ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Near-optimal sensor placement for linear inverse problems

Juri Ranieri

Joint work with: *Dr. Amina Chebira, Prof. Martin Vetterli, Prof. D. Atienza, Dr. Z. Chen, I. Dokmanic, A. Vincenzi, R. Zhang*

Sensing the real world

We experience the surrounding environment through sensors.

We have a set of natural sensors, i.e., eyes, ears, nose...

Sensing the real world

We experience the surrounding environment through sensors.

We have a set of natural sensors, i.e., eyes, ears, nose...

Tech devices are equipped with many sensors, providing an incredible amount of information about the real world.

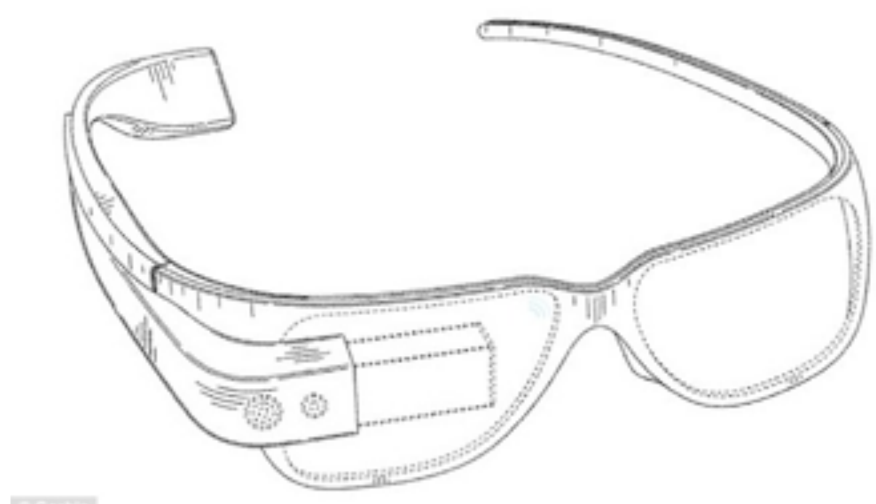

Inverse problems

We have access to an incredible amount of data.
How can we use it?

- Provide it to the end-user as measured,
- Store it on a server for future use,
- Use it to estimate other parameters of interest.

Inverse problems

We have access to an incredible amount of data.
How can we use it?

- Provide it to the end-user as measured,
- Store it on a server for future use,
- Use it to estimate other parameters of interest.

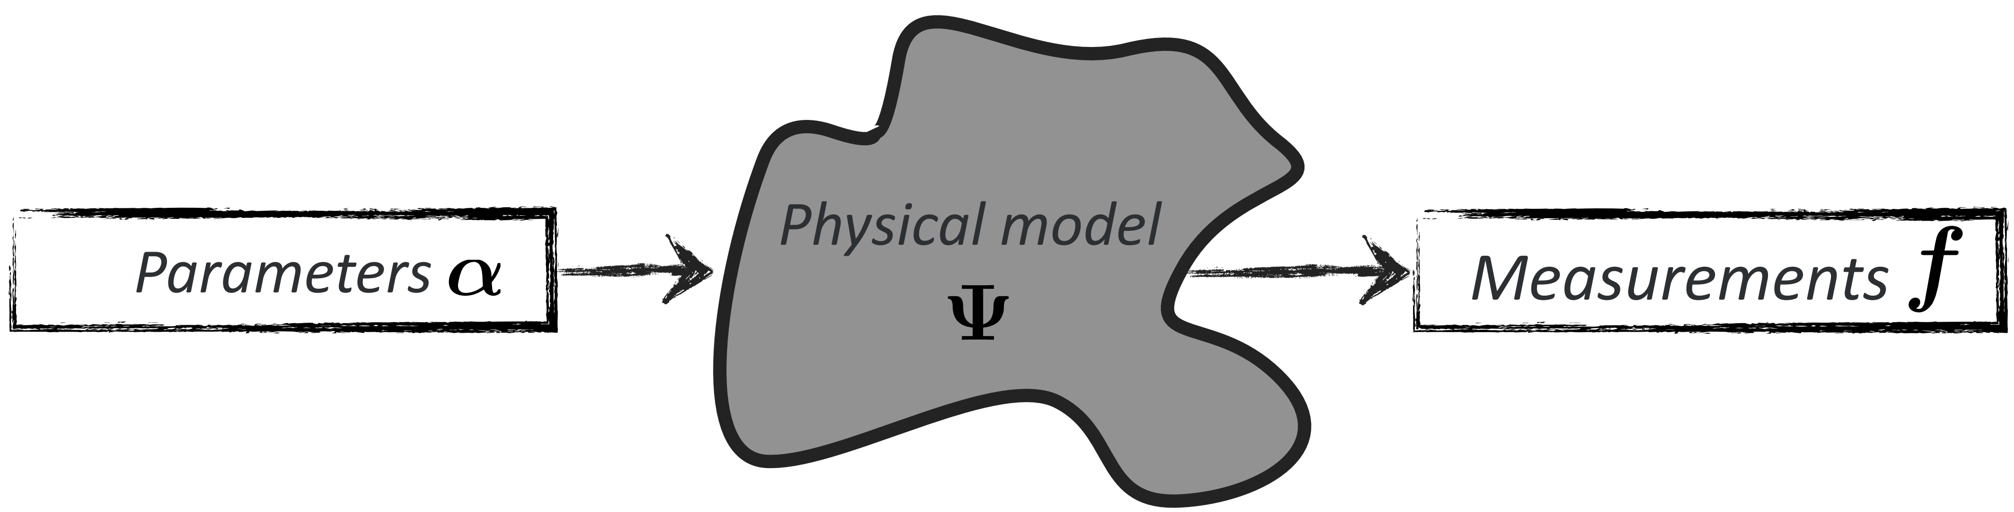

Inverse problems

We have access to an incredible amount of data.
How can we use it?

- Provide it to the end-user as measured,
- Store it on a server for future use,
- Use it to estimate other parameters of interest.

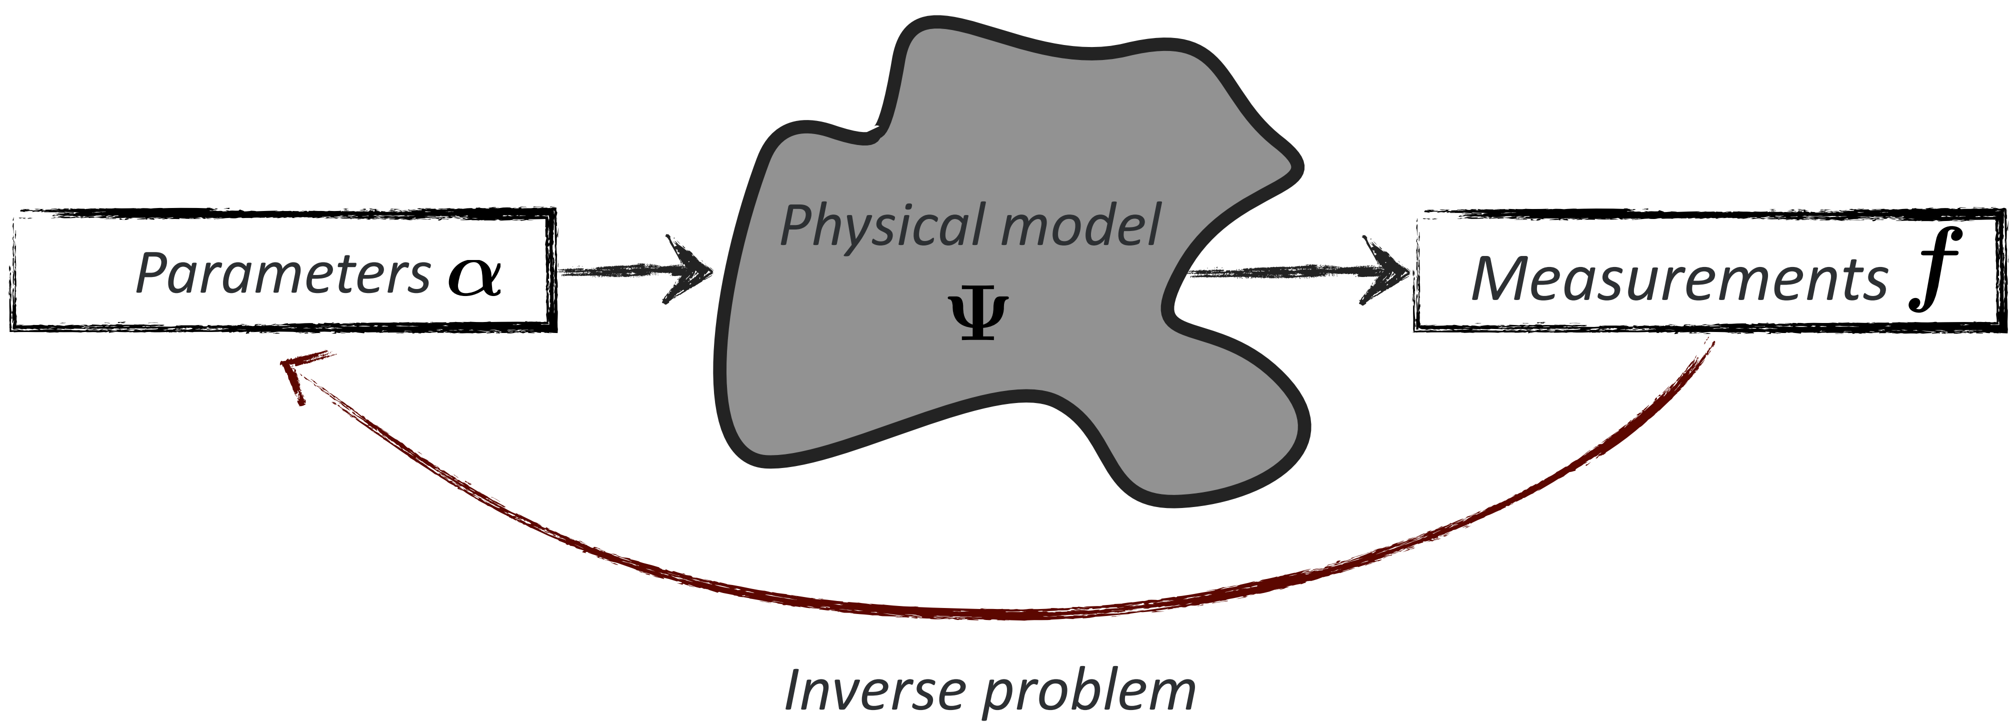

Inverse problems

We have access to an incredible amount of data.
How can we use it?

- Provide it to the end-user as measured,
- Store it on a server for future use,
- Use it to estimate other parameters of interest.

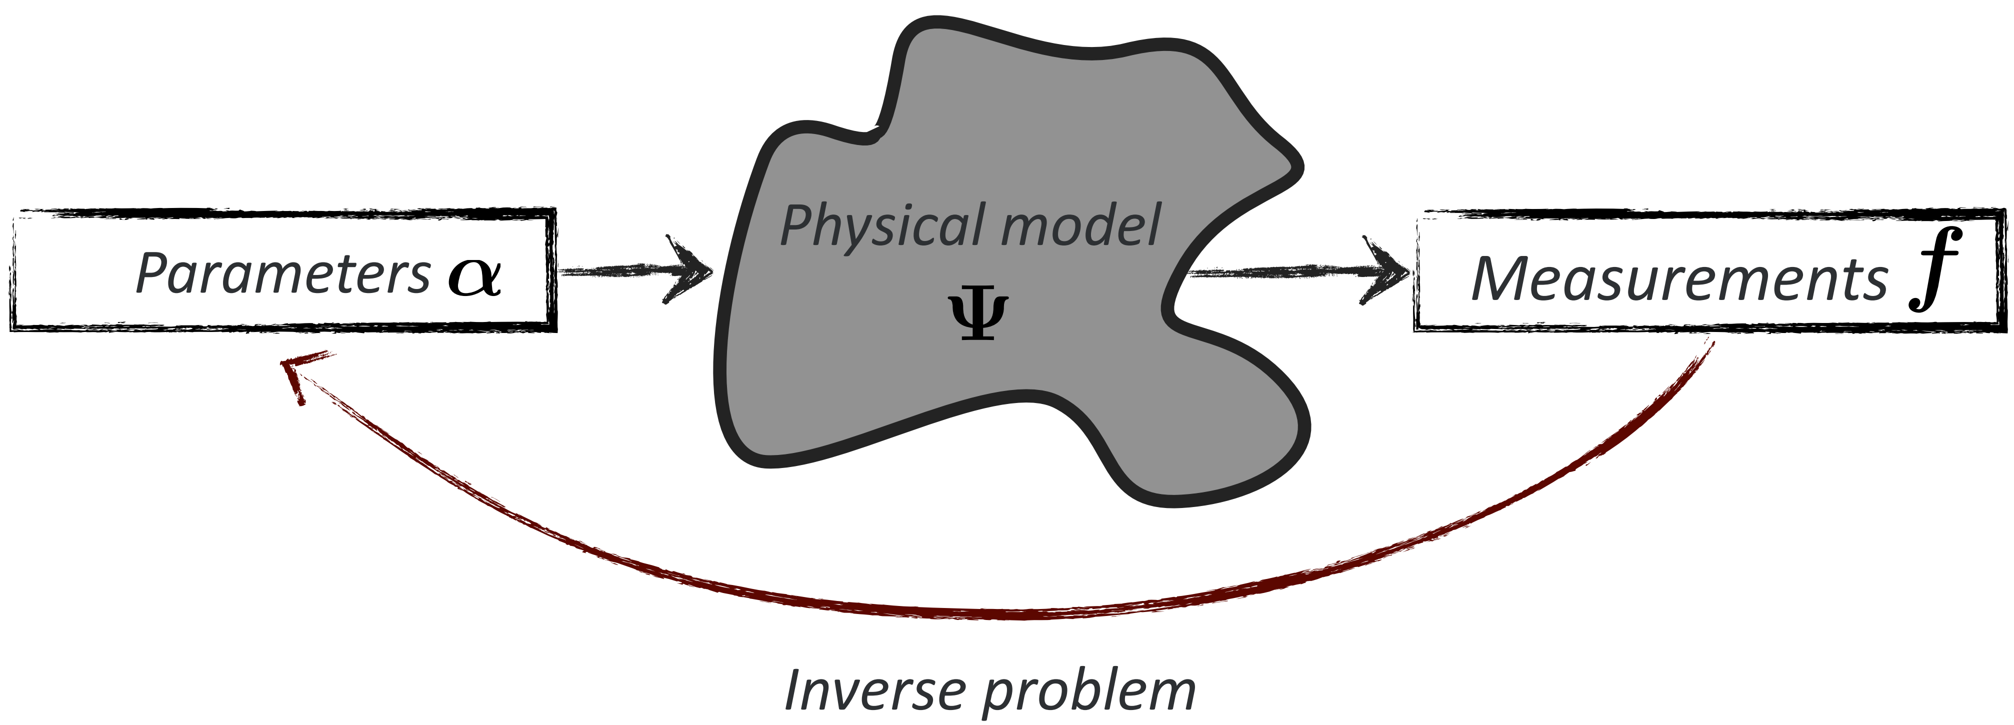

Inverse problem are variegated.
Signal processing problems are inverse problems.

