Using Server Architecture and Multi-Threaded Processors and Software to Time-Lock Multiple Data Streams in Time-Critical Physiological Experiments

Christopher M. Storey[1, 2] and Charles J. Robinson[1-4]

[1]Center for Rehabilitation, Engineering, Science and Technology, Clarkson University, Potsdam, NY / [2]Dept. of Bioengineering, Louisiana Tech Univ., Ruston, LA / [3]Dept. of Electrical and Computer Engineering, Clarkson University, Potsdam, NY / [4]Research Service, Syracuse VA Medical Center, Syracuse, NY.

Abstract

To ferret our subtle clues in the data that we collect, we need highly stable and reliable instrumentation, and a way to link disparate data streams together. In a typical laboratory such as ours, collecting a variety of physiological data in time-synced fashion has in the past required a number of desktop computers, each one handling a different aspect of data collection. This wealth of data required us to do laborious off-line file conversion and synchronizing of three main data streams. Since new technologies such as hyper-threading and dual-core technologies have brought enormous power to desktop systems, we opted to use a server-style system to be able to multitask in a network environment, but still maintain Windows XP as the operating system. The server platform provides us the means necessary to combine our various data collection schemes into a single unit, and for on-line data calibration and conversion, while still allowing us an easy transition from our previous hardware and software.

Introduction

A wealth of multi-dimensional data can now be collected during biomechanical studies of human motion and postural reactions to perturbation. These include biomechanical measures like AP and ML Centers-of-Pressure (COP), weight on platform versus weight supported by harness, horizontal ground reaction forces, head and foot accelerations in multiple dimensions, distributions of pressures under the foot, and joint and limb trajectories as measured by motion-capture marker systems. Multi-channel EMG data and psychophysical responses collected simultaneously add richness to any control model built.

Collecting all of this data in time-synced fashion in the past has required using a number of computers, each one handling a different aspect of data collection, where one computer serves as the master and triggers the other slave computers to begin data collection. Yet many of the vendors of propriety data collection interfaces (like a pressure mat or a motion capture system) supply there own proprietary software to act as a master, with other data slaved in. For large-scale observations, this feature results in a number of datasets collected over the same epoch, but does not necessarily create datasets that are time synchronized with one another. Using each device as a master, with data from other devices imported into each setup's data collection, the already large size of the data for a single experiment rapidly grows due to the storage of duplicate data in multiple systems. However, the datasets are not truly duplicate, as each was processed a different way, with possible variations in amplitude and noise and certainly offsets in timing.

Parallel processing requires systems capable of running thread simultaneously. While parallel processing required multiprocessor systems, new technologies such as hyper-threading and dual-core technologies have brought this power to many desktop and laptop systems. These technologies provide the ability to increase multiprocessor systems exponentially, which is the basis of many new server platforms on the market.

Time-Series Data

Like most research laboratories, our lab uses computers to control our experiments, collect data, analyze the results, and write up any resultant manuscripts. Our studies investigate the psychophysics of balance and postural control.

All people sway. We make very short perturbations of a platform on which a subjects stands that are generally of a length less than that a person's sway, as measured at the foot and ankle, and hence should be and are near the ability of a subject to detect them. Using psychophysical techniques, we apply peri-threshold perturbations to iteratively find the detection limit, and attempt to figure out what physiological or biomechanical variable(s) was (were) used by the subjects in making a correct detection, or led them to a false detection. Details of our experimental set-up and procedure have been described adequately elsewhere. Here we note that our Sliding Platform For Assessing Lower Limb Stability with Synced Tracking, EMG and Pressure measurement (SLIP-FALLS-STEPm) is a vibration-free, translating platform that rides on air bearings[1]. Our typical protocol has a subject picking in which of two 4s intervals that the perturbation occurred, with data collection occurring in thirty 15s windows. This protocol is repeated 2 to 4 times.

