How to review scientific papers

Don Porter

Portions courtesy Samarjit Chakraborty

Why do we write papers?

- Dissemination of research results to the scientific community
- Get credits for our work (required for graduating, getting a job, promotion, tenure, etc.)

06733298.pdf (page 1 of 9)

IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY, VOL. 22, NO. 6, NOVEMBER 2014 2337

Relaxing Signal Delay Constraints in Distributed Embedded Controllers

Dip Goswami, Reinhard Schneider, and Samarjit Chakraborty

***Abstract*— Embedded systems often involve transmitting feedback signals between multiple control tasks that are implemented on different electronic control units communicating via a shared bus. For ensuring stability and control performance, such designs require all control signals to be delivered within a specified deadline, which is ensured through appropriate timing or schedulability analysis. In this brief, we study controller design that allows control feedback signals to occasionally miss their deadlines. In particular, we provide analytical bounds on deadline misses such that the control loop retains its stability and meets its control performance requirements. We argue that such relaxation allows us to 1) use lower quality communication resources (e.g., event-triggered instead of time-triggered communication) and 2) provide more flexibility-e.g., use simulation-in communication timing analysis since analytical worst-case delay bounds for real-life communication protocols are often pessimistic. We illustrate this approach using the FlexRay communication protocol for distributed automotive control systems.**

***Index Terms*— Distributed embedded controllers, flexible delay constraints, FlexRay, timing analysis.**

Fig. 1. Distributed embedded controller.

I. Introduction

THE DESIGN of distributed control systems-where control tasks are implemented on different electronic control units (ECUs) communicating via a shared bus-typically require end-to-end timing guarantees from the embedded implementation platform. For example, in Fig. 1, the controller may be designed with the assumption of a specified maximum sensor-to-actuator delay, which includes the computation times of the software tasks and the signal transmission delays on the communication bus. Such delay assumptions are guaranteed using appropriate scheduling policies along with necessary timing or schedulability analysis. For many real-life communication protocols like CAN or FlexRay, providing tight analytical timing bounds is often difficult and leads to either pessimistic results or resource overprovisioning [2]. In many cases, this problem is circumvented using resource reservation techniques like time-triggered protocols that provide better timing guarantees and are easier to analyze. However, they are also usually more conservative and lead to poor resource utilization.

In contrast to this constraint on all control signals having to meet their deadlines, in this brief, we propose controller design techniques that allow control signals to occassionally miss their deadlines (or be dropped). Our main technical contribution is to provide an analytical bound on such deadline misses such that stability and control performance are nevertheless guaranteed. We believe that such relaxation may be exploited in the implementation platform design and analysis stages. For example, instead of using time-triggered communication that ensures all control signals are delivered in time, priority-based communication protocols may be used with certain signals missing their deadlines. Furthermore, bounds on deadline misses from our controller design stage may be used as an additional safety margin, thereby allowing less precise and hence less pessimistic, e.g., simulation based-timing analysis techniques for the implementation platform. It is worth mentioning that recently there has also been work on verification of communication schedules, which ensure that bounds on allowable control signal deadline misses are satisfied by the implementation platform [1], [8], [9], [16]. These results may be coupled with our proposed design to provide certifiable implementations where needed.

The issue of incorporating feedback signal delay into controller design has been widely studied by the networked control systems community [10]–[12]. One of the seminal results on bounding signal deadline misses while guaranteeing stability, along the lines of our work, may be found in [17]. Since then a number of other papers have reported various special cases of this result [7], [13], [15]. These results mostly provide bounds over an infinite horizon of samples (i.e., address asymptotic behavior), which are usually difficult to check or formally verify for an implementation platform. in addition, all such results are concerned with stability and not performance, which is important in real-life settings. We attempt to overcome these shortcomings; our proposed design method has better applicability, which we illustrate using a FlexRay based [5] distributed controller design example from the automotive domain.

Manuscript received May 23, 2013; revised December 5, 2013; accepted January 16, 2014. Manuscript received in final form January 19, 2014. Date of publication February 6, 2014; date of current version October 15, 2014. Recommended by Associate Editor A. Horch.

D. Goswami is with the Technische Universiteit Eindhoven, Eindhoven 5612, The Netherlands (e-mail: d.goswami@tue.nl).

R. Schneider and S. Chakraborty are with the Technische Universitaet München, Munich 85748, Germany (e-mail: reinhard.schneider@rcs.ei.tum.de; samarjit@tum.de).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCST.2014.2301795

What is a review?

