Beating Nyquist by randomized sampling

Felix J. Herrmann*
fherrmann@eos.ubc.ca

Joint work with Gang Tang, Reza Shahidi, Gilles Hennenfent, and Tim Lin

***Seismic Laboratory for Imaging & Modeling**
Department of Earth & Ocean Sciences
The University of British Columbia

slim.eos.ubc.ca

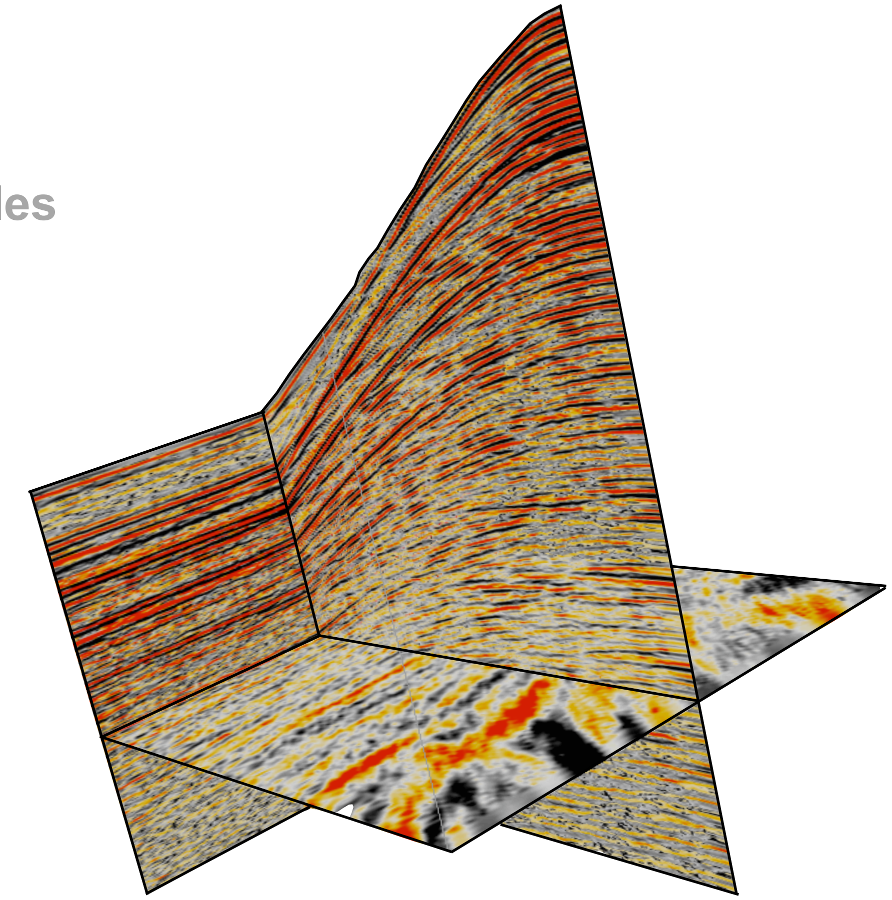

EAGE Workshop: Reconstruction, Recovery and Interpolation of Multi-dimensional Seismic Wave Fields
Amsterdam, June 8th, 2009

Relation to existing work

- filter-based methods [Spitz’91, Fomel’00]
  - convolve the incomplete data with a data-adaptive interpolating filter
- wavefield-operator-based methods [Canning and Gardner’96, Biondi et al.’98, Stolt’02]
  - explicitly include wave propagation
  - require knowledge of velocity model
  - computationally intensive
- transform-based methods [Sacchi et al.’98, Trad et al.’03, Zwartjes and Sacchi’07]
  - non-adaptive and fast
  - no explicit link with wave propagation
  - related to recent developments in Compressive Sensing (CS)

Motivation

- **Seismic data processing, modeling & inversion:**
  - firmly rooted in Nyquist's sampling paradigm for high-dimensional wavefields
  - too *pessimistic* for signals with *structure,* i.e, there exists some sparsifying transform (e.g. Fourier, curvelets)

- **Recent theoretical & hardware developments**
  - Alternative multiscale, localized & directional transform domains that compress seismic data
  - New nonlinear sampling theory that supersedes the overly pessimistic Nyquist sampling criterion
  - New autonomous data acquisition devices that allow for more flexibility during acquisition
  - New simultaneous & continuous recording

- **Extensions to higher-D through blue-noise sampling**

Motivation cont'd

- **Solution strategy:**
  - *leverage new* paradigm of *compressive sensing* (CS)
    - identify wavefield reconstruction from missing sources & receivers or from simultaneous acquisition as instances of CS
    - reduce acquisition, simulation, and inversion costs by ***randomization*** and deliberate ***subsampling***
  - recovery from sample **rates ≈ computational cost** *proportional* to **transform-domain sparsity** of *data* or *model*

- **Remove the "curse of dimensionality" by removing constructive aliases/interferences**
  - breaking the ***periodicity of regular sampling***
  - using ***incoherent sources***

- **Turn problem into a "simple" denoising problem ....**
  - use blue-noise sampling techniques from computer graphics community

Problem statement

Consider the following (severely) underdetermined system of linear equations

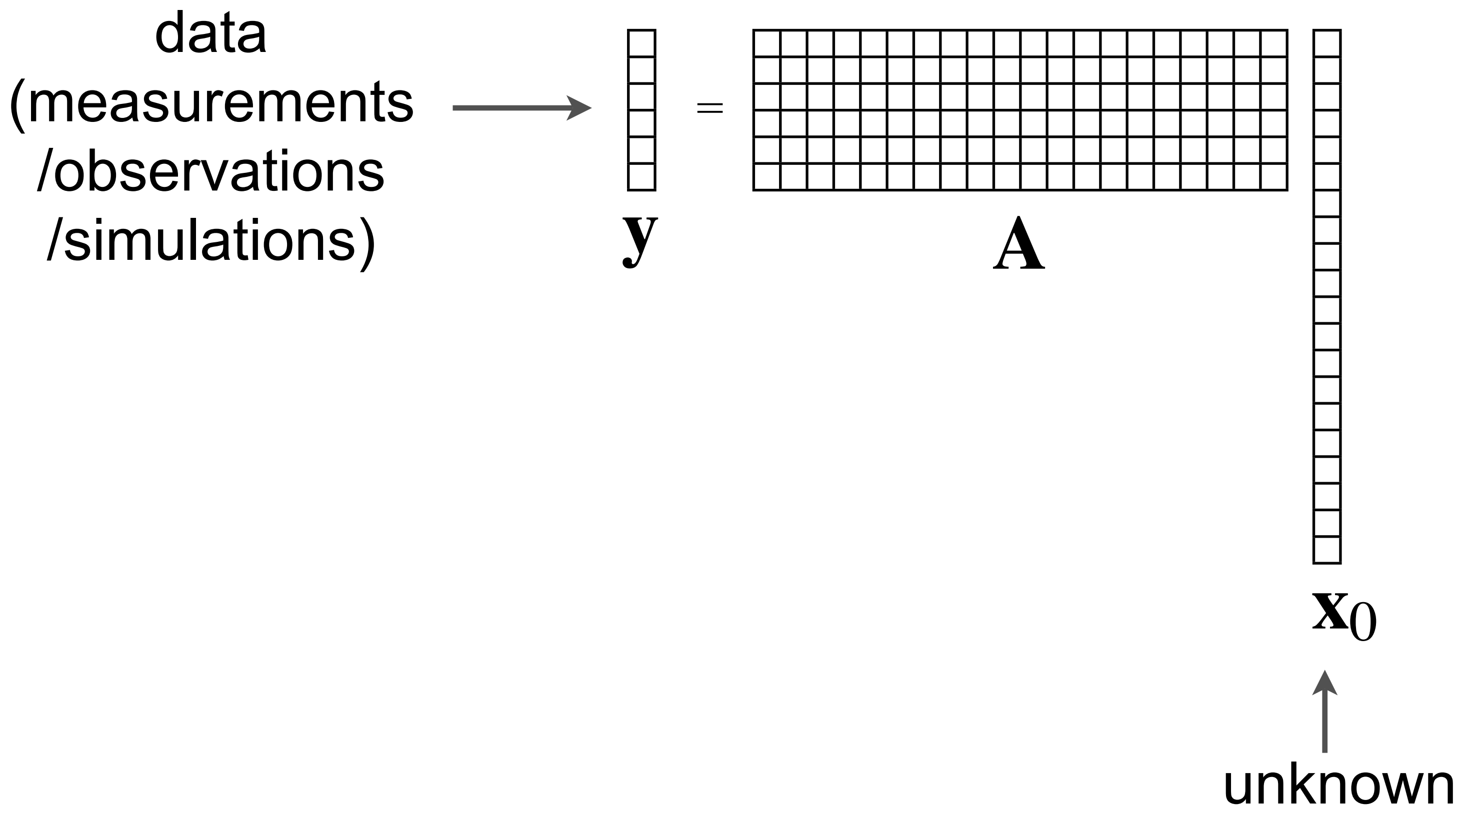

Is it possible to recover $\mathbf{x}_0$ accurately from $\mathbf{y}$?

Perfect recovery

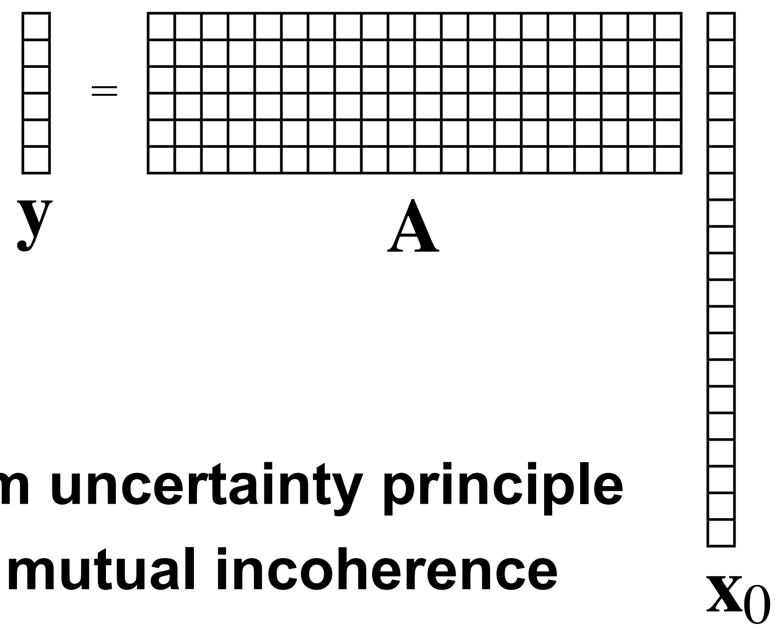

- conditions
  - **A** obeys the **uniform uncertainty principle**
  - ***randomized* A <=> mutual incoherence**
  - $\mathbf{x}_0$ is **sufficiently sparse**
- ***nonlinear*** recovery procedure:

$$\underbrace{\min_{\mathbf{x}} \|\mathbf{x}\|_1}_{\text{sparsity}} \quad \text{s.t.} \quad \underbrace{\mathbf{A}\mathbf{x} = \mathbf{y}}_{\text{perfect reconstruction}}$$

- performance
  - *S*-sparse vectors recovered from roughly on the order of *S* measurements (to within constant and *log* factors)

[Candès et al.'06]
[Donoho'06]

Simple example

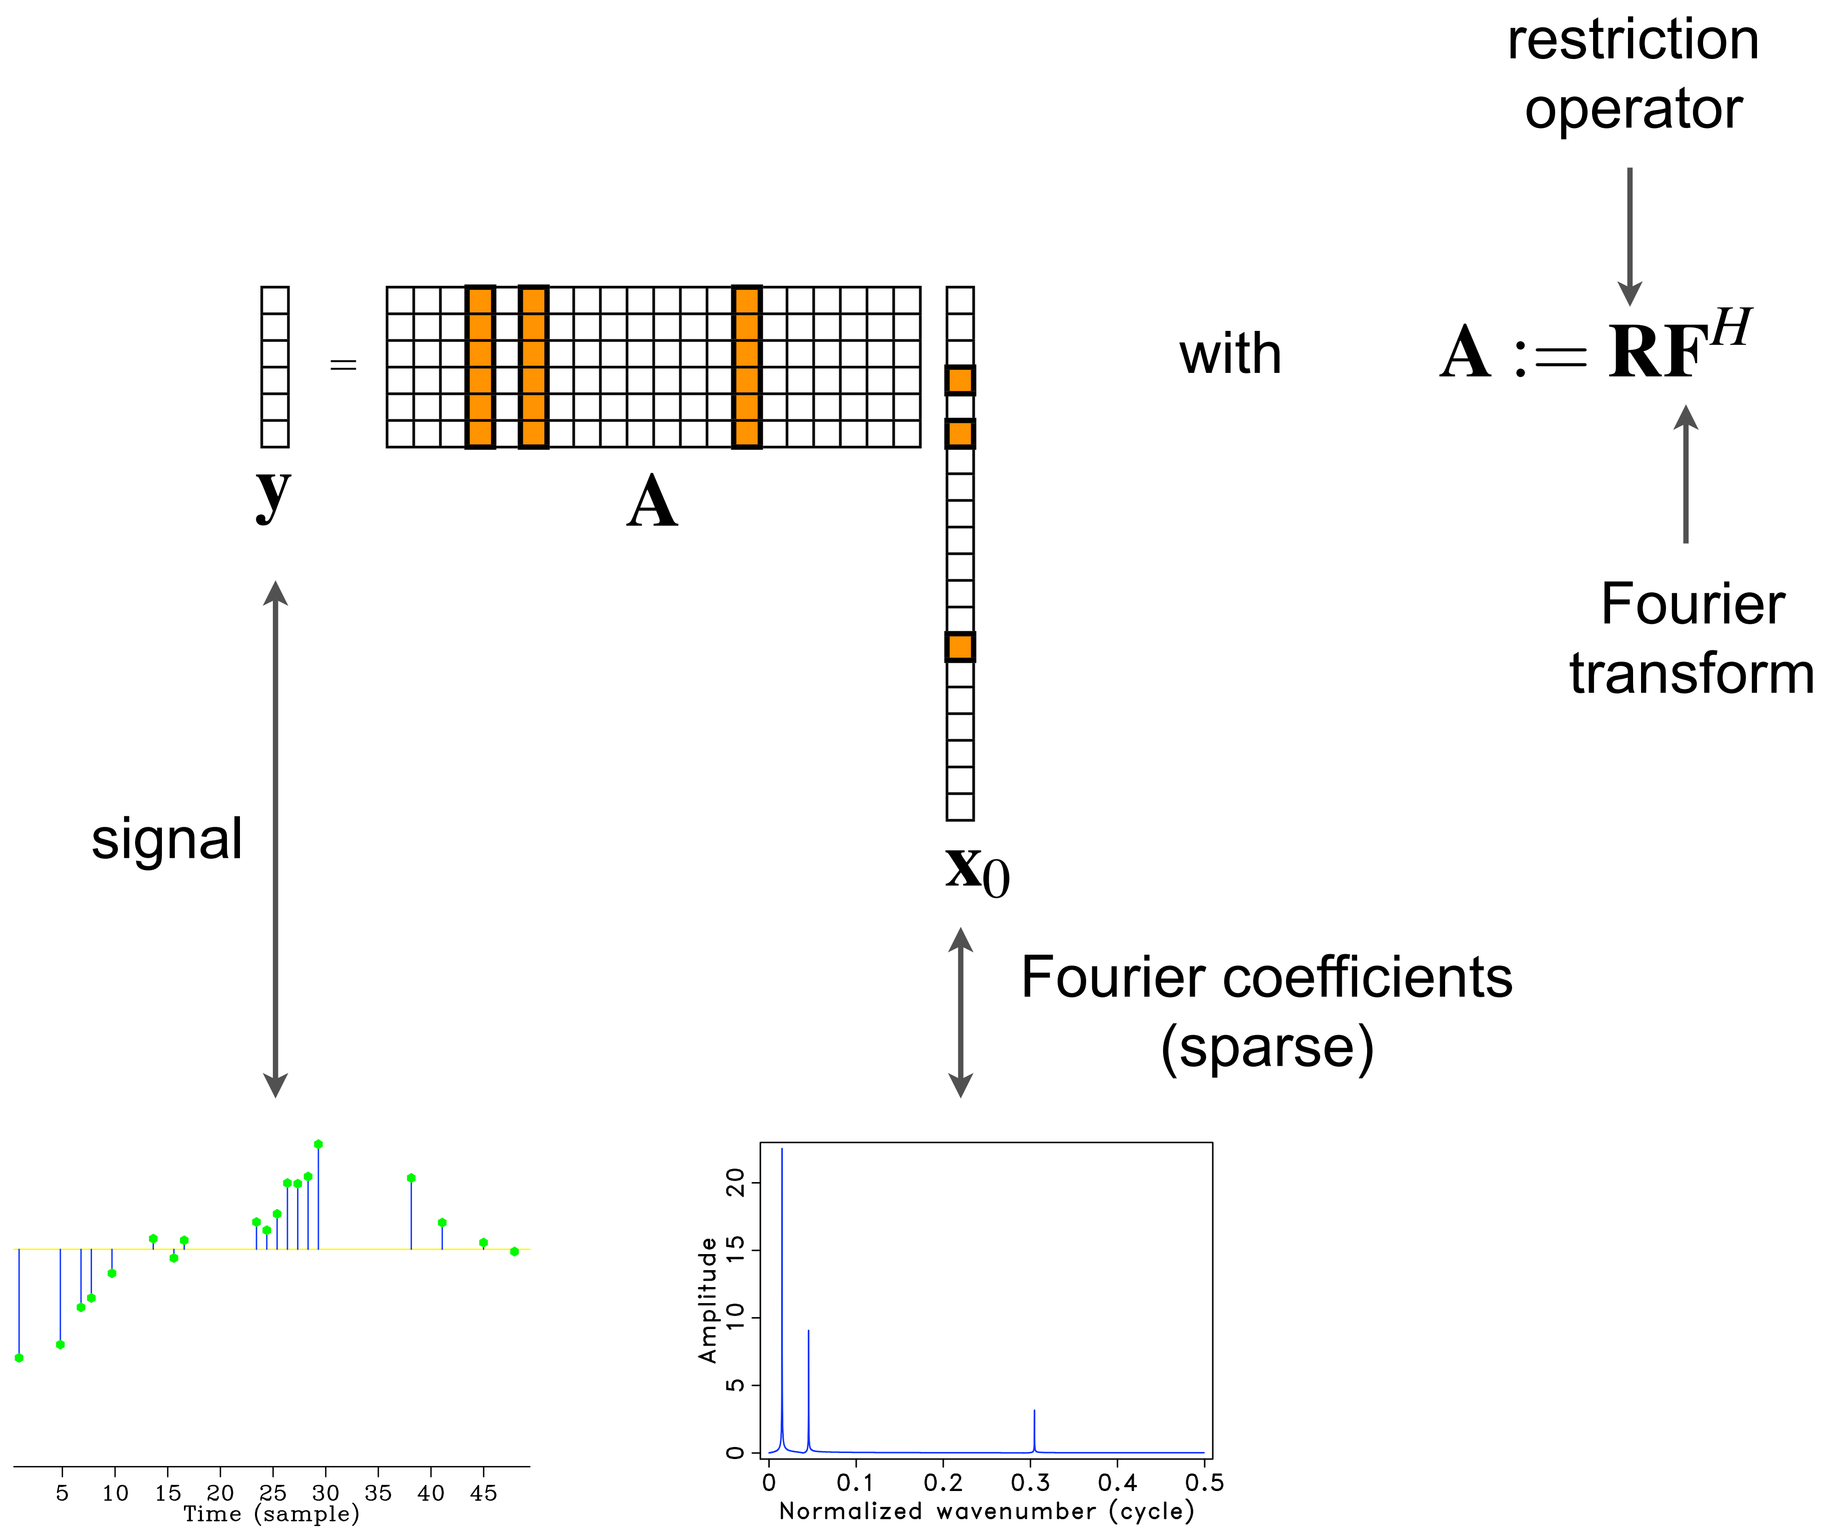

NAIVE sparsity-promoting recovery

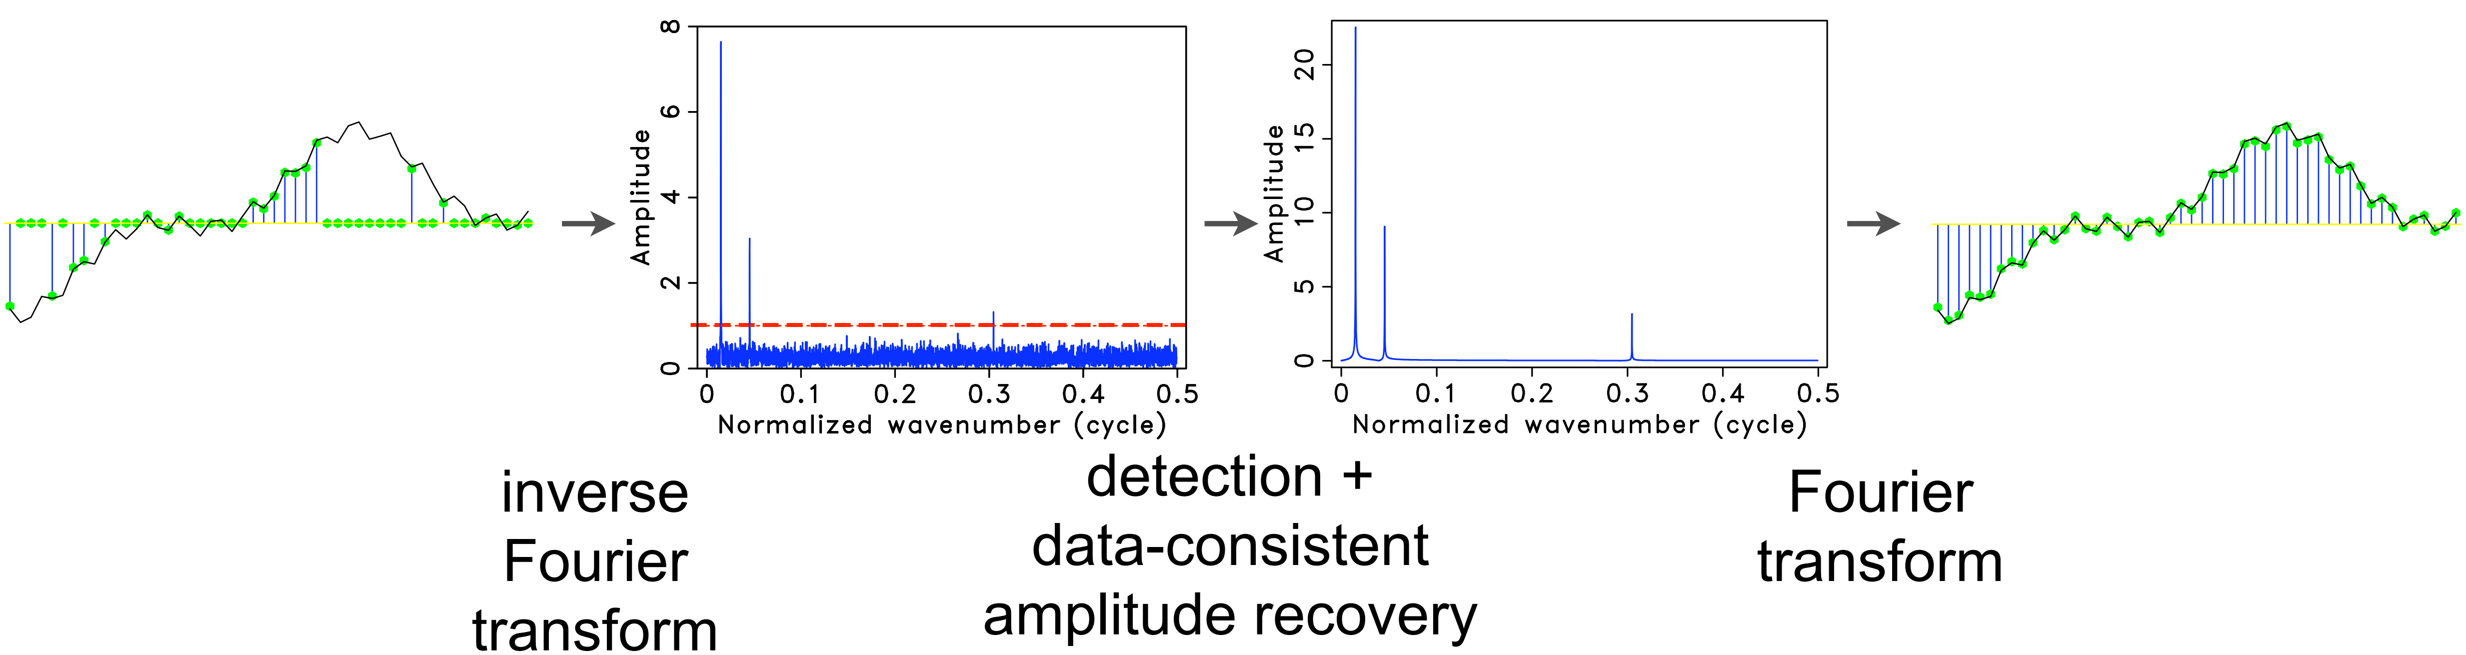

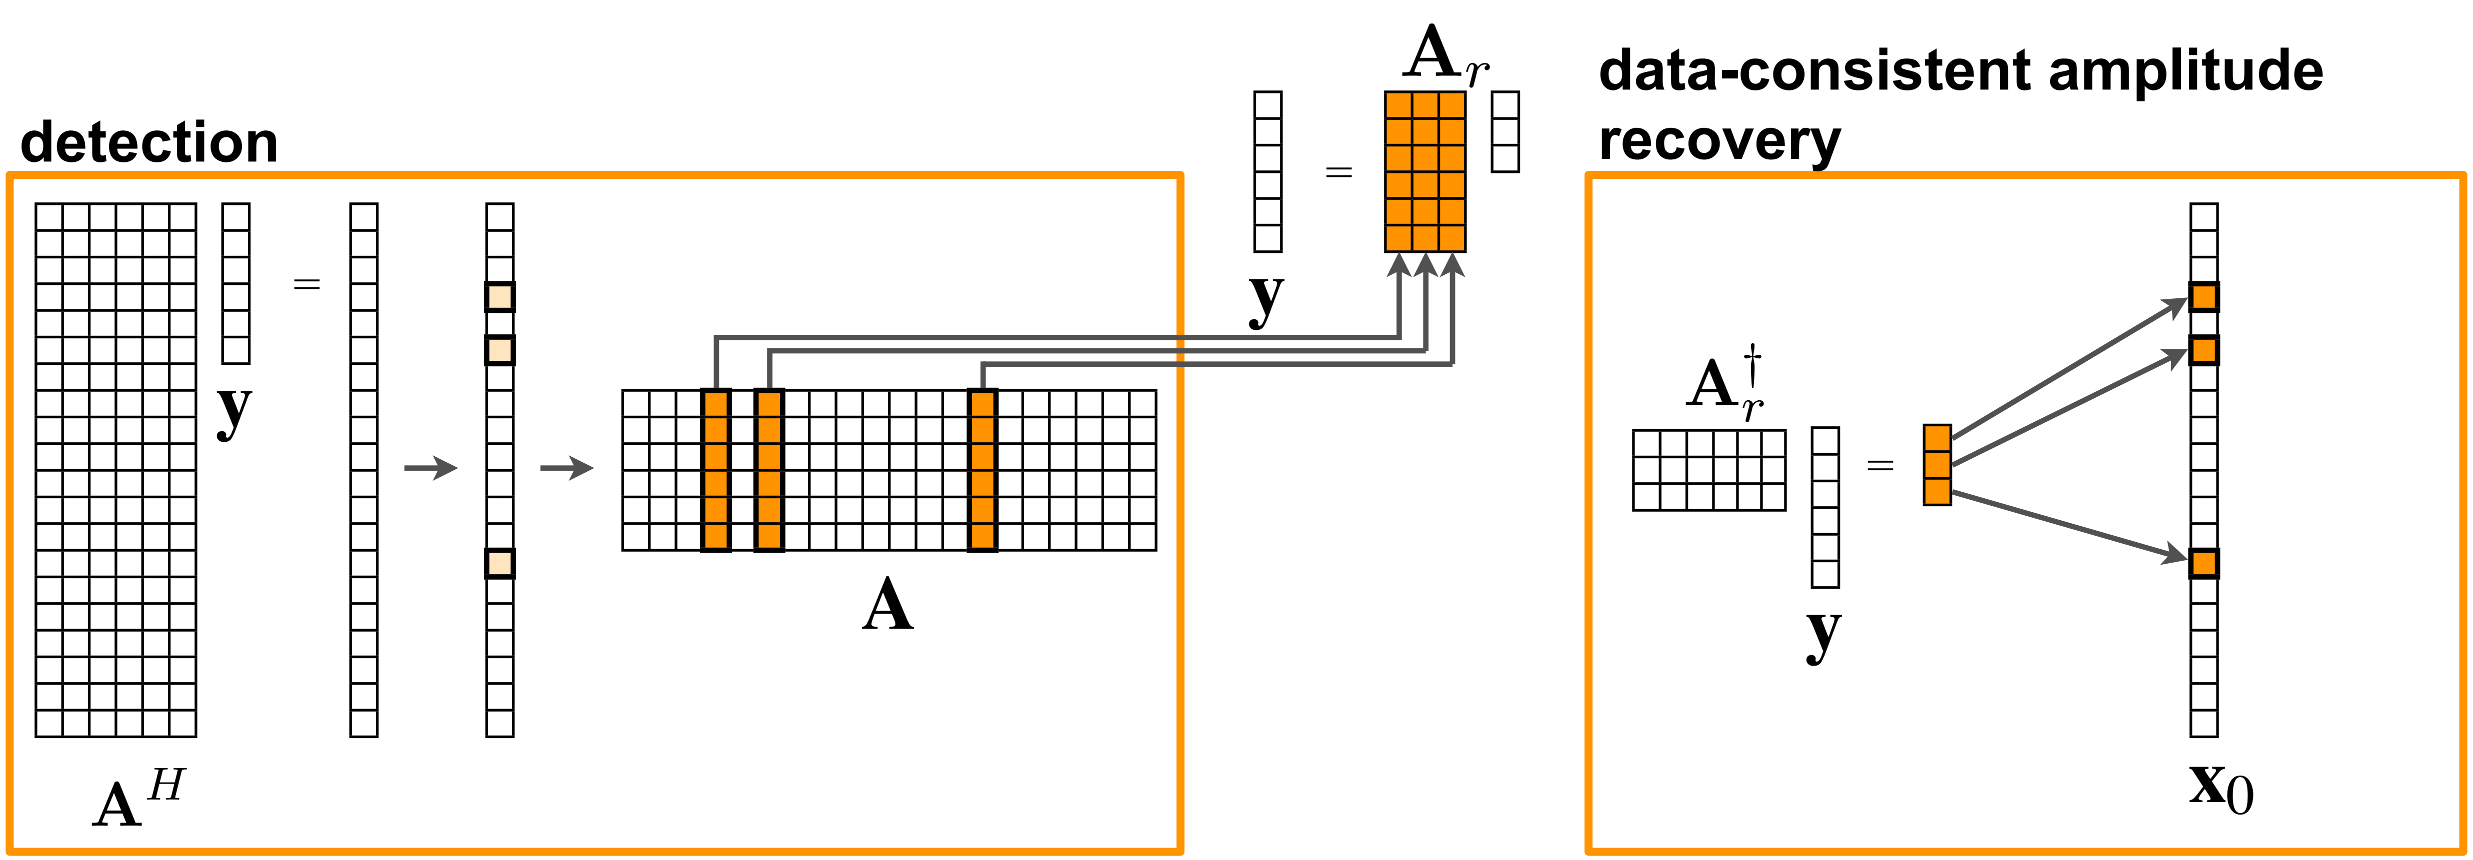

Undersampling “noise”

- “noise”
  - due to $\mathbf{A}^H\mathbf{A} \neq \mathbf{I}$
  - defined by $\mathbf{A}^H\mathbf{A}\mathbf{x}_0-\alpha\mathbf{x}_0 = \mathbf{A}^H\mathbf{y}-\alpha\mathbf{x}_0$

1 out of 2 | 1 out of 4 | 1 out of 6 | 1 out of 8

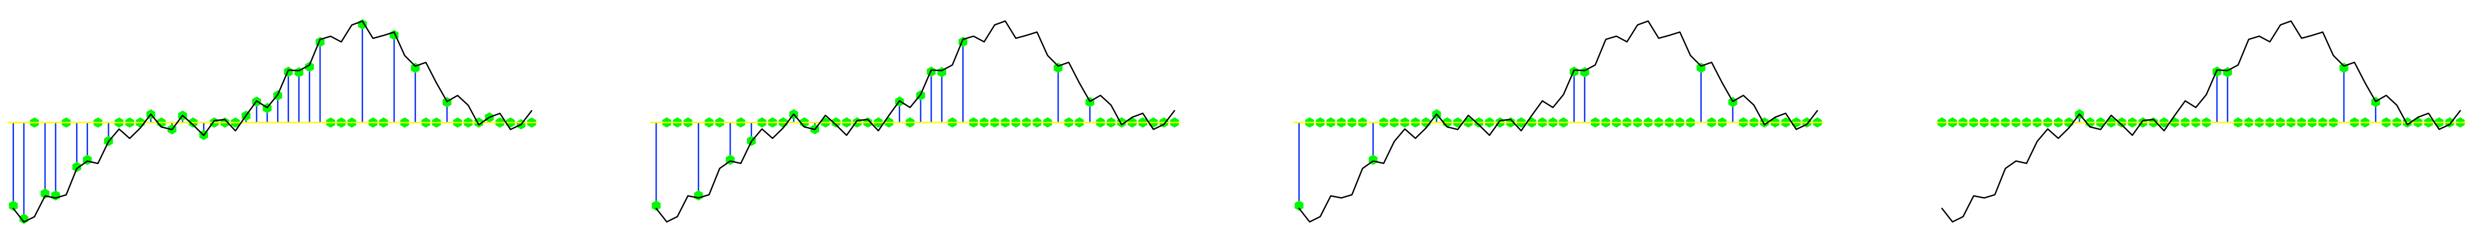

less acquired data →

3 detectable Fourier modes

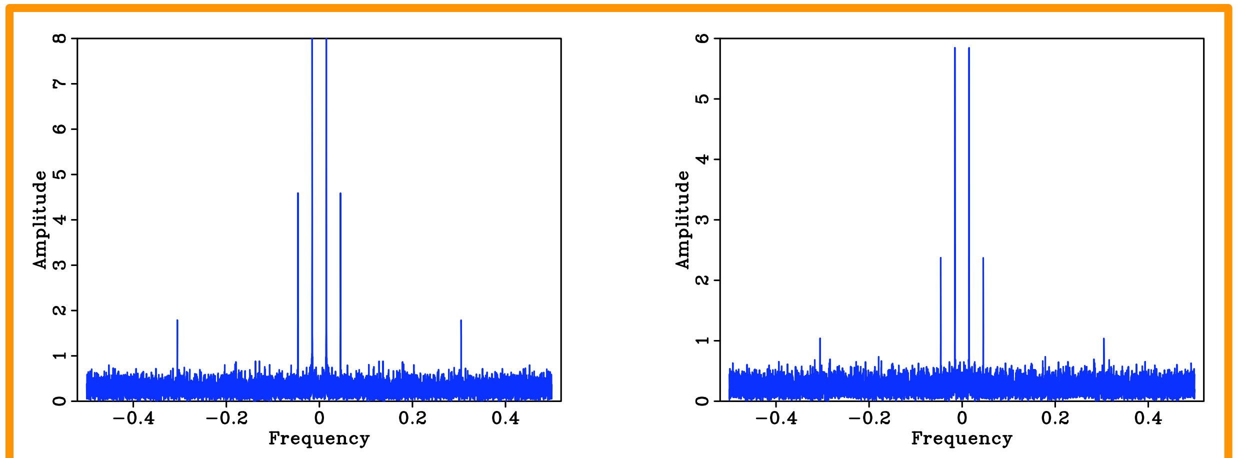

2 detectable Fourier modes

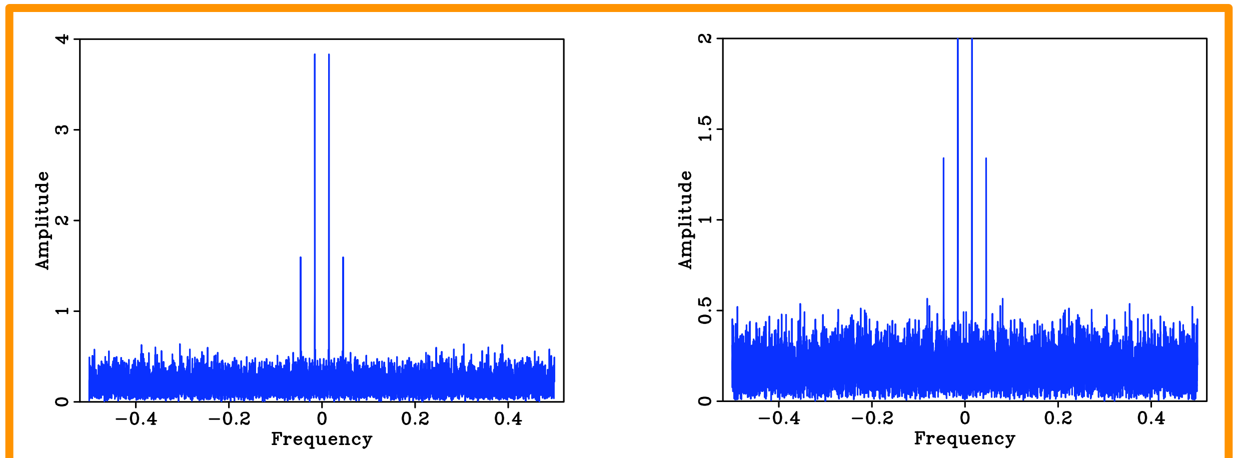

Wavefield sampling and nonlinear recovery

- *sparsifying transform*
  - typically **localized** in the time-space domain to handle the complexity of seismic data
    - curvelet transform (dyadic-parabolic partition of the f-k domain)
    - [windowed Fourier transform]
- *sampling scheme*
  - generates incoherent random undersampling “noise” in the sparsifying domain
  - do not create large gaps
    - because of the limited spatiotemporal extend of transform elements used for the reconstruction
- *sparsity-promoting solver*
  - requires few matrix-vector multiplications

