

ARCHITECTURE OF NEAR REAL-TIME MONITORING SYSTEMS FOR WATER DISTRIBUTION SYSTEMS


by
Sebastian K. Bögershausen

**Bachelor of Science in Civil Engineering
The George Washington University, 1997**

Submitted to the Department of Civil and Environmental Engineering in
Partial Fulfillment of the Requirements for the Degree of

**MASTER OF ENGINEERING
In Civil and Environmental Engineering**

at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2001



BARKER

©2001 Sebastian K. Bögershausen. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part.

Signature of Author: _____

Department of Civil and Environmental Engineering
May 11, 2001

Certified by: _____

George Kocur
Senior Lecturer, Civil and Environmental Engineering
Thesis Supervisor

Accepted by: _____

Oral Buyukozturk
Chairman, Departmental Committee on Graduate Studies

ARCHITECTURE OF NEAR REAL-TIME MONITORING SYSTEMS FOR WATER DISTRIBUTION SYSTEMS

by
Sebastian K. Bögershausen

Submitted on May 11, 2001, to the Department of Civil and
Environmental Engineering in Partial Fulfillment of the Requirements
for the Degree of

MASTER OF ENGINEERING
In Civil and Environmental Engineering

ABSTRACT

The town of Arlington, Massachusetts, is planning to improve its existing water meter reading process by installing wireless transmitters on the 12,200 meters throughout the town. A web-based Java software program was developed at MIT to manage the readings (as frequent as every ten minutes) from this monitoring system and perform analytical tasks to prepare billing reports, analyze usage patterns, detect leaks, notice water theft, and find deteriorating meters. This thesis describes the system architecture of this water distribution management software, how it compares to large regional supply systems, and what possibilities monitoring networks of this kind have in the future.

Thesis Advisor: George Kocur

Title: Senior Lecturer for Civil and Environmental Engineering

ACKNOWLEDGEMENTS

I am eternally grateful to those who made it possible for me to pursue and fulfill one of my dreams:

Col. & Mrs. Charles and Irene Hertel

Mr. Bob Gladstone

Dr. George Kocur, Massachusetts Institute of Technology

Dr. Majid Manzari, The George Washington University

Dr. Irving Shames, The George Washington University

And sincerest thanks to my friends and colleagues who worked on the Arlington Project with me:

Mr. Brad Butler, MIT

Mr. Mameet Khanolkar, MIT

Mr. Kent Larson, Town of Arlington Water Department

Mr. Marcis Kempe, MWRA

TABLE OF CONTENTS

1	INTRODUCTION	6
2	LITERATURE REVIEW	8
2.1	OVERVIEW	9
2.2	AUTOMATION LEVELS IN WATER SUPPLY SYSTEMS	9
2.3	MEASUREMENT AND CONTROL VARIABLES.....	10
2.4	ARCHITECTURE OF A WATER SUPPLY SYSTEM.....	11
2.5	USE OF TELEMETRY IN WATER MANAGEMENT	15
2.6	DISTRIBUTION MODELING.....	19
2.7	APPLICATIONS IN LEAK DETECTION PROCESSES	19
3	PERFORMANCE AND OPERATIONS OF THE MWRA MONITORING SYSTEM	25
3.1	SCOPE.....	25
3.2	HYDROLOGICAL MONITORING.....	26
3.3	BIOCHEMICAL MONITORING.....	28
3.4	SIMULATION/MODELING	30
3.5	MWRA'S SCADA SYSTEM.....	31
4	ARLINGTON REMOTE MONITORING SOFTWARE GENERAL DESCRIPTION	32
4.1	OVERVIEW	32
4.2	SOFTWARE FUNCTIONALITIES	33
4.2.1	<i>Bill Preparation.....</i>	<i>34</i>
4.2.2	<i>Meter Watch.....</i>	<i>34</i>
4.2.3	<i>Leak Detection</i>	<i>35</i>
4.2.4	<i>Usage Pattern Analysis.....</i>	<i>36</i>
4.2.5	<i>Meter Accuracy Validation.....</i>	<i>36</i>
4.2.6	<i>Water Theft Detection</i>	<i>37</i>
4.2.7	<i>Administrative Functions</i>	<i>38</i>
4.3	SOFTWARE IMPLEMENTATION AND APPLIED TECHNOLOGIES	38
5	SOFTWARE ARCHITECTURE AND FUNCTIONAL MODEL DESCRIPTION	40
5.1	USAGE SCENARIO	40
5.1.1	<i>User Profile</i>	<i>40</i>
5.1.2	<i>Administration of the Data Tables Use Case.....</i>	<i>40</i>
5.1.3	<i>Configuration and Operation of the Reporting Tools Use Case</i>	<i>43</i>
5.1.4	<i>Viewing of the Report Results Use Case.....</i>	<i>44</i>
5.1.5	<i>Bill Preparation Use Case</i>	<i>45</i>
5.1.6	<i>Usage Pattern Analysis Use Case.....</i>	<i>47</i>
5.1.7	<i>Meter Watch Functionality Use Case</i>	<i>49</i>
5.1.8	<i>Future Enhancements.....</i>	<i>50</i>
5.2	FUNCTIONAL MODEL AND DESCRIPTION.....	50
5.2.1	<i>Bill Preparation Functionality.....</i>	<i>50</i>
5.2.2	<i>Meter Watch Functionality.....</i>	<i>54</i>
5.2.3	<i>Usage Pattern Analysis.....</i>	<i>57</i>
5.2.4	<i>Database Access</i>	<i>59</i>
5.2.5	<i>Administrative Functions</i>	<i>63</i>

5.2.6	Software Interface Description	67
5.3	BEHAVIORAL MODEL AND DESCRIPTION	67
5.3.1	Description of Software Behavior	68
5.3.2	Process State Diagram	70
6	CONCLUSIONS	71
7	REFERENCES	73

LIST OF FIGURES

FIGURE 2-1	- DISTRIBUTION SYSTEM COVERAGE	12
FIGURE 2-2	- OVERALL CONTROL AND MANAGEMENT	14
FIGURE 2-3	- TYPICAL TELEMETRY SYSTEM USING RADIO COMMUNICATIONS	15
FIGURE 2-4	- COMPUTER CONTROL AND TELEMETRY	18
FIGURE 2-5	- SIMPLE TIME-LAGGED NETWORK PROBLEM	21
FIGURE 2-6	- SAMPLE NETWORK FOR TRANSIENT ANALYSIS	23
FIGURE 3-1	- MWRA FAILURE REPORTS	27
FIGURE 4-1	- HIGH-LEVEL ARCHITECTURE OF WATER SYSTEM	33
FIGURE 4-2	- SOFTWARE ARCHITECTURE	39
FIGURE 5-1	- USE CASE DIAGRAM FOR ADMINISTRATION OF METER TABLE	41
FIGURE 5-2	- REPORT CONFIG USE CASE	44
FIGURE 5-3	- RUN ANALYSIS USE CASE	44
FIGURE 5-4	- VIEW REPORTS USE CASE	45
FIGURE 5-5	- BILL PREPARATION USE CASE	46
FIGURE 5-6	- USAGE PATTERN ANALYSIS USE CASE	48
FIGURE 5-7	- METER WATCH USE CASE	49
FIGURE 5-8	- BILL PREP CLASS DIAGRAM	52
FIGURE 5-9	- BILL PREP SEQUENCE DIAGRAM	53
FIGURE 5-10	- METER WATCH CLASS DIAGRAM	55
FIGURE 5-11	- METER WATCH SEQUENCE DIAGRAM	56
FIGURE 5-12	- USAGE PATTERN CLASS DIAGRAM	57
FIGURE 5-13	- USAGE PATTERN SEQUENCE DIAGRAM	58
FIGURE 5-14	- DATABASE ACCESS CLASS DIAGRAM	60
FIGURE 5-15	- DATABASE ACCESS USING ONLY DBBEAN SEQUENCE DIAGRAM	61
FIGURE 5-16	- DATABASE ACCESS USING DBRESULTS SEQUENCE DIAGRAM	62
FIGURE 5-17	- CLASS DIAGRAM FOR ADMINISTRATIVE CLASSES	64
FIGURE 5-18	- UTILITY CLASSES DIAGRAM	65
FIGURE 5-19	- UPDATE METER SEQUENCE DIAGRAM	66
FIGURE 5-20	- PROCESS STATE DIAGRAM FOR METER AND METER READ	70

1 Introduction

The Department of Public Works of the town of Arlington, Massachusetts, decided in the year 2000 to improve their existing water meter reading process by installing wireless transmitters on all the roughly 12,200 meters throughout the town. The town's incentive to install this system was to improve the knowledge and understanding of the local water usage, optimize the billing service, and potentially lower the operating cost of the water department.

Three students from the Massachusetts Institute of Technology's Master of Engineering program in Information Technology for Civil Engineering worked over a period of nine months to build a functional software package to run this system. We met with the client (Arlington's Department of Public Works) to collect the requirements and gather all necessary information to be able to build the software.

Work began in September 2000 through requirements collection and documentation, client meetings and discussions with the regional water distribution agency, the Massachusetts Water Resources Authority (MWRA). By December 2000, we had developed a prototype that showed all features we intended to implement. In the month of January we substantially revised the system based on client feedback and designed each functional element of the system in further detail. The software code was written in February and March 2001 with final changes and enhancements being made over the following months.

The software is intended to provide functionality for the manager of a local water distribution system to capture and control the daily readings of a wireless meter reading network. This system provides the pathways to store data in retrievable locations and performs analytical functions to analyze usage patterns, detect leaks, search for a loss of meter accuracy, etc. Some of these functions require interactions with the regional water authority and potentially also their operations software.

This thesis seeks to explain the architecture of the local monitoring system - how it operates and interfaces with the regional system, how it compares to large-scale water operations software packages (SCADA systems – Supervisory Control And Data Acquisition systems) and what potential systems such as this have for future applications. It is logically divided into chapters that increase in technical depth and detail in relation to the system.

The first chapter describes the technical literature that is available today regarding water distribution software systems and what tasks these programs perform. Topics vary from hardware to software and hydrological modeling

systems. The next chapter is a description of the MWRA operations and performance monitoring systems; a discussion of the process, types and locations of monitoring systems that the MWRA employs, and how this relates to a local monitoring network.

Chapter four describes in general the intent of the software package developed at MIT and what the exact functionalities are that need to be implemented to have this system fully functional. It also covers the technologies used to create this program and how they are applied to the problems at hand. Chapter five follows up on this by describing in detail the application that has been developed, as well as exhibiting standard use cases and sequences.

2 Literature Review

Chapter Sources Overview

(Refer to References at the end of this thesis for exact publishing information)

2.2 Automation Levels in Water Supply Systems

Automation Levels in Water Supply Systems

by Edmundo Koelle

in

Water Supply Systems – State of the Art and Future Trends

Edited by E. Cabrera and F. Martínez

2.3 Measurement and Control Variables

Measurement and Control Variables

by Edmundo Koelle

in

Water Supply Systems – State of the Art and Future Trends

Edited by E. Cabrera and F. Martínez

2.4 Architecture of a Water Supply System

Public Water Supply – Models, Data and Operational Management

Dušan Obradovic and Peter Lonsdale

2.5 Use of Telemetry in Water Management

Operational Control of Water Systems

M.A. Brdys and B. Ulanicki

and

Communications in Advanced Technology in Water Management

Edited by K.F. Roberts, Written by J.G. Hurcom

2.6 Distribution Modeling

Operational Control of Water Systems

M.A. Brdys and B. Ulanicki

2.7 Applications of Distribution Modeling in Leak Detection Processes

Network monitoring and the algorithmic location of leaks under steady and unsteady conditions

by James A. Liggett in

Water Supply Systems – State of the Art and Future Trends

Edited by E. Cabrera and F. Martínez

And Analysis of Water Distribution Systems

Thomas M. Walski

and

Identifying and Reducing Losses in Water Distribution Systems

James W. Male

2.1 Overview

This chapter explores the existing technical literature that pertains to the architecture of a data management system of a water supply organization. This material is important for understanding the effect of computer hardware and software on the water industry and the possibilities that they provide. Although the literature ranges over a wide area of topics, the intent of this review is to provide a backdrop of understanding to the implications of a wireless water meter reading system and the software for its management.

2.2 Automation Levels in Water Supply Systems

An automated water supply system is a water distribution network that has an operational computer system that controls (or suggests actions to control) the flow and distribution of the water. Possible actions that can be controlled are such things as opening and closing valves, activating pumps, adjusting pressure by increasing tank fill levels, etc. In addition to controlling the parameters of the distribution network, the system also allows the interpretation of data that is collected from points in the network such as meters, pumps, tanks, etc.

The basic system architecture of an automated water supply network is composed of a number of elements that act in concert to perform the acquisition of data, the supervision of the system operation and diagnostic consumption reporting. A system that fulfills these needs is called a SCADA system – Supervisory Control and Data Acquisition system. The SCADA system generally has two physical components: remote stations and a central computer.

Remote Stations are installed at strategic points in the distribution network, such as pump stations, storage tanks, meters and flow control valves. Each remote terminal is composed of a microprocessor with analog and digital input and output interfaces and is connected to the local instruments and actuators. The remote unit provides the acquisition of data and its transmission to the central computer. It can also receive commands from the central computer and perform actions based on this input.

The central computer receives the signals transmitted by the remote station and processes them, producing graphical and/or numerical information for analysis by the operator.

In order to create an automated water supply network to control the distribution and not only provide operational supervision, it is necessary to develop computational programs that will process the information received by the central computer and propose a “Plan of Maneuvers”. The orders of this plan are then transmitted to the remote stations, where the corresponding action is performed.

The levels of automation for water supply systems may be classified as follows:

- Level 0* Non-programmed local automatic control (sensors)
- Level 1* Programmed local automatic control and automatic control valves
- Level 2* SCADA system without automatic control action
- Level 3* SCADA system with automatic control

At level 0 and level 1, the decision-making process depends entirely on the experience of a system operator. No decisions are made based on integrated information at a centralized location, but on a local basis instead. In liberally dimensioned water networks where demands are served by a small number of tanks, operation at level 0 or 1 can be enough for the necessities, although the system performance will not be optimal. However, potential side effects can be pressure variations and intermittent service, due to the (relative) lack of integrated control over the network.

Distribution of water in large urban centers is operated at level 2 and the operation of many of the control points in the network is activated remotely after much interpretive and simulation work has been performed. Automating this process of decision-making (or decision proposal) leads to level 3, which is the level that many water distribution authorities are currently in the process of developing or installing.

Even though the computational capabilities of computers today are certainly sufficient to perform highly complex distribution simulations and decision-support analysis, there are still difficulties in obtaining adequate physical data for the calibration of these models and thus the operator's interference is still necessary to make corrections to the proposed Plan of Maneuvers. Also, the problem of attempting to predict long-term usage is made severely more difficult due to meteorological conditions and local contingencies.

2.3 Measurement and Control Variables

The components that make up the measurement body of an automated supply network are the source of all the data that guides the distribution. The integration of these instrumentations and a SCADA system forms a controlled hydraulic network, so it is necessary to be aware of the measurements that need to be taken as well as what controls should be implemented.

The minimum set of measurements and controls to be made may be grouped in the following manner:

1. Measurements

- Tank levels
- Pressure in pipelines
- Flow rate
- Chlorine residual
- Turbidity
- Control valve degree of opening
- Local temperature

2. States of Apparatus

- Pump (on/off)
- Shut-off valve (closed/open)
- Remote terminal (connection enabled/disabled)
- Power (on/off)
- Control (local/remote)
- Safety parameters

3. Controls

- Pump (on/off)
- Valve (open/closed)
- Alarms
- Vibrations
- Temperatures
- Actuator

There are a variety of tools that allow the measurement and control of the distribution, and it is beyond the scope of this thesis to evaluate the differences between the types and manufacturers. For the purposes of the installation of a wireless meter reading system, the above list needs to be amended to include water meters and their states (registering/not registering).

2.4 Architecture of a Water Supply System

The control of even a relatively small distribution network can present a challenging problem to the designer, since there are several seemingly conflicting interests. For example, a system with many alternate paths and redundant supply sources provides a greater level of security of the supply, whereas a distribution tree with one source that branches out to ultimately serve every user is easier to manage since the flow is easier to ascertain and the number of pipes is reduced. In addition, not all systems only serve their customers, but also act as transmission lines to the next contiguous town or distribution system.

Although it can be difficult to integrate many of these conditions into one design, the most influential design parameters of a water supply system remain constant: Geographic contingencies determine well and tank locations and pressure zones and govern most of the system layout. A system of valves controls where and how the water will reach the customer and regulates the pressure zones. Between distribution districts, meters need to be installed to ascertain usage data for each district.

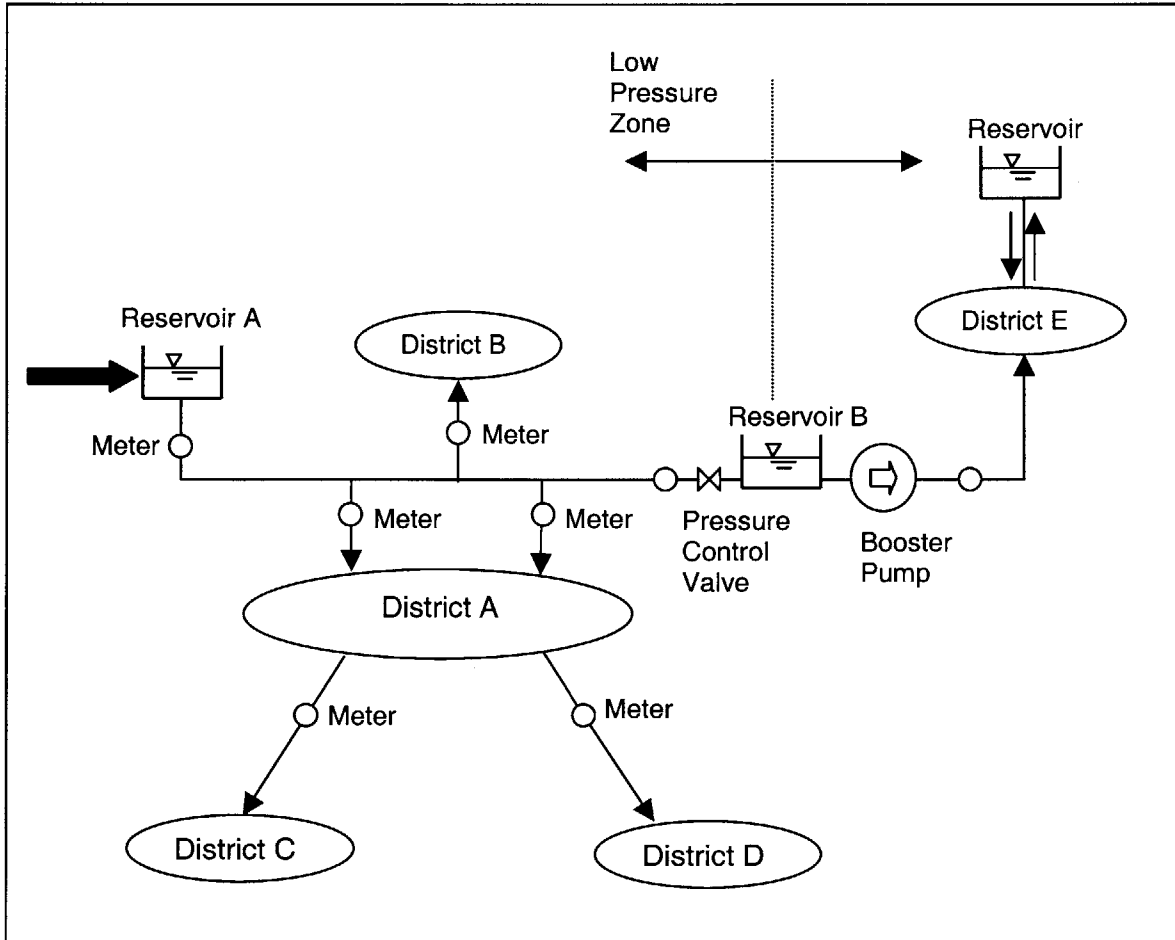


Figure 2-1 - Distribution System Coverage

Reproduced from *Public Water Supply – Models, Data and Operational Management* by Dušan Obradovic and Peter Lonsdale

Figure 2-1 shows part of an urban distribution network that requires a high and low pressure zone based on its geographic layout. Water flows from the supply source into the main reservoir A and from there to the lower reservoir B by gravity. A pressure-reducing valve controls the inflow into reservoir B. Several lower districts are served from this reservoir. A booster pump brings water up in elevation to district E and reservoir C.

All districts are monitored at the district level and at the consumer level. The readings from these meters, as well as data from the pressure valves, tank levels, etc. are transferred to an operational management system that allows the creation of a feedback loop of operational data and resulting actions and commands (see Figure 2-2).

