From **Wireless Sensor Networks: A Systems Perspective** (Artech House, 2005)

Sensor Webs in the Wild

Kevin A. Delin

NASA Jet Propulsion Laboratory

Pasadena, CA 91109

Introduction

In October 2001, a new era in wireless sensor systems began when the NASA Sensor Web, deployed at the Huntington Botanical Gardens in Southern California, went online. For the first time, it was possible for a person with nothing more than a computer, an Internet connection, and a standard browser to watch streaming, real-time data generated by an *ad hoc* wireless networked system permanently embedded in an outdoor environment. Wireless networked systems, Sensor Webs in particular, are only just beginning to change the ways in which we can sense, monitor and control large spatial areas.

The Sensor Web's capabilities are useful in a diverse set of outdoor applications ranging from precision agriculture to perimeter security to effluent tracking. Wireless networks of sensors are often marketed as replacements for running wire to sensing points. Naturally this holds true for the Sensor Web as well, where the individual pods communicate among themselves wirelessly. However, it is more significant that the Sensor Web, with its unique global information sharing protocol, forms a sophisticated sensing tapestry that can be draped over an environment. This Sensor Web approach allows for complex behaviors and operations, such as on-the-fly identification of anomalous or unexpected events, mapping vector fields from measured scalar values and interpreting them locally, and single-pod detection of critical events which then triggers changes in the global behavior of the Sensor Web. This chapter describes the Sensor Web, a technology enabled by the confluence of the massive computer and telecommunications markets. In addition, it discusses the design and deployment of such systems in the wild, away from the temperature- and humidity-controlled rooms characteristic of offices and factories.

The Sensor Web: A Different Type of Wireless Network

In 1997, the Sensor Web was conceived at the NASA Jet Propulsion Laboratory (JPL) to take advantage of the availability of increasingly inexpensive, yet sophisticated, mass consumer-market chips for the computer and telecommunication industries and the use of

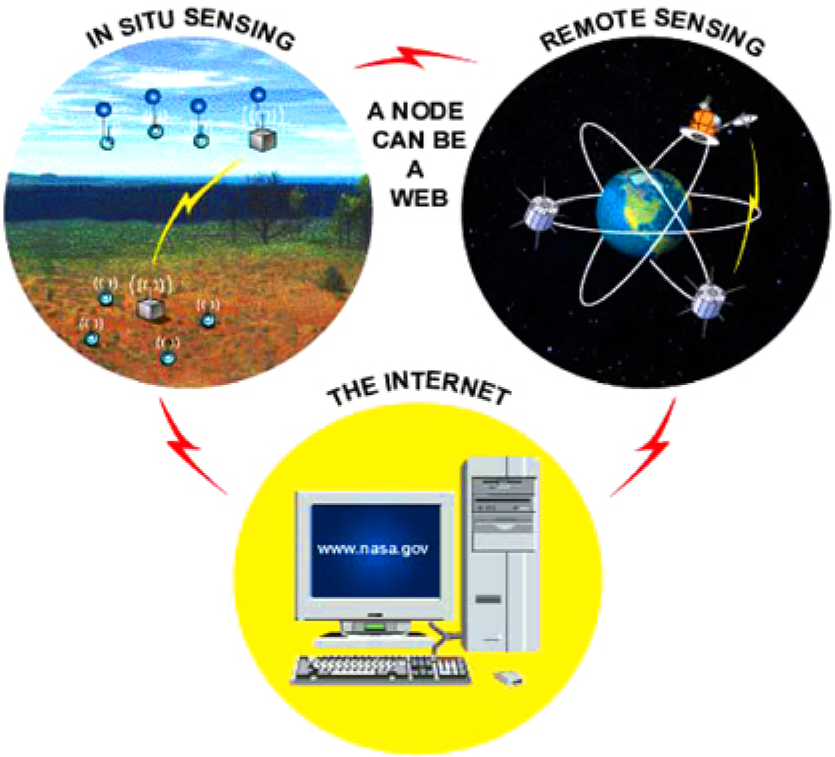

Figure 1: Generalized concept of the Sensor Web, including both orbital and terrestrial platforms. Note the recursive nature of the Sensor Web as individual nodes on a particular Sensor Web may be Sensor Webs themselves.

them to create platforms that share information among themselves and act in concert as a single system. This system would be embedded into an environment to monitor and even control it. The purpose of a Sensor Web system is to extract knowledge from the data it collects and use this knowledge to intelligently react and adapt to its surroundings. It links a remote end-user's cognizance with the observed environment.

In its most general form, the Sensor Web is a macro-instrument comprising a number of sensor platforms [1]. As shown in Figure 1, these platforms, or pods, can be orbital or terrestrial, fixed or mobile. Coordinated communication and interaction among the pods provides a spatio-temporal understanding of the environment. Specific portal pods provide end-user access points for command and information flow into and out of the Sensor Web. At NASA/JPL, the focus has been on *in situ* Sensor Webs, with the resulting system output viewed over the Internet.

