A Linear-Complexity EKF for Visual-Inertial Navigation with Loop Closures

Patrick Geneva, Kevin Eckenhoff, and Guoquan Huang

Motivation

- Want to leverage loop-closure information to perform drift-free visual-inertial navigation
- Naive use of keyframes for loop-closure increases filter complexity
- Achieve linear complexity while still allowing for frequent loop-closures

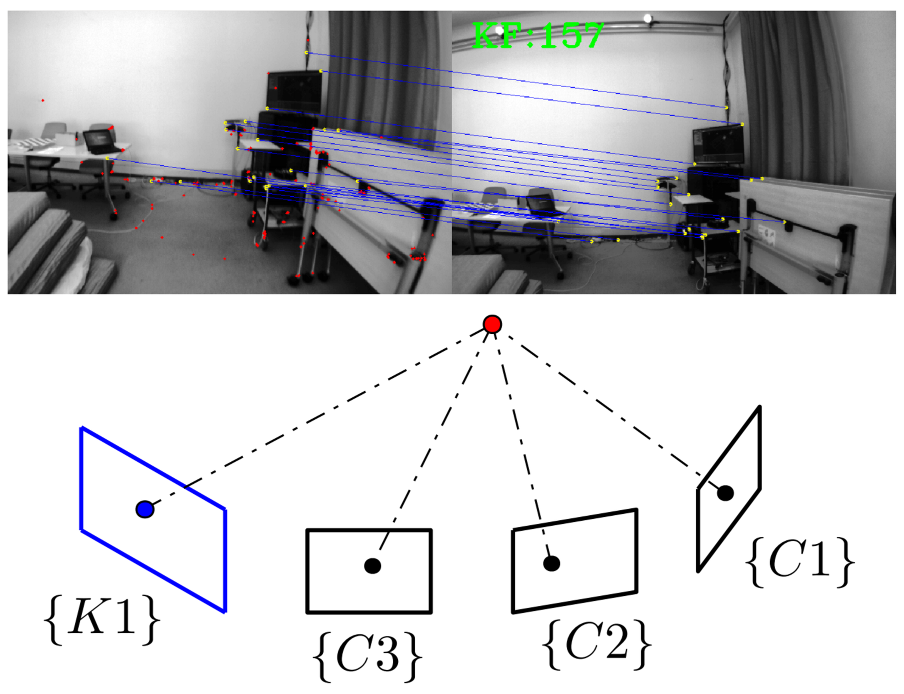

Figure 1. Illustration of keyframe-based loop-closure. Can leverage a match to a prior pose {K1} to indirectly correct the current poses {C1}-{C3}.

Contributions

- Leverage Schmidt-Kalman filter to reduce complexity of EKF update
- To obtain loop-closure measurements for current features, leverage 2D-to-2D matching to historical keyframes
- Due to loop-closures, system outperforms standard MSCKF with minimal computational overhead

Standard MSCKF

- Standard MSCKF [Mourikis 2007] estimates current IMU along with history of poses

$$\begin{bmatrix}\hat{\mathbf{x}}_A^+ \\ \hat{\mathbf{x}}_S^+\end{bmatrix} = \begin{bmatrix}\hat{\mathbf{x}}_A^- \\ \hat{\mathbf{x}}_S^-\end{bmatrix} + \begin{bmatrix}\mathbf{L}_A \\ \mathbf{L}_S\end{bmatrix}\mathbf{S}^{-1}\mathbf{r}$$

$$\begin{bmatrix}\mathbf{P}_{AA}^+ & \mathbf{P}_{AS}^+ \\ \mathbf{P}_{SA}^+ & \mathbf{P}_{SS}^+\end{bmatrix} = \begin{bmatrix}\mathbf{P}_{AA}^- & \mathbf{P}_{AS}^- \\ \mathbf{P}_{SA}^- & \mathbf{P}_{SS}^-\end{bmatrix} - \begin{bmatrix}\mathbf{L}_A\mathbf{S}^{-1}\mathbf{L}_A^\top & \mathbf{L}_A\mathbf{S}^{-1}\mathbf{L}_S^\top \\ \mathbf{L}_S\mathbf{S}^{-1}\mathbf{L}_A^\top & \mathbf{L}_S\mathbf{S}^{-1}\mathbf{L}_S^\top\end{bmatrix}$$

