

Lessons Learned Building TeamTrak: An Urban/Outdoor Mobile Testbed

Jeffrey Hemmes, Douglas Thain, Christian Poellabauer,
Christopher Moretti, Phil Snowberger, and Brendan McNutt
Department of Computer Science and Engineering
University of Notre Dame

{jhemmes, dthain, cpoellab, cmoretti, psnowber, bmcnutt}@nd.edu

Abstract

Much research in mobile networks relies on the use of simulations for evaluation purposes. While a number of powerful simulation tools have been developed for this purpose, only recently has the need for physical implementations of mobile systems and applications been widely accepted in the literature. In recognition of this need, and to further our research objectives in the area of wireless sensor networks and mobile cooperative systems, we have built the TeamTrak mobile testbed, which gives us real-world experience with research concepts as we develop them. Additionally, results from outdoor field tests are used to further enhance the capabilities of the testbed itself.

1 Introduction

The study of mobile networks has become very popular over the last several years, with an abundance of published work in the area of routing protocols, energy conservation, peer-to-peer searching algorithms, and other related areas. Most of this research relies on the use of modeling and simulation for evaluation. While simulations are valuable for demonstrating proof of concept and testing properties such as scalability, real-world tests often expose mistaken assumptions. To facilitate evaluation of wireless networks, we are developing the TeamTrak mobile testbed. Users carry TeamTrak devices enabled with various sensors and collect data which is then used to improve system design and implementation. While currently used for several ongoing research efforts at the University of Notre Dame, TeamTrak continues to evolve, and empirical data directly contributes to advancing the platform's capabilities.

TeamTrak consists of a heterogeneous collection of commodity laptop computers and PDAs connected over a wireless ad hoc network and able to receive positioning data through portable GPS receivers. Sharing of sensor data,

to include GPS position, is a fundamental part of the underlying network protocol, so every connected node in the system can determine the location of any other by examining routing information, and allows cooperating users to share data and network services without the benefit of fixed infrastructure.


2 TeamTrak Components

TeamTrak is a standalone application designed to build and run on multiple platforms. Our goal is to evaluate research ideas in mobile distributed computing without specialized or custom fabricated hardware, so we use inexpensive commodity equipment as much as possible. Our research prototype consists of 32 Lenovo X41 Thinkpad tablet computers running Windows XP and eight HP iPAQ hx2795b PDAs running Windows Mobile. A standard ANSI Z89.1 Class C safety helmet provides a convenient platform for mounting mobile sensor equipment.

The communication medium is wireless ethernet (IEEE 802.11b) in ad hoc mode (no base station). The standard Windows IP configuration is used: each node detects a network, then negotiates a link local RFC 3927 IP address. Although wireless ethernet is not optimized for outdoor peer-to-peer communication, it is supported by standard consumer electronics, and thus needs no specialized hardware.

Data sharing in TeamTrak is accomplished through the use of a simple distance-vector routing protocol similar to RIP [9]. At 1-second intervals, each node broadcasts the contents of its routing table, including sensor data for each entry, to all other connected nodes. The selection of a proactive routing protocol is primarily for simplicity and is not without some consequences of which we were aware when building the prototype. We also do not employ the multiple packet types found in DSDV [11]. Stabilizing the routing table initially requires some time as the data must propagate through the network using a flood-like mechanism. This further implies that clearing stale data involves delays as well, which leads to the well-known counting to infinity

¹This research is partially supported by DoD Defense University Research Instrumentation (DURIP) Grant Number W911NF-06-1-0120



32 tablets:
Lenovo X41

On Helmet:
Garmin GPS-18
Watchport Camera
PNI V2Xe Compass

On Foot:
SparkFun v5
Accelerometer



8 handhelds:
HP iPAQ
hx2795b

In Pocket:
HP Bluetooth
Nav System

Figure 1. The TeamTrak Hardware Testbed

TeamTrak consists of 32 fully-instrumented tablet kits and 8 lightweight kits. Each tablet kit consists of a Lenovo tablet connected to a helmet-mounted GPS, camera, and compass, and a foot-mounted accelerometer. Each lightweight kit consists of an HP PDA and Bluetooth GPS. Both run nearly identical software and interoperate via wireless Ethernet and Bluetooth.

