CHIME: The Canadian Hydrogen Intensity Mapping Experiment

Kendrick Smith (PI)

Cosmology on Safari, Feb 2017

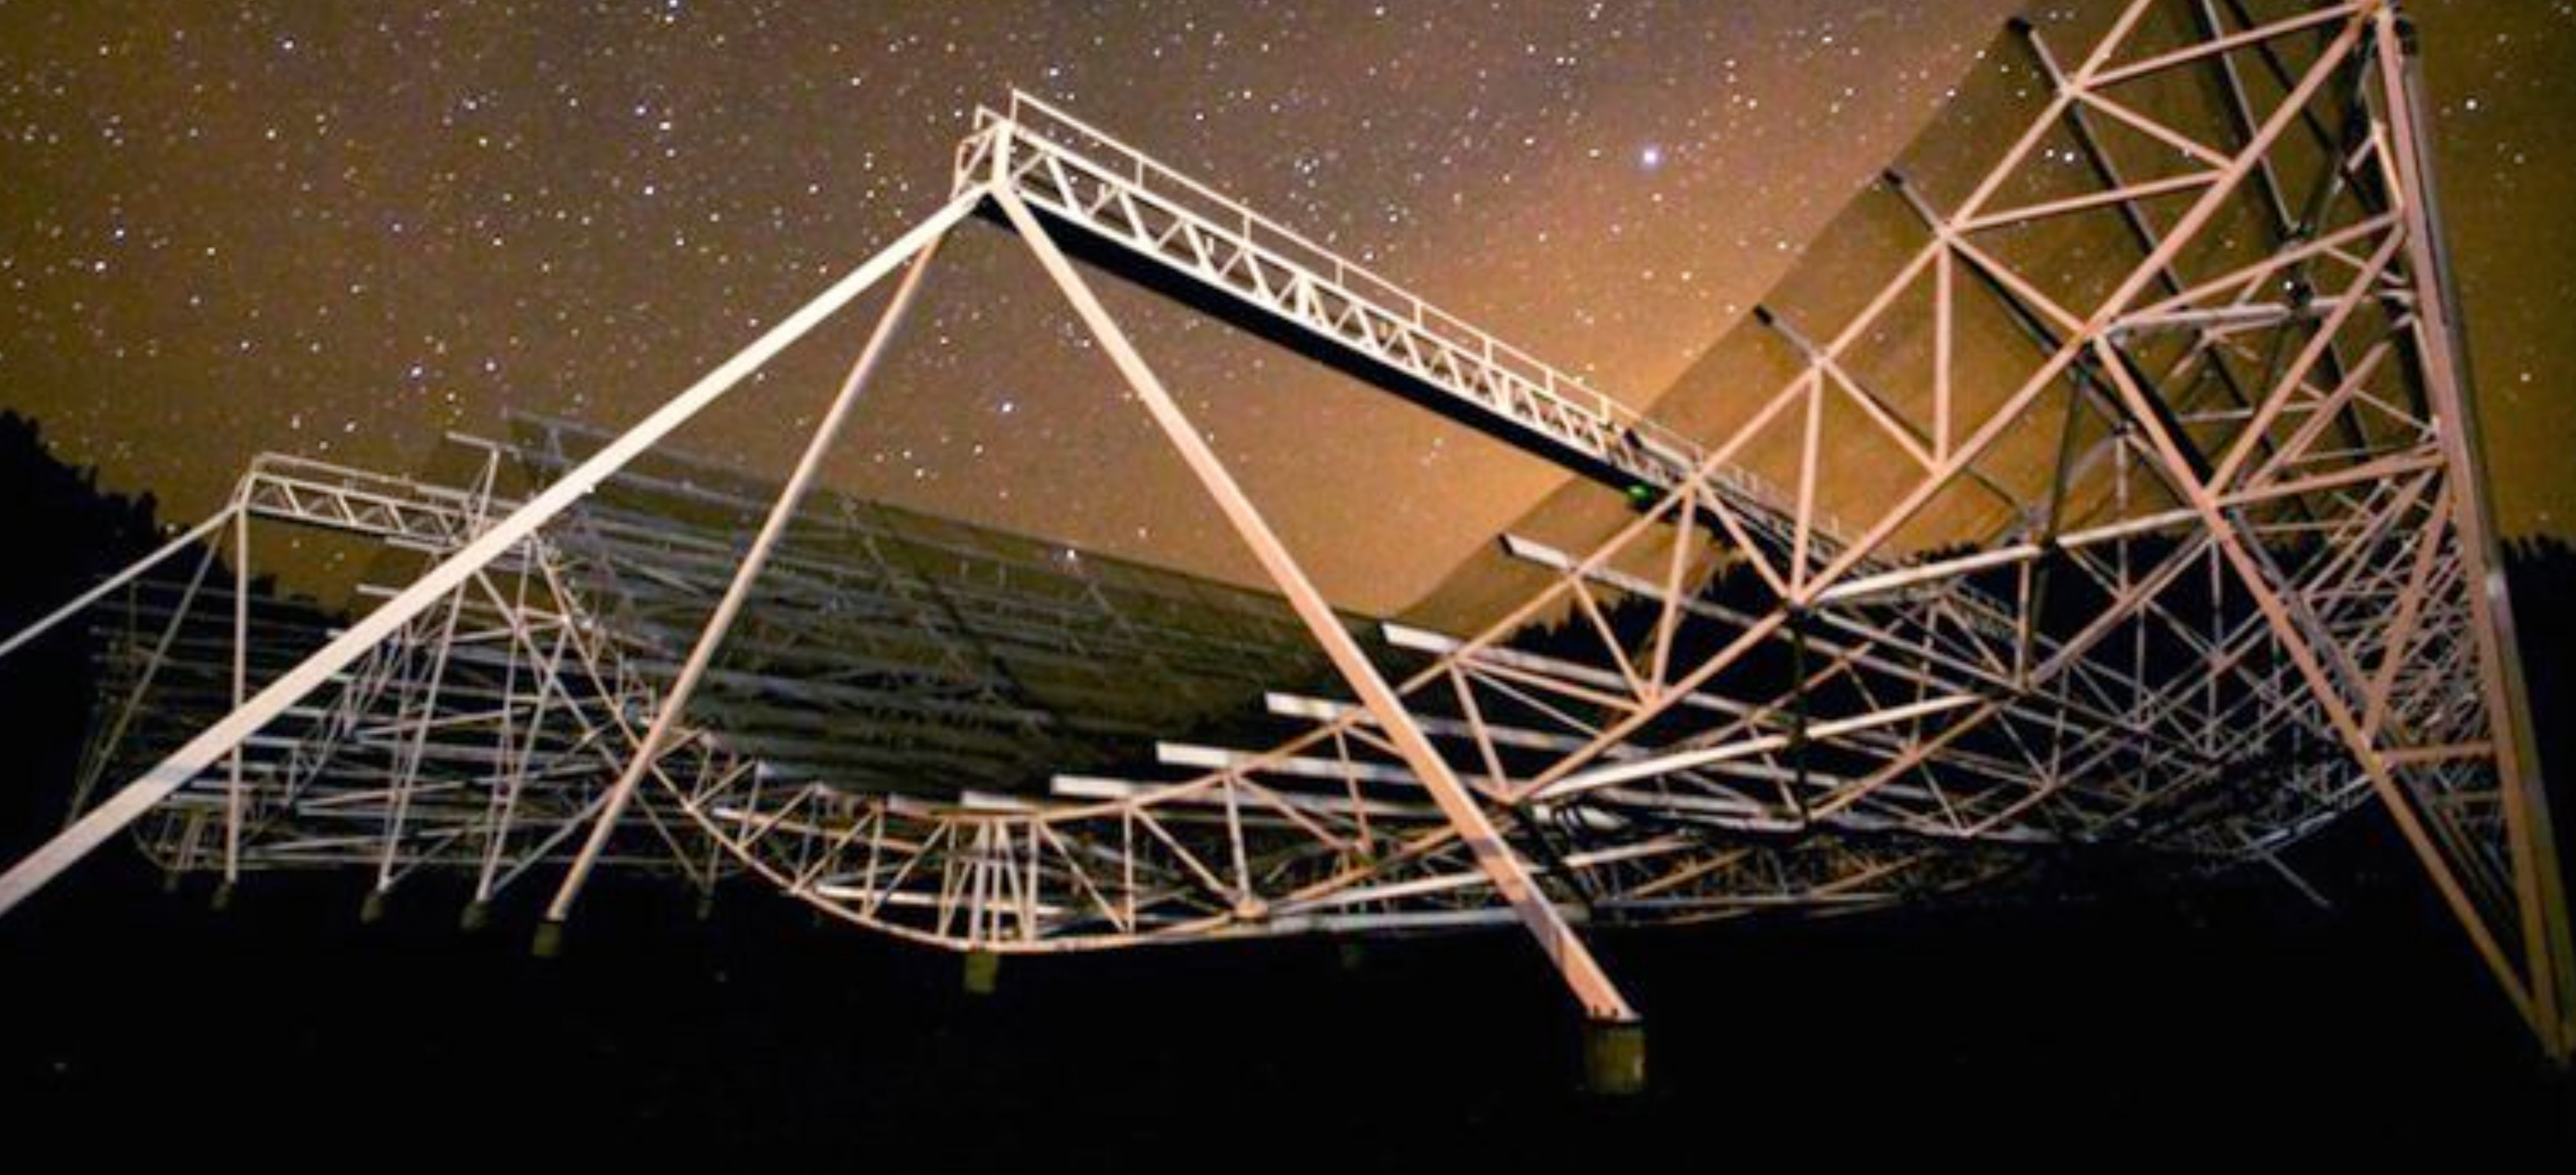

Single-feed radio telescope

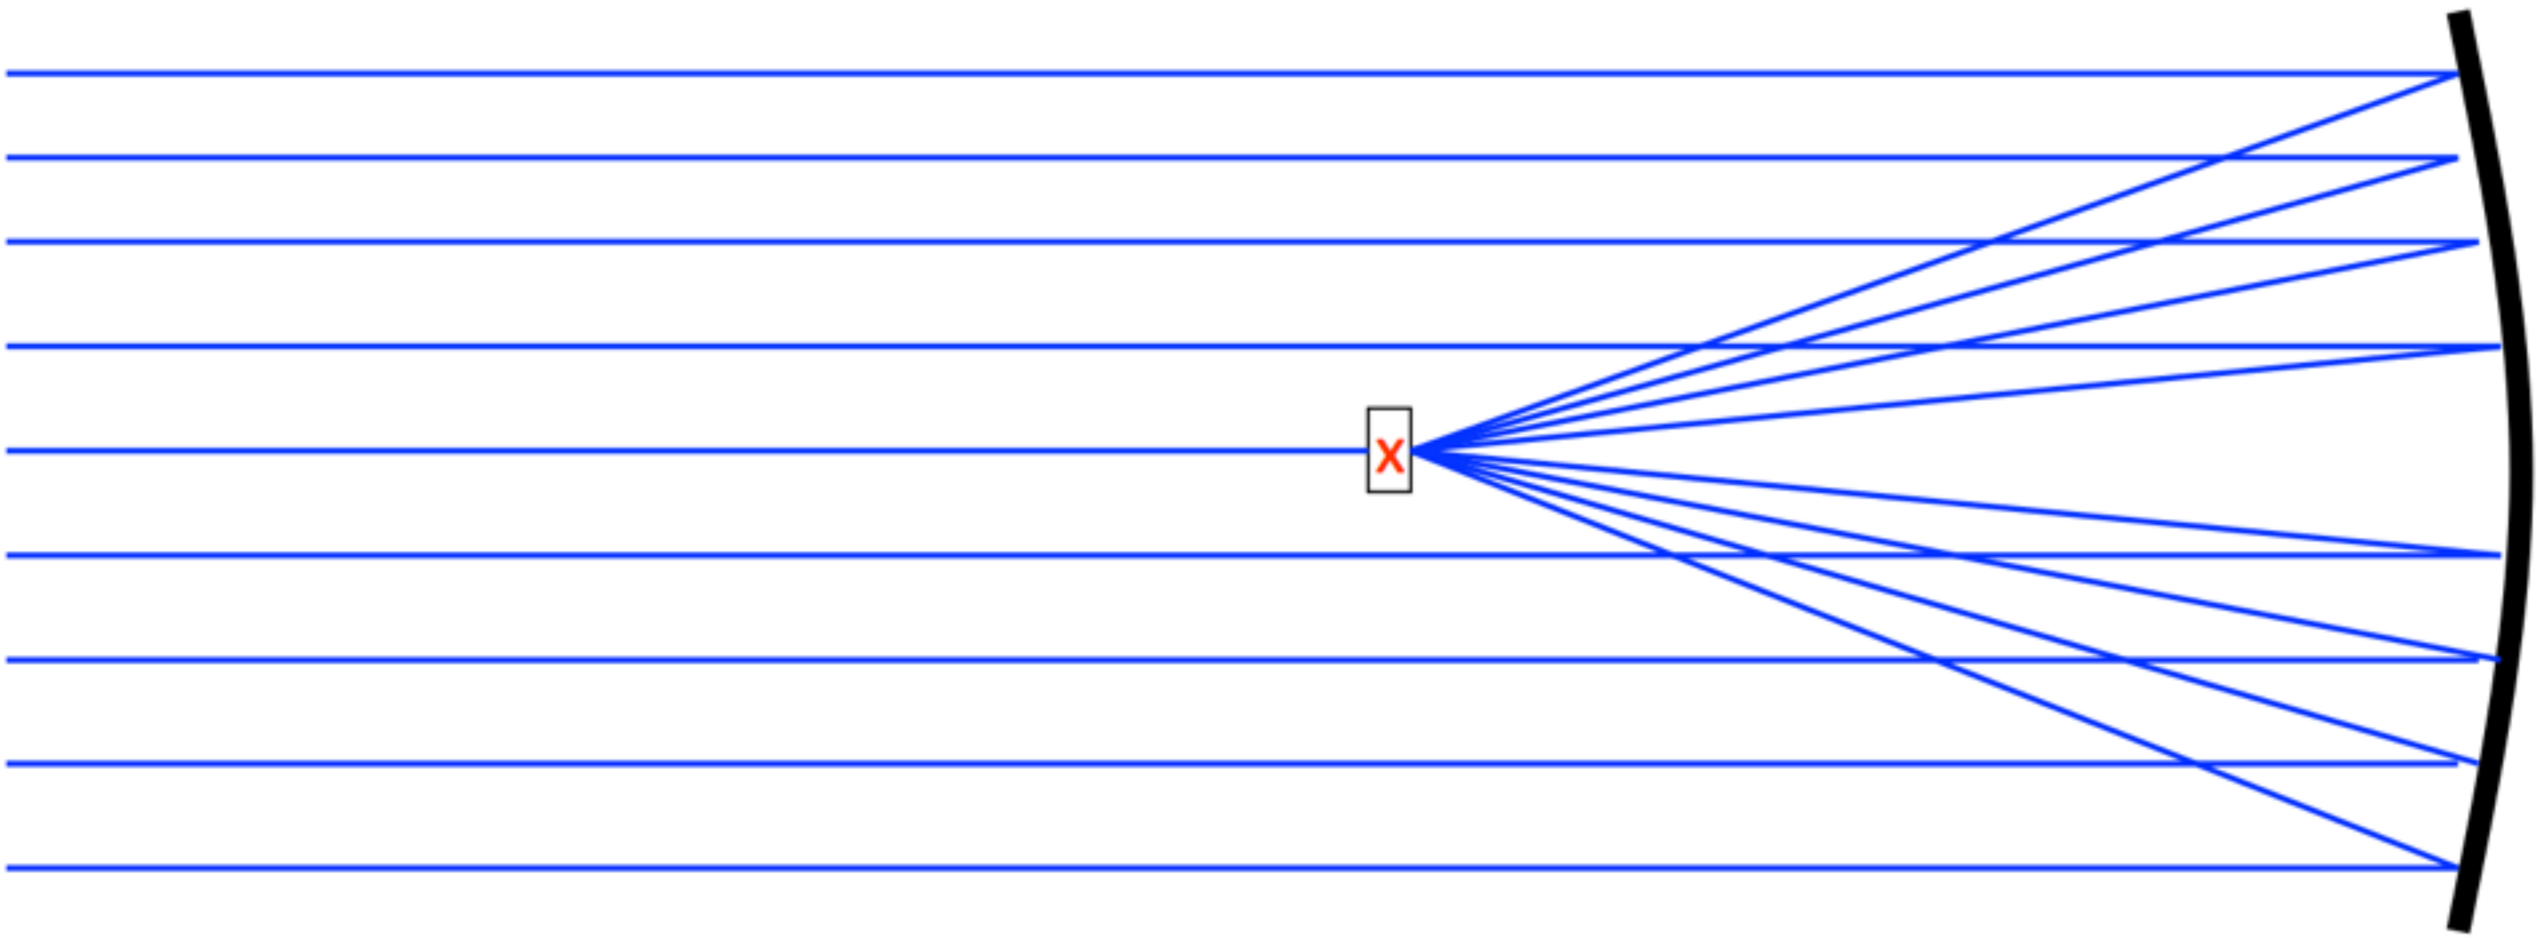

Focuses via physical delays: constructive interference only occurs for a specific direction on the sky

Phased-array interferometer

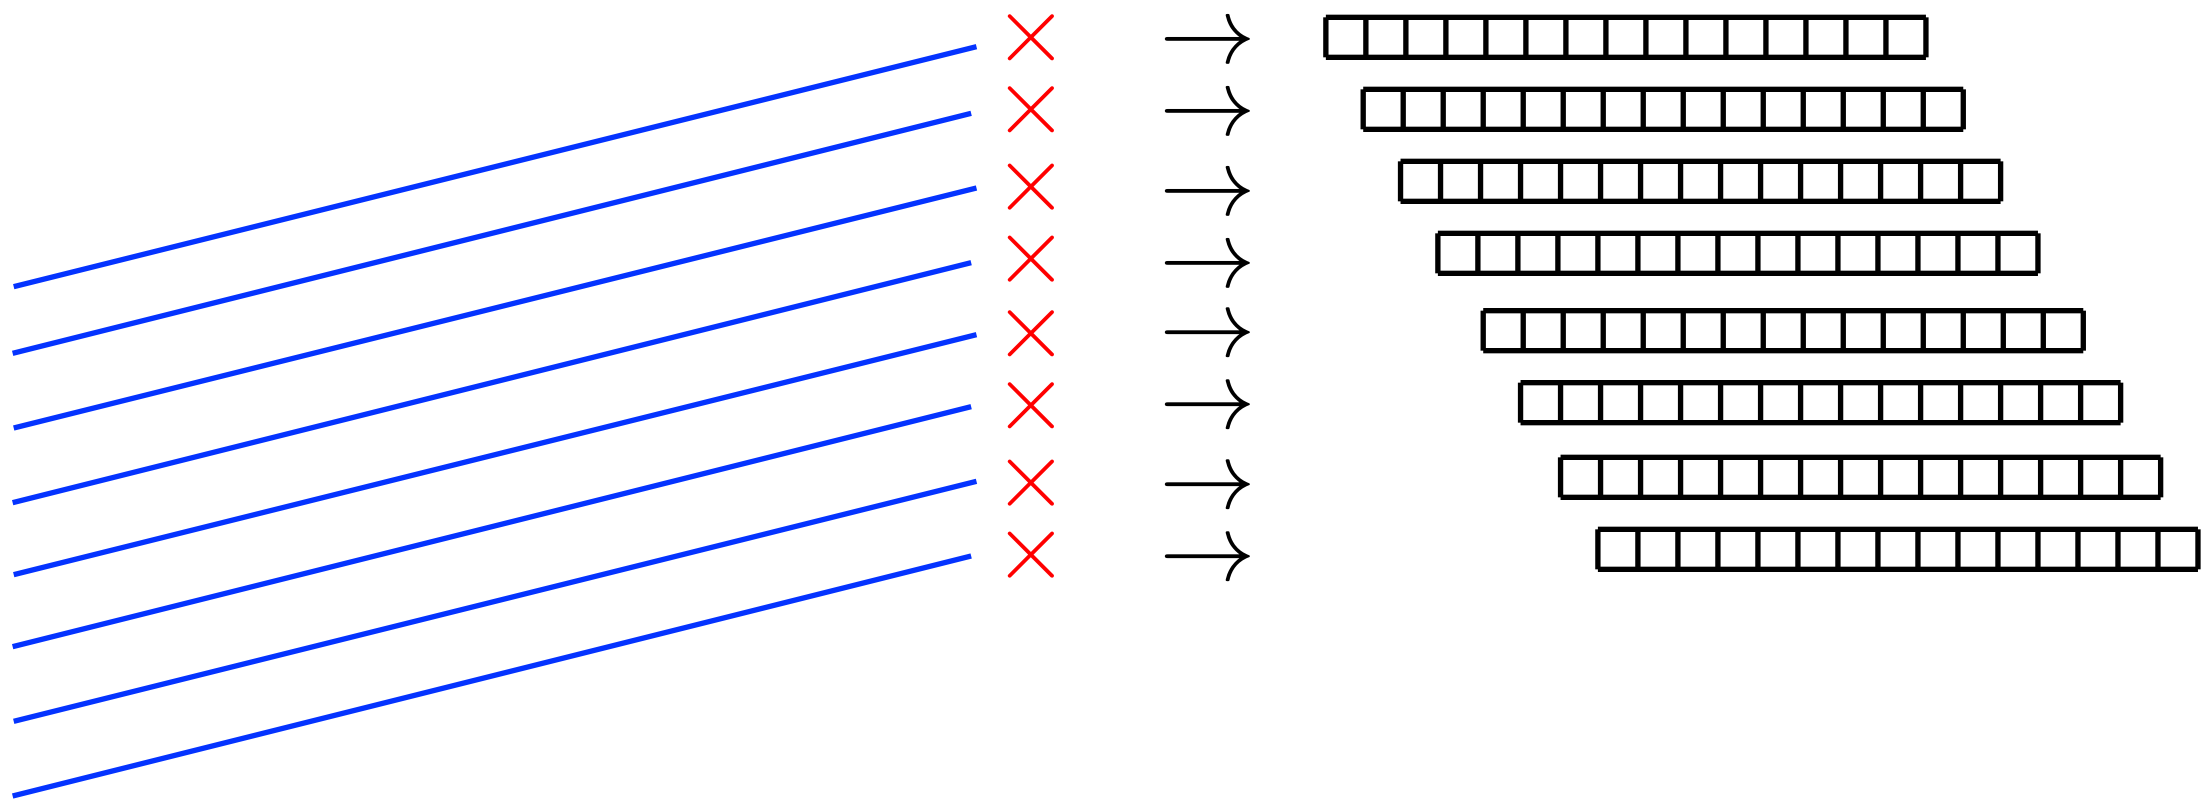

Dish is replaced by an array of antennas whose signals are digitized.

By summing signals with appropriate delays, can mimic the dish and focus on part of the sky.

Can "repoint" telescope by changing delays.

