# IEEE AESS Distinguished Lecture

## Distributed MHT and Extensions

**Dr. Stefano Coraluppi**
STR Chief Scientist
IEEE Fellow

Online, 5 November 2024

# Outline

- Multi-target tracking (MTT)
- Multiple-hypothesis tracking (MHT)
- Advanced centralized MHT
- Distributed MHT
- Simplified distributed MHT

# Multi-target tracking (MTT)

- We are given a sequence of *sets* of measurements, and are to determine a *set* of trajectories
  - Unknown, time-varying number of targets
- An intractable posterior probability distribution – both *computationally* and *conceptually*
  - $p(X^k|Z^k)$
  - Most approaches do not seek to evaluate this distribution; a rare exception is the JMPD by Kreucher et al. (2005);
  - Even if we could do so, what is the MTT output? There is a MAP estimation difficulty as noted by Mahler (2014)

# Some comments on MTT methods

- Some MTT methods focus primarily on multi-target filtering, with no automated track management
  - GNN, JPDA, PMHT
  - Track management is handled externally, e.g. *Integrated JPDA* (JIPDA)
- Some MTT methods directly address both track management and multi-target filtering
  - MHT, MCMC, PHD, BP
- There is a distinct literature on *identity management* (IM) algorithms that addresses the coupling in target identity estimates

# Linear Gaussian state evolution

- We often assume uncorrelated noise across dimensions
  - $X_{k+1} = A_k X_k + w_k, w_k \sim N(0, Q_k)$
- Nearly constant position
  - $A_k = 1, Q_k = q\Delta t_k, \Delta t_k = t_{k+1} - t_k$
- Nearly constant velocity
  - $A_k = \begin{bmatrix} 1 & \Delta t_k \\ 0 & 1 \end{bmatrix}, Q_k = q \begin{bmatrix} \frac{(\Delta t_k)^3}{3} & \frac{(\Delta t_k)^2}{2} \\ \frac{(\Delta t_k)^2}{2} & \Delta t_k \end{bmatrix}$
- Ornstein Uhlenbeck
  - $A_k = \exp(-\gamma \Delta t_k), Q_k = q \frac{1-\exp(-2\gamma\Delta t_k)}{2\gamma}$

# Linear Gaussian state evolution

Nearly-constant position


Nearly-constant velocity

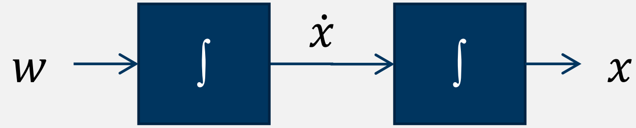


Nearly-constant acceleration


Ornstein-Uhlenbeck

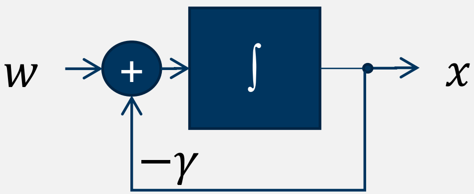


Integrated Ornstein-Uhlenbeck

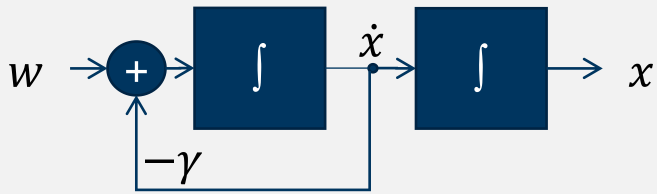


Singer model

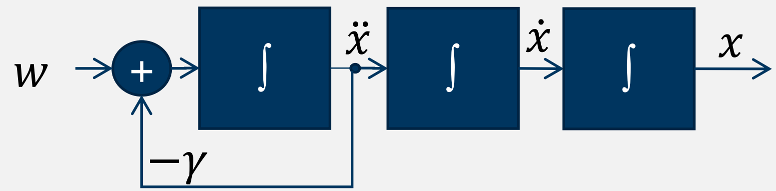


2nd order Ornstein-Uhlenbeck

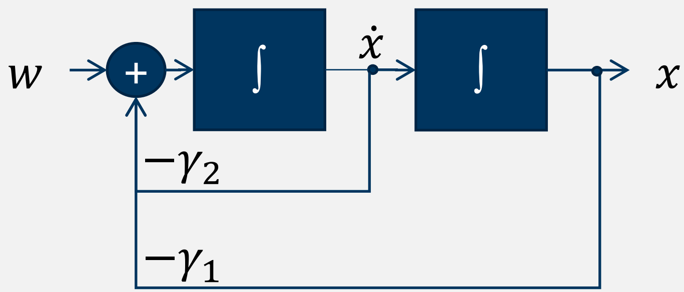

# Poisson target existence

- Initial time $t_0$. Discrete time sequence $t^k = (t_1, \ldots, t_k)$. $\Delta t_k = t_{k+1} - t_k$.
- Continuous-time birth-death process
  - Arrival rate $\lambda_b$, death rate $\lambda_\chi$
- Death probability
  - $p_\chi(\Delta t_k) = \int_{t_k}^{t_{k+1}} \lambda_\chi e^{-\lambda_\chi \tau} d\tau = 1 - e^{-\lambda_\chi \Delta t_k}$
- Poisson birth rate
  - $\mu_b(\Delta t_k) = \int_{t_k}^{t_{k+1}} \lambda_b e^{-\lambda_\chi (t_{k+1} - \tau)} d\tau = \frac{\lambda_b}{\lambda_\chi}\left(1 - e^{-\lambda_\chi \Delta t_k}\right)$
- Number of births in non-overlapping intervals are *independent*

# Sensor modeling

- Detection statistics
  - Point target assumption with detection probability $p_d$
  - Sensor Poisson clutter with mean $\Lambda$ (uniformly distributed in measurement volume)
- Measurement statistics
  - $z_k = g(X_k) + v_k$
  - Additive (Gaussian) noise

# Operational performance metrics

- *Multi-target tracking* (MTT) yields a set of track trajectories
  – How do we compare this to the set of true trajectories?
- Operational performance metrics
  – Multiple measures are of interest: coupled and non-exhaustive
  – Most measures are not metrics in a mathematical sense
- Scalar performance metrics
  – OSPA, GOSPA, and track-level GOSPA
  – Scalar assessment (albeit with some parameters)
  – Satisfies requirements to be a mathematical metric

# Completeness, purity, estimation error

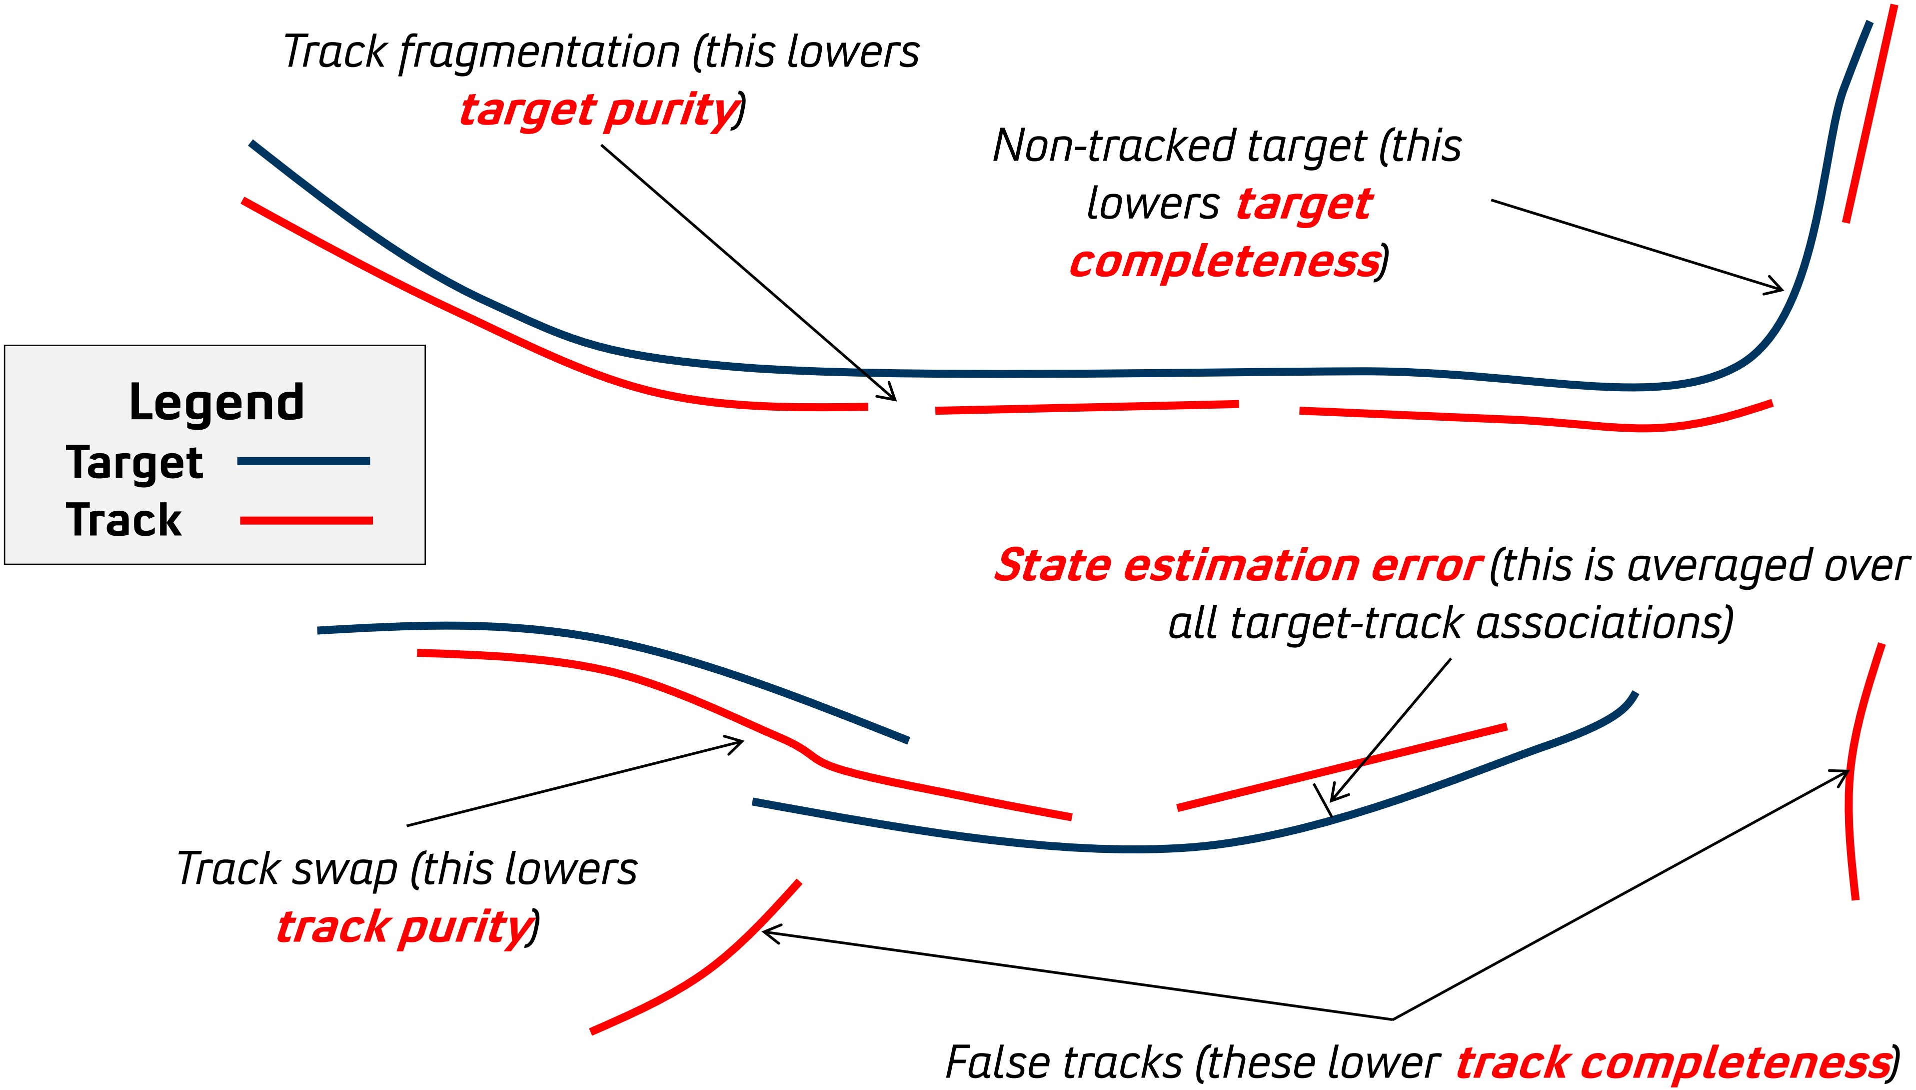

# Outline

- Multi-target tracking (MTT)
- Multiple-hypothesis tracking (MHT)
- Advanced centralized MHT
- Distributed MHT
- Simplified distributed MHT

# Multiple-hypothesis tracking (MHT)

- Multi-target tracking challenge: an intractable posterior probability distribution – *computationally* and *conceptually*
  - $p(X^k|Z^k)$
- Hybrid-state decomposition
  - $p(X^k|Z^k) = \sum_{q^k} p(X^k|Z^k, q^k)p(q^k|Z^k)$
- MHT approach uses *maximum a posteriori* (MAP) estimation
  - $\hat{q}^k = \arg\max_{q^k} p(q^k|Z^k)$ (MAP estimation here is well-posed)
  - $\hat{X}^k \approx \arg\max_{X^k} p(X^k|Z^k, \hat{q}^k)$
- Recursive formulation
  - $p(q^k|Z^k) = \frac{p(Z_k|Z^{k-1}, q^k)p(q_k|q^{k-1})p(q^{k-1}|Z^{k-1})}{p(Z_k|Z^{k-1})}$

# Track-oriented MHT recursion

- Global hypothesis recursion in factored form

$$
-\,p\left(q^k \middle| Z^k\right) = p\left(q^{k-1} \middle| Z^{k-1}\right) \cdot \frac{\left\{\frac{\exp(-\mu_b-\Lambda)\Lambda^r}{r!}\right\} \prod_{z_j \in Z_k} f_{fa}(z_j)}{p\left(Z_k \middle| Z^{k-1}\right)}
$$

Common to all global hypotheses

$$
\cdot\, p_\chi^\chi
$$

Target deaths

$$
\cdot \left((1-p_\chi)(1-p_d)\right)^{\tau-\chi-d}
$$

Target missed detections

$$
\cdot \prod_{j \in J_d} \frac{(1-p_\chi) p_d f\left(z_j \middle| Z^{k-1}, q^k\right)}{\Lambda f_{fa}(z_j)}
$$

Detection of previously existing targets

$$
\cdot \prod_{j \in J_b} \frac{p_d \mu_b f(z_j)}{\Lambda f_{fa}(z_j)}
$$

New target detections

# Example

- Online MHT identifies (approximately) the MAP data association solution over a sliding time window
  - Hypothesis forest depth is *n-scan* (set to one in example below)

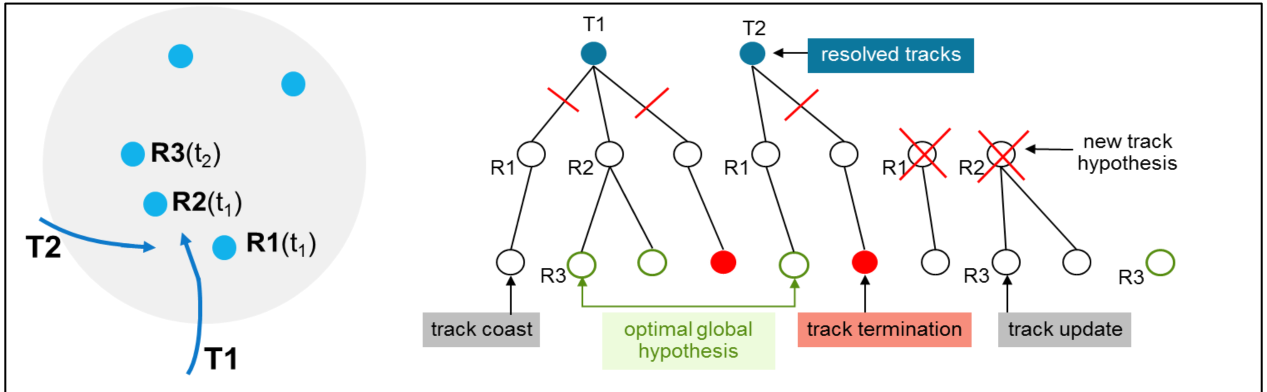


$$A = \underbrace{\begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \end{bmatrix}}_{\text{Track hypotheses}}$$