Wireless networks are not a new approach to environmental monitoring and it is common to find systems where remote sensors in the field communicate to central points for data processing in a star-network formation. *The Sensor Web, however, is a temporally synchronous, spatially amorphous network,* creating an embedded, distributed monitoring presence which provides a dynamic infrastructure for sensors. By eschewing a central point on the network, information flows everywhere throughout the instrument (see Figure 2).

So far, this sounds like a typical *ad hoc*, self-configuring network. Often the ideas of hopping information around such a network are framed in terms of the power advantage gained by doing so. While this advantage certainly exists, the Sensor Web concept goes a step further: The individual pods comprising a Sensor Web are not just elements that *can* communicate with one another; they are elements that *must* communicate with one another. Whereas wireless networks are typically discussed as confederations of individual elements (like computers connected to the Internet), *the Sensor Web is a single, autonomous, distributed instrument*. The pods of a Sensor Web are akin to the cells of a multi-cellular organism; the primary purpose for information flow over a Sensor Web is not about getting data to a gateway or an end-user, but rather to the rest of the Sensor Web itself.

By design, the Sensor Web spreads collected data and processed information throughout its entire network. As a result, there is no design criterion for routing as in other wireless systems. Routing, by definition, is the focused movement of information from one point to another. In contrast, data collected by a Sensor Web is spread everywhere, rendering meaningless the concept of routing. Instead, the communication protocol on a Sensor Web is relatively simple and is structured for both omni- and bi-directional information flows. Omni-directional communication implies no directed information flow, while bi-directional communication lets individual pods (and end-users) command other pods as

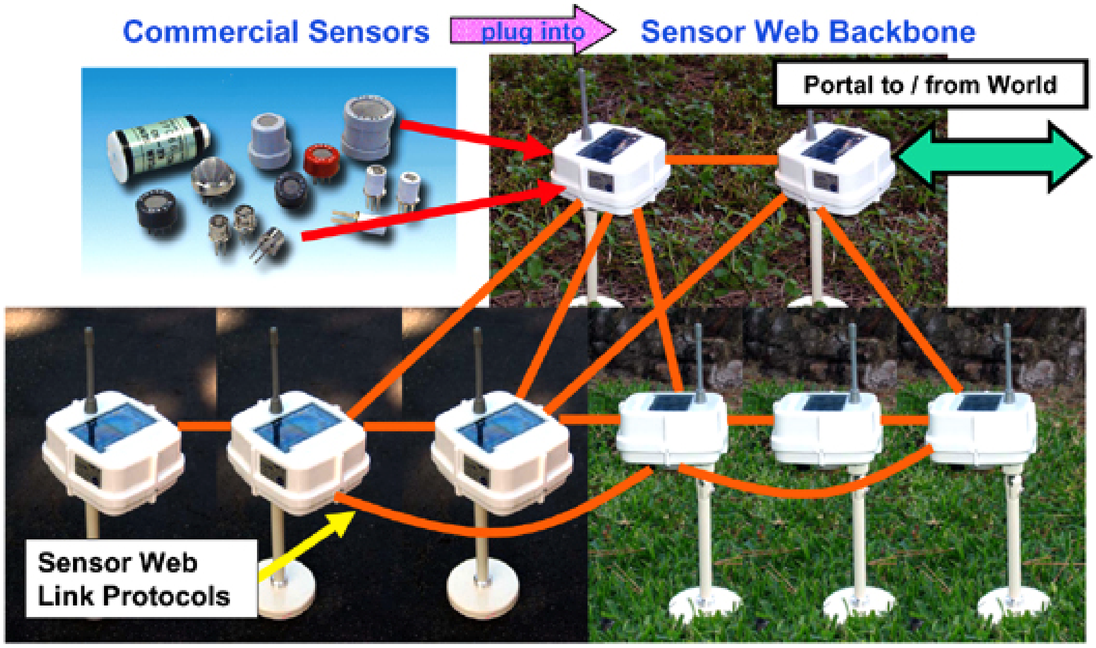

Figure 2: The Sensor Web forms an informational backbone that creates a dynamic infrastructure for the sensors in the Sensor Web pods. The total Sensor Web instrument includes both the individual pods and the space between them.

well as receive information from them. Consequently, information on the Sensor Web can result from four types of data: (a) raw data sensed at a specific pod, (b) post-processed sensed data from a pod or group of pods, (c) commands entered into the distributed instrument by an external end-user, and (d) commands entered into the distributed instrument by a pod itself. The Sensor Web processes this internal information, draws knowledge from it, and reacts to that knowledge.

