Proceedings of the 1998 IEEE
International Conference on Robotics & Automation
Leuven, Belgium • May 1998

Towards Force-Reflecting Teleoperation Over the Internet

Günter Niemeyer and Jean-Jacques E. Slotine

Nonlinear Systems Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139
gunter@ai.mit.edu, jjs@mit.edu

Abstract

This paper extends earlier results on stable force-reflecting teleoperation in the presence of significant time-delays to the case, frequent in practice, where the transmission delays are themselves varying with time in an unpredictable fashion. It shows that stability can be preserved through the systematic use of specially designed wave-variable filters. The resulting performance of the teleoperation system is illustrated in simulations, and is consistent with reasonable expectations on "ideal" behavior. The results may provide a practical tool for implementing force-reflecting teleoperation over the Internet.

1 Introduction

Teleoperation has enjoyed a rich history and has lead both to many practical applications and to a broad vision of interacting with environments far removed from the user (Sheridan, 1989). To provide a more complete interaction, force feedback is often included, since this information can considerably improve the user's ability to perform complex tasks (Sheridan, 1992).

However, by their very definition, teleoperation systems frequently experience significant time-delays in the communications between local and remote sites. Untreated, even small delays can lead to instability due to unwanted power generation in the communications. Through the introduction of the wave variable concept (Niemeyer and Slotine, 1991), based on a reformulation of the passivity formalism of (Anderson and Spong, 1989), systematic analysis tools can be developed to understand these problems. Moreover, *stable* force reflecting teleoperators can be designed and shaped to act as simple virtual tools when exposed to large delays. They are also transparent to the user when delays are below human reaction time (Niemeyer and Slotine, 1997a).

Existing results have concentrated on unknown but constant time delays. Here we extend these ideas to time-varying transmission delays. Two classes of practical problems come to mind. First, variable delays due to motion of the slave systems. For instance, space-based or underwater telerobotic applications involve moving vehicles and thus experience changing transmission times to and from the stationary operator. The resulting variations, however, are typically very slow and in practice can often be ignored.

Second, and significantly more interesting, are rapidly and possibly randomly varying transmission delays. This is the case, for instance, in satellite-based transmission through varying relay sites. Perhaps more intriguingly, this is also the case in the Internet, which has frequently been suggested as a means for creating teleoperation systems between a variety of remote sites, given the availability of ever less expensive force-reflecting interfaces. Information is transmitted in small packets and is routed in real-time through a possibly large number of intermediate stops. While average latencies may be low, the instantaneous delays may increase suddenly due to rerouting or other network traffic. In the extreme, the connection may be temporarily blocked. Such effects distort the signals (as seen with respect to time), can introduce high-frequency data, and can lead to instability if left untreated.

For example, examine in Figure 1 the observed trans-continental round trip delay times between MIT and California, a distance of roughly 4000 kilometers. The average latency of 0.1 seconds is comparable to the human reaction time, so that its effects can be made largely transparent to the user. However, the variation in the delay is strong, rapidly changing more than 50%, and contains many components near 10Hz. In a closed-loop system, such fluctuations may interact with the 0.1 second delay and cause stability problems.

The solution proposed in this paper exploits the

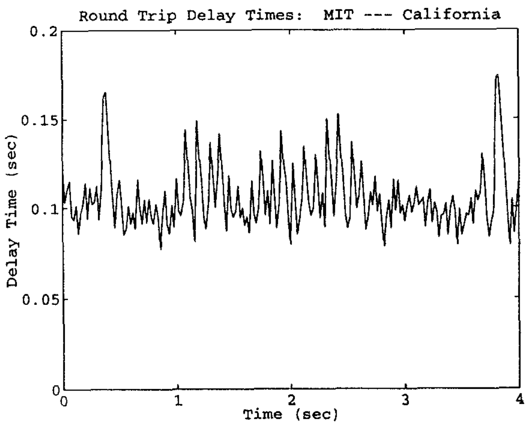

Figure 1: Observed round trip delay times between MIT and California, sending 50 data sets per second.

wave variable formalism to incorporate information as soon as it arrives, but modulating it if necessary so as to explicitly insure proper behavior. For a constant delay, the approach automatically simplifies to the standard wave variable method, which implicitly limits performance as needed to maintain stability at that delay (Niemeyer and Slotine, 1997b). For zero delay, it automatically simplifies to a classic force-reflecting configuration. As a result, the system feels soft when the instantaneous delay is large, then becomes crisp as the delay shrinks.

