

The Emergence of a Networking Primitive in Wireless Sensor Networks

By Philip Levis, Eric Brewer, David Culler, David Gay, Sam Madden, Neil Patel, Joe Polastre, Scott Shenker, Robert Szewczyk, and Alec Woo

Abstract

The wireless sensor network community approached networking abstractions as an open question, allowing answers to emerge with time and experience. The Trickle algorithm has become a basic mechanism used in numerous protocols and systems. Trickle brings nodes to eventual consistency quickly and efficiently while remaining remarkably robust to variations in network density, topology, and dynamics. Instead of flooding a network with packets, Trickle uses a “polite gossip” policy to control send rates so each node hears just enough packets to stay consistent. This simple mechanism enables Trickle to scale to 1000-fold changes in network density, reach consistency in seconds, and require only a few bytes of state yet impose a maintenance cost of a few sends an hour. Originally designed for disseminating new code, experience has shown Trickle to have much broader applicability, including route maintenance and neighbor discovery. This paper provides an overview of the research challenges wireless sensor networks face, describes the Trickle algorithm, and outlines several ways it is used today.

1. WIRELESS SENSOR NETWORKS

Although embedded sensing applications are extremely diverse, ranging from habitat and structural monitoring to vehicle tracking and shooter localization, the software and hardware architectures used by these systems are surprisingly similar. The typical architecture is embodied by the mote platforms, such as those shown in Figure 1. A microcontroller provides processing, program ROM, and data RAM, as well as analog-to-digital converters for sensor inputs, digital interfaces for connecting to other devices, and control outputs. Additional flash storage holds program images and data logs. A low-power CMOS radio provides a simple link layer. Support circuitry allows the system to enter a low-power sleep state, wake quickly, and respond to important events.

Four fundamental constraints shape wireless embedded system and network design: power supply, limited memory, the need for unattended operation, and the lossy and transient behavior of wireless communication. A typical power envelope for operating on batteries or harvesting requires a 600 μ W average power draw, with 1% of the time spent in a 60 mW active state and the remainder spent in a very low power 6 μ W passive state.

Maintaining a small memory footprint is a major requirement of algorithm design. Memory in low-cost, ultra-low-power devices does not track Moore’s Law. One indication of this is that microcontroller RAM costs three orders of magnitude more than PC SRAM and five orders more than PC DRAM. More importantly, SRAM leakage current, which grows with capacity, dictates overall standby power consumption and, hence, lifetime. Designs that provide large RAMs in conjunction with 32-bit processors go to great lengths to manage power. One concrete example of such nodes is the Sun SPOT,²⁰ which enters a low-power sleep state by writing RAM contents to flash. Restoring memory from flash on wakeup uses substantial power and takes considerable time. The alternative, taken in most sensor node designs, is to have just a few kilobytes of RAM. This, in turn, imposes limits on the storage complexity of network (and other) protocols, requiring routing tables, buffering, and caches be kept small. The historical trends of monetary and energy costs suggest these constraints are likely to last.

Wireless sensors are typically embedded in the physical environment associated with their application. Com-

Figure 1: EPIC, KMote, and Telos motes. Each has an 8MHz microcontroller, 10kB of RAM, 48kB of program flash, and a 250kbps radio.



munication connectivity varies due to environmental and electromagnetic factors, with the additional constraint that no human being will shepherd the device to a better setting, as with a cell phone or a laptop. The degree of the network at a node, i.e., the number of nodes in its communication neighborhood, is determined not by the desired network organization but by the physical device placement, which is often dictated by application requirements and physical constraints. There may be thousands of nodes in close proximity, or just a few. A single transmission may be received by many devices, so any retransmission, response, or even a simple acknowledgment, may cause huge contention, interference, and loss. Redundancy is essential for reliability, but it also can be a primary cause of loss.

This last point is one of the key observations that have emerged from a decade of development of networking abstractions for wireless sensor networks: the variety of network topologies and densities across which sensor network protocols must operate calls for a polite, density-aware, local retransmission scheme. This paper describes the Trickle algorithm, which uses such a communication pattern to provide an eventual consistency mechanism to protocols and services. In the past ten years, a key insight that has emerged from the wireless sensor network community is that many protocol problems can be reduced to maintaining eventual consistency. Correspondingly, Trickle has emerged as the core networking primitive at the heart of practical, efficient, and robust implementations of many sensor network protocols and systems. Before diving into the details of the Trickle, however, we review how core sensor networking protocols work and differ from conventional networking protocols, with the goal of exploring how a Trickle-like primitive satisfies some of their needs.