Beamforming interferometer

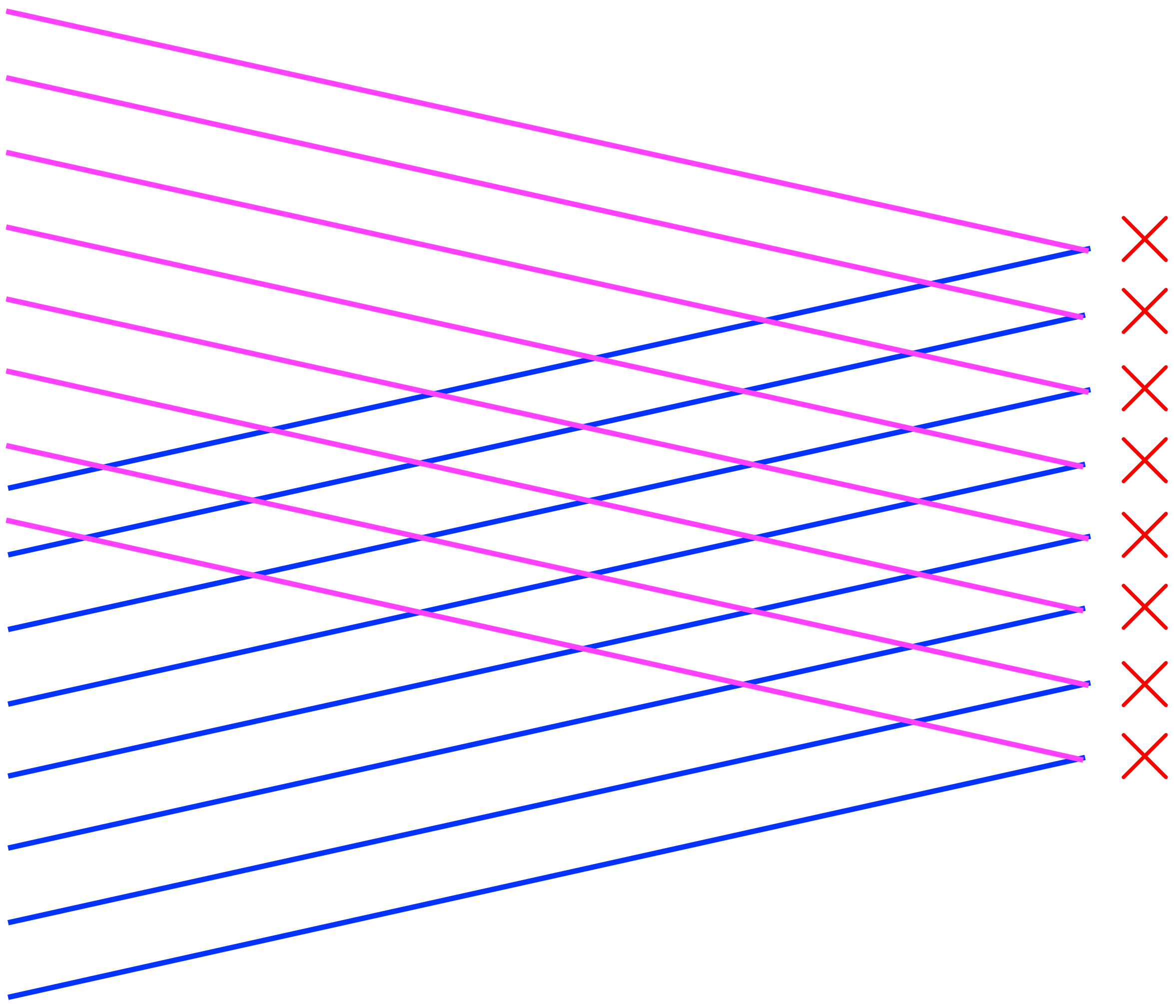

Copy the digitized signals and repeat the computation N times (in parallel). Equivalent to N single-feed telescopes with the same collecting area as the interferometer!

CHIME

Under construction now, first light expected this summer.

We have a 4 x 256 array of antennas. and enough computing power to form all 1024 independent beams in real time.

Raw sensitivity is the same as 1024 single-feed radio telescopes!

McGill

THE UNIVERSITY OF BRITISH COLUMBIA

UNIVERSITY OF TORONTO

DRAO

NRC · CNRC

CHIME beamforming, cartoon form

At any instantaneous observing time, each antenna sees a narrow strip on the sky ("primary beam").

By beamforming in software as previously described, we can make 1024 "formed" beams with size ~20 arcmin.

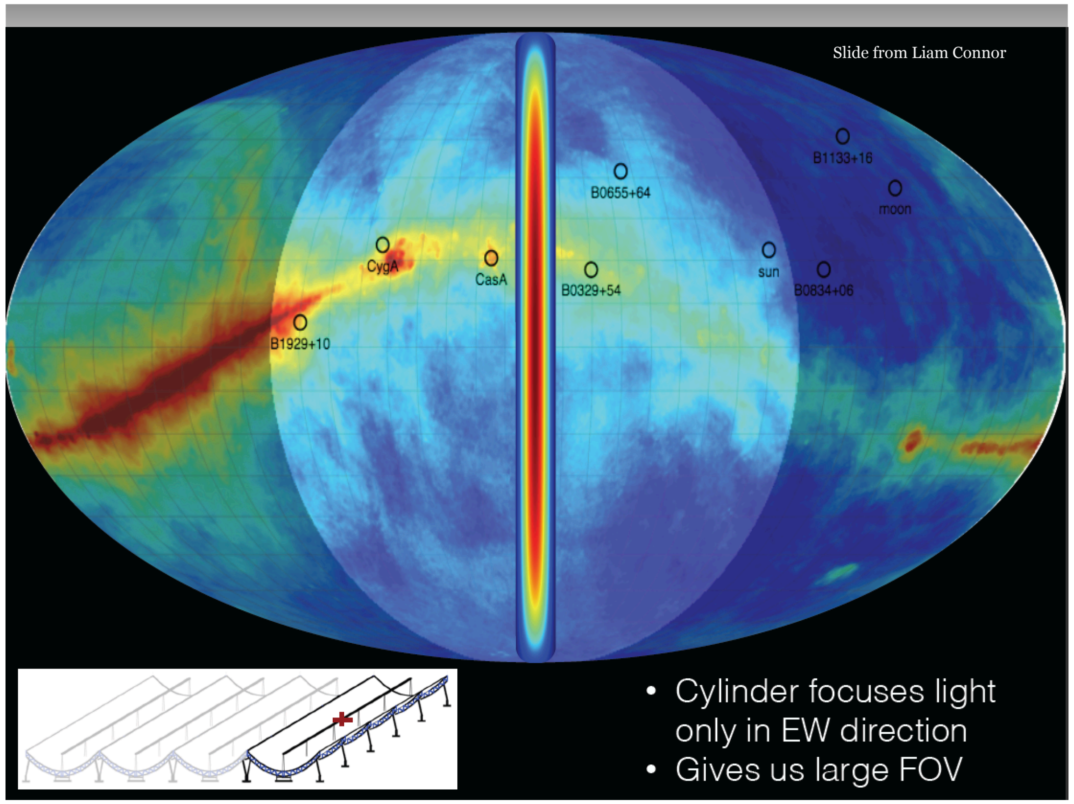

primary beam

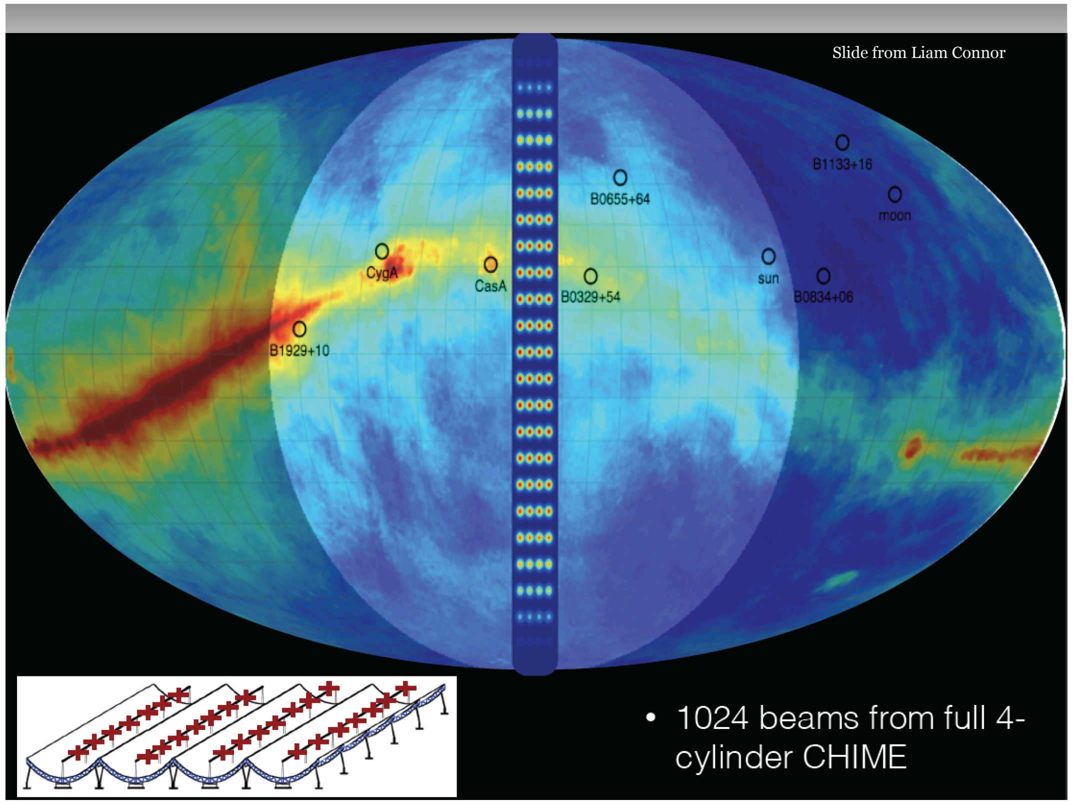

formed beams

CHIME beamforming, cartoon form

The primary beam is fixed in telescope coordinates, but as the Earth rotates, it sweeps over the full sky.

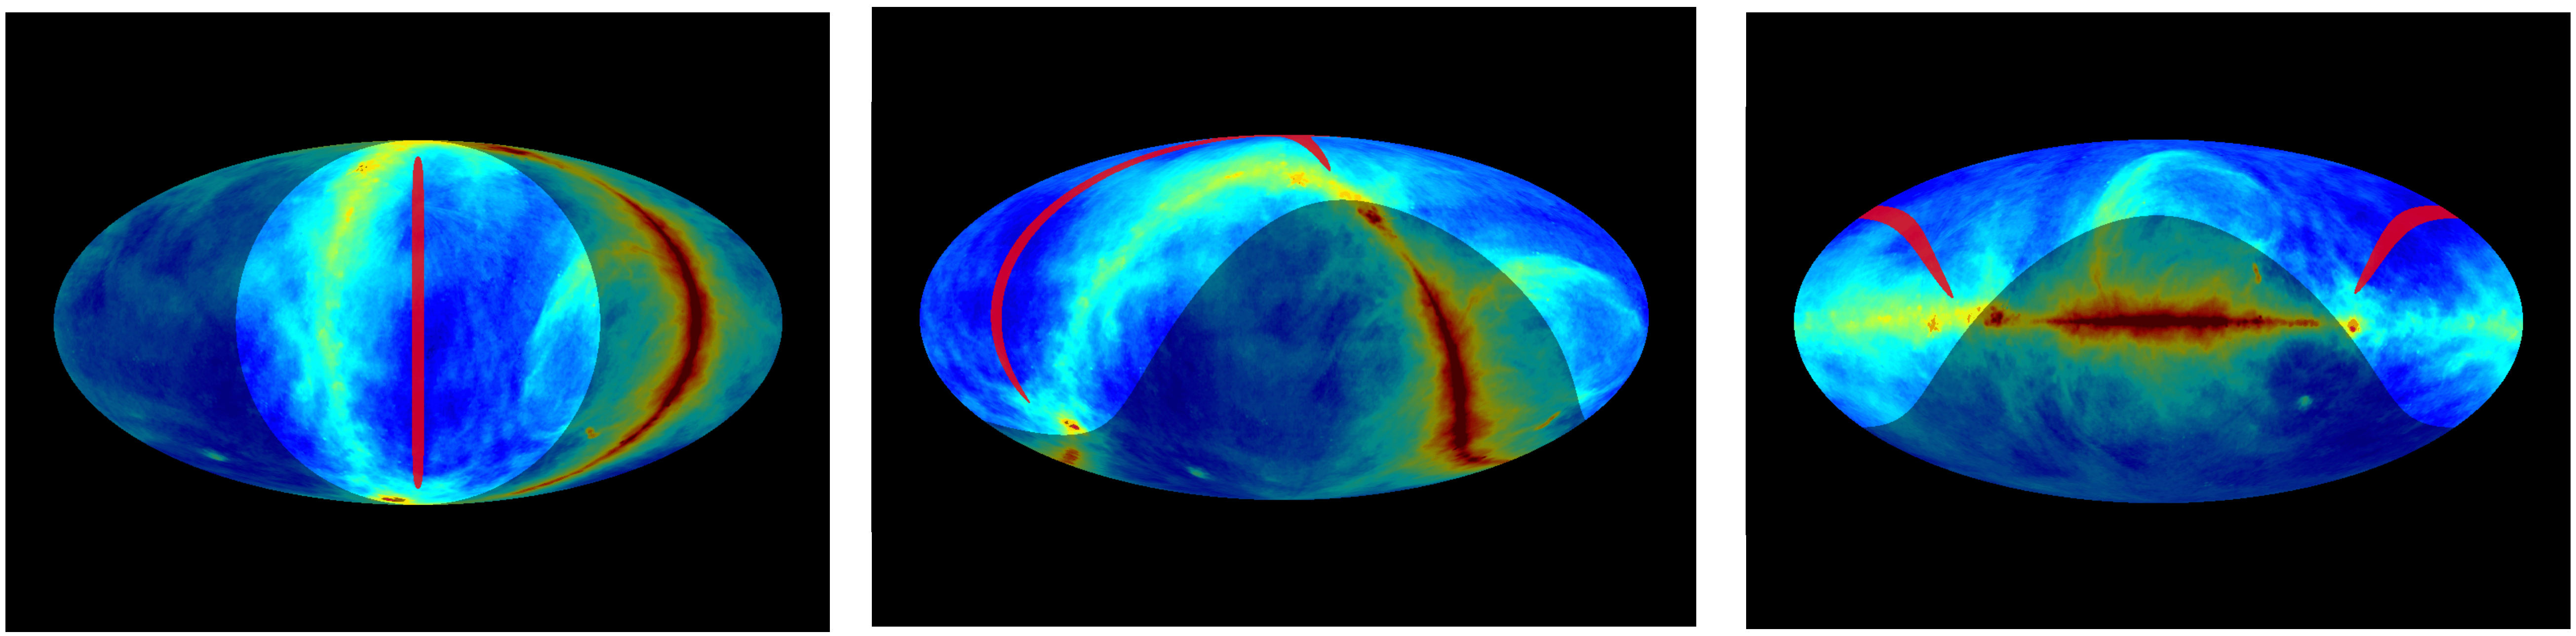

Every 24 hours, we get a sky map with:

- Angular resolution: 20 arcmin
- Sky coverage: half the sky
- Frequency range: 400-800 MHz.

  (We see neutral hydrogen at $z = 0.8$-2.5 via the 21-cm line)

Mapping speeds (back-of-envelope)

For many purposes, the statistical power of a radio telescope can be quantified by its mapping speed:

$$M \approx (\text{Collecting area } A) \times (\text{Number of beams}) \\ \times (\text{order-one factors})$$

	A	$N_{\rm beams}$	$M/(10^5\ \mathrm{m}^2)$
Parkes 64m	3200 m^2	13	0.41
Green Bank 100m	7850 m^2	7	0.55
Aricebo 300m	70000 m^2	7	4.9
FAST 500m	200000 m^2	19	38
CHIME	6400 m^2	1024	66

=  CHIME ?!

The catch

	A	N_{beams}	Mapping speed
FAST 500m	200000 m^2	19	38
CHIME	6400 m^2	1024	66

In principle, sensitivity is roughly proportional to mapping speed, but computational cost is proportional to N_{beams} (or worse).

What we have really done is move the difficulty from hardware to software.

CHIME hardware, cartoon form

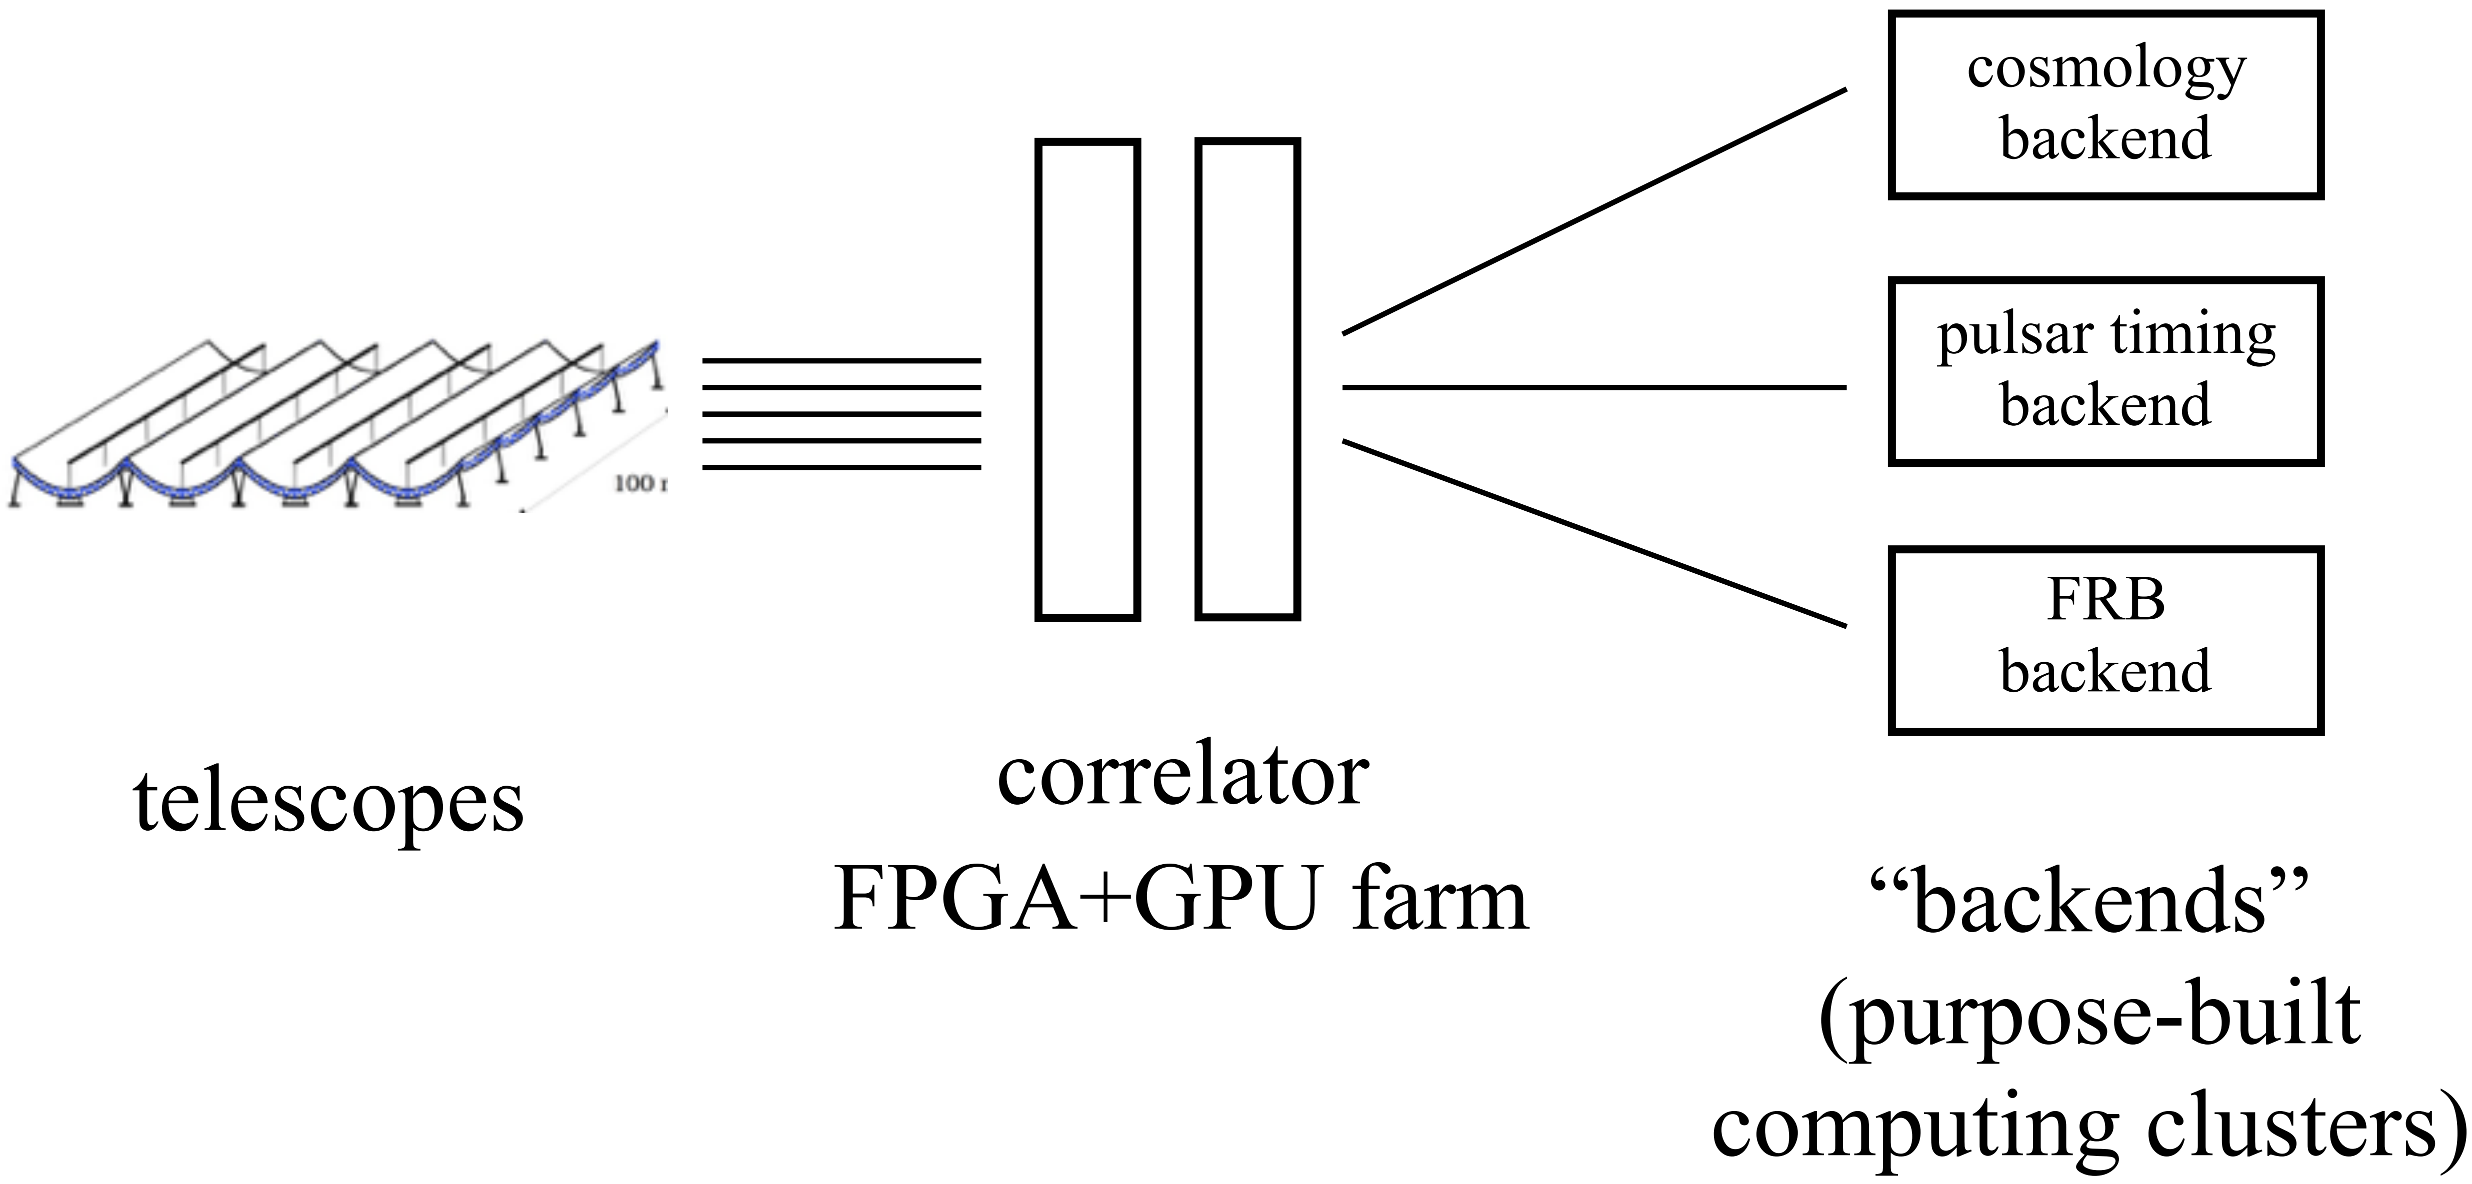

CHIME hardware, cartoon form

Each backend asks the correlator for a different data stream over the network. For example:

"Pulsar timing backend"

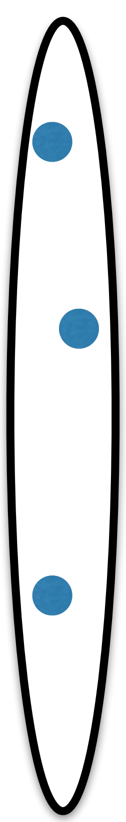

receives digitized electric field

with nanosecond sampling

at 10 specified sky locations

"FRB backend"

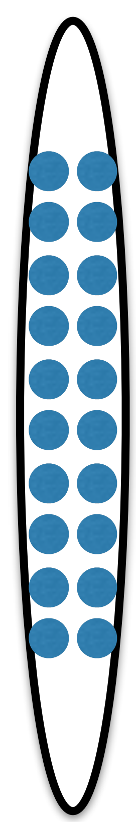

receives intensity, obtained from electric field by time-downsampling and polarization-averaging

with millisecond sampling

at a regular array of 1024 sky locations

To get a sense for the scale of the computational problems, consider the FRB backend. Every second it receives:

1024 beams
x 16384 frequency channels
x 1024 time samples

= $1.6\ 10^{10}$ numbers/second (1 petabyte/day)

To put this in perspective, simulating Gaussian random noise at this rate using the C++ standard library would require a dedicated 420-core cluster.