We first briefly review wave variables and their application to teleoperation in Section 2. We then detail the problems caused by variable time-delays and introduce this paper's suggested solution in Section 3. As the central part of the solution, the reconstruction filter is examined in more depth in Section 4. Finally, we demonstrate the complete system via simulation in Section 5, and offer some concluding remarks in Section 6.

2 Wave Variables

Wave variables are central to our developments and are briefly reviewed here. We refer the reader to (Niemeyer and Slotine, 1997a, 1997c) for a detailed discussion.

2.1 Definition

The key feature of wave variables is their encoding of velocity and force information. In particular, we define

$$\mathbf{u} = \frac{b\dot{\mathbf{x}} + \mathbf{F}}{\sqrt{2b}} \qquad \mathbf{v} = \frac{b\dot{\mathbf{x}} - \mathbf{F}}{\sqrt{2b}} \tag{1}$$

Here $\mathbf{u}$ denotes the forward or right moving wave, and $\mathbf{v}$ denotes the backward of left moving wave. The characteristic wave impedance b is a positive constant or a symmetric positive definite matrix and assumes the role of a tuning parameter, which allows matching a controller to a particular environment or task.

To clarify this definition, consider the layout of a teleoperator system shown in Figure 2. At the local (master) side, the velocity $\dot{\mathbf{x}}_m$ and force $\mathbf{F}_m$ information is combined to provide a command wave signal $\mathbf{u}_m$, which reaches the remote (slave) side after a delay T. There it is decoded into a velocity $\dot{\mathbf{x}}_s$ or force $\mathbf{F}_s$ command. The same process returns information back to the master side.

Note that the inherent combination of velocity and force data makes the system well suited for interaction with unknown environments. Indeed it behaves like a force controller when in contact with a rigid object and like a motion controller when in free space. The parameter b also allows online trade-off between the two quantities, fine tuning the behavior.

A wave signal itself is then best described as a '*move or push*' command, where the sign determines the direction. The receiving side will then either move or apply forces depending on the current situation. The returning wave will have the opposite sign, i.e. 'push back', if no motion was possible. Or it will have the same sign, i.e. 'move with', if motion occurred.

2.2 Passivity

The original motivation for introducing wave variables, is their effect on passivity. Indeed the power input becomes

$$P_{in} = \dot{\mathbf{x}}^T\mathbf{F} = \tfrac{1}{2}\mathbf{u}^T\mathbf{u} - \tfrac{1}{2}\mathbf{v}^T\mathbf{v} \tag{2}$$

where $\frac{1}{2}\mathbf{u}^T\mathbf{u}$ is the power flowing in the main forward direction and $\frac{1}{2}\mathbf{v}^T\mathbf{v}$ gives the power flowing back.

Now remember the condition for passivity, which requires more input energy than return energy

$$\int_0^t P_{in}\,\mathrm{d}\tau = \int_0^t \dot{\mathbf{x}}^T\mathbf{F}\,\mathrm{d}\tau \geq -E_{store}(0) \quad \forall t \geq 0 \tag{3}$$

where $E_{store}(0)$ denotes the initial stored energy (Slotine and Li, 1991).

In the wave domain, this condition becomes

$$\int_0^t \tfrac{1}{2}\mathbf{v}^T\mathbf{v}\,\mathrm{d}\tau \leq \int_0^t \tfrac{1}{2}\mathbf{u}^T\mathbf{u}\,\mathrm{d}\tau + E_{store}(0) \quad \forall t \geq 0 \tag{4}$$

Not surprisingly a system is passive if the energy in the returning wave $\mathbf{v}$ is limited to the energy provided by the ingoing wave $\mathbf{u}$ or stored initially.

