On Correcting Sewer Robots' Odometry Errors by Reasoning

Joachim Hertzberg, Frank Schönherr

GMD-SET, Schloss Birlinghoven, 53754 Sankt Augustin, Germany,

hertzberg|schoenherr@gmd.de

Abstract

When inspecting sewers, it is required that discovered objects like damages or previously unknown laterals are entered into the existing sewer map at the correct metric position. On the other hand, odometry conditions in sewers can be very poor, and they may vary greatly with slip, slope, and sewage current. The paper describes a simple reasoning method for correcting the odometry error according to local conditions. The method is based on correcting past measurements, and hence the corresponding new map entries, according to safe localization points at individual landmarks and to a model of the local odometry error or errors in between the localization points. A sketch of the intended application of the method in an autonomous, multi-segment, articulated sewer robot is given.

1   THE APPLICATION PROBLEM

Autonomous mobile robots live in space, must cope with space, and hence notoriously reason about space. Examples of such reasoning are path planning, trajectory planning, and navigation under robots' various perception impairments. In this paper, we are dealing with *localization,* i.e., the question: Given a map of the environment, where exactly am I? Some applications allow this problem to be solved satisfactorily by plastering the environment with beacons or by using external references like the Global Positioning System (GPS). Some applications don't: GPS signals, for example, are shielded in most buildings, and if the purpose of some service robot is to work in inaccessible areas, then installing beacons may be impossible or cause high cost. Much research is currently done in localization methods for mobile robots; we will address some related work *en passant.*

Our localization problem is somewhat particular, according to the particular application area: Sewers. The long-term background is the intention to build autonomous mobile robots for continuously inspecting inaccessible (i.e., < 1 m diameter) sewers in order to record problems like damages or hazardous substances. We do not further expose this area and the rationale for using autonomous robots; see [Hertzberg et al., 1998a] for an overview.

The domain characteristics of interest for this paper can be explained easily. A metric map exists for every sewer, specifying, among other things, manholes, main pipes and laterals. The map is correct with respect to manholes, mains, and most laterals. During inspection, a number of interesting features like damages (cracks, occlusions, grown-in tree roots) or other items (unrecorded laterals) have to be entered into the map. Our users require measured positions to be exact within a range of 10 cm. Sensor conditions for autonomous sewer robots are poor. Sewers are cylindric, narrow, dark, dirty, slippery, and wet. Pipes vary greatly in diameter, material and wearout. Odometry is particularly difficult because of the much varying, sometimes very high, degrees of slip, slope, and sewage current.

Conventional sewer inspection uses tethered camera platforms, where position is determined in the archaic (and somewhat error-prone) way of measuring the cable as it gets unreeled. Autonomous robots must do without, and so the intuitive problem is:

How can an autonomous sewer robot know with sufficient metric precision where it is when spotting a reportable feature, to enter it correctly into the map?

(1)

Our own previous work for topologically correct positioning [Hertzberg & Kirchner, 1996] and metrically correct map entries [Schönherr, Hertzberg, & Burgard, 1998] has made some very weak assumptions about the correctness of turning maneu-

vers in the sewer and about the correctness of object detection and classification; in consequence, we had to use probabilistic (POMDP) methods for position estimation that resemble those used, e.g., in hallways, where translation imprecision, drift, and rotation imprecision would blur the position estimation (e.g., [Koenig & Simmons, 1998]). For a new sewer robot platform, which we are developing in a joint research project [Cordes et al., 1997; Hertzberg et al., 1998b], turning errors are no longer likely to occur. As rotation and lateral drift are physically impossible in pipes, a simple localization method may be used that compensates more effectively for the translational localization error.

The next section describes the basic idea of the method. Sec. 3 sketches refinements of the model of the locally varying odometry error. Sec. 4 concludes.

2 THE BASIC METHOD

As in our previous probabilistic approaches [Hertzberg & Kirchner, 1996; Schönherr, Hertzberg, & Burgard, 1998], we assume that manholes in the sewer are detected safely and that they figure metrically (and therefore, topologically) correctly in the sewer map. This assumption is warranted in practice because manholes are the part of a sewer that can be seen and measured from street level, so they are easily accessible and known. Detection of manholes is safe because the sensoric change from a narrow pipe into a wide-open shaft just cannot be overlooked.

So a metric position error due to odometry errors occurs only in the passage through the pipe from manhole to manhole. To give an idea of its order of magnitude, experiments with the sewer robot platform KURT ([Kirchner & Hertzberg, 1997], Fig. 1(a)) with a very simple, calibrated odometry in a dry test sewer network have yielded a standard deviation of the odometry error of 17%, with outlyers beyond 30% occurring frequently [Schönherr, Hertzberg, & Burgard, 1998]. Figure 1(b) gives an idea of how wrong position estimation would be if it would rely exclusively on odometry under this error quantity.