To ferret our subtle clues in the data that we collect, we need highly stable and reliable instrumentation, and a way to link disparate data streams together. During the build-up years in our lab, we grew the experiment by adding hardware as a system. For instance, we started with a single computer for experimental control and data collection. We added another data collection computer when we introduced the measurement of the distribution of pressure under the foot from an HR TekMat. We also later added a motion capture system that used retro-reflective markers to trace the movement of the joints and other body reference points that could be caused by the perturbation. This system also had its own computer. We used digital out signals from the controlling computer to trigger data collection routines in the other two computers. Fig. 1 details the setup and the outputs.

In the past, this wealth of data required us to do laborious off-line file conversion and synchronizing of the three main data streams. We wished that all of this conversion could occur simultaneously during data collection, but the computers in our original implementation lacked the processing power to process data efficiently without testing delays. The outdated technology could not take advantage of multithreading to parallel process.

Specifications for a New Lab

With a move to another university, we have had the chance to set up a second SLIP-FALLS-STEPm research lab for fundamental studies, while maintaining the original lab in a clinical setting within the VA research service. Based on >10 yrs experience with the original SLIP-FALLS lab, we set down a series of specifications for the new lab:

1. The essential elements of the user interface needed to remain the same as seen from the clinical environment.
2. The command and control aspects of the platform had to be functionally equivalent to previous implementations, and previous code had to be reused when possible.
3. The operator should be provided a user-friendly interface with which to monitor the progress and output of all these processes in real-time during a testing sequence.
4. The number of channels of data collected by the FALLS protocol must be increased to allow for additional sensor and EMG inputs. The EMG channels were to be increased from the original four to a user-selectable between four and sixteen. The amount of support the safety harness provides to the subject should also be collected and calculated.
5. The motion analysis system should be upgraded from a single camera, 2-D system to a multi-camera, 3-D system.
6. The FALLS data should be immediately stored in engineering units, rather than in raw voltages that required post-possessing. EMG potentials should be converted on-line and

stored as RMS time-series data, with a further conversion to a percentage of that seen under maximal contraction if possible.

