

Configuring and managing a large-scale monitoring network Solving real world challenges for Ultra Low Powered and long-range wireless mesh networks

Laurent Maleysson & Christophe Dugas

CORONIS SYSTEMS

Le Millénaire 290 rue Alfred Nobel 34000 Montpellier France
{laurent.maleysson or christophe.dugas}@coronis-systems.com

Abstract

In creating wireless networking solutions suitable for deployment in harsh, unpredictable, and widespread environments, we were confronted with a series of problems as-yet unsolved by commercially available technologies. The purpose of this article is to describe how we addressed mission-critical customer requirements by developing a wireless technology explicitly for devices in Ultra Low Power (ULP) and Long-Range wireless mesh networks. The key endpoints in our target implementation are battery-operated devices located in hard-to-reach places, but which are nonetheless expected to offer a lifespan of several years without human intervention.

We provide an overview of the technical requirements for building ULP networks, with a focus on configuration and management (including patent pending self-configuration and dynamic-routing features).

This is followed by a case study of an existing 25,000 node wireless network deployed for an automatic meter reading (AMR) solution, and examples of provisioning individual nodes in complex real-world networks. We also describe how transmitting information about existing network hierarchy to new nodes not only preserves overall battery life in other network nodes, but also simplifies installation efforts significantly. The technology described here is particularly applicable to metering, telemetry, remote monitoring, and large-scale data collection solutions, while straightforwardly suited for personal and property security, medical surveillance, access control, lighting systems, as well as numerous industrial sectors.

With a strong background in utility metering systems, Coronis Systems provides ready-to-use wireless solutions for manufacturers, VARs, and integrators in the automatic remote metering and wireless sensor network industries.

1. Introduction

In many areas, applications require flexible architecture networks that run autonomously under potentially extreme operating conditions, revealing a whole new set of technical problems to resolve. For example, data collection and monitoring devices can require fixed installations of nodes with a battery life of several years, deployed in hostile and hard-to-reach locations that most wireless solutions cannot handle. The basic criteria for providing intervention-free networking are: ultra low power consumption, long-range

communication ability, and low cost. The right solution also requires reliability, easy installation, and remote network management. All of these features need to be part of the solution's underlying design.

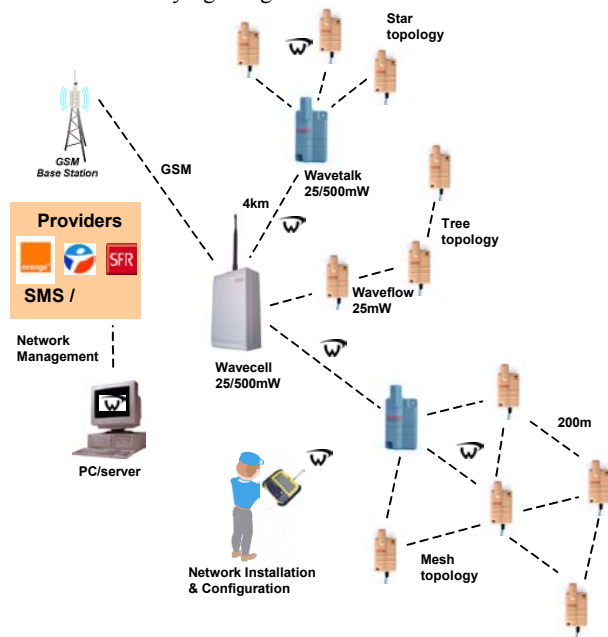


Figure 1: Optimal topology is automatically configured for each local group in large ULP networks

2. Meeting real-world needs with a ULP-optimized solution

For many kinds of wireless network applications, the lack of an appropriate networking technology makes it unrealistic to deploy timesaving and convenient solutions in the field. Often solutions that would benefit consumers and suppliers are subject to severe performance constraints, such as using battery-powered devices whose life spans are expected to exceed several years, and cost-effective solutions are next to impossible.

