

Conclusions from the European Roadmap on Control of Computing Systems

Karl-Erik Årzén*, Anders Robertsson, Dan Henriksson
Dept of Automatic Control, LTH
Lund University, Box 118, SE-221 00 Lund, Sweden

Mikael Johansson, Håkan Hjalmarsson, Karl Henrik Johansson
Dept of Signals, Sensors and Systems
Royal Institute of Technology, SE-100 44 Stockholm, Sweden

1. Background

The use of control-based methods for resource management in real-time computing and communication systems has gained a substantial interest recently. Applications areas include performance control of web-servers, dynamic resource management in embedded systems, traffic control in communication networks, transaction management in database servers, error control in software systems, and autonomic computing. Within the European EU/IST FP6 Network of Excellence ARTIST2 on Embedded System Design a roadmap on Control of Real-Time Computing Systems has recently been completed. The focus of the roadmap is how flexibility, adaptivity, performance and robustness can be achieved in a real-time computing or communication system through the use of control theory. The item that is controlled is in most cases the allocation of computing and communication resources, e.g., the distribution or scheduling of CPU time among different competing tasks, jobs, requests, or transactions, or the communication resources in a network. Due to this, control of computing systems also goes under the name of *feedback scheduling*.

The roadmap is divided into six research areas: control of server systems, control of CPU resources, control of communication networks, error control of software systems, feedback scheduling of control systems, and control middleware. For each area an overview is given and challenges for future research are stated. The aim of this position paper is to summarize the conclusions concerning these research challenges. In this paper, we will only cover the first four of the areas above. A preliminary version of the roadmap can be found on <http://www.control.lth.se/user/karlerik/roadmap1.pdf>

1.1 Motivation

Feedback-based approaches have always been used in engineering systems. One example is the flow and congestion control mechanisms in the TCP transport protocol. Typical of many applications of this type is that feedback control is used in a more or less *ad hoc* way without any connections to control theory. During the last 5-10 years this situation

has changed. Today control theory is beginning to be applied to real-time computing system in a more structured way. Dynamic models are used to describe how the performance or quality of service (QoS) depend on the resources at hand. The models are then analyzed to determine the fundamental performance limitations of the system. Based on the model and the specifications control design is performed. In some cases the analysis and design are based on optimization. A recent textbook, [Hellerstein *et al.*, 2004], addresses the question of how to introduce control theory for computer science students.

The publications in the area are rapidly increasing, see e.g., [Hellerstein *et al.*, 2005] and the references therein. However, so far most of the work presented in literature have been conducted by scientists working *either* in the real-time computing or telecommunication fields, *or* in the automatic control field. Unfortunately, this has sometimes led to erroneous models and strange results. In order to achieve good results a multi-disciplinary approach is necessary.

1.2 Modeling

Control of computing systems introduces new types of problems that are not present when controlling physical plants. A main problem is the lack of first principles models. When controlling a physical plant the laws of nature decide to a large degree the behaviour of the plant and can be used to derive dynamical models. A computing system, on the other hand, is a man-made artifact whose internal behaviour is not governed by any laws of nature, at least not on the macroscopic level. This means that it is, generally, not possible to derive any first principles models. One exception, where theoretical models are available is queuing theory [Kleinrock, 1975]. Queuing-theory models have also been used with some success in the design of computing-system controllers. A drawback with queuing models is that they in most cases only hold in the average case and that they assume certain statistical properties, e.g., Poisson traffic.

Computing systems are discrete-event dynamic systems (DEDS). This makes it natural to use a timed discrete-event modeling formalism, such as timed automata or timed Petri nets. This is, however, in many cases too fine-grained

¹Corresponding author. Email: karlerik@control.lth.se

and easily leads to state-space explosion. Another issue is the types of problems that these formalisms typically lend themselves to. Automata-based formalisms are well-suited for expressing and analyzing safety properties and blocking properties. These properties are, however, not the main concerns in performance control. Instead, issues such as stability, performance, and robustness are the main concerns. For these types of problems a time-driven approach is more natural. However, the lack of first principles knowledge necessitates a system identification-based approach, in which a discrete-time difference equation model is derived from measured input and outputs. The models and controllers derived in this way are based on periodic sampling. Although periodic controllers are, to a large extent, the approach that is mostly used in applications, it is from many respects more natural to invoke the controller in an event-driven fashion. For example, in a queue-length control problem it makes more sense to calculate a new control action when a request is queued or dequeued, or every n th enqueue/dequeue event, rather than periodically. A problem with aperiodic or event-based systems and aperiodic control of this type, though, is the lack of theory and tools for analysis and design.