A discrete model for physical fields

We consider a discretization of the physical field:

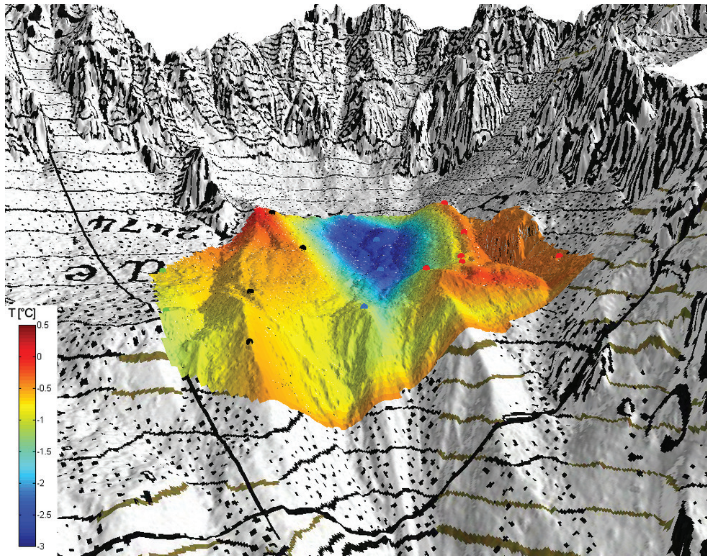

Environmental sensing

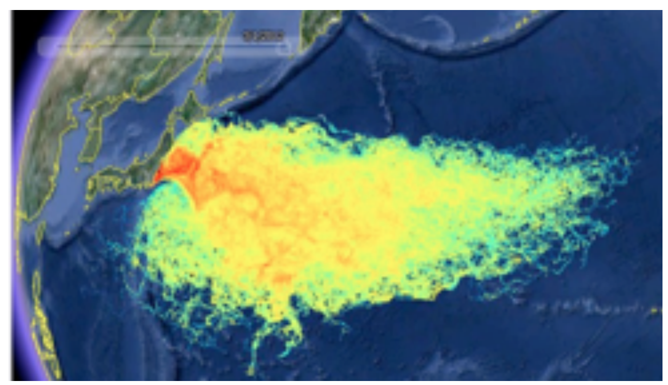

Pollution

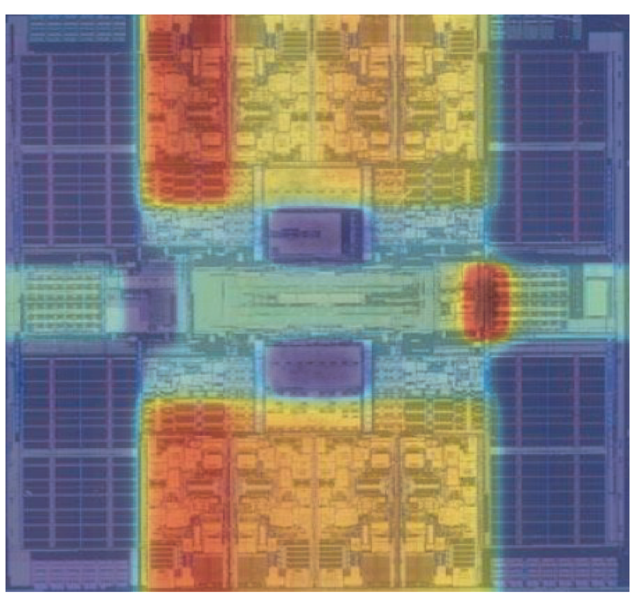

IC temperature

A discrete model for physical fields

We consider a discretization of the physical field:

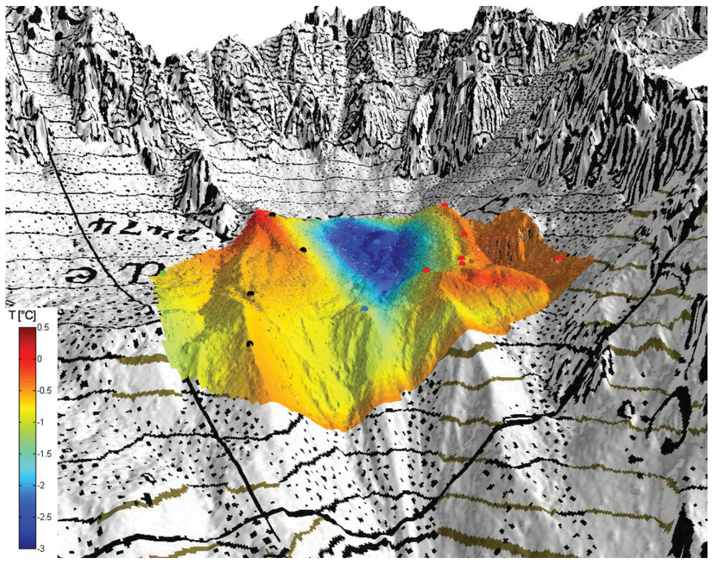

Environmental sensing

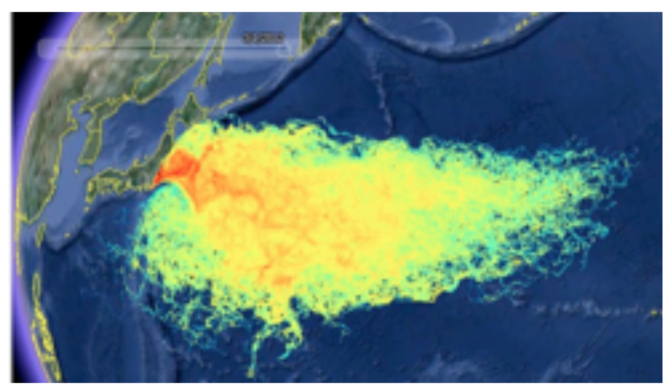

Pollution

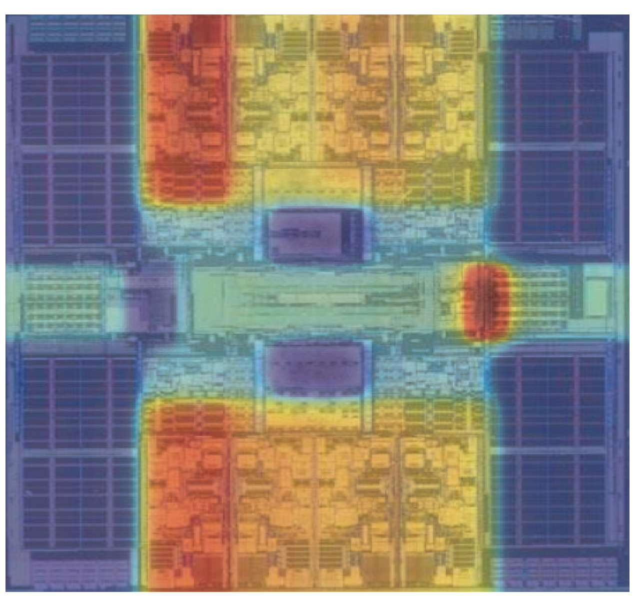

IC temperature

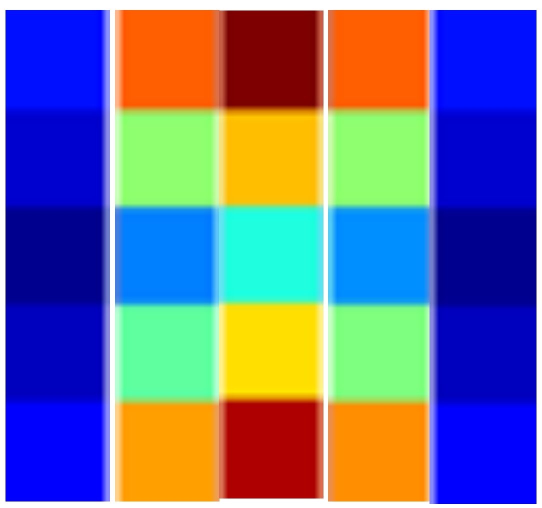

$\mathbf{f}$

A discrete model for physical fields

We consider a discretization of the physical field:

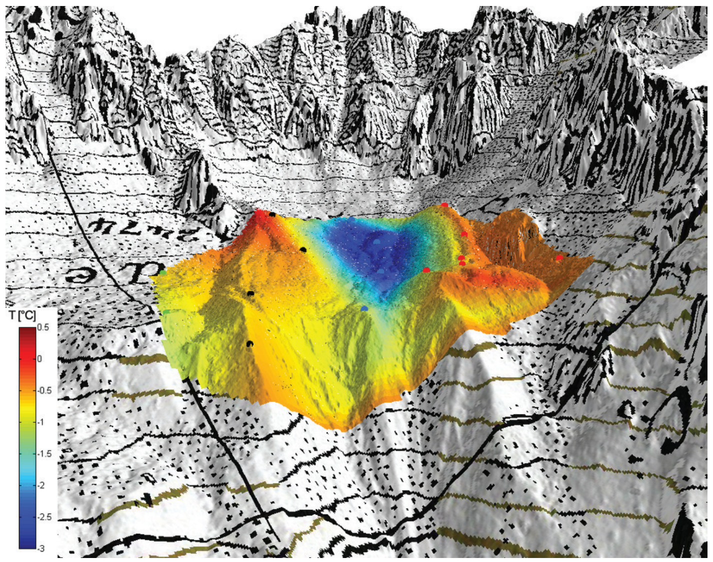

Environmental sensing

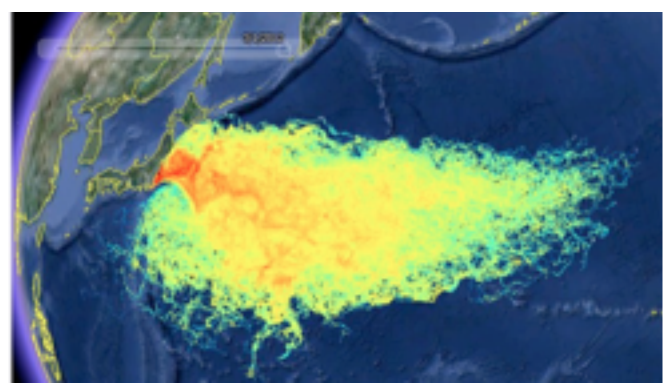

Pollution

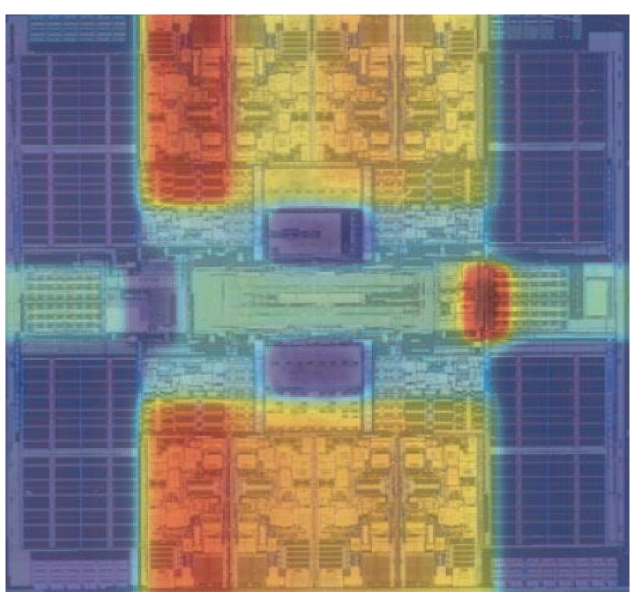

IC temperature

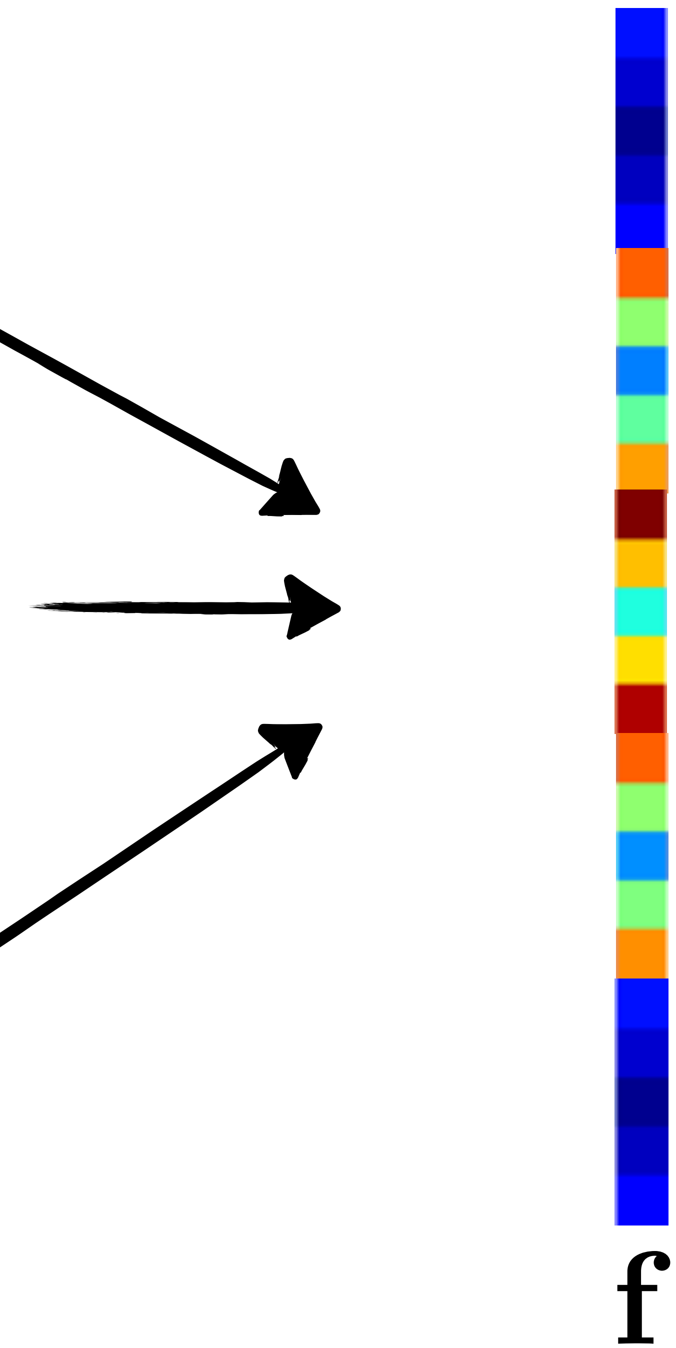

Solving a linear inverse problem

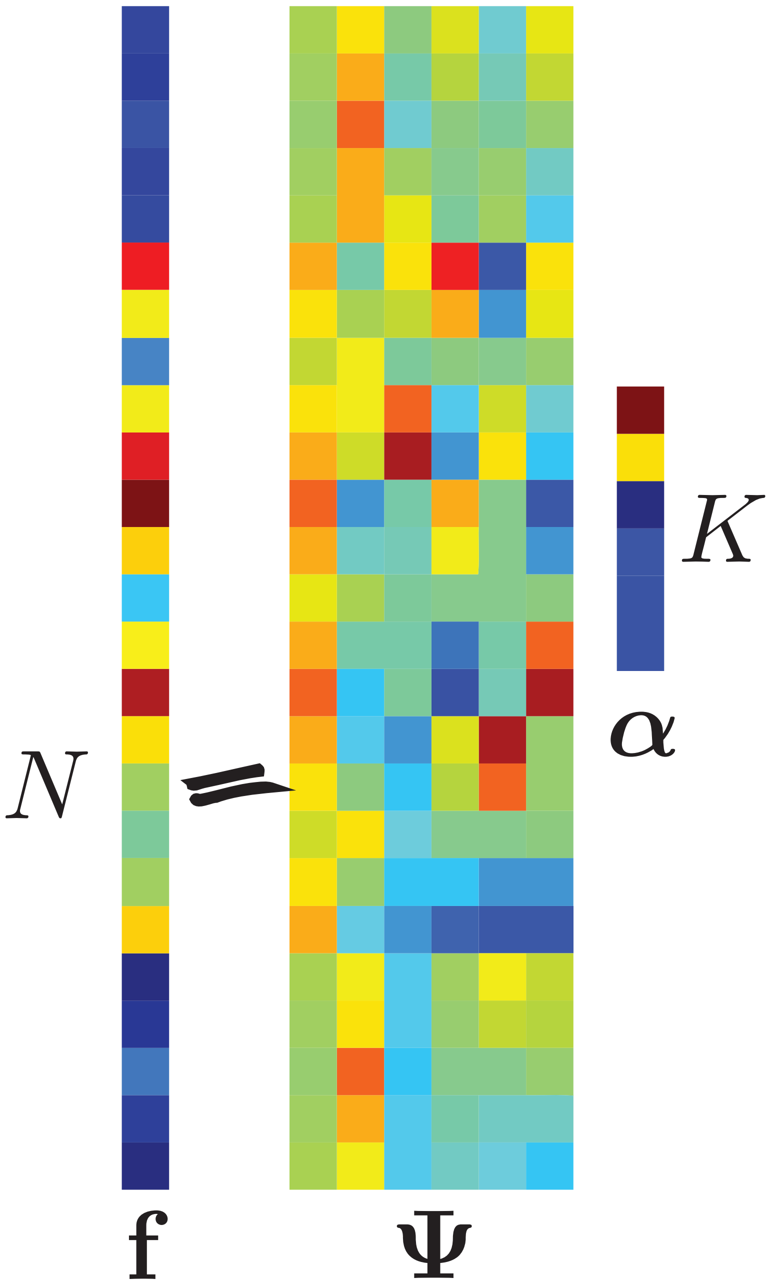

- Source localization.
- Data interpolation (low-dim representation).
- Boundary estimation.
- Parameter estimation.