2D discrete curvelets

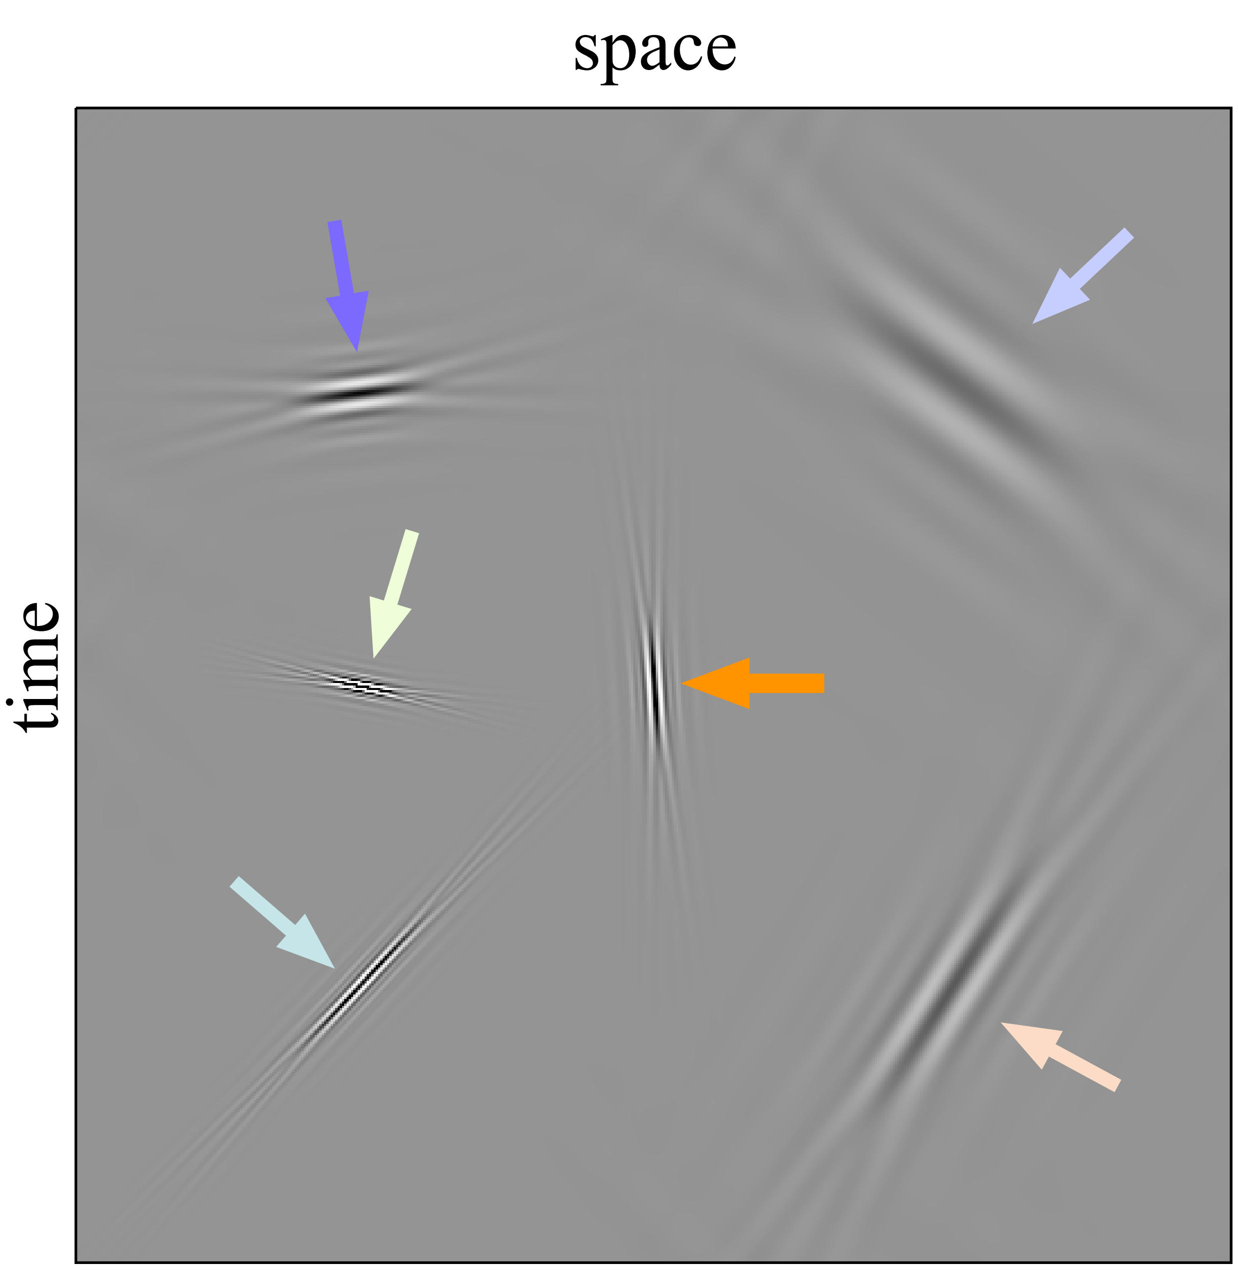

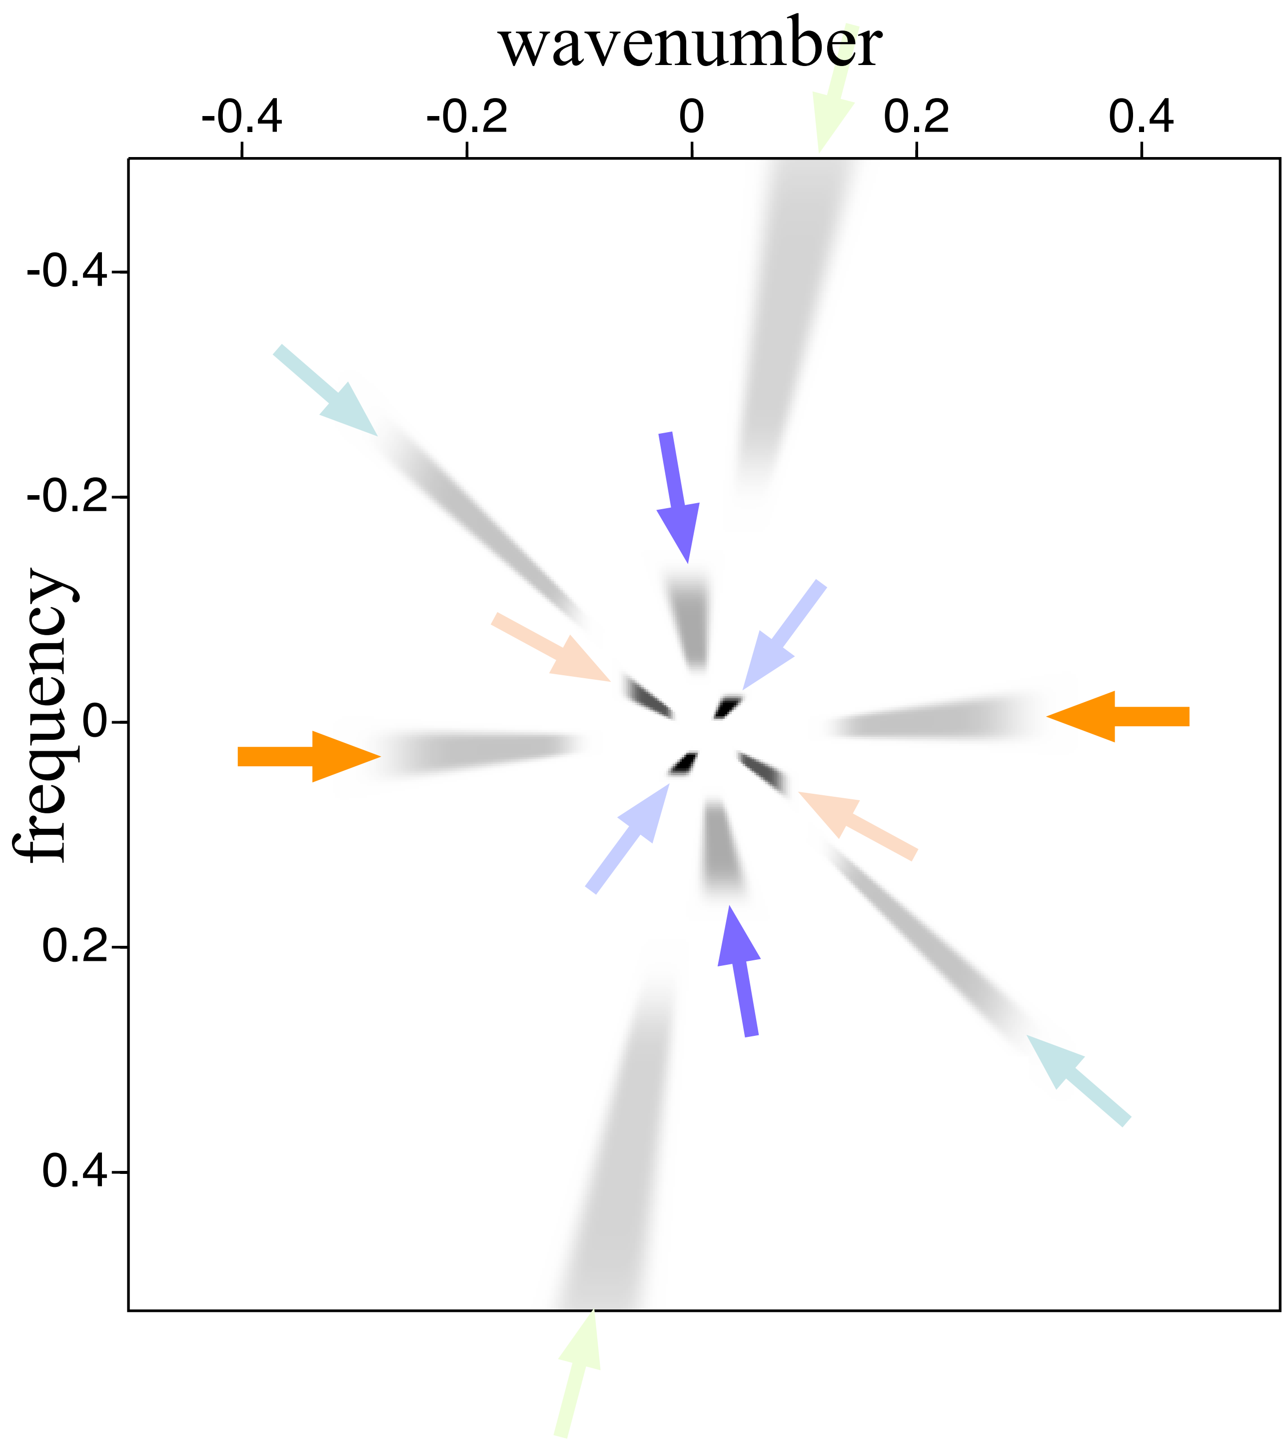

Fourier reconstruction

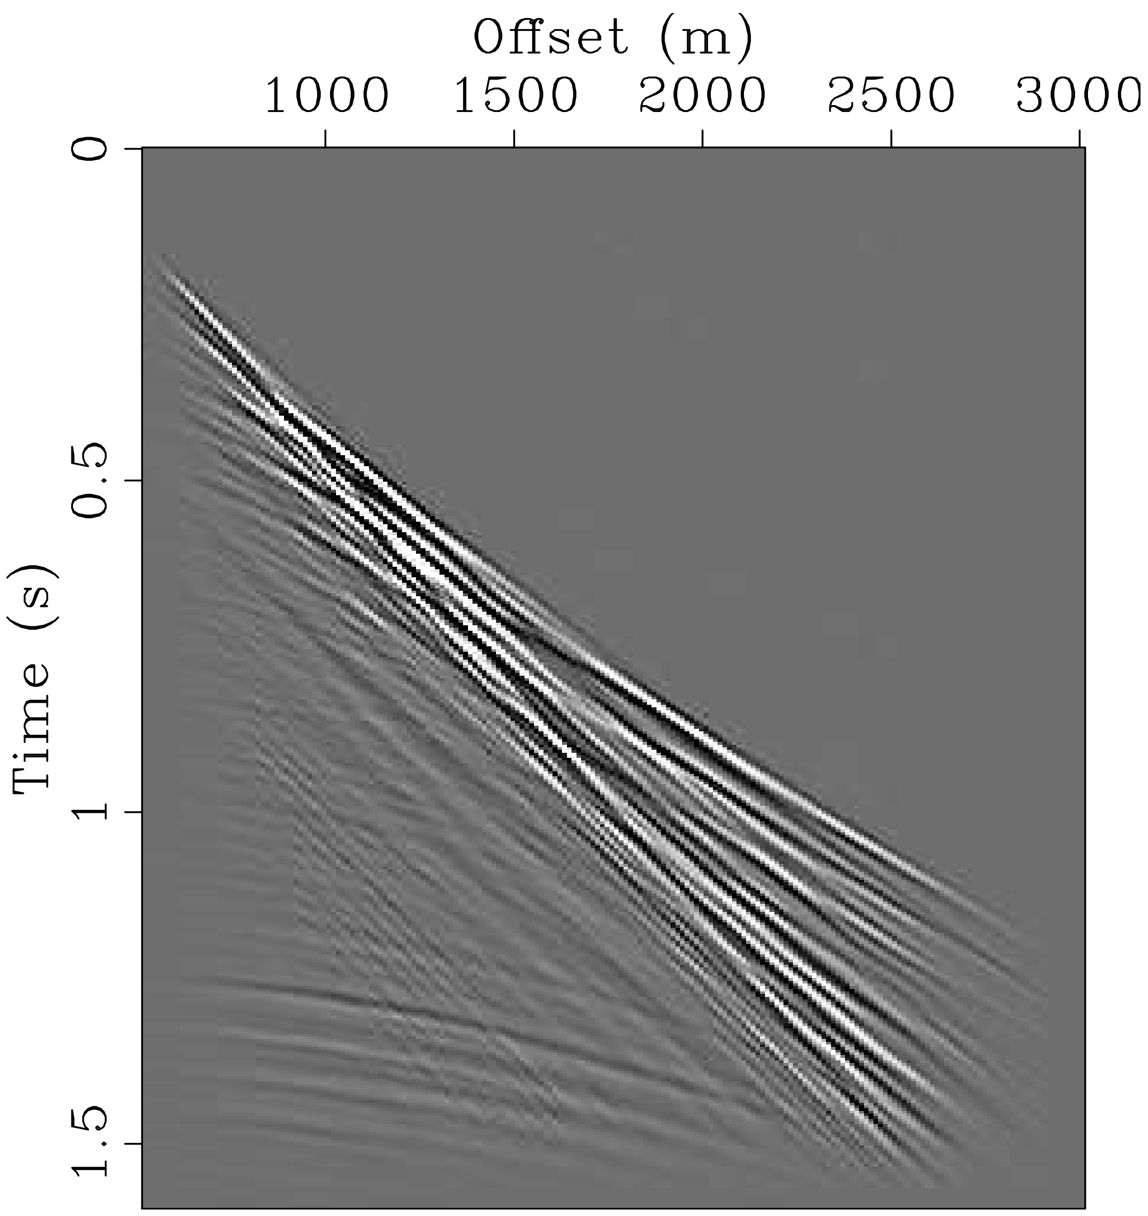

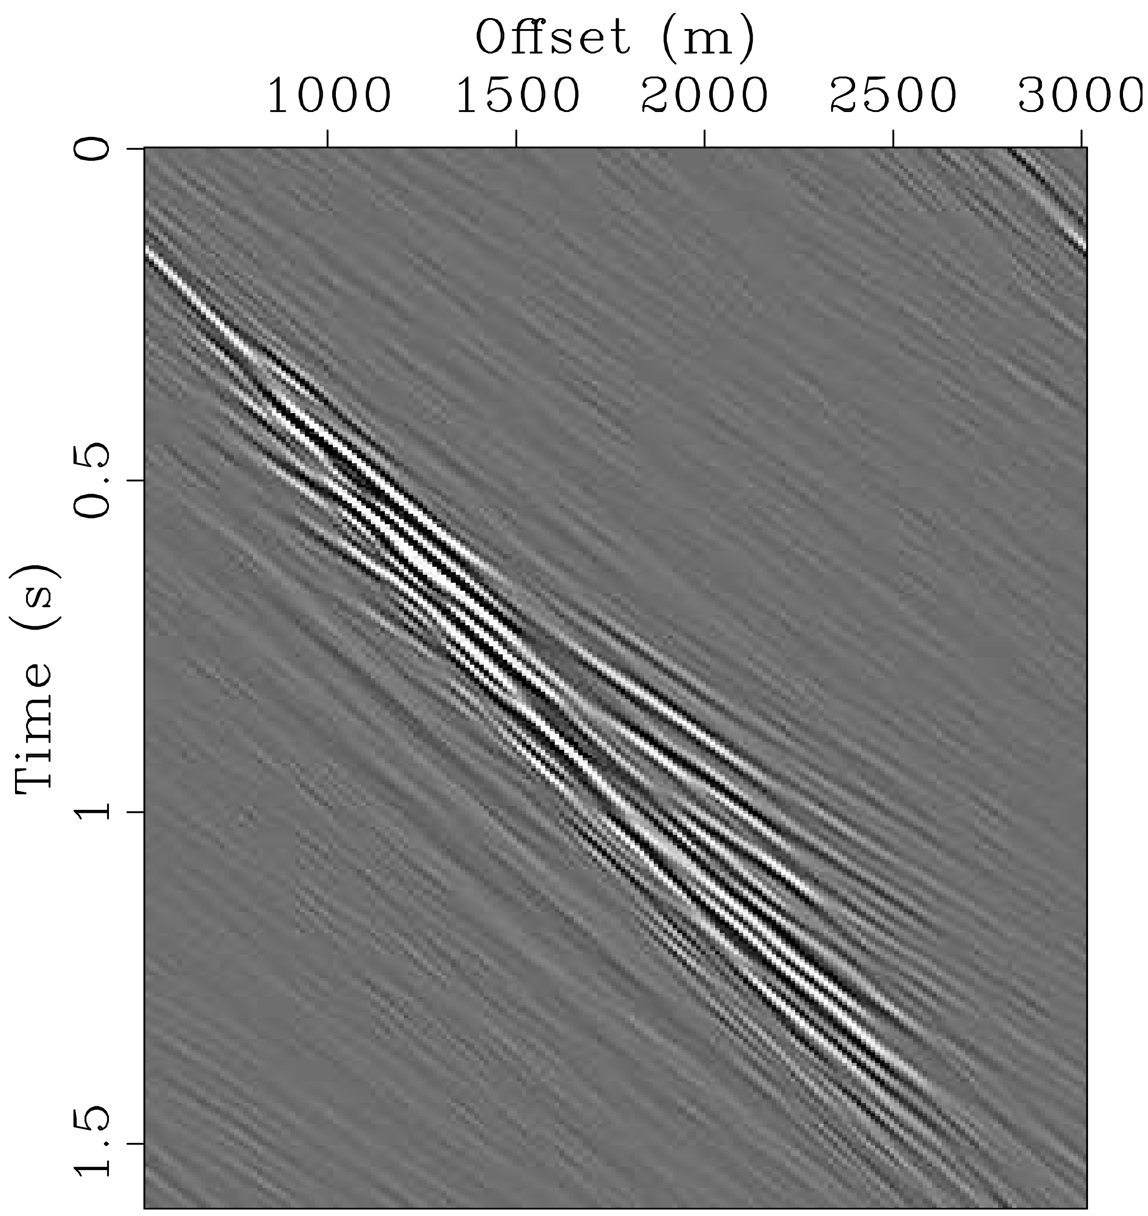

1 % of coefficients

Wavelet reconstruction

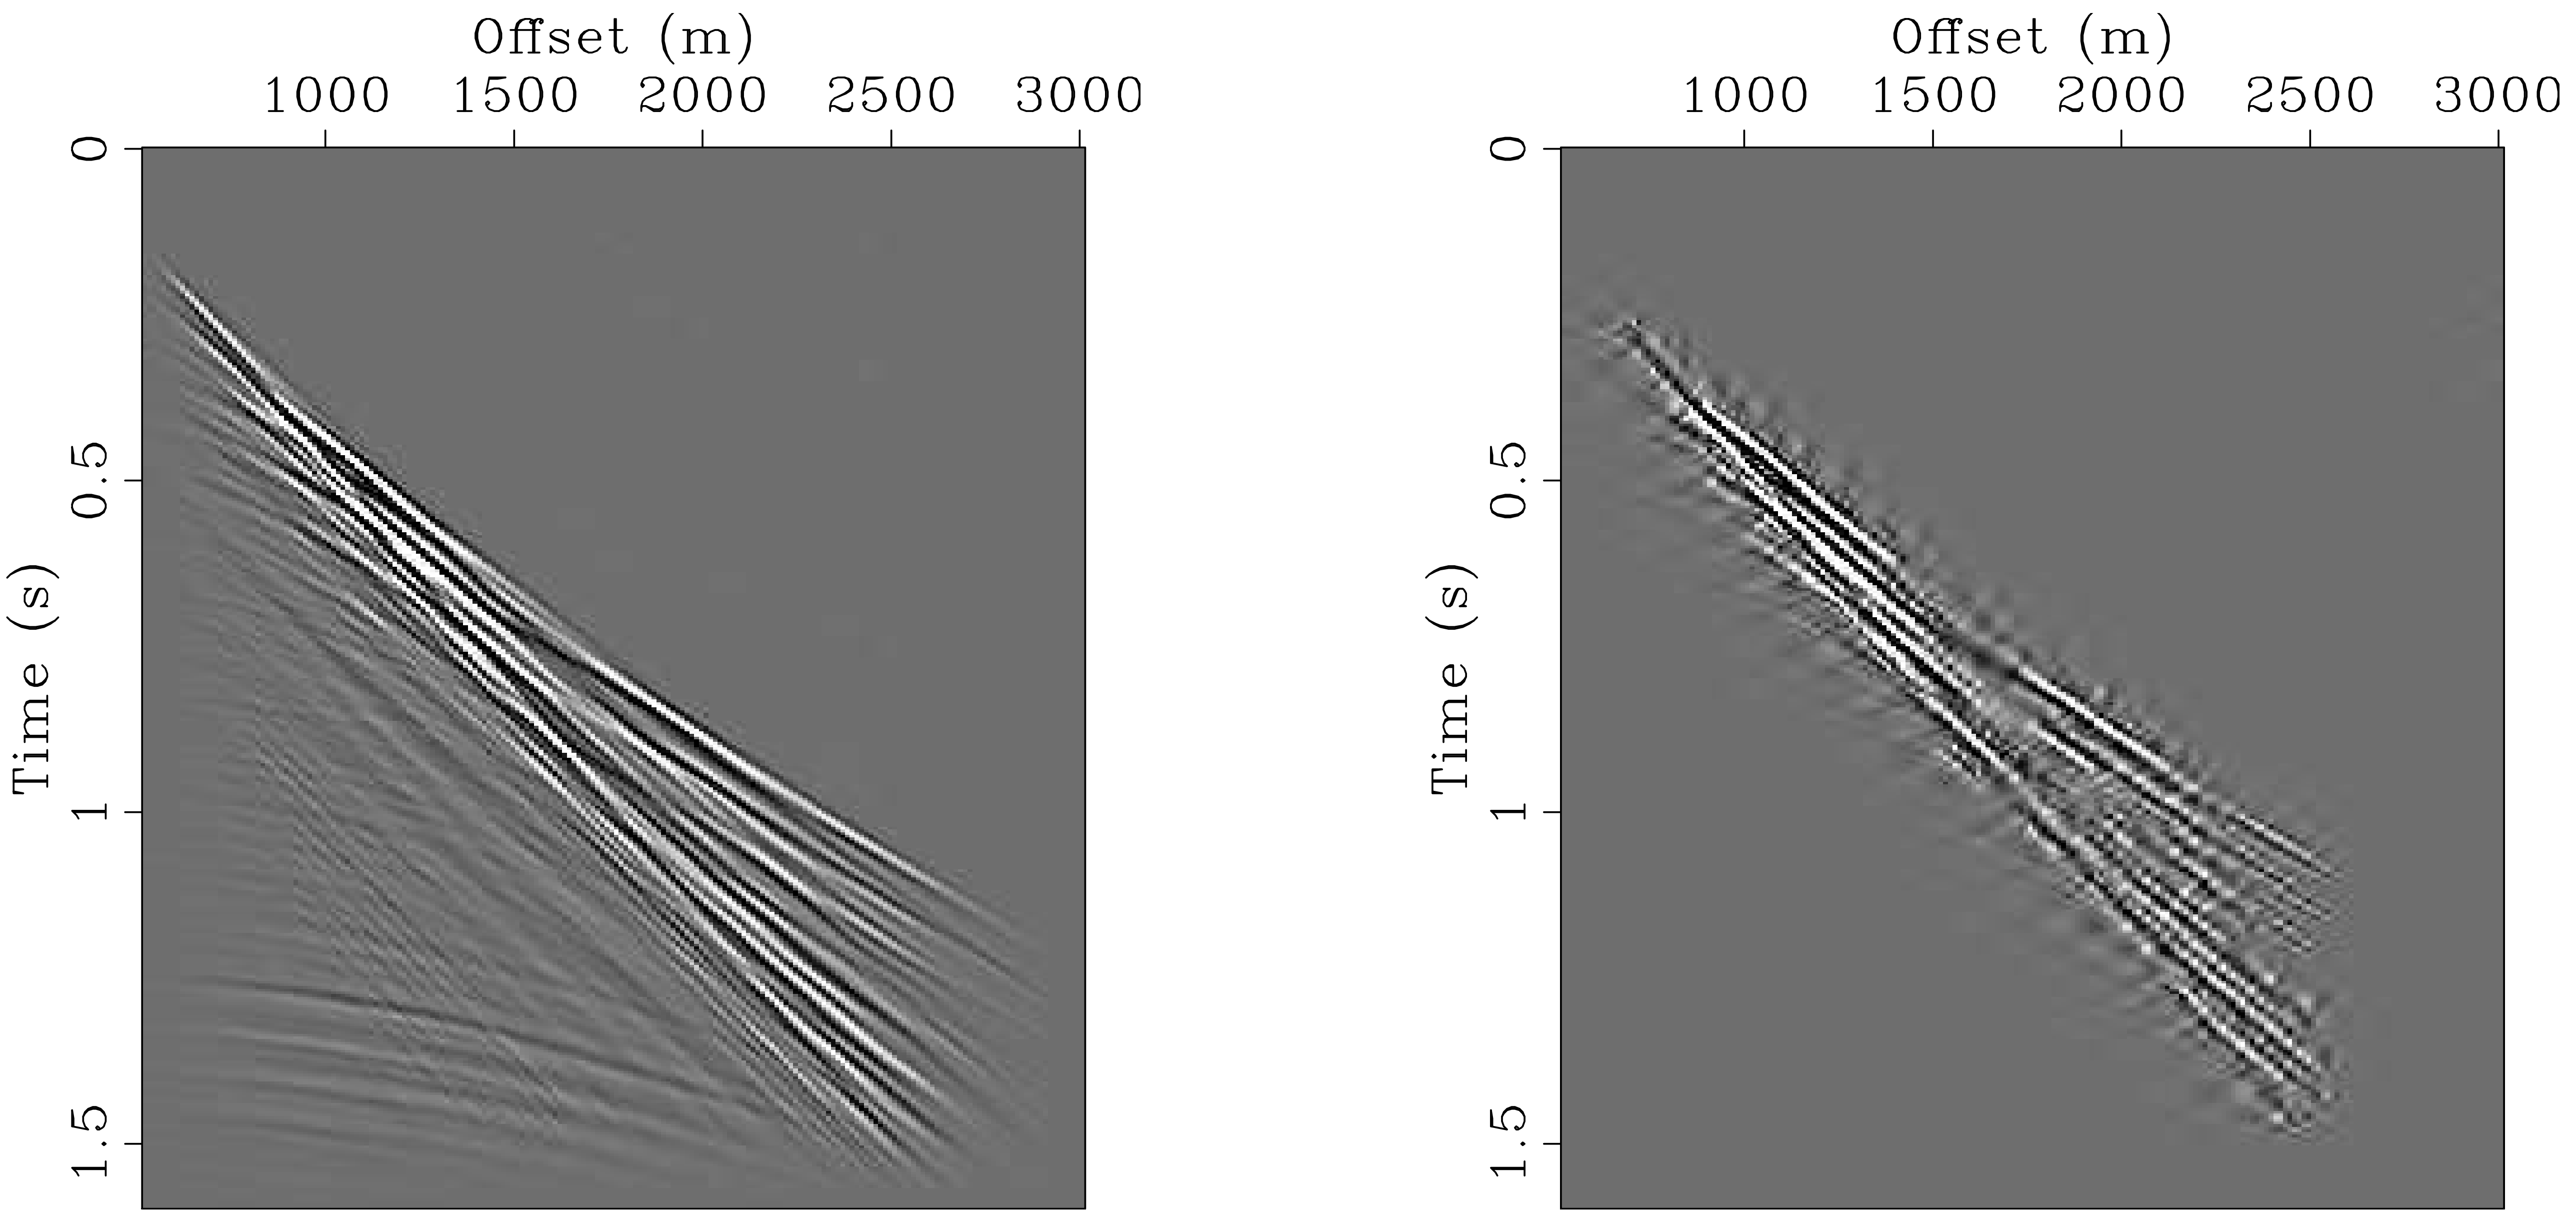

1 % of coefficients

Curvelet reconstruction

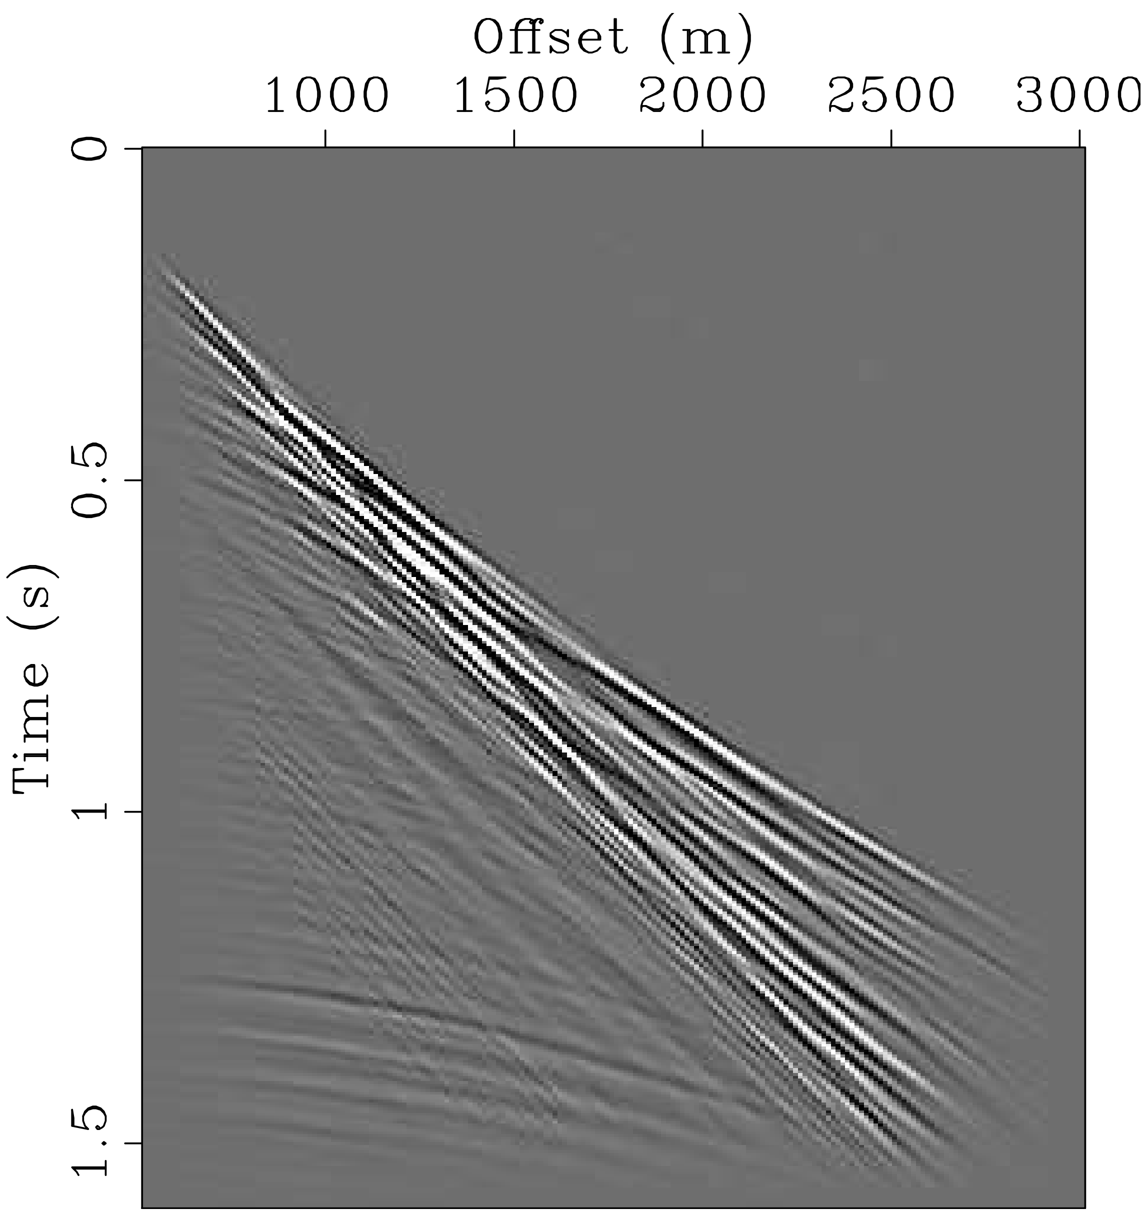

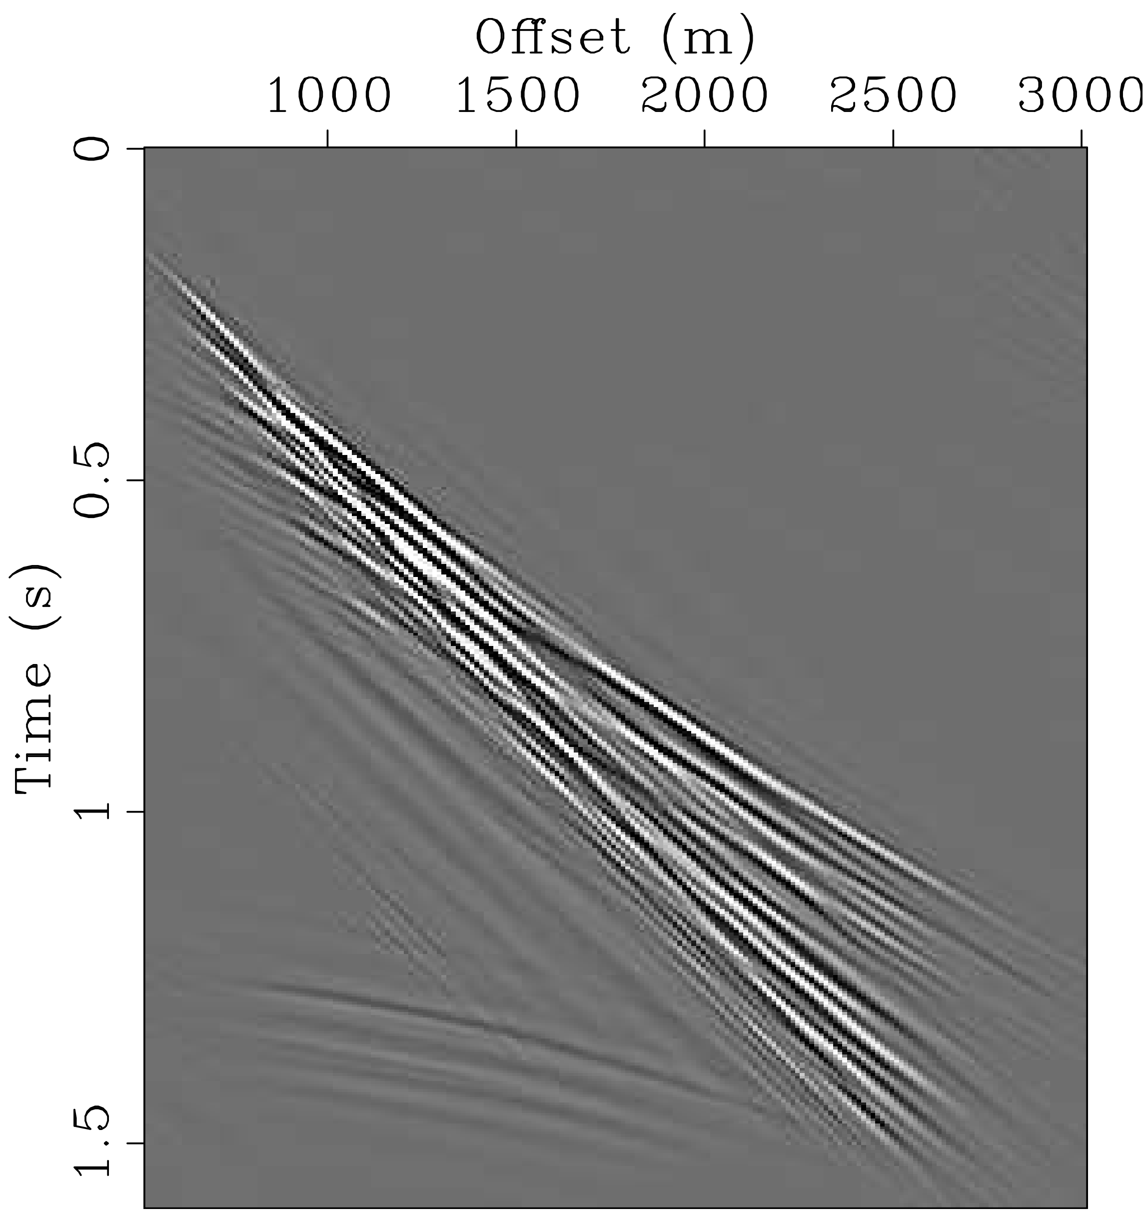

1 % of coefficients

Wavefield sampling and nonlinear recovery

- *sparsifying transform*
  - typically **localized** in the time-space domain to handle the complexity of seismic data
    - curvelet transform (dyadic-parabolic partition of the f-k domain)
    - [windowed Fourier transform]

- *sampling scheme*
  - generates incoherent random undersampling “noise” in the sparsifying domain
  - do not create large gaps
    - because of the limited spatiotemporal extend of transform elements used for the reconstruction

- *sparsity-promoting solver*
  - requires few matrix-vector multiplications

Localized transform elements & gap size

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} ||\mathbf{x}||_1 \quad \text{s.t.} \quad \mathbf{y} = \mathbf{A}\mathbf{x}$$