In some network installations, nodes can be widespread over great distances, requiring routers, gateways, and even cell phone network connections. ULP and long-range solutions are well adapted to citywide networks and large industrial sites, but issues related to network installation, management, and the reality of extreme power constraints can make it

impossible to create hi-density networks in these environments. The solution is a wireless network protocol that extends and enhances existing standards to build efficient networks for metering, home automation, temperature monitoring, and other types of data collection solutions.

Some solutions that appear to solve this problem today really do not meet all requirements for range, battery consumption, flexible topology, and network management. The challenge facing technology provider companies is how to serve Wireless Sensor Network markets with low-cost products that can handle low data volumes over extended distances, but where devices are often installed in hard-to-reach locations and only have a few micro-amps available power. Market demand has driven attempts to meet the challenges and balance tradeoffs presented by often-conflicting parameters.

Current radio frequency (RF) standards all have shortcomings for these solutions: Wi-Fi (costly, power consuming), Bluetooth (low power, short range – 10m), and ZigBee (much lower power but range still too short – 20 m). None address all of these constraints simultaneously.

Starting in 2000, Coronis Systems chose to bridge the gap between standards-based wireless communications and ultra-low-power (ULP) devices with its Wavenis technology, designed with Bluetooth extension capabilities. Because of differing technical and economic constraints, Coronis uses two distinct solutions with similar architectures: end-points with ULP, and access points with either just ULP or a Bluetooth/ULP dual-mode master. This strategy makes it possible to address the market for ULP communications while offering a way for Bluetooth companies to provide solutions to reach new market sectors.



Figure 2: Wavefront OEM card with Wavenis RF transceiver

With an installed base of nearly 400,000 ULP nodes (and another 500,000 on order!), Wavenis matches market reality and industrial deployment requirements. Coronis products include the Wavenis RF front-end transceiver and Wavenis protocol stack embedded in a low-power micro-controller, offering point-to-point, point-to-multipoint (broadcast, polling), and repeater functions for star, tree, and mesh network topologies. Wavenis implements narrowband FHSS techniques (50 kHz channel bandwidth @ 20 kbps achieving a very high link budget of 125dB with only 25mW output power – yielding a range of 1km line of sight & 300m indoors) with advanced data processing (Forward Error Correction with BCH (31,21) coding, data interleaving), Automatic Frequency Control (AFC), Automatic Sensitivity Control and QoS management (output power, RSSI, ...) with fast access time (max of 1s for metering apps) and average operating current as low as 10µA. Based on this technology core, Coronis’s ULP products (such as Waveflow – hereafter referred to as “WF” in diagrams) integrate a single battery

(3.6V / 3.4A.h / Li-So-Cl) and achieve a lifetime of 15 years at an extremely attractive price – as low as 20 Euros.

3. Case Study: Installing and configuring a 25,000 node ULP network

The hardest part about evaluating new technologies is seeing how they actually perform in the field. We will illustrate this article with a case study on a remote data collection application that currently monitors an entire city’s water meters. SAUR, the water utility provider for the city of Sables d’Olonne (France) has a fully automated network of 25,000 ULP nodes for gathering water consumption data via a wireless mesh network and transmitting it back to headquarters. It is the first large-scale remote-controllable, fixed wireless network of its type in Europe.

The premise of this network consists of placing pulse-detecting sensors on each water meter, transmitting counter data to a wireless module for storage, calculating consumption, then uploading data to a concentrator connected to the GSM cell phone network.




Figure 3: Devices are routed with successively decreasing QoS values (from T_0 to T_2) with the same route level of 1 (one level from the root).

4. Automatic network configuration while respecting ULP and quality constraints

Optimized embedded software plays a critical role in the overall solution. We will now focus on how automatic network configuration affects the viability of ULP networks. A new routing algorithm actually extended the energy resources on nodes in large-scale, hi-density wireless networks, while simplifying installation and configuration. For this we had to devise an algorithm to avoid exponential increase in communications when setting up network nodes.