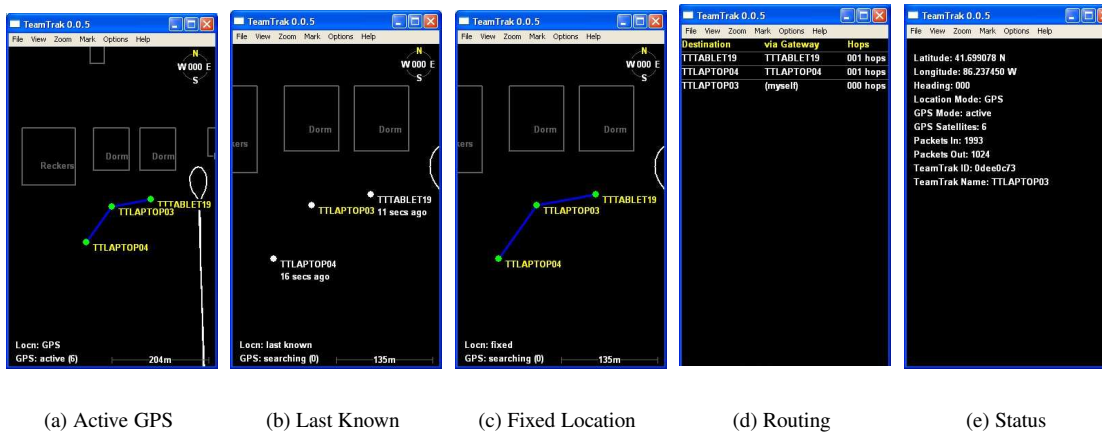


Figure 2. TeamTrak Display Modes

TeamTrak has several display modes that present state information about network. The scope display shows the current GPS location of the device on a map, its physical relation to other connected nodes, and the links between nodes. Figures 2(a), 2(b), and 2(c) show the scope display with an active GPS signal, the last-known location of disconnected nodes, and physical locations set manually, respectively. Figure 2(d) shows the routing table, and 2(e) the current status of the local device.

problem [7]. Our approach to handling stale data is to make it persistent until cleared from memory by the operator.

Augmenting the basic hardware platform are sensors connected by USB or serial port. GPS data are provided through Garmin GPS-18 USB GPS receivers, or in the case of the PDAs, an HP Bluetooth GPS receiver. In addition to GPS, the platform includes the PNI V2Xe digital 2-axis compass, the Watchport/V2 digital camera, and SparkFun SerAccel v5 digital accelerometer. As we continue to gain experience with sensor devices, we add them to the overall system architecture and incorporate the sensor data of interest into the communication protocol. Figure 1 shows the devices that comprise each TeamTrak node.

TeamTrak has several location modes, depending on availability of live GPS data or the specific application. If

a GPS receiver is connected and receiving a live signal, the display indicates such in the lower left corner, as shown in Figure 2(a). If the GPS signal is lost, the display shows the node at its most recent known location, as shown in Figure 2(b). Similarly, as nodes move and become disconnected, the display will continue to show their last known location, but changes the symbol used to represent each and also indicates the length of time since a packet was last received from each. Where no GPS signal can be obtained, a node may approximate its own location by averaging live GPS information from its neighbors within a single hop. Location may also be set manually for testing purposes or correcting GPS error in cases where the exact location of a node is known. In this case, the display indicates the location is fixed, as shown in Figure 2(c). Additionally, the dis-




Figure 3. Challenges of Sensors: Obstructed GPS Measurements

These graphs show the distribution of positions computed by a GPS device once per second over three days. The device had a clear view of the southern sky, but the northern sky was blocked by a building. Note that the device computes several “excursions” that deviate from the average by several hundred meters.

play shows compass heading for each node and is capable of displaying live video imagery from the digital camera.

3 Experience with Sensor Measurements

When we first embarked on TeamTrak, we quickly discovered that each sensor presented its own set of distinct problems in measurement and interpretation. In this section, we describe some of the practical difficulties of using commodity sensors.

