Argonne NATIONAL LABORATORY

... for a brighter future

Elegant Parallelization Progress Report

10/24/07

Yusong Wang Michael Borland

Hairong Shang Robert Soliday

Simulations with a Large Number of Particles

- Recent development
  - *Each slave is allocated memory only for the portion of particles it owns, instead of the memory required for all the particles*
  - *On a cluster of quad-cores, configured master to be run on a single node without sharing memory with other slave nodes to allow maximal number of particles to be simulated (master holds all the particles information for reading and writing)*
- Current capability and limitations
  - *Simulated with 60 million particles on apex cluster*
  - *Can't simulate with 100 million or more particles because of the memory bottleneck on the master*
  - *Reading a large particle file is very time-consuming (serial operation)*
- Planned future work
  - *Use parallel SDDS I/O to finish all the I/O operations without a central process (i.e. Master) holding all the particle information, which will improve both the capability and performance of **Pelegant***

Increasing Number of Particles is Necessary

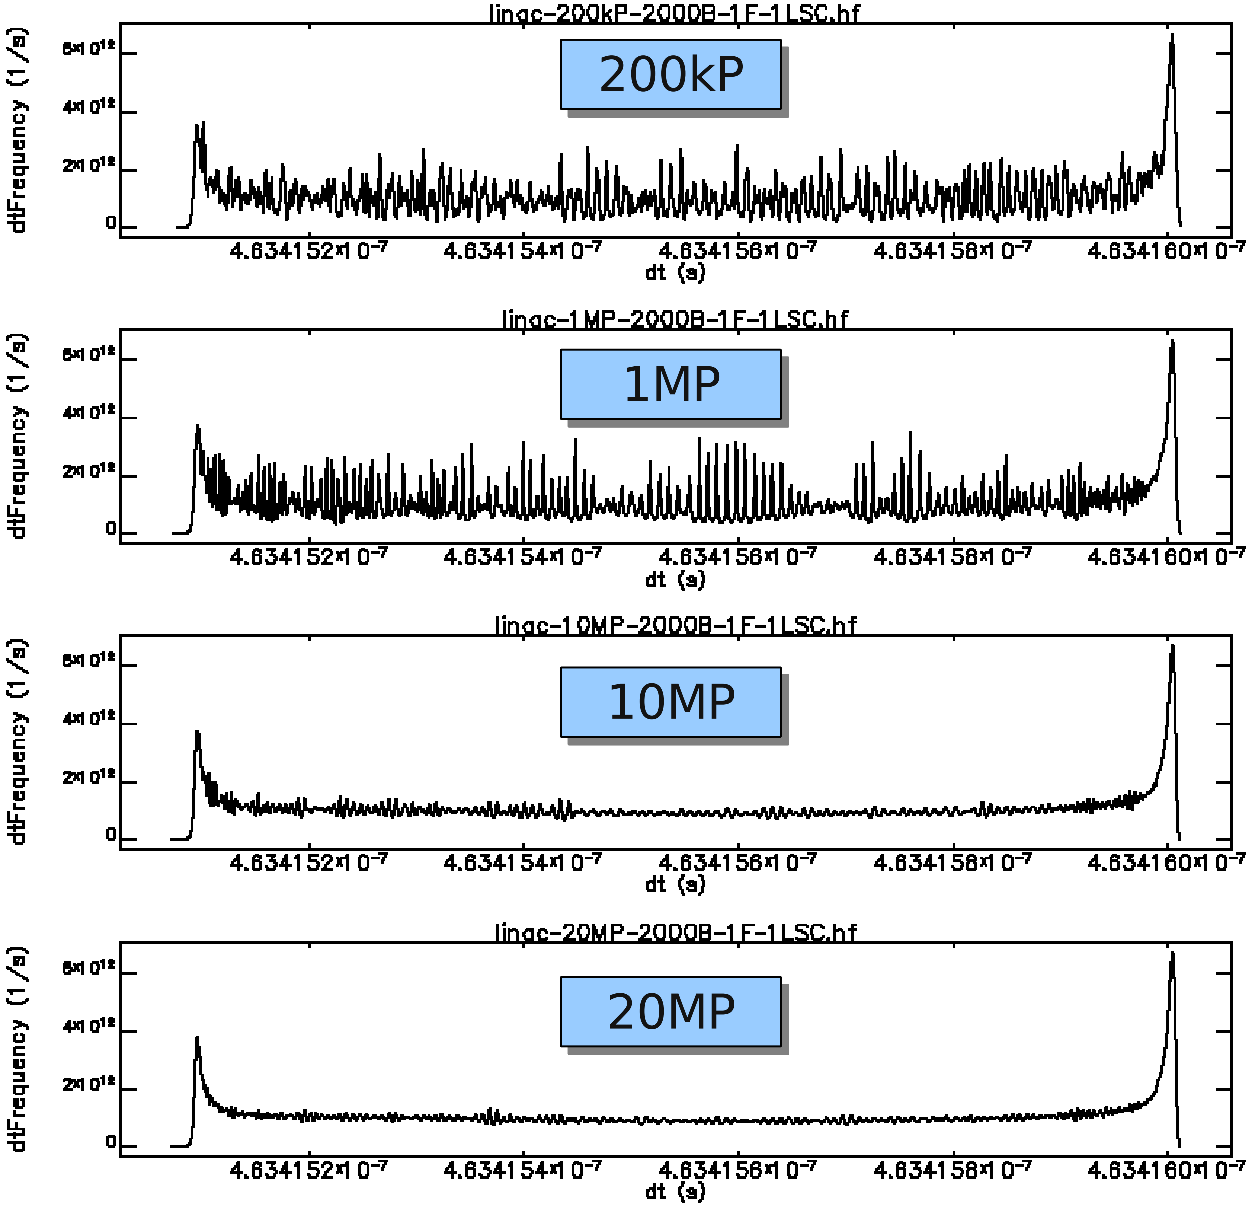

- Things start to look good for ~20 million particles
- We ended up using 60 MP in FERMI simulations for microbunching gain scan

Parallel SDDS I/O

- Motivations
  - *The computing power of **elegant** has been enhanced significantly because of recent parallelization and optimizations. The SDDS library with sequential execution become the bottleneck of the simulation for both memory and I/O operations*
- New parallel SDDS features