2. NETWORKING PROTOCOLS

Networking issues are at the core of embedded sensor network design because radio communication—listening, receiving, and transmitting—dominates the active energy budget and defines system lifetime. The standard energy cost metric for multihop protocols, in either link layer meshing or network layer routing, is communication cost, defined as the number of individual radio transmissions and receptions. One protocol is more efficient than another if it can provide equivalent performance (e.g., throughput, latency, delivery ratio) at a lower communication cost. Protocols focus on minimizing transmissions and making sure transmitted packets arrive successfully.

Almost all sensor network systems rely on two multihop protocols for their basic operation: a *collection* protocol for pulling data out of a network and a *dissemination* protocol for pushing data into a network through one or more distinguished nodes or egress routers. Many higher level protocols build on dissemination and collection. For example, reprogramming services such as Deluge⁹ use dissemination to deliver commands to change program images. Management layers²² and remote source-level debuggers²⁵ also use dissemination. Reliable transport protocols, such as RCRT,¹⁸ and rate control protocols such as

IFRC,¹⁹ operate on collection trees. Point-to-point routing schemes, such as S4,¹⁶ establish overlays over multiple parallel collection topologies.

While collection and dissemination have the opposite communication patterns (all-to-one vs. one-to-all) and differ in reliability (unreliable vs. reliable), both maintain eventually consistent shared state between nodes. The rest of this section provides a high-level overview of these two protocol classes. It provides details on the challenging problems they introduce, and how some of them can be solved through eventual consistency.

2.1. Pushing data in: dissemination

One problem sensor network administrators face is dynamically changing how a network collects data by changing the sampled sensors, the sampling rate, or even the code running on the nodes by disseminating the change to every node in a network. We begin with a discussion of dissemination protocols because they were the original impetus for Trickle and are its simplest application.

Early systems used packet floods to disseminate changes. Flooding protocols rebroadcast packets they receive. Flooding is very simple—often just a line or two of code—but has many problems. First, floods are unreliable. Inevitably, some nodes do not receive the packet, so users typically repeatedly flood until every node receives it. Second, in high density networks, many nodes end up rebroadcasting packets at the same time. These messages collide and cause a form of network collapse called a “broadcast storm.”¹⁷

Second-generation dissemination and network programming systems like Xnp³ and TinyDB¹⁵ use an adaptive flood combined with a protocol to request missing messages. Adaptive flooding uses an estimate of the node density to limit the flooding rate. The missing message protocol allows nodes to request the (hopefully few) missing messages from their neighbors. Unfortunately, getting such protocols to work well can be tricky, especially across a range of network densities and object sizes.

Another way to look at dissemination protocols is that they ensure that every node has an eventually consistent version of some shared state, such as the value of a configuration parameter or command. Data consistency is when all nodes have the same version of that state, and nodes resolve inconsistencies by updating neighbors to the newer version. Inductively, these definitions cause the network to converge on the most recent version. To disseminate a command, a system installs it on one node as a newer version and initiates the consistency protocol.

Casting dissemination as a data consistency problem means it does not provide full reliability. Eventual consistency only promises to deliver the most recent version to connected nodes. Disconnected nodes can and often do miss updates. In practice, however, this limitation is rarely problematic. An administrator who changes the data reporting rate three times then adds some new nodes and expects them to receive the most recent reporting rate, not all three. Similarly, when sending commands, users do not expect a new node to receive the entire history of all commands injected into a net-

work. A node that is disconnected for several minutes will still receive the most recent command when it reconnects, however.

Dissemination protocols succeed where flooding and its derivatives fail because they cast the problem of delivering data into maintaining data consistency among neighbors. This allows them to provide a very useful form of reliability in arbitrary topologies with no a priori topology knowledge or configuration. An effective dissemination protocol, however, needs to bring nodes up to date quickly while sending few packets when every node has the most recent version: this is correspondingly a requirement for the underlying consistency mechanism.