Global Positioning System. While initially testing the GPS device by simply walking around campus, we observed several anecdotal problems. First, the GPS is very sensitive to its placement relative to the body. A signal could not be obtained if the device was placed in a pocket or even held close to the laptop or handheld. Only by placing the device on the shoulder or on top of the head could a signal be obtained. Second, because nearby buildings interfere with signal reception despite having a clear view of most of the sky, an obstruction of part of the horizon would cause delays of 5-10 minutes in obtaining a GPS fix. Finally, while moving around campus with a GPS fix, the device would occasionally “jump” to locations as much as a kilometer away, and wander in that region for seconds or even minutes before returning to the proper location. Each of these problems occurred with both the Garmin GPS-18 and the HP iPAQ Navigation System designed for an automobile windshield.

To get a more quantifiable understanding of GPS variance under non-ideal conditions, we recorded the behavior of a GPS-18 unit once per second over three days. The unit was placed in the window of a building with a clear view of the southern sky, but, completely blocked from the northern sky; a view that would be common in an urban setting.

The results of this experiment are shown in Figure 3. The left graph shows a scatter plot of every position mea-

surement taken, centered on the mean (the absolute accuracy of the average measurement is unknown). The most striking feature of this plot is that the data are by no means distributed evenly. Large sequences of measurements drift across the average position, mostly in the north-south direction. In extreme cases, measurements are as much as 500m off of the average. Inspection of the data shows that these are not individual exceptions, but rather the measurement drifts to an extreme value, then drifts back to the average.

The center graph in Figure 3 shows the cumulative probability of a measurement’s distance from the mean (the right graph shows the same data at extreme ranges). As can be seen, a wide range of measurements are common. About 35% of measurements are within 10 meters, 90% within 50m, 98% within 100m, 99% within 150m, and 99.5% within 200m. Although most measurements are reasonably accurate, the fraction of seriously diverging measurements is large enough that it cannot be ignored. In a sufficiently large network of devices in non-ideal conditions, we must assume that perhaps 1-2% report an inaccurate GPS fix.

Accelerometer. Regardless of the accuracy or price of a GPS, it cannot be used in locations with no view of the sky, such as very dense urban areas or indoors. Previous work has suggested the use of accelerometers mounted on the surface of mobile robots, with the output integrated to give current velocity and change in position, which is then used for fine location. This works acceptably on a robot [8, 10], which has relatively consistent acceleration patterns. Initial testing of the accelerometer in this fashion for humans was not encouraging. Humans have no solid platform on which to mount the device, leaving an uncertain axis of motion. In addition, humans have “noisy” motions that lead to extraordinary integration errors. Even in controlled situations, computing position from the accelerometer for more than a few seconds results in errors on the order of kilometers.

However, the accelerometer can be used to detect spe-

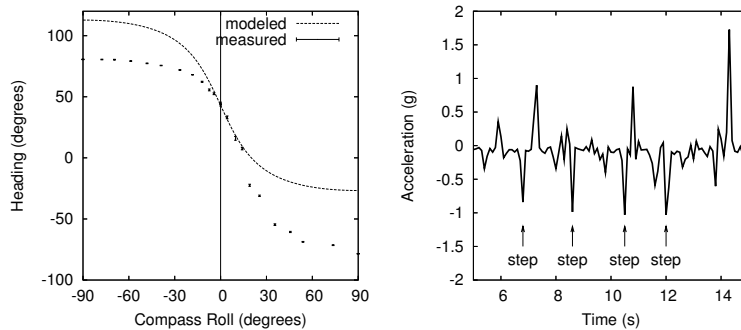


Figure 4. Challenges of Sensors: Compass and Accelerometer

The left graph shows the variation in compass reading as the compass is tilted through a range of angles without changing the heading. The right graph shows the reading from an accelerometer affixed to the shoe of a person walking.