  - *The basic reading and writing functions have been parallelized in the SDDS I/O based on the parallel I/O in MPI.*
  - *There is no difference for end users, as it will operate (read and write) on a single file as before.*
- Performance
  - *All tests were done on NFS file system, will test on PVFS in the near future.*
  - *It is able to read 200M particles (file size is 12G, and it took 340 seconds with 64 processors), which is impossible with serial SDDS.*

Parallel SDDS I/O

- Performance
  - *Results of reading 20M particles (1.2G file with 20M rows and 7 columns)*

Number Of Processors	Time Cost (seconds)	SpeedUp
1	101.81	1
2	39.18	2.6
4	27.33	3.73
8	20.36	5
16	19.15	5.32

- Future Plan
  - *Add feature of parallel reading/writing SDDS file in column order (instead of row major order as now) to improve the performance*
  - *Add parallel visualization feature for large simulation result*
  - *May use collective-I/O for better performance.*

Multi-bunch Simulations

- Challenges for parallelization
  - *The collective effect in the FTRFMODE element (and other elements) requires intensive communications between processors to calculate a global histogram from local histograms computed by each of processors.*
  - *E. g., in a simulation of multiple short-bunch beam, we need 4 million bins for the histogram on each processor, which will have 4 million*n_cpu*2*8 bytes messages to be passed. This will happen in every turn.*
  - *The gain of performance from particle decomposition can be vanished by these heavy communications. In the extreme case, the code can't even be run because of network saturation.*