Schmidt-MSCKF

- Set Kalman gain of “nuisance” parameters to zero to reduce complexity during update

$$\begin{bmatrix}\hat{\mathbf{x}}_A^+ \\ \hat{\mathbf{x}}_S^+\end{bmatrix} = \begin{bmatrix}\hat{\mathbf{x}}_A^- \\ \hat{\mathbf{x}}_S^-\end{bmatrix} + \begin{bmatrix}\mathbf{L}_A \\ \mathbf{0}\end{bmatrix}\mathbf{S}^{-1}\mathbf{r}$$

$$\begin{bmatrix}\mathbf{P}_{AA}^+ & \mathbf{P}_{AS}^+ \\ \mathbf{P}_{SA}^+ & \mathbf{P}_{SS}^+\end{bmatrix} = \begin{bmatrix}\mathbf{P}_{AA}^- & \mathbf{P}_{AS}^- \\ \mathbf{P}_{SA}^- & \mathbf{P}_{SS}^-\end{bmatrix} - \begin{bmatrix}\mathbf{L}_A\mathbf{S}^{-1}\mathbf{L}_A^\top & \mathbf{L}_A\mathbf{S}^{-1}\mathbf{L}_S^\top \\ \mathbf{L}_S\mathbf{S}^{-1}\mathbf{L}_A^\top & \mathbf{0}\end{bmatrix}$$

Loop-closure Measurements

- Measurements can be written as function of both states

$$\mathbf{r}'_f \simeq \mathbf{H}_{A_k}\tilde{\mathbf{x}}_{A_{k|k-1}} + \mathbf{H}_{S_k}\tilde{\mathbf{x}}_{S_{k|k-1}} + \mathbf{n}'_f$$

- Match current features to past keyframes to get additional loop-closure measurements (see Figure 1)

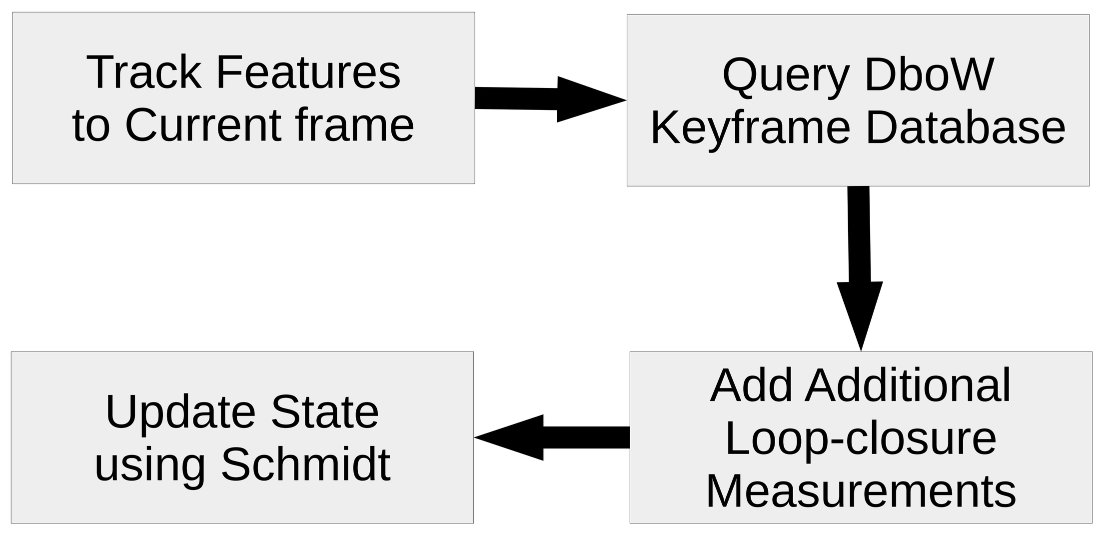

Complexity

- Cost reduces to O(n) due to only needing to update the cross-terms of the covariance

