Semiclassical Sampling and Linear Inverse Problems

Sampling meets Microlocal Analysis

Plamen Stefanov

PURDUE
UNIVERSITY

IAS 2019

The classical Shannon sampling theorem says that if $f(x)$ has a Fourier transform $\hat{f}$ supported in the box $[-B, B]^n$ (i.e., it is band-limited), then $f(x)$ is uniquely and stably determined by its samples $f(sk), k \in \mathbf{Z}^n$ if the sampling rate s satisfies $0 < s \le \pi/B$.
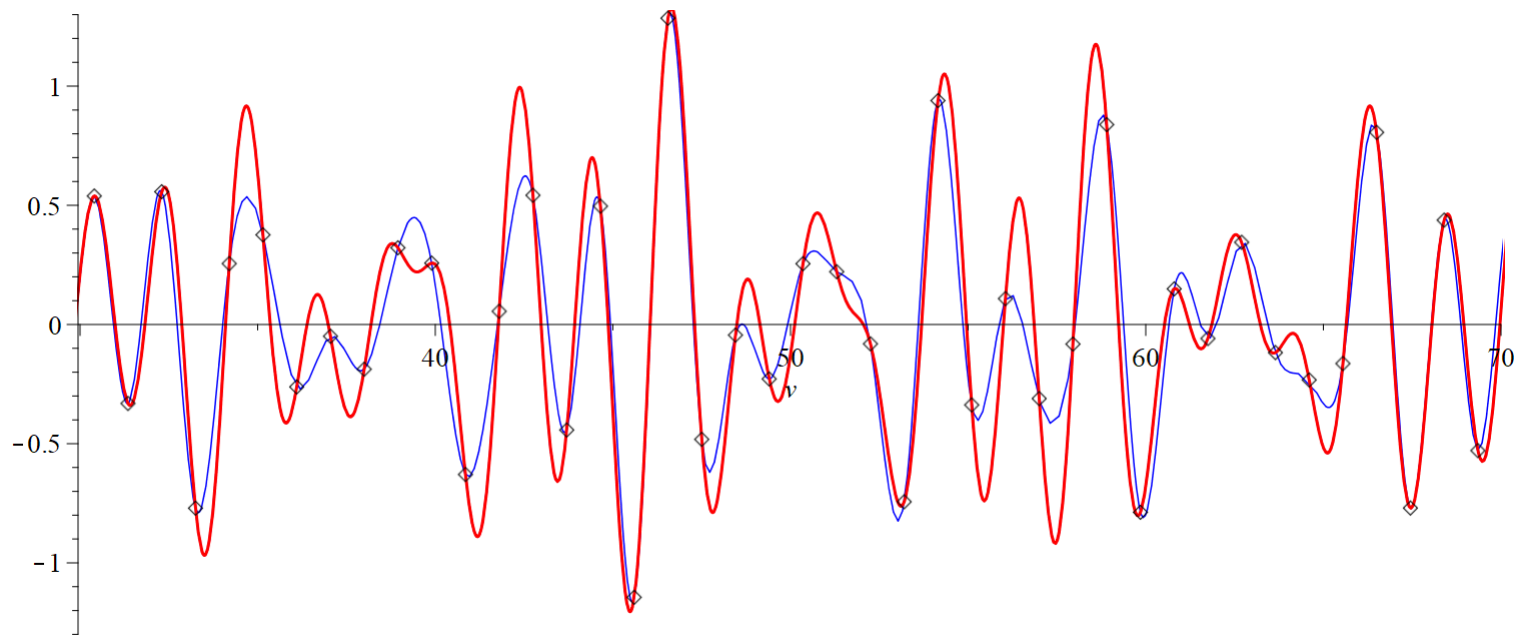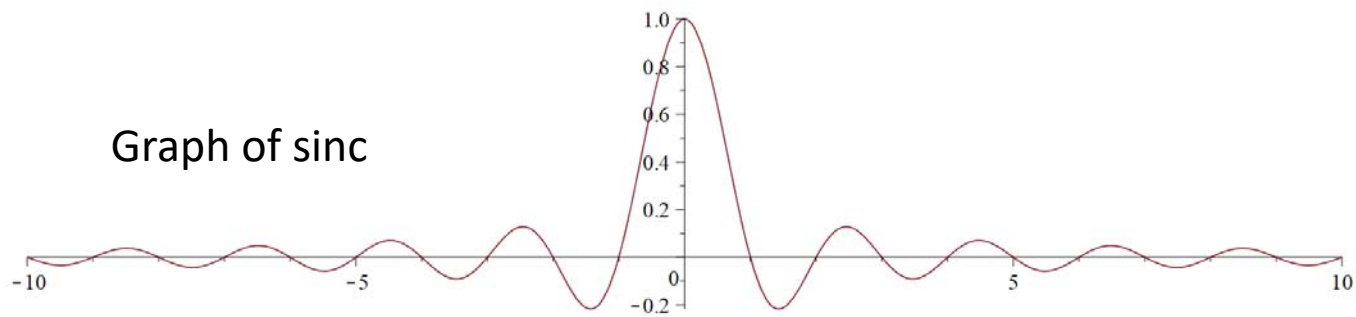
More precisely,

$$f(x) = \sum_{k \in \mathbf{Z}^n}^{\infty} f(sk) \chi\left(\frac{\pi}{s}(x - sk)\right), \qquad \chi(x) = \prod_j \operatorname{sinc}(x_j)$$
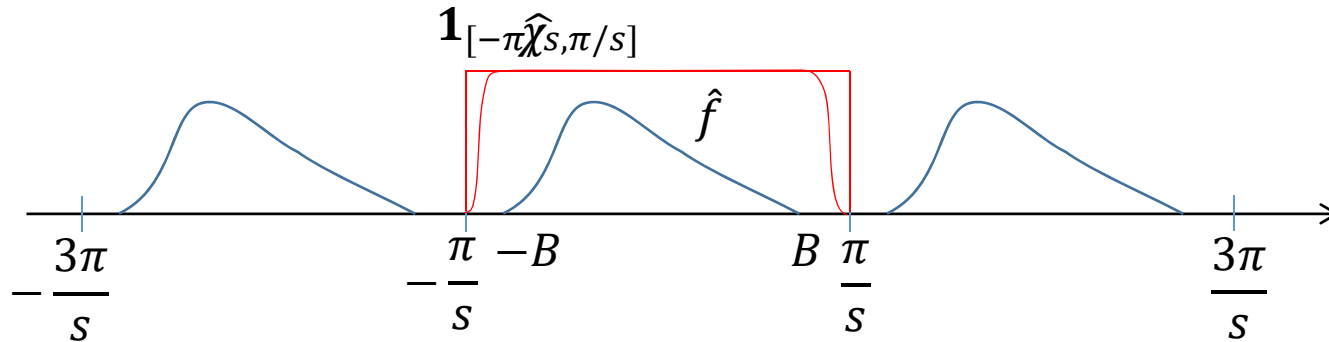
and (unitarity)

$$\|f\|^2 = s^n \sum_{k \in \mathbf{Z}^n}^{\infty} |f(sk)|^2 .$$

The proof is simple. Think of f as the inverse FT of $\hat{f}$. Then the samples $f(sk)$ are the Fourier coefficients of $\hat{f}$, more precisely of its periodic extension over the lattice $B\mathbf{Z}^n$.

Graph of sinc

The red curve is the function reconstructed, indistinguishable from the original. The blue curve is a cubic spline interpolation. Note that the reconstruction has a global character. Here $s = 0.95 * \text{Nyquist}$.

In the frequency domain, things look like this: Nyquist condition satisfied (no aliasing). Since $[-B, B] \subset [\pi/s, \pi/s]$,

$$\mathbf{1}_{[-\pi/s, \pi/s]}$$

$$\hat{f}$$

$$-\frac{3\pi}{s} \qquad -\frac{\pi}{s} \quad -B \qquad B \quad \frac{\pi}{s} \qquad \frac{3\pi}{s}$$

Take the $2\pi/s$-periodic extension $\hat{f}_{\text{ext}}$ of $\hat{f}$; then the samples are its Fourier coefficients:

$$\hat{f}_{\text{ext}}(\xi) = s^n \sum_k f(sk) e^{-i\xi k} .$$

Multiply that extension by $\mathbf{1}_{[-\pi/s, \pi/s]}(\xi)$ to get $\hat{f}$ back:

$$\hat{f}(\xi) = \mathbf{1}_{[-\pi/s, \pi/s]}(\xi) s^n \sum_k f(sk) e^{-i\xi k}.$$

Take $\mathcal{F}^{-1}$. The effect of the multiplication is the appearance of $\mathcal{F}^{-1}$ of $\mathbf{1}_{[-B, B]}(\xi)$ in a convolution, which is the sinc function.

If we have oversampling ($\text{supp}(f)$ strictly in $(-B, B)$), we can choose $\hat{\chi}$ smooth, hence χ will be in the Schwartz class unlike the sinc function.

