Nanjing University of Aeronautics and Astronautics (NUAA)

27 October 2017

Decision fusion with corrupted reports in multi-sensor networks

Mauro Barni

University of Siena, Italy

Summary

- Introduction and motivation
- Distributed detection in adversarial setting
- Asymptotic Information-theoretic analysis
- Decision fusion with byzantine nodes
  - Optimum decision fusion: a game-theoretic approach
  - A simplified approach based on message passing
- Conclusions and directions for future research

Distributed detection setup

$$\mathbf{x}_i = (x_{i,1}, x_{i,2} \ldots x_{i,m})$$

Observation vector available to i-th node

$$\mathbf{u}_i = (u_{i,1}, u_{i,2} \ldots u_{i,m})$$

Report sent to FC by i-th node

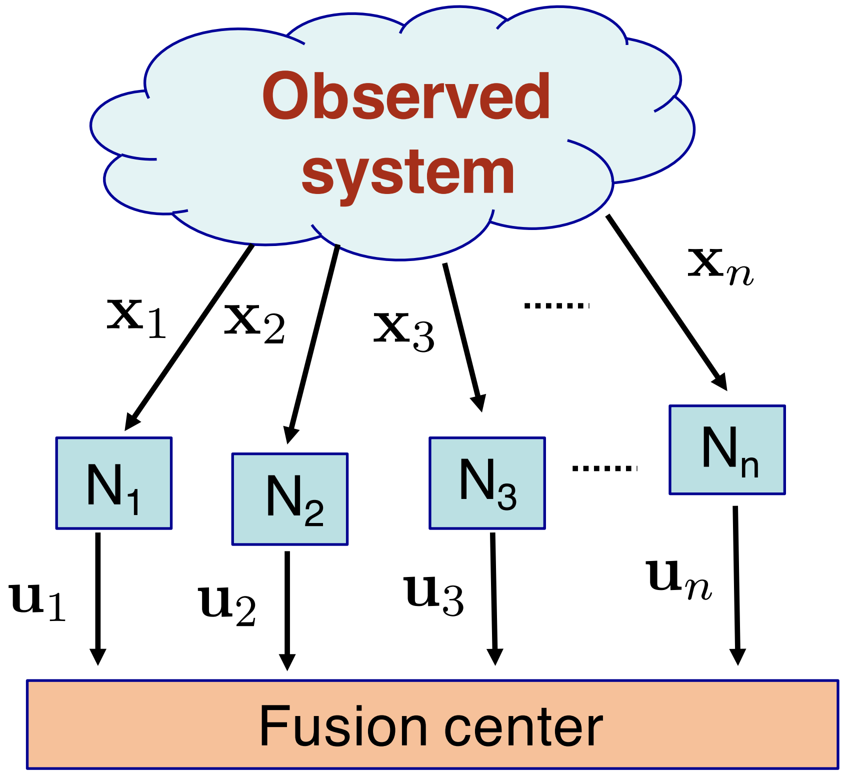

- FC performs a *Binary Hypothesis Test* on system state.
- The test often aims at detecting when the system exits a safe state S_0

A wide variety of applications

– Wireless sensor networks
– Spectrum sensing for cognitive radio
– Intrusion detection
– Network monitoring
– Anomaly detection
– Smart grid
– Social networks
– Reputation systems
– Multi-clue decision making

Distributed detection in adversarial setting

- An attacker may corrupt part of the system to induce a decision error
- Different versions:

Corrupted observations

Corrupted nodes

Corrupted reports

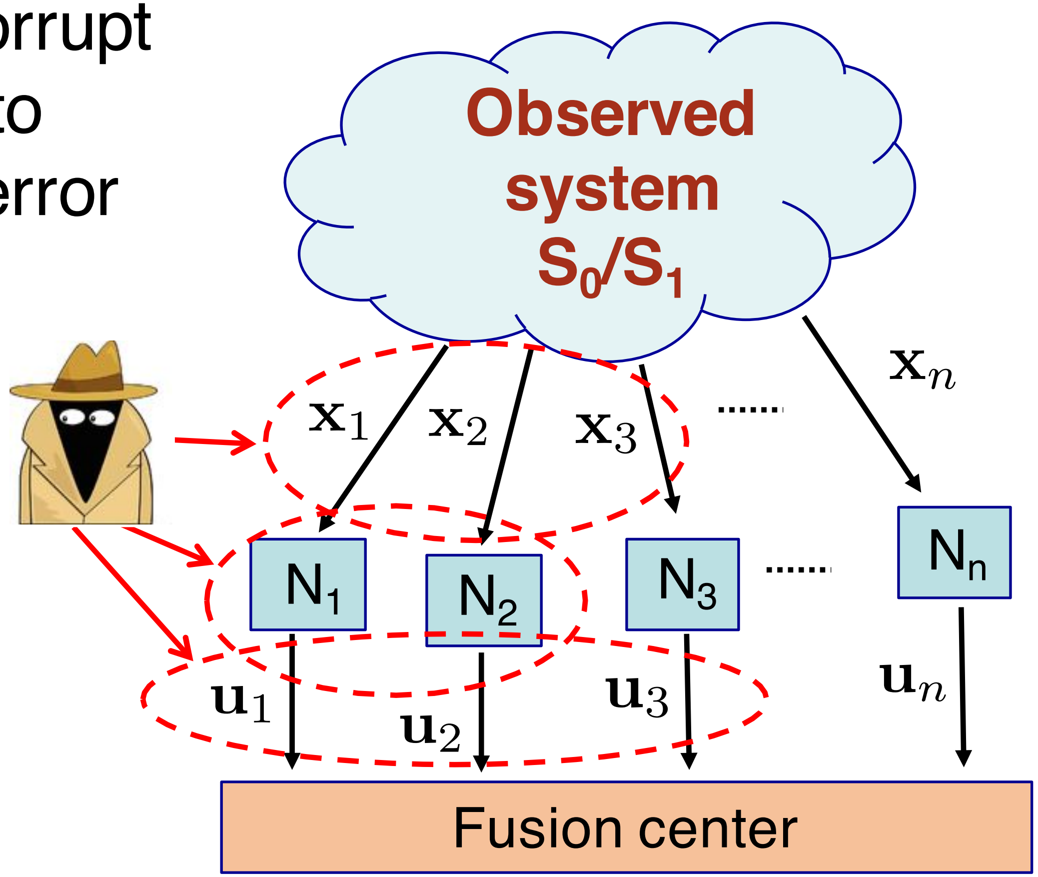

Asymptotic Information-theoretic analysis

Basic assumptions

- System state does not change over time
- Number of observations for each node goes to infinity ($m \rightarrow \infty$)
- **Game-theoretic approach**
- Similarity with SI game [1], solution provided in [2]

[1] M.Barni, B.Tondi, The Source Identification Game: an Information-Theoretic Perspective, IEEE Trans. on Information Forensics and Security, vol. 8, no. 3, pp. 450 –463, March 2013.

[2] M. Barni, B. Tondi, "Multiple-Observation Hypothesis Testing under Adversarial Conditions", Proc. of WIFS 2013, IEEE Int. Workshop on Information Forensics and Security, Ghuanzhou, China, 18-21 November 2013, pp. 91-96.

Game Theory in a nutshell

Two-player game

$G(S_1, S_2, u_1, u_2)$

$S_1 = \{s_{1,1}, s_{1,2} \dots s_{1,n1}\}$ Set of strategies available to first player

$S_2 = \{s_{2,1}, s_{2,2} \dots s_{n2}\}$ Set of strategies available to second player

$\mathrm{u}_1(s_{1,i}, s_{2,j})$ Payoff of first player for a given profile

$\mathrm{u}_2(s_{1,i}, s_{2,j})$ Payoff of second player for a given profile

Competitive (zero-sum) game

$\mathrm{u}_1(\cdot,\cdot) = -\mathrm{u}_2(\cdot,\cdot)$

Sequential vs strategic vs multiple moves games

Equilibrium

Optimal choices

In game theory we are interested in the optimal choices of rational players

(stricly) Dominant strategy

The best strategy regardless of the other player's move

$$u_1(s_1^*, s_2) > u_1(s_1, s_2) \quad \forall s_1 \in S_1 \quad \forall s_2 \in S_2$$

... then equilibrium is

$$(s_1^*, s_2^*) \text{ with } s_2^* \text{ such that}$$

$$u_2(s_1^*, s_2^*) \geq u_2(s_1^*, s_2) \quad \forall s_2 \in S_2$$

Equilibrium

Nash equilibrium

No player gets an advantage by changing his strategy assuming the other does not change his own

$$u_1(s_1^*, s_2^*) \geq u_1(s_1, s_2^*) \quad \forall s_1 \in S_1$$

$$u_2(s_1^*, s_2^*) \geq u_2(s_1^*, s_2) \quad \forall s_2 \in S_2$$

... and many others

- worst case assumption
- rationalizable equilibrium
- ...

The SI game (with multiple observations)

Payoff and structure of the game: ***Neyman-Pearson***

- *D* aims at minimizing the false negative error probability P_{fn} under the constraint that P_{fp} stays below a threshold.
- Omniscient A. He/she acts only under S1, his aim being the maximization of P_{fn}
- ➢ **Zero-sum game: $u_A = -u_D = P_{fn}$**

Space of D's strategies

- All detection regions based on *on first order (possibly joint) statistics*;
- *Asymptotic version* of the problem: constraint on asymptotic decay rate of P_{fp} ($P_{fp} < 2^{-\lambda m}$)

Several versions of the game

- D has full knowledge of system statistics and bases the decision on all the available information still relying on first order statistics

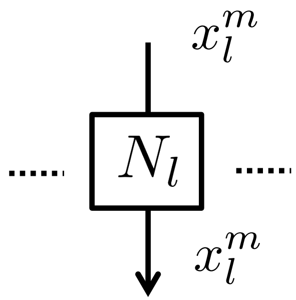

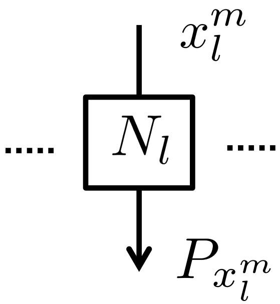

- D still has full knowledge of system statistics but observes only the marginals

- D has full knowledge of system statistics but decides by fusing local decisions

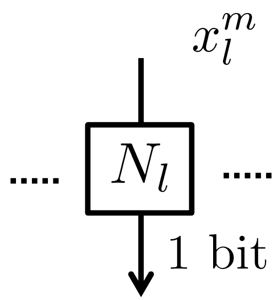

Some noticeable results proven in [2]

- The game theoretic formulation of the problem is dominance solvable
- Optimum fusion strategy checks if the joint empirical pmf of the observations is in accordance with the expected one. For the full statistics case we have

$$\Lambda_0^* = \left\{ \hat{P} \in \mathcal{P}_m : \mathcal{D}(\hat{P}||P_{\mathbf{x}}) < \lambda - |\mathcal{X}|^k \frac{log(m+1)}{m} \right\}$$

- The optimum fusion strategy does NOT pass from the identification of malevolent nodes
- Under certain assumptions, reliable decision is possible even in the presence of only one uncorrupted node

[2] M. Barni, B. Tondi, "Multiple-Observation Hypothesis Testing under Adversarial Conditions", Proc. of WIFS 2013, IEEE Int. Workshop on Information Forensics and Security, Ghuanzhou, China, 18-21 November 2013, pp. 91-96.

