Algorithms for the interpretation of continuous measurement of the slurry level in storage tanks

Fabrizio Mazzetto,[1] Pasqualina Sacco,[2] Aldo Calcante[3]

[1]Faculty of Science and Technology, Free University of Bolzano; [2]TerritoriumOnLine srl, Bolzano; [3]Department of Agricultural Engineering, University of Milano, Italy

Abstract

This paper describes a device for monitoring the slurry level in storage tanks. In particular, it describes the inference engine used within the StorEyes software. This software was developed by the Authors to interpret and manage data collected from measurements taken by ultrasonic or pressure sensors mounted on slurry tanks for continuous monitoring of the stored effluent level. The monitoring device consists of a data-logger that uses a GPRS device for the wireless transmission of data to a Farm Information System server, a buffer for the temporary storage of data in case of malfunction of the farm network, and a GPS receiver. The inference engine consists of a set of algorithms used to derive from raw measurements (usually in V) intelligible information about the loading/unloading events in storage facilities. The inference engine has been tested and calibrated through two years of acquisitions at two pilot farms with different slurry tank loading and unloading management systems.

Introduction

Since the 1960s, the continuous expansion of the livestock sector has led to the development of intensive livestock farms, characterized by a high number of cattle per hectare and a high consumption of extra-farm provenience forage.

This has inevitably led to an excessive production of zootechnical effluent compared to the available farmland, thus transforming a primary source of nutrients for crops in a low-value material, to be disposed of at the lowest possible cost (Mazzetto, 2009a). As a result, incorrect management has led to a distribution of effluent in the field that is not always sustainable and appropriate from an agronomical point of view. The main consequence of this incorrect effluent distribution is the pollution of surface water and groundwater, due to an excess of nitrates (Sommer, 2004). As confirmation of this, the European Environment Agency (EEA) has estimated that Italian agriculture accounts for over 60% of the release of nitrogen into surface waters and it has identified a close correlation between nitrate concentration and the intensity of national agricultural activity (Provolo, 2008). This situation has forced public administrations to put in place a strict preventive monitoring system of effluent distribution based on the subdivision of the territory into areas characterized by different levels of vulnerability. The first initiative of this kind was the Nitrates Directive (EEC, 1991) and the subsequent national and regional laws.

Following completion of the Directive in 2006, the Regione Lombardia (2007) passed a Regional Resolution, the Nitrate Directive, to extend the areas of vulnerability to a large part of the Po Valley and to the flood zones, as required by the hydro-geological plan for the Po River Basin.

In addition to drawing up the Operative Plans (POA/POAs) and the annual communication of PUAs/PUA by the farmers that were already laid down by the previous regulations, the 2007 Resolution provides for self-monitoring to verify the effective application of the aforesaid Operative Plans. A penalty system was put in place to deal with cases of misconduct on the part of the farmers. The framework highlights the need for farmers to have a monitoring system on their farms that allows for the automated collection of data related to the management of livestock effluents and the subsequent storage of the data in suitable databases. To do this, farmers must adopt a Farm Information System (FIS) able to store and process data related to the handling of zoo-technical effluents on the farm (McGechan, 2000).

The METAMORFOSI project, created by the Department of Agricultural Engineering of Milan, aims to automatically produce reports on slurry management using FISs and Territorial Information Systems (TISs) that are connected by wireless networks.

In particular, StorEyes is a software program developed within the framework of the METAMORFOSI Project (Mazzetto, 2009b,c) that handles the control of nitrogen flows in a specific territory, starting with monitoring the management of the animal effluents on the farm. In compliance with the Nitrate Directive, the overall project deals with the automatic monitoring on individual farms of both the storage facilities for zoo-technical effluents and their related spreading activities. In this way, extending the system to a congruous number of farms within a single area, it is possible: i) to control the potential entry of nitrogen from zoo-technical effluents in that territory through monitoring the storage facilities; ii) associate the withdrawals with the areas where the effluents are spread through the monitoring of the spreading machinery used.

The project has, therefore, identified methods and related IT tools that are capable of providing continuous monitoring of both the amount of effluent present at any moment in the farm storage facilities and the dynamics of the spreading machines. Nowadays, this is possible thanks to the remarkable advances achieved by the application of wireless data transmission networks in agriculture (McKinion, 2004; Wang, 2006). Complete and comprehensive control is only achieved when it is possible to integrate information obtained independently through these two monitoring initiatives.

In particular, the potential for constantly recording the volume of effluents present at each individual farm allows a very realistic estimate to be made of the nitrogen flows which enter and leave the territory, and which are related to the volumes of incoming and outgoing effluents at the individual farm storage points. Obviously, this is much simpler if it involves liquid or semi-liquid effluents (slurry waste); these are exactly the types of effluents that have been considered for use with the monitoring system here described.

A continuous measurement of the slurry level in livestock storage tanks also makes management of effluents easier in terms of planning regarding spreading or tank loading/unloading. To ensure precise control, a data acquisition frequency of at least one measurement per minute must be ensured. This is usually achieved through ultrasonic or pressure sensors appropriately installed in each storage tank. This frequency generates a relevant amount of data that will then be continuously processed to automatically interpret and register all the operational events performed by the farmer in each tank. For both practical and analytical reasons this cannot be done manually: in fact, variations in levels are often strongly affected by a high measurement noise that can hide actual events or propose false ones.

