This notebook analyzes 3D simulations in terms of grain growth, with either pure grain growth, grain growth with immobile particles or grain growth with mobile particles (pores), as described in section 4.5 of the paper.

First, we show the results from the performance table again:

	name	fe_rhs	sts_rhs	speedup
0	GG	53279	14288	3.728933
1	mobi	156288000	1365954	114.416737
2	immobi	14110453	1224604	11.522462

```
np.float64(0.045903766300502476)
```

Next, the temporal behaviour: The simulations generally yield the expected qualitative behaviour. Pure grain growth is reasonably close to $n = 2$, with the deviation mostly due to bad statistics towards the end of the simulation run. Immobile particles induce a limiting grain size somewhat above the Zener prediction. Finally, mobile particles increase the exponent, though not to the expected $n = 3$. We suspect that the transition from pure to particle-dragged grain growth simply takes more than (and hence larger domains) than employed here.

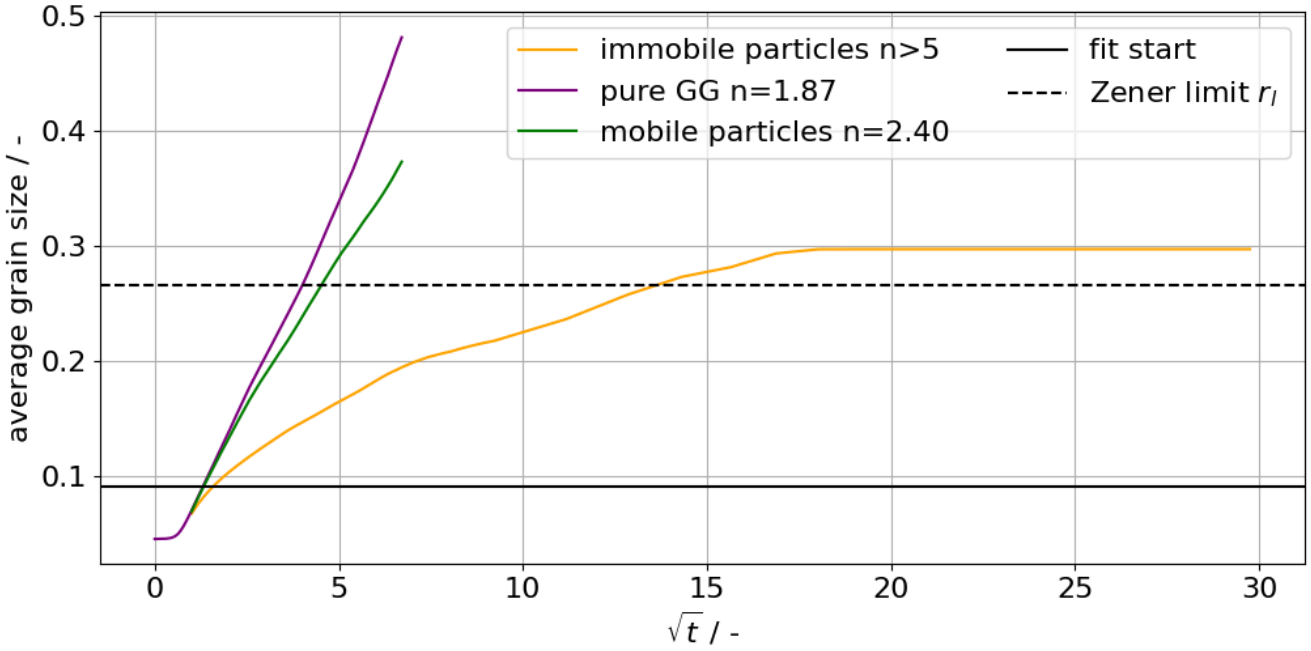


```
/tmp/ipykernel_24569/2104795248.py:2: RuntimeWarning: invalid value encountered in po
wer
  return (G0**n + A*x)**(1.0/n)
```

We also analyze the full step width and rejection output for these cases. Note that the simulations were ran from times 0->1, 1->5, 5->45, 45->85, 85->885, with restarts not using prior adaptive information but starting from the stable Euler step.

For pure grain growth we see pretty much no rejections and a steady growth of the allowable time step. This is mostly due to the interface regions growing smaller over time, which means more trivially solved regions, combined together with slower dynamics as the curvature reduces too. The initial rejection is simply since the initial conditions are sharp small time steps are required for accurate evolution. The spikes downwards are caused by I/O points.