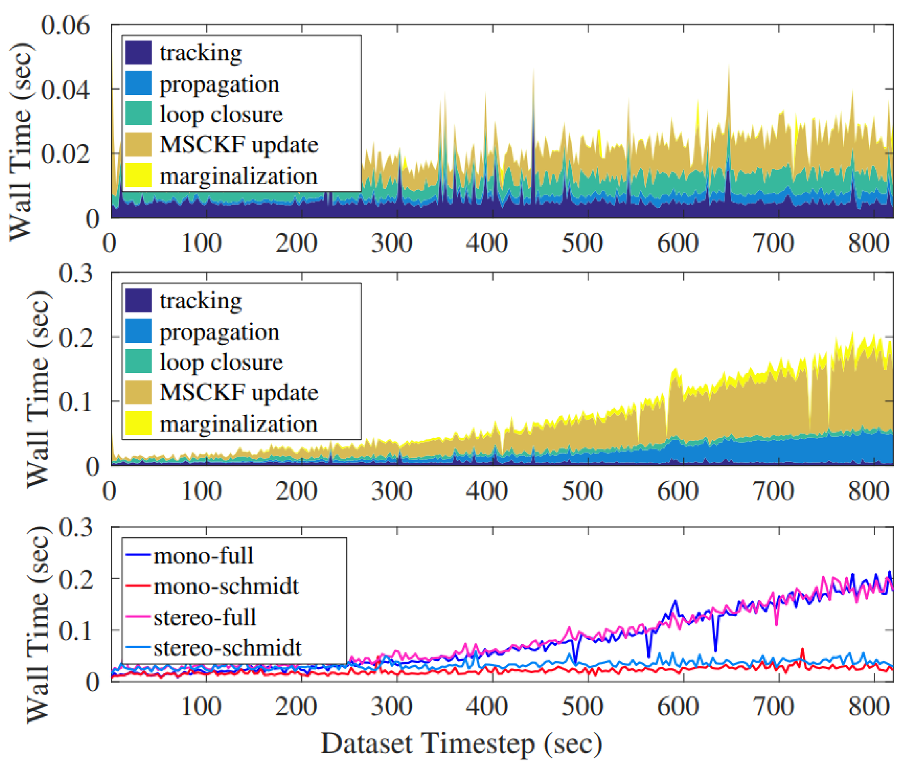

- Full covariance update quickly becomes computationally infeasible (middle of above figure)
- Schmidt remains realtime with over 400 keyframe poses by the end of the trajectory

