Finite Range Scalar Quantization for Compressive Sensing

Jason N. Laska laska@rice.edu
Petros T. Boufounos petrosb@merl.com
Richard G. Baraniuk richb@rice.edu

Presenter: Volkan Cevher vcevher@rice.edu

SAMPLING AND QUANTIZATION

Classical Sampling and Quantization

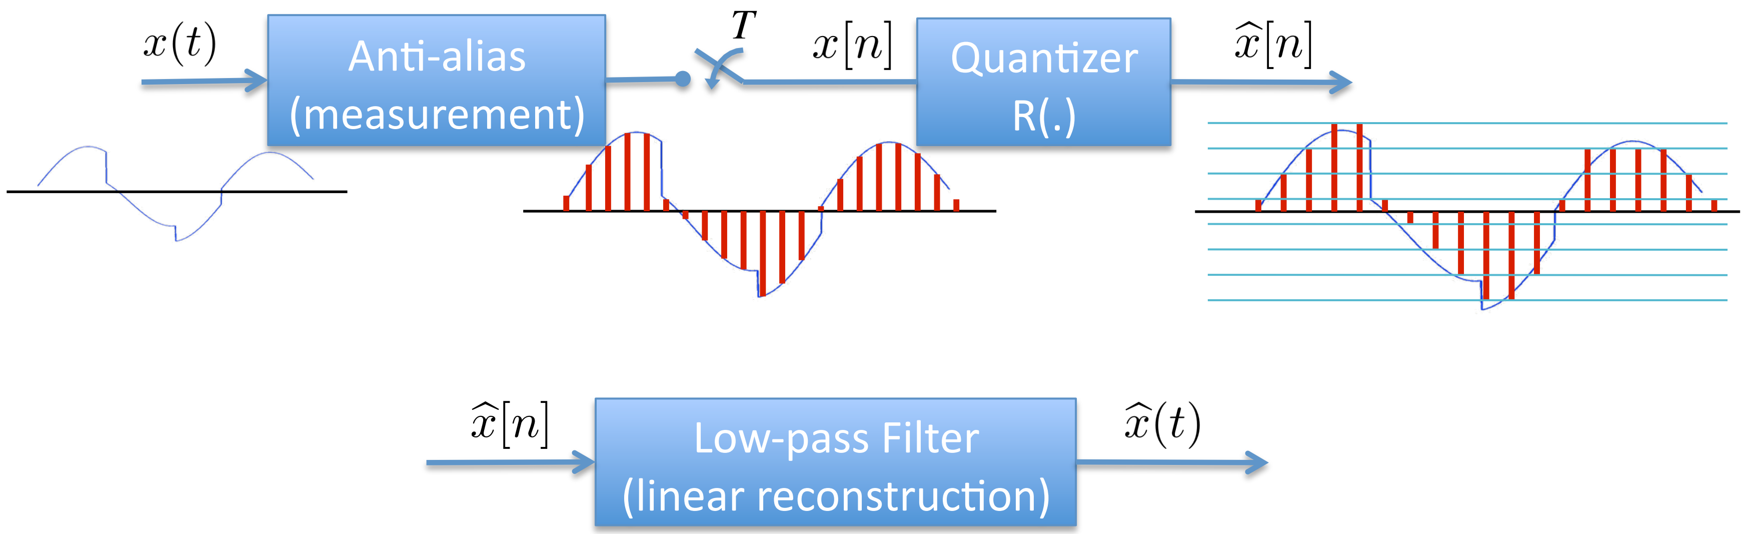

- **Sampling**: discretization in time
  - Lossless at the Nyquist rate
- **Quantization**: discretization in amplitude
  - Always lossy
- Need **both** for digital data acquisition

Quantizer Design

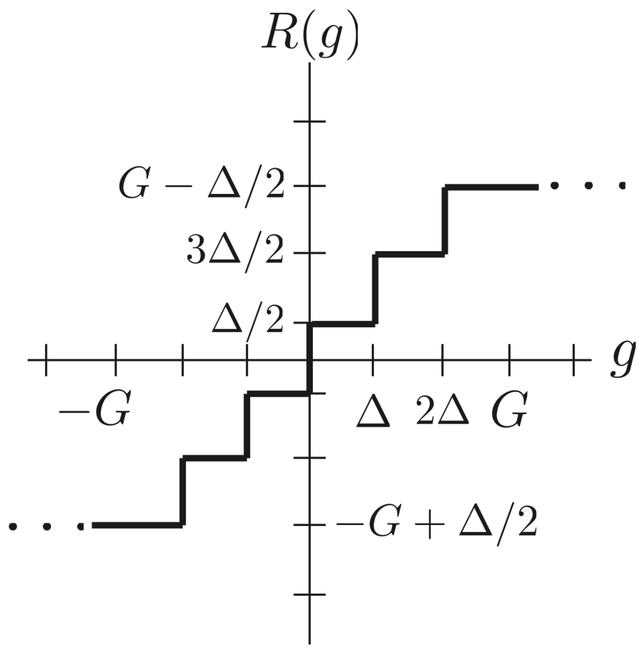

- Finite Range Quantizer
  - Finite range necessary
  - Hardware constraints
  - Bit-length constraints
- Design parameters:
  - G: Saturation level (threshold)
  - Δ: Quantization interval
  - B: Bit-rate (per coefficient)
- Design goals:
  - Smaller interval Δ: Less quantization error
  - Larger threshold G: Less saturation error
  - Trade-off: $\Delta=2^{-B}G$

Classical Quantizer Design

- Linear reconstruction penalizes saturation significantly
- Almost universal design principle:
  reduce or **eliminate** saturation
- Design heuristic:
  - Given: **bit budget** B bits/sample
  - Given: **maximum** signal **amplitude** A
  - Set: **threshold** $G \leq A$, i.e., $\Delta \geq 2^{-B+1}A$
- Drawback: **amplitude** A often **unknown**
  - In practice: Automatic Gain Control
  - Setting can be too **conservative**