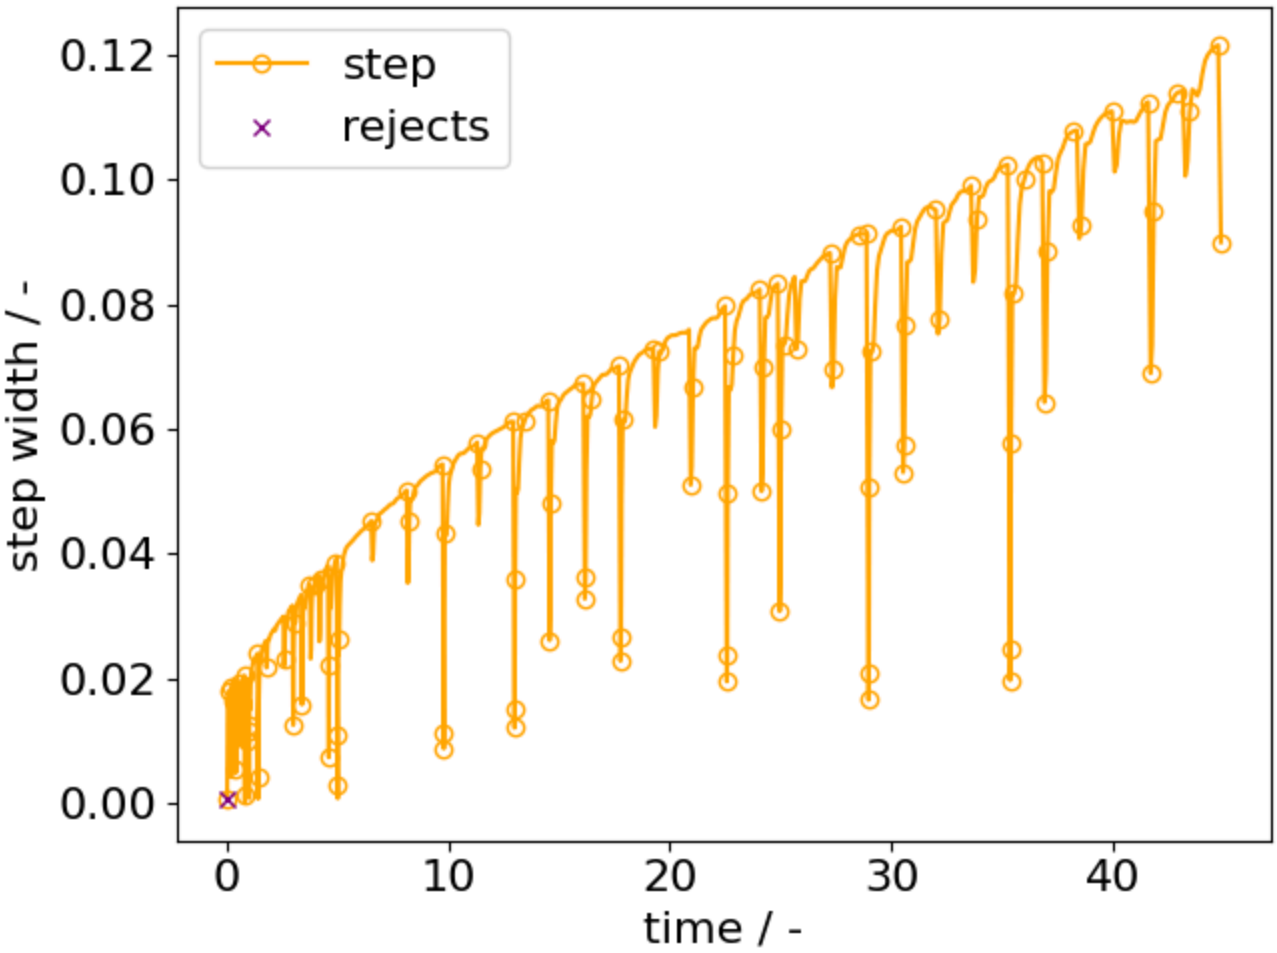


```
Text(0, 0.5, 'step width / -')
```

With immobile particles we see initialy a lot of change, but as more grain boundaries get locked out due to being satured with particles, the time step slowly stabilizes. Due to the overall slower dynamics there were no rejected steps; no initial step rejection occured as the initial guess (stable Euler step) was small enough to satisfy the tolerance. Note that the step width trajectory also mimics the stagnation of the grain size evolution.