Any data analysis on a timestream of this size will be an extremely difficult computational problem.

=  CHIME ?

=  CHIME + a lot of math, for some problems?

Fast radio bursts

Fast radio bursts (FRB's): a genuine mystery in astrophysics!

Occasionally, a bright (~1 Jy), narrow (~1 ms), non-repeating, highly dispersed radio pulse is observed.

"Dispersed" means that arrival time at frequency ν is delayed as (delay) $\propto \nu^{-2}$

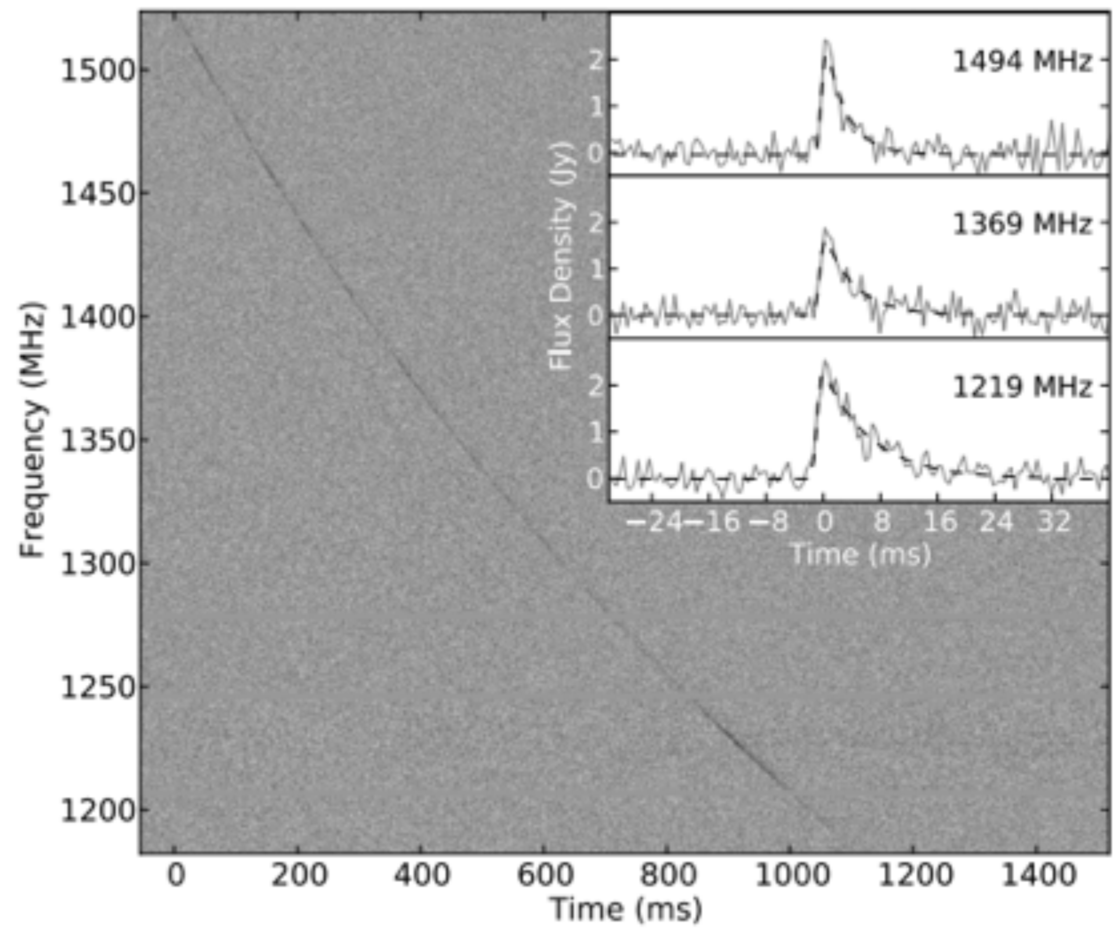

Fast radio bursts

The ν^{-2} delay is interpreted as dispersion due to an optically thin, cold, unmagnetized plasma of free electrons between source and observer. One can infer the dispersion measure (DM)

$$\mathrm{DM} = \int dx\, n_e(x)$$

The FRB's are a rare population of events with very high DM, suggesting that they may be at cosmological distances.

Recently one FRB (the "repeater") was shown to be at z=0.2!

Only ~20 FRB's have ever been observed. To test hypotheses, we need more data!

Fast radio bursts

For CHIME, the forecasted event rate is ~10 FRB's per day (!!)

However, achieving this event rate requires summing a hard computational problem. Need to sum over all straight lines:

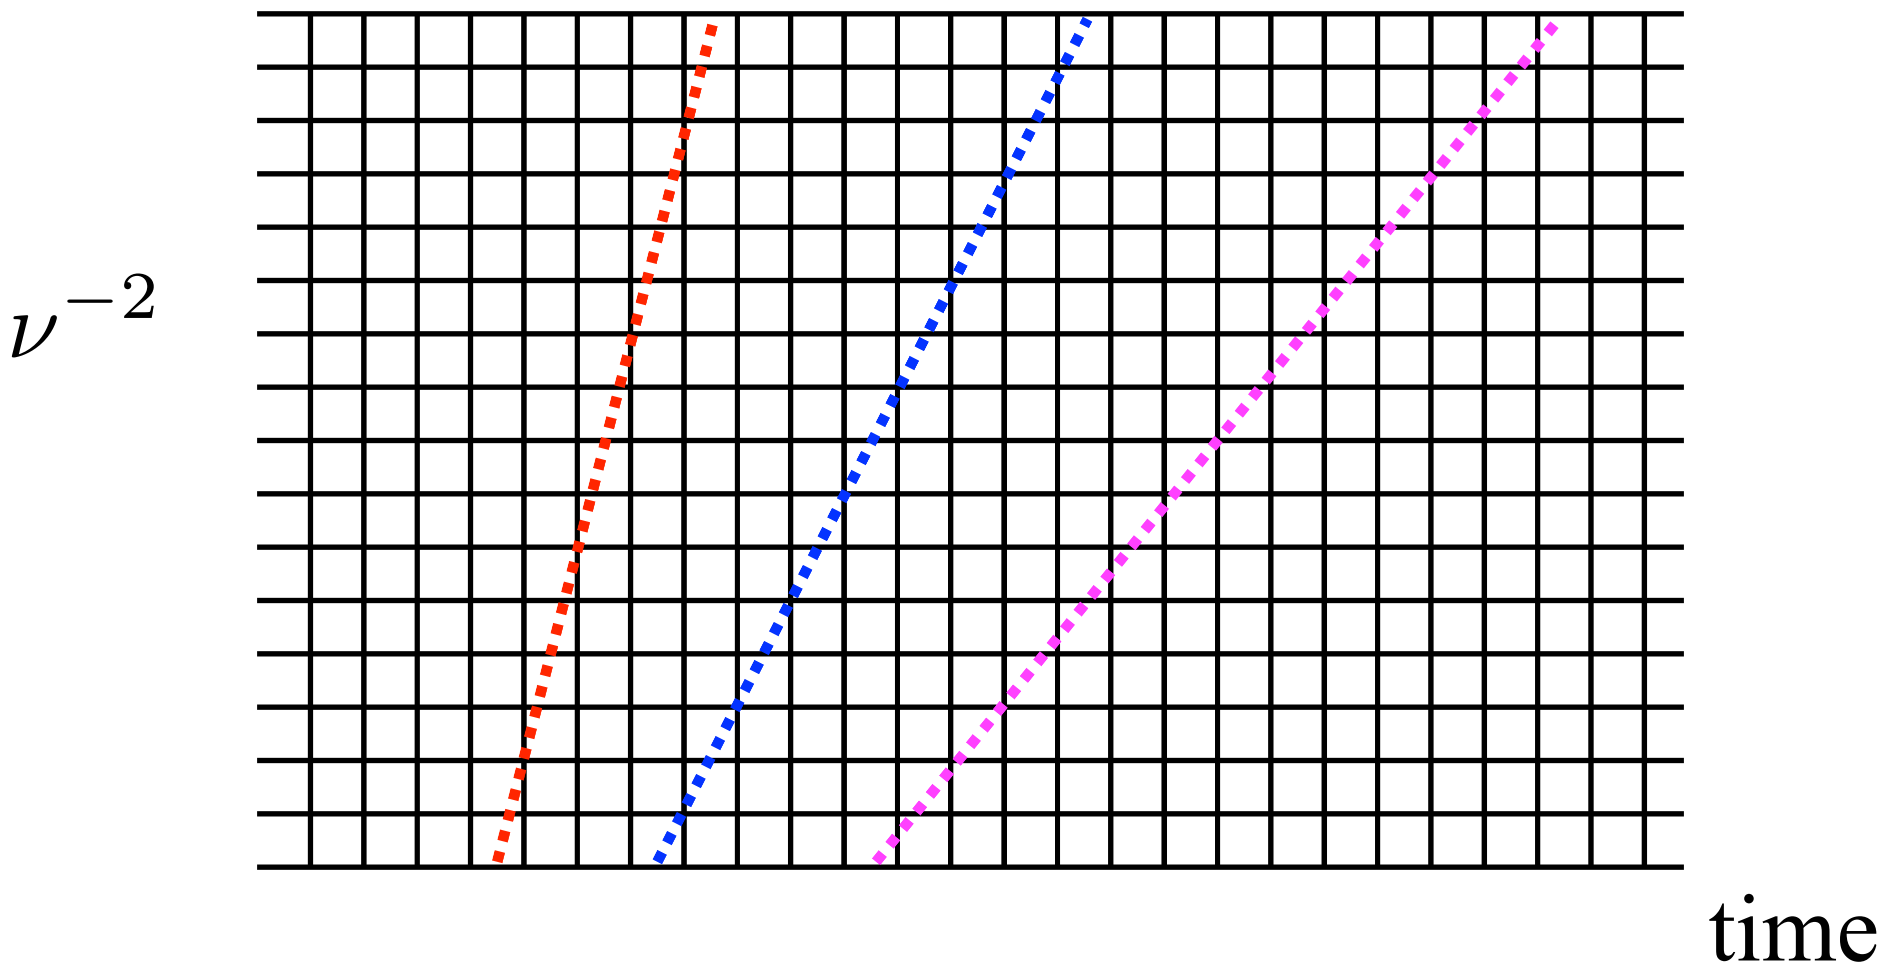

Fast radio bursts

Our algorithm ends up approximating each straight-line track by a jagged sum of samples. The sums are built up recursively as explained in the next few slides.

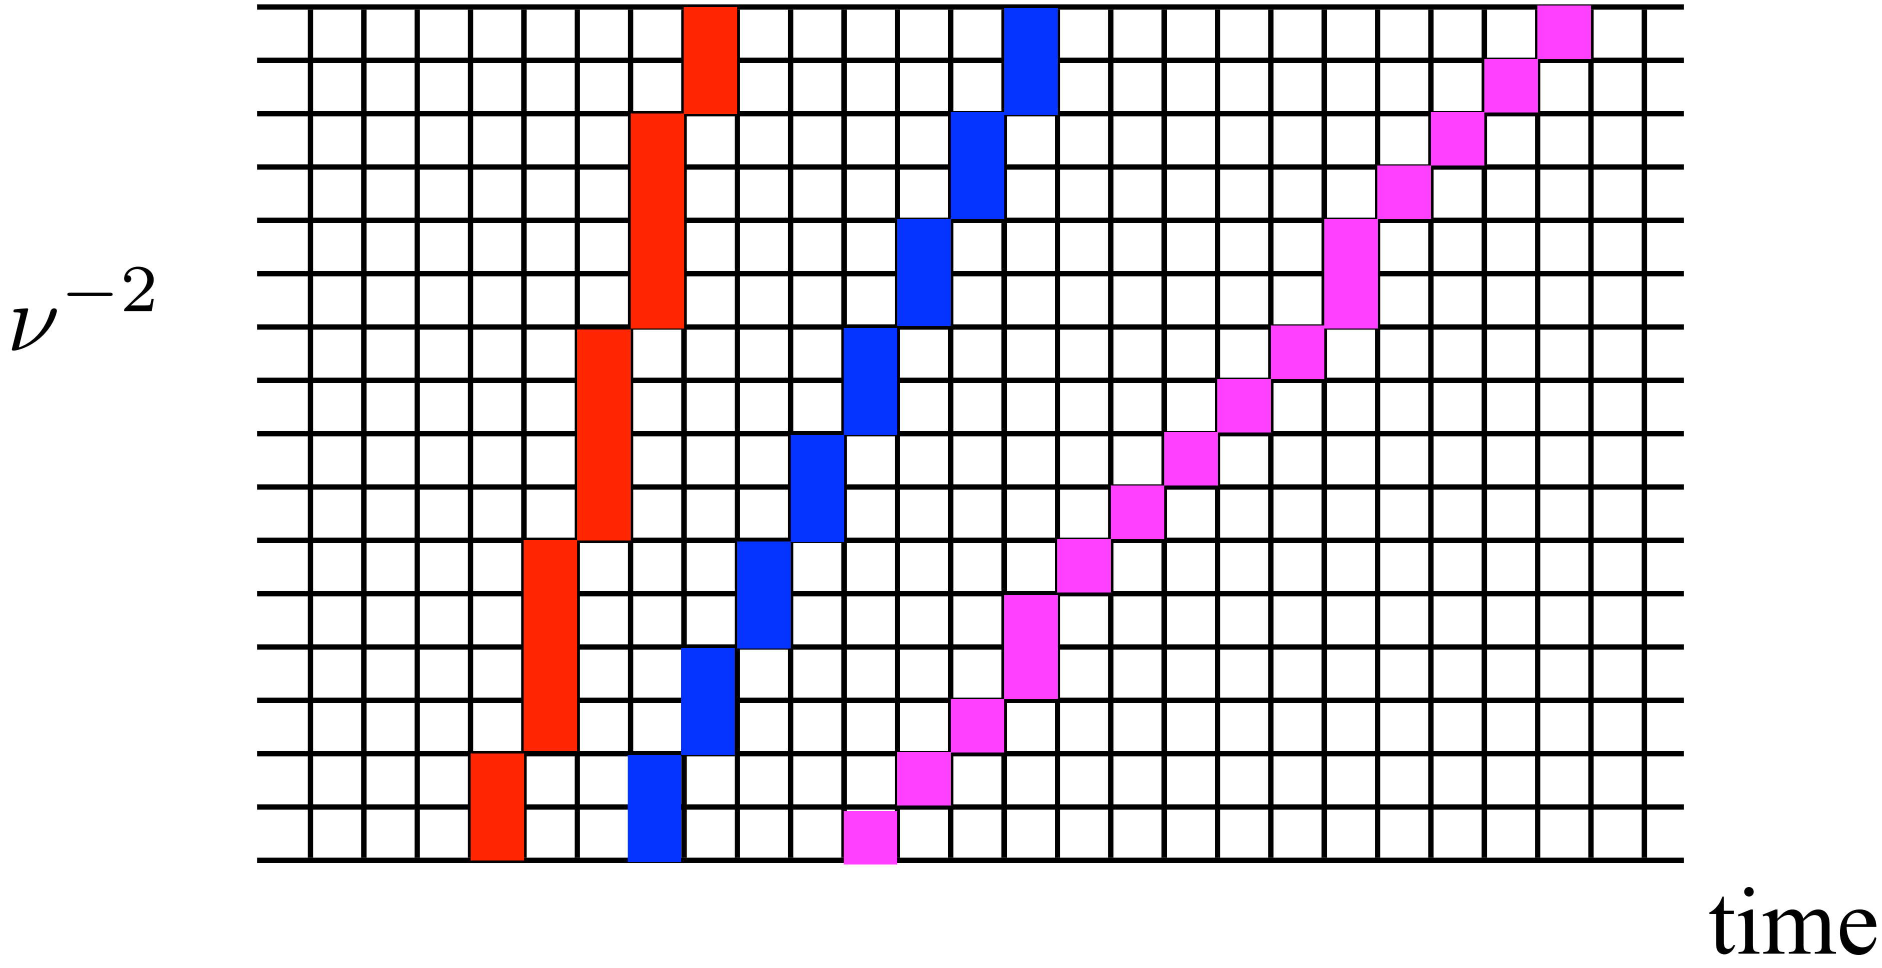

Fast radio bursts

First iteration: group channels in pairs. Within each pair, we form all “vertical” sums (blue) and “diagonal” sums (red).

Output is two arrays, each half the size of the input array.

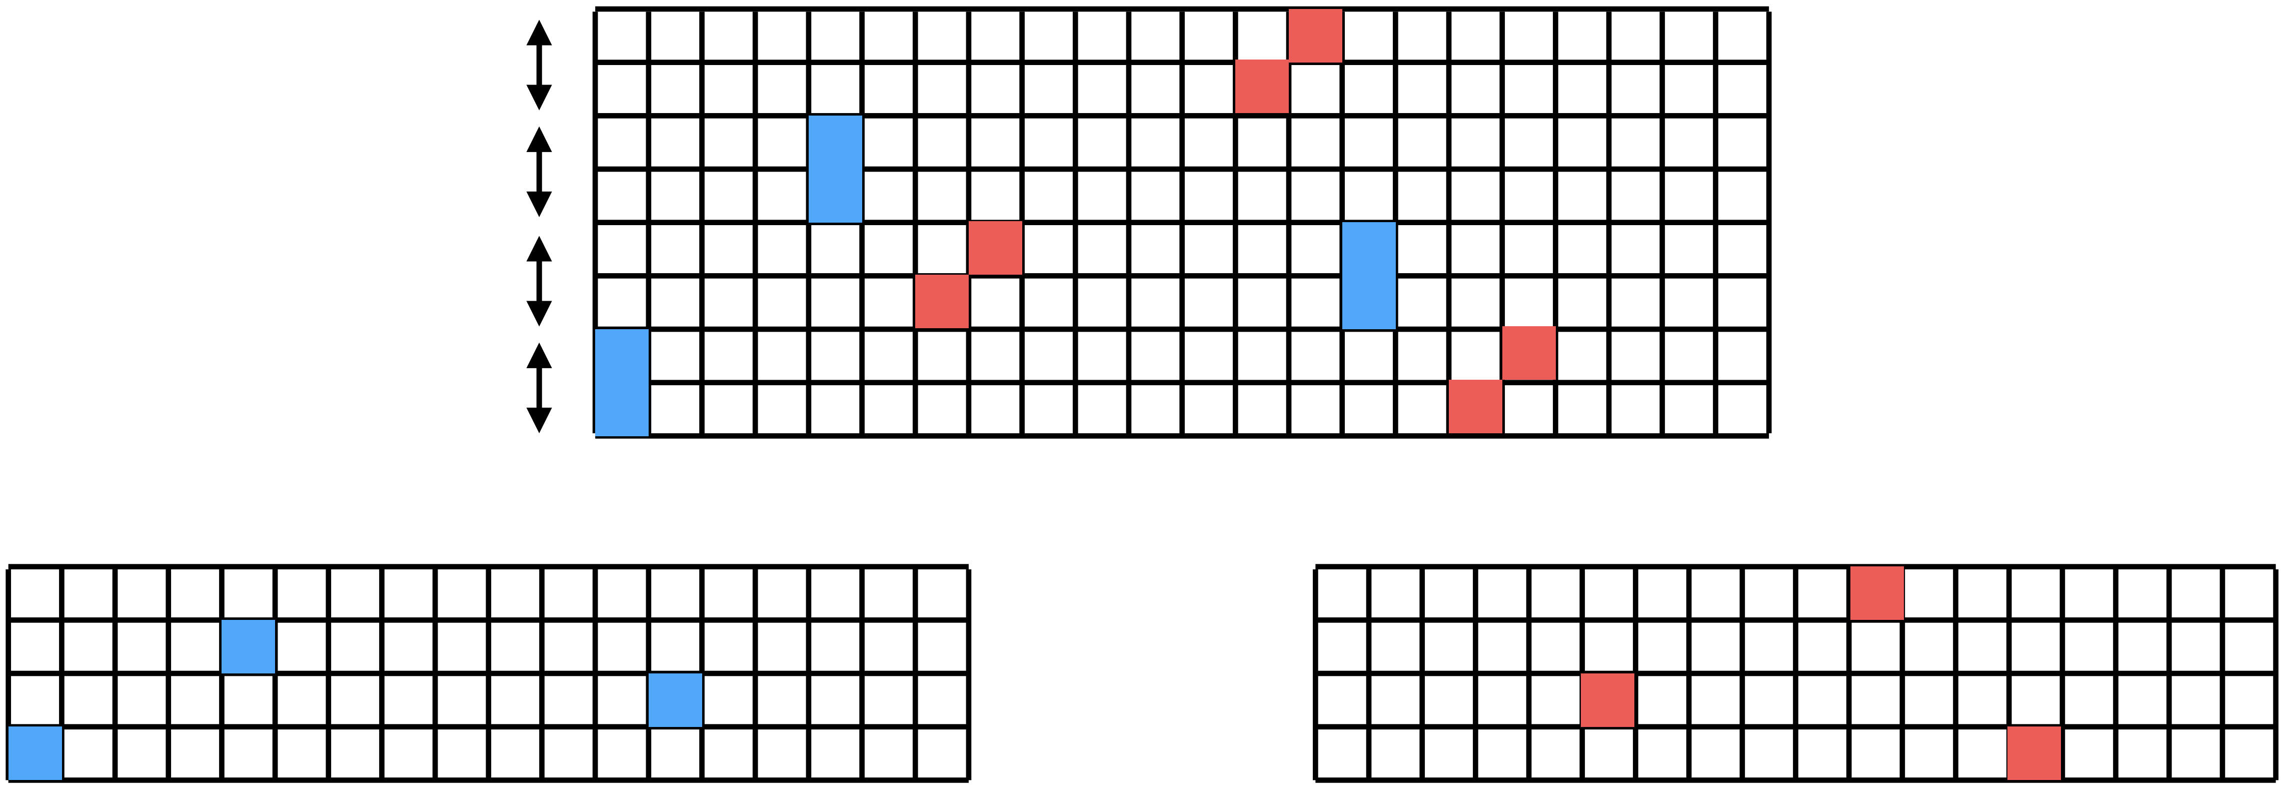

Fast radio bursts

Second iteration: sum pairs into "pairs of pairs".