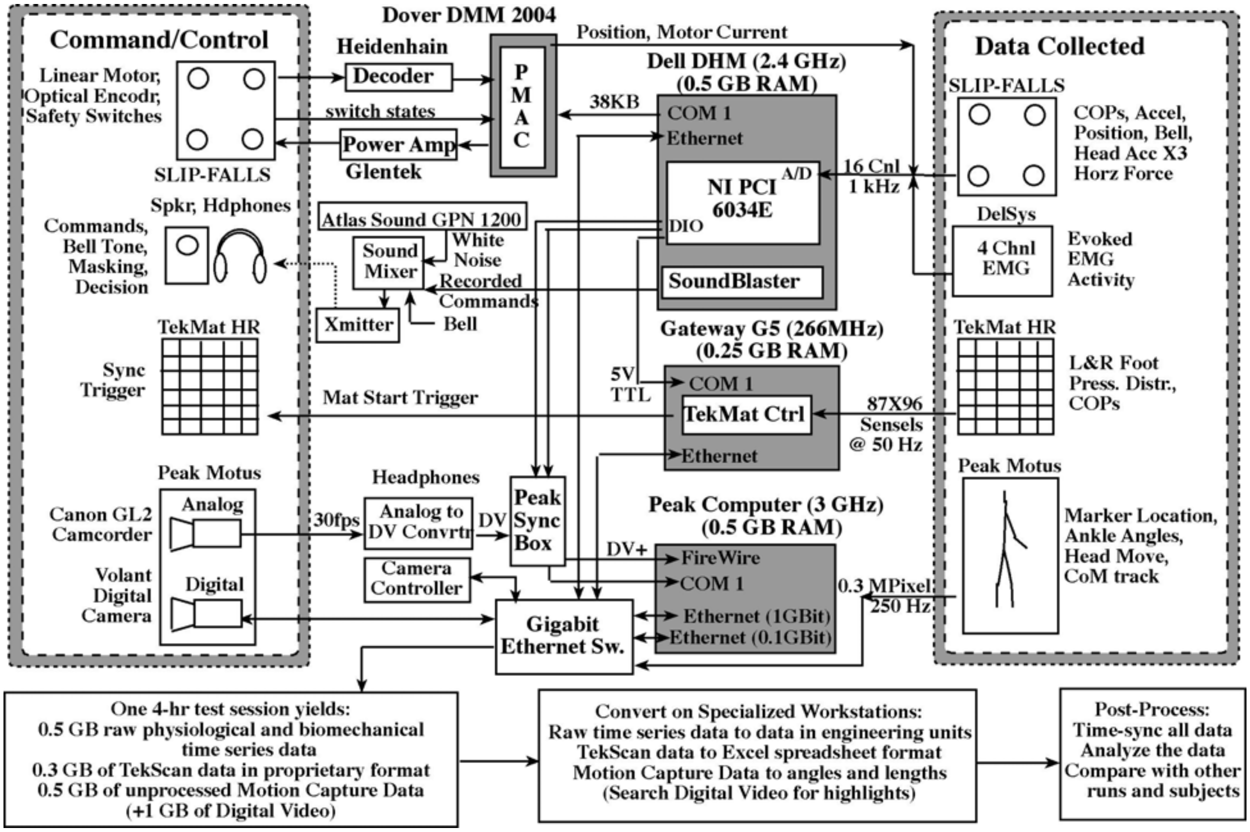

Figure 1: The original SLIP-FALLS-STEPm lab consists of a sliding platform (SLIP), hardware and software routines (FALLS) for collecting neurophysiological, biomechanical (Position, Acceleration, Centers-of-Pressure), EMG and psychophysical data, and equipment and software to simultaneously measure position markers, and foot pressure distributions (STEPm). The SLIP is controlled by a single board Programmable Multi-Axis Controller (PMAC) that receives commands from the FALLS computer (Dell DHM). This computer also relays pre-recording instructions to the subject via a SoundBlaster card and headphones. A Tekscan HRMat (TekMat) measures foot pressure distribution with an array of 87 by 96 sensels (4 per cm^2). A separate computer controls the TekMat. Its data collects starts with an external RS232 trigger. Motion capture of the location of retro-reflective markers is achieved by a single digital camera Peak-Motus system, along with a back-up analog camcorder. The peak computer is triggered also by the FALLS computer with an additional sync signal generated at the start of a platform move. A single 4-hr test session generates over 2 GB of data, all of which has to be processed offline after the completion of the experiment.

Solution

As our current setup was incapable of performing the required computations without testing delays, we focused on what was needed to meet these objectives. We needed a system that would remain under the Windows XP operating system to maintain current software and equipment drivers, which satisfied the first through third specifications. The remaining specifications require additional new or replacement equipment.

To meet the third specification we had to upgrade our multifunction data acquisition card (NI PCI 6034E) from a 16 analog inputs to 32 analog inputs (NI PCIe 6259M). We also achieved additional input be routing subject response (bell) to the digital inputs instead of counting peaks of analog input. To provide signal conditioning and signal access we used NI SC-2345 and NI BNC 2090 for I/O. The NI BNC has a dual functionality of allowing us access to EMG signals so they can be inputted into Peak Motion capture system. To acquire these signal in the Peak system and to meet the fifth specification the hardware had to be upgraded. Since the purchase of our previous system Peak-Motus was acquired by VICON, which allowed us to upgrade to VICON’s superior cameras and hardware while maintaining same user interface with updated Peak software.

For the final specification, the data had to be converted on the fly to engineering units without causing any testing delays. A computer was needed that could parallel process threads, and not only utilize preemptive multitasking. New compact multi-processor server technologies were our focus for a new FALLS computer. We decided to purchase a Gateway E-9515-R series server to meet the fourth specific.

Implementation

Traditionally, for computers to be able to multitask in a network environment required a server style operating system. Due to such a small consumer base, there was very poor hardware support, which has left a mark on those who had endeavored to utilize its power in the past. With the advent of Windows XP to the general consumer, which is based off Microsoft's original server platform, the ability for user-friendly server platform was launched. Windows XP was chosen due to its widespread use and familiarity. Yet it remained stymied by its everyday use on desktops that it could handle computers that are more powerful and efficient.

