Parallelization of Scan Matching for Robotic 3D Mapping

Andreas Nüchter
University of Osnabrück
Institute of Computer Science
Knowledge Based Systems Research Group

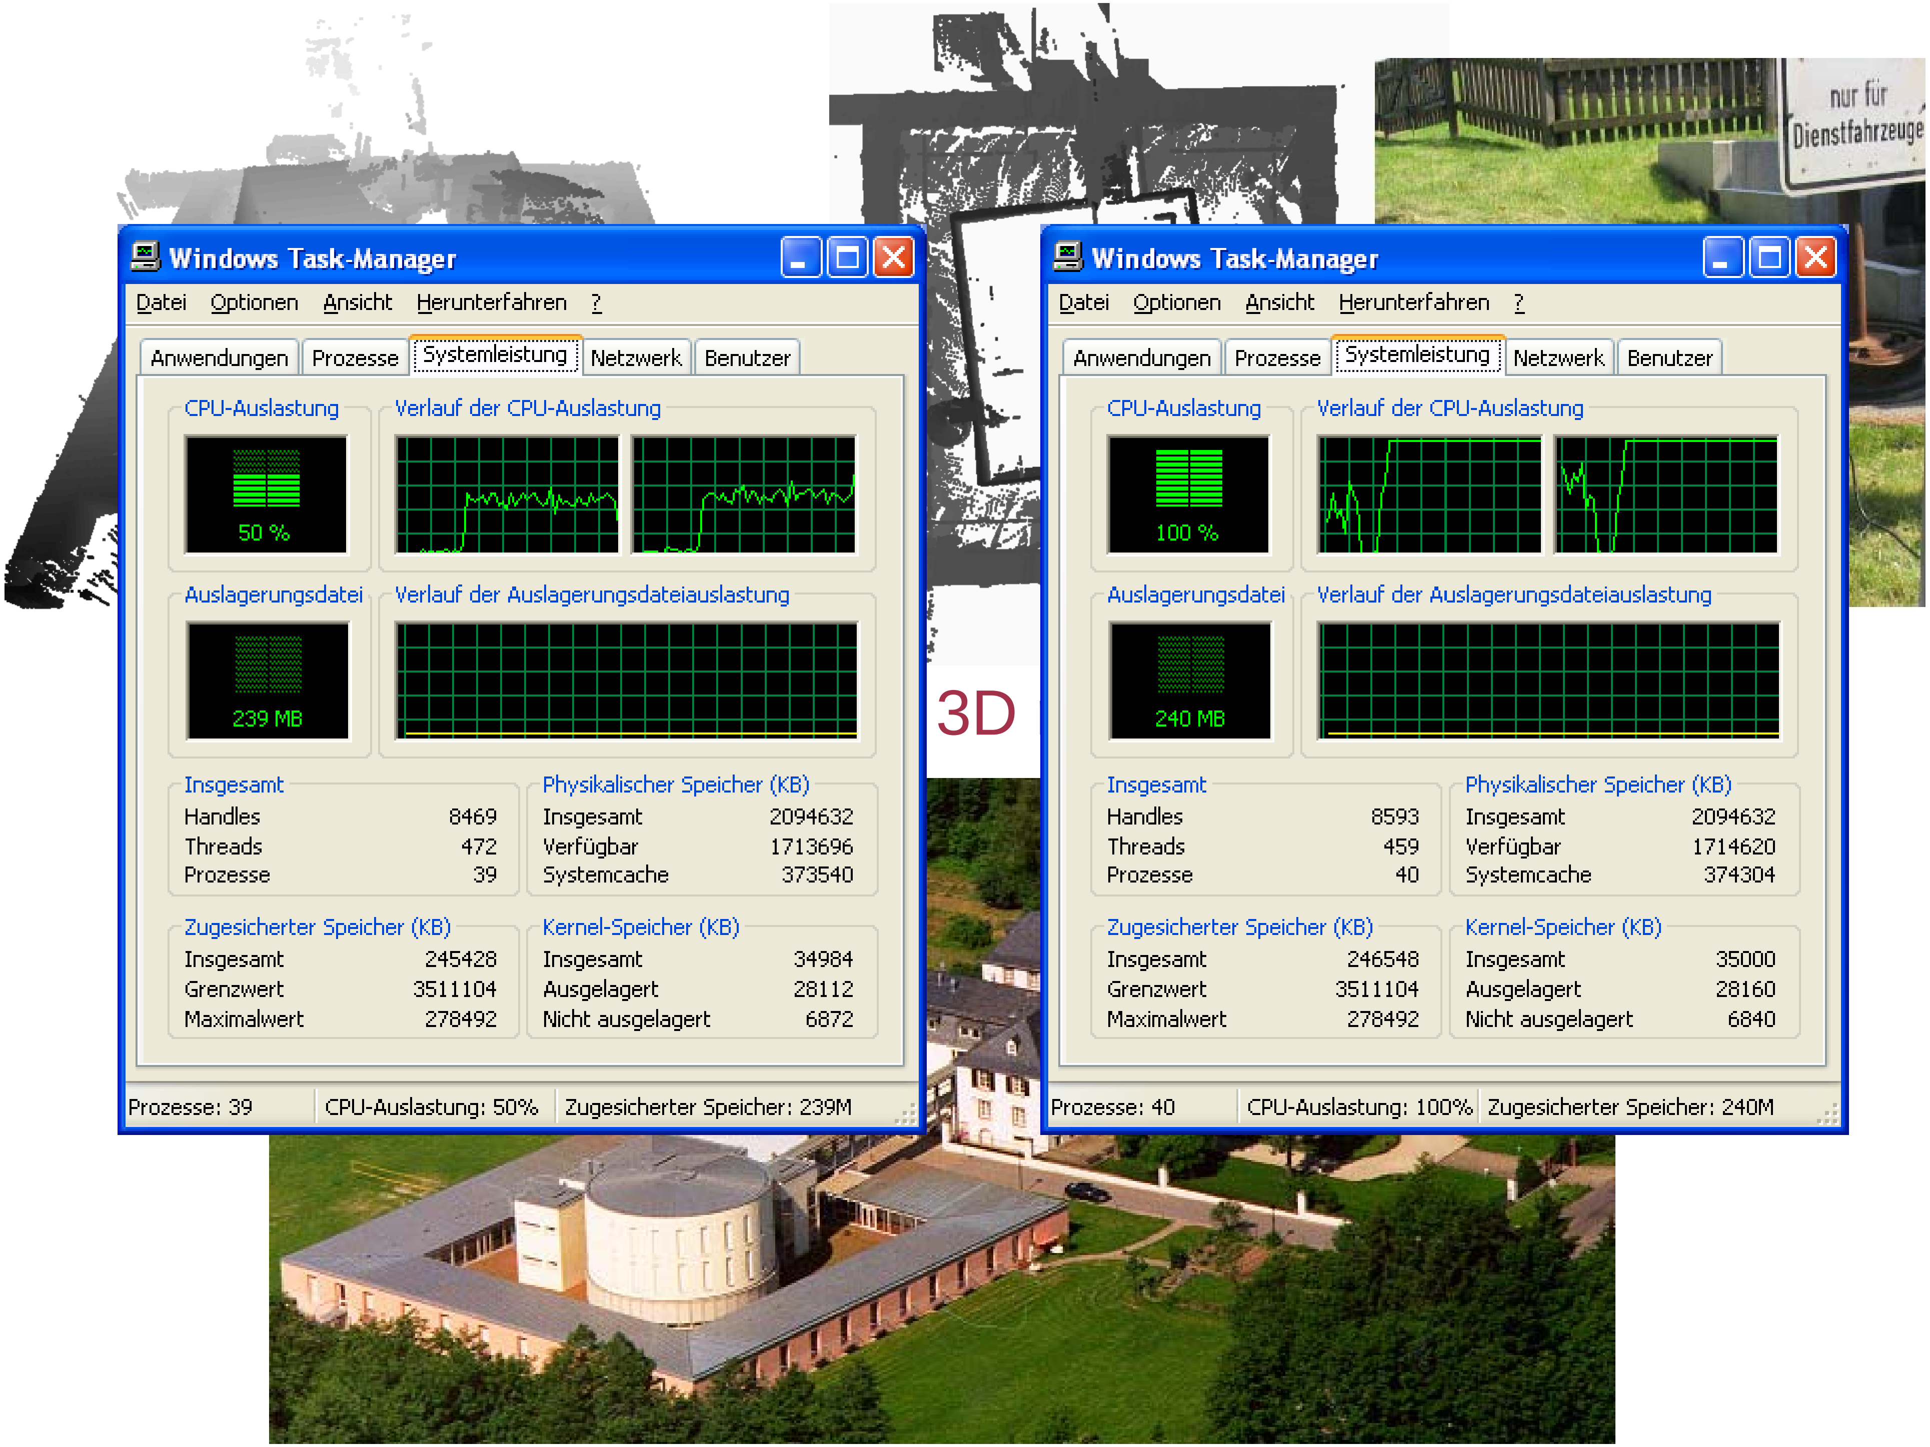
nur für
Dienstfahrzeuge
3D
Windows Task-Manager
Datei Optionen Ansicht Herunterfahren ?
Anwendungen Prozesse Systemleistung Netzwerk Benutzer
CPU-Auslastung
50 %
Verlauf der CPU-Auslastung
Auslagerungsdatei
239 MB
Verlauf der Auslagerungsdateiauslastung
Insgesamt
Handles 8469
Threads 472
Prozesse 39
Physikalischer Speicher (KB)
Insgesamt 2094632
Verfügbar 1713696
Systemcache 373540
Zugesicherter Speicher (KB)
Insgesamt 245428
Grenzwert 3511104
Maximalwert 278492
Kernel-Speicher (KB)
Insgesamt 34984
Ausgelagert 28112
Nicht ausgelagert 6872
Prozesse: 39
CPU-Auslastung: 50%
Zugesicherter Speicher: 239M
Windows Task-Manager
Datei Optionen Ansicht Herunterfahren ?
Anwendungen Prozesse Systemleistung Netzwerk Benutzer
CPU-Auslastung
100 %
Verlauf der CPU-Auslastung
Auslagerungsdatei
240 MB
Verlauf der Auslagerungsdateiauslastung
Insgesamt
Handles 8593
Threads 459
Prozesse 40
Physikalischer Speicher (KB)
Insgesamt 2094632
Verfügbar 1714620
Systemcache 374304
Zugesicherter Speicher (KB)
Insgesamt 246548
Grenzwert 3511104
Maximalwert 278492
Kernel-Speicher (KB)
Insgesamt 35000
Ausgelagert 28160
Nicht ausgelagert 6840
Prozesse: 40
CPU-Auslastung: 100%
Zugesicherter Speicher: 240M