2. Control of Server Systems

More and more business and services rely on Internet and server technology. Queue management is important in all servers, e.g., web servers. Server requests are stored in an input queue, the server or worker thread servicing the requests are stored in the ready queue or in different waiting queues, e.g., in order to access memory. Many aspects of the real-time performance of server systems can be inferred from the behaviour of queues.

In a queue it is the difference between the service rate and the arrival rate that determines the delay experienced by the requests. Two types of actuators can be used. An *enqueue* actuator influence the arrival rate of the queue. One example of this is *admission control*. A *dequeue* actuator instead influences the service rate of the requests. Examples of this type of actuator mechanisms are different forms of quality adaptation.

A queue can be modelled in various ways. Using queuing theory, several types of models can be developed. One example is Tipper's nonlinear flow model [Tipper and Sundareshan, 1990]. At a high level, a queue can be seen as an integrator. This can be modeled using, e.g., a difference equation and then analyzed with control theory. Both Tipper's model and integrator-based models can be used as the basis for control design, e.g., [Robertsson *et al.*, 2003].

Flow models of queuing systems approximate the steady-state behaviour of the queue and are typically more accurate the higher the load is on the server. However, for small and medium loads these types of models are less appropriate. An open question is how to combine queu-

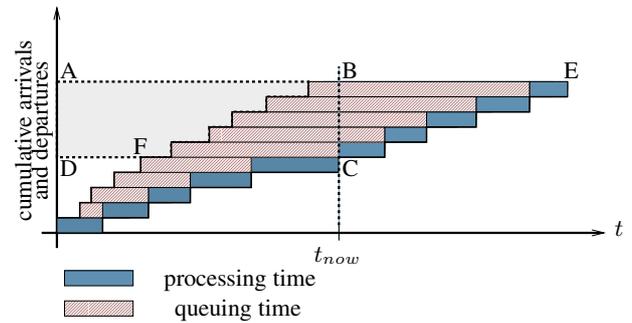


Figure 1 Server queuing and processing delay over time.

ing models with control-theoretic methods. A common approach in delay control is to use nonlinear models from queuing theory for feedforward combined with simple feedback control of PID type, e.g., [Sha *et al.*, 2002]. The aim of the feedforward path is to provide fast setpoint responses, whereas the role of the feedback controller is to compensate for disturbances and incorrect modeling assumptions. An example of the latter is incorrect assumptions about the stochastic nature of arrivals and departures.

In [Henriksson *et al.*, 2004] an improved feedforward scheme is presented, that makes no assumptions about the statistical properties of the traffic. Instead, it predicts future delays as a function of instantaneous measurements of the situation in the server queue. This includes current queue length and the arrival times of the queued requests, which are assumed to be recorded for use in the prediction, see Fig. 1. The basic idea with the predictor is to choose the service rate that achieves a desired average delay of the requests in the system taking into account their average queuing delay up until the current time. By continuously updating the predictor as requests enter and leave the queue, sudden variations are taken care of more rapidly than using the queuing-theoretic models. A similar approach can also be applied to admission control.

Much work has also been performed on multi-class queuing systems using priority queues. Here it is the ratio between the average delays of adjacent service classes that is subject to control.

2.1 Research Challenges

The main challenge in control of servers and software systems in general is to derive a unified theory and framework for performance control of queuing systems that combine elements from control theory and queuing theory and allow an integration of both time-driven liquid model formalisms and event-driven formalisms. Modeling plays a major role here. Which is the right or optimal abstraction level for this type of control problem is still a question with no clear answer. Models at different levels and types need to be combined. We also need better insight in how one should cor-

rectly abstract a real server by a suitable queuing model. It is further desirable to combine time-based models with event-based discrete models. Better understanding is needed for which models types that are best suited for a particular application. It is also possible that new models types must be derived for this type of problems.