4.1. Launching configuration to update default settings

Device configuration can be simplified by using intelligent algorithms inside the communication protocol itself. The algorithm’s efficiency is critical for obtaining networks that are manageable in real life, but which also stand up to the rigors of ultra low-power and long-range solutions. Though many types of devices can benefit from these features (from wrist-watches to lighting control and security access systems),

we will focus on the metering park example as we describe how the protocol’s dynamic routing ability simplifies network management and obtains the highest possible Quality of Service (QoS) without wasting battery life.

The downside to any radio technology that offers a relatively high link budget (frequently associated with long radio range) is the tendency for new nodes to query a large number of devices – in fact, all those within range! With very tight ULP constraints, one of the most critical issues was to find a way to avoid the exponential increase in unnecessary inter-device communications. There is nothing particularly innovative about an installation procedure in which new nodes use broadcast mode to locate neighboring nodes to find the best network route. However, our solution offers a combined, iterative use of three selection criteria for establishing the best route while generating minimum inter-device communications. Only devices that satisfy specified conditions will respond. These criteria are:

1. Minimum accepted Quality of Service (based on RSSI).
2. Device hierarchy in the network (creating a notion of network hierarchy with level 0 assigned to all “root” type devices, level 1 for nodes attached directly to the root, level 2 for nodes attached to level 1 nodes, level 3 for nodes attached to level 2 nodes, etc.).
3. The number of nodes already attached to a device (i.e. the number of nodes for which a device is already acting as a router, a service that impacts battery life).

Two situations may arise during installation. If a network technician already has information about a network’s topology, the installation procedure can be optimized by indicating initial search conditions. Without network topology information, the procedure runs through a dynamic configuration routine. Regardless of the initial conditions, we can describe configuration algorithm as follows:

1. New nodes seek a level “n” device with a minimum QoS of “QoS X”.
2. If a level “n” device is not found, continue iteration while gradually accepting lower QoS until minimum P_{max} or $RSSI_{min}$ levels are reached.
3. If an appropriate level “n” device is still not found, iterate “n” and start again.

Broadcast mode is used but it does not trigger unproductive traffic. New nodes only receive responses from devices with the right hierarchical level and a QoS better than, or equal to that requested. If several devices meet the selection criteria, the requestor chooses the one with the fewest devices being routed, or if that number is the same, the one with the best returned QoS

All devices should be physically in place prior to triggering the bottom-up configuration process via a handheld computer. As part of the power-saving scheme, nodes are instructed to configure themselves, rather than receiving a command from a root device upon detection. The route for each *root device* R_i is set to its own address $\langle R_i \rangle$, and QoS for root devices is set to the maximum value (e.g. 100). When new nodes are installed into a network, they are initially assigned a logical route level of -1, and QoS is set to 0 in nodes that are not yet

routed. All nodes that are not yet paired have their routes initialized as $\langle \text{NULL} \rangle$.

Note: in this section we use the term root devices to include actual network roots as well as other potential terminal devices and concentrators.

Routing begins at root devices, with their routes respectively set to their own addresses $\langle R_i \rangle$ and route levels set to 0. New nodes are then placed around the root devices. This is not mandatory, but routing will be more efficient when devices linked directly to a root are available for those requiring a more complex route. In reality, nodes are self-configuring and the order of installation is not important. Also, once installation is complete, networks configure and maintain themselves through periodic requests initiated by nodes.

4.2. Sending an installation frame to start the process

During provisioning, nodes engage in a *request-proposal-pairing* sequence in order to establish the best path to a root device. The best path of course, is the one where QoS is highest overall, even if it implies going through a router instead of taking a more direct route. The installation technician uses software on a wireless-enabled handheld computer to send an *installation frame* to a new node so that it broadcasts a targeted request for a routing assignment. The installation frame usually contains information about network topology already in place.