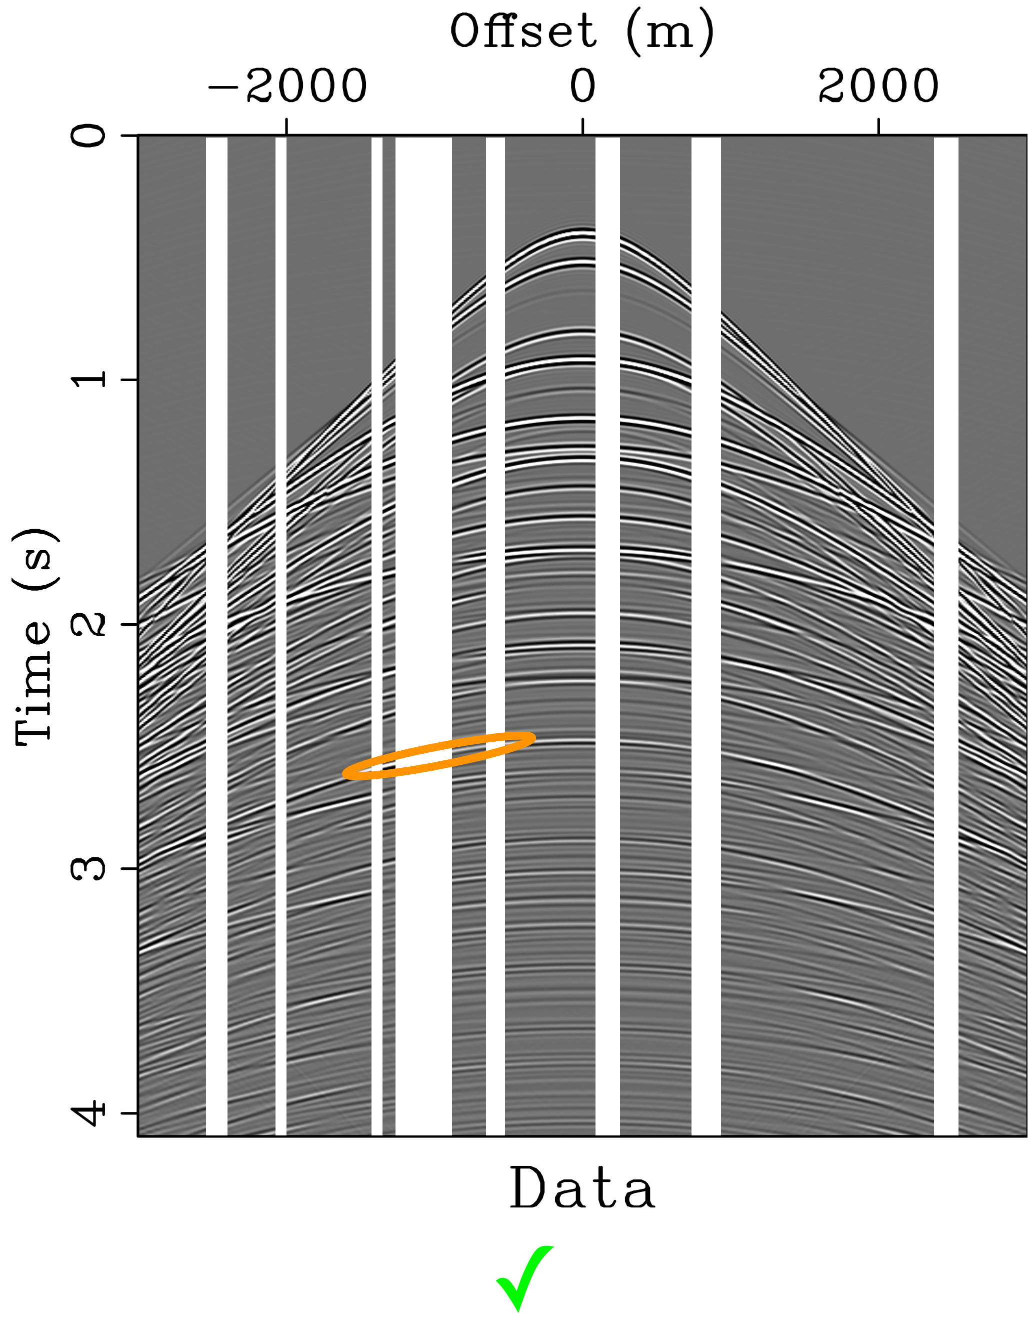

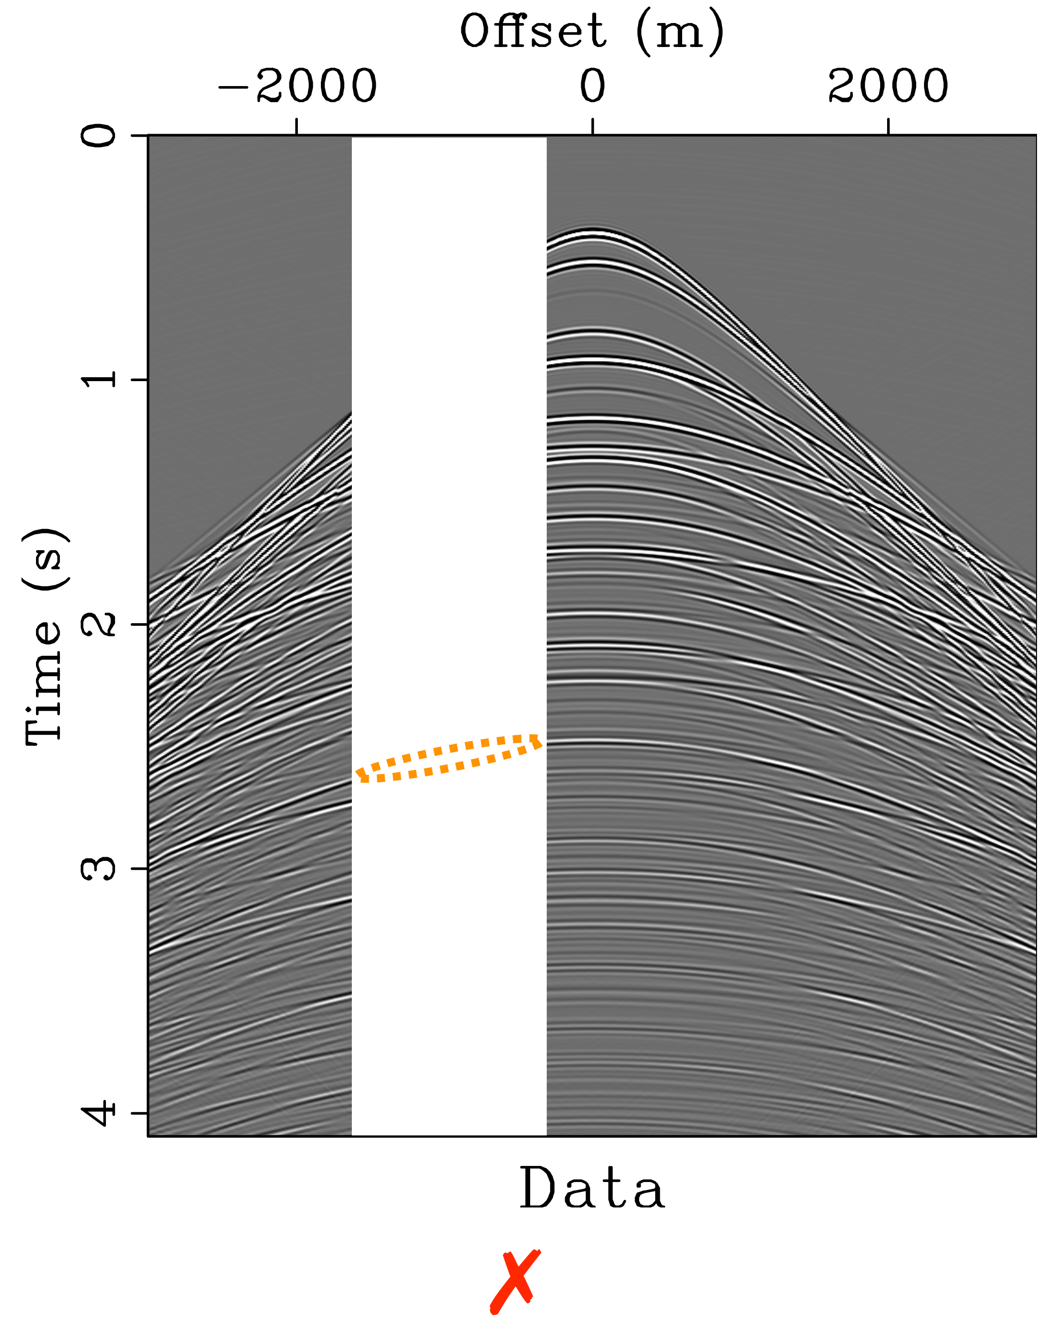

Discrete *randomized* jittered undersampling

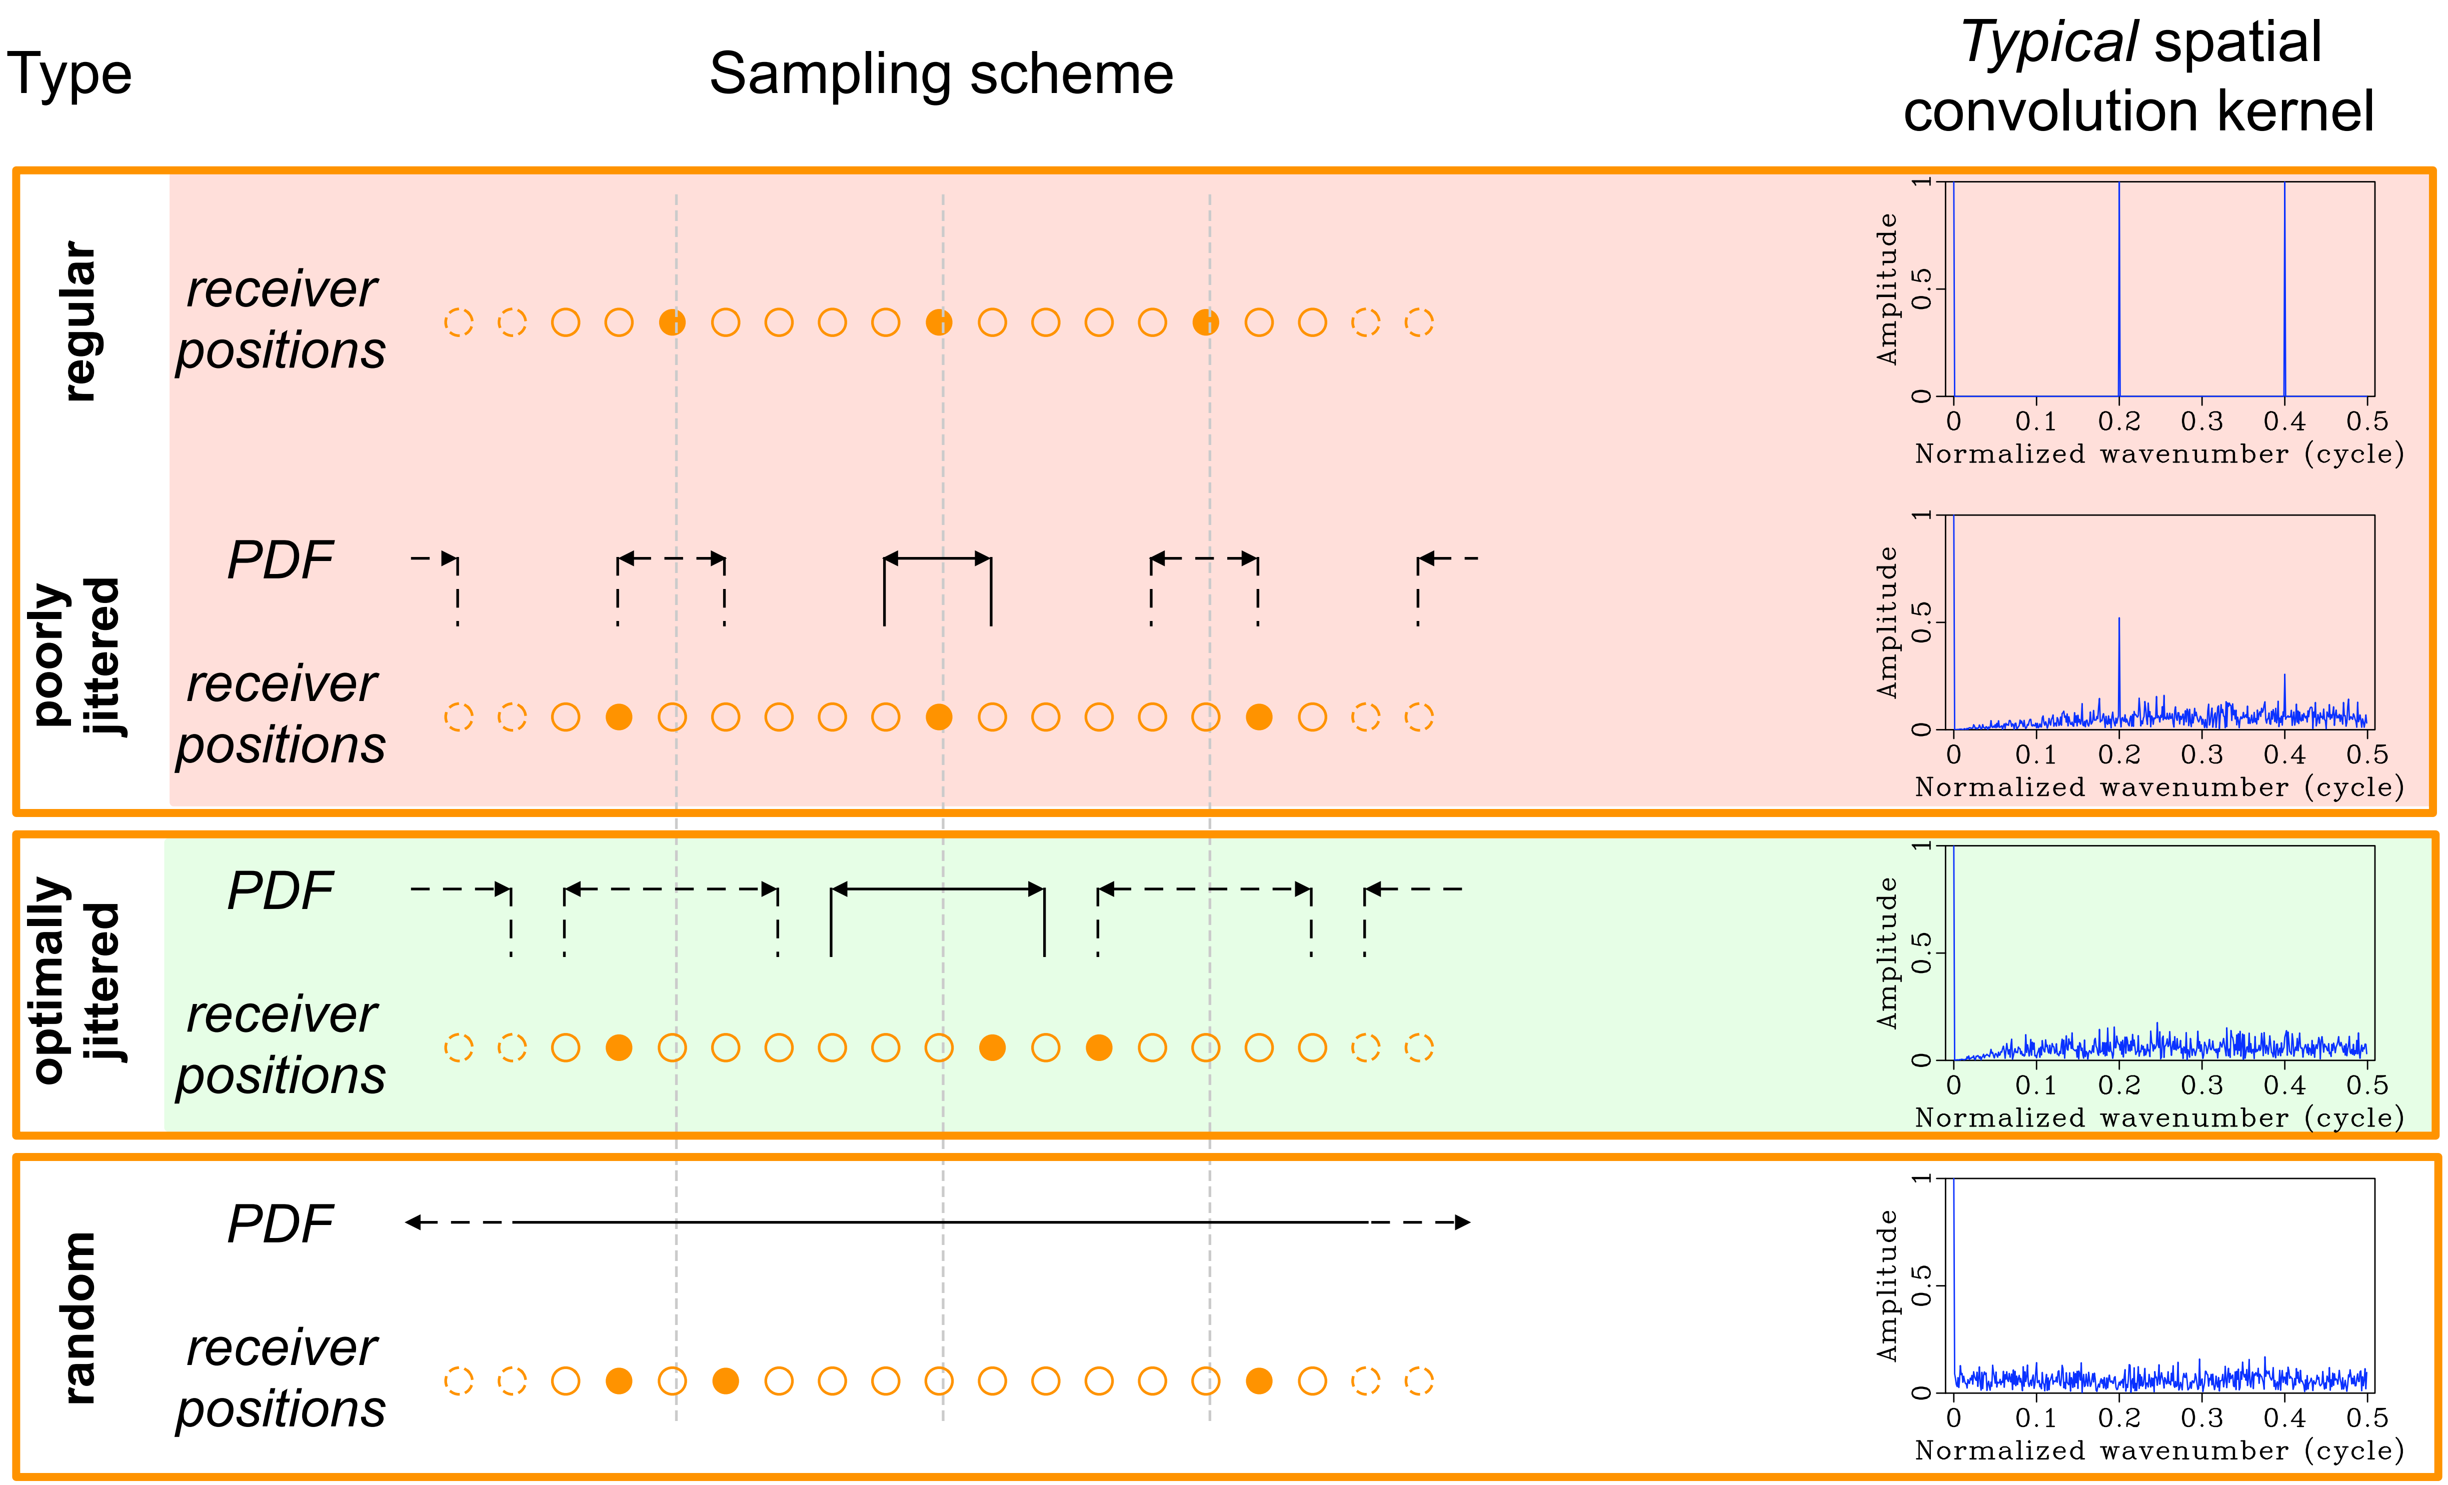

[Hennenfent and FJH ‘08]

Model

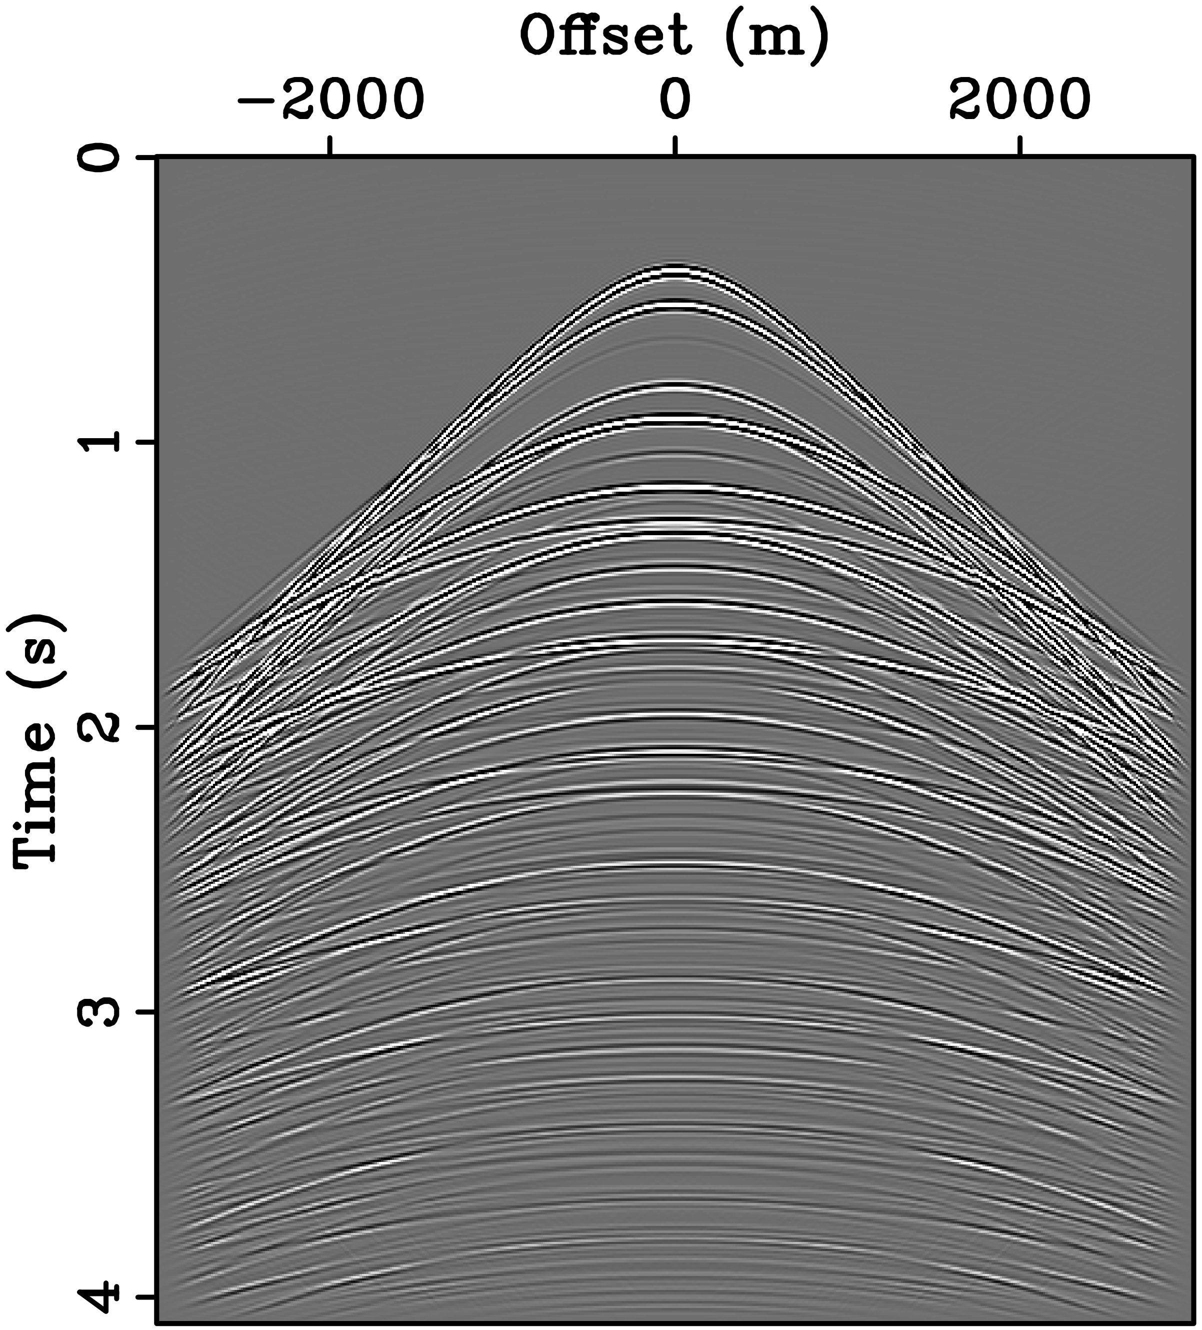
Offset (m)
−2000
0
2000
Time (s)
0
1
2
3
4

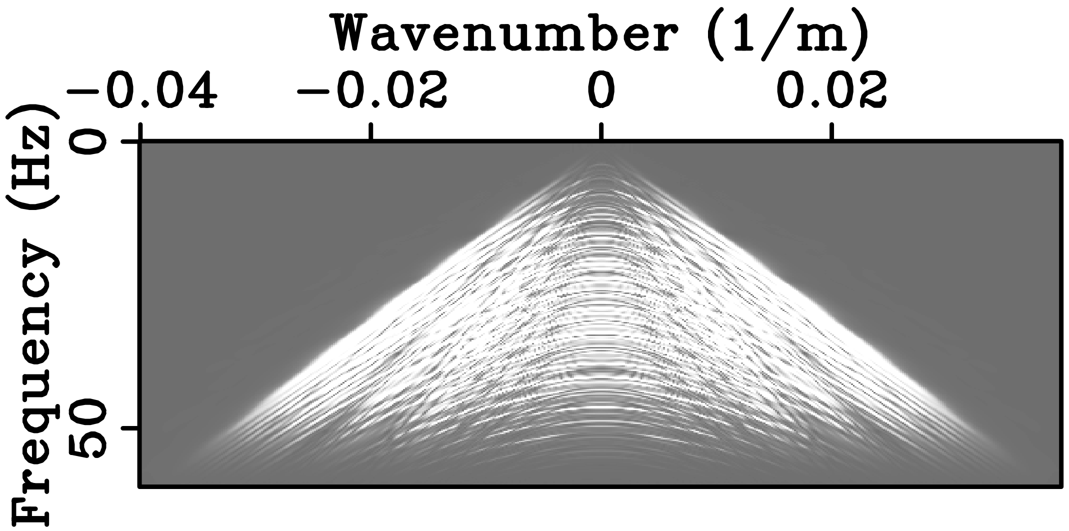
Wavenumber (1/m)
−0.04
−0.02
0
0.02
Frequency (Hz)
0
50

Random 3-fold undersampling

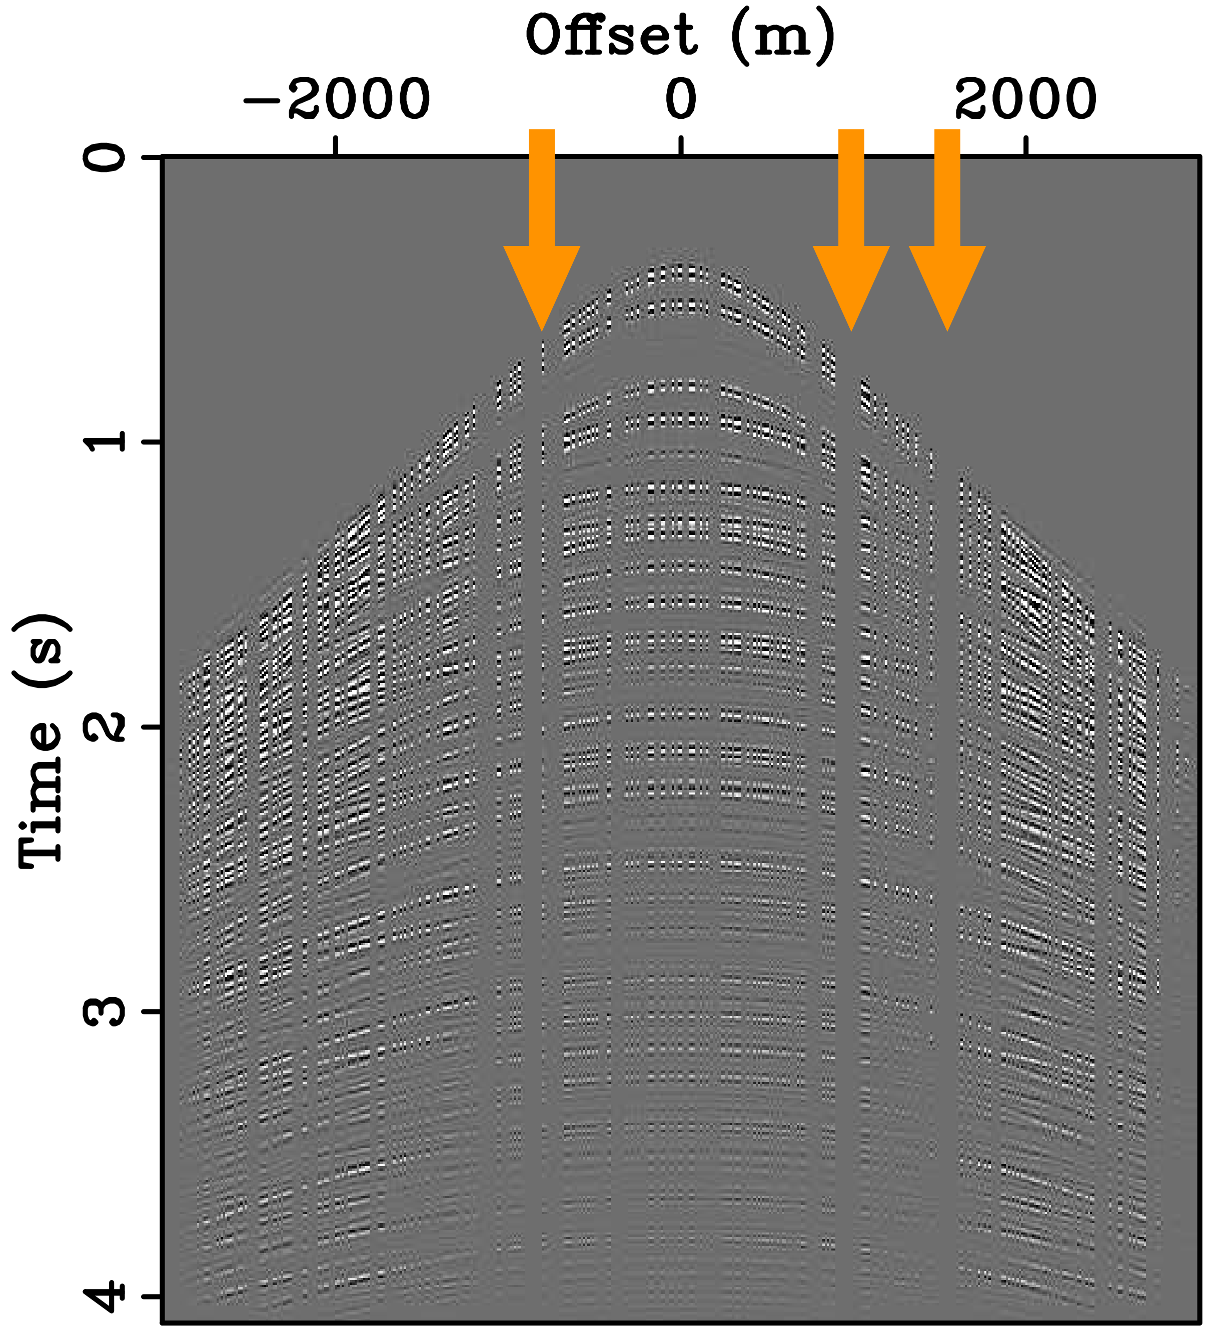

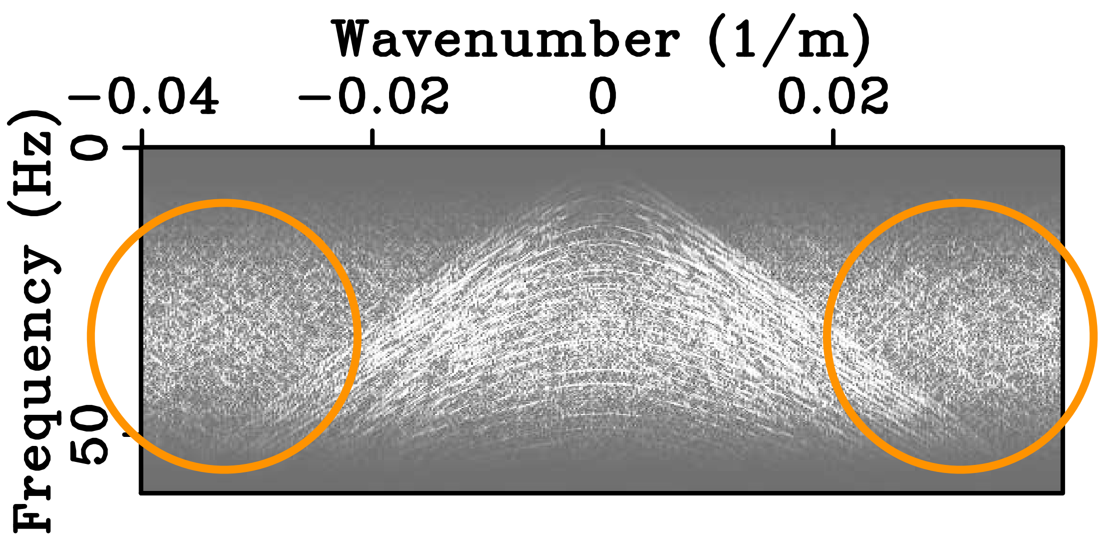

CRSI from random 3-fold undersampling

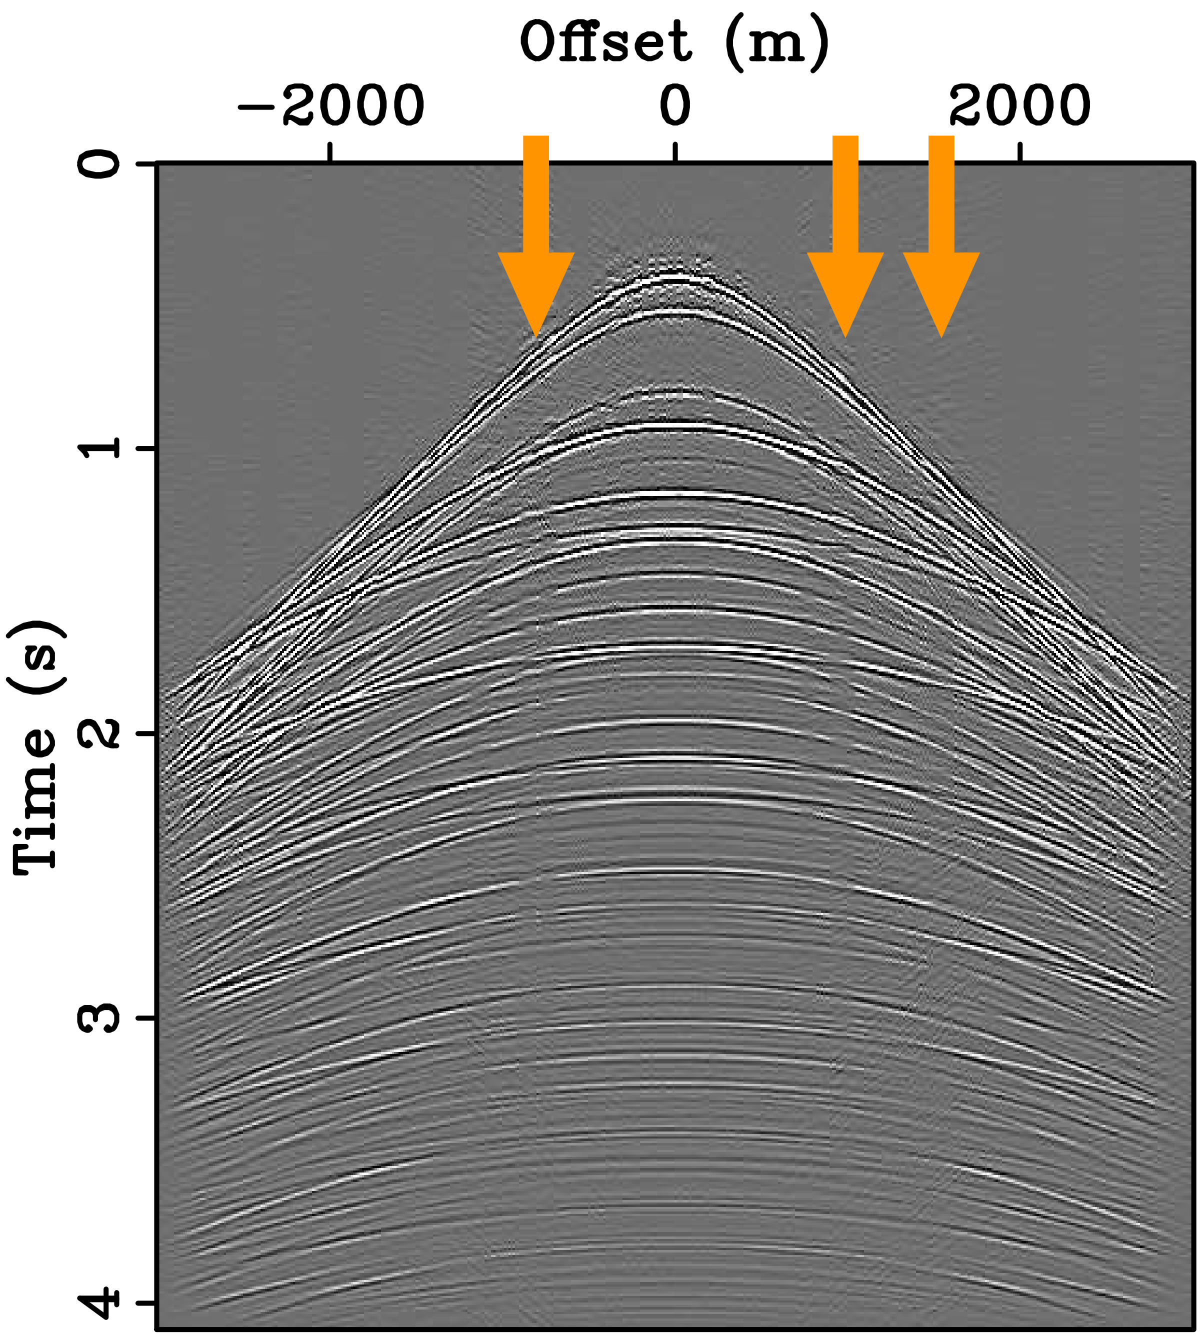

SNR = 9.72 dB

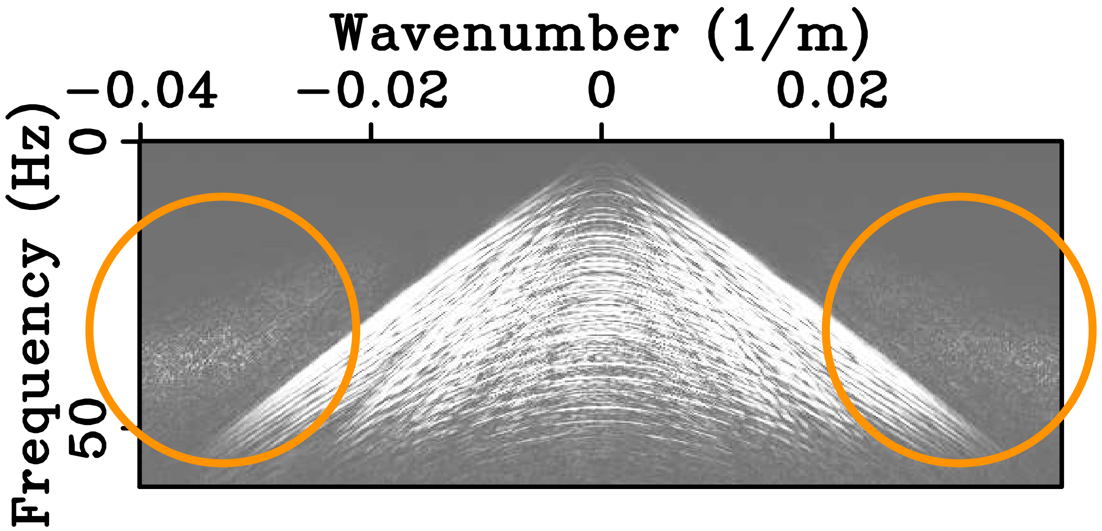

$$\mathrm{SNR} = 20 \times \log_{10}\left(\frac{\|\mathrm{model}\|_2}{\|\mathrm{reconstruction\ error}\|_2}\right)$$

