Discussion on: "Development and Experimental Verification of a Mobile Client-Centric Networked Controlled System"

Nathan A. Wedge, Ahmad Al-Hammouri, Michael S. Branicky, and Vincenzo Liberatore
Department of Electrical Engineering and Computer Science
Case Western Reserve University, Cleveland, Ohio 44106
{nathan.wedge, ahmad.al-hammouri, mb, vincenzo.liberatore}@case.edu

Summary of Contributions

The paper under discussion [12] presents the analytical results underlying the authors' design of a networked control system (NCS) operating over a cellular network. Their formulation fits the common framework [7, 14], in which there are transmission delays present in the sensor/controller and controller/actuator paths. One key difference between this work and existing literature relates to the use of a client-centric system, which is mandated by security issues. However, this seems to make little to no functional difference in the implementation or resulting analysis. The chief theoretical results of the paper are based upon linear matrix inequality (LMI) theory. In particular, the authors have provided an interesting, LMI-derived analytical result for ascertaining the stability of a networked control system with delays. Some existing work [4] has pursued similar derivations for systems with time-varying delays without applying the results directly to networked control systems; a complementary approach [1] provided LMIs that could be used to guarantee the stability of an undelayed networked control system under packet losses. Another key point in the work here is that the authors assume that the application of control signals and delivery of state information happens on the boundaries of sampling intervals. This approach is advantageous to their analytical approach, which augments the system with delayed state measurements up to a certain bound. Moreover, it provides results for situations in which the delay is greater than the sampling period, a scenario that is often neglected in NCS research, but fairly common in wide-area networks. In summary, their different viewpoints plus the development of an end-to-end system under quite restrictive bandwidth and delay constraints is of obvious value to the progress toward a larger vision of networked control.

The Bigger Picture

Networked control is central to the larger vision of *tele-epistemology*, which posits that knowledge can be acquired from a distance through network-mediated interaction [3]. The fundamental tele-epistemological tenets require that physical environments be controlled irrespectively of physical or communication distance. The tele-epistemological direction is substantially advanced by this paper [12] in that the (wireless) telephony network is arguably the most pervasive, widespread, and reliable network. Whereas most previous work in networked control had focused on Local Area Networks (LANs), such as those found in a single factory plant or in a peripheral interconnection bus, the approach in this paper can scale to practically arbitrary geographical distances. As a consequence, it is a significant extension of the networked control field.

Geographical scalability does require bridging long distances, and in practice such a capability is necessary but not sufficient. Geographic closeness gives rise to one distance function, but other critical metrics are induced, for example, by the interconnection topology—which results in geographical and network distances that are often unrelated. For example, a telephone GPRS is but one single network in the Internet. This type of interconnectivity has effectively created a *network of networks*, with the additional benefits afforded by flexibility, evolvability, and interoperability [2]. This paper is an excellent first step toward interconnected, geographically-scalable networked control, and the natural future steps should complement geographic distances with the inherently complex structure of the interconnection topology.

The network of networks architecture has a profound impact on networked control. First, an individual network can usually be configured to support

Quality-of-Service (QoS) assurances. However, QoS cannot be guaranteed in general in the Internet (see, for example, the seminal, Turing-award winning paper by Cerf and Kahn [2]). Specifically, a telephony GPRS provides hard bounds on jitter, loss rates, and bandwidth, but no such guarantees exist in the best-effort model that is prevalent in the Internet. In turn, poor QoS impacts the real-time properties of the feedback loop. Furthermore, the lack of QoS also implies relatively poor clock synchronization [11]. A striking consequence is that, in an asynchronous distributed system, there is no deterministic and fault-tolerant algorithm for two end-points to reach agreement on even one bit [5].

As for networked control, this lack of synchronization does not permit certain schemes upon which much previous work relied (e.g., [6, 7, 13, 14]). Furthermore, this previous work—primarily on LANs—indicates that *event-driven control* that takes immediate action upon receipt of the sensor data (by calculation and dispatch of a control signal) and similarly upon receipt of the control signal (by beginning actuation) would provide superior performance given an optimal control strategy [6]. Thus, much previous work has focused on the use of event-driven control. In contrast, the paper under discussion assumes that the application of control signals and delivery of state information happens on the boundaries of sampling intervals. In extensive *wide-area* network experiments, we have recently found that this more regular application of control signals can lead to better overall performance.

In summary, networked control in the Internet age is a rich and mostly unexplored research area. It is the natural future direction implied by this paper.

Some Implementation Issues

In addition to the issues above, the complexity of distributed control applications in general has increased substantially. Control software itself is becoming more complex, as measured by the number of its function points, for example. Moreover, control problems bring hard real-time constraints. Software Engineering has made impressive strides for such complex, real-time distributed applications (e.g., [10]); Computer Networks has seen the creation of real-time protocols (e.g., [8]). Although some of those parallel contributions address primarily the software development process and communication standards, many impact directly the communication timing, which in turn affects control stability and performance.