Contents

- State of the Art in Robotic Mapping
- 6D SLAM with ICP Scan Matching
- Parallelization
- 3D Mapping Examples

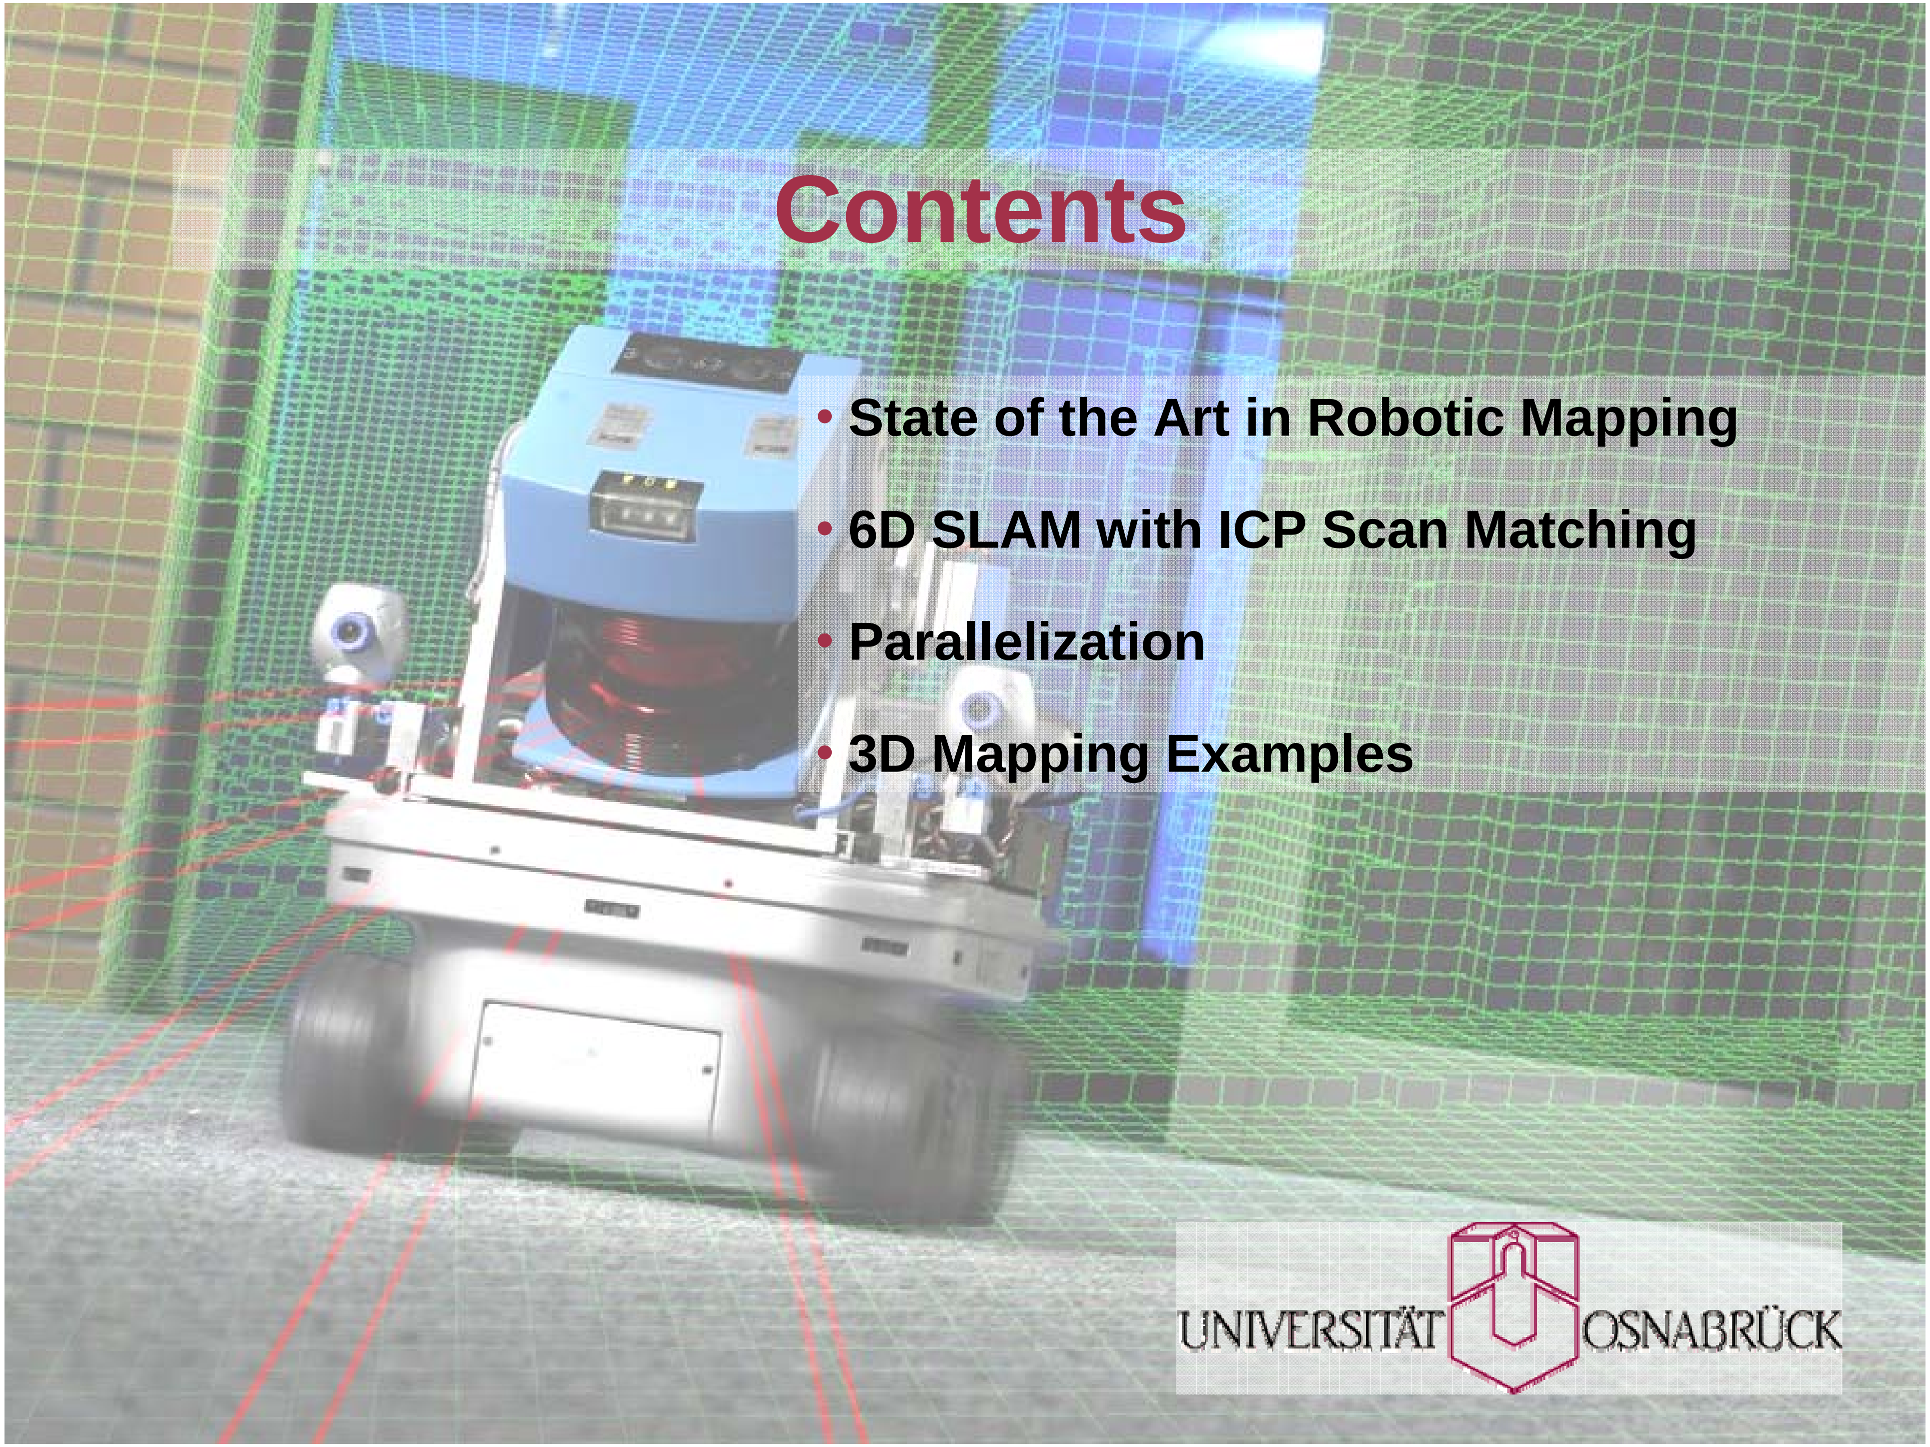

UNIVERSITÄT OSNABRÜCK

Contents

- **State of the Art in Robotic Mapping**
- **6D SLAM with ICP Scan Matching**
- **Parallelization**
- **3D Mapping Examples**

UNIVERSITÄT OSNABRÜCK

Simultaneous Localization and Mapping

- If one knows the pose (position and orientation) of a mobile robot precisely, then the sensor readings can be used to build a map.
- Unfortunately, pose measurements are always imprecise ☹
- The pose of a robot is easy to compute from sensor readings, given a map.

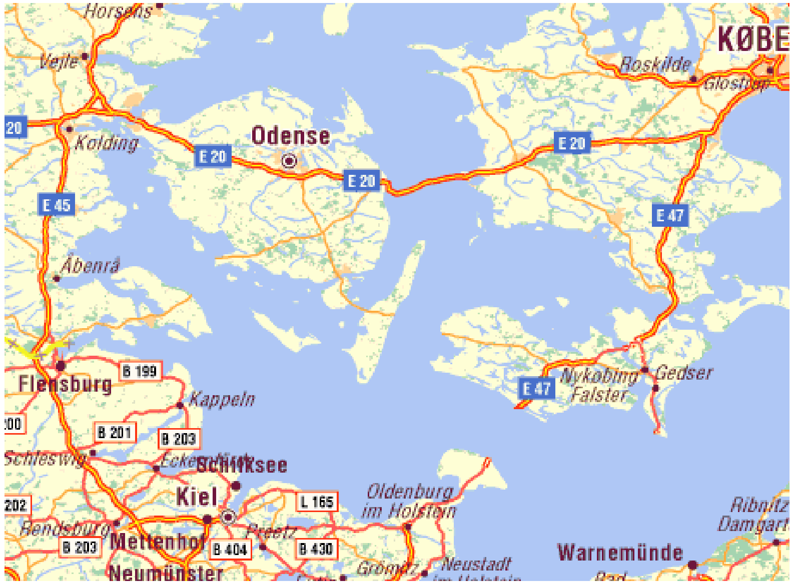

UNIVERSITÄT OSNABRÜCK

State of the Art in Robotic Mapping (1)

- Laser scanner are the state of the art sensors for metrical environment mapping
- Mapping based on scan matching (Lu, Milios)
