# ALIVE: Proton-Beam-Driven Plasma Collider

July 2024

J. Farmer , A. Pukhov, F. Willeke and A. Caldwell

# Outline

- Motivation
- Recap: short proton driver for wakefield acceleration
- Scaling to high energy
- Round witness beams
- Conclusion

Special emphasis on the challenges!

# Motivation

Currently available proton drivers have sufficicient energy to accelerate a witness bunch to the energy frontier in a *single stage*.

- inject once, at low energy
- higher average gradient
- potential to reuse existing infrastructure

However, while driver energy is a good start, it’s not sufficient.

# Proton challenge #1

At wavebreaking, plasma wakefields impart $mc^2/e$ per $1/k_p$.

Driver bunch should be $\sim 1/k_p$ long.

High gradients require short drivers,
but current proton beams are centimetres long.

# Proton solution #1a

High wakefields are still possible with long beams

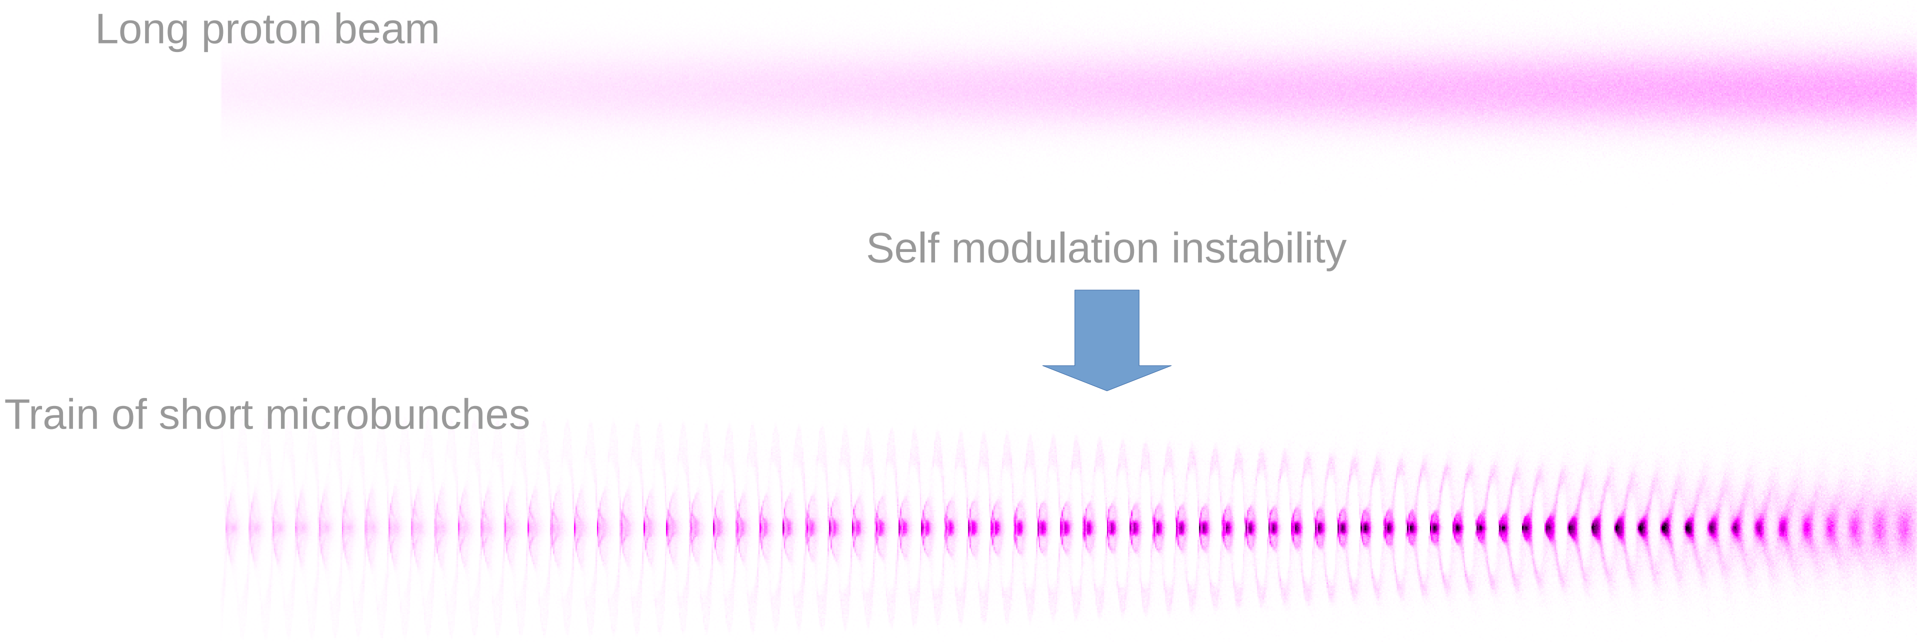

# Proton solution #1a

Why not build an AWAKE collider?

Luminosity

- SPS drive beam every ~30 seconds
- SMI limits efficiency (driver loss and saturation)

Not a criticism of AWAKE,
interesting studies possible for fixed-target, and lots of R&D.

# Proton solution #1a

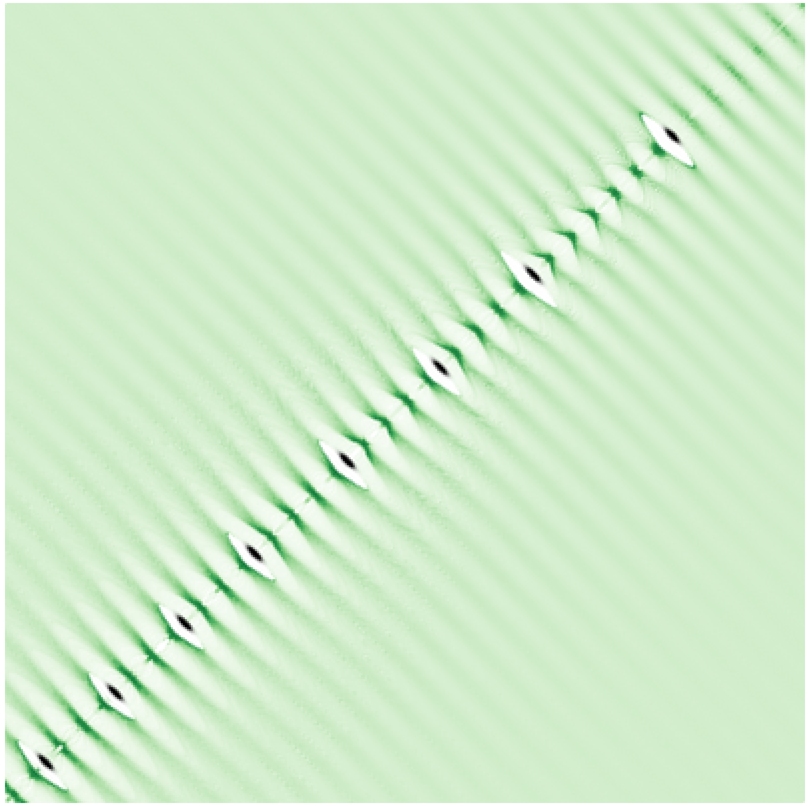

Injecting multiple witness bunches could help, but not sufficient

Wakefield regeneration in a plasma accelerator, PRR, *in press*.
https://arxiv.org/abs/2404.14175

# Proton solution #1b

Invent a new proton source.

Should have:

- sub-millimetre bunch length (gradient)
- ~100 µA average current (luminosity)
- Reasonable ($n_b$~$n_e$) charge density (gradient)

Original concept was to use fast-ramping synchrotron.
Better solutions may exist, see Ferdinand’s talk tomorrow.

# Recap: Higgs factory

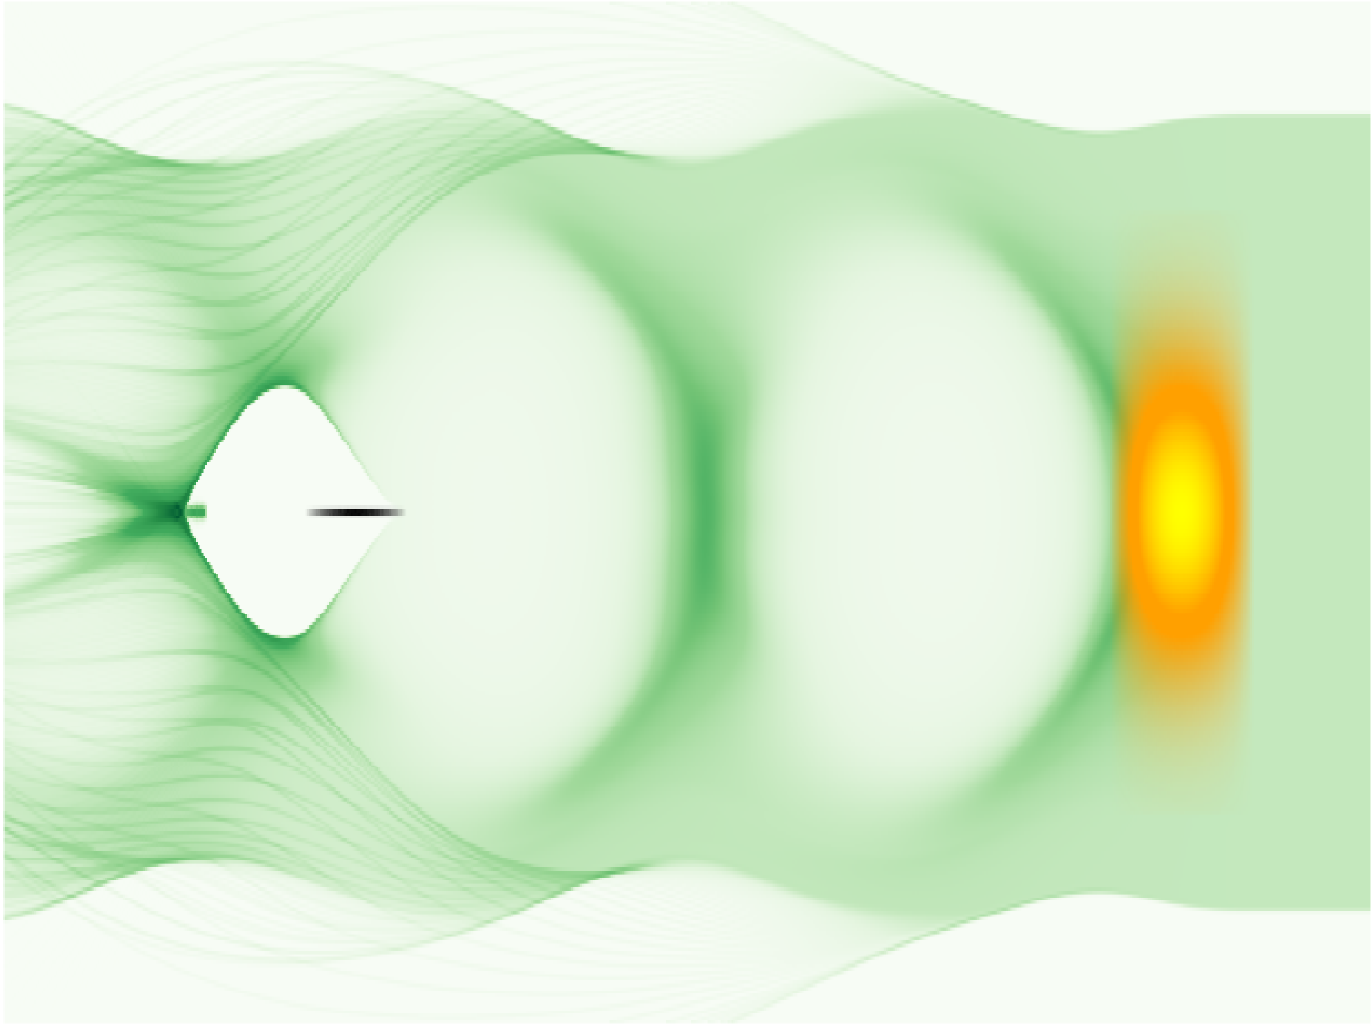

Short driver excites moderately nonlinear wake.

Witness electron bunch drives its own blowout (AWAKE alumni).

Olsen, Adli, Muggli (2018)

Positron acceleration possible (but not easy).

# Recap: Higgs factory

Initial proton driver chosen to generate suitable wakefields

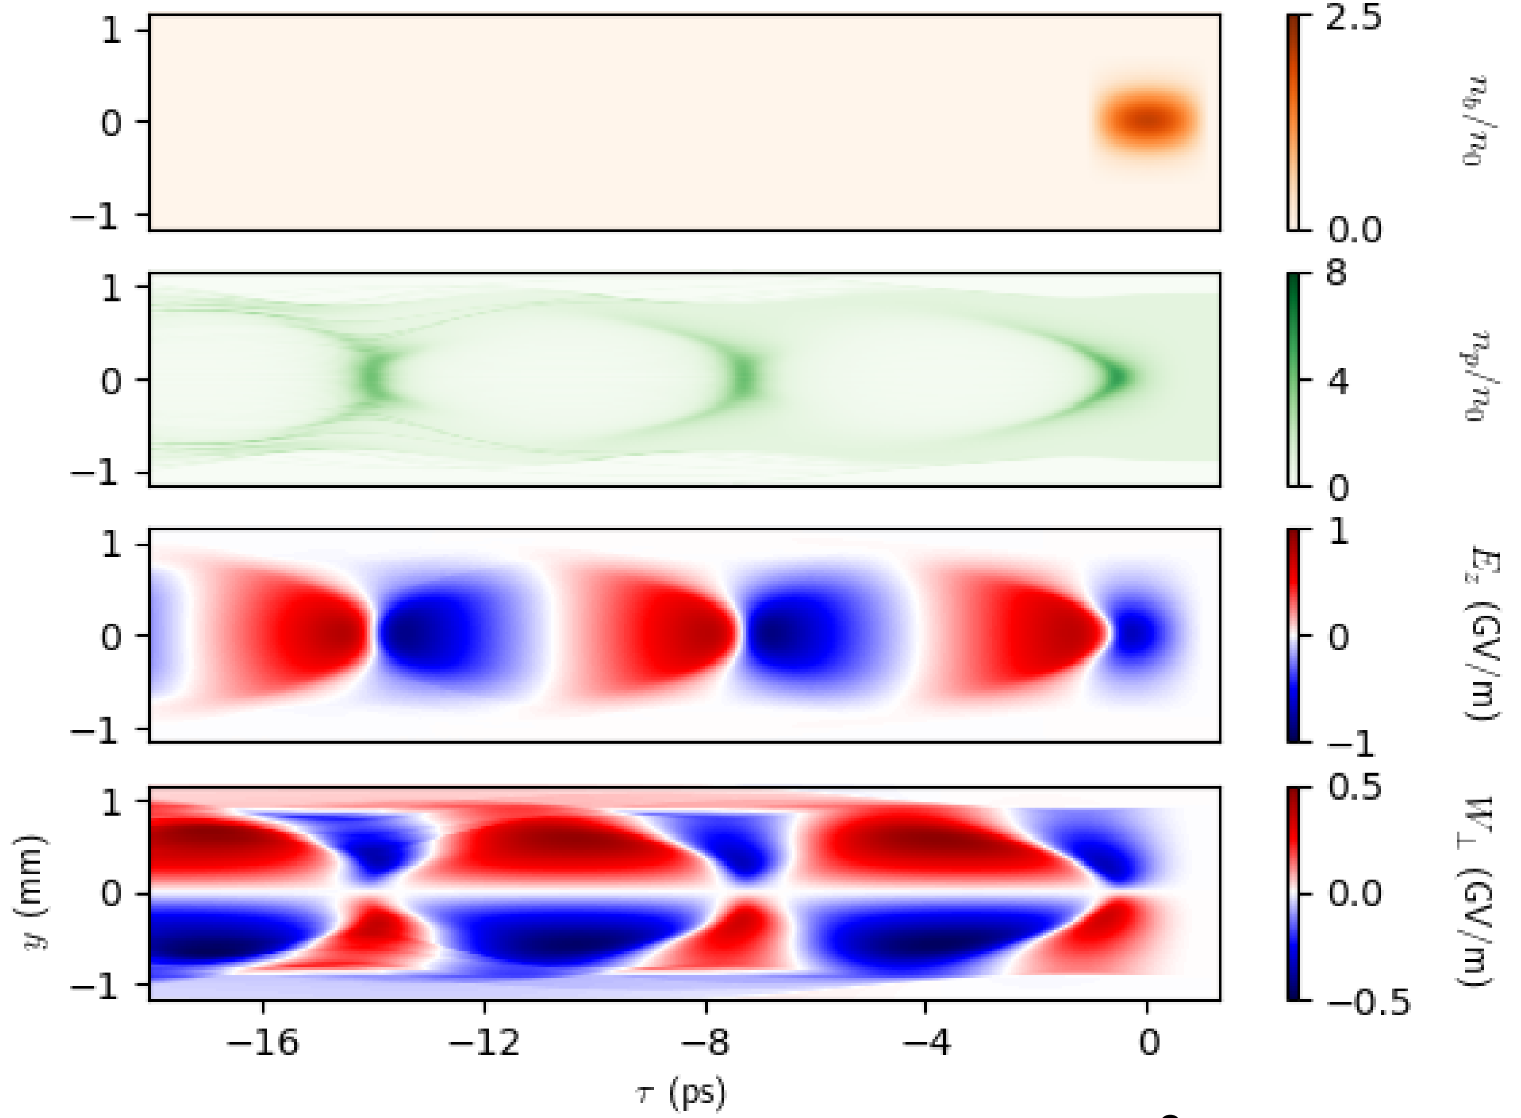


s=0 m

# Recap: Higgs factory

Initial proton driver chosen to generate suitable wakefields

Driver rapidly pinches (saturated plots)

