New Algorithms for Diagnosis of Multiple Faults

Stephen Boyd Dimitry Gorinevsky Argyris Zymnis

Electrical Engineering Department, Stanford University

AvSafe Tech. Conference, Washington D.C., 11/18/09

Where it Fits

- IVHM Milestone 1.2.2.2: *Develop Bayesian and Hybrid Reasoning Methods for State Estimation and Diagnosis.*
- demonstrated on ADAPT power system testbed
- method as described handles Boolean faults ($i.e.$, present or not) and continuous measurements
- extensions (not reported here) handle other cases
- additive fault signature model means gross faults cannot be handled

Our Goal, Approach, and Method

- we have
  - noisy measurements from a system
  - a set of faults that can occur *in any combination* (in particular, *multiple faults* can occur)
  - a statistical model of the measurements and fault occurences
- goal: develop a *scalable algorithm to diagnose which combination of faults has occurred*
- our approach: statistical (Bayesian)
- our method: based on recent developments in *convex optimization*

Outline

- statistical fault/measurement model
- maximum a posteriori (MAP) fault diagnosis
- relaxed MAP approach
- some examples
- quantized measurements
- conclusions

Measurement Residuals

- measurement/fault model: $z = F(x) + w$
  - $z \in \mathbf{R}^m$ is vector of m measurements
  - $x \in \{0, 1\}^n$ is fault pattern (2^n possible values)
  - $w \in \mathbf{R}^n$ is vector of random or unmodeled variables, noises
  - F is a complex function mapping faults and noises to measurements

- measurement *residuals*: $y = z - F(0)$

  difference of actual and predicted measurement, assuming
  - the model is correct
  - there is no noise

Linearized Residuals and Fault Model

- linearized residuals model: $y = \sum_{i=1}^{n} x_i a_i + v$
  - $a_i \in \mathbf{R}^n$ is *fault signature* (or pattern) for fault i
  - (noise free) residual is sum of fault patterns for faults that occur
  - can write as $y = Ax + v$, where $A \in \mathbf{R}^{m \times n}$ is *fault signature matrix*
- statistical model: we'll assume
  - noise $v \sim \mathcal{N}(0, \sigma^2 I)$
  - faults occur independently, with probabilities $p_1, \ldots, p_n$

  (but our approach works with more general assumptions . . .)
- yields a Bayes net with very specific structure

Bayesian Fault Diagnosis Formulation

- we are given
  - model parameters A, σ
  - prior fault probabilities $p \in \mathbf{R}^n$
  - the measurement y

- $P(x|y)$ is (posterior) probability that fault pattern x has occurred

- maximum a posteriori probability (MAP) estimate of fault pattern:
  - $x \in \{0, 1\}$ which maximizes $P(x|y)$
  - the most likely fault pattern, given the measurement
  - can add weights to trade-off false positives, false negatives, . . .

MAP Fault Diagnosis

- fault pattern loss:

$$\begin{aligned} l_y(x) &= \log P(0|y) - \log P(x|y) \\ &= \left(1/2\sigma^2\right)\left(\|y - Ax\|_2^2 - \|y\|_2^2\right) + \lambda^T x \end{aligned}$$

  with $\lambda_j = \log((1 - p_j)/p_j)$

- MAP estimate is solution of optimization problem

$$\begin{array}{ll} \text{minimize} & l_y(x) \\ \text{subject to} & x_j \in \{0, 1\} \end{array}$$

  with variable $x \in \{0, 1\}^n$

Ambiguity Group

- ambiguity group: set of fault patterns with nearly optimal loss
- these are also candidates for the true fault pattern
- far more useful in practice than just the MAP solution

Solving MAP Problem (Globally)

- sadly, you can't (in general); it's a (hard) Boolean quadratic program
- direct enumeration: evaluate $l_y(x)$ for all 2^n fault patterns
  - not practical for $n > 15$ or so
- branch and bound, other mixed integer QP methods
  - can be very slow
- not clear that global solution of MAP problem gives better estimation performance than a good heuristic solution of MAP problem