- Probabilistic theory of mapping using uncertain motion and sensor models (Kalman-Filter, Maximum Likelihood Estimation, Expactation Maximation).

⇨ 2D Data, 3D Poses

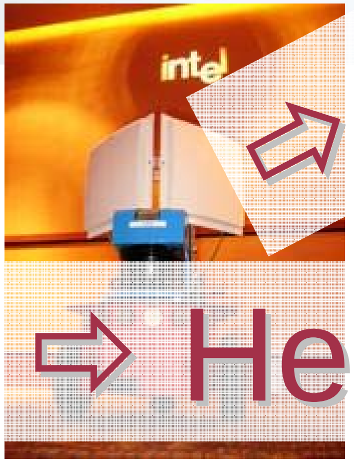

(2nd data set)

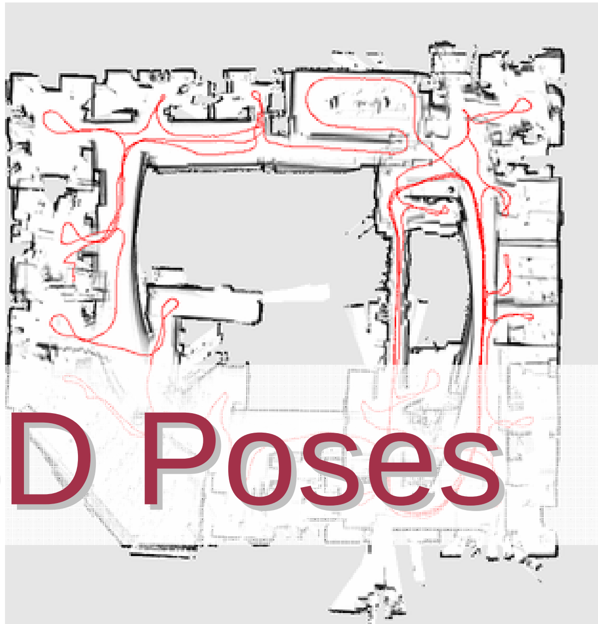

⇨ Here: 3D Data, 6D Poses

Related work: Freiburg (Burgard), UW (Fox), Stanford (Thrun)

State of the Art in Robotic Mapping (2)

		Dimensionality of pose representation: 3D	Dimensionality of pose representation: 6D
Sensor data	2D	2D mapping of planar sonar and laser scans. See (Thrun, 2002) for an overview.	3D mapping using a precise localization, considering the x,y,z-position and the roll yaw and pitch angle
Sensor data	3D	3D mapping using a planar localization method and, e.g., an upward looking laser scanner	3D mapping using 3D laser scanners or (stereo) cameras with pose estimates calculated from the sensor data