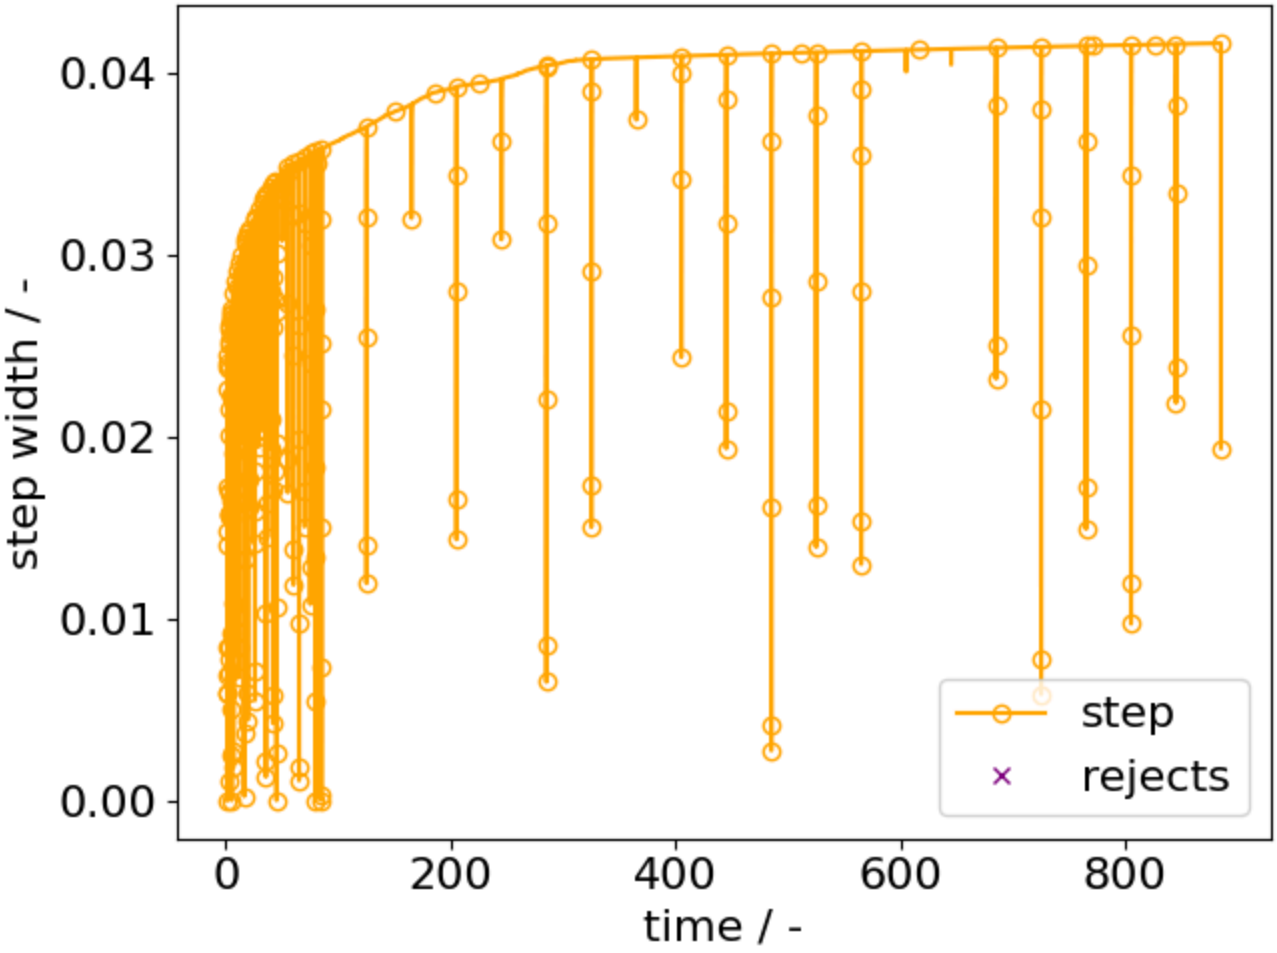


```
Text(0, 0.5, 'step width / -')
```

In contrast to the previous simulations the mobile pore case shows more interesting behaviour: We observe regular step rejection together with increasing step widths, with the latter for the previously stated reasons. The step rejections tend to happen with slight overshoots in the error (most with $e_n < 1.1$). We suspect that this originates from increased error from the concentration coupling which wasn't really stressed in the previous two cases, with concentration being effectively a passive background field.