The challenges for control are connected to the modelling challenges. How do we develop a control theory based on this type of models? The combination of time-driven control design with event-driven implementation is one major issue. In control in general and process control in particular, the characteristics of different types of control loops and control problems are well known and even in some cases formally categorized. Similarly a number of well-defined controller structures exist, e.g., cascade control and ratio control. The same type of classification is necessary also in control of computer systems. One possibility is to make use of ideas from design patterns to create well-defined patterns for server control problems. Large eCommerce servers are multi-tier systems consisting of web server front-ends, business logic in the intermediate layers, and database servers as back-ends. The overall system is a MIMO system where control is needed at several layers. Model-based predictive control (MPC) is an interesting possibility here. MPC also explicitly handles constraints on control signals and state variables, which is common in queueing problems, e.g., buffer size limits.

Our current notion of dynamics is based on the behaviour of physical systems, e.g., mechanical systems. It is not necessarily so that this type of dynamics also suits software systems. The same holds for stability. It is not completely clear what an unstable software system really means or what type of stability definitions that make sense. Related to this is the question of how we design or program software systems in such a way that they are observable and controllable. Which types of sensor and actuators makes most sense for this type of systems.

In order to make control of server systems applicable on a wider industrial scale it is necessary to have built-in support for this in operating systems and/or middleware. On which level this should be handled is not clear. Should there be a special POSIX/Control standard defined?

3. Control of CPU Resources

Feedback scheduling of CPU resources is an area where fairly much research has been performed, especially for embedded real-time systems. In feedback scheduling the allocation of CPU resources is based on a comparison of the actual resource consumption by, e.g., a set of tasks, with the desired resource consumption. The difference, or control error, is then used for deciding how the resources should be allocated to the different users. Feedback scheduling is primarily suited for applications with soft or adaptive

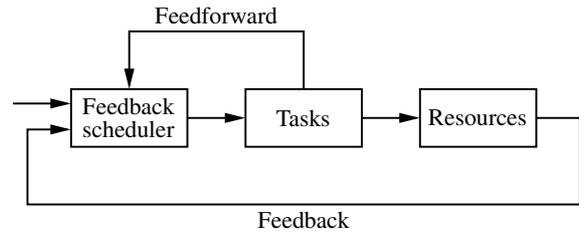


Figure 2 A general feedback scheduling system. The scheduler adjusts the tasks' demands based on feedback from the current use of critical resources. The tasks may also inform the scheduler that they are about to consume more resources (feedforward).

real-time requirements. This includes different types of multimedia applications, but also a large class of control applications. Feedback scheduling of CPU resources has strong relationships with the queue control employed for server systems and many of the results in one area can be directly applied in the other area.

An early result is given in [Stankovic *et al.*, 1999] where the Feedback Control EDF scheduling algorithm is presented. A PID controller is used to regulate the deadline miss-ratio for a set of soft real-time tasks with varying execution times, by adjusting their CPU utilization. The approach has later been extended with an additional PID controller that controls the CPU utilization.

Many scheduling techniques that allow QoS adaptation have been developed. An interesting mechanism for workload adjustments is given in [Buttazzo *et al.*, 1998], where an elastic task model for periodic tasks is presented. A large amount of feedback-based or adaptive global QoS management systems have also been proposed. Some examples are [Chu and Nahrstedt, 1999; Aparah, 1998]. In [Yuan and Nahrstedt, 2003] issues of QoS and energy savings are experimentally evaluated using the CPU scheduler GRACE-OS.

Control-based ideas have also been used for dynamic allocation of bandwidth in aperiodic task servers and for dynamic allocation of resource reservations in reservation-based scheduling. The main application area for these techniques is multimedia applications, e.g., streamed audio and video. The idea behind resource reservation is to explicitly control the computing resources assigned to a given activity (job, task, or application). Each activity receives a fraction (reservation), U_i , of the processor capacity and will behave as if it was executing alone on a slower, virtual processor. The motivation for feedback is the need to cope with incorrect reservations, to be able to reclaim unused resources and distribute them to more demanding tasks, and to be able to adjust to dynamic changes in resource requirements. Hence, a monitoring mechanism is needed to measure the actual demands and a feedback mechanism is needed to perform the reservation adaptation. Two types of feedback are possible. On a global, system-wide level a QoS controller adjusts

the size of the individual reservations given to the different activities based on the measured performance and resource utilization. On a task or activity level, local feedback is employed to adjust the resource requirements of the individual tasks based on the experienced QoS levels and the amount of resources available to the task, as decided by the global QoS controller. The local resource usage can be adjusted through rate adaptation, by executing the task at different service levels, and by job skipping.