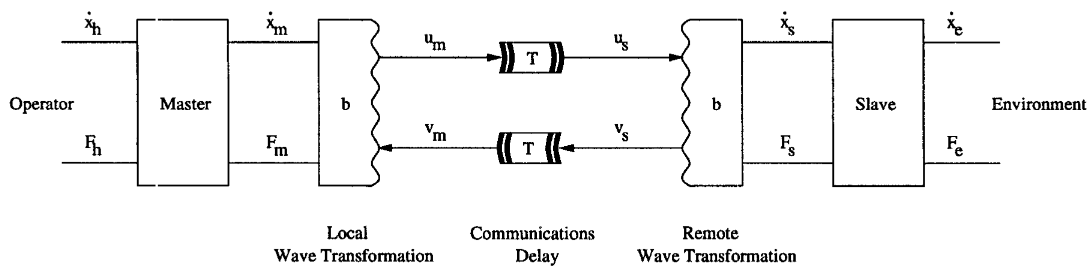

Figure 2: The wave based teleoperator transforms both local and remote information into wave variables before transmission to the other side.

Notice that each wave signal essentially contains its own power to execute the command, independent of its dual wave. In contrast, both power variables $(\dot{\mathbf{x}}, \mathbf{F})$ must be known in order to understand the power flow and passivity. It is this property and the new condition (4), which make wave variables robust to delays.

2.3 Time Delays

If each wave contains its own power and passivity only requires that the output or return wave be limited by the input or command wave regardless of phase, then delaying a wave signal does not alter passivity. Instead, it simply stores the energy in the wave for the delay time and releases it thereafter.

Consider again the basic teleoperator layout of Figure 2. Assume for now a constant delay in the communications between local and remote site. The total power input into the communications block at any point in time is given by

$$P_{in} = \dot{\mathbf{x}}_m^T \mathbf{F}_m - \dot{\mathbf{x}}_s^T \mathbf{F}_s \tag{5}$$

where the minus sign appears because power is considered positive while flowing in the main direction from left to right.

Substituting the wave transformation equations, we can also compute this power input as

$$P_{in} = \tfrac{1}{2}\mathbf{u}_m^T\mathbf{u}_m - \tfrac{1}{2}\mathbf{v}_m^T\mathbf{v}_m - \tfrac{1}{2}\mathbf{u}_s^T\mathbf{u}_s + \tfrac{1}{2}\mathbf{v}_s^T\mathbf{v}_s \tag{6}$$

where all variables are measured at the current time t.

But the communications transmits and delays the waves as

$$\mathbf{u}_s(t) = \mathbf{u}_m(t-T) \tag{7a}$$

$$\mathbf{v}_m(t) = \mathbf{v}_s(t-T) \tag{7b}$$

Substituting into (6) and integrating, we find that all input power is stored according to

$$\int_0^t P_{in}\,\mathrm{d}\tau = E_{store}(t) = \int_{t-T}^t \tfrac{1}{2}\mathbf{u}_m^T\mathbf{u}_m + \tfrac{1}{2}\mathbf{v}_s^T\mathbf{v}_s\,\mathrm{d}\tau \geq 0$$

assuming zero initial conditions. The wave energy in $\mathbf{u}_m$ and $\mathbf{v}_s$ is thus temporarily stored while the waves are in transit, making the communications not only passive but also lossless. This is independent of the delay time T, and does not require knowledge thereof.

2.4 Position Tracking

In the basic form, a wave based teleoperator transmits the wave signals, which encode both velocity and force, but do not contain any explicit position information. Position tracking is guaranteed only implicitly. Consider the deflection or position error between the master and slave sides.

$$\Delta\mathbf{x}(t) = \mathbf{x}_m(t) - \mathbf{x}_s(t) \tag{8}$$

$$= \frac{1}{\sqrt{2b}}\int_0^t \mathbf{u}_m(\tau) + \mathbf{v}_m(\tau) - \mathbf{u}_s(\tau) - \mathbf{v}_s(\tau)\,\mathrm{d}\tau$$

which is obtained by solving (1) for velocity and integrating. Substituting the delay equations (7), we have

$$\Delta\mathbf{x}(t) = \frac{1}{\sqrt{2b}}\int_{t-T}^t \mathbf{u}_m(\tau) - \mathbf{v}_s(\tau)\,\mathrm{d}\tau \tag{9}$$

which will reach zero when the system comes to rest and the wave commands are zero for the T seconds.