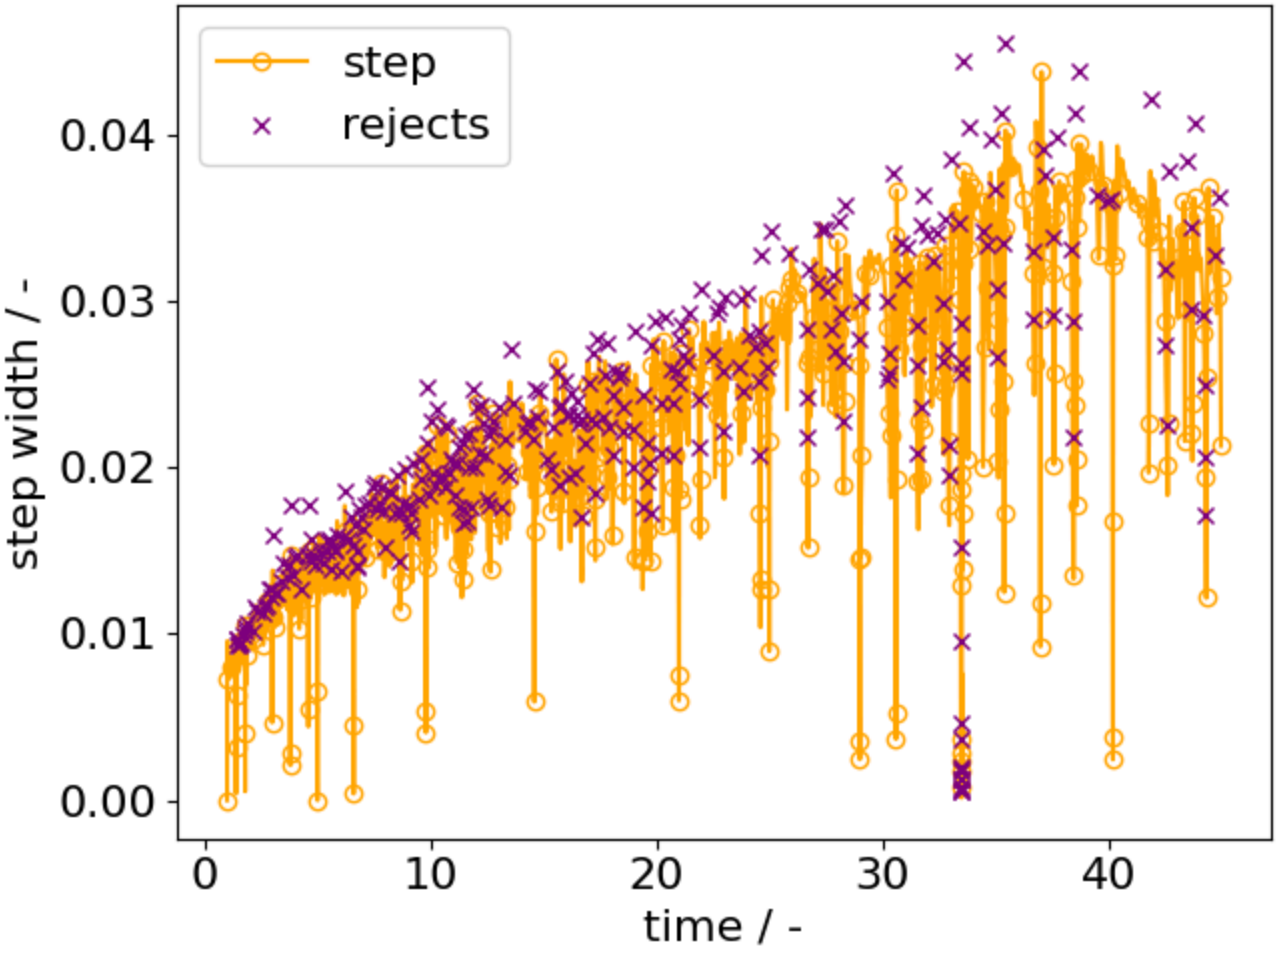


```
Text(0, 0.5, 'step width / -')
```

```
Rejection rate: 0.1326194398682043
```

Most rejections happen by barely going above 1. It's likely that the error estimate fails to be accurate enough for the coupled dynamics, although the phase-field should be barely changing per step due to the largely different timescales.