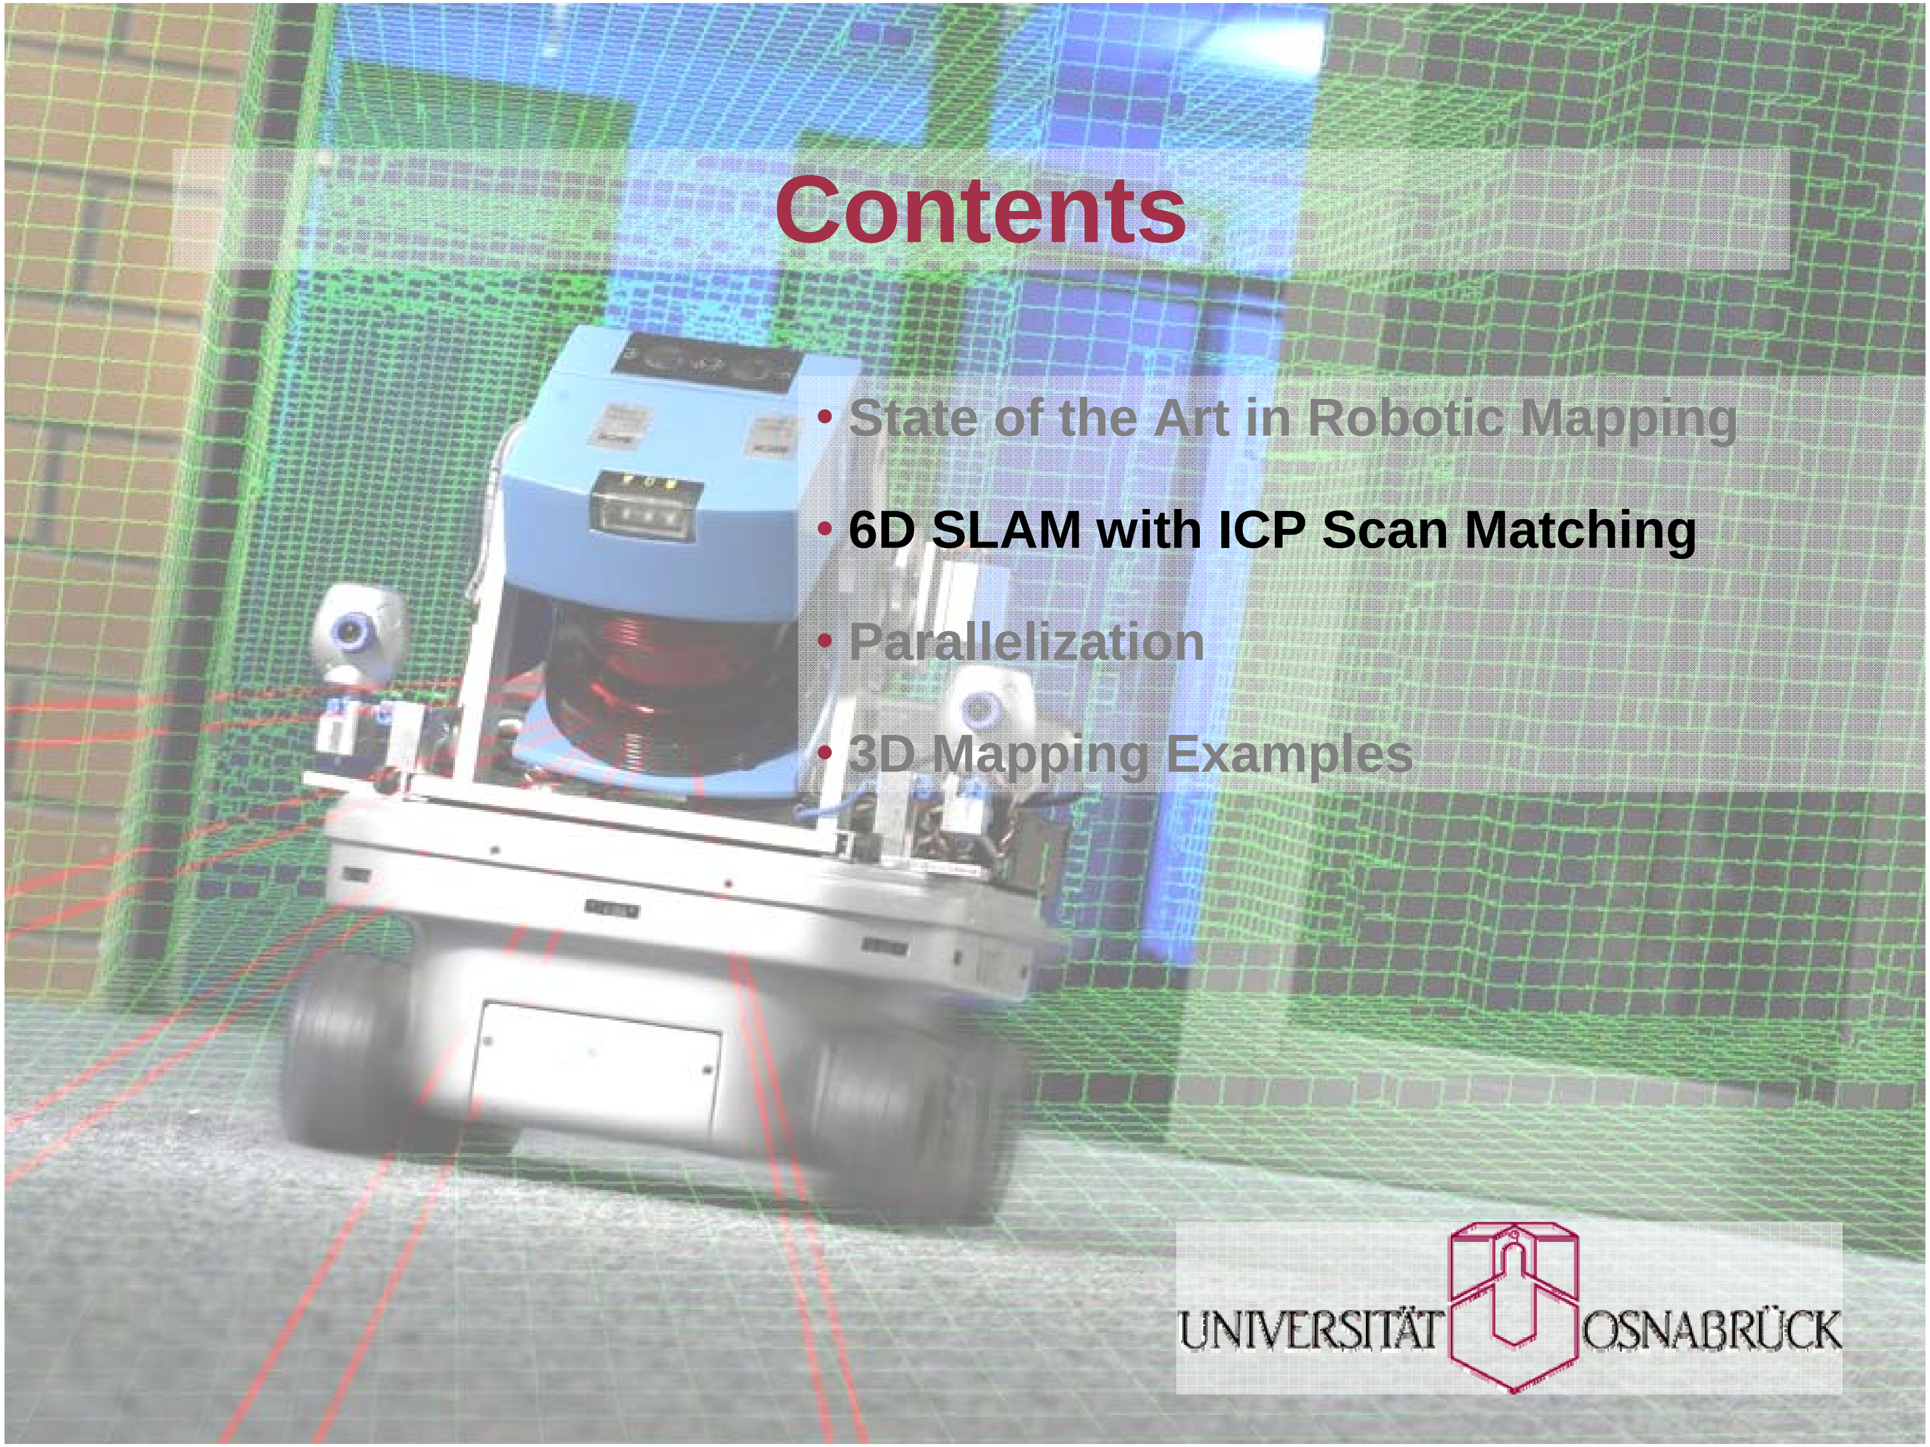

Contents

- State of the Art in Robotic Mapping
- **6D SLAM with ICP Scan Matching**
- Parallelization
- 3D Mapping Examples