As one specific example, real-time distributed applications, such as networked control, favor connectionless protocols like UDP, due to the low overhead and the potentially detrimental effect of the use of stale information (and therefore lack of usefulness for a mechanism that achieves reliability by retransmission). In this paper, though UDP is the mechanism for control signal delivery, the transmission of sensor data is accomplished over FTP (and hence, TCP). Justification for this selection is conspicuously absent, especially considering the packet loss scenarios that follow in the paper. Aside from choosing TCP as a transmission protocol for transferring sensor data, choosing FTP in particular—among other applications, e.g., plain TCP sockets—exaggerates the problem even more. This is true because even though FTP maintains a single persistent connection to communicate commands, it starts a new TCP connection for every data transfer request (i.e., every FTP `get` command) and closes the connection after the transmission completes [9]. Therefore, at least one round-trip time is wasted to set up such a connection before transferring data starts, see Figure 1.

Consequently, the choice of FTP both exposes signals to more delays/losses and imposes restrictions on deploying this implementation in applications where sensors usually have very limited resource constraints. This questionable protocol choice probably contributes to the poor performance of the cellular network in the experimental study. Though the authors make the point that voice calls are given higher priority than data calls (i.e., their control system), their measured delays seem extreme, and it is doubtful that any modern cellular client would tolerate data delays that average 18 seconds. Furthermore, their quoted data rates are small compared to the attainable rates mentioned. Barring the pos-

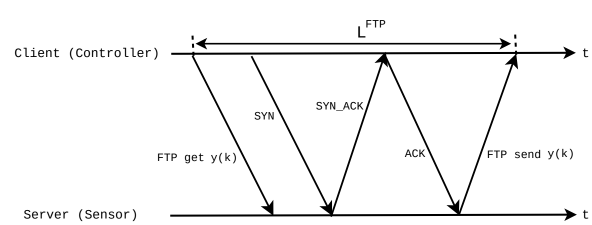

Figure 1: FTP data transfer. FTP-`get`, and FTP-`send` packets of the original paper decompose into FTP-`get`, SYN, SYN-ACK, ACK, and FTP-`send` packets.

sibility of extreme overhead (i.e., if packets contain only a few bytes of data, but there is significant overhead due to headers and connection establishment), this weakens the particular choice of cellular protocol from the bandwidth standpoint.

A related issue is the application of PCs in the control loop. One advantage of networked control systems in general is the potential to save on resources by not allocating control components at the location of the plant. However, the scheme presented used a PC capable of FTP service in close proximity to the plant, both to deliver sensor data and to implement the control commands. Thus, at least in the case of the specific example system, the amount of resources devoted to FTP service and network data transmission by the PC (or, later, presumably by a network-enabled chipset) begs the question of whether they should be leveraged to accomplish the local control task directly, with the client providing remote supervision (e.g., sending setpoints).

Conclusions

Overall, the authors have presented the particular concerns of control design over a cellular network medium and presented an implementation of such a networked control system. We pointed out a few issues with that implementation. However, their development of methodology for implementing control systems under quite restrictive bandwidth and delay constraints is of obvious value to the progress toward the larger tele-epistemological vision of globally-distributed control over networks of networks.

References

[1] B. Azimi-Sadjadi. Stability of networked control systems in the presence of packet losses. In *Proc. IEEE Conf. on Decision and Control*, pages 676–681, Maui, HI, December, 2003.

[2] V. Cerf and R. Kahn. A protocol for packet network interconnection. *IEEE Trans. on Communications Technology*, 22(5):627–641, 1974.

[3] K. Goldberg, editor. *The Robot in the Garden: Telerobotics and Telepistemology in the Age of the Internet*. MIT Press, Cambridge, MA, 2001.

[4] E.T. Jeung, D.C. Oh, J.H. Kim, and H.B. Park. Robust controller design for uncertain systems with time delays: LMI approach. *Automatica*, 32(9):1229–1231, 1998.

[5] N.A. Lynch. *Distributed Algorithms*, Morgan Kaufmann, San Francisco, CA, 1996, chap. 21.

[6] J. Nilsson. *Real-Time Control Systems with Delays*. Ph.D. thesis, Automatic Control, Lund Institute of Technology, Lund, Sweden, 1998.

[7] J. Nilsson, B. Bernhardsson, and B. Wittenmark. Stochastic analysis and control of real-time systems with random time delays. *Automatica*, 34(1):57–64, 1998.

[8] L.L. Peterson and B.S. Davie. *Computer Networks*, Morgan Kaufmann, San Francisco, CA, 2000, chap. 9.

[9] J. Postel and J.K. Reynolds. RFC 959: File transfer protocol, October 1985.

[10] E. Shokri and P. Shreu. Real-time distributed object computing—an emerging field. *Computer*, 33(6):36–44, June 2000.

[11] A.S. Tanenbaum and M. Van Steen. *Distributed Systems. Principles and Paradigms*, chapter 4. Prentice-Hall, Upper Saddle River, NJ, 2002.

[12] A. Tzes, G. Nikolakopoulos, and I. Koutroulis. Development and experimental verification of a mobile client-centric networked control system. *European Journal of Control*, 2005.

[13] W. Zhang. *Stability Analysis of Networked Control Systems*. Ph.D. thesis, Electrical Engineering and Computer Science, Case Western Reserve University, Cleveland, OH, 2001.

[14] W. Zhang, M.S. Branicky, and S.M. Phillips. Stability of networked control systems. *IEEE Control Systems Magazine*, 21(1):205–214, 2001.