Since there is no specific routing of information, all pods share everything with each other. After each measurement is taken, both raw and processed information from each pod are moved throughout the Sensor Web to all other pods before the next measurement is taken. Because the Sensor Web is a single, distributed instrument, its internal operations are synchronous from pod to pod (again in contrast to more common wireless networks). In this way, a total snapshot associated with that instant in time is available to all pods on the Sensor Web. This global data sharing allows each pod to be aware of situations beyond its specific location. Pods may therefore combine data across the Sensor Web to identify a moving front and determine its speed and direction – a task that a single-point measurement cannot accomplish. Pods may also use neighbors to examine the stochastic nature of their local measurements to determine whether or not the data collected are well-behaved. Such macroscopically coordinated data processing would not be as straightforward if each pod were semi-autonomous on the network, as in typical wireless sensor systems. There is a degree of stiffness to the information flow over the Sensor Web compared to the individually directed node-to-node information threads on more typical wireless systems. Sensor Web pods may be thought of as individual, synchronized pixels in a much larger instrument that can take snapshots at regular intervals of the entire environment in which it is embedded. Moreover, each pixel is simultaneously aware of the overall picture as well as its local readings.

Since there is no distinction between instructions originating from the end-users and those from other pods within the system, the Sensor Web is both a field-programmable and self-adapting instrument. Bi-directional communication also allows portal pods to be linked (by Internet or satellite) so that the end-user of a particular Sensor Web may be another web. As a result, a local Sensor Web can leap beyond the bounds of its own spatial confinement—an important consideration for very large-scale environmental studies.

Sensor Web Pods

A Sensor Web pod consists of five basic modules:

(1) The radio, which links each pod to its local neighborhood. The NASA/JPL Sensor Web pods use radios operating in the 900 MHz license-free Industrial, Science and Medical (ISM) band with an upper range of about 200 m or more. Communication occurs in bursts at 56 kbaud. (Implicitly, we assume none of the *in situ* Sensor Webs discussed here are deployed underwater, where acoustic modems severely limit bandwidth and communication range relative to these ISM radios.)

(2) The microcontroller, which contains the system's protocols, communicates with the attached sensors, and carries out data analysis as needed.

(3) The power system. The NASA/JPL system uses a battery pack with solar panels to keep the batteries charged. The combination of solar panels and micropower electronic design have kept Sensor Web pods operating in the field for years without requiring maintenance.

(4) The pod packaging. This key module is often overlooked, especially for Sensor Web applications in the wild. The package must be light, durable, inexpensive, and sealed against such elements as rain, snow, salty sprays, dust storms, and local fauna. In addition, it must provide for easy and rapid mounting.

(5) The sensor suite. This module is completely determined by the specific application. Ideally, the sensor suite will be the prime determining factor for the size, cost, and power requirements of a Sensor Web pod, making the Sensor Web infrastructure attractive for any application. What is considered an inexpensive or small Sensor Web pod in one application may not be viewed as such in another.

In the hyperbolic world of high-technology, engineering metrics are often based more on sound bites than sound principles. We have been conditioned by decades of experience with Moore's Law (and the technology revolution associated with it) to think that smaller is always better. There are certainly practical reasons for limiting the size of a Sensor Web pod. In an outdoor environment, smaller and lighter pods are easier to deploy since more can fit into, say, a backpack. However, shrinking pods to infinitesimal sizes is undesirable for a typical outdoor Sensor Web system. Consider the impact of size with respect to three key Sensor Web pod design issues: power, antenna size, and transducers.

An important design consideration for typical outdoor Sensor Webs is pod longevity. In many cases, deployment is only practical during certain seasons so intra-season maintenance must be avoided. As a result, maximizing the available power, by cleverly using batteries and/or energy harvesting, is critical. Batteries are often rated in terms of their energy density (watt-hours per unit volume), because cells can be added in series to increase total available voltage. The larger the volume of the Sensor Web pod, the more volume is available for power from any particular battery technology.

There are only two ways to maintain a given available battery power level while allowing the pod volume to shrink: improve the battery technology or reduce energy use within the pod. While there are numerous efforts to provide higher energy-density power sources than are typically available (e.g. lithium ion batteries), none are yet commercially available for consumer use. In addition, many experimental batteries have limited lifetimes. Moreover, any suitable battery technology must be essentially zeromaintenance and environmentally robust (especially to changes in temperature, both seasonal and diurnal). As for improving energy efficiency, the laws of physics require a certain power output to broadcast a given distance. Therefore, although one can lower the energy per bit involved in computation, the wireless communication puts a hard limit on how much energy will be required for the system to operate for a given pod-to-pod distance.

Now consider energy harvesting, which is typically accomplished via solar power charging secondary batteries. Here, too, the smaller the platform, the smaller the solar panels used to re-energize the system, and the smaller amount of energy that can be harvested for a given panel. Clearly, beyond a certain size, the smaller one designs a Sensor Web pod for a given set of operating parameters, the more one gives up in longevity with respect to available power.

