

Terra-Scope™ - a MEMS-based Vertical Seismic Array

Steven D. Glaser, Min Chen, Tom E. Oberheim
Center for Information Technology Research in the Interest of Society,
University of California, Berkeley
California, USA

ABSTRACT

The Terra-Scope™ system is an affordable 4-D down-hole seismic monitoring system based on independent, microprocessor-controlled sensor Pods. The pods are nominally 50 mm in diameter, and about 120 mm long. They are expected to cost approximately \$6000 each. An internal 16-bit, extremely low power MCU controls all aspects of instrumentation, eight programmable gain amplifiers, and local signal storage. Each pod measures 3-D acceleration, tilt, azimuth, temperature, and other parametric variables such as pore water pressure and pH. The following parameters are independently controllable at each pod: pre-trigger length, post-trigger length, trigger time stamp, sampling rate, trigger level, trigger parameters, non-volatile storage, and calibration and self-evaluation. Each Pod communicates over a standard digital bus (e.g. RS-485) through a complete web-based GUI interface, and has a power consumption of less than 400mW. Three-dimensional acceleration is measured by pure digital force-balance MEMS-based accelerometers. These accelerometers have a dynamic range of more than 115 dB and a frequency response from DC to 1000 Hz. The accelerometer chip uses a 5th order delta-sigma feedback loop to yield a noise floor of less than 30 ngrms/√Hz. Accelerations above 0.2 g are measured by a second set of MEMS-based accelerometers, giving a full 160 dB dynamic range. The prototype of the device is currently undergoing evaluation. The first array will be installed in the fall of 2005.

Keywords: vertical seismic array, MEMS-based sensor, and real-time system.

INTRODUCTION

The increasing availability of vertical seismic array data from around the world has broadened our ability to analyze wave propagation and site response in the near surface. With accelerometers at depth, we can trace the actual effects of the near surface materials on the propagating seismic waves. So-called vertical arrays have come online in several sites in California, Taiwan, and Japan. By comparing multiple downhole recordings and a related surface recording, one can observe how the waves change as they progress through the ground, encountering the materials in the soil profile. We are installing two prototype vertical seismic arrays directly on either side of the Hayward Fault on the UC Berkeley campus. The arrays will consist of eight pods spaced over the 35 m embedment depth, and will be unique in the world for its proximity to a major active fault, its integration into a local and regional seismic system, and its low installation cost due to the use of currently unfolding micro-electronics and computer science devices and methodologies.

Our array uses a powerful yet inexpensive and easy to install down-hole seismic tool called Terra-Scope™. Each station of the array consists of 3-component accelerometer units, tilt sensors, magnetometer, and pore pressure sensor. The sensors will be incorporated into an intelligent networked sensor Pod that includes upwards of 4 Mbytes of non-volatile memory, an 16-bit micro-controller, 62 channels of digital interfaces, 8 analog channels, an Ethernet client, power controller, and backup batteries. Since each station is a digital entity, all communications along the array can take place on a bi-directional backbone. The terminus of the array, at the top of each bore hole, will be a powerful Local Gateway that will serve as data aggregator, GPS-based timer and locator, power source, Ethernet server, and wireless data link. We estimate that each Pod will cost is about \$6000 using present-day costs.

The system uses a simple software solution based on a custom event-driven scheduler to multiplex the concurrent flows of information across each gateway, which is connected to a transceiver, a secondary storage device, a sensor oriented I/O system, and a power management subsystem. Compared with traditional data logging systems, networked sensors offer two major advantages: they can be retasked in the field and they can easily communicate with the rest of the system.

VERTICAL ARRAYS

A down hole seismographic array is a sequence of instruments installed at various depths in the ground in a set of adjacent boreholes, to record the ground motion (generally in terms of velocity or acceleration) at multiple depths and at the surface of a site during an earthquake. As the authors and others (e.g. Elagmal et al, 1996, Stiedl et al, 1996) have pointed out over the years, vertical array data provides a direct way to quantify the accuracy of ground motion predictions, the models used to make such predictions, and the applicability of our estimates of in situ soil properties.