Account for tracks T1 and T2 (rows 1–2)

Account for reports R1, R2, and R3 (rows 3–5)

objective $\max c'x$

subject to $Ax = b$

$x_i \in \{0,1\}$

(for each vector element)

$$b = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

# More to MHT than scoring recursion

- The MHT equations prescribe a batch optimal solution; practical solutions involve judicious simplification
  - Sliding-window processing
  - Decouple data association and track management

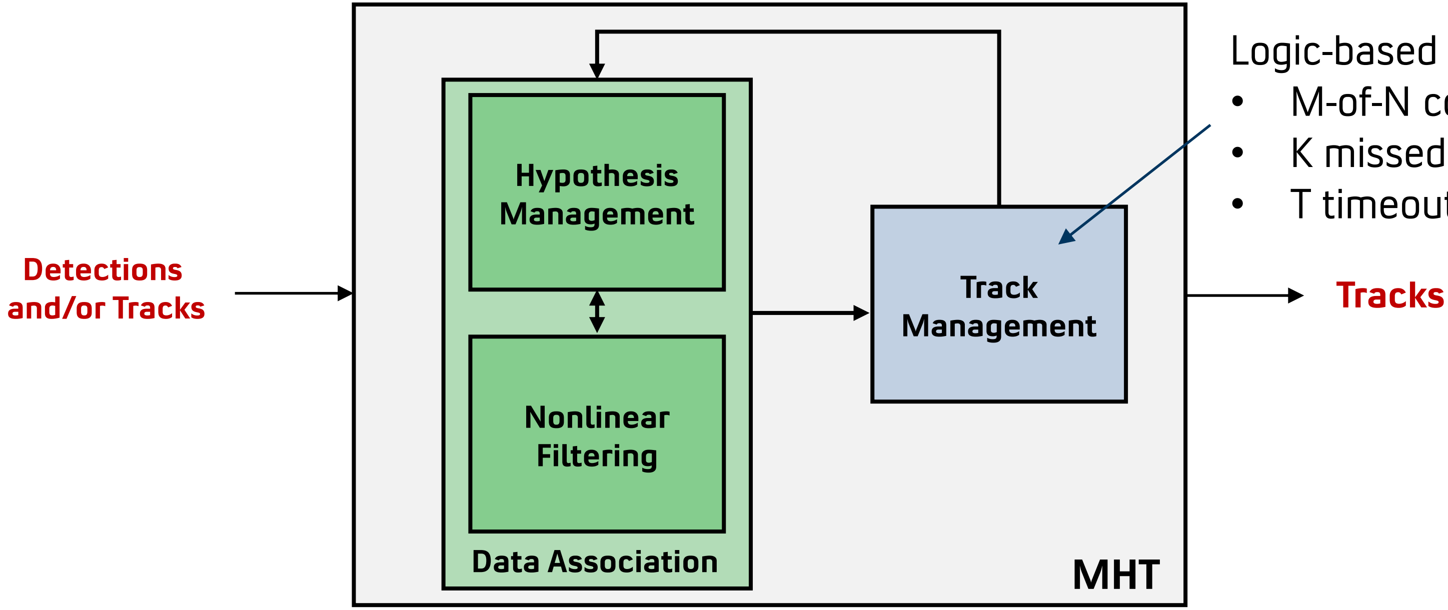


Logic-based track management
- M-of-N confirmation
- K missed detections
- T timeout

# In ideal settings things work very well

- Benign conditions
  - Data is well-matched to modeling assumptions
  - No complex detection or measurement effects (e.g. fading, bias errors, nonlinearity, limited observability, merged/repeated measurements)
  - No acquisition, processing, or communication delays
  - Manageable data volume with like sensors
- Performance trends match expectations
  - Benefit of increasing hypothesis depth
  - Benefit of increasing sensor data rate

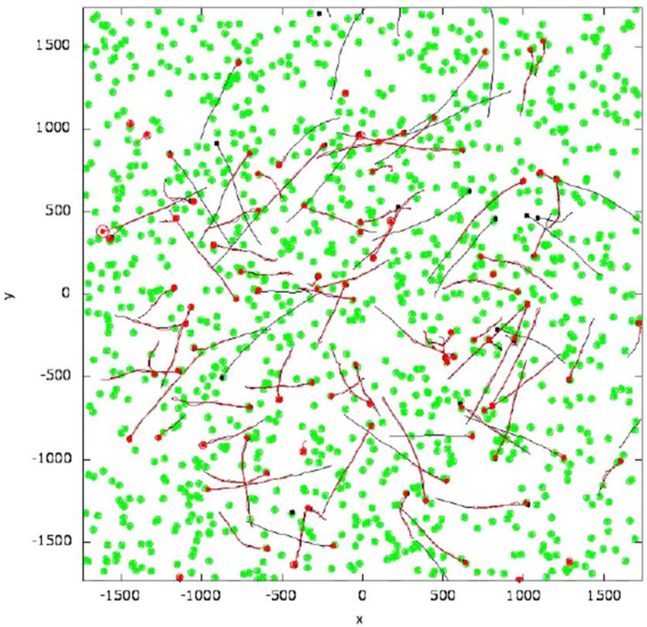

# Achieving performance requires more

**GMTI Tracking**

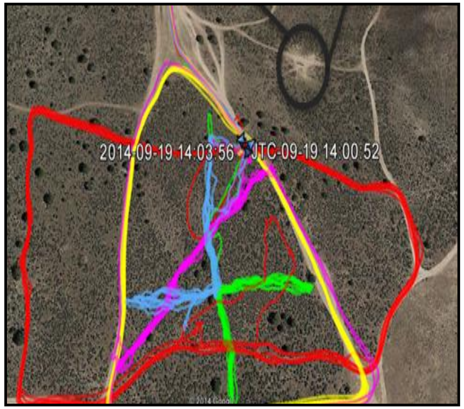


**OPIR Tracking**

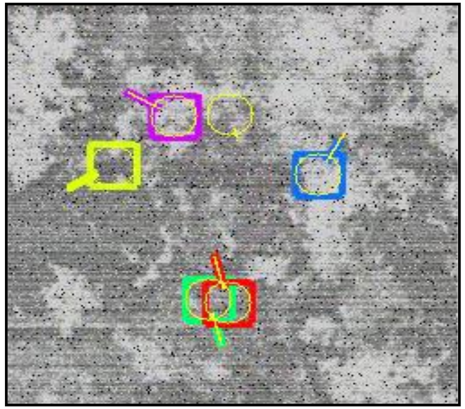

**Video Tracking**

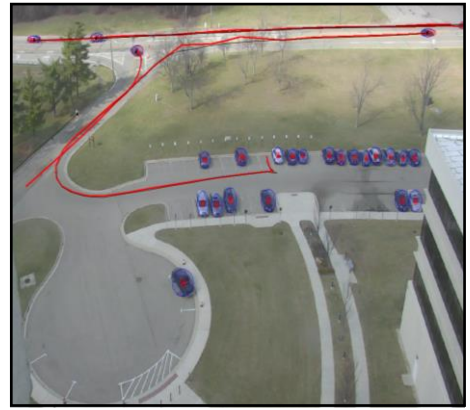

**Single Sensor**

**Air Surveillance**

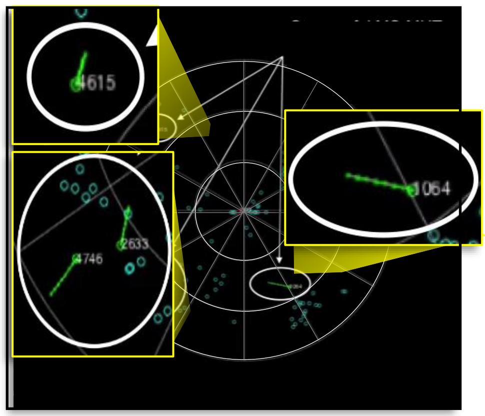


**Passive Sonar Tracking**

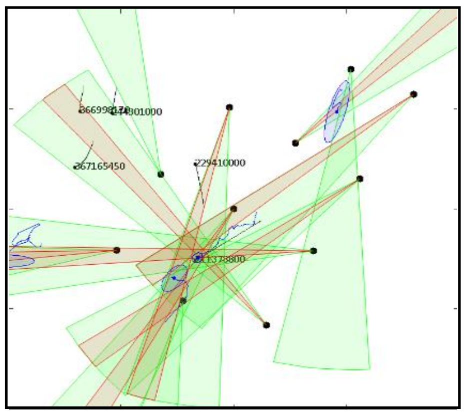

**Multi-INT Fusion**

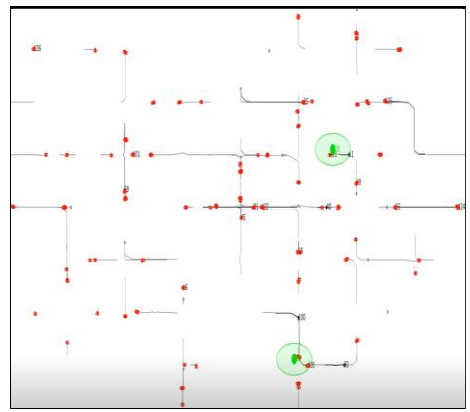

**Multiple Sensor**

Illustrations of STR *Multi-Stage MHT* (MS-MHT)

# Outline

- Multi-target tracking (MTT)
- Multiple-hypothesis tracking (MHT)
- Advanced centralized MHT
- Distributed MHT
- Simplified distributed MHT

# Nonlinear filtering

- Geometric (vs. ML) track initialization
  - AOA sensors: $J = \sum_{i=1,2} \left\|\tilde{X}_i\right\|^2 = \sum_{i=1,2} \|X - X_i\|^2$
  - TDOA sensors

*Coraluppi et al., FUSION 2021*

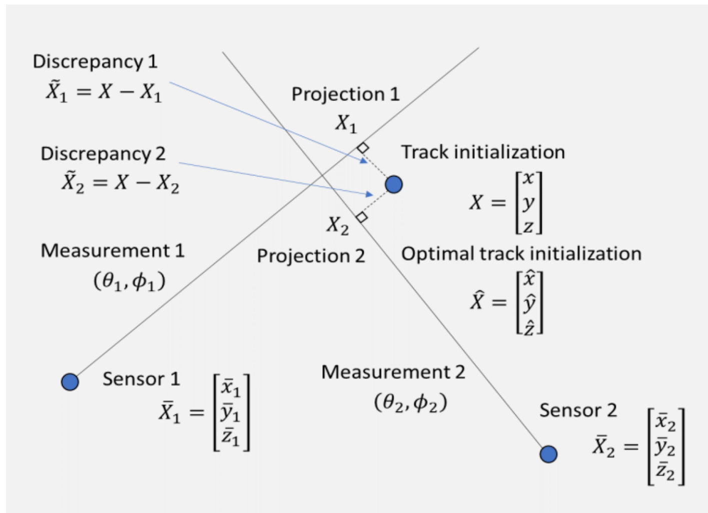


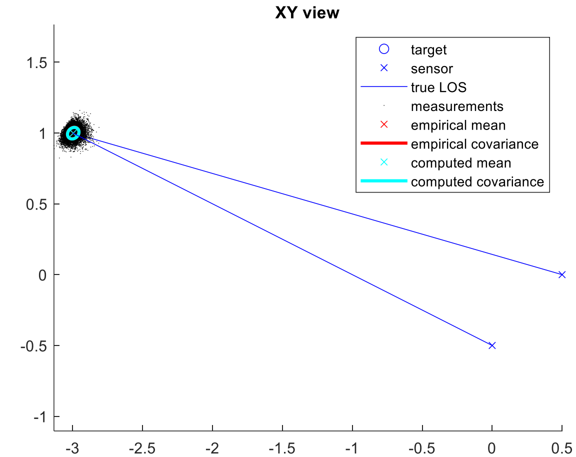


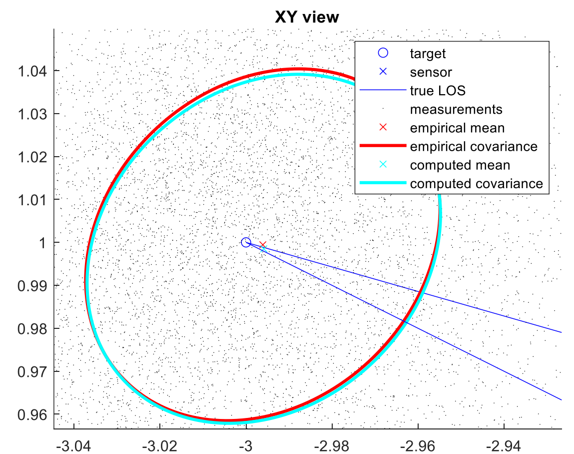


- Sequential EFK, multiple-model filtering, particle filtering

# Feature-aided tracking

- Feature-aided MHT
  - Enhanced likelihood ratio
  - Composite confirmation logic
  - State augmentation

