Investigating Multi-Agent Terrain-Aided Navigation

Eric Sevonius*†, Fredrik Gustafsson†, and Gustaf Hendeby†
* Saab Dynamics, Linköping, Sweden
e-mail: eric.sevonius1@saabgroup.com
† Dept. of Electrical Engineering, Linköping University, Linköping, Sweden
e-mail: firstname.lastname@liu.se

***Abstract*—Due to recent events that have demonstrated the vulnerabilities of *global navigation satellite systems* (GNSS) there has been an increased interest in alternative methods for localization. One traditional alternative method is *terrain-aided navigation* (TAN), where a platform localizes itself by measuring the terrain elevation and comparing it to a *digital elevation map* (DEM). While single-agent TAN has been extensively studied, multi-agent TAN remains less explored. We are addressing the multi-agent TAN problem with a focus on its properties and performance bounds. So far, the snapshot performance, *i.e* the performance of a solution from one time step, has been investigated. As a part of this investigation, an approximate *Cramér-Rao Lower Bound* (CRLB) has been derived and a simulation study conducted where the performance with respect to the altitude sensor accuracy, the group formation accuracy, the number of agents and their formation was evaluated. Notably, we observe that the solution is relatively insensitive to errors in agent position, suggesting that low-accuracy inertial navigation systems and distance sensors are sufficient for determining their positions. Increasing the number of agents beyond a few seems to have a large effect on both the efficiency and robustness of the estimator, which lessens as the number of agents increases. However, increasing the number of agents does not compensate for poor altitude sensor quality. Additionally, while spatial separation between agents is important for effective map utilization, further separation beyond a certain point does not enhance performance. These findings provide design guidelines for multi-agent TAN systems and identify areas for further research.**

Background

Recent events have shown that *global navigation satellite systems* (GNSS) are susceptible to interference by malicious actors. However, the need for high precision localization remains. Navigation using terrain elevation, commonly referred to as *terrain-aided navigation* (TAN), has been explored extensively. A typical application is an airplane measuring the terrain elevation using radar and comparing these measurements to an elevation database. An inherent difficulty in utilizing measurements of the elevation of the terrain is that a single measurement will correspond to multiple points in an elevation database. By collecting multiple measurements of the elevation at approximately known locations, the number of possible locations can ultimately be narrowed down significantly. This is typically achieved by moving across the terrain and estimating the relative displacement between the measurements using inertial sensors. The principle of TAN is illustrated in Fig. 1, where measurements at two positions separated by a known relative position d is sufficient to obtain a unique position by considering the intersection between the possible locations of the two measurements.

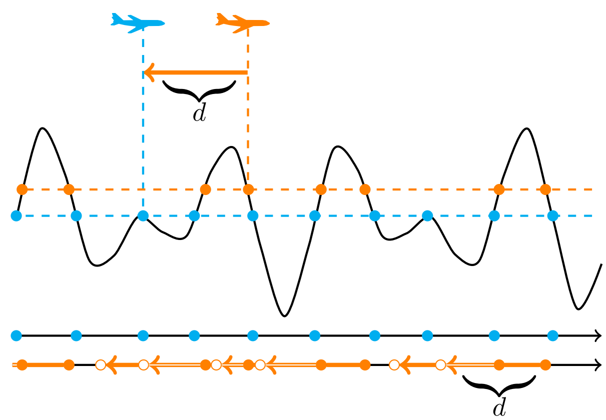

Fig. 1: Illustration of a TAN scenario. Two measurements of the terrain elevation are obtained at a known distance d apart.

Recently, cooperative localization has seen increased interest. A group of platforms equipped with sensors can leverage the information from its neighbors to improve the overall positioning of the whole group. In the TAN-setting, multiple agents can act as a sensor array and not only sample in time, but also spatially. If the relative positions between the sensors is approximately known (using e.g. sensors or estimates from external sources) then the setup is analogous to the scenario in Fig. 1, using two sensor platforms at a single time step. Some work on the topic of multi-agent TAN has been done [1, 2], where the focus has been on frameworks for decentralized estimation.

Consider an array of N sensors located at the relative positions d_n, $n = 0, \dots, N-1$. The first sensor is considered to constitute the origin of the array, *i.e*, $d_0 = 0$. At a single time step this results in the individual measurements

$$y_n = h(x + d_n + w_n) + e_n, \tag{1}$$

where $h(x)$ is the terrain database, $w_n \sim \mathcal{N}(0, \Sigma_w)$ is the uncertainty in the relative positions in the array and $e_n \sim \mathcal{N}(0, \sigma_e^2)$ is the error in the terrain altitude sensors. The configuration is shown in Fig. 2. Using this formalism, the multi-agent TAN problem can be stated as finding the first

This work was performed within the Competence Center SEDDIT (Sensor Informatics and Decision making for the DIgital Transformation), supported by Sweden's Innovation Agency within the research and innovation program Advanced digitalization.