So-called vertical arrays have come on-line in several sites in California (Baise and Glaser, 2000; de Alba and Faris, 1996; Archuleta et al., 1992), Taiwan (Tang et al., 1989), and Japan (Katayama et al., 1990). These arrays are changing the way we understand seismic ground motion by allowing the 3-D evaluation of seismic wave propagation (Baise and Glaser, 2001; Stidham et al., 1999; Elgamal et al., 1996; Abrahamson et al, 1991). Vertical arrays have traditionally been very expensive to install and maintain (Steidl and Nigbor, 2001), much like traditional structural seismic instrumentation. We envision a field-deployable array that can be installed in a variety of manners. Initially the arrays will be installed into uncased (or cased) boreholes. By the end of the project, the arrays will be inserted into the ground by ubiquitous cone penetrometer equipment, and the array elements being commercial available to the entire community. A cartoon of the scope of the system is shown in Figure 1.

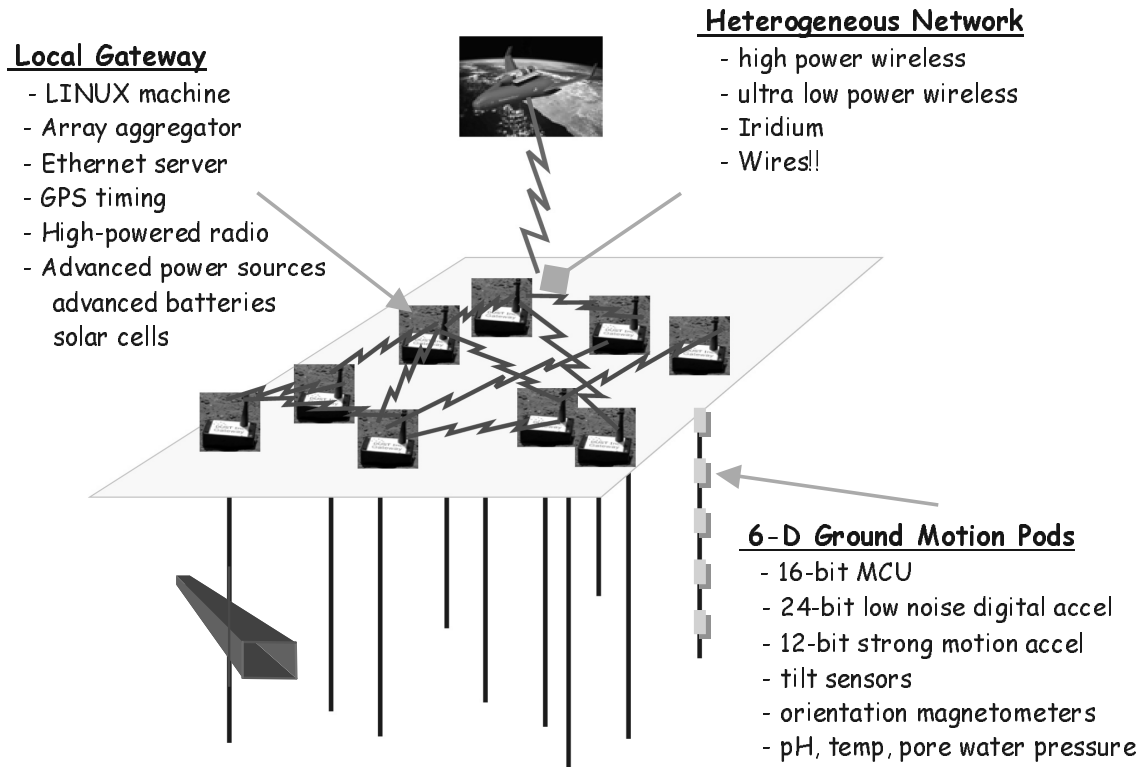


Figure 1 Cartoon of the vertical seismic array and Terra-Scope™ instrumentation

The low unit cost is made possible by utilizing new MEMS-based microsensors. MEMS (micro-electromechanical systems) devices are micro scaled machines that are manufactured using the same microfabrication techniques used for integrated circuits. The low unit costs of ICs are made possible by the extremely large volumes turned out per fabrication facility. The total investment per fab is astronomical, but tens of millions of each device are sold each year. Civil-oriented programs will never have the funding to sponsor the development and manufacture of a custom sensor that meets a researcher's every wish. We can, however, leverage the deep pockets of entities such as DARPA and the automotive industry. At present there are tens of MEMS-based accelerometers readily available offering a wide range of specifications. These range from the ADXL202e: \$10/each or less, very accurate, 9-bit dynamic range, to extremely low noise 115 dB fully digital devices that currently cost \$1000/each. We are also leveraging the many millions of dollars invested annually by DARPA in the smart dust and NEST programs on the Berkeley campus.