Leads to highly nonlinear wakefields

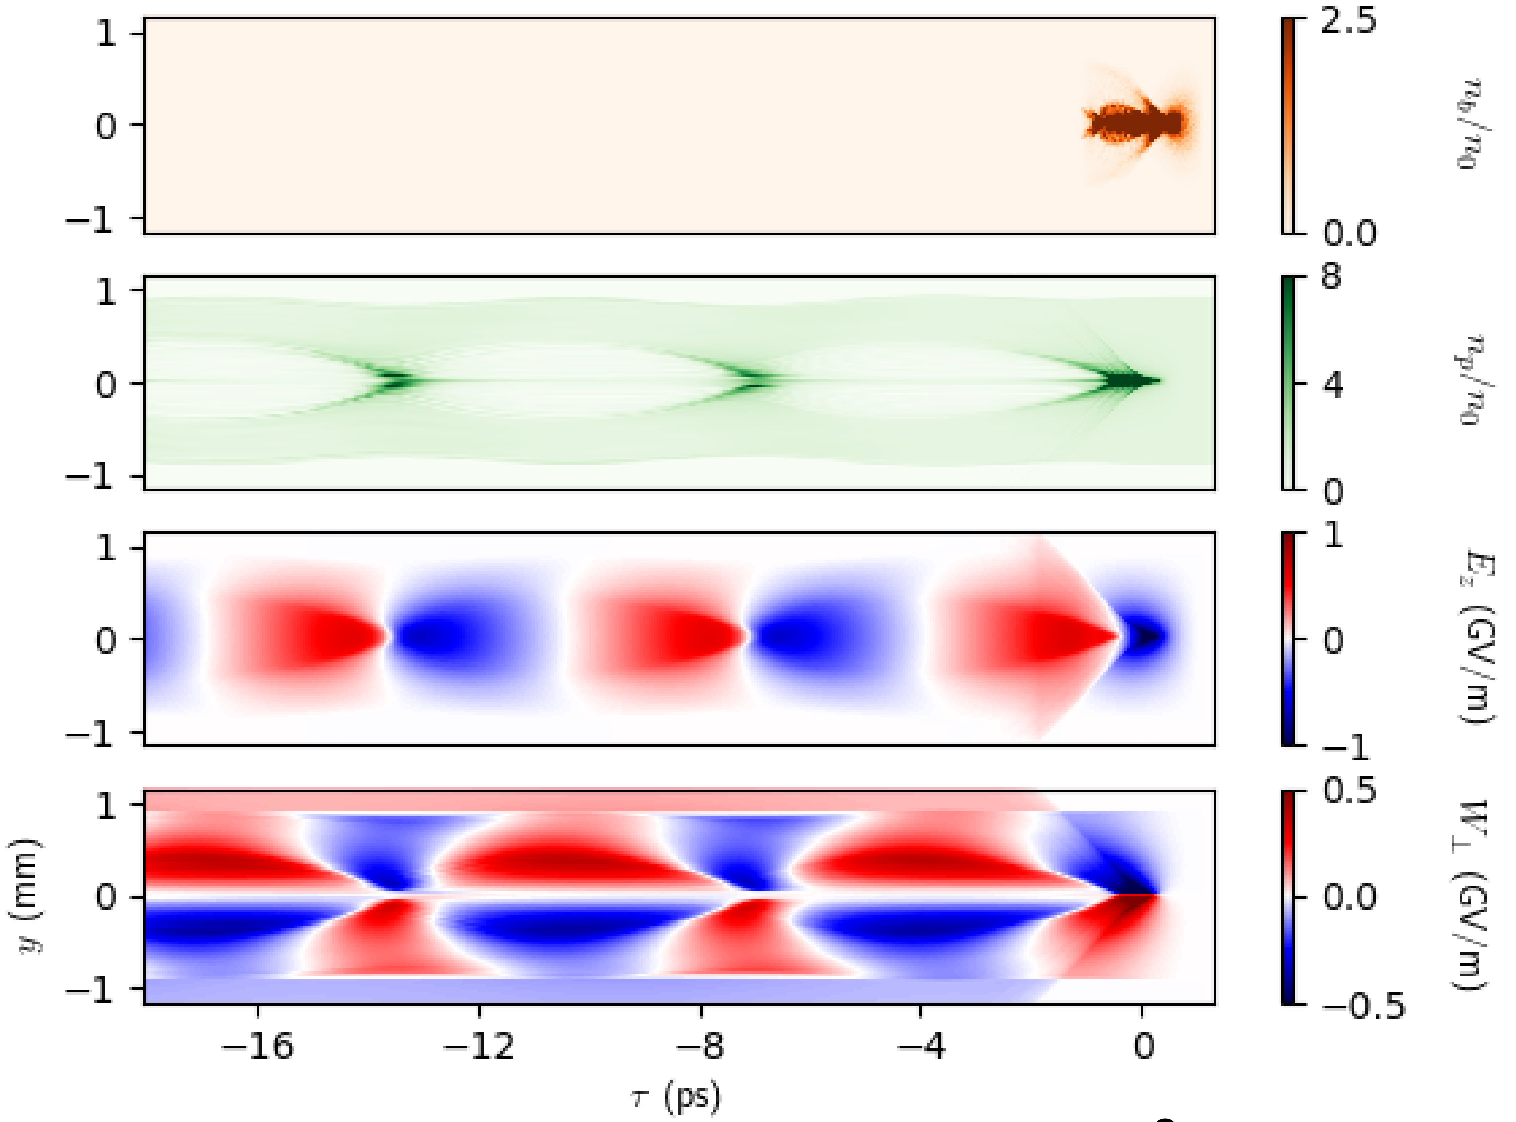


s=2 m

# Recap: Higgs factory

Initial proton driver chosen to generate suitable wakefields

Driver rapidly pinches

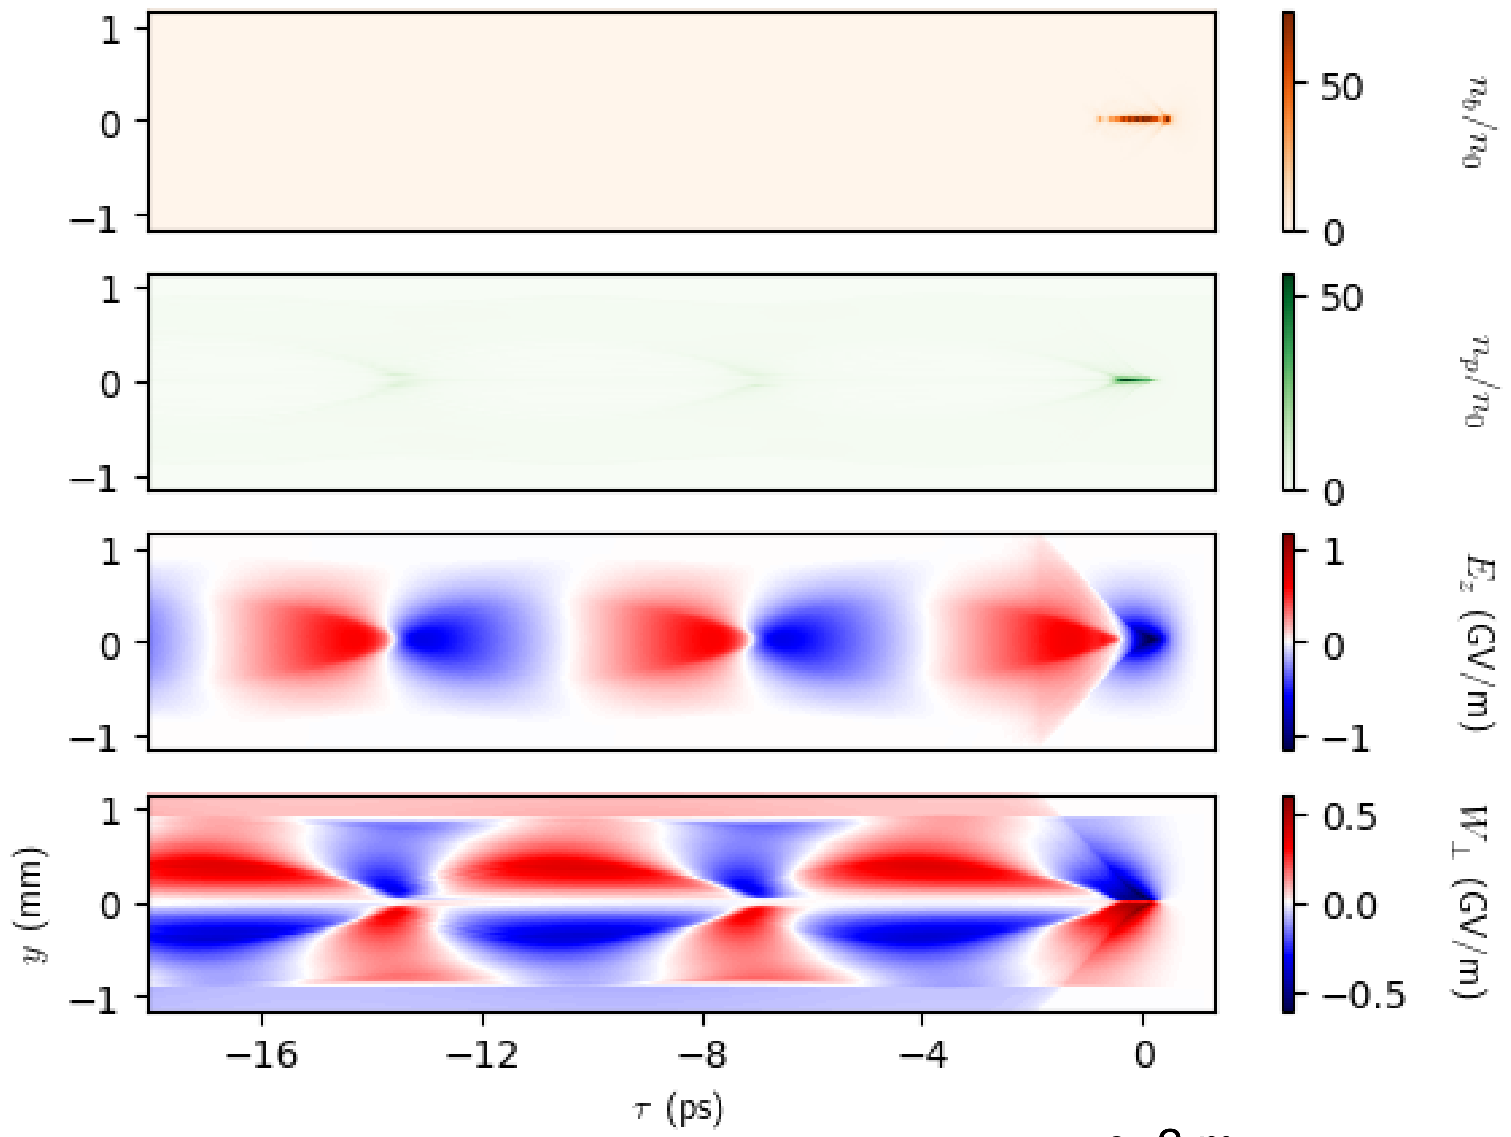


s=2 m

# Proton challenge #2

No blowout for positively charged bunches,
makes matching difficult.

Extreme pinch prevents positron acceleration,
nonlinear response reduces transformer ratio.

Good initial wakefields not sufficient,
need to counteract strong focussing wakefields.

# Proton solution #2

Tailored driver emittance to counter focussing fields.

- 2 μm at head
- initially constant
- rises linearly to 58 μm

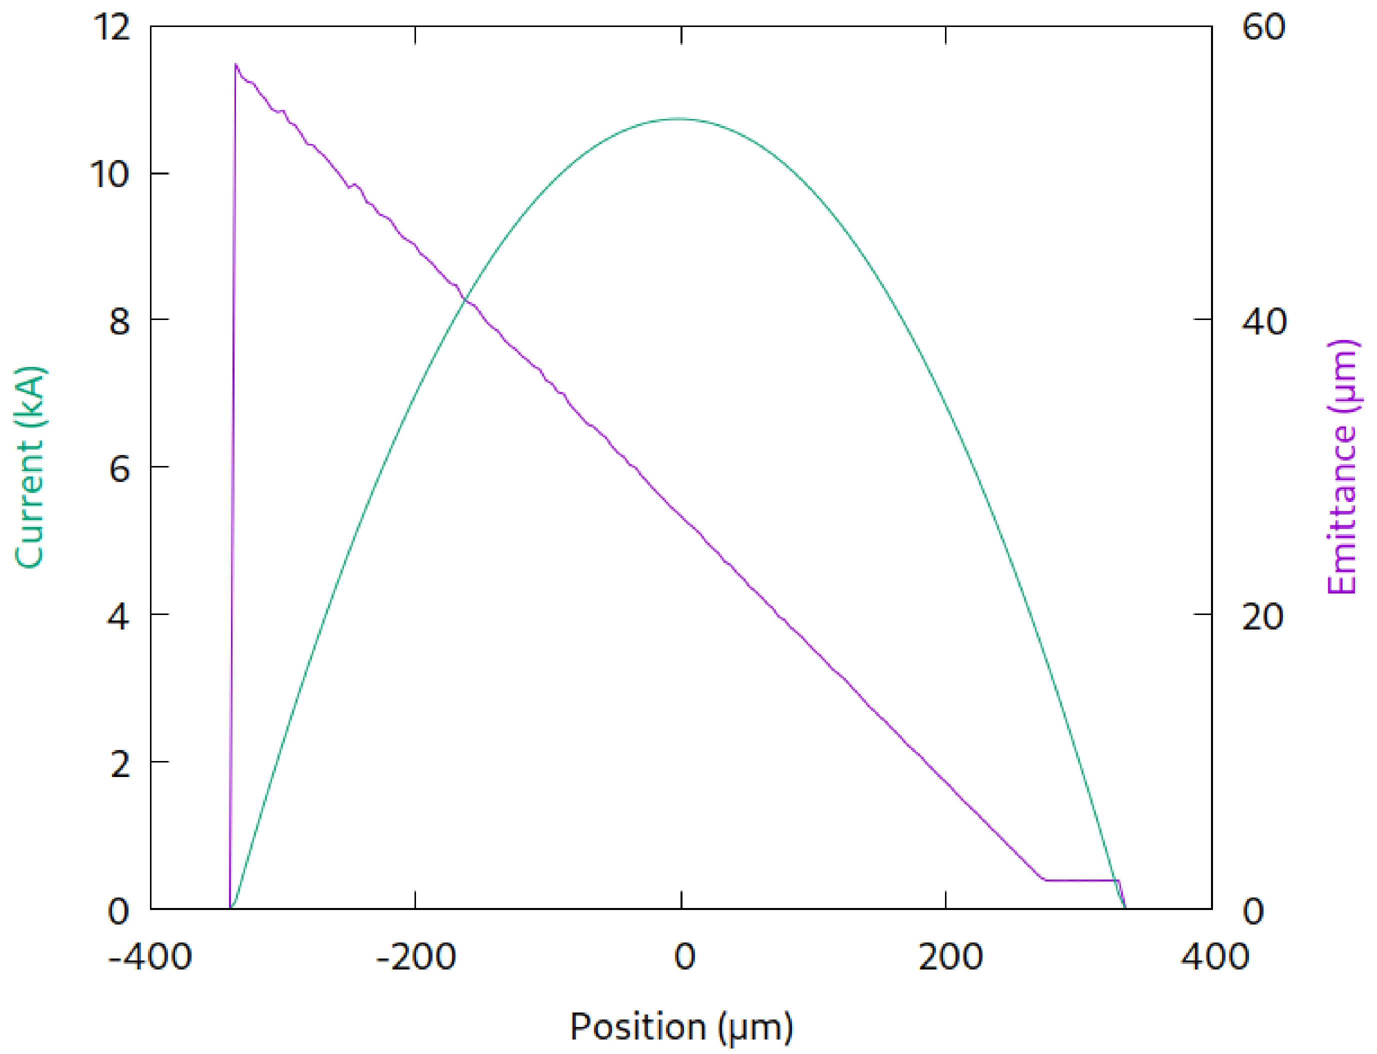

# Proton solution #2

How can we generate a tailored emittance profile?