In addition, the final slurry mass balance over a pre-defined time period (even in terms of kg of N) must be derived from those measurements as accurately as possible.

The StorEyes was used to provide all the analytical tools required to interpret these measurements. Therefore, the package provides the answers to all possible queries on the historical events (or N flows) occurring in the farm tanks. It also includes, first, an *inference engine* to transform raw data into intelligible management information.

From this point of view, StorEyes can be said to belong to the family of *inventory management programs*.

Materials and methods

Monitoring systems for effluent storages

A complete slurry waste tank monitoring system consists of: i) a data-logger device to be mounted directly on the tank, also including sensors to provide slurry waste level readings through different types of measurements; ii) a set of computing and *inferring procedures* used to produce information from the raw data collected; and iii) a user interface to facilitate access and the use of information in the control activities that are related to the management decision-making processes.

The data-logger (Car Securer device 4.00, Media Systems Ltd., Bulgaria) is connected to the sensor to measure the level of slurry. It consists of: i) a data recorder (two analog and four digital inputs); ii) a GPRS device for the wireless transmission of data to the FIS server; iii) a buffer for the temporary storage of data in case of malfunction of the farm network; iv) a GPS receiver. In particular, the GSM module allows: i) the connection via GPRS (max. speed 64 kbit/s) to the farm server; and ii) communication with the user via SMS to fulfill specific data requests or to signal possible malfunction.

The GPS receiver is only used to synchronize the data recorded through the internal clock. In this case, measurements of the position are, in fact, unnecessary.

The characteristics of the monitoring system depend on:
- the *type of storage tanks* (above ground tanks or tanks placed under the slatted floor) and their spatial location (this affects the number of sensors that can be managed by each data-logger);
- the t*ype of sensor employed* (ultrasonic or hydrostatic sensor).

The data-logger may be equipped with two types of sensors (Figure 1) depending on whether they measure: i) pressure by means of sensors (Jumo, dTrans P90, Germany, output 0-10V) to be immersed in the hydrostatic fluid to be measured so that the hydrostatic pressure acting on a sensitive membrane generates a voltage signal that is proportional to the effluent level (L, in *m*); and ii) *distance* (the distance between the sensor and the slurry surface) by means of acoustic sensors (Sick, UM 30, Germany, output 0-10V) placed above the free surface of the slurry, and emitting an ultrasonic pulse that enables the flight-time to be measured once the pulse is reflected by the slurry surface; this kind of US-sensor, therefore, generates a voltage signal proportional to the distance D (*m*) to be measured. As the distance HS is known, the final level is then obtained by the formula L = HS – D (Figure 1).

Both types of sensors require a preliminary on-site calibration. Logged data are continuously and automatically retrieved (a single measure per minute) via a wireless GPRS-transmission and transmitted to a central server with StorEyes installed, together with the databases (DB) relating to the basic *farm resources* (Res-DB = plant, machines, farm structures and land, data-loggers) and *recorded activities* (Tank-DB and Operation-DB). It checks the presence of new measured data once a day. These data are then treated and the following computations are then performed:

measured data A) → raw data B) → inferred data C) → use of information.

A procedures convert and filter data logged by sensors into intelligible (raw) data, *i.e.* the calibration equations are used first to convert electric measurements into level measurements that are then aggregated and cleaned up as much as possible to remove all forms of noise prior to being definitively expressed in terms of effluent volumes; to do this StorEyes recalls information on tank features from Res-DB.

B procedures then perform inferences to provide final average hourly volumes of slurry waste contained in the tanks and to identify effluent uploading and downloading events. Such events can later be confirmed, or not, by the user.

C procedures enable the farmer to access the inferred data both in table and graphic forms (volume *vs* time diagrams, at daily, weekly or monthly scales; Gantt diagrams of events). The results of the inference procedures that have to be kept are stored in the Inf-DB. A set of queries is also provided to allow for a large series of surveys and to investigate the above mentioned links between Tank-DB and Operation-DB events.

The possible settings for the use of the information collected from monitoring of the storage facilities for the zoo-technical effluents can
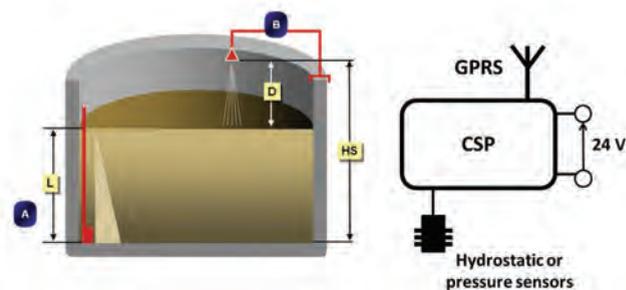

Figure 1. Types of sensors used to monitor the slurry level in storage tanks. (A) Pressure gauges (hydrostatic sensors). (B) Distance gauges (ultrasonic sensors) and diagram of the data-logging device for monitoring slurry level in storage tanks.

be found both on the farm and throughout the territory. For the moment, the principal analyses and assessments that can be performed with StorEyes are targeted at farm use.