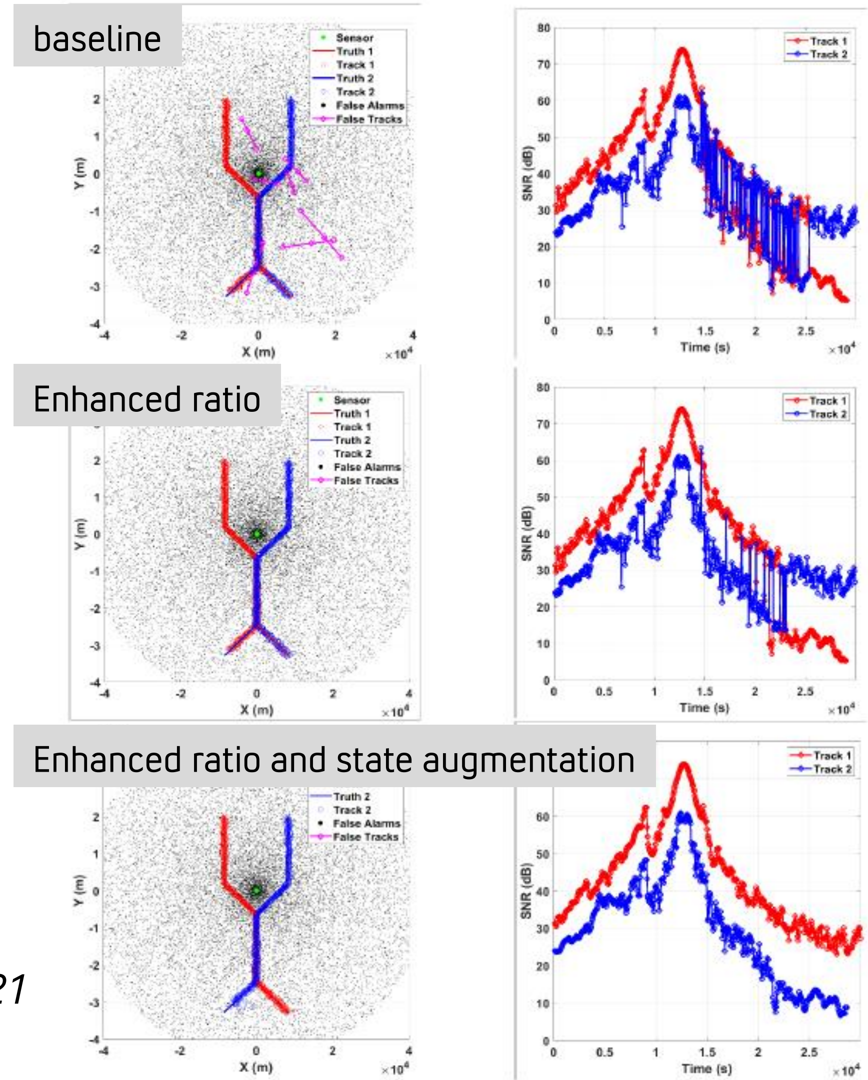


*Lexa et al., FUSION 2021*

# Hypothesis aggregation

- Alternative views of MAP estimation

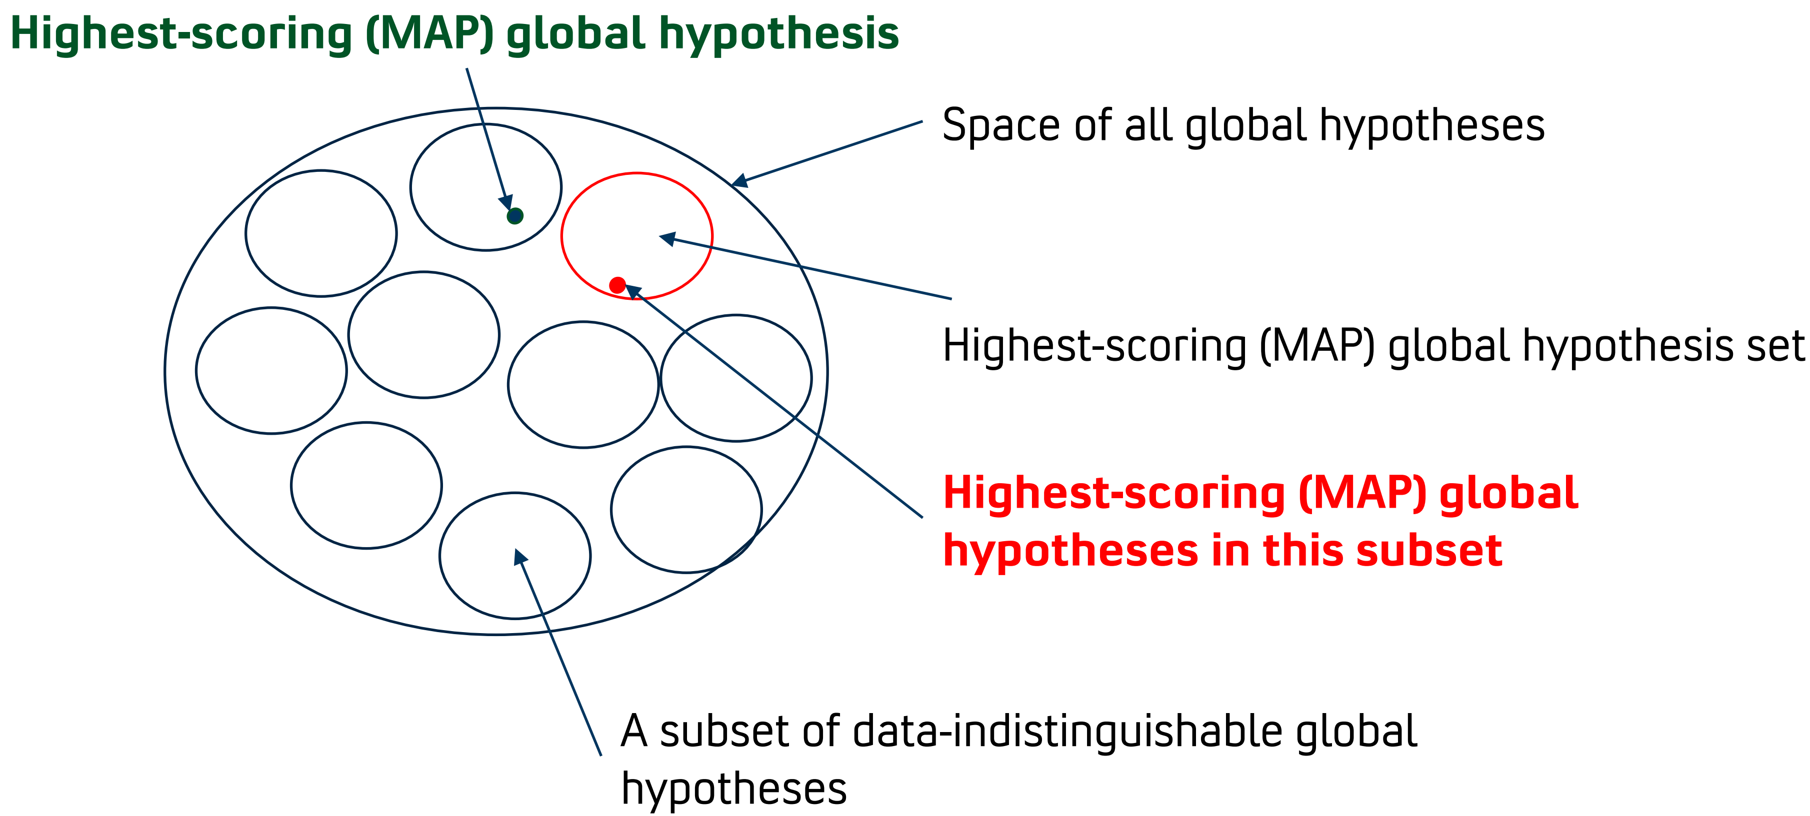

# Hypothesis aggregation

- Aggregate over data-indistinguishable global hypotheses

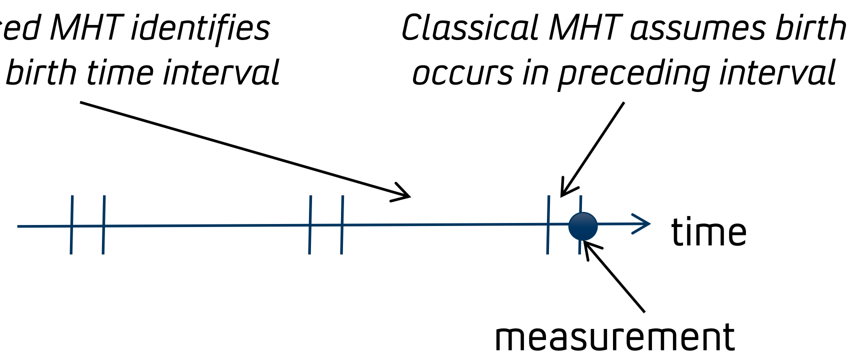


- Aggregate over similar global hypotheses
- Enhanced MHT recursion for merged and repeated measurements

*Coraluppi and Carthel, IEEE T-AES 2018*

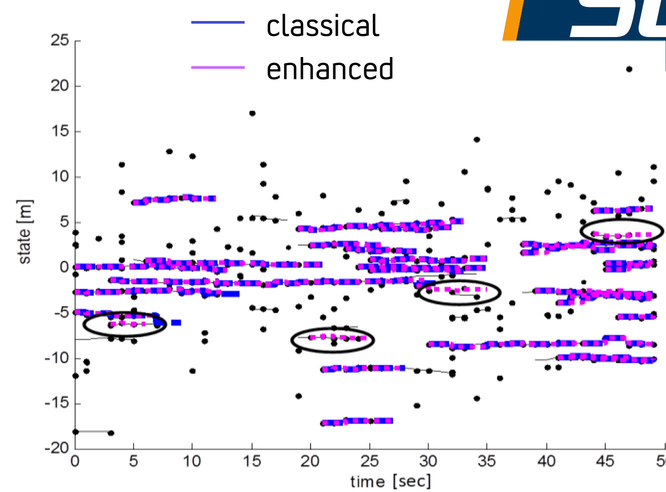


*Coraluppi and Carthel, IEEE T-AES 2014*

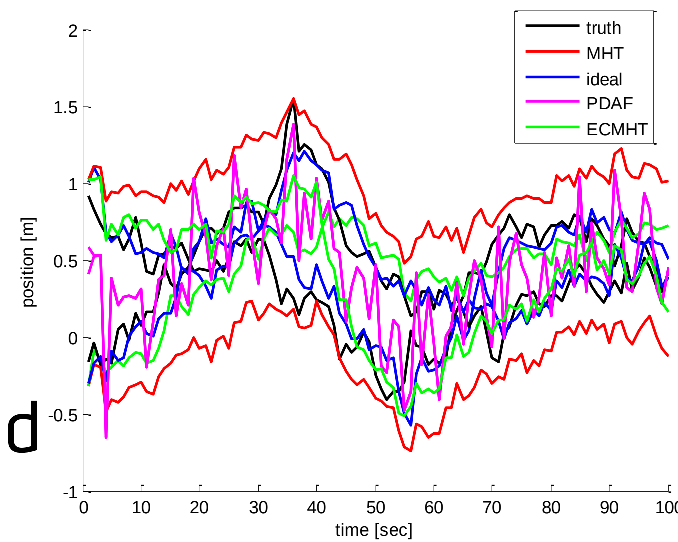


*Coraluppi and Carthel, IEEE Aerospace 2012*

# Context-aware tracking

- Types of context
  - Aggregate patterns life
  - Motion constraints
  - Partial knowledge of target objectives (avoid detection, achieve mission, etc.)

*Maritime traffic*

*Ground traffic*

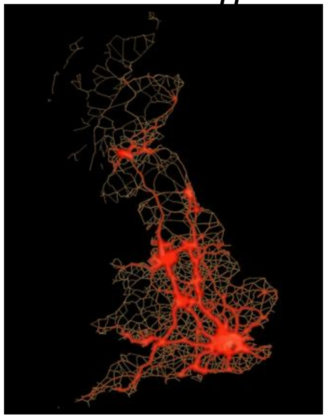

*Air traffic*

- Approaches
  - Modifying or augmenting sensor data
  - Modify target dynamics, including Gaussian-mixture filters and particle filters

# Single-model context-aware tracking

- All available context and prior information should be exploited
  - This is particularly beneficial in limited-coverage settings
- One approach to do so is to embed such information in a nominal trajectory