Data describing the current state and performance are captured and sent to the central operations center. This information is analyzed (preferably with a model of the system, see 2.6 - Distribution Modeling), a forecast of future developments is made and managers at various levels take appropriate action. Examples are:

- Long-term management – planning of new resources to meet future potential shortfalls
- Medium-term management – changing operating regimes to take advantage of cheaper tariffs
- Short-term management – redeployment of manpower and equipment to reduce leakage

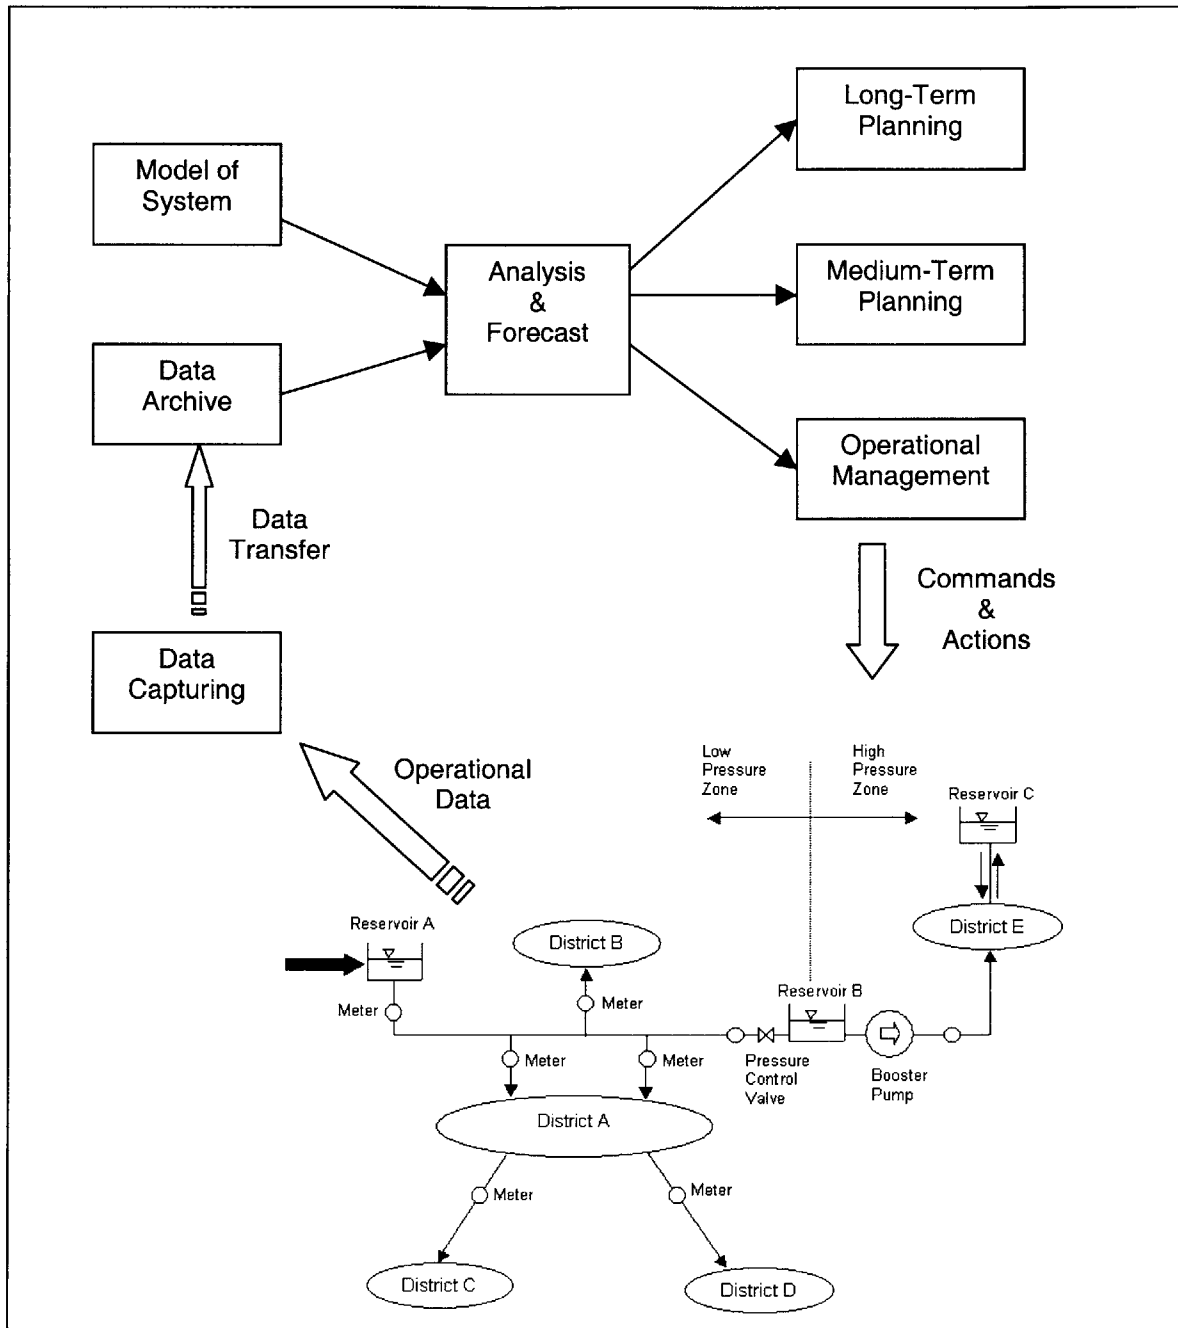


Figure 2-2 - Overall Control and Management

Reproduced from *Public Water Supply – Models, Data and Operational Management*
 by Dušan Obradovic and Peter Lonsdale

2.5 Use of Telemetry in Water Management

Water distribution networks are spatially distributed over a wide area and require communication links connecting distant sites with control computers. A collection of hardware and software that enables remote sensing and actuating is called a telemetry system. The origin of the data to be transferred by the communications system is in most cases sensors, which are installed in the supply system at key locations. The system state is continuously monitored and the collected information is stored, processed and finally used to decide on a control action. The control decisions are transformed into physical action by specialized actuators.

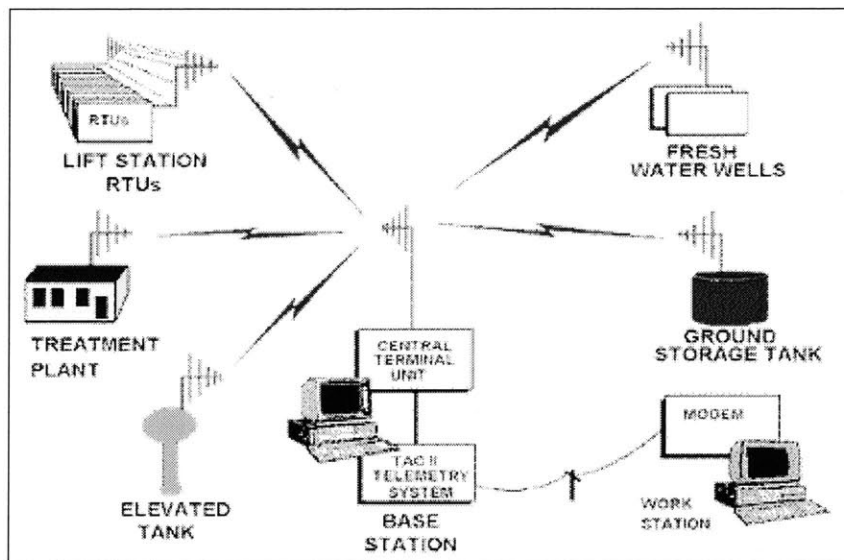


Figure 2-3 - Typical Telemetry System using Radio Communications

The areas in which telecommunications play a key role can be considered in three categories:

- Monitoring and control of the processes (water treatment/distribution)
- Operation of support processes (customer billing systems, infrastructure management systems, materials management systems)
- Managerial control of field operations

In the past, monitoring systems were designed to transfer data from monitoring locations to a central computer through the use of a modem and via public switched telephone network (PSTN). These developments commenced in the 1980's with relatively simple machines that dialed a master receiving station

following the initiation of an alarm condition, and then transferred digitized structures of site address and plant state. Rapid development continued throughout the 1980's, culminating in highly sophisticated intelligent outstations able to pass on data at a much faster rate and in more accurate and secure form.

PSTN provided a cost effective means of communication for plant data recovery where continuous real time supervision is not necessary and as such was suitable to service installations such as meters or apparatus that only needs to send exceptional changes.

Today, water monitoring systems only make limited use of the PSTN to transfer data. Three hardware layers characterize modern systems: The telemetry system, the local control room and the coordination level (see Figure 2-4). The telemetry system transmits data in both directions between the physical system and the computer using communication links. The computer and other auxiliary equipment in the control room are interconnected by a local area network (LAN). For large water networks, there is normally another decision level that communicates with the local control rooms through a wide area network (WAN). Machines on the LAN are connected through an Ethernet card and TCP/IP protocol, which allows all machines and monitors to communicate with each other. The top level in the decision hierarchy coordinates the activities of many local control centers.

An operator in the control room receives on-line measurements from the process and executes the control decisions. The operator must be provided with an efficient interface to obtain the system behavior and to implement control commands. The telemetry computer runs software for the data communication, data processing and the data presentation. The database plays the role of enabling access to information regarding the history of the water network behavior.

The communication units shown in Figure 2-4 provide an interface between the control computer and the transmission lines. Data are coded by modem for transmission purposes and sent into the line. On the other end of the connection they are decoded by a similar modem. The transmission using modems is not limited to PSTN lines, but these days can also occur over leased-lines, digital subscriber lines (DSL) and cable TV networks.

The major communication routes that connect the important sites of the water distribution system (large pump stations and treatment works) are implemented as high-speed data highways. To transmit high volumes of data, optical fibers are often employed. Measurements are taken at regular intervals and are transmitted to the telemetry database. Optical fibers are particularly attractive because of their high capacity and immunity to electromagnetic interference.

For less critical sites where quantity transmitted is less, medium speed communication links are sufficient, such as phone lines, radio transmitters or serial cables. The phone lines can be owned by a water company or leased from a telecommunications company. Leasing the phone lines is expensive, so normally these are only used for batch transmission at times of cheap telephone tariffs. Data are stored locally by a data logger and then sent as one batch to the central computer. Radio links are useful in rural areas for widely distributed water networks. Remote monitoring and control can be performed without cables or line rentals. Radio telemetry can be installed quickly and relocated as site requirements change. Wireless data technologies are becoming more viable and will most likely see their level of use increase, due to their great versatility and ease of use and maintenance. As the reliability of wireless data systems increases, so will their use.

The intelligent controllers and remote monitors shown in Figure 2-4 interact directly with the physical system. The intelligent controllers are equipped with processors and memory so that they can store measurements, execute commands sent from the decision centers or run their own programs. They have a number of serial I/O ports served by I/O cards. I/O cards handle the following functions and signals: A/D (analog-to-digital), D/A conversions, digital signals, analog signals, pulse input and stepper motor output, etc. The remote monitors perform only basic functions such as meter reading and actuating pumps and valves. They have analog and digital inputs and outputs to interact with the analog and digital instruments.

Early distributed control systems used proprietary hardware interfaces, protocols, modems and transmission parameters. Much effort has been made to achieve some standardization. Interface standardization facilitates both development and maintenance of control schemes. Standards define factors, such as circuit board design, cable connector pin assignment, and signal levels. The best recognized standards are: IEEE Standard 488, CAMAC, and EIA Standards. IEEE standards define the electromechanical aspects of parallel transmission. The CAMAC standard provides the interfaces for parallel, bit serial, and byte serial communication, and has found major use in the nuclear industry and process control applications. The Electronic Industries Association (EIA) standards RS232, RS422 and RS449 define bit serial synchronous and asynchronous environments and are used in data processing as well as in process control. These days, RS422 and IEEE488 are being displaced by direct LAN (Ethernet) connections.

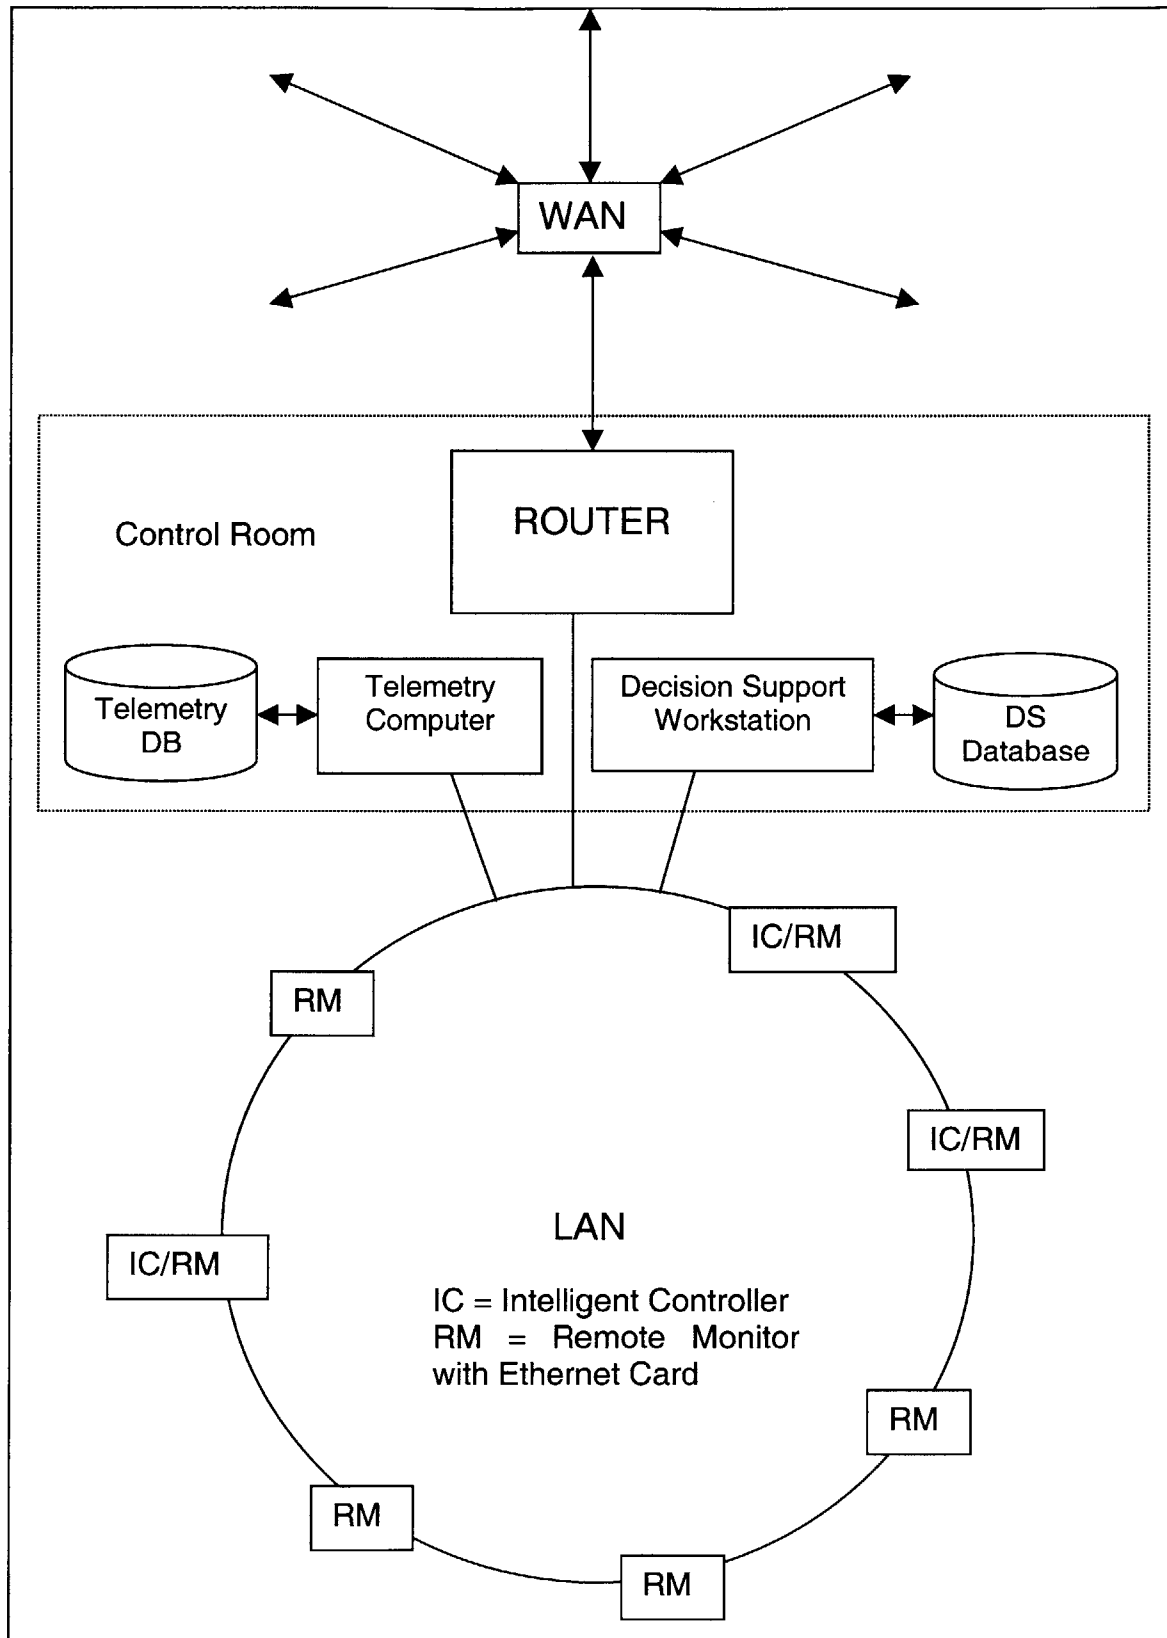


Figure 2-4 - Computer Control and Telemetry

2.6 Distribution Modeling

Algebraic relationships can be established to derive static models for pipes, valves and pump stations in a water distribution network. Differential equations serve to create dynamic models for more complex networks or reservoirs. There are passive elements (pipes and valves), which dissipate energy, and active elements (pumps), which supply energy to the network. An important aspect from the modeling point of view is that the pressure head decreases along the passive elements and is increased by the active elements. The structure of the network is described by a node-branch incidence matrix, where the rows correspond to the network nodes and the columns are associated with the branches of the network.

In general, an overall mathematical model of the network can be put either into nodal or loop form, depending on its independent parameters. A mixed nodal-loop version can also be derived. Most commercial packages today utilize nodal formulations.

From each node to the next one, a hydraulic relationship is created based on traditional fluid dynamics. Thus based on the pressure head at the reservoir and knowing all the losses that occur between the reservoir and a node in question, an estimate of the resulting head can be calculated. Creating a network of such relationships allows a more complex and therefore meaningful analysis. Such a simulation model can improve the performance of a remotely monitored system mainly in aiding the leak detection process as well as system verification to estimate whether use for a location is reasonable.

The model is derived in terms of unknown nodal heads and unknown branch flows. During simulation applications, heads at the non-reservoir nodes are unknown heads, while branch flows present unknown flows. In the case of estimation, a situation may be more general and depend upon available data. Data can be filled in at known nodes, as long as one can be sure that this additional data is consistent with the boundary conditions that govern the simulation.

2.7 Applications in Leak Detection Processes

Every water distribution system has leaks and their detection can be problematic as the distribution lines are often buried or otherwise inaccessible. However, it is of interest to perform leak detection and eradication because of possible increases in system performance and reduction in the operating cost. Methods that can be employed to detect leaks are:

- Periodic pressure surveys attempt to find leaks by calculating an expected friction (or other) head loss through the pipe, which can be compared to the measured on-site head.
- Listening devices for sonic leak detection listen to the sound of the flow through a pipe. When water escapes from an orifice, it causes a vibration in the 500 – 800 Hz range. This sound travels along the pipe wall and can be heard at a considerable distance by an observer with the proper equipment.
- Infrared surveys search for escaped water by looking at a temperature profile of the area around the pipe.
- Ground-penetrating radar searches for motion in the vicinity of a buried pipe
- Analysis of the changes in the ratio of daytime to nighttime consumption allows a numerical comparison to estimate if there are any losses through the pipe.

A monitoring system that captures information about the pressure and flow through the network can greatly aid in the detection of leaks. Flow meters and pressure gauges monitor the system continuously and provide a very large amount of data. Based on this data, the health of the distribution system can be analyzed.

There are two possible methods of analysis: Steady state and non-steady-state. A steady state analysis assumes that the supply network is entirely static, all pressures are at design levels, and that there is no other flow in the system other than the flow in question. A non-steady state analysis prepares a more realistic estimate of the network by allowing the pressures and flows to vary throughout the entire system, while still trying to ascertain the pressure and flow conditions at a specific location. The design and operation of water distribution systems has traditionally been done by steady state techniques. For design purposes steady state methods are adequate, but a large distribution system never really attains equilibrium in operation. Thus, for monitoring purposes, a transient analysis is more applicable. Also the analysis is an inverse analysis, meaning that we know the output from measurements and wish to find the system characteristics that produced that output.

James A. Liggett describes in *Network monitoring and the algorithmic location of leaks under steady and unsteady conditions* a computer program written to solve a transient state distribution network problem and searches for leaks. This program takes simultaneous differential equations and solves them based on data received from the monitoring system and boundary conditions. Each equation describes the state of a node in the distribution network. A node can be a key point where flow parameters like velocity and pressure are monitored, or simply an intermediary location for which flow characteristics are calculated.

The four methods that this program employs are outlined below.

1. Forward calculations for flows and pressures

This mode is the standard analysis for transient states in looped and branched networks. It can compute the results of pumps coming online, valve closures, etc.

2. Time lagged calculations for leak areas

Information from advanced time is used to compute what must have occurred at a previous time. As the simplest such problem, consider Figure 2-5. Normal boundary conditions are given at the left end of the pipe but no conditions are given at the right end. A pressure measurement is made at some intermediate point. A transient calculation can be carried forward using the left boundary condition and the pressure measurement in the normal manner. The conditions at the right end can be computed by carrying information in the positive time direction along the forward characteristic and in the negative time direction along the backward characteristic.




Figure 2-5 - Simple Time-lagged Network Problem

Reproduced from Network monitoring and the algorithmic location of leaks under steady and unsteady conditions by James A. Liggett

3. Inverse calculation for calibration and/or leak areas

The inverse calculation makes use of the forward calculation, which can be in the time-lagged mode. This function takes data for the parameters at known nodes and attempts to find a distribution pattern that fits this data. For the purposes of finding leaks, the transient calculation has two large advantages over the steady state calculation: (a) Pressure waves are less affected by friction than the general flow and thus the precise friction values become less important to the calculation (the computation is less sensitive to friction) and (b) Most networks operate the majority of the time in a transient mode so that the monitoring operation does not need to wait for stationary periods to take pertinent readings for this analysis.

4. Location of sudden breaks

This part of the program acts as a monitoring function for the system. The algorithm is triggered by a pressure jump of predetermined magnitude over a short period of time in any of the measurement locations. The time of passage of the surge past several pressure gauges can often locate the position of the origin of the surge given the wave speed in the pipes in a manner similar to that used to locate earthquakes. A node can then be placed in that position (if the surge is not already located at a node), and the system is reconfigured to calculate water loss.