The inference engine: an overview

The B procedures transform raw data into intelligible data. Therefore, they only represent the inference engine of this application.

The inference engine performs a sort of day by day interpolation process treating each *daily raw data set* separately. The scope is to reconstruct a simplified profile of the daily level pattern that includes only a set of *stabilized levels*, also called *benchmark levels*. Such benchmarks include a set of some level-values against which the measured raw data are linkable with greater frequency; less than an acceptable error variance. So, benchmarks aim to eliminate measurement noise, proposing stable fitted levels with changes in discontinuous steps only in the case of occurrence of loading/unloading events. In other words, *a change in a benchmark value can be caused only by a discrete, sudden change in the actual effluent levels.* Therefore, the daily benchmark analysis only allows discontinuous events to be identified. A multiple-day analysis must also be performed to be able to identify slow, progressive variations in level values.

The *single-day benchmark interpolation* is not carried out through a conventional statistic regression; in fact, this cannot meet the need to identify loading/unloading events. It is rather based upon the following calculation steps:

S1. identification of *all the possible* daily benchmarks, even filtered with respect to the need to obtain stable profiles over time;

S2. *replacement of each raw datum with its closest benchmark* (identifying preliminary discontinuities in the daily benchmark pattern); and

S3. application of a *final filter* to confirm (or reject) the previous replacement on the base of a deviation test.

Discontinuity points are then forced to correspond to storage *loading or unloading events*, *i.e.* to those processes that can be associated to injection or withdrawal activities on slurry storage, respectively (Figure 2). The daily analysis also identifies *presumed events*. These could occur during possible prolonged recording interruptions along the day, resulting in some holes in the availability of raw data over time.
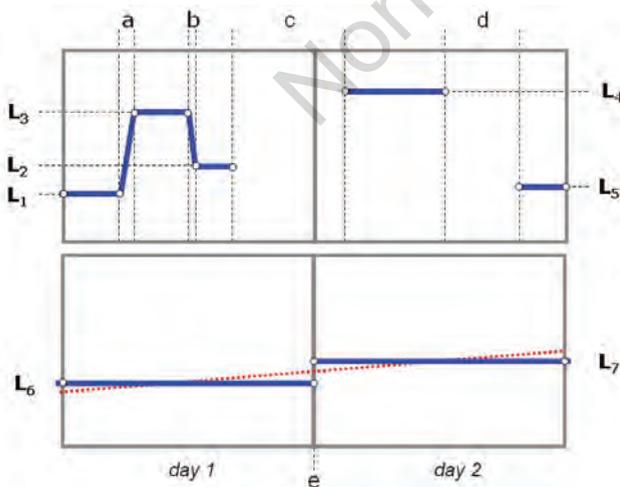

Figure 2. Types of events occurring on the inferred level values of an effluent storage. Values from L_1 to L_7 are benchmark levels. a, loading; b, unloading; c, compensative, presumed loading; d, presumed unloading; e, compensative loading.

The subsequent *multiple-days analysis* aims to identify the so-called *compensatory events*, *i.e.* an event that causes a slow and gradual change in the tank level, not identifiable by instantaneous events. A compensatory event, therefore, may be due either to activities related to slow and continuous load fills (*e.g.* when the load is the result of the daily periodic cleaning of the aisle housing), or to various phenomena not directly linked to storage management activities, but rather to random disturbing events (gas production by fermentation, convection, growth of weeds on the effluent surface, etc.).

The interruption of logging that occurs between two or more consecutive days is also treated as a compensatory event.

Details on the algorithms

The goal of the inference engine is to replace a sequence-set G_0 of original raw data (available over a period of D days) with a new sequence set E of revised values according to which the storage management events can be inferred more clearly. A first *single-day analysis* is performed iterati-vely for each day $d \in D$ on a subset $G_{0d} \subseteq G_0$:

$$G_{0d} = \{ g_{0dj} \mid d \in D \}$$

$$j = 1..N_J, \quad \text{with} \quad j \in N \text{ and } g_{0dj} \in R$$

The whole computation sequence foresees the steps mentioned above:

$$G_{0d} \xrightarrow{S1} C_d \xrightarrow{S2} E_{ad} \xrightarrow{S3} E_d$$

(being $E_d \subseteq E$, as well)

(*Daily raw data → benchmark levels → interpreted data → inferred data*).

All the elements of both G_0 and E (and related subsets) are associated with time instants. As the system works on a discrete time basis, the timeline set T_d is simply expressed as:

$$T_d = \{ j \mid j \leq N_J = 1440 \cdot \phi \}, \quad \text{with} \quad j \in N$$

where N_J = maximum number of daily measurements, and ϕ = frequency (*min^{-1}*). The experience of the 2-year project period highlighted that a frequency $\phi = 1$ min $^{-1}$ is an appropriate compromise to obtain satisfactory inferences without an excessive redundancy of data (N_J = 1440). S1-step performs the construction of the C_d set that is usually composed of a small number of elements, such that:

$$C_d = \{ c_{dr} \mid d \in D, c_{di} \in G'_{0d} \},$$

$$r = 1.. N_C, \quad \text{with} \quad r \in N \text{ and } N_C << N_J.$$

In short: i) each element c_{dr} identifies a *benchmark* storage level; ii) the number of elements of N_c is generally very low depending ultimately on the number of management events that cause a change in level during the day; iii) the elements c_{dr} are not associated with time instants (their link to T_d will be performed in the following step); and iv) the G_{0d} set is empty only when a complete blackout in the recording of data occurs. Of course, in such situations, it is not possible to apply any form of inference; therefore, it will be $C_d = \{0\}$.