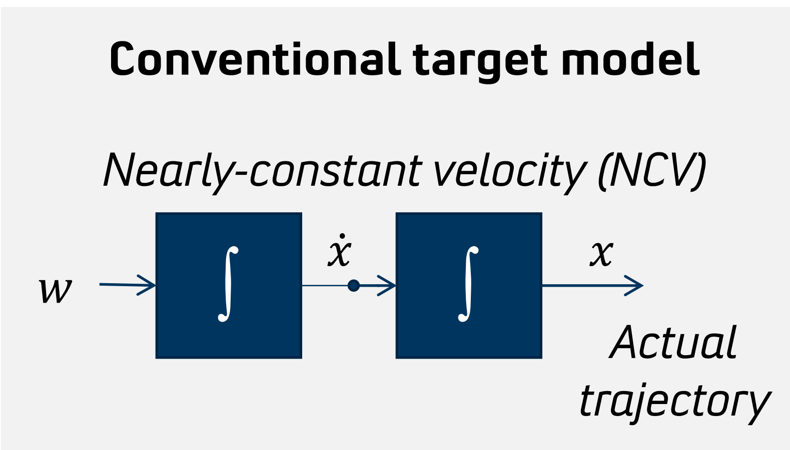


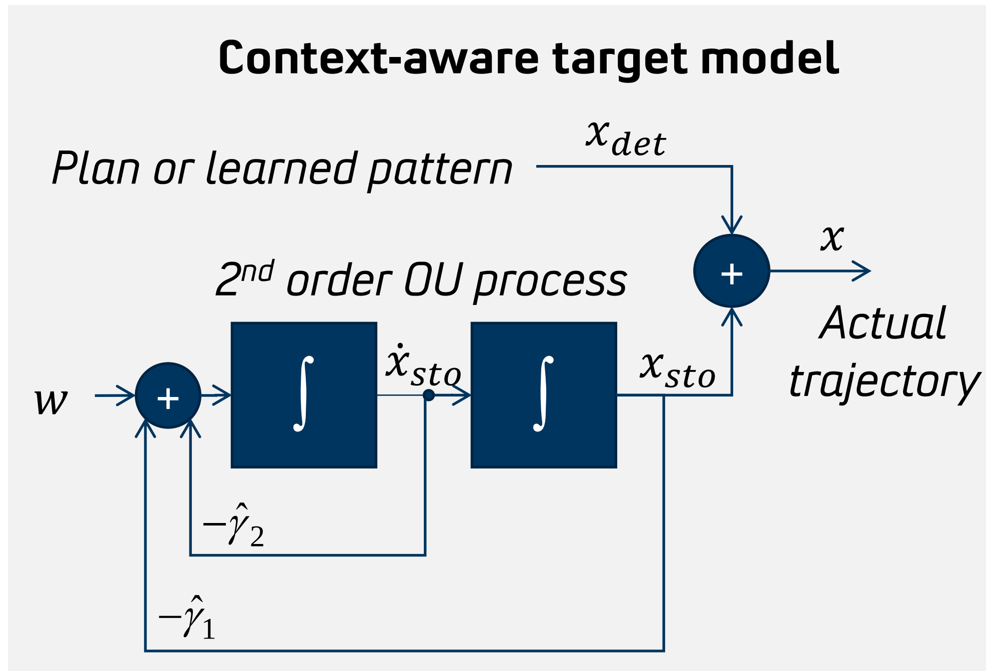

# Conventional processing

1. Good tracking in area of good coverage

2. Track extrapolation where data is absent (tracks are not terminated due to known coverage gap)

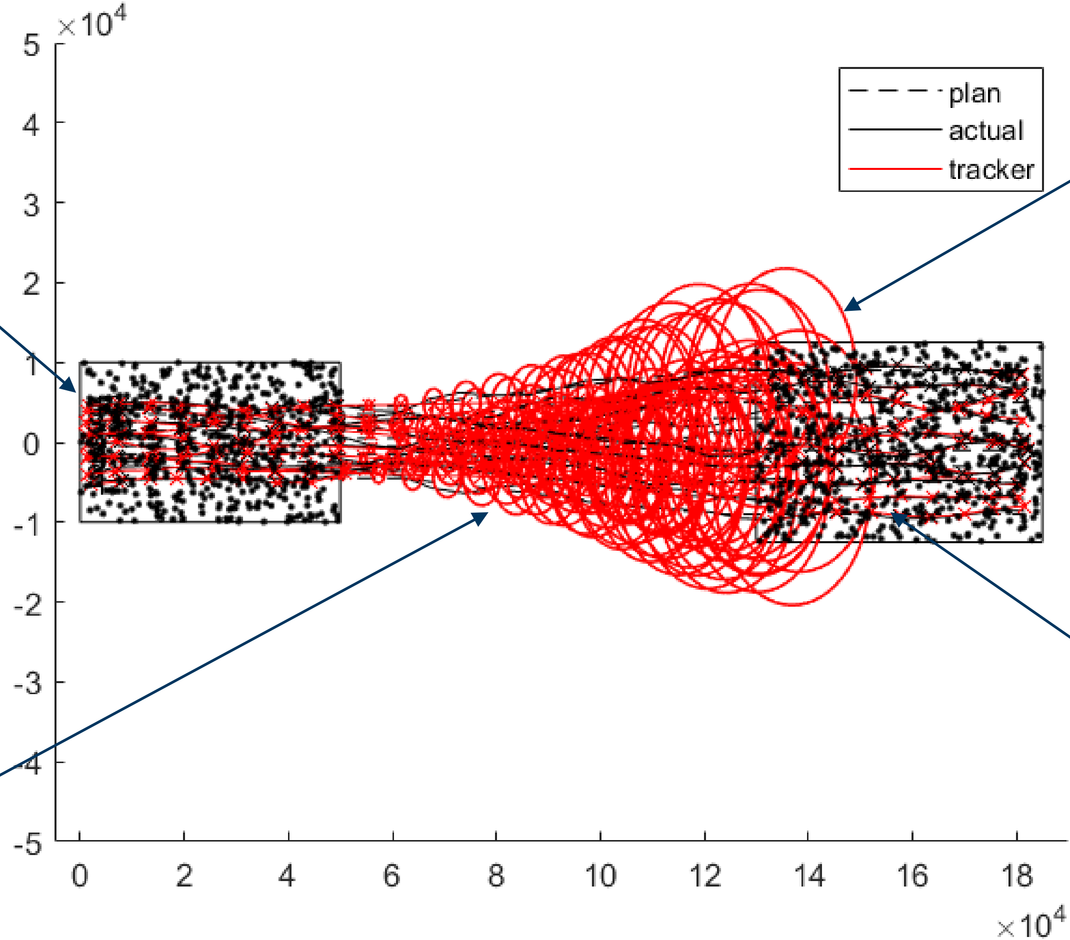


3. Error ellipses exhibit cubic growth in positional uncertainty (consistent with NCV model)

4. Target reacquisition (with mostly incorrect association, despite principled MHT likelihood scoring)

# Context-aware processing

1. Good tracking in area of good coverage

2. Track extrapolation where data is absent maintains bounded uncertainty and exploits nominal flight-plan information

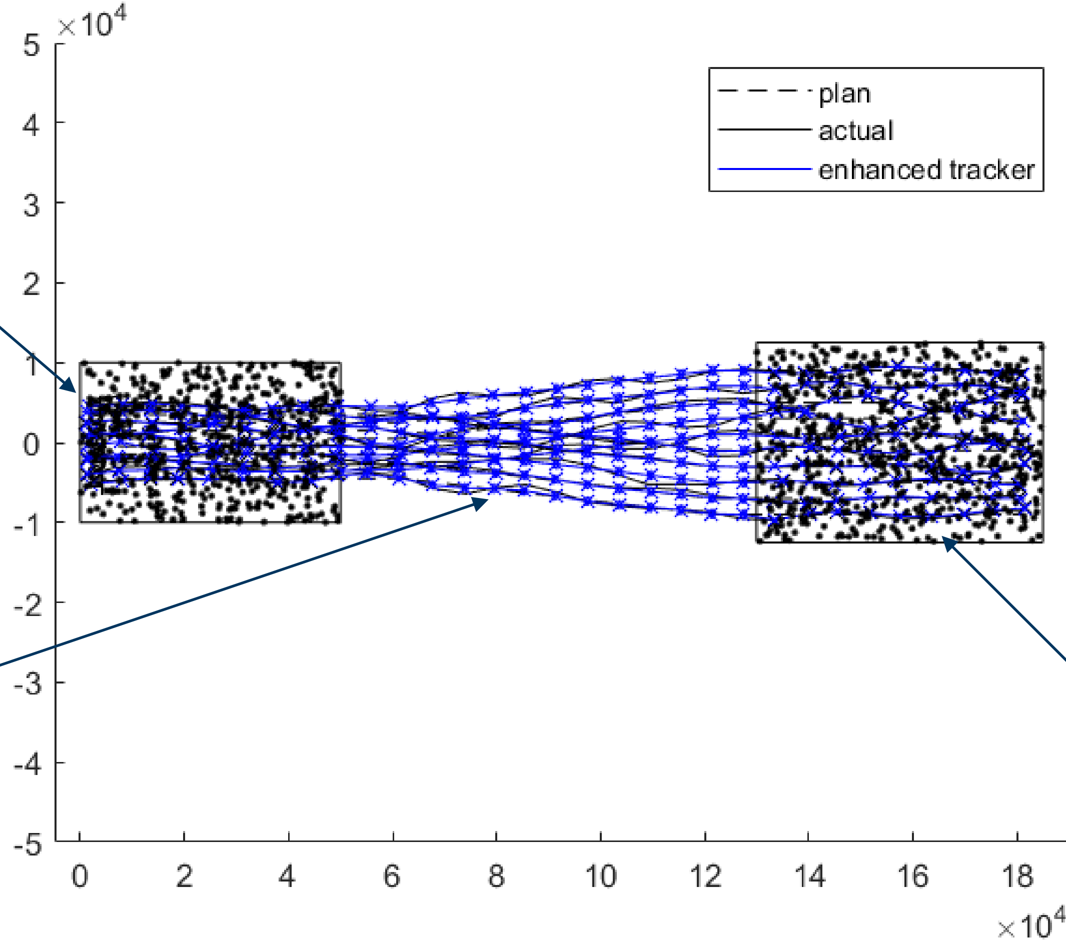


3. High-accuracy target reacquisition

# Single-model context aware tracking

- In simulations with N targets and large coverage gaps, *probability of correct association* (PCA) with data association processing degrades to the performance lower bound
  - $PCA \approx \frac{1}{N}$
- Use of generalized Ornstein-Uhlenbeck filters leads to a substantial performance gain
  - When coverage is good, tracking performance is similar to conventional processing
  - When coverage is poor, track filters give more weight to context, maintaining bounded and statistically-consistent uncertainty
  - Resulting target reacquisition is greatly facilitated

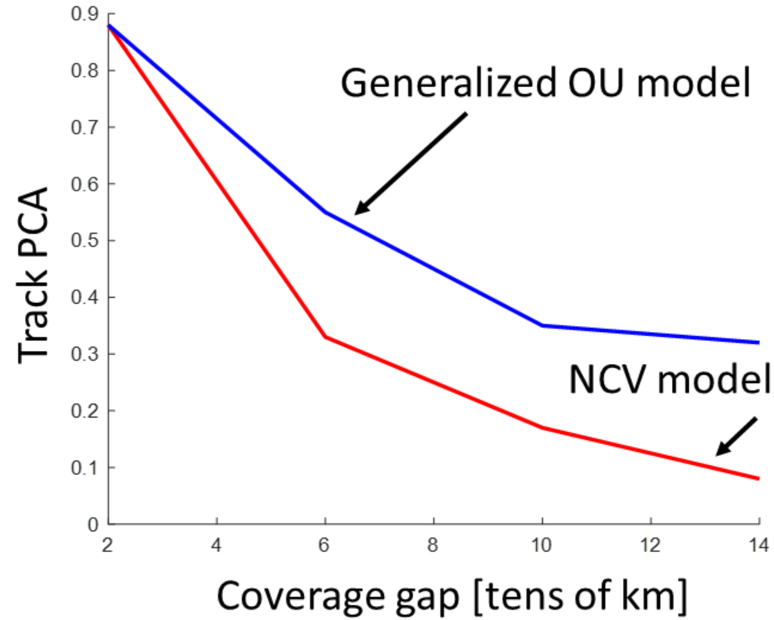

# Multiple-model context-aware tracking

- Uncertainty is decoupled across multiple filter modes

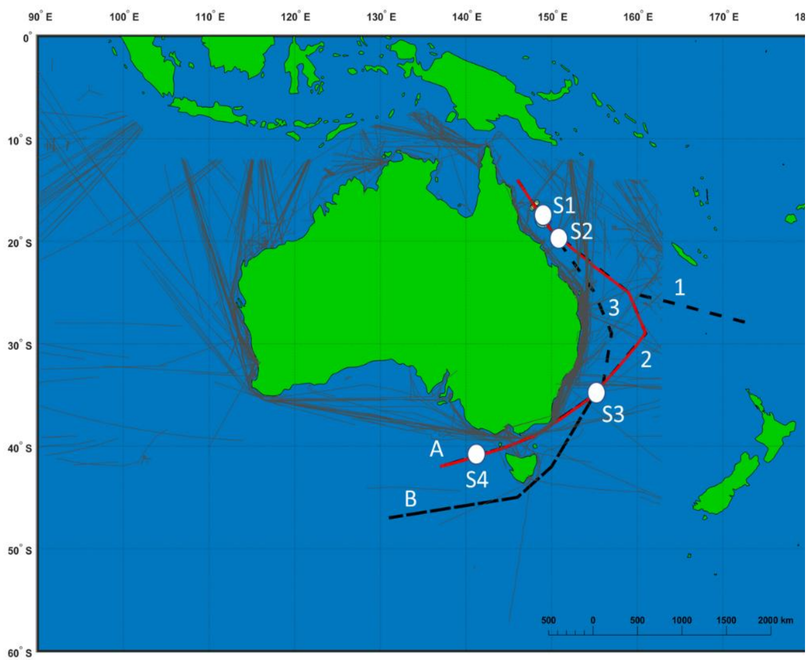


*Coraluppi et al., AEROCONF 2021*

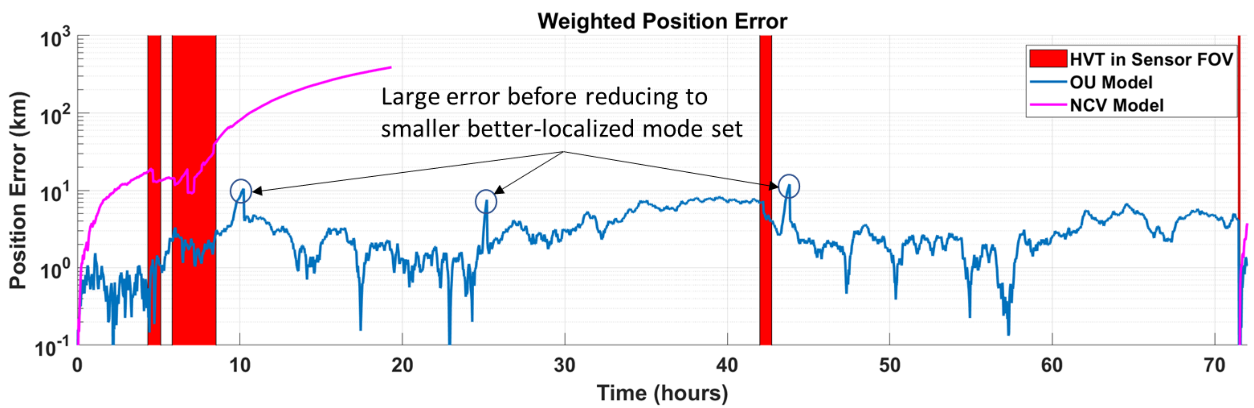


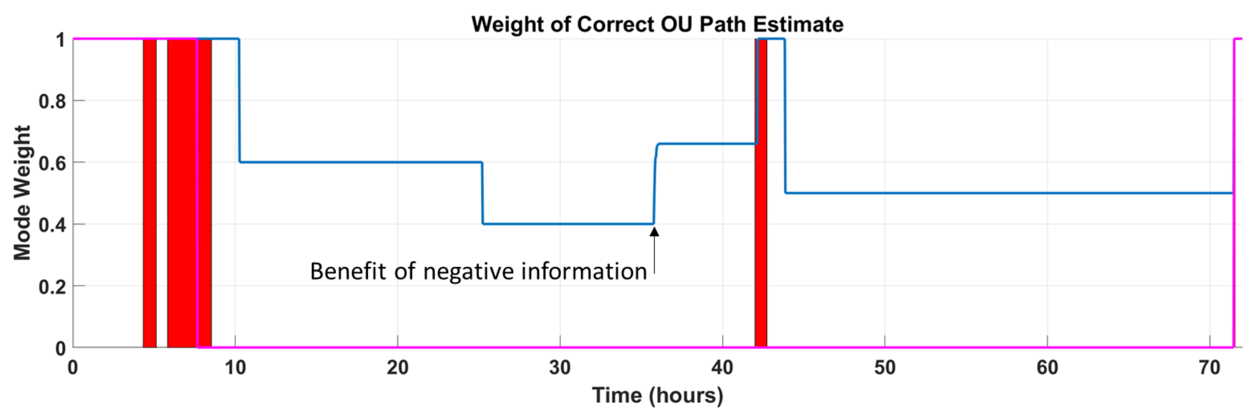

# Outline

- Multi-target tracking (MTT)
- Multiple-hypothesis tracking (MHT)
- Advanced centralized MHT
- Distributed MHT
- Simplified distributed MHT

# Distributed tracking

- Centralized tracking is optimal given no processing constraints
- Empirical advantages of distributed tracking
  - Robust to detection fading and measurement biases
  - Effective for limited-observability sensors or highly-disparate sensors (*track before fuse*)
  - Robust to high frame rate (*track before track*) or large networks of degraded sensors (*fuse before track*)
  - Enables asynchronous data association

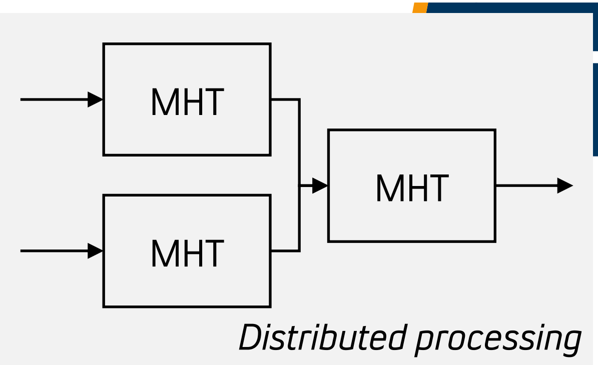


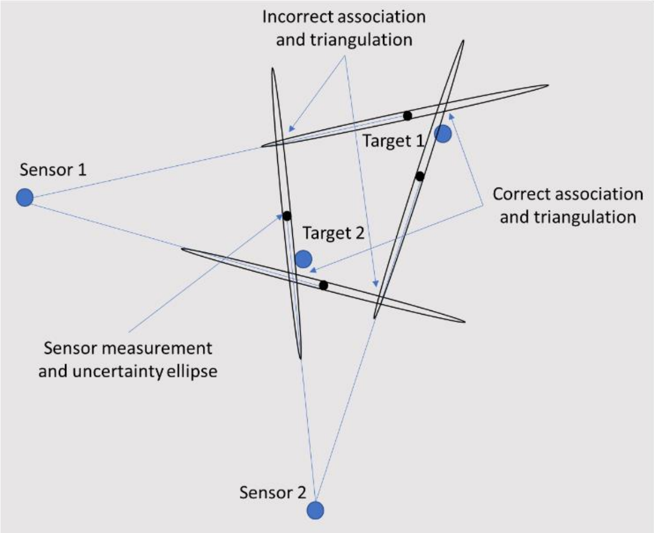


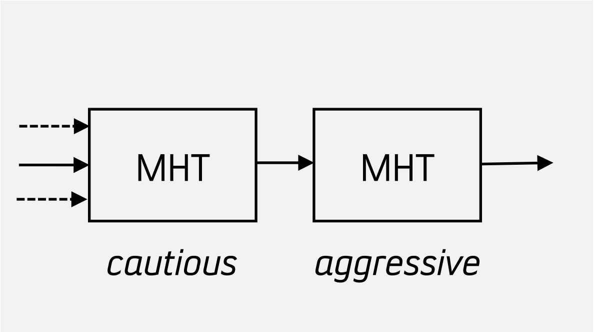

# Avoiding distributed estimation

- Objects of interest are sequences of associated measurements
  - Enables modularity, avoids track-correlation issues, allows for (static) measurement fusion, allows for stage-specific target & sensor models
- All modules perform object association
  - Breakage logic provides robustness
- Flexible connections between modules are possible
  - Track breakage, track suppression, multi-look processing, etc.