Antennas are also directly related to platform size. Again, the laws of physics dictate the appropriate antenna geometry for a given operating frequency to ensure a proper impedance match into the radiated space. As a result, while the on-board processor and radio electronics may shrink, the antenna may not if a particular communication range is required. Without proper coupling, radiation efficiency is reduced and radio power must be increased to maintain range. We therefore find that, since most outdoor Sensor Webs require pod-to-pod ranges of at least tens of meters, indiscriminate shrinkage of the individual antennas clearly compromises the telecommunication subsystem.

Lastly, consider the sensors themselves. Many sensors used in outdoor field applications, though compact and inexpensive, are not Micro Electro-Mechanical System (MEMS) devices and therefore cannot be integrated into the Sensor Web pod at the chip level. Examples include gas sensors (which often require a certain volume of sample gas) and seismometers (which require a certain mass for appropriate mechanical resonant frequencies). As a result, for a wide variety of Sensor Web applications in an outdoor environment, the sensors will be additional components added to the basic pod platform. Clearly, there is little to be gained by continually shrinking the platform if the sensors themselves remain the limiting size element. Moreover, shrinking the platform may actually complicate the design if it becomes difficult to integrate the sensors into the pod.
As shown in Figure 3, the NASA Sensor Web pods have been developed in several sizes, ranging from that of a gumball to that of a couple of decks of playing cards. Significantly, the gumball-sized pod dates back to 1998 [1], demonstrating that, even then, it was relatively easy to make small platforms so long as only simple measured parameters (i.e. temperature, humidity, etc.) and short pod-to-pod communication distances (i.e. order of meters) were required. Such small pods are ideal for building or factory monitoring, but less practical for outdoor environments for the reasons discussed above. From this discussion, it is apparent that, while smaller pods are desirable, shrinking pods beyond a certain point leads to diminishing returns.

Figure 3: Various Sensor Web pods. Top Left: Functioning Sensor Web 1.0 pod, circa 1998. Note the small size which includes antenna, battery, and temperature and light sensors. Top Right: Sensor Web 3.1 pod deployed at the Huntington Botanical Gardens, circa 2002. It is about the size of two decks of playing cards. The pod is mud spattered from rain and watering and has a chewed antenna. Subterranean sensors (soil moisture and temperature) can be seen going into the ground. Bottom Left: Sensor Web 3.1 pod deployed less than a quarter mile from Space Shuttle Launch Pad 39A at NASA Kennedy Space Center (note shuttle on pad). This Sensor Web measured environmental conditions in the lagoons surrounding the launch site. Bottom Right: A Sensor Web 5.0 pod, circa 2004. This new generation of Sensor Web pods is more compact and more power efficient than previous ones, a direct result of exploiting Moore's Law in its design.

Sensor Web Deployments

With the objective of performing genuine *in situ* monitoring, the NASA/JPL Sensor Webs Project has aggressively fielded many instruments. Sensor Webs have been deployed in a large variety of demanding real-world locations for lengthy periods of many months or even years. Here, we summarize some representative deployments. A more complete list, along with real-time, streaming data from current deployments, is available over the Internet [2].

Huntington Botanical Gardens, San Marino, CA: Sensor Webs have been at the Huntington Botanical Gardens in San Marino, CA since the deployment of Sensor Web 2.0 in June 2000. Sensor Web 3.0, the first permanent wireless sensor network system to provide continuous real-time streaming data to users over the Internet, was deployed in October 2001 in a nursery area, remaining there until August 2002. At that point, it was replaced by Sensor Web 3.1 which also expanded the spatial extent of the coverage into the public areas of the gardens. In all these systems, canonical botanical parameters, including light levels, air temperature, and air humidity, are measured by the pods in 5 minute intervals.

Significantly, the deployment of Sensor Web 3.1 occurred in two stages. Pods 0 (the mother or portal pod) through 11 were deployed in August 2002. The purpose of this initial deployment was to string the pods out over a large area and determine the robustness of the connectivity. It typically took 4 or 5 hops for the data to move to and from the extreme points on the network, several hundred meters apart. (Shorter duration system tests have demonstrated that 12 or more hops are possible with present protocols.) The second stage of deployment occurred in late January 2003. In contrast to the first stage, it was performed with extensive input from the Huntington staff to help monitor their key areas of interest. As a result, the pod placement is more confined. As anticipated by the Sensor Web protocol design, the second set of pods seamlessly integrated with the first set and the new 19-pod system coalesced within a few measurement cycles. This was the second demonstration of augmenting an existing and functioning Sensor Web deployment, the first having occurred at Lancaster Farms in Virginia in the spring of 2002.

