

<https://j.mp/nextpnr-try-hard-mode>

```
nextpnr --try-hard=N --stop-on-success=[true|false]
```

When `--try-hard` mode is enabled, nextpnr will do **N** place and route runs (using different initial seeds). These runs will be done in parallel, using as many CPUs as available (IE on a 96 CPU machine, it will run up to 96 in parallel). At the end, the bitstream with the higher possible fmax will be returned with some statistical information about the runs including;

- (a) The distribution of fmax across runs including the faster, slowest and average values.
- (b) The distribution of run completing times including the faster, slowest and average values.

When `--try-hard` mode is enabled with `--stop-on-success=true` then when any run successfully meets the timing constraints, all other runs are immediately terminated and the run succeeding result is returned.

Question

Is this better **implemented as a wrapper** around nextpnr or **directly into nextpnr**?