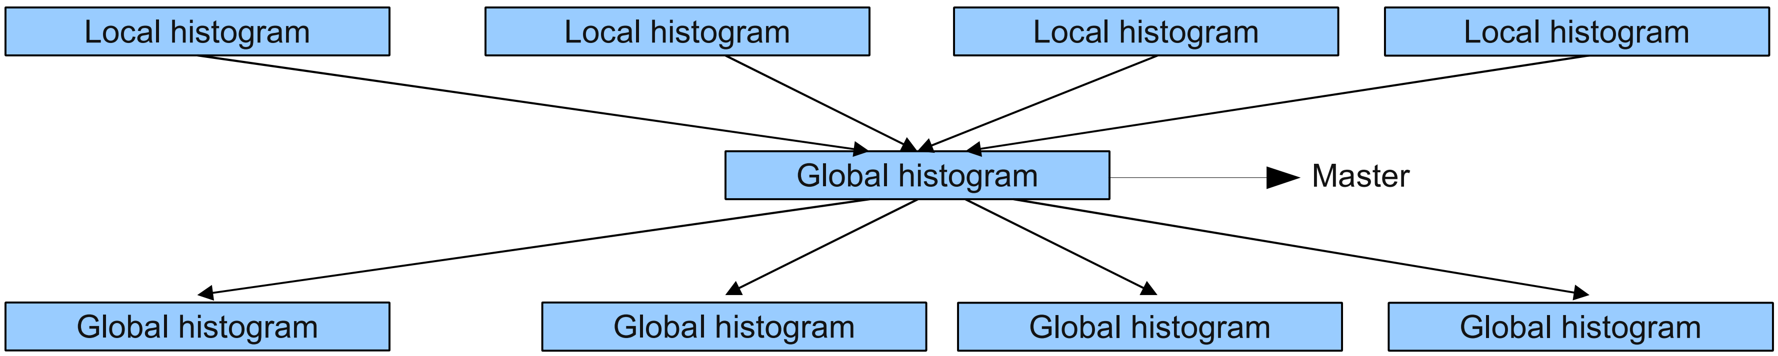

Multi-bunch Simulations (Continued)

- Parallel implementation
  - *We observed, in practice, a lot of bins in the histogram are just empty, i.e., the information are useless in those empty bins.*
  - *By passing the information only in the non-empty bins with their real indices in the global array, the communication size was reduced to ~ 1/1000 of the original (the whole histogram) size.*