TERRA-SCOPE™

Networking Gateway

Compared with traditional data logging systems, networked sensors offer two major advantages: they can be retasked in the field and they can easily communicate with the rest of the system. In situ retasking allows the operator to change system characteristics based on the analysis of the initial results. Suppose that initially we want to collect localized ground motions; however after the initial interpretation of the data we might desire to change the scale granularity and utilize the entire array as a single sensor through beamforming. Individual sensor nodes can communicate and coordinate with one another. If appropriate, the network can perform in-network aggregation. At the head of each array sits a local gateway (LG), about a 75 mm cube. It serves as the host of the absolute time master clock (either from GPS or WWV) as well as an Ethernet host to store and distribute data and information out to the world. The LG is based on a variable clocked RISC-based, low power (1.5 amp) LINUX machine with large static memory and an embedded Ethernet server. The WAN connection is wireless (e.g., direct radio link, two-way satellite) or wired as is convenient. The LG will aggregate and process the multiple data streams from array stations, and either push the data onto the web or store it until queried by the main server. Included in the LG is a GPS which will provide exact timing for all the array stations, as well as providing accurate array location. Local batteries are charged by an attached solar array, and when possible a hard-wired powered Ethernet link will provide ample power for the entire array. The system architecture addresses the possibility of disconnection at every level. Each layer (sensor nodes, gateways, base stations) has some persistent storage which protects against data loss in case of power outage. Each layer also provides data management services. While many types of communication can be unreliable, when it comes to data collection, long-latency is preferable to data loss.