Windows XP has both great hardware and software support, but if you try to purchase it with a server from major computer manufacturers, they will turn you down, or let you purchase separately with no support. They want you to purchase one of their newly branded server operating systems, which keep with the old tradition of having poor hardware and support for the everyday user and researcher. These server operating systems are expensive, which comes at the cost of paying per user license that allows for a true multi-user environment. Although this provides a limitation to the consumer, a multi-user environment would be a seldom-used feature in the lab environment. Since companies want the consumer to pay hundreds to thousands of dollars more for official server operating systems, they place limits on the software so it can only use a certain amount of the computer's resources.

Windows XP has a limit of two physical processors, but thanks to new technologies in the central processing unit (CPU), the limitation has become less stifling. Our new server class machine in Fig. 2 is composed of two Intel Xeon 2.8 GHz Dual-Core Processors. Each core also contains hyper-threading technology that is similar to dual-core but shares resources. Therefore, the software limitation imposed on us is met since we only have two physical processors. That limitation is surpassed by the fact that we have eight logical processors on which programs run.

To take advantage of extra processors software today is multithreaded, which translates into breaking up the program into smaller operations that can run independently and asynchronously. New multithreaded programs are able to push the processing envelope by distributing the work-load across all the logical CPUs. For our SLIP-FALLS-STEPm platform, we use LabVIEW® to run our experiment, record data, process data, and synch with other research systems. LabVIEW® provides a nice graphical programming interface so novice programmers can use it. It also allows for the flexibility in advanced programming for creating threaded applications, with communication streams between each, and for communicating with third party software.

The majority of our data analysis is performed in the Matlab package. With LabVIEW 8.0, you have an easy to script object to communicate and process data in Matlab. Given that we integrated and threaded our data acquisition and analysis, we have virtually eliminated offline processing time. In addition, Matlab can take advantage of Intel's Extended Memory 64-bit Technology (EM64T). The EM64T permits us to run Matlab 64-bit on our server, which also requires a 64-bit operating system (Windows XP 64-bit). With 64-bit software, the EM64T allows one to address over 4GB of memory to which will dramatically decrease the processing time by removing hard drive reads and writes due to virtual memory usage. In addition, the EM64T provides 64 bits of precision for accurate calculations. Nonetheless, we still have not reached the potential of our server. Therefore, we decided to run simultaneously the data acquisition hardware and software for Tekscan foot-pressure mat in the server.