Optimally-jittered 3-fold undersampling

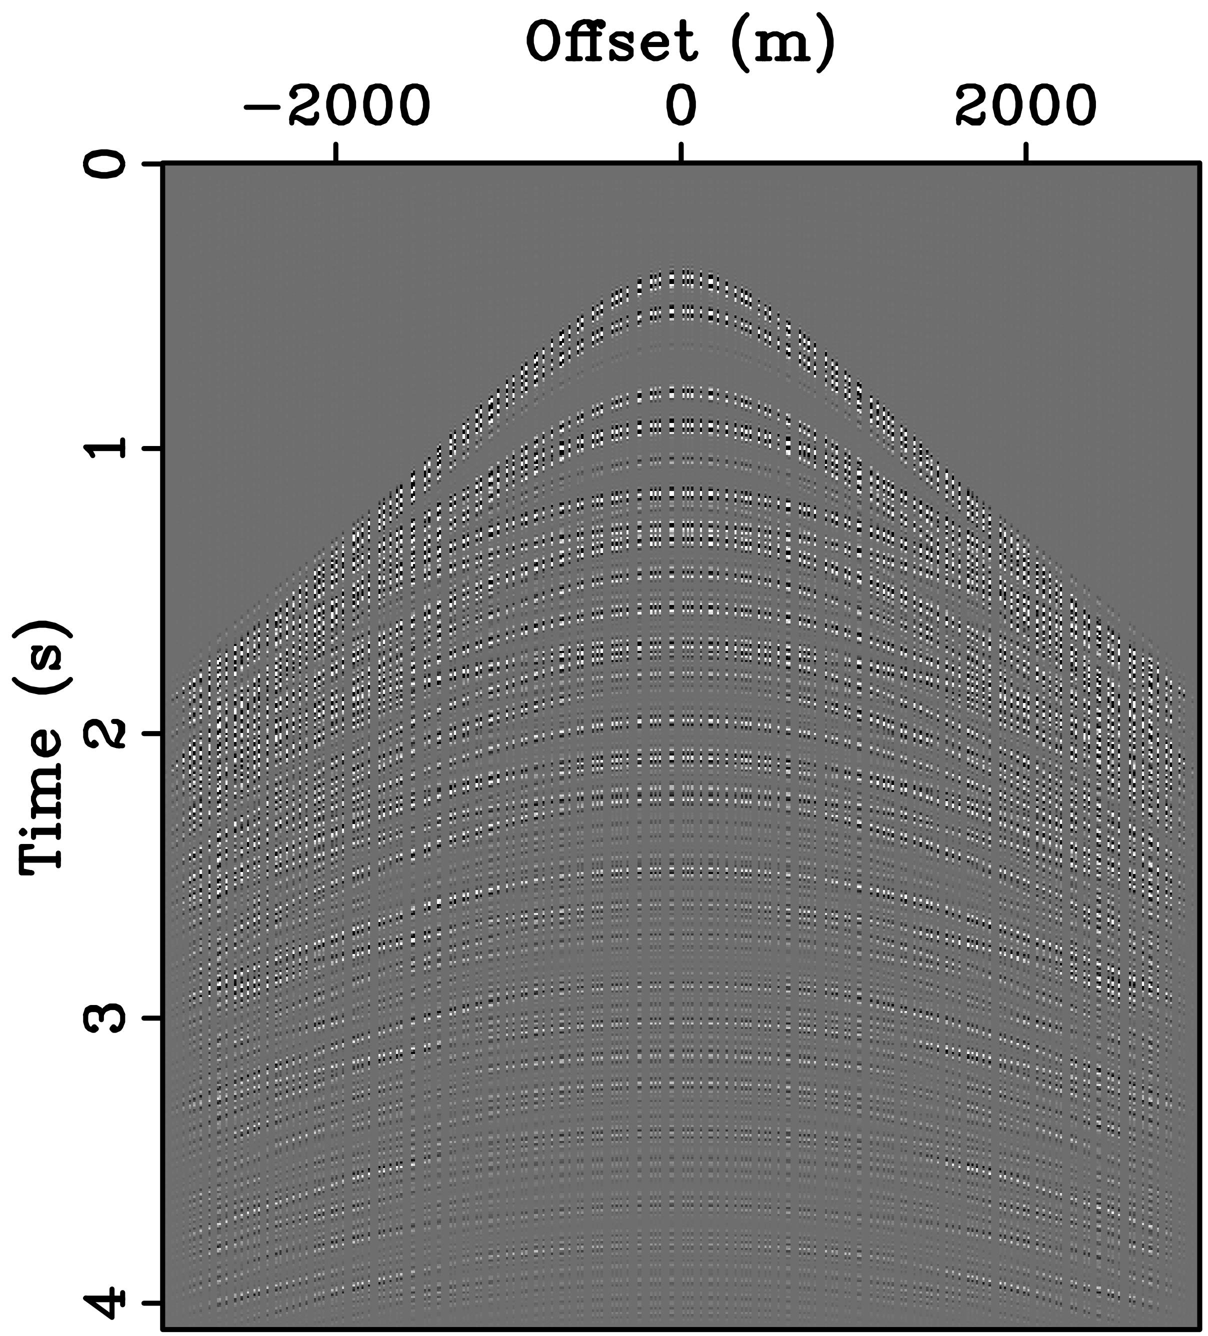

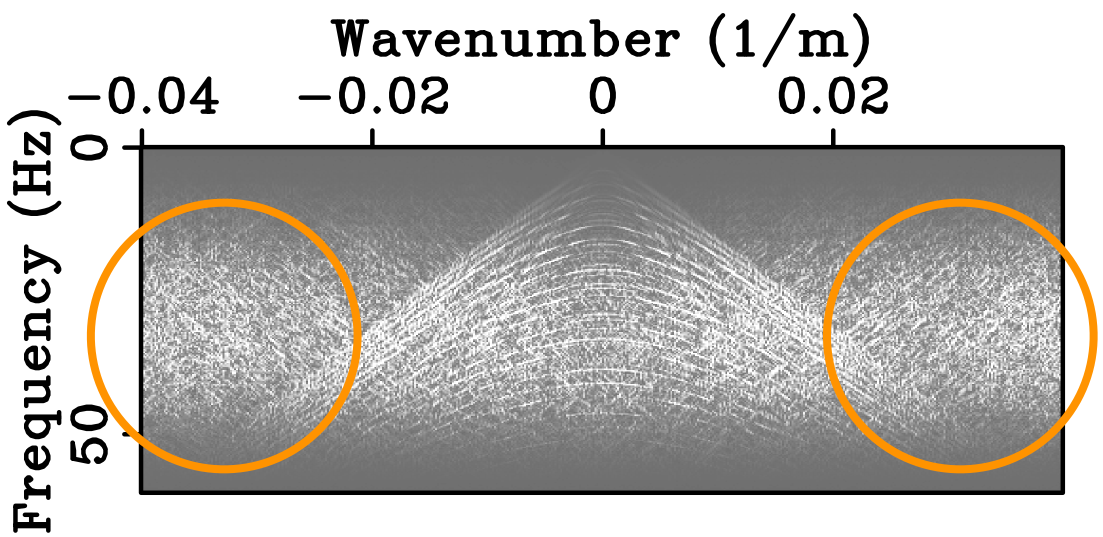

CRSI from opt.-jittered 3-fold undersampling

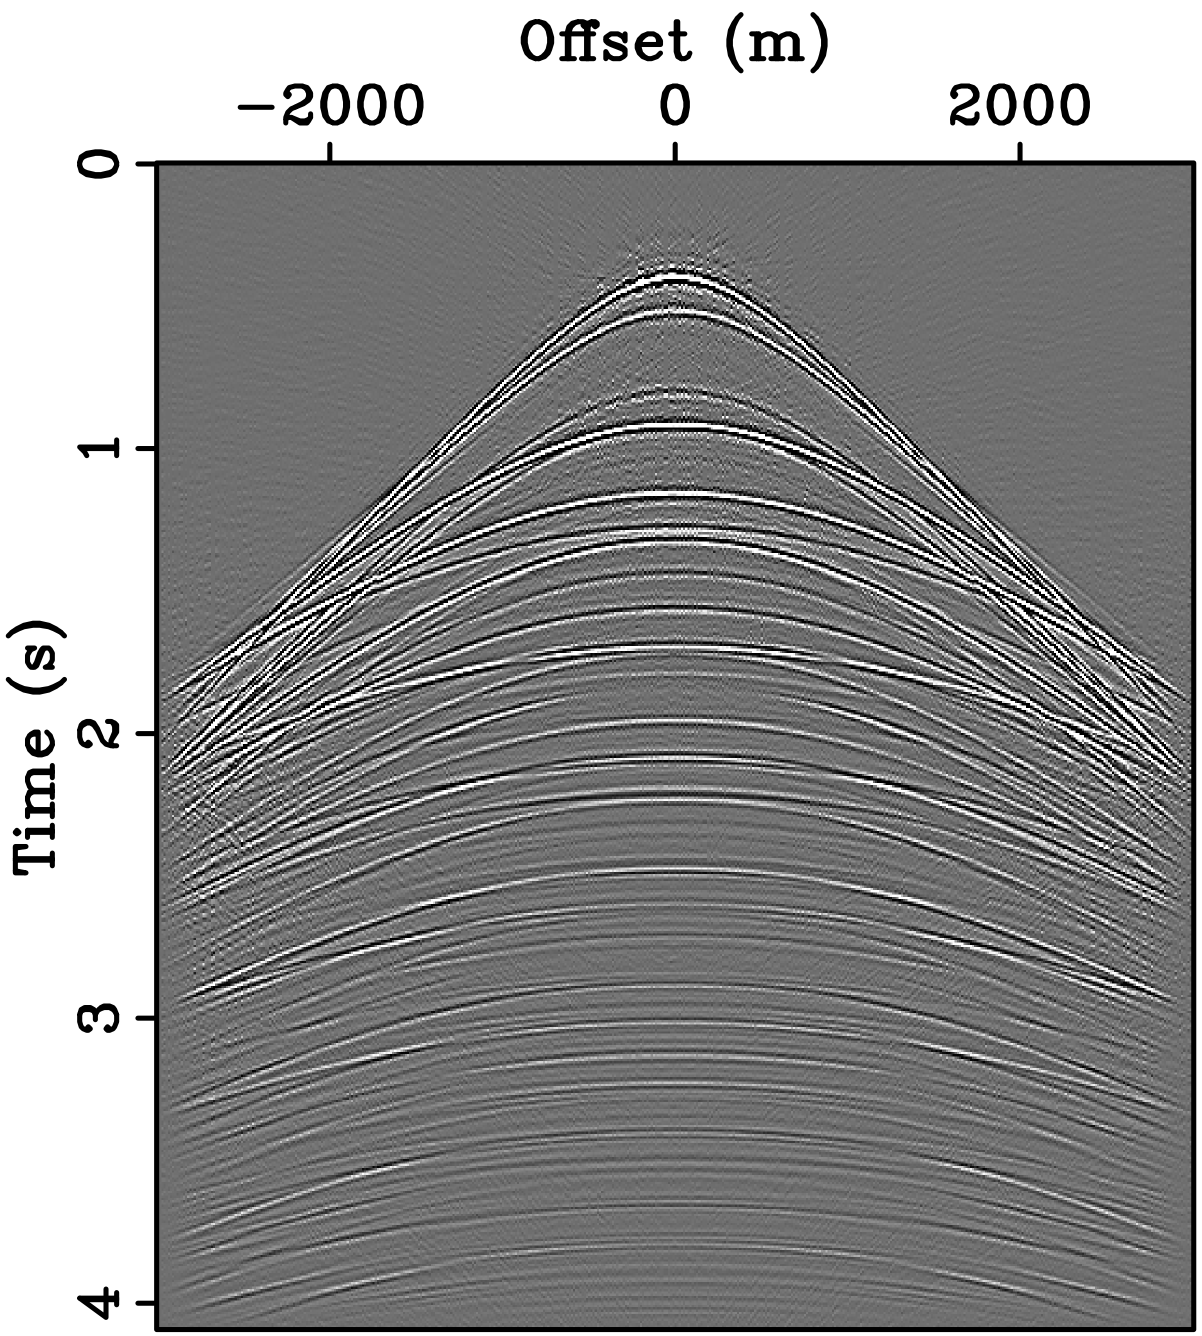

SNR = 10.42 dB

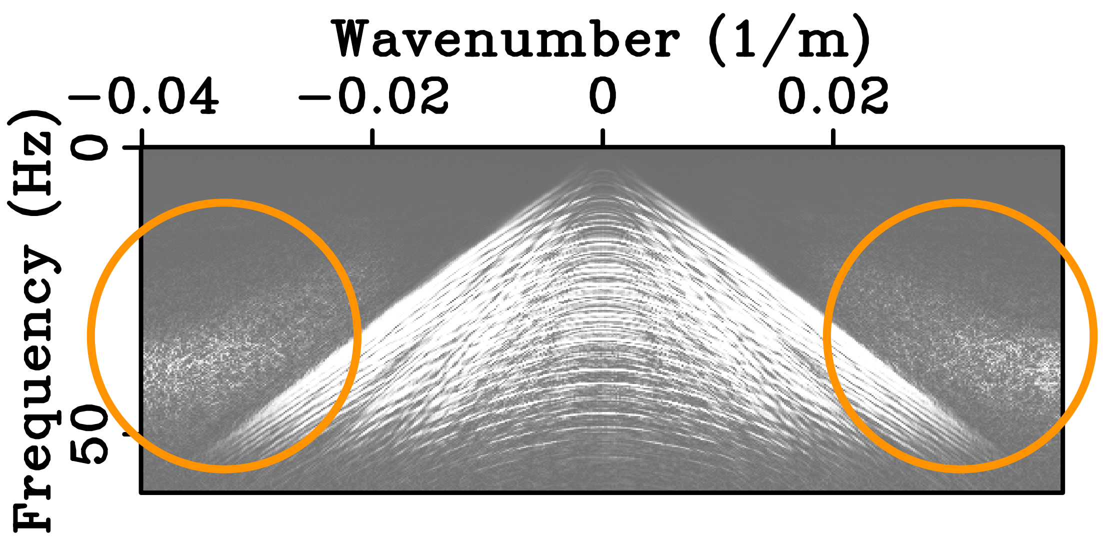

Question

- **Question: What is better? Having missing *single-source* or *missing randomized simultaneous* experiments?**

- Comparison between different undersampling strategies for source experiments:
  - ***Randomized*** jittered shot positions
  - ***Randomized*** simultaneous shots

Model

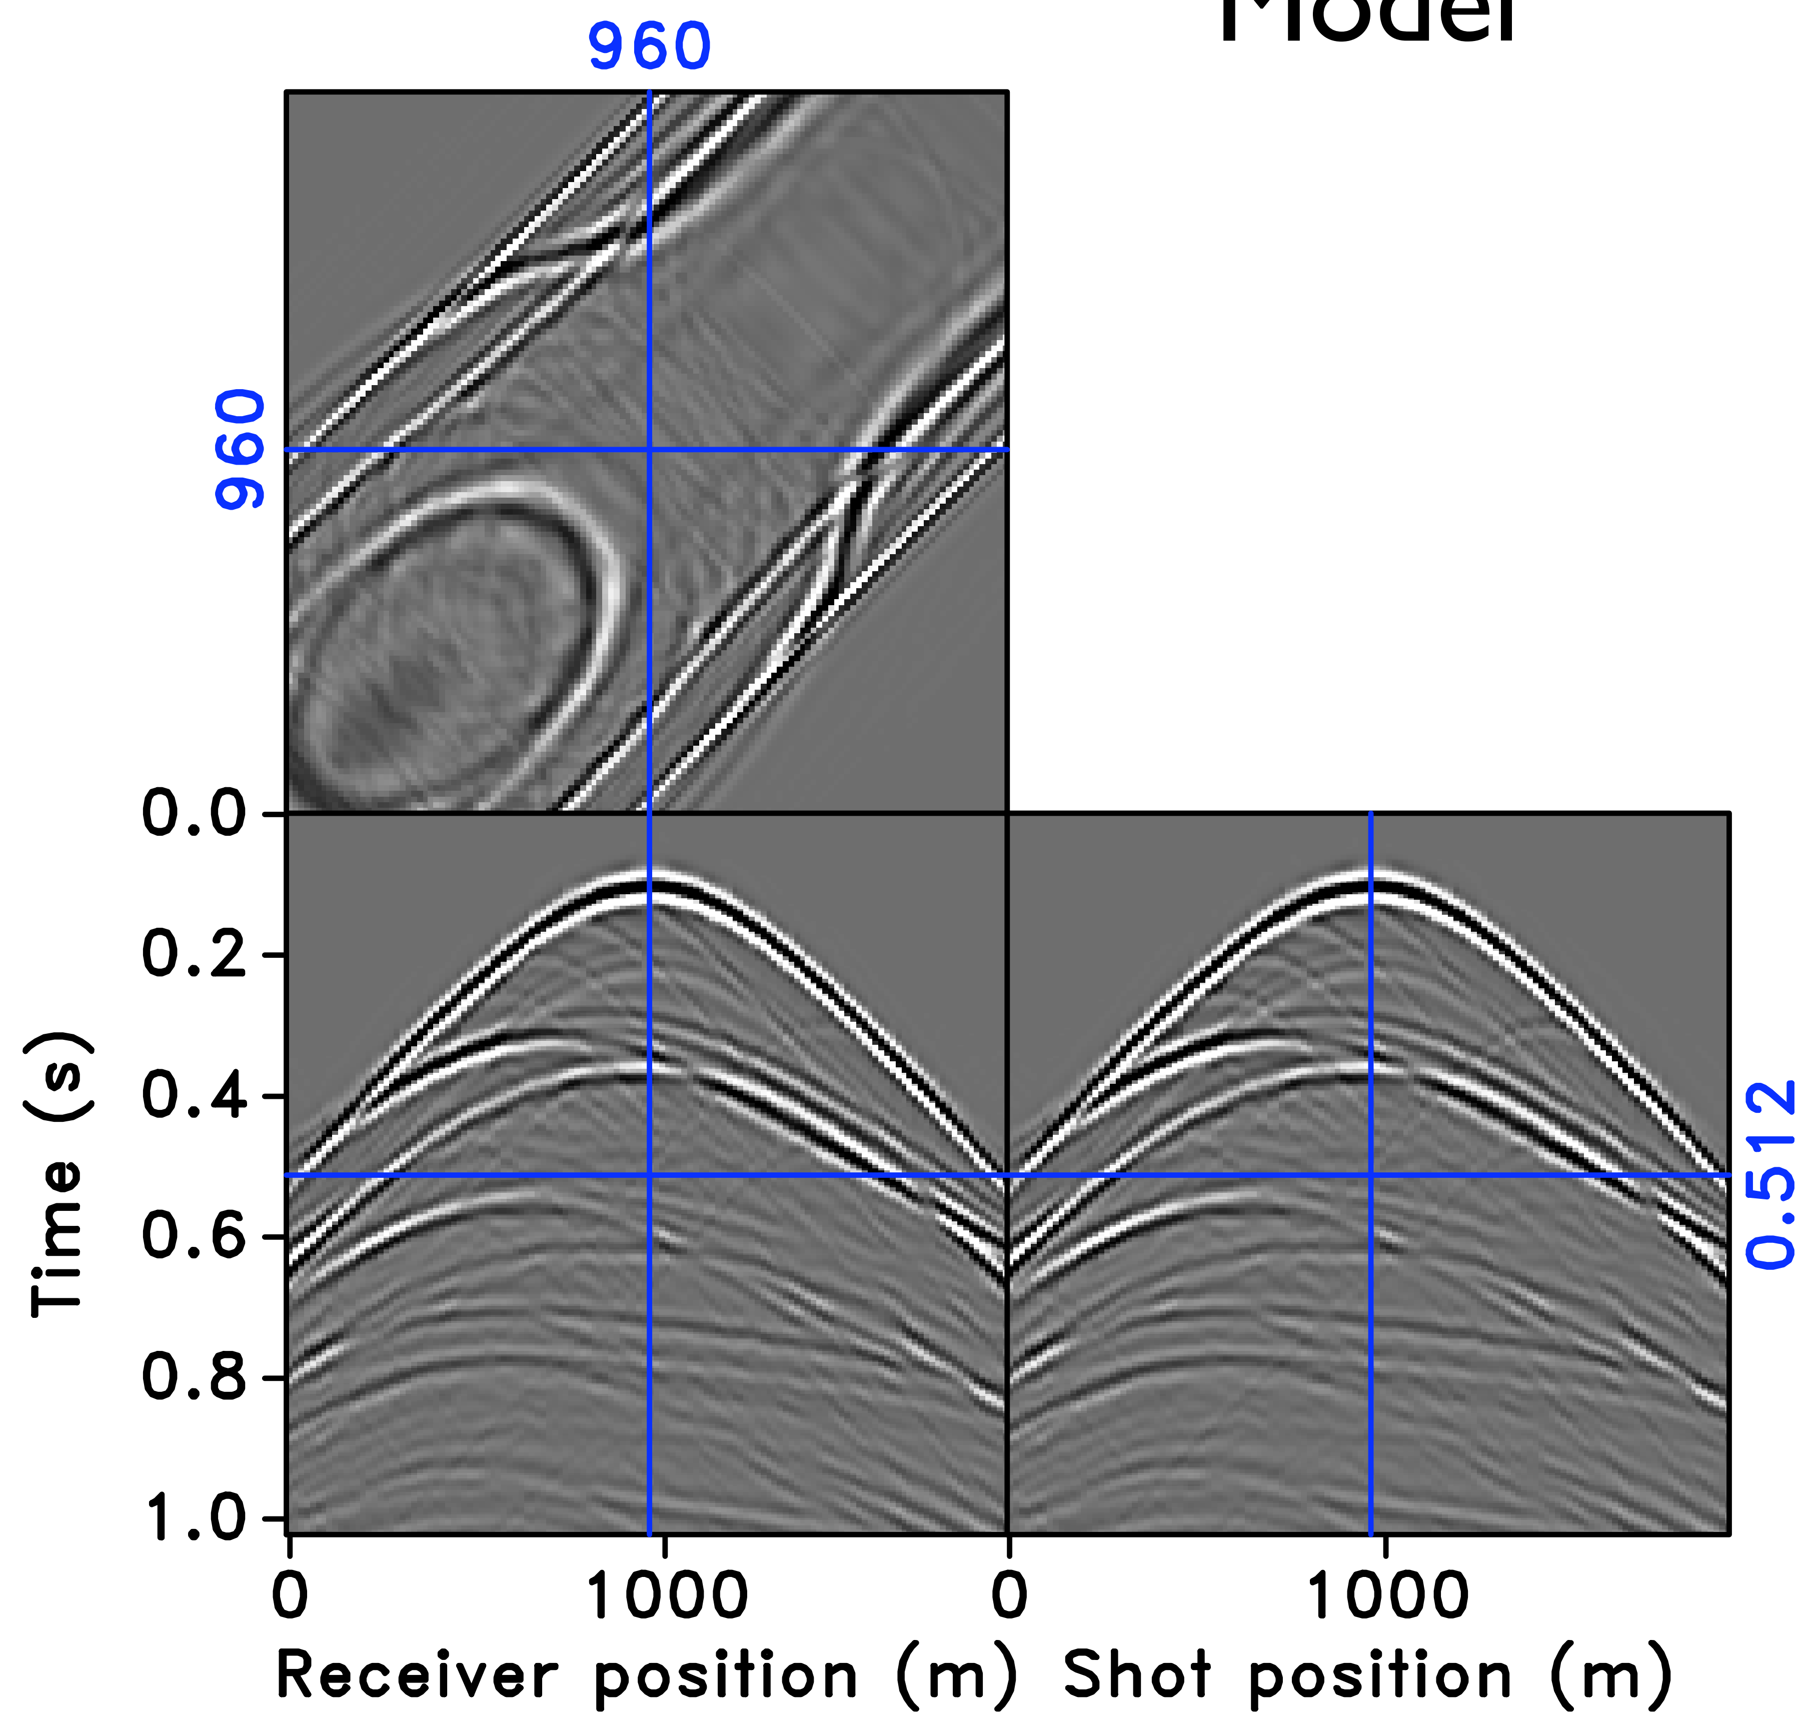

960
960
0.512
0.0
0.2
0.4
0.6
0.8
1.0
Time (s)
0
1000
0
1000
Receiver position (m)
Shot position (m)

