On the Scalability of Cooperative Time Synchronization in Pulse-Connected Networks

Sergio D. Servetto

School of Electrical and Computer Engineering — Cornell University
http://cn.ece.cornell.edu/

(Joint work with An-swol Hu.)

Acknowledgements

Graduate students:

- Ron Dabora, *An-swol Hu*, Georgios N. Lilis, Mingbo Zhao (Cornell/ECE).

Sources of support:

- *Fundamental Performance Limits of Large-Scale Sensor Networks*.
  NSF CAREER award CCR-0238271.
- *The Reachback Channel in Wireless Sensor Networks*.
  NSF SENSORS grant CCR-0330059. Joint with T. Berger, L. Tong, S. Wicker.
- *Self-Configuring Sensor Networks for Disaster Prevention, Mitigation and Recovery*. NSF ITR grant ANR-0325556. Joint with Cornell ECE, CEE and Economics faculty, and staff at NYS Wadsworth Labs.

Outline

- Limits of Sums of Randomly Shifted Pulses.
- Cooperative Time Synchronization and Distributed Modulation.
- Summary and Conclusions.

Outline

- Limits of Sums of Randomly Shifted Pulses.
- Cooperative Time Synchronization and Distributed Modulation.
- Summary and Conclusions.

Limits of Sums of Randomly Shifted Pulses

- Start with a concrete example: all nodes want to communicate a common message over a single multiple access channel.
- The signals at the output of an asynchronous, analog, unfaded, power constrained, bandwidth unconstrained Gaussian MAC have the form:

$$R_N(t) = A_N(t) + n(t) = c_N \sum_{i=1}^{N} p(t - \psi_i) + n(t),$$

where c_N is a constant, $[\psi_1 ... \psi_N]$ is a vector of random delays, and $n(t)$ is Gaussian noise.

Any communication task involves computing some $f(R_N(t))$... so, what can we say about its statistics, for different choices of f?

Limits of Sums of Randomly Shifted Pulses

Example: take a *linear* functional $f : \mathcal{L}^2(\mathbb{R}) \to \mathbb{R}$ (e.g., a matched filter), $[\psi_1 ... \psi_N]$ an *independent* random vector, and N large. Then:

$$f(R_\infty(t)) \quad = \quad \lim_{N\to\infty} c_N \sum_{i=1}^{N} \underbrace{f(p(t-\psi_i))}_{\text{all independent}} + f(n(t)) \quad = \quad \ell_\infty \sim \mathcal{N}(\mu, \sigma^2),$$

by the CLT. How can we detect a signal based on these observations?

- A hypothesis test for different variances $\Rightarrow$ *Crude, low performance.*
- Estimate delays ψ_i as in multipath $\Rightarrow$ *Not feasible with N large.*
- Optimal multiuser detection $\Rightarrow$ *NP-Hard, intractable with N large.*

With *correlated* delays however, we could control the distribution of ℓ_∞...

Limits of Sums of Randomly Shifted Pulses

Our take on this problem:

> *Define a distributed estimation problem (each node computes an estimate of its own ψ_i — correlations due to all nodes estimating the same thing), so that the shape of $A_\infty(t)$ can be controlled.*