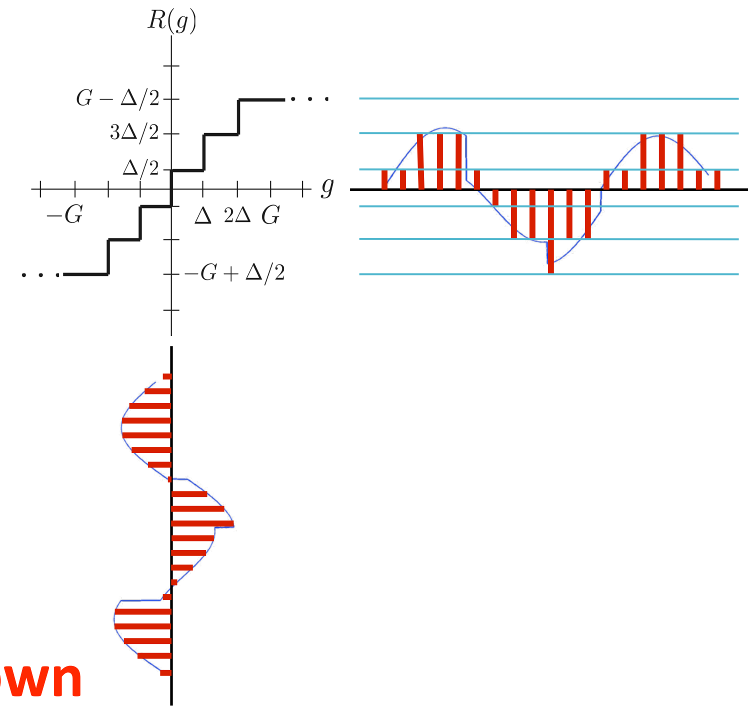

COMPRESSIVE SENSING

Compressive Sensing (CS)

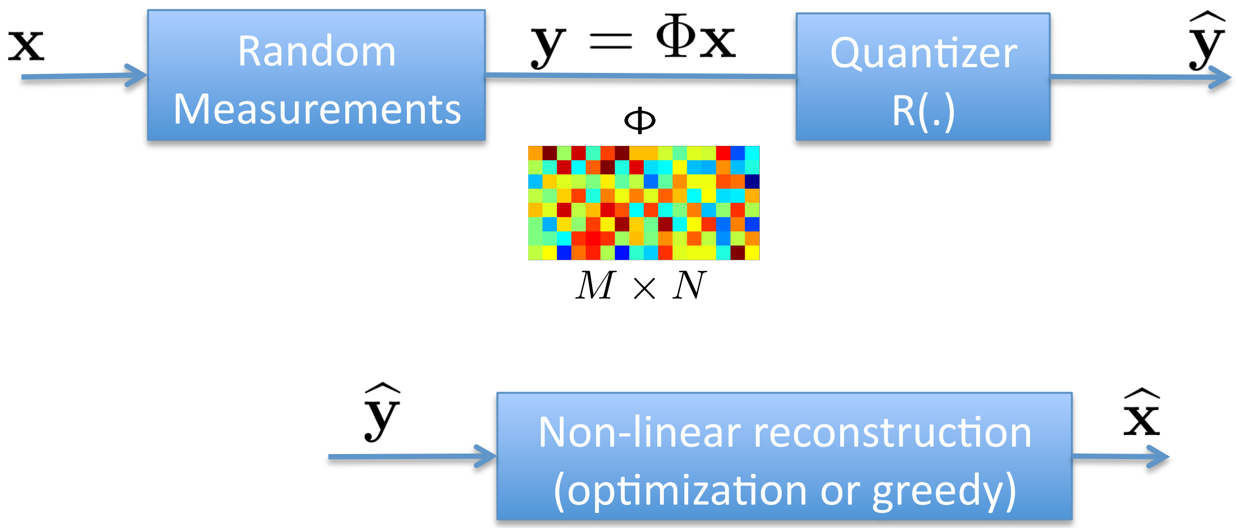

- **Compressive Sampling**: requires fewer measurements
  - Lossless for sparse signals
- **Quantization**: discretization in amplitude
  - Always lossy, still necessary
- Need **both** for digital data acquisition

CS Measurement Model

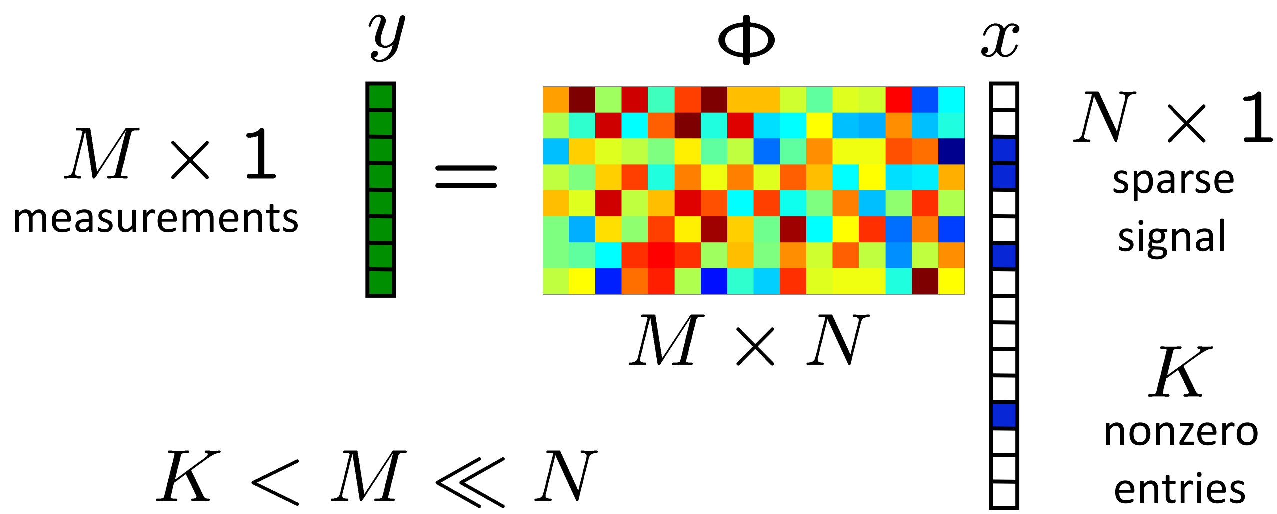

- x is K-sparse or K-compressible
- Φ **random**, satisfies a *restricted isometry property (RIP)*