Interpolate

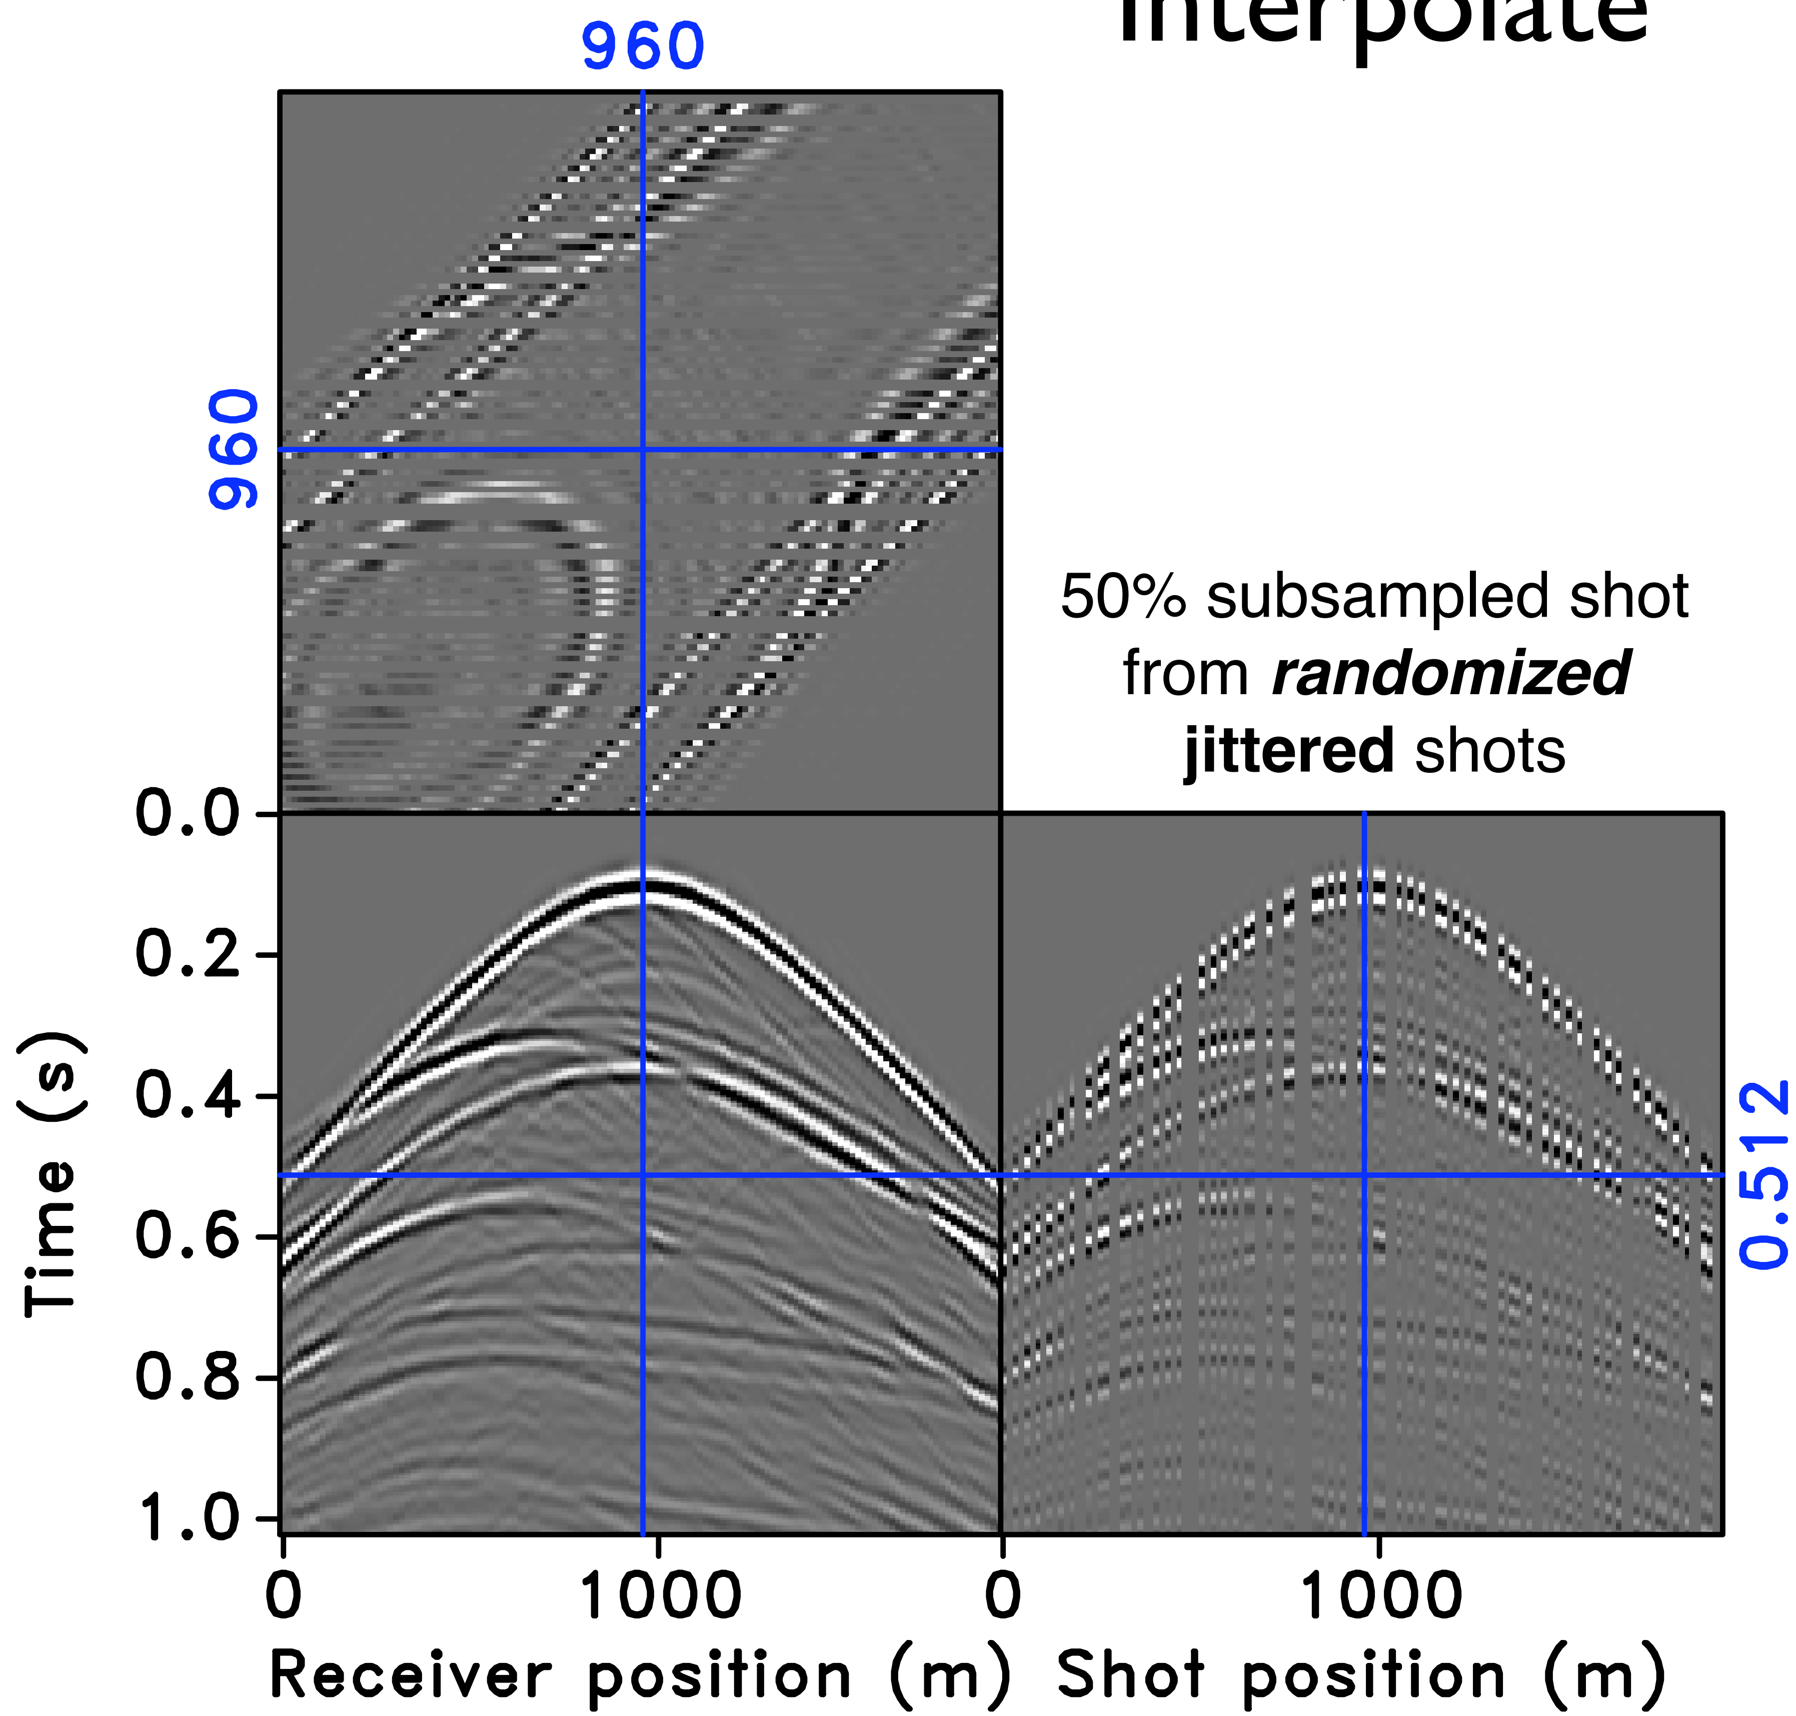

50% subsampled shot from ***randomized*** **jittered** shots

Interpolate

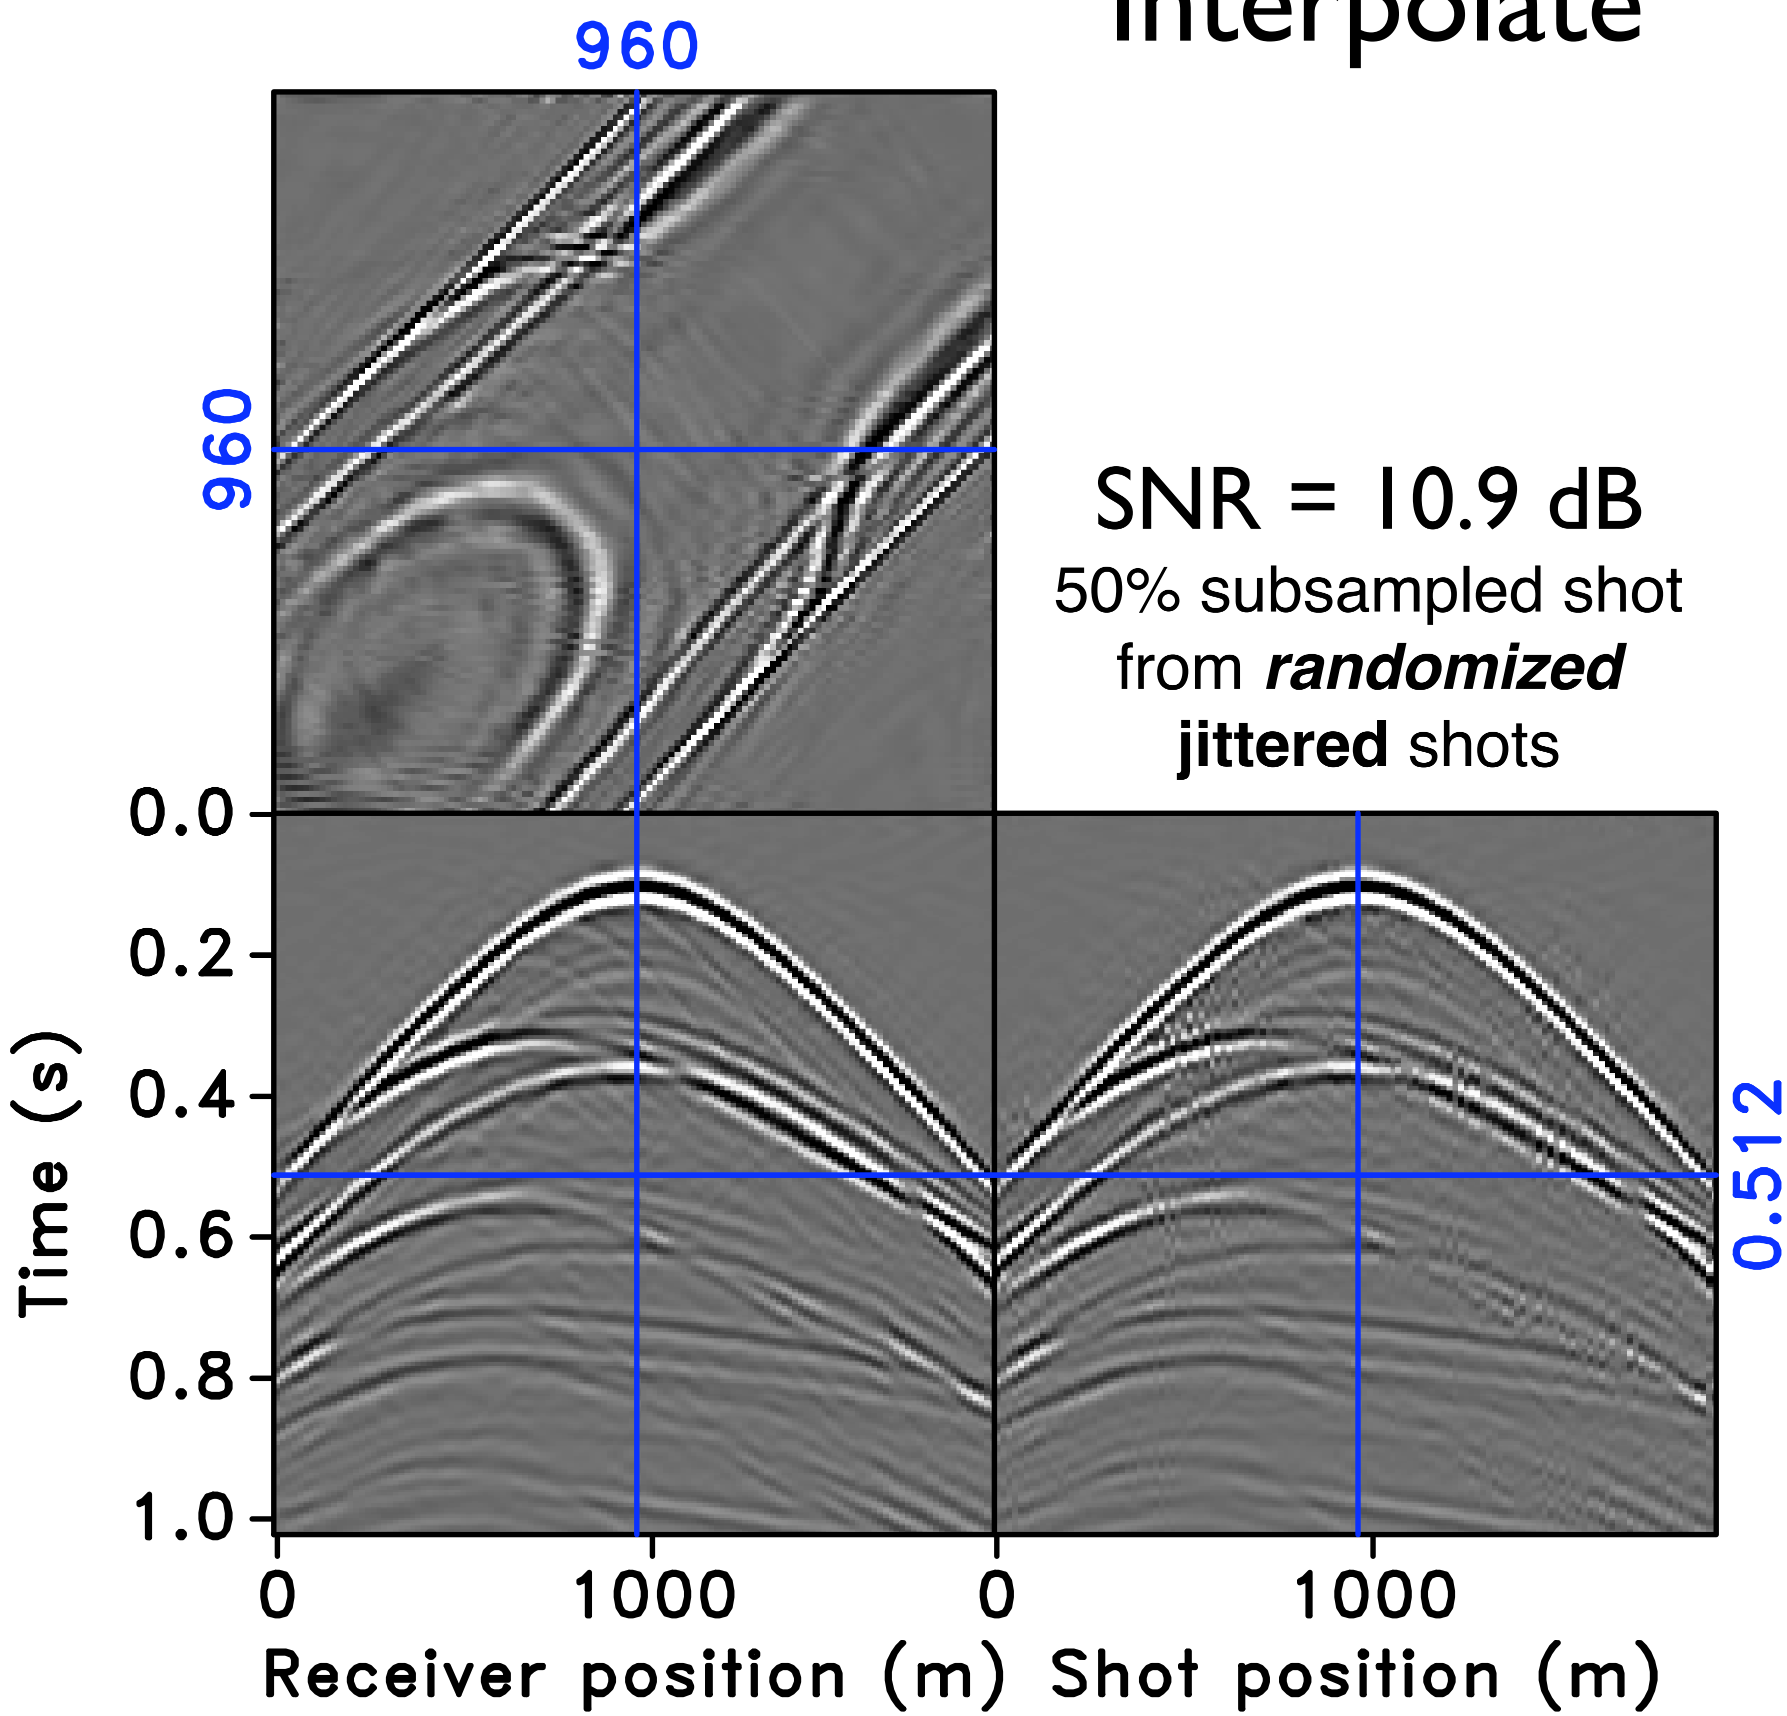

SNR = 10.9 dB

50% subsampled shot from ***randomized*** **jittered** shots

Simultaneous & continuous sources

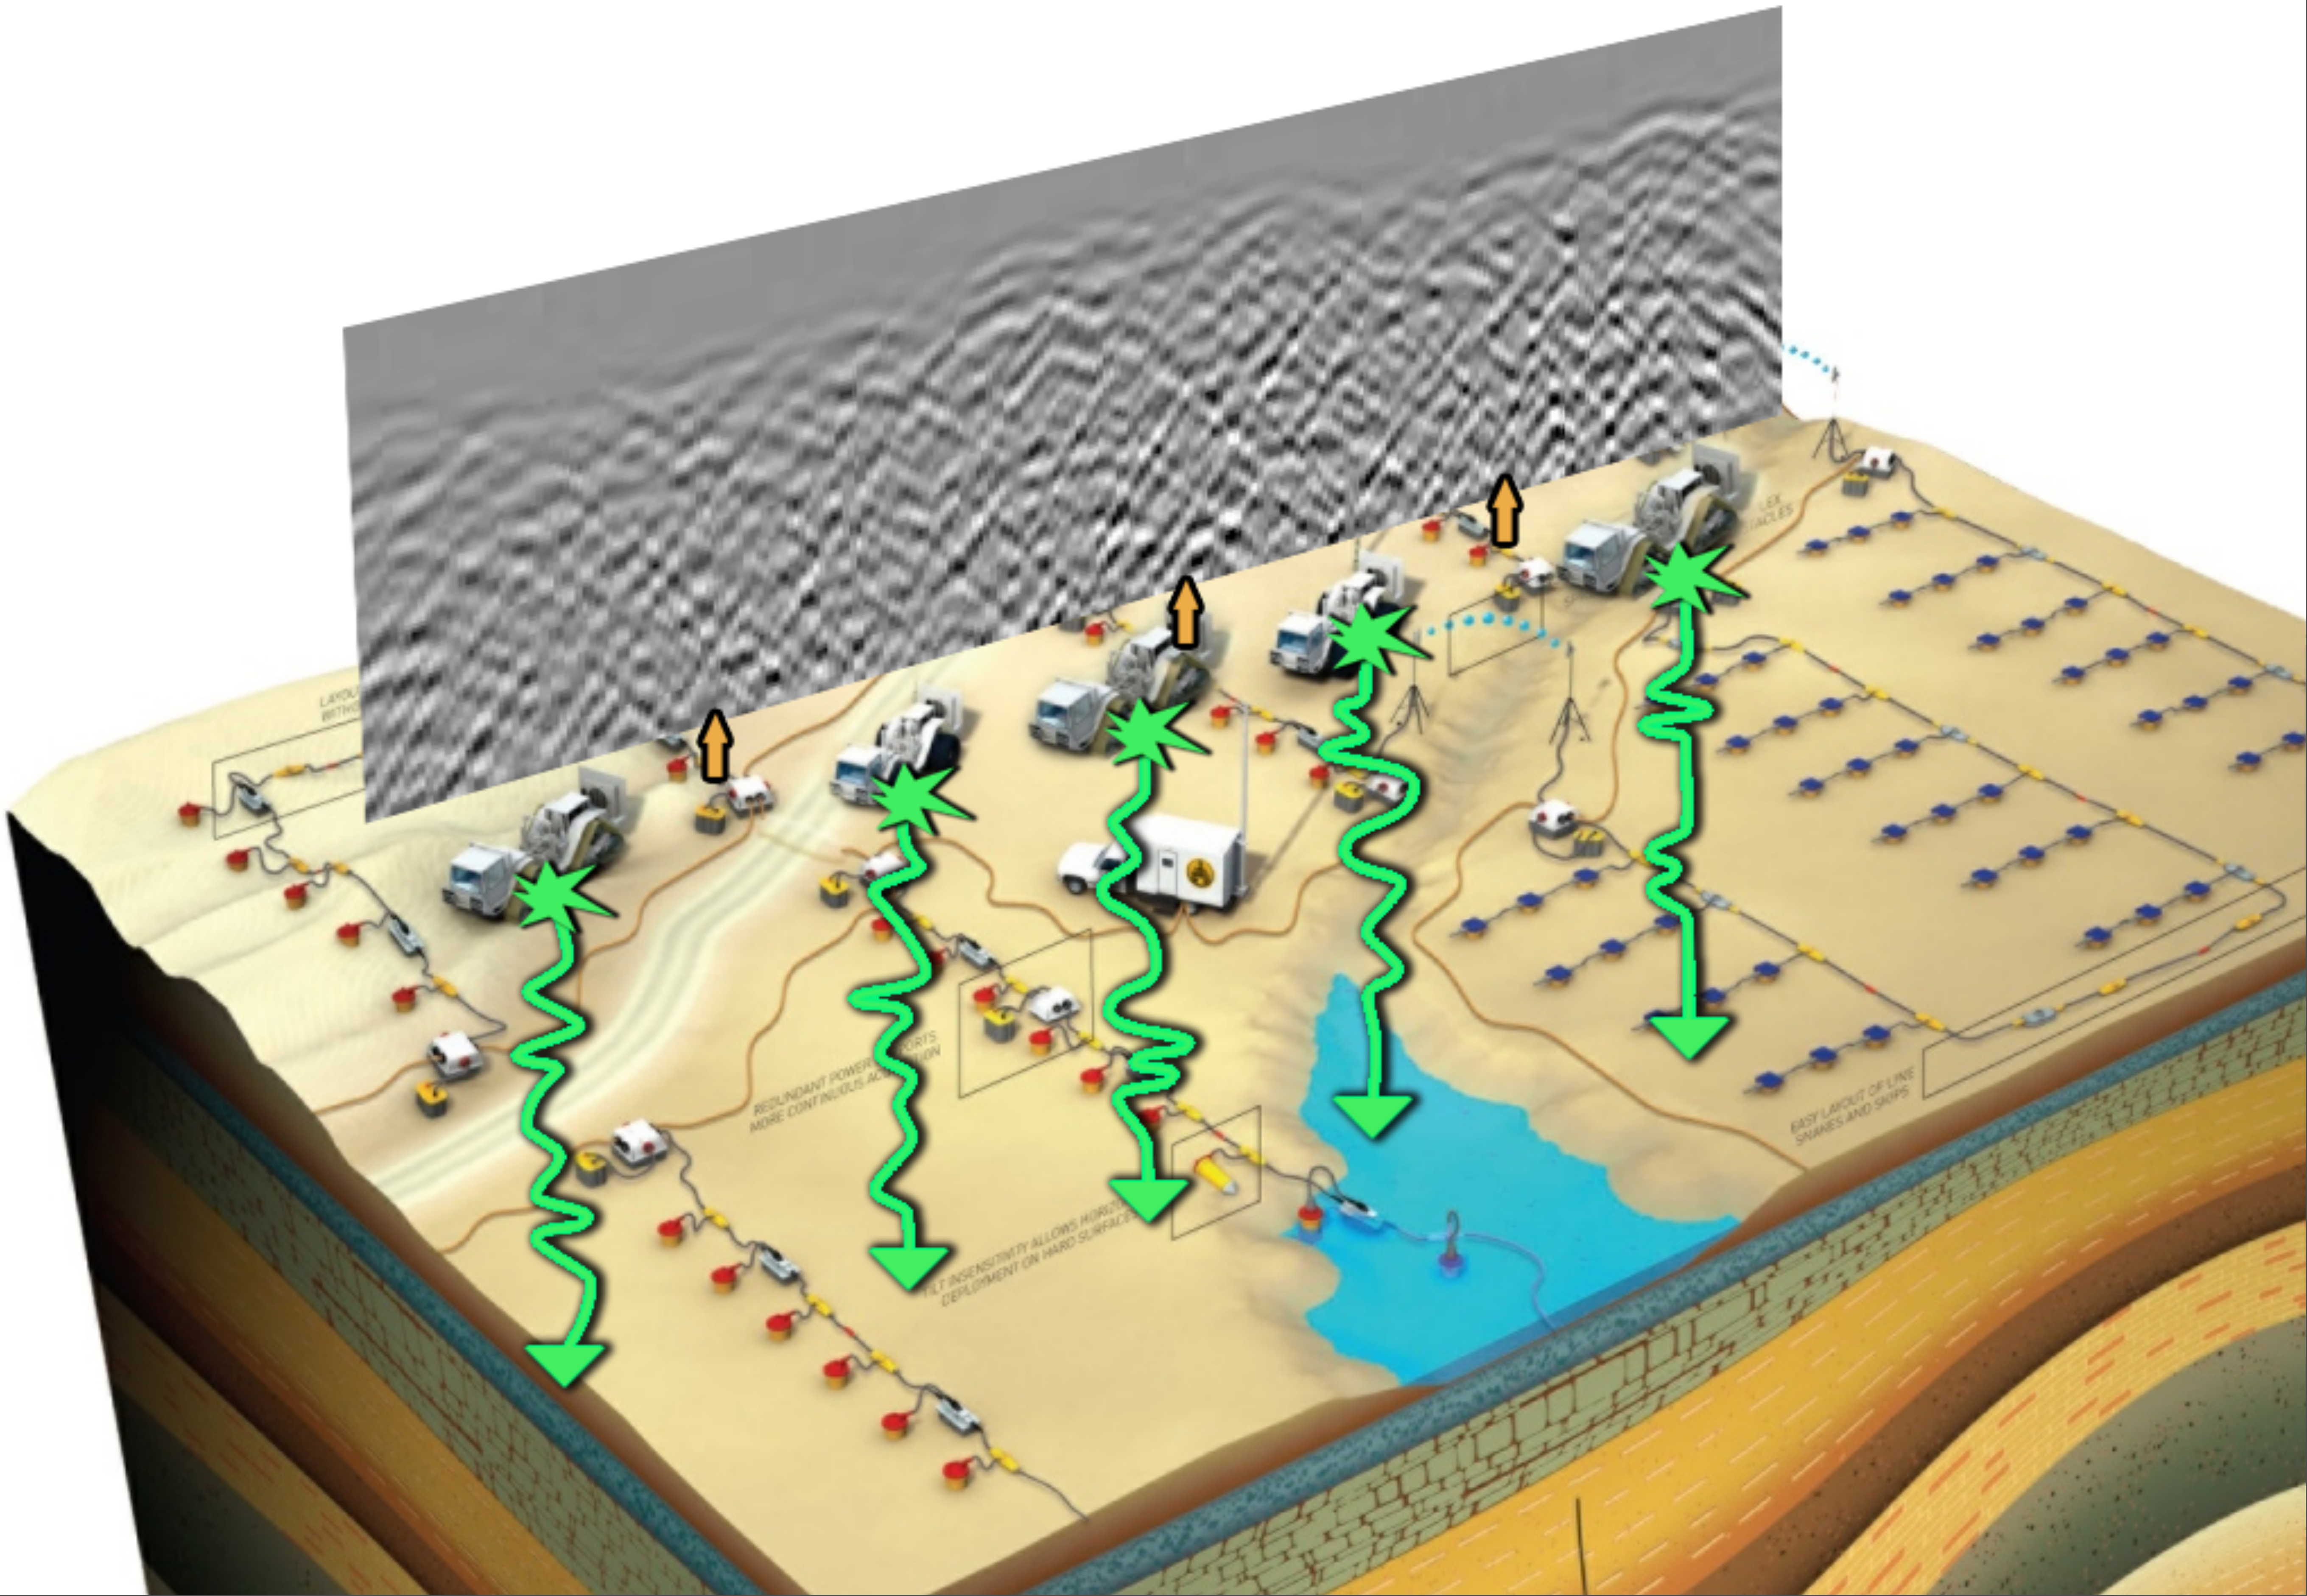

Demultiplex

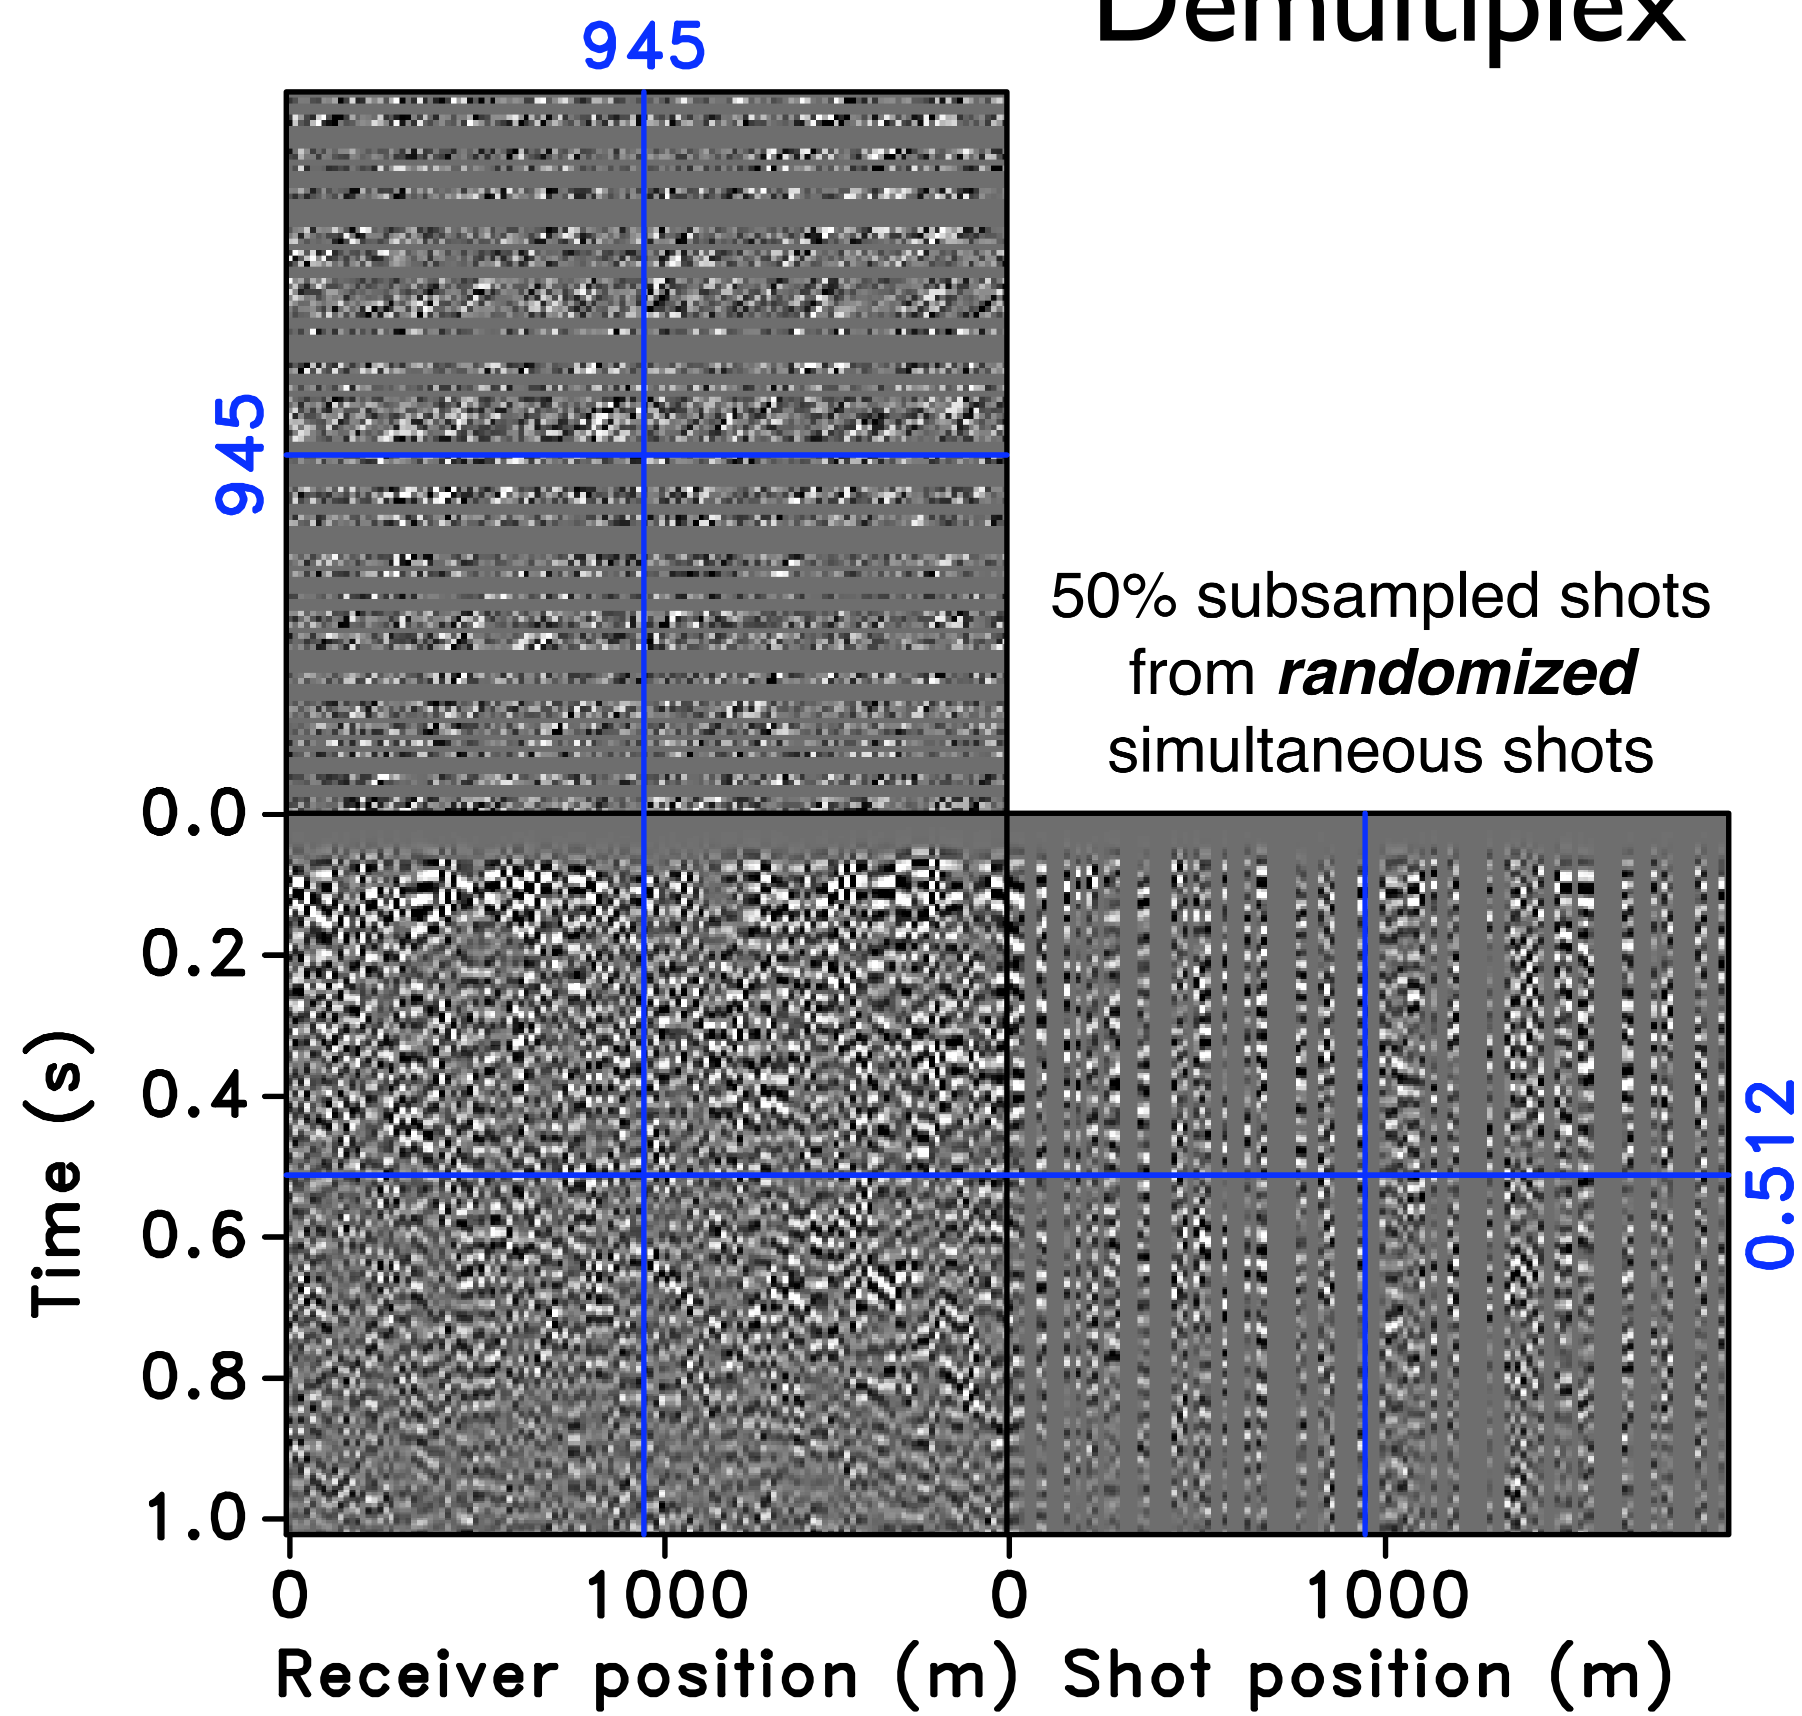

Interpolate

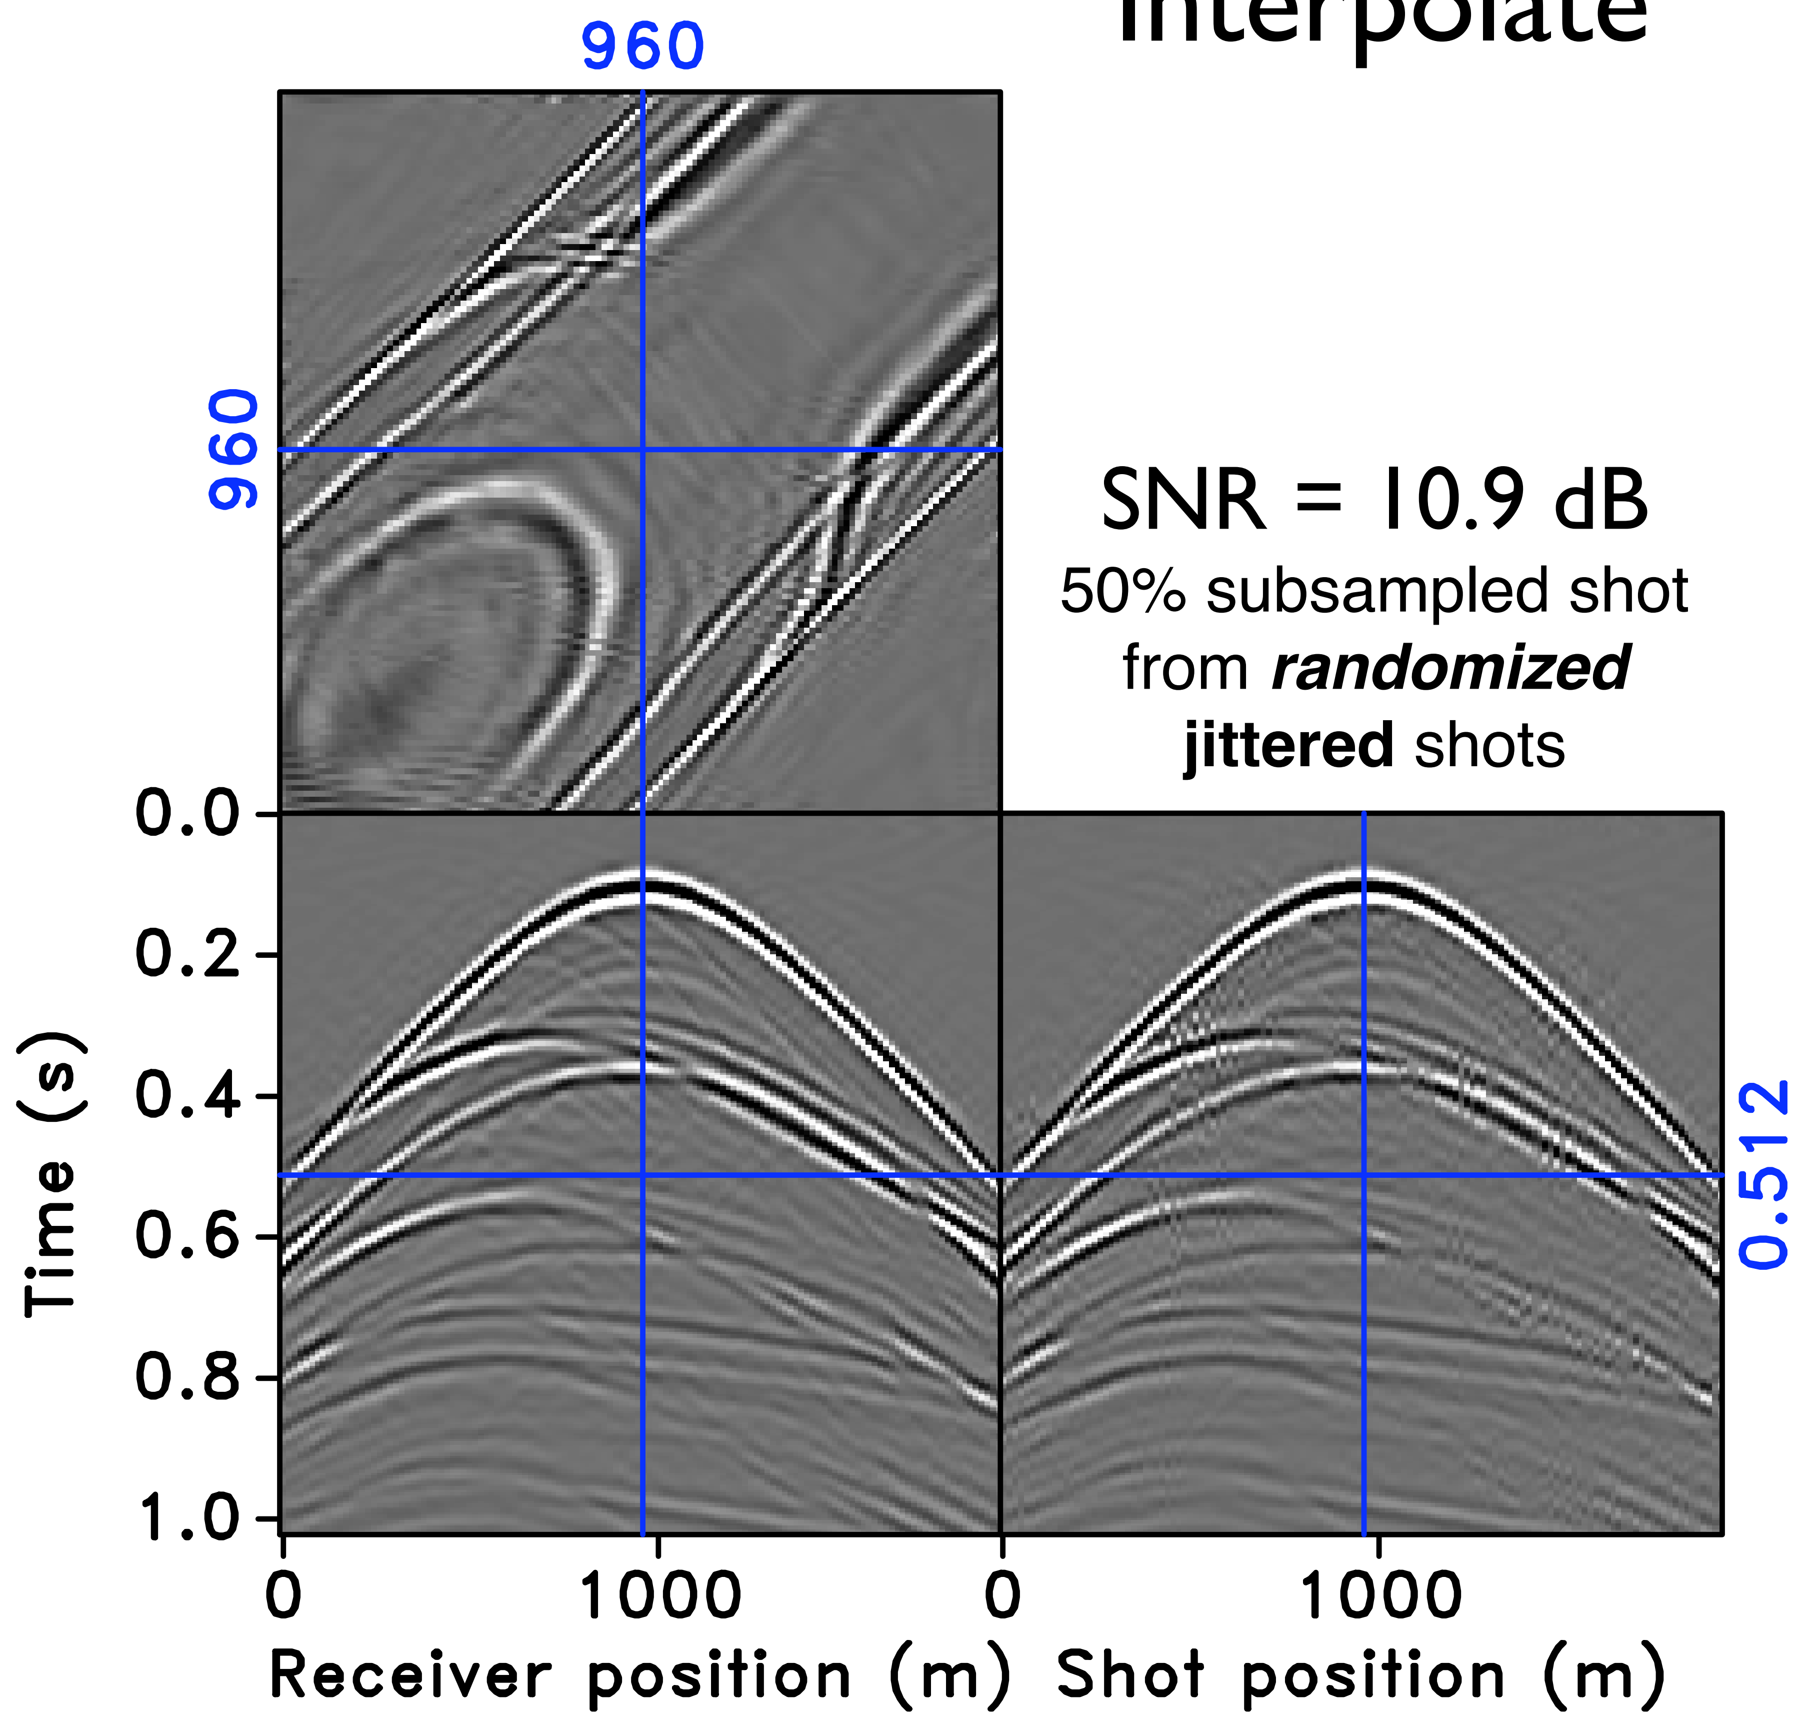

SNR = 10.9 dB

50% subsampled shot from ***randomized*** **jittered** shots

Demultiplex

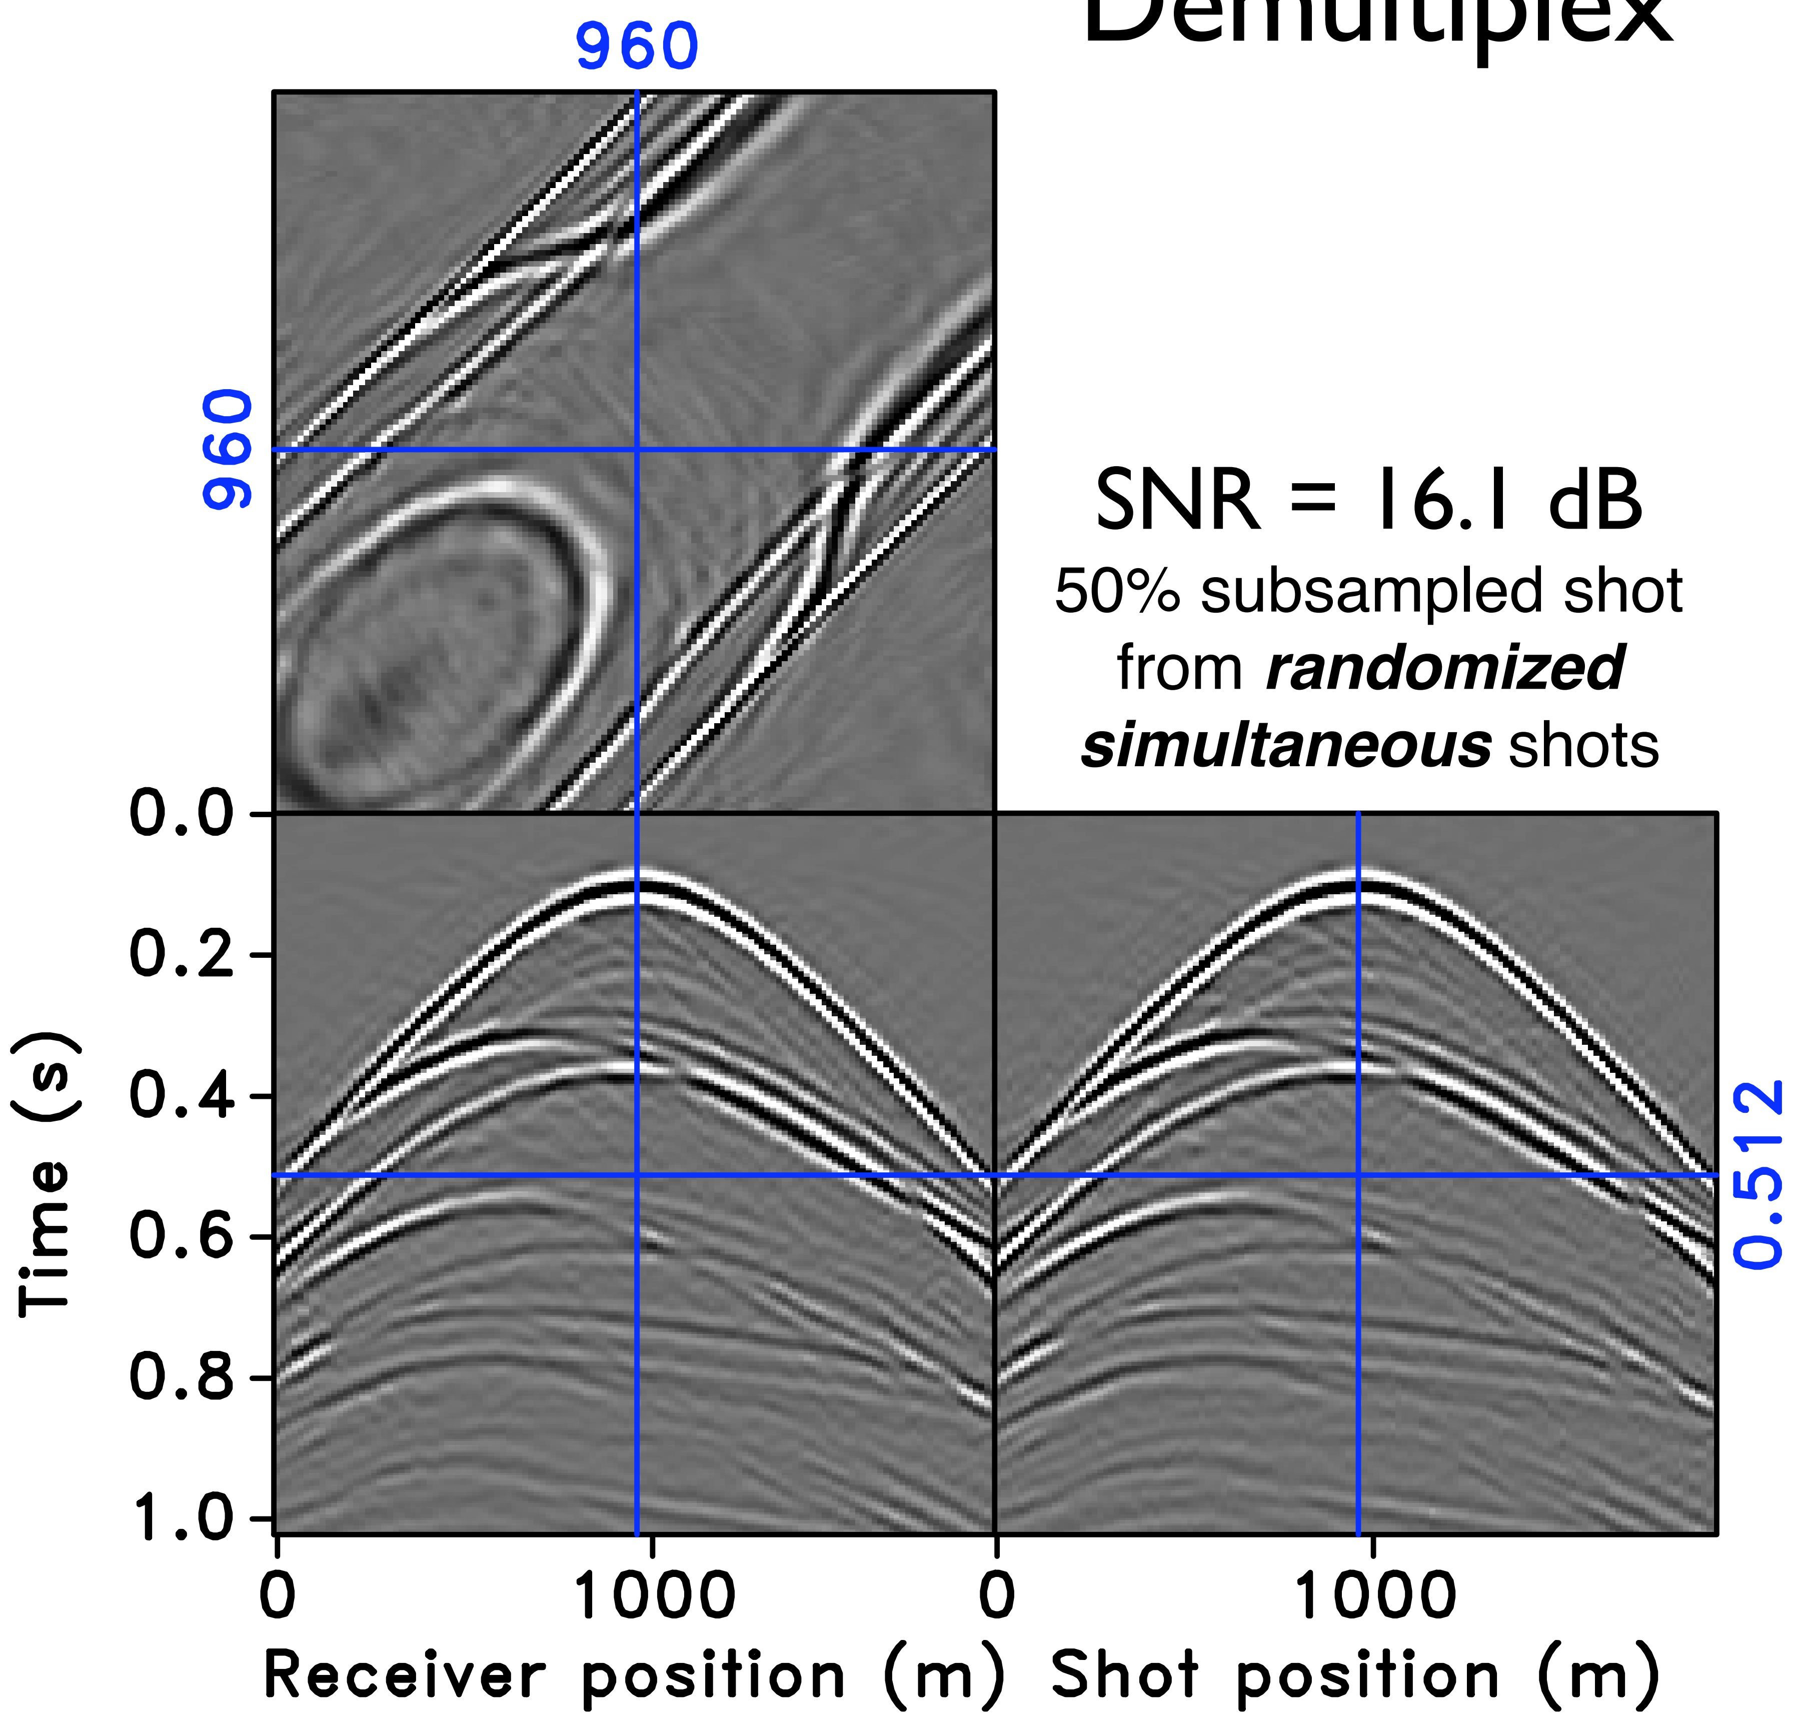

SNR = 16.1 dB

50% subsampled shot from ***randomized simultaneous*** shots

Model

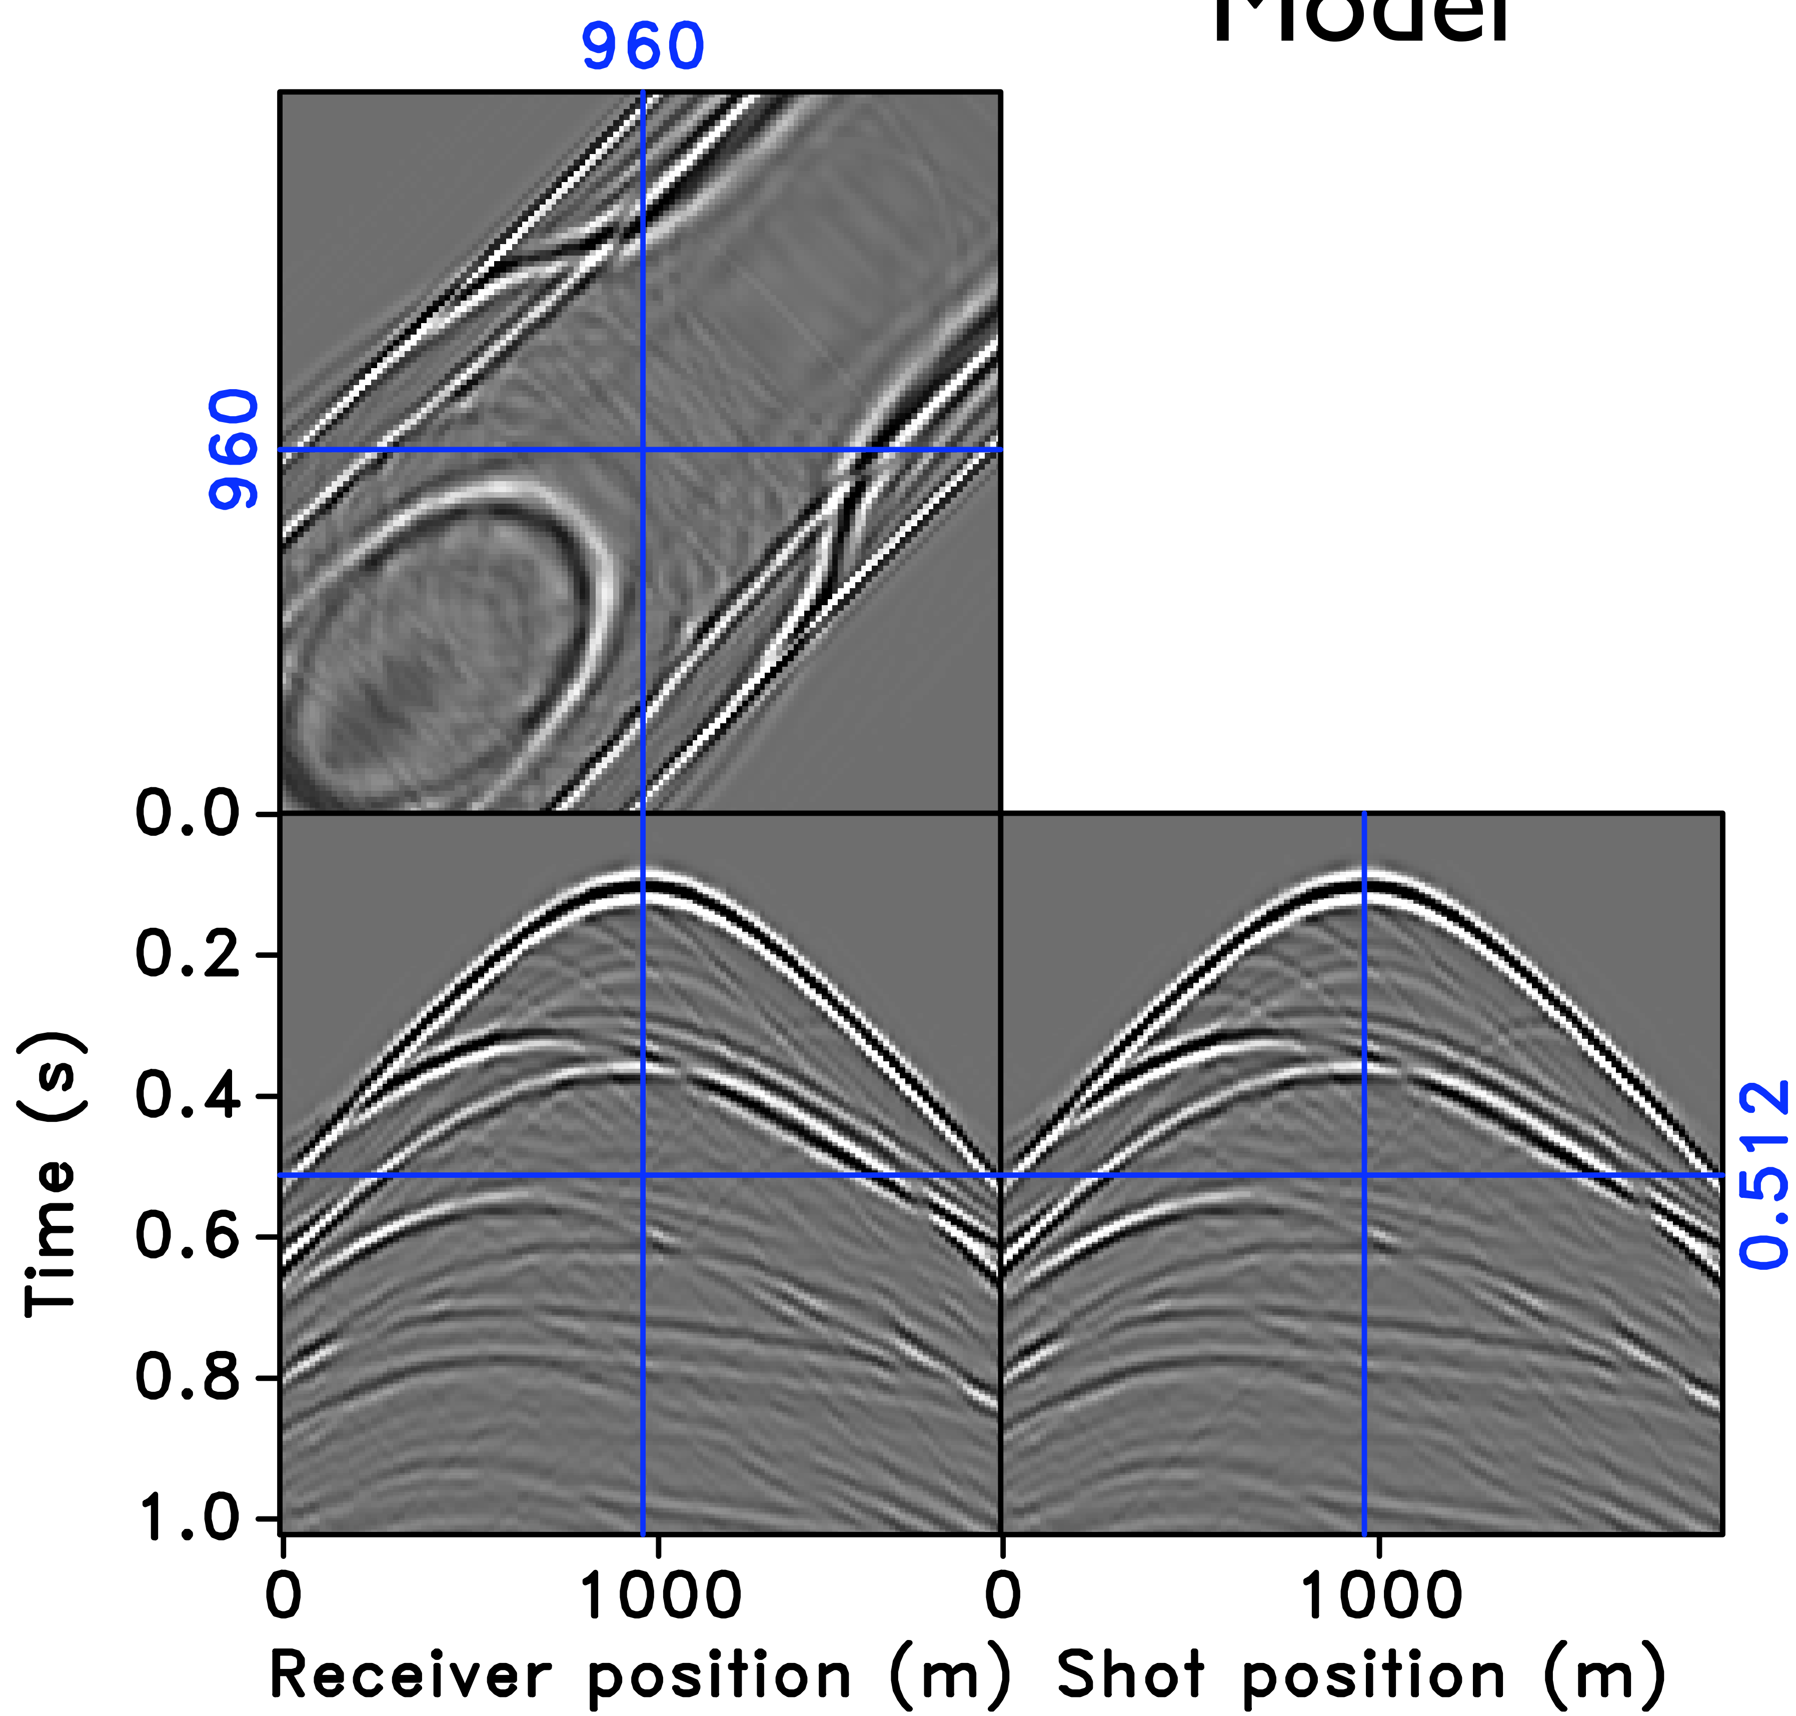

960
960
0.0
0.2
0.4
0.6
0.8
1.0
Time (s)
0.512
0
1000
0
1000
Receiver position (m)
Shot position (m)

Regular vs uniform randomized 2D sampling

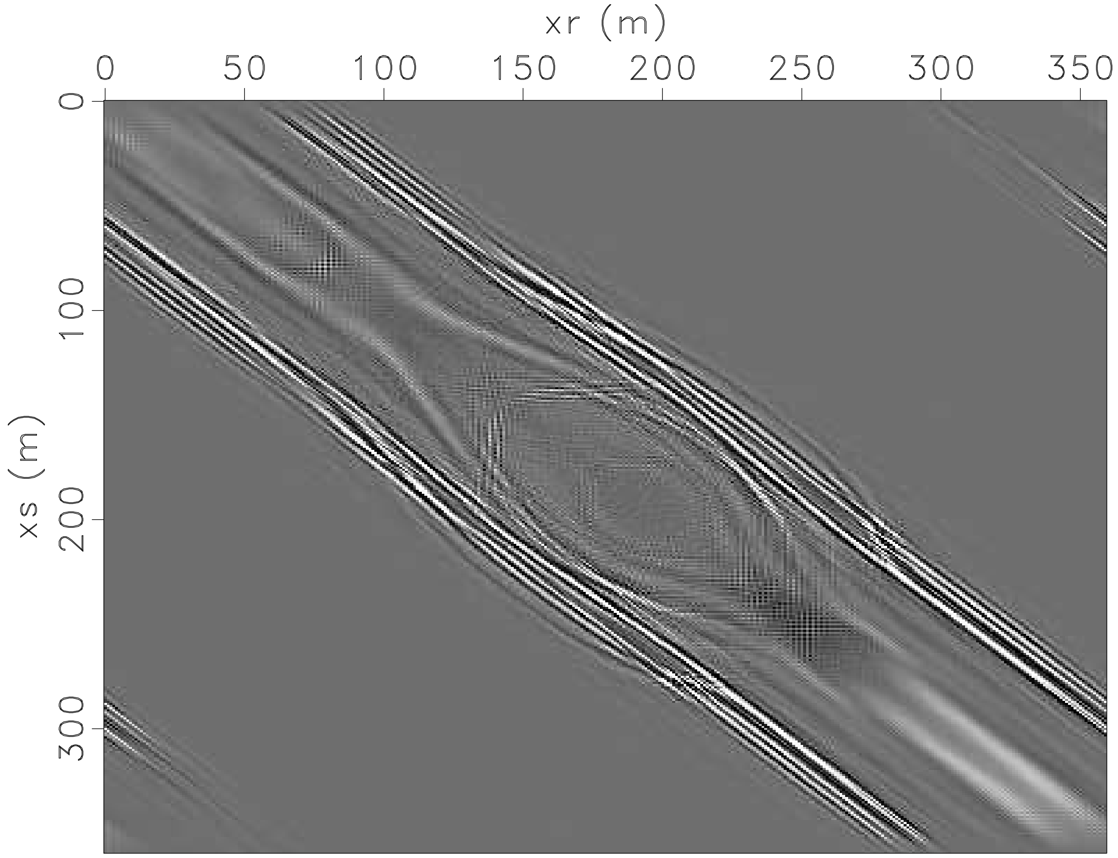

CRSI reconstruction from regular 2-D sampling (25% of data taken)
SNR: 4.161 dB

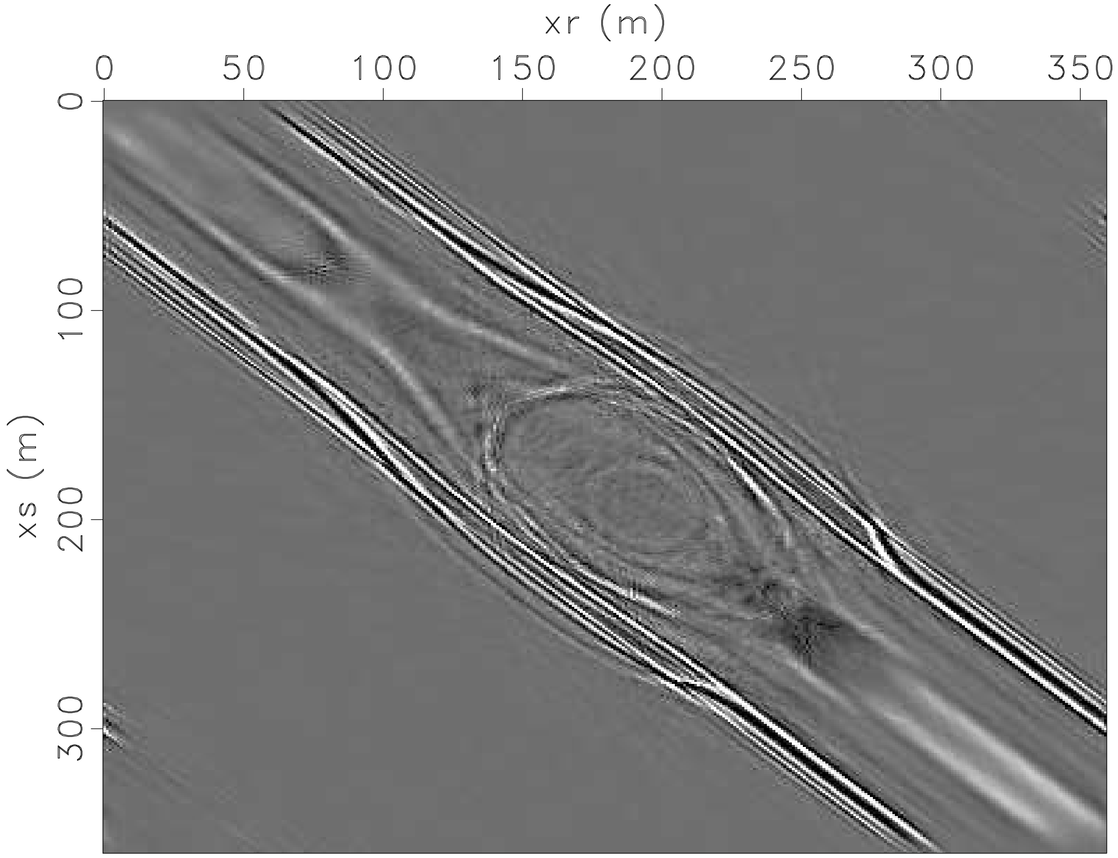