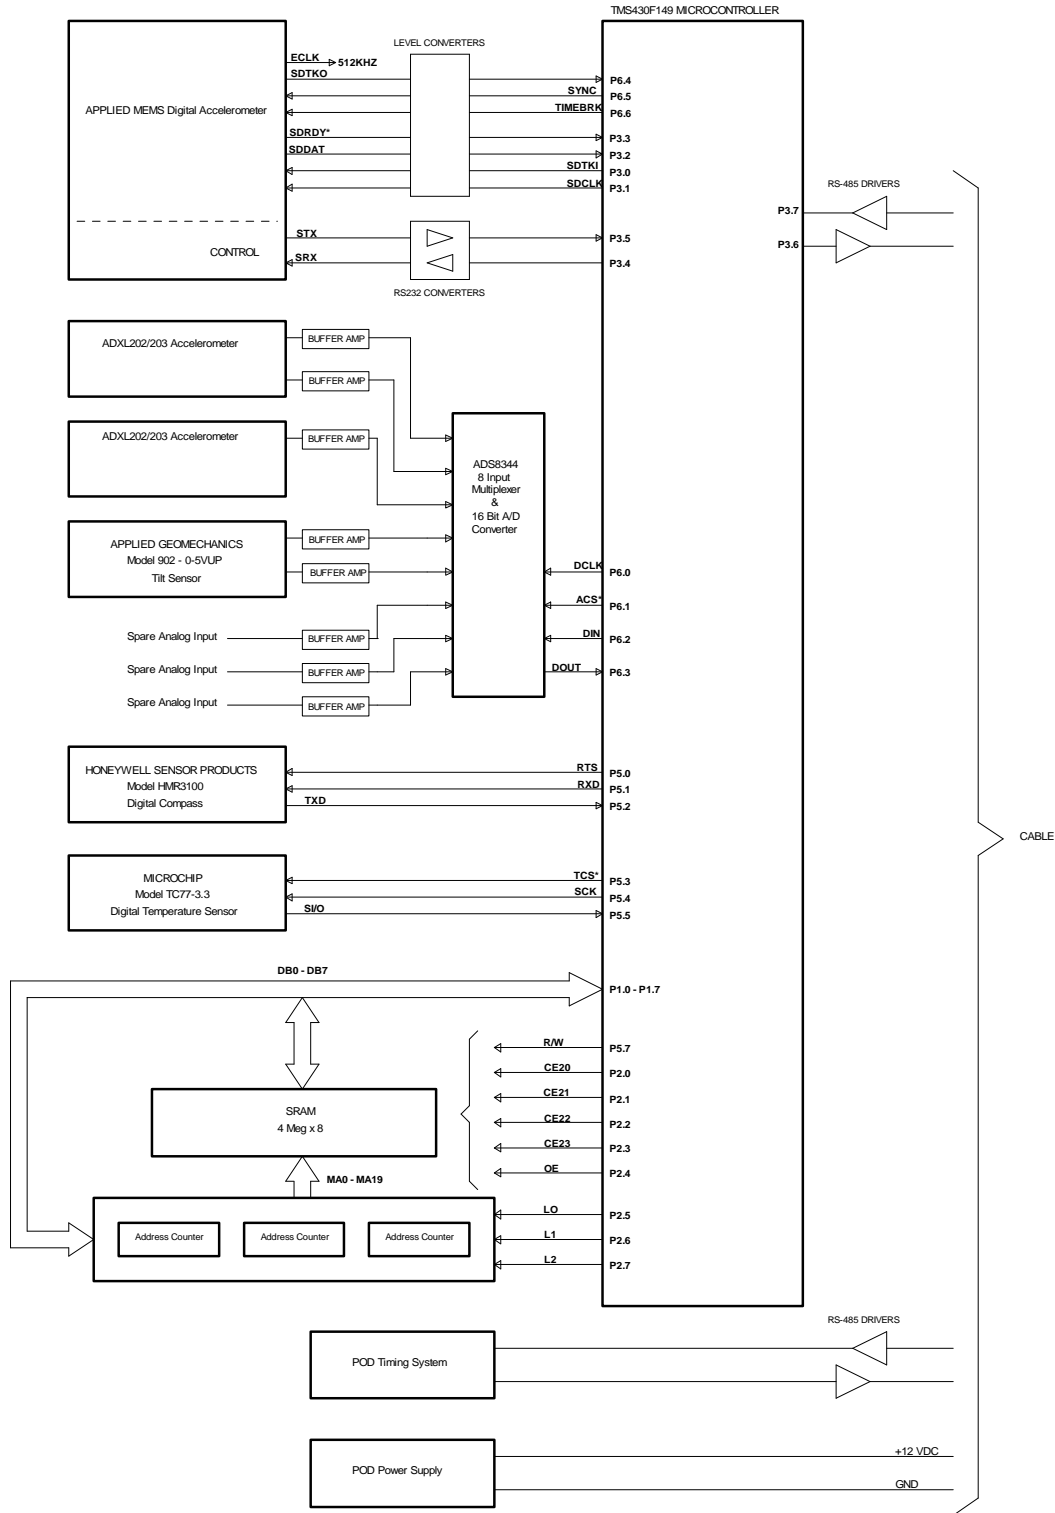
Pods

The "Pod" is the active independent down-hole measuring device that is the heart of Terra-Scope, shown as a block diagram in Figure 2. Each Pod is an independent, MCU-controlled agent with four or more Mb of non-volatile memory. The system is designed to implement the following

- Integration of advanced technology accelerometers (30 ngrms/ $\sqrt{\text{Hz}}$ noise floor, 24-bit direct digital); second set of accelerometers extends dynamic range to 1.7g;
- Tilt (pitch and roll) with 0.003 degree repeatability;
- Azimuth ~ 0.1 degree repeatability;
- Real-time clock within 0.5 ms;
- Parametric measurands include temperature, pH, pore water pressure, etc.;
- Record all dynamic variables at a digitization rate of 250 Hz for 1 min. before trigger and 2 after;
- Local Gateway integrates solar power, improved batteries, GPS timing, and Ethernet server;
- Fully dynamic networking, real-time reprogramming and peer-to-peer sensor fusion.