Frequency channels have now been merged in quadruples. Within each quadruple, there are four possible sums.

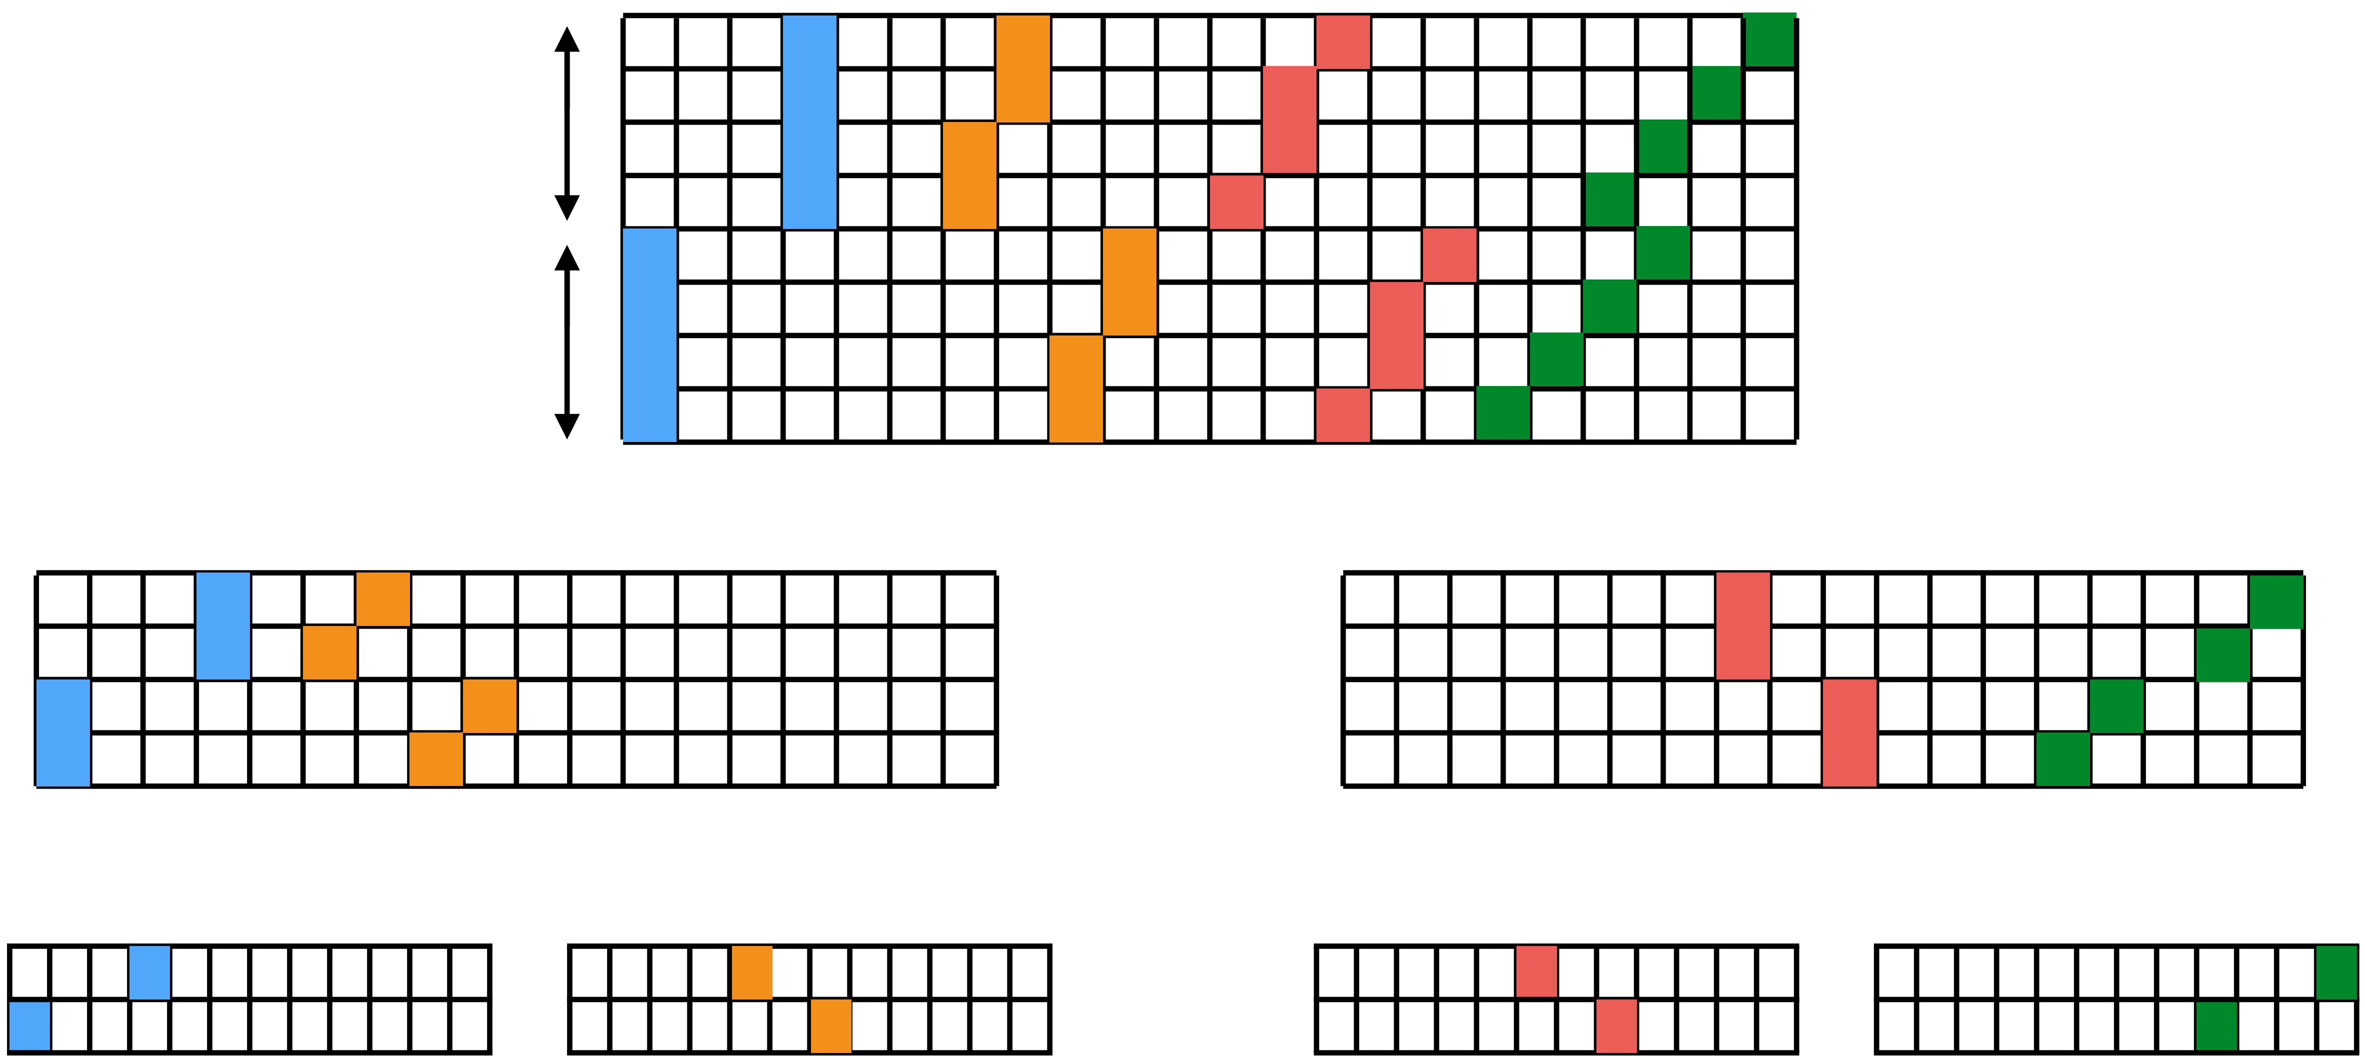

Fast radio bursts

Last iteration: all channels summed.

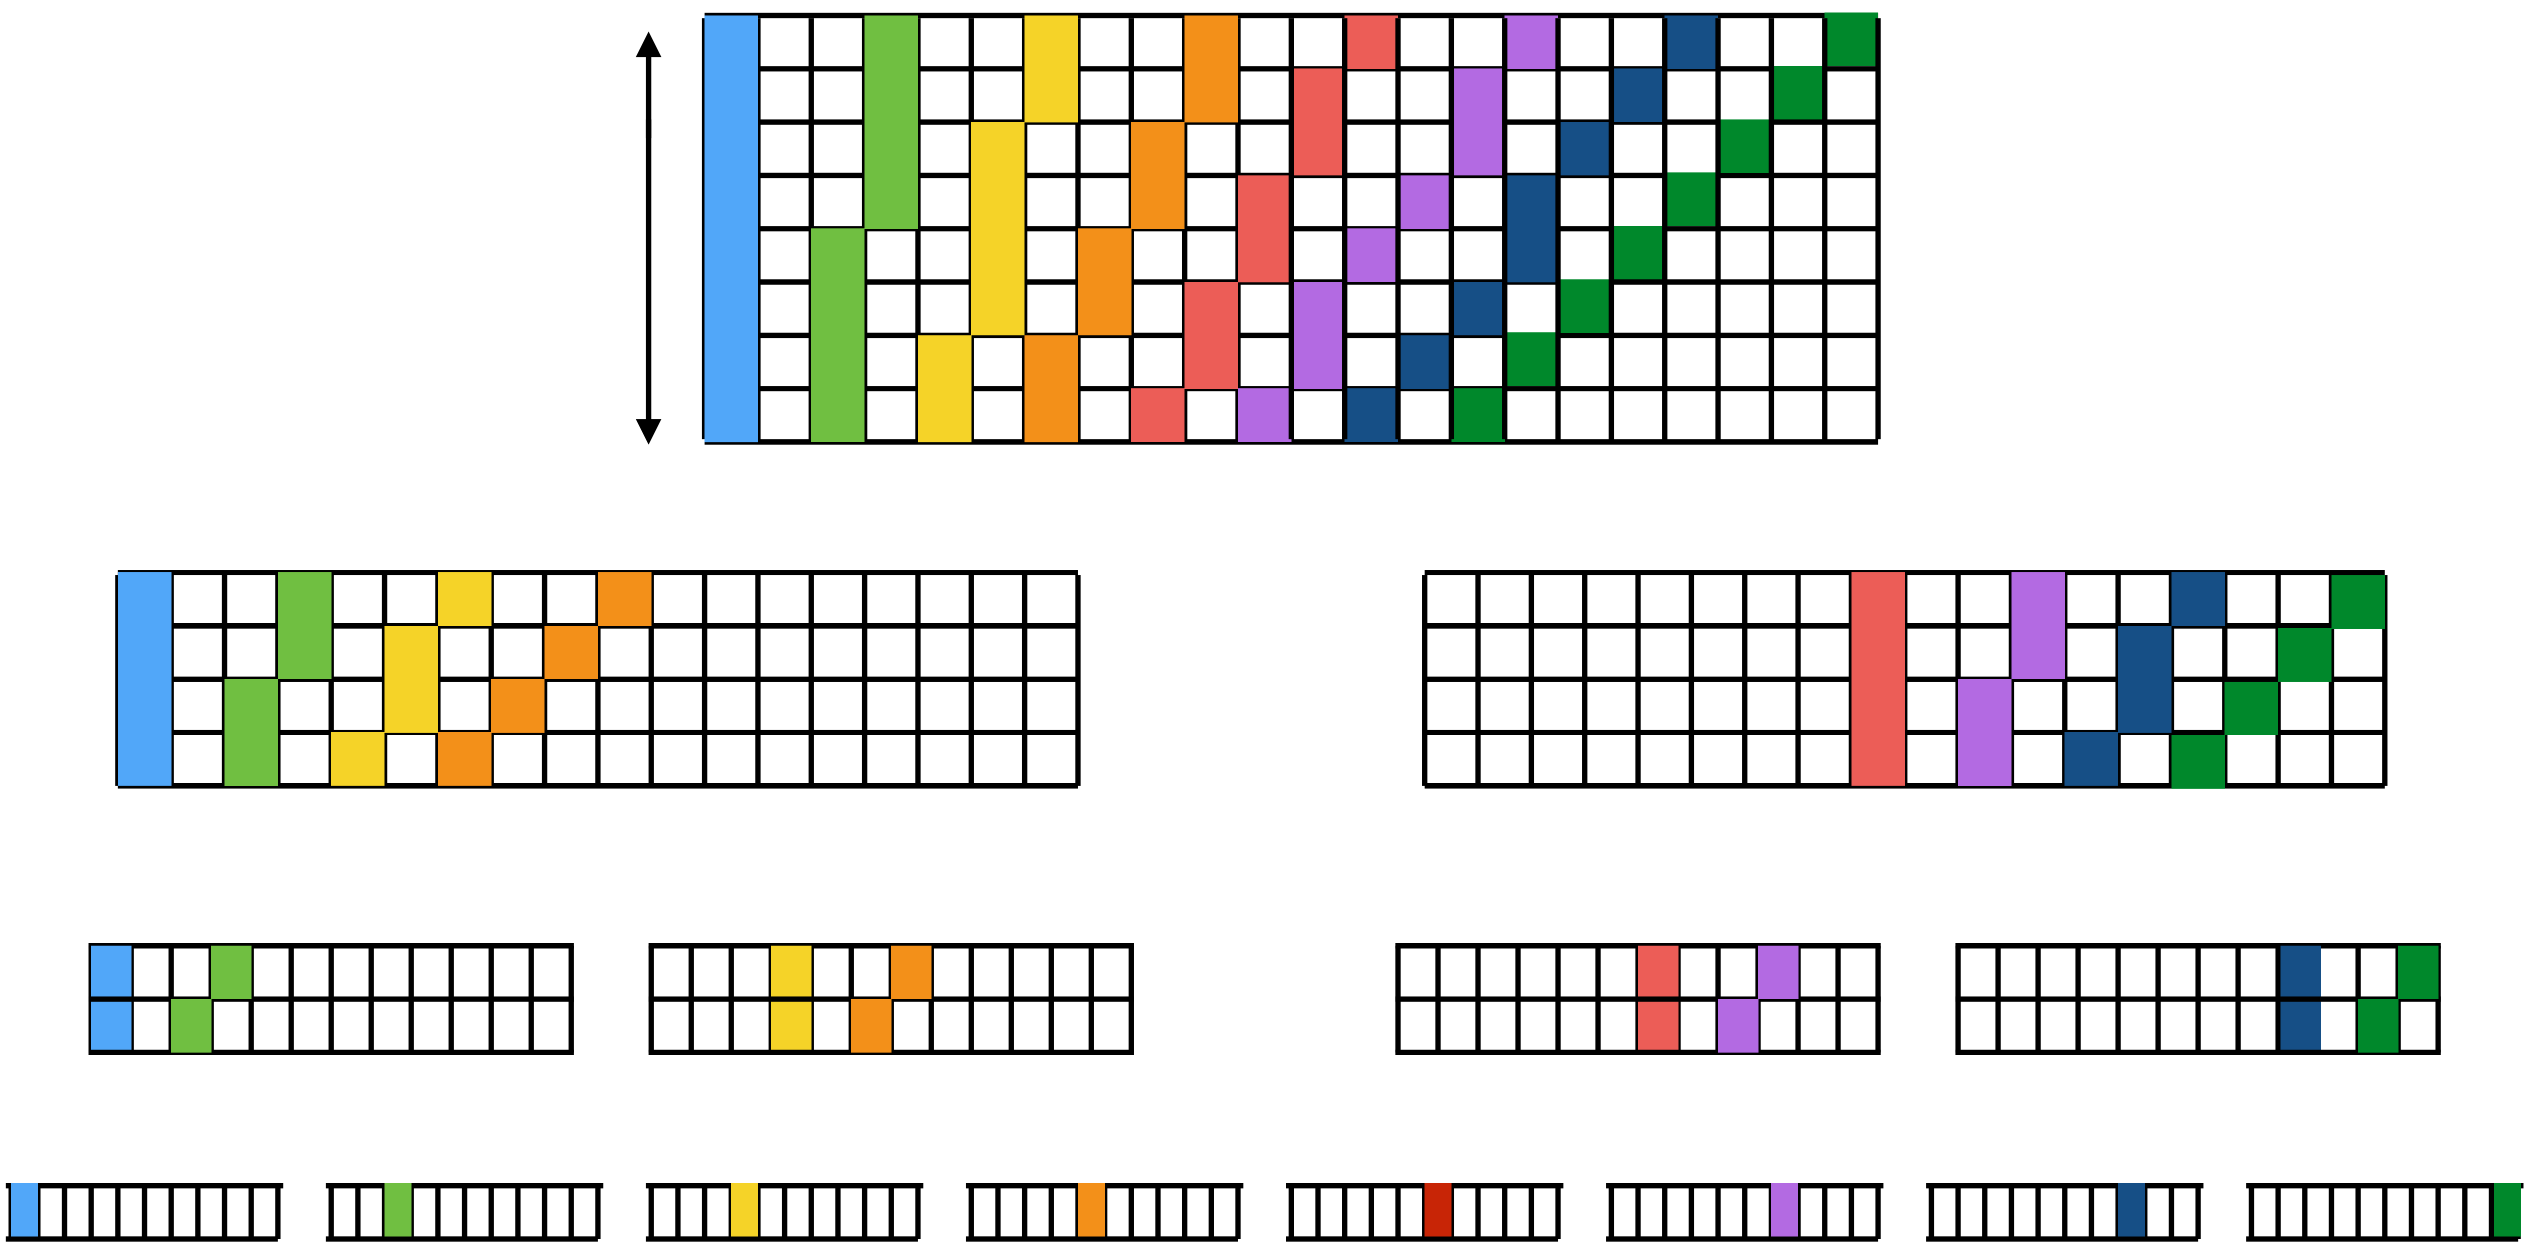

Fast radio bursts

- Lots of technical details to get right! E.g. algorithm can be made close to to statistically optimal, but only if "hidden" choices are made correctly.
- Appears to have a memory bandwidth bottleneck, but this can be solved with some tricks ("blocking" the algorithm)
- To run fast enough for CHIME, needs a lot of low-level black magic (assembly language kernels, etc!)
- Bottom line: we can do a near-optimal CHIME FRB search on a ~1200-core cluster, for around ~$300K.
- We are building this cluster now, and expect to find ~10 FRB's per day, starting this summer!
- Public code coming soon…

=  CHIME for FRB's?

=  CHIME for FRB’s? Yes!

Searching for pulsars

In principle, CHIME has the sensitivity to find many new pulsars!
"One GBNCC per day"

However, a blind pulsar search in a CHIME-sized dataset is an unsolved problem. Existing algorithms are much too slow.

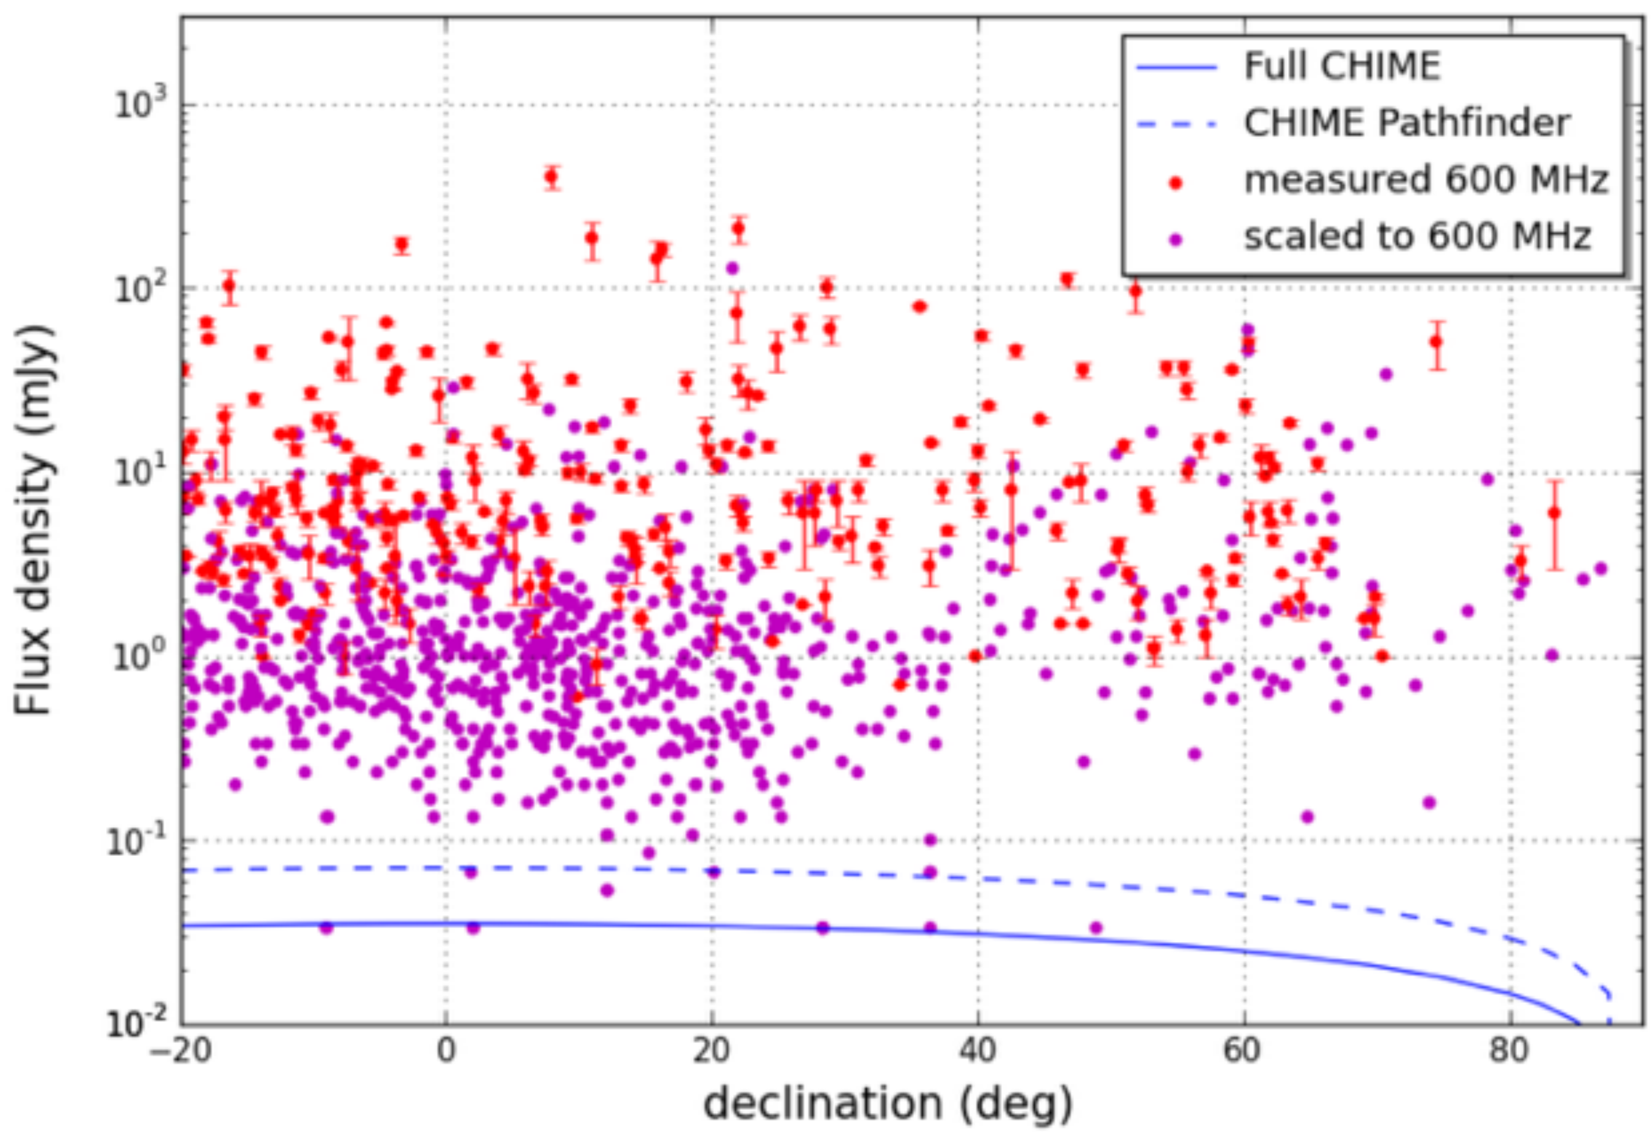

The pulsar search problem

Searching for a quasiperiodic series of pulses in a noisy intensity timestream I(t).