Solving a linear inverse problem

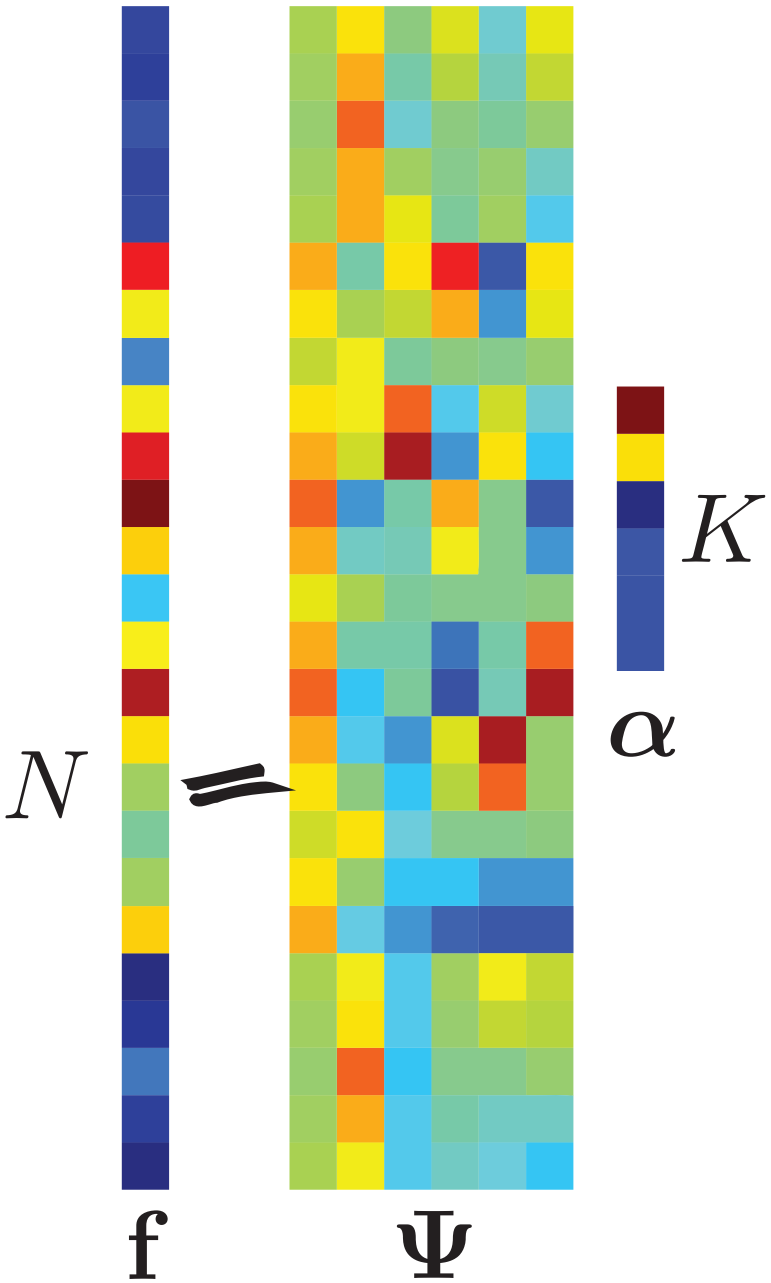

- Source localization.
- Data interpolation (low-dim representation).
- Boundary estimation.
- Parameter estimation.

We aim at precisely estimating α.

Sensing is expensive

Sensing is generally expensive and maybe technically difficult.

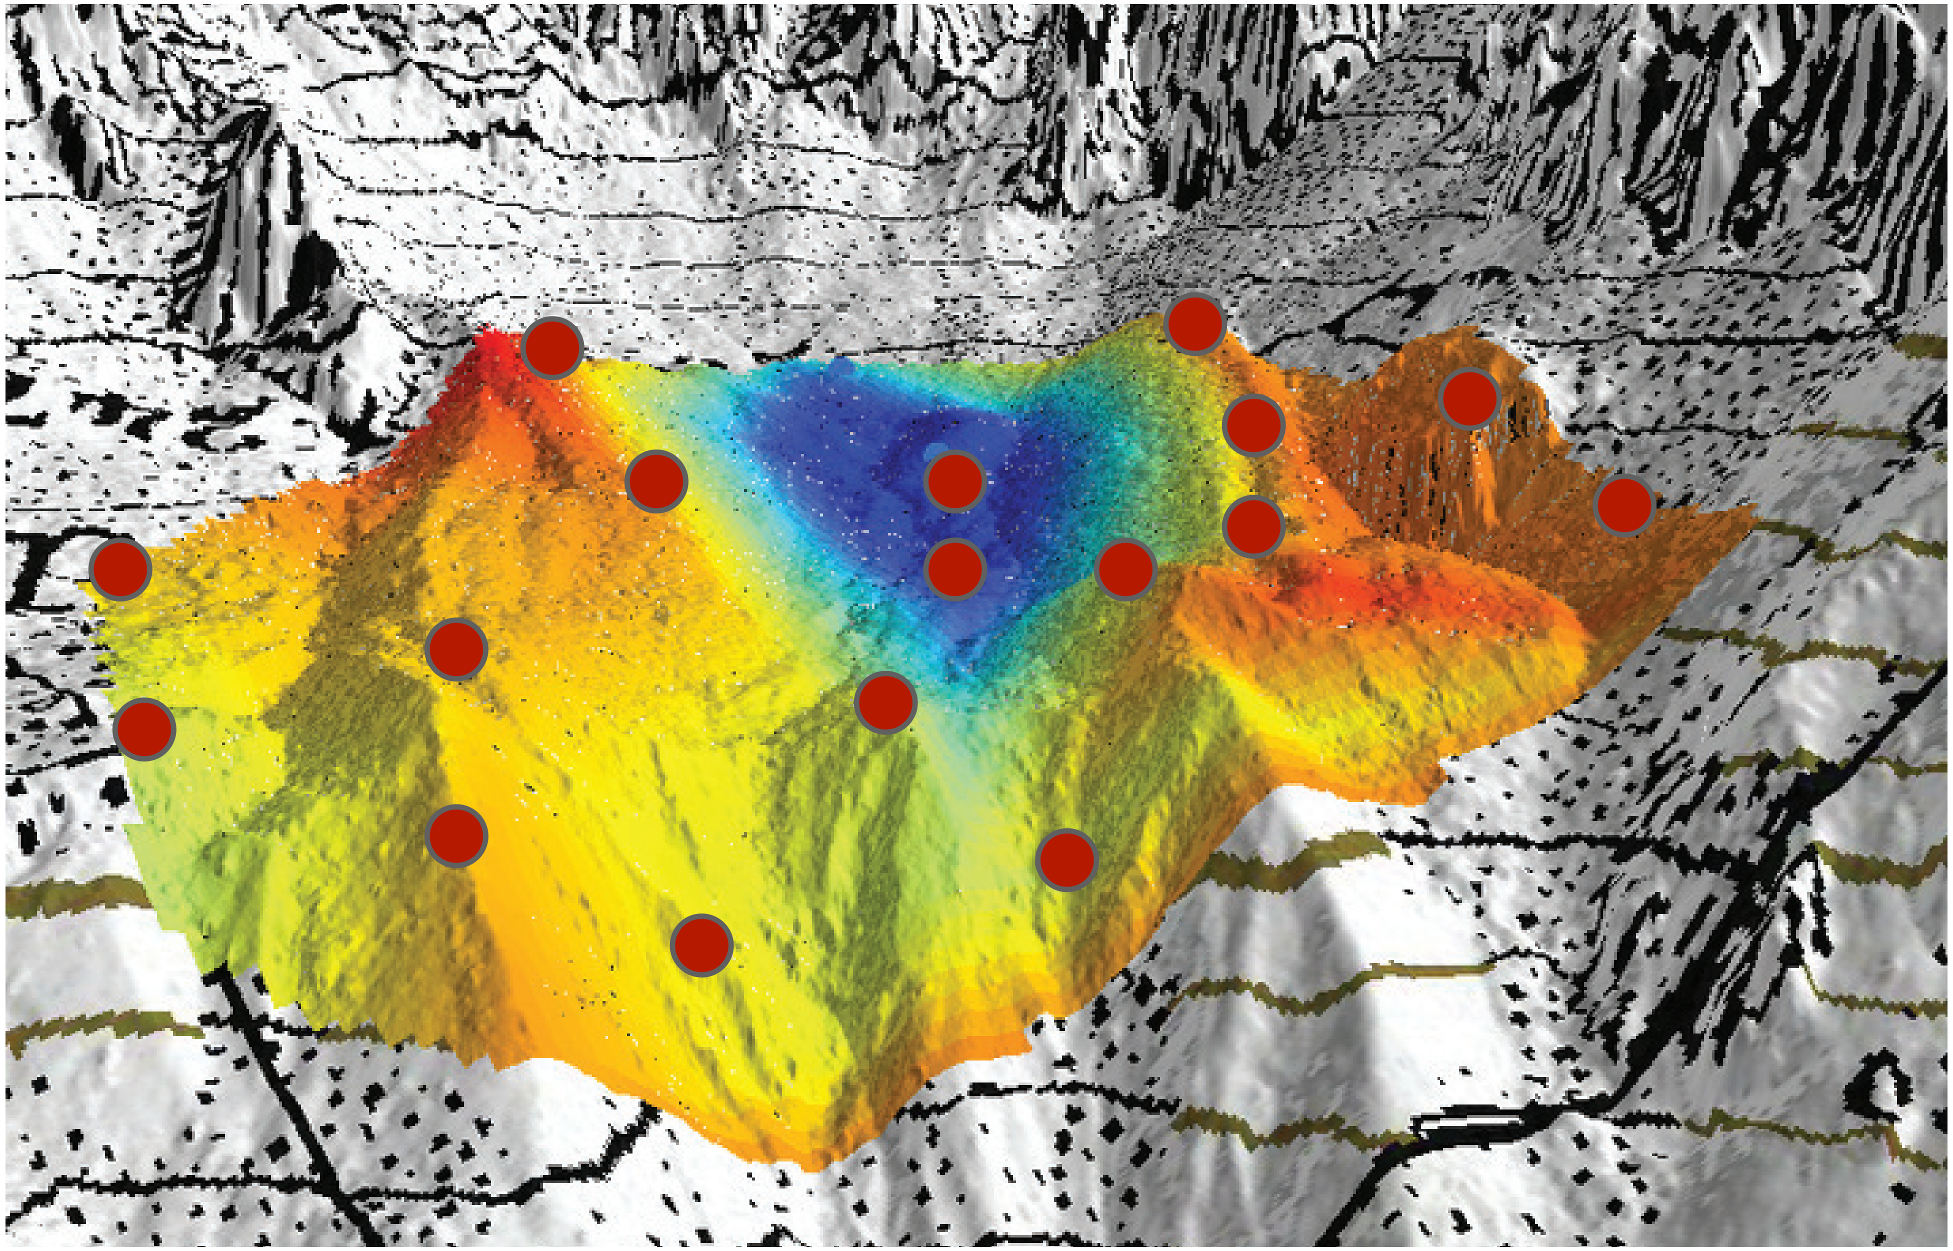

Where do we place the sensors to get the maximum information?

Where do we place the sensors?

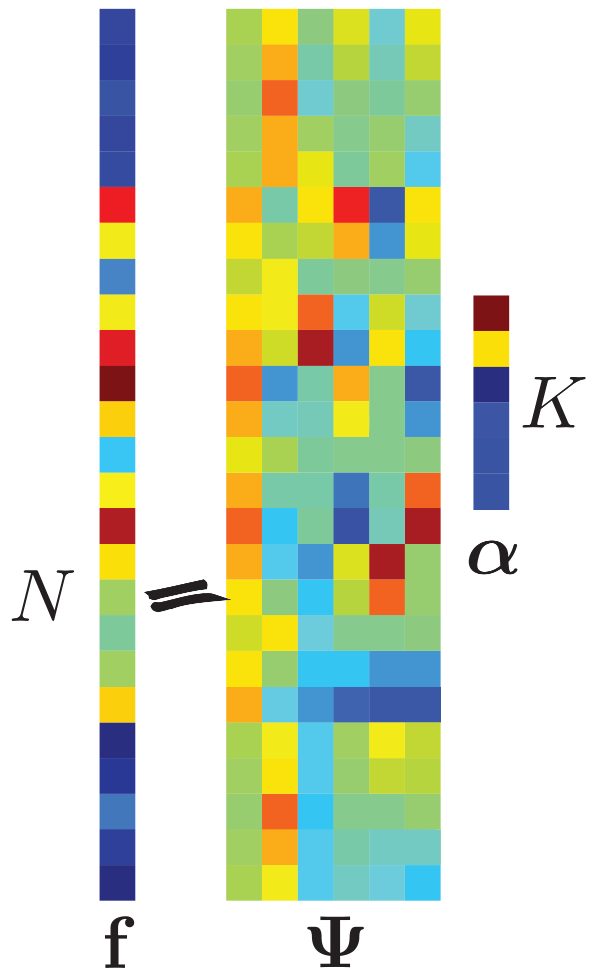

Where do we place the sensors?

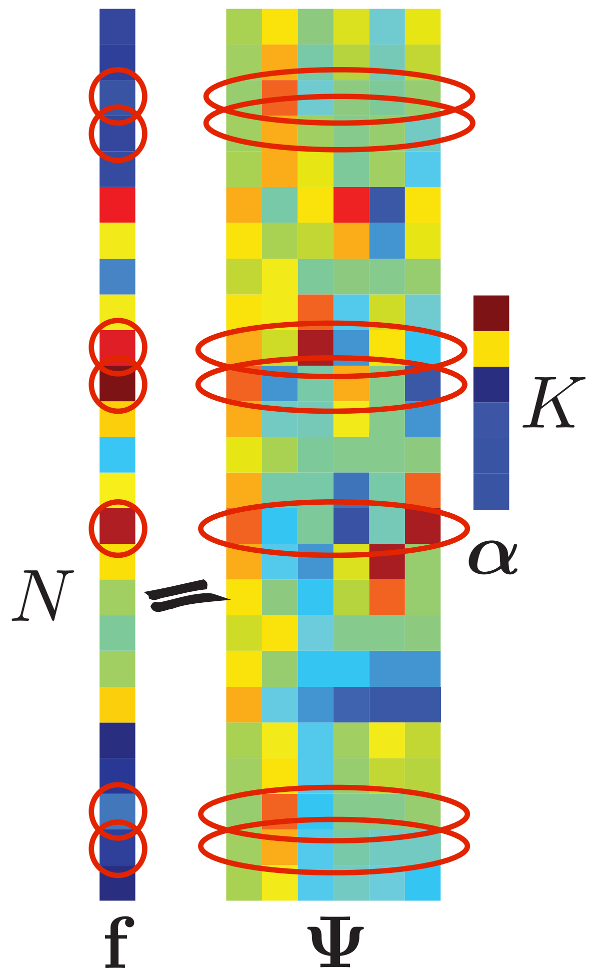

Where do we place the sensors?