In case of partial blackout days (*i.e.* with still enough data to trace a partial storage level pattern), and in those in which no loading or unloading events occur, the inference calculates at least the benchmark-level c_{d1}, that can, therefore, be seen as a sort of *average daily level.*

The construction algorithm C_d performs a time-independent calculation based on the search for levels which occur with high frequencies

over the day, together with the application of some filters for removing outliers. The latter are first identified through the following test applied between two adjacent measurements over time:

$$\text{Hyp:} \quad \left| g_{0d,j} - g_{0d,j-1} \right| > k \cdot \Delta L_{REF}$$

$$\text{Hyp:} \quad \begin{cases} \text{True} \implies g'_{0d,j} = R_L \cdot \text{round}(g_{0d,j-1}/R_L) \\ \text{False} \implies g'_{0d,j} = R_L \cdot \text{round}(g_{0d,j}/R_L) \end{cases}$$

$$\Delta L_{REF} = 100 \cdot (Q_{max}/A) \quad \text{(in cm} \cdot \text{min}^{-1})$$

where ΔL_{REF} is a warning level reference variation and k is a dimensionless filtering coefficient. The other new variables depend on the construction features of the storage facilities: A = surface of the storage tank (m^2); Q_{max} = maximum slurry flow rate (m$^3 \cdot$ min^{-1}) occurring in loading or unloading events; R_L = minimum resolution for reliable level measurements (cm). By default, it is set $k=3$. In any case, the higher the noise in the G_{0d} set, the higher the k-value to be set up prior to running the inference procedure. It was shown that a value of $R_L= 1$ cm is a reliable resolution limit for the systems in hand. Therefore, for example, in a plant with a storage surface of 200 m^2 and a maximum slurry withdrawal flow rate of 180 m^{3}h^{-1} (i.e. $Q_{max} = 3$ m^3 min^{-1}) we have ΔL_{REF}=1.5 cm $\cdot$ min^{-1}; so each measurement with a deviation greater than $3 \cdot 1.5$=4.5 cm with respect to the previous one is replaced by the latter. Such a test even performs a rounding procedure thus providing a final G'$_{0d}$ set with integer level values (in cm). A frequency spectrum analysis of the daily levels is then carried out only on this new set. In short, such an analysis: i) identifies the levels with a daily frequency greater than a threshold dF_{REF} (by default set to dF$_{REF}$ =2%). These will be then eligible to be classified as benchmarks; ii) includes all adjacent eligible values into a single group; and iii) assumes that the level benchmark to be selected as representative for each group in
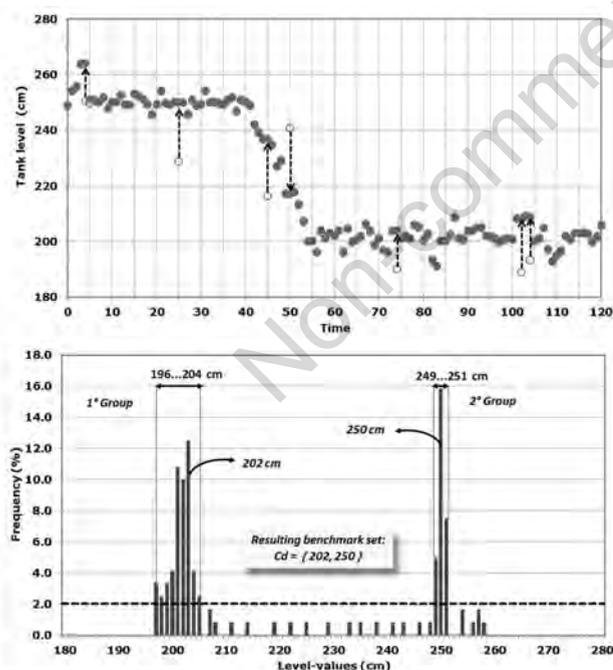
hand is the level with the highest frequency within the group itself. A simplified example of the procedure is shown in Figure 3, where we obtained C$_d$ = {202, 250} with Nc=2 benchmarks. In the following S2-step (C$_d \to$E$_{ad}$) each g$_{0dj}$ measurement is first converted into a new e$_{adj}$ value (e$_{adj} \in$ E$_{ad}$) derived from the benchmarks in Cd according to the search for a minimum deviation:

$$e_{adj} = c_{dr*} \; , \; \text{with} \; r^*: \min_{j} \left(\left| g'_{0dj} - c_{dr} \right| \right), \; \forall r$$