Solving MAP Problem (Approximately)

- 'approximately' means: $\hat{x}$ may not have (globally) minimum loss
- simple (partial) enumeration methods (widely used!)
  - enumerate $l_y(x)$ for all n single-fault patterns
  - enumerate $l_y(x)$ for all $n(n-1)/2$ double-fault patterns
- local optimization, $e.g.$,
  - start from $x = 0$
  - flip bits one at a time
  - accept any changes that decrease l_y
- these methods can work well, in some (simple) cases

Our Method

- form (computationally tractable) *convex relaxation* of MAP problem
- solve relaxation using fast custom method
- round relaxed estimates
- carry out simple local optimization

efficient, scales to large problems, and seems to work very well, even in challenging cases

QP Relaxation of MAP Problem

$$\begin{array}{ll} \text{minimize} & l_y(x) \\ \text{subject to} & 0 \leq x_j \leq 1 \end{array}$$

- allow variables x_i to be *between* 0 (false) and 1 (true) (sometimes called 'soft decisions')
- a convex QP, easily solved
- solution z gives lower bound on loss, $i.e.$,

$$l_y(x) \geq l_y(z) \text{ for all } x \in \{0,1\}^n$$

- if $z \in \{0,1\}^n$, it's MAP optimal
- if $\lambda \geq (1/\sigma^2)A^T y$, $x = 0$ is MAP optimal

Rounding the Relaxed Solution

to get fault estimate (and ambiguity group) from z:

- for each $t \in (0, 1)$, round all $z_i \geq t$ to one, and all $z_i < t$ to zero (can be done efficiently by sorting z_i)
- evaluate $l_y(\hat{x})$
- pattern with smallest $l_y(\hat{x})$ is our (relaxed MAP) estimate x^{rmap}
- patterns with nearly smallest $l_y(\hat{x})$ form (approximate) ambiguity group

Local Optimization

- start with list of K patterns with smallest values found after rounding, sorted by loss (this is our initial ambiguity group estimate)

$$x^{(1)}, \ldots, x^{(K)}$$

- cycle through index j
  - tentatively replace $x_j^{(1)}$ with $1 - x_j^{(1)}$ and evaluate loss
  - if this pattern has lower loss than $x^{(K)}$, add it to list, delete $x^{(K)}$
  - in particular, if this pattern has lower loss than $x^{(1)}$, it replaces $x^{(1)}$ as the least loss pattern found so far
  - continue until there are no changes over one sweep of index j

Approximate MAP Estimation

approximate relaxed MAP problem:

$$\text{minimize} \quad l_y(x) + \kappa\psi(x) = l_y(x) - \kappa \sum_{j=1}^n \log(x_j(1 - x_j))$$

with implicit constraint $x_i \in (0, 1)$

- parameter $\kappa > 0$ controls quality of approximation
- solution is no more than $2n\kappa$ suboptimal for RMAP problem
- an unconstrained smooth convex problem, readily solved by Newton type methods, even for very large problems

Small Example

- $m = 50$ measurements
- $n = 100$ potential faults
- A_{ij} chosen from $\mathcal{N}(0, 1)$ distribution
- $p_j = 0.05$ (so expected number of faults is 5)
- $\sigma = 1$, corresponding to signal-to-noise (SNR) ≈ 2.2
- runtime around 20ms for RMAP, 2s for SDP, 20min for SDP$+$
- all relaxations recover MAP solution
  (verified using branch & bound, after 110 iterations)
- over many runs, probability of exact fault detection is 95%