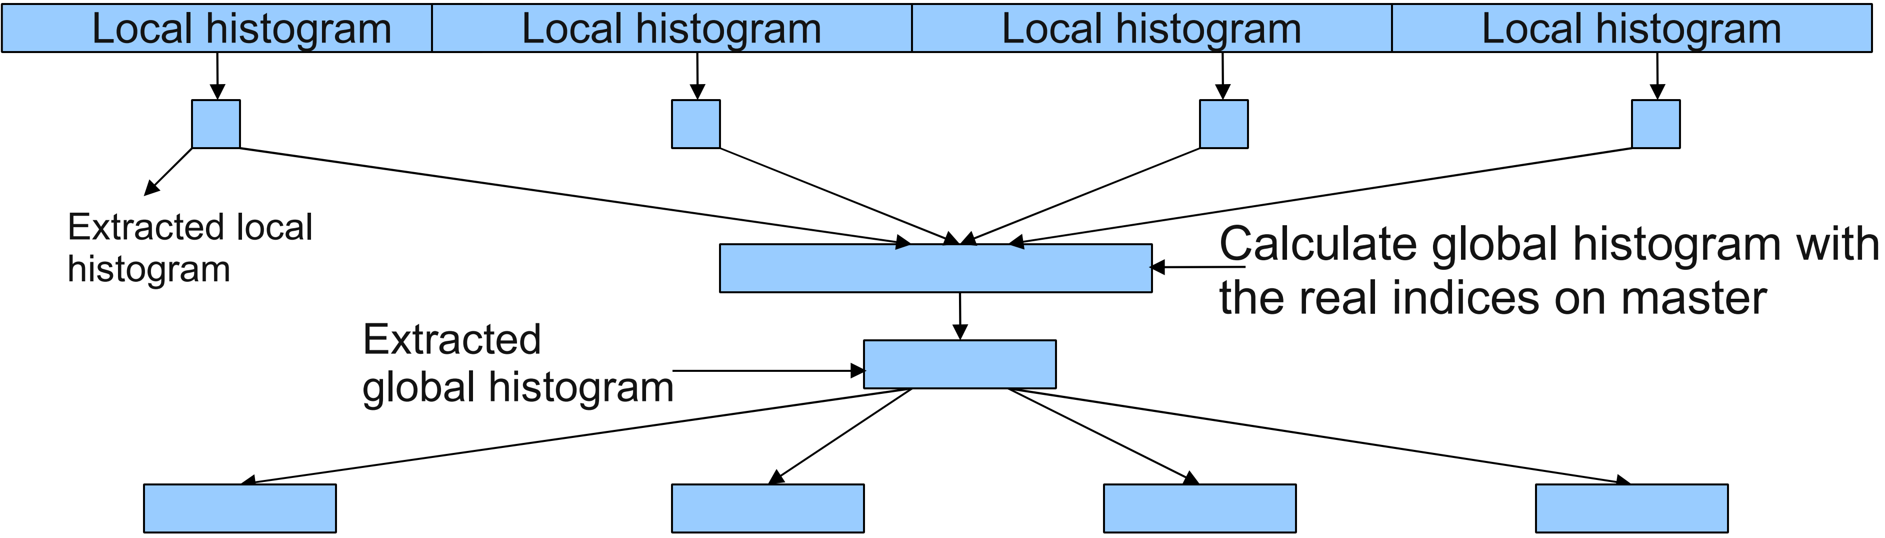

Multi-bunch simulations (Continued)

- Other optimizations
  - *Profiled the code with jumpshot and found the initialization spent a big portion of time for the ftrfmode element. After replaced with calloc function the time, tracking time was reduced significantly*
  - *Parallelized the voltage calculation for different modes, which could reach linear speedup if the number of mode is equal to the number of CPUs, as voltage for each mode can be calculated independently*
  - *Parallelized the watch_parameter operation: the slaves will be responsible for calculating its local part of contributions and sending the statistical result to the master for writing;*
- Performance and future work
  - *Simulation of 1.5 millions particles with 4 million bins running 30000 turns: Improved **elegant** (1CPU): 41:37:08; **Pelegant** (on 24 CPUs): 02:37:53*
  - *Duplicate this model to other collective-effect elements in **elegant** for multi-bunch simulations*

Other Improvements

- Parallelization
  - *Parallelized the LSCDRIFT element*
  - *Parallelized the histogram and improved the efficiency by 20 ~ 30 % for an LSC Oscillation example*
  - *Added savitzkyGolay implementation for all the orders in the time domain, which reduced time spent on a simulation with impedance study from 31 hours to 38 minutes on weed with 10 processors. Simulations including CSRCSBEN and WAKE elements with savitzkyGolay smoothing will benefit from this improvement significantly.*