System Vision

We are basing our device on the concept of an instrumentation processor - an "open source" solution, if you will, to embedded networked sensing. This intelligent platform will accept any sensor (analog or digital) as an input signal, process and buffer the signals, and interface with any commercial-off-the-shelf (COTS) radio module, embedded system, or computer. The main driver for development of this technology are the many problems that researchers have had in using the Mote devices available from vendors such as Crossbow Inc., and Dust Inc. These devices are in practice



5765-41 V. 1 (p.4 of 9) / Color: No / Format: A4 (Letter) / Date: 2005-02-06 16:31:18

Figure 1. Block diagram of Pod

File		
POD - BORE HOLE DATA ACQUISITION SYSTEM		
Size	Document Number	Rev
		1.0
Date		
Thursday, December 30, 2004 Sheet 1 of 1		

SPIE USE: ___ DB Check, ___ Prod Check, Notes:

sole-source packages, and the user is beholden to the manufacturer to make them work properly. Experience using the Motes by relatively competent users (e.g. Berkeley, LBNL, UCSD, LANL, MIT, Seneca Inc., Shinkawa Sensor Technologies, etc.) has shown that the devices most often do not work as advertised. It is our opinion that a scientist or engineer who wants to use the wireless sensor node technologies should not have to earn a Computer Science Ph.D. in order to do so. What is readily available in COTS form, with many varieties and alternatives are MEMS-based sensors and radio frequency (rf) data transmission links. The user will select a wired or wireless link that best serves their need, at a price they can afford. In fact, advanced processing of complex sensor data is the space from which commercial developers are moving away from. Ember Inc. calls itself “experts in embedded wireless networking,” offering customers low-cost, low-power, and high performance best-in breed radio chips, embedded networking software optimized for control applications, and straightforward evaluation, development, and network management tools and services. The business plan of Dust Inc. has moved towards the data streaming model, providing the link for simple data such as thermostats or security. The science/engineering user is not a market driver and will always be left out.

System Design

The system under consideration will be used to measure seismic signals traveling through soil and rock. Given that these are very dispersive materials, the frequencies of interest are very low, up to about fifty Hz at the most. To cover any eventualities and to allow using the systems to perform other types of monitoring, the design criteria is to allow digitization rates of up to 250 samples per second, or a Δt of 4 ms. For a digital system running at 4 MHz, this data rate can be considered “slow,” and many tasks ordered and completed between samples. Figure 3 depicts the system architecture.

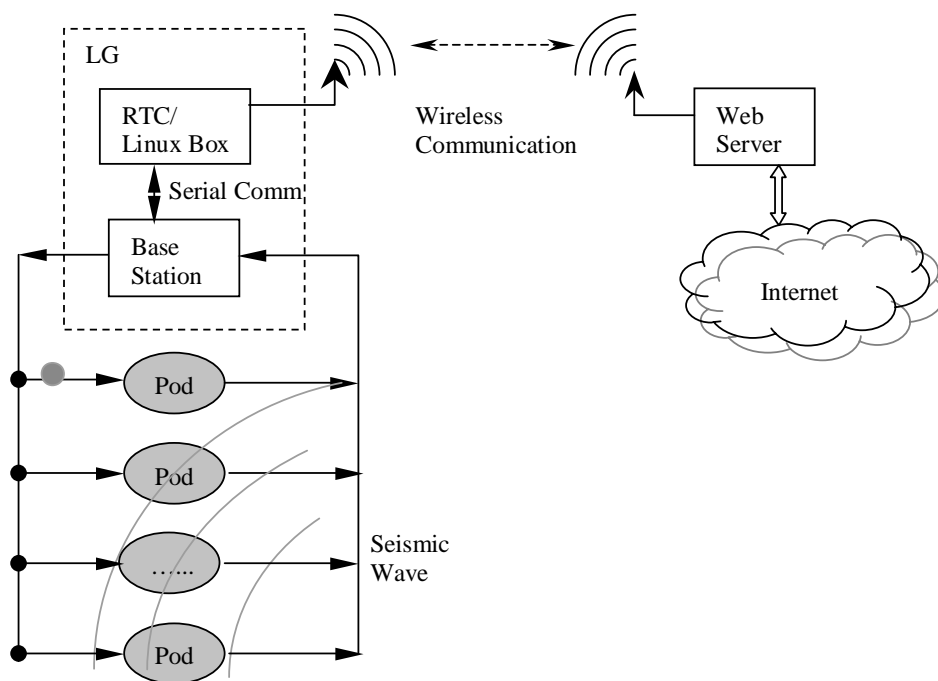


Figure 3. System Architecture

The system is designed so that a second three-minute event can be recorded to SRAM while the first segment downloads independent of system operation. The nominal one minute pre-trigger segment insures that motion occurring before strong shaking is captured. A nominal two-minute post-trigger recorded segment insures that all motions and