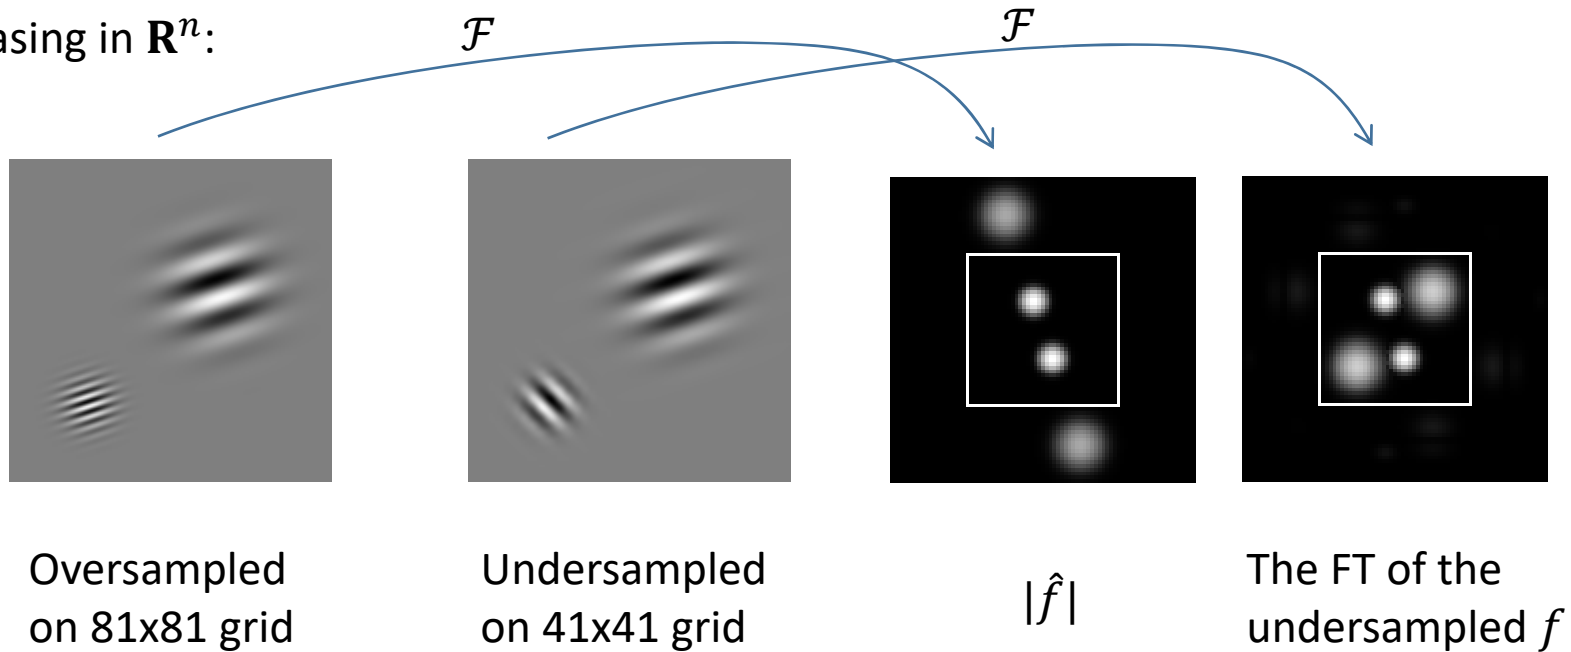**Aliasing:** If the Nyquist condition is not satisfied:

$$\mathbf{1}_{[-B,B]}$$

$$\hat{f}$$

$-3B \qquad -B \qquad\qquad B \qquad\qquad 3B$
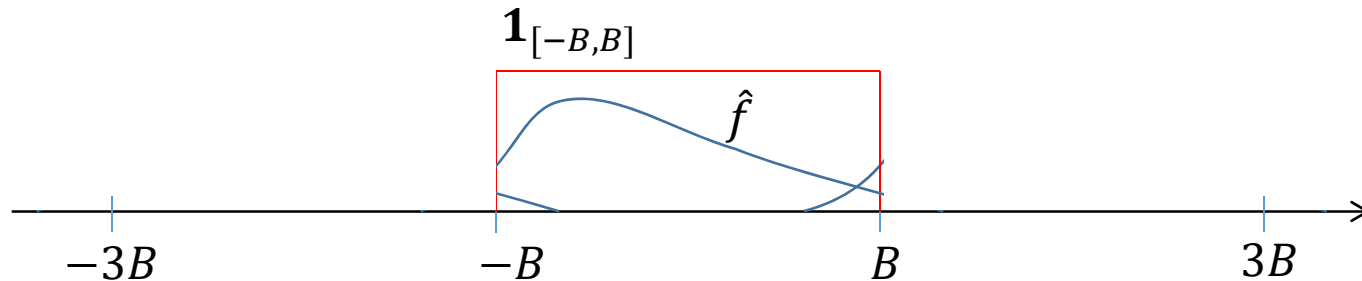
The $2B$-periodic extension has overlapping segments.

If we restrict back to $[-B, B]$, we get $\hat{f}$ truncated plus other stuff, i.e., we get $\hat{f}$ modulo $2B$. Its inverse FT is not f anymore.

This creates aliasing. Frequencies get shifted. If we have a smooth $\hat{\chi}$ instead of $\mathbf{1}_{[-B,B]}$, we get an FIO, actually (in the sense below).

Aliasing in $\mathbf{R}^n$:

Oversampled on 81x81 grid

Undersampled on 41x41 grid

$|\hat{f}|$

The FT of the undersampled f

The original consist of two patterns: one higher frequency than the other. First, we sample and reconstruct properly.

Next, we undersample the higher frequency pattern but still sample properly the lower frequency one. The reconstruction changes the direction and the frequency of the undersampled pattern.

The Fourier transforms demonstrate the shifting (folding) of the frequencies.

We are interested in sampling the data g $= Af$ with A linear, g given, $f =$? We assume that A is an FIO, which is true very often: in integral geometry, thermo-acoustic tomography, etc.

(i) Sampling Af: Having an estimate of supp $\hat{f}$ (the smallest detail of f), how dense should we sample Af?

(ii) Resolution limit on f, given the sampling rate of Af.

(iii) Aliasing: if we undersample Af, what kind of aliasing we get for f?

(iv) Averaged measurements/anti-aliasing: if we blur the data (either because the detectors are not points and average already or because we want to avoid aliasing), what do we get for f?

To answer (i), one may say – estimate supp $\widehat{Af}$ knowing supp $\hat{f}$ and apply the sampling theorem, we are done. The problem is that this is not straightforward. Af may not be even band limited if f is.

We look at the problem as an asymptotic one. The "smallest detail" is a small parameter tending to 0. Rescale ξ to ξ/h (i.e., $\xi = \eta/h$) with $0 < h \ll 1$. This makes it a semi-classical problem!

If $|\xi| \leq B/h$, then for the rescaled ξ we have $|\xi| \leq B$. We want to work in the phase space, i.e., to account for the x dependence as well. So band limited now means that $f_h(x)$ has a compact $\mathrm{WF}_h(f)$.

This brings us to the first problem we need to study:

(v) Semi-classical sampling. How to sample $f_h(x)$ with a compact $\mathrm{WF}_h(f)$?

- Uniform sampling on rectangular or non-rectangular lattices?
- Non-uniform sampling?
- How many sampling points are enough?

Theorem (semi-classical sampling):

Let $f_h \in C_0^\infty(\Omega)$ with $\mathrm{WF}_h(f) \subset \Omega \times [-B, B]^n$. Then for $s \leq \pi/B$,

$$f_h(x) = \sum_{k \in \mathbf{Z}^n} f(shk) \chi\left(\frac{\pi}{sh}(x - shk)\right) + O_s(h^\infty),$$

where χ is a product of sinc functions. Parseval's equality holds, too, up to $O(h^\infty)$.

Just a rescaled classical version, with error estimates. The condition on $\mathrm{WF}_h(f)$ is a condition on $\mathcal{F}_h f$ modulo $O(h^\infty)$.

The step size is sh with $s \leq \pi/B$.

As above, if $\mathrm{WF}_h(f) \subset \Omega \times (-B, B)^n$ (oversampling), χ can be made rapidly decreasing.

We call the projection $\Sigma_h(f)$ of $\mathrm{WF}_h(f)$ onto ξ the frequency set of f.
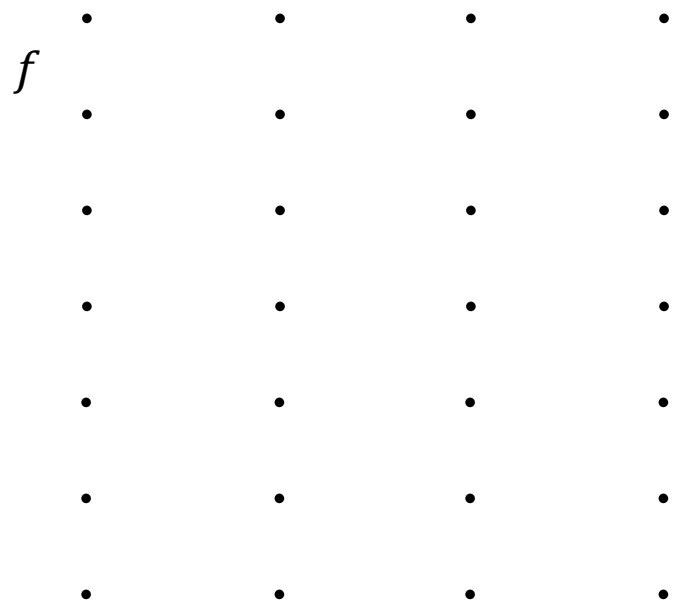
(v) Semi-classical sampling
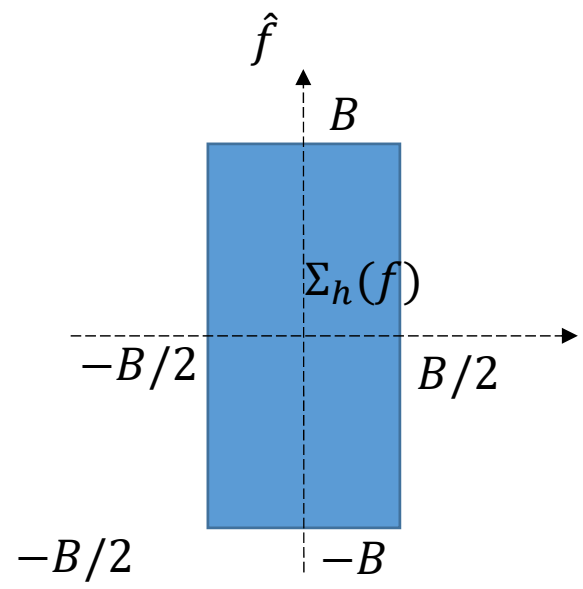
f