Φ has RIP of order $2K$ with constant ϵ
If there exists ϵ s.t. for all $2K$-sparse x:

$$(1-\epsilon)\|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1+\epsilon)\|x\|_2^2$$

- $M = O(K \log N/K)$

CS Reconstruction

- Reconstruction using **sparse approximation**:
  - Find sparsest $\mathbf{x}$ such that $\mathbf{y} \approx \Phi\mathbf{x}$
- Convex optimization approach:
  - Minimize l_1 norm: e.g.,

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \|\mathbf{y} - \Phi\mathbf{x}\|_2 \leq \epsilon$$

- Greedy algorithms approach:
  - MP, OMP, ROMP, StOMP, CoSaMP, …
  - PYAMP (Pick Your Acronym Matching Pursuit)

CS Measurement Hallmarks

- All measurements are **democratic**
  - They all contribute the same to the reconstruction
  - Can drop some, and it does not matter which ones

- Measurements are **i.i.d**
  - The distribution of Φ entries is the reason

- In the limit, measurements look **Gaussian**
  - Central Limit theorem because of averaging

- **Asymmetric** processing
  - Cheap measurement, expensive reconstruction

QUANTIZATION AND CS

Compressive Sensing and Quantization

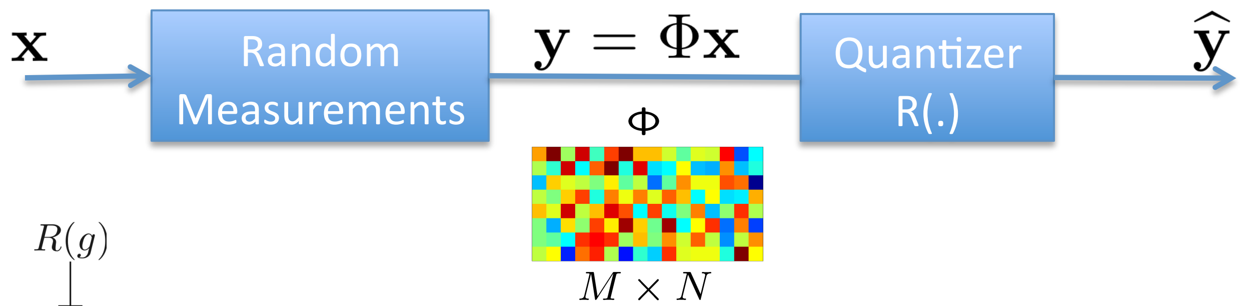

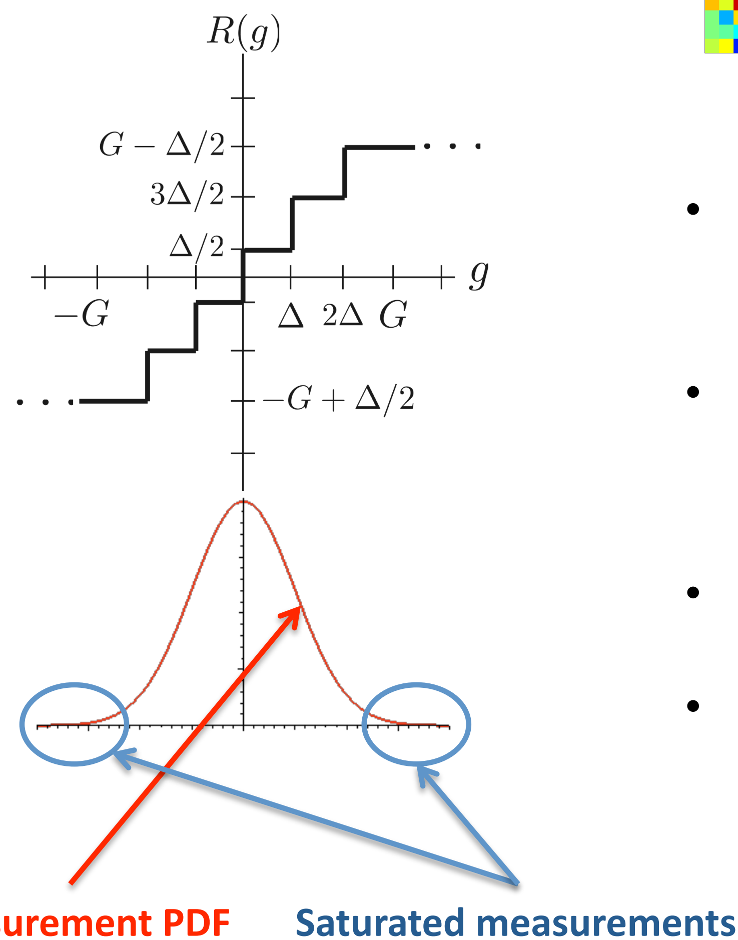

- Given:
  - **Bit budget** B bits/sample
  - **Signal norm** $\|\mathbf{x}\|_2$
- Set quantization **threshold** G
  - Implicitly sets quantization interval $\Delta = 2^{-B+1}G$
  - Implicitly sets saturation rate at $2Q(G/\|x\|_2)$
- ~~Classical heuristic: Set G large (avoid saturation)~~

  Wrong! Will revisit!
- Note:
  - varying G equivalent to fixing G and varying signal amplification
  - $Q(\cdot)$ denotes the tail of the Gaussian distribution

Reconstruction

- Convex optimization:

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \|\widehat{\mathbf{y}} - \Phi\mathbf{x}\|_2 \leq \epsilon$$

- ϵ bounds the quantization error

- Question: What should ϵ be set to?
  - "Straightforward" if no saturation
    - upper bound on quantization error
    - use a quantization noise model
    - l_p, $p>2$ constraints instead [e.g. Jaques et al. '09, others]
  - Saturation can make error **large** or **unbounded**

- How to deal with saturated measurements?
  - Treat them as **any other measurement**
  - Treat them as missing data and **ignore them**
  - Treat them as providing **different information**