Besides requiring a constant delay, this argument is also questionable, as it assumes perfect numerical integration. Indeed an explicit position command for either master or slave can only be obtained by decoding the wave signals into a velocity command and integrating.

2.5 Wave Integrals

To avoid the numerical integration step and provide explicit position feedback, we can use and transmit the wave integrals in parallel with the wave signals themselves. Indeed, just as the wave signals encode

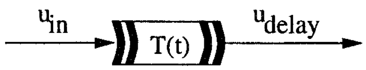

Figure 3: Variable time delay in the wave domain.

velocity and force, their integrals encode position and momentum information.

The integrated wave variables are defined as

$$\mathbf{U}(t) = \int_0^t \mathbf{u}(\tau)\,\mathrm{d}\tau = \frac{b\mathbf{x}+\mathbf{p}}{\sqrt{2b}} \qquad (10a)$$

$$\mathbf{V}(t) = \int_0^t \mathbf{v}(\tau)\,\mathrm{d}\tau = \frac{b\mathbf{x}-\mathbf{p}}{\sqrt{2b}} \qquad (10b)$$

where $\mathbf{x}$ denotes position and $\mathbf{p}$ denotes momentum, which is the integral of force

$$\mathbf{p} = \int_0^t \mathbf{F}\,\mathrm{d}\tau \qquad (11)$$

Note in many cases, we place little or no importance on the actual momentum value and it can often be eliminated from the system if so desired. Also the addition of wave integrals does not change the passivity arguments. It only provides an explicit feedback path for what is already theoretically guaranteed.

3 Variable Time Delays

We now focus our attention on variable delay as illustrated in Figure 3. Based on the previous discussion, we examine just an isolated single wave delay. The overall system remains passive if this element stays passive, i.e. if its output energy is limited by its input energy. And the position tracking is guaranteed its output wave integral tracks its input integral. As such the forward and return delays may be different and are handled separately by duplicating the following efforts for both transmission paths.

3.1 Untreated Variable Delays

Let us first understand the effect of a variable delay left untreated. Thus we use

$$u_{out}(t) = u_{in}(t-T(t)) = u_{in}(t_s(t)) \qquad (12)$$

where t_s is the sample-time, for which the input value is currently presented at the output. The difference between the current time t and the corresponding input sample time $t_s(t)$ is the delay $T(t)$.

As the delay varies, the wave signal is distorted. Indeed, if the delay increases, the sample time changes only slowly and the input values are held longer. Hence the signal is stretched. In the extreme, if the sample time becomes constant and the delay grows as fast as time itself, the output also becomes constant. Note we assume that the order of the wave signal is preserved, i.e. that data arrives at the remote site in the same order it is transmitted. This implies that the sample time will never go backwards and the delay time can not increase faster than time itself,

$$\dot{T} \leq 1 \qquad (13)$$

In contrast, if the delay shortens, the sample time and hence the output signal change more rapidly. In essence the signal is compressed. Here the extreme case can lead to shock-waves where multiple data samples arrive at the remote site simultaneously. This implies discontinuities and a jump in the output signal.

The delay variation and the corresponding changes in the wave signal may easily effect the system, if it is in any way correlated to the wave signal itself. For example, remember that the wave signal is interpreted as a 'push' command. If the signal is expanded during a positive push and compressed during a negative command, the output will be biased in the position direction.

More formally, both the wave integral, which determines position tracking, and the wave energy, which determines passivity, are no longer conserved.

$$U_{out}(t) = \int_0^t u_{out}(\tau)\,\mathrm{d}\tau \neq U_{in}(t-T(t)) \qquad (14a)$$

$$E_{out}(t) = \int_0^t u_{out}^2(\tau)\,\mathrm{d}\tau \neq E_{in}(t-T(t)) \qquad (14b)$$

Thus neither position tracking nor passivity are guaranteed. Also as part of the signal compression/expansion, the frequency content changes which may produce other unexpected effects.

3.2 Integral Transmissions

The above discussion illustrates that using the distorted wave signal based on (12) does not produce the desired results, as it does not preserve the wave integral or energy. As both of these quantities are central to the stability and performance of wave based systems, we propose the following solution, illustrated in Figure 4.