$sh \leq \dfrac{\pi h}{B}$

$\hat{f}$

B

$\Sigma_h(f)$

$-B$ B

$-B$
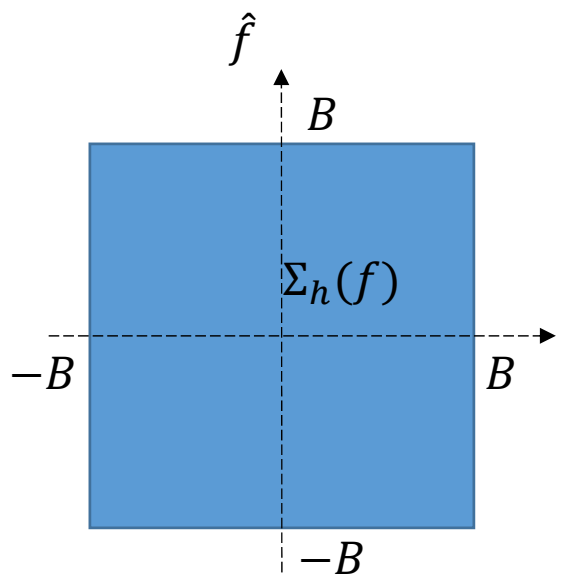
The sampling lattice   The Fourier domain

(v) Semi-classical sampling

f

$\hat{f}$

B

$\Sigma_h(f)$

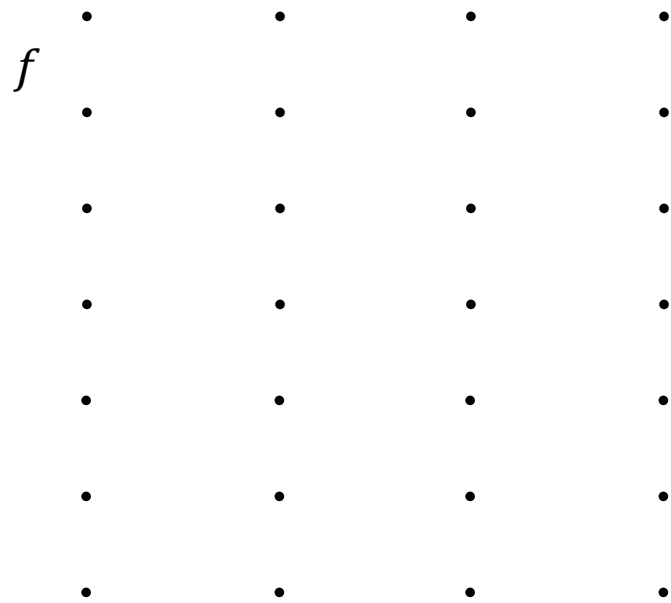$-B/2$ $B/2$

$sh \leq \dfrac{\pi h}{B}$

$-B/2$ $-B$

The sampling lattice   The Fourier domain

The lattice could be rectangular

(v) Semi-classical sampling

f

$$sh \le \frac{2\pi h}{B}$$

$\hat{f}$

B

$-B/2$ $B/2$

$\Sigma_h(f)$

$-B$ B
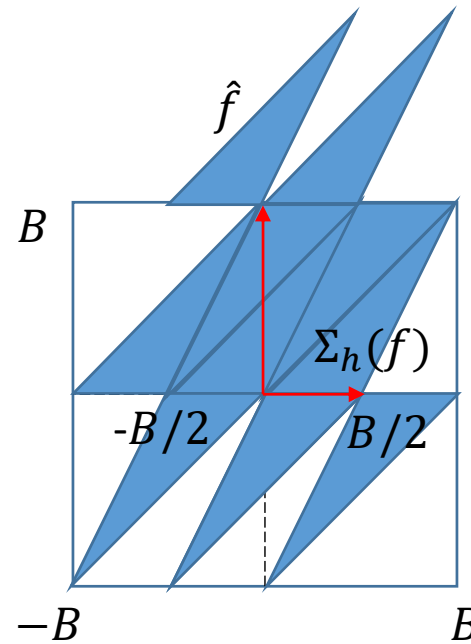
The sampling lattice   The Fourier domain

The previous sampling lattice (rescaled by 2) works for $\Sigma_h(f)$ supported like this. The reason is that we can tile the plane with (non-intersecting) shifts generated by the vectors $(B/2,0)$ and $(0,B)$. Then $\hat{\chi}$ has to be supported there.

One can use linear transformations. If shifts of $\Sigma_h(f)$ under the translations $\xi \mapsto \xi + 2\pi(W^*)^{-1}$ do not intersect, where $\det W \neq 0$, then we can sample over the lattice $x_k = shWk$, $k \in \mathbf{Z}^n$ with $s \leq 1$.

This is a periodic but not a rectangular lattice (a parallelogram one).

One can reduce the number of samples by taking some partition of unity χ_j and in each one, sampling according to $\Sigma_h(\chi_j f)$. This gets us closer to the important question of **non-uniform sampling**.

In the classical case, there is a well cited (780 google scholar citations) paper by LANDAU in Acta, saying, among the rest, that if we can sample uniquely and L^2- stably some f over a possibly non-uniform set of points then the density of that set must have a Weyl type of lower bound proportional to the Lebesgue measure of supp $\hat{f}$. This links sampling to spectral theory.

We have the following microlocal analog of this result.

Theorem. Let $\{x_j(h)\}, j = 1, \dots N(h)$ be a set of points in $\boldsymbol{R}^n$. Let $K \subset T^*\boldsymbol{R}^n$ be a compact set. If

$$\|f\|^2 \leq \frac{C}{N(h)} \sum_{j=1}^{N(h)} \left| f_h\left(x_j(h)\right) \right|^2 + O(h^\infty)$$

for every f_h with $\mathrm{WF}_h(f) \subset K$, then

$$N(h) \geq (2\pi h)^{-n} \mathrm{Vol}(K^{\mathrm{int}})(1 + o(h)).$$

The main novelty here, besides the semi-classical setting, is that we relate the number of points needed to the phase volume of f_h, i.e., to $\mathrm{Vol}\left(\mathrm{WF}_h(f)\right)$ instead of $\mathrm{Vol}(\mathrm{supp}\,\hat{f})$ (multiplied by $\mathrm{Vol}(\mathrm{supp}\,f)$).

The proof uses the spectral asymptotics (DIMASSI & SJÖSTRAND) for the smoothened version of the semiclassical ΨDO $\mathbf{1}_K(x,\xi)$.

In the uniform cases above, we have equality.

(i) How to sample Af when A is an FIO and $\mathrm{WF}_h(f) \subset \Omega \times \{|\xi| \leq B\}$ (for example)?

In other words, we know an a priori bound of the smallest detail f has. How to sample Af?

We need to fit $\mathrm{WF}_h(Af)$ in the smallest product $\Omega' \times K$ with K a box or a parallelogram or a domain satisfying the tiling property w.r.t. some lattice. Then we sample on the reciprocal one.

Now, if A is a semi-classical FIO or if $\mathrm{WF}_h(Af)$ was the classical WF set, then $\mathrm{WF}_h(Af)$ is related by $\mathrm{WF}_h(f)$ by the canonical relation C of A. It turns out that, aside from the zero section, this is also true in this case

$$\mathrm{WF}_h(Af) \setminus 0 \subset C \circ \mathrm{WF}_h(f) \setminus 0.$$

Then the sampling of Af is determined by the geometry of

$$C \circ (\Omega \times \{|\xi| \leq B\}).$$

(ii) Resolution limit on f, given the sampling rate of Af.

(ii) Let us say we sample Af at same (limited) rate. This is what is done in practice. What is the smallest detail of f we can possibly recover?

The way this is asked, it depends on whether A is invertible, microlocally or not, in the first place.

Let us say we sample Af on Ω' at a rate sh. Then if

$$\mathrm{WF}_h(f) \subset \mathrm{C}^{-1} \circ \left(\Omega' \times [-B', B']^{n'} \right)$$

we captured all "detail" of Af and did not lose any resolution about f is A is microlocally invertible in the first place. If A is elliptic, then we can actually recover all the detail of f. If the above inclusion fails, there is aliasing.

Note that the set on the right is not necessarily a product. This means resolution changing with location and direction.

(iii) Let us say that Af is undersampled. How would that affect f?

From now on, we assume a reconstruction with $\hat{\chi}(\xi)$ smooth (so it is a s.c. symbol); then $\chi \in \mathcal{S}$. An inspection of the reconstruction formula says that instead of f_h, we get

$$Gf_h = \sum_{k \in \mathbf{Z}^n} G_k f_h \, ,$$

with G_k a semi-classical FIO with a canonical relation

$$S_k : (x, \xi) \mapsto (x, \xi + 2\pi k/s)$$

(and amplitude $\hat{\chi}(s\xi/\pi)$ when the phase is $\phi = (x - y) \cdot \xi - 2\pi i k \cdot y/s$).

The aliasing leaves x in place but shifts the frequencies around (and changes their directions) as we saw above.

This characterizes aliasing as an h-FIO. The question however is what does aliasing of Af do to f (not Af)?

Let us say that A is elliptic and associated to a local diffeo. Then aliasing of the data gives us the following reconstruction

$$\sum_k A^{-1} G_k A f$$

instead of f (if there is no aliasing, only $k = 0$ is there and $G_0 = \text{Id}$). The result is a sum of h-FIOs with canonical relations

$$C^{-1} \circ S_k \circ C.$$

S_k are frequency shifts but the composition is not, in general! We can expect shifts in the x variable as well!