Figure 4: Device A benefits from the best offered QoS

4.3. Registering nodes via an assignment request

Next the node broadcasts an *assignment request* – a data frame that nodes use to request a route. The only devices that respond are those that can offer the requested route level and a QoS better than, or equal to the minimum requested. For example, a **Device B** cannot answer an assignment request from **Device A** when:

- **Device B** is not yet installed (logical route level is -1)
- **Device B** is installed but its route level is *higher* than requested (indicating that **Device A** is seeking a more direct route to the root). If a more direct route is not found, **Device A** will probably raise the requested route level to match that provided by **Device B**.
- **Device B** is installed but its route level is *lower* than requested (a rare scenario in reality, indicating that **Device A** had already sent a request to **Device B**, but received or transmitted a QoS that was too low).
- **Device B** offers the right route level, but the assignment request is received with a lower QoS than that requested.

During the route assignment process, the requested QoS takes RSSI into account. A device will not answer when RSSI is lower than the requested QoS requires, thus favoring devices with the highest possible link budget. If no proper responses are received, the requesting node can adjust its criteria and try again with either a lower QoS or a higher route level (i.e. accepting devices farther removed from the root) as routers.

In high density networks we usually try to install nodes with successively decreasing QoS values, while keeping the requested route level as close to 0 as possible (seeking a direct route to the root while maintaining acceptable QoS).

4.4. Choosing the best route via an assignment proposal

Device B_i receiving an *assignment request* from **Device A** will answer when level and QoS conditions are met. The response is an *assignment proposal* embedded in a data frame for **Device A** to analyze. This frame is sent using CSMA (RTS/CTS) to prevent collision if several devices respond.

Device B_i provides **Device A** with information about its RF QoS to the root. In an *assignment proposal*, QoS parameters can indicate: RSSI of the *assignment request* it received; QoS between **Device B_i** and the root; The number of devices already routed by **Device B_i**. Based on these parameters, **Device A** compares the QoS for each answering **Device B_i** and selects the one that best matches the requested criteria.

Prior to any installation, the QoS for a root device is set to the maximum value, and to zero for all other devices. If several devices return *assignment proposals* with QoS at the maximum value, selection will first be based on RSSI, then on the number of nodes being routed (to avoid overload).

4.5. Pairing

Pairing occurs after **Device A** chooses the best **Device B**. In this process, the route between **Device A** and the root device (terminal device or concentrator) must be setup on **Device A**; and the route to **Device A** must be setup on the root device. Since our protocol uses *connection-oriented* RF links, the routes between devices and intermediate nodes are known when communication is initiated.

Note that when **Device A** selects **Device B_i**, a route is already defined between **Device B_i** and a root device. If **Device B_i** is a root device itself, then the route is a direct link <B_i>. Otherwise, the route for **Device A** becomes <A|R_i>, where <R_i> is the route between **Device B_i** and the root.

In the first step of pairing, **Device A** requests **Device B_i** for the route <B_i> to the root device. **Device B_i** returns its own stored route. Then **Device A** sends a message to the root device via **Device B_i**. Route notification may be explicit or implicit, as route information is carried in LLC and MAC layers. In less optimized protocol stacks information can be handled at the L-PDU level.

This mechanism is one of the basic factors for ensuring easy network management, both during installation and later on for maintenance. This design is the foundation for a configuration process that works in the field and respects performance requirements of ultra low power wireless mesh networks.

5. Provisioning individual ULP nodes using the Wavenis algorithm

Most discovery algorithms lead to exponentially increasing traffic, which is detrimental to power consumption. Existing self-configuring and dynamic routing algorithm proposals for large-scale outdoor WSNs [see REFERENCES] are inappropriate where hundreds or thousands of Ultra Low Power / Long Life sensors have to be installed and monitored. Solutions that lead to an inordinately high number of data frames before routes are established, even within a single cluster, cannot satisfy customer requirements in the long run.