Decision fusion with Byzantines

Decision fusion with Byzantines

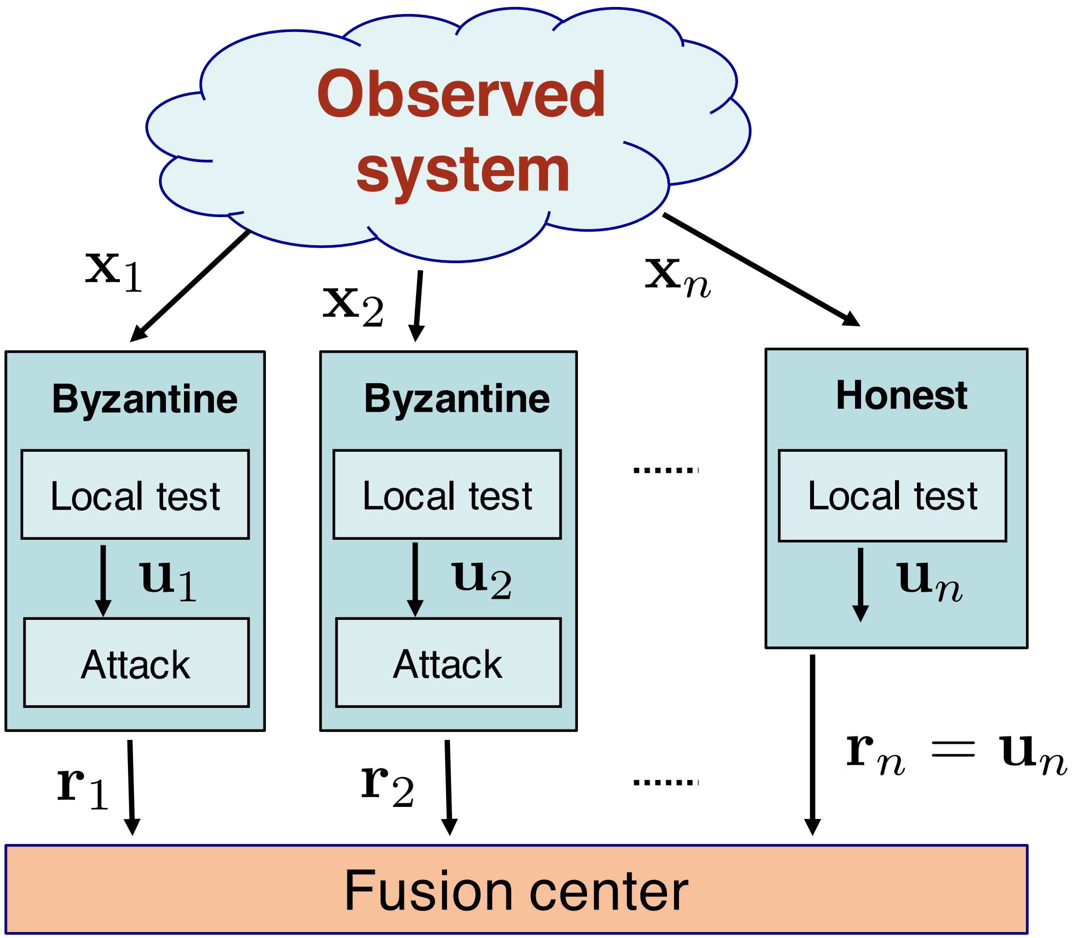

- **Now system state changes over time**
- The fusion center makes its choice based on the results of the local decisions made at the nodes
- Global decision on *m* states
- Corrupted nodes (called Byzantines [3]) may submit wrong reports

[3] A. Vempaty, L. Tong, P. Varshney, “Distributed Inference with Byzantine Data”, Signal Processing Magazine, vol. 30, no. 5, September 2013

Possible approaches

- Byzantines isolation
  - A. S. Rawat, P. Anand, H. Chen, and P. K. Varshney, “Collaborative spectrum sensing in the presence of Byzantine attacks in cognitive radio networks,” IEEE Trans. Signal Process., vol. 59, no. 2, pp. 774–786, Feb. 2011.
  - A. Abrardo, M. Barni, K. Kallas, and B. Tondi, “Decision fusion with corrupted reports in multi-sensor networks: A game-theoretic approach,” in Proc. IEEE Conf. Decision Control (CDC), Los Angeles, CA, USA, Dec. 2014, pp. 505–510.
- Byzantine-tolerant schemes
  - M. Gagrani, P. Sharma, S. Iyengar, V. Nadendla, A. Vempaty, H. Chen, and P. Varshney, “On noise-enhanced distributed inference in the presence of Byzan-tines,” in Proc. 49th Annu. Allerton Conf. Communications Control Comput-ing, Sept. 2011, pp. 1222–1229.
- **Optimum fusion**

System and attack model

- Equiprobable independent system states

$$P_{S_i}(0) = P_{S_i}(1) = 0.5$$

- Constant and independent local decisions errors
- Symmetric local decision errors

$$\varepsilon = \Pr(U_{i,j} \neq S_j)$$

- Byzantines flip local decision with probability P_{mal}

$$\Pr(U_{i,j} \neq R_{i,j} \mid \text{node is Byzantine}) = P_{mal}$$

- Byzantines flip decisions independently of each other (non cooperative malicious nodes) and on subsequent states
- Nodes status and Byzantines' strategy do not change over time