Several Sensor Web pods in the current deployment measure soil temperature and moisture with modular subterranean probes (see Figures 3 and 4). The Huntington Gardens staff has used the system to remotely monitor the state of their greenhouses and to ensure that watering (from both sprinkler and rainfall patterns) is uniform across various areas. Rains in February 2003 showed the surprising variations in these parameters, even on scales of a meter. This suggests using Sensor Webs to provide continuous resource management information on both micro- (of order 1 meter) and meso- (of order 100 meters) spatial scales that is unattainable from most remote (e.g. satellite or airborne) measurements.

The Huntington Gardens deployment covers a large number of micro-environments. Sensor Web pods are located in greenhouses, in the open sun, and in shady areas, while the portal pod remained in a temperature-controlled room. As a result, the Sensor Web, a single instrument system, was subjected to significant differential heating, a stress not encountered in typical indoor deployments. Nevertheless, it is possible to compensate the electronics to handle this issue while maintaining system integrity and synchronicity. In addition, one pod was found with an antenna chewed, probably by a bird or coyote. This abuse did not interrupt pod functionality. Sensor Web 3.1 at the Huntington has run continuously for over 2 years showing the robustness of both the system protocols and pod packaging.

Antarctic Wilderness: In cooperation with researchers hunting for meteorites, a Sensor Web was deployed on the East Antarctic Ice Sheet [3]. International treaty prevents leaving unattended equipment in the remote fields, so this deployment was limited to 3 weeks starting in December 2002. Nevertheless, the Sensor Web had to prove itself in an extreme test environment with low-temperatures, constant wind, and dry air. In addition, the deployment was made under very hard conditions with researchers fully suited up for

Figure 4: Sensor Web 3.1 pods 12 (rear, upper center of picture) and 13 (foreground, lower left corner) deployed at the Huntington Gardens around cycads. One cycad trunk is partially seen on right side of picture, the fronds of a second cycad are seen just to the right of Sensor Web pod 12. Subterranean soil temperature and moisture are attached to both pods. Differing subterranean soil moisture conditions between these two pods alerted gardeners that the two plants were not receiving identical watering. It was discovered that ground brush around one of the cycads was acting as a sponge, retarding water entry into the ground, and was subsequently removed.

the bitter weather as shown in Figure 5.

The deployment was a test of the system in preparation for extended studies of biological activities in cryogenic environments, especially those on Mars. The more hostile the environmental conditions, the more widely distributed any biological blooms are expected to be in both time and space. The continual monitoring presence of the Sensor Web, coupled with its abilities to do event-triggered sensing, make it an ideal instrument ideal for such studies. For example, if one pod detects that its local conditions are more favorable to a thaw which might activate a bloom, the entire Sensor Web could begin making measurements more frequently to better understand the short-term event. Thus most data are accumulated only when a significant event occurs, further conserving power and memory resources on the Sensor Web. In addition, the Sensor Web provided data across the region that minimized stochastic components resulting from specific pod placement.

The 14-pod Sensor Web was distributed about the home base site over a distance of 2 km. Parameters such as air and soil temperature, humidity, and light flux were measured across the Sensor Web at 5 minute intervals. System operations could be monitored

Figure 5: Sensor Web 3.0 deployment on East Antarctic Ice Sheet. Bamboo is used as the mounting stake for its known resilience to the harsh Antarctic conditions. Wires connect the pod to the temperature sensors at the surface. Note the second pod located in the snowmobile in lower right corner.

continuously at the home base site. Typical temperatures over the deployment were lower than –10 °C with extremes below –20 °C. From this deployment, it was discovered that temperature differences between sensors exposed to the air and those buried in a moraine could differ consistently by 10 °C or more. The Sensor Web system not only performed well under the harsh, dry conditions, but was also easy to set up. (As is often the case, no members of the NASA Sensor Web team assisted in the deployment.) Under severe weather conditions, such issues are critical for mission success. Significantly, these pods were identical to the ones that were deployed at the Huntington Gardens; there were no special modifications made for the different environment.

Flood Basins, Tucson, AZ: We have deployed a Sensor Web at the Central Avra Valley Storage and Recovery Project (CAVSARP) facility located west of Tucson, AZ [4]. The facility is located in a desert environment in the semi-arid Southwest United States where the artificial recharge basins experience repeated flood cycles. The controlled flooding conditions at the CAVSARP facility are ideal for the investigation of various hydrologic processes. There are several technology-related reasons for this site choice as well. The CAVSARP facility allows us to continue our efforts to develop the Sensor Web as a tool for the study of spatio-temporal phenomena. For example, the Sensor Web can track the moving flood front, follow the infiltration of water into the ground, and provide information to map and characterize the lateral and vertical extent of the floodwaters. Moreover, the extreme temperature variations of the Arizona desert (both diurnal and seasonal) provide yet another test of the Sensor Web's robustness. Again, this Sensor Web hardware was essentially identical to that previously deployed with no special modifications made for the new environment.