Small Example: Results from Single Run

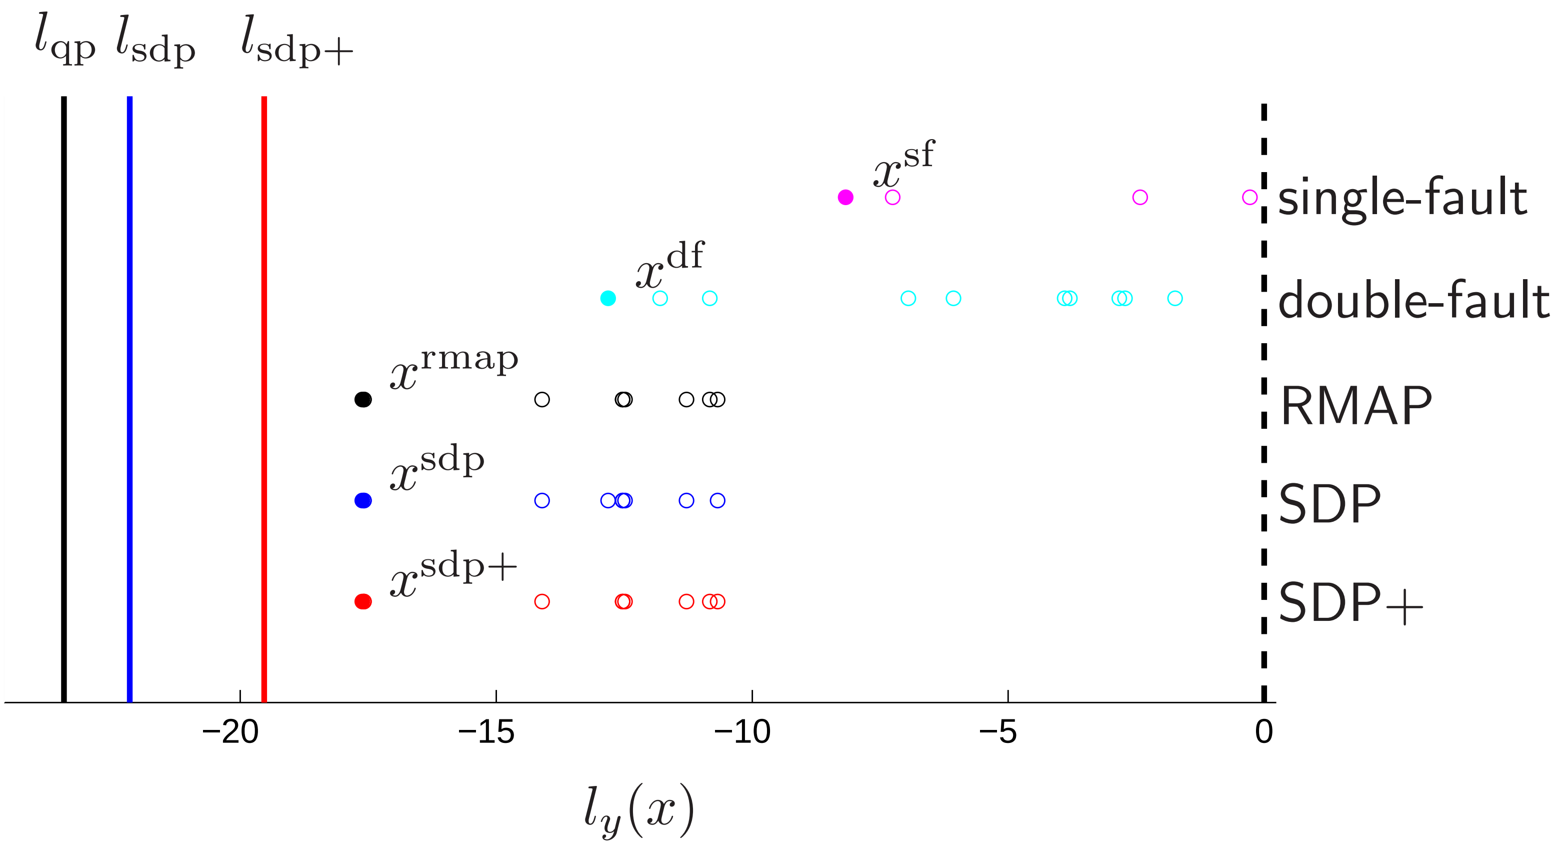

Performance vs SNR: Top Hit Percentage

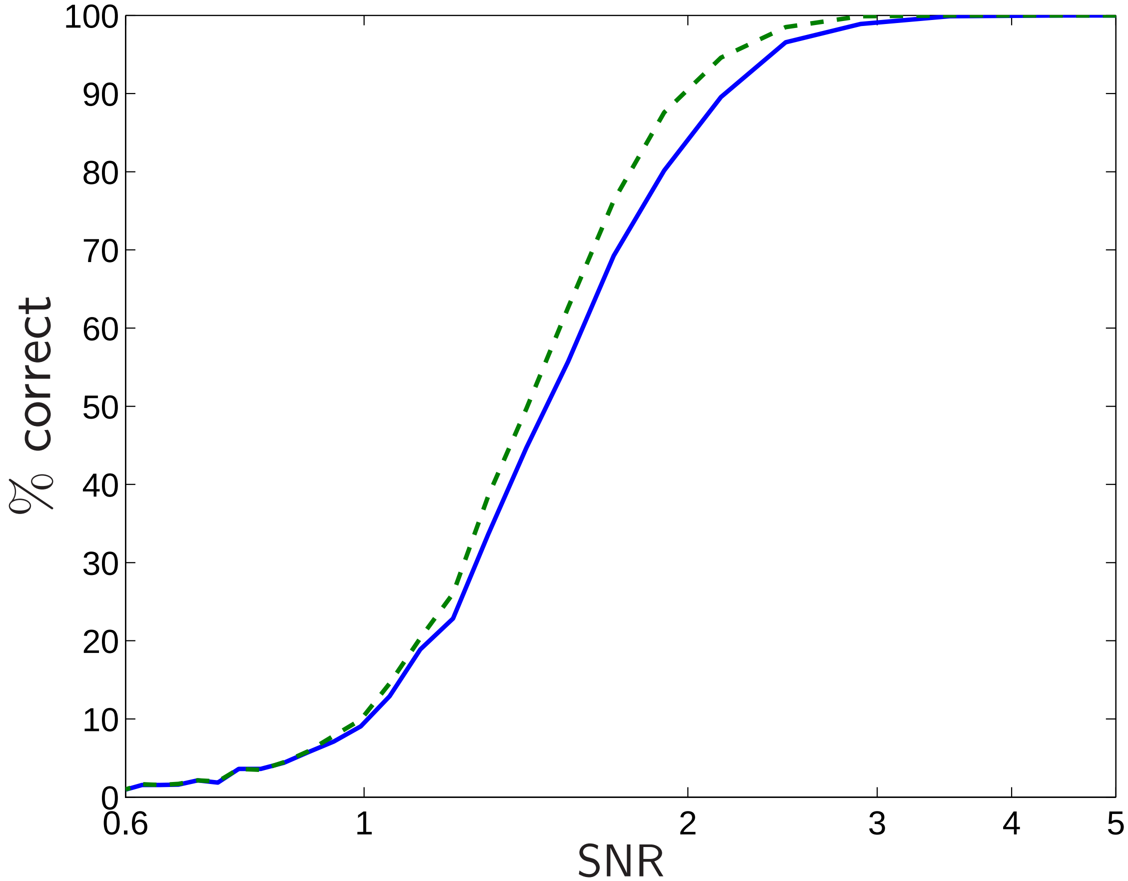

blue: RMAP; green: RMAP+LO

Performance vs SNR: Top 10 Hit Percentage