UNIVERSITÄT OSNABRÜCK

The ICP-Algorithm

Scan registration Put two independent scans into one frame of reference

Iterative Closest Point algorithm [Besl/McKay 1992]

For prior point set M ("model set") and data set D

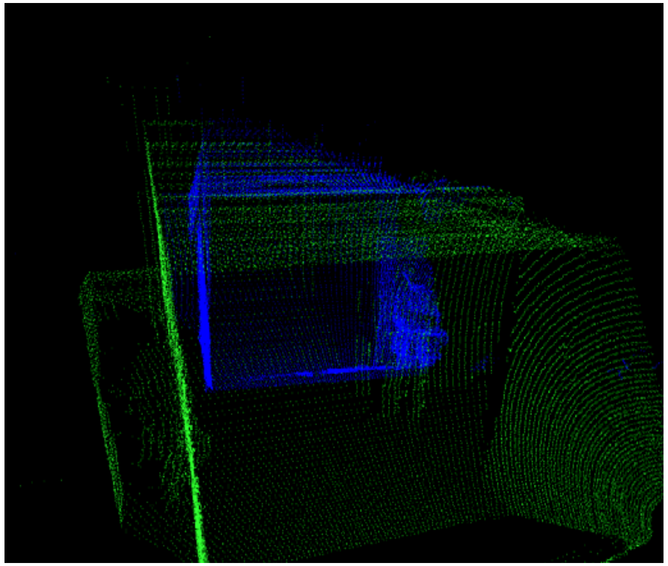

1. Select point correspondences $w_{i,j}$ in {0,1}
2. Minimize for rotation $\mathbf{R}$, translation $\mathbf{t}$

$$E(\mathbf{R},\mathbf{t}) = \sum_{i=1}^{N_m}\sum_{j=1}^{N_d} w_{i,j} \left|\left|\mathbf{m}_i - (\mathbf{R}\mathbf{d}_j + \mathbf{t})\right|\right|^2$$

3. Iterate **1.** and **2.**

SVD-based calculation of rotation

• works in 3 translation plus 3 rotation dimensions
$\Rightarrow$ 6D SLAM with closed loop detection and global relaxation.

Global Relaxation

- Consecutive Scan Matching with ICP results in an erronous map, since small matching errors sum up.

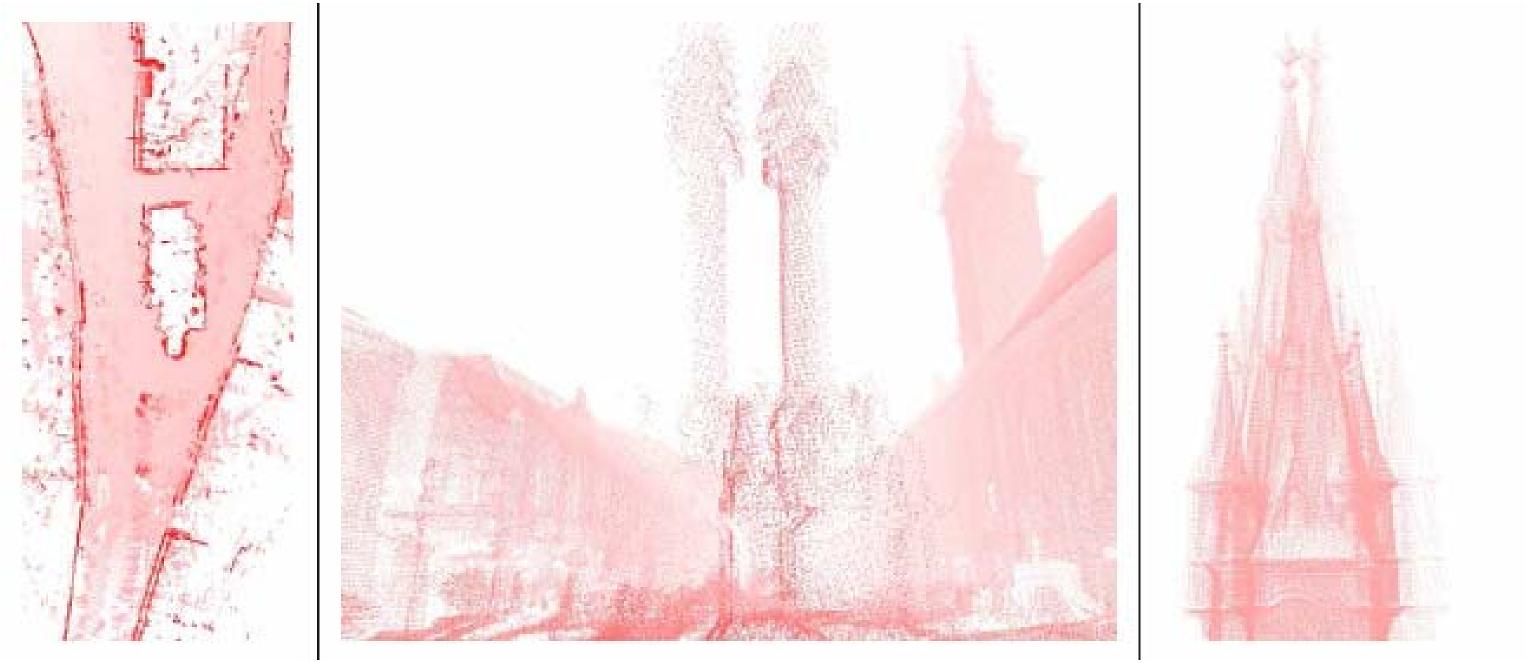

Global Relaxation

- Consecutive Scan Matching with ICP results in an erronous map, since small matching errors sum up.

Global Relaxation

- Consecutive Scan Matching with ICP results in an erronous map, since small matching errors sum up.

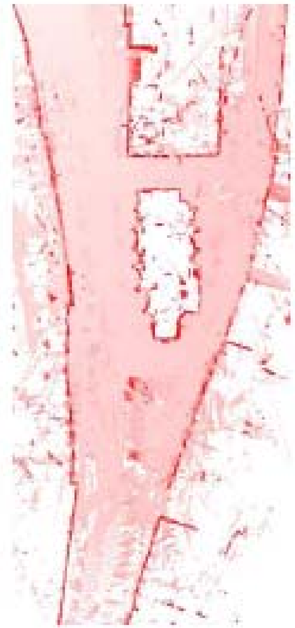

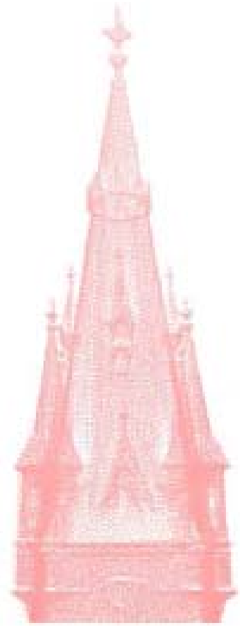

Global Relaxation

- Consecutive Scan Matching with ICP results in an erronous map, since small matching errors sum up.

⇨ Replace the ICP error function by a global one, i.e.,

$$D_{i,j} = X_i - X_j$$

$$W = \sum_{(i,j)} (D_{i,j} - \bar{D}_{i,j})^T C_{i,j}^{-1} (D_{i,j} - \bar{D}_{i,j})$$

where $\bar{D}_{i,j} = D_{i,j} + \Delta D_{i,j}$ models random Gaussian noise, added to the unkown exact pose $D_{i,j}$ and $C_{i,j}$ the co-variance matrix of the overlapping scans computed from **closest point pairs**.