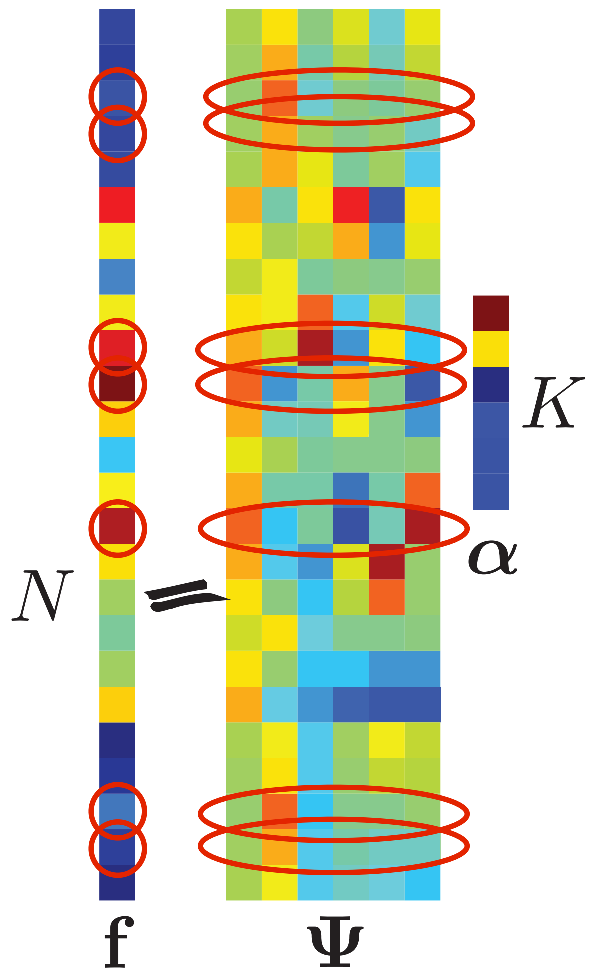

pick $L \ll N$ rows!

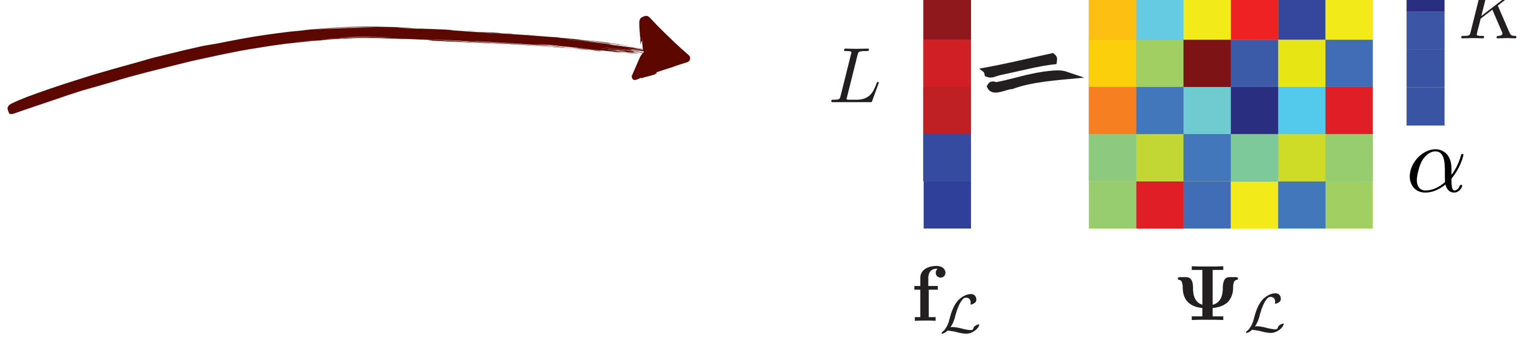

Sensor placement finds $\mathcal{L}$.

Where do we place the sensors?

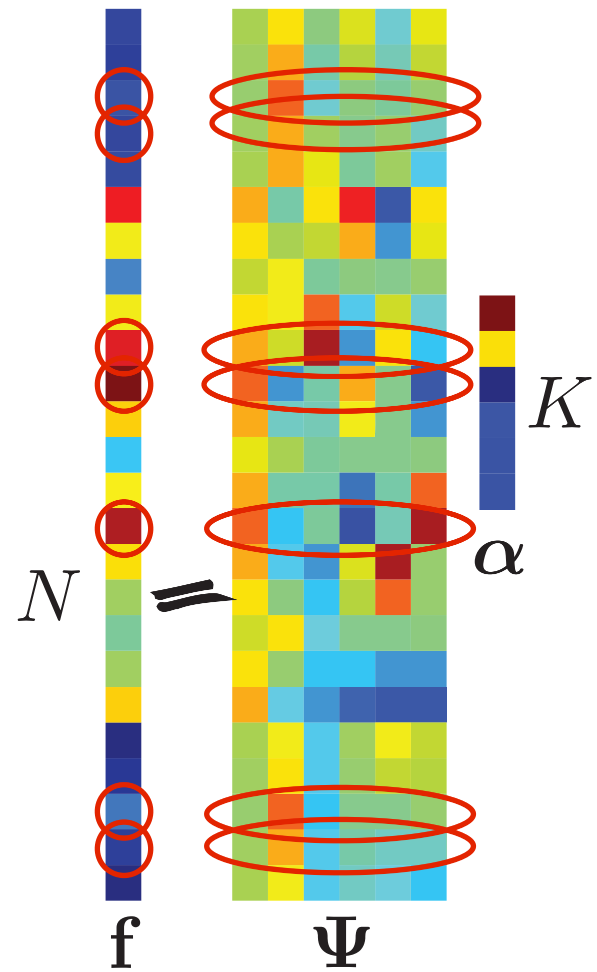

pick $L \ll N$ rows!

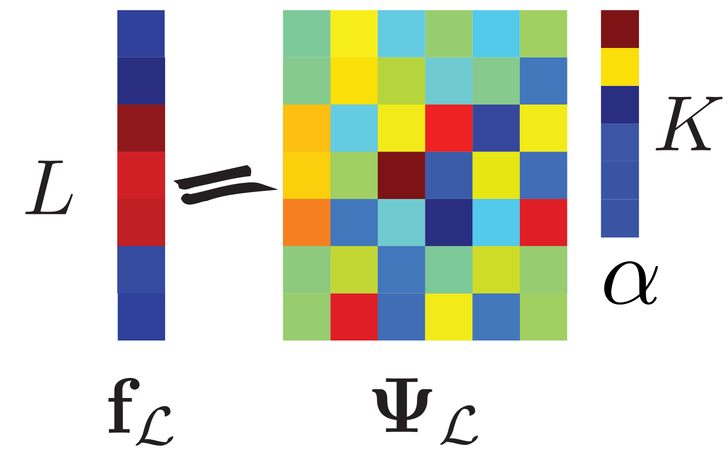

Sensor placement finds $\mathcal{L}$.

Problem: choose the placement $\mathcal{L}$ to minimize the MSE of $\boldsymbol{\alpha}$.

$$\mathcal{L} = \arg\min_{\mathcal{A}} \|\boldsymbol{\alpha} - \widehat{\boldsymbol{\alpha}}(\boldsymbol{f}_{\mathcal{A}})\|_2 \qquad |\mathcal{A}| = L$$

Where do we place the sensors?

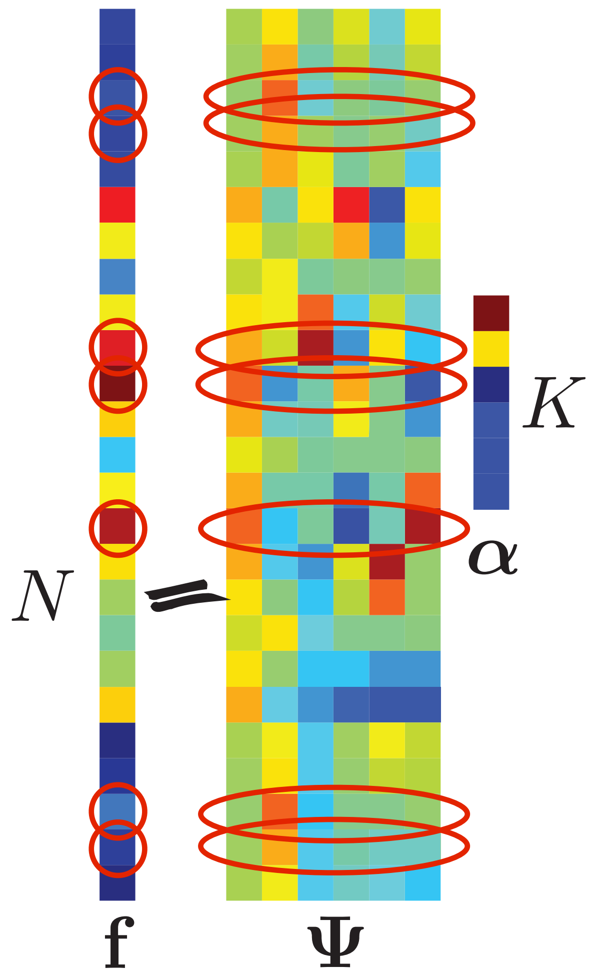

pick $L \ll N$ rows!

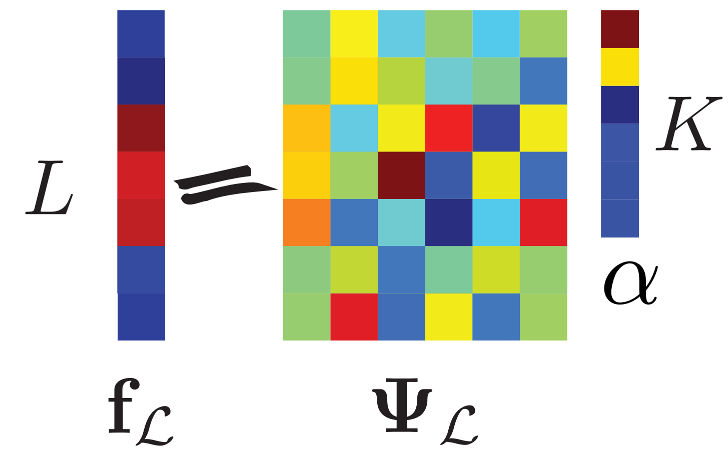

Sensor placement finds $\mathcal{L}$.

Problem: choose the placement $\mathcal{L}$ to minimize the MSE of $\boldsymbol{\alpha}$.

$$\mathcal{L} = \arg\min_{\mathcal{A}} \|\boldsymbol{\alpha} - \widehat{\boldsymbol{\alpha}}(\boldsymbol{f}_{\mathcal{A}})\|_2 \qquad |\mathcal{A}| = L$$

Subset selection: NP-hard! [Das 2008]

Classic approximation strategy: go greedy!

Greedy algorithms:
at each iteration, pick the best local choice.

- No guarantees about the distance between the greedy and the global optimal solution
- Polynomial time (if the cost function can be efficiently computed)

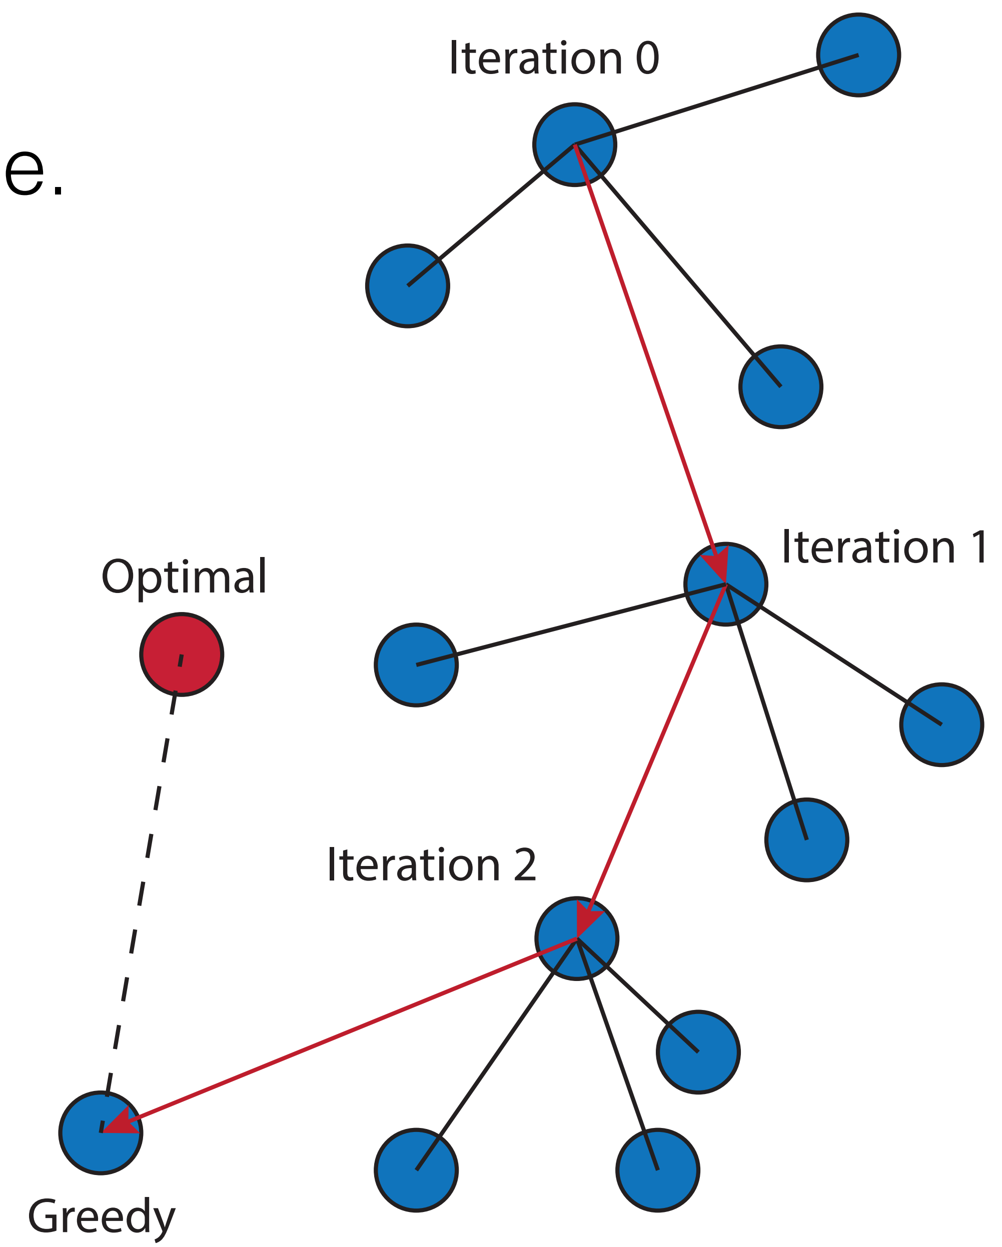

Classic approximation strategy: go greedy!

Greedy algorithms:
at each iteration, pick the best local choice.

