


Practical Compressed Sensing: modern data acquisition and signal processing

Thesis by
Stephen R. Becker

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2011
(Defended April 7, 2011)

© 2011
Stephen R. Becker
All Rights Reserved

Acknowledgements

It has been my privilege to work with highly talented and motivated researchers over the past few years. Much of the thesis work presented here was inspired or due to direct collaboration with these people and there would be no thesis were it not for them.

The most significant influence has been from my adviser, Professor Emmanuel Candès, who has supported me in many ways since I first came to campus. I owe him not just for scientific ideas, but for guiding me in the world of academics. I'm especially impressed by his integrity, his enthusiasm for the beauty of math, and his desire to use math to improve technology such as MRI. Like many great teachers, he leads by example.

It was a pleasure to work on the RMPI project with Dr. Emilio Sovero and Dr. Eric Nakamura from the Northrop Grumman Corporation; Dr. Michael Grant; and Professors Michael Wakin, Justin Romberg, Azita Emami, and Emmanuel Candès. Most importantly, thanks to my fellow student Juhwan Yoo from the MICS lab at Caltech. He is primarily responsible for the implementation of the RMPI chip. I'm grateful to his dedication and friendship over the past three years. Matthew Loh of the MICS lab also deserves recognition for his significant contributions to the project.

On my optimization endeavors, I've benefited greatly from interactions with many people. A great source of knowledge and friendship has been my co-authors, Jérôme Bobin, Michael Grant, and Emmanuel. Thanks as well to scholars in the field who have at some point contributed to my projects either explicitly or through informative discussions. I'm indebted to Patrick Combettes, Marc Teboulle, Joel Tropp, and Emmanuel for writing letters on my behalf, and I thank the Fondation Sciences Mathématiques de Paris for their support. I was fortunate to attend an IPAM long-workshop "Modern Trends in Optimization and Its Application" in Fall 2010 at UCLA; thanks to Russ Coflisch, Amber Puha, and the NSF for making this happen, and to Lieven Vandenberghe and Emmanuel in particular for the scientific organization. This was a fantastic opportunity for a graduate student. Similarly, I am very grateful to Alexandre d'Aspremont, Francis Bach, and Martin Wainwright for organizing and then inviting me to the "Sparse Statistics, Optimization and Machine Learning" conference in Banff, January 2011.

In the Applied & Computational Math (ACM) department at Caltech, I'm always impressed by the level of thought in the graduate students. Thanks to all of them for friendship, discussions, and

the occasional “ACM-TEA” seminar. The faculty in ACM has always been excellent. Sheila Shull and Sydney Garstang have been incredibly helpful to me and all the other students and faculty. Special thanks to Alex Gittens and Michael McCoy for academic discussions, as well as the non-academic discussions during the commute to UCLA.

Many thanks to my committee members: Lieven Vandenberghe, Babak Hassibi, Azita Emami, Joel Tropp, and Emmanuel Candès. Joel especially has been instrumental to my academic development over the past three years, and it’s clear he cares deeply about the students at Caltech. The Candès group, past and present, has been a great source of role models for a young graduate student. Thanks to Ewout van den Berg, Laurent Demanet, Hannes Helgason, Ben Recht, Justin Romberg, Mike Wakin, Lexing Ying, Paige Randall, Peter Stobbe, Jérôme Bobin, and Xiaodong Li. I’m thankful for the years of camaraderie and friendship with Yaniv Plan. I’m also grateful to my undergraduate adviser, Francis Starr, who started me on my academic path.

My friends at Caltech have made my stay here very enjoyable. In particular, Sean Tulin and Justus Brevik were willing to be in close quarters with me for 3 weeks, so that says something about their friendship. My parents, Mary and Charlie, and brother, Andy, have always been supportive, and they have helped incredibly by shaping me as a person and always encouraging my academic pursuits.

Finally, I owe my happiness to my wife Kimberly. Thank you, Kim.

Abstract

Since 2004, the field of compressed sensing has grown quickly and seen tremendous interest because it provides a theoretically sound and computationally tractable method to stably recover signals by sampling *at the information rate*. This thesis presents in detail the design of one of the world's first compressed sensing hardware devices, the random modulation pre-integrator (RMPI). The RMPI is an analog-to-digital converter (ADC) that bypasses a current limitation in ADC technology and achieves an unprecedented 8 effective number of bits over a bandwidth of 2.5 GHz. Subtle but important design considerations are discussed, and state-of-the-art reconstruction techniques are presented.

