Attempt to simulate the sensitivity of GEO 600 to beam jitter using WaveProp

Roland Schilling

In the course of increasing the sensitivity of GEO 600 there also came up the idea to reduce the mode-cleaning effect in order to increase the input light power at MPR. In this context the question was raised how sensitive to beam jitter GEO 600 is, i. e. how much mode-cleaning actually is needed.

Therefore, I recently made an attempt to answer this question by trying a simulation using WaveProp. This required only moderate modifications to the GEO 600 simulation program I already used in 2002. (I reported about this at the GEO-meeting 2002 in Glasgow.) The program does roughly the following:

It starts with the light field in front of MPR, transmits it through MPR and interferes it with the light coming back from BS (which initially is an empty field). The interference phase is optimized automatically. Then it is split (the distance MPR – BS is assumed to be zero) according to the BS data and one part is reflected at BS, the other transmitted through BS. The reflected part is brought to interference with the beam coming from the south and transmitted through BS (again initially zero), the transmitted part is brought to interference with the beam coming from the south and reflected at BS (also zero initially). Here the basic phase offset is the same as the one obtained at MPR, but with a certain value added for the detuning of the SR cavity. Then both parts are propagated to their respective far mirrors, reflected there, propagated to the near mirrors, reflected again, a.s.o. back to the BS. Here again two interferences take place forming the two beams traveling towards MSR and MPR. Also the distance MSR – BS is assumed to be zero, but I think, in view of the fact that the Rayleigh length is something like 150 m this doesn't make any difference. The distances to the far mirrors are corrected accordingly. I have taken into account the real mirror data as well as the thermal lens in the east arm caused by the beam-splitter.

In a first step I run this program with a well aligned input beam. After 20 000 round trips, corresponding to ≈ 160 ms in real time, the fields have settled, as can be seen in Fig. 1. This is done with 1 W input power and takes about $3\frac{1}{2}$ min on the QND here in Hannover. At the end of this run the settled values of the phase offsets needed for constructive interference at MPR and for a dark fringe at the south port as well as the light fields running towards MPR and returning from MSR are recorded in separate files. This allows to restart the simulation at exactly the point reached after 20 000 round trips, thus saving computing time.

The next step then is to do such a restart, but with a slightly modified program. Now, the interference phases are not allowed to be optimized anymore, but frozen to the value obtained at the initial run. This corresponds to an extremely low unity-gain frequency for the MPR and dark-fringe servos, i. e. all dynamic actions are uninfluenced by this servos. The input beam is not perfectly aligned anymore, but wiggled sinusoidally in the input angle by a certain amount, the *angular beam jitter*, and at a certain frequency. For, e. g., a frequency of 500 Hz we only need to calculate 250 round trips in order to cover one period, since one round trip corresponds to $\approx 8\,\mu$s.

I tried three frequencies (250 Hz, 500 Hz and 1 kHz) and two values for the amplitude of the beam tilt (1 μrad and 0.5 μrad). The results are shown in the figures on the following pages. Each plot contains four curves: In cyan is shown the tilt of the input beam, in red the deviation of the center of the beam (coming from BS) at MFE, in blue the deviation from the dark-fringe phase (corrupting the GW signal) and in magenta the deviation from the resonance phase of the PR cavity. Please note the different Y-scales on the respective upper and lower plots.

I also have repeated the whole procedure with a certain, fixed dark-fringe phase offset (1 mrad) as is used for DC-lock or DC-readout, but without any significant difference.

I have quite some difficulties to understand and interpret the outcome of this simulation. I suppose that the high-frequent ripples on the traces are an artifact caused by the iterative nature of the simulation. The amplitude of these ripples slowly decays by about a factor of two over 800 round trips.

There are mainly two peculiar things one can deduce from the plots:

Looking at the red traces (the spot position at MFE) one can conclude that the amplitude of the displacement at MFE is proportional to tilt amplitude and, strangely enough, also proportional the tilt frequency. From the blue curves (the dark-fringe phase deviation) one can see that there obviously is a quadratic transfer function involved: the frequency is twice tilt frequency and the amplitude is proportional to the square of the tilt amplitude. But the big surprise is their independence from the frequency!

Can anybody explain this strange behavior? Is it all nonsense and caused by the deficiencies of the simulation or in my approach to tackle the problem?

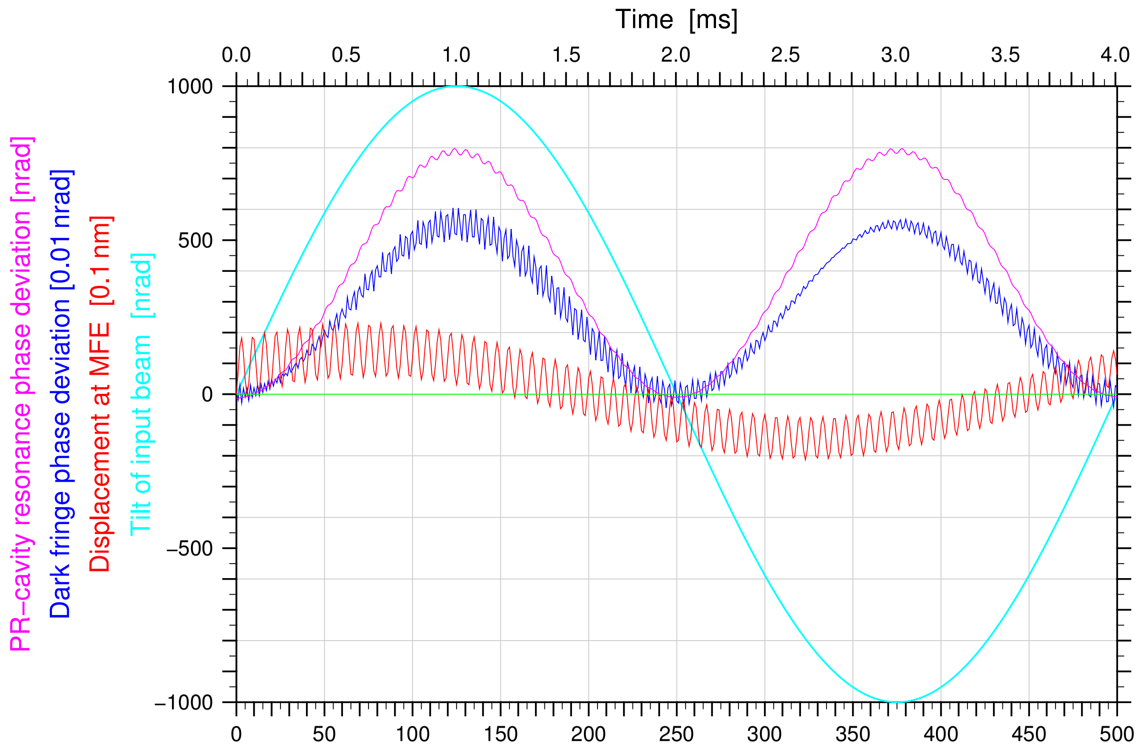
Time [ms]
0.0
0.5
1.0
1.5
2.0
2.5
3.0
3.5
4.0
1000
500
0
−500
−1000
0
50
100
150
200
250
300
350
400
450
500
PR−cavity resonance phase deviation [nrad]
Dark fringe phase deviation [0.01 nrad]
Displacement at MFE [0.1 nm]
Tilt of input beam [nrad]