- No guarantees about the distance between the greedy and the global optimal solution
- Polynomial time (if the cost function can be efficiently computed)

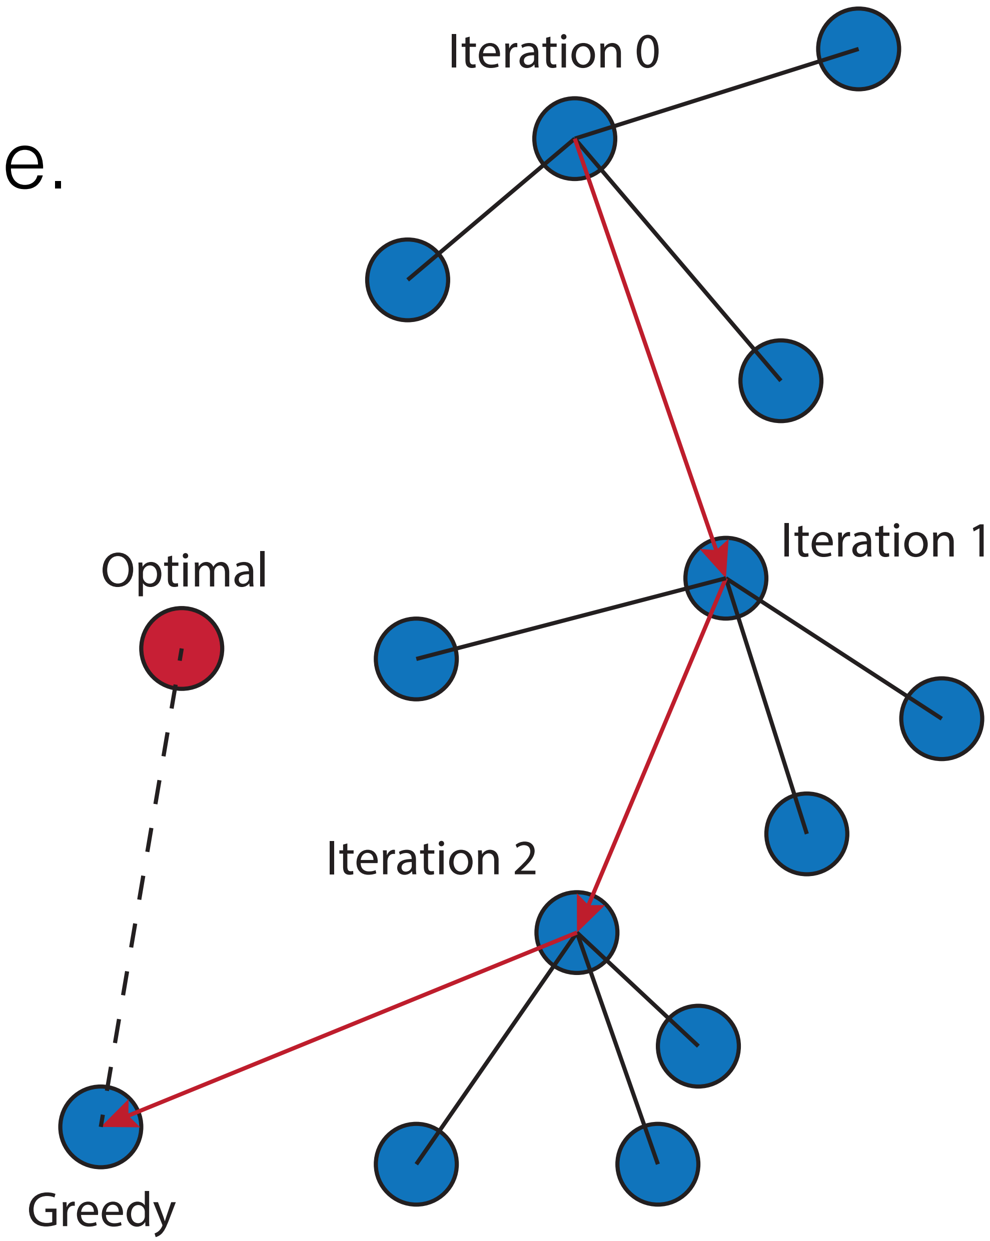

Can we optimize directly the MSE?
MSE greedy minimization is usually inefficient [Das 2008] and slow.

Frame Potential as a proxy of the MSE

Frame Potential (FP) as a cost function: $\mathrm{FP}(\boldsymbol{\Psi}_{\mathcal{L}}) = \sum_{i,j \in \mathcal{L}} |\langle \psi_i, \psi_j \rangle|^2.$

Frame Potential as a proxy of the MSE

Frame Potential (FP) as a cost function: $\mathrm{FP}(\boldsymbol{\Psi}_{\mathcal{L}}) = \sum_{i,j \in \mathcal{L}} |\langle \psi_i, \psi_j \rangle|^2.$

- FP is a measure of the closeness to orthogonality,
- The minimizers of the FP are UN tight frames [Casazza 2009],
- Minimizing the FP induces a minimization of the MSE.

Frame Potential as a proxy of the MSE

Frame Potential (FP) as a cost function: $\mathrm{FP}(\mathbf{\Psi}_{\mathcal{L}}) = \sum_{i,j \in \mathcal{L}} |\langle \psi_i, \psi_j \rangle|^2.$

- FP is a measure of the closeness to orthogonality,
- The minimizers of the FP are UN tight frames [Casazza 2009],
- Minimizing the FP induces a minimization of the MSE.

Physical interpretation [Fickus 2006]:
vectors repulse each other like electrons under the Coulomb force.

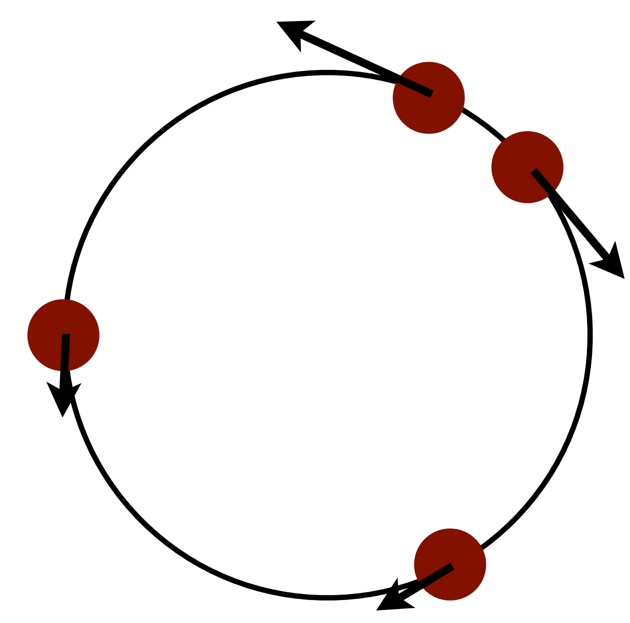

Frame Potential as a proxy of the MSE

Frame Potential (FP) as a cost function: $\mathrm{FP}(\boldsymbol{\Psi}_{\mathcal{L}}) = \sum_{i,j \in \mathcal{L}} |\langle \psi_i, \psi_j \rangle|^2.$

- FP is a measure of the closeness to orthogonality,
- The minimizers of the FP are UN tight frames [Casazza 2009],
- Minimizing the FP induces a minimization of the MSE.

Physical interpretation [Fickus 2006]:
vectors repulse each other like electrons under the Coulomb force.

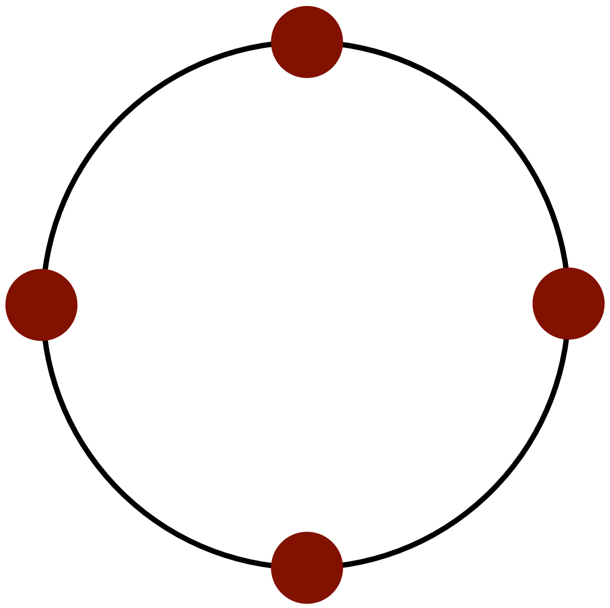

Greedy algorithms and submodularity

Theorem [Nemhauser 1978]: Consider a greedy algorithm maximizing a submodular, normalized, monotonically increasing set function $g(\cdot)$.
Then, the greedy solution is near-optimal:

$$g(S_{\text{opt}}) \geq g(S_{\text{greedy}}) \geq \left(1 - \frac{1}{e}\right) g(S_{\text{opt}})$$

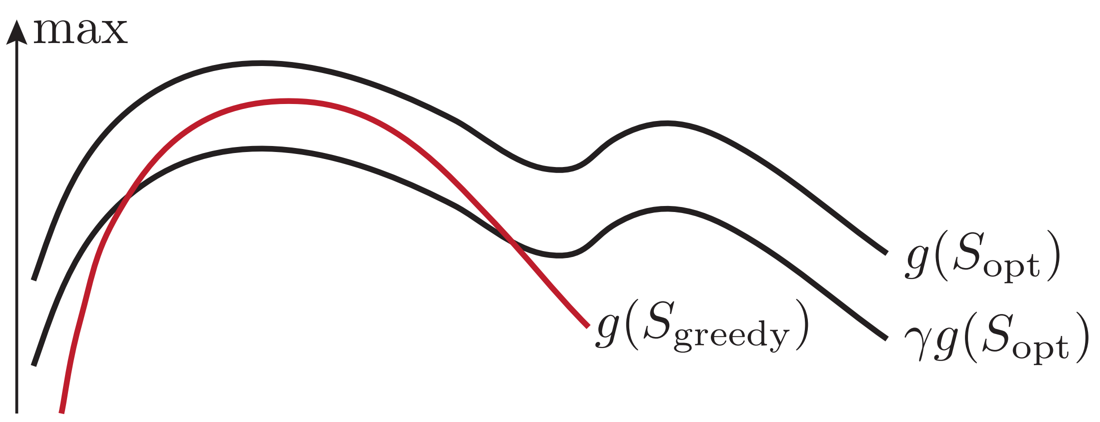

- Submodularity ~ concept of diminishing returns.

Greedy algorithms and submodularity

Theorem [Nemhauser 1978]: Consider a greedy algorithm maximizing a submodular, normalized, monotonically increasing set function $g(\cdot)$.
Then, the greedy solution is near-optimal:

$$g(S_{\text{opt}}) \geq g(S_{\text{greedy}}) \geq \left(1 - \frac{1}{e}\right) g(S_{\text{opt}})$$

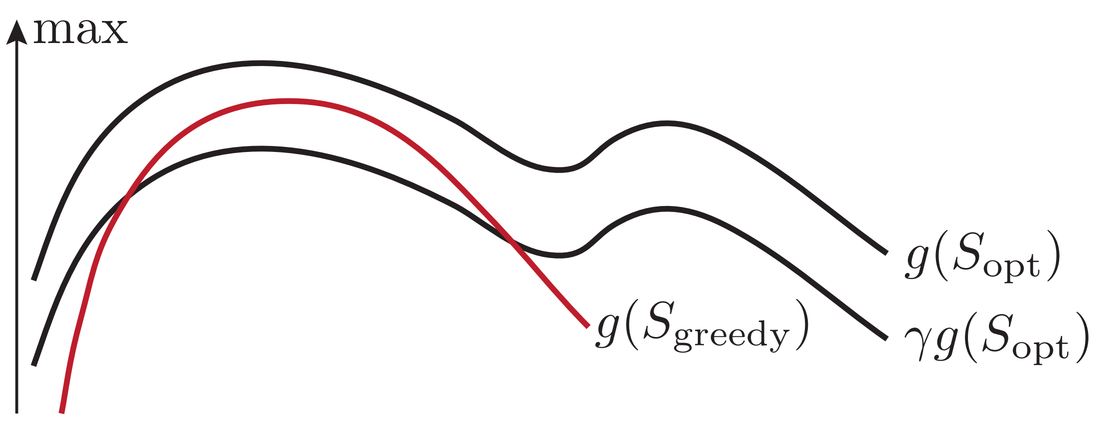

- Submodularity ~ concept of diminishing returns.
- Frame Potential is submodular