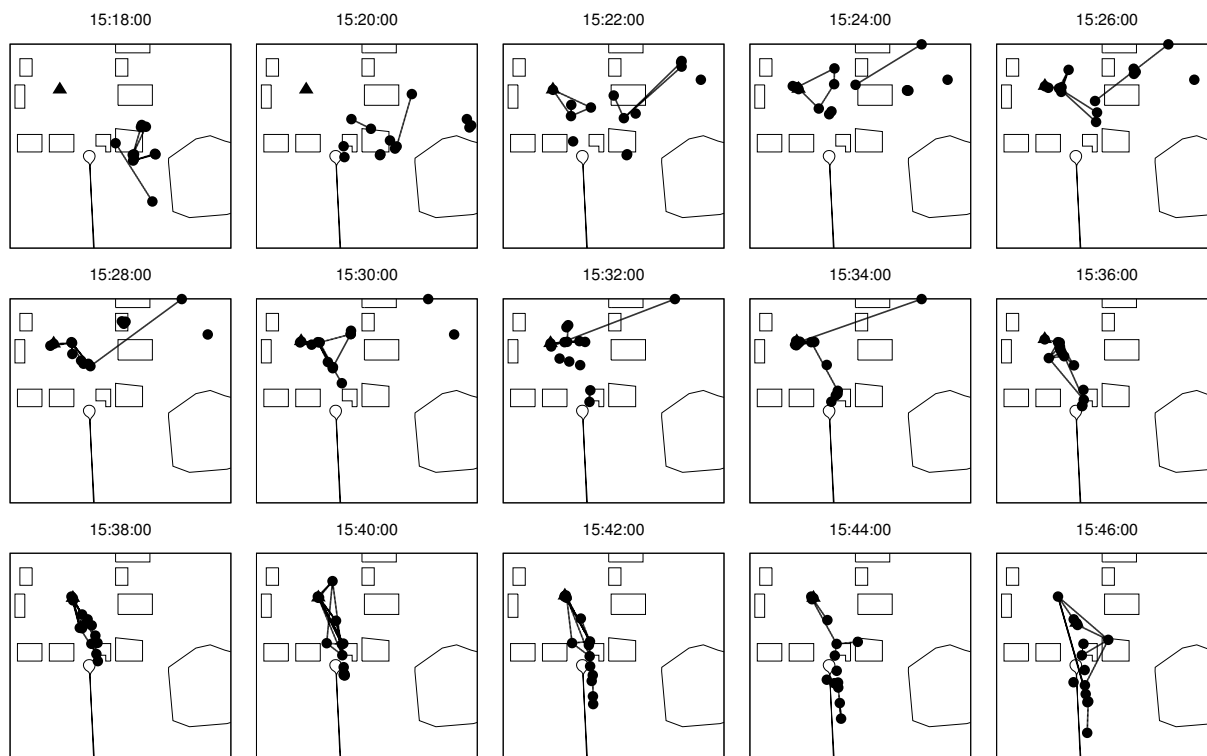


Figure 5. Time Lapse Display of a TeamTrak Network

A time-lapse display of a field test by an undergraduate class. Each box shows the system state at 2-minute intervals. Dots indicate mobile nodes. Heavy lines indicate nodes that are in direct contact via wireless ethernet. Lighter lines indicate campus landmarks. The scale of each box is about 1km on each side. Initially, all nodes are searching for the “rabbit” indicated by a triangle in the upper left corner. By 15:32:00, all have found the rabbit, either directly or indirectly. At this point, they attempt to form the longest unbroken chain possible, extending south.

cific motions of the body. Attached to a person's shoelaces it can be used to measure individual steps. Coupled with a compass (next section), we are optimistic that accelerometers can be used to provide location updates when GPS is not available. This will be explored in future work.

Digital Compass. While the GPS device provides position data, it does not determine the direction a unit is facing. Using a digital compass, we may determine direction, and use this along with the accelerometer to compute position changes. In addition, direction data can be collected and used to infer areas of focus, annotate camera images, or inform wireless coverage areas. Unlike the GPS signal, magnetic heading is detectable in nearly any indoor or outdoor situation.

When held in a stable flat position at various angles, the V2Xe compass was found to have a measurement error of about one degree. The difficulty comes in the measurement stability with respect to rolling motions that do not affect heading. A two-axis compass will register an error as it is rolled. Figure 4 shows the magnitude of this roll when held at a heading of 45 degrees. Rolling motion must be kept to 10 degrees or less, otherwise the error is considerable.