Process

First the geometry of the system must be resolved. Aside from general topology, the resolution of geometry includes the identification of boundary conditions and measurement points. In preparation for calculation, the program may place nodes in the system that are not specified by the user for the purpose of controlling the Courant number. (The Courant number is a non-dimensional number used in computational fluid dynamics that is defined as flow velocity multiplied by time step divided by unit length.) The method of characteristics uses temporal interpolation to resolve the differences in time of the end points of the characteristics at the nodes. To maintain an accurate solution with minimum diffusion, an attempt is made to keep to Courant number in the neighborhood of one.

The program next determines the paths. These are the routes of shortest wave travel times between the boundary nodes or the measurement point and all other nodes in the system. The paths are used both to determine time lags and to locate surges in the break location subroutine.

The time lags for each node are calculated. The time lag depends on the distance of the node from the nearest boundary condition or measurement point, the number of unknowns at the node and the time lags of each of the neighboring nodes. The equation formulation follows in which the location of each term in a coefficient matrix is found. The path determination, time lags and equation formulation form the parts of the program that contain the logic to proceed with the transient calculation.

Once the time lags have been calculated and the simultaneous equations established, an analysis can be run to perform the time-lagged problem for leaks or the inverse problem for leaks and/or calibration.

As an example, consider the flow network shown in Figure 2-6. The inflow in node 1 is pressure dependent and the inflows in nodes 5 and 6 are constant. The outflow in node 3 is also pressure dependent. This flow network is translated into

a set of simultaneous equations, using the four indicated flows as boundary conditions and with an equation for every node. Readings from the monitoring system can either provide additional data to be taken into account in the calculations, or they can serve as data against which to compare the results of a calculation based purely on the boundary conditions.

A primary difficulty with transient analysis of actual pipe networks is that the boundary conditions are seldom known and the discharge is not monitored on a continuous basis. However, given enough transient data, a calibration program should be able to estimate those discharges.

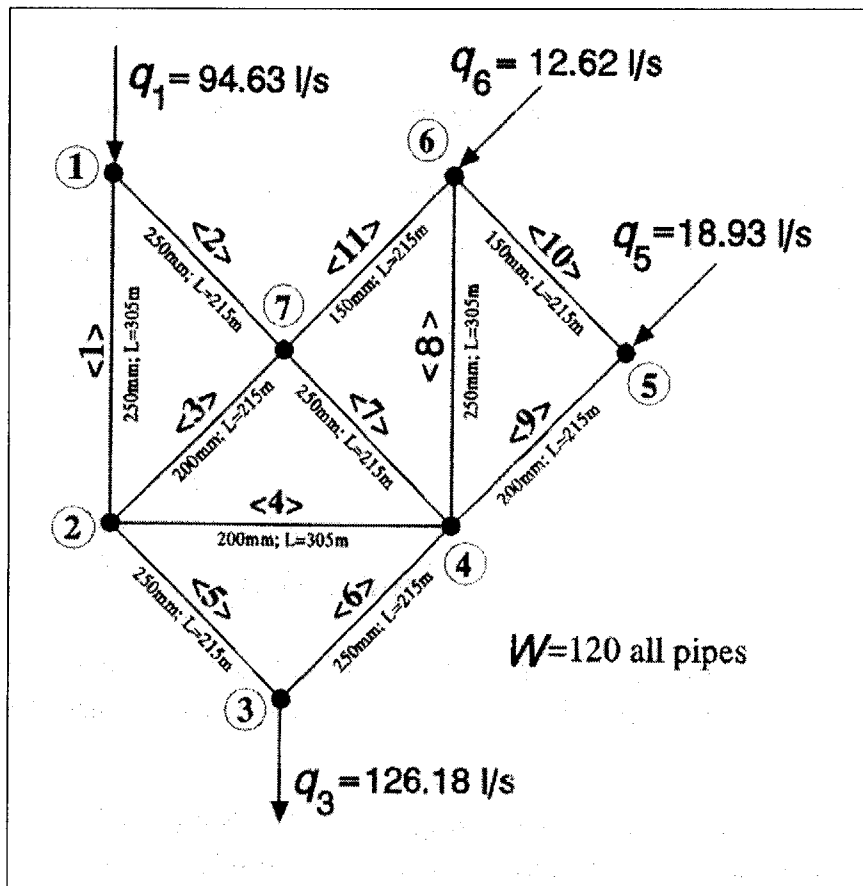


Figure 2-6 - Sample Network for Transient Analysis

From *Network monitoring and the algorithmic location of leaks under steady and unsteady conditions* by James A. Liggett

The results of the analysis are values for the pipe friction and leak area (size of leak), which can be plotted against actual data to determine similar trends. For each node, a profile can be established which indicates the performance. Since most leaks in distribution systems are the locations of pipe joints, the profiles

(and nodes) are frequently created at pipe joint locations, in addition to the measuring nodes. The program can perform several iterations of each calculation to attempt to meet actual data from the monitoring systems and thus create relatively reliable results.

3 Performance and Operations of the MWRA Monitoring System

The town of Arlington receives its water from the regional water resources authority and so interfaces with their distribution network at several key points. The regional authority has its own monitoring network that is currently totally independent from the system that we have built. However, it is not unreasonable to assume that at some time in the future, information from these two data and distribution networks will be exchanged to create a more efficient system. For this purpose, the author interviewed the regional water resources authority to determine their operating standards and how, where and when their monitoring system works.

3.1 Scope

The current system the Massachusetts Water Resources Authority (MWRA) employs to manage and operate their distribution network is logically divided between the drinking water distribution and sewerage disposal operations. For the purpose of this document, the focus will mainly be on the drinking water distribution aspects. Where technology implications overlap or complement in an important manner between the two systems, this will be taken into account.

The monitoring of the drinking water distribution occurs in two categories:

- Hydrological Parameters
- Biochemical Parameters

The hydrological parameters are monitored for billing and performance analysis purposes, while the biochemical properties of the water are measured to ensure the health of the public and validate treatment operations. MWRA employees can view the results of the monitoring and measurement in near-real time by accessing a centralized server. Using an application called “Process Book”, the data can be viewed either graphically or numerically and under varying time aspects. From the captured data, weekly and monthly reports are generated and circulated through the various departments to ensure an agency-wide understanding of the current system performance.

This chapter describes the scope of the monitoring, how wireless monitoring is used, and how this large regional system interfaces with a local distribution network.

3.2 Hydrological Monitoring

Hydrological monitoring measures the status of a water distribution system in terms of:

- Flow
- Pressure
- Velocity
- Precipitation
- Reservoir levels

Monitoring is the reporting of any of these parameters to a SCADA system (Supervisory Control and Data Acquisition), while metering records the quantity delivered to a specific location. Metering occurs for two reasons: Custody transfer of the water and system performance analysis. The custody transfer meters are located at the entrance into a town's distribution network and are read as frequently as every two seconds. The MWRA has roughly 160 meters to monitor, most of which are at custody transfer locations to towns. Accordingly, most the meters are large commercial meters, which are more susceptible to accuracy deterioration than the smaller meters used for individual households.

From a management standpoint, knowing the consumption of a town is important since it serves as an indicator for the pressure requirements of a system. Pressure monitoring may provide instantaneous information, however, it is necessary to be able to anticipate a consumption rate for a distribution main in order to adequately size reservoirs, tanks, pipes, flows, etc. to provide the pressure necessary to meet usage and flow velocity requirements.

Meters, pressure monitors and reservoir level sensors report their readings over leased lines from the public switched telephone network (PSTN). Although a reading can be prompted at any given time (data pull), the standard operation is a data push to the server at pre-programmed time intervals. The largest problem the MWRA has encountered in operating their meter reading system is failure of the leased lines (see Figure 3-1). The causes for this vary; however, a frequent problem is that since the MWRA's leased lines are not used for telephone purposes, they do not have a dial tone. When telephone company employees are searching for empty lines to make a new connection, often they simply "listen" to the line and upon not hearing a dial tone, they assume that the line is unused and put a new connection on that line.

The MWRA issues two internal documents to keep track of system status and performance. The "Weekly Water Quality Report" and "Monthly Management Indicator Schedule" assemble data to enable operations staff to identify water source fluctuations, monitor process performance, and identify unusual patterns or anomalies. In regards to the metering operations, both of these documents

include detailed information regarding reservoir levels and precipitation, both of which are direct contributors to system pressure and flow conditions. Since the amount of water metered is dependent, among other things, on the water pressure, operations personnel relies on these documents to give an indication of how to react to current situations in the distribution network.

Throughout the MWRA's distribution system, there are pressure alarms that are customized for each location. If the pressure falls below a pre-determined value, the sensor will set off an alarm at the centralized operations control center. Although the MWRA's authority does not extend into a town's distribution network, it is possible for them to notice an unusual pressure drop in a distribution main further down the line than their apparatus reaches. Under these circumstances, the town water resources manager could be contacted for further action.

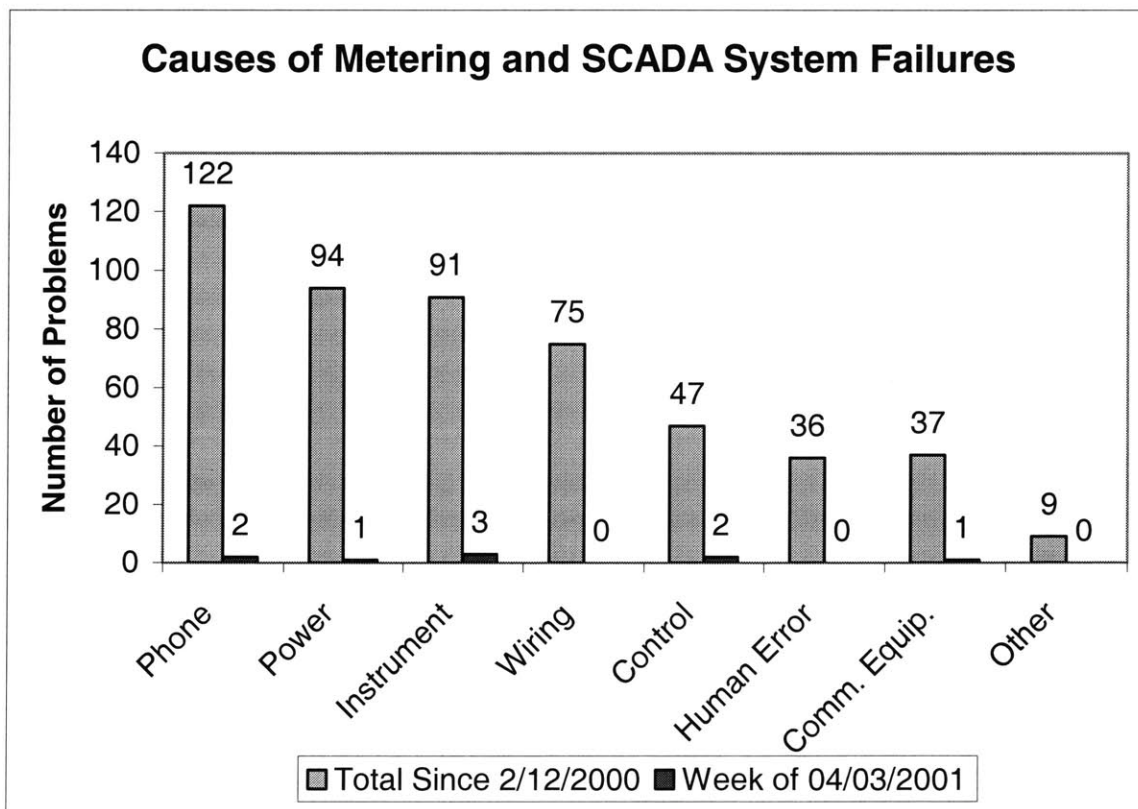


Figure 3-1 - MWRA Failure Reports

Reproduced from MWRA's "Weekly Water Field Operations Report",
 March 30, 2001, p.7

3.3 Biochemical Monitoring

Biochemical monitoring occurs on every possible parameter that could affect the health of the public. Since the MWRA deals not only with the supply of the drinking water, but also the treatment of sewage, the water quality indicators that must be monitored extend past ensuring drinking water quality to also ensuring that effluent from their treatment facilities is of acceptable quality. The apparatus and processes of determining the water quality parameters are beyond the scope of this document, however, some meaningful information regarding the monitoring of these parameters has been collected.

The six general groups of water quality indicators are:

- Microbial
- Turbidity & Algae
- Corrosiveness (pH and alkalinity)
- Disinfectant residual
- Disinfectant by-products
- Mineral Analysis

Some of the specific parameters that are monitored are:

- Temperature
- PH & alkalinity
- Chlorine residual
- Turbidity
- Dissolved oxygen
- Total dissolved solids
- Fecal coliform levels
- Algae
- Giardia cysts
- Trihalomethanes (THM)
- Haloacetic Acids (HAA)
- Total organic carbon (TOC)
- Odor & taste (non-electronic monitoring)

The effect of an introduction of chemicals into a distribution network cannot be monitored immediately, but must be performed downstream after some mixing has taken place and the concentration of treatment chemical has balanced throughout the flow. For this reason, chemical monitors are frequently installed downstream from the point of introduction. The amount of the treatment chemical used needs to be proportionate to the flow, so it is also necessary to be monitoring the flow either directly at or shortly before the treatment location.

Most chemical sensors are some form of a probe that is permanently immersed in the flow and registers the concentration of a substance. An inherent problem is the build up of clogs and biofilm on the sensor, which can result in inaccurate readings. Because of this, there is some reluctance to put a lot of trust in the readings from chemical monitoring systems, unless they are frequently and

thoroughly maintained. In a system as large as the MWRA's, it can be difficult to ensure maintenance of such a nature.

Historical information is considered necessary to acquire a more reliable profile of the water quality. An experienced operations manager needs to evaluate the report of the concentrations of all dissolved chemicals as well as the quantity of the treatment chemicals used. Judging by the amounts of treatment chemicals used in the past under similar circumstances, he/she can ascertain how much trust can be given to the monitored profile.

The MWRA has a branch that deals specifically with industrial effluent and treatment plant effluent monitoring called TRAC (toxic reduction and control). For the most part, their duty is to locate dissolved and particulate metals in both sludge and effluent. Since sludge is frequently used as fertilizer, the presence of such substances needs to be monitored. One of the largest sources of metal contaminants is the use of molybdenum as a corrosion inhibitor in water that flows through cooling towers for large air-conditioning systems. The hotter the season, the higher the concentrations of molybdenum in the incoming water of a treatment plant will be.

The hardware used to transmit data from biosensors to the central operations control center is either leased lines from PSTN, hardwired cables (local in treatment facilities) or radio transmission. There have been attempts to install systems that use PSTN without leased lines, which instead use a modem to dial up the central control center and report a possible alarm status. However, these attempts have been unsuccessful since there were so many system failures that ultimately significantly reduced trust in the system. Some monitors did not report alarm status due to the power source running out (in the case of a battery driven unit), others reported alarm status erroneously. The MWRA is currently searching for a more reliable system that features a constant power source.

Corrosion monitoring of drinking water as well as sewage has been a crucial to the system performance determination, since it gives an indication of how soon maintenance work will need to be performed. This is an important consideration for the MWRA, and so an efficient and reliable system has been installed to monitor the aggressiveness of the environment. In addition to chemical sensors to acquire the pH and alkalinity of the water, this system includes sensors on the monitoring equipment to keep track of the performance of apparatus. Parameters such as temperature and vibration of pumps and meters are monitored to give an indication of system performance.

3.4 Simulation/Modeling

The MWRA uses simulations of their water operations for both of their networks, the distribution and sewage networks. Modeling is used for the hydrological design and planning of future conditions or system expansions. For the sewage network, it is also used to estimate the current system state and performance, since there is fewer accurate flow data available. Monitoring sewage flow is complicated since large suspended solids in the water can collide with measuring apparatus and break or clog the machinery.

Modeling the current performance of the network is useful when a measuring device is malfunctioning or other unknown factors become important. The known measurements and boundary conditions are used to create a simulation of the network and estimate unknown values at another point. Although known readings for flow are taken into account, the MWRA does not base their simulation entirely on this data. Measurement inaccuracies as well as erroneous information on the degree to which a valve is opened could easily produce inaccurate results from the simulation. To avoid these problems, the MWRA's modeling software calculates the pressure gradient throughout the network and bases the resulting estimates on this gradient instead. Taking into account the boundary conditions and the flow data, the simulation can neglect flow readings that are not within a reasonable threshold of where they should be based on the valve and reservoir conditions.

For the purposes of modeling a hypothetical situation such as partial network shutdowns due to maintenance, a proposed system expansion, or a rain storm, the MWRA collects data from their flow meters and attempts to use it directly (rather than calculate a gradient). Historical data can provide useful information on a town's usage or seasonal variations in inflow and outflow, and for estimating future conditions at a distinct location, operational managers at the MWRA tend to put more trust in this data than the results of a simulation. The collected data is fed into a model that calculates a system response and allows the manager to take appropriate action. Many times construction schedules and plans for line diversions to accommodate large inflows are made based on these simulations.

The MWRA's simulations are also used to design system expansions and perform optimization studies. Parameters that need to be taken into account for these applications include seasonal water table variations and precipitation data, flood plain modeling and backwater curve analysis. The network needs to be designed to meet safety conditions, but all within reason. Design for a 1000-year flood might be safe, but not cost efficient, while design for a 1-year flood is cost efficient, but not safe. The resulting design will thus have to take into account that at (rare) times, the network will be totally overloaded and unable to drain the water shed quickly enough to avoid flooding. An analysis of how much water the

system can handle and the magnitude and time-frame of build-up of overflow behind the system produces a backwater curve that allows the estimation of how long it will take to drain a flood and how high a flood level can be. In addition, the simulation can attempt to find an optimal path for draining the watershed, based on current network load and paths.

3.5 MWRA's SCADA System

The MWRA has a SCADA (Supervisory Control And Data Acquisition) system that runs on a central server and is administered by their MIS (Management Information Systems) department. The system runs the engineering and monitoring applications and allows the staff to access data and applications remotely. It manages both the sewage and drinking water operations and monitoring. There are two Oracle databases on the server; one holds the monitoring data for the drinking water and sewage systems, while the second database holds the TRAC (Toxic Reduction And Control) data.

The tool that staff uses to view the data on the SCADA system is called Process Book and utilizes a tab-based navigation system to view the results of the different biochemical and hydrological monitoring parameters. Process Book pulls together data sources from various different locations on the server to display the results and supports instant querying of some sources (most measurement parameters are updated every two seconds). For staff who need to work from their home, Process Book allows remote access with a dial-up modem.

The MWRA plans to make their SCADA system accessible over the Internet, which they anticipate to be useful to town water managers throughout eastern Massachusetts. A web site will display selected parameters to the MWRA customer who accesses it through a web browser. The system data are available in near-real time (the system updates every 15 seconds) and are viewable in number and graphic format.

The MWRA SCADA system is still under construction and it not anticipated to be fully completed until 2007. The process of designing the system began in 1986, but a full proposal was not completed until 1993. Installation began on pilot projects in 1996 and since then slowly but surely all apparatus and departments are switching over to the new system. This process is taking long because the operational area and complexity of the MWRA system is so large and in many cases customized applications need to be build to manage the data sources and transfer.

4 Arlington Remote Monitoring Software General Description

4.1 Overview

The Department of Public Works of the town of Arlington, Massachusetts, decided in the year 2000 to improve their existing water meter reading process by installing wireless transmitters on all the roughly 12,200 meters throughout the town. The town's incentive to install this system was to improve the knowledge and understanding of the local water usage, optimize the billing service, and potentially lower the operating cost of the water department.

Prior to the installation of the wireless monitoring system, the town received their readings by having a town employee connect a handheld meter-reading gun to the meter twice a year. This collected one reading that related six month's worth of flow. Upon installation of the transmitters, the town will be capable of taking readings as frequently as every 10 minutes, allowing a much higher accuracy in its water supply related analytical and administrative tasks.

The only type of reading the town monitors is the flow through each meter. There is currently no plan to install any other types of reading systems (pressure, biochemical composition, etc.) and for the purposes of the town, there will most likely not be a need for any other type of metering service. Although the Arlington's water supply network is divided into pressure zones, the pressure is maintained and monitored by the regional water resources authority and so relieves the town of Arlington from having to monitor these parameters. In addition to the consumer meters, the town will also monitor the flow through pressure relief valves that allow cross flow between the pressure zones.

Water meters register the flow (normally an analog reading) and a Hexagram STAR AMR® wireless transmitter sends the reading to a cellular tower. There are five of these towers placed on public buildings throughout the town that provide the cellular coverage. The readings are stored on a Hexagram server until they are downloaded to a MySQL database, which in turn resides on a server that is dedicated to the water management system. Also on this server is the analytical software required to perform several interpretive functions (see below). The town Treasurer's current billing system, the ICS (Integrated Collection System), operates an Informix database that receives monthly data transfers from the MySQL database through the analytical software.