FrameSense

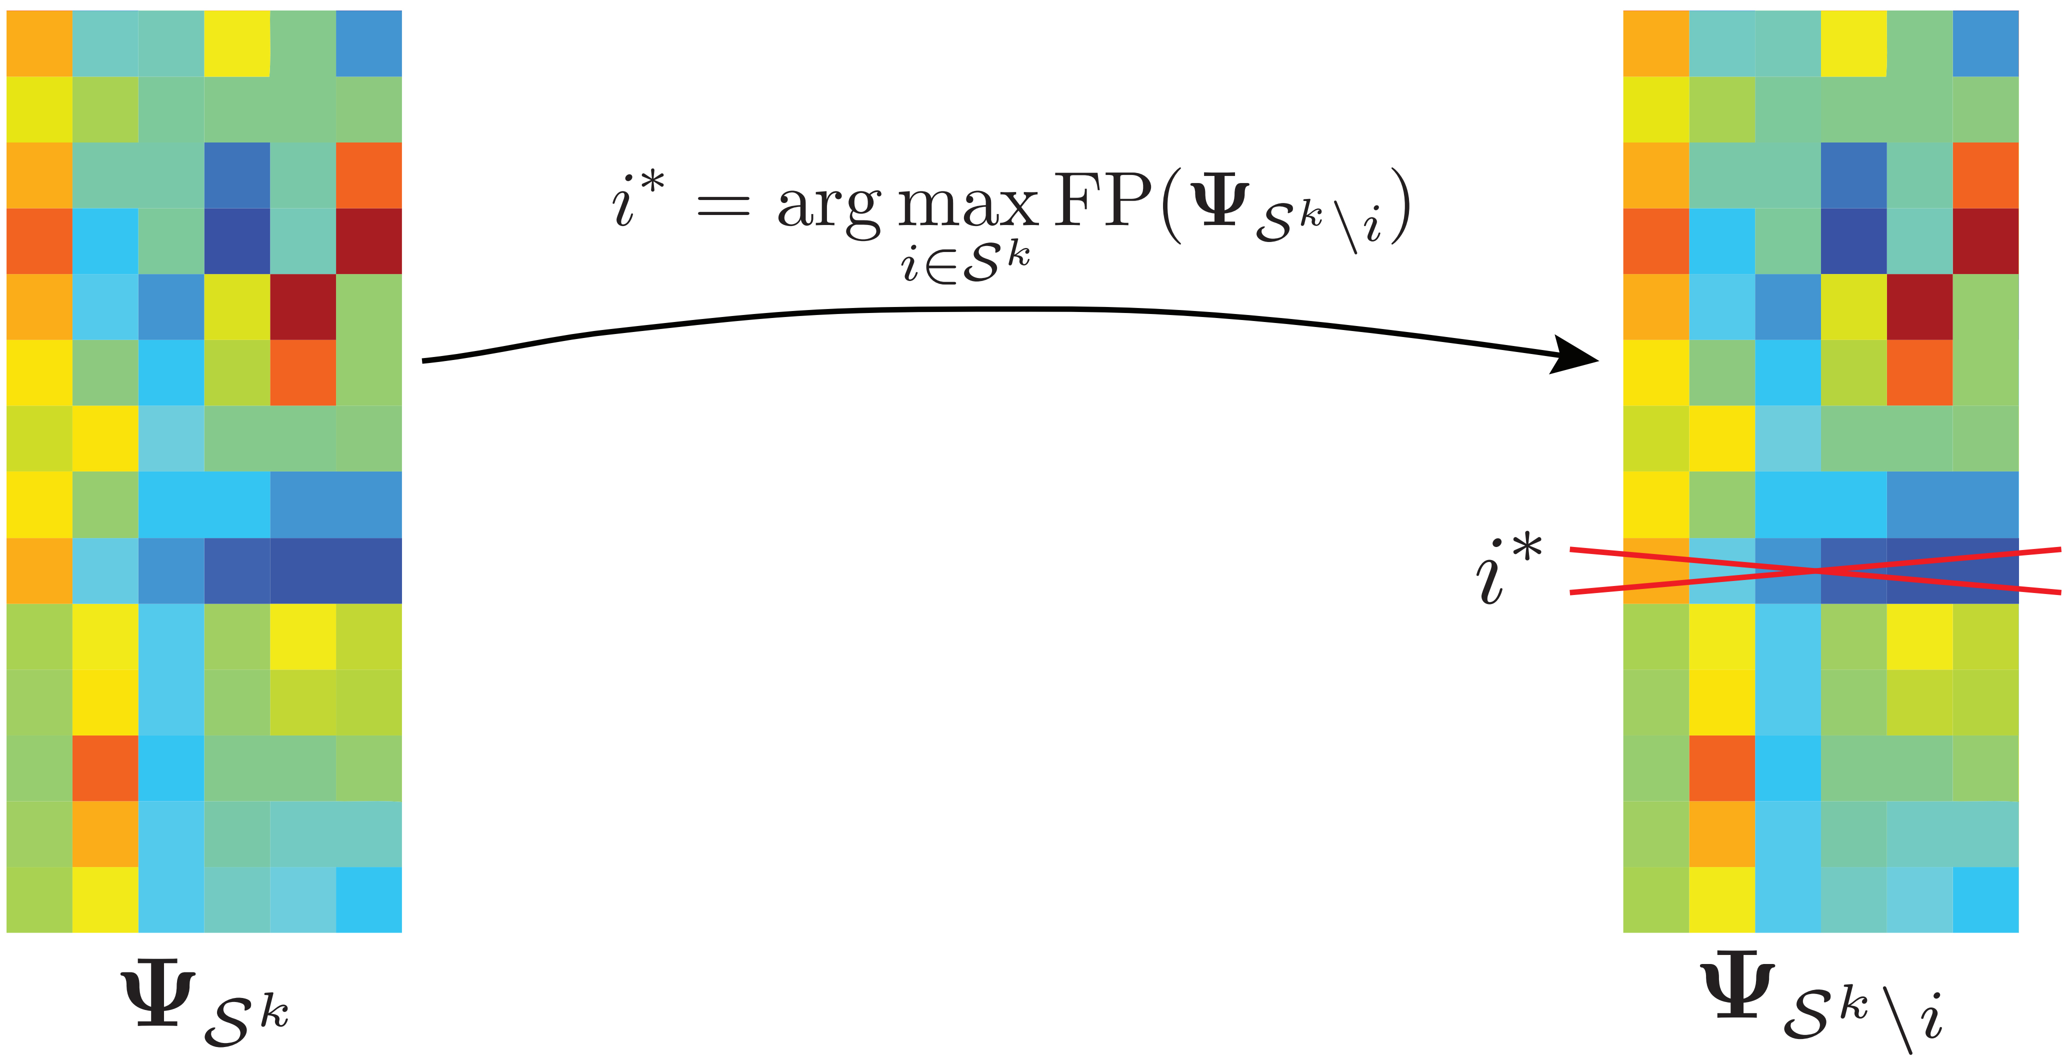

Greedy worst-out sensor selection:

- At the k-th iteration, we remove the row maximizing the FP of $\boldsymbol{\Psi}_{\mathcal{S}^k}$,
- After $N - L$ iterations, the sensor placement is $\mathcal{L} = \mathcal{S}^{N-L}$.

FrameSense is near-optimal w.r.t. MSE

Our strategy:

- FP is submodular,
- FrameSense is near-optimal w.r.t. FP,
- We derive LB and UB of the MSE w.r.t. FP,
- FrameSense is near-optimal w.r.t. MSE.

FrameSense is near-optimal w.r.t. MSE

Our strategy:

- FP is submodular,
- FrameSense is near-optimal w.r.t. FP,
- We derive LB and UB of the MSE w.r.t. FP,
- FrameSense is near-optimal w.r.t. MSE.

Theorem [Ranieri C. V. 2013]:
Given $\boldsymbol{\Psi} \in \mathbb{R}^{N \times K}$ and $L \geq K$ sensors, FrameSense is near-optimal w.r.t. MSE (under certain conditions on $\boldsymbol{\Psi}$):

$$\mathrm{MSE}(\boldsymbol{\Psi}_{\mathcal{L}}) \leq \gamma\, \mathrm{MSE}(\boldsymbol{\Psi}_{\mathrm{opt}}),$$

where the approximation factor depends on the spectrum and the norms of the rows of $\boldsymbol{\Psi}$.

$\boldsymbol{\Psi}_{\mathcal{L}}$ and $\boldsymbol{\Psi}_{\mathrm{opt}}$ are the greedy and the optimal solution, respectively.

FrameSense is near-optimal w.r.t. MSE

Our strategy:

- FP is submodular,
- FrameSense is near-optimal w.r.t. FP,
- We derive LB and UB of the MSE w.r.t. FP,
- FrameSense is near-optimal w.r.t. MSE.

Theorem

Given $\boldsymbol{\Psi}$ … Sense is
near-opti … $\boldsymbol{\Psi}$):

where the … ctrum and
the norms

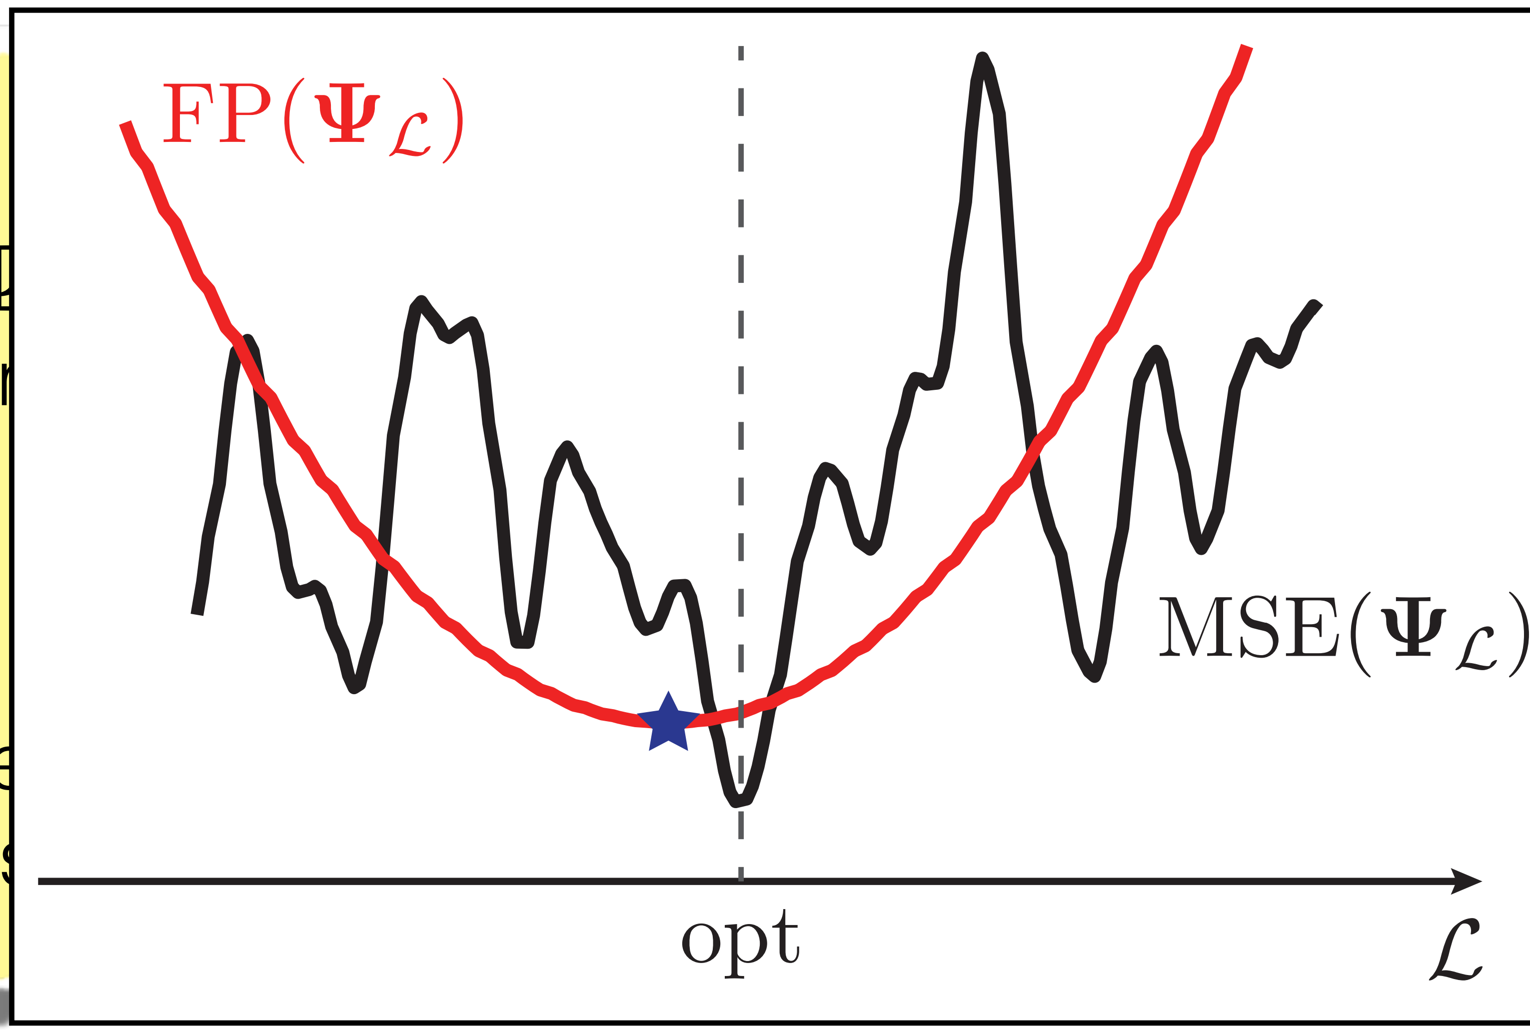

$\boldsymbol{\Psi}_{\mathcal{L}}$ and $\boldsymbol{\Psi}_{\mathrm{opt}}$ are the greedy and the optimal solution, respectively.

Proposed algorithm

FrameSense (polytime, no heuristics, guarantees):

- Greedy algorithm optimizing the Frame Potential (FP),
- Near-optimal w.r.t. the MSE,
- State-of-the-art performance w.r.t. the MSE,
- Low computational complexity.