3.1 Research Challenges

In addition to several of the challenges for server systems, the following items are important for control of CPU resources. Multiprocessor systems will become common in the near future also for certain embedded applications. So far very little of the control-based methods to CPU resource management have been applied to multiprocessor systems. Power saving is becoming increasingly important in all computer applications, including server systems. Adjusting the CPU speed using, e.g., Dynamic Voltage Scaling (DVS) techniques, is an alternative way of adjusting the service requirements of a task. Minimizing the power consumption is also an important goal in itself for many networked embedded systems, e.g., sensor networks. The joint optimization problem of minimizing energy while still meeting real-time constraints already today receives considerable attention from the research community. However, it is an important area also for the future. Resource management in distributed systems where an activity spans multiple nodes is also an important issue. How do we adapt the resources individually in the different nodes in order to obtain a good global behaviour, e.g., acceptable end-to-end response times?

Hierarchical resource allocation schemes based on dynamic reservations in combination with local feedback control loops for the individual tasks is an interesting and promising approach where more research is needed. How do we enforce the notion of virtual CPUs that execute within a real CPU with, possibly, different scheduling policies, and where the share that each virtual CPU receives of the total CPU resources is dynamically adjusted based on resource requirements and availability?

One of the goals of feedback scheduling is to better make use of scarce resources. If this should be achieved it requires that the feedback scheduling mechanism itself does not consume too much resources. Hence, efficient feedback scheduling mechanisms are of great importance.

4. Control of Communication Networks

The success of the Internet as a worldwide information carrying network can be attributed to the feedback mechanisms that control the data transfer in the transport layer of the IP stack. These algorithms have historically managed to dis-

tribute network resources among contending users in a sufficiently fair and resource-efficient way. An explanation to this is that the control is allocated at the end-systems (users) and hence obey a decentralized structure. Furthermore, together with the source control, buffers have played a key-role during the evolution of the Internet. Since end-users base control action on limited, corrupt and delayed information; buffers are used at links inside the network to smooth out errors in the control, hence making the system more robust. Auxiliary control from the network interior has also been introduced by “intelligent” links that mark or drop packets depending on the traffic load. This is referred to as Active Queue Management (AQM).

Historically, congestion control algorithms have been designed by computer scientists outside the framework of control theory. The tremendous complexity of the Internet makes it extremely difficult to model and analyze, and it has been questioned if mathematical theory can offer any major improvements in this area. Recently, however, significant progress in the theoretical understanding of network congestion control has been made following seminal work by Kelly and coworkers [Kelly *et al.*, 1998]. The key is to work at the correct level of aggregation, which is fluid flow models with validity at longer time-scales than the round-trip time. By explicitly modeling the congestion measure signal fed back to sources, posing the network flow control as an optimization problem where the objective is to maximize the total source utility, it is shown that the rate control problem can be solved in a completely decentralized manner [Kelly *et al.*, 1998; Low and Lapsley, 1999]. This assumes that each source has a (concave) utility function of its rate.

To ensure that the system will reach and maintain a favorable equilibrium, it is important to assess the dynamical properties, such as stability and convergence, of the schemes. Stability of the basic schemes, which allow dynamic rate control and static marking, or dynamic queue management schemes and static source rate control, was established already in [Kelly *et al.*, 1998; Low and Lapsley, 1999] but under idealized settings. A unifying framework for global stability of congestion control laws based on passivity has been proposed in [Wen and Arcak, 2004].

The above results have all ignored the effect of network delay, and assumed that price information is available instantaneously at the source, that the sources take immediate action, and that the new rates affect the link prices instantaneously. However, stability of the protocols in equilibrium depends critically on the feedback delay. Recent research therefore focuses on source- and link control laws that guarantee stability for more general network configurations and delay distributions.

Wireless networks are specially interesting from a resource control point of view. Whereas the link capacities in wireline networks are fixed, the capacities of wireless

links can be adjusted by the allocation of communication resources, such as transmit powers, bandwidths, or time-slot fractions, to different links. Adjusting the resource allocation changes the link capacities, influences the optimal routing of data flows, and alters the total utility of the network. Hence, optimal network operation can only be achieved by coordinating the operation across the networking stack. This is often referred to as *cross-layer optimization*.