Optimum fusion rule

If all the parameters of the system are known the optimum decision rule at the FC can be derived as follows

$$s^{m,*} = \arg\max_{s^m} P(s^m|\mathbf{r})$$ MAP estimate

$$s^{m,*} = \arg\max_{s^m} P(\mathbf{r}|s^m)$$ ML estimate

$$s^{m,*} = \arg\max_{s^m} \sum_{a^n} P(\mathbf{r}|a^n, s^m) P(a^n)$$

$$= \arg\max_{s^m} \sum_{a^n} \left(\prod_{i=1}^{n} P(\mathbf{r}_i|a_i, s^m) \right) P(a^n)$$

$$= \arg\max_{s^m} \sum_{a^n} \left(\prod_{i=1}^{n} \prod_{j=1}^{m} P(r_{ij}|a_i, s_j) \right) P(a^n).$$

s^m = sequence of system states

a^n = vector with states of nodes

Optimum fusion rule

$\delta = \varepsilon(1 - P_{mal}) + (1 - \varepsilon)P_{mal}$ Prob that FC receives a wrong report

$m_{eq}(i)$ Number of times for which the report of node *i* is equal to the state

$$s^{m,*} = \arg\max_{s^m} \sum_{a^n} \left(\prod_{i:a_i=0} (1-\varepsilon)^{m_{eq}(i)} \varepsilon^{m-m_{eq}(i)} \prod_{i:a_i=1} (1-\delta)^{m_{eq}(i)} \delta^{m-m_{eq}(i)} \right) P(a^n)$$

To go on it is necessary to make some assumptions on the distribution of byzantine nodes across the network: $P(a^n)$

Byzantines distribution

1. **Unconstrained maximum entropy distribution**

   Letting P_{mal} = 1 forces the mutual information between *S* and *R* to zero making any meaningful decision impossible

2. **Constrained maximum entropy distribution, fixed $E[N_B]$**

   Entropy is maximized by assuming i.i.d. node states with

$$\alpha = Pr(A_i = 1) = E[N_B]$$

$$\arg\max_{s^m} \prod_{i=1}^{n} \left[(1-\alpha)(1-\varepsilon)^{m_{eq}(i)} \varepsilon^{m-m_{eq}(i)} + \alpha(1-\delta)^{m_{eq}(i)} \delta^{m-m_{eq}(i)} \right]$$

The complexity of the optimum fusion rule is linear in *n* and exponential in *m*

Byzantines distribution

3. Constrained maximum entropy distribution, $N_B < n/2$

Equiprobable a^n (only those for which $N_B < n/2$)

Complexity of optimum fusion rule is exponential in *m* and quadratic in *n* (dynamic programming [3])

[4] A. Abrardo, M. Barni, K. Kallas, B. Tondi, "A Game-Theoretic Framework for Optimum Decision Fusion in the Presence of Byzantines", *IEEE Trans. Information Forensics and Security*, vol.11, no. 6, 2016

$$s^{m,*} = \arg\max_{s^m} \sum_{I \in \mathcal{I}_{n_B}} \left(\prod_{i \in I} (1-\delta)^{m_{eq}(i)} \delta^{m-m_{eq}(i)} \prod_{i \in \mathcal{I} \setminus I} (1-\varepsilon)^{m_{eq}(i)} \varepsilon^{m-m_{eq}(i)} \right)$$

A game theoretic perspective

- Application of the optimum fusion rule requires that the FC knows P_{mal}
- Large values of P_{mal} are more effective in inducing a decision error
- If byzantine nodes are identified $P_{mal} = 1$ does not make any harm
- With $P_{mal} = 0.5$ we have $I(S,R) = 0$

- Which value of P_{mal} should the Byzantines choose?
- How can the FC know the vale of P_{mal} ?

- **We adopt a game-theoretic perspective**

Decision fusion with Byzantines game

Two-player game (Byzantines collectively playing as a single player)

$$\mathcal{S}_B = \{P_{mal}^{B} \in [0,1]\}$$
$$\mathcal{S}_{FC} = \{P_{mal}^{FC} \in [0,1]\}$$

Payoff equal to error probability at the fusion center

Strategic game

Computation of the equilibrium point

- Run simulations by quantizing the set of strategies

$$P_{mal} = \{0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$$

- Length of observation window *m* plays a major role
- We run simulations with small and medium values of *m*
- Show results for $n = 20$, $\varepsilon = 0.1$
  - $m = 4$
  - $m = 10$

Small *m*, independent node states

P^B_{mal}/P^{FC}_{mal}	0.5	0.6	0.7	0.8	0.9	1.0
0.5	0.33	0.37	0.44	0.58	0.73	0.85
0.6	0.60	0.54	0.59	0.70	0.80	1.14
0.7	1.38	1.20	1.19	1.24	1.29	2.40
0.8	3.88	3.56	3.36	3.31	3.35	6.03
0.9	9.93	9.61	9.57	9.55	9.54	11.96
1.0	20.33	20.98	21.70	21.90	21.84	**19.19**