Sensing the temperature of a processor

One of the applications where FrameSense shines…

Application: temperature sensing

A modern 8 core microprocessor

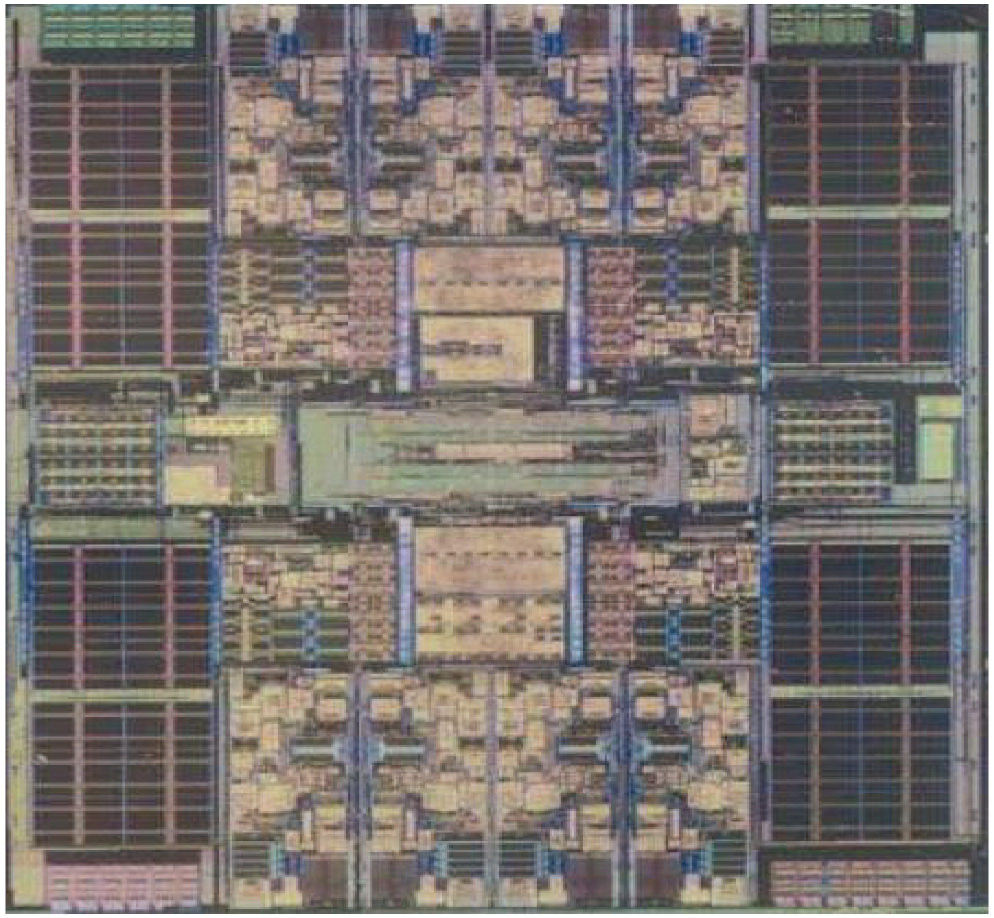

Application: temperature sensing

A modern 8 core microprocessor

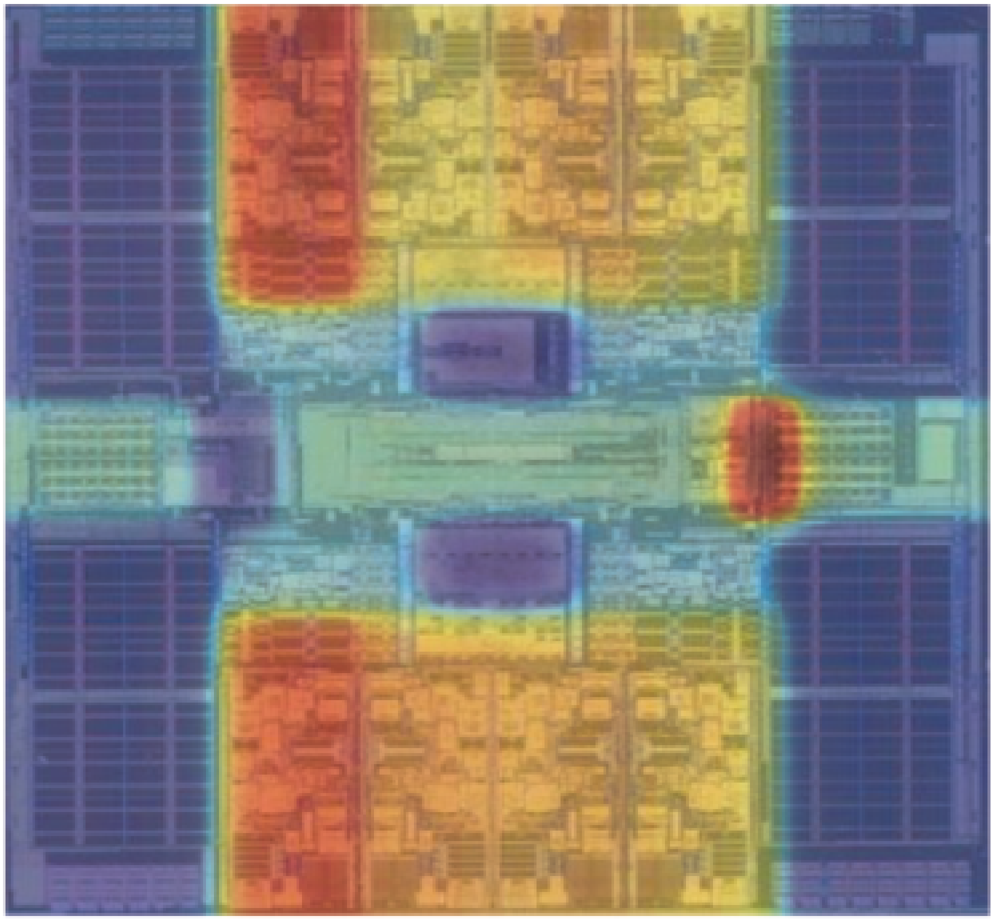

Application: temperature sensing

A modern 8 core microprocessor

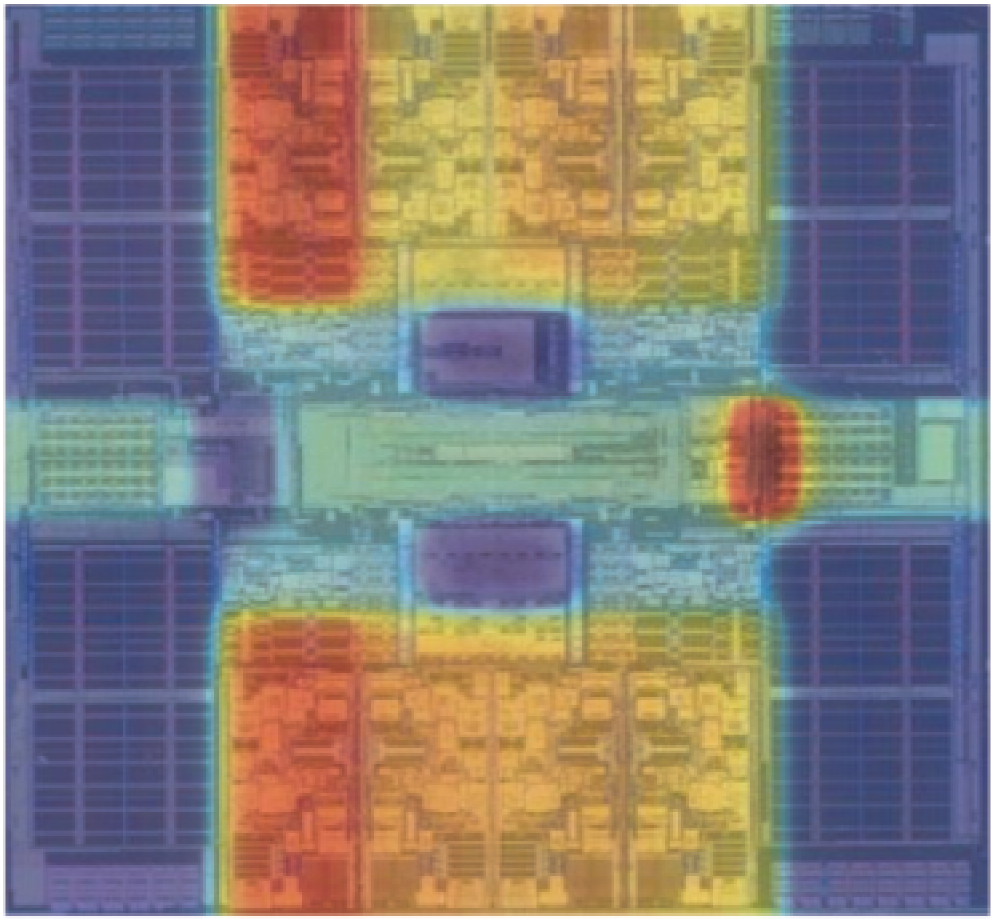

- Thermal stress: failures, reduced performance, increased power consumption, mechanical stress.
- Temperature information is desirable to optimize workload.
- Temperature cannot be sensed everywhere.

Problem statement

Objectives:

- Design an algorithm to recover the entire thermal map from few measurements.
- Design a sensor placement algorithm to minimize the reconstruction error.

Problem statement

Objectives:

- Design an algorithm to recover the entire thermal map from few measurements.
- Design a sensor placement algorithm to minimize the reconstruction error.

Given: a set of thermal distributions, representing the workload of the processor.

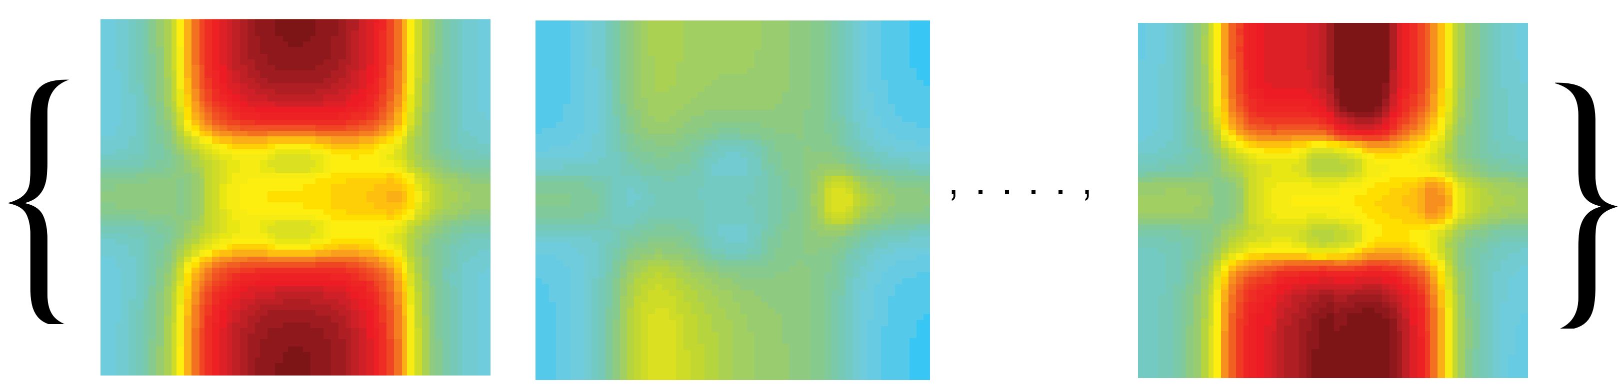

Low-dimensional linear model

We learn $\boldsymbol{\Psi}$ using PCA [Ranieri V.C.A.V. 2012]

$\boldsymbol{f} = \alpha_1$ 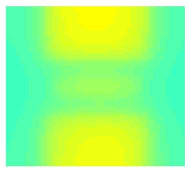 $+\alpha_2$ 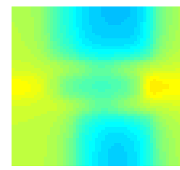 $+\alpha_3$ 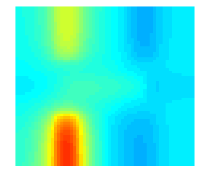 $+ \ldots$

$+\alpha_4$ 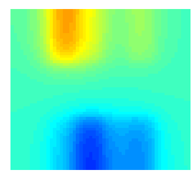 $+ \ldots + \alpha_K$ 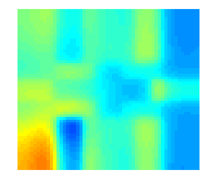

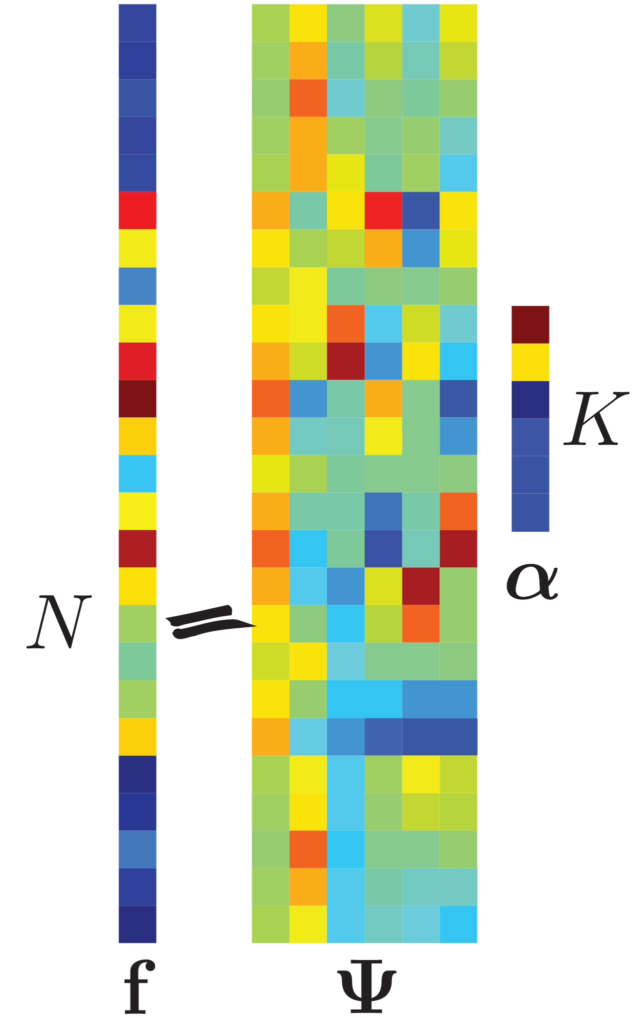

Low-dimensional linear model

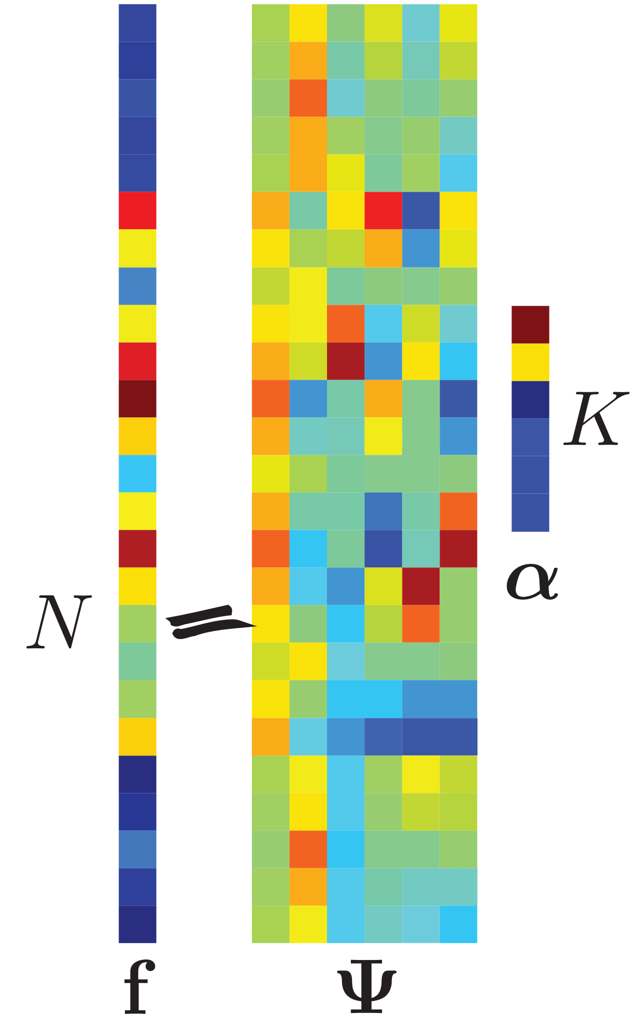

We learn $\boldsymbol{\Psi}$ using PCA [Ranieri V.C.A.V. 2012]

$\boldsymbol{f} = \alpha_1$ 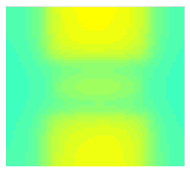 $+\alpha_2$ 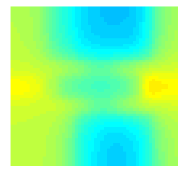 $+\alpha_3$ 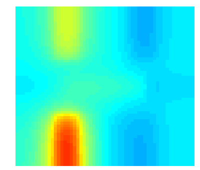 $+ \ldots$

$+\alpha_4$ 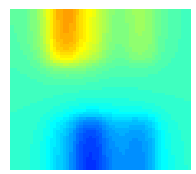 $+ \ldots + \alpha_K$ 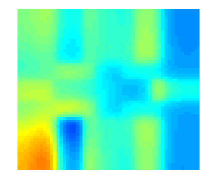

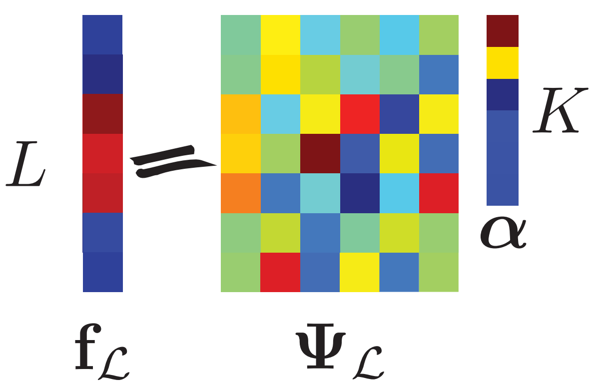

Recover the parameters from few measurements to recover the thermal map.

Low-dimensional linear model

We learn $\boldsymbol{\Psi}$ using PCA [Ranieri V.C.A.V. 2012]

$\boldsymbol{f} = \alpha_1$ 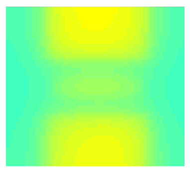 $+\alpha_2$ 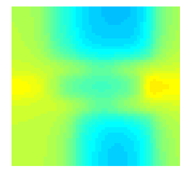 $+\alpha_3$ 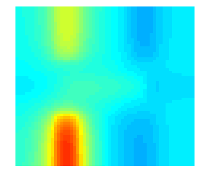 $+ \ldots$

$+\alpha_4$ 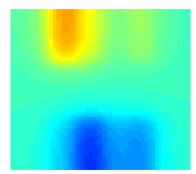 $+ \ldots + \alpha_K$ 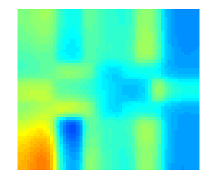

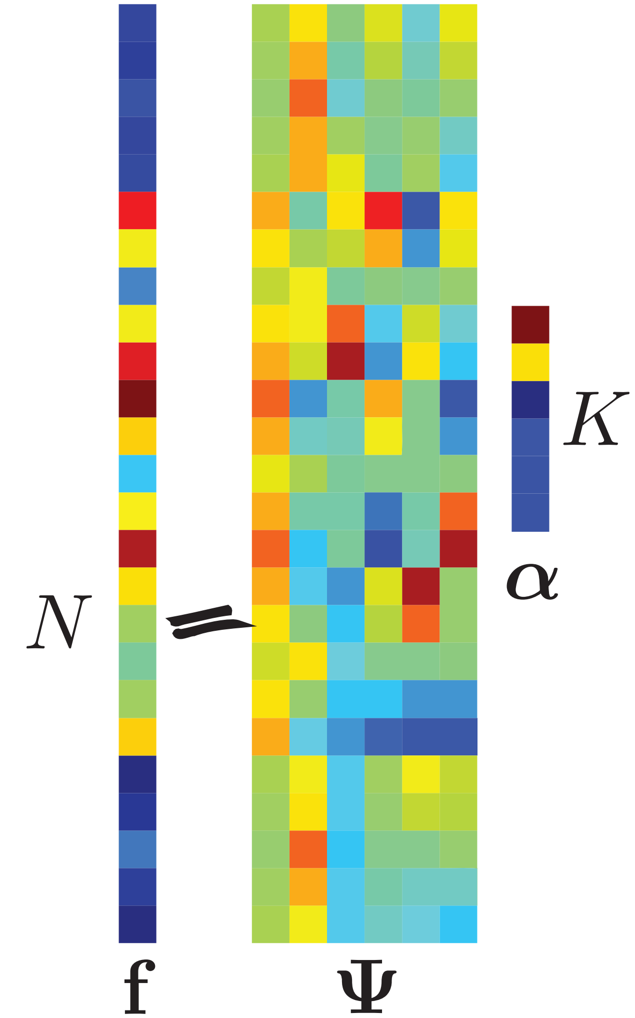

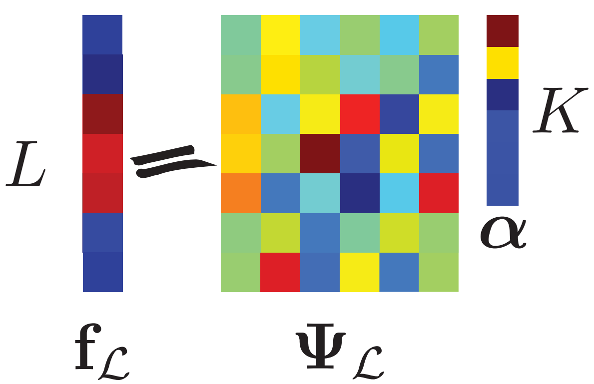

Recover the parameters from few measurements to recover the thermal map.