A basic question is whether it is worthwhile to introduce advanced resource management and coordination schemes. One way of attacking this problem is to try to determine the *information-theoretic capacity*, which includes optimization over all possible modulation and coding schemes and involves many of the unsolved problems of network information theory. An alternative approach is to focus on *network layer capacity*, where coding and modulation schemes are fixed, and one optimizes over some critical parameters, such as power allocations and scheduling decisions.

4.1 Research Challenges

Control-based approaches in communication networks is a very large research field, in particular if wireless systems, e.g., sensor networks, are included. In order to be able to control the network performance it is necessary to measure and modify the network parameters. The current ISO-OSI stack layer is not ideally supported for cross layer designs where information from the lower layers must be made available at the application layer and where the application layer must be able to modify the behaviour of the lower layer protocols dynamically. Hence, new protocols and protocol models are needed that simplify this.

Theories and engineering principles for dynamically allocating resources in wireless ad hoc networks to ensure quality of service are needed for a wide range of applications. One interesting suggestion is to have a formal, possibly optimization-based, theory for the design of network protocols based on a model of the underlying network and a specification of the application requirements.

While the use of mathematical decomposition techniques as guiding principle for organizational design is well-known in economics and operational research (e.g., [Holmberg, 1995]), the application of such ideas to networked systems has just begun to appear [Chen *et al.*, 2005; Chiang, 2005; Lin and Shroff, 2004; Johansson and Johansson, 2006].

Breaking up the layered structure of the networking stack may also have negative consequences, partially in terms of maintenance and compatibility issues but also in terms of the resulting performance. In particular, it has been observed that cross-layer coordination protocols can introduce dependency relations and unintended interactions [Kwadia and Kumar, 2005]: in some situations, adaption mechanisms in different layers can start working against each other, leading to worse practical performance than in

a layered network. It is thus important to develop control-theoretic tools for analyzing protocol dynamics in order to guarantee stable and efficient overall behaviour.

The control of network performance often requires access to network state variables, such as available bandwidth, round-trip times, and packet loss. These variables are typically not immediately available, but must be estimated from other quantities. The design of reliable and efficient estimators for network state is thus instrumental for many applications, and requires the development of simple and flexible models of network dynamics together with the associated advances in estimation theory.

Improving congestion control is intimately linked to the quality of the used models. The development of accurate fluid flow models will help understanding the limitations communication networks are subject to and provide a basis for new control laws. Present fluid flow models disregard important system aspects and very little has been done in terms of experimental validation of the proposed models.

5. Error Control of Software Systems

The development of completely defect-free complex software systems is extremely difficult, if not impossible. At the same time several large existing software systems are remarkably stable and reliable in the presence of thousands or maybe millions of residual software bugs, e.g., the telecom networks or the WWW system of systems. Hence, rather than focusing the development effort on trying to eliminate all bugs at design time it is important to develop methods that allow us to develop safe and stable software systems that still can utilize COTS-quality software components with a considerable amount of residual bugs. Hence, the focus should be on detection and recovery from software errors at run-time, in addition to elimination of software errors at design-time.

The idea behind error control of software is to use techniques from feedback control in order to detect malfunctioning software components and, in that case fall back on, a well-tested core software component that is able to provide the basic application service with guarantees on performance and safety. Hence, the basic idea assumes that a certain amount of defect-free components are available, that can be used to implement the fall-back safety core service. The second key idea is to always design your system to have a simple and well-formed dependency tree, with a minimal number of dependency relations among components. This is necessary in order to be able to identify the core services and keep them small. The background to several of the key ideas of the area is given in [Sha, 2001].

In software error control, our view of what control is has to be broadened substantially. Control is normally concerned with the temporal behaviour of systems. The ideas behind software error control are, however, not restricted

to the temporal behaviour. The same approach can in principle also be used for applications that only contain functional requirements. In this case software error control has strong relationships to techniques that are commonly associated with fault tolerance, e.g., hardware and software redundancy and diversity through replication and N-version programming. However, the principles behind software error control have so far mainly been applied to reactive applications, i.e., avionics control systems.

5.1 Research Challenges

The major challenge is to develop a new paradigm for software stability control, based on an integration of concepts from fault-tolerant computing and control, that is applicable to a wide range of application types. The number of documented examples where software error control has been applied is small. In order to increase the understanding for the subject and to develop the necessary methods and theory, more documented applications must be developed. The relationships to the methods within the traditional fault tolerance area must also be clarified. It is further necessary to investigate for which application types, other than feedback control, the approach is suitable.