This high-level system architecture is shown below:

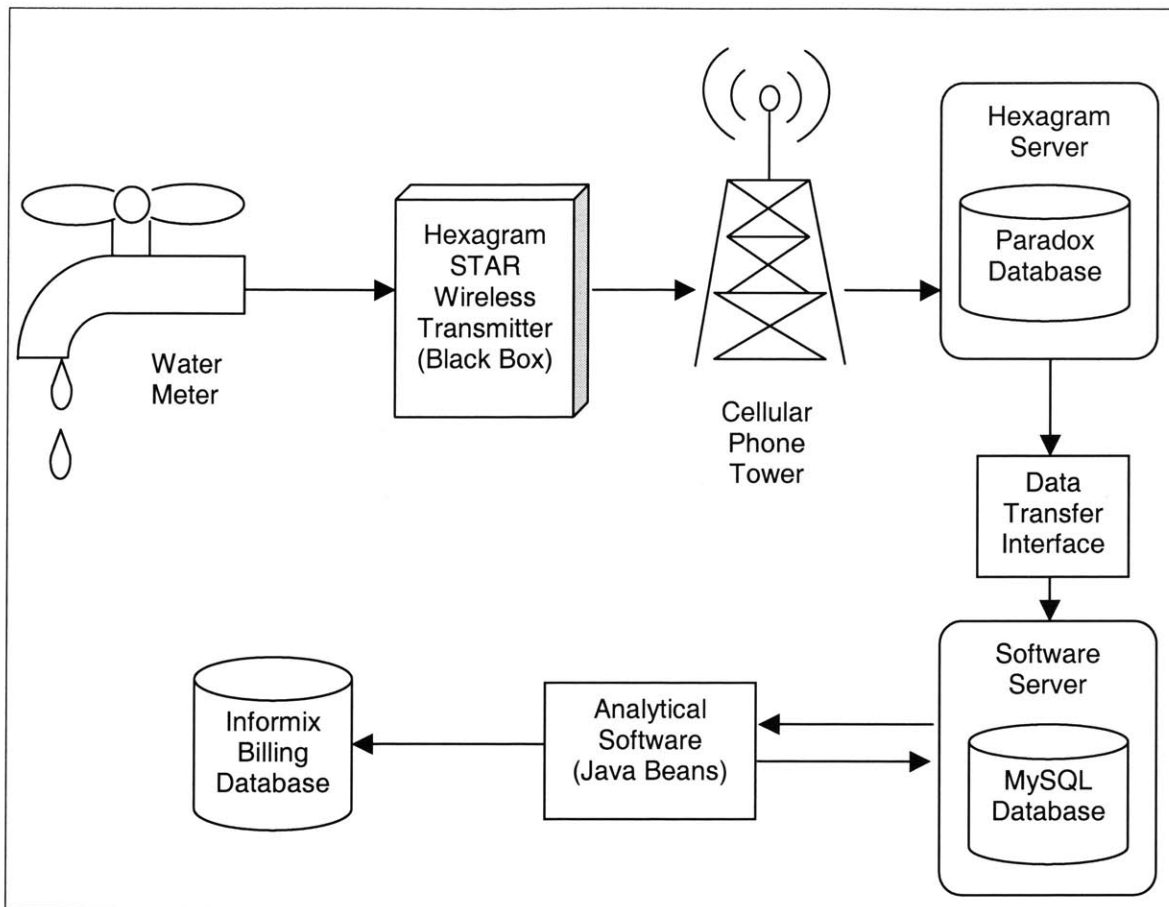


Figure 4-1 - High-Level Architecture of Water System

The Hexagram system is set to a default of two readings for each meter per day (although this can be modified under some circumstances). The readings are sent to the Hexagram server and are downloaded to the MySQL database twice a day as well. The Water Department can then perform a variety of analytical and administrative tasks with this data in order to get a good grasp of the town water use.

4.2 Software Functionalities

The main functionalities that the analytical software implements are:

- Bill Preparation
- Meter Watch
- Leak Detection
- Usage Pattern Analysis
- Meter Accuracy Validation
- Water Theft Detection
- Administrative Functions

4.2.1 Bill Preparation

In order for the Treasurer to prepare billing statements, a monthly reading for each account needs to be sent to the Integrated Collection System (ICS). This data set should be clear of all incorrect readings, missing readings or otherwise inoperable values. The billing analysis function allows the user to review the readings and detect suspicious or missing readings and take corrective action to ensure data integrity.

Upon selecting a time period over which the data integrity check is to be run, the software will scan through every reading for every account during that time period and attempt to find singularities that meet a preset condition, such as:

- Missing readings
- Estimated readings
- Abnormally high readings
- Abnormally low readings

If the system finds any readings that meet these criteria, the user is prompted to correct this situation or indicate to ignore it. Once the user decides all meter information is correct and accurate, the daily readings are aggregated into a monthly total for each meter. The monthly meter information is then sent to the ICS for billing purposes.

In order to find missing readings, the date/time stamp of the meter reading table is used. If the time gap between one reading and the next is larger than 14 hours, for example, a missing reading has occurred, since meters are supposed to be read every 12 hours.

Estimated readings shall have a flag column in the meter read table. Every time an estimated reading is entered into the system, that flag is set as well. Searching for estimated reading simply lists the readings that have this flag set.

Searching for abnormally high or low readings compares the average daily use in one time period to the average daily use in another time period and checks whether the difference is beyond a threshold. For example, the average daily use in the first week of a month is compared to the average daily use of the previous four weeks and the difference is checked against the threshold.

4.2.2 Meter Watch

The software allows the user to monitor a specific meter. When a meter is placed on “watch”, all the readings for that meter within a time period (specified in the

“Meter Watch Config” Web page in the software) will be listed for the user to see and interpret. This functionality is intended for the experienced water resources manager of the town to be able to search through an account to inspect its correctness. Possible reasons for placing a meter on watch could be suspicion of water theft, unusual usage patterns, and security precautions for vacationing citizens.

4.2.3 Leak Detection

There are two possible types of leaks that this system checks for: The first is a leak between the source of the water and the consumer meter, and the second is a leak on the other of the consumer meter, the “house side”. The source of the water flowing into the Arlington supply network is the Massachusetts Water Resources Authority, or MWRA. The two types of leak detection methods have thus been termed:

- MWRA Comparison
- Individual Leak Detection

MWRA Comparison

Arlington is subdivided into three pressure zones: the high, low and intermediate pressure zones. The MWRA can provide a reading of how much water has flowed through each of their meters that serve the respective pressure zones. The Arlington Department of Public Works can then use the MWRA Comparison functionality to aggregate all the use of the Arlington meters by pressure zone and compare these values to the MWRA flow. It needs to be taken into account, however, that there are certain unmetered and thus unaccounted for uses, such as fire hydrants.

Periodic monitoring of the relationship between the Arlington meters and the MWRA meters should allow experienced users to detect abnormal conditions in the flow between the source and the consumer. Although the software can aggregate and display the flow through the town, human experience is still necessary to interpret and respond to unusual usage patterns and leaks. For this purpose, historical data can be shown alongside current data in graphs that indicate the flow conditions. The user can set a threshold value as well, which lets the system notify the user when the difference between the MWRA flows and the Arlington flows is above or below a certain percentage.

Individual Leak Detection

The individual leak analysis attempts to identify leaks on the house side of the meter. The system will analyze individual meter readings to find suspect activity that may be the result of a leak. The possible scenarios that the system automatically checks for are a drastic sudden increase in use and an increase in average daily use above a certain percentage.

Should an account meet the drastic sudden increase criteria, the Department of Public Works may take action immediately and contact the customer. If the threshold for increase in average daily use is exceeded, that account can be put on a night-time watch, in which the meter is read every 15 minutes between the hours of 1 a.m. and 4 a.m.

Theoretically, if there is no leak on the house side of the meter, the reading should be zero when the customer does not use any water. Since it is reasonable to assume that in at least one 15 minute period between the hours of 1 a.m. and 4 a.m. no use will occur, the night-time watch can be very effective at ruling out suspect leaks. Should the results of this watch be conclusive, the customer may be notified of a potential leak at their location.

4.2.4 Usage Pattern Analysis

This functionality provides the Water Department with a tool to aggregate groups of meters by some of their parameters and view their usage over a specified time period, in order to determine usage patterns. Parameters that can be used to filter groups apart are such things as meter size, pressure zone, and customer type (commercial or residential).

The system displays the use of all the meters in the group aggregated together and calculates the average use over a user defined time period. The user is able to search for trends in the usage of meter groups and utilize this information when estimating a reading or trying to establish pricing guidelines. For example, the program could sum up all the use for all 5/8" diameter residential meters in the month of April and calculate the average use per day in that time period, then compare that value to the average daily use in the previous month(s).

4.2.5 Meter Accuracy Validation

As a water meter ages, it has a tendency to under-register. This process depends on the quality of the water, specifically the amount of colloids and the corrosiveness. Water meter manufacturers provide guidelines how to detect or

estimate this effect through the means of deterioration curves that describe the meter's performance over time with respect to a given water quality.

In order to find a deteriorated meter, the system searches for continually smaller readings over a long period of time. Naturally, changes in tenancy that result in a different number of consumers per meter make this difficult, since a new usage history needs to be established before the deterioration trend can be determined. The user has the option of entering a threshold and a time period over which to run the analysis. The program compares the usage of one time period to the usage of previous time periods of the same lengths and attempts to find downward trends above the threshold. Theoretically, if the time period is long enough and a sufficient amount of data is available, a deteriorated meter should be detectable.

4.2.6 Water Theft Detection

Although it would be nice if this functionality were not needed, it is a necessary appendage, since some people will try to get their water service for free. Tricks used to do this include circumventing the meter by placing an auxiliary pipe around it, tampering with the meter, or turning the meter around to reduce the reading. However, with a little savvy engineering much of this tomfoolery can be detected.

An analysis is run on the collected data to find meters with suspicious activity. The program searches for four types of behavior:

- **Continuous zero usage**
Search for all meters with a period of continuous zero usage in a defined time period
- **Negative readings**
Search for readings that are lower than the previous reading (by a user defined threshold)
- **Broken seals**
When the maintenance staff finds a meter with a broken seal, this is indicated in the database. The user is capable of retrieving an accounts history of broken seals to search for a recurring breakage.
- **Loss of physical connection**
The system queries the database for each occurrence of the loss of physical connection of the wireless device. When the wireless device is disconnected from the meter, the device will send a signal to the Hexagram server indicating this action. Each broken seal event per meter will be recorded in the results during the defined time period.

Meters exhibiting any of the above types of behavior will be written to a results table when the analysis is run. The user may then take appropriate action for each meter, if necessary. The available actions include editing information in the database and placing meters on a more detailed watch.

4.2.7 Administrative Functions

In addition to the methods that allow performing analytical tasks with the collected data, the users of this system need to also be able to administer the data source. New accounts will need to be created, as well as old ones deleted, for example. For this purpose, the software has a number of administrative functions that allow the modification of:

- Account information
- Meter information
- Meter sizes
- Property types
- Hexagram readings
- Meter readings
- MWRA readings

4.3 Software Implementation and Applied Technologies

The entire user-interface is programmed in Hypertext Mark-Up Language (HTML) in order create an environment that can easily be adapted in the future to let the town's residents access their account information over the Internet. The HTML files are resident on an Apache Tomcat Server running a Red Hat Linux operating system. The software architecture is outlined in Figure 4-2.

Within the HTML pages are embedded Java Server Pages (JSP) that send requests for queries to the database via a Java Bean. The bean uses the Java Database Connectivity (JDBC) interface to send Structured Query Language (SQL) queries to the MySQL database. The results of the query are then sent back to the JSP and displayed.

For a closer description of the Arlington Remote Monitoring Software user interface, please refer to *Investigating the J2EE Software Architecture for Infrastructure Monitoring: A Water Metering Case Study* by Mameet Khanolkar.

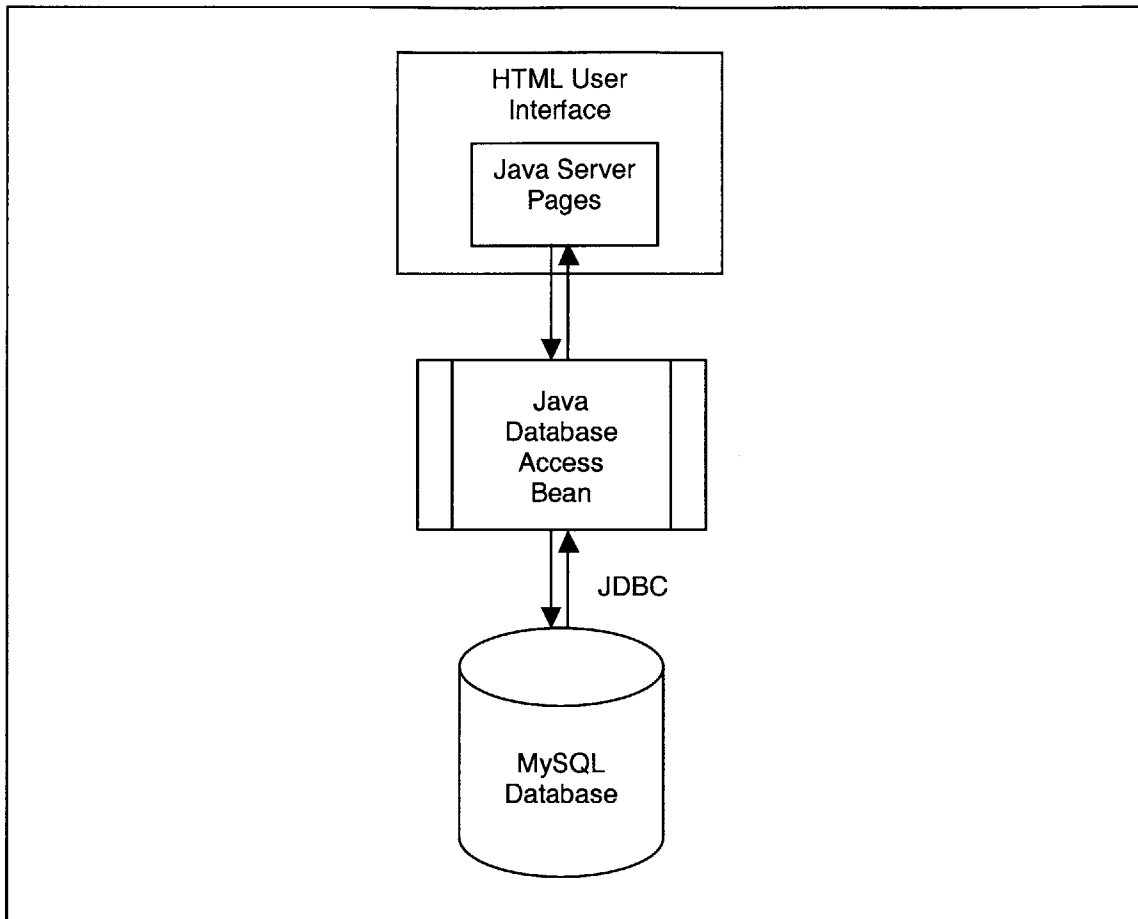


Figure 4-2 - Software Architecture

To illustrate: The user wants to determine the status of an account and navigates from a main menu screen to the account administration functionality. An HTML page displays a form that shows all fields needed to clearly identify a specific account and prompts the user to fill in that information. Upon entering this data and clicking on the "Submit" button, the data is sent to a JSP page that relates the entered information to the actual column names in the database and sends a request to the Java Bean, in the form of a SQL statement that searches for an account with the information entered by the user. The Java Bean opens a connection to the database, and sends the request on. The database performs the query, finds the account in question and sends the results back to the bean. The bean passes the results on to the JSP and closes the database connection. Now the JSP is processed by the server and sends back to the user a simple HTML page that displays the account information the user requested.

5 Software Architecture and Functional Model Description

This chapter explains the architecture of the Arlington, Massachusetts, Remote Water Meter Reading software as built by Bradford Butler, Mameet Khanolkar and the author. For a thorough description of the database design, please refer to *The Business Transformation Effects of Information Technology* by Bradford Butler; for an elaborate description of the workings of the Java Server Pages and Java Beans, please refer to *Investigating the J2EE Software Architecture for Infrastructure Monitoring: A Water Metering Case Study* by Mameet Khanolkar.

5.1 Usage Scenario

5.1.1 User Profile

The software to manage the daily readings from the wireless meter reading system was developed exclusively for the water department of the Department of Public Works of the town of Arlington, Massachusetts. The only anticipated users of this software program are the town water manager and other municipal engineering personnel. Although it is foreseeable that portions of the functional site will be available to the public and that the software will be open-source, no specific accommodations have been made for these circumstances.

Use Cases

Use cases will be described for the implemented functionalities of the software program:

- Administration of the data tables of the water software
- Configuration and operation of the reporting tools
- Viewing of the report results for the system functionalities
- Bill Preparation
- Usage Pattern Analysis
- Meter Watch Functionality

5.1.2 Administration of the Data Tables Use Case

Three tables constitute the main information for operating the monitoring system. The Account table holds all of the information about a user account, which can be associated with one or more meters in the Meter table. The readings for the meters are stored in the MeterRead table.

These tables are administered using Java server pages that connect to the database on the back end of the system. Four standard functions are available, regardless of table type or content: Adding, deleting, updating and finding a record in the database. A sample use-case diagram is shown for the Meter table; the diagram for all other tables in the database is identical.

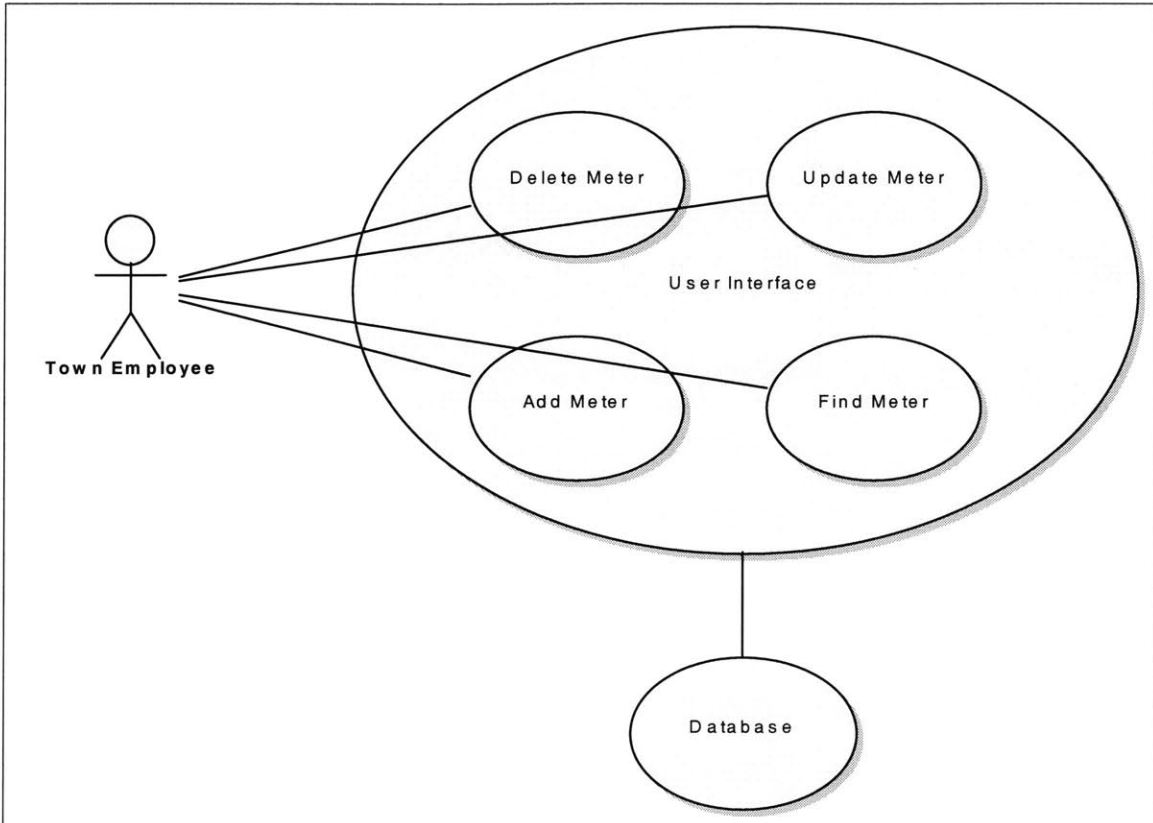


Figure 5-1 - Use Case Diagram for Administration of Meter Table

5.1.2.1 Administrative Functions Narrative

The process of administering any of the tables in the database is identical for all of them, so only one sample process will be described here.

From the main menu screen, the user navigates to the Admin Meter administrative functionality and is presented with a menu of options on the left hand side of the screen. These options are:

- Find
- Add
- Update
- Delete

To the right of this menu is a default function – the “Find Meter” function. The user can either proceed with finding a meter, or select any of the other options from the menu.

Find Meter

To find a meter, the user can enter information to query the Meter table for in any of the provided fields. These fields correspond to the columns in the Meter table in the database. For the case of a Meter Find, these fields are:

- Account Number
- Suffix
- ARB Number
- Meter Size
- Meter Type
- Seq Number
- Location
- Meter Number
- X Multiple
- Date Installed
- Read Instructions
- Status (Active or Deleted)

The user enters information in any or all of the fields and clicks on the “Submit” button at the bottom of the screen to start the search. The database is then queried for meters that fit the conditions as the user has entered and if any such meters are found, their information is displayed on the screen.

Add Meter

The user clicks on the “Add” button on the navigation menu. He/She is then presented with the same list of parameters that are listed in the “Find” functionality above. However, in this case all fields except “Read Instructions” must be filled out. This data will be entered directly into the database, so the user should not enter partial information at this stage, but fully enter all available data on the meter.

Once the user has finished entering the meter data, he/she clicks the “Submit” button at the bottom of the screen and the following Java Server Page performs some error checking, such as ensuring that the account number and suffix entered do not exist already and that important fields have not been left blank. If there are no problems with the data as entered, a connection to the database is established and the Meter table is amended. The user receives a confirmation screen that displays the information as it has just been entered.

Update Meter

The user clicks on the “Update” button on the navigation menu. He/She is then presented with the same list of parameters that are listed in the “Find” functionality above. Once the user has finished entering the meter data, he/she clicks the “Submit” button at the bottom of the screen and a connection is established to the database to retrieve the meter information.

The following screen presents the user with again a list of input boxes, however, the current information for the meter entered in the previous step is filled in these fields. The user now has the option to change any of the information in these fields. Once done with this, he/she clicks the “Submit” button at the bottom of the screen. The following Java Server Page performs some error checking, such as ensuring that the account number and suffix entered do not exist already and that important fields have not been left blank. If there are no problems with the data as entered, a connection to the database is established and the Meter table is updated. The user receives a confirmation screen that displays the information as it has just been entered.

Delete Meter

The user clicks on the “Delete” button on the navigation menu. He/She is then presented with the same list of parameters that are listed in the “Find” functionality above. Once the user has finished entering the meter data, he/she clicks the “Submit” button at the bottom of the screen and a connection is established to the database to retrieve the meter information.

The following screen presents the user with the data for the meter just entered in the previous step. He/She is asked to confirm whether this is the meter to delete. Upon clicking on the “Delete” button, a connection to the database is opened and the “Status” field for the appropriate meter is set to “Del”, indicating that this meter has been deleted. The user receives a confirmation screen that displays that this meter has been deleted.