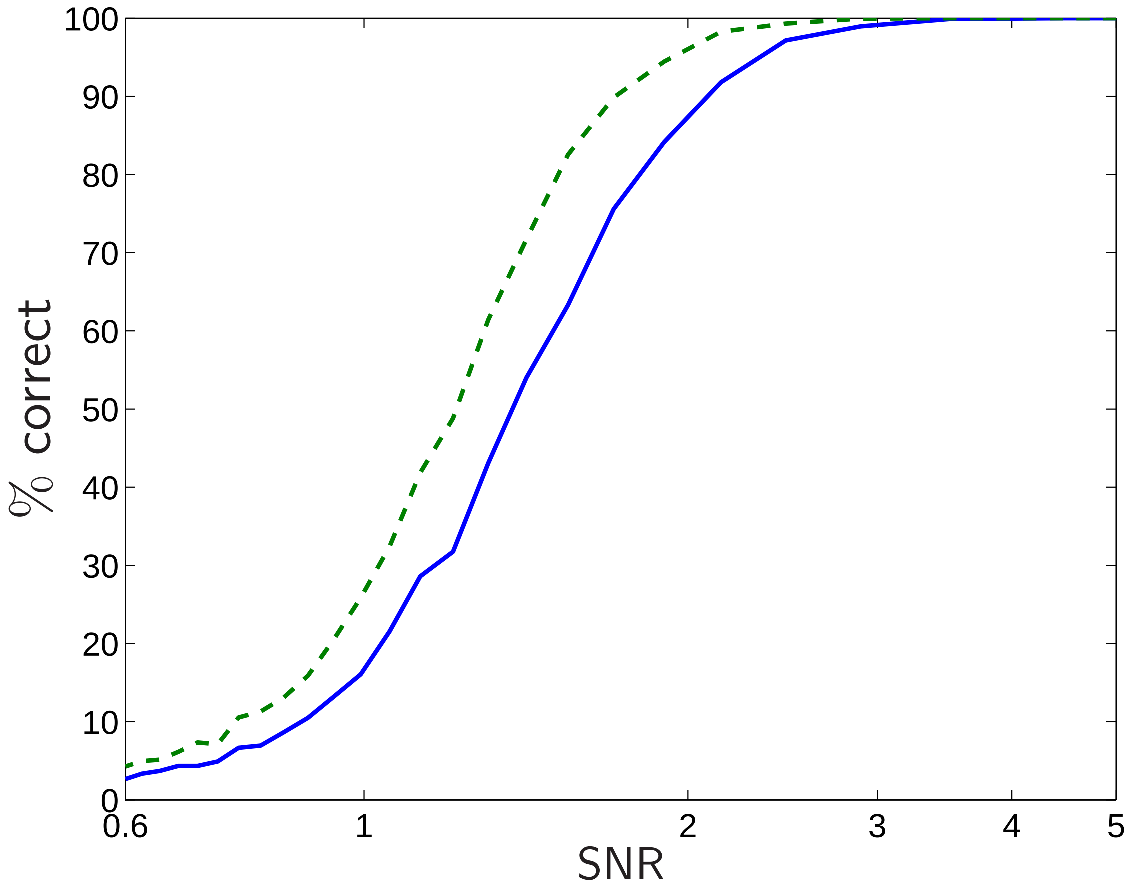

blue: RMAP; green: RMAP+LO

Larger Example

- $m = 10000$ measurements
- $n = 2000$ potential faults
- A is sparse, with 1% of entries nonzero; nonzero entries $\mathcal{N}(0,1)$ (so, each measurement affected by around 20 faults)
- $p_j = 0.002$ (so, around 4 faults on average)
- $\sigma = 1.5$, corresponding to SNR ≈ 0.5
- runtime about 20s (Matlab implementation; would be much faster in C)
- over many runs, probability of exact fault detection is 92%