Dropping Measurements

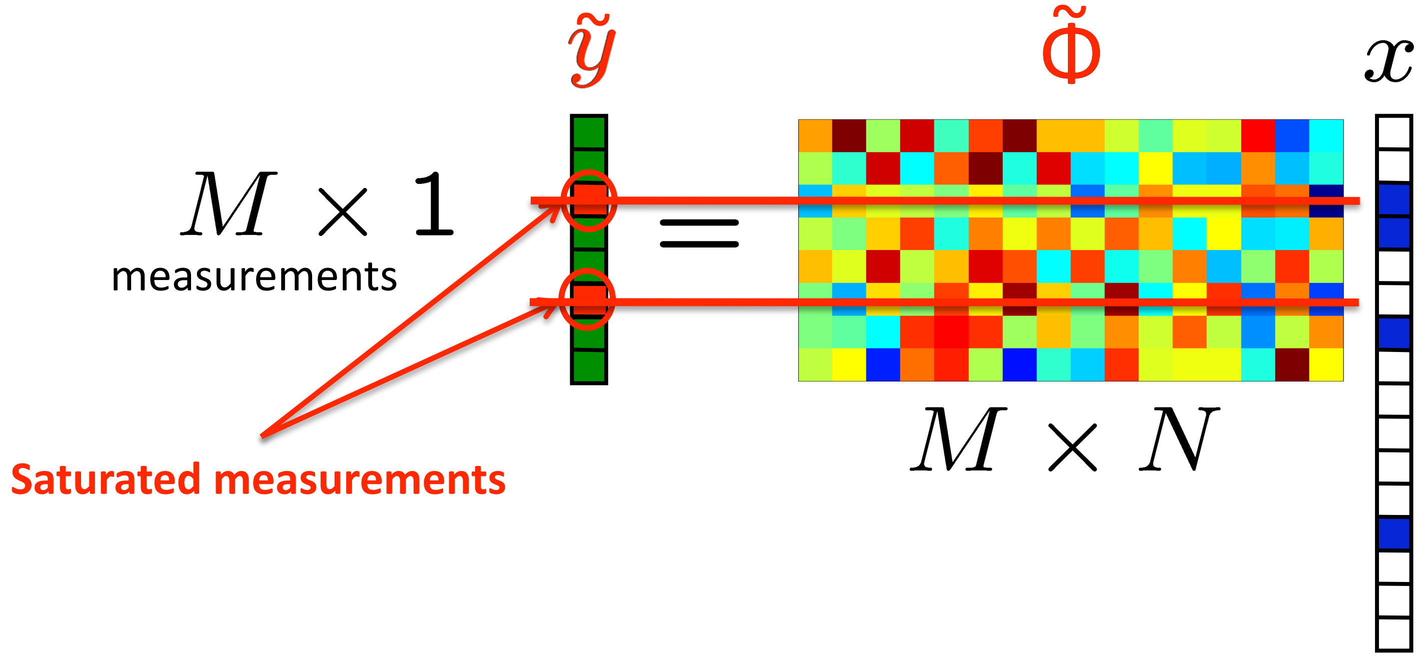

- Measurements are **democratic**
  - They are all equally important
  - We can **drop** the **saturated** ones

- Run the standard optimization on the remaining ones

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \|\widetilde{\mathbf{y}} - \widetilde{\Phi}\mathbf{x}\|_2 \leq \epsilon$$

- The remaining $\widetilde{\Phi}$ still satisfies RIP (as long as we don't drop too many)

Quantization Trade-Off

- **Increasing** saturation level G:
  - We **keep more measurements**, i.e., better reconstruction
  - We have **larger quantization interval**, i.e., worse reconstruction

- **Decreasing** saturation level G:
  - We **drop more measurements**, i.e., worse reconstruction
  - We have **smaller quantization interval**, i.e., better reconstruction

- A trade-off exists, **contrary** to the classical heuristics

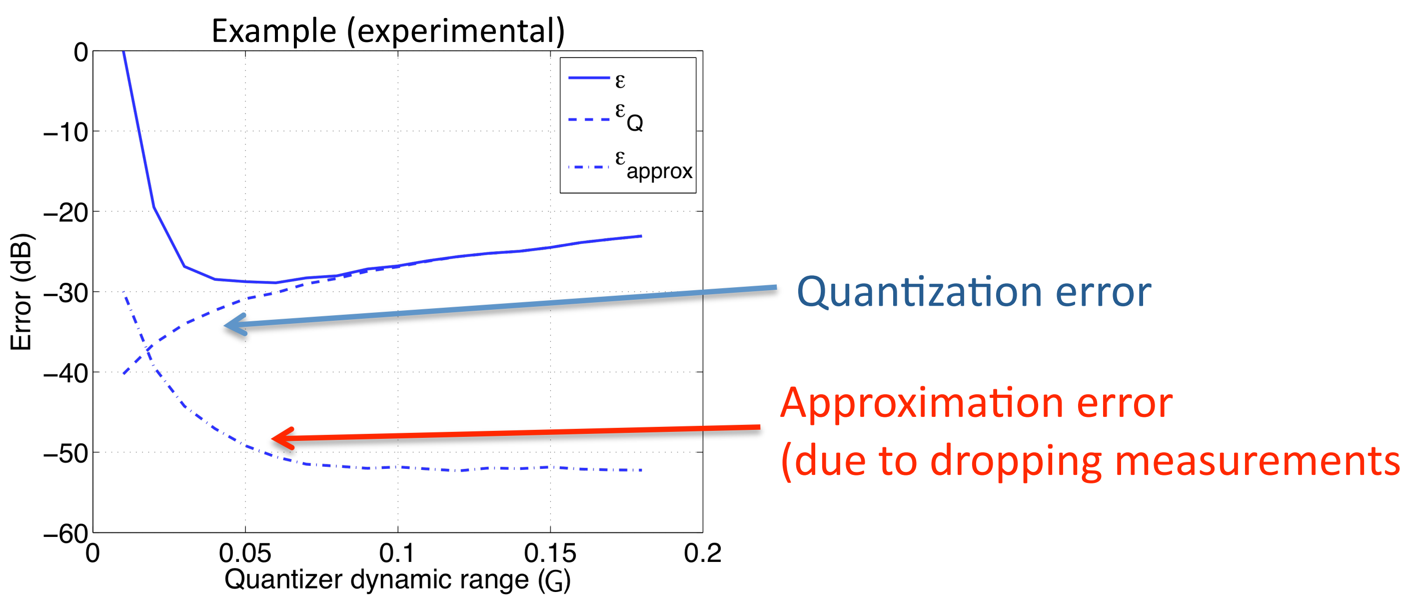

Experimental Results

Conventional approach (treat saturated data as measurements)