aftershocks are recorded. Each gateway is the command center for an array; multiple hardware pods are suspended on a bus cable at different levels, synchronized with the global timer from gateway, and sample data and communicate with base station. Two buses provide both physical connection and communication channel between hardware pods and base station. The base station is a single board machine with Linux operating system. Its main responsibility is to (1) send the real time clock to the pods, (2) connect the users/operators and the system via a web server to allow remote and effortless interaction over the Internet, (3) connect the pods by sending and relaying commands, and (4) and receiving event data from the Pods and storing the data into the database. This will be accomplished by building two applications to run concurrently on each gateway, one to control the Pods and arbitrate tasks, and a GUI web application to interact with the user. At all times the gateway will send regular timing packets to the pods over the bus or power lines.

In order to maximize the number of pods that can be attached to a single gateway, we will utilize two independent bus lines, one for message transmission and the other for data return. Upon start-up the gateway will perform initialization and synchronization of the array Pods. Each Pod lives an independent existence, with its MCU controlling the local sensors and setting a programmable trigger threshold. The 16-bit MCU is also responsible for controlling and segmenting the local four megabytes of SRAM, pre- and post-trigger data lengths, and trigger type. Each Pod waits primed and ready for an event. Upon a local trigger, a message is sent to the gateway over bus 2 (uplink) and the gateway sends a trigger command to all other Pods over bus1 (downlink), and to other gateways over RF or wired network. Data is stored directly to the local SRAM along with time stamp. At the end of the “event,” the gateway can upload the signals from each Pod in an orderly manner over bus 2, and log to the database. Finally, the gateway takes care of arbitration of the token for the Pod communication.

Based on the system architecture, we implement a microcontroller-based real-time embedded system. We pay special attention on four design issues:

1. Time Synchronization

Time security is the prerequisite to get accurately synchronized sample data for the system. We use master-slave clock mode to synchronize the timing. As Figure 4 shows, the base station owns the master clock, with one additional data line separate from communication line connected to the clocks of the slaves nodes (pods). All the pods share the master clock. The master send time tick message at a certain interval, The clock on each pod will update their time with the real time tick from base station in order to avoid the timer drifts.




Figure 4. Time Synchronization

All pods will be interrupted to response the tick information in order to be synchronized.

2. Real-Time Constraint

One of most challenging point for the system is the high-speed sampling rate which is up to 250 samples/s for six to nine data channels. Beside that, our pods also need to meet the communication and data processing demands. These means that in every 4 ms interval, the pods and base station must finish data collection, processing, and two way data exchanges.

In order to meet the above design demands, we designed a custom scheduler, similar to a multitasking operating system, to manage internal (Pod) and external (system) behavior in an orderly manner without an embedded real-time operating

system support. The cooperative scheduler is a shared-clock scheduler, which means pods and base station share the same clock resource and every 4 ms, the base station send a tick with the data message to begin the cycle of scheduler, the pods then reset its own timer for 2 ms period, for the first 2 ms, the pods collect data, for the second 2 ms, they will execute the corresponding command in the message and communicate with base station.




Figure 5. System Time Pattern

The scheduler imposes a real-time requirement to both pods and base station, in this hard real-time system, since they need execute all the processing during half the 4ms interval, and must finish the execution before the next tick comes, otherwise the scheduler pattern will be destroyed. Careful software performance evaluation must be executed in order that the time constraints be met.

3. Asynchronized Communication

The pattern of real-time response of the system must be kept in order to keep the synchronization for all pods. However, the Linux box interacts with base station through two independent channels. In order to insure that the time distribution of base station not be disrupted by any request from data request, we disable the interrupt on the serial communication channel so that it can operate asynchronously.

On the other hand, when the pods transport the data to the base station, the Linux machine should not interrupt the base station, so buffers must be used in base station to hold commands until a free cycle is reached. The usage of buffers is described in Figure 6.




Figure 6. Buffer in Base Station

All the commands from the Linux engine will be loaded into the buffer, so the microcontroller can fetch the command when it send the time tick. When the pods send the data, it will save into buffer, so that when the scheduler come to the second half of the cycle, which is the time for communicating with Linux box, it will fetch the data.

4. Time-Division Arbitration

All pods must get a token in order to communicate through the shared bus. The base station, on the other hand, is the arbitrator of the token, so we let base station assign the token without immediately responding to pod's acknowledgement.