Most likely:
with difficulty

BUT emittance is initially constant before growing monotonically

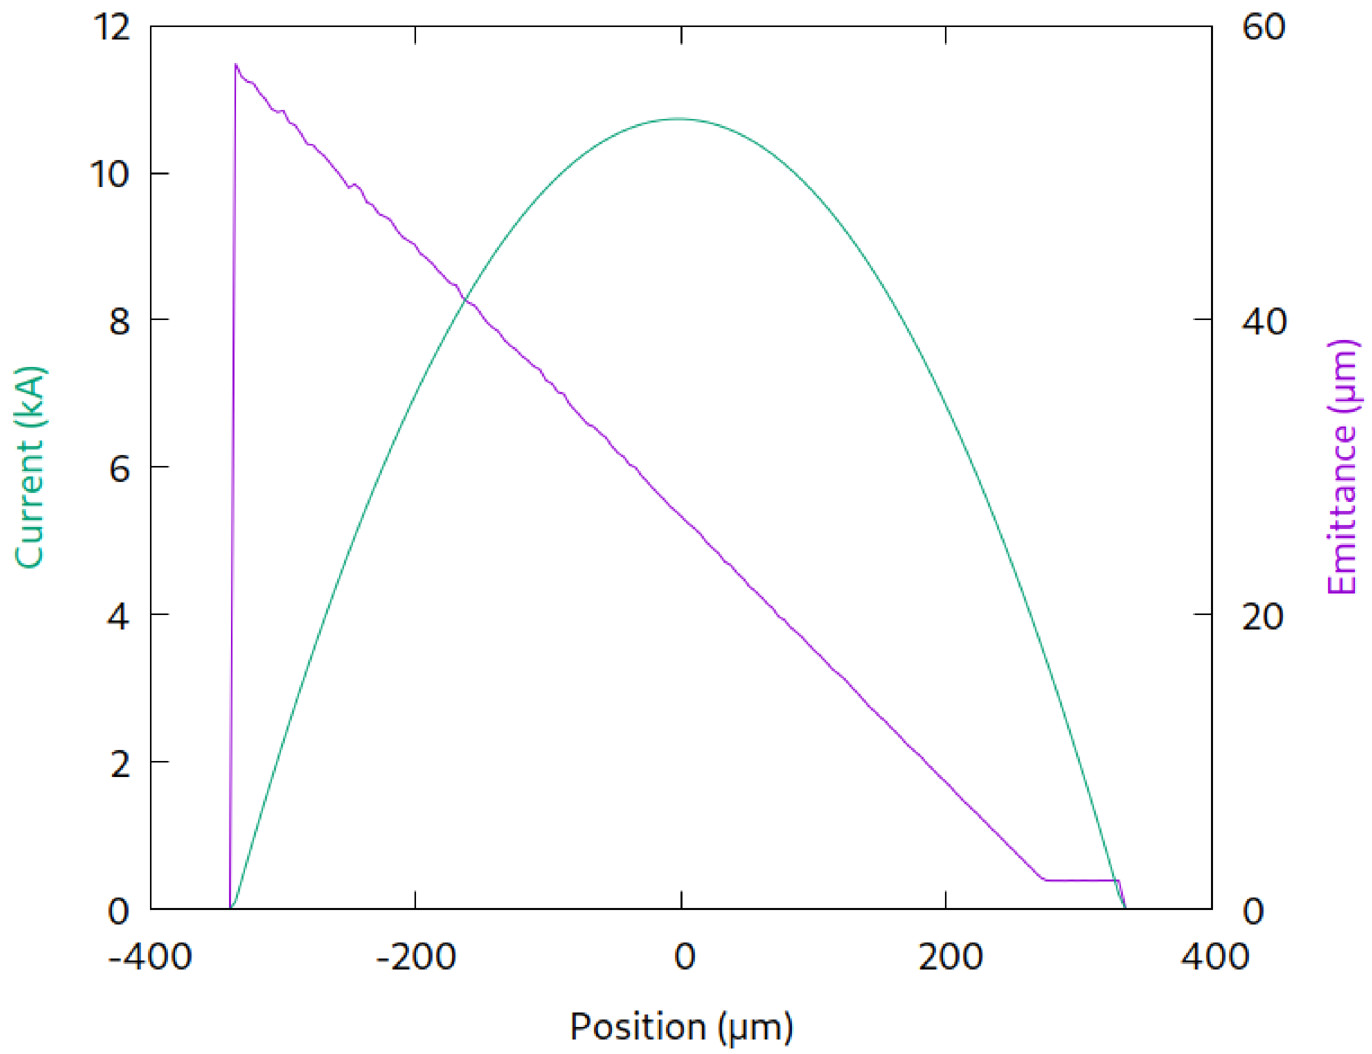

# Proton solution #2

How can we generate a tailored emittance profile?

Most likely:
with difficulty

BUT emittance is initially constant before growing monotonically

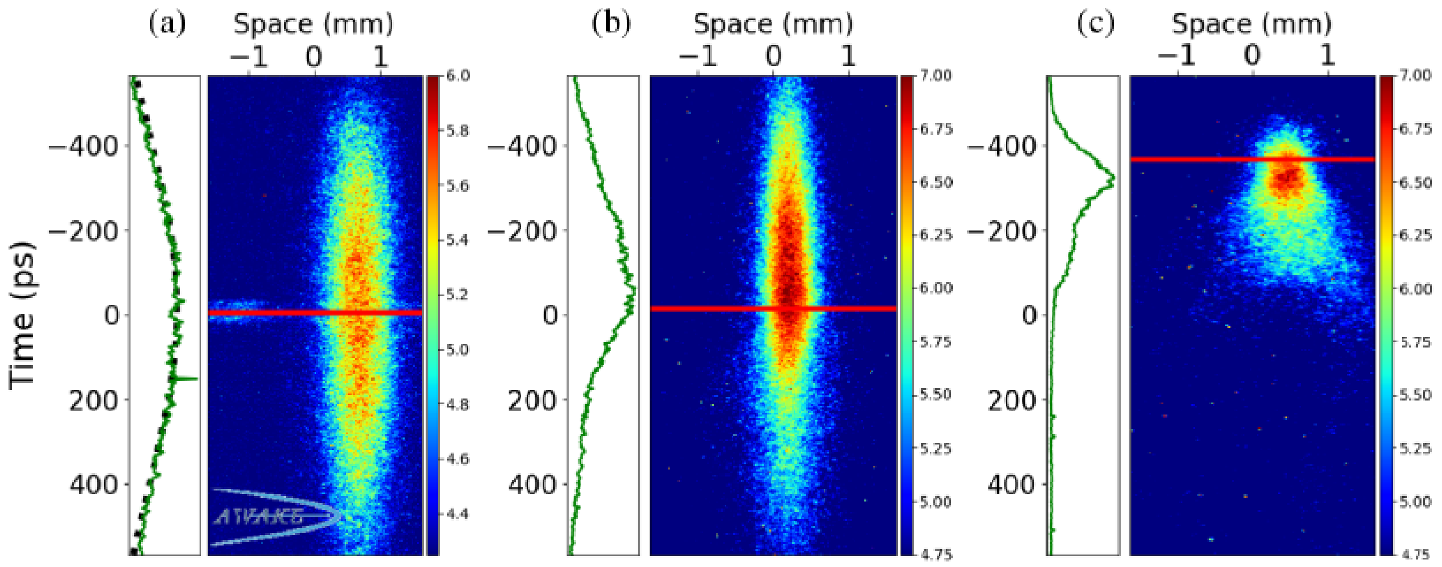


[AWAKE Collaboration, PRL (2019)]

Harness plasma instabilities?

# Proton solution #2

Initial proton driver chosen to generate suitable wakefields

Tailored profile stable for 2 m

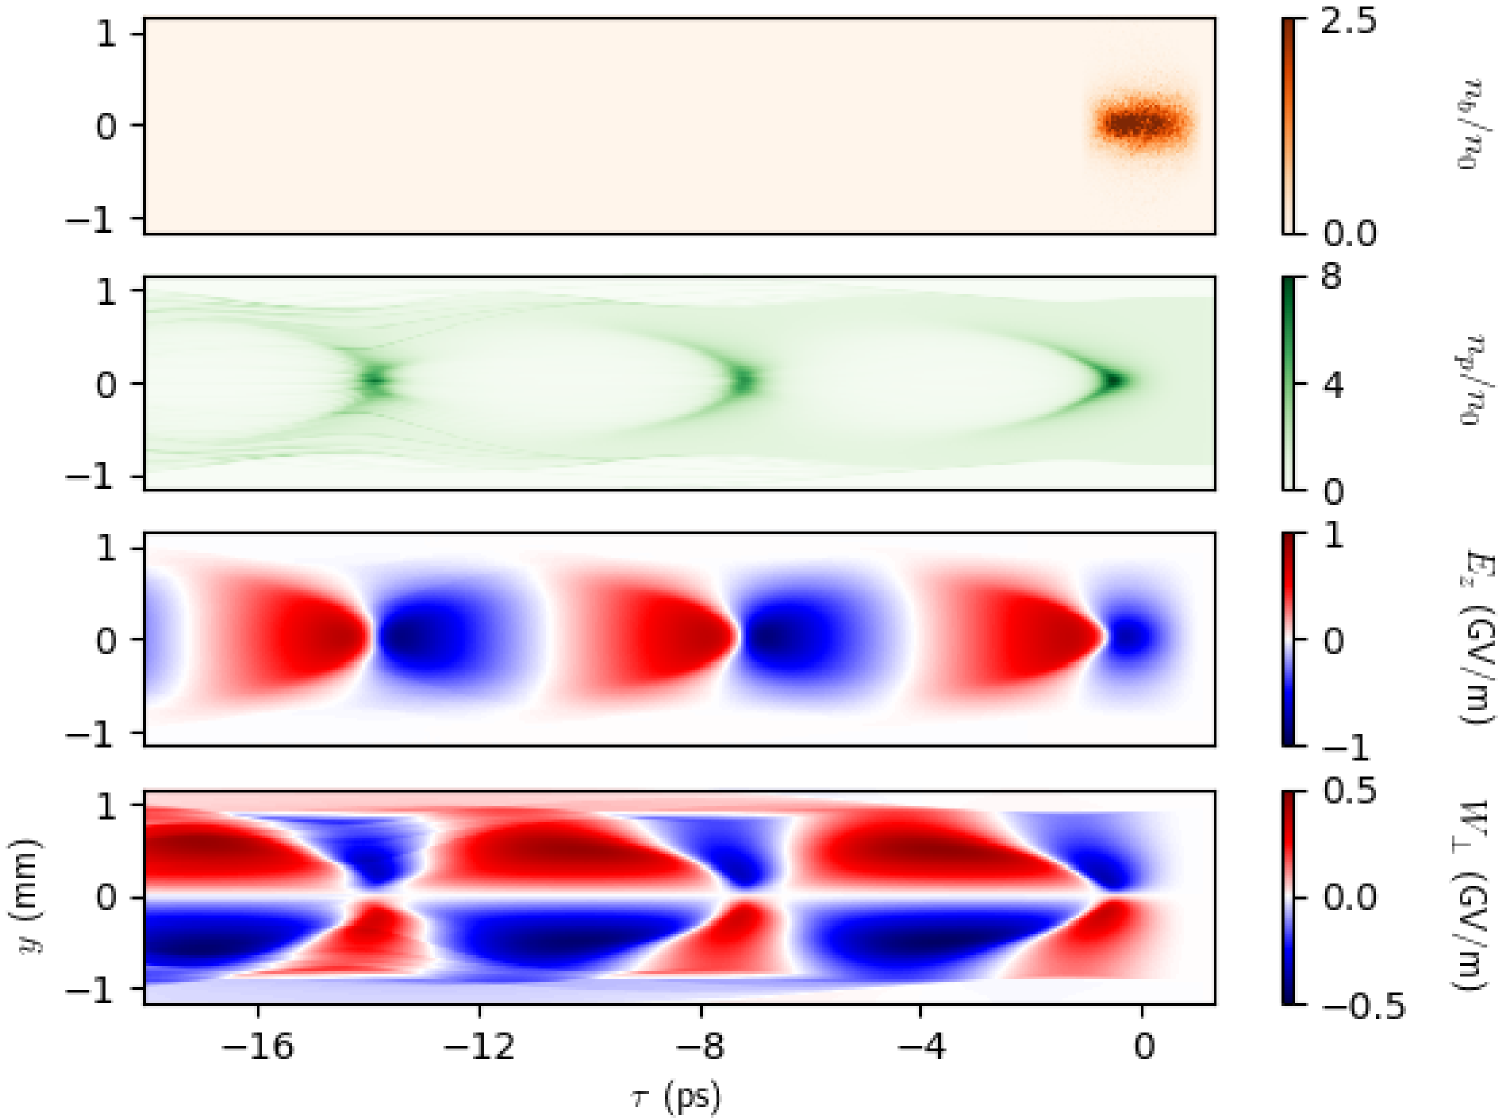


s=2 m

# Proton solution #2

Initial proton driver chosen to generate suitable wakefields

Tailored profile stable for 20 m

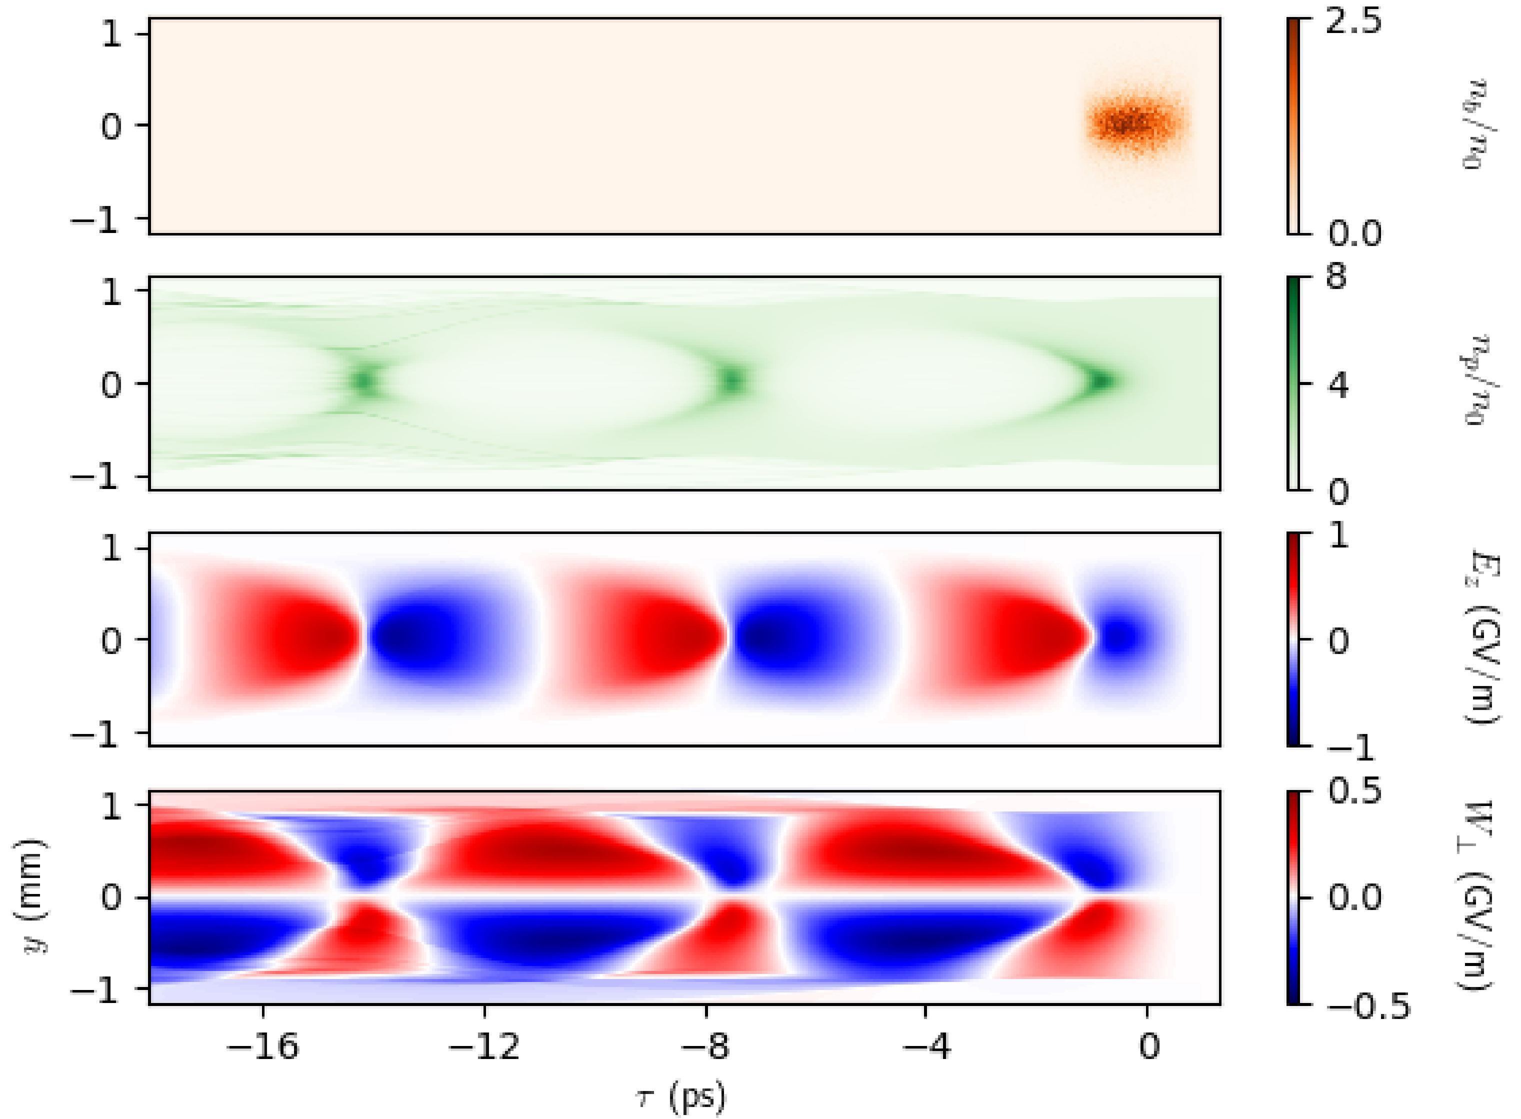


s=20 m

# Proton solution #2

Initial proton driver chosen to generate suitable wakefields

Tailored profile stable for 200 m

Acceleration limited by longitudinal dispersion.