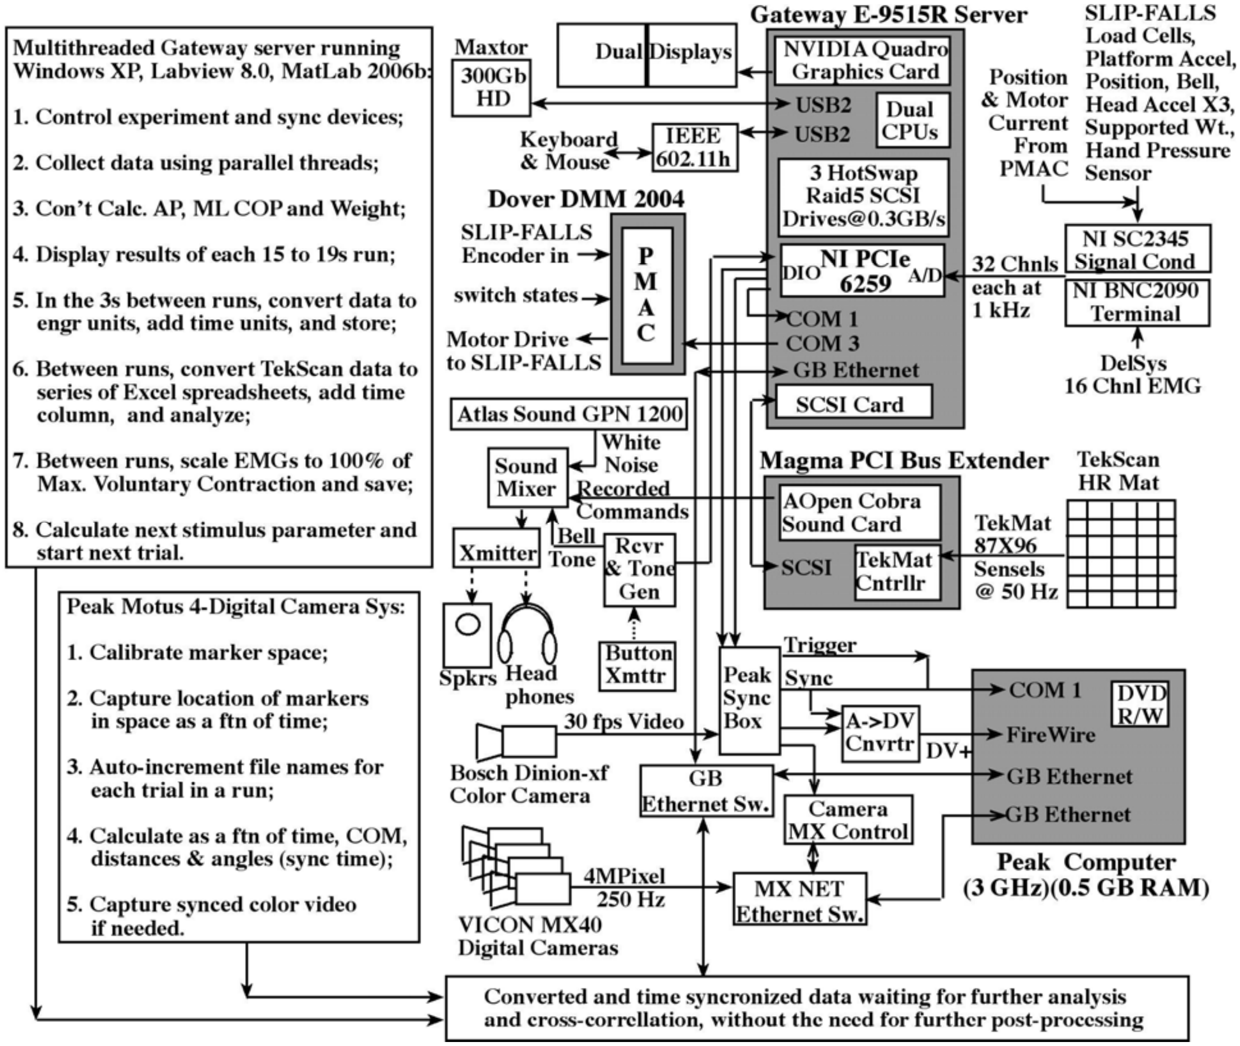

Multithreaded Gateway server running Windows XP, Labview 8.0, MatLab 2006b:

1. Control experiment and sync devices;
2. Collect data using parallel threads;
3. Con't Calc. AP, ML COP and Weight;
4. Display results of each 15 to 19s run;
5. In the 3s between runs, convert data to engr units, add time units, and store;
6. Between runs, convert TekScan data to series of Excel spreadsheets, add time column, and analyze;
7. Between runs, scale EMGs to 100% of Max. Voluntary Contraction and save;
8. Calculate next stimulus parameter and start next trial.

Peak Motus 4-Digital Camera Sys:

1. Calibrate marker space;
2. Capture location of markers in space as a ftn of time;
3. Auto-increment file names for each trial in a run;
4. Calculate as a ftn of time, COM, distances & angles (sync time);
5. Capture synced color video if needed.