This section provides examples to show how network configuration is taken into account by the protocol itself, alleviating network managers of an otherwise complex and potentially insurmountable challenge. We continue to use the example of building a wireless mesh network such as described in our 25,000-node case study.

5.1. Adding new nodes to a simple wireless network

Let us start with a Coronis Wavecell root device (concentrator) named WC1, installed and running normally. Its route level is 0 and QoS is 100. WF1 is an already-installed Waveflow node (wireless meter monitor), paired to WC1. Here is the sequence of events that occur when Waveflow node WF2 is added to the network:

1. The handheld computer is used to send an installation frame to WF2 via PPP (wirelessly).
2. WF2 broadcasts an *assignment request* seeking a level 0 device, with a minimum QoS of Q_x. It then goes into receive mode and waits for *assignment proposals*.
3. If it receives a QoS Q_X > Q_x, WC1 (level 0) returns an assignment proposal. Otherwise there is no feedback and WF1 stays in receive mode until timeout expires.
4. WF2 pairs with WC1 if QoS and route conditions are met.

Here WF1 will not respond to the assignment request because its route level is 1 and WF2 requested a level 0. After pairing, WF2's route level becomes 1. Then both WF1 and WF2 could respond to an assignment request with a level of 1.




Figure 5: WF2 establishes a direct route to WC1 after detecting the desired QoS

If WC1 does not respond to WF2 because its QoS is too low, WF2 notifies the handheld that connection was not possible. Although marginally slower, PPP-CS transmission uses RTS/CTS to avoid collisions. This not a critical issue as it is

only used during installation. The installation frame includes information related to optimizing the network

5.2. Using repeaters to build a network

The network gets deeper as more nodes are added. In this next scenario, root devices (WC1 and WC2) are installed and running normally, and node WF1 is already paired to WC1. Node WF2 is already paired to WC2. Here is a possible sequence of events for installing WF3, which is too far away from the root devices to establish a direct route with them:

1. Handheld sends installation frame to WF3.
2. WF3 starts by broadcasting an *assignment request* seeking a level 0 device, with a minimum QoS of Qx.
 - WC1 does not receive the request, or responds with a lower QoS, and is ignored by WF3.
 - WC2 does not receive the request, or responds with a lower QoS, and is ignored by WF3.
 - Receive mode on WF3 ends without any *assignment proposal* when timeout expires.
3. WF3 broadcasts a new *assignment request*, this time with a Level of 1, and a minimum QoS of Qx'.
 - WF1 does not receive the request or responds with a lower QoS, and is ignored by WF3.
 - WF2 does not receive the request or responds with a lower QoS, and ignored by WF3.
 - Receive mode on WF3 ends without any *assignment proposal* when timeout expires.
4. WF3 broadcasts a third *assignment request*, for a level 1 device and a minimum QoS of Qx".
 - WF1 does not receive the request or responds with a lower QoS and is ignored by WF3.
5. WF2 returns an *assignment proposal* via PPP-CS, with a received QoS better or equal to Qx".
6. WF3 pairs with WF2.
7. WF2 notifies WC2 of the route to WF3.

If there is no feedback from the level 1 devices WF_i in step (4.), WF3 responds that connection was not possible, or it issues a new assignment request with a higher route level.