"Quasiperiodic": frequency slowly varies with time, e.g. due to Doppler shift in a binary system.

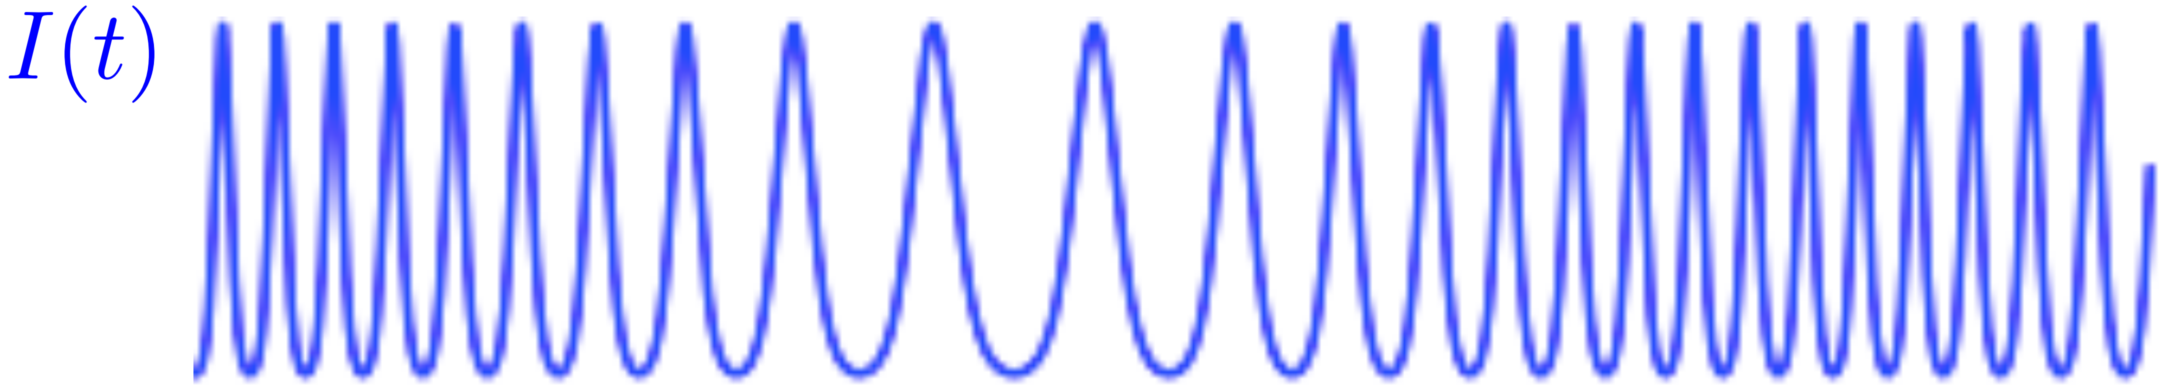

The phase model $\Phi(t)$ is a dimensionless function of time, such that pulses appear when Φ is a multiple of 2π .

Example: regular pulsar, linearly evolving phase $\Phi(t) = \omega t$

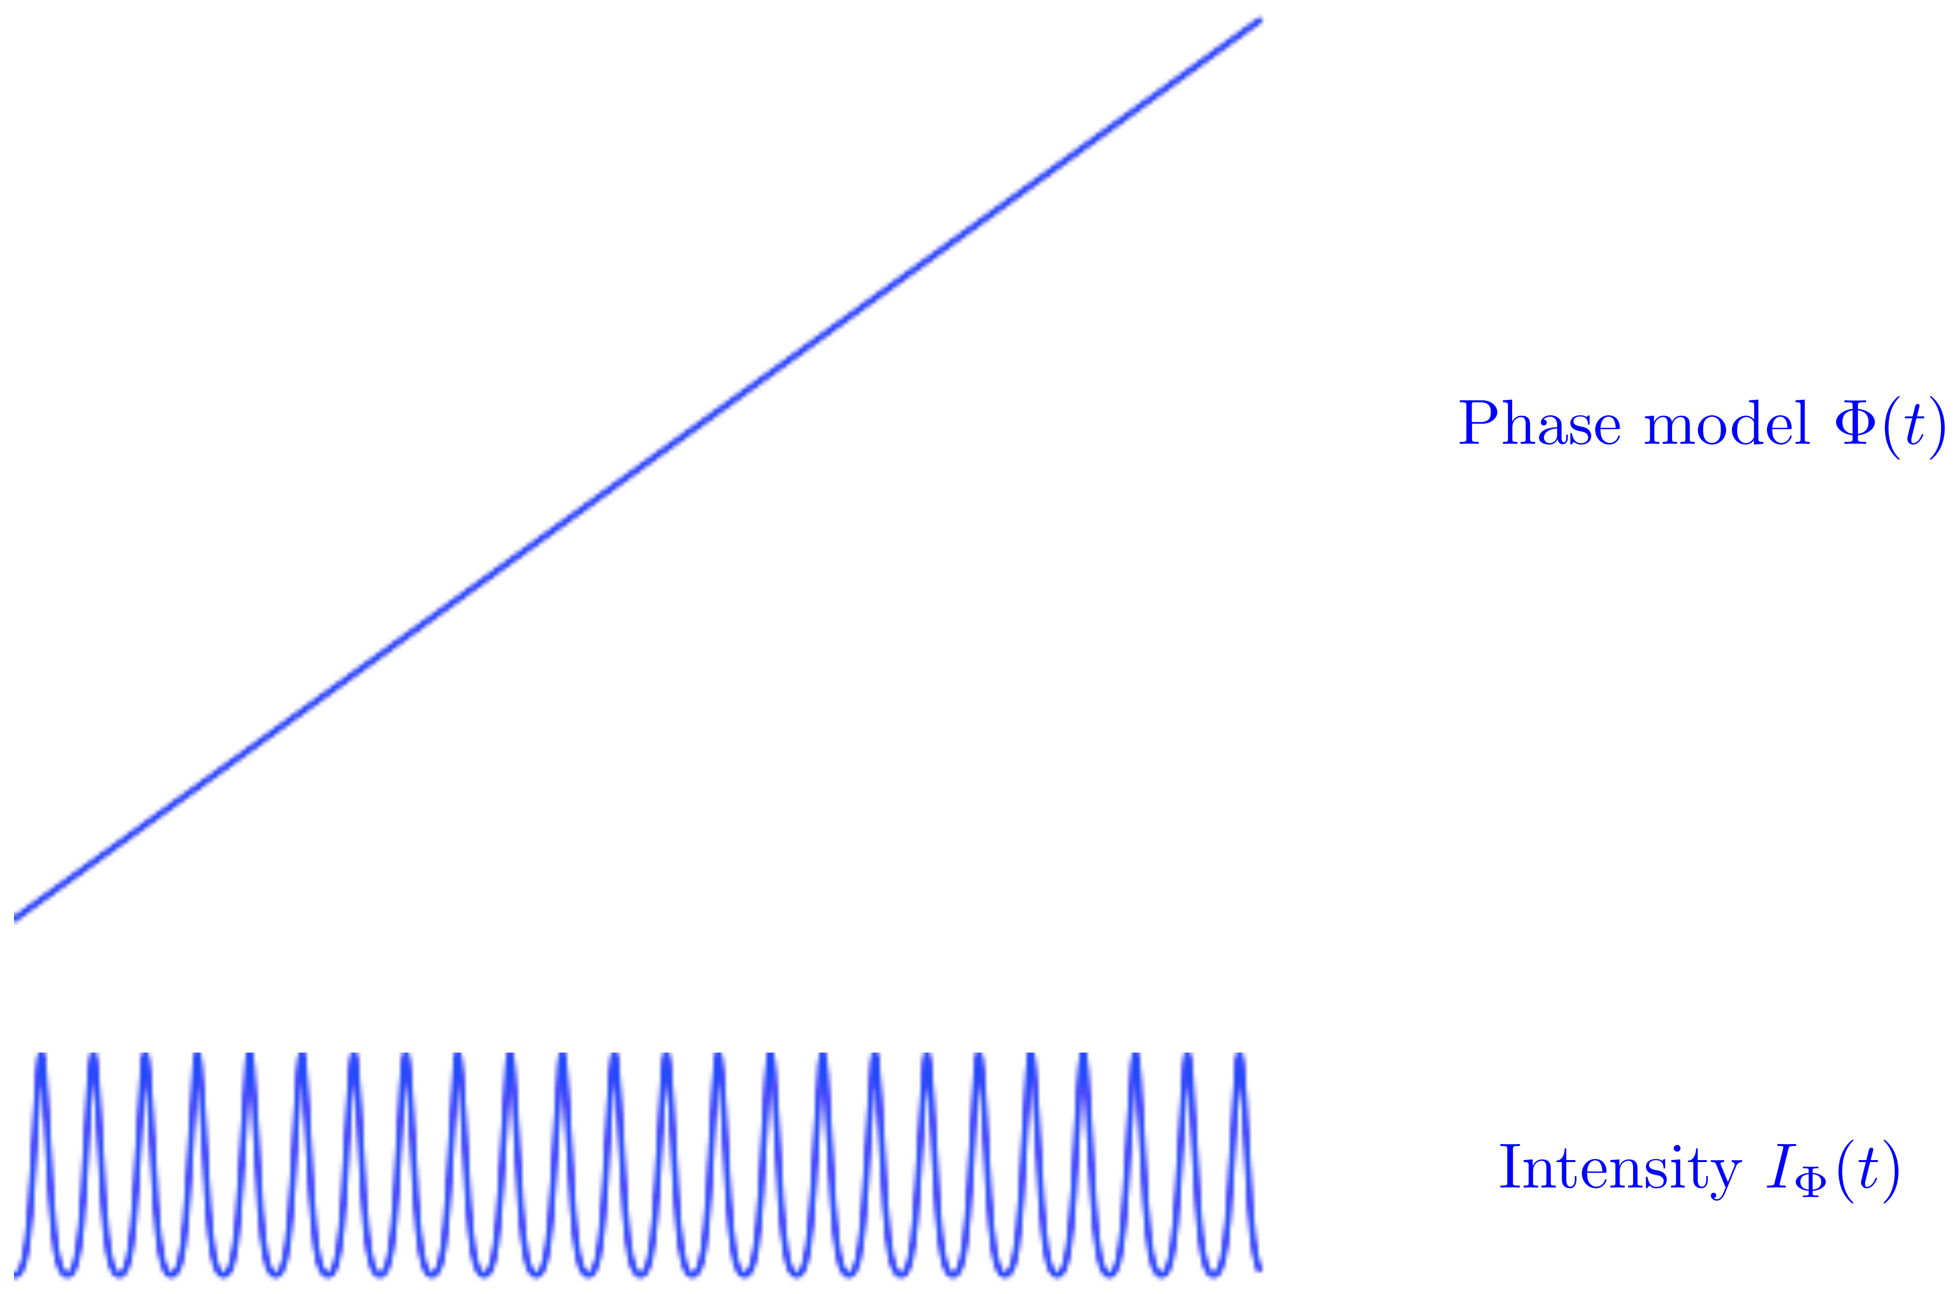

Decelerating pulsar $\Phi(t) = \omega t - \alpha t^2$

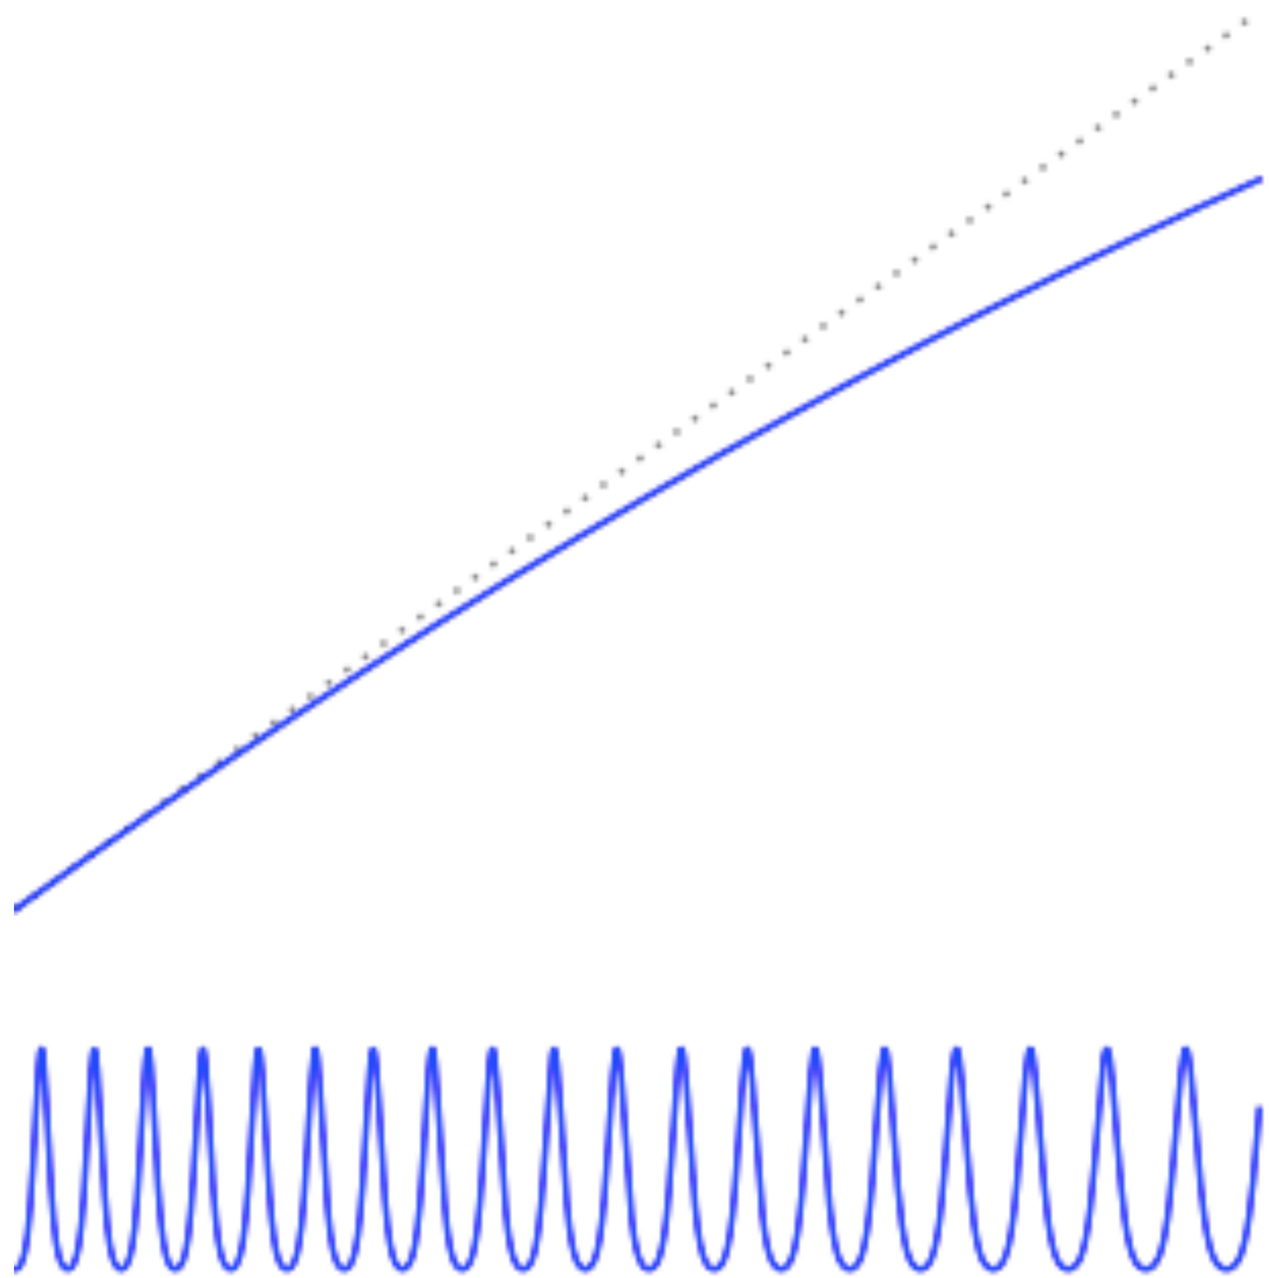

Phase model $\Phi(t)$

Intensity $I_\Phi(t)$

Binary pulsar (detailed phase model contains many parameters!)