5.1.3 Configuration and Operation of the Reporting Tools Use Case

To view analytical reports of the meter reading data, reporting tools are created for every system functionality. Each report has a set of parameters that configure what data is searched for when the analysis is run. For this purpose, the operation of the software is divided into two logical steps: First the user configures the parameters for a certain report type and runs the report; and second the user views the results of the report.

5.1.3.1 Reporting Tools Narrative

- The user navigates to the report configuration screen where he/she selects what analysis shall be run and what the parameters for that analysis are
- The user selects the “Submit” button which saves the entries made into the report configuration and returns the user to the main menu screen

Schematically, the use case description of this process is as follows:

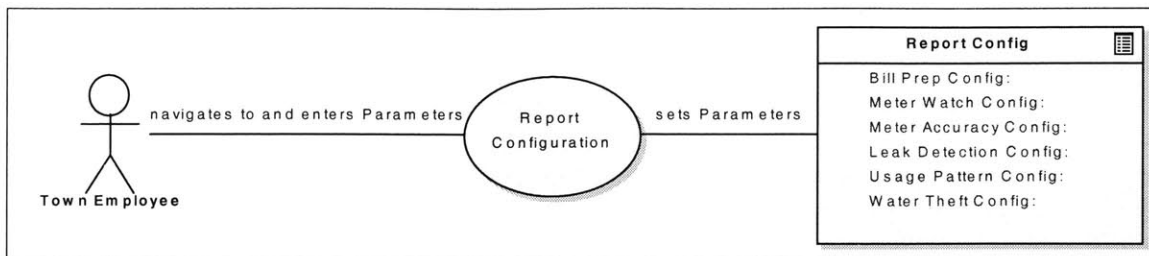


Figure 5-2 - Report Config Use Case

5.1.4 Viewing of the Report Results Use Case

Report Results Narrative

- The user selects “Run Analysis” which performs the functions selected in the configuration parameters page
- The user receives a confirmation screen indicating what reports will be generated
- The user selects “Run Analysis” and the methods for the appropriate report are called
- Once the processes have finished, the user is returned to the main menu screen
- The user selects any of the results sections for the functionalities which have just been run

Schematically, the use case description of this process is as follows:

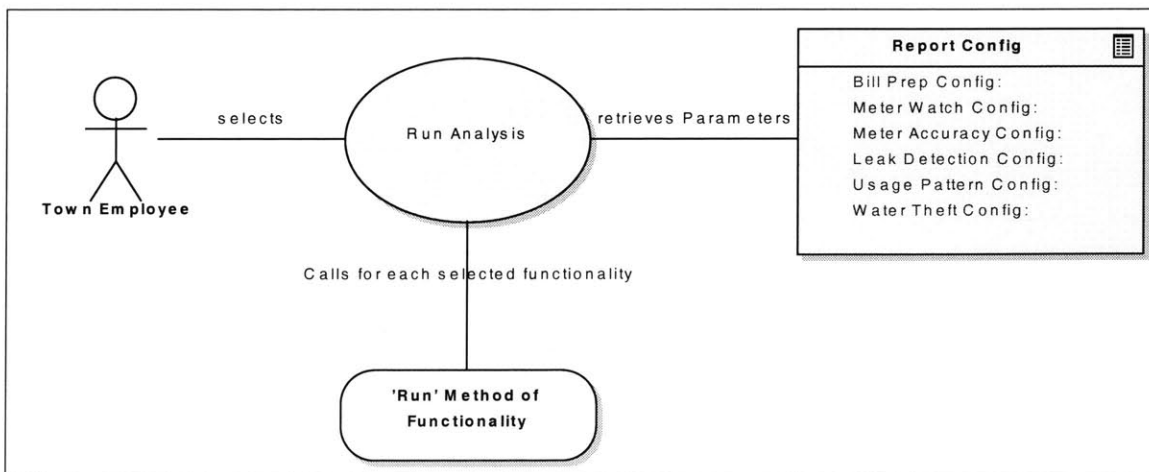


Figure 5-3 - Run Analysis Use Case

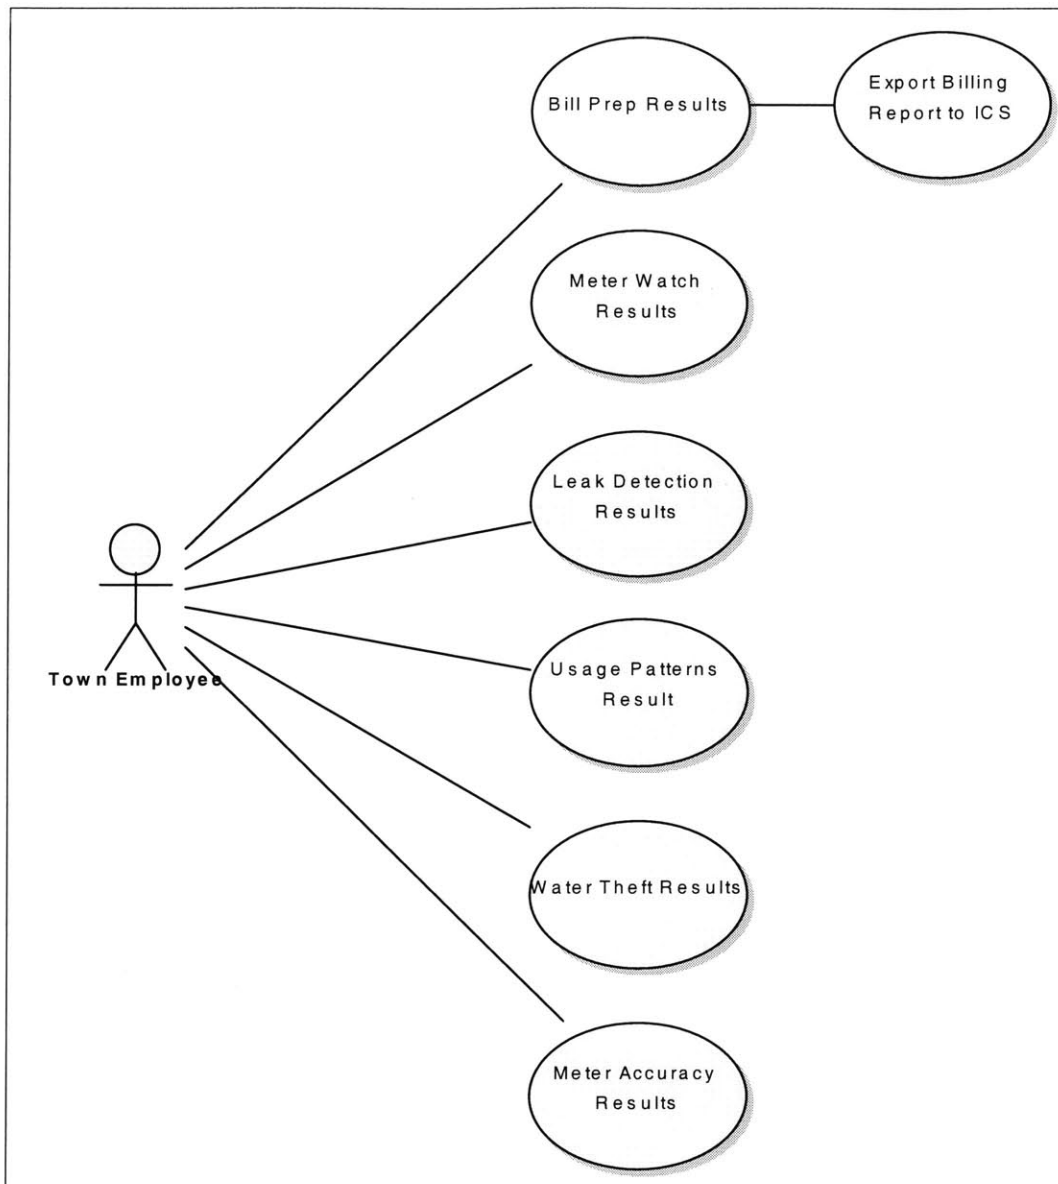


Figure 5-4 - View Reports Use Case

5.1.5 Bill Preparation Use Case

The bill preparation process searches for trends in the Meter Read table that suggest:

- Missing readings
- Estimated readings
- An abnormal increase in use
- An abnormal decrease in use

The main elements of the Bill Preparation functionality are described below:

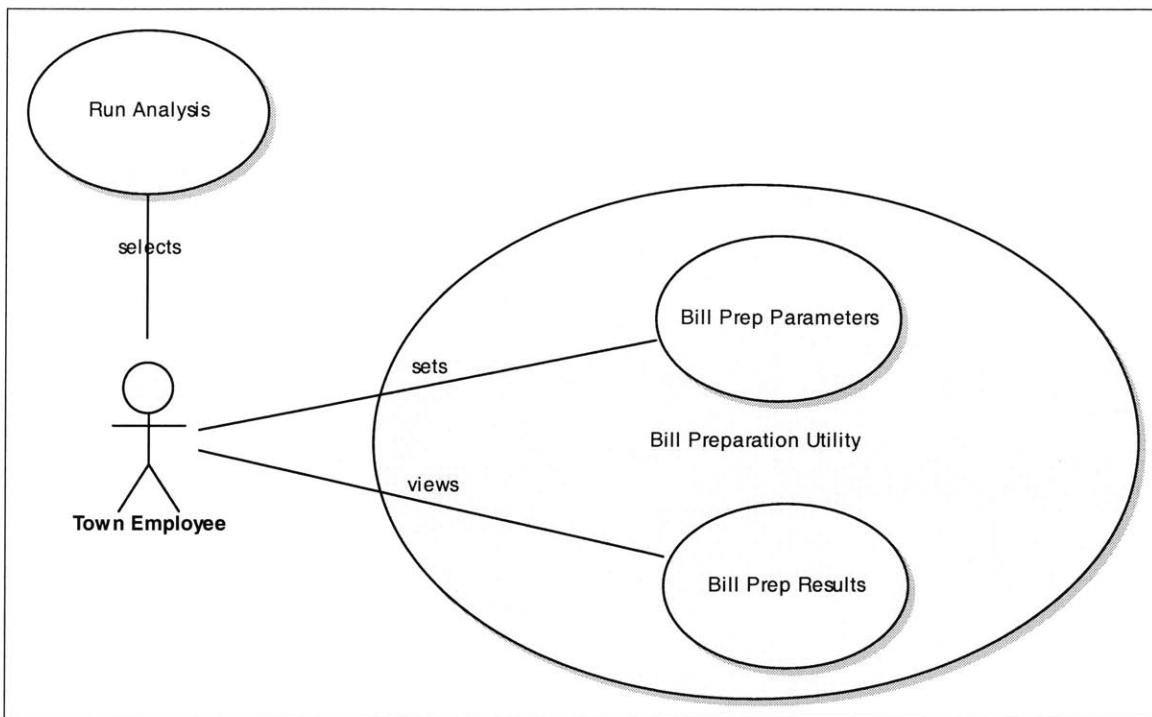


Figure 5-5 - Bill Preparation Use Case

5.1.5.1 Bill Preparation Narrative

The user navigates to the Report Configuration screen and enters the analysis end date. He/she then selects what functions of the Bill Prep analysis are to be run and sets the parameters for that analysis. The parameters that the functionalities take are:

Missing Readings:

- Time Period
The duration over which missing readings are searched for, counting back from the analysis date
- Missing Interval
The maximum amount of time between two readings. If the time between two readings is larger than the missing interval, a missing reading is reported.

Estimated Readings:

- Time Period
The duration of time over which estimated readings are searched for, counting back from the analysis date.

Abnormal Increase:

- **Current Time Period**
The time period in which the abnormal increase is to be found.
- **Past Time Period**
The time over which the average daily use is calculated. This use then serves as the baseline for the comparison for detecting an increase.
- **Percent Increase**
Only the meters whose average daily use in the Current Time Period is larger than the average daily use in the Past Time Period by at least this percentage are reported.

Abnormal Decrease

- **Current Time Period**
The time period in which the abnormal decrease is to be found.
- **Past Time Period**
The time over which the average daily use is calculated. This use then serves as the baseline for the comparison for detecting a decrease.
- **Percent Increase**
Only the meters whose average daily use in the Current Time Period is smaller than the average daily use in the Past Time Period by at least this percentage are reported.

Once these parameters are set, the user clicks on the “Submit” button and is returned to the main menu screen. After clicking on the “Run Analysis” button and receiving a confirmation screen, the analyses are performed based on the parameters just set in the Report Config screen. Once the processes have finished running, the user is again returned to the main menu screen. He/she now clicks on the Bill Prep Report option and is presented with the results of the analysis in tabular format, broken down by analysis type. Each meter is listed separately and a link is provided for the user to view the details of the respective meter. These details are in Microsoft Excel format, which easily enables the user to graph the data if necessary.

5.1.6 Usage Pattern Analysis Use Case

The Usage Pattern analysis aggregates groups of meters together and reports their water usage over a specified time period. The user navigates to the Report Configuration screen and sets the following parameters for the Usage Pattern analysis:

- **Meter Size**
Five checkboxes are provided to select the types of meter sizes that shall be reported on. Currently the available meter sizes are 5/8”, 1”, 1 ½”, 2”, and 3” diameter.

- Account Type
Meters can be grouped by their account type: Residential or commercial. Check boxes are provided for the selection of each type.
- Time Period
The time period over which the meter readings shall be aggregated, based on the end date of the analysis and counting backwards.

The Usage Pattern use case is shown in Figure 5-6:

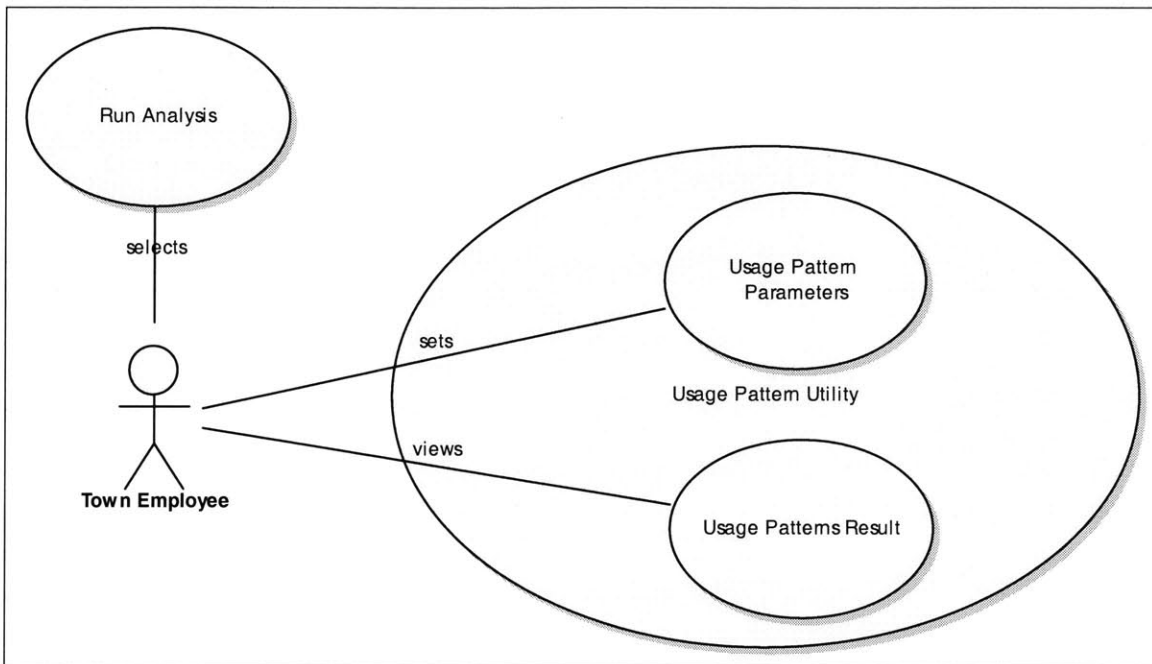


Figure 5-6 - Usage Pattern Analysis Use Case

5.1.6.1 Usage Pattern Narrative

Upon having set the meter group as desired, the user clicks on the “Submit” button at the bottom of the Report Config screen and is returned to the main screen.

After clicking on the “Run Analysis” button and receiving a confirmation screen, the analyses are performed based on the parameters just set in the Report Config screen. Once the processes have finished running, the user is again returned to the main menu screen. He/she now clicks on the Usage Pattern Report option and is presented with the results of the analysis in tabular format. The data that are listed are:

- The average daily use for the entire group over the time period

- The maximum use of any meter aggregated together over the time period
- The minimum use of any meter aggregated together over the time period
- The maximum average daily use for any one meter over the time period
- The minimum average daily use for any one meter over the time period

5.1.7 Meter Watch Functionality Use Case

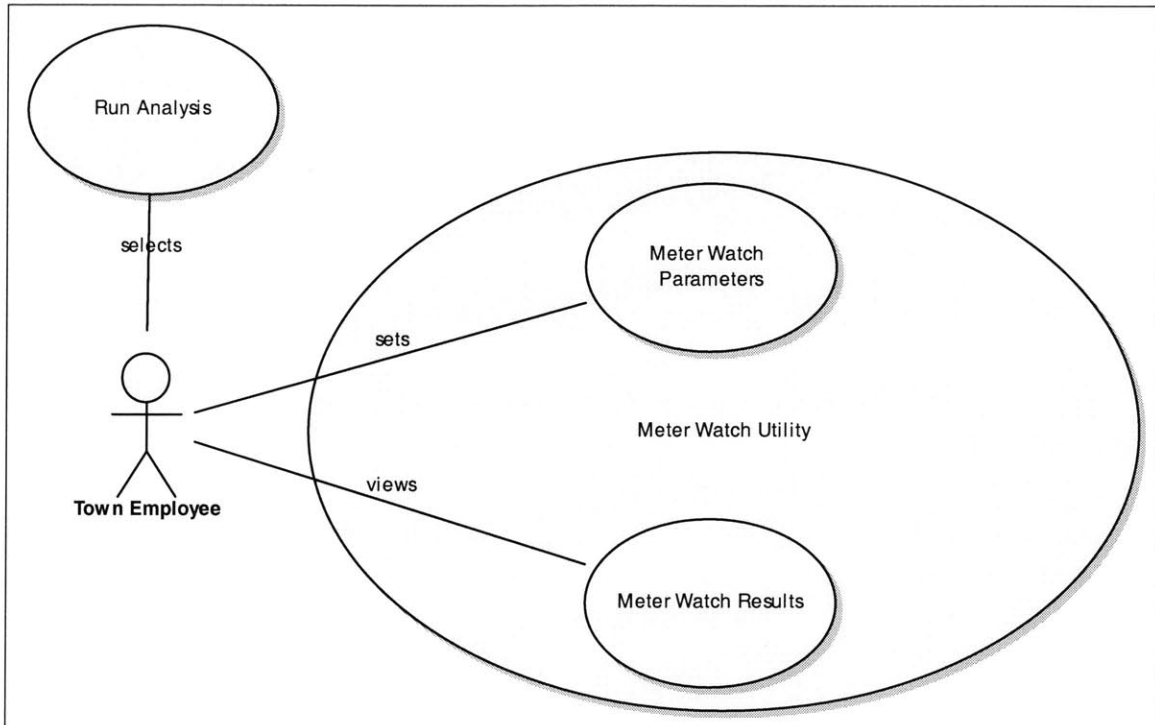


Figure 5-7 - Meter Watch Use Case

The Meter Watch process monitors the readings for a specified meter in the Meter Read table. The user navigates to the Meter Watch Configuration screen where he/she selects the meter(s) to be placed on watch by entering:

- Account Number
- Suffix
- Time Period (in days)
The number of days for which the meter readings will be reported, counting backwards from the analysis end date.
- Reason
A predefined reason for placing the meter on watch can be selected from a drop-down menu
- Comments
Any additional comments the user wishes to attach to the watch of this meter

Upon clicking on “Submit”, the user is returned to the main menu screen. To set the end date of the analysis, the user navigates to the Report Configuration screen and enters the analysis end date. He/she then selects the Meter Watch analysis check box and clicks “Submit” to set these parameters and is returned to the main menu screen.

After clicking on the “Run Analysis” button and receiving a confirmation screen, the analyses are performed based on the parameters just set in the Report Config screen. Once the processes have finished running, the user is again returned to the main menu screen. He/she now clicks on the Meter Watch Report option and is presented with the results of the analysis in tabular format, broken down by analysis type. Each meter is listed separately and a link is provided for the user to view the details of the respective meter. These details are in Microsoft Excel format, which easily enables the user to graph the data if necessary.

5.1.8 Future Enhancements

Although the software is currently not set up for an “Ad Hoc” query that the user can enter into the system, this functionality was originally intended and can potentially still be of use. The idea is to have the user enter an SQL statement that is sent to the MySQL database and the results are displayed in the same environment as the results of currently existing functionalities. Currently this is possible only if the user accesses the database directly on the server.

5.2 Functional Model and Description

This section will explain in detail the functions that have been implemented in the Arlington water meter monitoring system. These are:

- Bill Preparation
- Meter Watch
- Usage Pattern Analysis
- Database Access
- Administrative Functions

5.2.1 Bill Preparation Functionality

For a detailed description of the intent of the bill preparation functionality, please refer to Section 4.2.1. The name of this functionality is frequently shortened to simply “Bill Prep”. The classes of the Bill Prep functionality relate as shown in Figure 5-8.