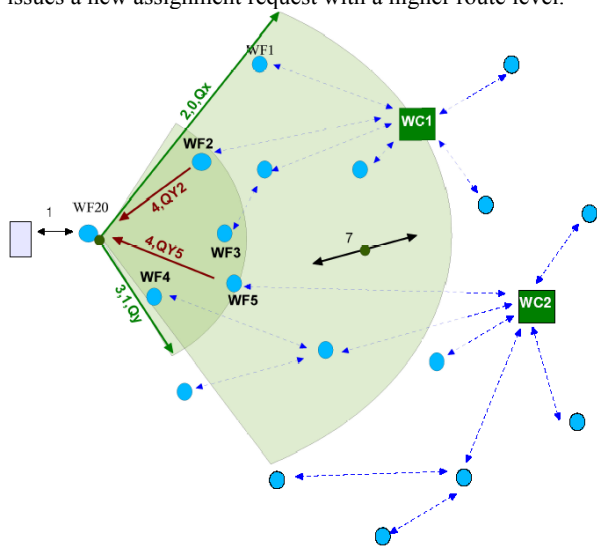


Figure 6: Building complex wireless mesh networks automatically

5.3. Network under construction

This last example shows how new nodes can still be added efficiently in increasingly complex networks, such as the one shown in figure 6. Here is a sample sequence of events when installing a 20th ULP device called WF20:

1. An installation frame is sent to WF20 via PPP from the handheld computer. Installation data also provides Qx, the required minimal QoS. Qx is based on the QoS levels in other installed devices.
2. WF20 broadcasts an *assignment request* seeking a level 0 device with a minimum QoS of Qx.
 - WC1 receives the request but QoS is lower than Qx, and it does not respond.
 - Other WC_i are out of range and cannot respond.
 - Other WF_i do not respond, either because they are already installed and their Levels are > 0, or they are not installed yet and their levels are -1 (routes are never established via devices with levels of -1).
 - Receive mode on WF20 ends without any *assignment proposal* when timeout elapses.
 - WF20 does not start a new *assignment request* with lower QoS because of information provided in the installation data frame.
3. Instead, WF20 broadcasts an *assignment request* with a level of 1 and a minimum QoS Qy.
 - WF3 and WF4 receive the request but do not respond because they are level 2 and level 1 was requested.
4. WF2 and WF5 return an *assignment proposal* via PPP-CS, with QY2 and QY5 respectively.
5. WF20 selects the best QoS and performs pairing with WF2 if QY2 > QY5.

QoS takes into account the RSSI level of the received frame, but also the number of devices already routed through it. With these metrics it is possible to find the best trade-off between communication reliability and power savings, and to dispatch repeater load equally throughout the network.




Figure 7: Management of QoS is also combined with step-by-step adjustments to output power

Managing QoS steps in assignment requests

In high-density networks, a large number of devices may be able to respond to an *assignment request*. Responses using schemes like RTS/CTS are too time consuming in this configuration and not very efficient. You can limit the number of potential responses to a request by starting with a very high QoS level, and allowing hops through repeaters with poor QoS only when there is no alternative.

Live management of wireless networks

During network installation, concentrators store a complete routing table of all network nodes attached to them. The table contains the physical address of each node as well as any intermediary nodes (other ULP nodes and/or repeaters). Network nodes also store their own routes for reaching a root or concentrator, with the physical address of the root or concentrator and all intermediary points.

When technicians go on-site to update a network, re-configuration and system refresh procedures are exactly the same as described earlier, with *installation frames* sent via handheld computer to the network nodes that need updating. The network also contains other necessary failsafes. For example, node breakdown is detected either by the supervisor when attempting to establish a connection, or the nodes themselves can realize that they were unable to connect to their root or concentrator to transmit a spontaneous alarm. In the first case, the supervisor can establish new routes to bypass the defective node. This can be done remotely through administration software or by sending a technician on-site.

When a node itself detects a problem after several attempts to establish communication, it will automatically launch its re-installation procedure and seek a new route by broadcasting new assignment requests as described earlier. This critical “network self-healing” process makes it possible to eliminate manual intervention and automate maintenance procedures.

Real-world results

The dynamic routing algorithm described in this article has been compared to other well-known approaches such as AODV and GRAd, both traditionally proposed with Zigbee stacks. It was our experience that the new algorithm generates from 10 to 100 times fewer communication exchanges for installation and setup. With network maintenance and self-configuration an ongoing process in real-world installation, this significantly extends the lifespan of end-points. A mathematical comparison is provided to demonstrate the higher performance of Wavenis solutions compared to others.