Due to KURT's simple odometry (which is essentially dead-reckoning), this error is certainly extreme and could be reduced by technical means. However, according to our experience, there seems to be no practical and affordable way of making sewer robot odometry really precise within the required range: 10 cm error in typical distances of up to 50 m, which is 0.2%.

In consequence, we do not tackle directly the problem (1) as stated, but do a little reasoning instead: we live with an odometry error and correct *ex post*
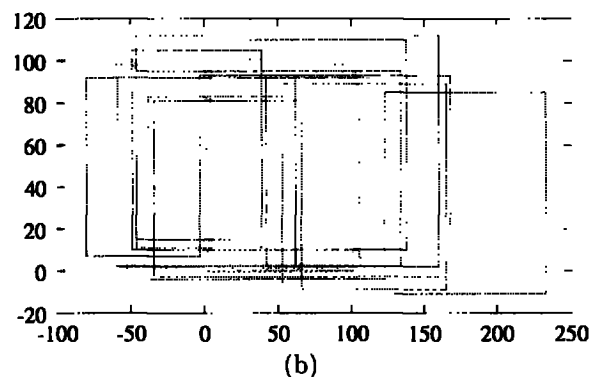

(a)

(b)

Figure 1: The sewer test platform KURT [Kirchner & Hertzberg, 1997] turning in a manhole of a dry test sewer (a), and error of calibrated odometry of KURT in a tour through an 8-shaped part (grey shade) of the test sewer (b). All turns in the tour were performed correctly. Distances given in decimeters; the tour started at the (0,0) coordinate.

the position information acquired between two places whose metric positions are known, i.e., between two manholes. The simplest way of doing this assumes that the odometry error is equally distributed over the just-performed passage from manhole to manhole: Compare the measured length of travel with the known pipe length; propagate the resulting overall error back to measured positions, correcting the position by the overall error percentage. Table 1 presents some sample data of single pipe runs of KURT taken from those used to produce Fig. 1, where the respective pipes contained inlets that KURT detected correctly. Trace no. 33 is an example of a typical "good" trial, the positioning error being reduced to 14 cm in 500 cm, or 2.8%. Trace 35 is among our best performance examples, with the error reduced from 46.8% to 0.6%. The explanation for the large reduction is simple. During the pass through the pipe, a huge

Trace	Pipe	Pipe length		Inlet position				
					Measured		Corrected	
		Actual	Measured	Actual	Pos	% Error	Pos	% Error
33	5 (S→N)	1000	903	500	439	12.2	486	2.8
35	1 (N→S)	900	1312	800	1174	46.8	805	0.6
28	5 (N→S)	1000	948	500	499	0.2	525	5.2
74	8 (W→O)	1000	860	500	370	26.0	430	14.0

Table 1: Samples of data from runs with KURT through test sewer pipes containing inlets. Inlet positions are corrected by the overall error in pipe length. All metric data in cm.

odometry error has accumulated (probably due to the high frequency of tilt correction maneuvers), blurring completely the positioning. However, as the inlet is quite close (100 cm) to the pipe end, the odometry error accumulated until meeting the inlet is mostly identical to the overall error, which gets washed out by the normalization. This effect is typical for the technique of backpropagating the correct position information.

The basic correction method is a heuristic, and heuristics may fail. Trace 28 is an example, where a basically correct value gets blurred. However, the overall performance is a success, given its utter simplicity. Averaging over 17 passes through pipes, the mean error percentage dropped from 16.6% to 7.5%. The performance is overshadowed by mediocre runs like no. 74, which shows a clear improvement, but is still bad beyond tolerability.

The basic position correction idea is borrowed from Hidden Markov Model (HMM, [Rabiner & Juang, 1986]) localization approaches as in [Thrun, Fox, & Burgard, 1998; Schönherr, Hertzberg, & Burgard, 1998], where secure information about the recent place is propagated back into past position estimations and into the own state transition model. The approach presented here is simpler, owing to stronger assumptions (correct object recognition) that it makes about the robot in its environment. In particular, we do not use discrete positions and do not have to cope with the probabilistic model of action outcomes. In consequence, we have to maintain no probability distribution over discrete object positions, but get along with one maximally likely position per detected object.