- Vetting mechanism to judge the soundness and value of a paper
- Done by a peer group (a set of established experts in that area)
- The decision to “accept” or “reject” the paper depends on the reviews

A typical reviewer

- Time restrictions - not ideal, but this is reality
- Voluntary service to the scientific community.
- If there are 4 reviews per paper then you should review 4x the number of papers you write, for the system to be sustainable

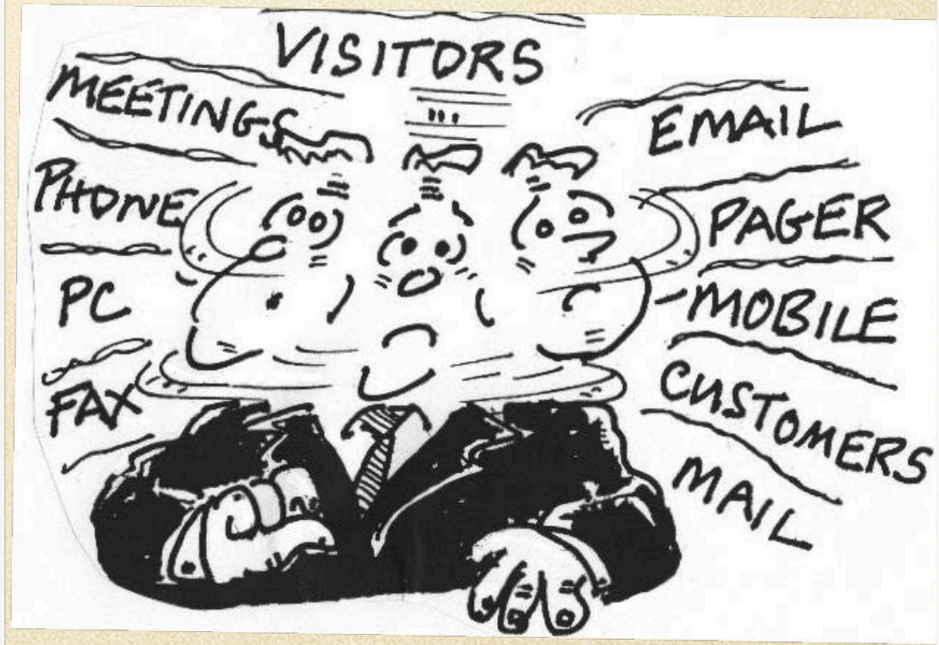

Reviewing Mechanisms

- Conferences:
  - Typically a one-time program committee (PC)
  - May meet in person (great networking)
    - Shifting to zoom ($$ + CO2 footprint)
  - Good conferences typically accept <20-30% of submissions
    - Most of what you read will be worse than what you get in a grad seminar
- Journals and books:
  - Typically a smaller editorial board
  - Solicits one-off, expert reviews
- Grants: standing or one-time panels, depending on agency

Anonymity

- Single-Blind Reviewing: Author names visible, reviewer names hidden (why?)
- Double-Blind: Author and reviewer names hidden
- Reviewer names usually visible to other reviewers (why/why not?)

Confidentiality

- Material under review is confidential
- **NEVER:**
  - Share or redistribute a draft you review
  - Borrow ideas from a submission
- As a reviewer, you should also expect your identity to be held in confidence

Why do reviews? and do them well?

- Trade volunteer work for credibility
  - PC lists are the “cool kids club”
- Social networking among PC members
  - Waaaaay more face time with leaders in field than at a conference
  - I get the most insight into how someone thinks from reading their reviews
- Learn things!
- Get better at writing papers
  (i.e., better understand your audience - sound familiar?)

Start Practicing Now!

- Reviewing is a key professional meta-skill
  - Takes practice and feedback
- Your advisor likely reviews papers - help them!
  - I often invite my students to “shadow review”

Reviews Have 2 Goals

- Your review needs to BOTH:
  - Explain to the PC why the paper should be accepted or rejected
  - Explain to the authors how to improve the paper

    (ideally, to the level you would argue for acceptance, or more strongly if you are already positive; occasionally, an idea is truly unfixable)

Giving Editorial Advice

- If at all possible, try to give a clear accept or reject rating
- And clearly explain why!
  - Most conferences use a 5 point scale
    - 1 = Reject, 5 = Strong Accept
  - Average paper review a 2.5, with variance .2 (made-up, “truthy” numbers). Why?
  - Do such reviews contribute “signal”?
    - “Strong Accept” == triple cheeseburger
- “Confidential Comments” a good place to be blunt/clarify

Author Feedback

- All papers have flaws
- Follow the Golden Rule:
  Write a review you would like to receive
  - Be constructive - actionable suggestions >> complaints
- For papers you want to reject:
  What would it take to convince you?
- Depersonalize criticism: “The paper” vs. “The authors”
- Don’t ascribe motive to flaws (e.g., lazy authors)

Author Feedback (2)

- Justify your criticisms
  - Saying something is not novel **requires** a citation
- Gut check: Can all of your requests fit in the space limit?
- Is it clear what feedback is the most important and what are nits?

Author Responses

- If authors have a response period:
  - Explicitly list the “pivot questions” for the work, if possible (why?)

Tricky Issues

- What if the paper is out of your area of expertise?
  - Always rate your expertise
  - As a representative member of community, you can still assess general interest level and clarity
  - Editor/Chair needs to know if all reviews are low confidence (why?)
- What if you don’t understand the paper at all?
  - Ok to say so
  - But consider reasons: writing vs your background?
    - In principle, you can read more related work

Avoiding Cringe

- What if a relevant paper doesn’t cite your work?
- Should you review a paper on a topic you are working on?
- What if you suspect a paper is written by a friend (or someone you find challenging)?