Chapter 5

Conclusion

Since each chapter has its own conclusion, we conclude this thesis not with a summary, but with possibilities for future work. Most of the issues raised by NESTA are addressed by the TFOCS algorithm, so we do not discuss extensions of NESTA.

5.1 Future of A2I

The current chip is expected to be fully calibrated by Summer 2011, and this will give many new insights into the design. The RMPI is necessarily a prototype chip and is programmable in many ways, allowing us to modify it as necessary; future versions will be more streamlined and have better performance. Since calibration is on-going, this is an active area of research, and new methods are being tested weekly. The end-goal is not just an accurate calibration for academic purposes, but a calibration scheme that is robust, easy, and fast.

The biggest advance in an A2I system would be a novel architecture that is specialized to receive a specific signal class. For example, a pulse *detection* system that only needs to detect the presence of a pulse, or perhaps only estimate a frequency, allows much more freedom than a system that tries to reconstruct the input signal exactly. Putting a matched-filter system on-chip would be a major step in this direction. Another ambitious target would be to implement special purpose coding matrices, such as expander graphs, and decode on-chip in real-time. It may be possible to make an adaptive system that uses a crude detection first-step in order to give information, in real-time, to a second-stage system which "zooms" to the frequencies of interest.

Another hardware improvement would be to use different types of chipping signals. This is governed by both theoretical considerations (a random sum of shifted sinc-like functions would be nice) and by practical considerations, since the chipping signal should put as little strain on the hardware as possible, and be as robust as possible to clock jitter. We note recent work on improving the basic ± 1 chipping signal [HBC11, ME09b], and the work of the modulated wide-band converter [MEE09] which has many good ideas.

For reconstruction, the NESTA and TFOCS algorithms have made large scale reconstructions possible at an academic level, but real-time reconstruction is still far away. The best method to achieve ultra-fast processing is with the compressive matched filter [ERW11], though the performance degrades when the input signal is complicated, since the parameter space becomes large so the algorithm slows down.

To improve the dynamic range of the system, it seems a tighter signal model is needed. Work on model-based compressive sensing [BCDH10,Cev09] suggests tree-based models can improve recovery, though the exact amount of improvement is likely to depend strongly on the specific signal class involved. Using a reweighted scheme with weights picked in accordance with a hierarchical structure, such as in [DWB08], is an obvious way to apply these models to ℓ_1 techniques without having to use greedy methods. Other ways to improve the system performance is to exploit all prior knowledge of the errors, such as estimating the covariance matrix of the noise, and using iterative algorithms to reduce the effect of non-linearities and jitter [WG11].

Other work on recovery is to make the system capable of studying a continuous stream of data, rather than a single window at a time. To this end, a windowing scheme will be essential, and the exact details still need to be worked out.

5.2 Improvements to TFOCS

Our biggest goal for TFOCS is to develop a robust convergence theory of the accelerated continuation scheme. This is not of just theoretical interest, but is also useful in practice since it may suggest what value to set the smoothing parameter μ, and determine what the stopping tolerance of the inner iterations should be. Ultimately, the user should not need to provide any parameters. The specialized line search has already obviated the need for a step-size parameter, and we are close to finding an automatic way to determine when to restart the algorithm.

Certain special problems, such as matrix completion, provide many avenues for future research. The SVD computation in matrix completion only requires the singular vectors that correspond to singular values above a given threshold, so the entire SVD is not needed. We predict that randomized linear algebra will play a big role in speeding up computations. Due to non-negligible errors, optimization algorithms need to gracefully accept approximate computation of gradients and still guarantee convergence.

We plan to continue support of the software package and incorporate new solvers, such as BFGS and Douglas-Rachford. "Splitting methods" such as Douglas-Rachford, though not well-known in some communities, may be very efficient at certain problems. Other planned improvements include automatic diagonal scaling to improve performance on "stiff" problems, and perhaps incorporating the solvers into a modeling environment like CVX [GB10] or porting the software to a free language like R or Python.