(Video 1) (Video 2) (Video 3)

UNIVERSITÄT OSNABRÜCK

Contents

- State of the Art in Robotic Mapping
- 6D SLAM with ICP Scan Matching
- **Parallelization**
- 3D Mapping Examples

UNIVERSITÄT OSNABRÜCK

Parallel Transformation Computation

- Computing the optimal rotation

  $\mathbf{R} = \mathbf{V}\mathbf{U}^T$ from the correletion matrix $\mathbf{H} = \mathbf{U}\Lambda\mathbf{V}^T$ via SVD

$$\mathbf{H} = \sum_{i=1}^{N} \mathbf{m}_i'^T \mathbf{d}_i' = \sum_{i=1}^{N} (\mathbf{m}_i - \mathbf{c}_m)(\mathbf{d}_i - \mathbf{c}_d)^T$$
$$= \begin{pmatrix} S_{xx} & S_{xy} & S_{xz} \\ S_{yx} & S_{yy} & S_{yz} \\ S_{zx} & S_{zy} & S_{zz} \end{pmatrix},$$

  with $S_{xx} = \sum_{i=1}^{N} m'_{ix} d'_{ix}, \; S_{xy} = \sum_{i=1}^{N} m'_{ix} d'_{iy}, \; \ldots$

- Idea: Compute closest points in parallel and transfer point pairs back to master

⇨ This takes too long…

Transfer an intermediate correlation matrix back to master!

[13] C. Langis, M. Greenspan, and G. Godin. The parallel iterative closest point algorithm. In *Proc. IEEE 3DIM*, Quebec City, Canada, June 2001.

Parallel Computation of Point Pairs (1)

- Single threaded computation of Point Pairs is accomplished by *k*-d trees
- We have shared memory ⇨ No need to duplicate the points
  - Parallel constuction a *k*-d tree

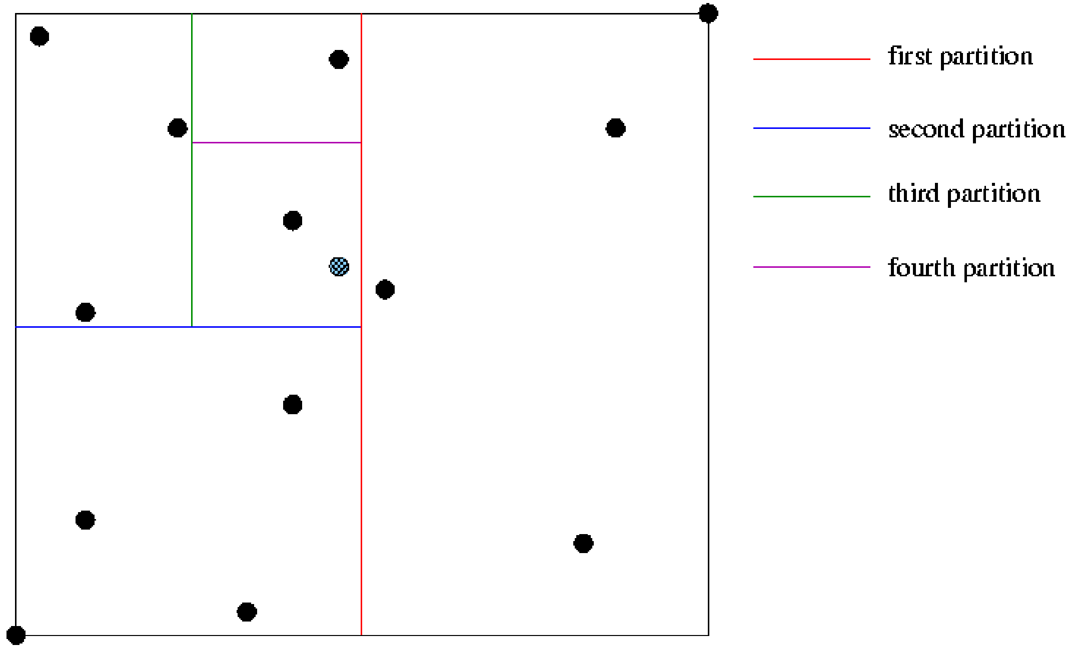

  - Parallel searching in *k*-d trees

Parallel Computation of Point Pairs (2)

```
class KDParams {
public:
  /** pointer to the closest point.
    *  size = 4 bytes of 32 bit machines
    */
  double *closest;
  /** distance to the closest point.
    * size = 8 bytes
    */
  double closest_d2;
  /** pointer to the point,
    * size = 4 bytes of 32 bit machines
    */
  double *p;
  /** expand to 128 bytes to avoid
    * false-sharing, 16 bytes from above
    * + 28*4 bytes = 128 bytes */
  int padding[28];
};
```

```
class KDtree {
  // [snip]
public:
   __declspec (align(16)) \
      static KDParams params[MAX_OPENMP_NUM_THREADS];
};
```

- *k*-d trees use back-tracking to determine the closest point and global variables

⇨ Encapsulate global variables in a class

- Pay attention to padding to avoid conflicts with the hardware cache

Parallel Computation of Point Pairs (3)

- Since Lu / Milios style 6D SLAM is also based on closest point pairs, we can immediately apply our ICP improvement
- Better: Compute point pairs in parallel of every link of the GraphSLAM graph.
  - It turns out, you don't need to change the *k*-d tree search structure.

How fast are we now?

1. 3D data sets acquired by Kurt3D (small 11 scans, single loop, office environment)

   Speedup: 23,8%

algorithm	one thread	two threads	four threads
ICP	12688	9821	9750
LUM	703	610	453
total	13391	10431	10203

2. 3D data sets acquired by Kurt3D (large, e.g., Dagstuhl castle, 83 scans, single loop)

   Speedup: 37,3%

3. Countinously acquired 3D data by Leibnitz Univ. Hannover (large, 1000 scans, trajectory up to 1 km, multiple loops)

   Speedup: 24,4%

4. Riegl data set (11 3D scans, single loop, high resolution, wide apex angle)

   Speedup: 23,7% (overall)
   48,2% (ICP alone)

number of 3D scans	one thread	two threads	four threads
2 (ICP)	20750	10985	10938
3 (ICP)	41984	21750	21750
6 (ICP)	134031	77390	77047
11 (ICP)	369828	218125	210515
11 (LUM)	794110	690023	678111
total	1163938	908148	888626

Up to 70% on a 4 processor dual core Itanium-2 (Osnabrück's compute server)