(iv) What if we sample $\phi_h * A f_h$ with $\phi_h(x) = \epsilon^{-n}\phi(x/\epsilon)$?

Let's say that A is associated with a local diffeo C. The operator $g \mapsto \phi_h * g$ is an h-ΨDO with a principal symbol $\hat{\phi}$. Then by Egorov's theorem,

$$\phi_h * A f_h = A P_h f_h + O(h^\infty) f_h,$$

where P_h is an h-ΨDO with a principal symbol $\hat{\phi} \circ C$. Depending on $\hat{\phi}$, this could blur f, so we are back to the problem (i) but we recover a blurred version $P_h f_h$ of f_h.

This works if ϕ_h changes from a sampling point to another; then the samples became samples of $Q_h A f_h$ with Q_h an h-PDO. We can still apply Egorov's theorem.

Finally, we can solve the inverse problem within the inverse problem: choose Q_h so that P_h becomes what we want, say convolution with a fixed kernel.

We present some applications. Let $\mathcal{R}_\kappa$ be the weighted Radon transform in 2D

$$\mathcal{R}_\kappa f(\omega, p) = \int_{x \cdot \omega = p} \kappa(x, \omega) f(x) dS_x.$$

Write $\omega(\varphi) = (\cos \varphi, \sin \varphi)$. It is well known that $\mathcal{R}_\kappa$ is an FIO with canonical relation $C = C_- \cup C_+$, where

$$C_\pm(x, \xi) = (\underbrace{\arg(\pm\xi)}_{\varphi}, \underbrace{\pm x \cdot \xi/|\xi|}_{p}, \underbrace{-x \cdot \xi^\perp}_{\hat\varphi}, \underbrace{\pm|\xi|}_{\hat p}).$$

Each $C_\pm$ is a diffeo, with inverse $(x, \xi) = C_\pm^{-1}(\varphi, p, \hat\varphi, \hat p)$ given by

$$x = p\omega(\varphi) - (\hat\varphi/\hat p)\omega^\perp(\varphi), \qquad \xi = \hat p\, \omega(\varphi).$$

Sampling $\mathcal{R}$ (with $\kappa = 1$) was studied by BRACEWELL, RATTEY & LINDGREN and NATTERER. The proofs are not complete because they require two-parameter asymptotics of Bessel functions. The results are actually asymptotic even though they are not stated as such. We will prove them for the weighted $\mathcal{R}_\kappa$.

Assume ${\rm WF}_h(f) \subset \{|x| < R, |\xi| < B\}$. Take C of that and project it to the $(\hat{\varphi}, \hat{p})$ variables. We get
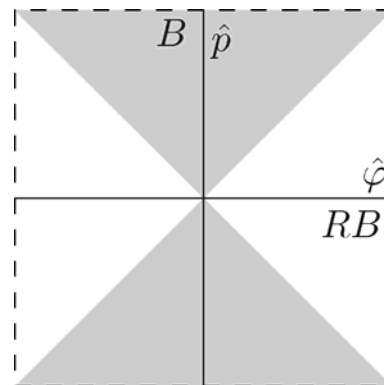
$$\Sigma_h(\mathcal{R}_\kappa f)$$

This is an elementary calculation knowing C. It is the same set obtained by the authors above.

The smallest bounding box is

$$[-RB, RB] \times [-B, B].$$

This determines relative sampling rates (before multiplying by h)

$$s_\varphi \leq \frac{\pi}{RB}, \qquad s_p \leq \frac{\pi}{B}$$
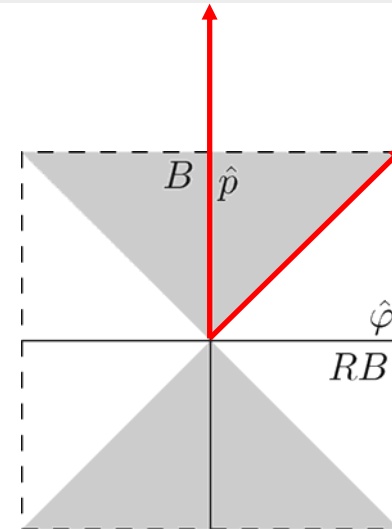
The sampling points are $8/\pi$ times twice the phase volume of f (since C is a double cover, "twice" is the right target).

To get a more efficient sampling set, notice that the plane can be tiled by translations by (RB, B) and $(0, 2B)$.

Thus we can sample on the parallelogram grid $shW\mathbf{Z}^n$ with any $s \leq 1$ and

$$W = \frac{\pi}{RB}\begin{pmatrix} 2 & -1 \\ 0 & R \end{pmatrix}.$$

Those are $4/\pi$ times the phase volume x2. Ideally, it would be equal to it.

Remember, $\Sigma_h(\mathcal{R}_\kappa f)$ is just a projection of $\mathrm{WF}_h(\mathcal{R}_\kappa f)$. We can do better than this. Here is $\mathcal{R}f$ (with $\kappa = 0$) for f consisting of 6 small Gaussians. Then we superimpose $\mathrm{WF}_h(\mathcal{R}f)$ (the projection) on it over each point.
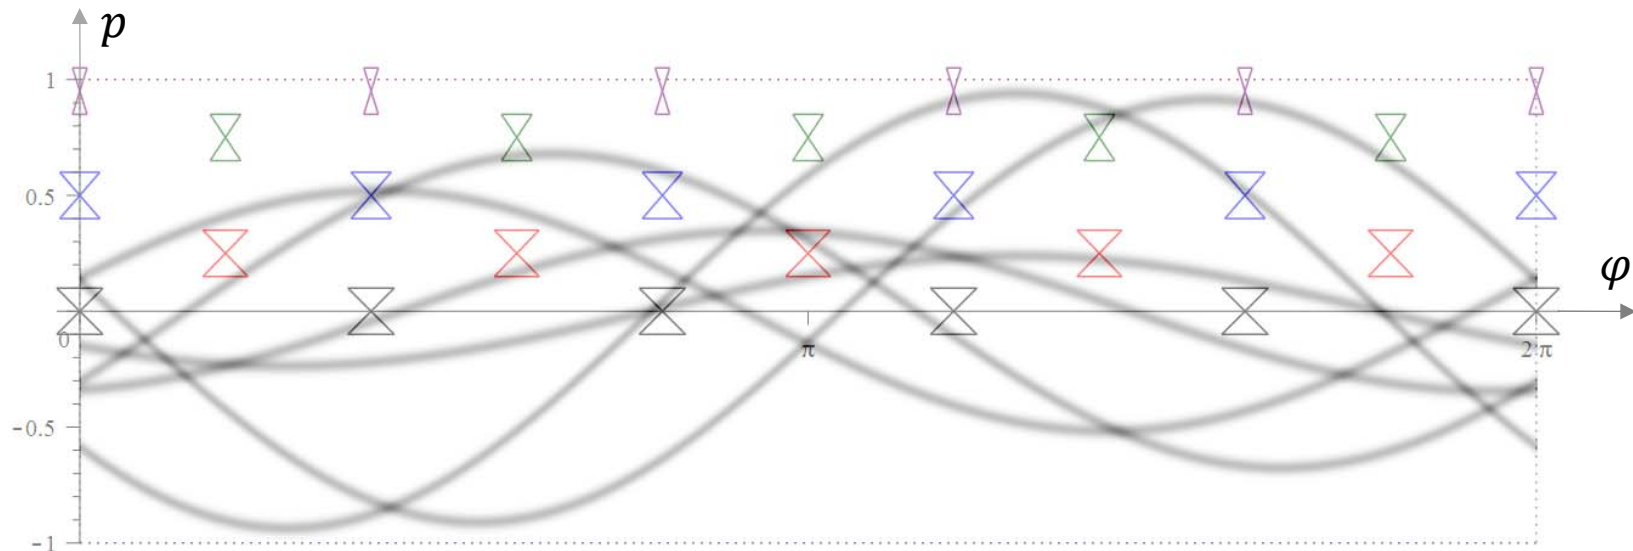
We see that the cone shape of $\Sigma_h(\mathcal{R}f)$ above is the union of all those cones, and it is also what happens at $p = 0$. If we sum up (integrate) the areas of those "cones", we get the sharp sampling count = 2x that for f. To get close to that, we need to split the rectangle into small horizontal strips and sample differently in each one of them.

Let s_φ and s_p be fixed. How would that affect f? From C, we have

$$|x \cdot \xi^\perp| \leq \pi/s_\varphi, \qquad |\xi| \leq \pi/s_p.$$

Circular lines well resolved. Radial not.