The procedure is then repeated on time ranges extended to 5 min, within which the benchmark with the higher frequency is selected to be definitively associated to all the measurements belonging to the same time range. This is to ensure that profile of the levels being reconstructed have greater stability. At this stage of the inference, such a profile is still limited to a sequence of horizontal lines: a change in benchmarks highlights a discontinuity point in the storage level profiles. It can be then used to detect the presence of a loading/unloading activity. Further logical filters check the consistency of points of discontinuity before confirming each activity. In the final S3-step (E$_{ad} \to$E$_d$), the inference is completed creating a new set of fitted measurements E$_d$, that can fully reproduce a stable daily tank level profile, including also the transitory phase of each loading/unloading activity. In proximity of the i-th discontinuity point, $j=p_i$ data are interpolated by a regression that describes the transitional phase through a linear variation of the levels. A priority index (in the range 0÷100) is assigned to the data that are being interpolated. It gradually increases approaching symmetrically pi where it has the maximum priority value (100). The interpolation gives the slope of the transitional profile, together with the estimate of the start and finish of the loading or unloading operation. In order to optimize spaces of system memory, only a subset SEd (such that SEd $\subseteq$ Ed) is permanently stored in the Inf-DB:

$$S_{Ed} = \{se_{dq} \mid d \in D, se_{dq} \in Ed\} \, ,$$
$$q = 1..N_E \; \text{with} \; q \in N \; \text{and} \; N_E << N_J$$
$$N_E = 2 \cdot N_c - 1$$

S$_{Ed}$ is only composed of the elements required to reconstruct the daily pattern of storage level throughout a linear interpolation between two consecutive elements. SE$_d$ is ordered with respect to time (it is also linked to T$_d$,), so its elements refer to specific instants over the day and are selected to correspond with: i) the first useful measurement for the day in hand; and ii) any sudden change in level benchmarks.

In the example in Figure 4 we have:

$$SE_d = \{250, 250, 202\}, \; NE = 2 \cdot 2 - 1 = 3.$$

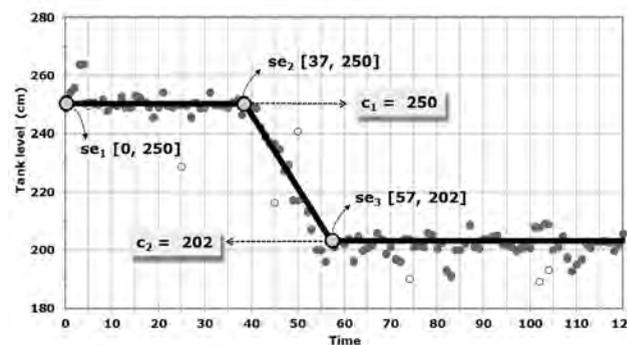It is then possible to clearly identify the storage management events:

Figure 3. Simplified example of the construction of the daily benchmark set C$_d$. The upper panel provides the original measurements G$_{0d}$ that for greater graphic clarity were limited to N$_J$ =120 ϕ= 12 min^{-1}). Empty circles are here related to outliers and the arrows indicate the replacements provided by the filtering test. The related spectrum frequency analysis is presented below with a threshold of dR$_{REF}$=2%.

Figure 4. Example of interpolation of raw-data levels with benchmarks (continued from Figure 3).

$(j_1 \div j_2) : se_1 = se_2 \rightarrow$ stationary trend : no activity

$(j_2 \div j_3) : se_2 = se_3 \rightarrow$ decreasing trend : withdrawal

$(j_3 \div j_{end}) : se_3 = se_{end} \rightarrow$ stationary trend : no activity

Based upon the above parameters and the features of the effluent at hand, the mass and nitrogen balance of the storage facility can then finally be calculated and reported over a required time range.

Testing the algorithms at farms

The inference engine was implemented within the StorEyes software package that was then applied and tested during a 2-year research period carried out on two livestock farms in Northern Italy. They were characterized by different methods of management of their slurry tank facilities.

The former (Farm 1) has two cylindrical $700 m^3$-tanks (map diameter 16 m; max. storage height 3.5 m) managed through a moderate number of loading and unloading events. The tanks are annexed to a herd of 100 dairy cows from which the liquid effluents (excrements mixed with washing water) are temporary collected into small underground storage tanks located at the head of the paddock areas. When these tanks are full, the effluent is definitively pumped inside the main tank. The withdrawal of slurry from the tank is usually driven by another pump that delivers the effluent up to the fields by conveying it into an underground pipeline. The distribution on the field is finally carried out by a tractor coupled to a spreader equipped with a distributor pipe (connected on-site to the underground pipeline) progressively rolled out in the field. Only rarely a towed $10 m^3$-slurry tank is also used for spreading.

The second farm (Farm 2) has a completely different management situation: the slurry tanks are underground and placed below the feeding area of the cowshed. This has a slotted concrete floor that allows the passage of the manure produced by the animals directly into the tank or through which the manure is scraped during cleaning. Here, therefore, the tank is continuously and progressively filled without any sudden uploading event. The storage facilities foresee two rectangular $300m3$ tanks (max. storage height 1.5 m). These are periodically downloaded into two towed slurry tanks (10 and 12 m^3, respectively) that are used for finally spreading the effluent on the fields.