$\alpha = 0.4$, n = 20

m = 4

$P_e \times 10^2$

P^B_{mal}/P^{FC}_{mal}	0.5	0.6	0.7	0.8	0.9	1.0
0.5	0.62	0.69	0.86	1.34	1.70	1.57
0.6	1.23	1.15	1.26	1.84	2.18	2.38
0.7	2.94	2.64	2.57	3.00	3.14	5.33
0.8	7.89	7.39	7.03	6.74	6.81	12.73
0.9	18.45	17.94	17.63	17.08	17.07	22.78
1.0	34.39	34.62	34.84	36.66	36.61	**33.14**

$\alpha = 0.45$, n = 20

m = 4

$P_e \times 10^2$

Small *m*, fixed number of Byzantines

P^B_{mal}/P^{FC}_{mal}	0.5	0.6	0.7	0.8	0.9	1.0
0.5	**3.80**	**3.80**	4.60	7.60	12.0	29.0
0.6	3.60	3.45	3.90	5.20	8.0	17.0
0.7	3.45	2.80	2.80	3.10	4.40	8.75
0.8	4.10	2.85	2.15	2.05	2.25	3.25
0.9	3.55	2.05	1.40	0.95	0.70	0.75
1.0	2.05	0.90	0.35	0.15	0.05	0.05

$N_B = 6$, $n = 20$

$m = 4$

$P_e \times 10^4$

P^B_{mal}/P^{FC}_{mal}	0.5	0.6	0.7	0.8	0.9	1.0
0.5	0.22	0.24	0.33	0.63	1.41	4.13
0.6	0.27	0.24	0.27	0.41	0.78	2.03
0.7	0.32	0.24	0.23	0.26	0.37	0.82
0.8	0.54	0.45	0.39	0.36	0.41	0.59
0.9	2.04	1.87	1.76	1.58	1.56	1.66
1.0	9.48	8.76	8.37	6.72	5.88	**5.51**

$N_B = 9$, $n = 20$

$m = 4$

$P_e \times 10^2$

Small *m*, fixed number of Byzantines

P^B_{mal}/P^{FC}_{mal}	0.5	0.6	0.7	0.8	0.9	1.0
0.5	1.2	1.4	1.9	3.1	6.3	18.9
0.6	1.5	1.4	1.4	2.0	3.7	10.0
0.7	1.4	1.1	0.945	1.1	1.7	4.0
0.8	1.4	0.95	0.715	0.58	0.675	1.2
0.9	2.1	1.4	0.995	0.745	0.71	0.78
1.0	7.3	5.7	5.3	3.7	3.0	2.9

$N_B = 8$, $n = 20$

$m = 4$

$P_e \times 10^4$

Nash equilibrium exists only in mixed strategies

	0.5	0.6	0.7	0.8	0.9	1.0
$P(P^B_{mal})$	0.179	0	0	0	0	0.821
$P(P^{FC}_{mal})$	0	0	0	0.844	0.156	0
$P^*_e = 3.8e-4$						

Medium *m*, independent node states

P^B_{mal}/P^{FC}_{mal}	0.5	0.6	0.7	0.8	0.9	1.0
0.5	0.11	0.13	0.19	0.73	2.16	0.68
0.6	0.11	8.32e-2	9.96e-2	0.26	0.67	1.30
0.7	0.18	7.66e-2	6.62e-2	9.52e-2	0.18	4.87
0.8	1.10	0.60	0.33	0.24	0.28	10.41
0.9	5.77	4.75	3.95	3.53	3.41	13.44
1.0	20.41	21.26	22.65	24.27	26.21	**18.72**

$\alpha = 0.4$, $n = 20$

$m = 10$

$P_e \times 10^2$

P^B_{mal}/P^{FC}_{mal}	0.5	0.6	0.7	0.8	0.9	1.0
0.5	0.20	0.23	0.47	2.88	10.92	1.26
0.6	0.22	0.18	0.24	0.80	2.85	2.93
0.7	0.50	0.19	0.15	0.23	0.65	10.64
0.8	2.61	1.24	0.63	0.41	0.59	20.65
0.9	11.74	9.28	7.08	5.65	5.21	25.85
1.0	34.25	34.94	36.01	37.74	39.87	**33.17**

$\alpha = 0.45$, $n = 20$

$m = 10$

$P_e \times 10^2$

Medium *m*, fixed number of Byzantines

P_{mal}^{B}/P_{mal}^{FC}	0.5	0.6	0.7	0.8	0.9	1.0
0.5	**1.22**	**1.22**	1.40	2.20	5.06	11.0
0.6	1.12	0.94	1.02	1.26	2.56	5.34
0.7	1.22	0.58	0.56	0.64	0.98	2.06
0.8	1.22	0.36	0.32	0.28	0.30	0.56
0.9	1.40	0.20	0.18	0.16	0.10	0.18
1.0	1.52	0.14	0.14	0.10	6e-2	4e-2

$N_B = 6$, $n = 20$

$m = 10$

$P_e \times 10^4$

For $N_B = 8$ and $N_B = 9$, a Nash equilibrium exists only in mixed strategies

	0.5	0.6	0.7	0.8	0.9	1.0
$P(P_{mal}^{B})$	0.4995	0	0	0	0	0.5005
$P(P_{mal}^{FC})$	0	0	0.66	0.34	0	0
$P_e^* = 1.58e-3$						

$N_B = 9$, $n = 20$

$m = 10$

Performance at the equilibrium

	Maj	HardIS	SoftIS	OPT
Independent nodes, $\alpha = 0.3$	0.073	0.048	0.041	0.035
Independent nodes, $\alpha = 0.4$	0.239	0.211	0.201	0.192
Independent nodes, $\alpha = 0.45$	0.362	0.344	0.338	0.331
Fixed n. of nodes $n_B = 6$	0.017	0.002	6.2e-4	3.8e-4
Fixed n. of nodes $n_B = 8$	0.125	0.044	0.016	0.004
Fixed n. of nodes $n_B = 9$	0.279	0.186	0.125	0.055
Max entropy with $N_B < n/2$	0.154	0.086	0.052	0.021
Max entropy with $N_B < n/3$	0.0041	5e-4	2.15e-4	1.9e-4

$\varepsilon = 0.1$

$n = 20$

$m = 4$

[Maj] Majority rule

[HardIS] A. S. Rawat, P. Anand, H. Chen, and P. K. Varshney, “Collaborative spectrum sensing in the presence of Byzantine attacks in cognitive radio networks,” IEEE Trans. Signal Process., vol. 59, no. 2, pp. 774–786, Feb. 2011.