Converted and time syncronized data waiting for further analysis and cross-correllation, without the need for further post-processing

Figure 2: New equipment set up. A dual core, multi-processor Gateway server running Windows XP, LabVIEW 8.0 and MatLab 2006b is the new FALLS computer. The Magma PCI bus extender attached to it via a SCSI connection allows the use of vender cards with the older PCI bus structure. The graphics card in the computer supports up to four simultaneous monitors. Three hot-swappable 200GB drives are configured as a RAID 5 set, to provide for data collection redundancy. The A/D card is expanded to 32 channels, with 16 now coming via cable from a 16 channel Delsys EMG amplifier. The Tekscan HR Mat controller PCI card no longer resides in a separate machine, so that data can now be better time- synced. The motion capture system is upgraded to a 4-camera system, each at 4 MPixel at 250 Hz, along with faster CPU. If desired, an analog video record of the test can be acquired. The output of the new system is such that no post-processing is required before correlative analyses can be carried out.

This addition brought us our first limitation on our server. The latest revision (3.0) to the PCI/PCX standard no longer contains a 5V connection. Some older cards (audio and Tekscan PCI cards) are set up for the old 5V protocol, which required us to develop a work-around. We installed a rack-mounted PCI bus extension system by Magma, which allowed up to 4 PCI devices to share a single PCI slot in the server and be backward-compatible PCI slots. The bus extension worked well with the audio card (used for subject commands) and Tekscan card (used for foot-pressure data acquisition) allowing us to incorporate both in our server configuration.

The small physical size of our server (form factor 2U of a rack enclosure) cuts down the volume of the equipment need for the testing system, which is aided by having low-profile PCI ports. We used low-profile PCI slots for a serial port (RS232) expansion card and SCSI 320 Mb/sec hot swappable RAID 5. The extra serial ports allowed us to control multiple pieces experimental hardware (Dover DMM 2004 and Tekscan HR Mat) simultaneously. Using RAID 5 for disk storage gives great data protection with only minimal loss of space as opposed to mirroring the hard drive. By striping the data across the hard drives with a parity bit, it enables the user to rebuild a hard drive's data completely if one crashes for high data security.

Server hardware was not designed for any flashy graphics cards, and there is no set high-speed graphic bus to use. However, the server has the new PCIe standard that many high-powered video cards currently use today. For our test monitoring, we chose the workstation class video card by NVIDIA because it gives is the ability to monitor all test parameters simultaneously since it has the ability to run up to four digital monitors. Also on the PCIe bus, we have our National Instruments data acquisition card. It is wired to two external hubs that allow for signal conditioning and data line access. These modules allow for synchronizing to both internal (Tekscan) and external (Vicon-Peak) software for total integration of data collection.

The Vicon-Peak system is a three-dimensional marker based camera system. The digital input and output (DIO) ports of the National Instruments DIO allow for triggering and synchronizing the video capture data to test events. Syncing is needed since there could be a delay before the cameras began recording. Syncing also allows the three-dimensional motion capture data to be aligned with the other data acquired by the server.

We have improved upon the data acquisition parameters of the original setup. Originally all signals were conditioned by separate external Daytronics signal conditioning modules with numeric displays. Now, a National Instruments SCC system is used that enables us to do individual two-stage signal conditioning on each line if needed (e.g., strain gage conditioning, followed by low pass filtering). We have used their SC breadboard modules on some signals to build our own circuitry to remove large DC offset voltages in some of our accelerometer signals. The LabVIEW driver software that comes with the SCC takes care of gain and offset calibration so that data is sent from the SCC already in calibrated engineering units (i.e., mm), rather than in raw numbers. This automated scaling and unit conversion occurring at data collection decreases the need for post-processing, and partially addresses our design criteria 6. With the SCC, we upgraded our data acquisition card from 16 analog inputs to 32, with the SCC taking the lower 16 channels, and EMG inputs the upper 16. The new 16-channel Delsys Bagnoli EMG amplifier has a 50-pin output connector that interfaces directly to a NI BNC breakout box (BNC2090) that handles the upper 16 channels, but that also allows us access to these signals, as well as providing the DIO outputs. With the upgraded EMG system, we can acquire inputs from eight bilateral muscle groups on the body. Changes will be able to be monitored not only in the muscle groups about the ankle as done now, but also the thigh, trunk, and neck muscles.