A single basin, measuring approximately 700×2400 ft^2, was strategically outfitted with 13 pods, their number and placement being determined by science requirements rather than technological limitations. The pods were mounted on stakes to elevate them above the flood waters, which can rise as high as 7 ft (see Figure 6). (While the pods themselves are water-tight, pod-to-pod radio communication would not be possible if they were submerged.) Each pod, in addition to collecting air temperature, humidity, and light levels, also collects two soil moisture readings (one at the surface and one 0.5 m below) and a surface soil temperature reading. This is accomplished by wires that run from the pod into the ground. Measurements are made at 5 minute intervals with the results being fed to the Internet in real-time (via the portal pod 0).

This Sensor Web has been collecting data since its deployment in November 2003. The real-time data stream is available via the Internet [2], with a sample screen-capture shown in Figure 7. Unlike remote techniques, which can only observe the basins for relatively short durations on finite schedules, the Sensor Web's data stream provides continuous information for tracking surface water motion and ground infiltration. The spatial and temporal patterns of wetting and drying can thus be fully monitored with results incorporated into hydrological models and compared with space-based and airborne investigations [5]. This analysis is ongoing. As a result, this Sensor Web can both augment and ground-truth the remote data traditionally used in hydrologic studies.

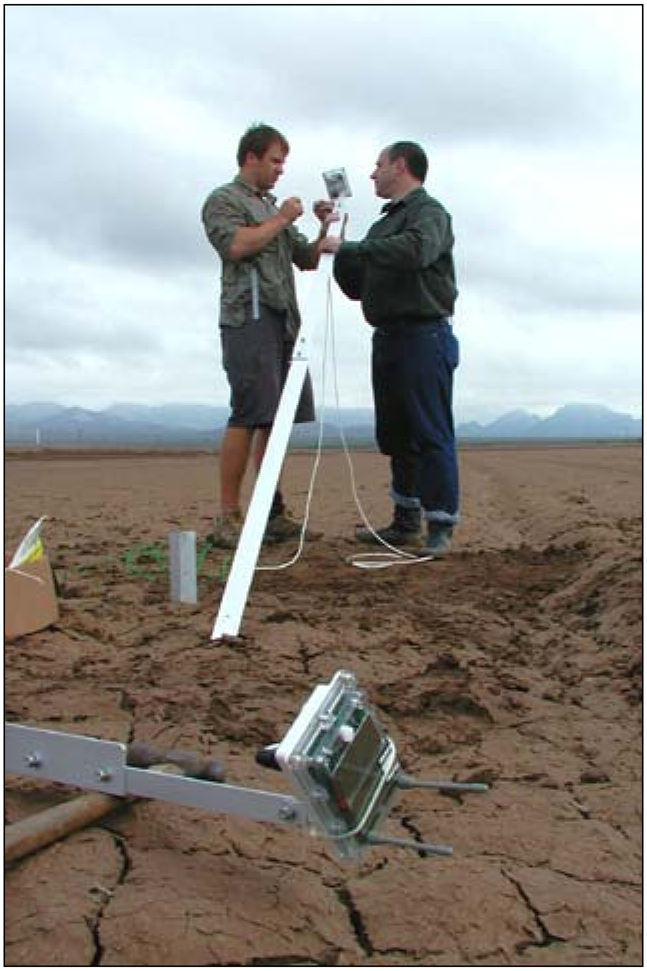

Figure 6: NASA/JPL team members deploy a Sensor Web 3.2 pod in a flooding basin outside Tucson, AZ. Extended stands allow the pod to stay operational above water during a flooding event.

The repeatable nature of the flooding/drying dynamics is apparent from the Sensor Web data shown in Figure 7. The soil moisture measurements are made with Watermark sensors [6] in which electrodes embedded in a granular matrix have a lower resistance when the surrounding soil is wetter. As a result, the raw data reveals the motion of the flooding water as sharp drops in resistance. (The diurnal cycles seen in the raw data are sensor artifacts and can be corrected with soil temperature measurements [7].) It only takes a few hours for the flood front to traverse the basin from the inlet in the northwest corner to the basin center, but a much longer time (about 20 hours) to reach the basin's southwest and southeast corners. Note, too, that the water reaches the southern border relatively evenly (as indicated by pods 1 and 11), which is expected from basin construction. Moreover, it is also clear from Figure 7 that the drying front traverses the reverse route, albeit at a much slower speed. Not surprisingly, the surface dries more thoroughly than the deeper portions of the ground.