The Bill Prep classes have the following functions:

- BPIncrease and BPDecrease search for an abnormal increase or decrease in usage and populate the BPRResults table in the database accordingly
- BPEst searches for estimated readings and enters the results into the database accordingly
- BPMiss searches for missing readings and enters the results into the database accordingly
- Bill Prep Report Java Server Page displays the results in HTML format by querying the database for the BPRResults table
- MBill creates a monthly billing report to be sent to the ICS (Integrated Collection System, the town of Arlington's billing database)
- Details retrieves the details for a meter that is specified by mouse-click from the listing of meters on the Bill Prep Report Java Server Page
- Report Bean stores all the report parameters as defined by the user in the Report Config function

5.2.1.1 Bill Prep Class Diagram

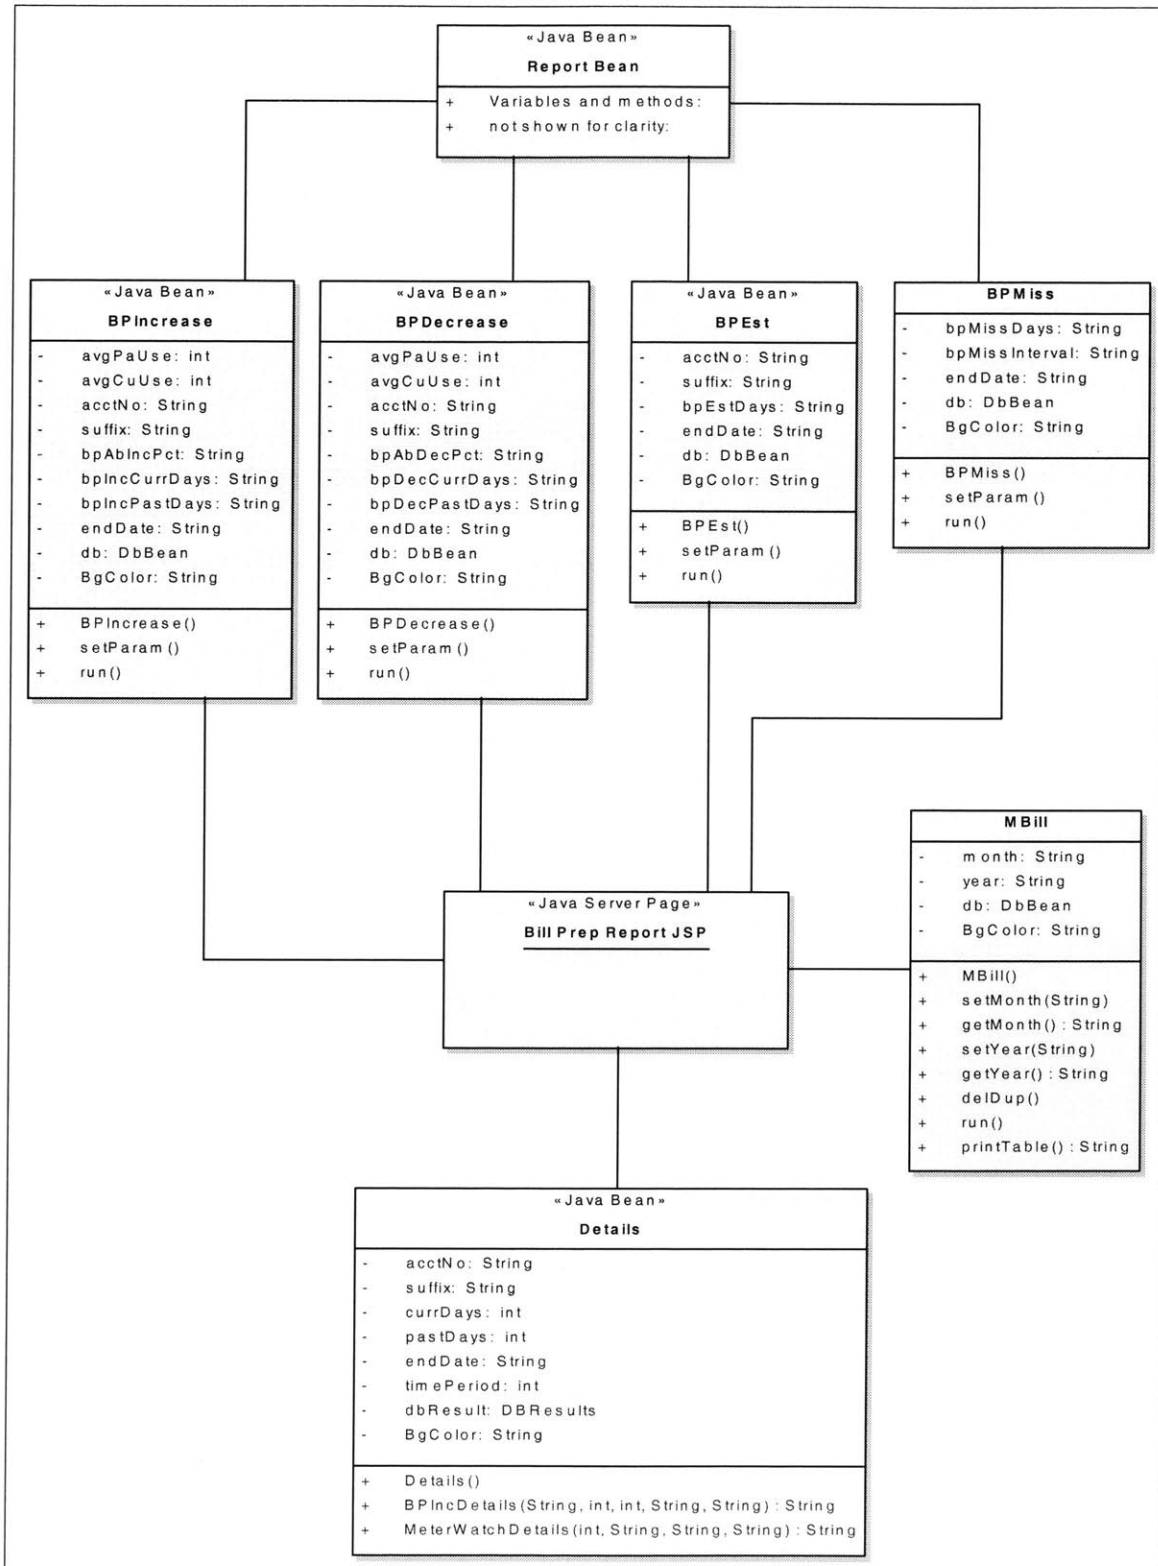


Figure 5-8 - Bill Prep Class Diagram

5.2.1.2 Bill Prep Sequence Diagram

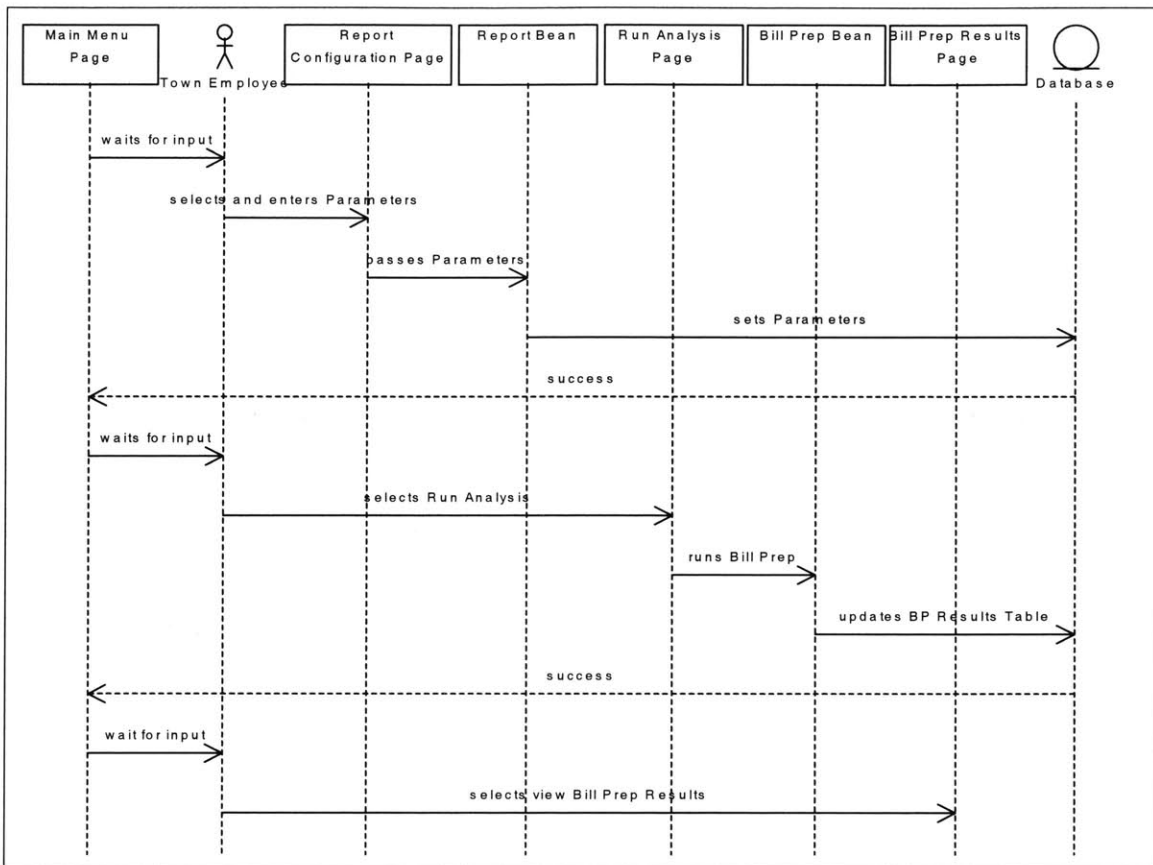


Figure 5-9 - Bill Prep Sequence Diagram

5.2.1.3 Bill Prep Interface Description

The Bill Prep classes for Abnormal Increase, Abnormal Decrease, Estimated and Missing Readings have two publicly accessible methods: set Param() and run(). The first takes the parameters from the Report Config table in the database and sets them in the object; the second connects the object to the database and performs the analysis based on the configured parameters. The class Mbill which creates the monthly billing report has separate methods for setting the parameters for the analysis, but again has a run() function to prepare the report.

The Details class has a function BPIncDetails() which takes as arguments the account number, suffix and dates for the bill prep analysis and returns all the meter readings for that meter on those dates.

The java server pages coordinate the calling of the methods, setting of parameters, etc. In some cases, the JSP's access the database prior to calling the methods, such as in the case of `BPIncDetails()`, where the JSP retrieves the dates for the details from the report configuration table prior to calling this method.

5.2.2 Meter Watch Functionality

For a description of the intent of the Meter Watch functionality, please refer to Section 4.2.2. How the classes of the Meter Watch functionality relate is shown in Figure 5-10.

The Meter Watch classes have the following functions:

- Meter Watch aggregates water readings for specified meters over user-defined time periods (as set in Report Config). Data are taken from the Meter Read table.
- Meter Watch Java Server Pages displays the results in HTML format by querying the database for the MeterWatchResults table
- Details retrieves the details for a meter that is specified by mouse-click from the listing of meters on the Bill Prep Report Java Server Page
- Report Bean stores all the report parameters as defined by the user in the Report Config function

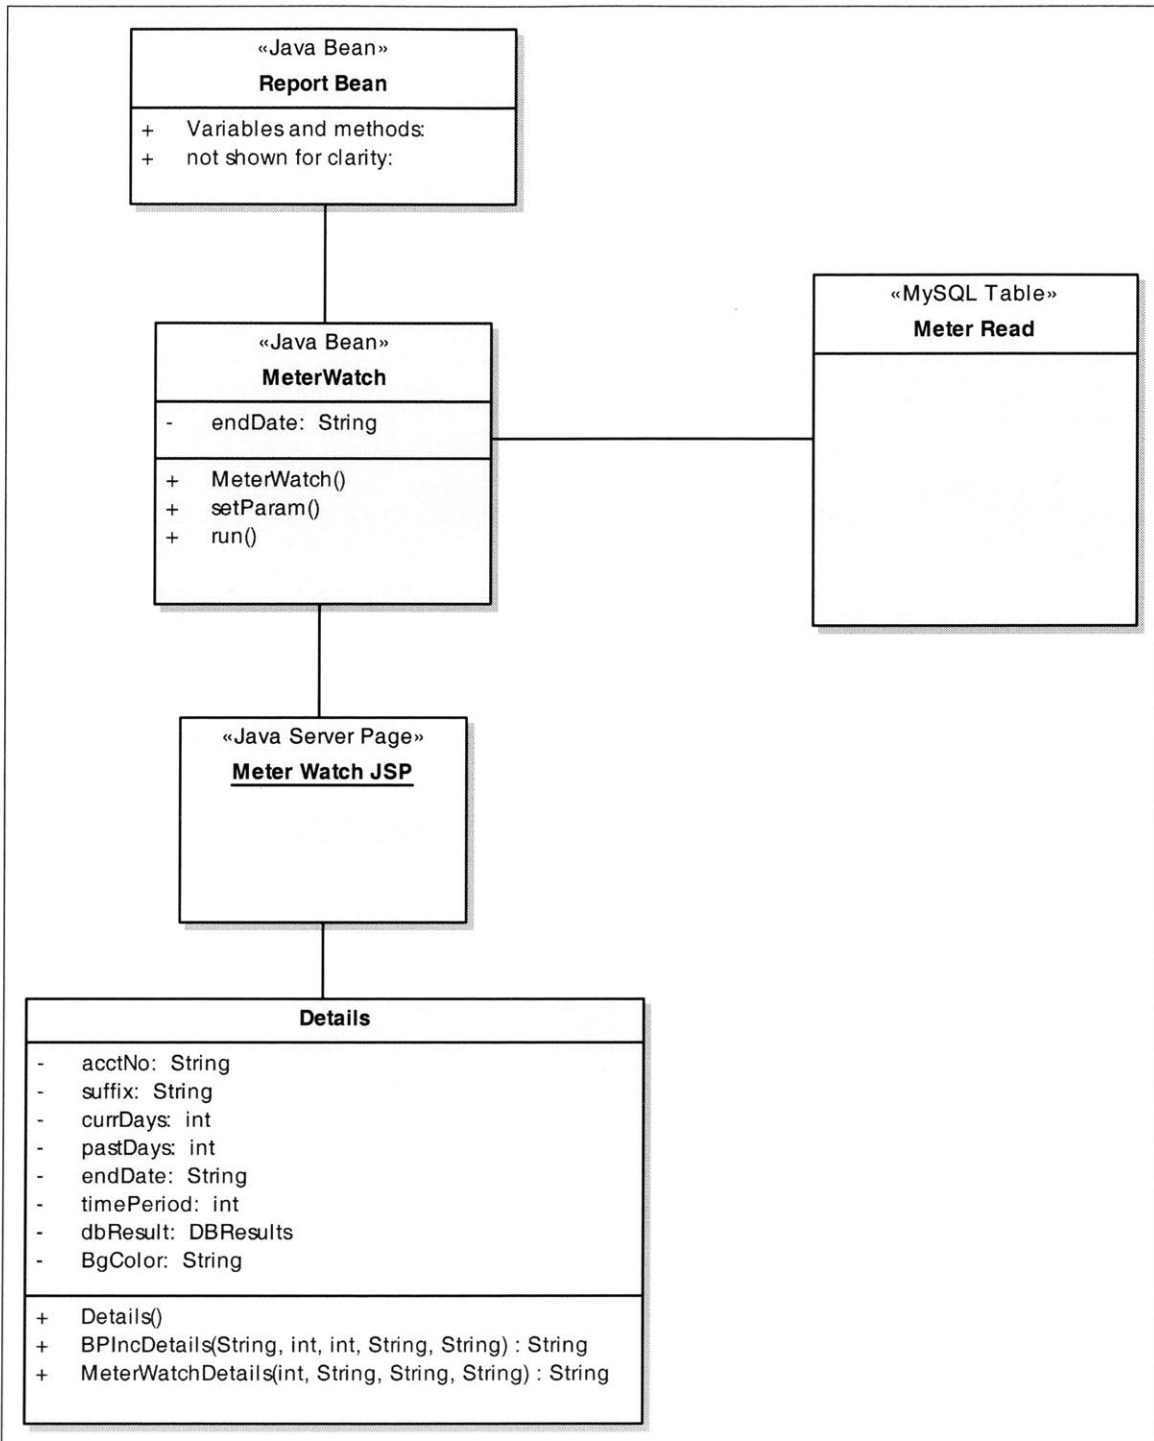


Figure 5-10 - Meter Watch Class Diagram

5.2.2.1 Meter Watch Sequence Diagram

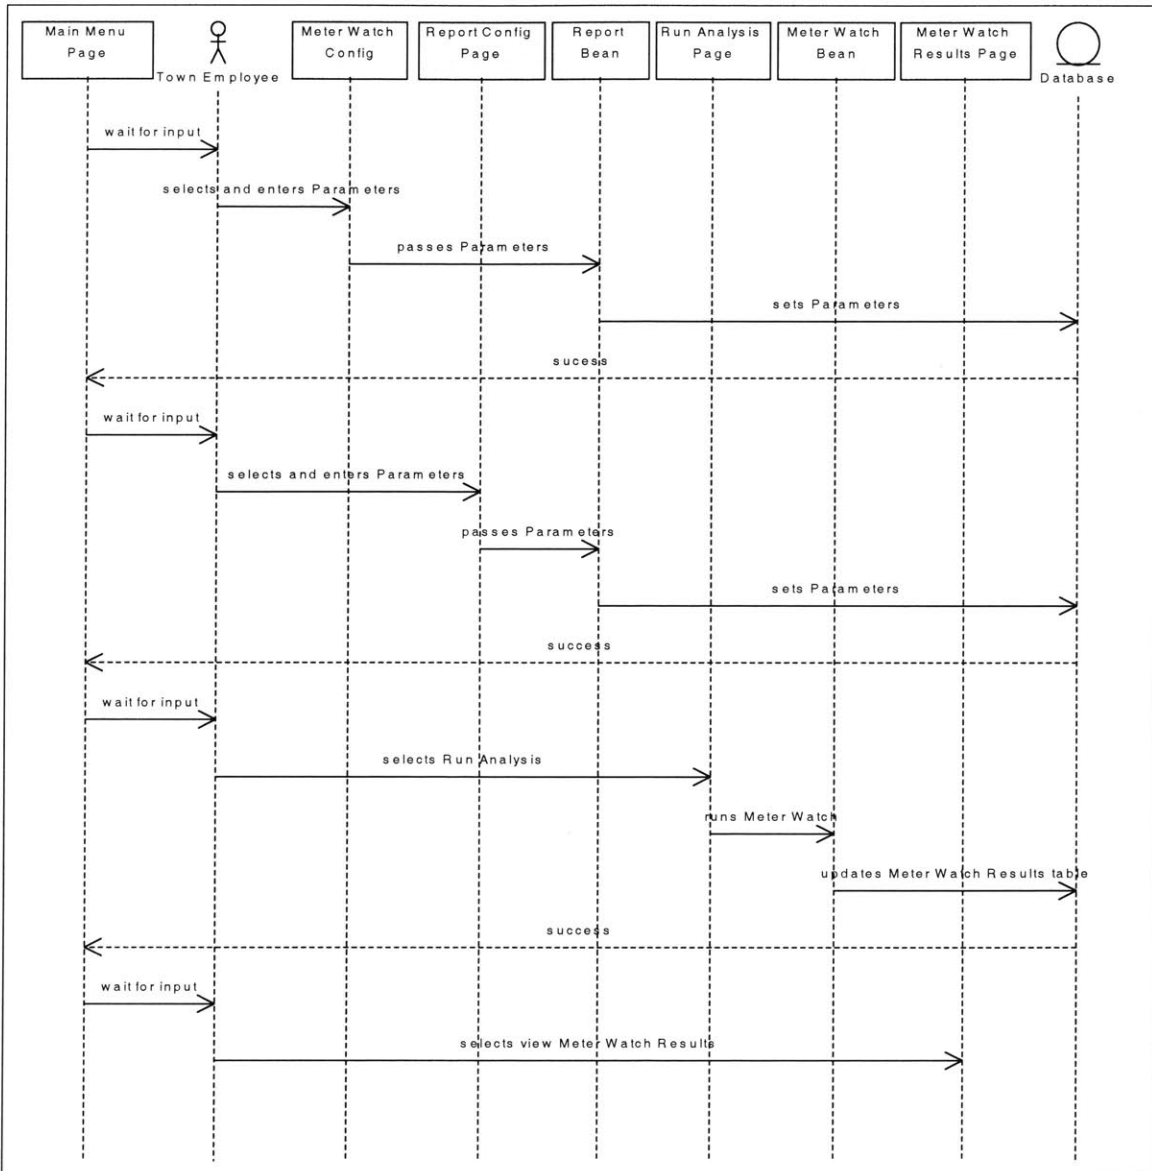


Figure 5-11 - Meter Watch Sequence Diagram

5.2.2.2 Meter Watch Interface Description

The Meter Watch class has two publicly accessible methods: set Param() and run(). The first takes the parameters for from the Report Config table in the database and sets them in the object; the second connects the object to the database and performs the analysis based on the configured parameters.

The Details class has a function MeterWatchDetails() which takes as arguments the account number, suffix and time period for the Meter Watch analysis and returns all the meter readings for that meter on those dates.

The Java server pages coordinate the calling of the methods, setting of parameters, etc. In some cases, the JSP's access the database prior to calling the methods, such as in the case of MeterWatchDetails(), where the JSP retrieves the dates for the details from the report configuration table prior to calling this method.