Our accelerometers (3 on the head and 1 on the platform) now have a peak-to-peak output of ±20 m/s^2. The signal conditioning provides a gain to allow for 1mV per mm/s^2 output. We maintain the original four load cells of the original force plate. We also still collect the position and motor current (shear force) from the Dover controller. Another change that we have implemented is the using the DIO for our subject acknowledgment signal. The DIO provides much better method of recognizing a subject's response than analog input with peak detectors. More offline processing is alleviated with the advent of global virtual channels in the LabVIEW software.

Discussion

Why did we choose a server over a desktop PC or workstation? First, what actually defines a server? Is it the operating system that it runs, that has "server" in the title? Is it the number of processors? Instead, is it simply any thing that is overly large and bulky or sleek and stylish that cannot be referred to as a desktop or laptop? Even people who are experience computers cringe at the word server due to the lack of support in that environment. From our standpoint, a server is a computer that maximizes the processing power per cubic inch of space it occupies, while not being singled for solitary use, and "serving" several people and purposes at once to offload the burden from desktop machines.

The server's compact size and rack-mountable design make it convenient to house and organize cables out-of-sight. Therefore, there is less clutter in the lab workspace and less confusion of equipment and function. The low profile PCI standard and PCI bus expansions also allow for the server to maximize the space it occupies without creating airflow problem that could lead to system failure.

We needed a computer that would not be outdated be new software in the near future. The extendable memory and hard drive system will provide increased capacity for years to come. The EM64T allows us to adapt to the new trend toward 64-bit computing. The EM64T permits us to perform tasks that would normally require time on supercomputers, and provides the capacity to upgrade RAM memory to 16 GB from the current 3 GB if needed for intensive analysis and simulations. The hot-swappable hard drive ensures low down time, but allows for almost plug and play expansion for up to six hard drives. This will ensure that we have plenty of capacity for subject data from future tests.

However, the main issue we had with the server was its cost. The price is reasonable given that it replaces two workstation PCs. The rack-mountable setup has allowed us to streamline our electronics in the lab and cut down on the clutter and confusion caused by the use of several PCs for testing. In addition, the new experiment protocol could not run without significant delays when implemented on a 2.4 GHz desktop PC. Therefore, even though our server cost around 3 times the price of a high-end PC, the benefits out weighed cost due to higher productivity of the lab.

Conclusion

The technologies exist to create a single system, which is capable of replacing multiple conventional PCs in current lab setups. The server platform provides the means necessary to combine these technologies into a single unit. By incorporating a server system into our experimental setup, our lab is able to spend more time analyzing usable data, writing, and developing new ideas for research, which allows for a more productive research environment.

Reference

[1] Robinson, C. J., Purucker, M. C., and Faulkner, L. W., 1998, "Design, control, and characterization of a sliding linear investigative platform for analyzing lower limb stability (SLIP-FALLS)," Rehabilitation Engineering, IEEE Transactions on [see also IEEE Trans. on Neural Systems and Rehabilitation], 6(3), pp. 334-350.

Biographical Information

CHRISTOPHER M. STOREY, MS, received his BS degree in computer science from Tulane University in 2003 and his MS in neuroscience from Tulane University in 2004. He is completing his doctorate in biomedical engineering at Louisiana Tech University. He is currently finishing his research on postural control at the Center for Rehabilitation Engineering, Science and Technology at Clarkson University's SLIP-FALLS-STEPm lab.

CHARLES J. ROBINSON, DSc, PE, received his DSc degree in Electrical Engineering from Washington University in 1979. He is the founding Director of Clarkson University's Center for Rehabilitation Engineering, Science, and Technology and holds the Schulman Chair of Electrical and Computer Engineering there. He is a Senior Rehabilitation Research Career Scientist at the Syracuse VA Medical Center. He is a Fellow of the IEEE and of AIMBE.