2.2. Pulling data out: collection

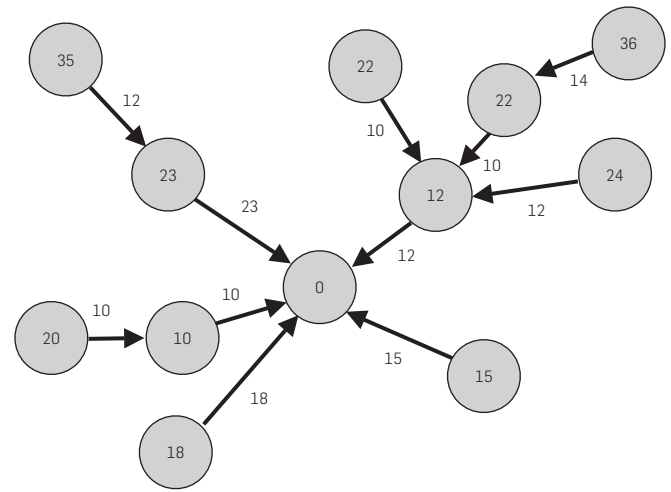
As the typical sensor network goal is to report observations on a remote environment, it is not surprising that data collection is the earliest and most studied class of protocol. There are many collection protocol variations, similar to how there are many versions of TCP. These differences aside, all commonly used collection protocols provide unreliable datagram delivery to a collection point using a minimum-cost routing tree. Following the general goal of layer 3 protocols, cost is typically measured in terms of expected transmissions, or ETX;² nodes send packets on the route that requires the fewest transmissions to reach a collection point.

The earliest collection protocol, directed diffusion, proposed dynamically setting up collection trees based on data-specific node requests.¹⁰ Early experiences with low-power wireless, however, led many deployments to move towards a much simpler and less general approach, where each node decides on a single next hop for all forwarded data traffic, thereby creating routing trees to fixed collection points. The network builds this tree by establishing a routing cost gradient. A collection point has a cost of 0. A node calculates the cost of each of its candidate next hops as the cost of that node plus the cost of the link to it. Inductively, a node's cost is the sum of the costs of the links in its route. Figure 2 illustrates an example topology.

Collection variations boil down to how they quantify and calculate link costs, the number of links they maintain, how they propagate changes in link state amongst nodes, and how frequently they re-evaluate link costs and switch parents. Early protocols used hop-counts⁸ as a link cost metric, similar to MANET protocols such as AODV and DSDV; second-generation protocols such as MintRoute²⁴ and Srcr² estimated the transmissions per delivery on a link using periodic broadcasts; third-generation protocols, such as MultiHopLQI, added physical layer signal quality to the metric; current generation collection protocols, such as Collection Tree Protocol (CTP), unify these approaches, drawing on information from multiple layers.⁶

Most collection layers operate as anycast protocols. A network can have multiple data collection points, and collection automatically routes to the closest one. As there is only one destination—any collection point—the required routing state can be independent of network density and size. Most protocols use a small, fixed-size table of candidate next

Figure 2: Sample collection tree, showing per-link and node costs. The cost of a node is its next hop's cost plus the cost of the link.



hops. They also attempt to strike a balance between route stability and churn to discover new, possibly better parents by switching parents infrequently and using damping mechanisms to limit the rate of change.

As collection protocols have improved and become better at choosing routes, reducing control traffic has become an increasingly important component of efficiency. While nodes can piggyback some control information on data packets, they need to send link-layer broadcasts to their local neighbors to advertise their presence and routing cost. Choosing how often to send these advertisements introduces a difficult design tension. A slow rate imposes a low overhead, but limits how quickly the tree can adapt to failures or link changes, making its data traffic less efficient. A fast rate imposes a higher overhead, but leads to an agile tree that can more accurately find the best route to use.

This tension is especially challenging when a network only collects data in response to events, and so can go through periods of high and low data rates. Having a high control rate during periods of low traffic is highly inefficient, while having a low control rate during periods of high traffic makes the tree unable to react quickly enough to changes. When starting a burst of transmissions, a node may find that link costs have changed substantially necessitating a change in its route and, as a result, advertised routing cost. Changes in costs need to propagate quickly, or the topology can easily form routing loops. For example, if a link's cost increases significantly, then a node may choose one of its children as its next hop. Since the protocol state must be independent of the topology, a node cannot avoid this by simply enumerating its children (constraining tree in-degree to a constant leads to inefficient, circuitous topologies in dense networks).

Current protocols, such as CTP²¹ and ArchRock's routing layer,¹ resolve this tension by reducing the routing gradient as a data consistency problem. The gradient is consistent as long as children have a higher cost than their parent. An inconsistency can arise when costs change enough to violate

this constraint. As long as routing costs are stable, nodes can assume the gradient is consistent and avoid exchanging unnecessary packets.

2.3. A general mechanism

The examples above described how two very different protocols can both address a design tension by reducing a problem to maintaining data consistency. Both examples place the same requirements on a data consistency mechanism: it needs to resolve inconsistencies quickly, send few packets when data is consistent, and require very little state. The Trickle algorithm, discussed in the next section, meets these three requirements.