[SoftIS] A. Abrardo, M. Barni, K. Kallas, and B. Tondi, “Decision fusion with corrupted reports in multi-sensor networks: A game-theoretic approach,” in Proc. IEEE Conf. Decision Control (CDC), Los Angeles, CA, USA, Dec. 2014, pp. 505–510.

Nearly-optimum decision fusion

- Complexity prevents the use of optimum decision fusion for large *m*
- Use of message passing (MP) to develop a fast nearly optimum detector at the FC
- MP is a nearly optimum iterative optimization procedure based computation on graphs theory
- The MP-based algorithm allows to extend our results to cases with large observation windows [5]

[5] A. Abrardo, M. Barni, K. Kallas, B. Tondi, "A Message Passing Approach for Decision Fusion in Adversarial Multi-Sensor Networks", *Information Fusion*, vol. 40, March 2018, pp. 101-111

Results (small m)

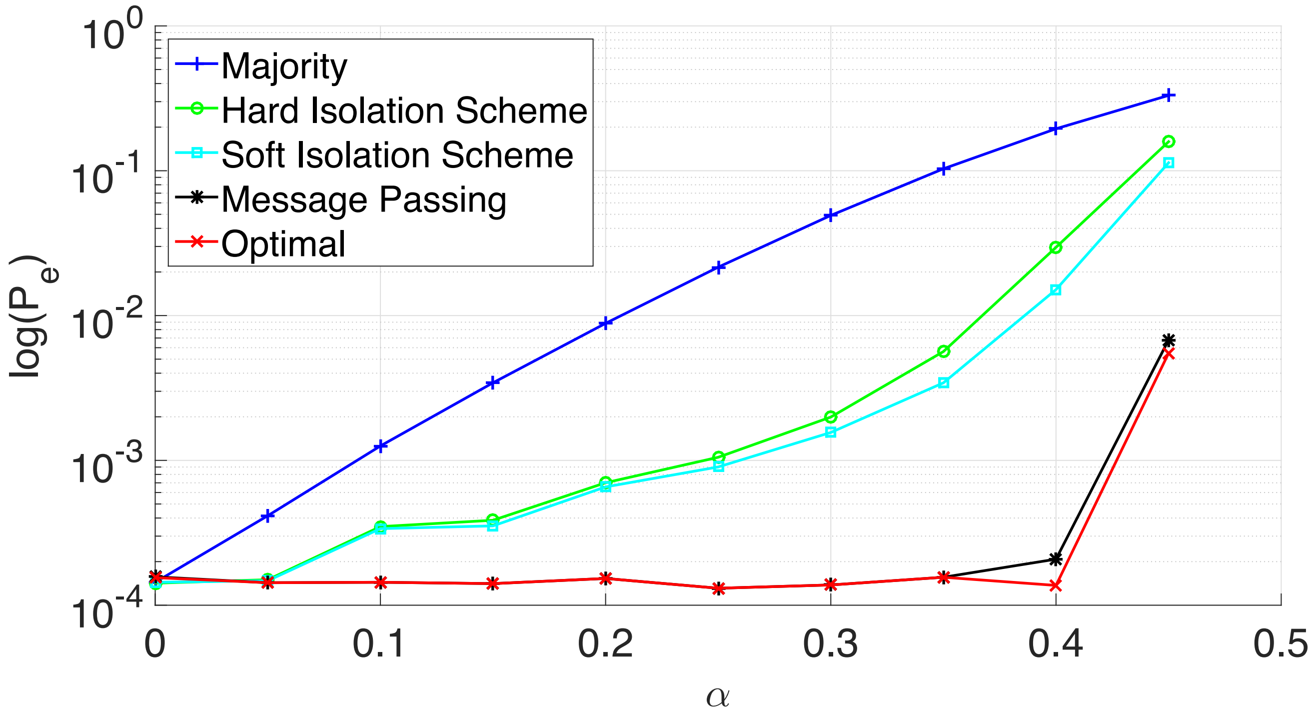

$\varepsilon = 0.15$

$n = 20$

$m = 10$

$P_{mal} = 1$

$\rho = 0.5$

We can now evaluate the performance also when *m* is large and for markovian sources