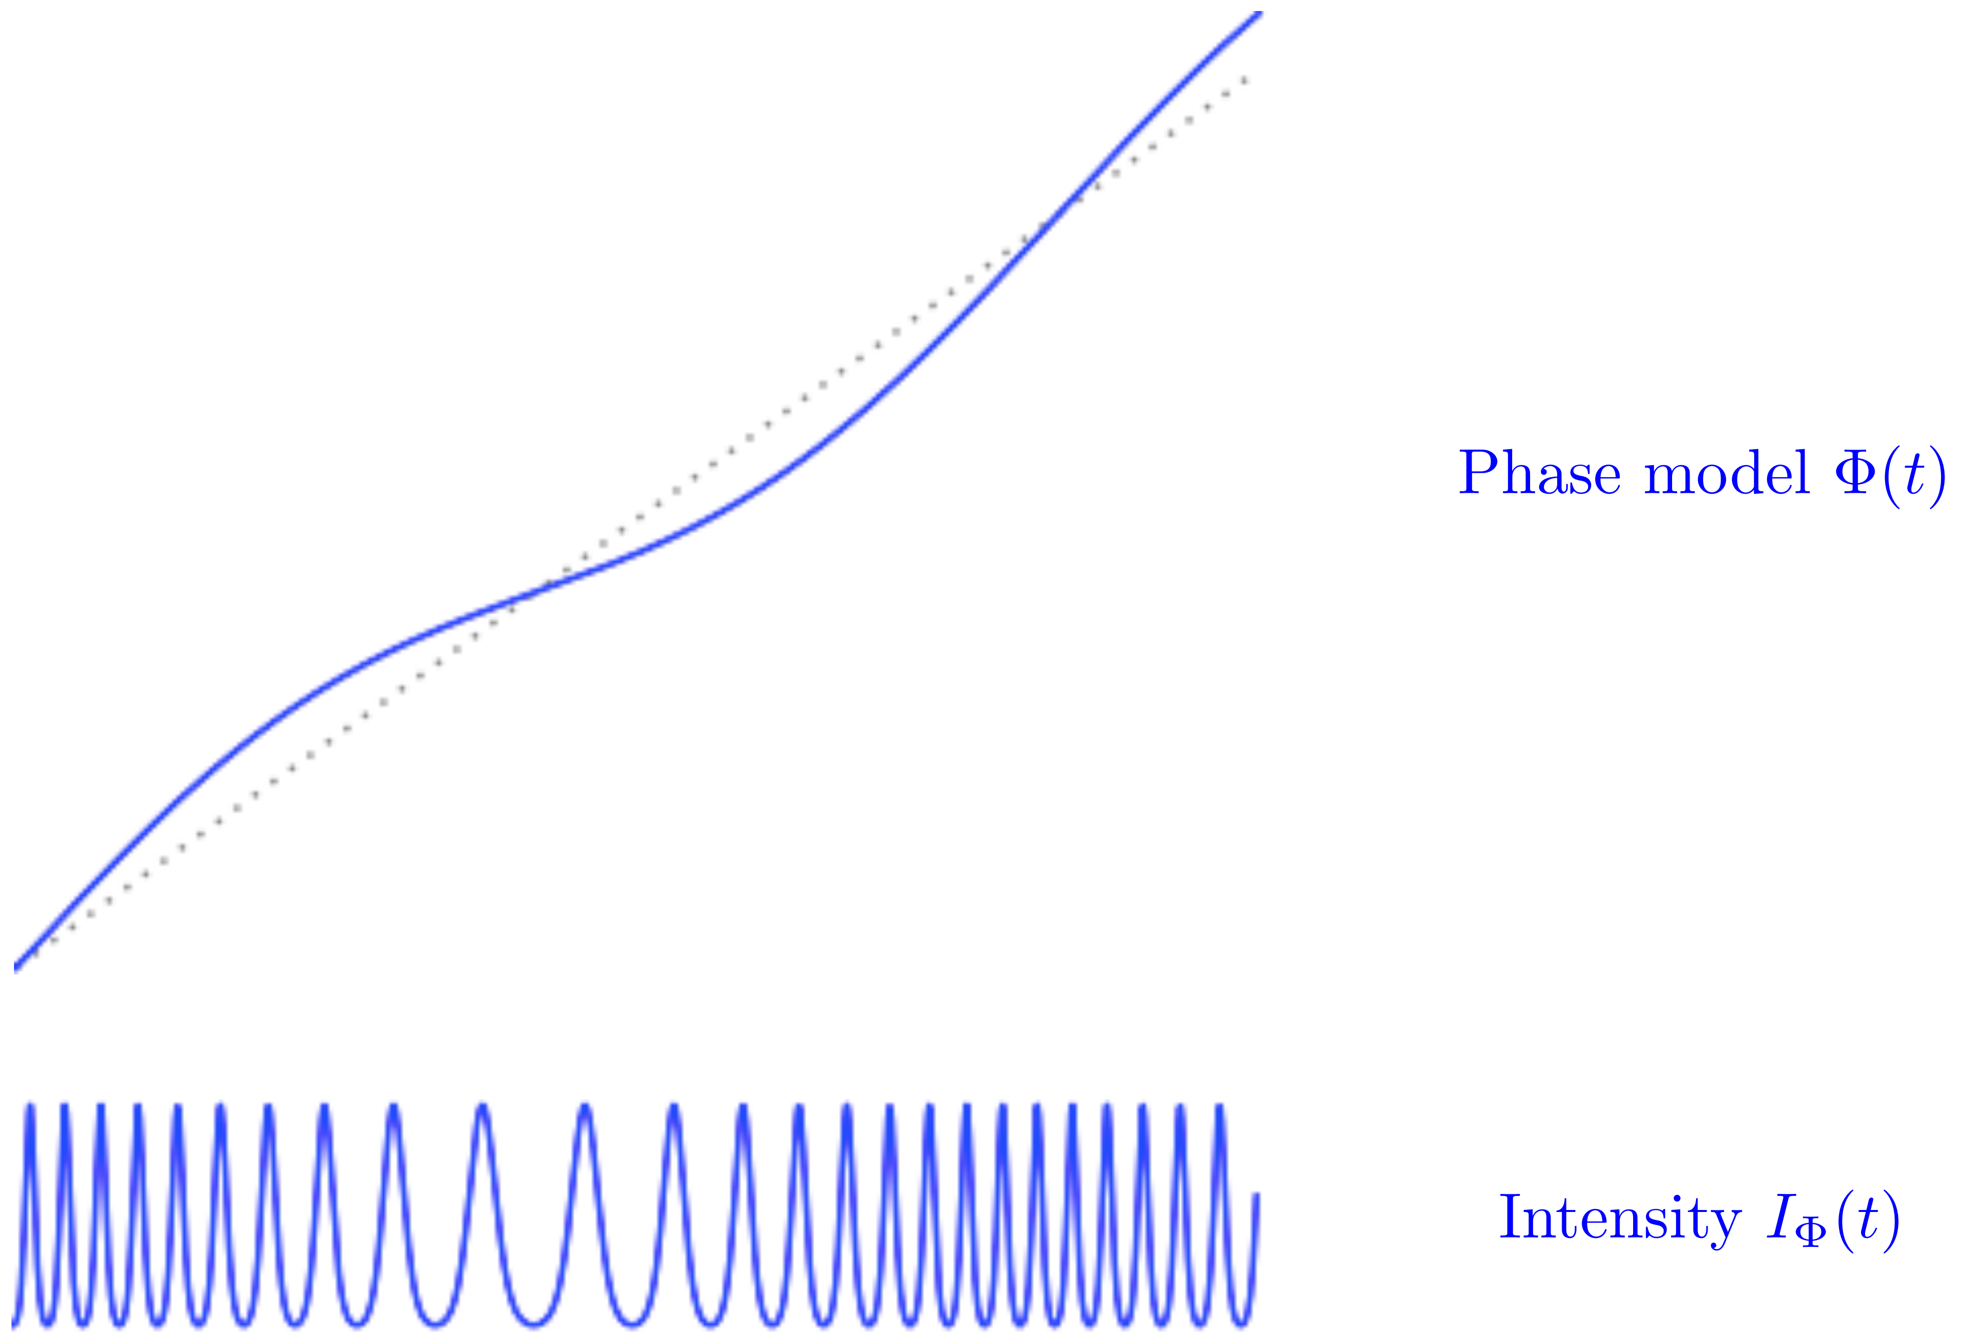

Fast coherent search

The optimal search algorithm is a "coherent search": loop over all possible phase models Φ and compute

$$\hat{\mathcal{E}}[\Phi] = \int d(t) I_\Phi(t)\, dt \quad \text{where} \quad \left\{ \begin{array}{rcl} d(t) & = & \text{data timestream} \\ I_\Phi(t) & = & \text{model timestream} \end{array} \right.$$

Brute-force computational cost is O(ST), where:

S = size of the search space (# of independent phase models)
T = size of timestream

Fast coherent search algorithms (KMS, arxiv:1610.06381) can do this search with computational cost O(S)!

[Current practice is to do a suboptimal search with cost O(S log T)]

Why pulsar search is so hard

The size S of the search space is a very rapidly increasing function of the timestream size T.

Example: consider modeling $\Phi(t)$ by a low-order polynomial.

In CHIME, S ~ 10^{30}! New ideas are needed.

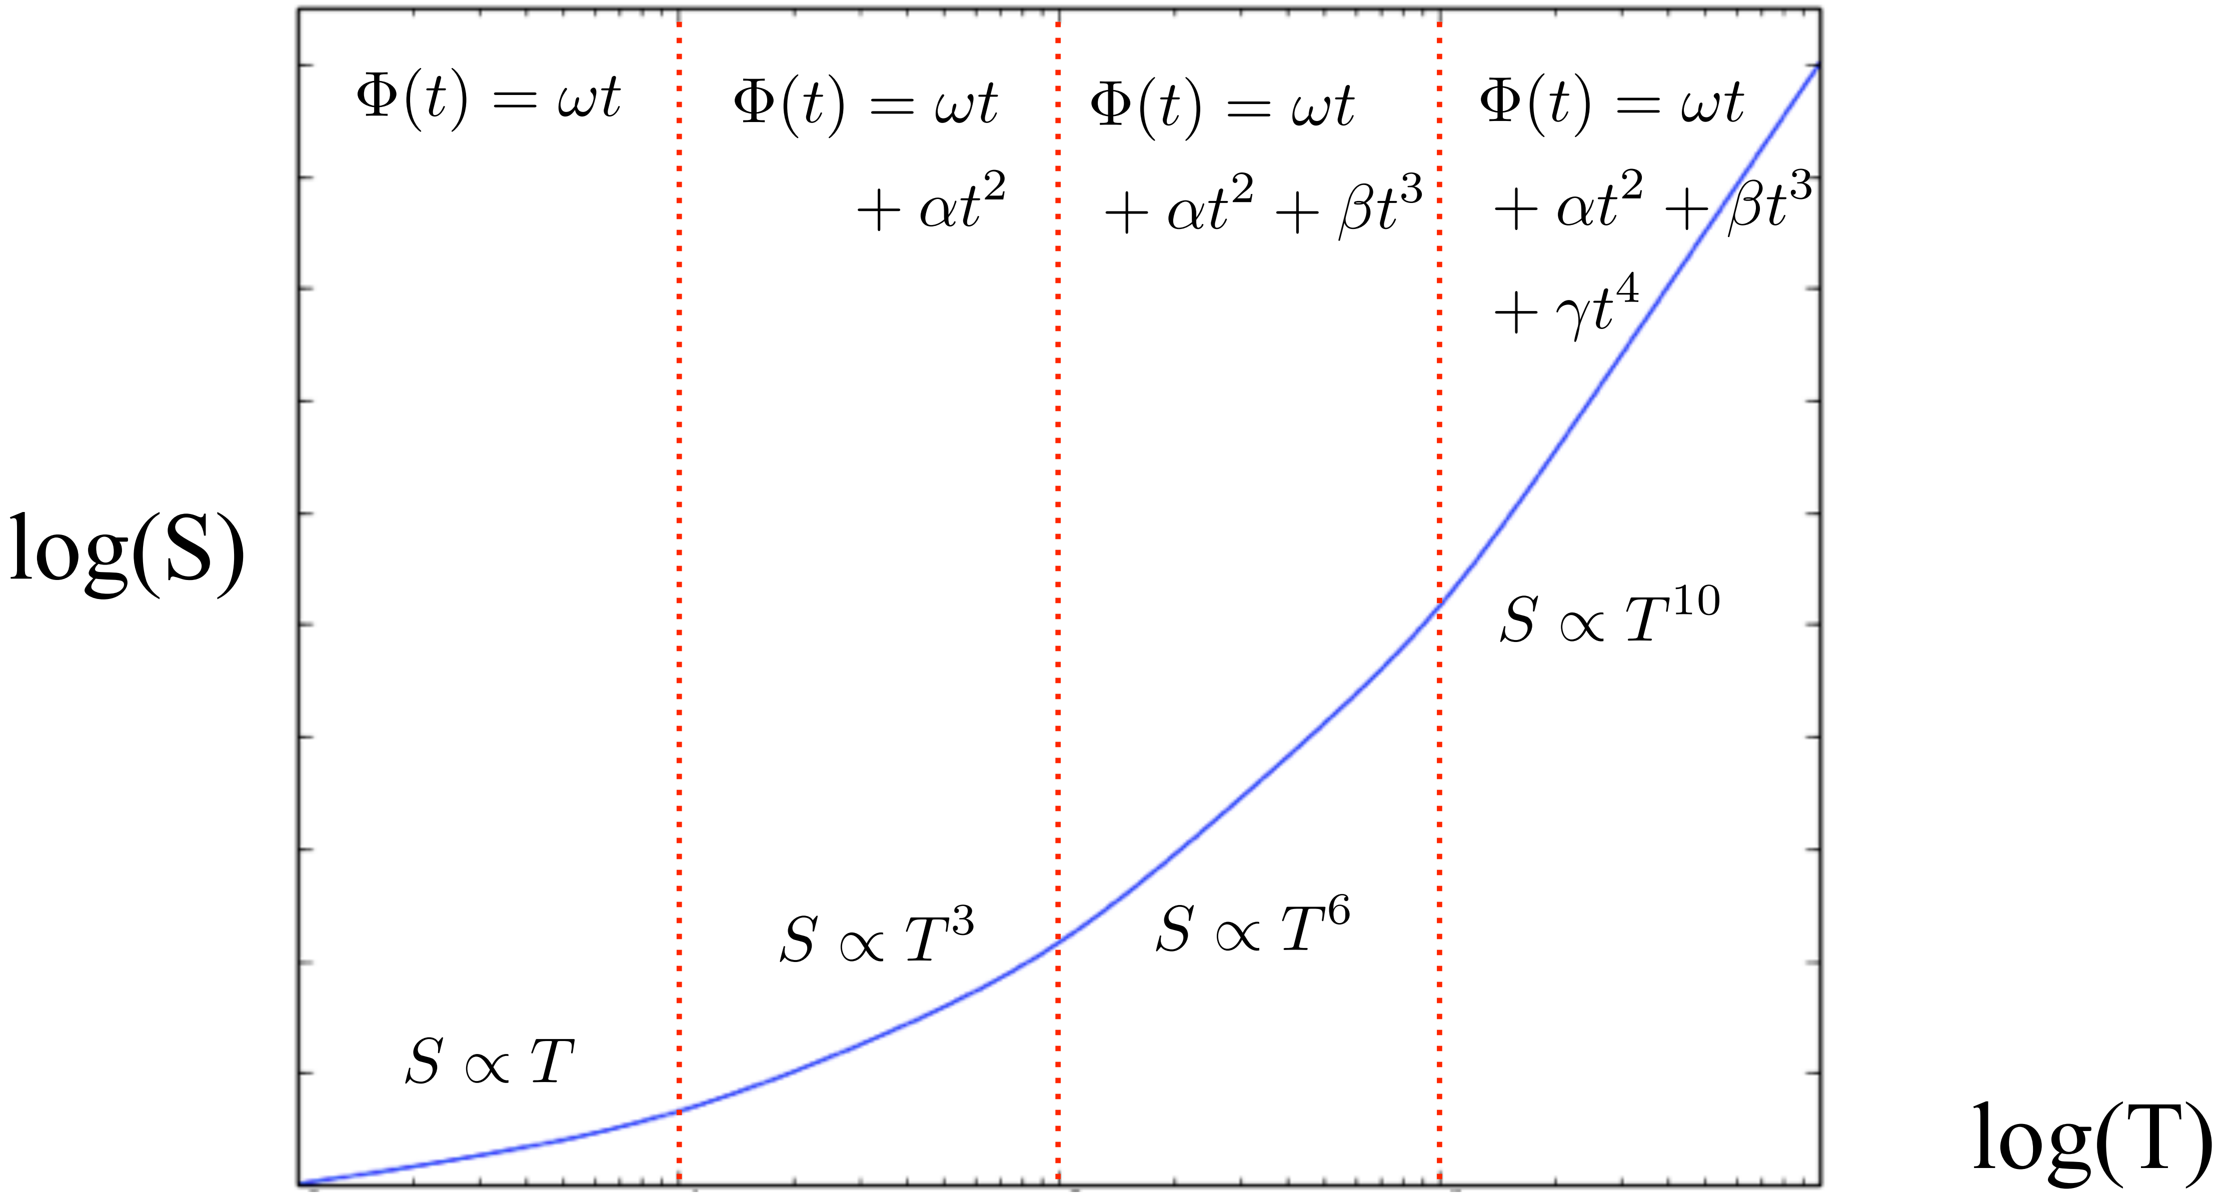

Semicoherent search

Divide the timestream into chunks of a fixed size. Consider "jagged" phase models whose acceleration $\ddot{\Phi}$ is constant in each chunk, but changes by $\pm\epsilon$ at chunk boundaries (Φ, $\dot{\Phi}$ evolve continuously)

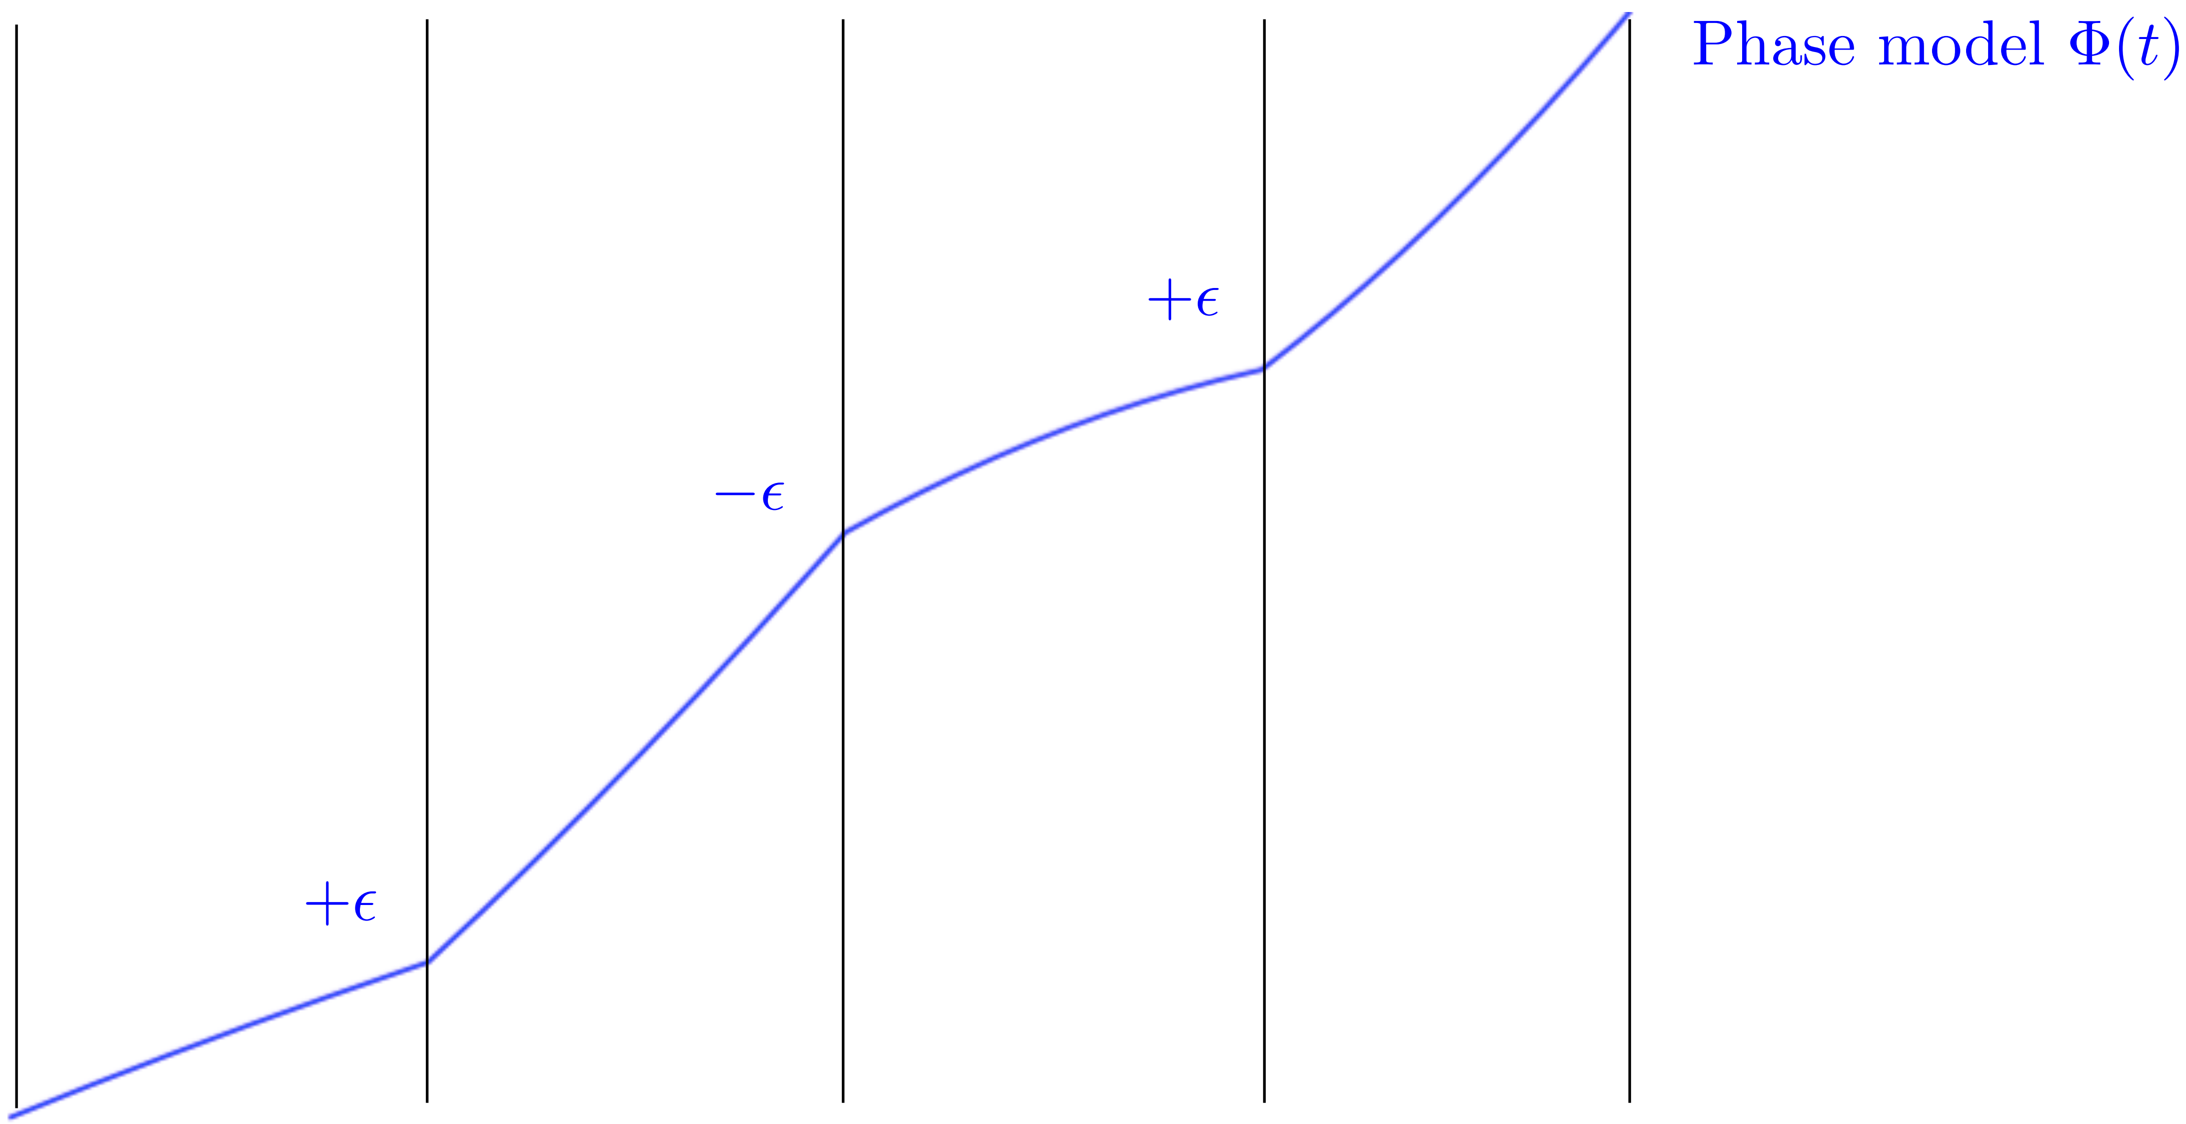

For initial data $(\Phi_0, \dot{\Phi}_0,,\ddot{\Phi}_0)$ define a statistic $\hat{\mathcal{H}}(\Phi_0, \dot{\Phi}_0, \ddot{\Phi}_0)$ which sums over 2^N jagged paths with initial conditions $(\Phi_0, \dot{\Phi}_0, \ddot{\Phi}_0)$

$$\hat{\mathcal{H}} = \sum_{\Phi(t)} \exp(r\hat{\mathcal{E}}[\Phi])$$

where $\hat{\mathcal{E}}[\Phi] = \int dt\, d(t) I_\Phi(t)$ is the coherent search statistic.