6. Conclusion

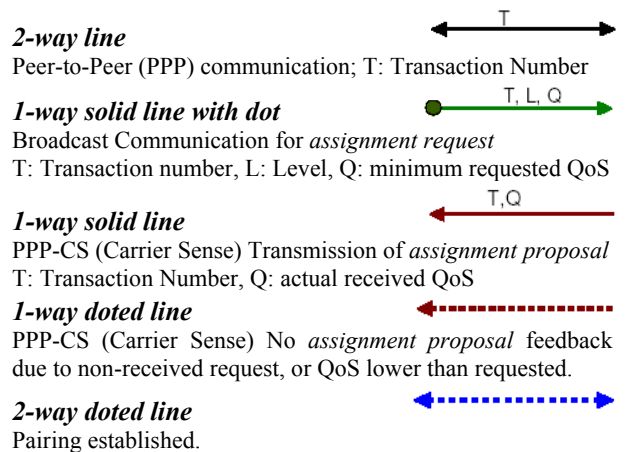
The overall Wavenis solution, including ULP wireless products, flexible WSN topology, and dynamic routing protocol enables very large, efficient wireless mesh networks with self-configuration and self-healing capabilities using a technology that does not adversely affect battery life in autonomous devices. By limiting the proliferation of data exchange required by resource-limited network nodes, and by providing other key network routing features on low-power embedded RF modules, this solution ensures the long-term viability and reliability of Ultra Low Power / Long Range networks, even very large ones like the 25,000 node network we used as an example, and other even larger sites currently being installed. The implications for real-world installations are significant as data monitoring and remote control solutions are increasingly implemented in a variety of scenarios, including wireless telemetry, automatic monitoring, metering solutions and smart environments such as homes, hospitals, and buildings.

7. Appendix

7.1. Terms used in this paper

- Transaction Number** indicates sequence order.
- Terminal/Root device** – the device to which new nodes must be linked.
- Level** – designates the difference between a new node (level = -1) and a terminal or root device (Level = 0).
- QoS (Quality of Service)** – takes into account radio link between new nodes and a terminal device or root.
- Pairing** – logical link between two devices.
- Route** – list of devices used to reach terminal device or root.
- WF (Waveflow)** – Coronis end-point product by designed for utility metering (water, gas, electricity, heat).
- WC (Wavecell)** – Coronis GSM/Wavenis gateway.

7.2. Graphic keys



8. References

- [1] Seema Bandyopadhyay and Edward J.Coyle, *An Energy Efficient Hierarchical Clustering Algorithm for Wireless Sensor Networks*. Purdue University, IN, USA - 2003 IEEE.
- [2] Bhaskar Krishnamachari, Yasser Mourtada, and Stephen Wicker, *The Energy-Robustness Tradeoff for Routing in Wireless Sensor Networks*. University of Southern California, Los Angeles, CA, USA – 2003 IEEE.
- [3] Jim Chou, Dragan Petrovic and Kannan Ramchandran, *Tracking and Exploiting Correlations in Dense Sensor Networks*. University of California, Berkeley, CA, USA – 2002 IEEE.
- [4] Babak Hassibi and Amir F. Dana, *On the power efficiency of sensory and ad-hoc wireless networks*. California Institute of Technology, Pasadena, CA, USA - 2002 IEEE.
- [5] Maggie Xiaoyan Cheng, Jianhua Sun, Manki Min and Ding-Zhu Du, *Energy-efficient Broadcast and Multicast Routing in Ad Hoc Wireless Networks*. University of Minnesota, USA – 2003 IEEE.
- [6] H. Nakakita, K. Yamaguchi, M. Hashimoto, T. Saito and M. Sakurai, *A Study on Secure Wireless Networks Consisting of Home Appliances*. 2003 IEEE.
- [7] Chih-fan Hsin and Mingyan Liu, *A Distributed Monitoring Mechanism for Wireless Sensor Networks*. University of Michigan, USA, 2002.