CRSI reconstruction from randomized 2-D sampling (25% of data taken)
SNR: 9.979 dB

Regular vs randomized sampling - residuals

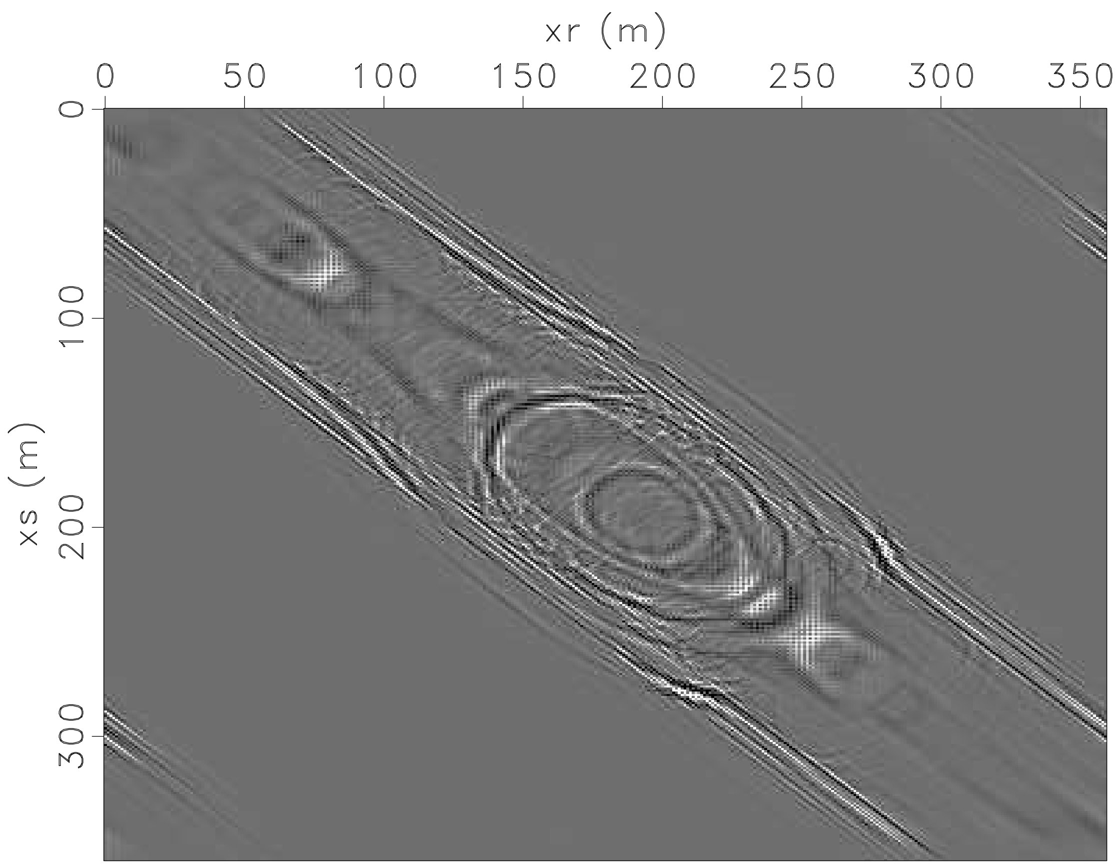

Regular sampling residual

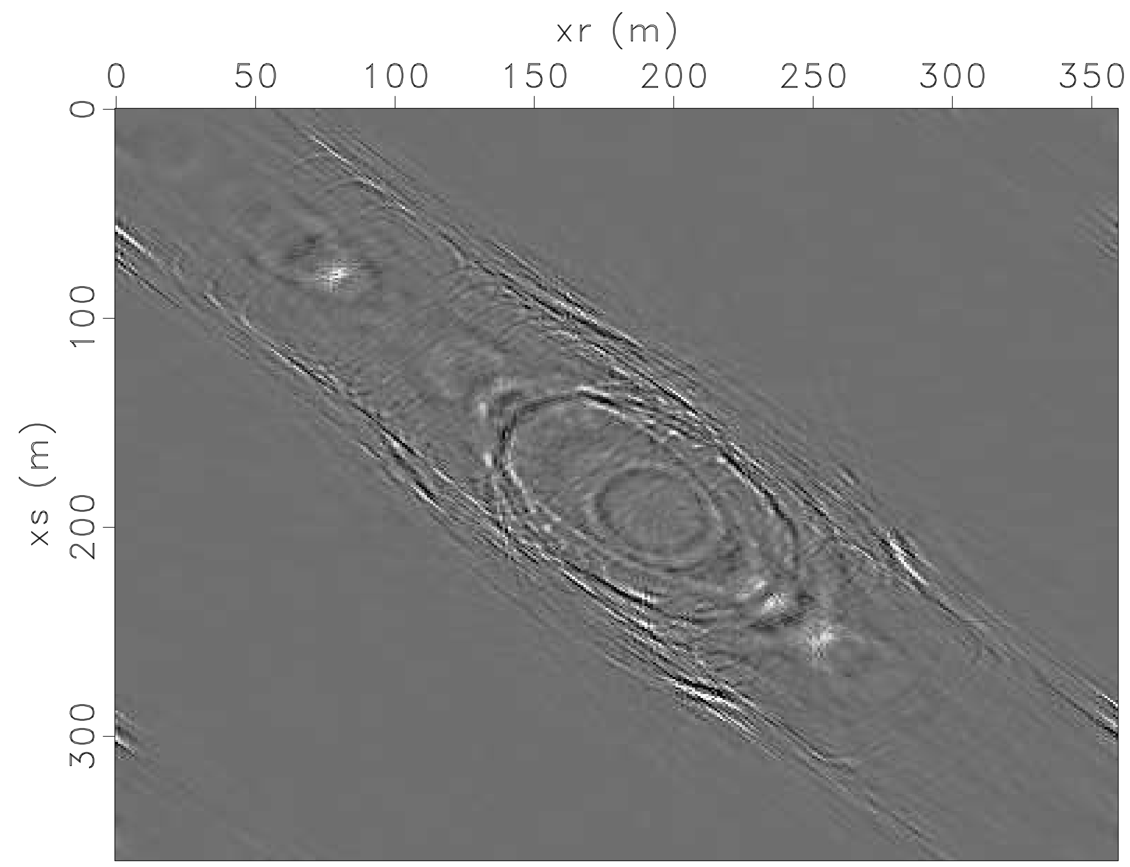

Random sampling residual

Regular vs. irregular sampling - freq. domain

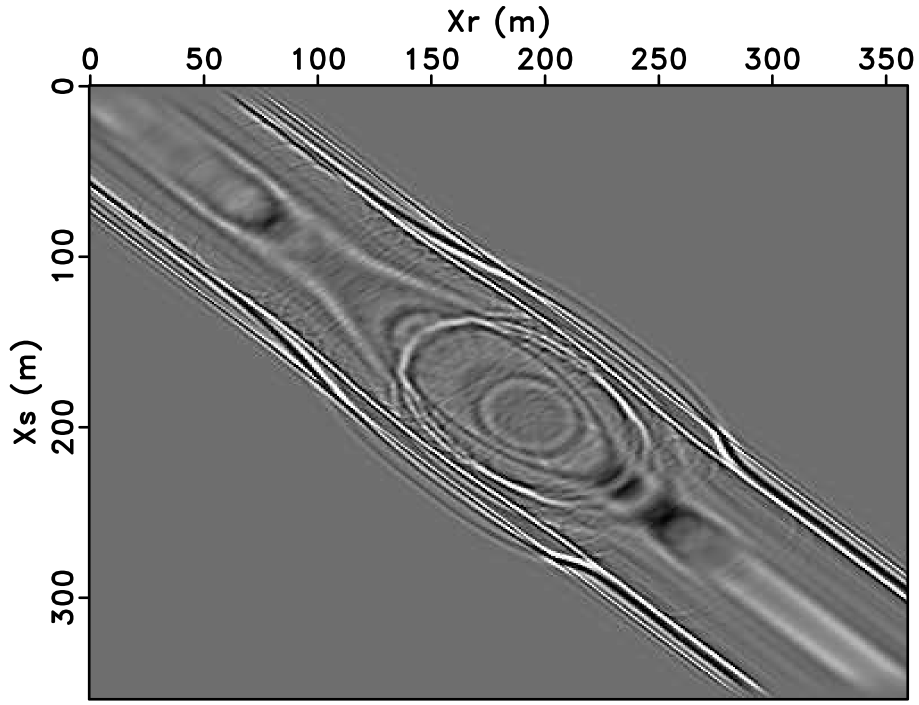

Original model

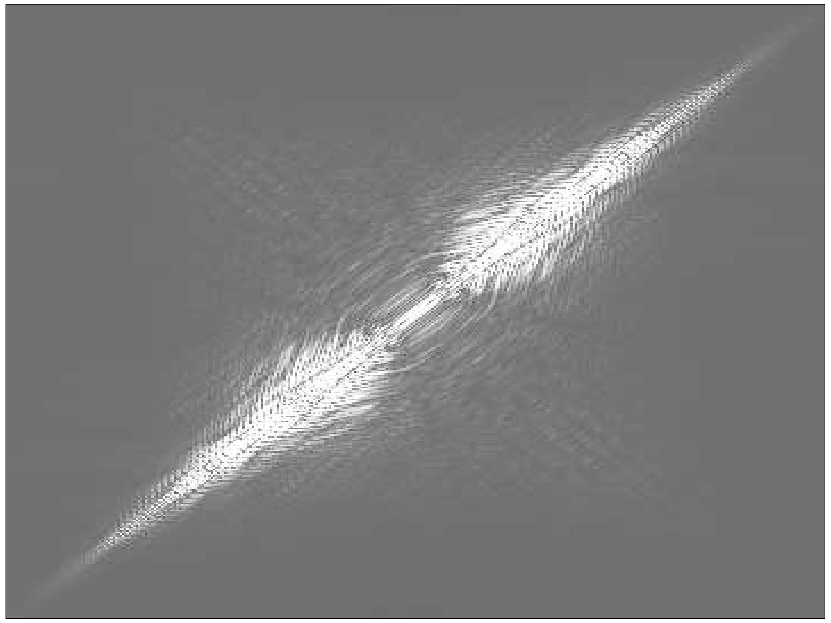

Original model spectrum

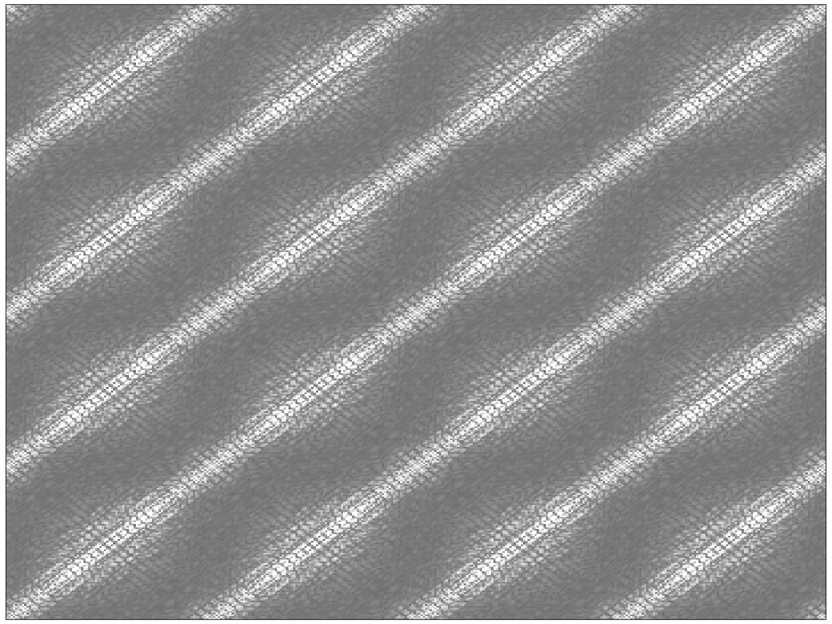

Reg. undersampled spectrum

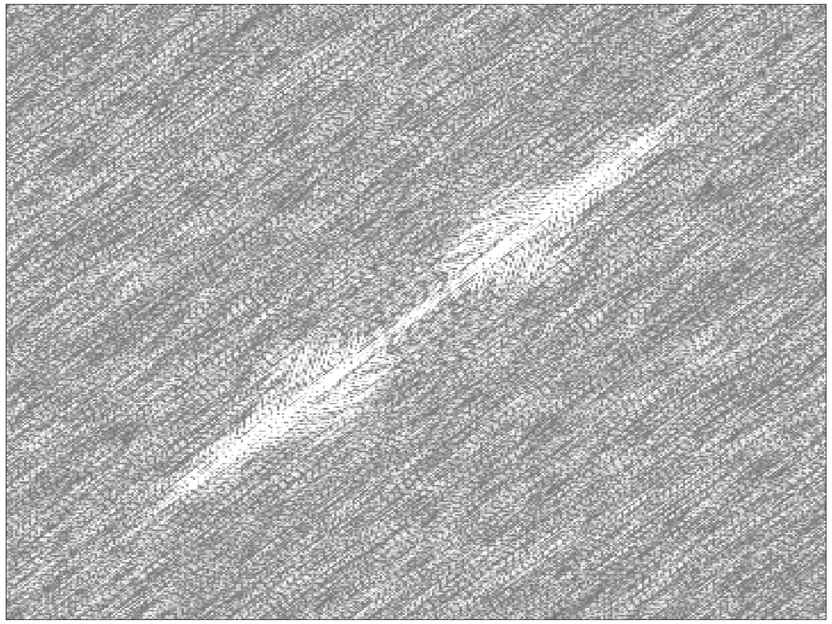

Irreg. undersampled spectrum

2-D discrete random *jittered* sampling

Type	Sampling scheme	Mask's spectra	Samples' spectra
Square	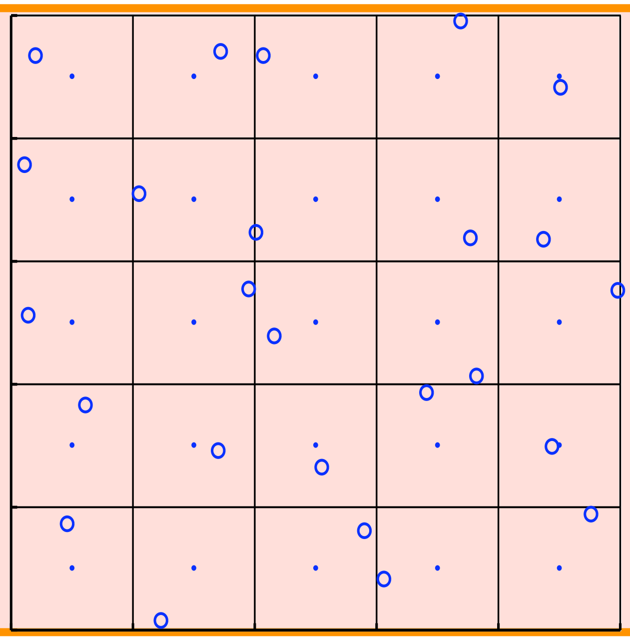	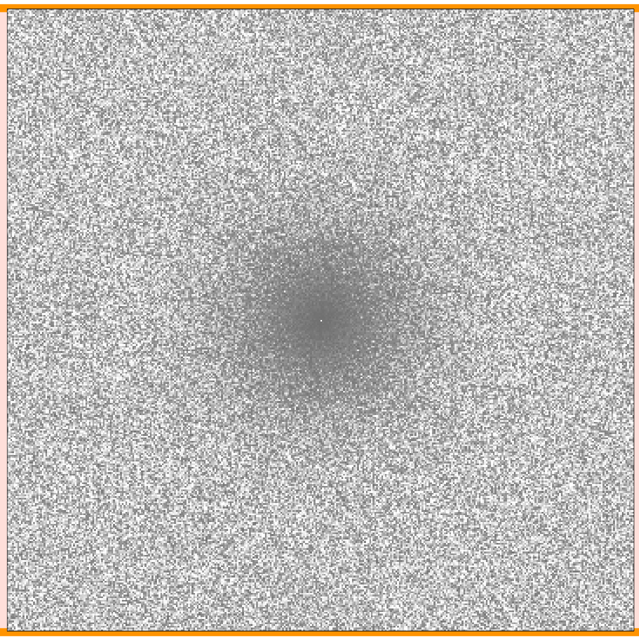	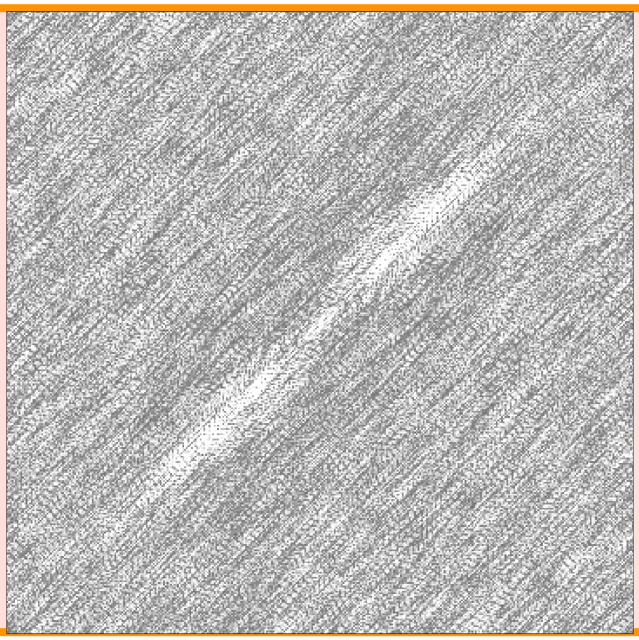
Hexagon	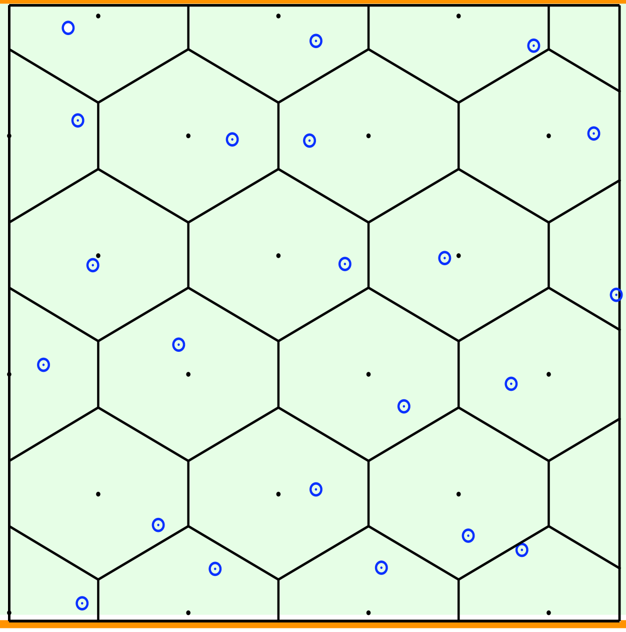	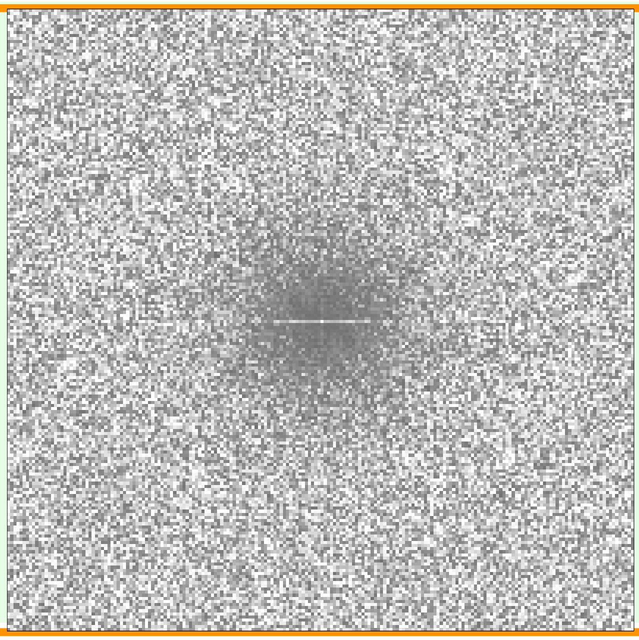	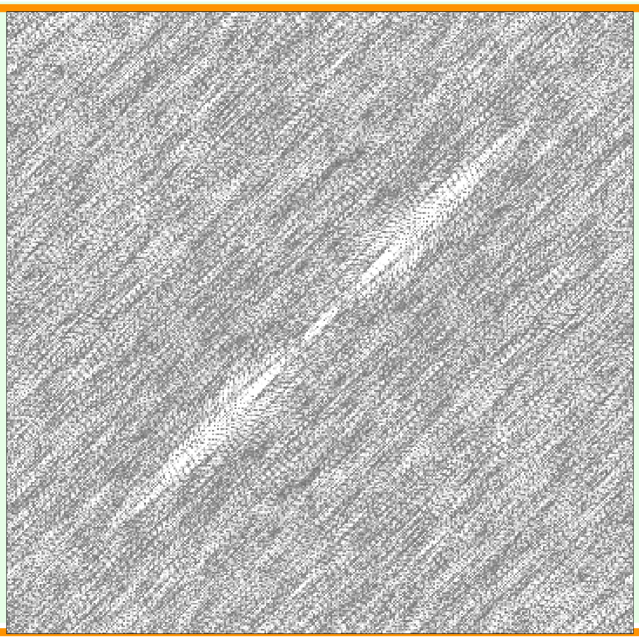

2-D discrete random jittered subsampling

- Cartesian & hexagonal jittered reconstructions almost the same.