By its nature, object mapping has a flavor of map learning (e.g., [Thrun, Fox, & Burgard, 1998]), yet is more simple and more specialized owing to the prior ground map information that we provide. Like in our previous work [Schönherr, Hertzberg, & Burgard, 1998], we propose to take the cautious approach of keeping the ground map as is. During an inspection

run, we produce an overlay of this map containing detected objects and their corrected positions, which is to be reported as the inspection protocol. Whether or not something of this information should lead to updating the ground map is not at a sewer robot's discretion.

3 TOWARDS REFINING THE ODOMETRY ERROR MODEL

The basic method is suitable for filtering out systematic odometry errors that occur constantly over the passage between two calibration points. Sewage current, slope, and material conditions are examples for error sources in a sewer, which can be assumed to remain constant between two manholes by construction regulations. But there is more, as demonstrated by the residual error after correction in Table 1.

Accelerations and decelerations, stops and goes are irregular sources of odometry error that cannot be expected to be evenly distributed over a pipe. What is worse, they may happen at any time in passing a pipe: To prevent toppling over, a sewer robot has to correct tilt, which results in discontinuities in the driving pattern. For small corrections, driving a slight curve may suffice; larger corrections may involve a full stop and pull-back maneuver. Odometry is impaired in either case.

Exactly how far this goes cannot be said in general. It depends on the individual robot's mechanics and odometry. The point is: The overall correction of the odometry values as in Sec. 2 must get corrected by a set of error terms that depend on special driving maneuvers performed between the start of the recent pipe passage, the to-be-measured object sighting(s), and the end of the pipe passage. So the overall error err between two calibration points (manholes) is a composition of the constant odometry error err_c occurring over the whole passage and a sum of individual error terms corresponding to the types of the

sequence of correction maneuvers $m^{(i)}$ performed over the passage:

$$err = err_c + \sum_i err_{m^{(i)}}$$

After arriving at a calibration point, err is known as the difference between the nominal and the factual odometry value.

A completed trace T of the recent passage has the structure

$$T = \langle m^{(i)}, o^{(j)}, err, l \rangle$$

where $o^{(j)}$ are the object sightings including the measured odometry values, the $m^{(i)}$ also include the respective odometry values, and l is the known true length of the passage. To estimate the true object positions, we have to estimate the posterior values of err_c and the $err_{m^{(i)}}$, given err; calculating the position estimations of the true $o^{(j)}$ from the trace is then straightforward.

We propose to approximate the error err_m induced by a correction maneuver m as depending at most on the information contained in the recent trace T, i.e.,

$$err_m = f_m(T)$$

for some function f_m that has to be determined. As mentioned above, the most odometry-relevant environment parameters current, slip, and slope can be assumed to be constant over a pipe connecting neighboring manholes, i.e., over the passage corresponding to some T. Assuming that under these homogeneity conditions, two maneuvers of the same type would yield the same error, it would even suffice for a first approximation to consider of T only the relative error err/l for defining err_m. Moreover, there are correction maneuvers that can be approximated by a constant perturbation f_m; the path curvature resulting from correcting a light tilt is an example.

The concrete set of correction maneuvers as well as the f_m contribution of each maneuver depends on the concrete robot mechanics, kinematics, sensor equipment, and motor control. It has to be determined empirically. For the KURT platform, for example, the following maneuver types are candidates for modeling:

Stop-deliberately: The robot stops in the pipe (e.g., to check an observation).

Accelerate-deliberately: After stopping, the robot resumes its passage.

Turn-to-correct-tilt: The robot initiates to correct a light tilt by a light turn. When in an appropriate position, it will later resume normal straight passage.

Stop-to-correct-tilt: In a position of dangerous tilt, the robot performs a full stop.

Pull-back-and-stop: After the previous maneuver, the robot sets back (possibly turning lightly against the tilt) until it hits the pipe bottom, where is stops to resume its passage with an accelerate-deliberately maneuver.

At the time of finishing this paper, we are not yet able to provide the f_m estimations for these maneuvers in KURT. In particular, the traces used in Sec. 2 do not contain all relevant information about correction maneuvers, as they were recorded without the idea in mind of considering them. So this is an obvious point for immediate future work.

4 CONCLUSION

The contribution of this paper is the method for "soft" correction of possibly severe odometry errors under certain conditions, which concern the environment, the ground map, and the reliability of landmark and object detection and turning maneuvers. The method promises to be of much practical value for sewer inspection, but can be transferred to other inspection domains with like assumptions and requirements. Historically, it has emerged from an HMM mapping method [Schönherr, Hertzberg, & Burgard, 1998], but is sufficiently simple to be understood without this background. Localization methods based on a grid representation of space—metrical or topological—like, e.g., [Koenig & Simmons, 1998; Thrun et al., 1998], are not well suited to achieve the metric precision of positioning observations as required for our application: On the one hand, using a grid of sufficiently fine grain would yield huge cell sets and impractical representations of motion imprecisions in different environment contexts; on the other hand, this motion imprecision of sewer robots can be handled more easily. Note that we can assume observation certainty concerning manhole detection, i.e., concerning possible turning points, so that our application problem lacks another important reason for using grid-based probabilistic localization approaches.