3. TRICKLE

The Trickle algorithm establishes a density-aware local broadcast with an underlying consistency model that guides when a node communicates. When a node's data does not agree with its neighbors, it communicates quickly to resolve the inconsistency. When nodes agree, they slow their communication rate exponentially, such that in a stable state nodes send at most a few packets per hour. Instead of flooding a network with packets, the algorithm controls the send rate so each node hears a small trickle of packets, just enough to stay consistent. Furthermore, by relying only on local broadcasts, Trickle handles network repopulation, is robust to network transience, loss, and disconnection, and requires very little state (implementations use 4–11 bytes).

While Trickle was originally designed for reprogramming protocols (where the data is the code of the program being updated), experience has shown it to be a powerful mechanism that can be applied to wide range of protocol design problems. For example, routing protocols can use Trickle to ensure that nodes in a given neighborhood have consistent, loop-free routes. When the topology is consistent, nodes occasionally gossip to check that they still agree, and when the topology changes they gossip more frequently, until they reach consistency again.

For the purpose of clearly explaining the reasons behind Trickle's design, all of the experimental results in this section are from simulation, in some cases very high-level abstract simulators. In practice, Trickle's simplicity means it works remarkably well in the far more challenging and difficult real world. The original Trickle paper,¹³ as well as Deluge⁹ and DIP¹⁴ report experimental results from real networks.

3.1. Algorithm

Trickle's basic mechanism is a randomized, suppressive broadcast. A Trickle has a time interval of length τ and a redundancy constant k . At the beginning of an interval, a node sets a timer t in the range of $\frac{\tau}{2}, \tau$. When this timer fires, the node decides whether to broadcast a packet containing metadata for detecting inconsistencies. This decision is based on what packets the node heard in the interval before t . A Trickle maintains a counter c , which it initializes to 0 at the beginning of each interval. Every time a node hears a Trickle broadcast that is consistent with its own state, it increments c . When it reaches time t , the Trickle broadcasts

Figure 3: Trickle parameters and variables.

τ	Communication interval length
T	Timer value in range $\frac{\tau}{2}, \tau$
C	Communication counter
K	Redundancy constant
τ_l	Smallest τ
τ_h	Largest τ

if $c < k$. Randomizing t spreads transmission load over a single-hop neighborhood, as nodes take turns being the first node to decide whether to transmit. Figure 3 summarizes Trickle's parameters.

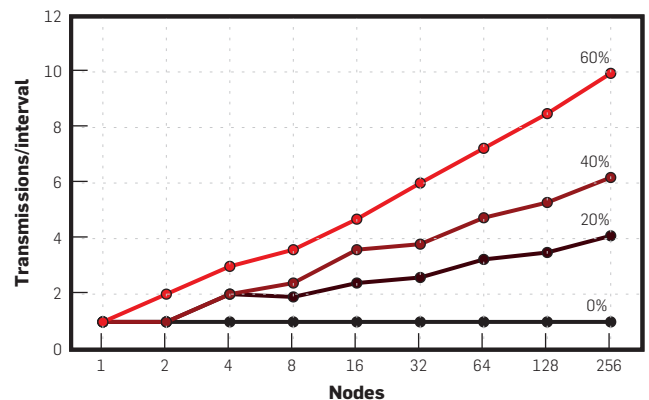
3.2. Scalability

Transmitting only if $c < k$ makes a Trickle density aware, as it limits the transmission rate over a region of the network to a factor of k . In practice, the transmission load a node observes over an interval is $O(k \cdot \log(d))$, where d is the network density. The base of the logarithm depends on the packet loss rate PLR : it is $\frac{1}{1-PLR}$.

This logarithmic behavior represents the probability that a single node misses a number of transmissions. For example, with a 10% loss rate, there is a 10% chance that a node will miss a single packet. If a node misses a packet, it will transmit, resulting in two transmissions. There is correspondingly a 1% chance a node will miss two packets from other nodes, leading to three transmissions. In the extreme case of a 100% loss rate, each node is by itself: transmissions scale linearly.

Figure 4 shows this scaling. The number of transmissions scales logarithmically with density and the slope line (base of the logarithm) depends on the loss rate. These results come from a Trickle-specific algorithmic simulator we implemented to explore the algorithm's behavior under controlled conditions. Consisting of little more than an event queue, this simulator allows configuration of all of Trickle's parameters, run duration, and the boot time of nodes. It models a uniform packet loss rate (same for all links) across a single hop network. Its output is a packet send count.