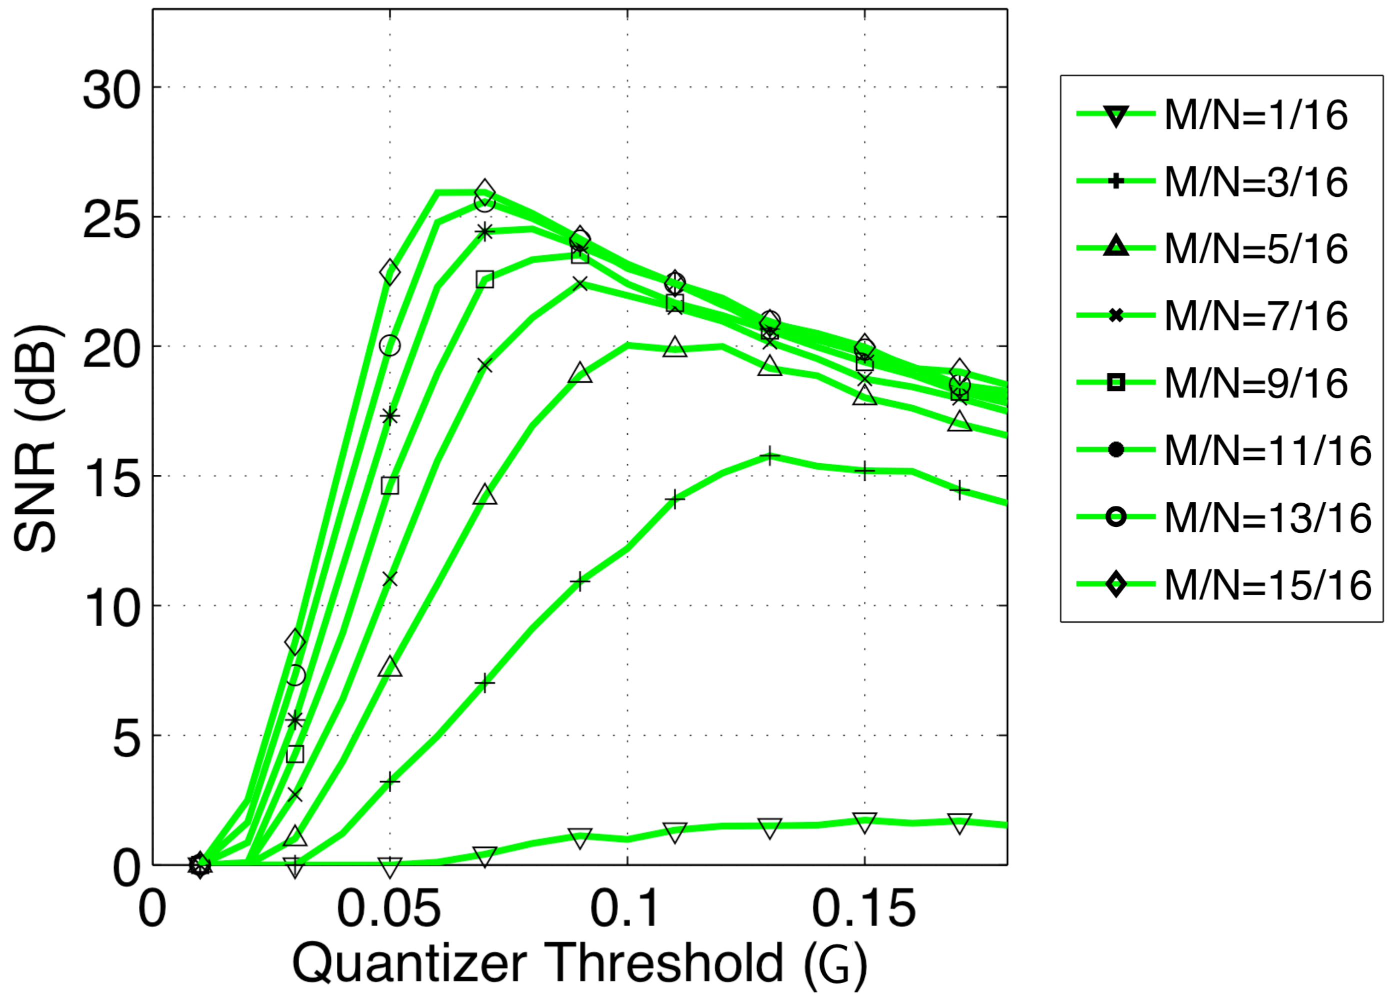

(a)

Experimental Results

New approach (discard saturated measurements)