The first one tells us that s_φ does not restrict resolution much in the center ($x \approx 0$) or when $\xi \parallel x$. On the other hand, s_p restricts resolution uniformly (which is consistent with the intertwining property of d_p^2 and Δ_x).

best resolution

uniform blur

plus non-local artifacts
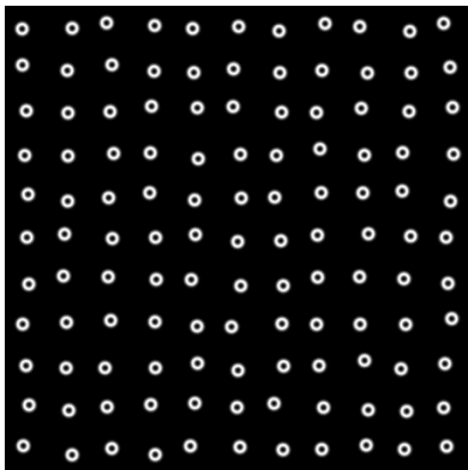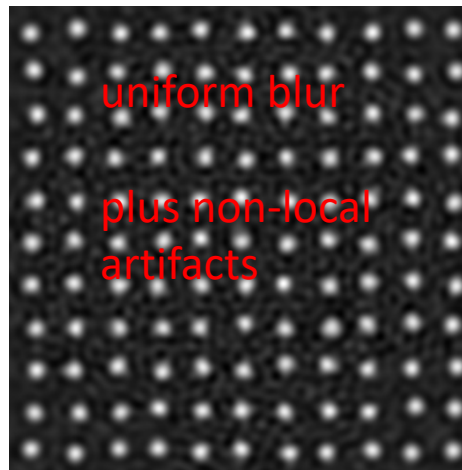
original	Undersampled in p and reconstructed	Undersampled in φ and reconstructed
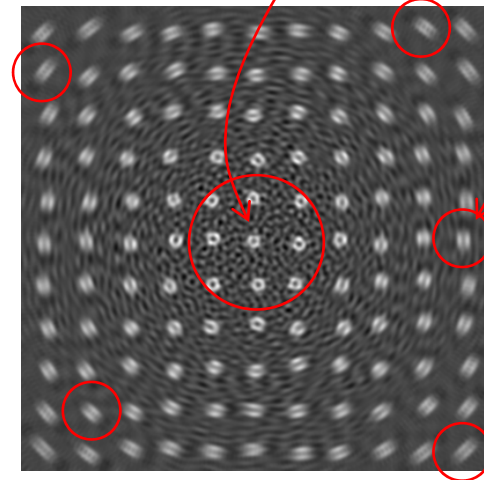
There were a lot of artifacts in those images due to aliasing of the data. As we discovered, the reconstructed image is an *h*-FIO of the original (a sum of such).

Assume we <u>undersample angularly</u>. The canonical relation *on the data* in this case is

$$S_k : (\hat{p}, \hat{\varphi}) \;\; \mapsto \;\; (\hat{p}, \hat{\varphi} + 2\pi k/s_\varphi).$$

No aliasing corresponds to $k = 0$; and typically, in addition, we will see $k = \pm 1$ only. This affects f as a sum of *h*-FIOs with canonical relations

$$(x, \xi) \;\mapsto\; C_\pm^{-1} \circ S_k \circ C_\pm(x, \xi) = \left(x_{\text{shifted}} := x - \frac{2\pi k}{s_\varphi} \frac{\xi^\perp}{|\xi|^2}, \xi \right)$$

when the projection of x_{shifted} to $\xi^\perp$ is in $(-\pi/s_\varphi, \pi/s_\varphi)$. So we see that x shifts along $\xi^\perp$ but ξ does not! This is in some sense the opposite to what happens in classical aliasing.

p

φ

reconstruction original   undersampled $\mathcal{R}f$ (aliased) $\mathcal{R}f$

In other words, the data $\mathcal{R}f$ is aliased as usual: the effect is local and the frequencies are shifted. The reconstructed f is affected differently: the frequencies are preserved but the phantom is shifted!

The Fourier transform of $\mathcal{R}f$

and of the aliased $\mathcal{R}f$.

The frequencies of $\mathcal{R}f$ shift left and right.

Assume we <u>undersample in the p variable</u>. A similar computations shows that both x and ξ shift.
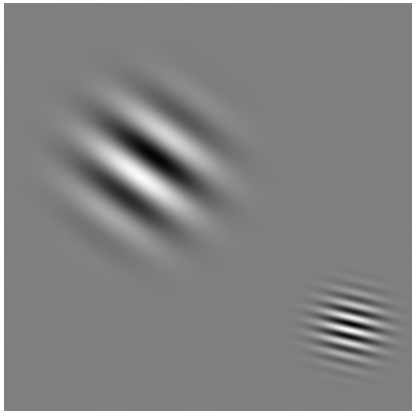
reconstruction original
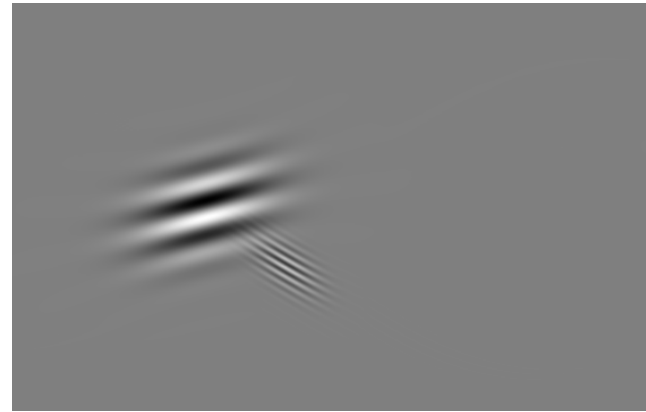
$\mathcal{R}f$ undersampled in p $\mathcal{R}f$

The higher frequency pattern disappears! It actually shifts out of the computational domain.

Blur the data: $Q_h \mathcal{R}_\kappa f$ with Q_h an h-ΨDO with a compact support and principal symbol $q_0(x, \xi)$. Egorov's thm: f is blurred by

$$p_0(x, \xi) = \frac{1}{2} q_0 \circ C_+ + \frac{1}{2} q_0 \circ C_-.$$

If Q_h is a convolution, i.e., $q_0 = \psi(a|\hat{\varphi}|^2 + b|\hat{p}|^2)$, then

$$p_0(x, \xi) = \psi(a|\xi|^2 + b|x \cdot \xi^\perp|^2).$$

The blur is uniform iff $b = 0$, i.e., we average w.r.t. p only. Well known.

If $a = 0$, the blur is position and direction dependent.
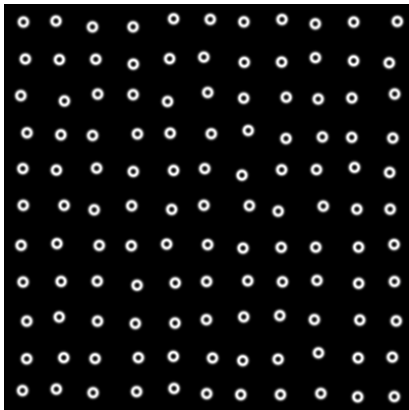
original $\mathcal{R}f$ averaged in p $\mathcal{R}f$ averaged in φ
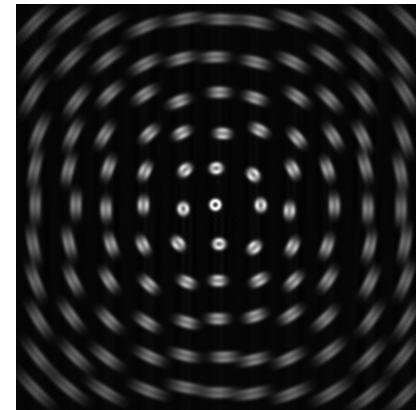
Example: The Radon transform in "fan-beam geometry"
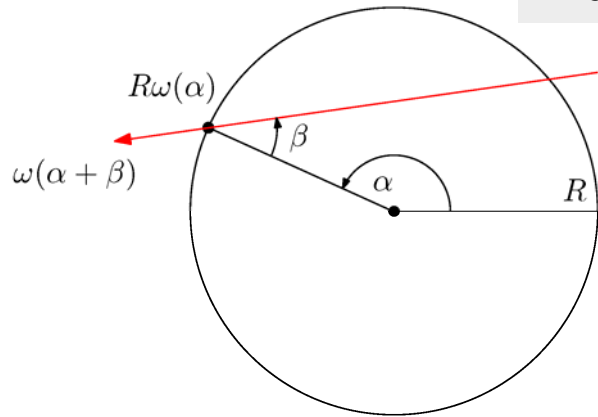
Lines parameterized by initial point $R\omega(\alpha)$ on the circle $|x| = R$ and direction making angle β with the radial ray.

This is diffeomorphic to the parallel geometry case.

One can compute $C = C_- \cup C_+$ as before. Assume as before

$$\mathrm{WF}_h(f) \subset \{|x| < R, |\xi| < B\}.$$

Then $\Sigma_h(\mathcal{R}_\kappa f)$ for all such f can be computed as before. There is a new twist here: $C_-(x, \xi)$ may be aliased, $C_+(x, \xi)$ may be not!