Instead of transmitting the wave signal itself through the delay and then integrating, transmit both the wave integral and wave energy explicitly:

$$U_{delay}(t) = U_{in}(t-T(t)) = \int_0^{t-T(t)} u_{in}(\tau)\,\mathrm{d}\tau \qquad (15a)$$

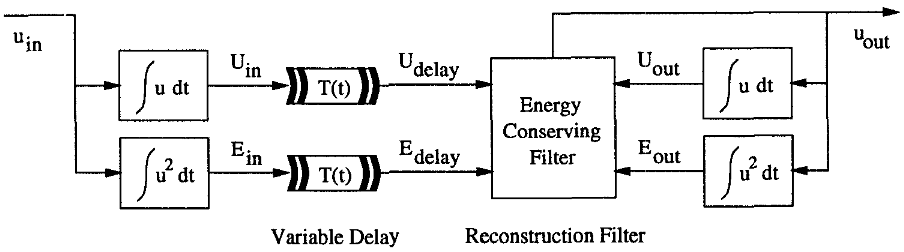

Figure 4: Transmitting the wave integral and energy, and reconstructing the output wave signal based on these quantities can overcome the damaging effects of variable delay distortions.

$$E_{delay}(t) = E_{in}(t-T(t)) = \int_0^{t-T(t)} u_{in}^2(\tau)\,\mathrm{d}\tau \quad (15b)$$

The integration process thus remains consistent, though the resulting values are delayed. Also compute the equivalent quantities for the output wave signal $U_{out}(t)$ and $E_{out}(t)$ from (14). Then explicitly reconstruct the output wave signal such that its integral tracks the delayed input integral

$$U_{out}(t) \to U_{delay}(t) \quad (16a)$$

while using only the available energy

$$E_{out}(t) \leq E_{delay}(t) \quad (16b)$$

The following section details this process.

Such an explicit reconstruction has several advantages. First, passivity is guaranteed by the definition of the system and is independent of the actual delay and or fluctuations thereof. We can build a passive and stable teleoperator on top of such communications. Second, explicit use of the wave integrals provides explicit position feedback. To this end, the wave integrals should be computed directly from position measurements.

If the delay is constant, no distortion is present and the output of such a reconstruction filter should equal the delayed original wave input. But should the delay fluctuates, the input may change more rapidly than can be matched with the incoming energy. In such cases, the filter is forced to smooth the signal to conserve energy. Much like the wave based controllers, we see the system introduce an automatic performance limitation to remain passive.

4 Reconstruction Filter

Many alternatives are possible to the filter problem defined by (16). To better understand our solution, let us first examine the responses to impulse inputs. While the system is not linear, in the sense that the sum of two inputs does not necessarily produce the sum of the two individual outputs, such responses illustrate the basic behavior as a function of the system parameters. Also impulse inputs appear in real problems if transmission is temporarily blocked and the built-up data is released altogether.

4.1 Impulse Response

First define the wave 'distance to go' $U(t)$ and energy reserve $E(t)$ available to the filter as

$$U(t) = U_{delay}(t) - U_{out}(t) \quad (17a)$$
$$E(t) = E_{delay}(t) - E_{out}(t) \geq 0 \quad (17b)$$

Our solution takes the form

$$u_{out}(t) = \begin{cases} \alpha \dfrac{E(t)}{U(t)} & \text{if } U(t) \neq 0 \\ 0 & \text{if } U(t) = 0 \end{cases} \quad (18)$$

where α is a tunable parameter.

Such filters have several interesting properties. For zero input, they maintain the fixed ratio

$$\frac{E(t)}{U(t)^\alpha} = \mathrm{c} \quad (19)$$

where c is a constant determined by initial conditions. Indeed,

$$\frac{\mathrm{d}}{\mathrm{d}t}\frac{E(t)}{U(t)^\alpha} = \frac{\dot{E}\,U - \alpha\,\dot{U}\,E}{U^{\alpha+1}} = \frac{-u_{out}^2\,U + \alpha\,u_{out}\,E}{U^{\alpha+1}} = 0$$