We use FrameSense to place the sensors.
K optimized given the noise level.

Performance evaluation

Reconstruction results on the 8-cores Niagara.

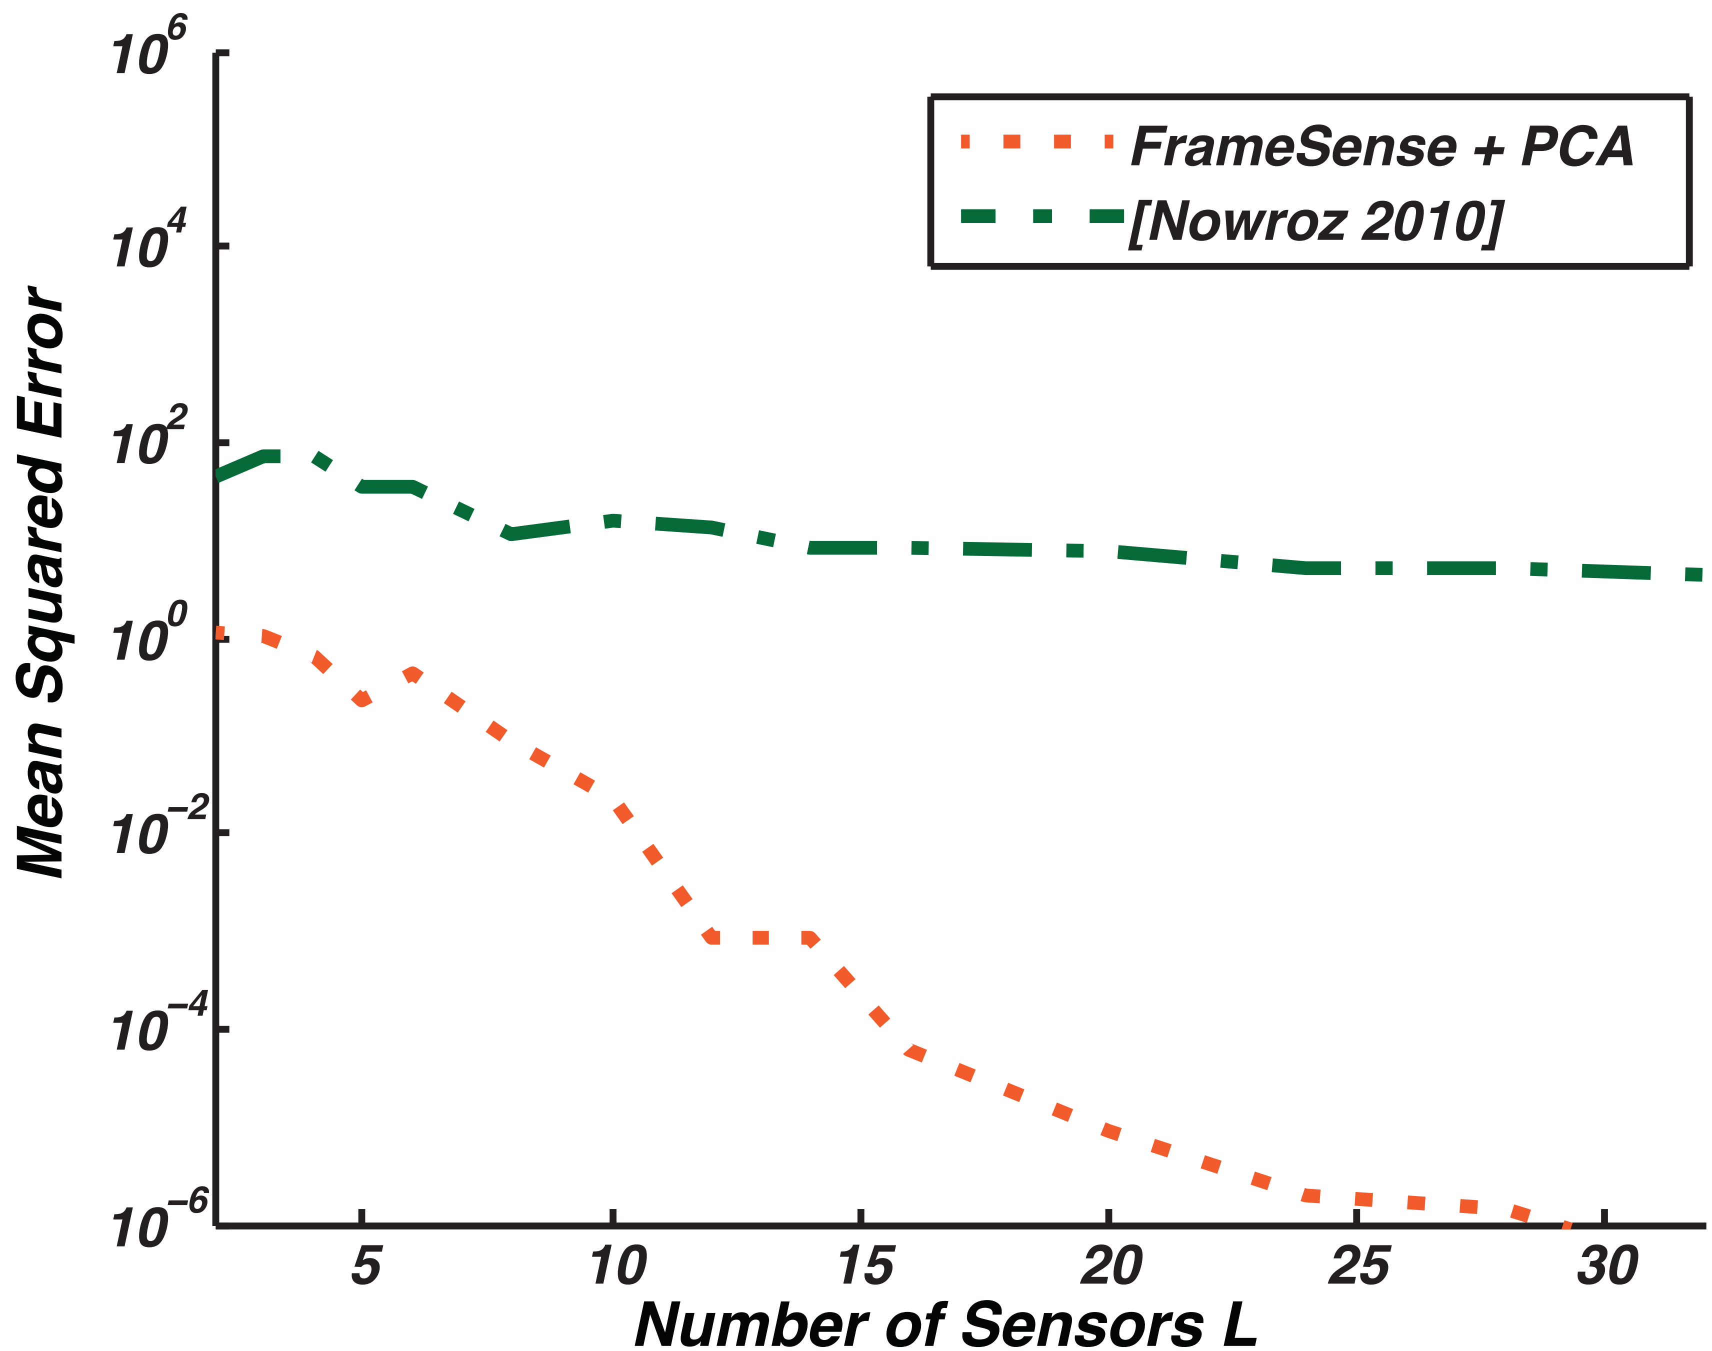

Performance evaluation

Reconstruction results on the 8-cores Niagara.

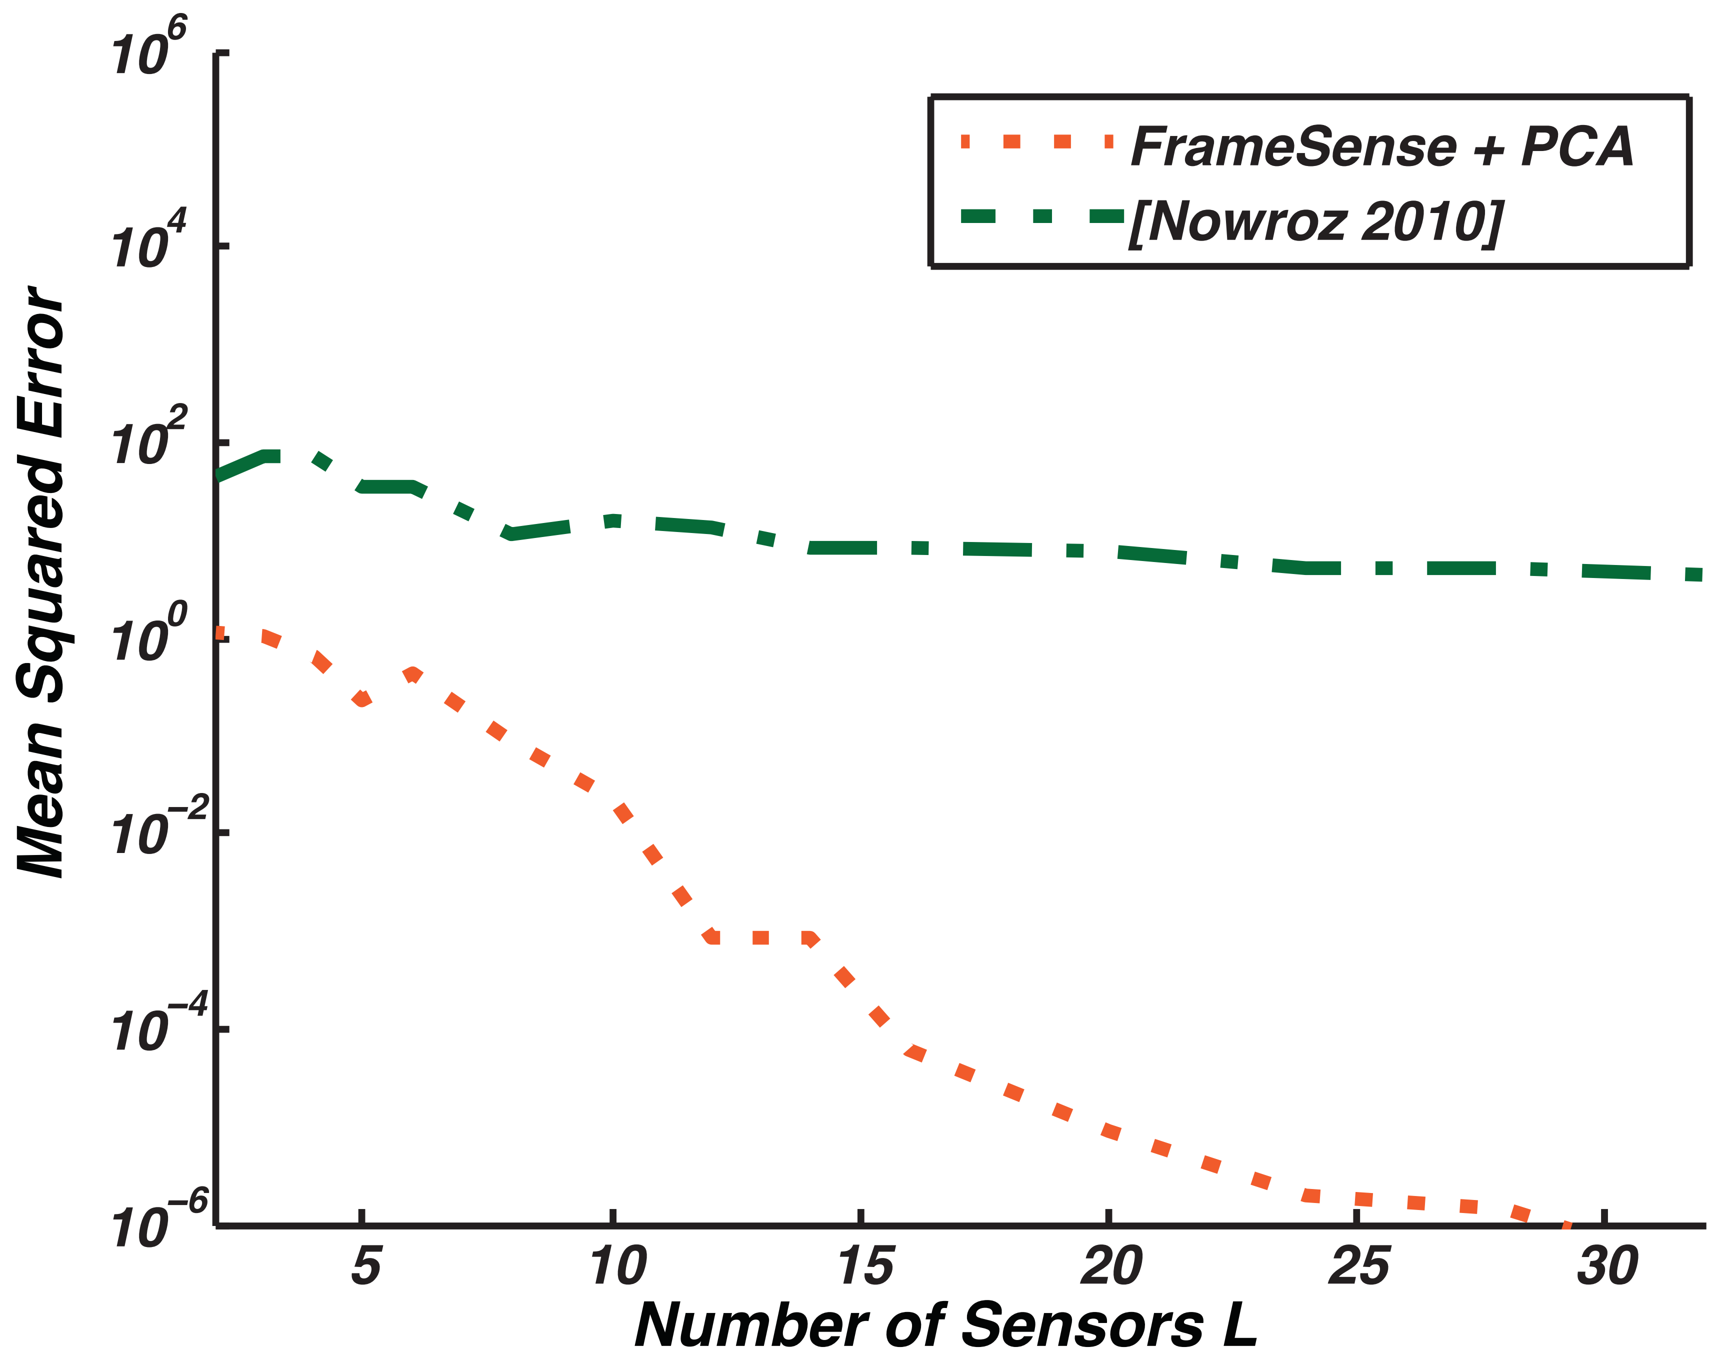

Similar results on a 64-cores STM architecture.

Results and future work

FrameSense (TSP 2014):

- A greedy algorithm based on the frame potential,
- First near-optimal algorithm w.r.t. MSE,
- Computationally efficient,
- State-of-the-art performance.

Applications

- Thermal monitoring of many-core processors (DAC 2012, TCOMP 2015),
- DASS: distributed adaptive sampling scheduling (TCOMM 2014).

Extensions

- Source placement for linear forward problems,
- Union of subspaces (EUSIPCO 2014).

Future work

- Sensor optimization for control theory,
- Tomographic sensing.

Thanks for your attention!
Questions?