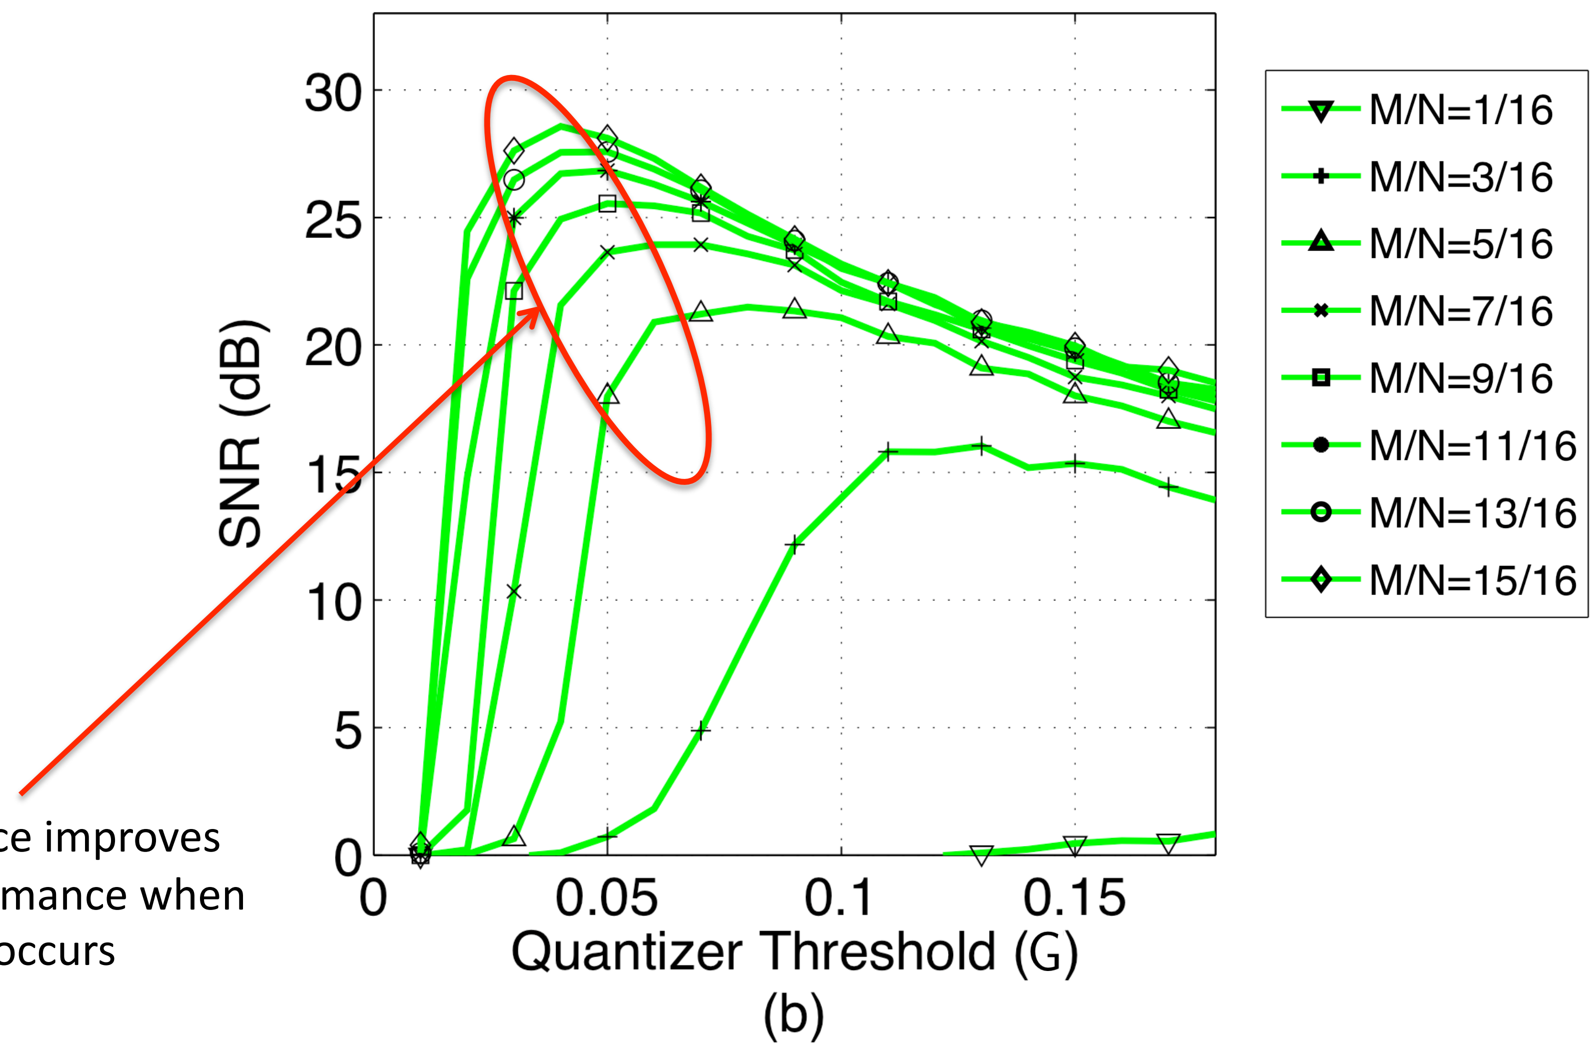

Performance improves
Best performance when saturation occurs

Experimental Results

Side by side comparison

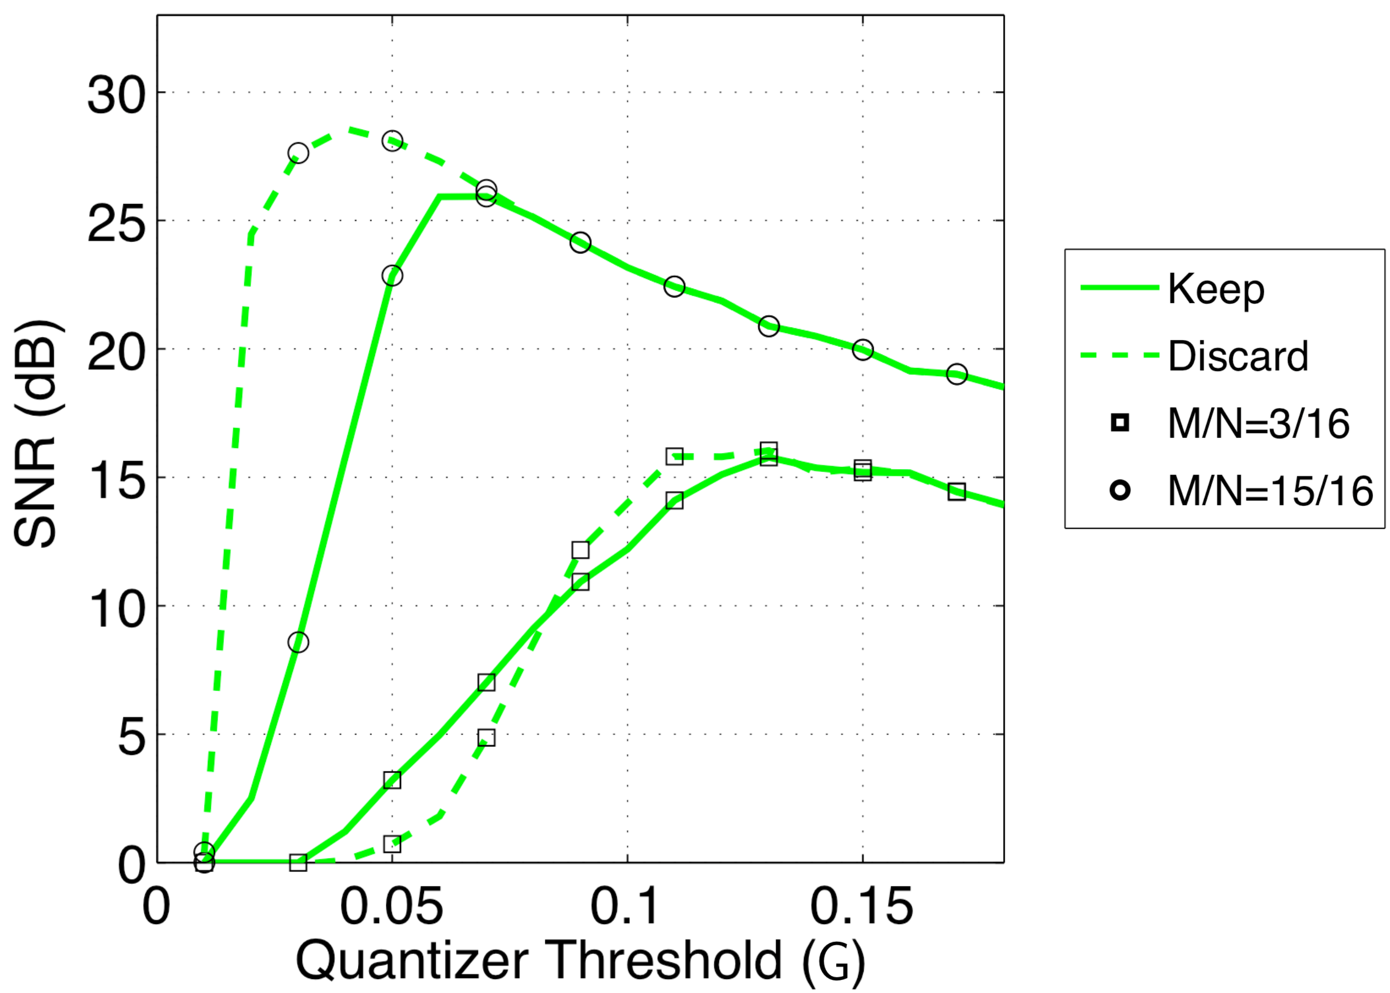

(c)