It will first send the token to pod 1, then wait the pod's response. The pod receiving the token will commence uploading the data, the rest of pods keep silent. The base station will arbitrate the tokens in the order of pod number. Once the base station changes the token to the next pod, all other pods notice and update their communication status. With the time division mechanism, we can implement a fair share.

Figure 7 shows how the pods respond to notification. If the red dot shown in figure represents the event occurrence, we see that all the pods detect the event at different time. Because the seismic wave moves from the earth center to the ground, the pod at the bottom detects first, and there is some delay before all the upper pods detect. Since all the pods




Figure 7. Time Division Arbitrations

have an allocated time slot to talk to the base station, the one who detects and get the token earliest will send the notification to the base station and will be received. The local gateway will ignore all other notification. In this figure, the notification which the arrow sends will be received.

CONCLUSIONS

The Terra-Scope™ system is an affordable 4-D down-hole seismic monitoring system based on independent, microprocessor-controlled sensor Pods. Each pod measures 3-D acceleration, tilt, azimuth, temperature, and other parametric variables such as pore water pressure and pH and etc. A local gateway serves as a controller center for the communication between multiple pods, it is also a data storage and gateway to the internet. The down hole array monitoring requires a hard real time property for the embedded system. A microcontroller-based real time embedded system is carefully designed in order to meet the real time constraints. The prototype of the device is currently undergoing evaluation, and further experiments are expected. The first array will be installed in the fall of 2005.

REFERENCES

1. Acar, C., and Shkel, A.M. (2004). Structural design and experimental characterization of torsional micromachined gyroscopes with non-resonant drive mode, *Journal of Micromechanics & Microengineering*, **14:1**, 15-25.
2. Applied MEMS, (2002). Si-Flex Force-Balance Servo Accelerometers, Model SF1500-UNLD

3. Archuleta, R.J., Seale, S.H., Sangas, P.V., Baker, L.M., and Swain, S.T.,(1992). Garner Valley Downhole Array of Accelerometers: instrumentation and preliminary data analysis. *Bulletin of the Seismological Society of America*, **82:4**, 1592-1621.
4. de Alba, P., and Faris, J.R. (1998). Treasure Island, California, Deep Instrumentation Array. *Proceedings, Second International Symposium on the Effects of Surface Geology on Seismic Motion*, vol. 1, 201 - 208.
5. Baise, L.G., and Glaser, S.D. (2000). Repeatability of Site Response Estimates Made Using System Identification, *Bulletin of the Seismological Society of America*, **90:4**
6. Baise, L.G., Glaser, S.D., and Sugano, T. (2001). Consistency of Dynamic Site Response at Port Island, *Earthquake Engineering and Structural Dynamics*, **30:6**, 803-818.
7. Elgamal, A.-W., Zeghal, M., Parra, E., Gunturi, R., Tang, H.T., and Stepp, J.C. (1996). Identification and Modeling of Earthquake Ground Response-I. Site Amplification. *Soil Dynamics And EarthquakeEngineering*, **15**, 499-522.
8. Katayama, T., Yamazaki, F., Nagata, S., and Lu, L. (1990). *Development of Strong Motion Database for the Chiba Seismometer Array*. Earthquake Disaster Mitigation Engineering, Inst. of Ind. Science, Univ. of Tokyo. Report No. 90-1 (14)
9. Steidl, J., Tumarkin, A.G., and Archuleta, R. (1996). What is a Reference Site?, *Bull. Seism. Soc. America*, **86**, 1733-1748.
10. Stidham, C. M. Antolik, A., Dreger, D., Larsen, S., and Romanowicz, B. (1999). Three-Dimensional Structure Influences on the Strong-Motion Wavefield of the Loma Prieta Earthquake. *Bull. Seism. Soc. America*, **89:5**, 1184-1202.
11. Tang, H.T., Tang, Y.K., Stepp, J.C., Wall, I.B., Lin, E., Cheng, S.C., Lee, S.K., and Hsiau, H.M. (1989). EPRI/TPC Large-Scale Seismic Experiment at Lotung, Taiwan. *Proceedings: EPRI/NRC/TPC workshop on seismic soil-structure interaction analysis techniques using data from Lotung, Taiwan, Report NP-6154*, Palo Alto: EPRI.