Applications of Pelegant

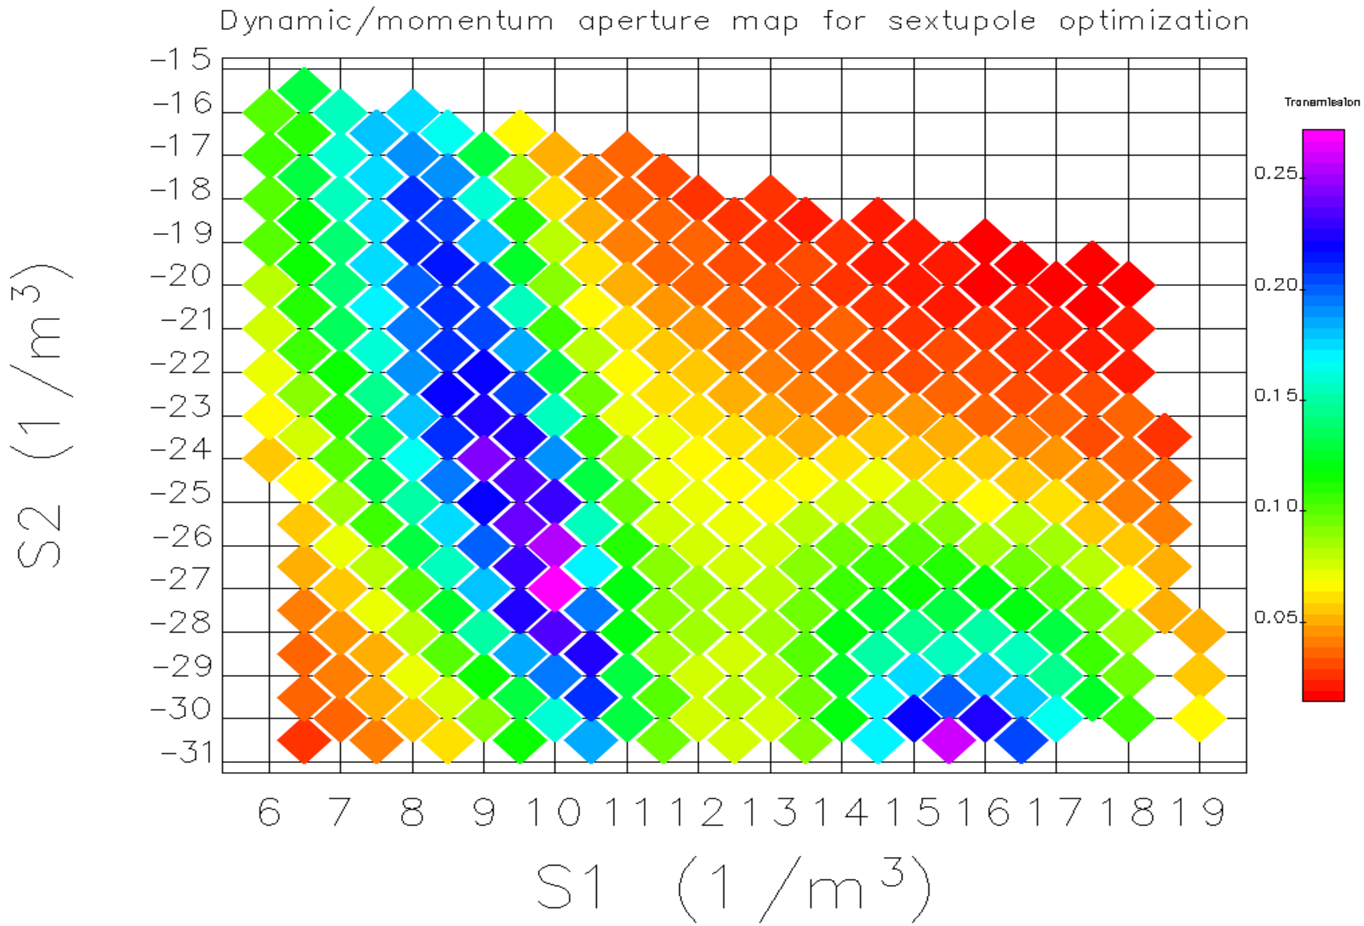

APS storage ring modeled with elegant (M. Borland et al.)

- This is an example of the use of parallel elegant to optimize the dynamic and momentum aperture of the APS storage ring. Application of this technique has significantly improved lifetime issues that arose in APS operations.

Applications of Pelegant (continued)