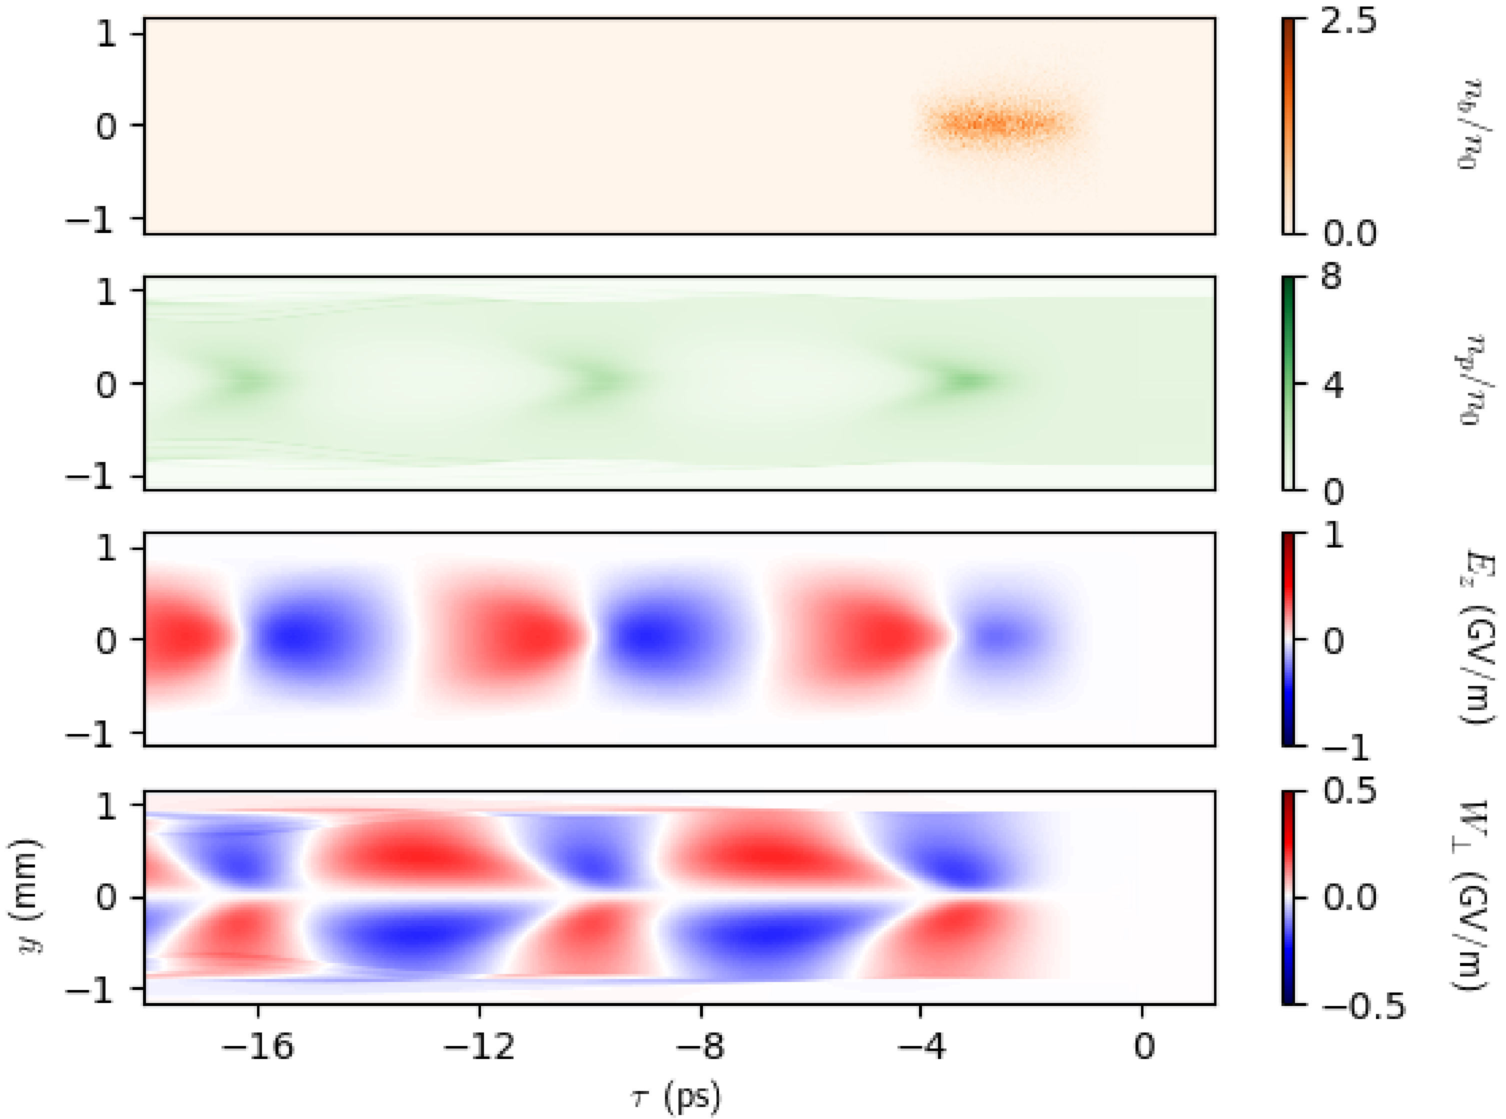


s=200 m

# Recap: Higgs factory

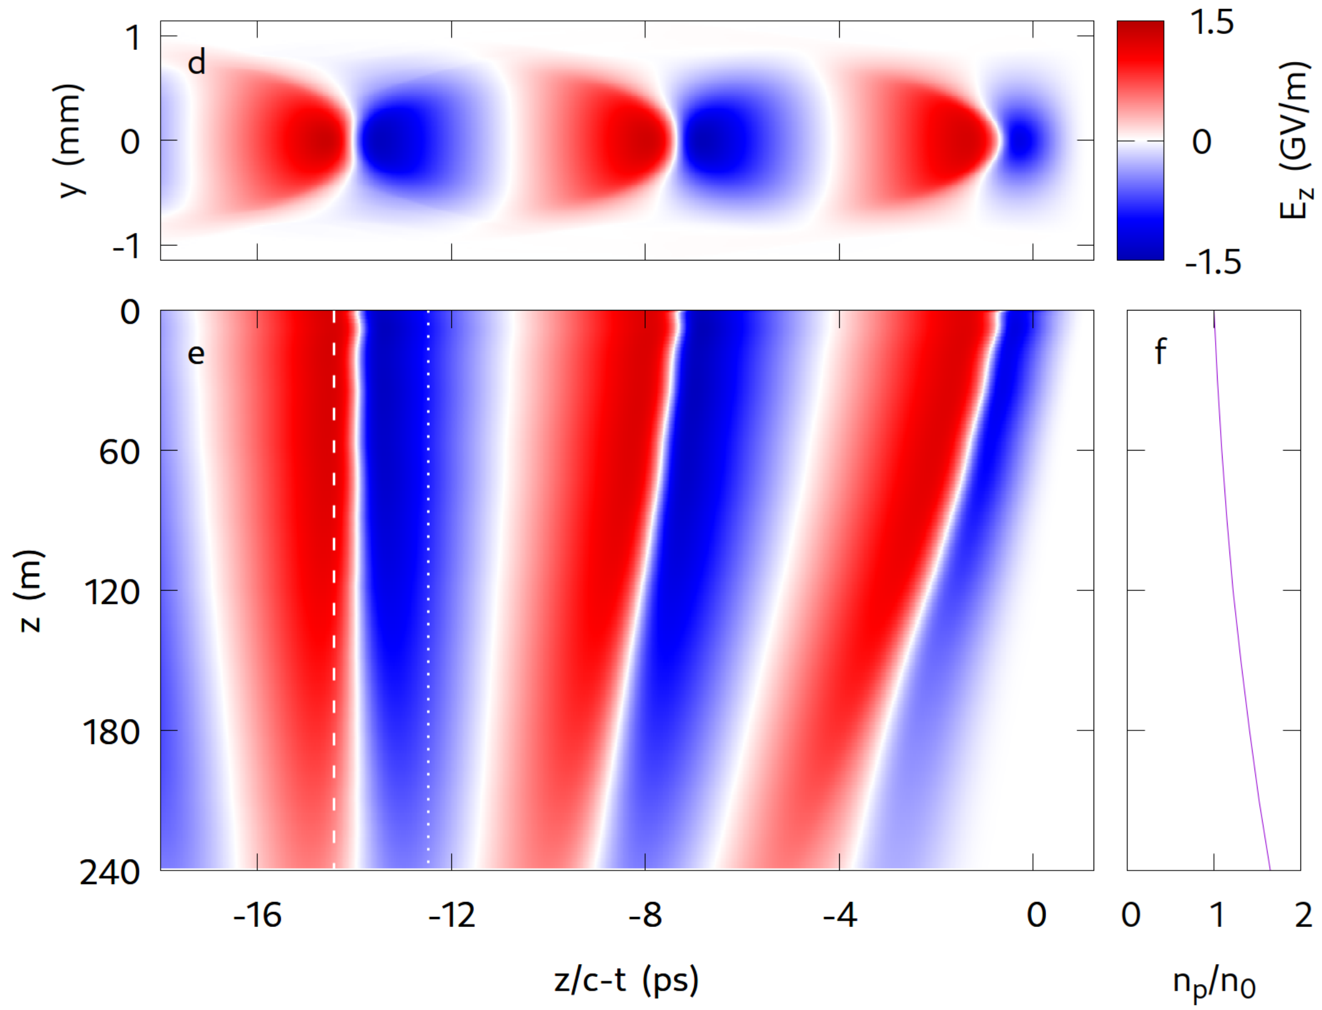


$10^{11}$ protons at 400 GeV
accelerate
$10^{10}$ e+/e- to 125 GeV.

Farmer et al. (2024)

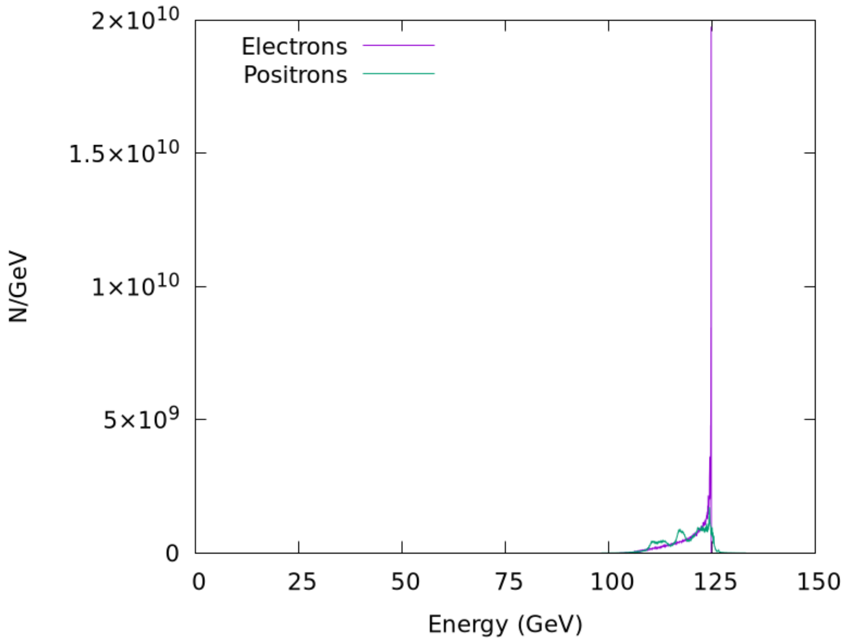


Dephasing avoided using density gradient.

# Recap: Higgs factory

Paper provides a proof-of-principle, and identifies key issues:

- need appropate driver, tailored emittance
- need scheme for high-quality positrons
- need energy recovery, cooling, witness beam generation [...]

Parabolic driver and witness, plenty of room for optimisation:

- shape driver – reduce longitudinal dispersion, higher efficiency from driver to plasma.
- shape witness – control beamloading, higher efficiency from plasma to witness.

# Witness self-matching

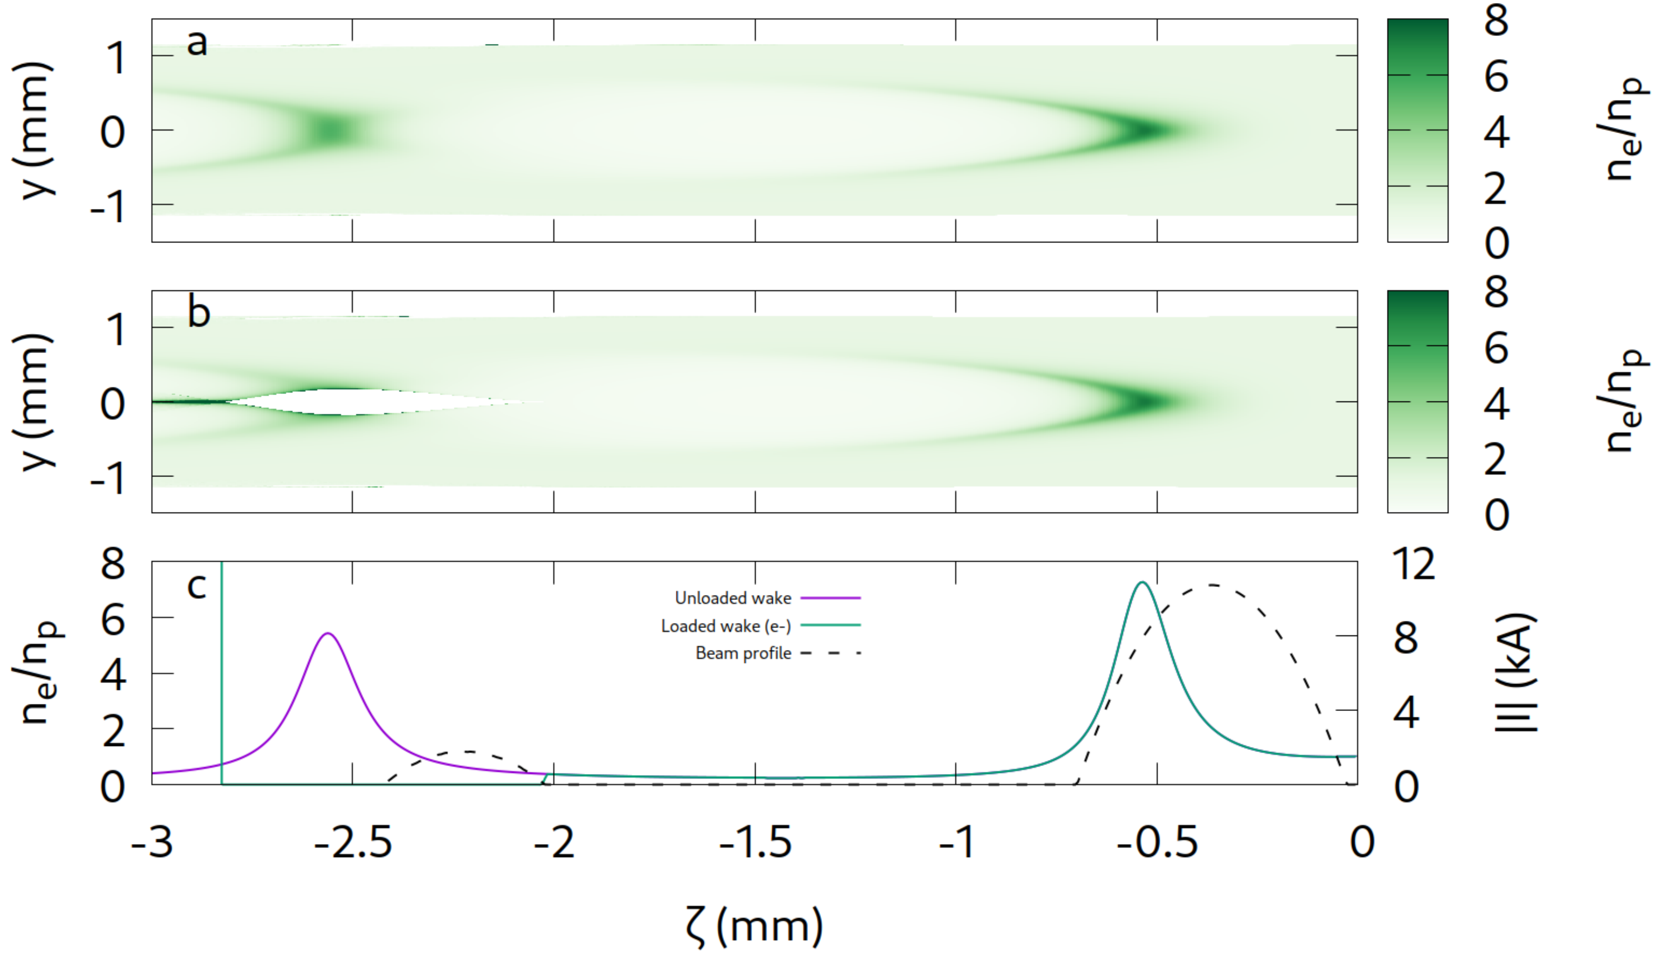


2D geometry (LCODE), frozen driver, electron witness.