Thus for an impulse response, $u_{out}(t)$ may be written as

$$u_{out}(t) = \begin{cases} \alpha\,\mathrm{c}\,U(t)^{\alpha-1} & \text{if } U(t) \neq 0 \\ 0 & \text{if } U(t) = 0 \end{cases} \quad (20)$$

which is continuous at $U(t) = 0$ for $\alpha > 1$ and bounded for $\alpha = 1$. Furthermore we have

$$\frac{\mathrm{d}}{\mathrm{d}t}U(t) = -u_{out}(t) = -\alpha\frac{E(t)}{U(t)} = -\alpha\,\mathrm{c}\,U(t)^{\alpha-1} \quad (21)$$

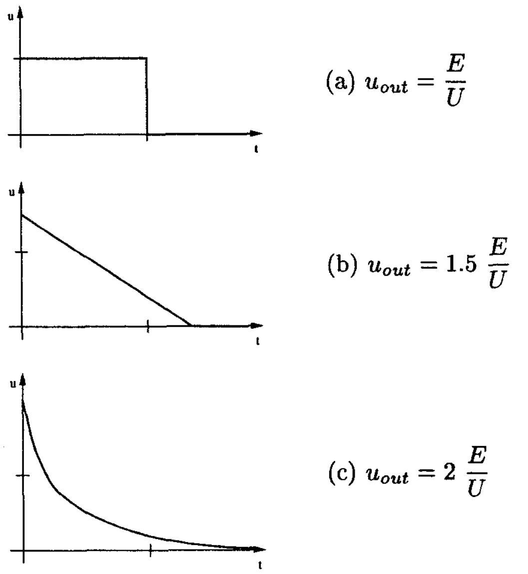

Figure 5: Impulse responses for the reconstruction filter under a variety of parameter settings.

Consequently, both the distance to go $U(t)$ and the remaining energy $E(t)$ reach zero. Given the first order nature, $U(t)$ will not cross the origin during this convergence process and $E(t)$ remains positive.

Varying the parameter α selects how quick versus smooth the convergence process is. Figure 5 shows the impulse responses for the values of $\alpha = 1$, $\alpha = 1.5$, and $\alpha = 2$. In the first case, the response is constant until the goal is reached and all energy has been used. This provides the fastest time to reach the goal, but also contains discontinuities which may be disruptive in practice.

The response is linear for $\alpha = 1.5$ and exponential for $\alpha = 2$. We will constrain the following developments to the latter case. The response is given by

$$u(t) = \frac{2\,E_0}{U_0}\,e^{-\lambda t} \tag{22}$$

where the bandwidth λ is determined as

$$\lambda = \frac{2\,E}{U^2} \tag{23}$$

Notice that the speed of the response depends on the amount of energy E available for a given distance U. The more energy, the faster the response. This also suggests a practical advantage: To limit the frequency content coming out of the reconstruction filter, we can saturate the stored energy reserve via

$$E(t) \leq \tfrac{1}{2}\,\lambda_{\max}\,U(t)^2 \tag{24}$$

4.2 Discrete Implementation

In practice, with a changing input $U_{delay}(t)$, the filter will not reach zero but continue to track the input with the first order behavior of (21). Limiting the energy according to (24) further smoothes the signal.

Nevertheless, when $U(t)$ and $E(t)$ approach zero the division in (18) is hard to compute. In addition a finite sampling rate approximation may make $E(t)$ negative in violation of (16). We therefore suggest a discrete implementation which accounts for a finite sample rate.

Rather than approximating the continuous derivatives, we rederive the equations starting with the requirement

$$\frac{E_{n+1}}{U_{n+1}^2} = \frac{E_n}{U_n^2} = \text{constant}$$

for a zero input. This leads to the output

$$u_{out} = \begin{cases} 0 & \text{if} \quad U = 0 \\ \dfrac{U}{\Delta t} & \text{if} \quad U^2 \leq E\,\Delta t \\ \dfrac{2EU}{U^2 + E\Delta t} & \text{if} \quad U^2 > E\,\Delta t \end{cases} \tag{25}$$

where Δt is the discrete sample time step.

The first case appears if there is no further energy or distance to go. The second case is new and accounts for the possibility that the convergence happens within one time step. Indeed in this case, explicitly reset $E = 0$ as the excess energy is unused. Finally the third case is adjusted to account for the zero order hold between samples.