Larger Example: Results from Single Run

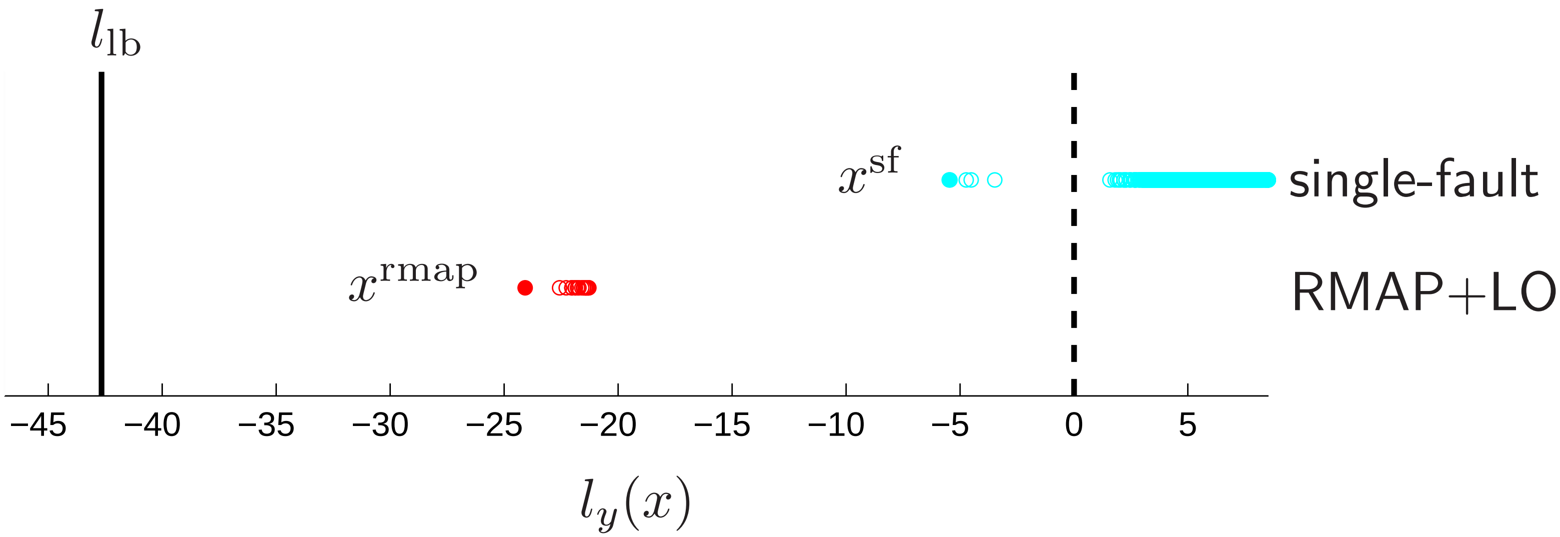

Conclusions

- when only one fault occurs, fault estimation is 'easy'
- when multiple faults can occur, fault estimation is hard
- we've proposed a new method for this case, leveraging convex optimization
- the method is efficient, scalable, and seems to work well
- extension to quantized measurements also works well, even with very coarse quantization

Extensions & Variations

- quantized measurements
- more sophisticated local optimization
- correlated Gaussian measurement noise
- Laplacian, other noise distributions
- (logical) constraints on faults, $e.g.$, $z_i \leq z_j$
- more general measurement nonlinearity
  (use linearized $A(x)$ in Newton method)
- dynamic (time-varying) case

References

- *Relaxed Maximum a Posteriori Fault Identification*,
  A. Zymnis, S. Boyd, and D. Gorinevsky,
  Signal Processing 2009

- *Mixed Linear System Estimation and Identification*,
  A. Zymnis, S. Boyd, and D. Gorinevsky,
  to appear, *Signal Processing 2010*

- *Mixed State Estimation for a Linear Gaussian Markov Model*,
  A. Zymnis, S. Boyd, and D. Gorinevsky,
  to appear, *IEEE Signal Processing Letters*

(full source code for all available on-line)