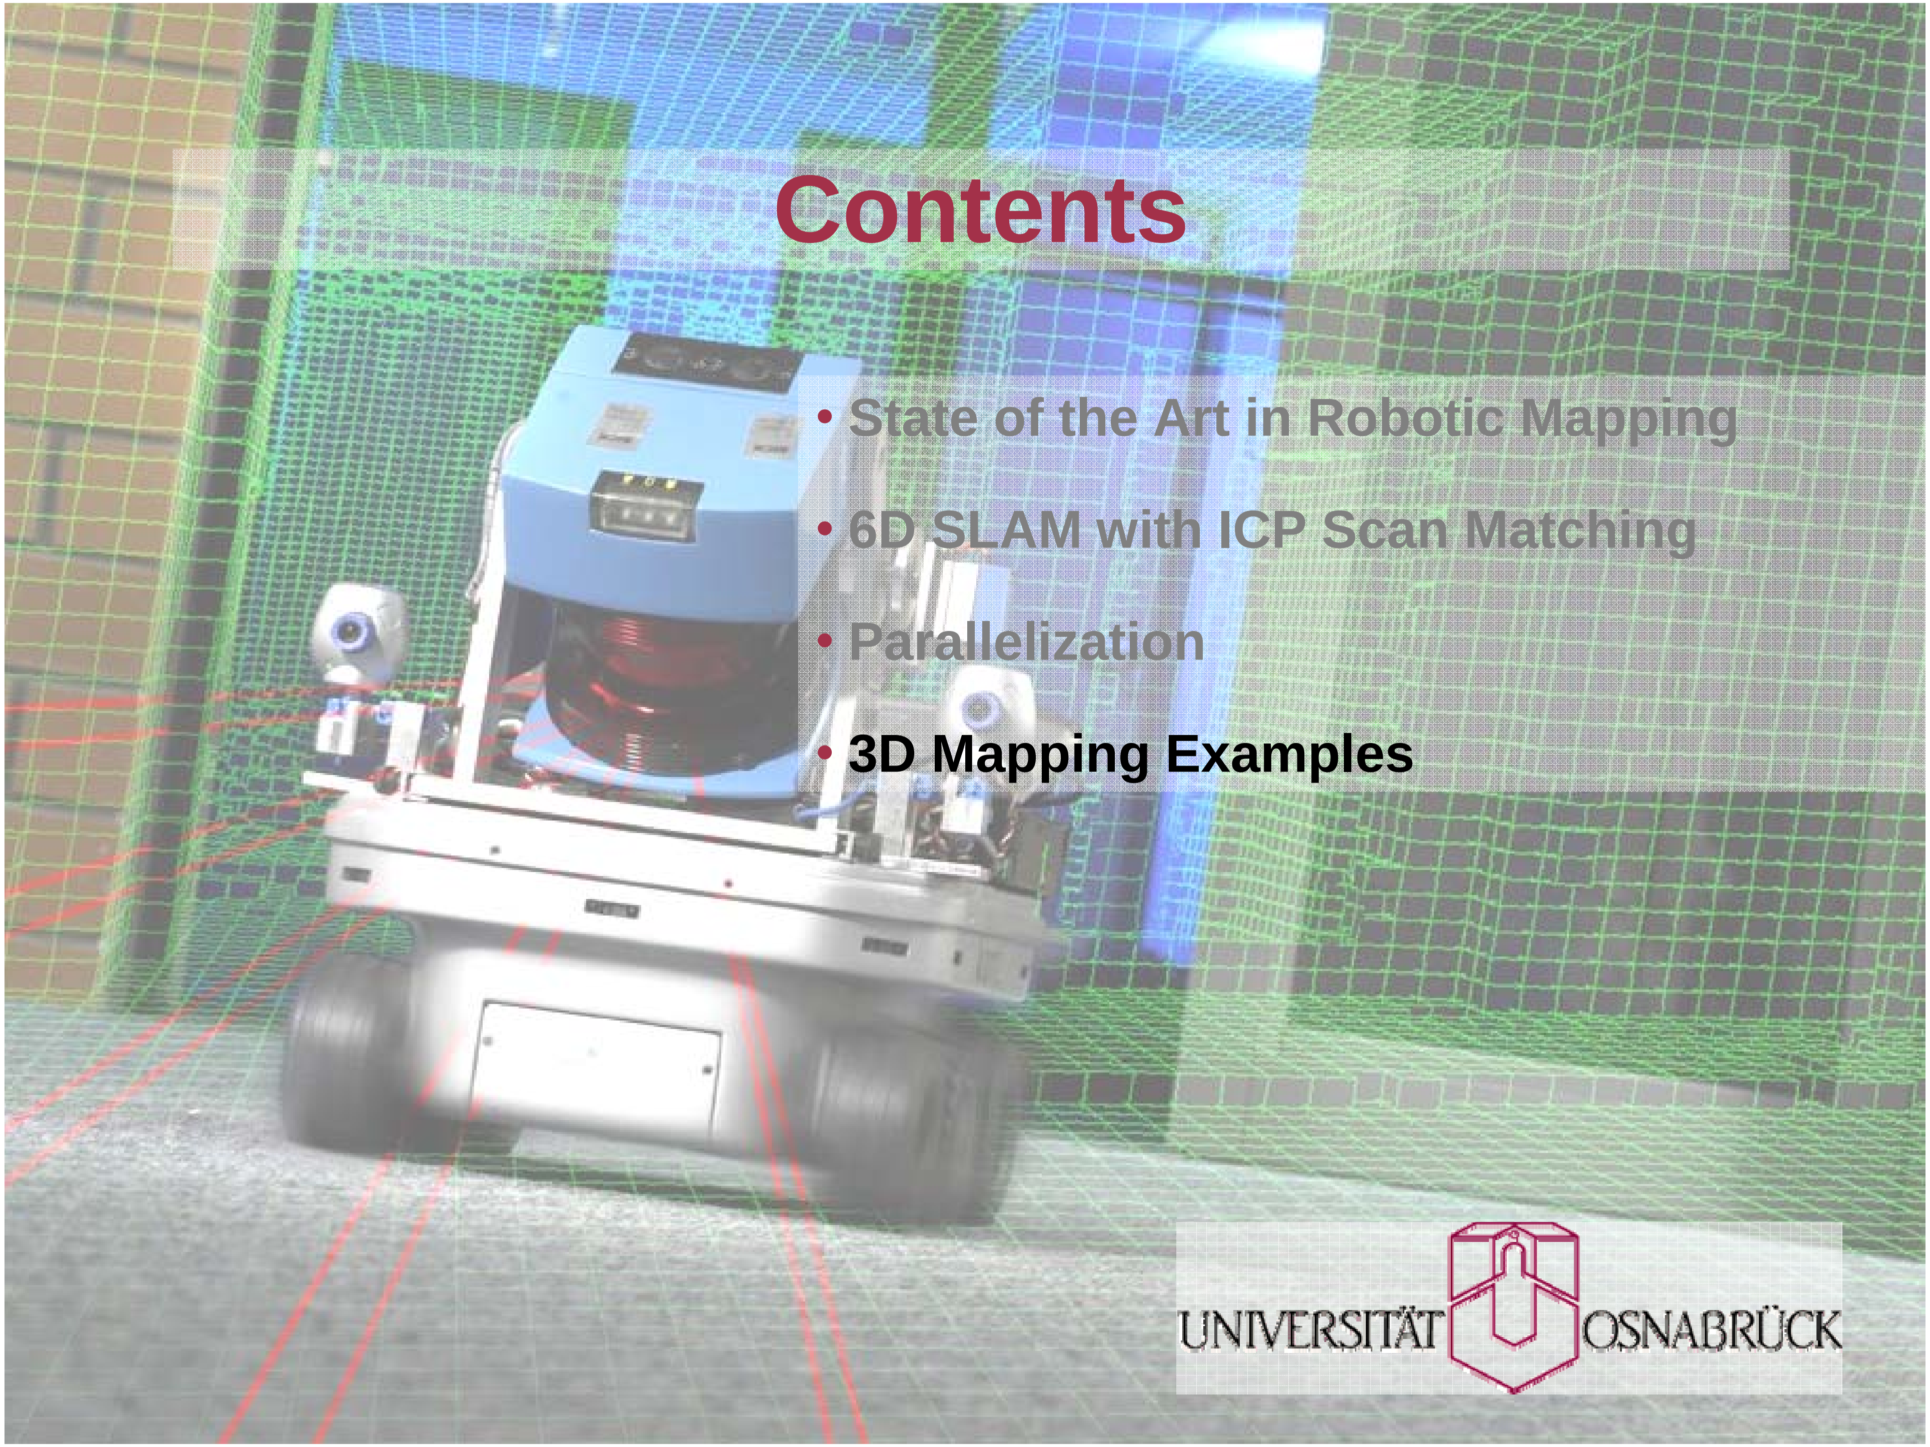

Contents

- State of the Art in Robotic Mapping
- 6D SLAM with ICP Scan Matching
- Parallelization
- **3D Mapping Examples**