Alternative: Exploit Implicit Information

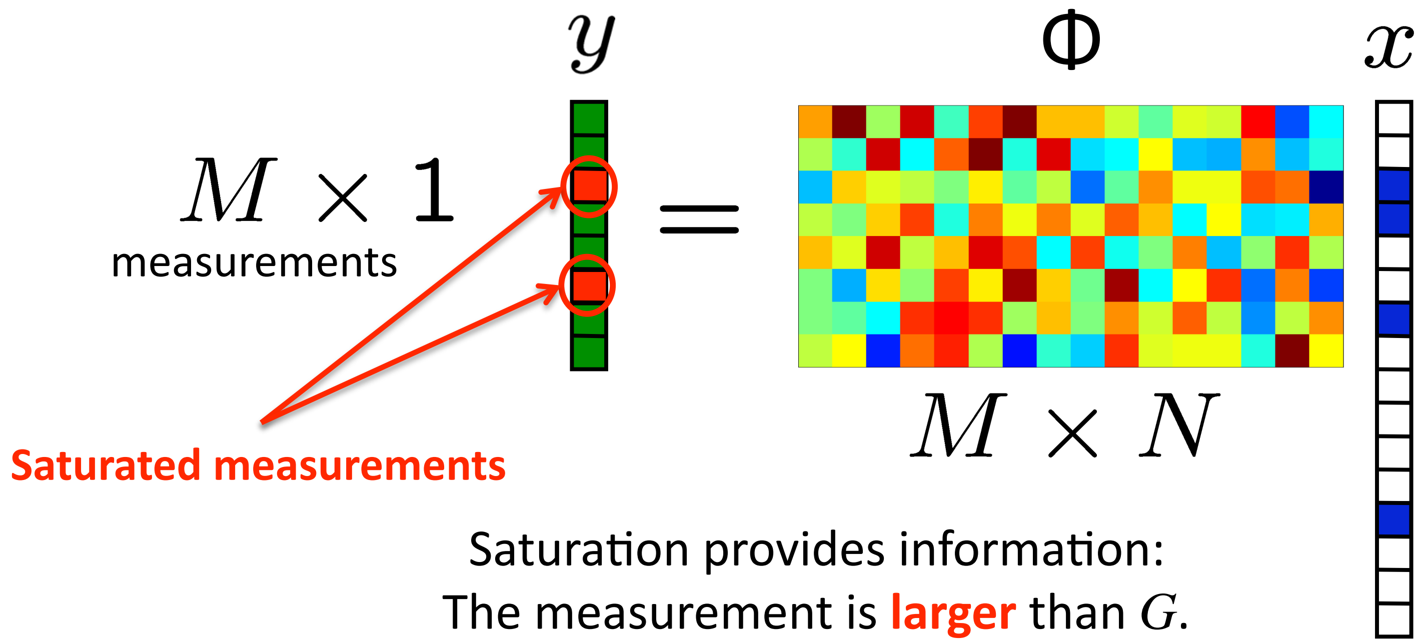

Saturation provides information:
The measurement is **larger** than G.

Treat measurement as a **constraint**!

