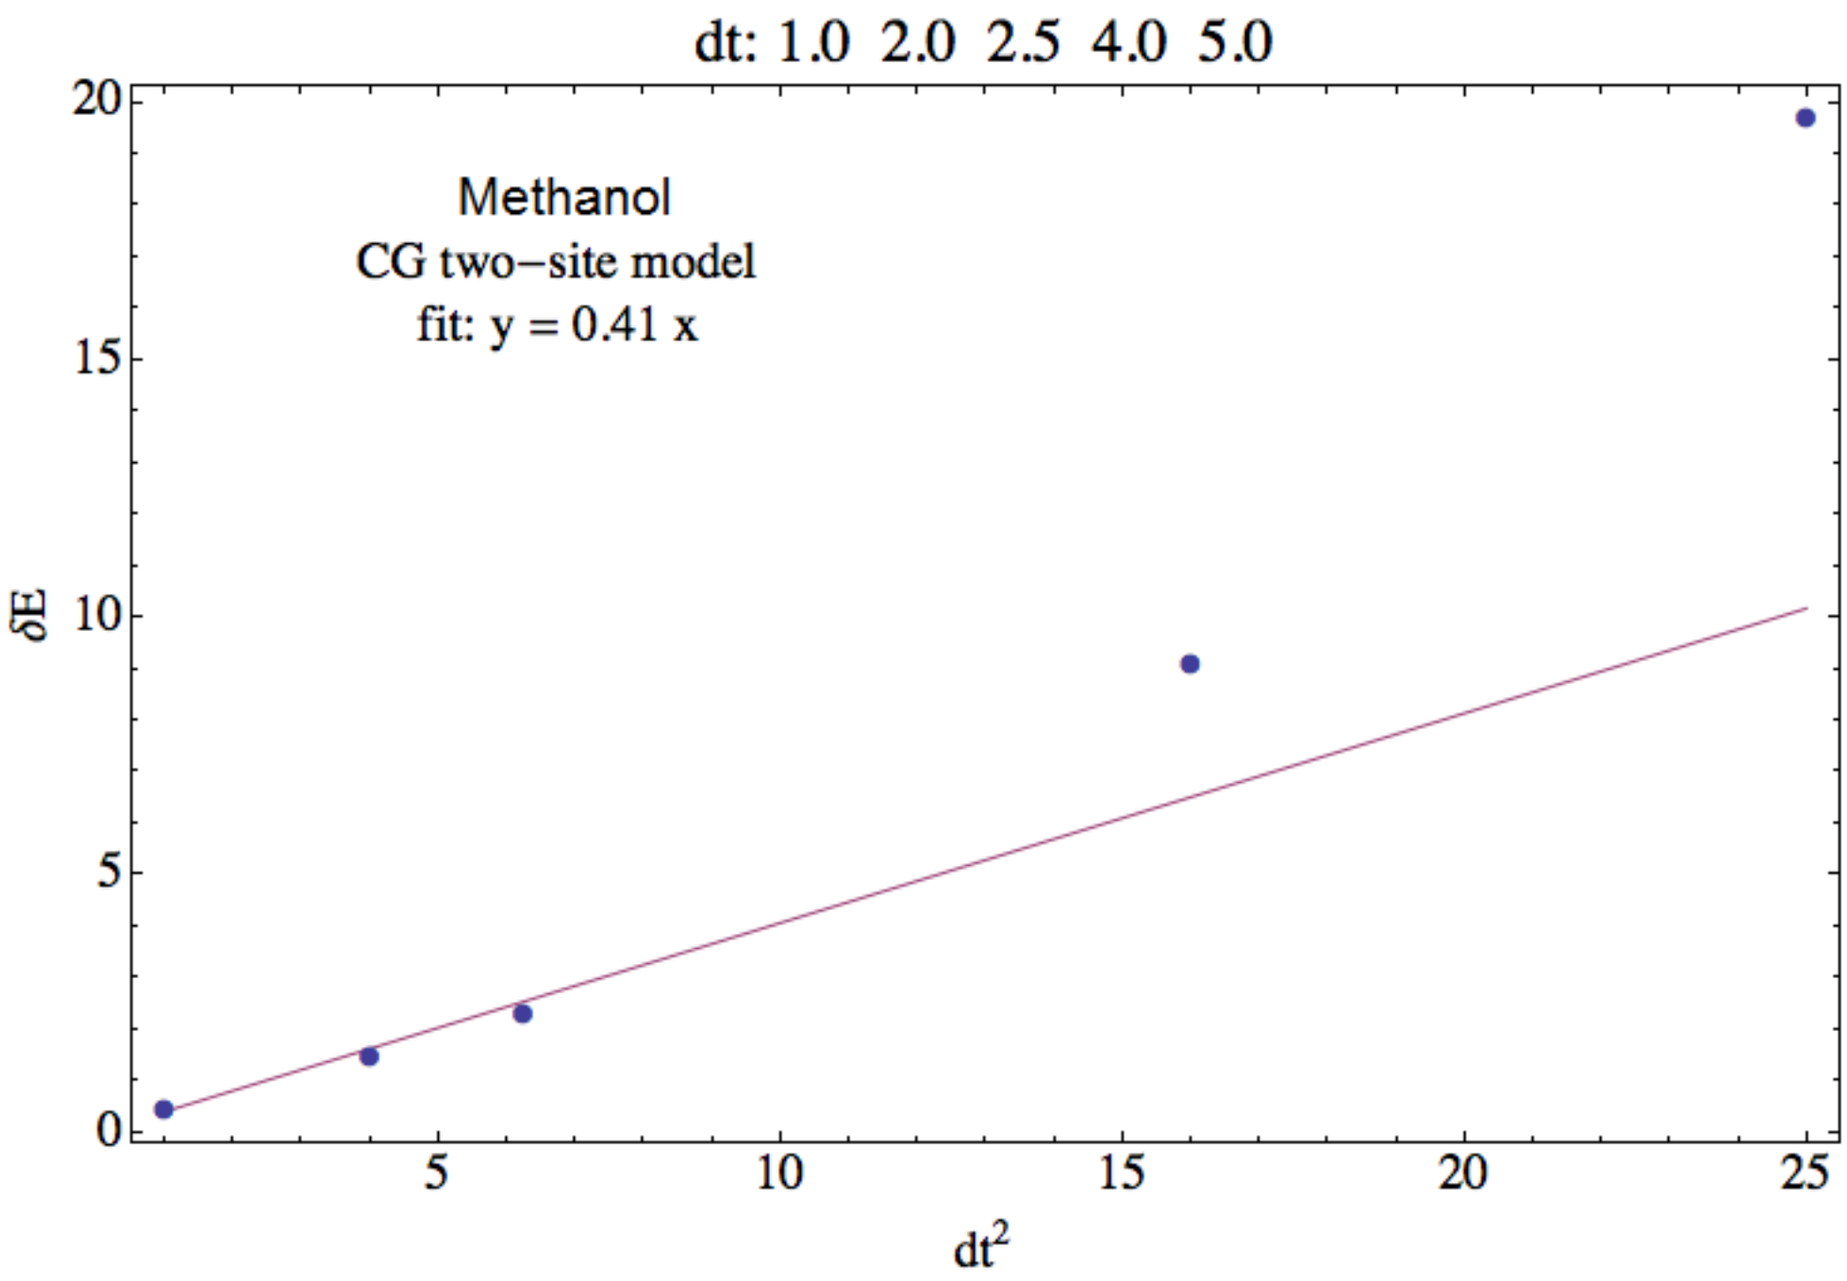
dt: 1.0 2.0 2.5 4.0 5.0
Methanol
CG two−site model
fit: y = 0.41 x
δE
dt²
0
5
10
15
20
25