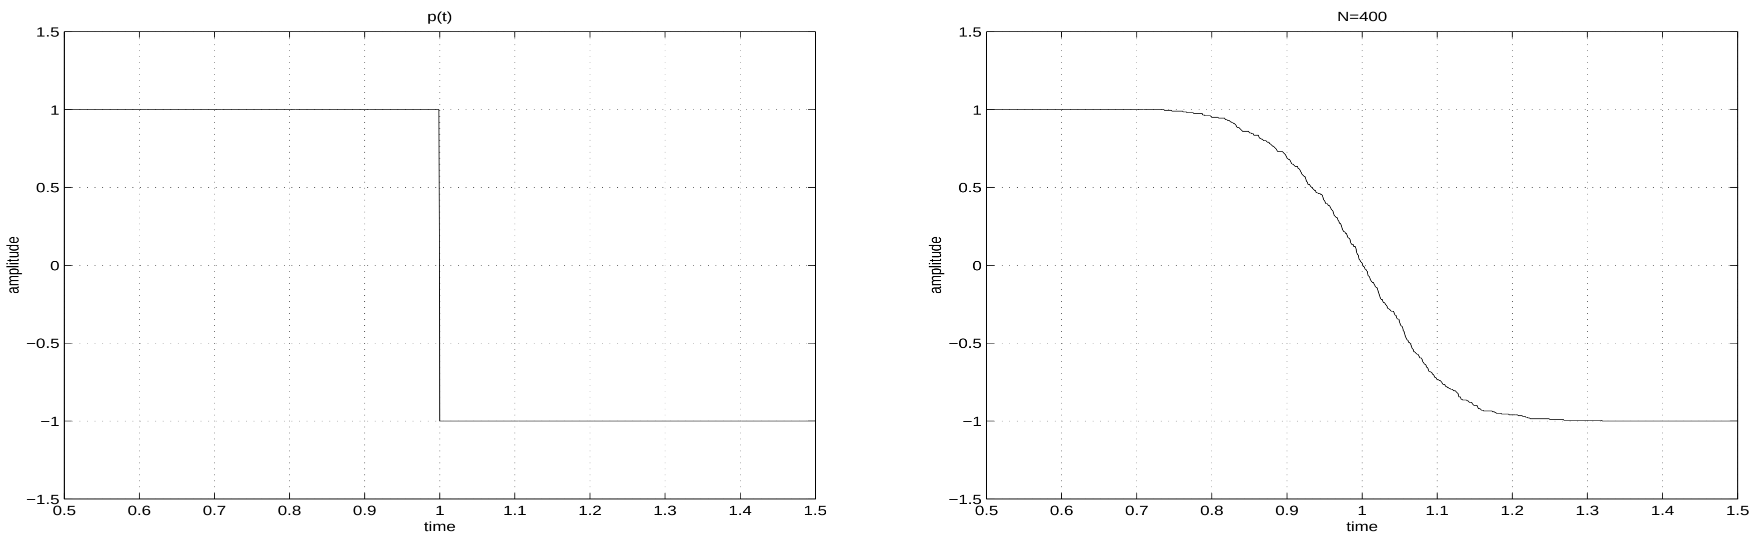

Limits of Sums of Randomly Shifted Pulses

- If p is continuous (except for at most a countable but not dense number of points), and bounded, then $A_\infty(t)$ is continuous.
- Complete characterization of the roots $\zeta = \{t : A_\infty(t) = 0\}$.
- Exact characterizations of the sets $\{t : A_\infty(t) > 0\}$, $\{t : A_\infty(t) < 0\}$.