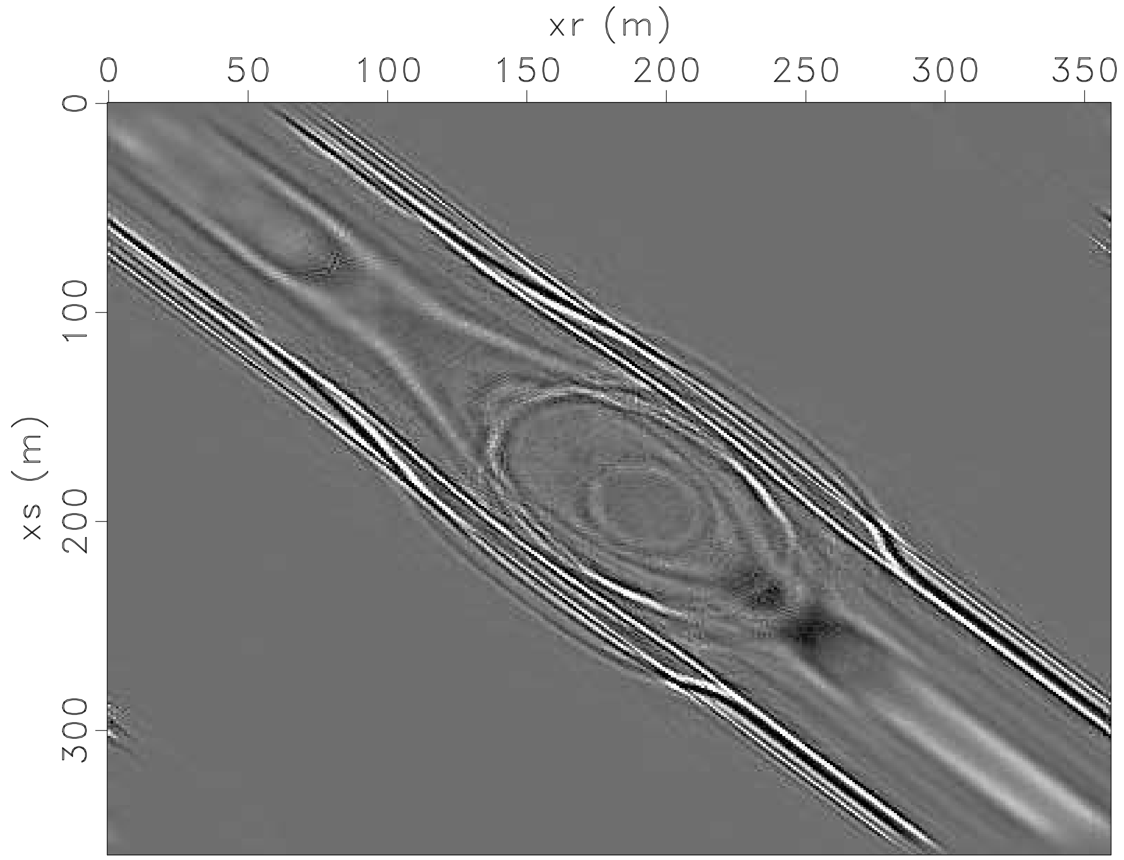

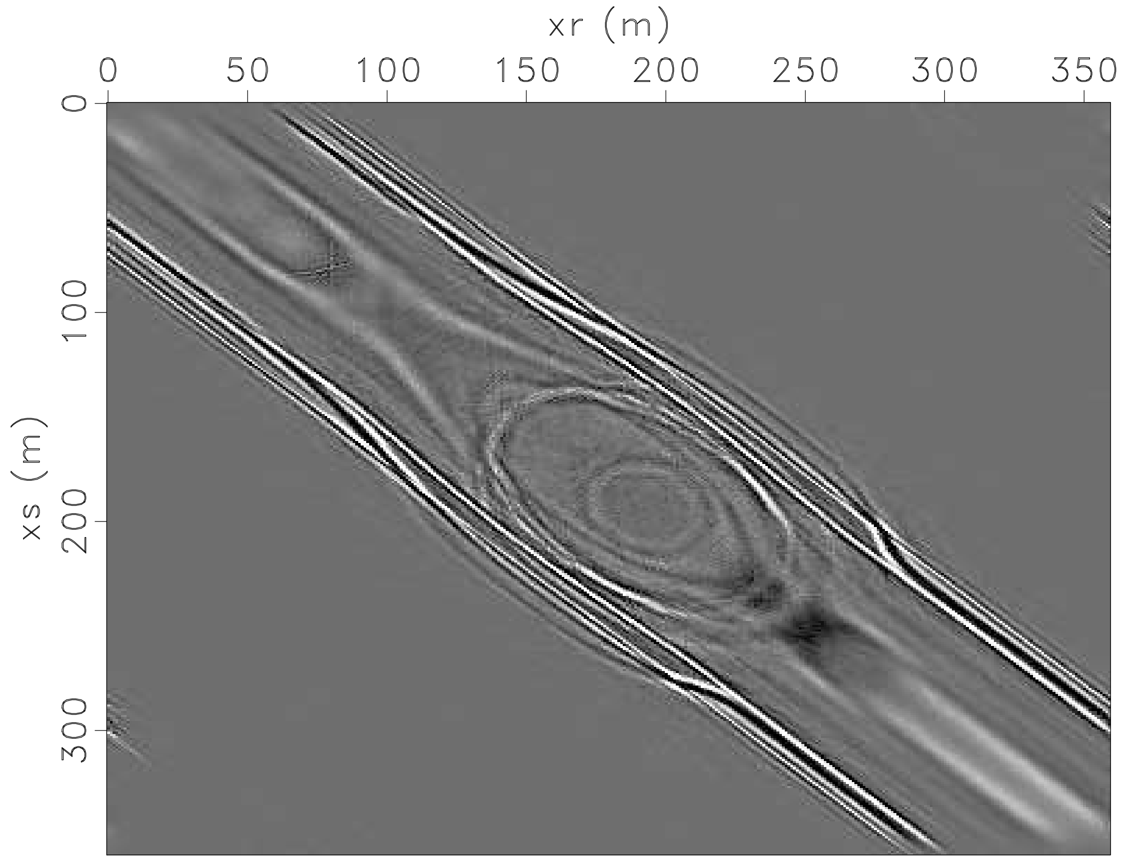

CRSI Recovery (Cartesian)
SNR = 10.820

CRSI Recovery (Hexagonal)
SNR = 10.904

2-D discrete jittered subsampling

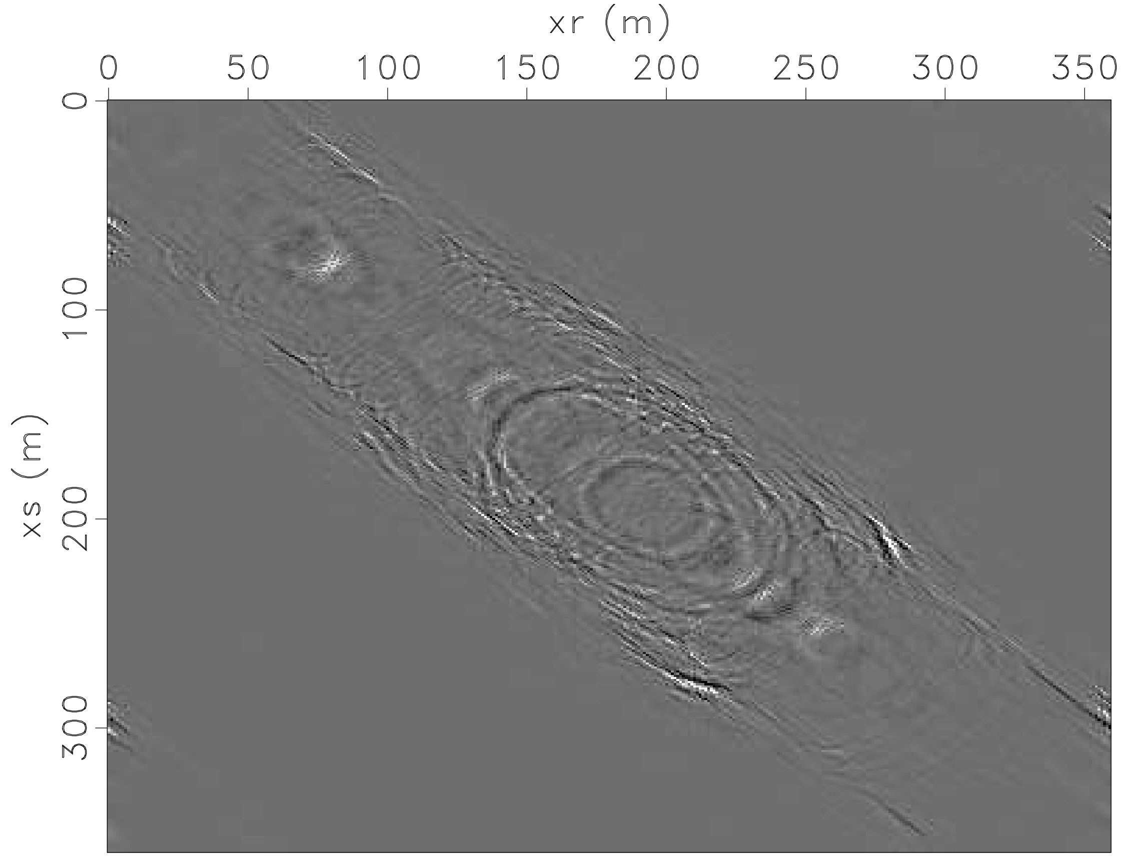

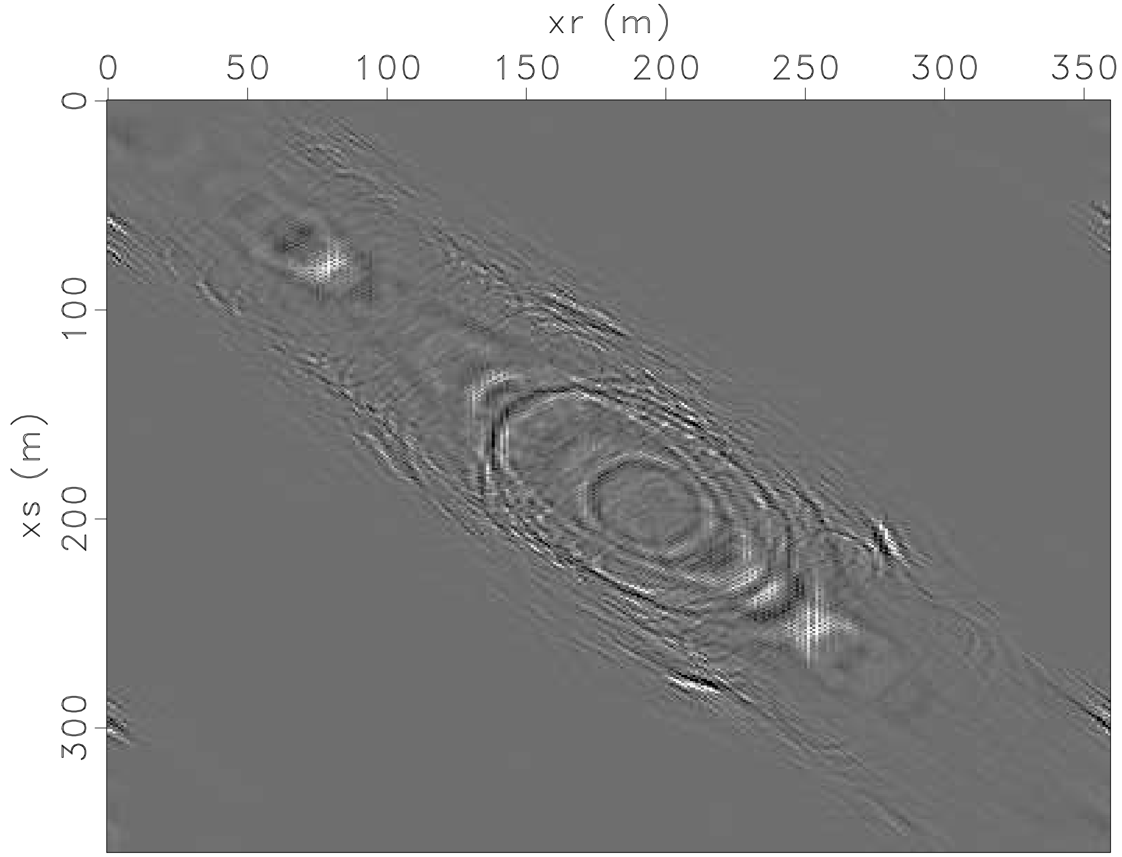

CRSI Residual (Cartesian)
SNR = 10.820

CRSI Residual (Hexagonal)
SNR = 10.904

Blue-noise spectra from 2D sampling methods

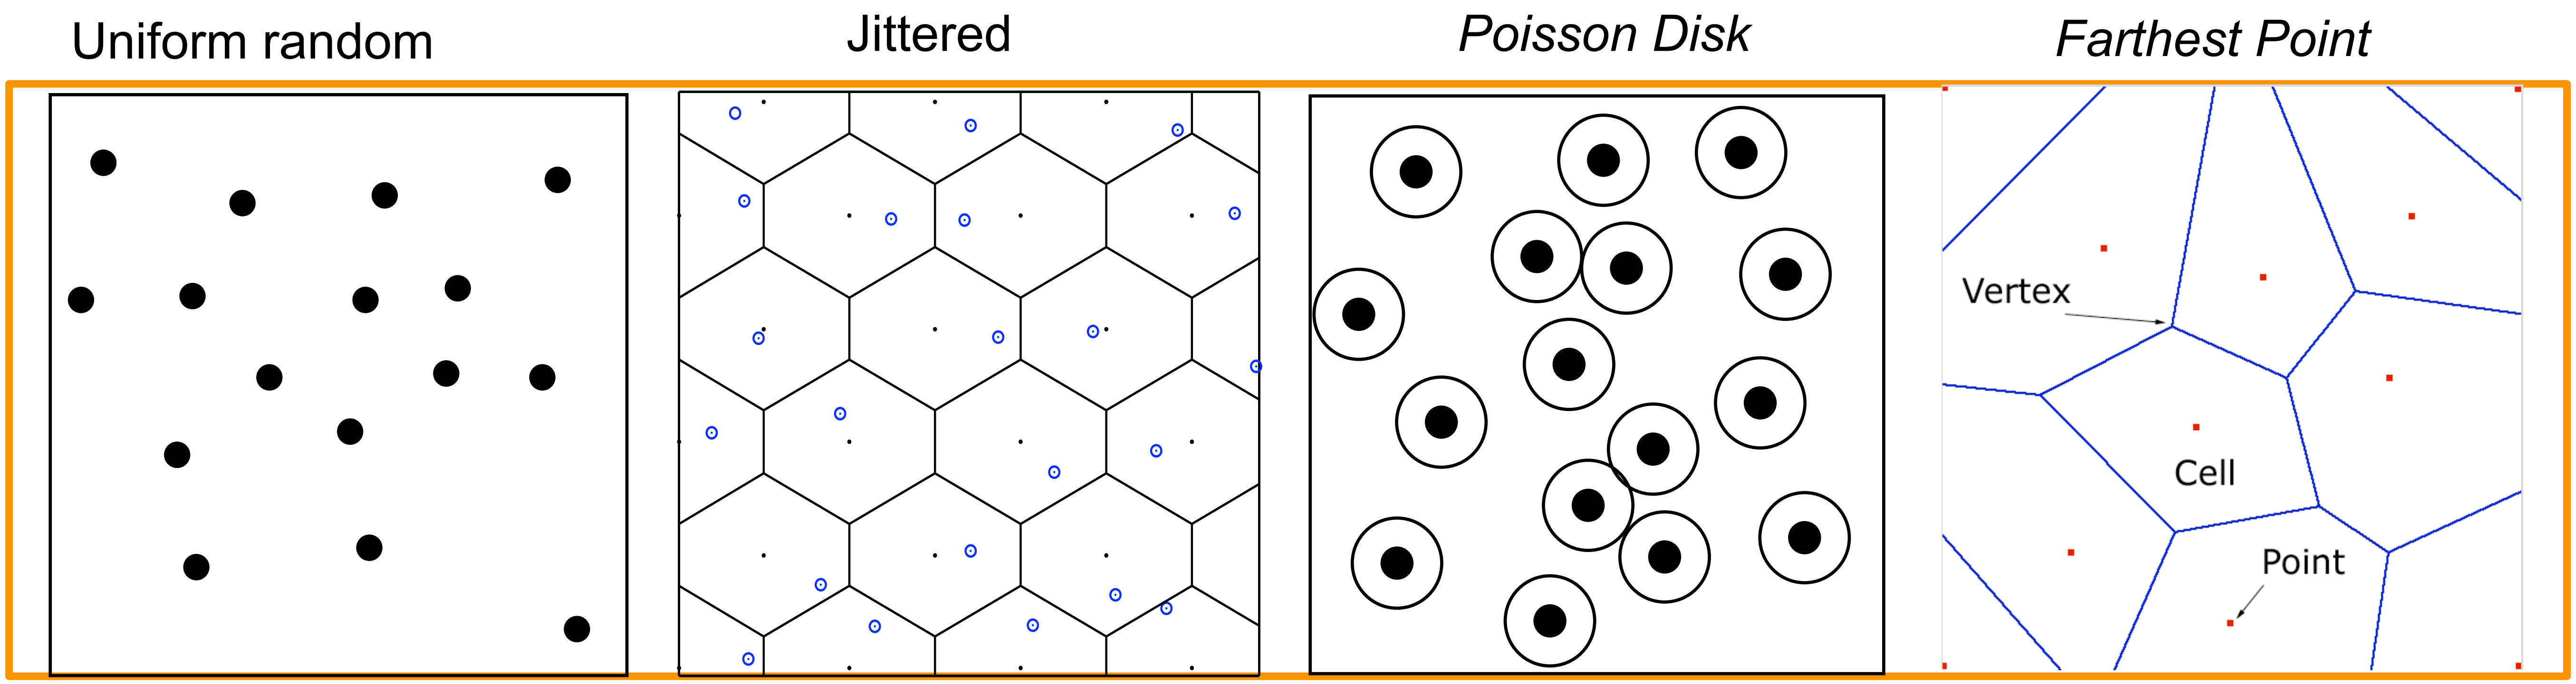

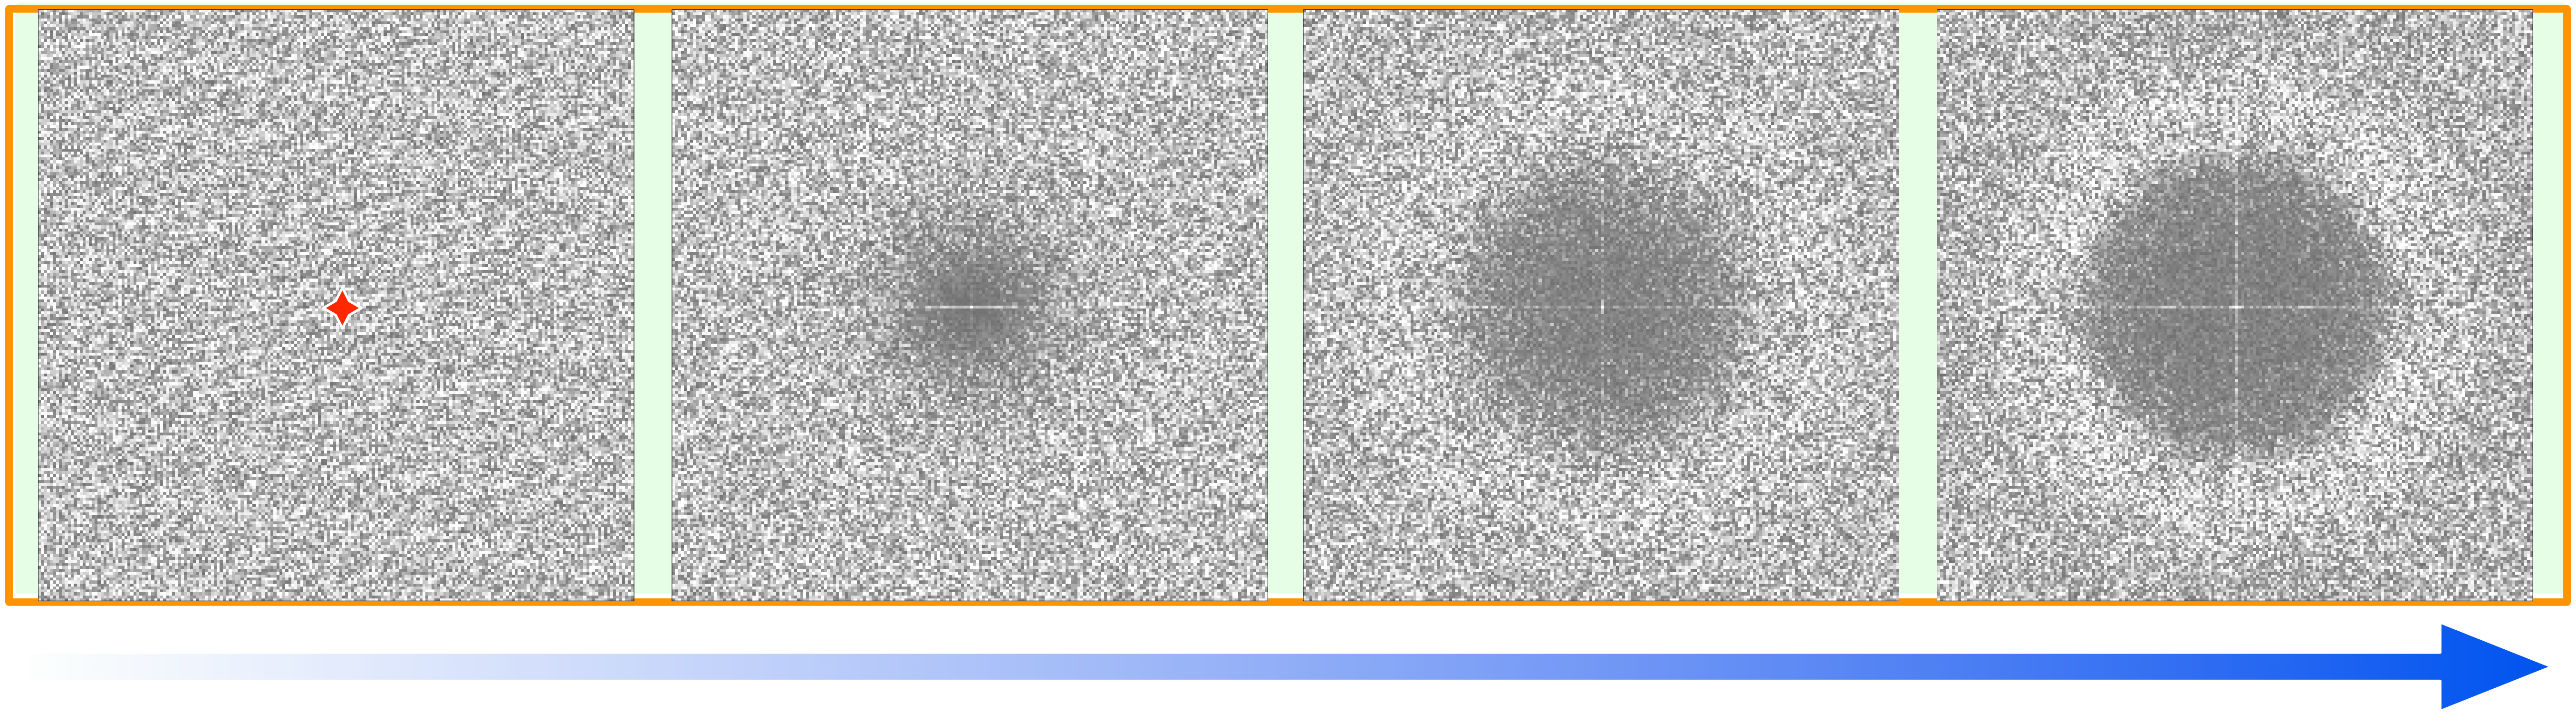

Spectra become increasingly *"blue"*

Randomized 2D uniform vs jittered

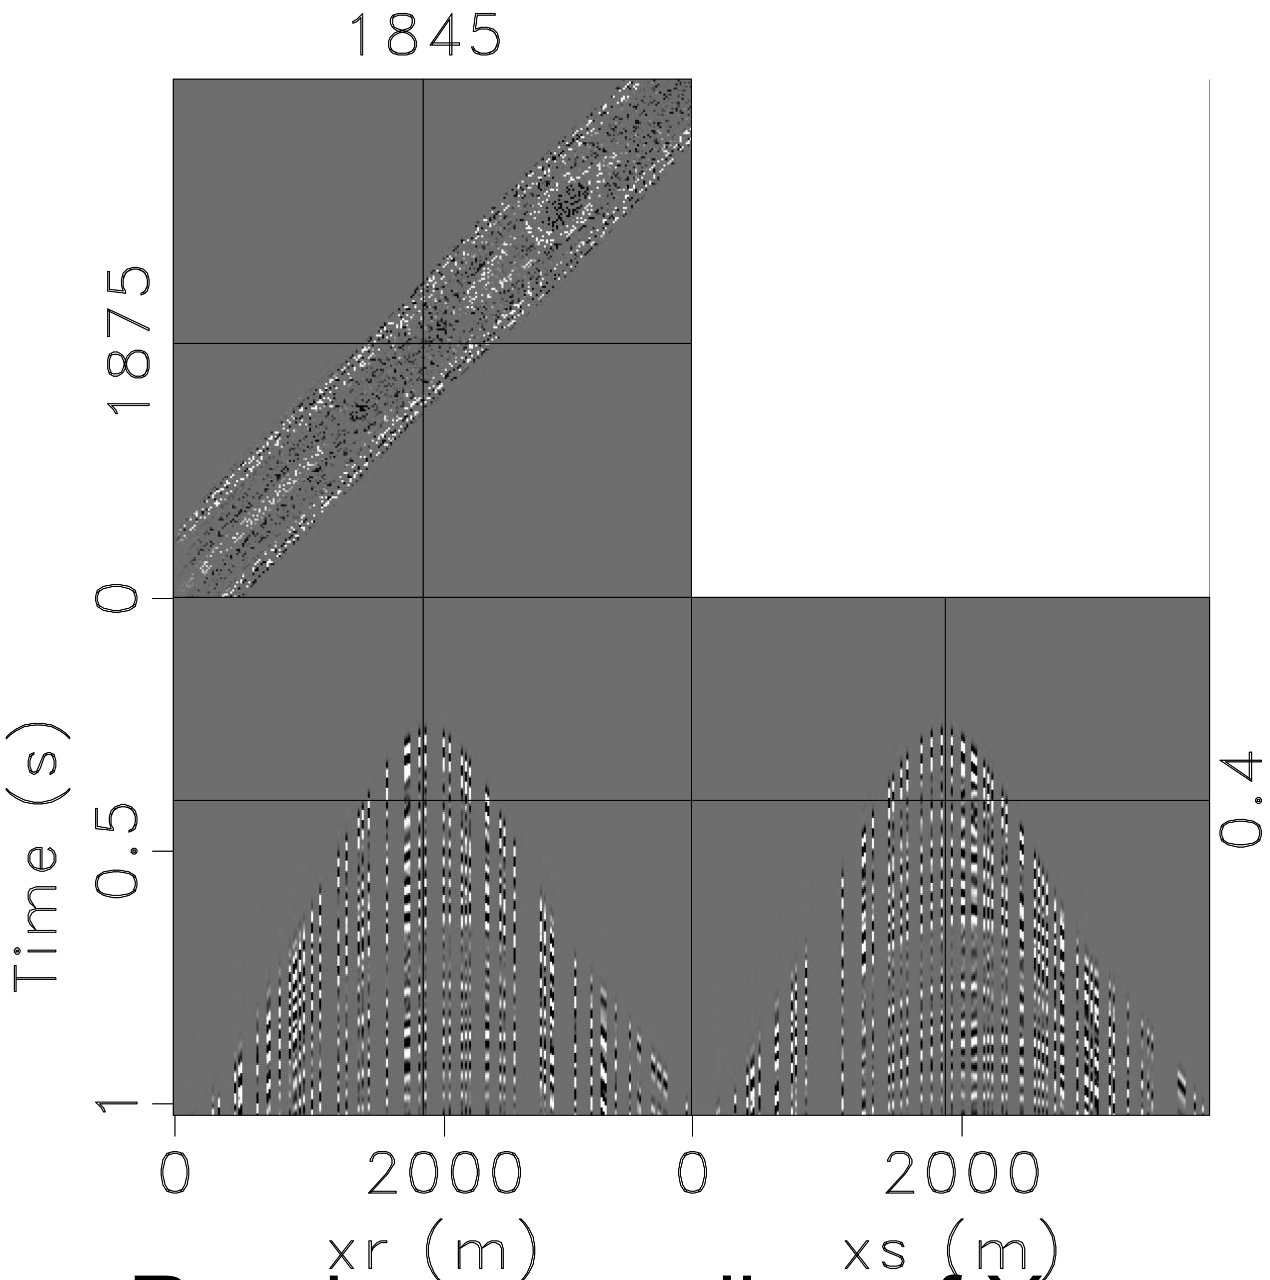

Random sampling of X-spread data (25% total data sampled)

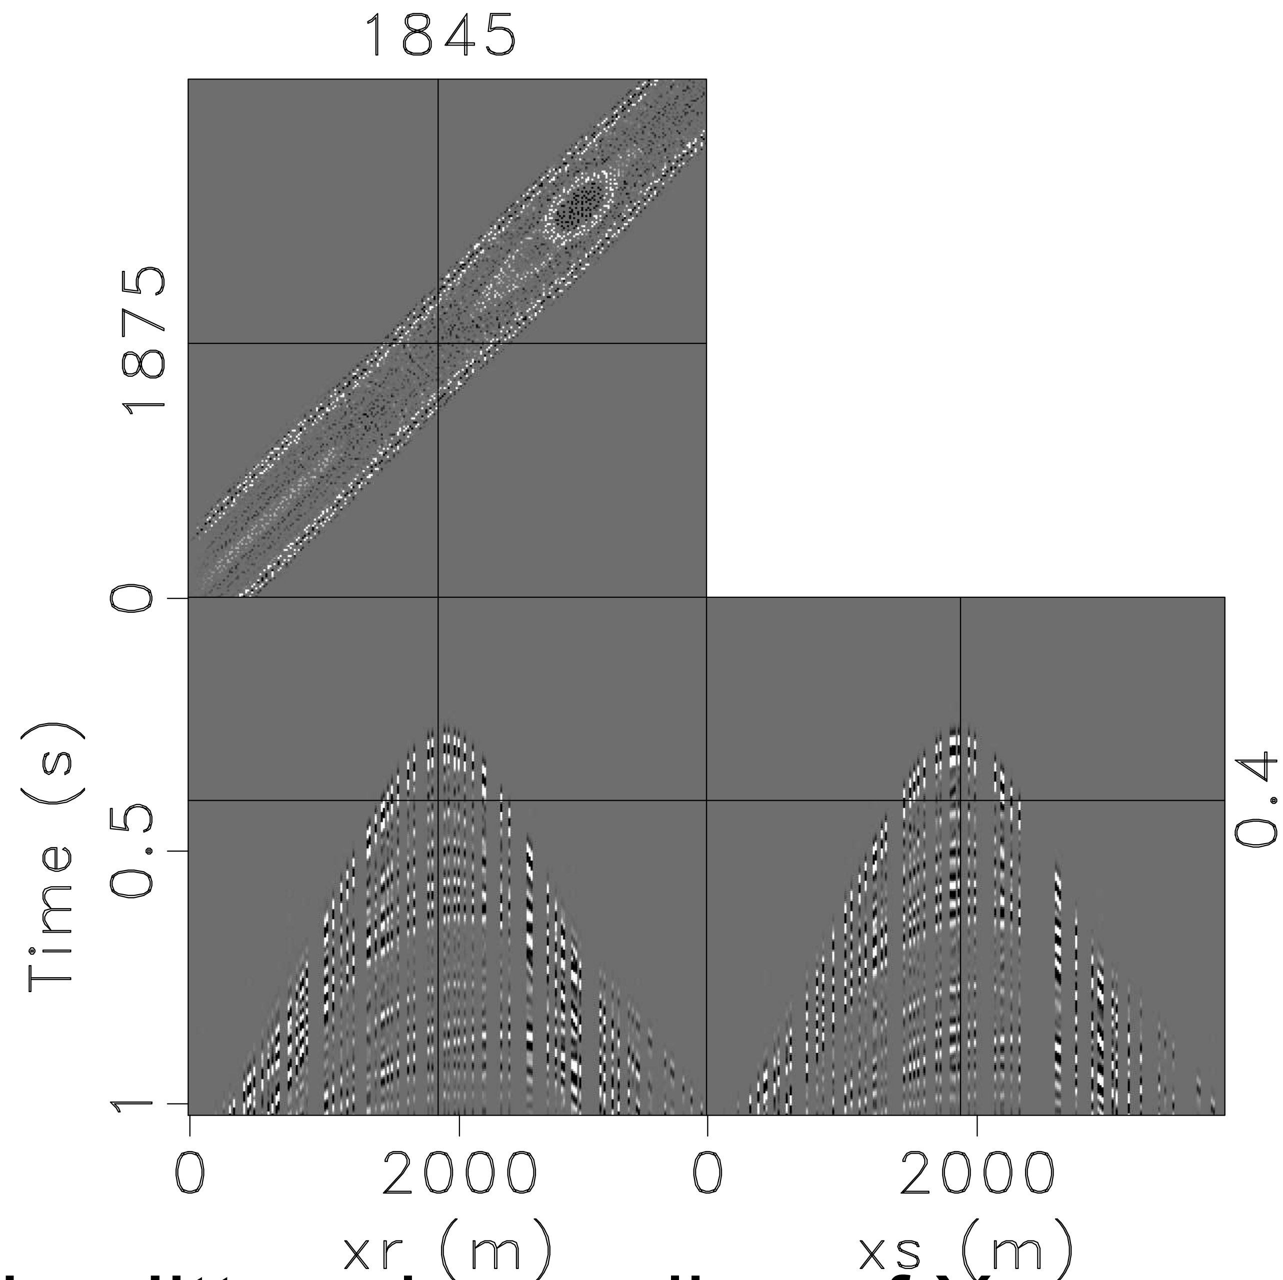

Jittered sampling of X-spread data (25% sampled)