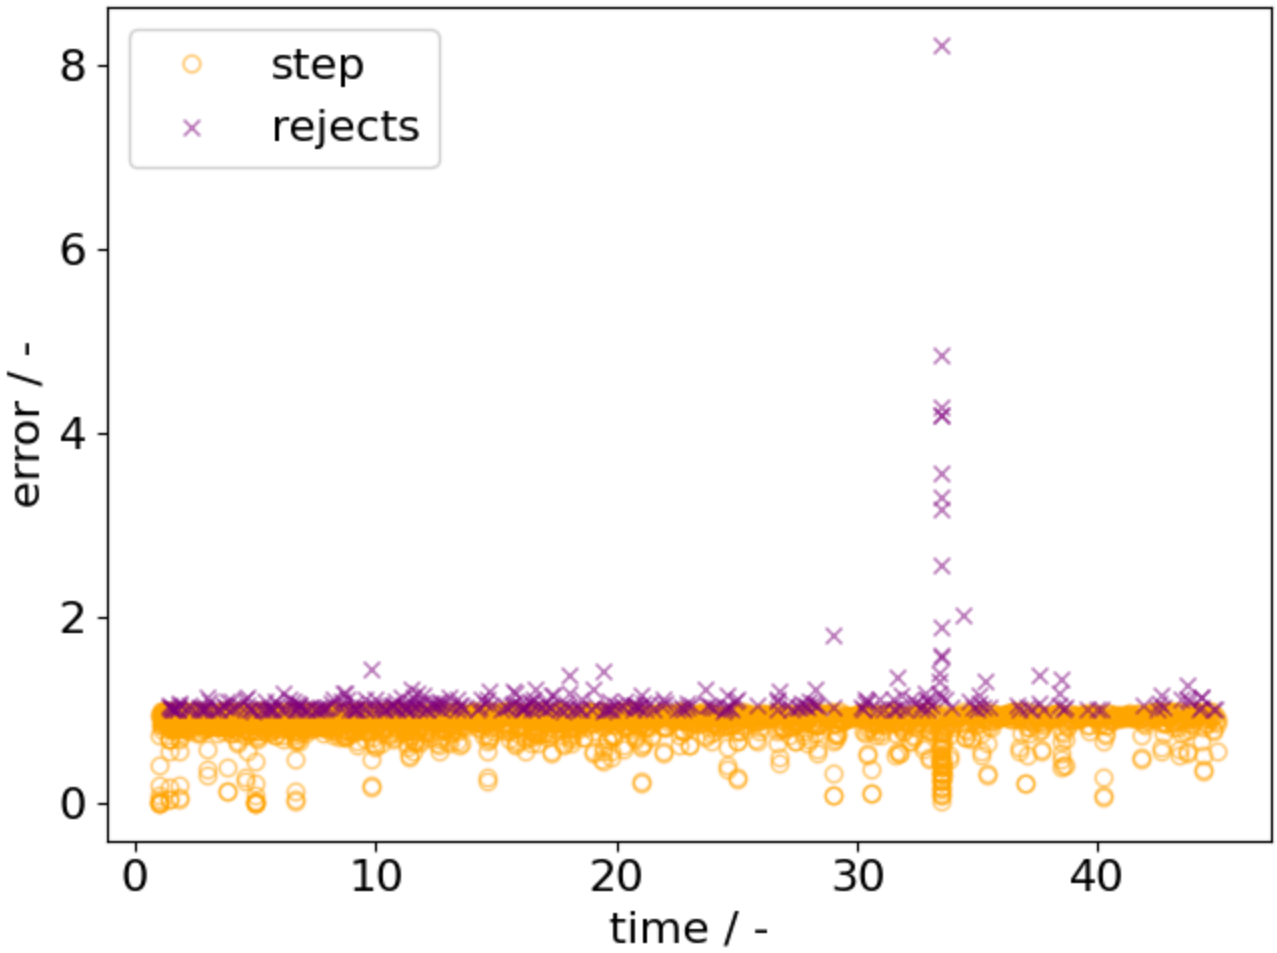


```
Median and average error for rejected steps: 1.045916477945483 1.168826061595156
```