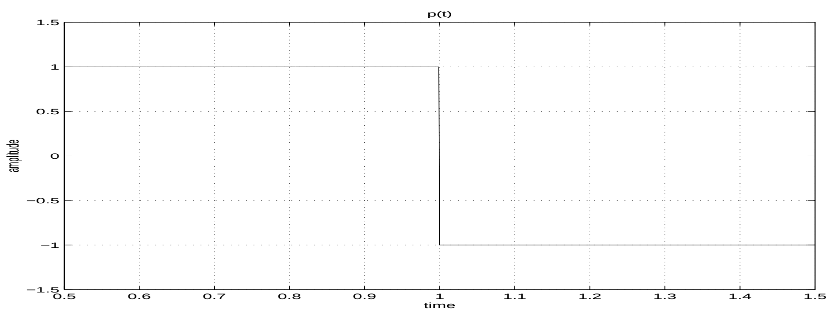

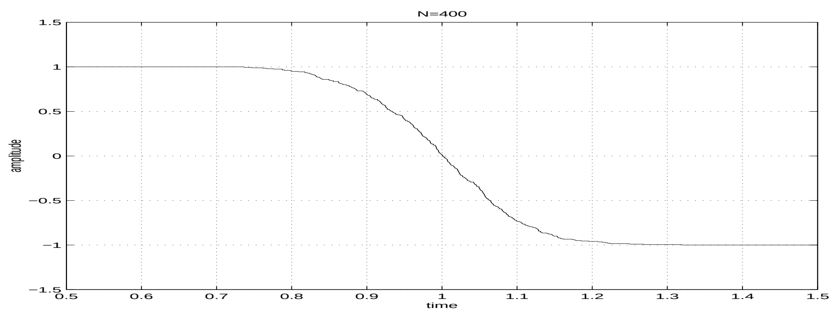

A. Hu, S. D. Servetto. *Algorithmic Aspects of the Time Synchronization Problem in Dense Sensor Networks*. In ACM/Kluwer Journal on Mobile Networks and Applications (MONET), 10:491-503, September 2005. Special Issue on Wireless Sensor Networks, with selected (and revised) papers from ACM WSNA 2003. *Invited Paper.*

A. Hu, S. D. Servetto. *On the Scalability of Cooperative Time Synchronization in Pulse-Connected Networks*. IEEE Trans. Inform. Theory; to appear.

Main (Simple) Elements of the Proof

- The Law of Large Numbers: for some random process X_t, for N large,

$$c_N \sum_{i=1}^{N} f(p(t - \psi_i)) \approx E(X_t).$$

- Use the shapes of f and p to provide bounds (pointwise, for fixed t), showing where $E(X_t) > 0$, $= 0$, < 0.
- Show that for "well behaved" statistics of the estimation error and for "not too discontinuous" pulses, the distributions of X_t and of $X_{t+\epsilon}$ are close.
- Use the fact that integrals are continuous functions of their arguments.

Long, tedious, but conceptually simple proofs. What is interesting is not the proofs in themselves, but what we can build on top of these results...

Outline

- Limits of Sums of Randomly Shifted Pulses.
- Cooperative Time Synchronization and Distributed Modulation.
- Summary and Conclusions.

Cooperative Time Synchronization: Problem Setup

- Random placement of N nodes on a unit square, each equipped with a local oscillator $c_i(t)$.
- An arbitrary node ("node 1") has an operational counter $s_1(t)$, that increments at integer values of $c_1(t)$ (tick at each period).
- *Goal: define $s_i(t)$ at each node such that they all tick at the same time.*

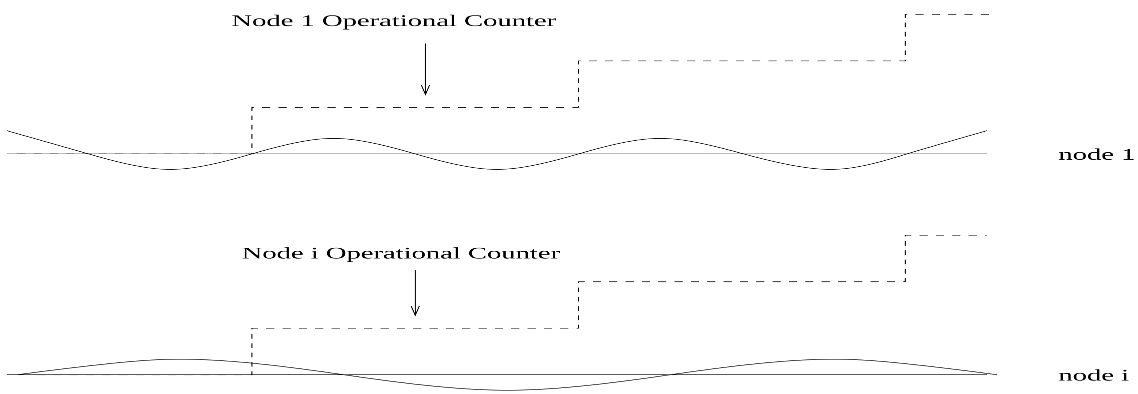

Cooperative Time Synchronization: Node Mechanics

```
TimeSync (observation length m > 1)

observe pulse arrival time;
if (first observed pulse)
  { receive packet and set operation counter; };
while (pulse arrived)
  { if (m or more arrival times in memory)
      { keep only m most recent and discard
          all other arrival times;
        use last m arrivals to estimate next
          pulse time t_p;
        transmit scaled pulse p(t) at time t_p;
      };
    observe pulse arrival time;
  };
```