Randomized 2D uniform vs jittered - reconstruction

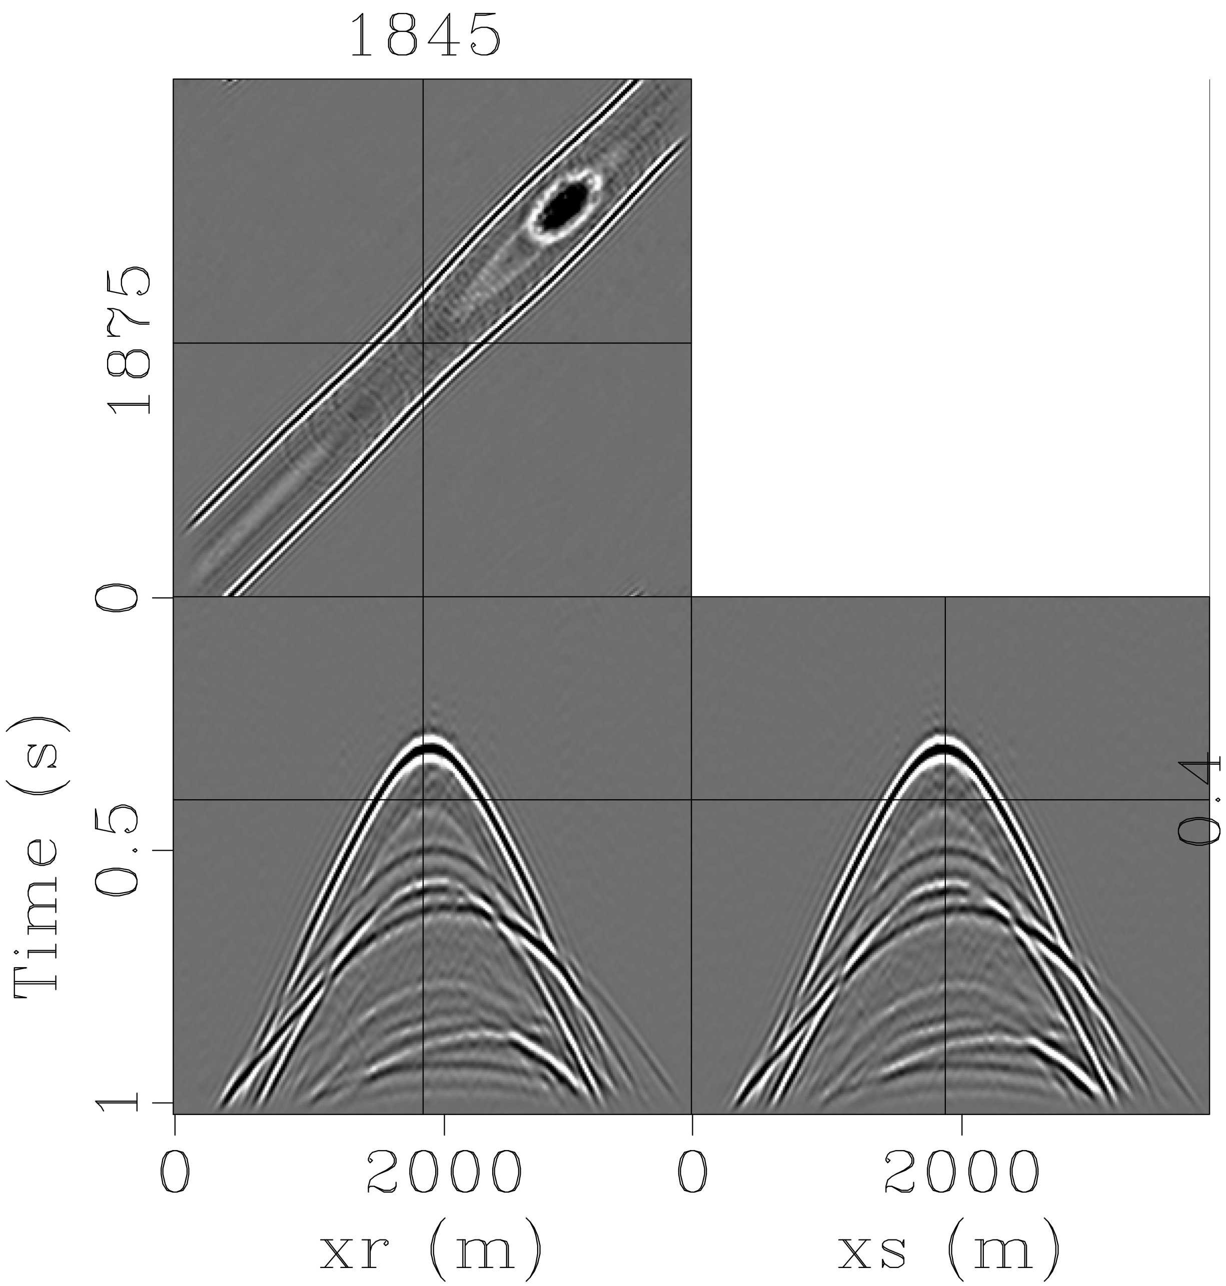

CRSI reconstruction from uniform samples
SNR=8.134

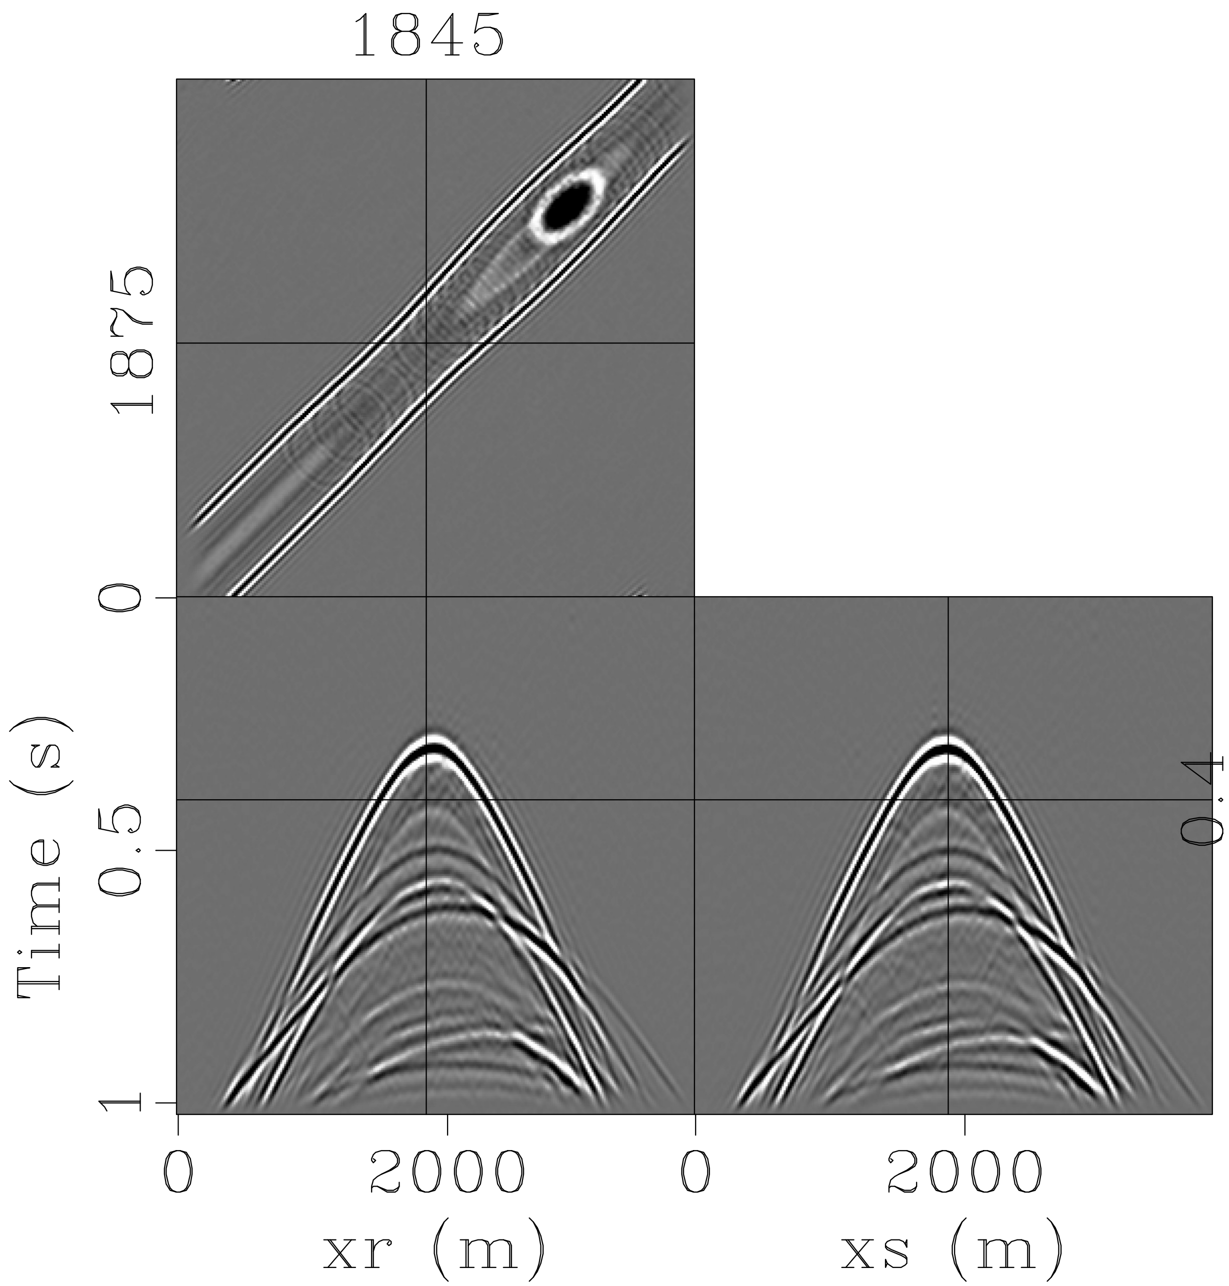

CRSI recon., 2D jittered (hexagonal) samples
SNR=8.434

Randomized 2D uniform vs jittered - residues

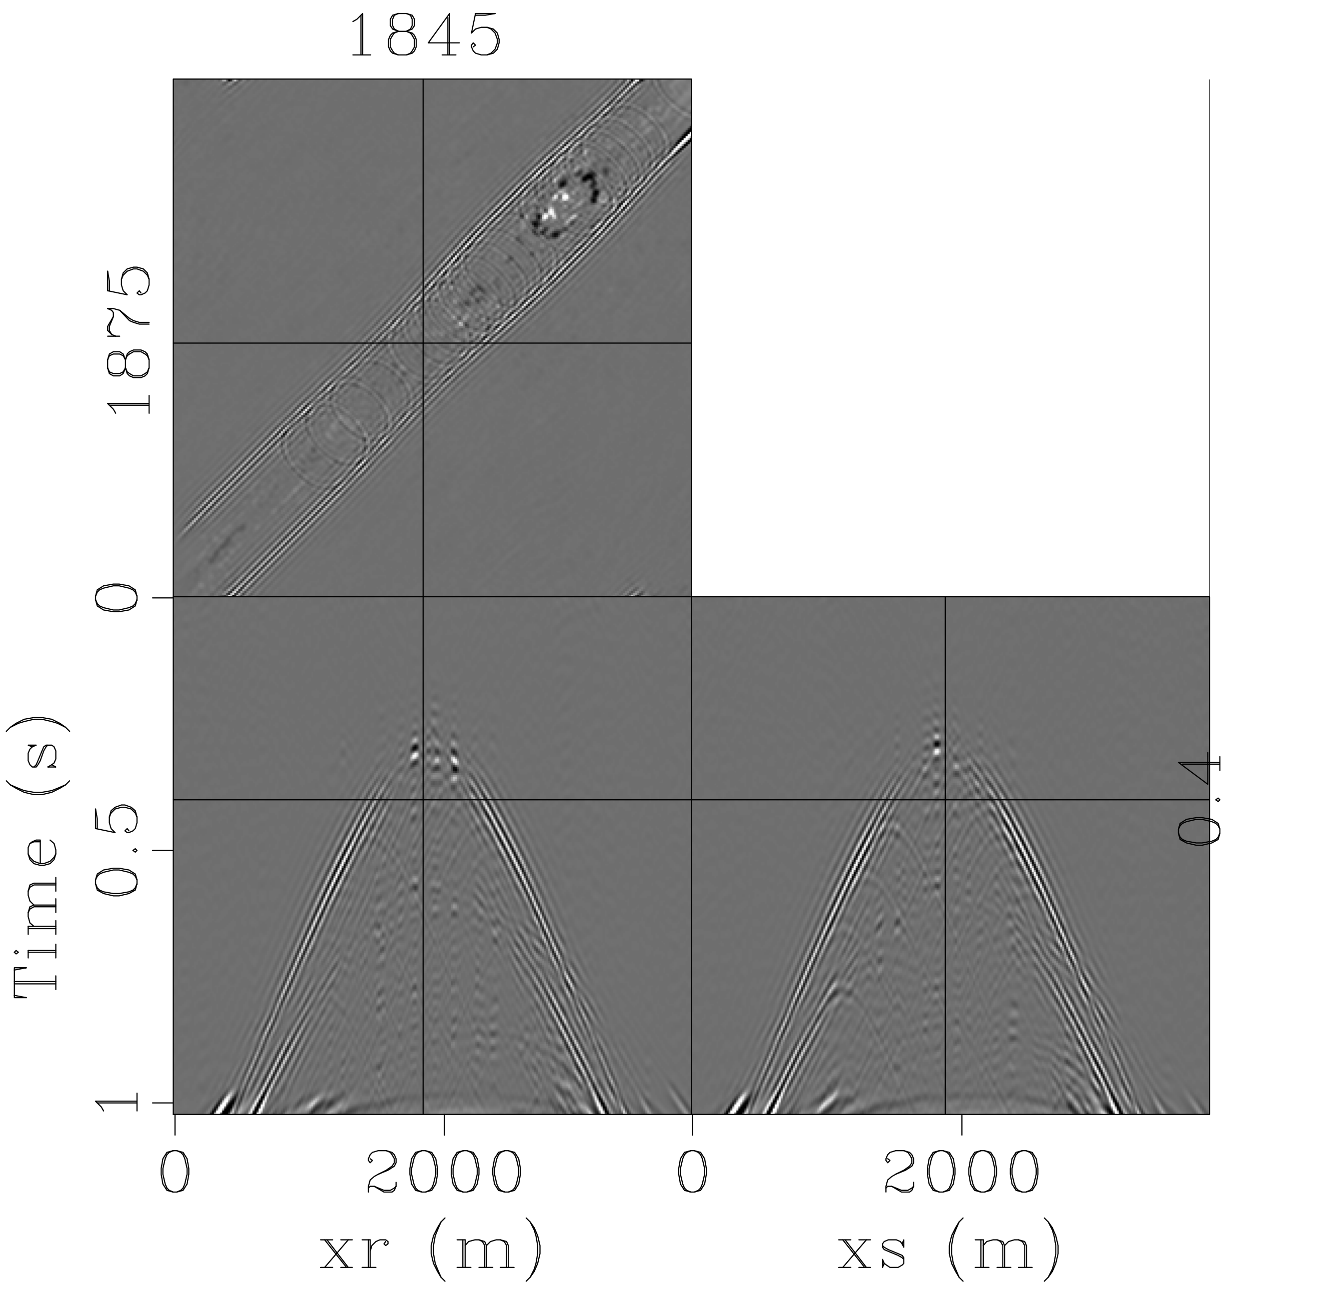

CRSI recon. residual from random samples

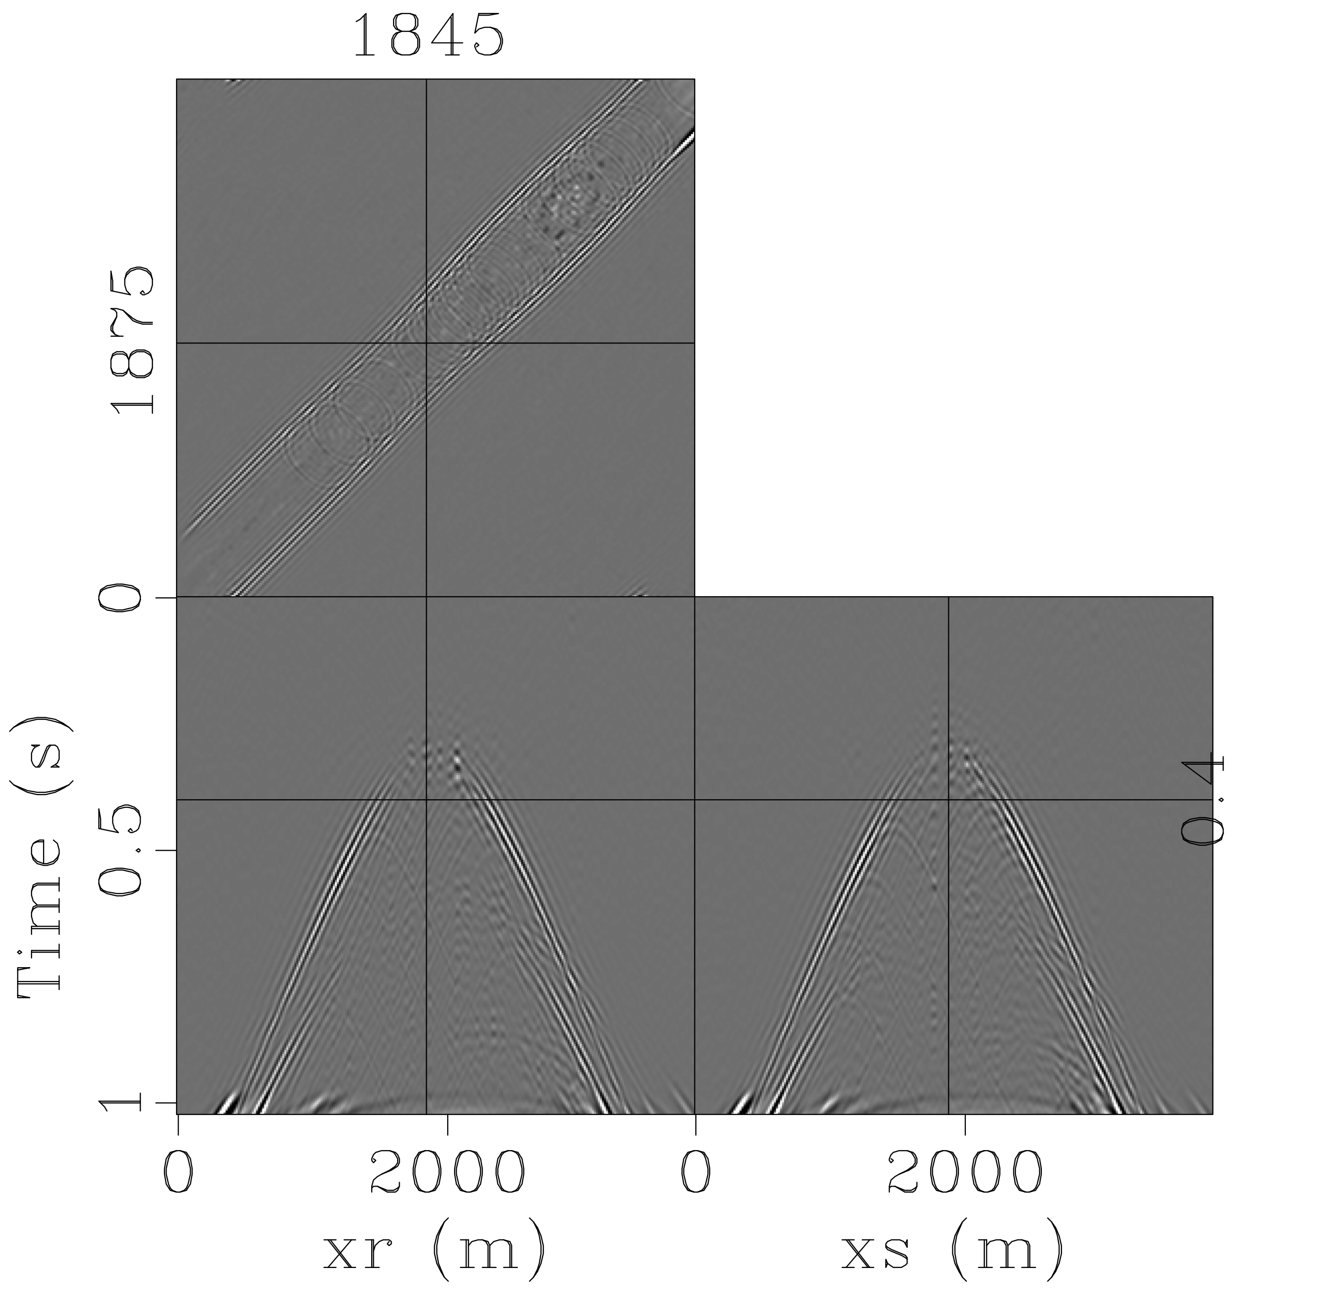

CRSI recon. residual from jittered (hexagonal) samples

Farthest point vs Poisson disk - reconstruction

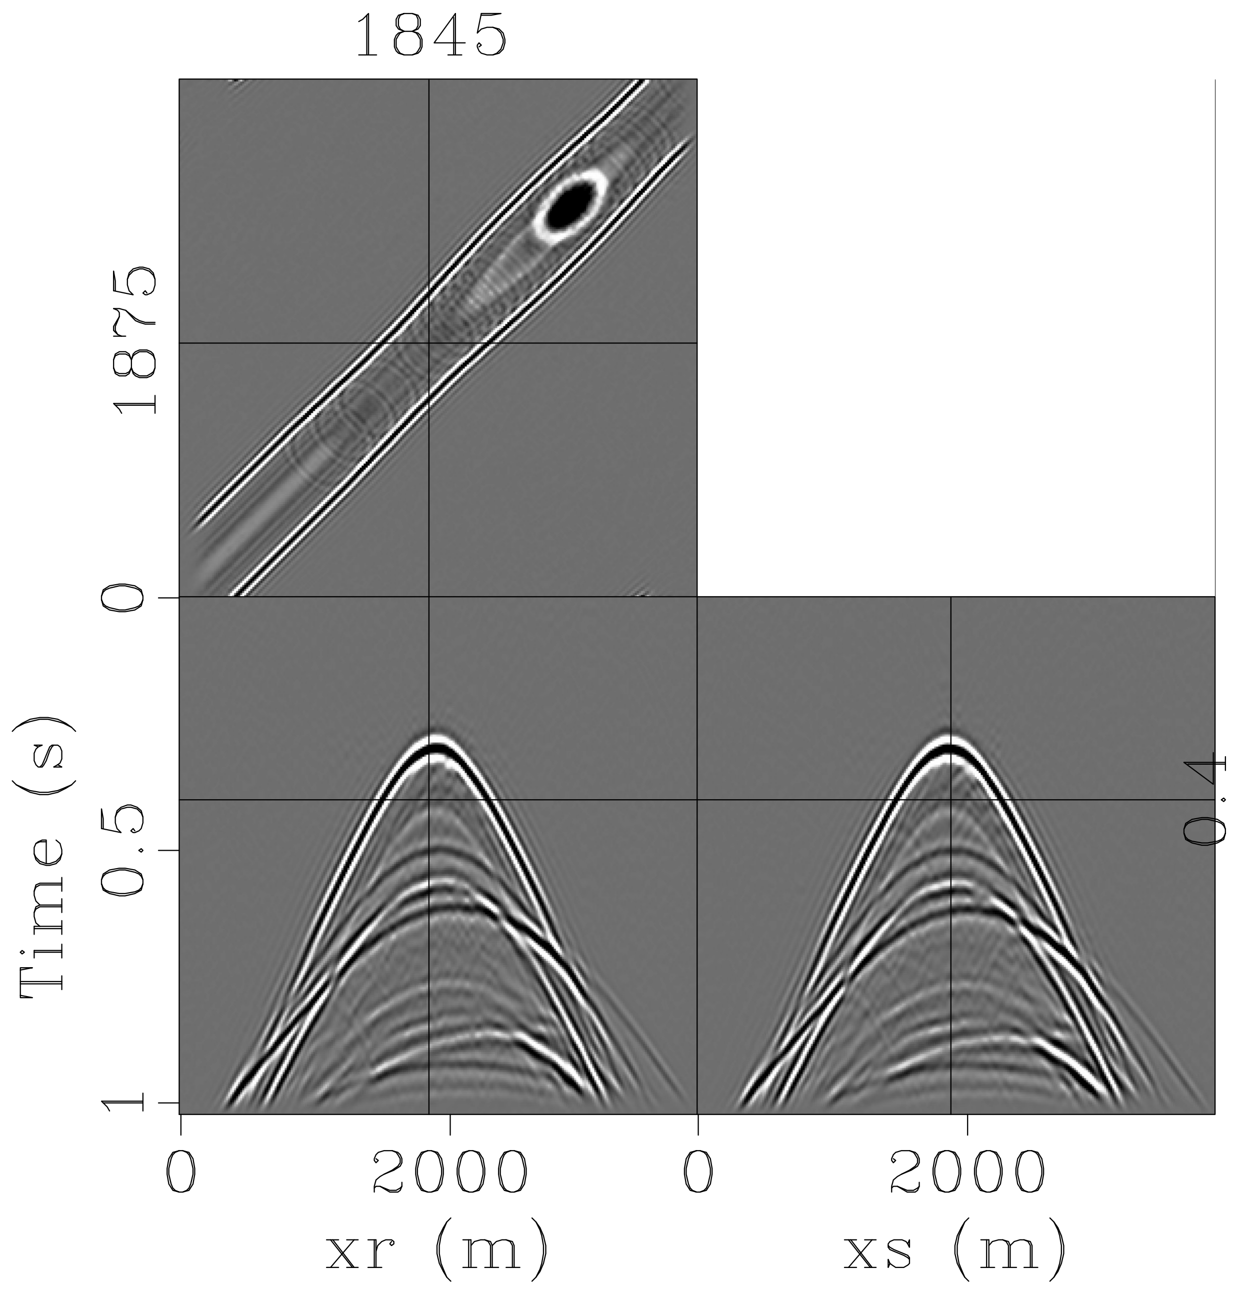

CRSI reconstruction from Farthest Point samples, SNR=8.496

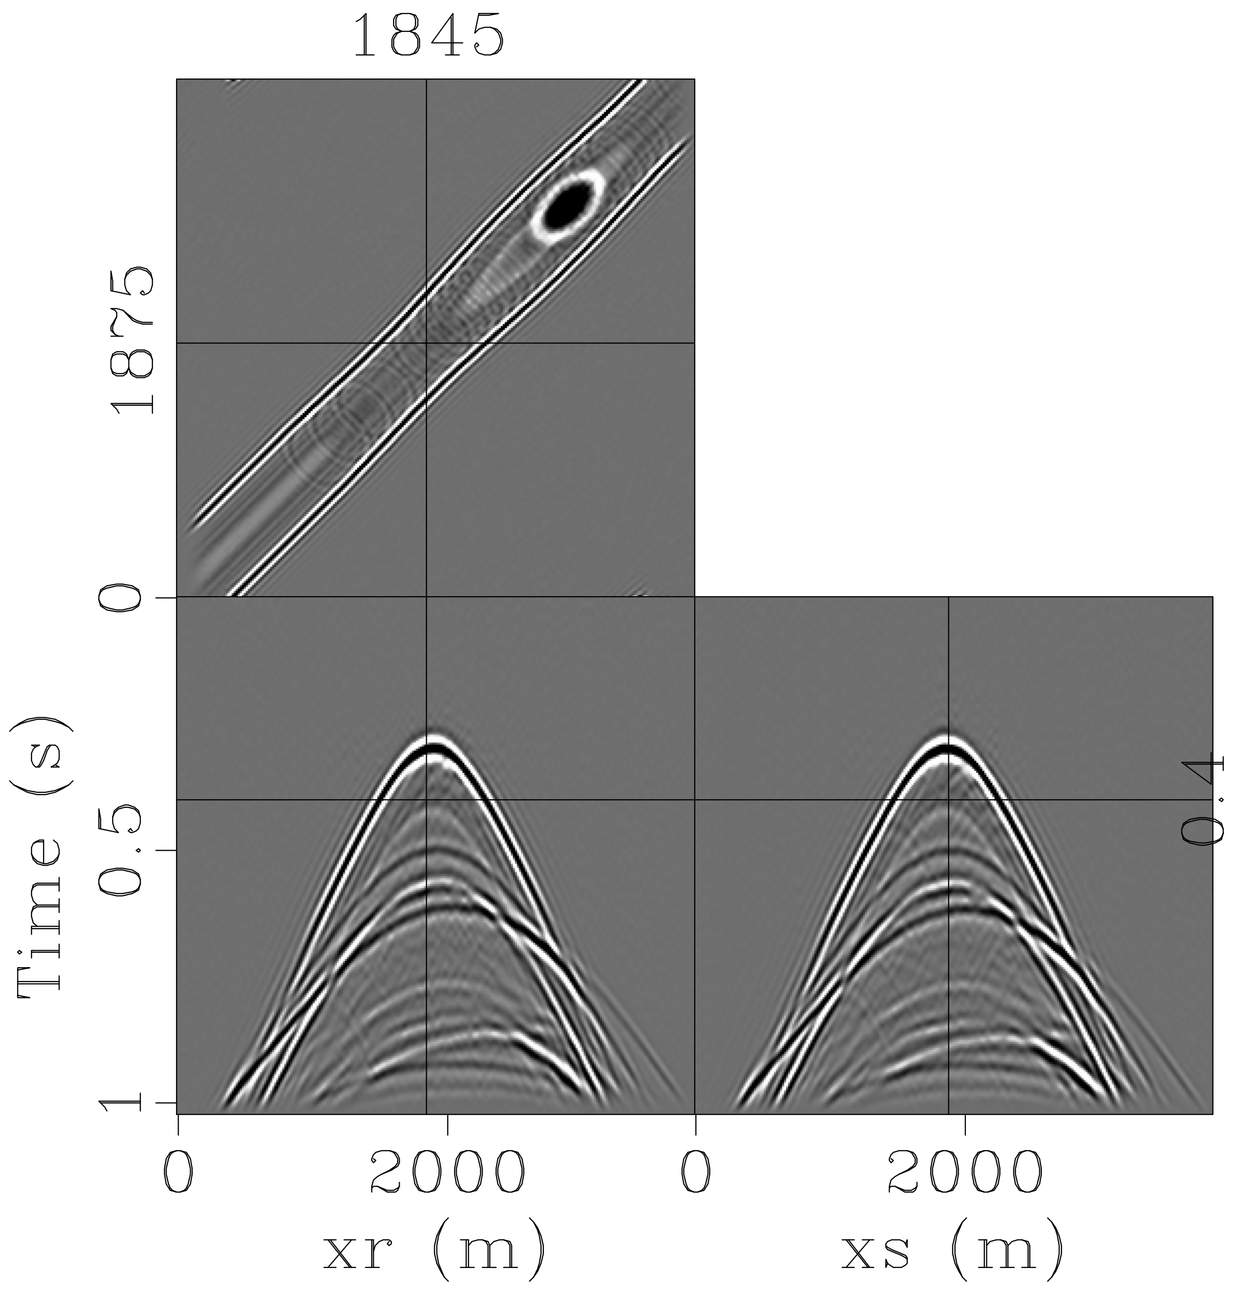

CRSI recon. from Poisson Disk samples SNR=8.483

Farthest point vs Poisson disk - residual

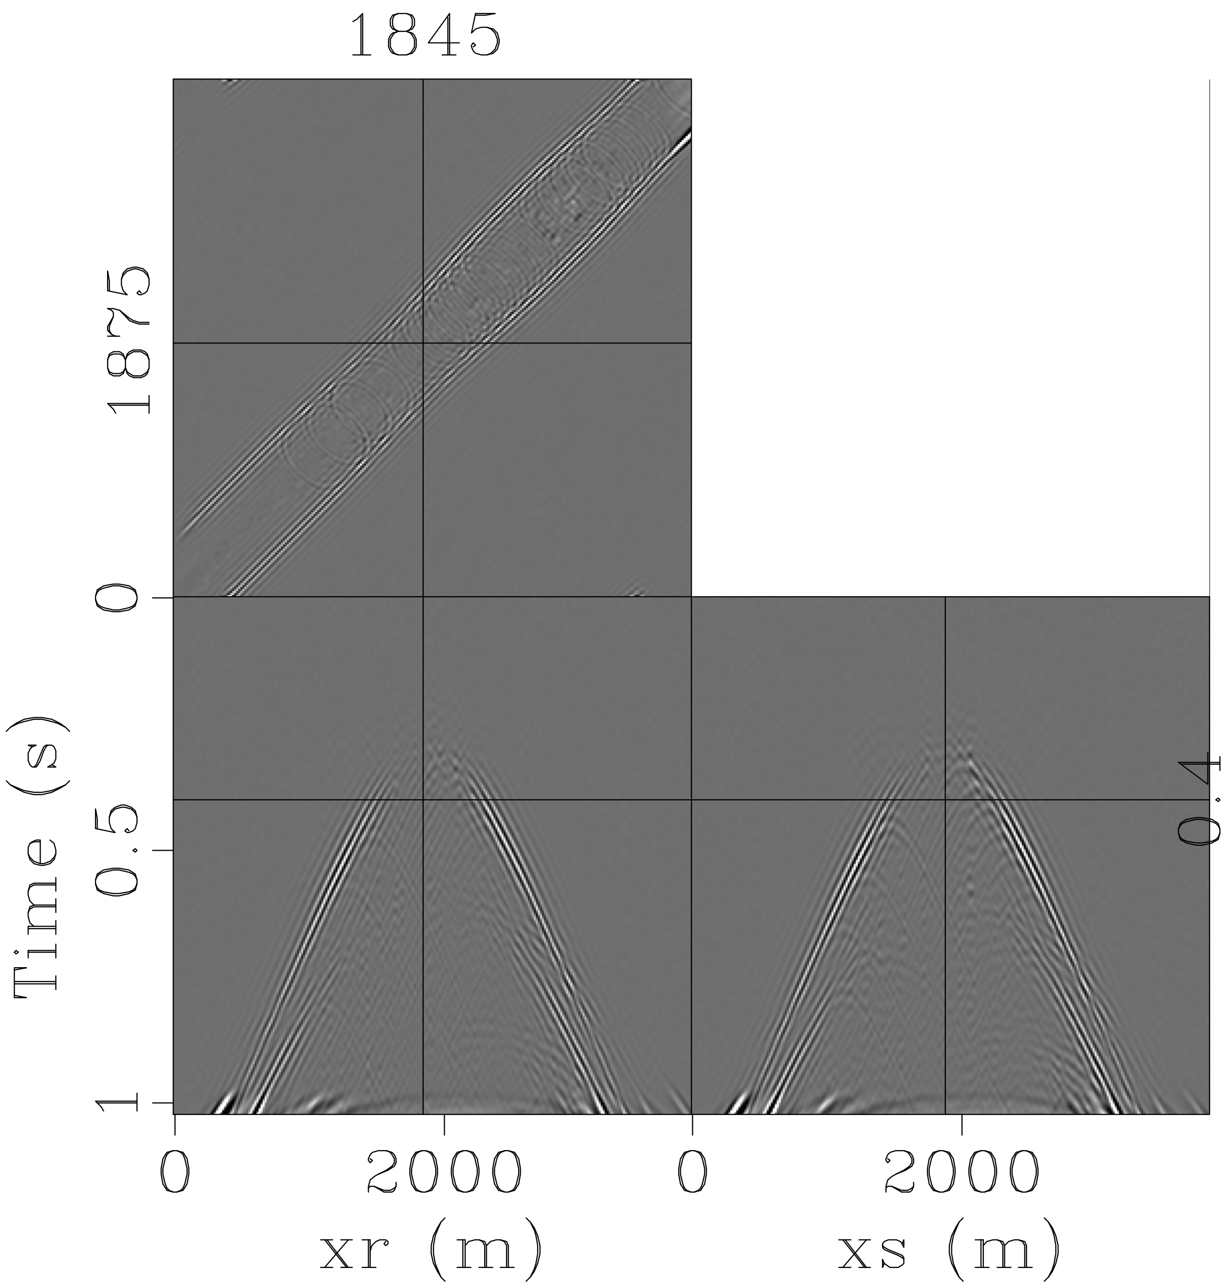

CRSI recon. residual from
Farthest Point samples

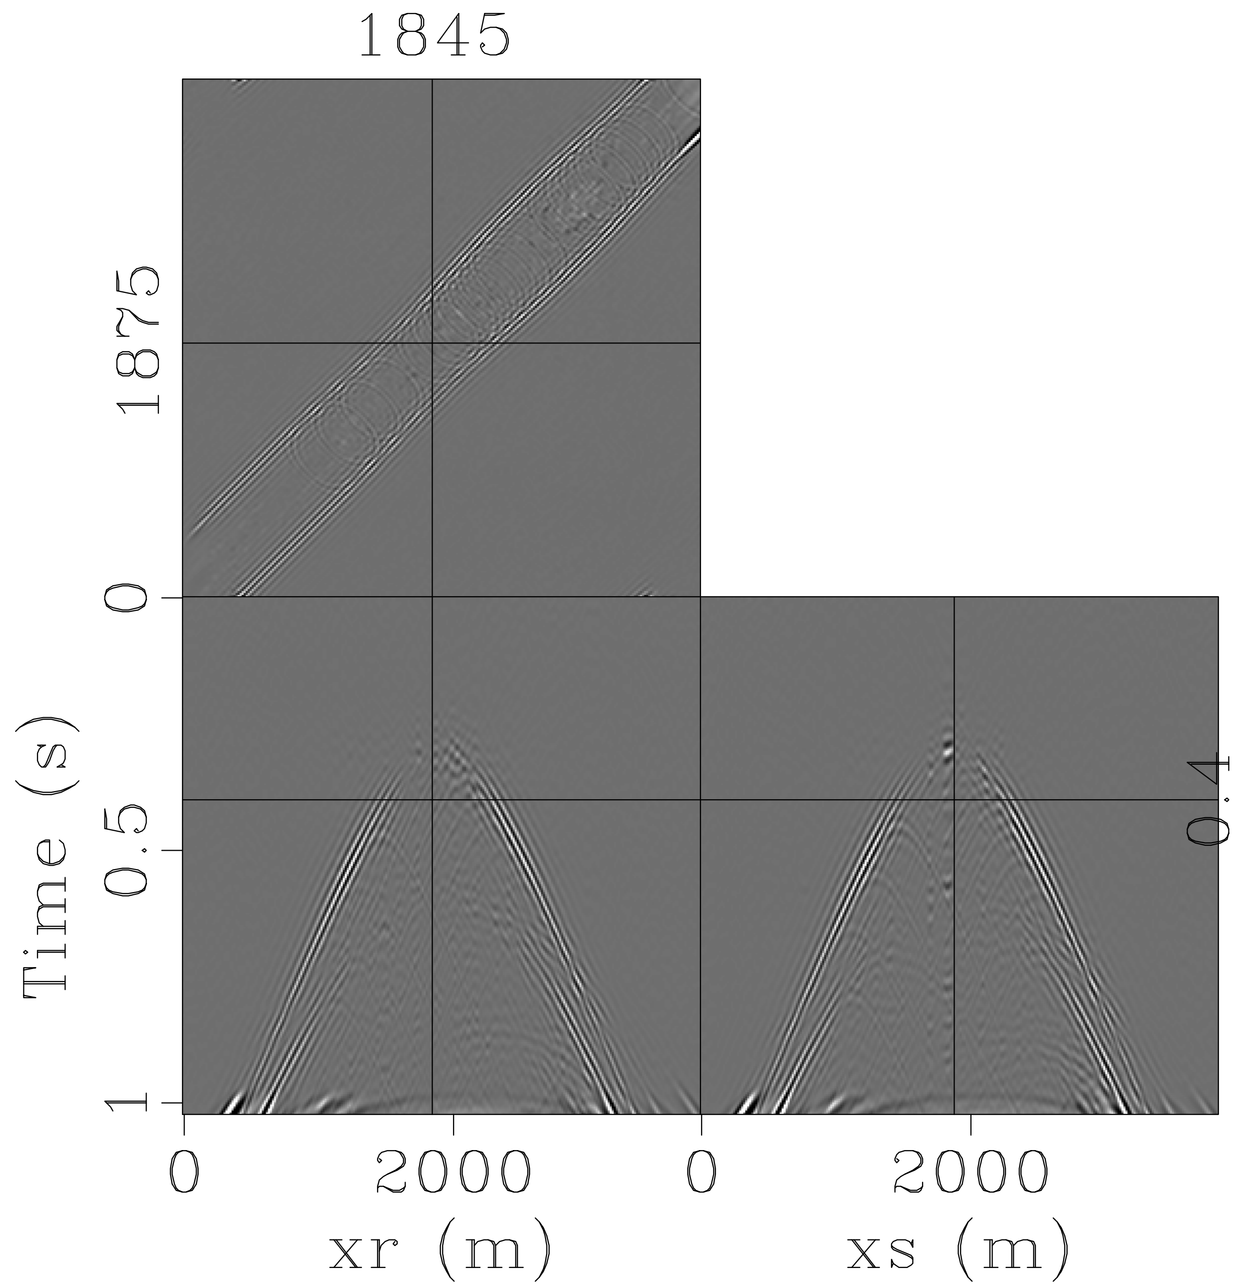

CRSI recon. residual from
Poisson Disk samples

Observation & extensions

- Findings from 1D jittered sampling extend to higher dimensions
  - randomized is better than regular subsampling
  - Cartesian versus hexagonal sampling are equivalent for optimal jittered sampling
  - Furthest point and Poisson sampling lead to similar results

- Gap-size control
  - jittered sampling gives explicit control max distance between adjacent samples
  - farthest point and Poisson disk also have bounds but not explicit

- Future extensions
  - variable density sampling
  - ungridded
  - exploring symmetry (e.g. reciprocity)

Conclusions

- **Randomization** is essential for recovery from incomplete data

- Good **randomized** sampling
  - with blue-noise characteristics give good curvelet recovery
  - with simultaneous sources gives excellent curvelet recovery

- **Randomization** leads to
  - "acquisition" of *smaller* data volumes that carry the ***same*** *information* or
  - to ***improved*** *inferences* from data using the *same* resources

- **Bottom line: acquisition costs** are **no** longer determined by the ***size*** of the ***discretization*** but by ***transform-domain sparsity*** of the sampled wavefield **...**

Acknowledgments

- E. van den Berg and M. P. Friedlander for *SPGL1* (www.cs.ubc.ca/labs/scl/spgl1) & *Sparco* (www.cs.ubc.ca/labs/scl/sparco)
- Sergey Fomel and Yang Liu for Madagascar (rsf.sf.net)
- E. Candes and the Curvelab team

This work was carried out as part of the Collaborative Research & Development (CRD) grant DNOISE (334810-05) funded by the Natural Science and Engineering Research Council (NSERC) and matching contributions from BG, BP, Chevron, ExxonMobil and Shell. FJH would also like to thank the Technische University for their hospitality.

slim.eos.ubc.ca

and... Thank you!