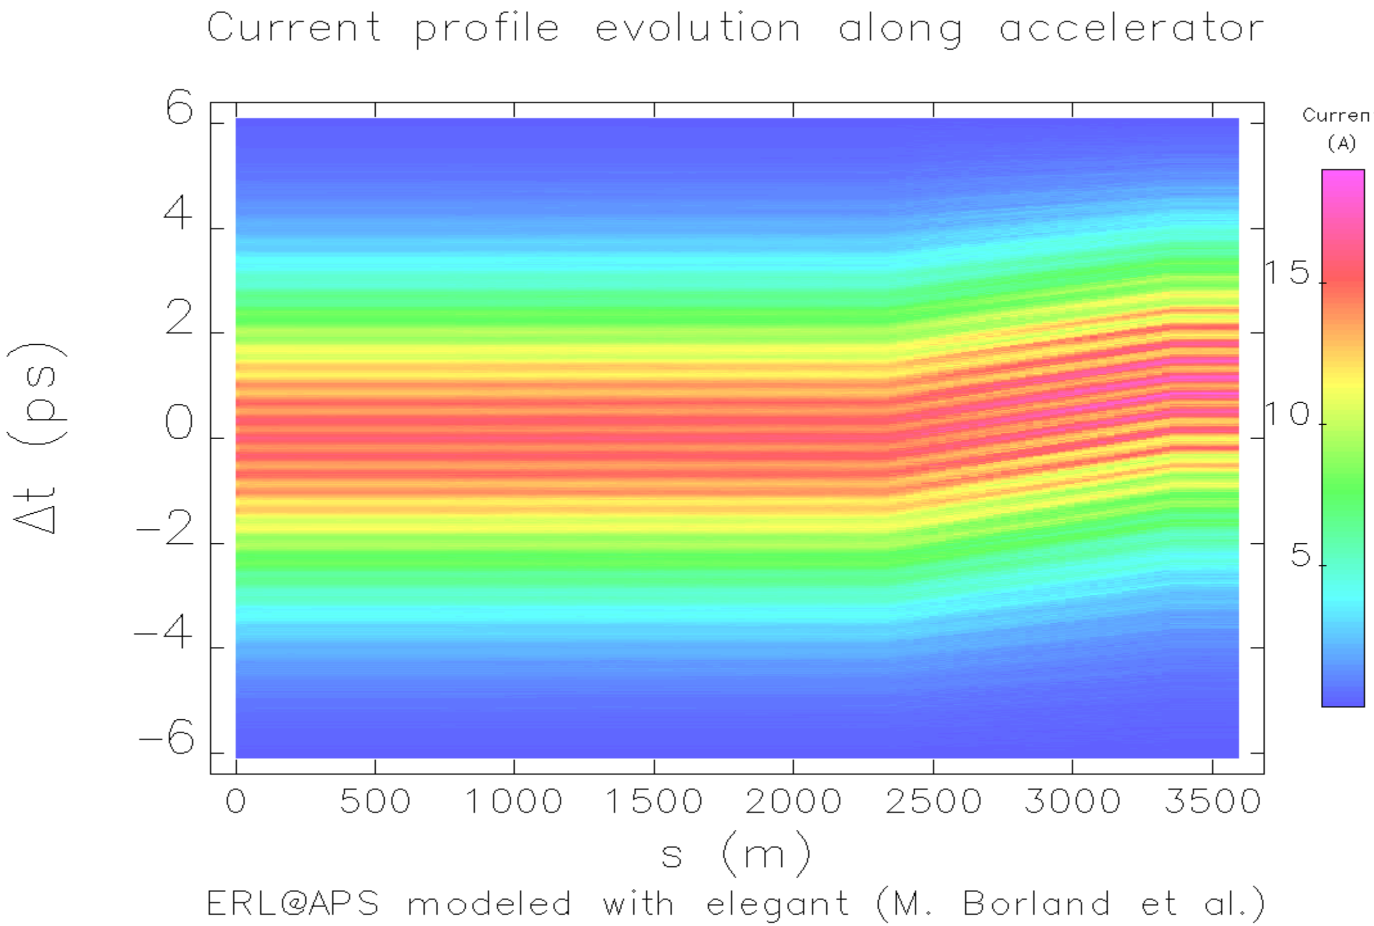

- This image shows the evolution of a 100-micron longitudinal density modulation imposed on an electron bunch traveling through the Outfield ERL@APS lattice. The simulation shows the relatively modest increase in the depth of the modulation as a result of longitudinal space charge and coherent synchrotron radiation.