A Simple Idea for Cooperative Time Synchronization

For a period $T > 0$, define estimates $[\psi_1 ... \psi_N]$ such that

$$\zeta \;=\; \{t : A_\infty(t) = 0\} \;=\; \{nT : n \in \mathbb{Z}\}.$$

Then ζ provides a clock that can be heard simultaneously by all nodes.

A. Hu, S. D. Servetto. *Asymptotically Optimal Time Synchronization in Dense Sensor Networks*. In Proc. ACM WSNA, 2003.

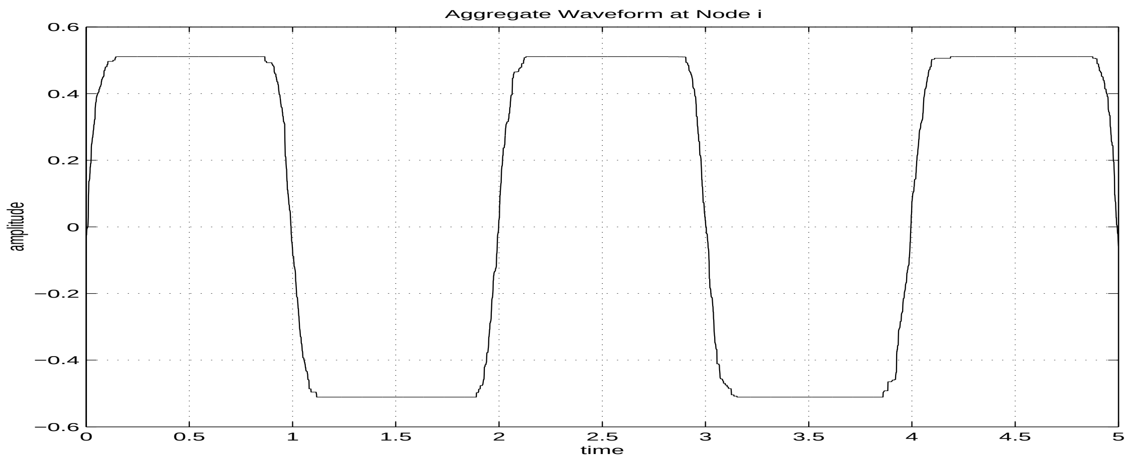

Asymptotically Optimal Time Synchronization

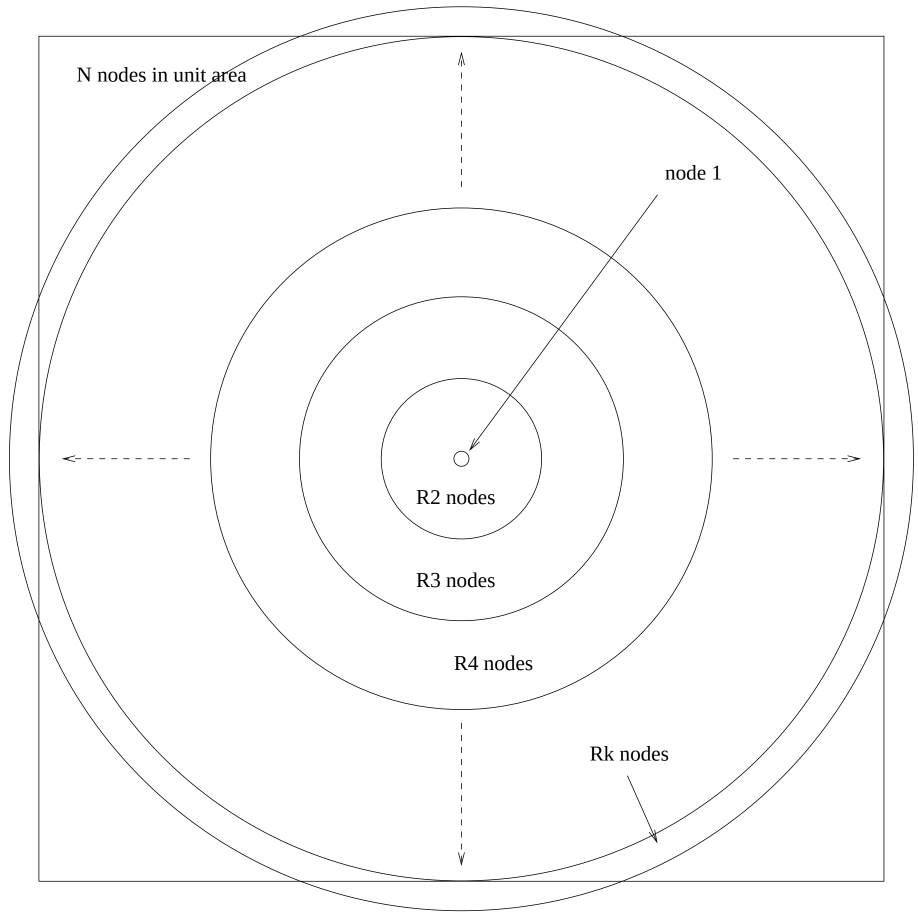

A Simple Idea for Distributed Modulation

Distributed carrier waves: given a sequence of bits $(b_n)_{n\in\mathbb{Z}}$ (known at all nodes), a symbol time T, and a repetition factor $\kappa \in \mathbb{N}^+$, make

$$\zeta \;=\; \{t : A_\infty(t) = 0\} \;=\; \bigcup_{n\in\mathbb{Z}} \begin{cases} \{nT\}, & \text{if } b_n = 0 \\ \{(n+\frac{k}{\kappa})T : k = 0\ldots\kappa-1\}, & \text{if } b_n = 1. \end{cases}$$

Then ζ provides an encoding of the bits $(b_n)_{n\in\mathbb{Z}}$ to a far receiver.