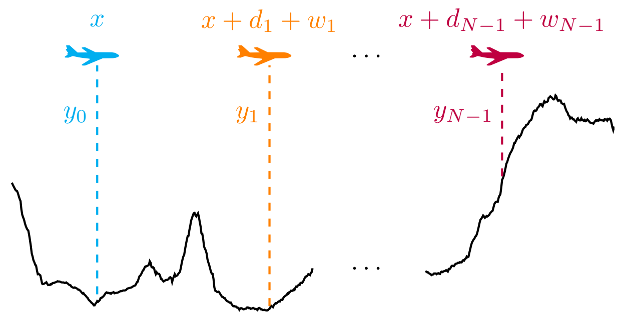

Fig. 2: Illustration of multi-agent TAN. The pieces of terrain are a part of a real digital elevation map of Sweden, Markhöjdmodell Nedladdning, grid 1+. ©Lantmäteriet.

sensor position x that best matches $y_1, \ldots, y_{N-1}$ in $h(x)$, or equivalently

$$\hat{x} = \arg \min_x V(x), \tag{2}$$

where $V(x)$ quantifies how well $y_1, \ldots, y_{N-1}$ match $h(x)$.

In [3], we examine what properties affect performance in a multi-agent TAN system. Specifically, we examine how the quality of altimeter sensors and relative distance estimates affect the performance of the estimator. These are factors that can be affected through, *e.g.*, using more expensive sensors. Additionally, we examine how the composition of the array affects performance by considering the effect of varying the number of sensors as well as the mean distance between them.

Results

In [3] we derived an approximate CRLB and conduct a simulation study for multi-agent TAN. The simulation study considered positioning an array of sensors in a (for simplicity) one-dimensional environment using a snapshot estimate. A real digital elevation map of Sweden was used to simulate height measurements, which is partially shown in Fig. 2. To investigate the properties of the estimate, four key parameters were identified and the influence of varying them investigated: the number of sensor platforms N, the mean relative distance between the platforms d, the size of the error of the elevation measurement σ_e^2, and the size of the error of the estimated relative position between the sensor platforms σ_w^2. The parameters were varied about a nominal set of parameters: $\sigma_e^2 = 1, \sigma_w^2 = 25, N = 10, d = 20$. The evaluation consisted of determining the robustness and efficiency of the estimator using Monte Carlo trials evaluated over a grid over the map. The robustness was evaluated by considering the fraction of successful trials, where a trial was deemed unsuccessful if its error was larger than 50 m. The efficiency of the estimator was evaluated using the *root mean square error* (RMSE) with respect to the successful trials. The results when varying N and σ_e are presented in Fig. 3 and Fig. 4 respectively.

One interesting observation that was made is that the solution is relatively insensitive to errors in the relative positions of the agents, which indicates that relatively low accuracy

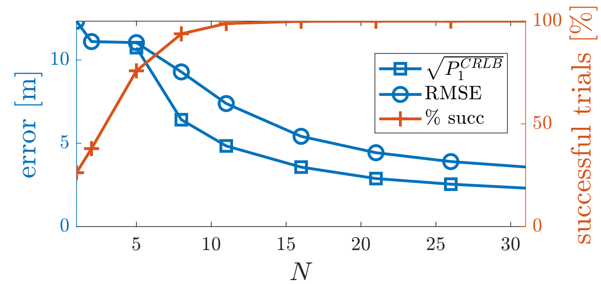

Fig. 3: Results for tests when varying the number of sensors N. Here the values of $\sqrt{P_1^{\text{CRLB}}}$ are omitted for $N < 6$.

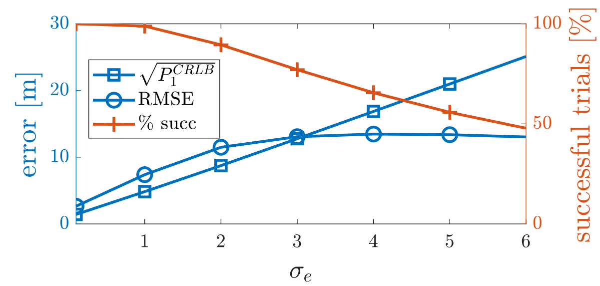

Fig. 4: Results for tests when varying the standard deviation of the error of the altimeters σ_e.

inertial navigation systems and distance sensors should suffice when determining the agent configuration. The CRLB and the results indicate that there is a lot of performance to be gained by increasing the number of agents beyond one, but that the effect is less significant asymptotically. However, it was also observed that to increase the number of agents is not a solution to low quality altitude sensors. That is, adding more agents is not a remedy for poor altitude sensors. It was also shown that to be efficient, the agents must be separated in space to utilize the variations in the map, but that once enough separation is obtained, separating the agents more will not improve the localization performance further.

Further Work

For future work we would like to consider the connection between the terrain and the performance closer. By considering the terrain as a realization from a stochastic process, it should be possible to connect the findings in this paper to the properties in the terrain, such as the correlation.

References

[1] André R. Braga, Marcelo G. S. Bruno, Emre Özkan, Carsten Fritsche, and Fredrik Gustafsson. Cooperative Terrain Based Navigation and coverage identification using consensus. In *2015 18th International Conference on Information Fusion (Fusion)*, pages 1190–1197, 2015.

[2] Hallysson Oliveira, Stiven Schwanz Dias, and Marcelo Gomes da Silva Bruno. Cooperative Terrain Navigation Using Hybrid GMM/SMC Message Passing on Factor Graphs. *IEEE Transactions on Aerospace and Electronic Systems*, 56(5):3958–3970, 2020.

[3] Eric Sevonius, Fredrik Gustafsson, and Gustaf Hendeby. Exploring the Properties of Multi-Agent Terrain-Aided Navigation (under review). In *2025 28th International Conference on Information Fusion (FUSION)*, 2025.