Experiments

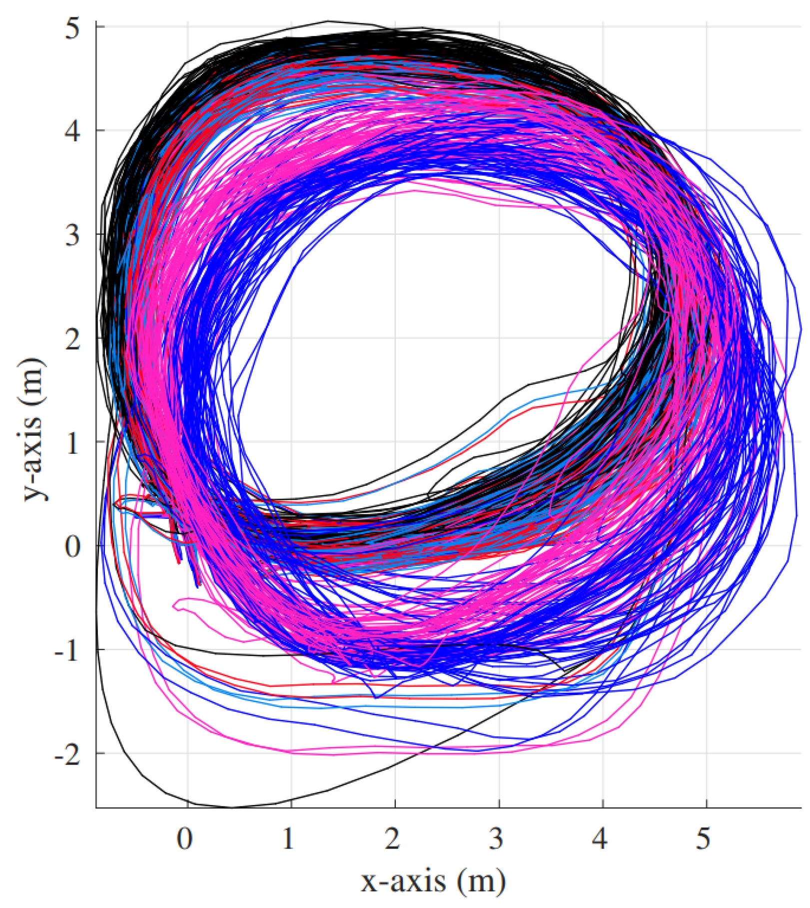

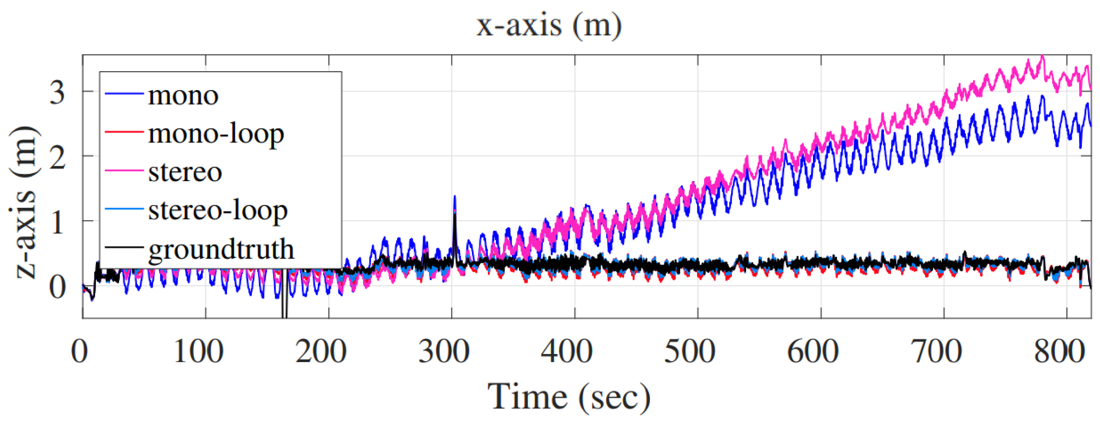

- Position RMSE error (meters) averaged over 30 runs

	MSCKF	MSCKF w. Full Loops	Schmidt-MSCKF
Monocular	1.660	0.243	0.266
Stereo	1.586	0.171	0.213

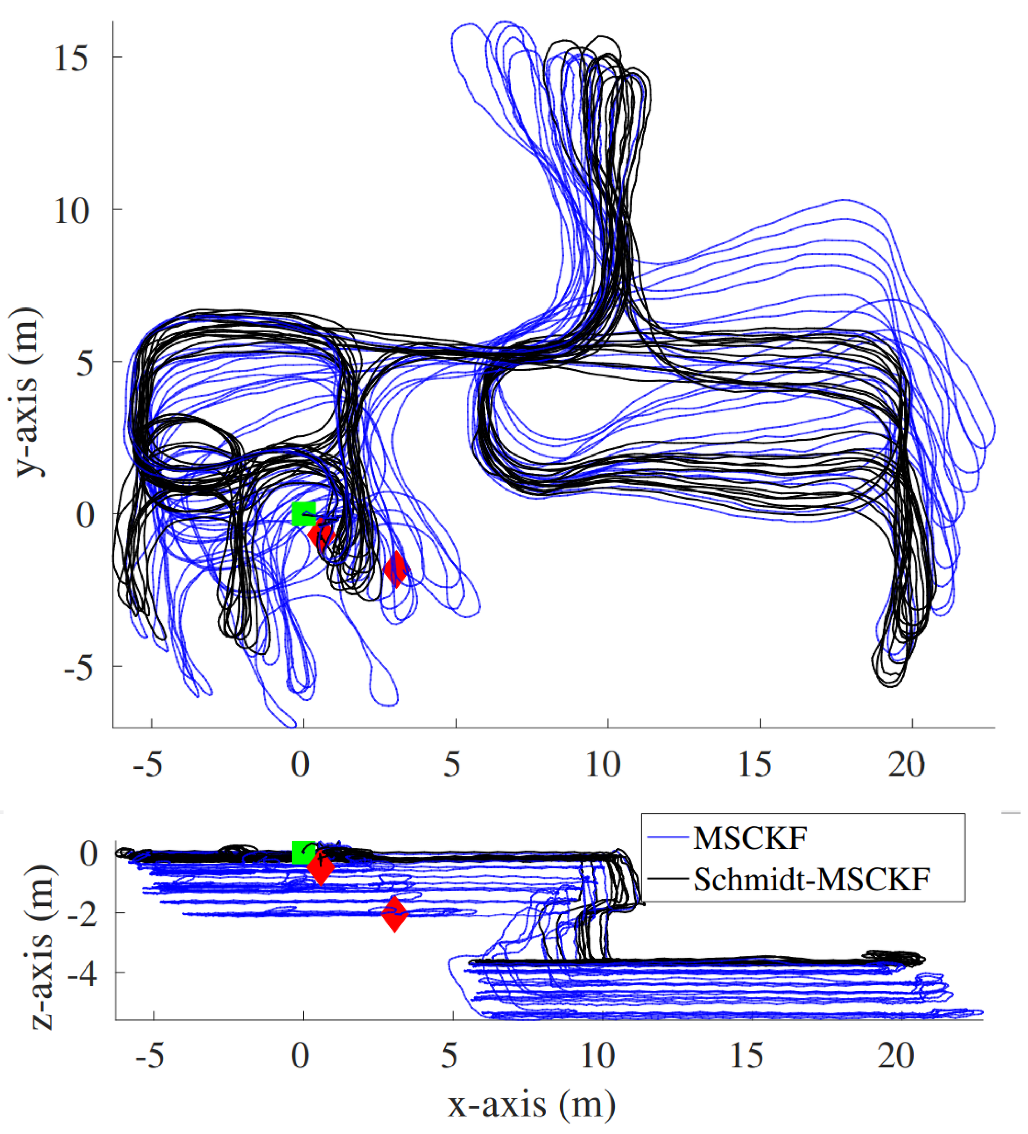

- Evaluated on two datasets (1.2km and 1.5km)
- Show an improvement over the standard MSCKF with performance close to full covariance updating
- Limited drift over each dataset while still remaining realtime