A. Hu, S. D. Servetto. *dFSK: Distributed FSK Modulation in Dense Sensor Networks*. In Proc. IEEE ICC, 2004.

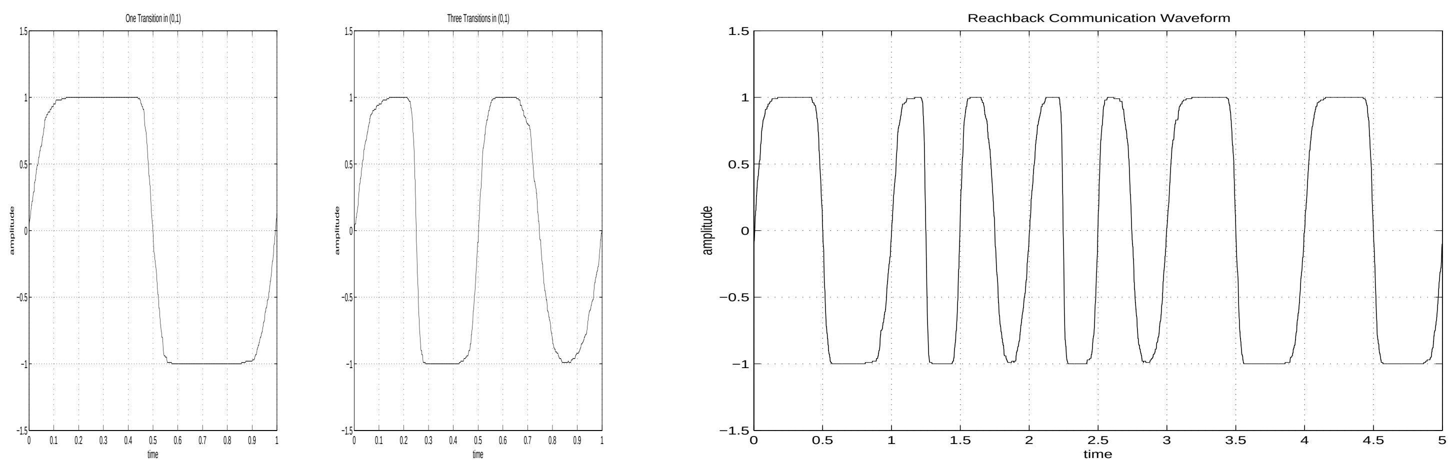

Outline

- Limits of Sums of Randomly Shifted Pulses.
- Cooperative Time Synchronization and Distributed Modulation.
- Summary and Conclusions.

Summary

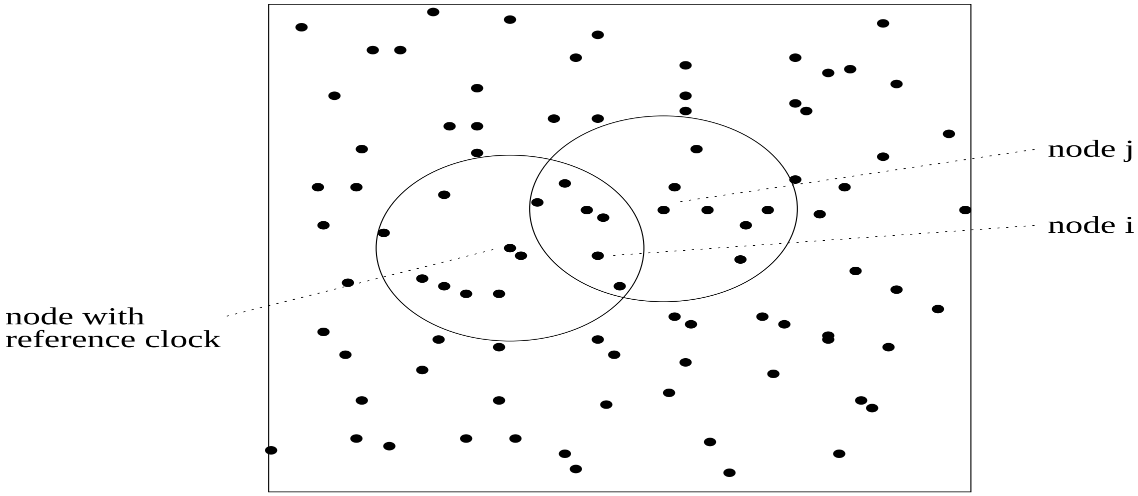

- Synchronization computing only time averages does *not* scale well.
- "Clean up data", by computing averages over time *and* space.
- New proof methods for convergence to synchrony, purely probabilistic, containing *most* classical Mirollo-Strogatz results as special case.

R. Mirollo, S. Strogatz. *Synchronization of Pulse-Coupled Biological Oscillators*. SIAM J. Appl. Math., 50(6):1645-1662, 1990.

Current/Future Work

- Analysis for large but finite networks – large deviation methods.
- Protocol implementation and debugging (on an *acoustic* array).
- Synchronization without an orderly startup phase, to eliminate complexities derived from having to deal with single points of failure.
- Implementation on asynchronous low-power circuits.
- Formal derivation and performance analysis for detectors.
- Bandpass description, plus determination of impact of timesync errors on performance as the frequency of the distributed carrier increases.

Conclusions

Two observations:

- At its origin, computing was done on big mainframes. But mainframes lost the “cost/performance” evolution race to smaller networked PCs.
- Radio design principles today are still conceptually similar to those of a mainframe: improve performance by adding complexity to a centralized system – no such thing as a “massively distributed radio” exists (yet...).

Main conclusion (hope?):

Massively distributed communication systems seem within reach!

Main Corollary...

http://cn.ece.cornell.edu/