Results and discussion

Monitoring the levels of a liquid stored in a tank apparently seems a very simple task; and this is actually true when the tank stores homogeneous liquids, irrespective of the volumes stored and their degree of viscosity. However, this task becomes very complex in practice when we have to deal with heterogeneous substances, often semisolid and subjected to phenomena of both decantation and fermentation, usually stored in environments that are in their turn affected by a high organizational complexity. During the experimental test period, the main problems we found in this respect dealt with: i) the difficulty of maintaining correctly function of the sensors adopted for taking the measurements; ii) the frequent interruption of the electricity supply in the cowsheds, with the consequent risk of data loss; iii) the difficulty of ensuring regular inspection and maintenance interventions; iv) the presence of disturbing phenomena that caused systematic fluctuations in the measurements to be taken due to the presence of convective motions in the stored masses or of weeds growing on their surface; and v) the disturbances that were generated by excessive viscosity of the substrate, during the loading/unloading activities, with further noise increase, and problems in detecting the timing of the events. In this framework, the perform-

ances of the inference engine, together with the use of the whole StorEyes package, were first tested to assess its *reliability* in *detecting* all types of events related to the management of the storage facilities, and in *estimating* the related mass balances with acceptable accuracy. This, of course, implies the ability to interpolate appropriately, even using a simplified and standardized pattern, the daily profiles of levels recorded.

In almost all cases, the StorEyes engine clearly identified the events that actually occurred, while also giving additional details on the way the event was performed. The case studies of the situations detected is very broad and only some of the most important examples can be discussed here.

Figure 5 reports some of the situations related to the activities carried out on Farm 1. Figure 5A shows a daily distribution carried out through the underground pipeline: the unloading event here appears as a progressive reduction in levels that is usually prolonged over time (up to 3-4 h per turn, with the operation that is executed at a very low speed, <0.5 m/s). Figure 5B describes a similar situation but with slightly shorter withdrawal turns, each anticipated by a loading event. Figure 5C shows a series of effluent distributions performed by slurry tanks; here the unloading event is characterized by a level pattern profile with shorter and more rapid variations over time. The ability of the system to detect such short events is evident-
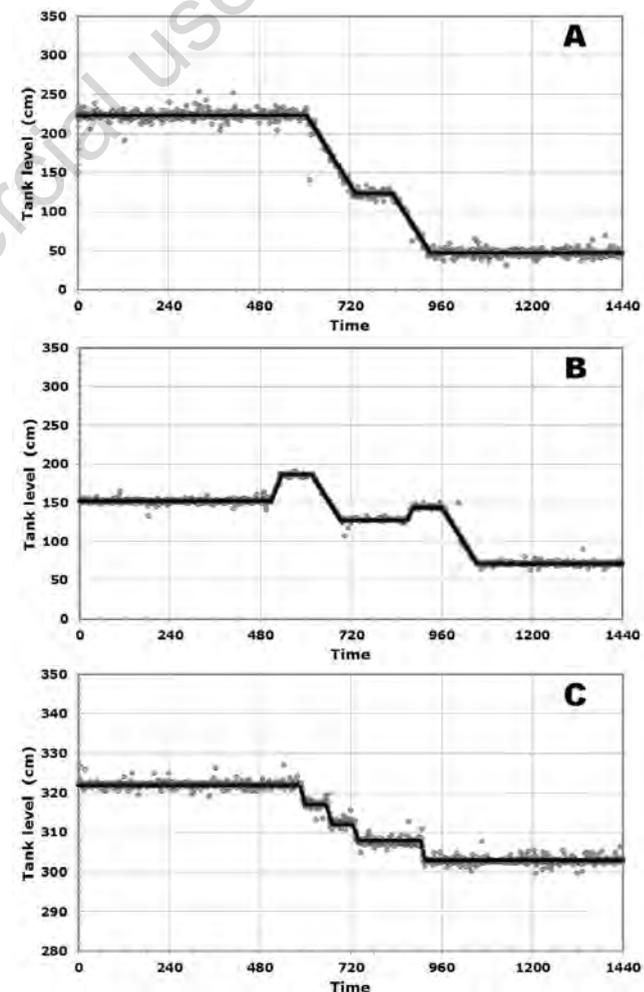

Figure 5. Examples of daily inferences carried out on storage facilities of Farm 1. Black lines show fitted results of the inference engine.

ly influenced by the noise disturbance of raw data. Despite the unfavorable conditions shown here in Figure 5A and C, the inference analysis was successfully performed.

Figure 6 provides similar case studies monitored on Farm 2. Figure 6A refers to a typical day with no withdrawal activities and shows the slight, progressive increase of the level in the tank. The trend is fitted by a single horizontal line representing the only benchmark identified for the day in hand. On such a farm, the mass of the total slurry input can only be estimated by performing a multiple day analysis providing the evaluation of compensatory loading events over a given period. Figure 6B, on the other hand, describes a withdrawal activity carried out in the early evening: this was actually two slurry tanks being loaded consecutively. However, the waiting time between the two loads was too short (approx. 15 min) and as such was not detected by the inference engine which grouped together the two unloading events in a single activity. In this situation, the final total mass remains correct but the computation of the time required to perform the operation is affected by errors. The higher the variability of the original data and the shorter the interval between two consecutive events, the lower the probability is of being able to distinguish between the two events themselves.