f

$\mathcal{R}f$

$|\widehat{\mathcal{R}f}|$
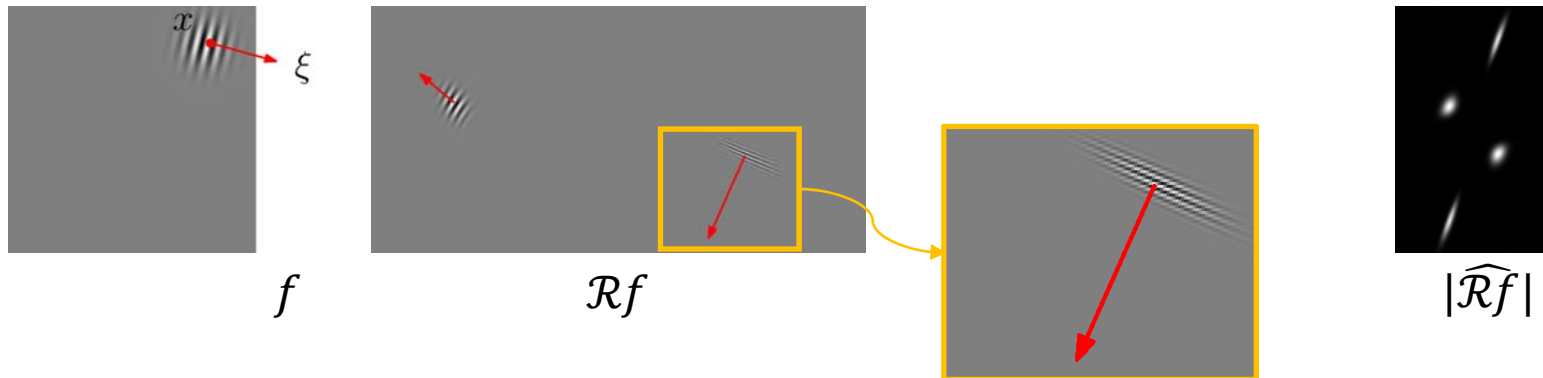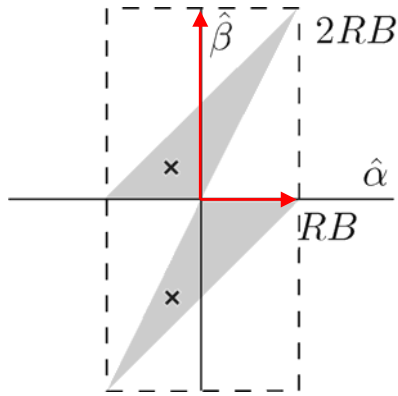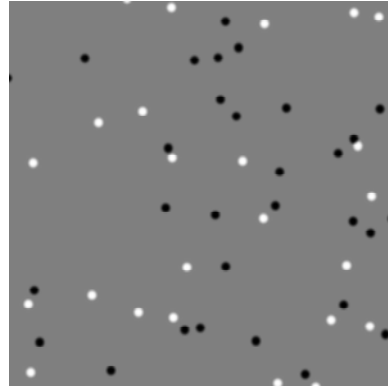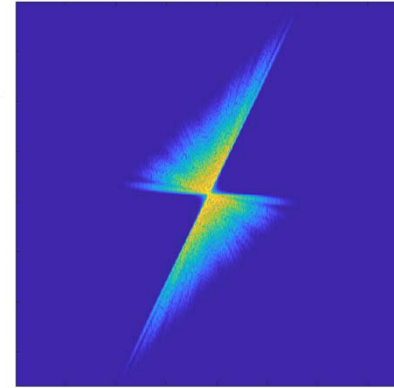
$$\Sigma_h(\mathcal{R}f) \qquad\qquad \text{Example: } f \qquad\qquad \text{Computed } |\widehat{\mathcal{R}f}|$$

The double rectangle above determines the sampling requirements on a rectangular grid:

$$s_\alpha \leq \frac{\pi}{RB}, \qquad s_\beta \leq \frac{\pi}{2RB}.$$

This is $\pi \times$ the parallel case (on a rectangular grid), and $8 \times$ more than the Weyl optimum.

We can tile the plane with translations by $(RB, 0)$ and $(0, 2RB)$ as we showed above. This doubles the step sizes. We still have twice the optimum number.

A better representation of $\mathrm{WF}_h(\mathcal{R}f)$:

$\Sigma_h(\mathcal{R}f)$
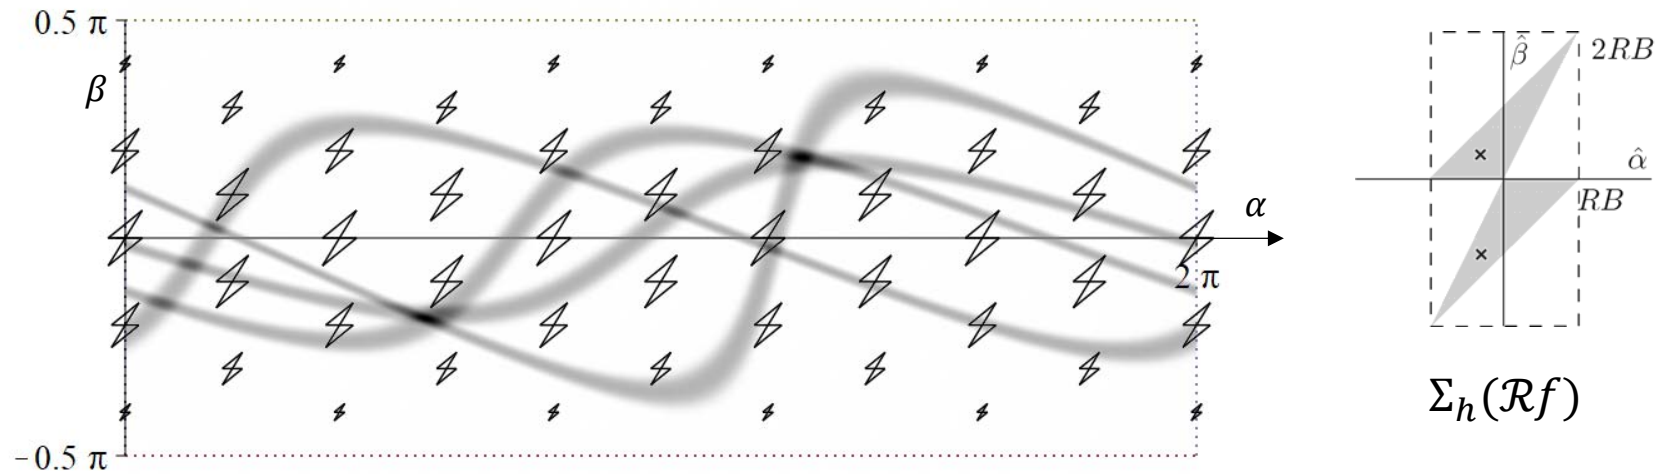
$\mathcal{R}f$ with $\mathrm{WF}_h(\mathcal{R}f)\big|_{(\alpha,\beta)}$ at a few points.

The figure we got earlier is the "worst scenario case": the union of all those. The greatest sampling requirements are near $\beta = 0$. One can sample over small parallel strips with different steps. This would bring us asymptotically close to the Weyl minimum.

Assume we sample $\mathcal{R}_\kappa f$ with relative rates s_α and s_β. What does this imply for f?
We need to understand

$$C^{-1}\left\{(\alpha,\beta) \in [0,2\pi) \times \left[-\frac{\pi}{2},\frac{\pi}{2}\right]; |\hat\alpha| \leq \frac{\pi}{s_\alpha}, |\hat\beta| \leq \frac{\pi}{s_\beta}\right\}.$$
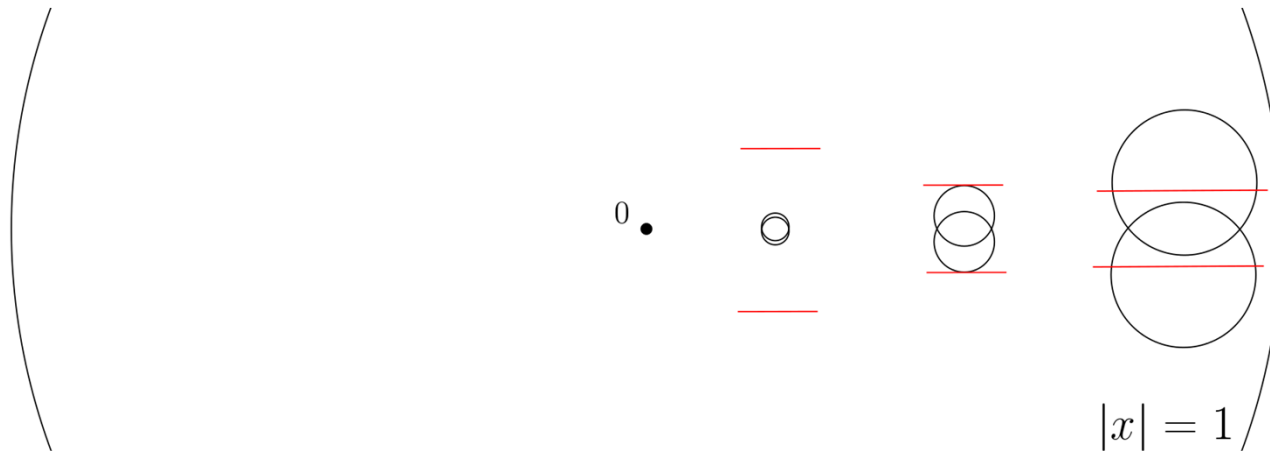
The diagram is rotationally symmetric.

The red bars represent the frequency cutoff caused by s_α.

The black circles correspond to s_β. The frequencies in *the intersection* of the circles are the non-aliased ones. The reason: C^{-1} is 2-to-1. As a result: there is more resolution in the data, if extracted properly!

If the sampling rates s_α and s_β are not small enough, aliasing occurs a before. A direct computation shows that x shifts along the direction of $\xi^\perp$ and ξ changes magnitude but not direction.