5 Simulation Results

The resulting behavior is illustrated by the following simulation, which is chosen to be fairly extreme in its large and sudden delay variation. Consider a standard wave-based teleoperation setup. The remote (slave) side is contacting a rigid object. On the local (master) side, the operator is inputing a constant force together with some damping. The total delay is constant at 0.1 second, distributed between the forward and return path.

At time $t = 0.5$ the return transmission (from slave to master) is blocked until time $t = 1.0$. For a half second, no signals are received on the master side. The slave side is unaffected. Thereafter the communications is unblocked and the entire data is instantly

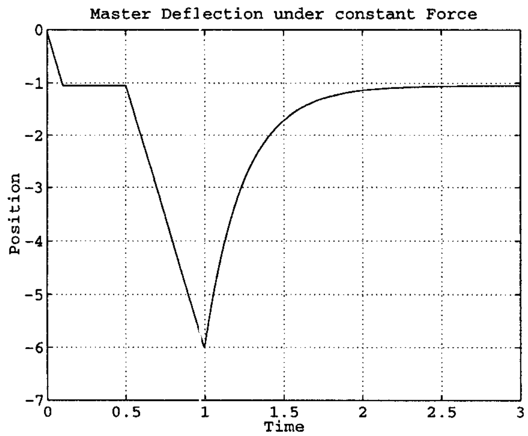

Figure 6: Master deflection under constant force and experiencing temporary transmission black out.

available. For the remainder of the time, the delay is again a constant 0.1 second.

Figure 6 shows the resulting master deflection. For the first 0.1 seconds, no signal is returning and the master is forced to deflect. After the regular delay time, the response from the remote site is available and the position holds constant, where the deflection is based on the equivalent system stiffness b/T. At $t = 0.5$, as the return signal disappears again, the position is once again forced to deflect to indicate the lack of knowledge. Finally, at $t = 1.0$ the situation returns to normal and the master can slowly reclaim its original position.

Note that the black-out time of 0.5 seconds directly translates to the time constant for the recovery. Shorter black-outs imply faster recoveries.

6 Concluding Remarks

The expansion and popularity of the Internet may allow force-reflecting teleoperation to achieve its full potential of letting users not only see, but physically interact with many distinct remote sites from the comfort of their local PC. The large fluctuations in time-delays for such network applications create unique and interesting problems, but we believe wave variables are well suited to these situations. With minimal computation and no advance knowledge of the transmission characteristics, wave variable filters can enable network based force-reflecting teleoperation, transparently during normal operation and degrading gracefully when the network is overloaded.

Acknowledgements

This report describes research done at the Nonlinear Systems Laboratory and the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the work has been provided in part by NASA/JPL.

References

Anderson, R. J. and Spong, M. W. (1989). Bilateral control of teleoperators with time delay. *IEEE Transactions on Automatic Control*, 34(5):494–501.

Niemeyer, G. (1996). *Using Wave Variables in Time Delayed Force Reflecting Teleoperation.* PhD thesis, MIT, Cambridge, MA.

Niemeyer, G. and Slotine, J.-J. E. (1991). Stable adaptive teleoperation. *IEEE Journal of Oceanographic Engineering*, 16(1):152–162.

Niemeyer, G. and Slotine, J.-J. E. (1997a). Designing force reflecting teleoperators with large time delays to appear as virtual tools. In *Proc. of 1997 IEEE Int. Conf. on Robotics and Automation*, pages 2212–2218.

Niemeyer, G. and Slotine, J.-J. E. (1997b). Telemanipulation with time delays. to be published in International Journal of Robotics Research.

Niemeyer, G. and Slotine, J.-J. E. (1997c). Using wave variables for system analysis and robot control. In *Proc. of 1997 IEEE Int. Conf. on Robotics and Automation*, pages 1619–1625.

Sheridan, T. B. (1989). Telerobotics. *Automatica*, 25(4):487–507.

Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control.* The MIT Press, Cambridge, MA.

Slotine, J.-J. E. and Li, W. (1991). *Applied Nonlinear Control.* Prentice Hall, Englewood Cliffs, New Jersey.

Slotine, J.-J. E. and Niemeyer, G. (1993). Telerobotic system. U.S. Patent #5266875.