Inspired by the need for a fast method to solve reconstruction problems for the RMPI, we develop two efficient large-scale optimization methods, NESTA and TFOCS, that are applicable to a wide range of other problems, such as image denoising and deblurring, MRI reconstruction, and matrix completion (including the famous Netflix problem). While many algorithms solve unconstrained ℓ_1 problems, NESTA and TFOCS can solve the *constrained* form of ℓ_1 minimization, and allow weighted norms. In addition to ℓ_1 minimization problems such as the LASSO, both NESTA and TFOCS solve total-variation minimization problem. TFOCS also solves the Dantzig selector and most variants of the nuclear norm minimization problem. A common theme in both NESTA and TFOCS is the use of smoothing techniques, which make the problem tractable, and the use of optimal first-order methods that have an accelerated convergence rate yet have the same cost per iteration as gradient descent. The conic dual methodology is introduced in TFOCS and proves to be extremely flexible, covering such generic problems as linear programming, quadratic programming, and semi-definite programming. A novel continuation scheme is presented, and it is shown that the Dantzig selector benefits from an exact-penalty property. Both NESTA and TFOCS are released as software packages available freely for academic use.

Contents

Acknowledgements	iii
Abstract	v
1 Introduction	1
1.1 Applications	5
1.1.1 Problems of interest	10
1.2 Historical development	14
1.2.1 Classical signal processing	15
1.2.2 Estimation and the rise of alternatives to least-squares	16
1.2.3 Leading up to compressed sensing	18
1.2.4 Compressed Sensing	19
1.3 The need for the RMPI	24
1.4 Principles of the RMPI	28
1.4.1 The NUS	28
1.4.2 The RMPI	29
1.5 Optimization background	32
1.6 Reading guide	39
2 RMPI	40
2.1 Introduction	41
2.1.1 Signal class	41
2.2 The design	46
2.2.1 Basic design	46
2.2.2 Theoretical performance	49
2.2.2.1 Input noise and channelized receivers	50
2.2.3 Related literature	53
2.2.3.1 Other RMPI systems	53
2.2.3.2 Related systems	57

2.3	Modeling the system	60
2.3.1	Simple models	62
2.3.2	SPICE	64
2.3.3	Simulink	65
2.3.4	Calibration	67
2.3.5	Phase blind calibration	70
2.4	General design considerations	73
2.4.1	Test signals	74
2.4.2	Error metrics	75
2.4.3	Number of channels	76
2.5	Chip sequence	80
2.5.1	Spectral properties of the chip sequence	81
2.5.1.1	Infinite period	81
2.5.1.2	Finite period	82
2.5.2	Chip design considerations	85
2.5.2.1	Chip sequence rate	86
2.5.2.2	Chip sequence period	87
2.5.2.3	Case study: test of NG chip sequence	92
2.6	Integration	94
2.6.1	General constraints	94
2.6.1.1	Northrop Grumman integrator design	95
2.6.1.2	Multipole systems	97
2.7	Recovery	103
2.7.1	Matched filter	103
2.7.2	ℓ_1 recovery	108
2.7.2.1	Analysis versus synthesis	108
2.7.2.2	Dictionary choice	110
2.7.2.3	Reweighting	114
2.7.2.4	Debiasing	118
2.7.2.5	Non-linearity correction	119
2.7.2.6	Windowing	123
2.7.2.7	Further improvements	125
2.8	Results	127
2.8.1	Non-idealities	128
2.8.1.1	Noise	128
2.8.1.2	Jitter	129

2.8.1.3	Quantization	130
2.8.1.4	Cross-talk	131
2.8.1.5	Clipping	131
2.8.1.6	Combining non-idealities	135
2.8.2	Simulation results	137
2.8.2.1	Single pulse	138
2.8.2.2	Two pulses	141
2.8.2.3	Comparison	143
2.8.3	Hardware	144
2.8.3.1	NG InP version	144
2.8.3.2	Version 1	145
2.8.3.3	Version 2	145
2.9	Recommendations	147
3	NESTA	149
3.1	Introduction	149
3.1.1	Contributions	151
3.1.2	Organization of the chapter and notations	152
3.2	NESTA	154
3.2.1	Nesterov's method to minimize smooth convex functions	154
3.3	Application to compressed sensing	155
3.3.1	NESTA	155
3.3.2	Updating y_k	157
3.3.3	Updating z_k	158
3.3.4	Computational complexity	159
3.3.5	Parameter selection	160
3.3.6	Accelerating NESTA with continuation	161
3.3.7	Some theoretical considerations	164
3.4	Accurate optimization	166
3.4.1	Is NESTA accurate?	166
3.4.2	Setting up a reference algorithm for accuracy tests	167
3.4.3	The smoothing parameter μ and NESTA's accuracy	168
3.5	Numerical comparisons	170
3.5.1	State-of-the-art methods	171
3.5.1.1	NESTA	171
3.5.1.2	Gradient projections for sparse reconstruction (GPSR)	172