However, apart from these circumstances, the StorEyes engine proved itself able to interpolate more than adequately the original data of storage levels. When comparing G_{0d} *vs* E_d data sets (Figure 7), we always find results with a slope and an intercept that are never significantly different from 1 and 0, respectively ($P_{99}<0.01$ for both).

In addition, in almost all cases, StorEyes clearly identified the events that actually took place, while also giving additional details of the way the event was performed. The software package has also been used in such a way as to provide the information to farm users clearly and precisely (Figures 8 and 9).

Altogether, apart from the situations with prolonged blackout registration periods, satisfactory results were always obtained. Indeed:
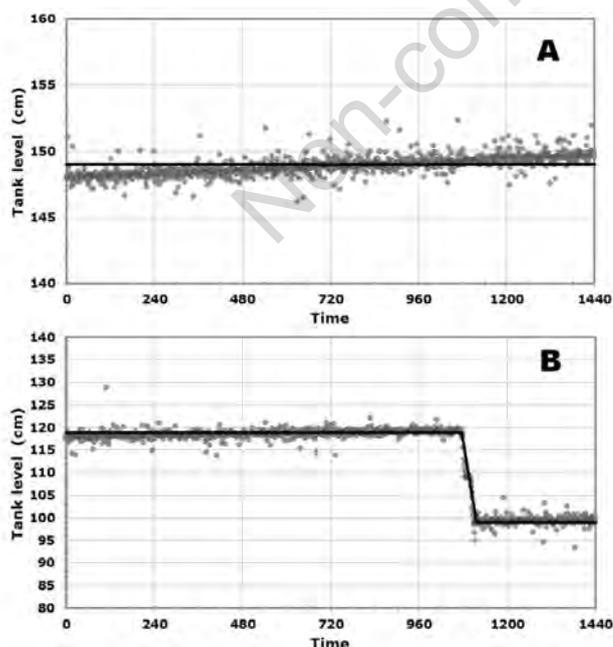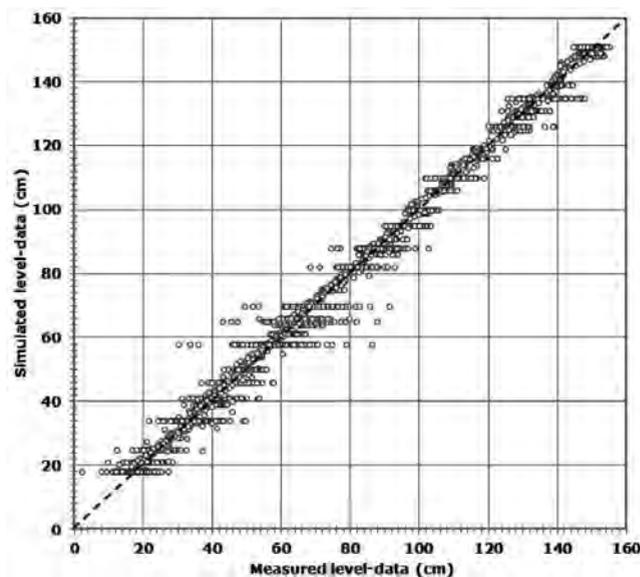

Figure 7. Comparison between observed (G_{0d}) and inferred data (E_d) performed on a set of measurements taken on Farm 2 over 42 days in the summer (July-August 2009).

Figure 6. Example of daily inferences carried out on storage facilities of Farm 2. Black lines show fitted results of the inference engine.

Figure 8. Example of monthly report provided by StorEyes according to its inference engine. 1) Raw data. 2) Raw data fitted by level benchmarks. 3) Report with event details for the period considered. Data refer to activities carried out at Farm 1.
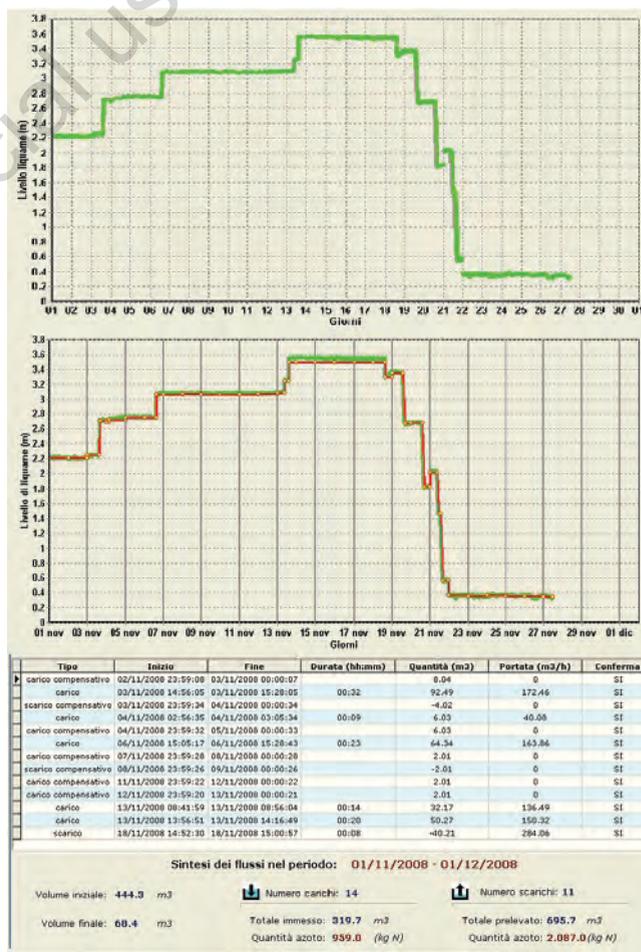
Figure 9. Example of monthly report provided by StorEyes according to its inference engine. 1) Raw data. 2) Raw data fitted by level benchmarks. 3) Report with event details for the period considered. Data refer to activities carried out on Farm 2.