6. Acknowledgements

This work has been funded by the EU/IST FP6 NoE ARTIST2. We gratefully acknowledge the valuable comments of the reviewers.

7. References

- Aparah, D. (1998): "Adaptive resource management in a multimedia operating system." In *Proceedings of the 8th International Workshop on Network and Operating System Support for Digital Audio and Video*.
- Buttazzo, G., G. Lipari, and L. Abeni (1998): "Elastic task model for adaptive rate control." In *Proc. 19th IEEE Real-Time Systems Symposium*, pp. 286–295.
- Chen, L., S. Low, and J. Doyle (2005): "Joint congestion control and media access control design for ad hoc wireless networks." In *Proc. of IEEE Infocom*. IEEE, Miami, FL.
- Chiang, M. (2005): "Balancing transport and physical layers in wireless multihop networks: Jointly optimal congestion control and power control." *IEEE JSAC*, **23:1**, pp. 104–116.
- Chu, H. and K. Nahrstedt (1999): "CPU service classes for multimedia applications." In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*.
- Hellerstein, J., Y. Diao, S. Parekh, and D. Tilbury (2005): "Control engineering for computing systems." *IEEE Control Systems Magazine*, **25:6**, pp. 56–68.
- Hellerstein, J. L., Y. Diao, S. Parekh, and D. M. Tilbury (2004): *Feedback Control of Computing Systems*. John Wiley.
- Henriksson, D., Y. Lu, and T. Abdelzaher (2004): "Improved prediction for web server delay control." In *Proceedings of the 16th Euromicro Conference on Real-Time Systems (ECRTS 04)*. Catania, Sicily, Italy.
- Holmberg, K. (1995): *Design Models for Hierarchical Organizations: Computation, Information, and Decentralization*, chapter Primal and dual decomposition as organizational design: Price and/or resource directive decomposition, pp. 61–92. Kluwer Academic Publishers.
- Johansson, B. and M. Johansson (2006): "Mathematical decomposition techniques for distributed cross-layer optimization of data networks." *IEEE Journal on Selected Areas in Communications*, November. (to appear).
- Kelly, F., A. Maulloo, and D. Tan (1998): "Rate control in communication networks: shadow prices, proportional fairness and stability." *Journal of the Operational Research Society*, **49**, pp. 237–252.
- Kleinrock, L. (1975): *Theory, Volume 1, Queuing Systems*. Wiley-Interscience.
- Kwadia, V. and P. R. Kumar (2005): "A cautionary perspective on cross layer design." *W. IEEE Wireless Communication*, **12:1**, pp. 3–11.
- Lin, X. and N. B. Shroff (2004): "Joint rate control and scheduling in multihop wireless networks." In *Proc. of IEEE Conference on Decision and Control*. IEEE, Paradise Island, Bahamas.
- Low, S. H. and D. E. Lapsley (1999): "Optimization flow control – I: Basic algorithm and convergence." *IEEE/ACM Transactions on Networking*, **7:6**, pp. 861–874.
- Robertsson, A., B. Wittenmark, and M. Kihl (2003): "Analysis and design of admission control in web-server systems." In *Proceedings of the 2003 American Control Conference (ACC'03)*, pp. 254–259. Denver, Colorado, USA.
- Sha, L. (2001): "Using simplicity to control complexity." *IEEE Software*, **18:4**.
- Sha, L., X. Liu, Y. Lu, and T. Abdelzaher (2002): "Queuing model based network server performance control." In *IEEE Real-Time Systems Symposium*.
- Stankovic, J. A., C. Lu, S. H. Son, and G. Tao (1999): "The case for feedback control real-time scheduling." In *Proc. 11th Euromicro Conference on Real-Time Systems*, pp. 11–20.
- Tipper, D. and M. K. Sundareshan (1990): "Numerical models for modeling computer networks under nonstationary conditions." *IEEE Journal on Selected Areas in Communications*, **8:9**, pp. 1682–1695.
- Wen, J. T. and M. Arcak (2004): "A unifying passivity framework for network flow control." *IEEE Transactions on Automatic Control*, **49:2**, pp. 162–174.
- Yuan, W. and K. Nahrstedt (2003): "Energy-efficient soft real-time CPU scheduling for mobile multimedia systems." In *19th ACM Symposium on Operating Systems Principles*. Bolton Landing, NY.