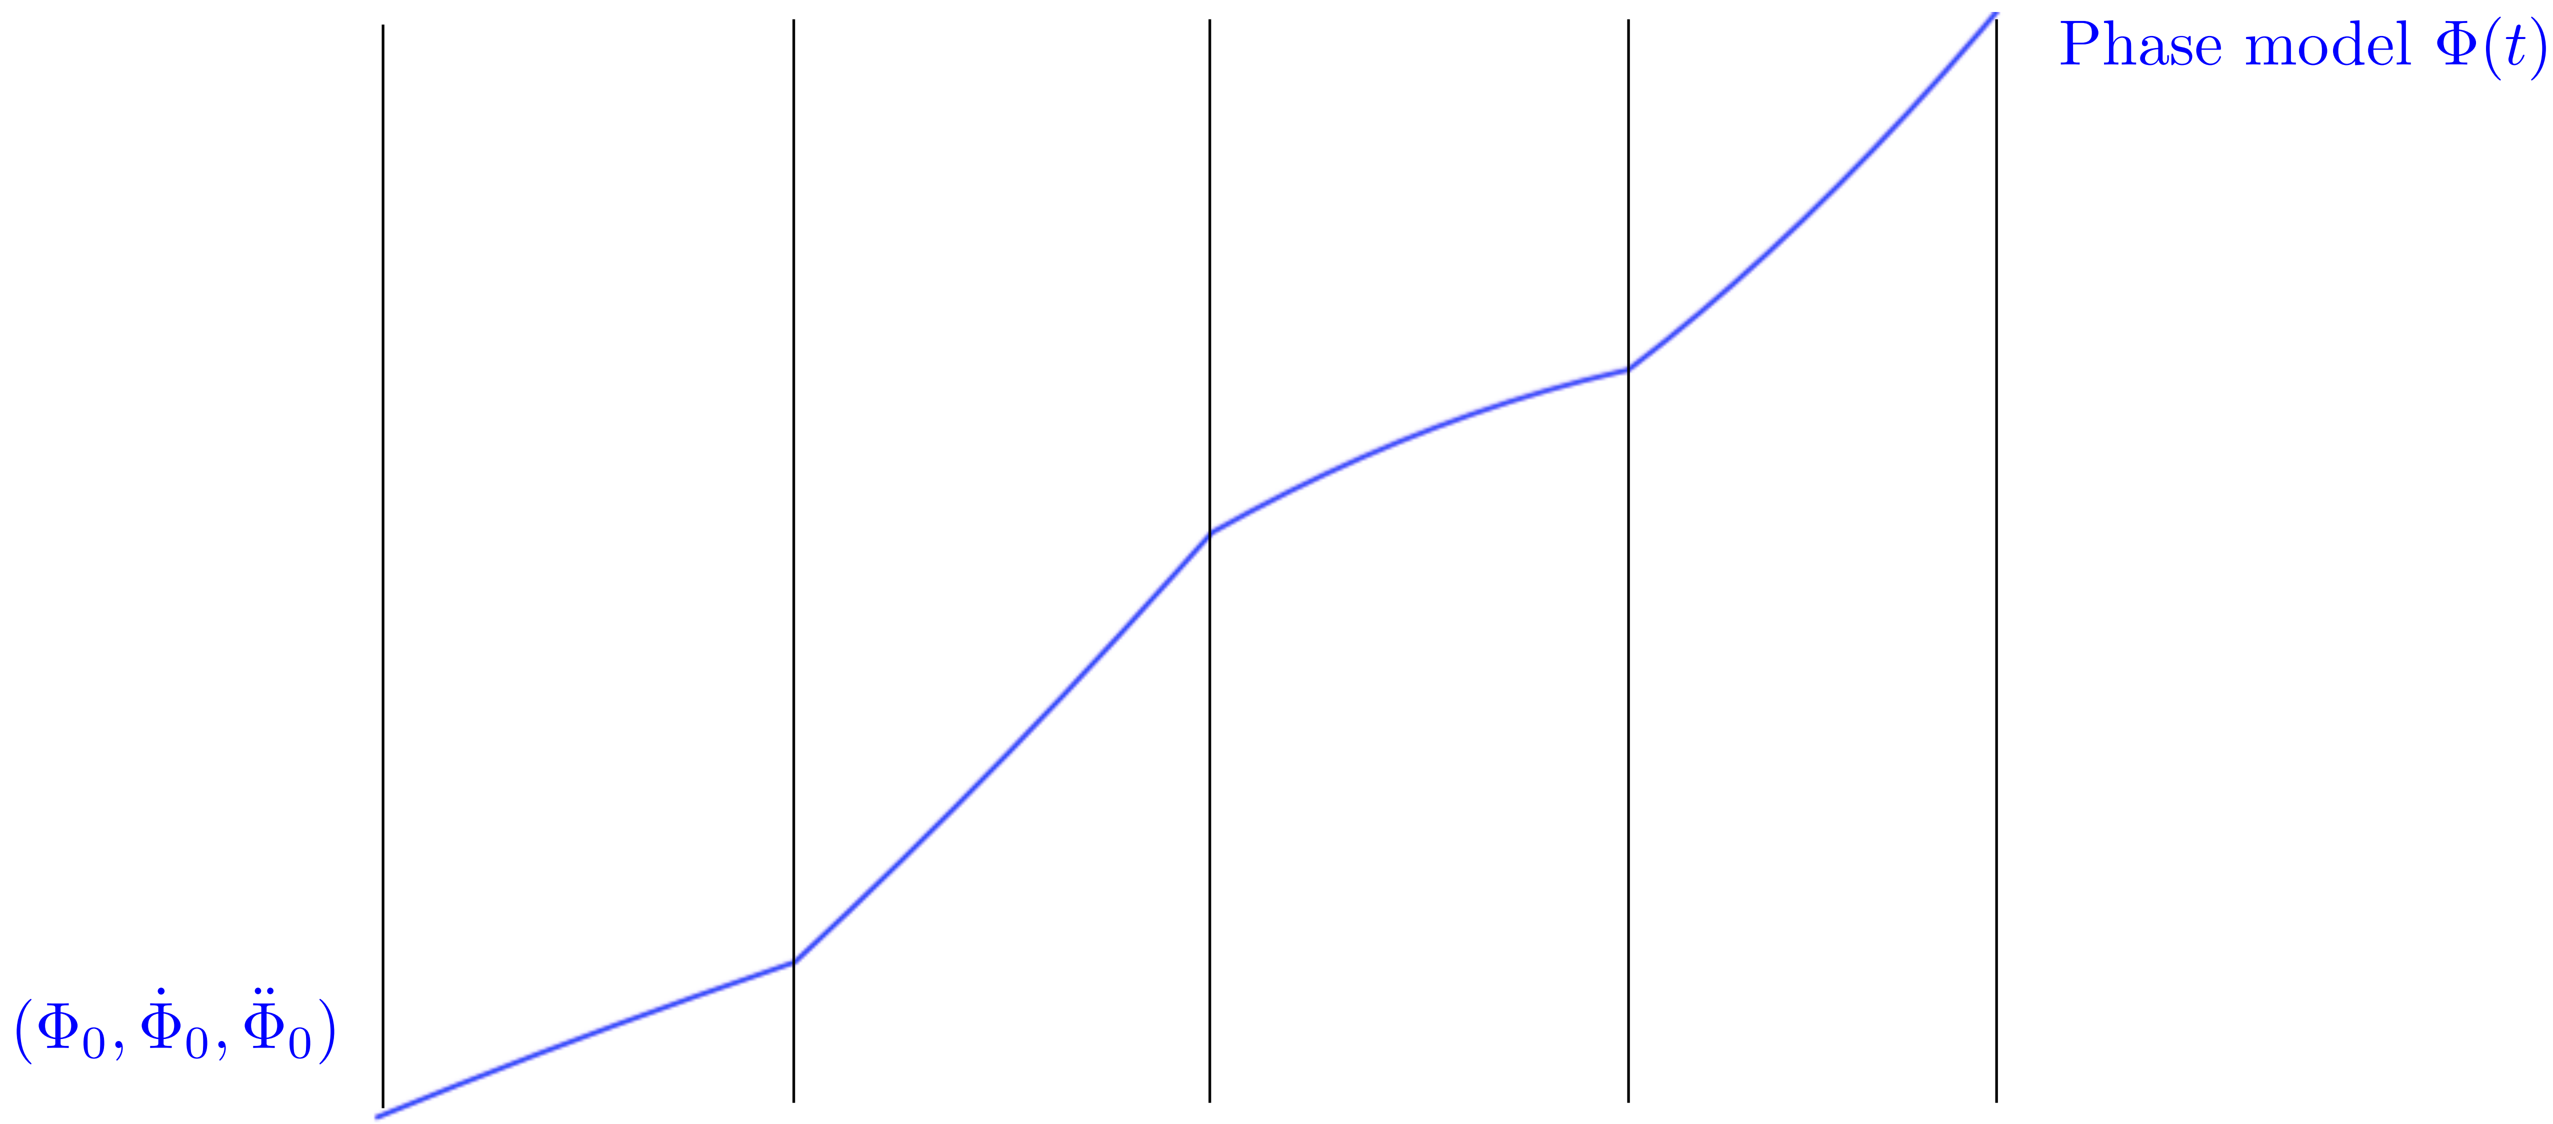

The $\hat{\mathcal{H}}$-search is fully coherent within each chunk. On longer timescales it sums over all ways of connecting coherent subsearches in a consistent way. ("jagged paths")

Key fact: $\hat{\mathcal{H}}$ is computable in O(N) time, using recursion relations. Cost grows linearly with timestream size T, provided $T \gg T_{\mathrm{chunk}}$

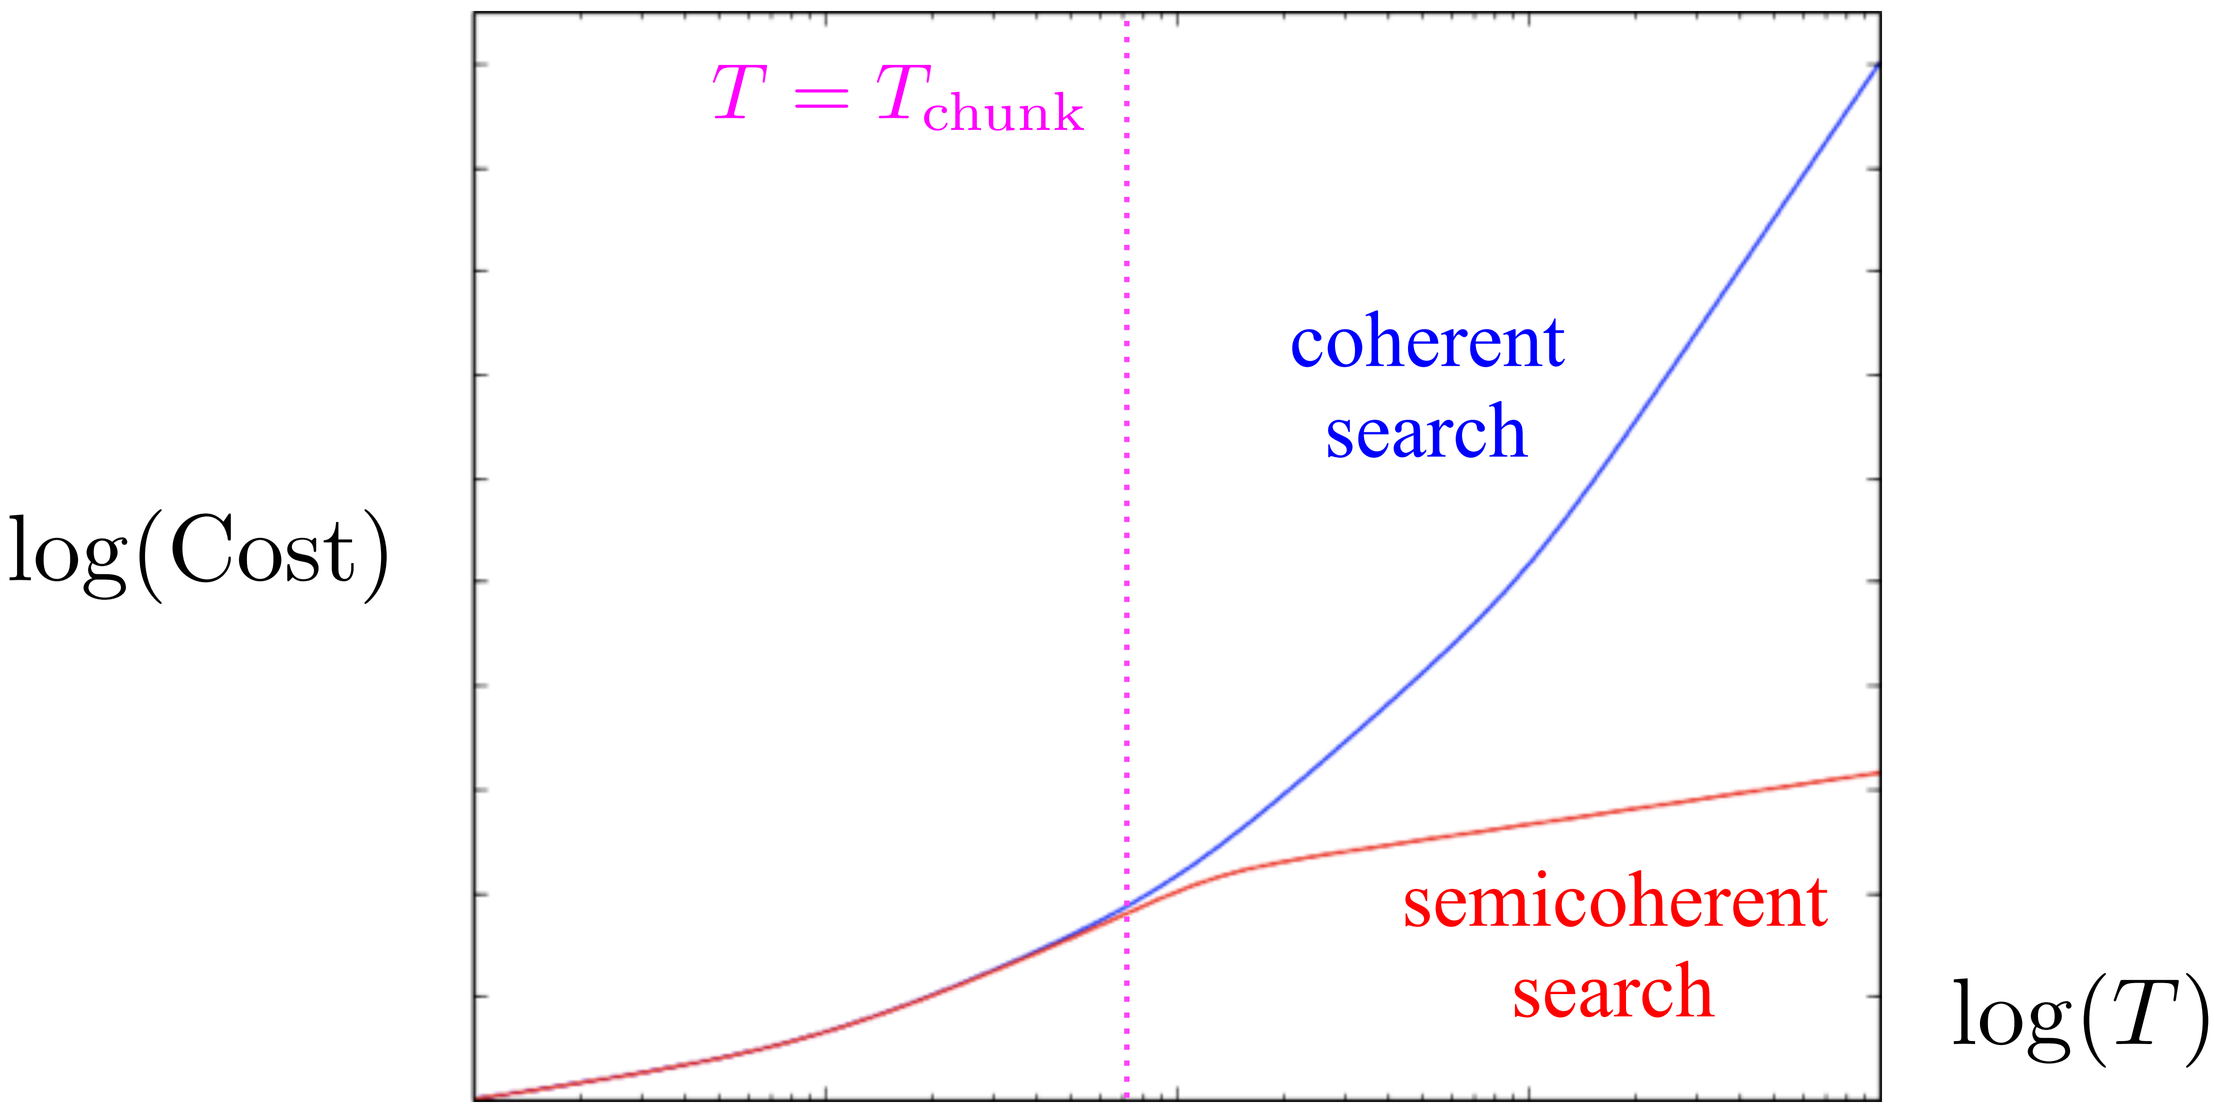

Semicoherent search

- A crucial question: how suboptimal is the semicoherent search?
- In simple examples, nearly optimal! E.g. for ~20% N_{chunks}=64, only ~20% suboptimal, but many orders of magnitude faster.
- These algorithms seem very interesting for upcoming contiguous-timestream experiments such as SKA.
- In CHIME, there is a new detail: discontiguous daily observations. We haven't studied this case yet, but we're working on it!

=  CHIME for pulsar search?

=  CHIME for pulsar search?

Looks promising, but we don't know yet!

Timing known pulsars

Daily timing of known pulsars can contribute to global efforts to detect gravity waves using pulsar networks.

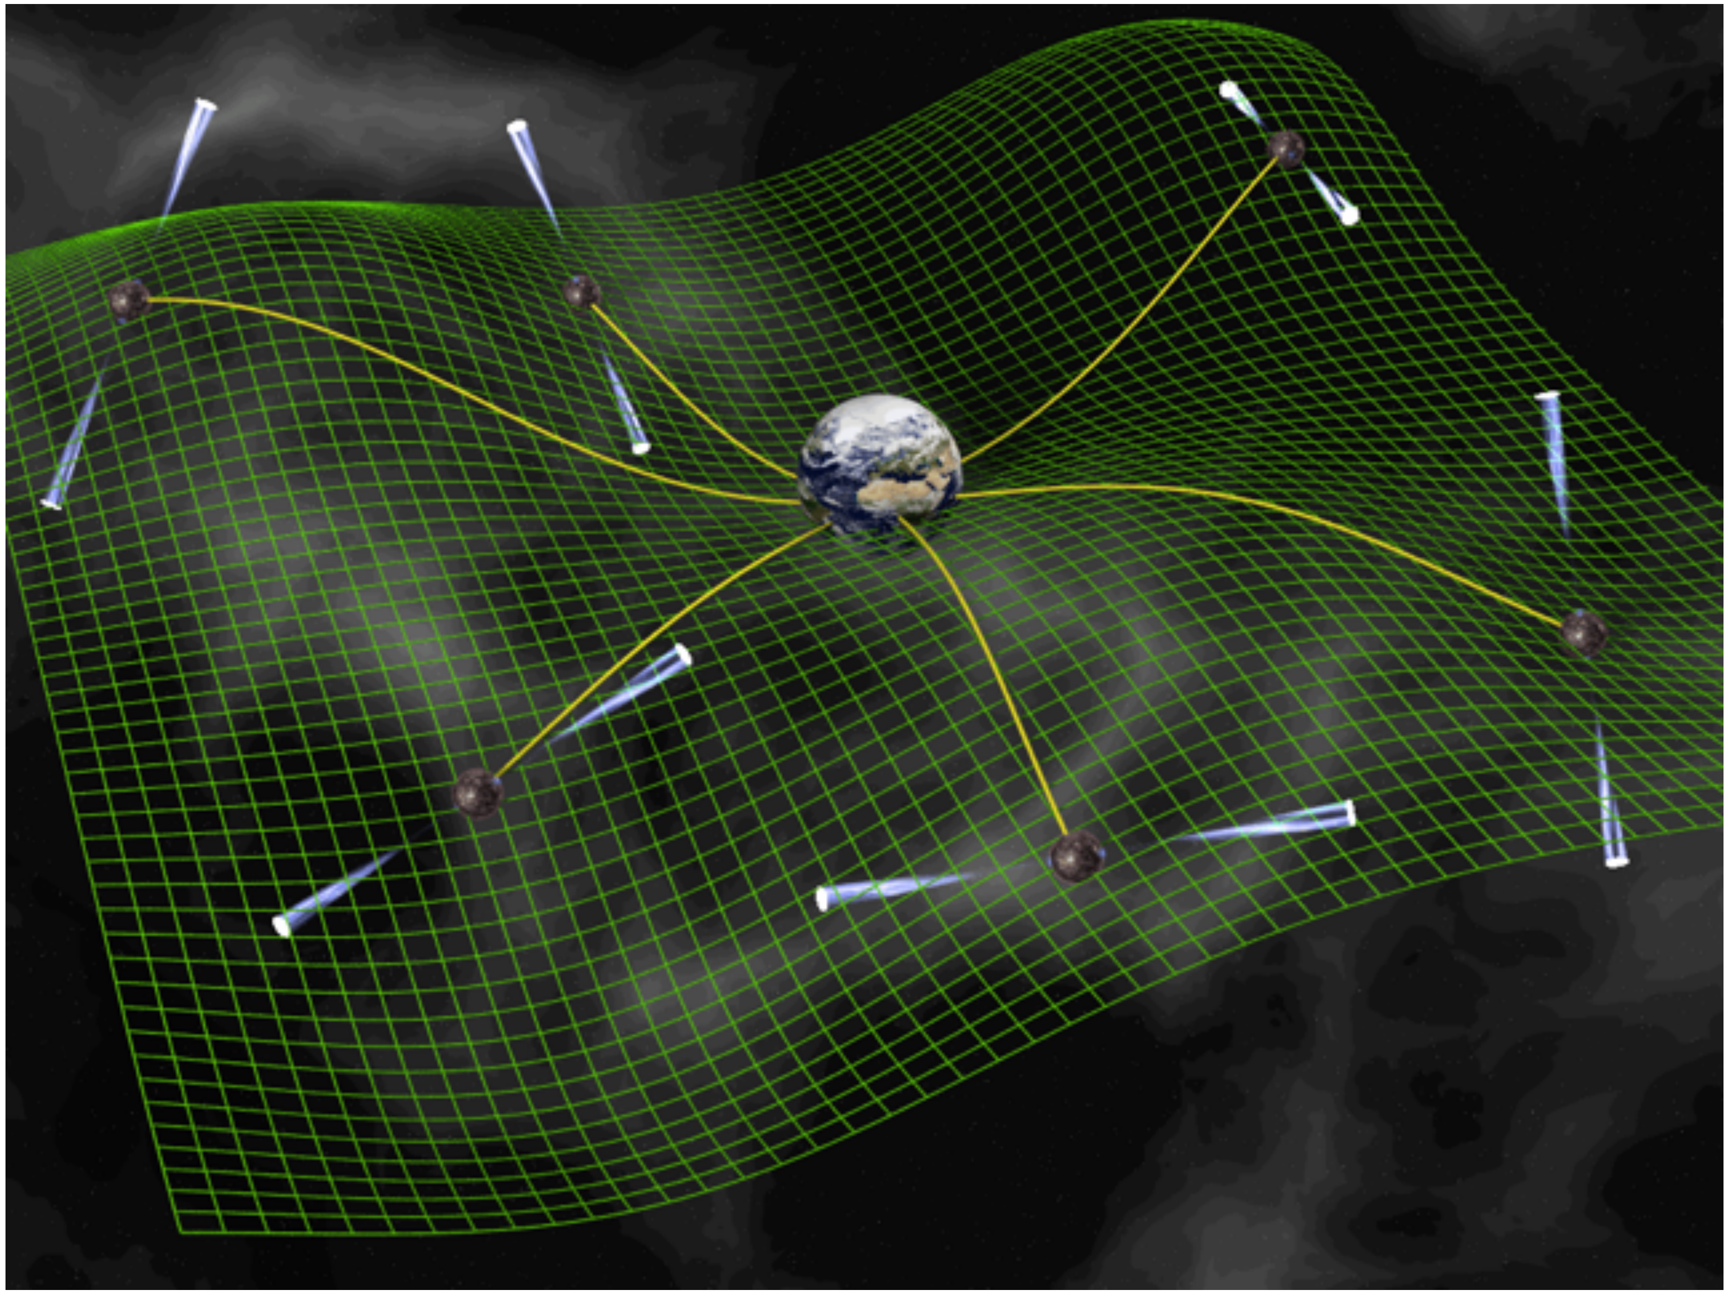

Cosmology via the 21-cm line

Neutral hydrogen (HI) has a long-lived emission line at λ_0=21cm

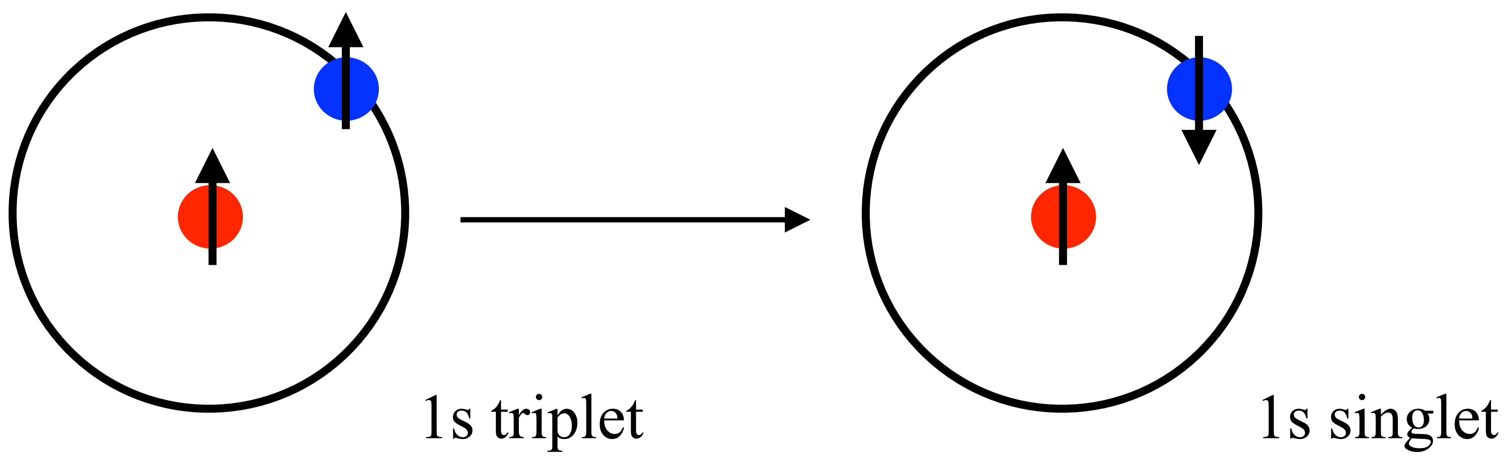

We observe the intensity of this emission as a function of sky angles θ,ϕ and wavelength $\lambda_{obs} = (21\text{ cm})(1+z)$.