Applications of Pelegant

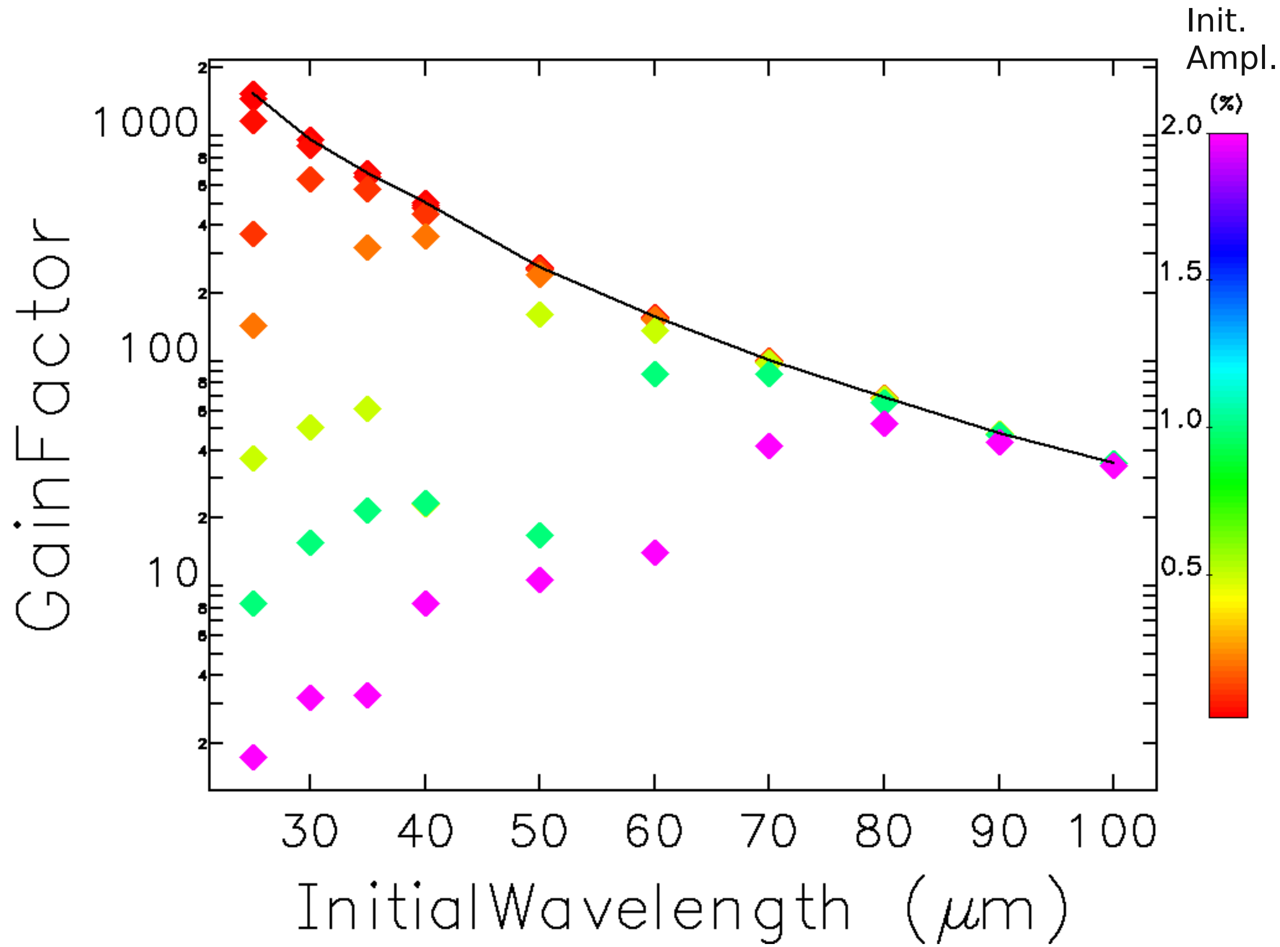

- This image shows the CSR+LSC gain for FERMI@ELETTRA as a function of the initial modulation wavelength. Decreasing the initial density modulation amplitude to 0.013% is necessary to avoid saturation and get the true linear gain. Requires 30~60 M particles.