# Witness self-matching

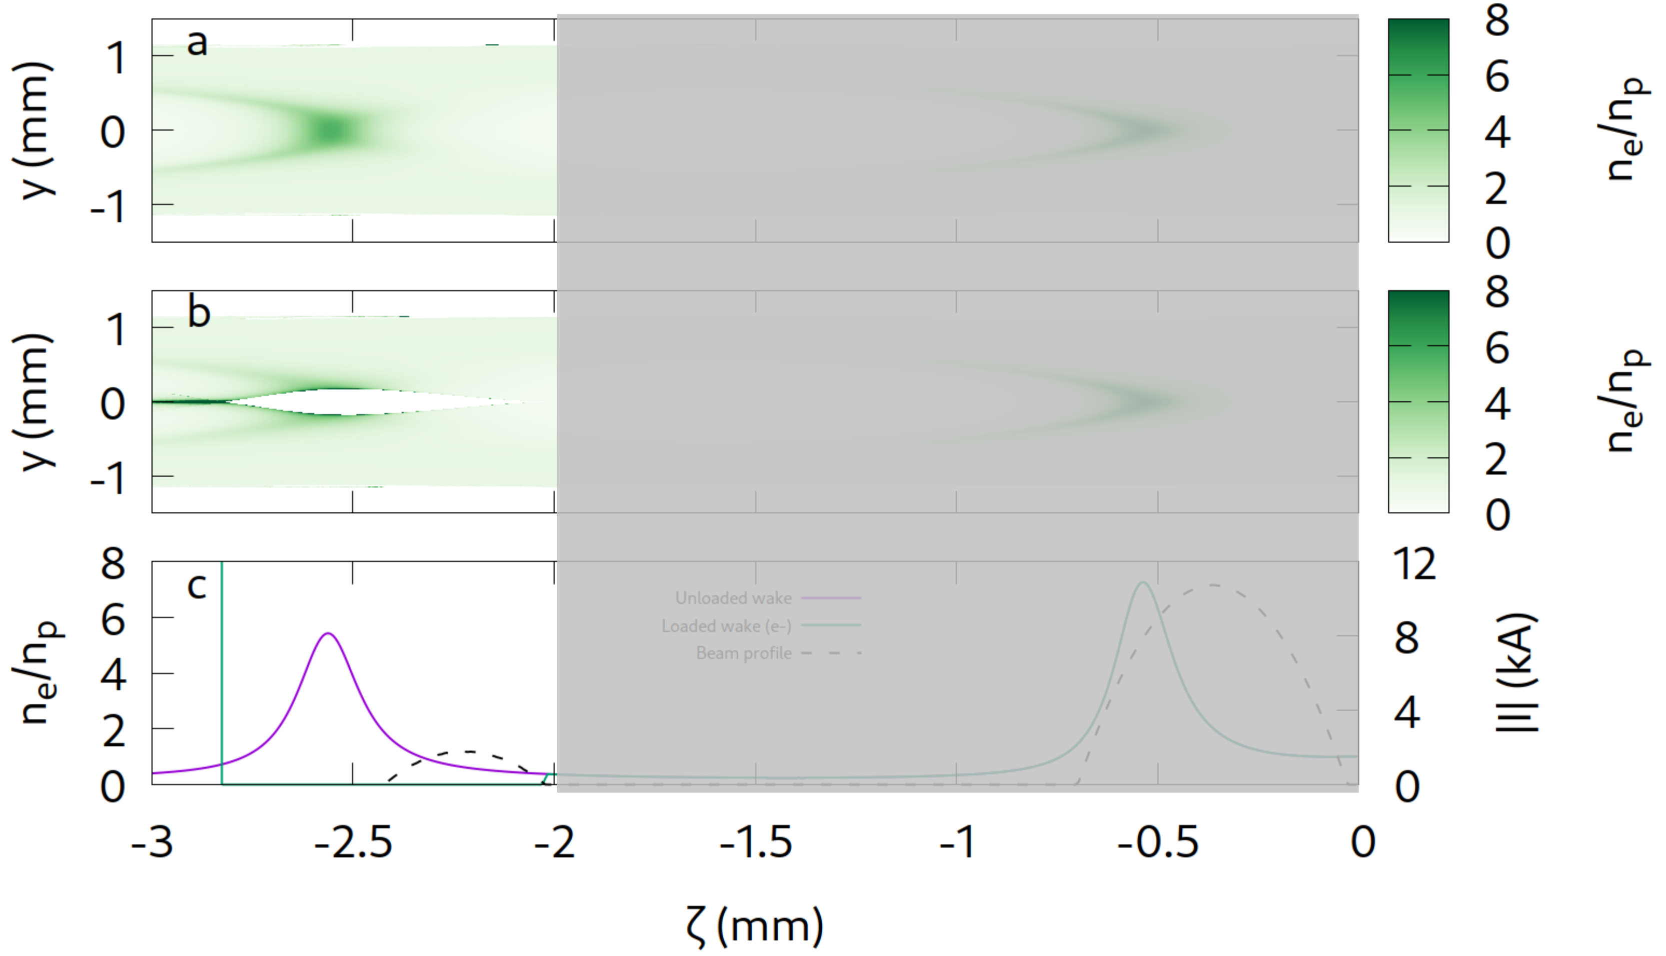


2D geometry (LCODE), frozen driver, electron witness.

# Witness self-matching

Energy gain
(trivial for frozen driver)

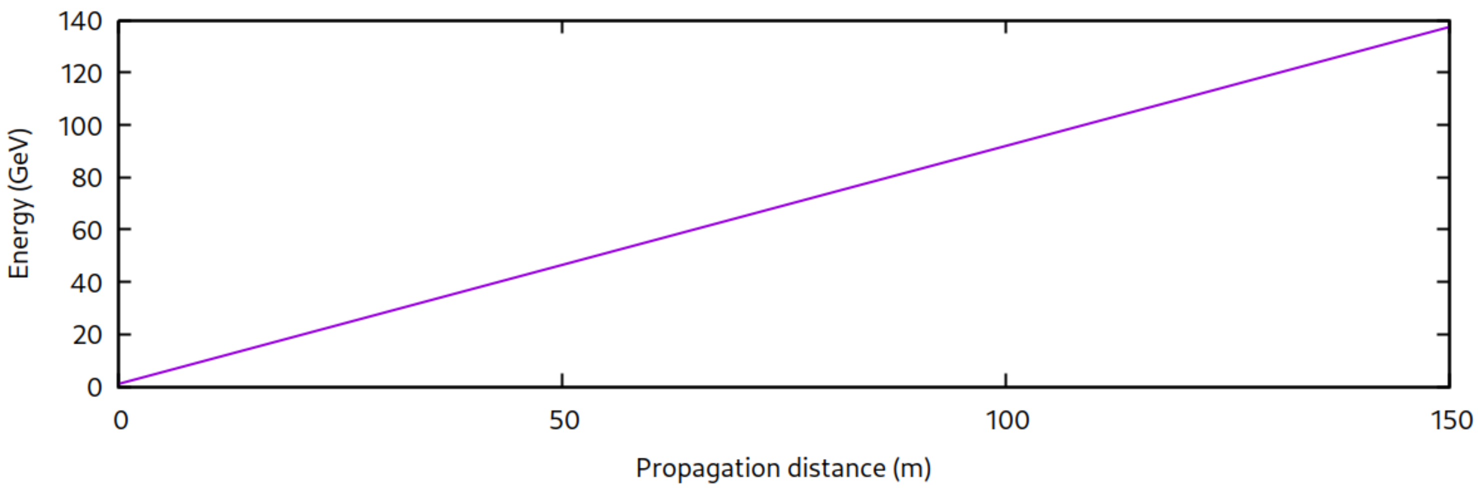


Emittance growth
due to ion motion
(lithium)

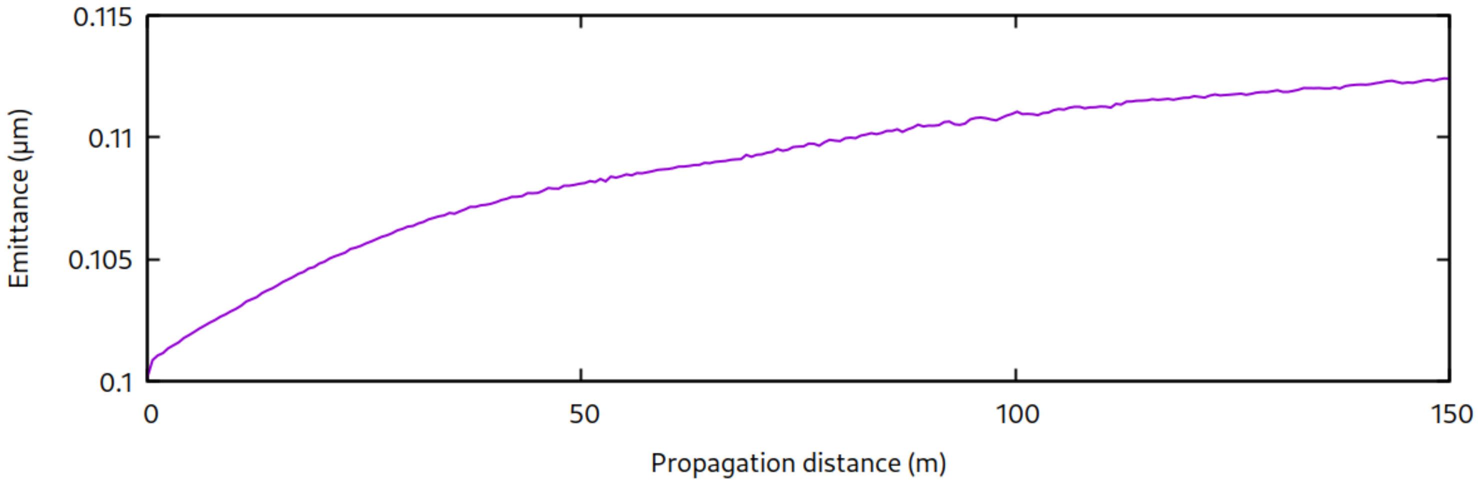

# Witness self-matching

Adiabatic focussing of witness during acceleration
1µm → 0.23 µm

Focussing field becomes increasingly nonlinear

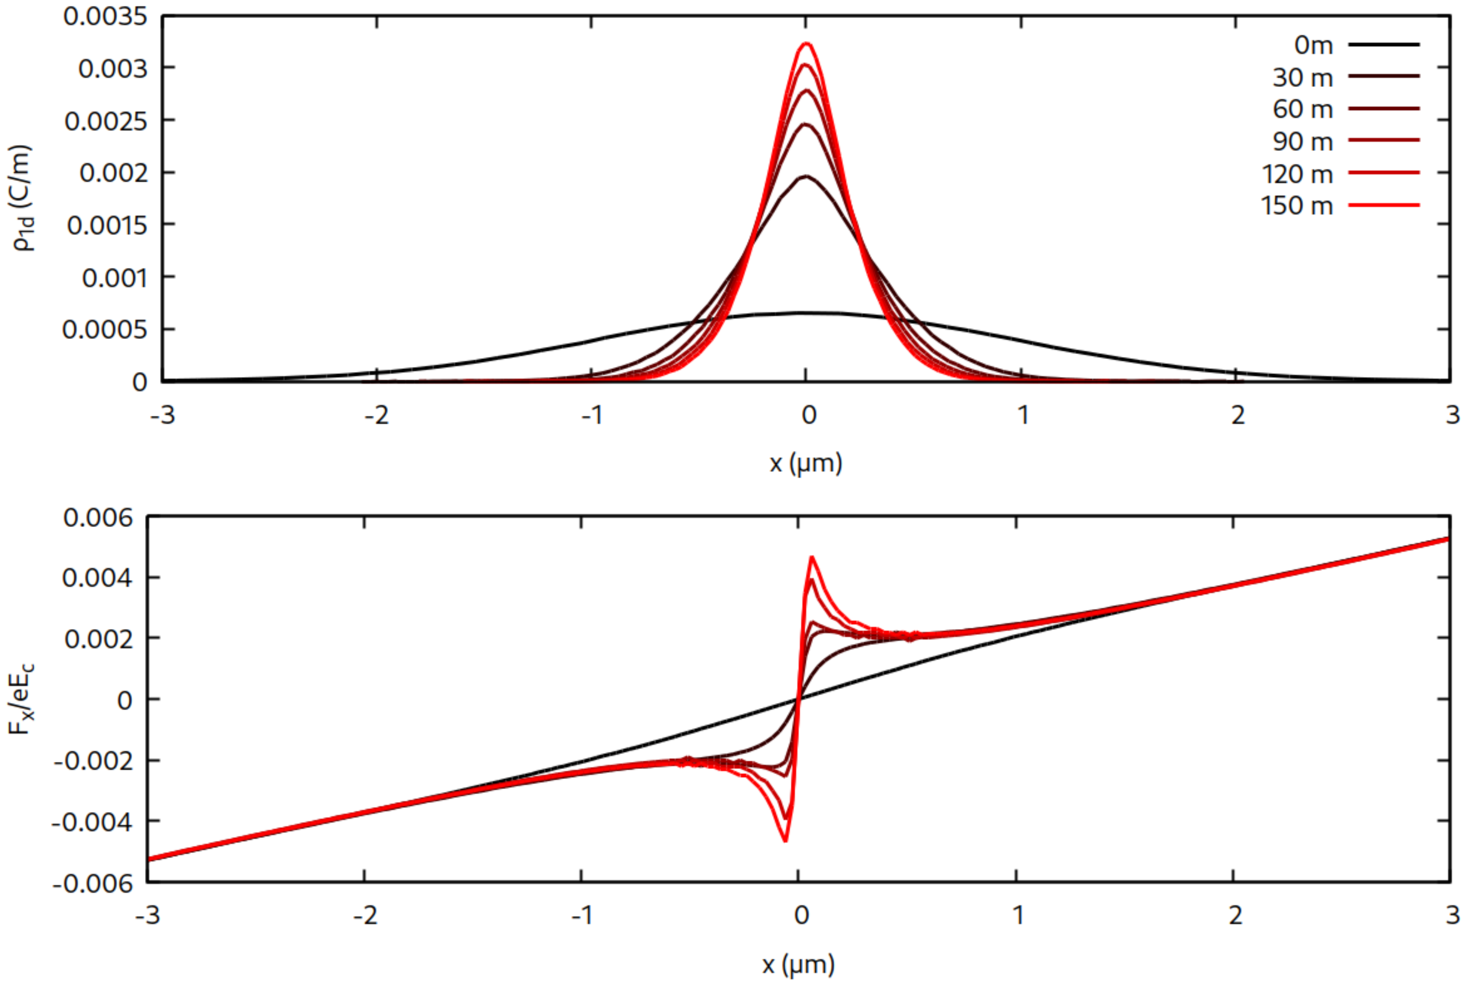

# Witness self-matching

Different longitudinal slices of the witness have different profiles

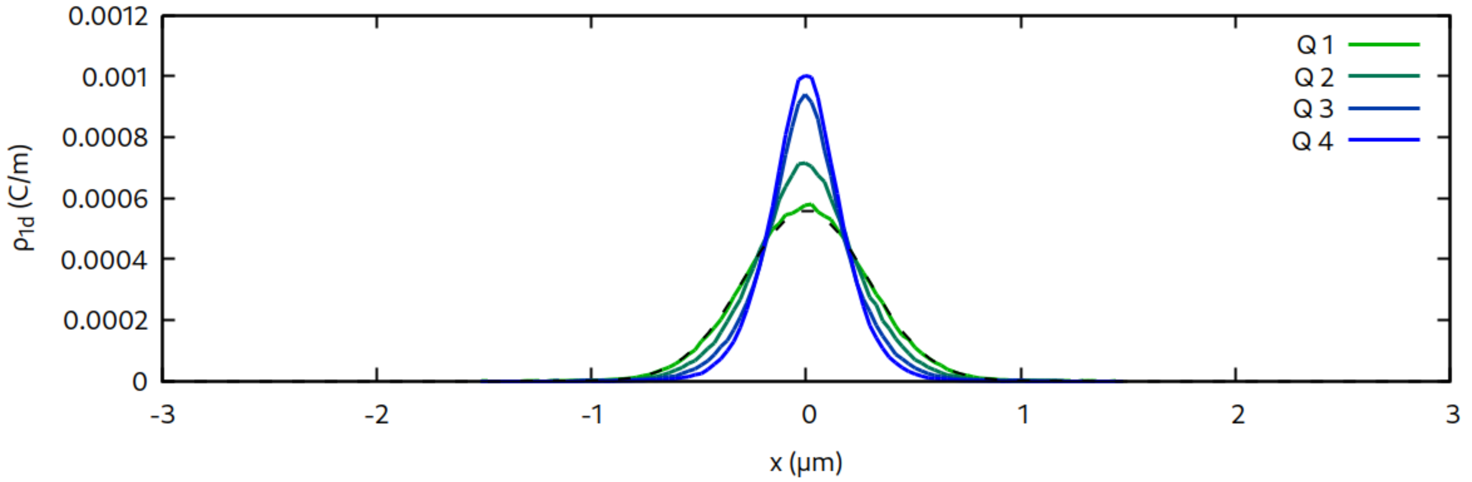


Head-to-tail variation of focussing fields.