Blur the data: $Q_h \mathcal{R}_\kappa f$ with Q_h an h-ΨDO with a compact support and principal symbol $q_0(x, \xi)$. Egorov's thm: f is blurred by
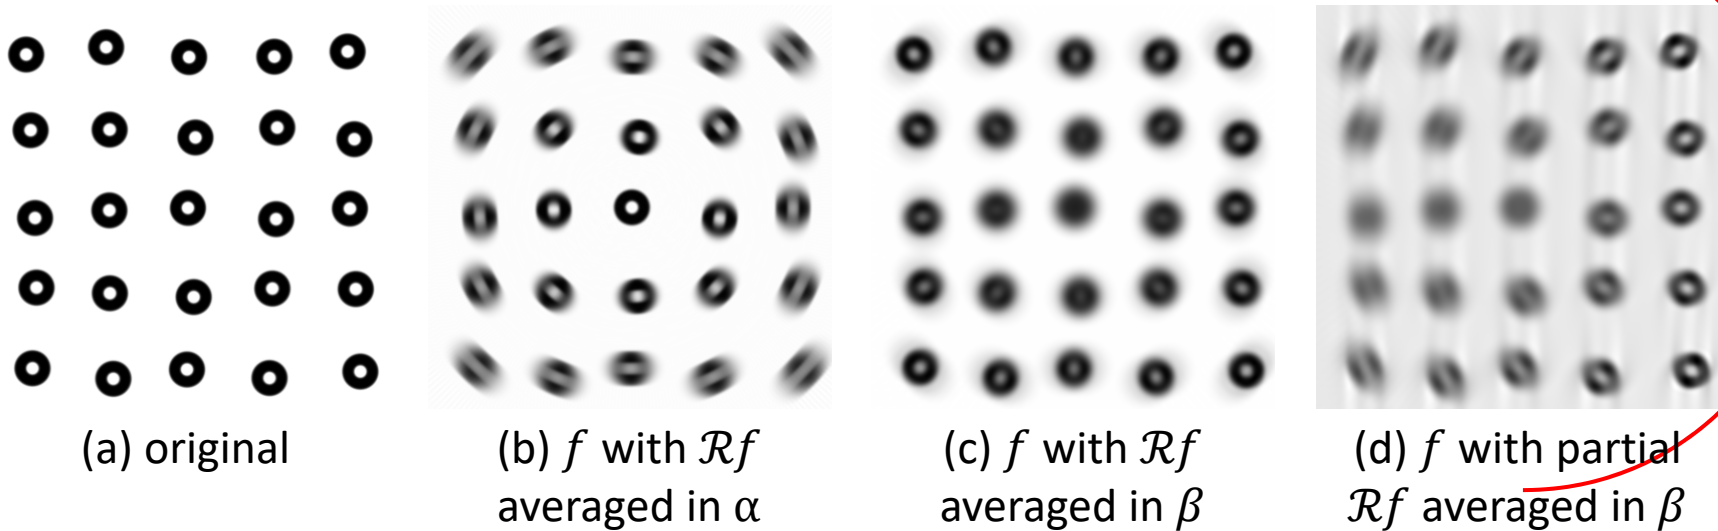
$$p_0(x, \xi) = \frac{1}{2} q_0 \circ C_+ + \frac{1}{2} q_0 \circ C_-.$$

If Q_h is a convolution, i.e., $q_0 = \psi(a|\hat{\varphi}|^2 + b|\hat{p}|^2)$, then

$$p_0(x, \xi) = \frac{1}{2} \psi \left(a|x \cdot \xi^\perp|^2 + b \left| x \cdot \xi^\perp \textcircled{+} \sqrt{R^2|\xi|^2 - (x \cdot \xi)^2} \right|^2 \right)$$

$$+ \frac{1}{2} \psi \left(a|x \cdot \xi^\perp|^2 + b \left| x \cdot \xi^\perp \textcircled{-} \sqrt{R^2|\xi|^2 - (x \cdot \xi)^2} \right|^2 \right).$$

The levels sets are as in the resolution analysis. Unlike in the parallel geometry case, those terms are different! They blur differently and then the results are added.

(a) original

(b) f with $\mathcal{R}f$ averaged in α

(c) f with $\mathcal{R}f$ averaged in β

(d) f with partial $\mathcal{R}f$ averaged in β

(b) Radial blur. Corresponds to the red bars.

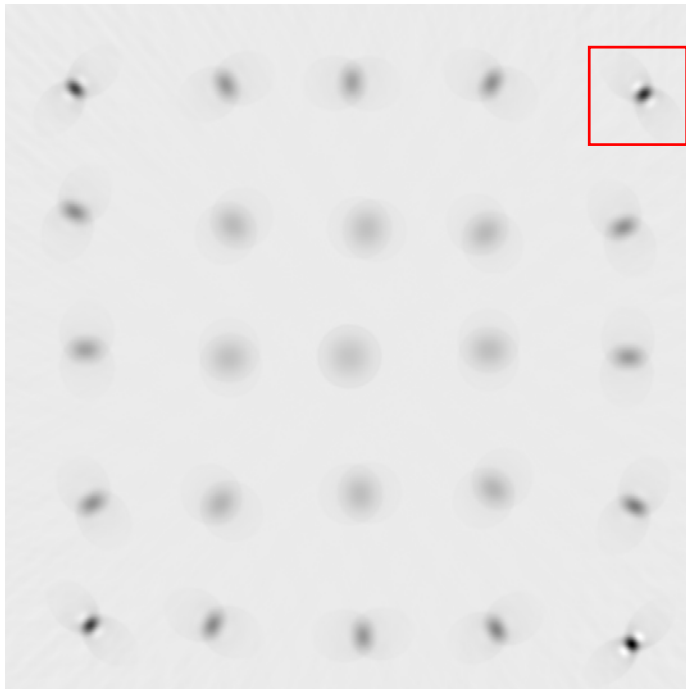(c) Sum of two blurs. Corresponds to the circles.

(d) Source restricted to the right semi-circle. We get better resolution near the source (direction dependent). Corresponds to one of the circles only.
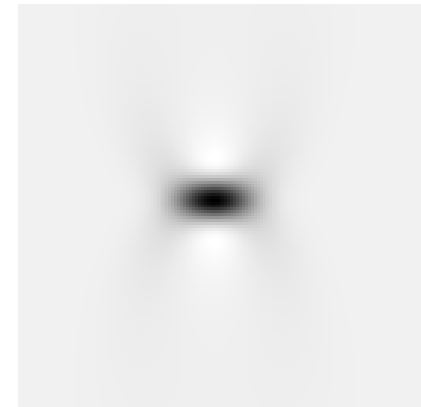

$|x| = 1$
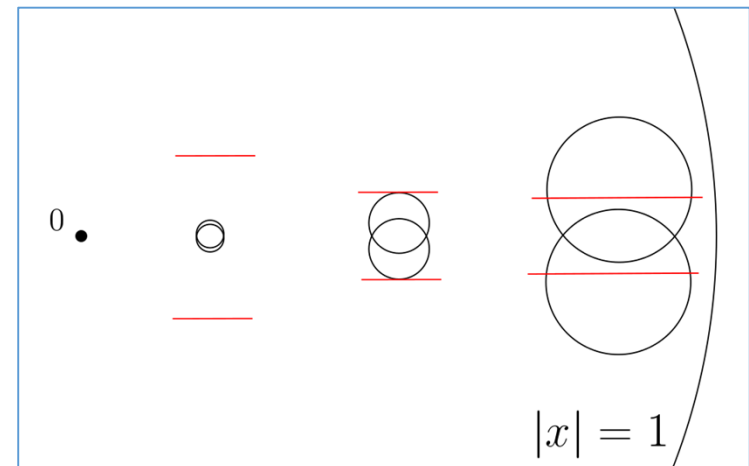
The resolution diagram

To understand better the effect we see in (c) above, here is another test function: a 5×5 array of almost delta-like Gaussians.
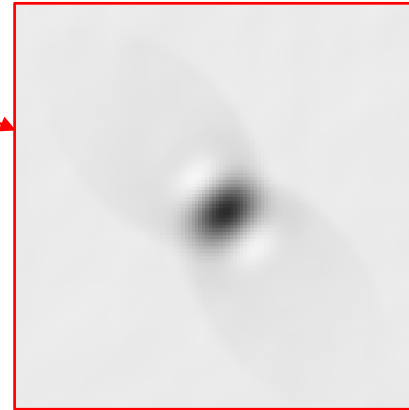
A magnified crop

A density plot of the computed symbol at $(0.88R, 0)$

Reconstructed f with $\mathcal{R}f$
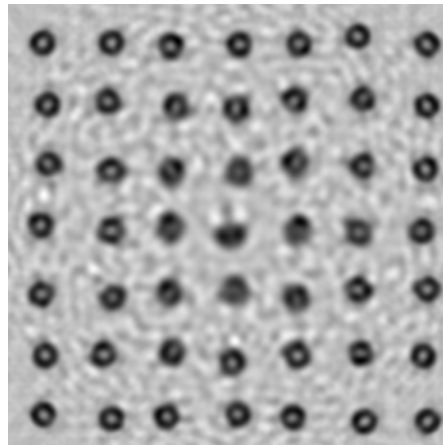averaged in β,
i.e., data: $\psi(\beta) * \mathcal{R}f$

$$p_0(x, \xi) = \frac{1}{2}\psi\left(\left|x \cdot \xi^\perp + \sqrt{R^2|\xi|^2 - (x \cdot \xi)^2}\right|^2\right)$$
$$+ \frac{1}{2}\psi\left(\left|x \cdot \xi^\perp - \sqrt{R^2|\xi|^2 - (x \cdot \xi)^2}\right|^2\right)$$
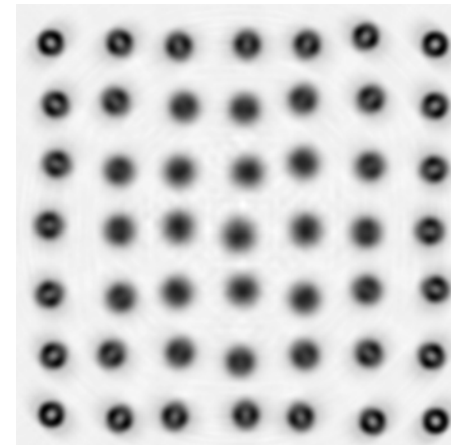
Bluring $\mathcal{R}_\kappa f$ can be used to avoid aliasing.
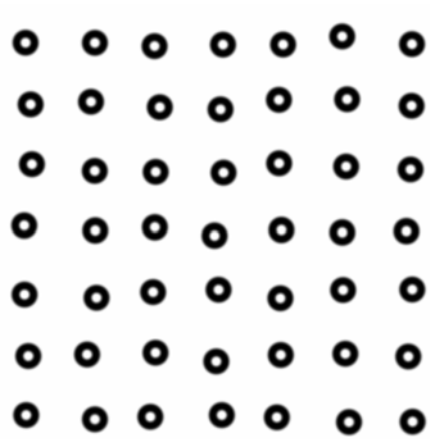
| original | f with $\mathcal{R}f$ undersampled in β | f with $\mathcal{R}f$ blurred in β, then sampled |

A lot of aliasing artifacts spreading everywhere

The anti-aliasing removes most of the artifacts but softens the image a bit more.