5.2.3 Usage Pattern Analysis

For a description of the intent of the Usage Pattern functionality, please refer to Section 4.2.4. The classes of the Usage Pattern functionality relate as follows:

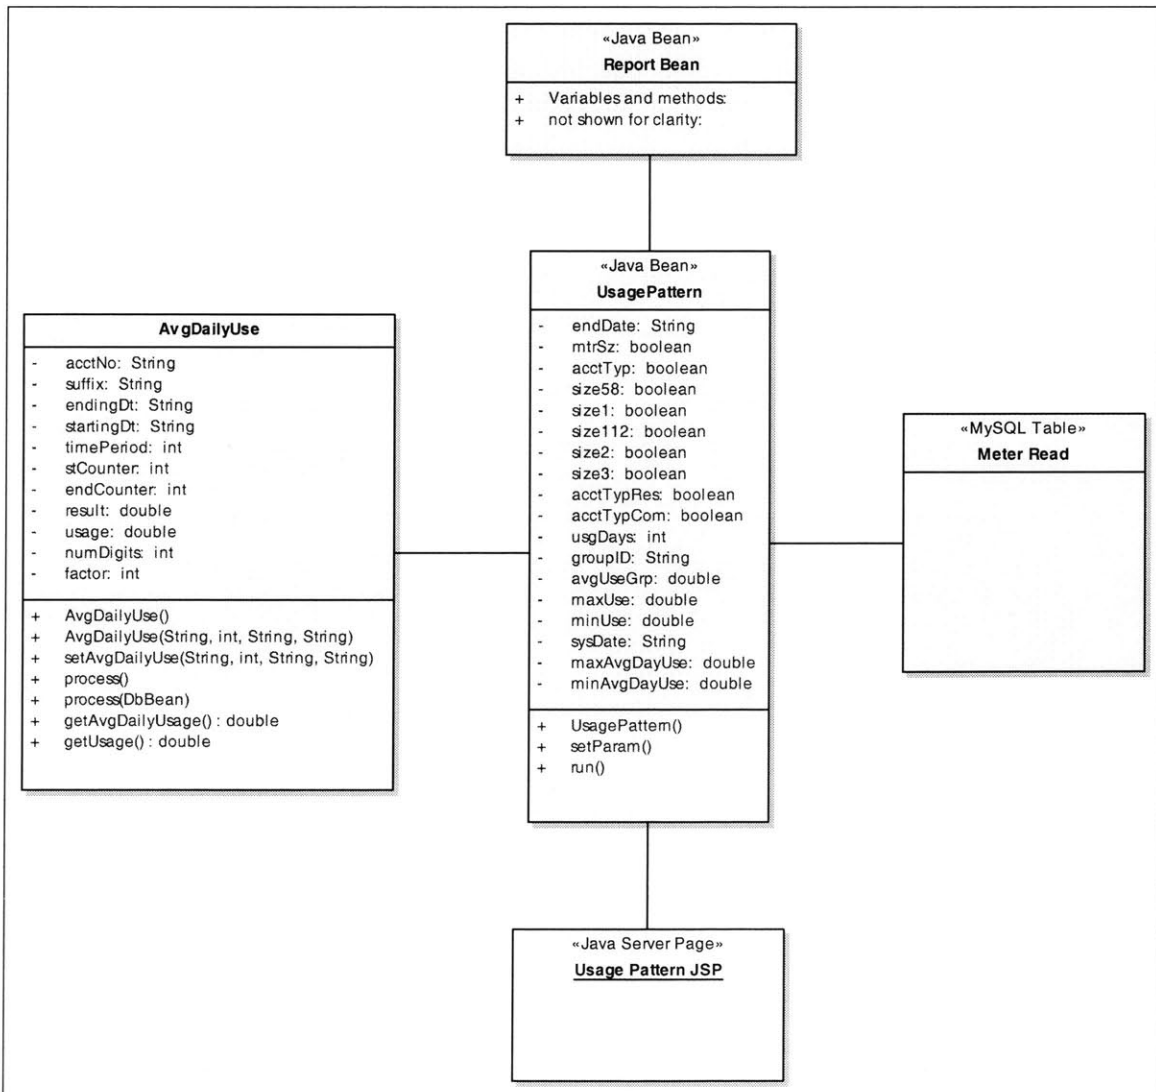


Figure 5-12 - Usage Pattern Class Diagram

5.2.3.1 Usage Pattern Sequence Diagram

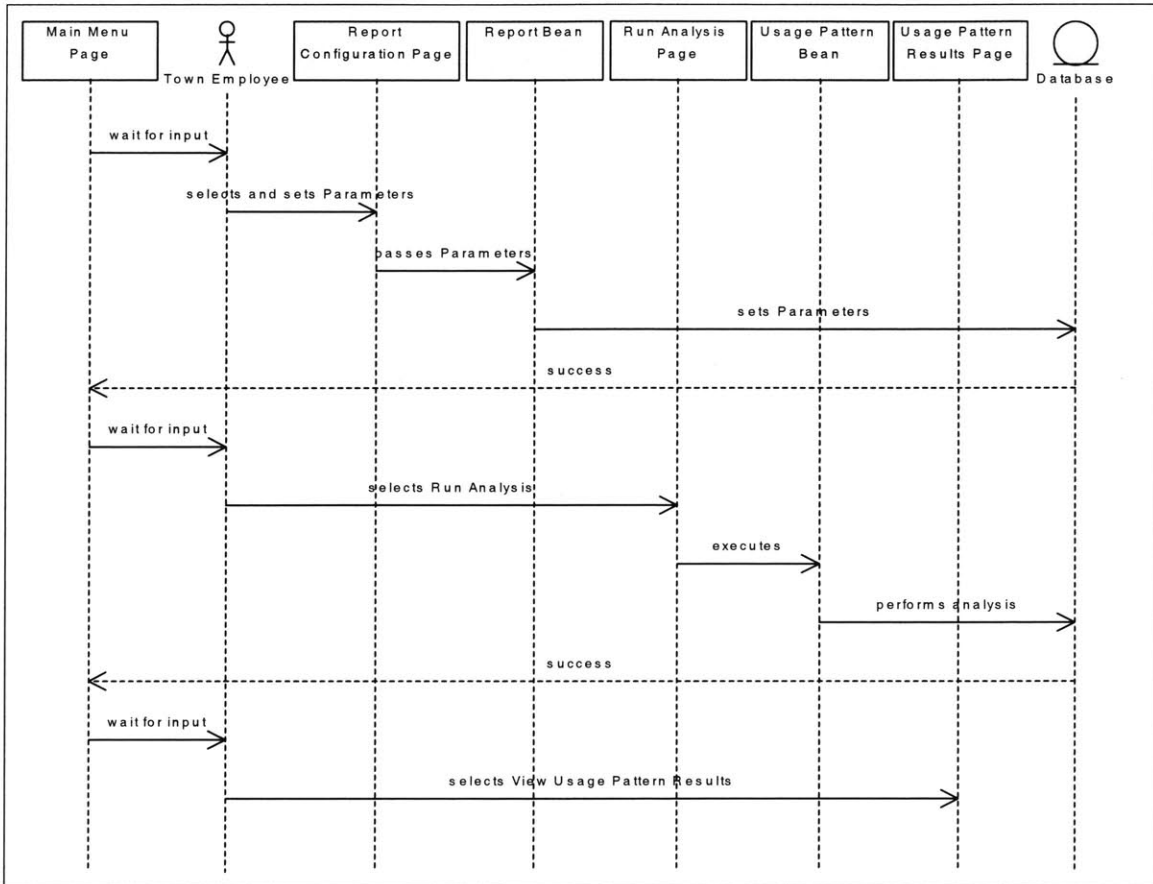


Figure 5-13 - Usage Pattern Sequence Diagram

5.2.3.2 Usage Pattern Interface Description

The Usage Pattern class has two publicly accessible methods: set Param() and run(). The first takes the parameters for from the Report Config table in the database and sets them in the object; the second connects the object to the database and performs the analysis based on the configured parameters.

5.2.3.3 Usage Analysis Performance Issues

Although initially the intent was to have the possibility to group meters by their type (rotating disc or turbine), this option was dropped from the current version of the software. This had to be done because the data available for design only consisted of the manufacturer name, but not the measurement apparatus type.

In addition, the option to classify meters by their size is entirely hard-coded into the software. The report configuration screen does not query the database for the distinct meter sizes and then present the user with the option to categorize a group by those meter sizes. The same situation exists in the creation of the "group code" which is dynamically created in the Usage Pattern Java Bean. This code is used as a recipe for the query to know what meters to aggregate together. The code consists, in part, of the names of the meter sizes, and these again are not taken from the database, but rather are taken from the Java Server Page for the report configuration. Although it is unlikely that these relatively standard meter (and pipe) sizes will change, it is certainly possible and this presents a system limitation. Similar issues could arise with the account type (residential/commercial).

5.2.4 Database Access

In order to access the database, two connection classes exist. The DBResults class (adapted from Marty Hall's book Core Servlets and JavaServer Pages, see references) creates a very simple interface to run a query on the database, receive a table as result and translate this table into an HTML table. The DbBean class creates an interface to send update statements or other utility commands to the database. The DbBean class is frequently used when several successive operations need to be performed on the database, which needs to be done because the current version of MySQL does not support views or nested queries.



Figure 5-14 - Database Access Class Diagram

5.2.4.1 Database Access Narrative

The database access classes are called upon every time a SQL statement is sent to the database, regardless of whether this is done from within a Java Server Page or Java Bean. The general process is as follows:

Access using only a DbBean object:

- A String variable is initialized and its value is set to the SQL statement that is to be executed
- A DbBean object is instantiated
- If a query is to be executed that produces a result:
 - A ResultSet is initialized and set to null. A ResultSet is a java.sql data type that holds the tabular results of a database query.
 - The results set is filled by calling the executeResults method of DbBean with the SQL String as argument
- If query is to be executed that does not produce any results (such as update, add or delete statements):
 - The appropriate methods of the DbBean object are called

Access using a DBResults object:

- A String variable is initialized and its value is set to the SQL statement that is to be executed
- A DBResults object is instantiated
- A DbBean object is instantiated
- The DBResults object is filled by calling the getQueryResults method of DbBean with the SQL String as argument
- The DBResults object can now be translated into an HTML table by calling the toHTMLTable method or toLinkHTMLTable method

5.2.4.2 Database Access Sequence Diagram

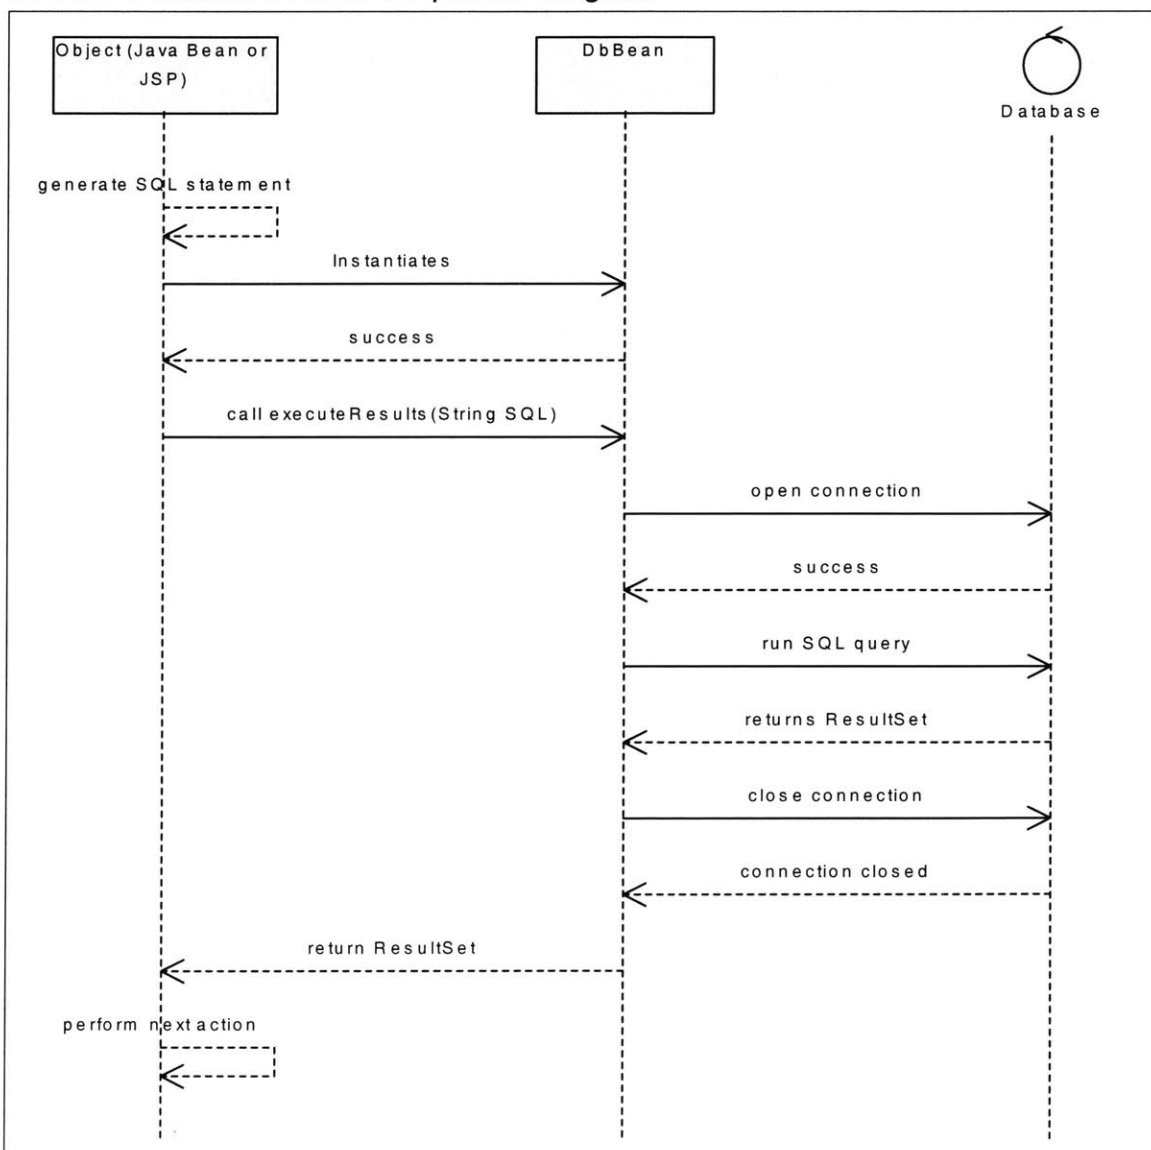


Figure 5-15 - Database Access Using only DbBean Sequence Diagram

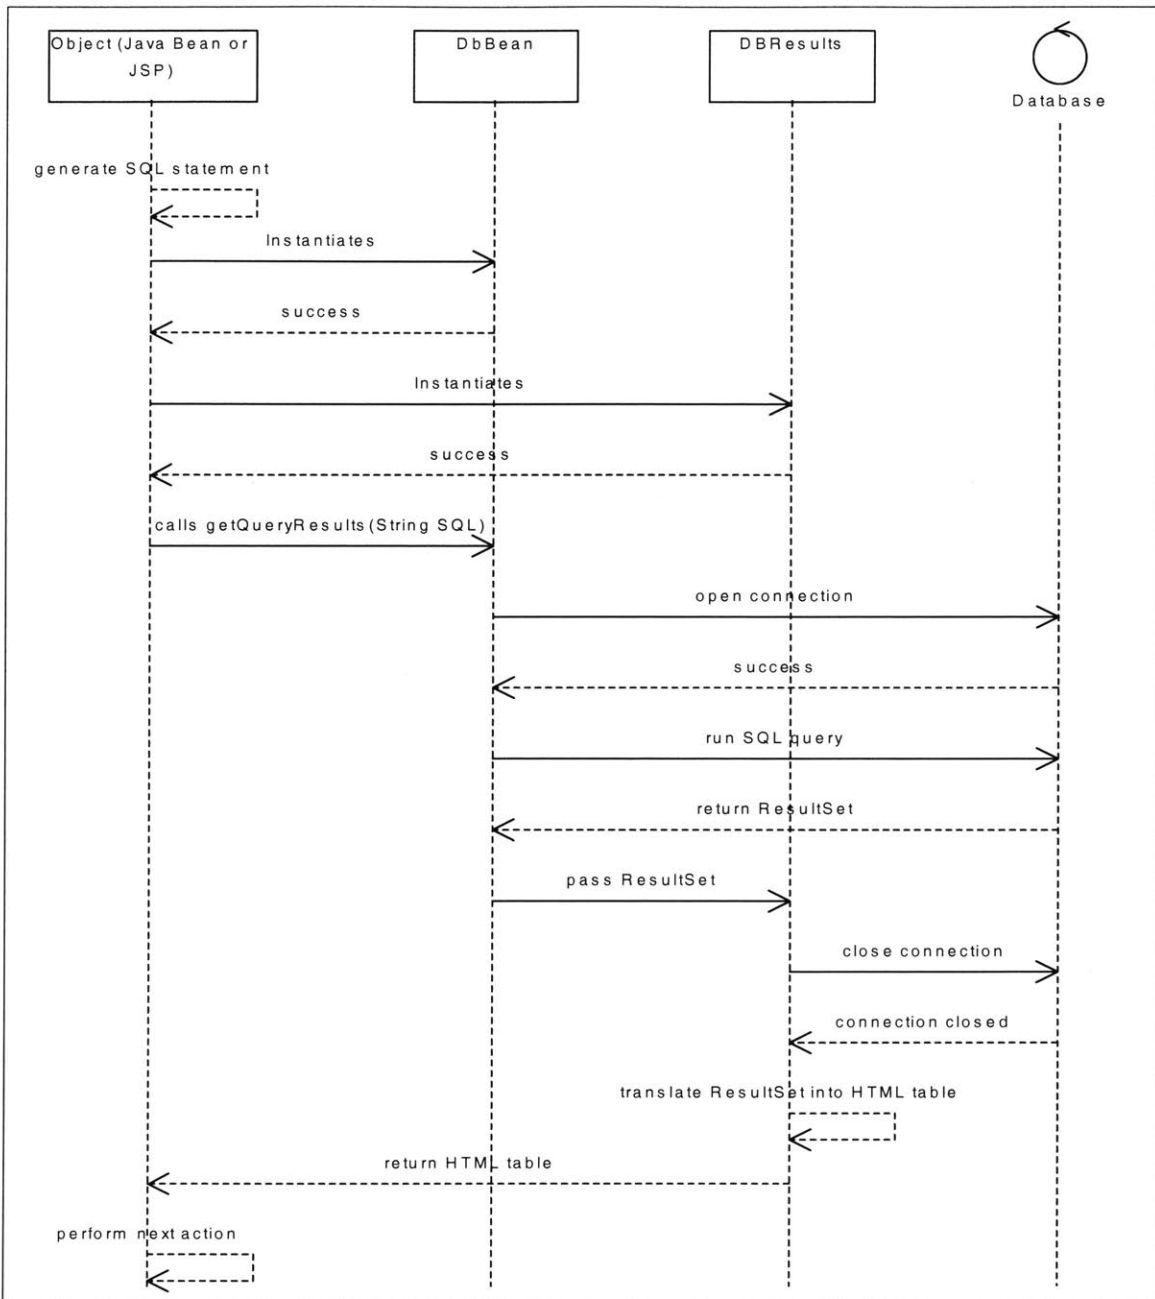


Figure 5-16 - Database Access using DBResults Sequence Diagram

5.2.4.3 Database Access Interface Description

The class DbBean has several publicly accessible functions:

- Close() and connect() govern the connection to the database

- `execAdd(String)`, `execUpdate(String)`, `execDelete(String)` and `execResults(String)` run SQL statements on the database through an open connection and respectively add, update, or delete a record, or retrieve a result.
- `getQueryResults(String, boolean)` runs a SQL statement on an already opened connection returns the results to a `DBResults` object, which then closes the connection.

5.2.5 Administrative Functions

Classes are needed to administer the tables in the database. The possible functions that these classes must perform are:

- Add a record
- Delete a record
- Update a record
- Find and display a record

In addition, referential integrity must be checked, since MySQL does not support this inherently. For the portion of the system that has been implemented up to now, the only tables that need to be administered are `Meter`, `MeterRead`, `Account` and `MeterWatch`. For the integrity checks, the only validations that need to be made are:

- `Meter`
Does an account exist for the meter to be associated with and does a meter with the designation entered already exist?
- `MeterRead`
Do the meter and account exist for the data entered?
- `MeterWatch`
Does the meter and suffix exist?

The administrative classes support functionality to search for any record by only partial information. That is, the system finds all accounts with street address "Mason Street" and "Massachusetts Ave" when only "Mas" were entered in the search field for street address.

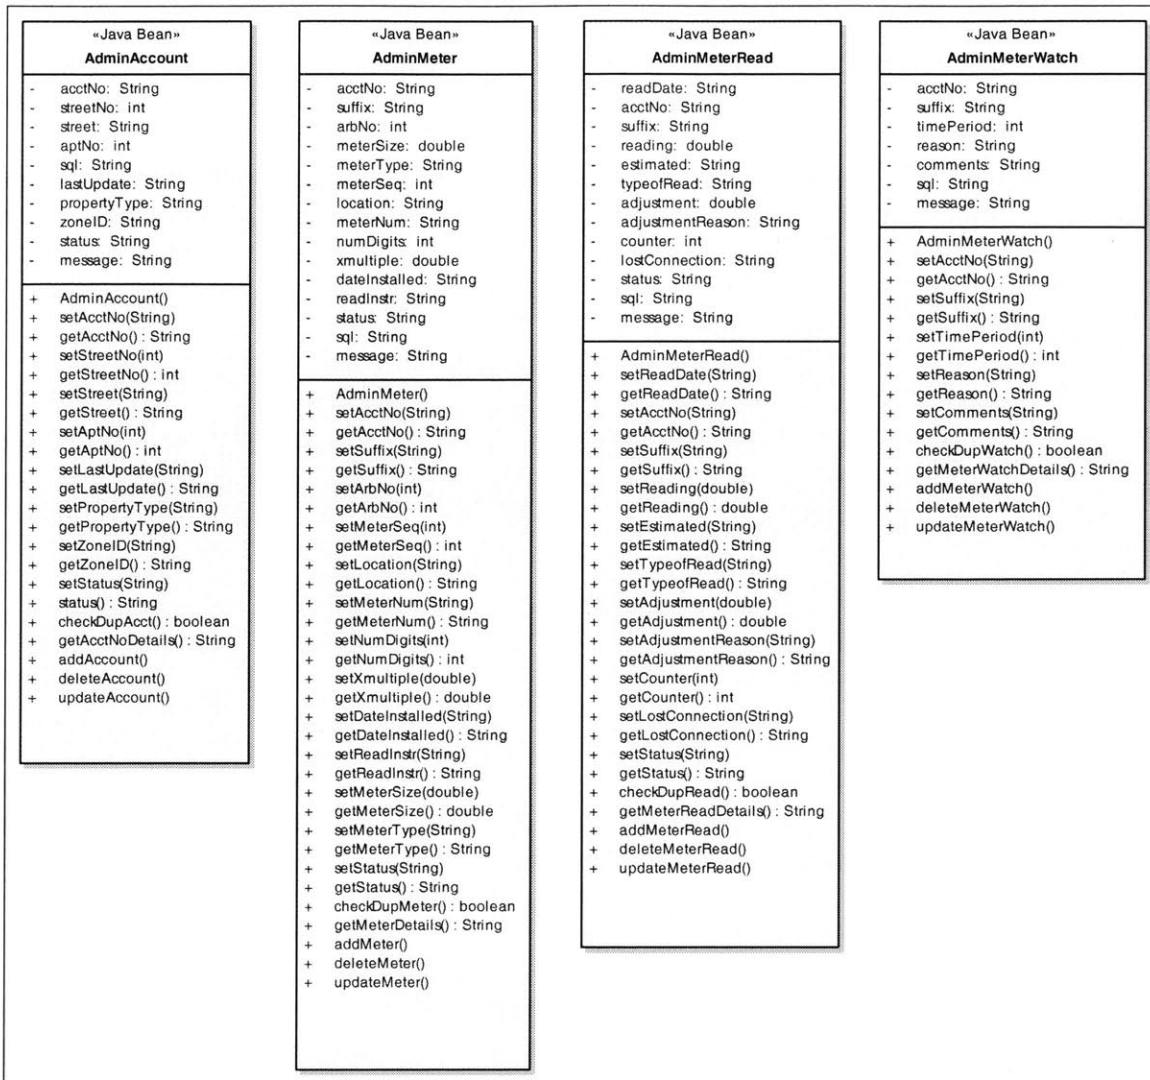


Figure 5-17 - Class Diagram for Administrative Classes

In addition to the administrative classes above, there are three utility classes that assist many of the functionalities in performing their jobs.