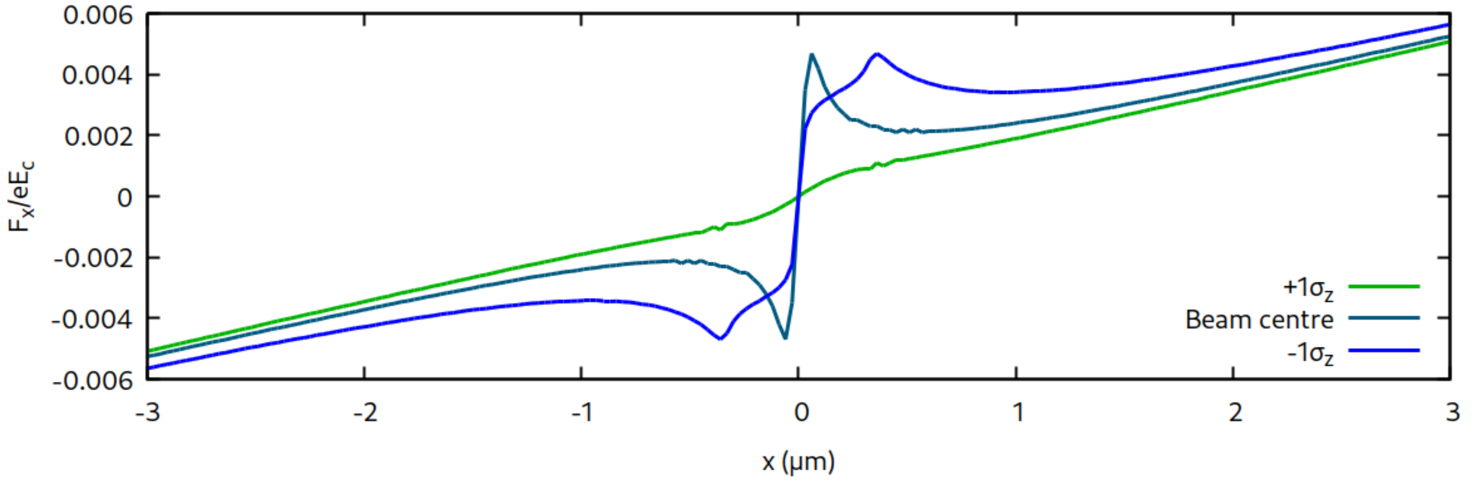

# Witness self-matching

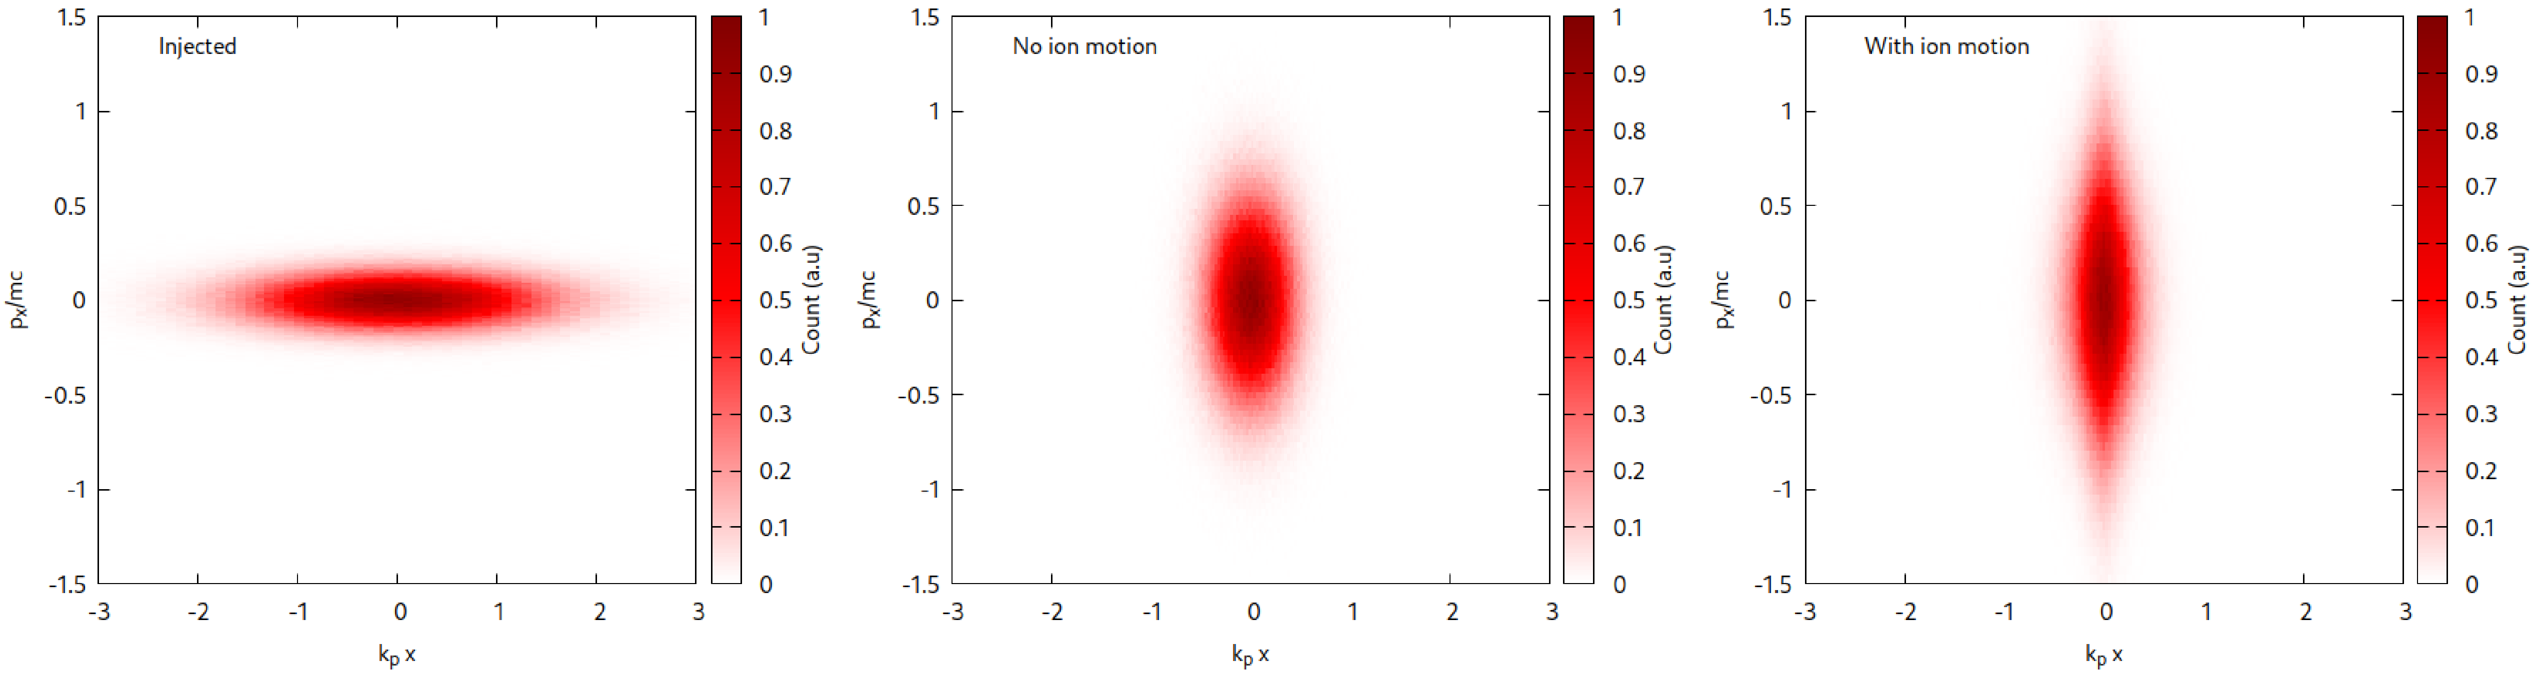


Suggests adiabatic focussing allows witness to self-match to nonlinear focussing fields

# Witness self-matching

Acceleration in a single stage has an unforseen benefit:

Witness bunch can match to the slow evolution of the focussing fields due to ion motion.

Results a head-to-tail variation in the transverse profile.

Farmer et al. (2024b)

# Scaling to high energy

Acceleration limited by longitudinal dispersion of the drive beam.

Witness energy grows *better than linearly* with driver energy.

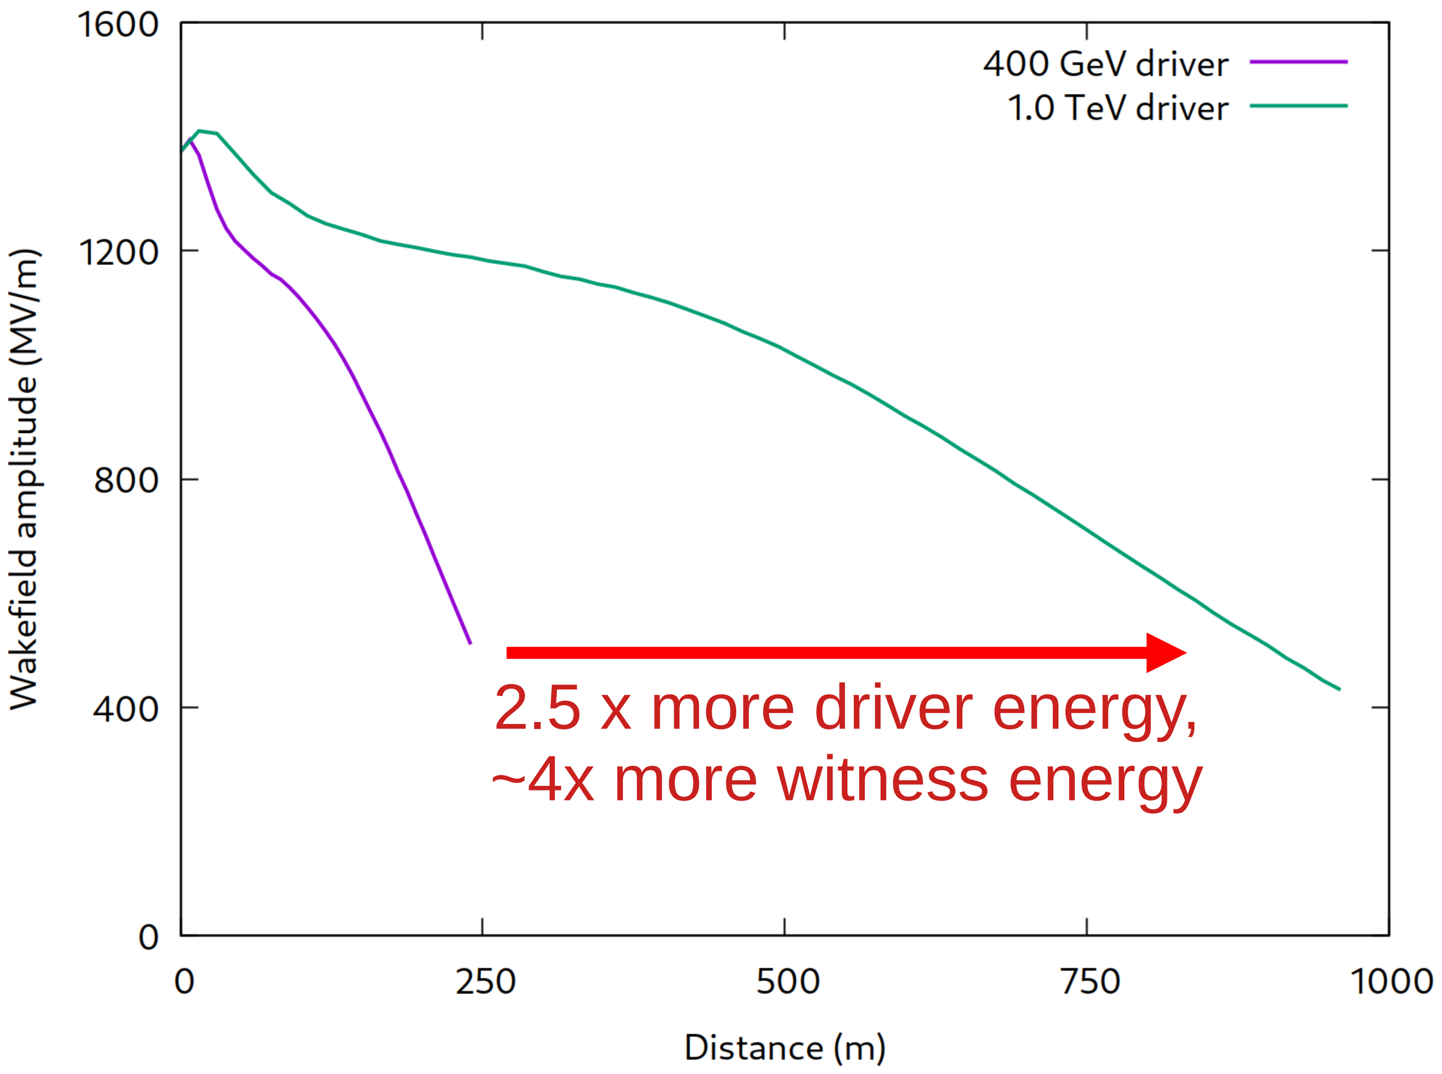

# Scaling to high energy

At low (up to 1 TeV) driver energy, acceleration length is limited by longitudinal dispersion of the drive beam.

- better than linear scaling with energy ($\gamma^{3/2}$)
- for higher driver energy, erosion of the driver head plays a larger role.