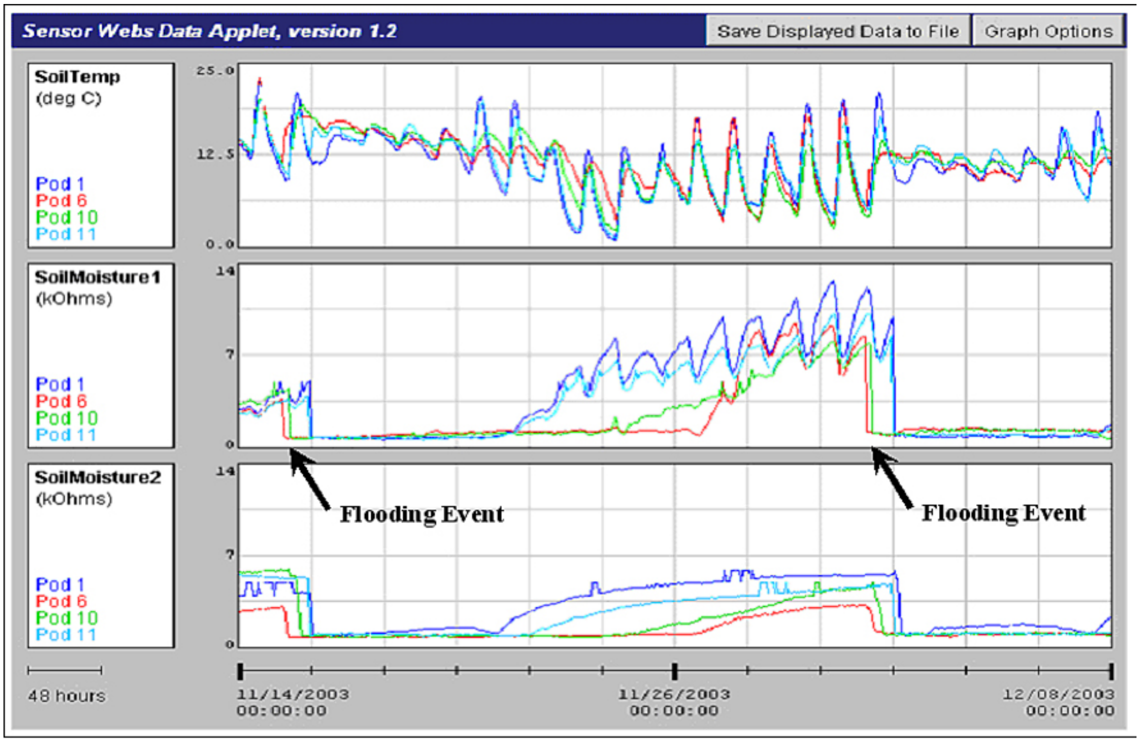

Figure 17.7: Screen-capture of Internet data from CAVSARP facility. Graphs (top to bottom): surface temperature (°C), surface moisture, and soil moisture at 0.5 m depth (relative units; lower values imply wetter soil). Diurnal cycles in moisture measurements are artifacts and can be corrected with local soil temperature. When the basin floods, Sensor Web pod 6 (in red, basin inlet, northwest corner) is the first to detect the water front, followed by pod 10 (in green, basin center), pod 11 (in light blue, southeast basin corner, diagonal from pod 6), and, lastly, pod 1 (in dark blue, southwest basin corner). The direction of inundation can be determined from this data. Note that the data also show that the basin dries out in reverse order of flooding. The traces correlate with water discharge into basin, inundation, infiltration, drying, and the beginning of another cycle.

Deployment Operations

From the experience of deploying Sensor Webs in a multitude of environments with varying conditions, it is apparent that the ease with which the system is deployed is just as critical for acceptance by end-users as are its technological aspects. With the exception of applications in battlefield theaters, most outdoor Sensor Web applications require the system to be deployed in manner that does not harm the monitored environment. For example, end users have expressed concerns that if Sensor Web pods are too small, local fauna may try to ingest them and choke. End-users also want to avoid littering local environments with hundreds of pieces of microelectronic gear. It is therefore worthwhile to consider Sensor Web technology from an operations point of view.

Most applications require tracking specific pod locations to a high precision. It is therefore highly unlikely that any wireless sensor network will simply be sprinkled over large areas. In addition, coupling sensors into the environment will also prevent such a passive deployment. For example, neither subterranean nor seismic sensors can be deployed by a sprinkling technique, as both require laborious efforts for appropriate sensor mounting. There are also applications, particularly those involving agriculture, where pod placements must be compatible with existing operations, such as harvesting. For example, Sensor Web pods must be mounted out of the way of threshers. Consequently, the mounting and placement of Sensor Web pods will be an active operation and likely to be done by hand in most instances.

In addition, we have encountered few applications in the wild where an average pod-to-pod distance is less than 10 meters. It is therefore necessary for the Sensor Web pods to have a wireless range limited only by the government-mandated ISM band specifications. A more limited wireless communication would not be suitable if the density required to maintain network connectivity exceeds the desired measurement density required to monitor the region. Therefore, a critical metric is the *Sensor Web pod density* rather than the more usually specified pod number. Ideally, the density of a Sensor Web should be just greater than that required by the application itself, to ensure connectivity and internal communication redundancy within the instrument.