The resulting 3D map traces cosmological structure.

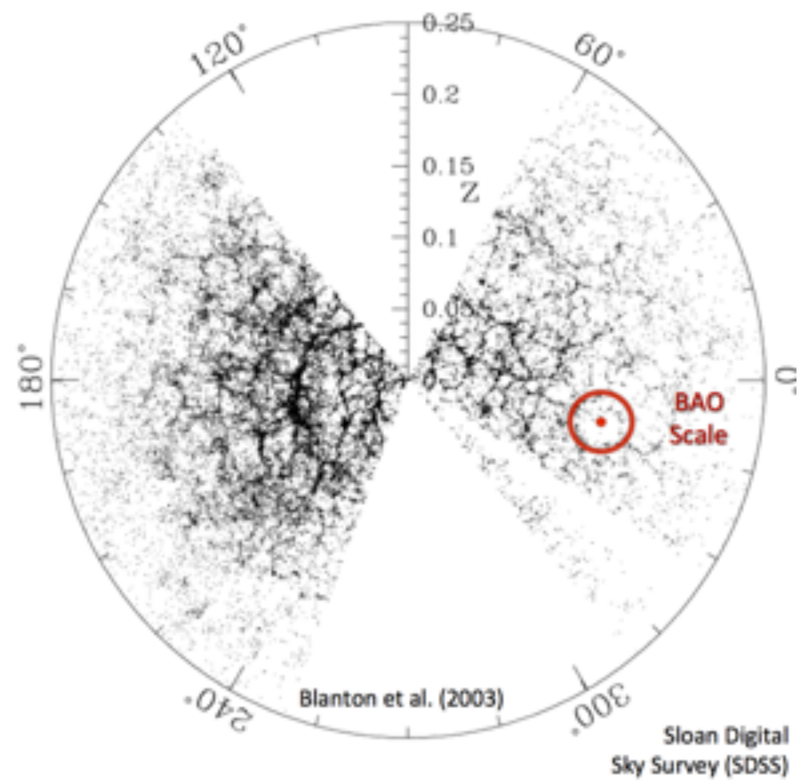

Spectroscopic galaxies: number density $n(\theta,\phi,z)$ traces large-scale structure.

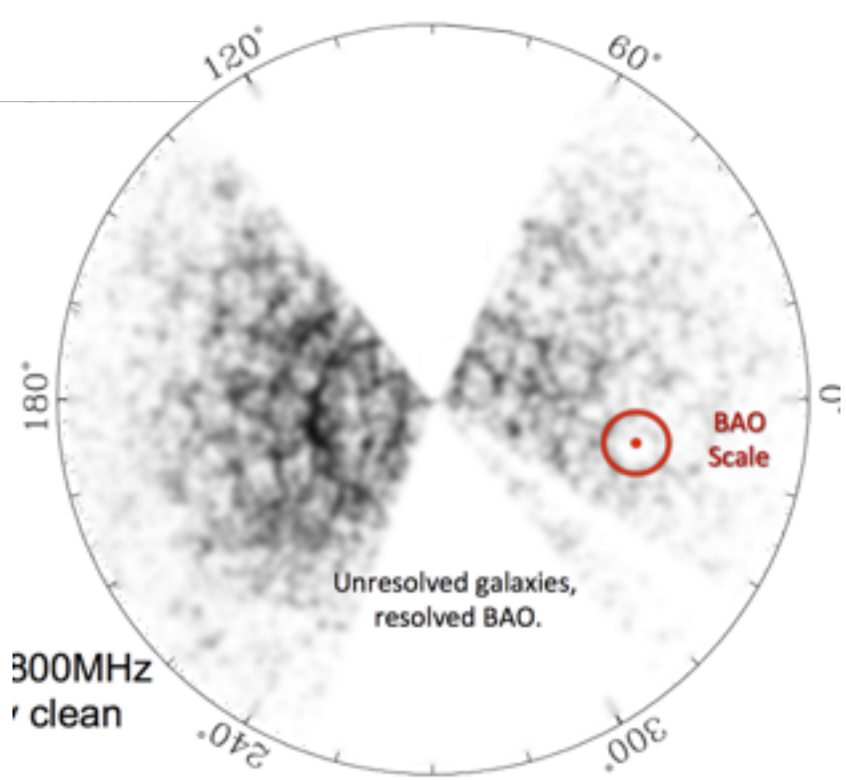

21-cm intensity mapping: brightness temperature $T(\theta,\phi,z)$ traces LSS.

CHIME:

redshift range $0.8 \leq z \leq 2.5$

radial resolution $\Delta z = 0.002$ (~5 Mpc)

angular resolution 0.3 deg (~20 Mpc)

Baryon acoustic oscillations

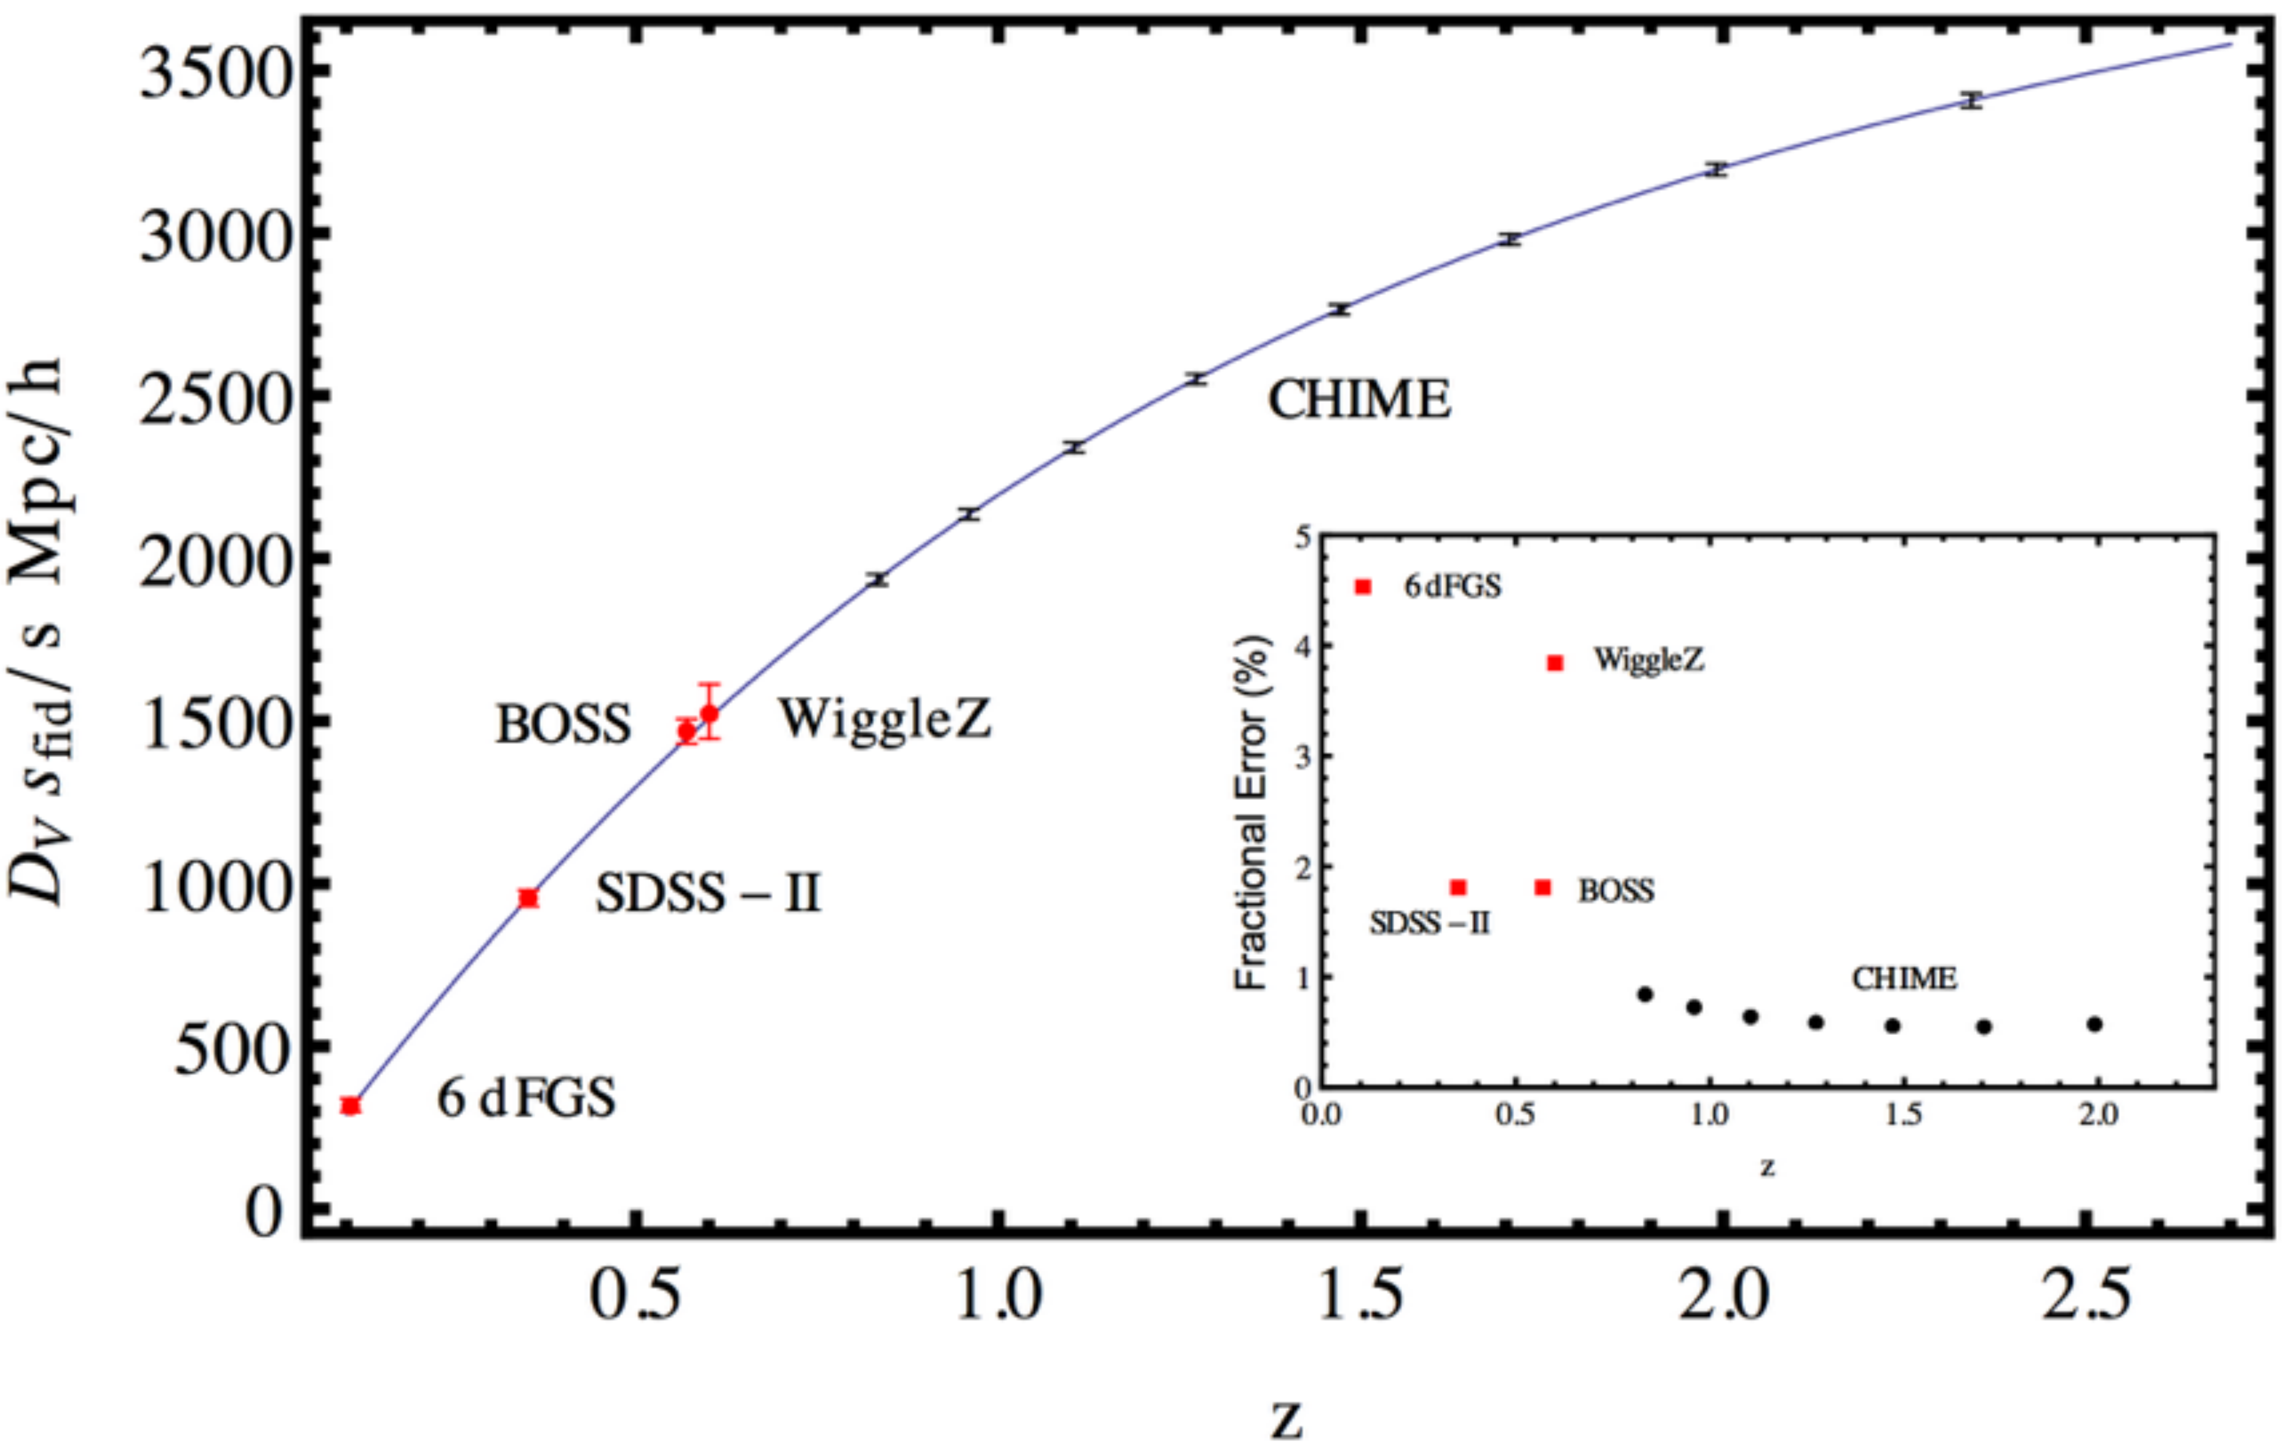

Why CHIME is so interesting

The CHIME design is a breakthrough if mapping speed per dollar is the metric.

Main advance on hardware side: using inexpensive commodity hardware in ways that haven't been done before.

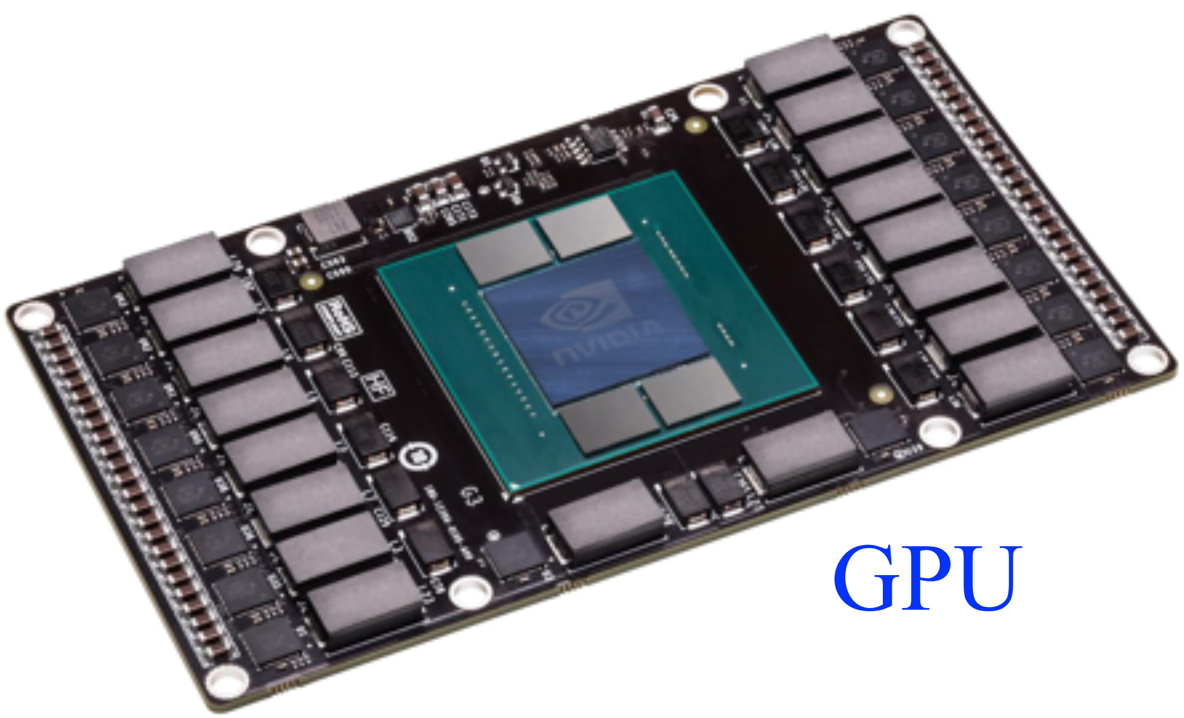

GPU

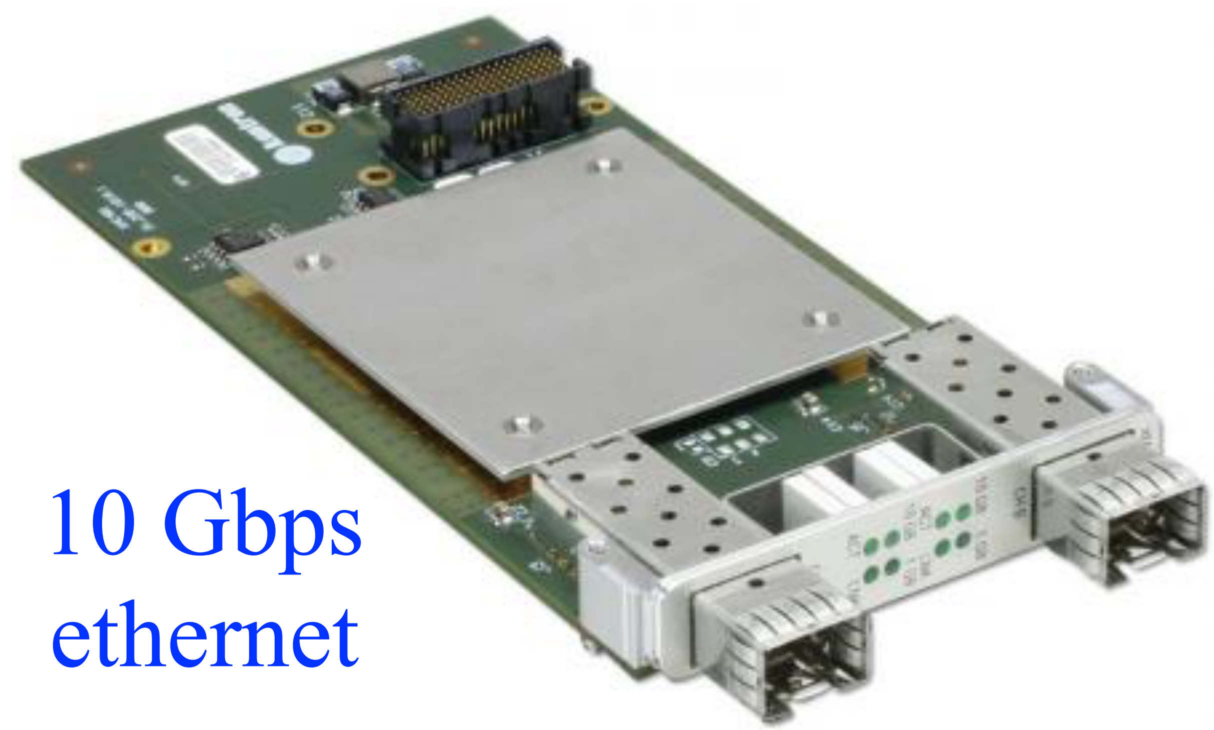

10 Gbps ethernet

Why CHIME is so interesting

Consider the cost of scaling up CHIME.
Suppose we increase collecting area A at fixed antenna density.

$$\text{Cost of computing} = \mathcal{O}(A^2 e^{-T/T_{\rm Moore}})$$
$$\text{Cost of everything else} = \mathcal{O}(A)$$
$$\text{Mapping speed } M = \mathcal{O}(A^2)$$

So total cost depends on mapping speed as

$$\text{Cost} = \begin{cases} \mathcal{O}(M e^{-T/T_{\rm Moore}}) & \text{if computation-dominated} \\ \mathcal{O}(M^{1/2}) & \text{otherwise} \end{cases}$$

A very shallow dependence!!

Why CHIME is so interesting

CHIME is potentially transformative in at least four areas:

- Fast radio bursts
- Searching for new pulsars
- Cosmology via the 21-cm hydrogen line
- Searching for gravity waves by timing known pulsars

The hardware is very inexpensive (< $10m) and also scales up very cheaply! Is radio astronomy the next big thing?

My perspective: it all depends on how much we can improve the algorithms.


Thanks!