Thermo-acoustic Tomography

Model:

$$\begin{cases} u_{tt} - c^2(x)\Delta u = 0 \quad \text{in } \mathbf{R} \times \mathbf{R}^n, \\[2mm] u|_{t=0} = f(x), \quad u_t|_{t=0} = 0, \end{cases}$$
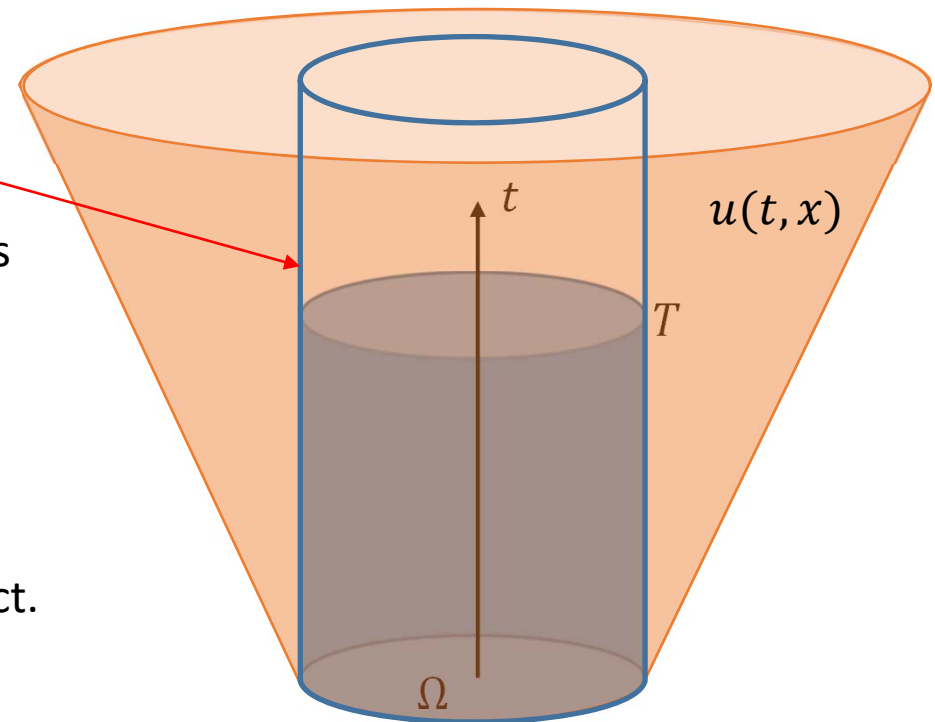
where $u(t,x)$ is the sound pressure, $c(x)$ is the sound speed equal to 1 outside Ω.

We measure $\Lambda f = u|_{[0,T] \times \partial\Omega}$.

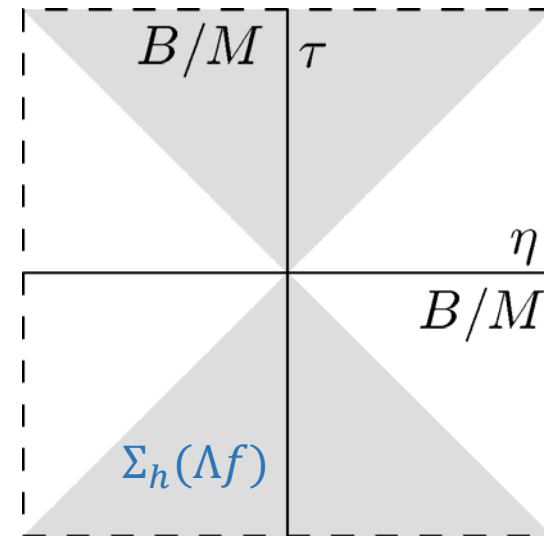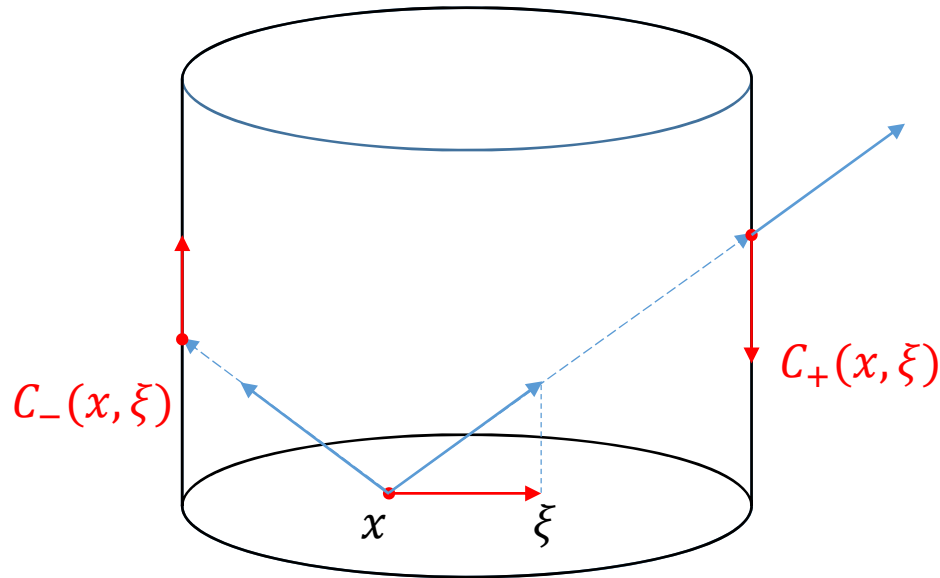There is no actual boundary, the waves propagate to the whole space.

Clearly, Λ is a classical FIO.

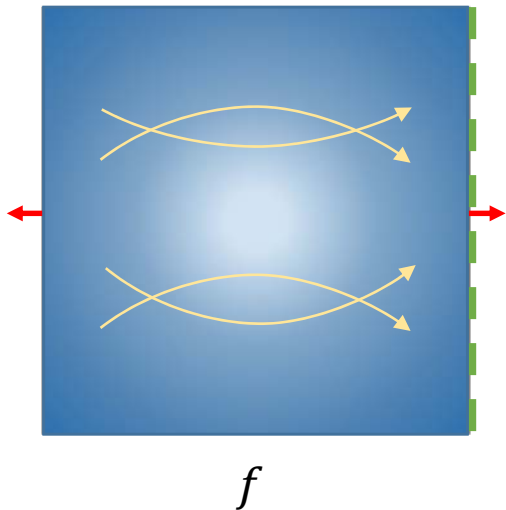Another studied model is when there are Neumann b.c. and the waves reflect. Λ is still a classical FIO.

Canonical relation $C = C_- \cup C_+$.

For every $(t, x) \in (0, T) \times \partial\Omega$, the range of the possible $C_\pm(x, \xi)$ (when $|\xi| \leq B$) is the characteristic cone projected there. Their union gives $\Sigma_h(\Lambda f)$, which is the cone above. Here, $M = \max c$ because we have $M = 1$ if we use the metric $c^{-2} dx^2$ on $T^*\Omega$ but we have $|\xi| \leq B$ with $|\xi|$ Euclidean (as determined by the sampling rate).
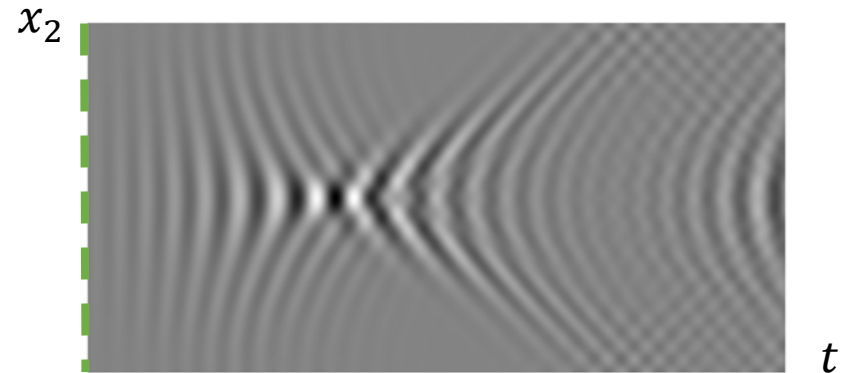
The induced metric on $\partial\Omega$ would results in another constant M' but when $\partial\Omega$ is piecewise flat, $M' = 1$.

Thermo-acoustic Tomography

It turns out that the geometry plays no role! Only $\max c$ and the shape of $\partial\Omega$ matter!

x_2

t

Slow region inside, caustics. Waves focus on $\partial\Omega$. Low sampling/ discretization requirements.

f

Compute Λf in the model with reflections. The r.h.s. shown.

x_2

t
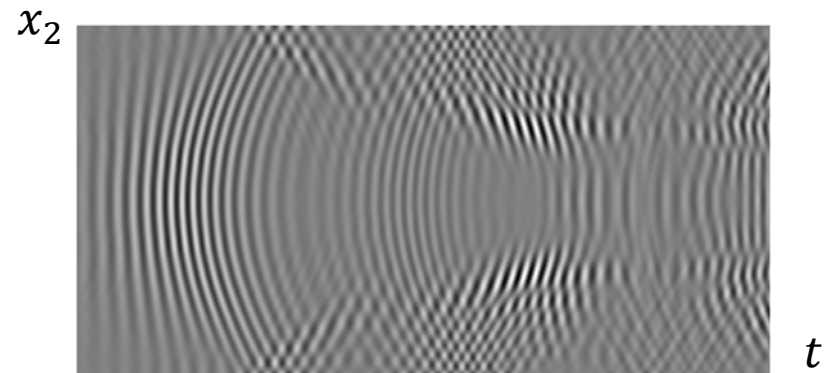
Fast region inside, no caustics (before reflections). Much higher frequencies, needs finer sampling.