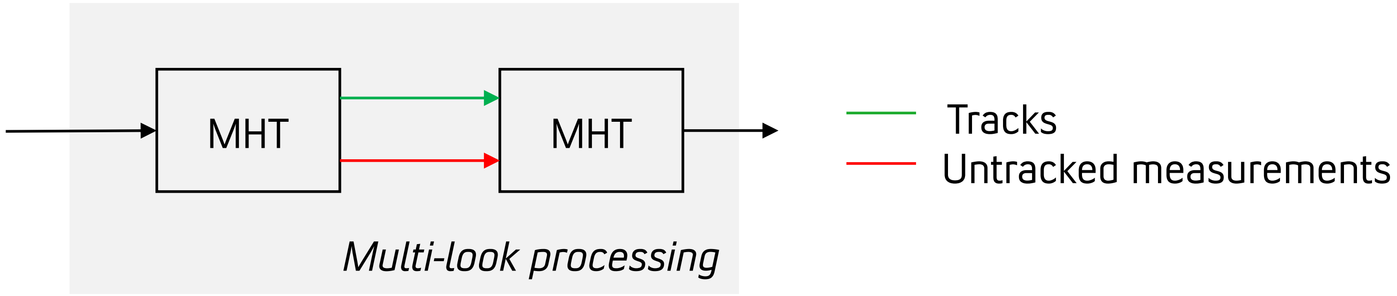

# A single-sensor GMTI tracking example

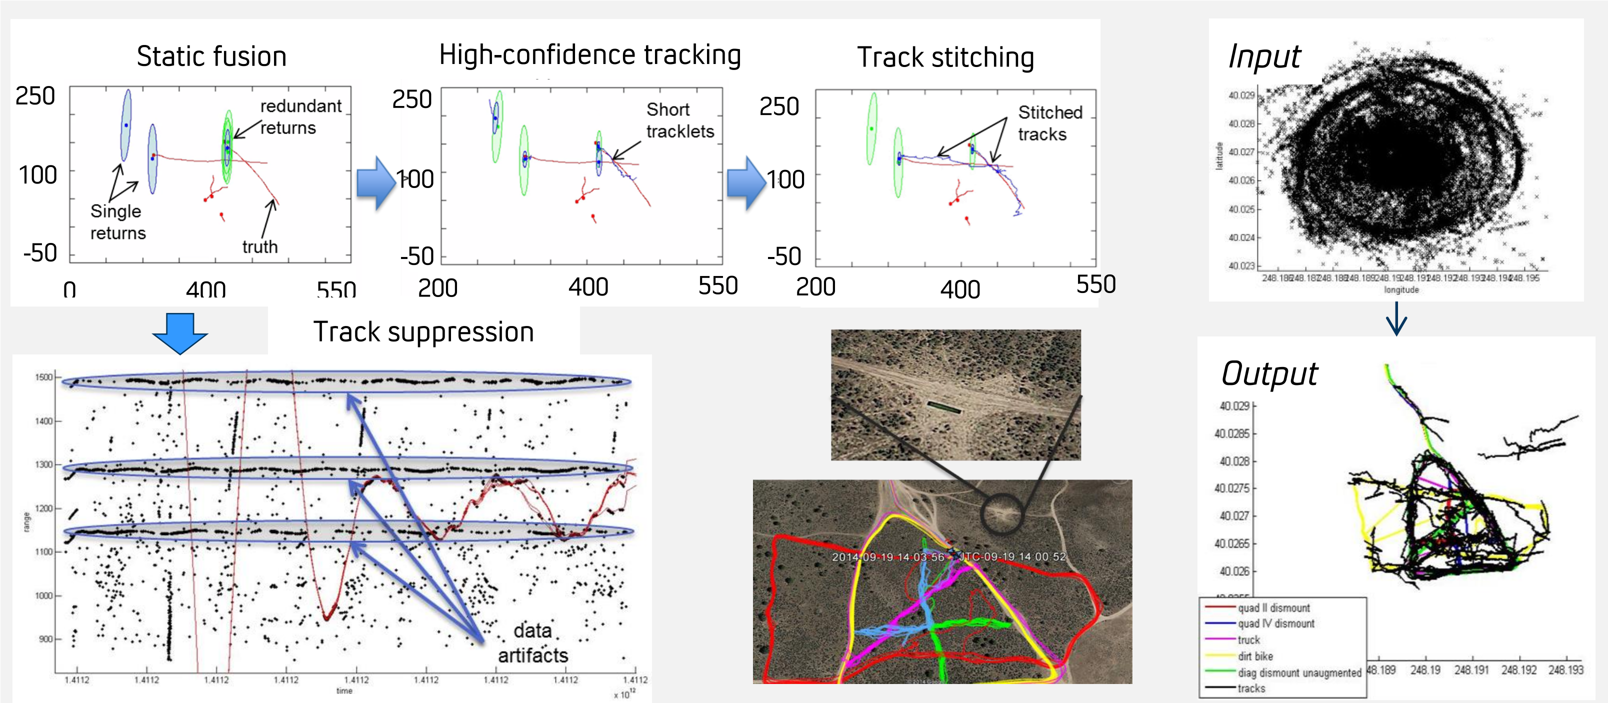


Coraluppi et al., NSSDF 2016

- Excellent benchmark results compared to other solutions

| Metric | MHT | Non-MHT | STR MS-MHT |
|---|---|---|---|
| Combined Target Completeness | 0.51 | 0.28 | 0.63 |
| Number of Tracks | 116 | 117 | 87 |
| Average Track Life (frames) | 172.97 | 116.60 | 273.85 |
| Combined Track Completeness | 0.85 | 0.68 | 0.89 |
| Average Spuriousness per Frame | 0.40 | 0.59 | 0.34 |

# Hierarchical tracking

- Relevant to passive sonar & radar
  - Address observability and repeated-measurement challenges
  - Decouple challenges of clutter suppression, localization, and target-level tracking
  - Allow for architectural tradeoffs

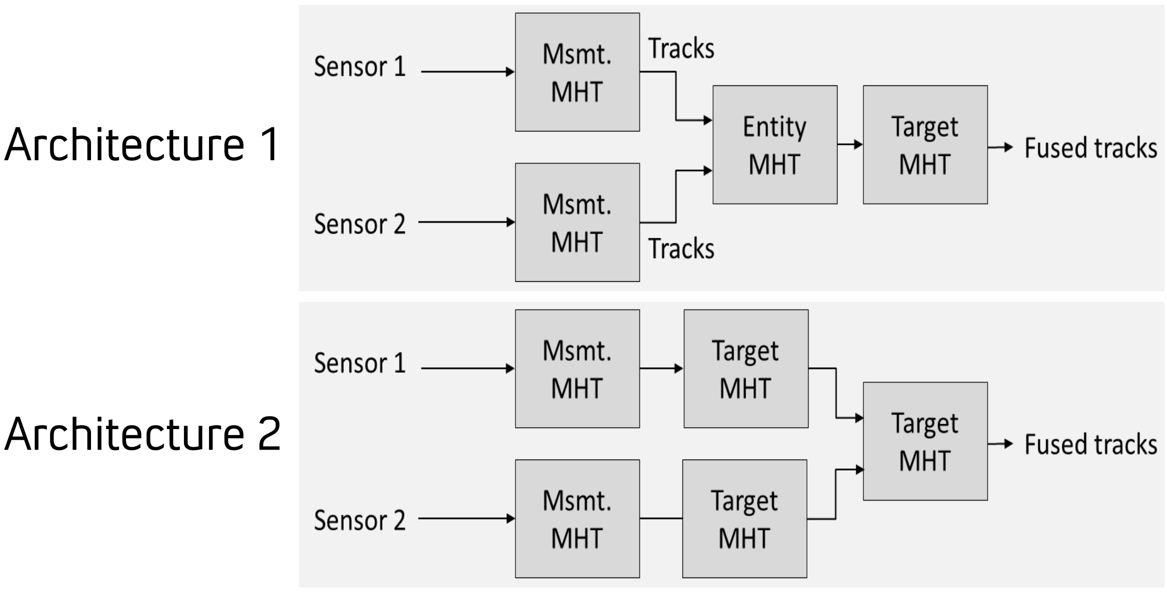


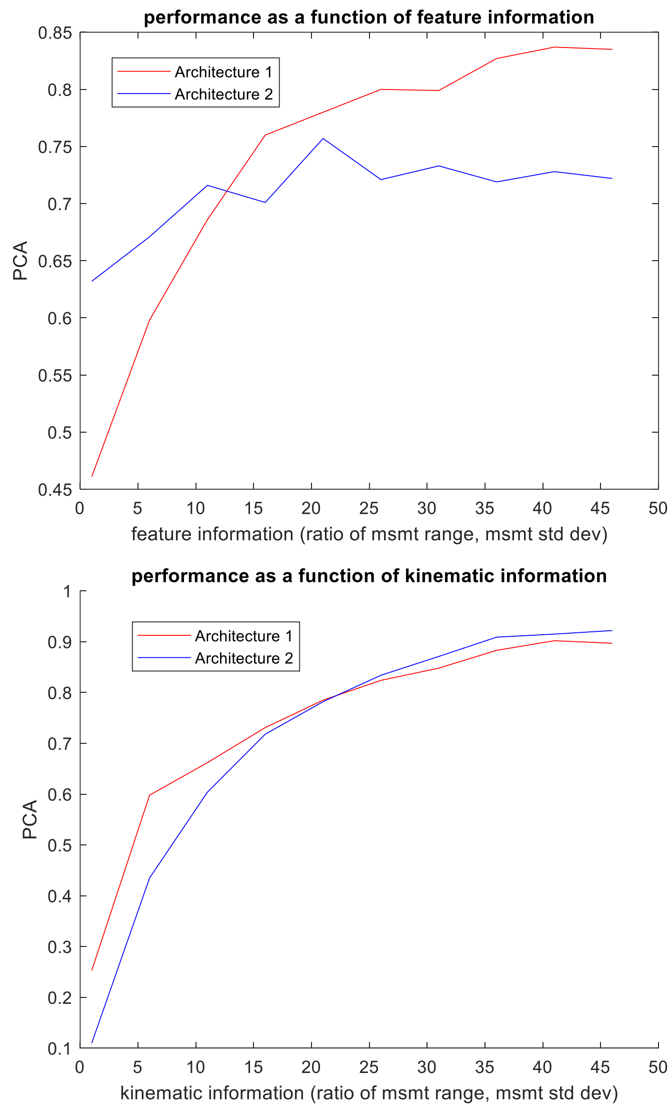


*Coraluppi et al., FUSION 2021*

# Multi-INT fusion

- Distributed MHT is insufficient
  - Asynchronous MHT provides some improvement (Coraluppi et al, IEEE Aerospace 2016)
- Hypothesis-oriented MHT can be simplified in many ways
  - Hybrid scheme exploits elements of MHT and classical *graph-based tracking* (GBT)
  - The essential simplifying assumption in GBT:

$$f(z_k|z^{k-1}) \approx f(z_k|z_{k-1})$$

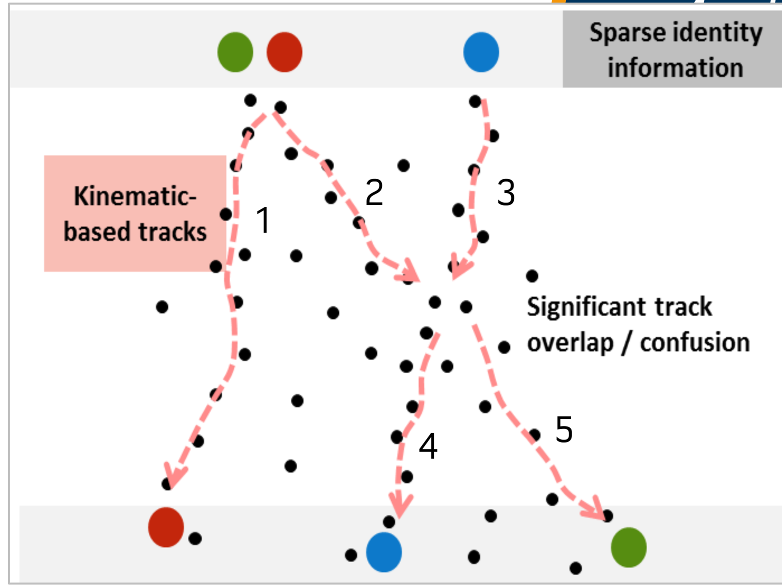


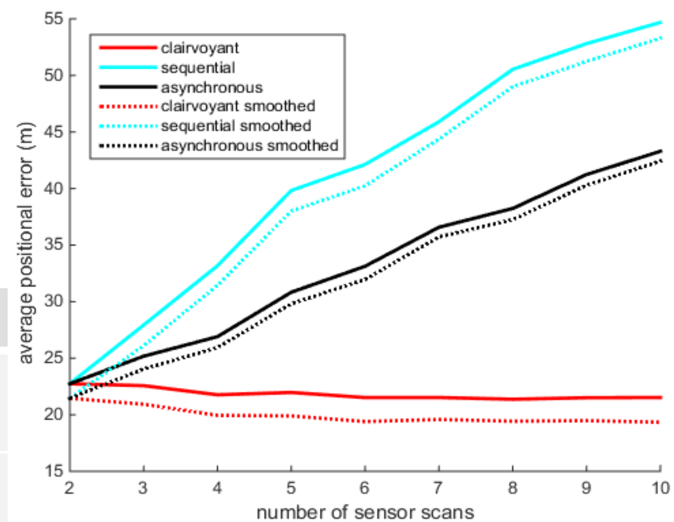


| Assumptions | General | Partially Markov data | Markov data |
|---|---|---|---|
| General | HO-MHT (Reid 1979) | *not investigated* | Viterbi (Wolf et al. 1989) |
| Poisson targets and clutter | TO-MHT (Kurien 1990) | Graph-based fusion (GBF) (Coraluppi et al. FUSION 2016) | GBT (D. Castañón 1990) (G. Castañón et al. 2011) |

# GBF Illustration

- Scenario: color indicates identity data

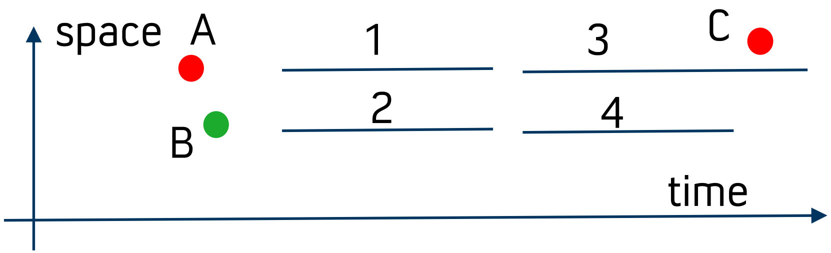


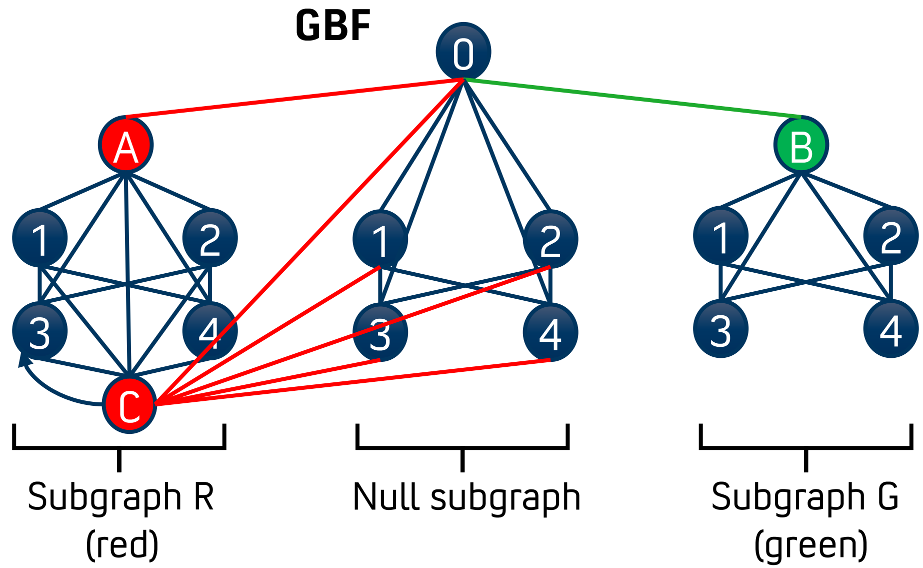


- GBF representation is compact and nearly lossless

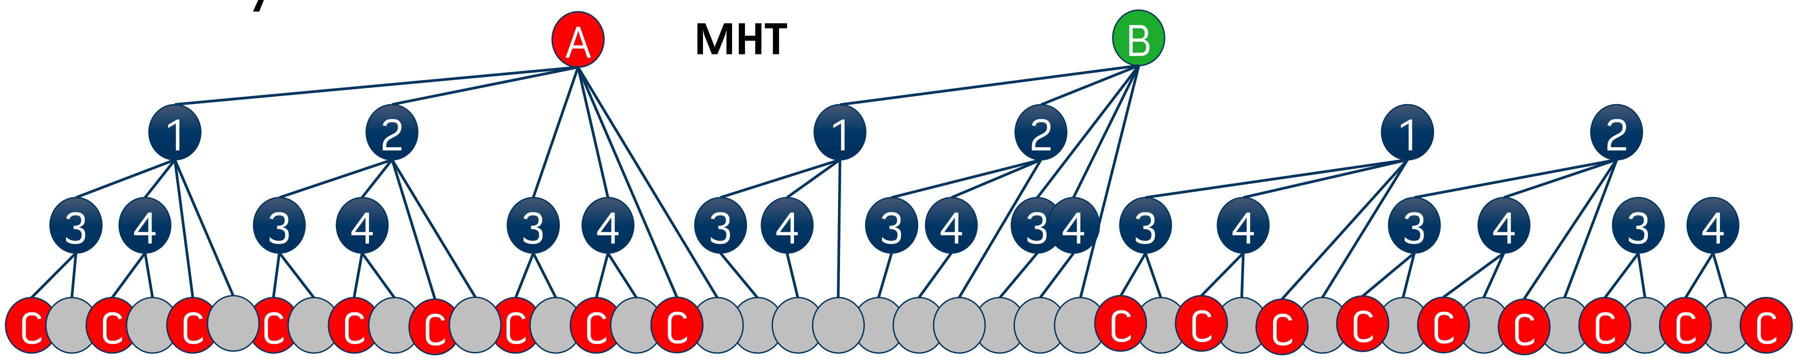

# GBF vs. MHT performance (notional)

- MHT has slightly better performance that GBF for a given n-scan
- GBF has lower computational effort than MHT for a given n-scan