a) all events were identified correctly; b) the values of the mass transferred has been estimated by figures very close to reality; and c) all changes in the level of accidental events not directly related to management activities (such as convection of the stored mass, rising gasses, prolonged rain, etc.), which were in fact very rare, have always been correctly interpreted and quantified as *compensatory events*, thus not altering the final total mass.

During the research period, attention was also given to practical issues such as the use of the software, also in combination with the other management packages used within the project, and in particular with that developed to monitor the machines used for effluent spreading. This allowed the total mass also to be combined with the mass distribution maps based on land use.

Although these aspects have not yet been subjected to a proper method of scientific investigation, the day by day experience at the two pilot farms allowed us to observe that the software: i) is able to meet user needs through its easy user-friendly interfaces and comprehensive graphical supports; ii) is highly suitable for integration into the current farm management routines; and iii) is an easy to understand intuitive package which only needs a few hours training. In fact, the computations performed are not particularly complicated.

Conclusions

In conclusion, the algorithm employed by the inference engine usually produced very satisfactory results, being able to detect and quantify autonomously all the relevant events related to the management of effluent storage facilities. The usefulness of these analyses is evident both as stand-alone application (*i.e.* at each moment it is possible to calculate and control the total effluent mass in a given time range), and as an application that integrates such an analysis with on-field monitoring of spreading activities to be carried out by tractors equipped with GPS and data-loggers. This solution allows field work registers to be automatically recorded (farm work-books), thus allowing the effluent mass transfers observed in the storage systems to be linked to the spatial information related to their distribution. The most important problems encountered during our research on storage monitoring activities were not directly due to the functionality of the inference engine, but rather to the performance of the various hardware components used (sensors, data-loggers, data transmission devices and server). Many improvements were made to the algorithms of the inference engine in order to overcome these hardware problems in an attempt to make the quality of the inference calculation as independent as possible of hardware component function. The ongoing research is now mainly concentrating on improving the reliability of all the hardware devices.

References

European Commission (EEC), 1991. Commission Decision of 12 December 1991 concerning the protection of waters against pollution caused by nitrates from agricultural sources, 91/676/EEC. In: Official Journal, L 375, 19/12/1991 pp. 1-8.

Mazzetto F, Calcante A, 2009a. Design, development and early testing of a device for the monitoring of zootechnical effluent using raingun distribution systems. J. Agric. Eng. 2:9-17.

Mazzetto F, Calcante A, Sacco P, Salomoni F, Landonio S, 2009b. Monitoring and remote controlling of slurry waste distribution activities for a sustainable management of livestock farms: the METAMORFOSI Project. Proc. 33rd CIOSTA Int. Conf., Reggio Calabria, Italy, vol. 2, pp 903-907.

Mazzetto F, Sacco P, 2009c. A software application package for monitoring and controlling slurry animal waste in storage tanks at livestock farms. Proc. 33rd CIOSTA Int. Conf., Reggio Calabria, Italy, vol. 2, pp 891-895.

McGechan MB, Lewis DR, 2000. Watercourse pollution due to surface runoff following slurry spreading. Part 2: Decision Support to minimize Pollution. J. Agric. Eng. Res. 75:429-447.

McKinion JM, Turner SB, Willers JL, Read JJ, Jenkins JN, McDade J, 2004. Wireless technology and satellite internet access for high-speed whole farm connectivity in precision agriculture. Agr. Syst. 81:201-212.

Provolo G, Riva E, Serù S, 2008. Gestione e utilizzazione dell'azoto di origine zootecnica: soluzioni tecnologiche e impiantistiche. Quaderni della Ricerca, Regione Lombardia 93:4-10.

Regione Lombardia, 2007. Regional Resolution of 21 November 2007, No 8/5868, 2nd Extr.Suppl., pp. 1-40.

Sommer SG, Hansen MN, Søgaard HT, 2004. Infiltration of slurry and ammonia volatilisation, Biosyst. Eng. 88:359-367.

Wang N, Zhang N, Wang MH, 2006. Wireless sensors in agriculture and food industry. Recent development and future perspectives. Comput. Electron. Agr. 50:1-14.