- AvgDailyUse calculates the average daily use of a specified account over a specified time period
- CheckRefIntegrity queries tables to determine if a given record already exists or not. The method returns true if referential integrity is upheld (record does not exist) and false if referential integrity is violated (record does not exist)
- DateMath uses Java's pre-existing Gregorian calendar to calculate dates, based on a given date and a time period to be added or subtracted

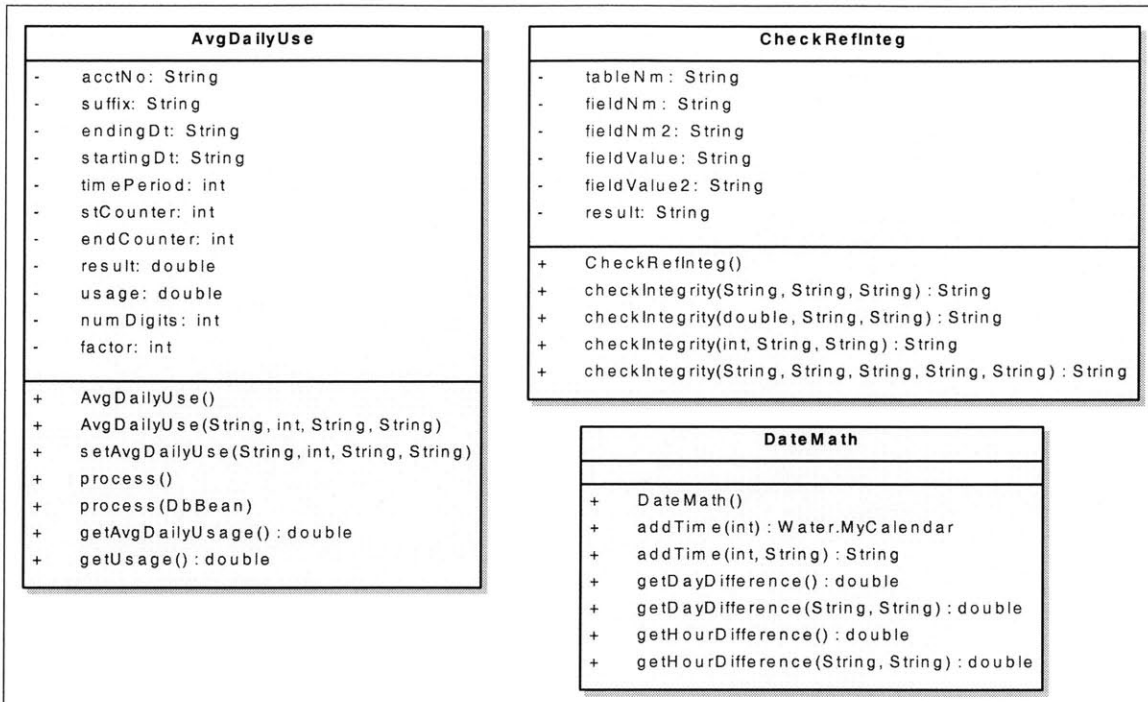


Figure 5-18 - Utility Classes Diagram

5.2.5.1 Administrative Functions Sequence Diagram

For the administrative functions' sequence diagrams, only one has been created, since these diagrams are relatively simple and repetitive. The diagram shown below represents the most complex of the administrative functionalities: The update of an existing record in the database, as shown on the example of the Meter table.

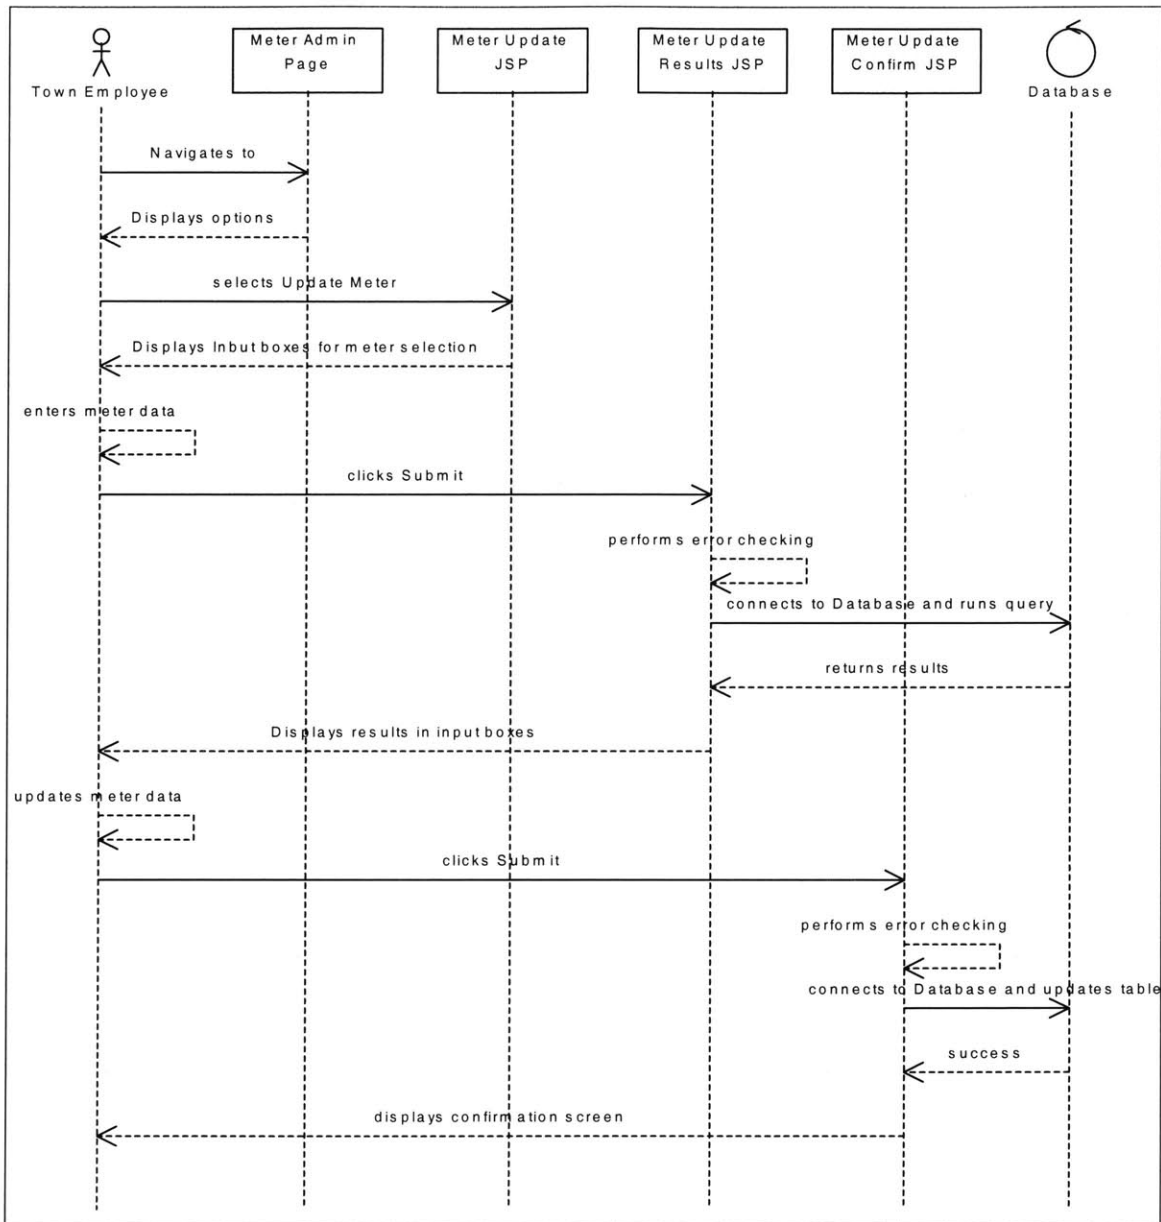


Figure 5-19 - Update Meter Sequence Diagram

5.2.5.2 Administrative Functions Interface Description

The administrative classes have public methods for adding, updating and deleting existing records in the database. These methods are generally called from the Java Server Pages that perform each of these functions. They also have the CheckDup() method, which performs error checking to ensure referential integrity of the entered data.

5.2.6 Software Interface Description

5.2.6.1 External Machine Interfaces

The software for the Arlington Remote Meter Reading System is resident on an Apache Server that will be located in the Arlington intranet. Machines on the local area network will be able to access this server through Ethernet, using universal resource locators (URL's) to access specific system functionalities. Due to some programming restraints from the user-interface, the external machine must be running Microsoft Internet Explorer to properly display the software files.

It is foreseeable that this system will eventually be accessible to residents of the town of Arlington over the Internet. Customers could then check their account information and water usage using their web browser. A current limitation of this is that the software does not include any security measures that allow different permissions for different users. Any user that connects to the main menu screen from any terminal has full access to the system.

5.2.6.2 Human Interface

The software has a graphical user interface (GUI) that is entirely programmed in Hypertext Mark-Up Language (HTML). This allows the user to navigate through the software using the familiar mouse point-and-click interface that is a universal standard today. The user must be using Microsoft Internet Explorer to access the software files, since the user interface is laid out to function properly using this browser. Although it was not the intent of the developers to limit this application to one type of browser, it proved to be difficult to create a totally browser independent software package.

Common software features have been used for the selection of options and entry of data, such as drop-down menus, input boxes and check boxes. Any user that is familiar with the World-Wide-Web (www) or any Microsoft Windows or Office application will not encounter any difficulty using this system.

5.3 Behavioral Model and Description

The software package as currently implemented has three main instances whose state and behavior warrants closer inspection: Account, Meter and Meter Read. These three aspects of a customers record at the town's water department encapsulate all necessary information to run a functional system.

For the purpose of this discussion, a closer inspection of the behavior of the Meter and Meter Read dimensions will be made. The Account information has

only the states 'Active' and 'Deleted' for the purpose of the water management software, since the billing functionality is separate from this system. Including such behavior for billing as 'Delinquent' and 'Current' is beyond the current software implementation and also beyond the scope of this thesis.

5.3.1 Description of Software Behavior

The Meter and Meter Read entities are governed by a set of events that modify their state. A closer look is taken at the events and the resulting states in the following paragraphs.

5.3.1.1 Process Events

Events of the Meter object:

Initially a meter is undefined and the event of defining a meter creates a meter that is associated with an account. This meter will be in standard operation until one of the following events occur:

- The meter stops registering
 - The meter is either fixed and returned to normal operation OR
 - The meter is broken and thus deleted from the database (and replaced)
- The meter seal is broken
- The meter is suspect of under-registration due to deterioration
 - The meter is either fixed and returned to normal operation OR
 - The meter is broken and thus deleted from the database (and replaced)
- The meter is deleted due to reasons other than being broken or deteriorated

In addition, a meter can be placed on meter watch. This does not change the operation of the meter itself, but periodically monitors the meters performance to ascertain whether it is properly functioning.

Events of the Meter Read object:

Initially a meter read is undefined until the system updates the database tables with a reading. This reading is current and correct unless user intervention occurs due to unreasonable values.

- A Meter Watch on a meter can find suspicious readings (abnormally high or low readings or zero readings). This would most likely prompt the user to replace the reading with an estimated reading or delete the reading.

- If the Meter Watch detects a missing reading, the user can enter an estimated reading to fill in the gap.
- The meter read is deleted due to other reasons

5.3.1.2 Process States

Meter States

- Active
 - Standard operation
 - Broken Seal
 - Under-registering
 - Not registering
 - (Meter Watch)
- Deleted

Meter Read States

- Active
 - Standard operation
 - Estimated Read
 - Suspicious Read
- Deleted

5.3.2 *Process State Diagram*

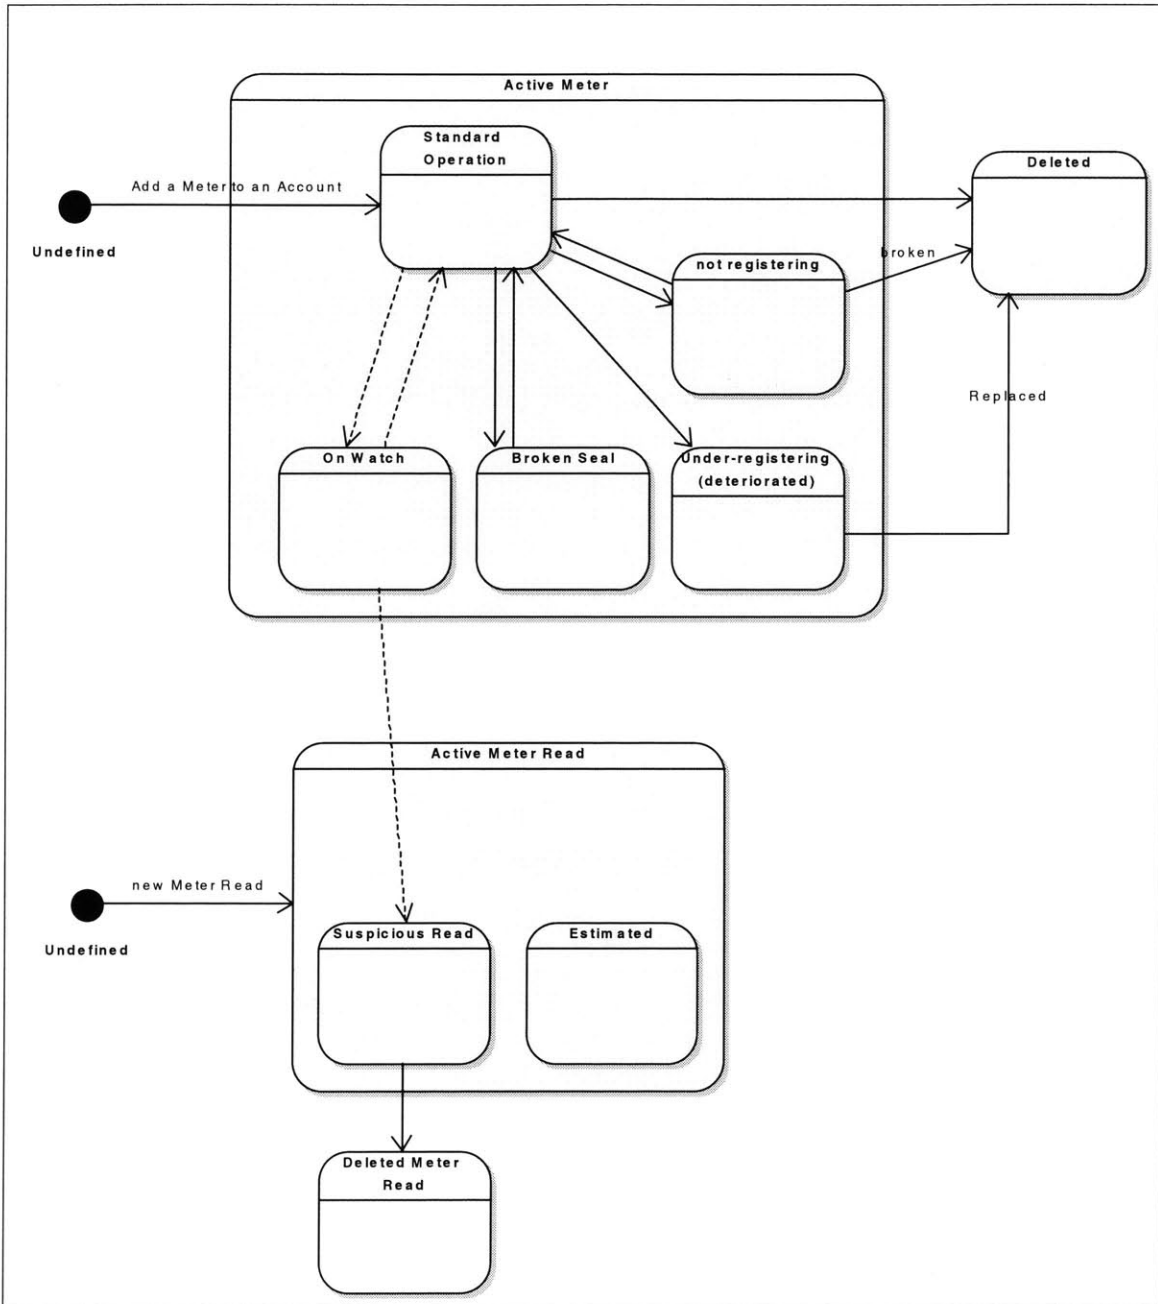


Figure 5-20 - Process State Diagram for Meter and Meter Read

6 Conclusions

The research team has implemented the essential parts of the Arlington Remote Monitoring System. These are the administrative components for the Account, Meter and Meter Read dimensions, and three of six functionalities: Bill Preparation, Meter Watch and Usage Pattern analysis. This allows the system to function and was a reasonable schedule for what was essentially a three-month implementation time with three developers spending about 20 hours a week on coding and designing the software. It was known from the beginning that the scope of this project was large and that most likely a team of three students would not be able to finish the entire system, so there was no pressure to do so.

Although the system encompasses more functionality than we were able to build, this research project will continue in the future and it is anticipated that next year's Master of Engineering students will build the remainder of the system. Depending on the number of students who will want to work on this project, it is certainly possible that the system will be expanded beyond the functionalities that our team envisioned, as the work that will be left to do would not employ more than at most two students.

The portion of the software we have built functions efficiently and correctly (although formal quality assurance testing has not been performed). The goals for this year's project have been achieved, in that the client has a program to read, maintain and interpret the data received from the Hexagram monitoring system server. The software effectively uses all technologies we anticipated learning and applying: Java Beans, Java Server Pages, SQL, Red Hat Linux and the many applications that we used for building the software (FrontPage, MS IDE, Visual Studio, JCreator, etc.).

At completion of this year's project, our team presented the client with the software (in addition to periodic review meetings that were held during the course of the year). During this presentation, every student found one thing or another about the system that could stand revision to make the software more user-friendly. Some aspects of the user interface are counter-intuitive and will require some acclimatization on behalf of the user before the system can be put to use productively. To alleviate this situation, we have inserted help files and on-screen instructions where necessary. However, this situation presents an area for possible future improvement.

Possible future additions could create a Meter Watch functionality that screens through the readings from the Hexagram server as they come in and immediately report any suspect activity. This could be useful for early leak detection, as well as for an envisioned security system that could notify a resident when water is being used at their residence when they are not at home. The applications for a system of this sort are varied and include home security and peace-of-mind for

people who take care of senior citizens. (Already an appliance in Japan reports to the beeper of concerned children whenever their parents use the teakettle.)

There is currently no way for the software to configure the read times of the Hexagram system. This will have to be done separately and certainly the integration of this task with the decision support software would be beneficial to the user. One of the main functionalities that the client will desire after having a good grip on the data from the monitoring system will be to adjust the reading times and frequencies. This will make the system as a whole much easier to use, rather than running a report on one system, interpreting the results and then switching to a separate system to enter the dates and times of specific readings to be taken.

An interface will have to be created between the Hexagram server and the decision support system server. This task is relatively simple, but because this year's research team did not have access to the data format that the Hexagram system will send its readings in, nor any information on the hardware, this task was left to be done in the future. The town of Arlington expects the hardware for the monitoring system to be installed in the next year to two years, so plenty time remains for integration tasks such as this.

7 References

- [1] Brdys, M.A., Ulanicki, B. Operational Control of Water Systems: Structures, algorithms and applications. Hertfordshire (UK): Prentice Hall, 1994
- [2] Butler, Bradford. The Business Transformation Effects of Information Technology. Massachusetts Institute of Technology: Master of Engineering Thesis in Civil and Environmental Engineering, 2001
- [3] Cabrera, E., Martínez, F. Water Supply Systems – State of the Art and Future Trends. Southampton (UK): Computational Mechanics Publications, 1993
- [4] Hall, Marty. Core Servlets and Java Server Pages. Upper Saddle River, NJ: Prentice Hall PTR, 2000
- [5] Khanolkar, Mameet. Investigating the J2EE Software Architecture for Infrastructure Monitoring: A Water Metering Case Study. Massachusetts Institute of Technology: Master of Engineering Thesis in Civil and Environmental Engineering, 2001
- [6] Liggett, James A. “Network monitoring and the algorithmic location of leaks under steady and unsteady conditions.” Water Supply Systems – State of the Art and Future Trends. Southampton (UK): Computational Mechanics Publications, 1993
- [7] Male, James W., Noss, Richard R., Moore, I. Christina. Identifying and Reducing Losses in Water Distribution Systems. Park Ridge, NJ: Noyes Publications, 1985
- [8] Obradovic, Dušan, Lonsdale, Peter. Public Water Supply – Models, Data and Operational Management. London (UK): E & FN Spon, 1998
- [9] Roberts, K. F. Advanced Technology in Water Management. London (UK): Thomas Telford, 1991
- [10] Walski, Thomas M. Analysis of Water Distribution Systems. New York, NY: Van Nostrand Reinhold Company, 1984