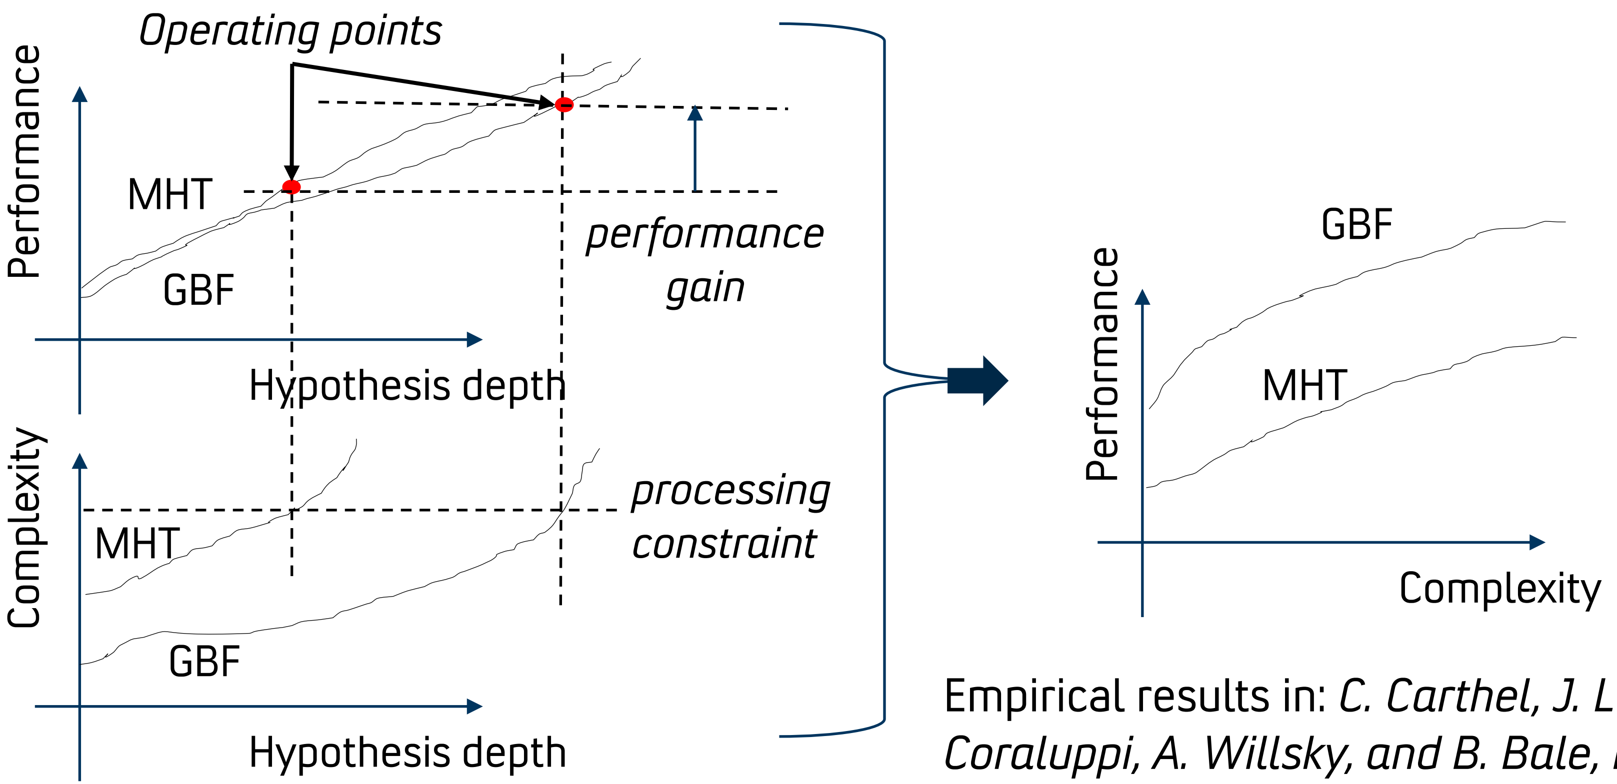


Empirical results in: *C. Carthel, J. LeNoach, S. Coraluppi, A. Willsky, and B. Bale, FUSION 2020*

# Outline

- Multi-target tracking (MTT)
- Multiple-hypothesis tracking (MHT)
- Advanced centralized MHT
- Distributed MHT
- Simplified distributed MHT

# Issue #1: upstream errors

- MHT I generates IA, IB
- MHT II generates IIA, IIB
- MHT III initially forms IIIA, IIIB, only to detect anomaly
    - Simple corrective action incurs fragmentation