3.5.1.3	Sparse reconstruction by separable approximation (SpaRSA)	172
3.5.1.4	ℓ_1 regularized least-squares (l1_ls)	172
3.5.1.5	Spectral projected gradient (SPGL1)	173
3.5.1.6	Fixed point continuation method (FPC)	173
3.5.1.7	FPC active set (FPC-AS)	173
3.5.1.8	Bregman	174
3.5.1.9	Fast iterative soft-thresholding algorithm (FISTA)	174
3.5.2	Constrained versus unconstrained minimization	175
3.5.3	Experimental protocol	175
3.5.4	Numerical results	176
3.5.4.1	The case of exactly sparse signals	176
3.5.4.2	Approximately sparse signals	178
3.6	An all-purpose algorithm	180
3.6.1	Non-standard sparse reconstruction: ℓ_1 analysis	180
3.6.2	Numerical results for non-standard ℓ_1 minimization	182
3.6.3	Total-variation minimization	183
3.6.4	Numerical results for TV minimization	185
3.7	Handling non-projectors	189
3.7.1	Revisiting the projector case	190
3.7.2	Non-projectors for $\epsilon = 0$ case	190
3.7.3	Non-projectors for $\epsilon > 0$ case	191
3.8	Discussion	193
3.8.1	Extensions	193
3.8.2	Software	195
4	TFOCS	196
4.1	Introduction	197
4.1.1	Motivation	197
4.1.2	The literature	199
4.1.3	Our approach	201
4.1.3.1	Conic formulation	201
4.1.3.2	Dualization	202
4.1.3.3	Smoothing	203
4.1.3.4	First-order methods	203
4.1.4	Contributions	204
4.1.5	Software	205

4.1.6	Organization of the chapter	206
4.2	Conic formulations	206
4.2.1	Alternate forms	206
4.2.2	The dual	207
4.2.3	The differentiable case	208
4.2.4	Smoothing	209
4.2.5	Composite forms	210
4.2.6	Projections	211
4.3	A novel algorithm for the Dantzig selector	212
4.3.1	The conic form	212
4.3.2	Smooth approximation	213
4.3.3	Implementation	214
4.3.4	Exact penalty	215
4.3.5	Alternative models	217
4.4	Further instantiations	218
4.4.1	A generic algorithm	218
4.4.2	The LASSO	219
4.4.3	Nuclear-norm minimization	220
4.4.4	ℓ_1 -analysis	221
4.4.5	Total-variation minimization	223
4.4.6	Combining ℓ_1 analysis and total-variation minimization	224
4.5	Implementing first-order methods	224
4.5.1	Introduction	225
4.5.2	The variants	225
4.5.3	Step size adaptation	227
4.5.4	Linear operator structure	229
4.5.5	Accelerated continuation	230
4.5.6	Strong convexity	233
4.6	Dual-function formulation	235
4.6.1	Background	235
4.6.2	Fenchel dual formulation	236
4.6.3	Convergence	239
4.6.3.1	Convergence when f is smooth	239
4.6.3.2	Convergence of inner iteration	240
4.6.4	Convergence of outer iteration	241
4.6.4.1	Overall analysis	243

4.7	Numerical experiments	244
4.7.1	Dantzig selector: comparing first-order variants	244
4.7.2	LASSO: comparison with SPGL1	246
4.7.3	Wavelet analysis with total-variation	247
4.7.4	Matrix completion: expensive projections	249
4.7.5	ℓ_1 -analysis	252
4.8	Extensions	256
4.8.1	Automatic restart	256
4.8.2	Specialized solvers for certain problems	256
4.8.2.1	Noiseless basis pursuit	256
4.8.2.2	Conic problems in standard form	257
4.8.2.3	Matrix completion problems	259
4.9	Software: TFOCS	260
4.10	Discussion	262
4.11	Appendix: exact penalty	263
4.12	Appendix: creating a synthetic test problem	265
5	Conclusion	268
5.1	Future of A2I	268
5.2	Improvements to TFOCS	269
	Bibliography	271