The methods used to mount the Sensor Web pods depend not only on the application but also the particular field site. Pod placement very close to the ground can limit transmission distance. Nevertheless, the Huntington Garden pods are within 10 inches of the ground and have sufficient communication power to keep an adequate pod-to-pod distance. Often, for logistical reasons, the Sensor Web pods tend to be mounted higher off the ground, with the attendant benefit of increasing the wireless range. Local terrain is rarely level which also tends to increase transmission distances. We have typically used posts (for horizontal surfaces) and brackets (for vertical surfaces) to mount the pods. These types of mounts are both small enough and light enough to bring into the field yet provide proper robustness for fixing the Sensor Web pods rigidly in place for long durations.

Sensor Web pod placement is probably the most application-specific issue: each deployment is different depending on local terrain, even if applications are for identical purposes. It is always necessary to have a clear understanding of the terrain to ensure proper pod density for both measurement and web connectivity. General purpose algorithms for "optimal" pod placement, which, by definition, do not account for local geometry or the specifics of an experiment, tend to be of limited use. In fact, pod placement has had the least systematic study and will likely remain an open issue for each end-user community rather than the technologists developing the Sensor Web systems.

The Future

The primary focus of Sensor Web development thus far has been to demonstrate that the technology is stable, robust, and attractive to potential end-users. For a user community to adopt it, however, the Sensor Web needs to be more than just well-engineered; it must also be easily deployed and maintained and provide valuable output. The overall simplicity of the Sensor Web system as an operational instrument is demonstrated by the fact that most Sensor Webs are deployed and operated in a variety of environments without requiring assistance from the NASA/JPL team.

Having demonstrated the Sensor Web's core capabilities with a myriad of deployments, we are now moving Sensor Web development into a new phase, focusing as much on applications as technology. The continuous, virtual monitoring and reacting capabilities have wide ranging uses for resource management, pollutant tracking, and perimeter monitoring. This step requires us to take full advantage of the large-scale awareness already built into the Sensor Web protocols. In this way, the output of, for example, a hydrologic Sensor Web would not consist of a collection of scalar measurements (soil moisture) but rather a single vector (water motion) with pod-to-pod data fusion occurring within the Sensor Web itself. This will truly give the Sensor Web the capacity to make sophisticated, autonomous decisions. As awareness of this unique distributed instrument and its capabilities spreads into other user communities, the Sensor Web is expected to become a dominant wireless sensor network architecture.

Acknowledgements

I wish to acknowledge the significant contributions of Scott Burleigh, Shannon Jackson, Dave Johnson, Ken Manatt, Myche McAuley, and Richard Woodrow in developing the Sensor Web instrument. Fielding the Sensor Webs would not have been possible without the active participation of my field site collaborators including Julio Barrantes, Greg Bonito, Eric Bender, Renee Brown, Nancy Chabot, Scott Collins, James Dohm, Jim Folsom, Ralph Harvey, Felipe Ip, Doug Moore, Danielle Rudeen, Jonathan Sax, and Theresa Trunnelle. The Sensor Webs described here were constructed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA.

References

[1] Delin, K.A., "The Sensor Web: A Macro-Instrument for Coordinated Sensing," *Sensors,* Vol. 2, 2002, pp. 275-280.

[2] http://sensorwebs.jpl.nasa.gov/

[3] Delin, K.A., R.P. Harvey, N.A. Chabot, S.P. Jackson, Mike Adams, D.W. Johnson, and J.T. Britton, "Sensor Web in Antarctica: Developing an Intelligent, Autonomous Platform for Locating Biological Flourishes in Cryogenic Environments," In *Proceedings of the 34th Lunar and Planetary Science Conference*, Houston, TX, March 2003.

[4] Delin, K.A., S.P. Jackson, D.W. Johnson, S.C. Burleigh, R.R. Woodrow, M. McAuley, J.T. Britton, J.M. Dohm, T.P.A. Ferré, Felipe Ip, D.F. Rucker, and V.R. Baker, "Sensor Web for Spatio-Temporal Monitoring of a Hydrological Environment," In *Proceedings of the 35th Lunar and Planetary Science Conference*, League City, TX, March 2004.

[5] Rucker, D.F., J.M. Dohm, T.P.A. Ferré, F. Ip, V.R. Baker, A.G. Davies, R. Castano, S. Chien, B. Cichy, T.C. Doggett, R. Greeley, and R. Sherwood, "Central Avra Valley Storage and Recovery Project (CAVSARP) Site, Tucson, Arizona: Floodwater and Soil Moisture Investigations with Extraterrestrial Applications," In *Proceedings of the 35th Lunar and Planetary Science Conference*, League City, TX, March 2004.

[6] Available from Irrometer Co., http://www.irrometer.com/

[7] Shock, C.C., "Soil Water Potential Measurement by Granular Matrix Sensors," In *The Encyclopedia of Water Science*, pp. 899-903, B.A. Stewart, and T.A. Howell (eds.), Marcel Dekker, 2003.