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1$$

$$\text{s.t. } \|\widetilde{\mathbf{y}} - \widetilde{\Phi}\mathbf{x}\|_2 \leq \epsilon \qquad \text{Unsaturated}$$

$$\Phi^{+}\mathbf{x} \geq G - \Delta/2 \qquad \text{Positive Saturation}$$

$$\Phi^{-}\mathbf{x} \leq -G + \Delta/2 \qquad \text{Negative Saturation}$$

Experimental Results (preliminary, paper in preparation)

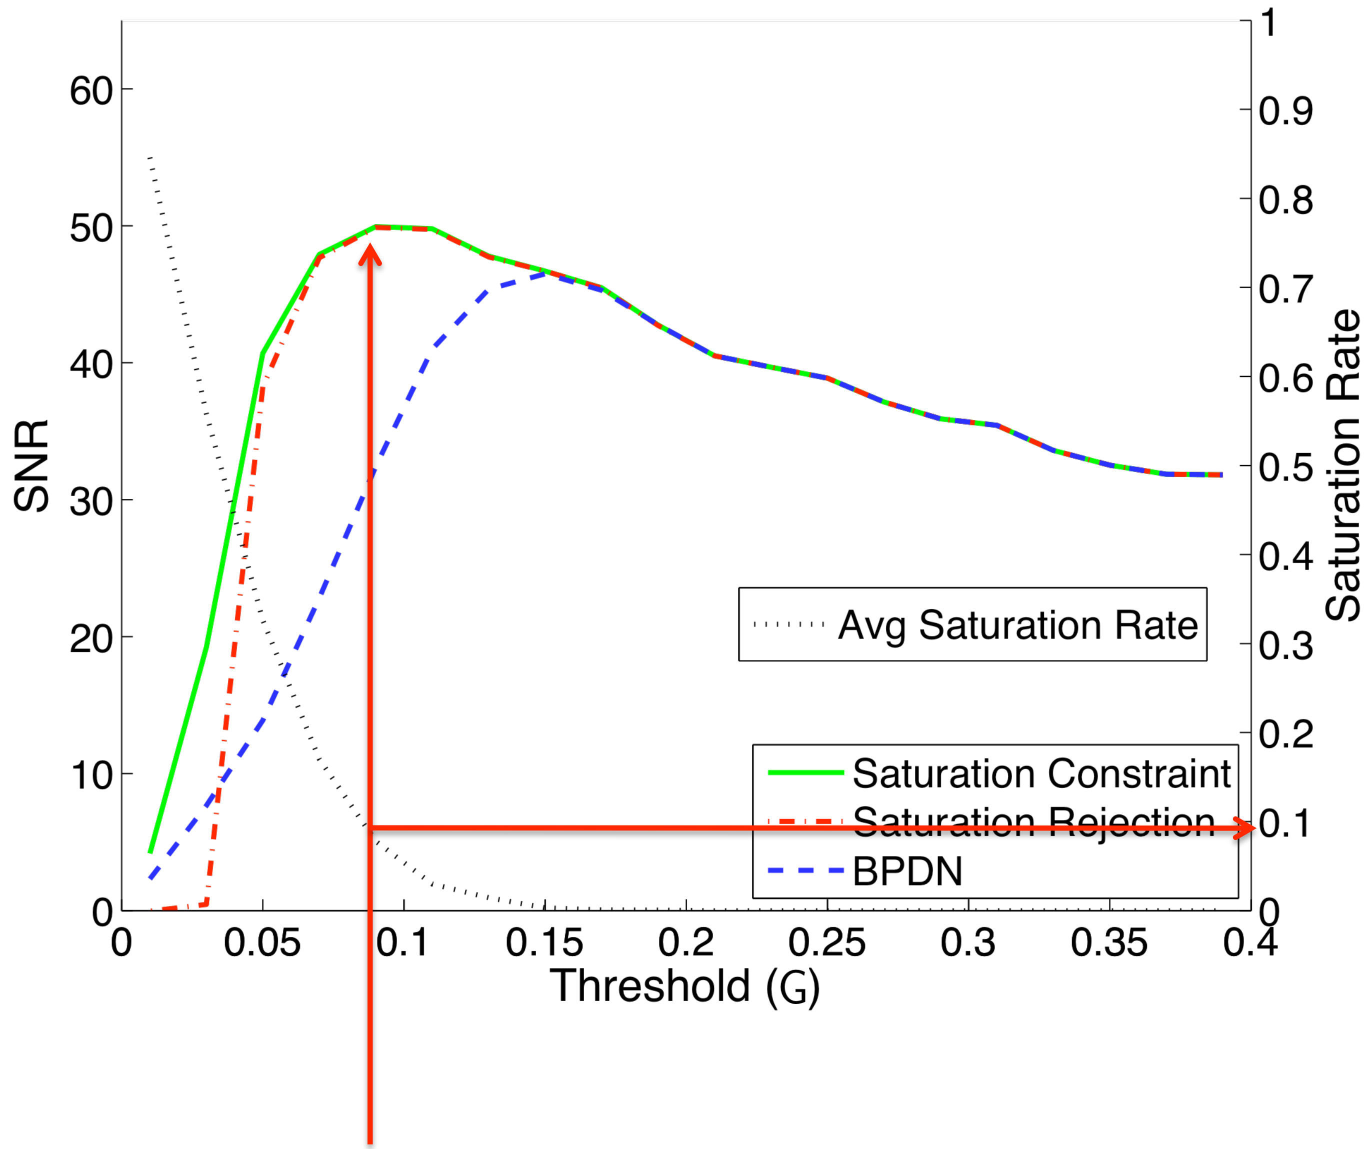

Note: optimal performance **requires** 10% saturation

Conclusions

- CS allows us to radically rethink quantization
- Randomized measurements **condition** the data
  - Data now looks **i.i.d. Gaussian**
  - Measurements are **democratic**
- Non-linear reconstruction provides **flexibility**
  - Can **drop** measurements
  - Can use measurements as **constraints**
- Optimal strategy: **saturate**!
  - Lower threshold $\Rightarrow$ lower quantization error
  - Penalty for dropping data is small
  - Suggests an automatic gain control strategy

Questions

Authors:

{laska,petrosb,richb}@rice edu

Rice CS repository:

dsp.rice.edu/cs