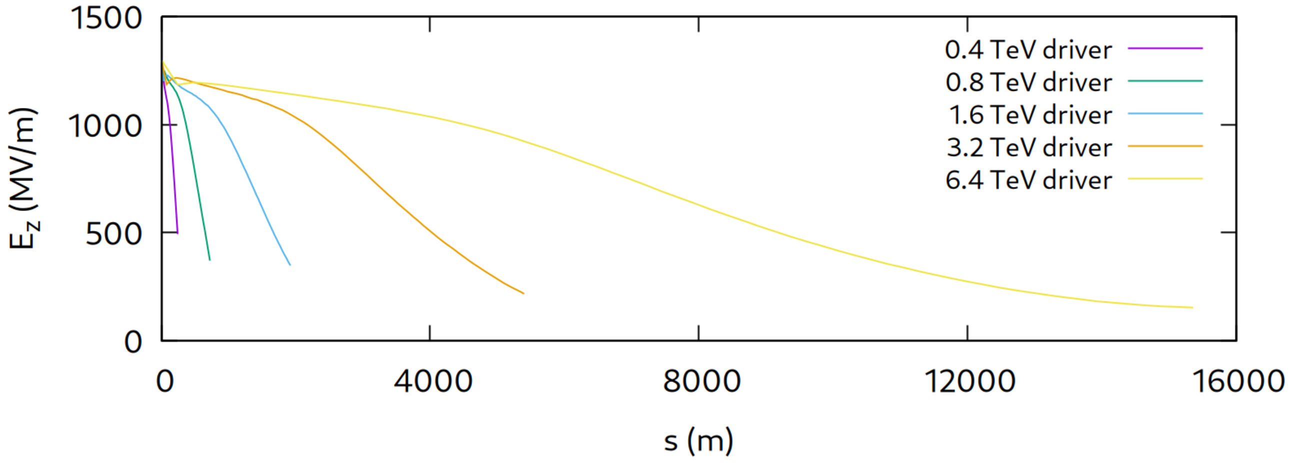


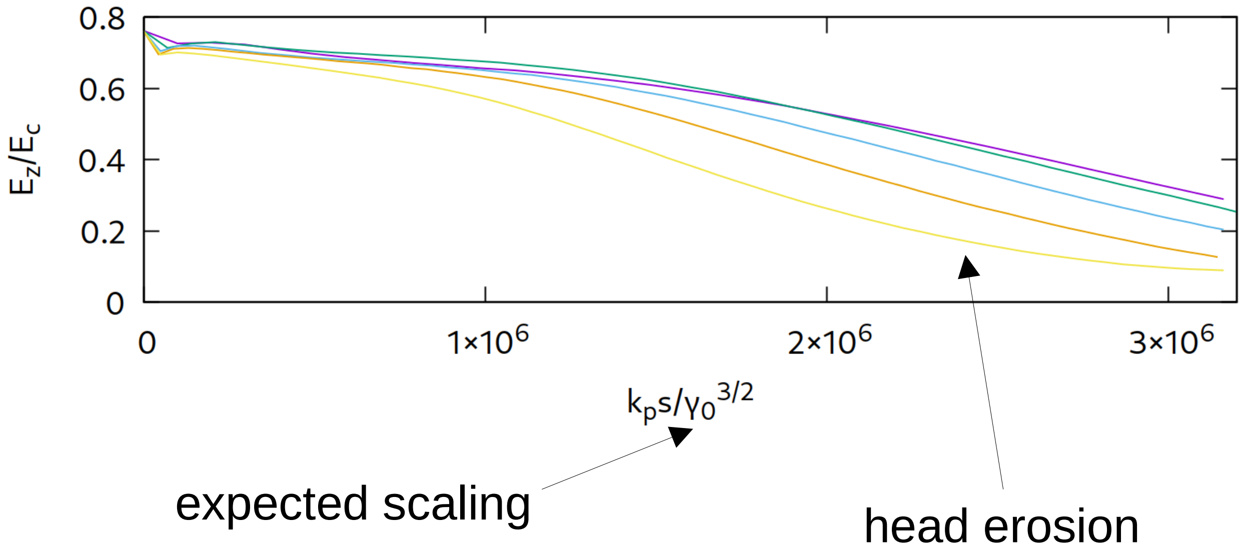

# Scaling to high energy

6.8 TeV driver (LHC energy) allows acceleration of a witness bunch to 5 TeV.

Caveats apply.

Paths for optimisation

- shape driver
- shape witness
- external focussing

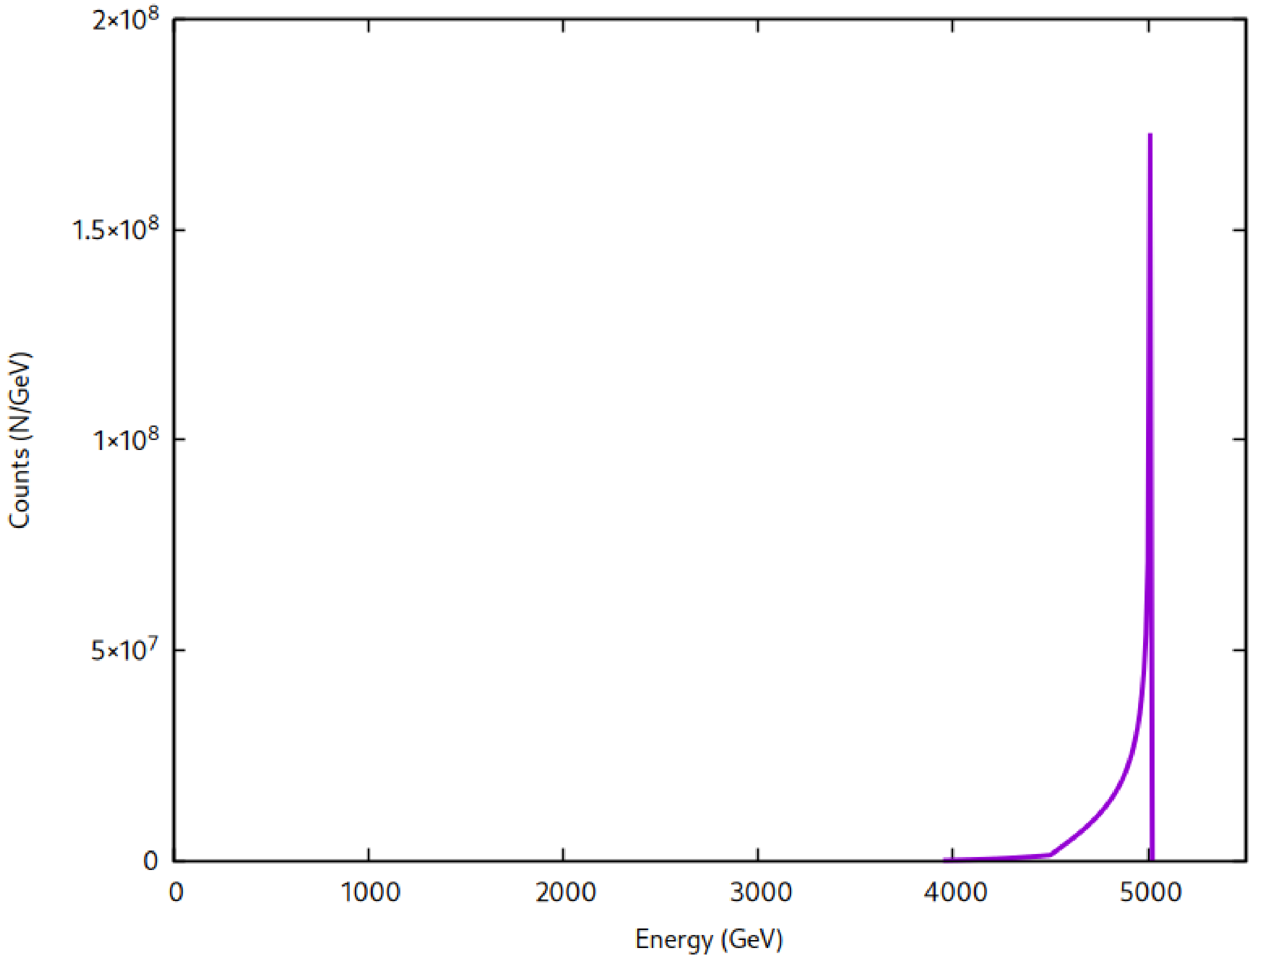

# Scaling to high energy

- CLIC at CERN

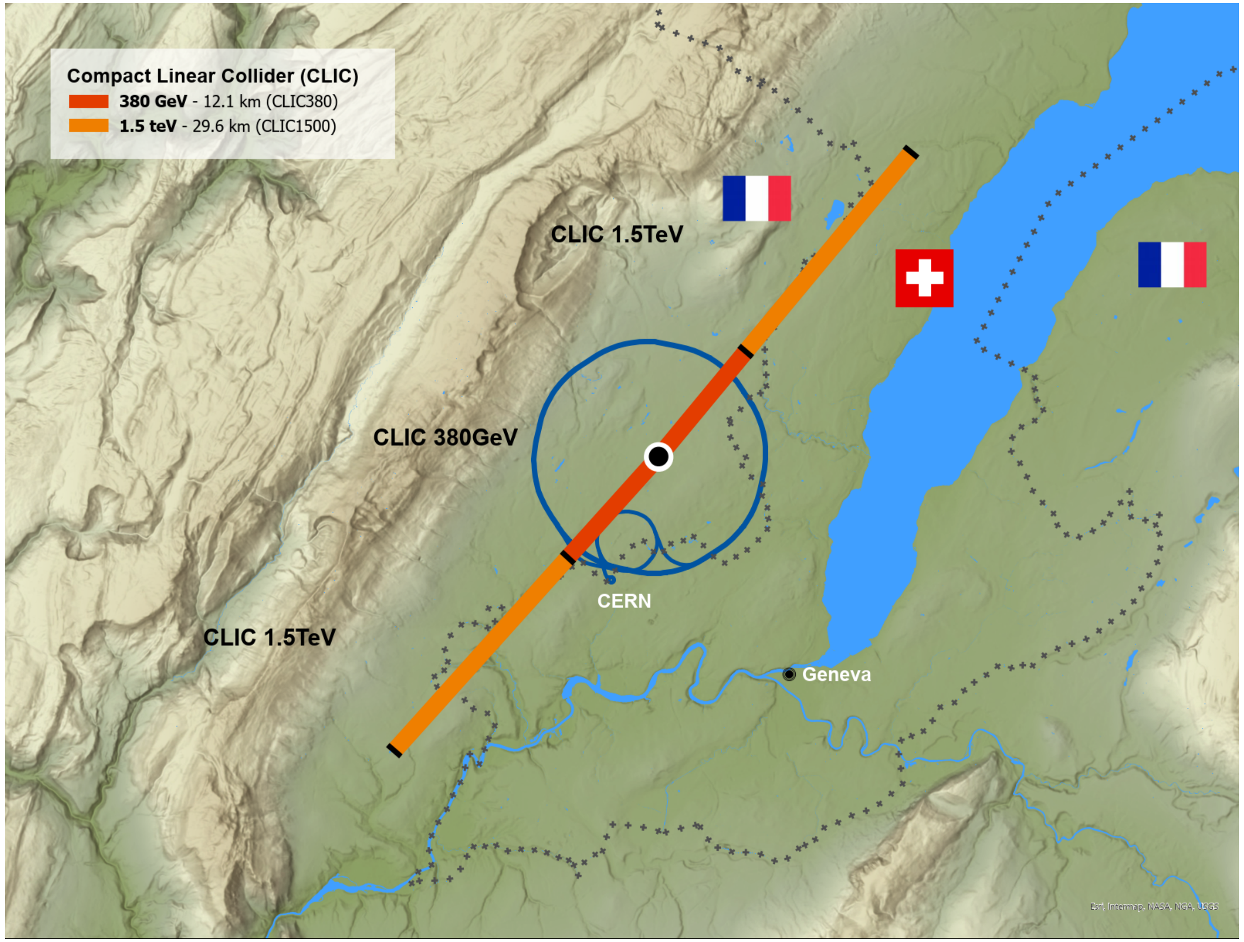


- ALiVE upgrade path

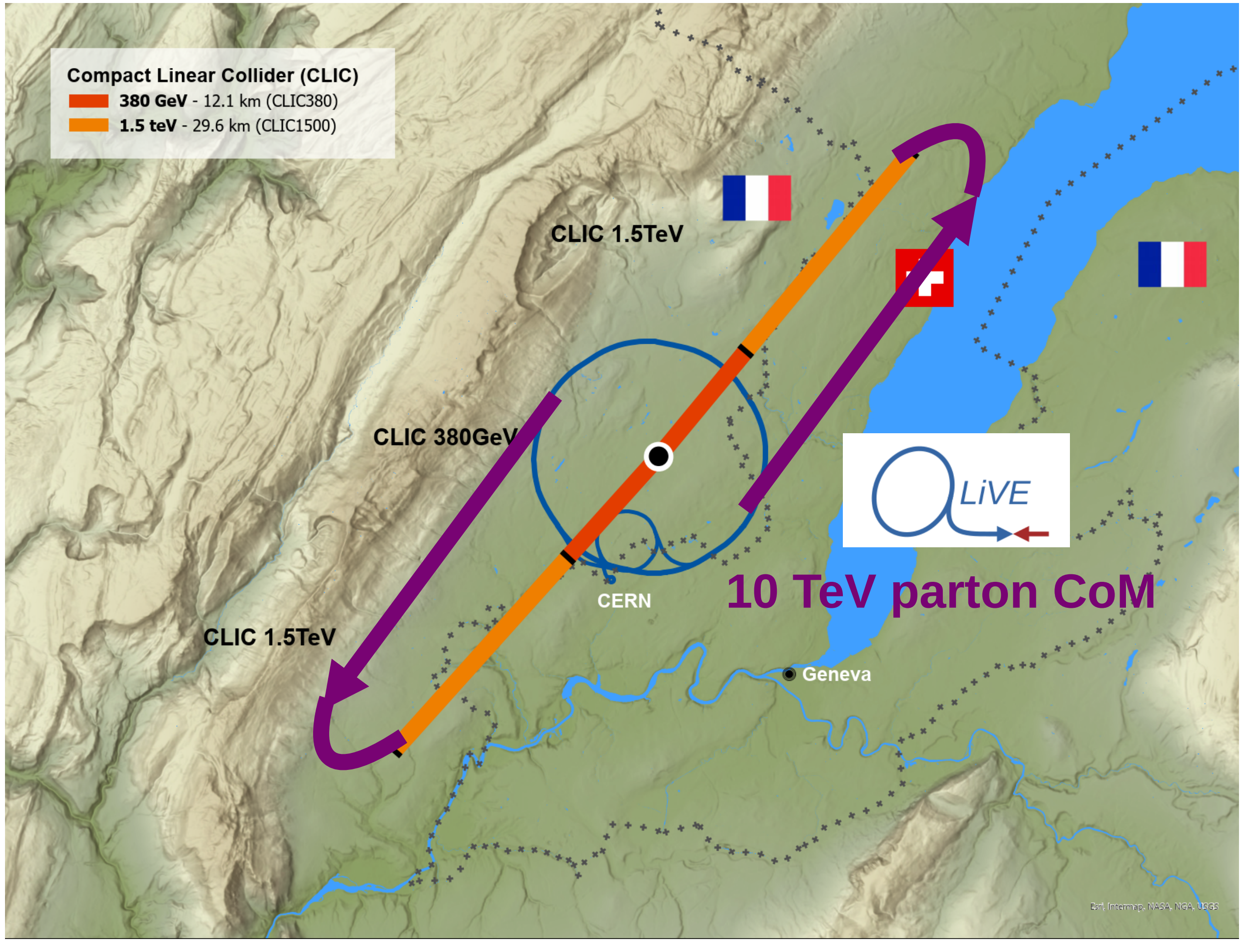

# Conclusions

Proof-of-principle for a proton-driven Higgs factory helps identify key challenges.

Self-matching of the witness bunch helps to limit emittance growth in a single stage.

Since protons aren't *that* relativistic, everything on the plasma side gets easier at high energy.

# Protons: pros

Single stage

- higher average gradient
- inject once at relatively low energy
- adiabatic self-matching

Potential to reuse existing infrastructure

Better than linear scaling with driver energy

Moderately nonlinear wakefields allow for positron acceleration, but still not a solution for high quality beams.