Results

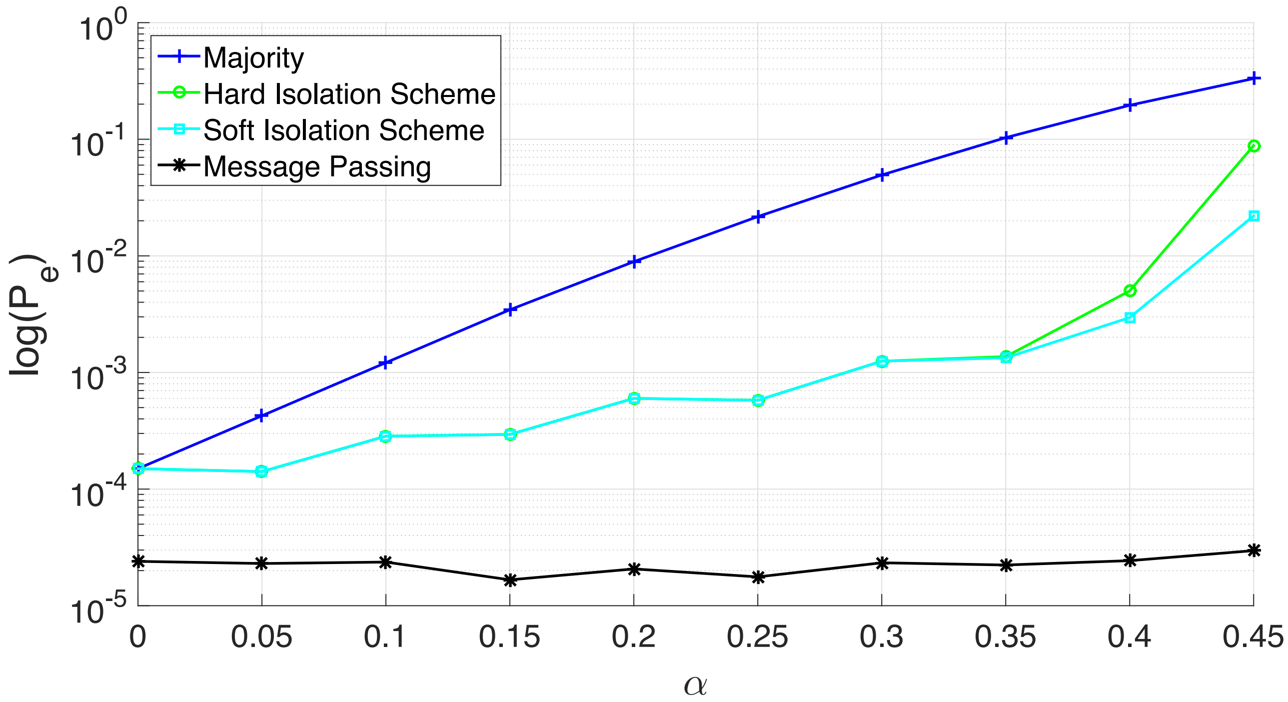

$\varepsilon = 0.15$

$n = 20$

$m = 30$

$P_{mal} = 1$

$\rho = 0.95$

For large m the optimum detector can not be applied. The performance of MP-fusion remain very good

Results: optimum attack strategy

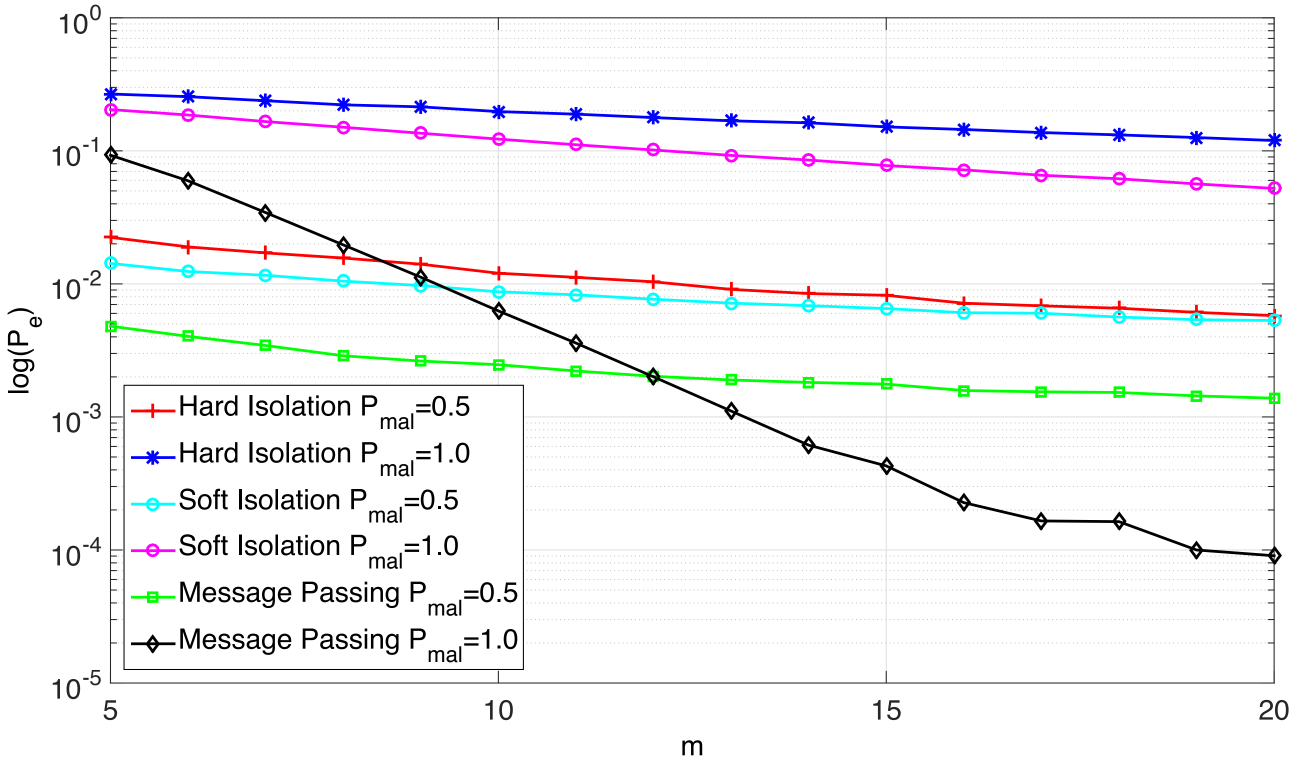

$\varepsilon = 0.15$

$n = 20$

$\alpha = 0.45$

$\rho = 0.5$

The tendency of passing from $P_{mal} = 1$ to $P_{mal} = 0.5$ for large values of m is confirmed (for nearly optimum decision fusion)