4 Experience with Field Tests

Based on common positive experience with wireless access using infrastructure networks, we were optimistic about the possibility of ad hoc networks. In initial tests with two participants, we achieved ranges of about 100m between hosts. However, these ranges were not consistent: sometimes a nodes could communicate at a distance around the corner of a building, while sometimes two nodes standing together could not communicate.

We conducted a series of field tests with an undergraduate distributed systems class, described in [6]. Note that our goal in these tests is **not** a systematic study of any one component, but rather to gain experience with the system, uncover practical implementation issues, and suggest problems for further study.

Figure 5 shows a time-lapse display of a recent test. In the figure, each node is shown as a dot, the target as a triangle, and active network connections as heavy lines between nodes. Lighter lines indicate buildings and roads. We asked 12 participants to perform a simple activity: find a "rabbit" node hidden on campus, and communicate its geographic location to the rest of the group via the routing protocol. As each participant found the rabbit, they were instructed to form the longest unbroken chain possible, extending south.

- **15:18:00** - The searchers were dispatched from the southeast corner of the map a few minutes earlier. Although in a relatively compact group, the network has already partitioned. Some proximate nodes are partitioned, while others farther apart can communicate.

- **15:22:00** - The network briefly becomes whole, then partitioned again. The western group has discovered the rabbit, while the eastern group continues northeast.
- **15:28:00** - Most nodes have discovered the rabbit, some indirectly. Note that one node has a very long network connection; in this case, the person in question was actually close to the rabbit, but the GPS was reporting an incorrect position until 15:36:00.
- **15:34:00** - All nodes have begun forming a long chain.
- **15:38:00** - The network is very strongly connected.
- **15:44:00** - The network has stretched into a chain. Various links in the chain connect and disconnect. Most people are standing in place and turning in either direction to look at their peers. Due to asymmetric network coverage, turning causes breaks in the network.

From this exercise, we may draw several observations:

- Despite the practical problems of wireless ethernet observed above, it can be used for a complex network of a fairly large geographic scale. The final network state at 15:46:00 is about 750m long.
- The relationship between wireless connectivity and distance is nontrivial. Connections can sometimes be made between widely separated nodes, and sometimes cannot between adjacent nodes, which poses problems for routing or location algorithms that assume a relationship between connection and location.
- Network partitions are even more common than we expected. Obviously, the network will partition when a group splits to undertake a search. However, even when standing in a line with the explicit goal of staying connected, we observed rapid partitions and reconnections occurring in several different ways. Any network protocol should be highly responsive to changes.

5 Lessons Learned and Future Work

Based on our experiences building, testing, and deploying TeamTrak, we identify and discuss three areas requiring further study and development of novel protocols and techniques that go beyond the state of the art.

GPS Needs Assistance in Urban Environments at Human Scale. Positioning with commodity GPS receivers alone is probably inadequate for most urban applications due to the limitations of the environment and the inherent capability of commercial receivers. Our test location is a university campus with buildings spaced much farther apart than one finds in dense urban settings, and we found numerous

places in which obtaining a GPS fix was practically impossible. Attempting to obtain a fix in specific areas at different times and days produced similarly poor results. Once a fix was obtained, our testers experienced no significant difficulty maintaining it, generally only losing signal due to movement around obstructions or other ephemeral factors. However, initially acquiring a fix may prove unacceptably lengthy for many applications.

Location awareness is exploited to improve application performance, reliability, or usability [12, 4, 5] in wireless and mobile systems, and GPS is commonly used to identify a mobile device's location. The experiments executed with the TeamTrak testbed indicate that GPS information is both unreliable and inaccurate (see Figure 3 in Section 3), which has been addressed in previous work on robust location and navigation systems [1, 3, 2]. However, we believe that existing location-aware devices can benefit from not only knowing current position but also direction and speed of themselves and other devices. For example, in military or rescue operations, users might enter *critical zones*, such as areas of increased risk to humans and/or mobile devices. Here, it may be desirable to *offload* important information from a mobile device to (a) protect information from capture and (b) prevent loss in case of device failure.