# Protons: cons

Needs short proton driver at high rep-rate

Protons are heavy

- acceleration limited by longitudinal dispersion

Protons are positively charged

- lower transformer ratio
- needs tailored emittance

No staging means limited opportunities for correction

- external focussing
- stability becomes very important

# ~ Fin ~

# ~ Backups ~

# Footprint

Fits on the Fermilab site
(P5 review)

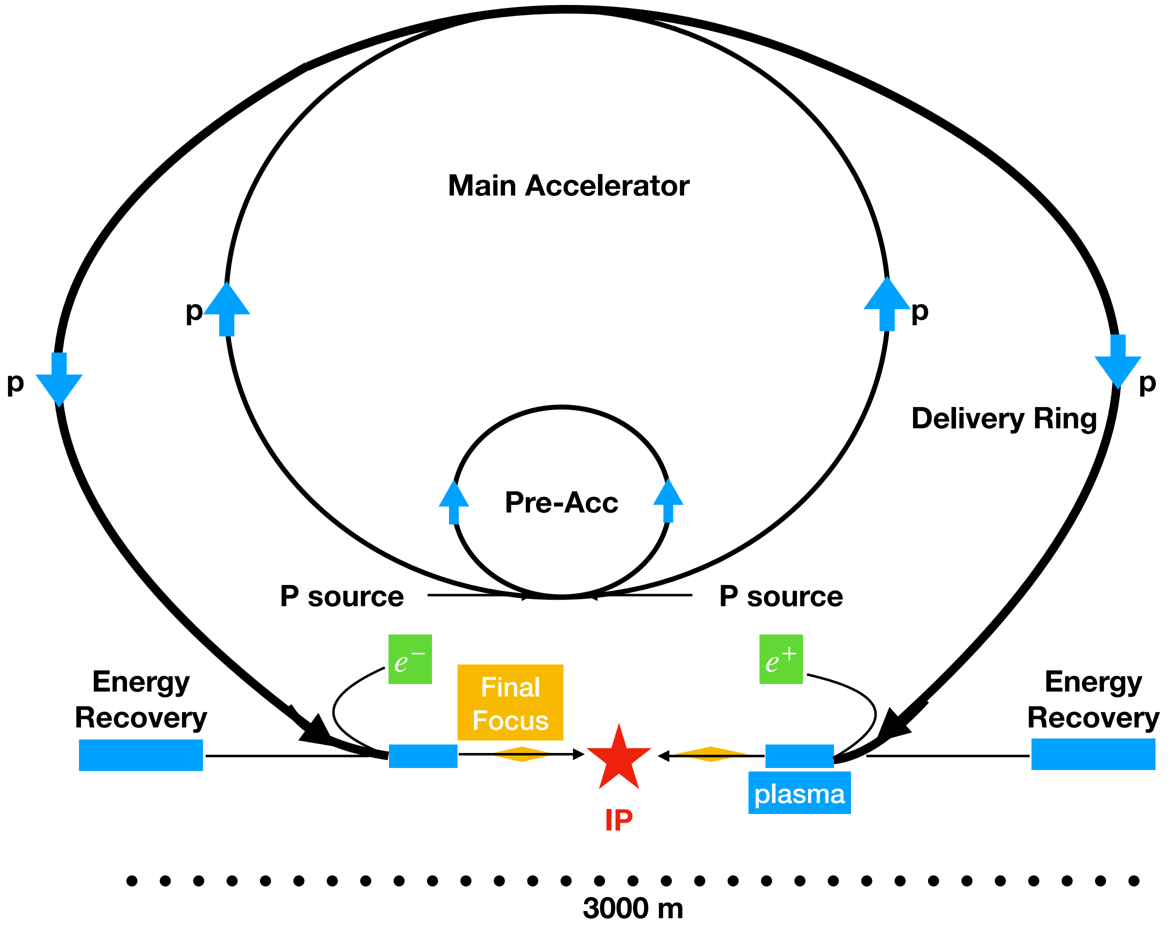

# Picking the driver: efficiency

Optimal length
for proton driver

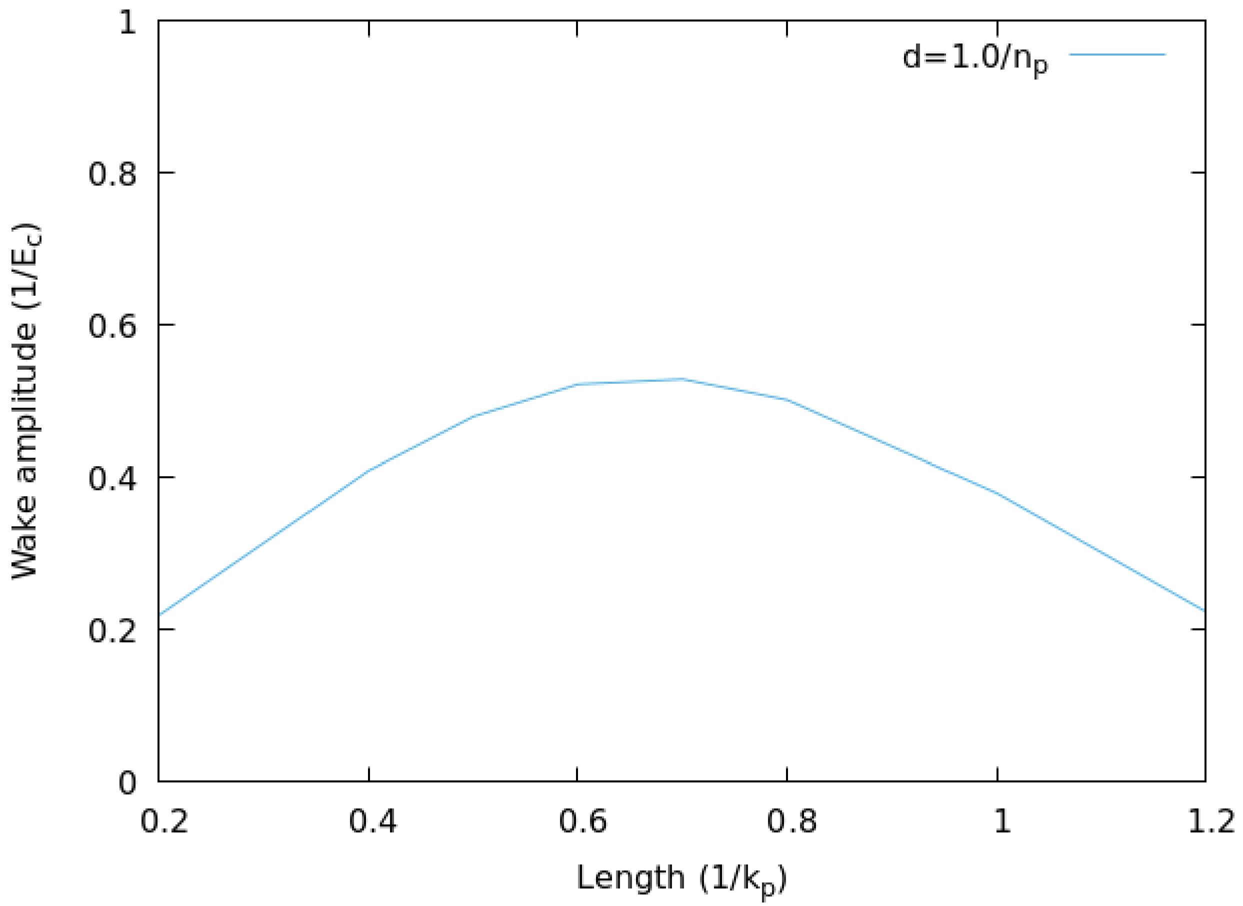

# Picking the driver: efficiency

Optimal length
for proton driver
depends on charge density.

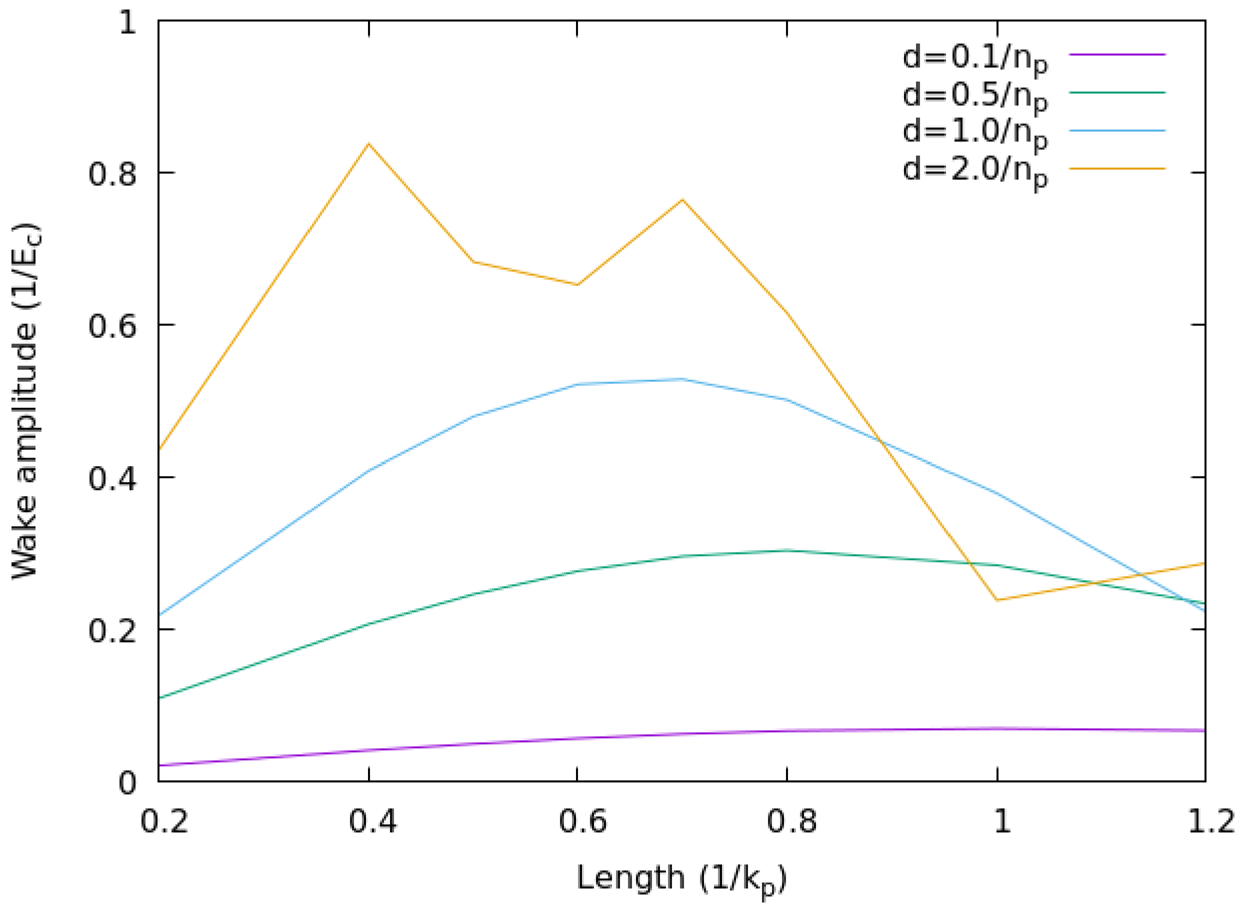

# Picking the driver: efficiency

Everything scales with plasma frequency

$1x10^{11}$ protons gives

- plasma density $3x10^{14}$ $cm^{-3}$
- driver length 150 µm
- Initial wakefields ~ 0.8GV/m

Pick 10% driver energy spread for “realistic” longitudinal emittance

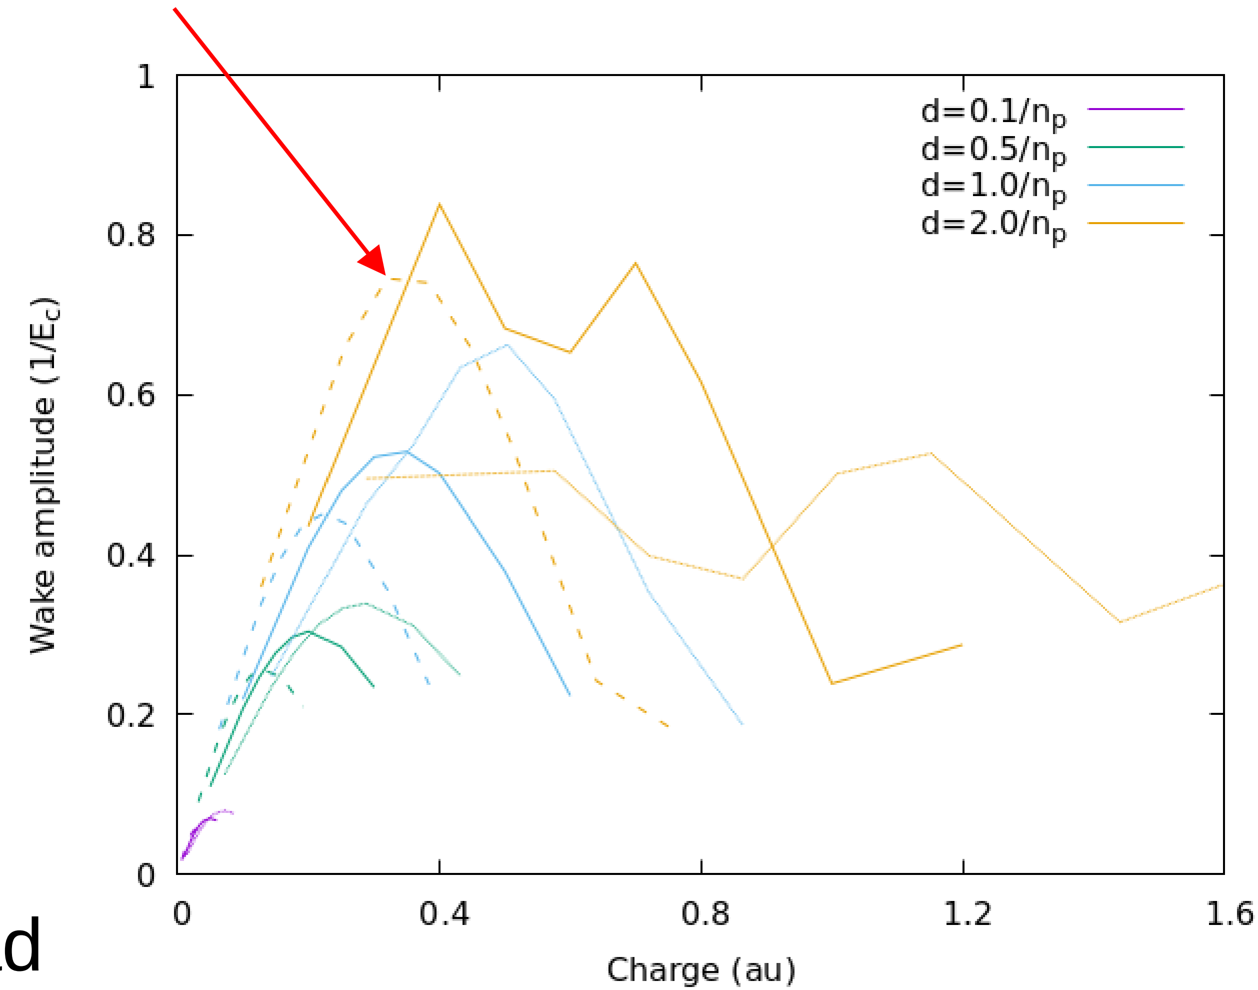


Dashed line: $k_pr = 0.8$
Solid line: $k_pr = 1.0$
Dotted line: $k_pr = 1.2$

# Cooling

Witness with 10% driver charge
absorbs ~20% of wakefield energy

Witness with 20% driver charge
absorbs ~40% of wakefield energy

Assume acceleration over 240m,
gives required cooling as 12.5 kW/m

# Cooling

Moderately nonlinear wakefields retain their structure after loading.

Could use a second witness bunch to “mop up” excess wakefield

# Luminosity

Combine everything:

- Assume proton beams at 5 Hz, with 1000 bunches per beam
- Assume witness beams with 20% driver charge, 100 nm emittance*, ILC optics, and negligible energy spread

*Flat beams should be investigated

and this scheme is competitive:
$1.7\text{x}10^{34}$ cm$^{-2}$s$^{-1}$

Farmer et al. (2024)

| Proton Accelerator Parameter | Symbol | Unit | Value |
|---|---|---|---|
| Proton energy | $E_p$ | GeV | 400 |
| Refill Time | $\tau$ | s | 0.2 |
| Bunch population | $N_p$ | $10^{10}$ | 10 |
| Number of bunches | $n$ | | 1000 |
| Longitudinal RMS | $\sigma_z$ | µm | 150 |
| Transverse RMS | $\sigma_{x,y}$ | µm | 240 |
| Normalized transverse emittance | $\epsilon_{T,p}$ | µm | $3-75$ µm |
| Power Usage | P | MW | 150 |
| **Plasma Parameters** | **Symbol** | **Unit** | **Value** |
| $e^-$ cell Length | $L_{e^-}$ | m | 240 |
| $e^+$ cell Length | $L_{e^+}$ | m | 240 |
| density - upstream | $n_p$ | $10^{14}$ cm$^{-3}$ | 3.2 |
| density - downstream | $n_p$ | $10^{14}$ cm$^{-3}$ | 5.2 |
| **e$^{\pm}$ Bunch Parameters** | **Symbol** | **Unit** | **Value** |
| Injection Energy | $E_{e,in}$ | GeV | 1 |
| Final Energy | $E_e$ | GeV | 125 |
| Bunch population | $N_{e\pm}$ | $10^{10}$ | 2 |
| Normalized transverse emittance | $\epsilon_{T,e}$ | nm | 100 |
| Hor. beta fn. | $\beta_x^*$ | mm | 13 |
| Ver. beta fn. | $\beta_y^*$ | mm | 0.41 |
| Hor. IP size. | $\sigma_x^*$ | nm | 73 |
| Ver. IP size. | $\sigma_y^*$ | nm | 13 |
| **$e^-e^+$ Collider Parameter** | **Symbol** | **Unit** | **Value** |
| Center-of-Mass Energy | $E_{\text{cm}}$ | GeV | 250 |
| Average Collision Rate | $f$ | kHz | 5 |
| Luminosity | $\mathcal{L}$ | cm$^{-2}$s$^{-1}$ | $1.7\times10^{34}$ |

# Flat vs round beams

Take CLIC parameters
at 380 GeV CoM
(Cilento, 2021)

Three cases:

- flat beams (CLIC baseline)
- round beams (geometric mean)
- Flat-focus
  (round emittance, keep beta*)