Synchronized attack

- Using a synchronized attack may increase significantly the effectiveness of the attack
- We assume that the Byzantines share the values assumed by a local source of randomness **q** = (q_1, q_2 … q_m)
- The optimum fusion rule can be easily derived by incorporating the value assumed by the local randomness into the maximization

$$s_i^* = \arg \max_{s_i \in \{0.1\}} \sum_{\{\mathbf{sqa}\} \backslash s_i} \prod_{i,j} p(r_{ij}|s_i, q_i, a_j) \prod_h p(s_h|s_{h-1}) \prod_k p(q_k|q_{k-1}) \prod_l p(a_l)$$

- Which can be implemented again by exploiting the sum product MP algorithm [6]

[6] A. Abrardo, M. Barni, K. Kallas, B. Tondi, "A Message Passing Approach for Decision Fusion of Hidden-Markov Observations in the presence of Synchronized Attacks", Proc. of MMEDIA17, 9-th Int. Conf. on Advances in Multimedia, April 23-27, 2017, Venice, Italy.

Results

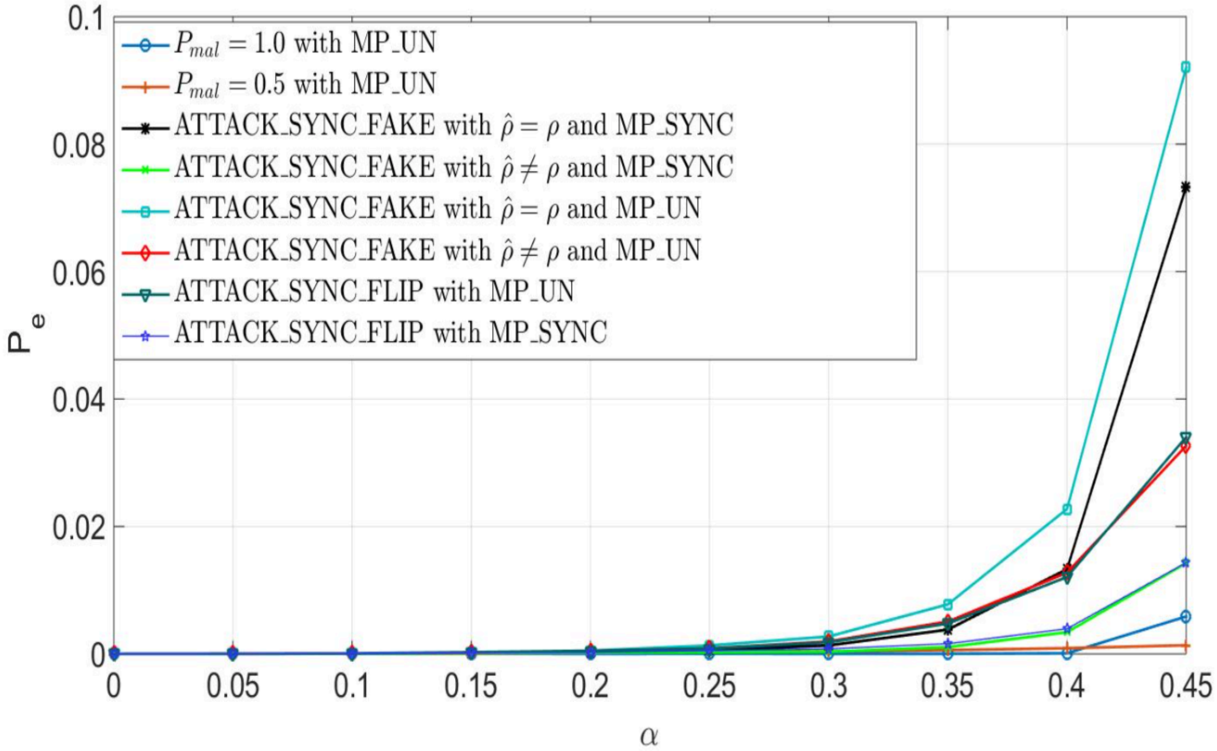

$\varepsilon = 0.15$

$n = 20$

$m = 10$

$\rho = \{0.5, 0.95\}$

The synchronized attack is by far more powerful than the asynchronous one. Game-theoretic analysis still on-going.

Conclusions and future research

- The case studied here is only an oversimplified example
- Many interesting extensions are possible:
  - Time varying attacks
  - Allow communication among Byzantines
  - Non-binary reports
  - Coalition games
  - …
- Distributed detection
  - K. Kallas, B. Tondi, M. Barni, “Consensus Algorithm with Censored Data for Distributed Detection with Corrupted Measurements: A Game-Theoretic Approach”, *Proc. of GameSec 2016, Conference on Decision and Game Theory for Security,* November 2-4, 2016, New York, NY, USA

Conclusions and future research

- Application to real cases
  - Network monitoring
  - Wireless sensor networks
    - Surveillance
    - Drone detection
    - …
  - Social networks
    - Crowdcomputing
- Implementation in testbed

Thank you
for your attention