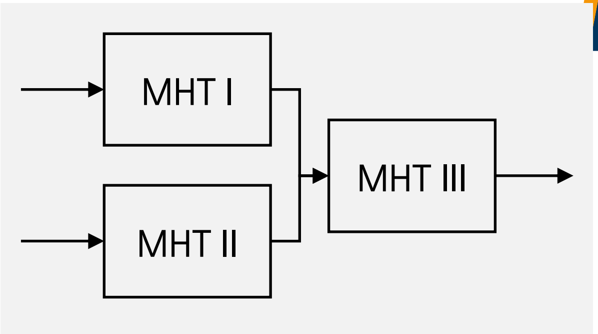


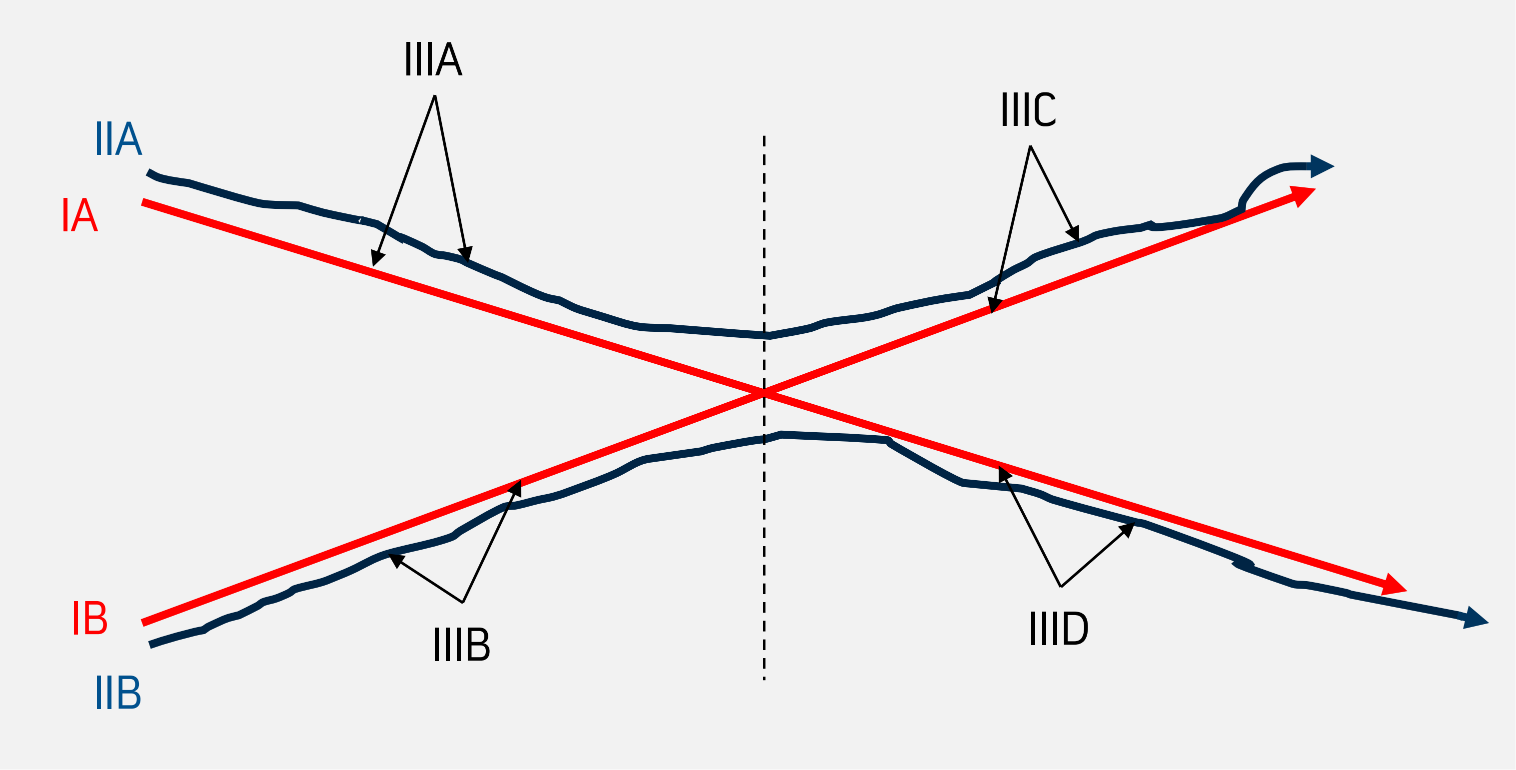

# Issue #2: cautious first stage

- Difficult to enforce track breaks at relevant times

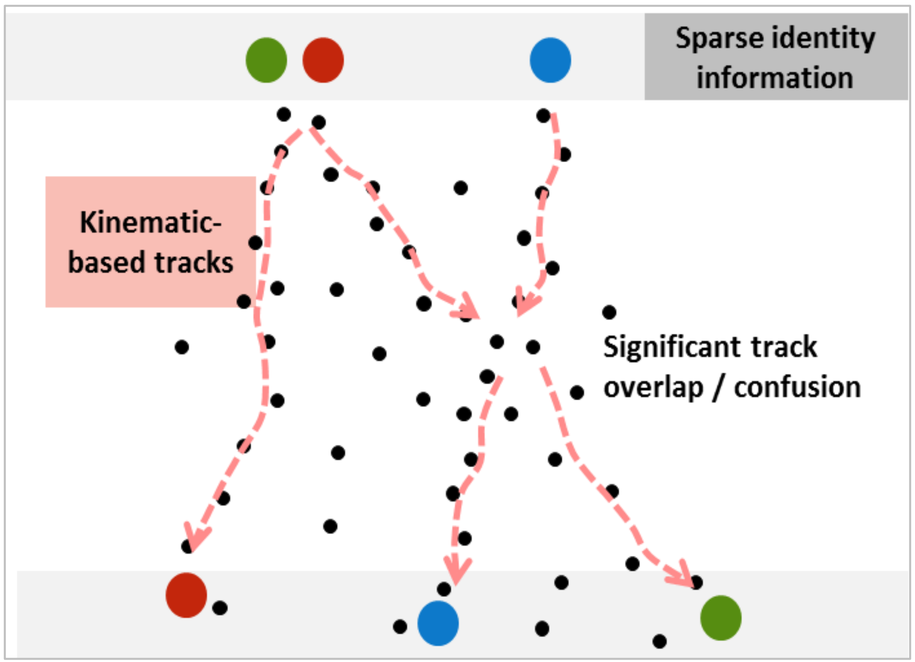


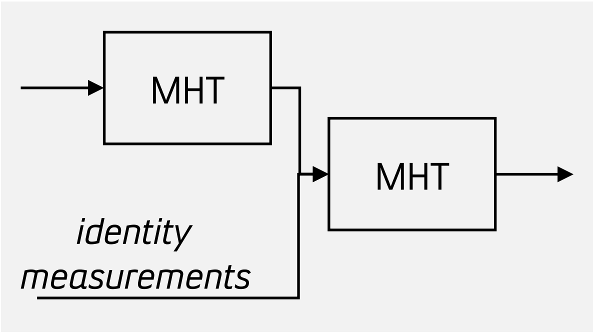

# Issue #3: variable quality sensors

- Need for principled track fusion logic

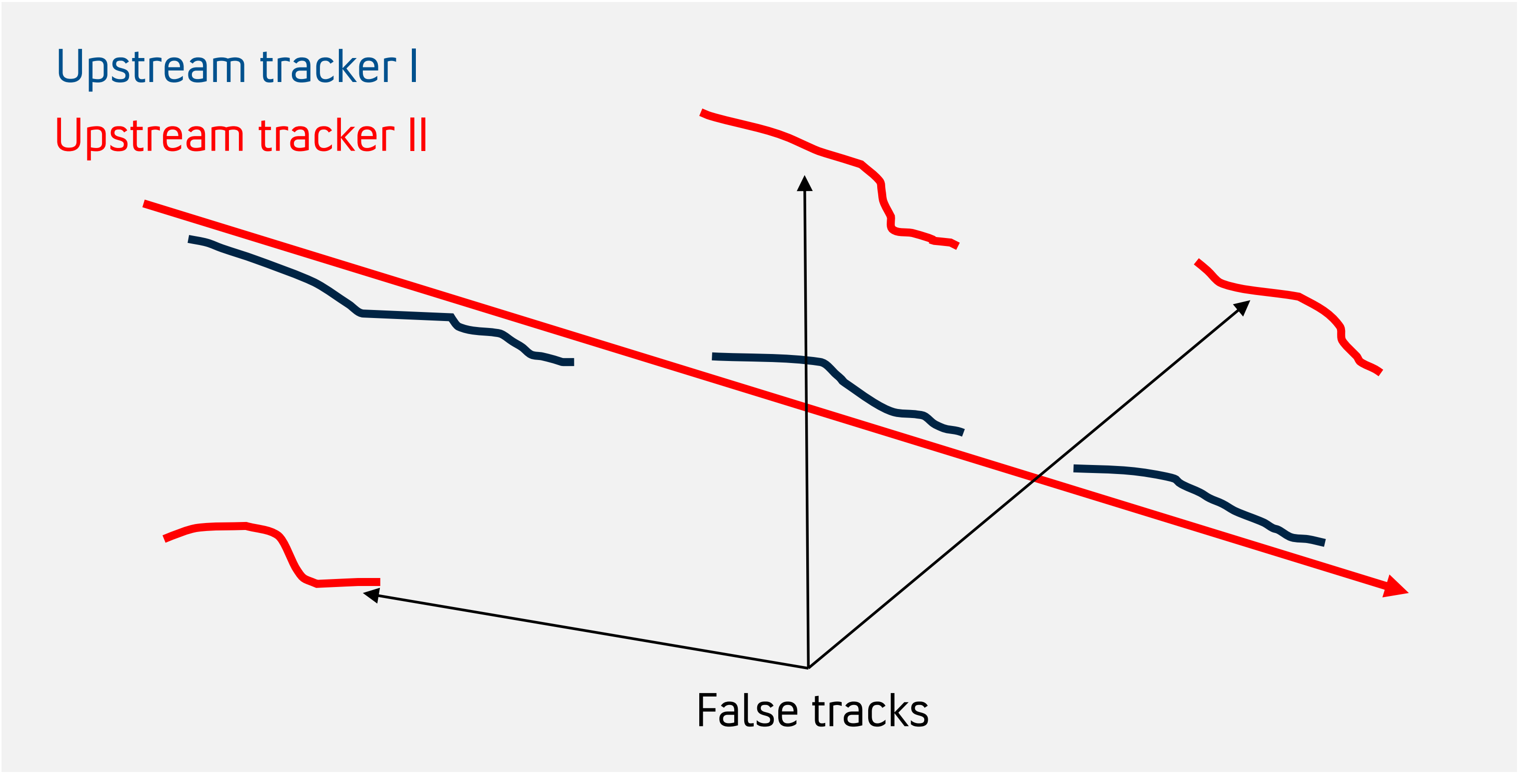

# Issue #4: non-fusion decisions

- How to revisit a non-fusion decision

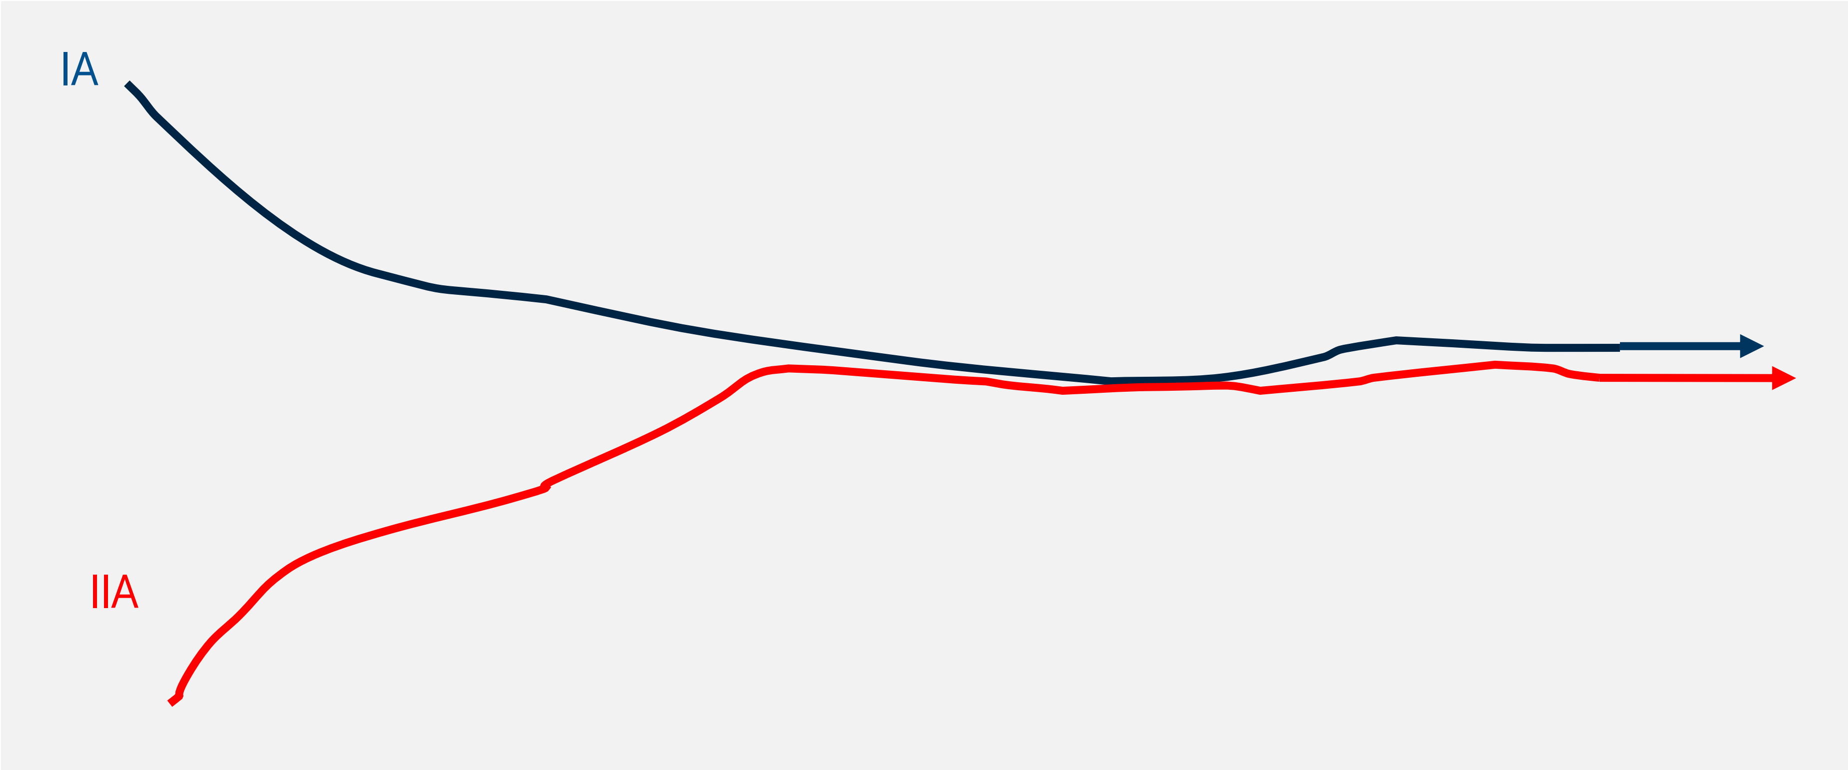

# Issue #5: upstream fragmentation

- How to relax association constraints

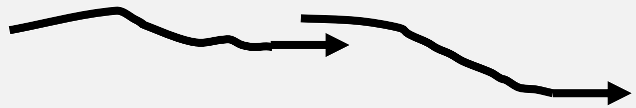

Temporally-overlapping tracks

# Issue #6: wide-area tracking

- How to contend with multiple uses of the same measurement

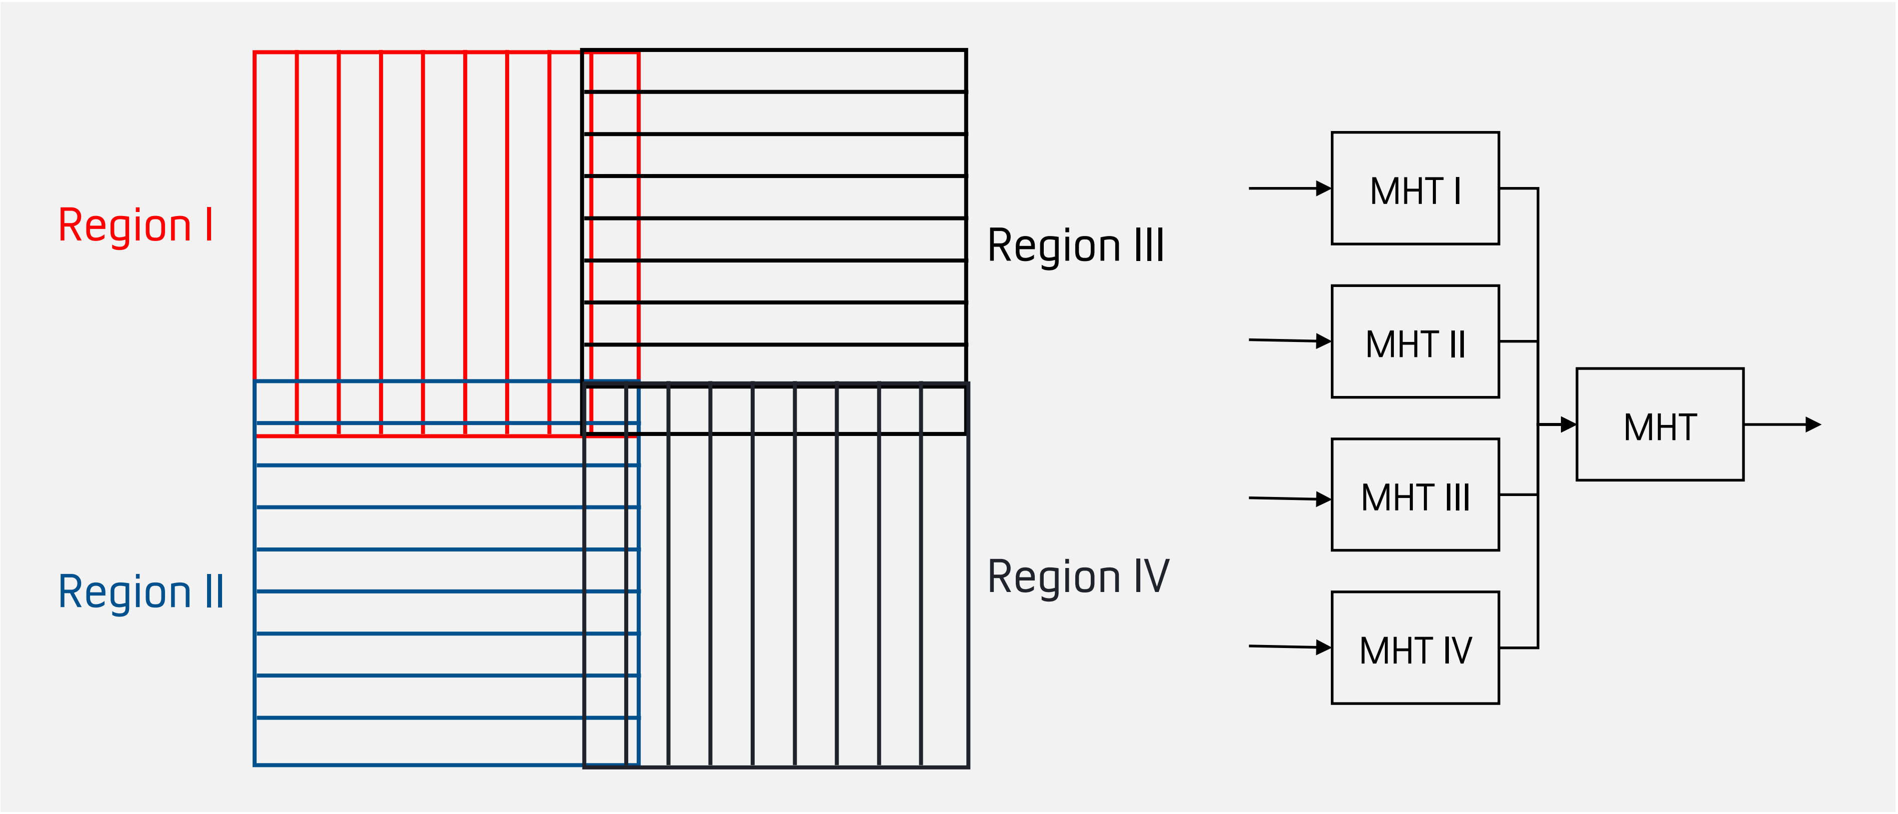

# Issue #7: distributed sensor networks

- How to contend with incompatible association decisions

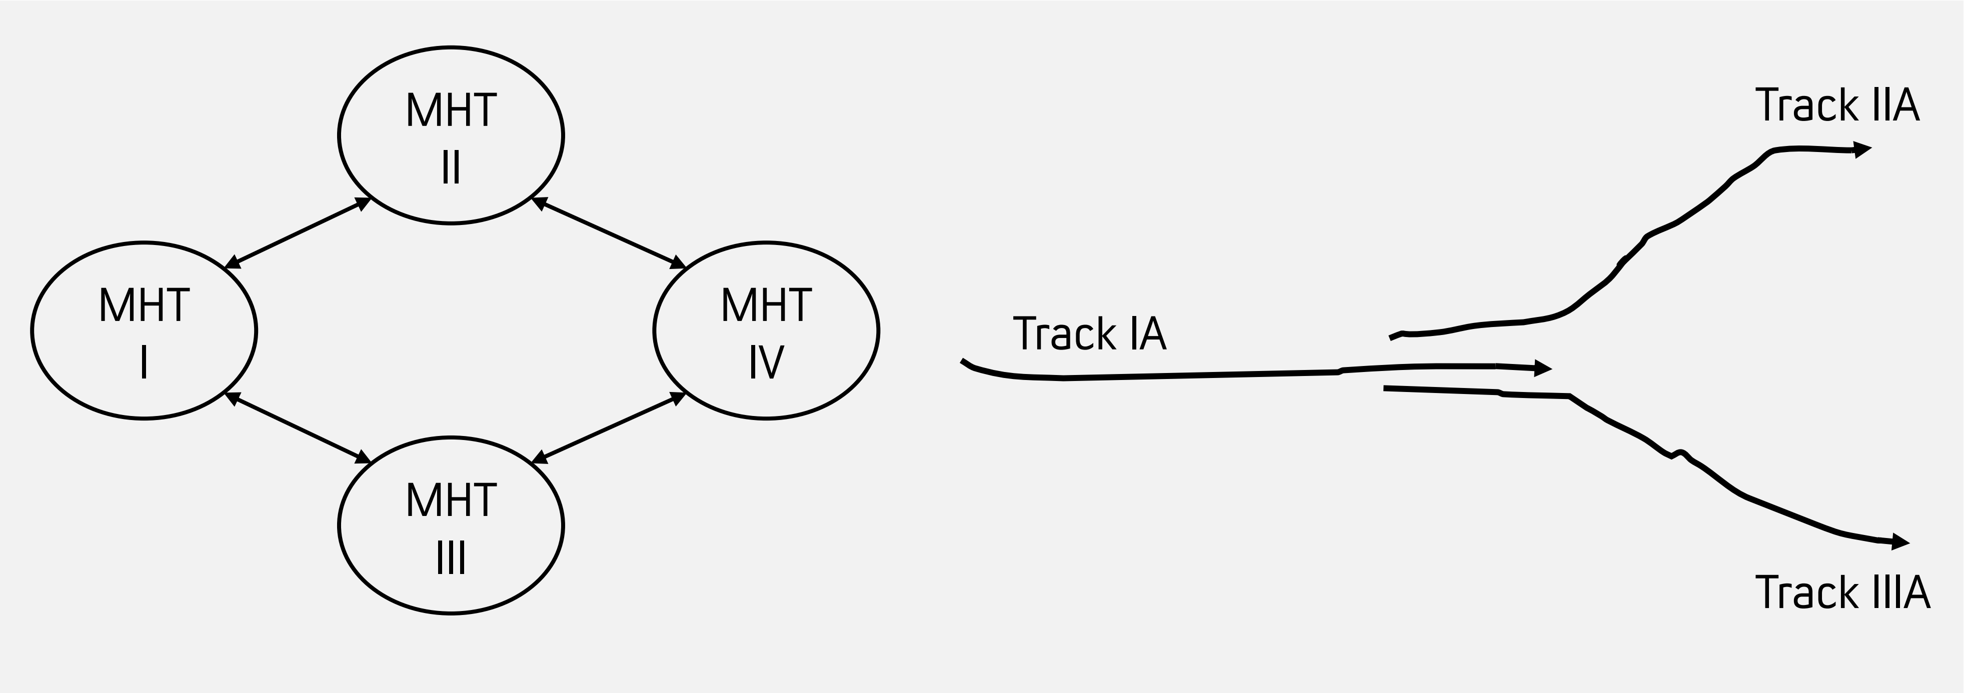

# Simplified distributed tracking

- The issues identified above can be mitigated by ignoring upstream association decisions
- A potentially better approach: introduce track label as a feature state in downstream tracking
  - One such state for each upstream tracker
  - $\tilde{X}_k = \begin{bmatrix} X_k \\ L_{I,k} \\ L_{II,k} \end{bmatrix}$