Human Operators Play an Important Role. Human operators can partially compensate for poor network connectivity for some applications. In our rabbit chase exercise, users tended to prefer maintaining connectivity with other users, and they frequently would move towards other nodes to reestablish lost connections. Behavior such as this can be exploited in mobile applications to augment network protocols, although much work remains to be done.

Sensors Can Compensate for Poor Network Connectivity. Consistency of ad hoc network connectivity proved to be much worse than originally anticipated. Using standard wireless ethernet, connectivity between nodes is at best haphazard. Despite the network limitations, displaying the last known locations of nodes proved helpful when trying to locate a particular node. Even a very transient connection in which few routing packets are exchanged is sufficient for determining locations. Since each node propagates information for all nodes with which it has connected, users working in concert can quickly converge on a particular location. How long stale data should be maintained is unclear, but users may remove it if the display becomes cluttered.

To account for variations in connectivity and signal strengths, we need wireless communication and routing protocols that are aware of the limitations of reduced coverage. Predictors, such as the current position and bearing of users of mobile devices, can be used to provide coverage maps on the fly, which in turn can adapt communications to account for predicted increases or decreases in connectivity. Additional sensor data can be used to achieve this.

6 Conclusions

The TeamTrak mobile testbed is our first step in building cooperative mobile systems and researching the use of various sensor data in such an environment. During the course of system development, we have made a number of observations about the use of off-the-shelf sensor devices, software, and hardware platforms that both pose significant challenges yet offer exciting possibilities for future expansion and integration. In many cases, empirical evaluation of components has changed our assumptions about the capabilities of such hardware and the feasibility of capturing and sharing meaningful sensor data in a mobile environment, and led us to identify open research problems and areas for future direction.

References

- [1] J. Buriat, T. Fraichard, and O. Aycard. Robust Navigation using Markov Models. In *IEEE-RSJ International Conference on Intelligent Robots and Systems*, August 2005.
- [2] J. Castro, V. Santos, and M. I. Ribeiro. A Multi-Loop Robust Navigation Architecture for Mobile Robots. In *IEEE International Conference on Robotics and Automation*, May 1998.
- [3] C. Gentile and L. Klein-Berndt. Robust Location Using System Dynamics and Motion Constraints. In *IEEE International Conference on Communications*, June 2004.
- [4] A. Haghighat, C. Lopes, T. Givargis, and A. Mandal. Location-Aware Web System. In *Workshop on Building Software for Pervasive Computing at the Object-Oriented Programming, Systems, Languages and Applications*, October 2004.
- [5] T. He, J. Stankovic, C. Lu, and T. Abdelzaher. SPEED: A Stateless Protocol for Real-Time Communication in Sensor Networks. In *International Conference on Distributed Computing Systems (ICDCS)*, 2003.
- [6] J. Hemmes, D. Thain, and C. Poellabauer. Work in Progress: Integrating Undergraduate Research and Education via the TeamTrak Mobile Computing Framework. In *Frontiers in Education Conference*, October 2007.
- [7] A. Leon-Garcia and I. Widjaja. *Communication Networks: Fundamental Concepts and Key Architectures*. McGraw-Hill Higher Education, 2000.
- [8] H. Liu and G. Pang. Accelerometer for Mobile Robot Positioning. In *IEEE Industry Applications Conference*, July 1999.
- [9] G. Malkin. RIP Version 2. IETF RFC 2453, November 1998.
- [10] G. Pang and H. Liu. Evaluation of a Low-cost MEMS Accelerometer for Distance Measurement. *Journal of Intelligent and Robotic Systems*, 30(3):249–265, March 2001.
- [11] C. Perkins and P. Bhagwat. Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers. In *ACM SIGCOMM '94 Conference on Communications Architectures, Protocols and Applications*, August 1994.
- [12] P. Ramanathan. Location-Centric Approach for Collaborative Target Detection, Classification, and Tracking. In *IEEE CAS Workshop on Wireless Communication and Networking (Invited)*, September 2002.