Odometry errors can certainly be reduced by constructive means on the robot's side, and for practical purposes, all cheap and simple such means for reducing it in the first place should be applied. In particular, a more sophisticated, multi-segment articulated sewer robot platform, which we are currently developing [Cordes et al., 1997; Hertzberg et al., 1998b], promises to yield much more exact odometry values owing to the high degree of redundancy of odometry measurements that its multi-segment architecture fa-

cilitates. However, our experience tells there is few hope to reduce it down to zero. Instead of using sophisticated technical equipment, we prefer to use simple and cheap algorithms for handling the problem. Enhancements of the basic method as described in Sec. 3 are possible and necessary with respect to approximating odometry errors induced by special driving maneuvers in concrete robots.

References

[Cordes et al., 1997] Cordes, S.; Berns, K.; Eberl, M.; Ilg, W.; and Suna, R. 1997. Autonomous sewer inspection with a wheeled, multiarticulated robot. *Robotics and Autonomous Systems* 21:123-135.

[Hertzberg & Kirchner, 1996] Hertzberg, J., and Kirchner, F. 1996. Landmark-based autonomous navigation in sewerage pipes. In *Proc. 1st Euromicro Workshop on Advanced Mobile Robots* (EUROBOT'96), 68-73. IEEE Press.

[Hertzberg et al., 1998a] Hertzberg, J.; Christaller, T.; Kirchner, F.; Licht, U.; and Rome, E. 1998a. Sewer robotics. In Pfeifer, R.; Blumberg, B.; Meyer, J.-A.; and Wilson, S., eds., *From Animals to Animats 5 – Proc. 5th Intl. Conf. on Simulation of Adaptive Behavior (SAB-98)*, 427–436. MIT Press.

[Hertzberg et al., 1998b] Hertzberg, J.; Christaller, T.; Kirchner, F.; Licht, U.; and Rome, E. 1998b. MAKRO – Serviceroboter im Kanalbetrieb. In Lenz, J., ed., *Globalisierung der Märkte und internationale Arbeitsteilung – auch im Rohrleitungsbau?*, 643-652. Essen: Vulkan-Verlag.

[Kirchner & Hertzberg, 1997] Kirchner, F., and Hertzberg, J. 1997. A prototype study of an autonomous robot platform for sewerage system maintenance. *J. Autonomous Robots* 4(4):319–331.

[Koenig & Simmons, 1998] Koenig, S., and Simmons, R. 1998. Xavier: A robot navigation architecture based on partially observable Markov decision process models. In *[Kortenkamp, Bonasso, & Murphy, 1998]*. AAAI Press/The MIT Press. chapter 4, 91-124.

[Kortenkamp, Bonasso, & Murphy, 1998] Kortenkamp, D.; Bonasso, R.; and Murphy, R., eds. 1998. *Artificial Intelligence and Mobile Robots*. Menlo Park: AAAI Press/The MIT Press.

[Rabiner & Juang, 1986] Rabiner, L., and Juang, B. 1986. An introduction to Hidden Markov Models. *IEEE ASSP Magazine* 1:4-16.

[Schönherr, Hertzberg, & Burgard, 1998] Schönherr, F.; Hertzberg, J.; and Burgard, W. 1998. Probabilistic mapping of unexpected objects by a mobile robot. In Orphanoudakis, S.; Trahanias, P.; Crowley, J.; and Katevas, N., eds., *Proc. Computer Vision and Mobile Robots Workshop (CVMR-98)*, 46 53.

[Thrun et al., 1998] Thrun, S.; Bücken, A.; Burgard, W.; Fox, D.; Fröhlinghaus, T.; Hennig, D.; Hoffmann, T.; Krell, M.; and Schmidt, T. 1998. Map learning and high-speed navigation in RHINO. In *[Kortenkamp, Bonasso, & Murphy, 1998]*. AAAI Press/The MIT Press. chapter 1, 21-52.

[Thrun, Fox, & Burgard, 1998] Thrun, S.; Fox, D.; and Burgard, W. 1998. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning* 31:29-53. Joint issue with J. Autonomous Robots, 5:253-271.