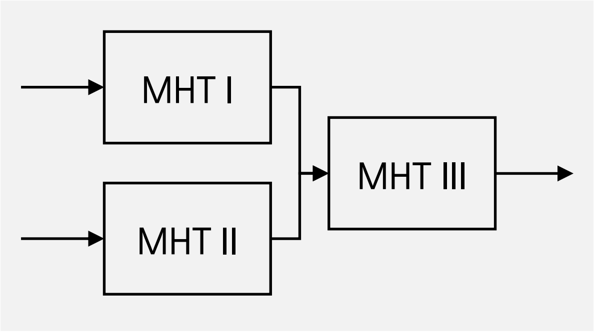


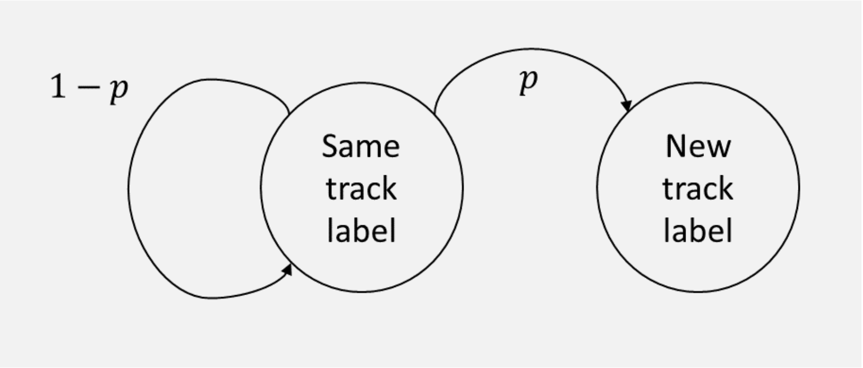

# Recovery from upstream error

- Input track breakage is account for in fused-track score
- No fused-track fragmentation

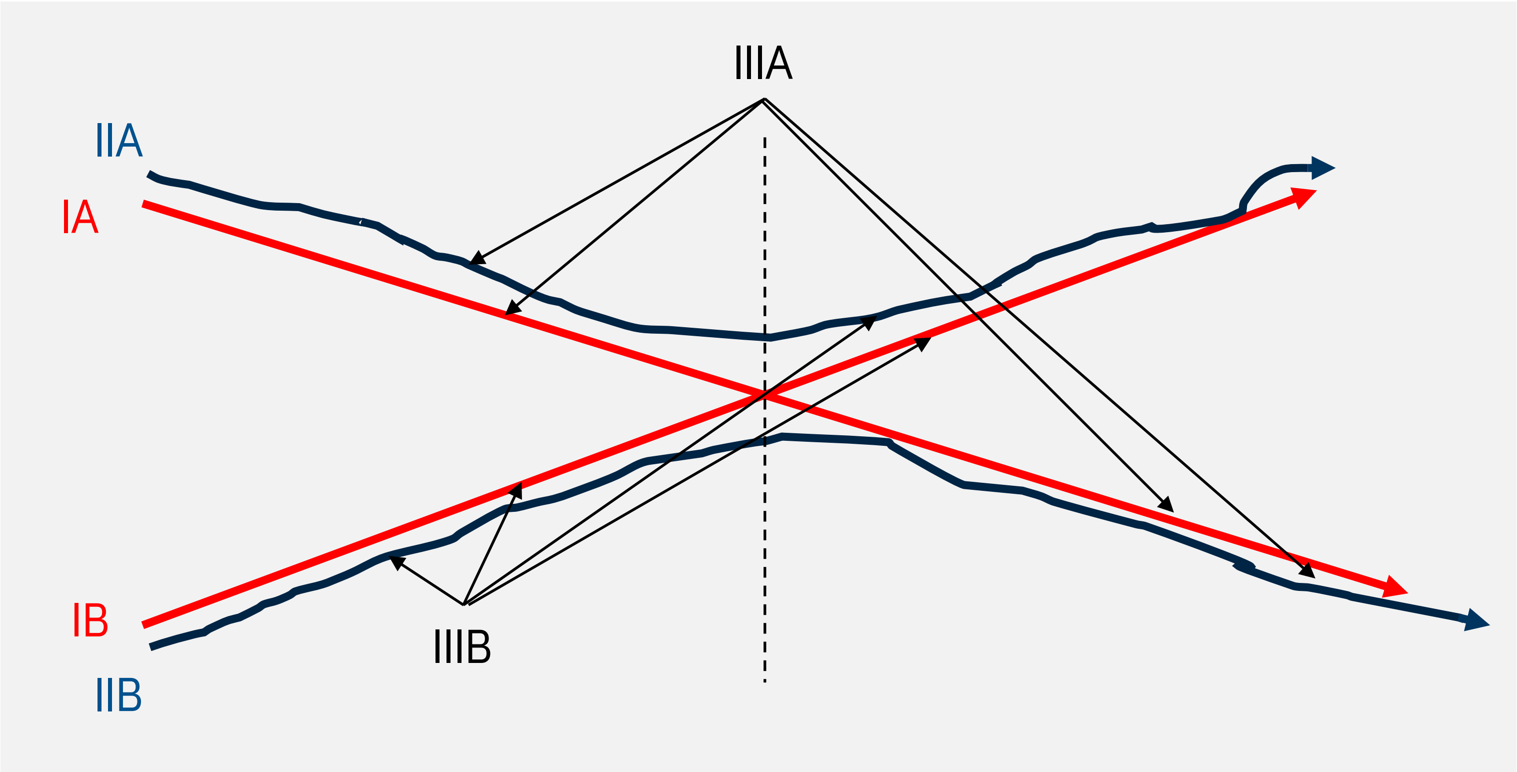

# Fusion of fragmented tracks

- No hard constraint to invalidate fused-track continuity

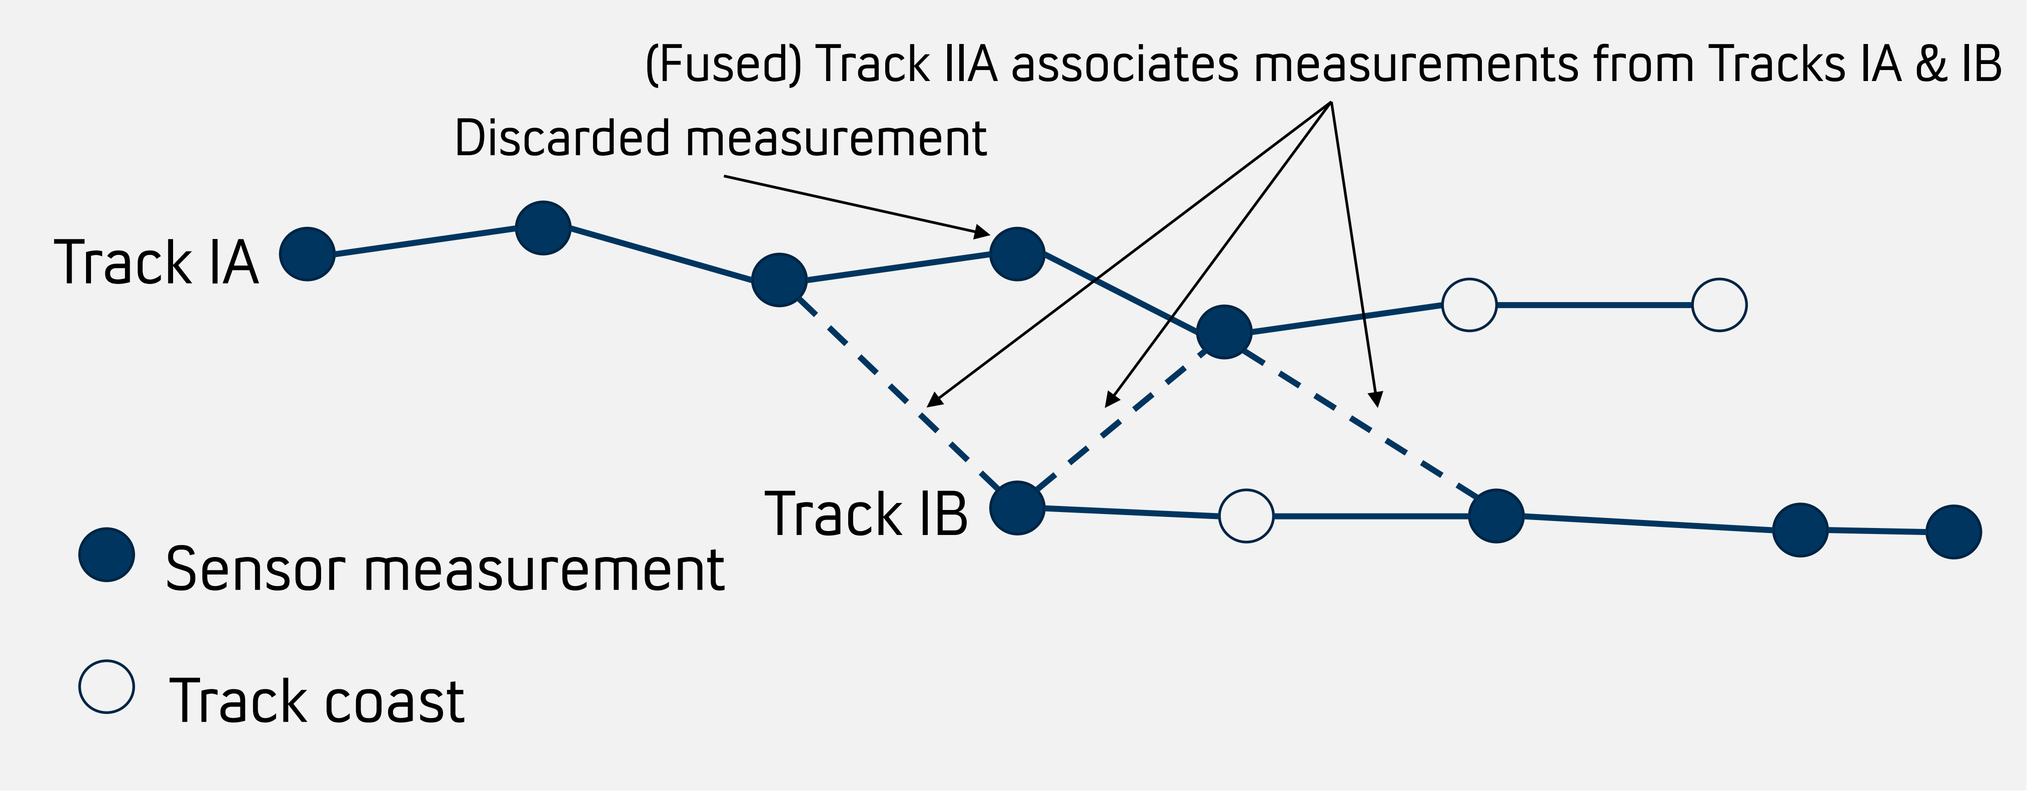

# Distributed tracker modeling

- Leverage *Mori Chang Chong* (MCC) exponential model for *probability of correct association* (PCA)
- Assume process noise governed by two-state Markov chain with $q_k \in \{q_{low}, q_{high}\}$

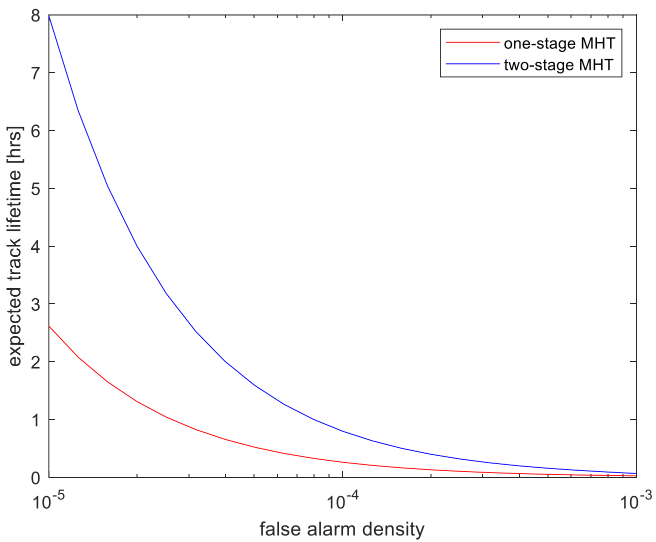


*S. Coraluppi et al., FUSION 2024*

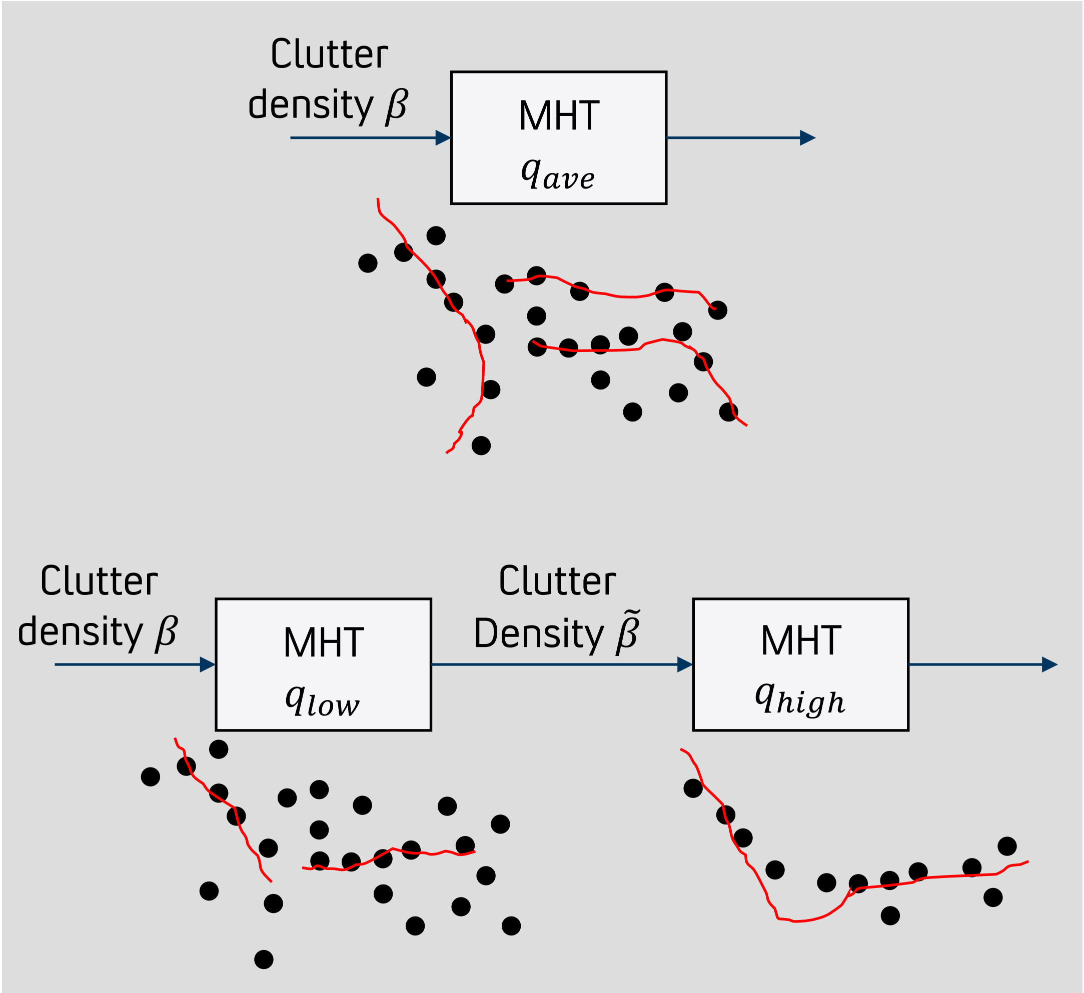

# Summary

- *Multi-target tracking* (MTT) is an essential component of many automated surveillance capabilities
  - *Multiple-hypothesis tracking* (MHT) is the leading operational approach to MTT
  - *Graph-based tracking* (GBT) is a fast, approximate version of MHT
- There is significant empirical evidence of the performance and robustness benefits of distributed MHT
  - There is a ***large design space*** of fusion architecture, connectivity, and modules
  - ***Calibration*** is required for downstream processing stages
  - Some ***undesirable effects*** are due to the hard constraints posed by upstream data association decisions
- A simplified distributed tracking paradigm relaxes upstream data association hard constraints
  - Track label may be posed as a ***target feature state***