UNIVERSITÄT OSNABRÜCK

3D Mapping of Urban Environments

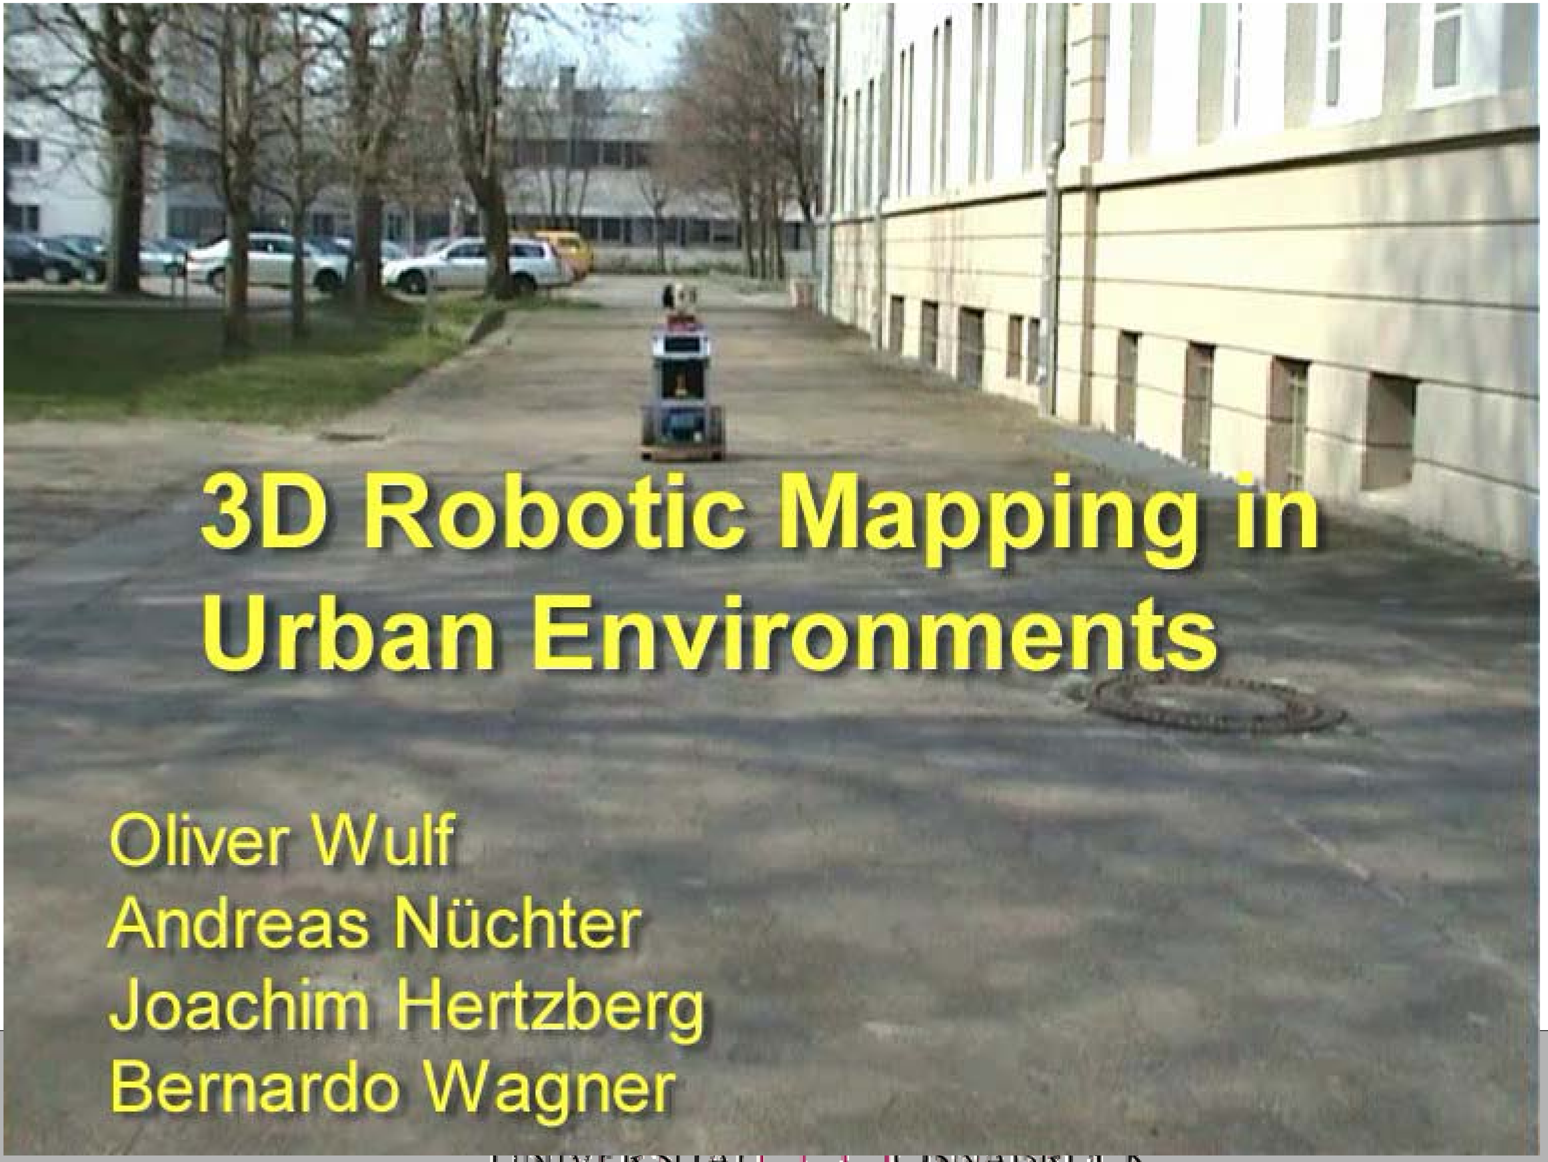

3D Mapping of Urban Environments

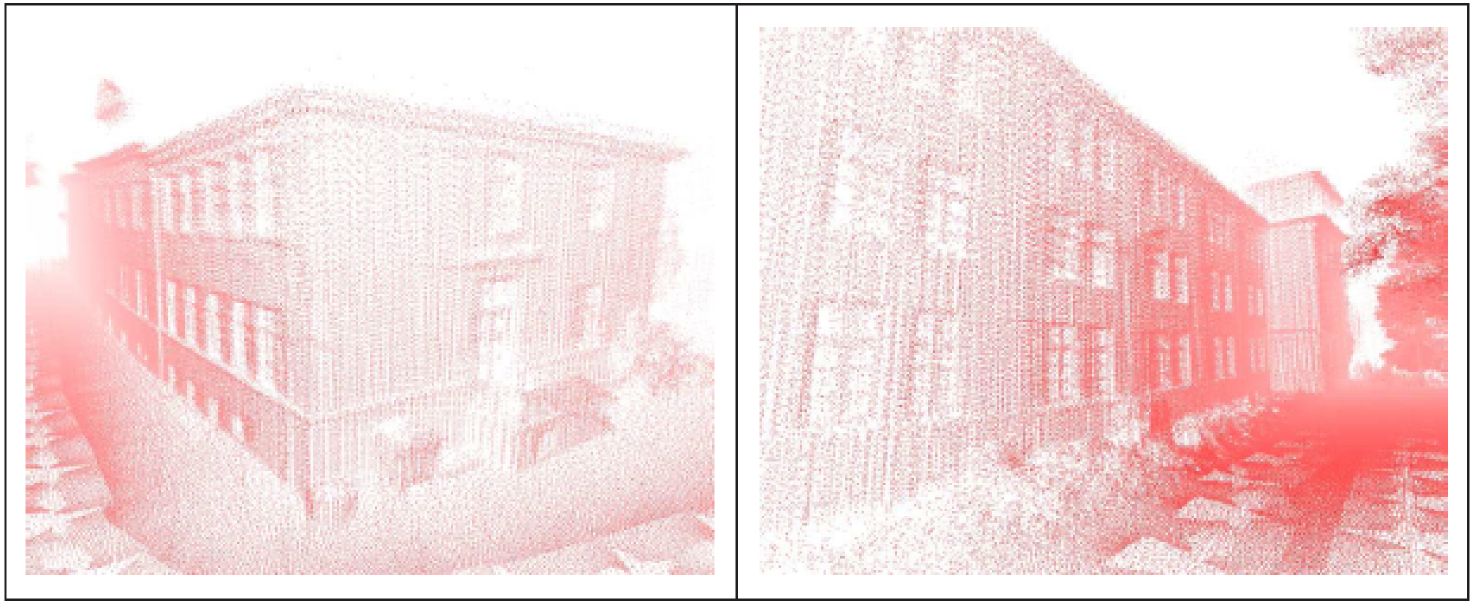

We acknowledge Oliver Wulf and Bernardo Wagner (Leibniz Universität Hannover) for providing the data set.

(Video)

3D Mapping of Urban Environments

We acknowledge Nikolaus Studnicka (RIEGL Laser measurement Systems GmbH, Horn) for providing the data set.

(Video courtesy Riegl)