Figure 4: Trickle's transmissions per interval scales logarithmically with density. The base of the logarithm is a function of the packet loss rate (the percentages)



3.3. Synchronization

The scaling shown in Figure 4 assumes that all nodes are synchronized, such that the intervals during which they are awake and listening to their radios line up perfectly. Inevitably, this kind of time synchronization imposes a communication, and therefore energy, overhead. While some networks can provide time synchronization to Trickle, others cannot. Therefore, Trickle is designed to work in both the presence and absence of synchronization.

Trickle chooses t in the range of $(\frac{\tau}{2}, \tau]$ rather than $(0, \tau]$ because the latter causes the transmission load in unsynchronized networks to scale with $O(\sqrt{d})$. This undesirable scaling occurs due to the *short listen problem*, where some subset of motes gossip soon after the beginning of their interval. They listen for only a short time, before anyone else has a chance to speak up. This is not a problem if all of the intervals are synchronized, since the first gossip will quiet everyone else. However, if nodes are not synchronized, a node may start its interval just after another node's broadcast, resulting in missed messages and increased transmission load.

Unlike loss, where the extra $O(\log(d))$ transmissions keep the worst case node that missed several packets up to date, the additional transmissions due to the short listen problem are completely wasteful. Choosing t in the range of $(\frac{\tau}{2}, \tau]$ removes this problem: it defines a "listen-only" period of the first half of an interval. A listening period improves scalability by enforcing a simple constraint. If sending a message guarantees a silent period of some time T that is independent of density, then the send rate is bounded above (independent of the density). When a mote transmits, it suppresses all other nodes for at least the length of the listening period. Figure 5 shows how a listen period of $\frac{\tau}{2}$ bounds the total sends in a lossless single-hop network to be $2k$. With loss, transmissions scale as $O(2k \cdot \log(d))$ per interval, returning scalability to the $O(\log(d))$ goal.

3.4. Controlling τ

A large τ (gossiping interval) leads to a low communication overhead, but propagates information slowly. Conversely,

Figure 5: Without a listen-only period, Trickle's transmissions scale with a square root of the density when intervals are not synchronized. With a listen-only period of duration $\frac{\tau}{2}$, the transmissions per interval asymptotically approach $2k$. The black line shows how Trickle scales when intervals are synchronized. These results are from lossless networks.

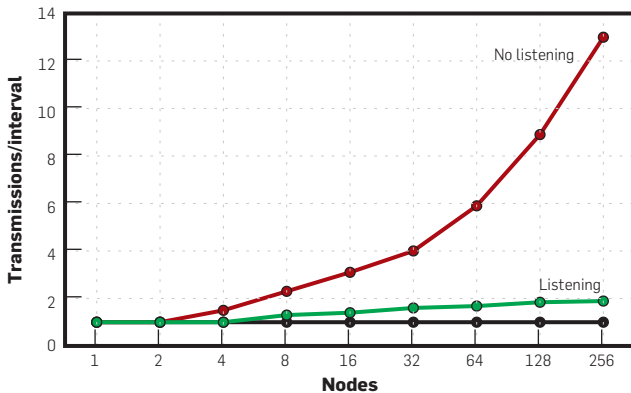


Figure 6: Trickle pseudocode.

Event	Action
τ Expires	Double τ , up to τ_h . Reset c , pick a new t^*
t Expires	If $c < k$, transmit
Receive consistent data	Increment c
Receive inconsistent data	Set t to τ . Reset c , pick a new t

* t is picked from the range $[\frac{\tau}{2}, \tau]$

a small τ imposes a higher communication overhead, but propagates data more quickly. These two goals, rapid propagation and low overhead, are fundamentally at odds: the former requires communication to be frequent, while the latter requires it to be infrequent.

By dynamically scaling τ , Trickle can quickly make data consistent with a very small cost. τ has a lower bound, τ_l , and an upper bound τ_h . When τ expires without a node receiving a new update, τ doubles, up to a maximum of τ_h . When a node detects a data inconsistency (e.g., a newer version number in dissemination, a gradient constraint violation in collection), it resets τ to be τ_l .

Essentially, when there is nothing new to say, motes gossip infrequently: τ is set to τ_h . However, as soon as a mote hears something new, it gossips more frequently, so those who have not heard the new data receive it quickly. The chatter then dies down, as τ grows from τ_l to τ_h .

By adjusting τ in this way, Trickle can get the best of both worlds: rapid consistency and low overhead when the network is consistent. The cost per inconsistency (shrinking τ) is approximately $\log(\frac{\tau_h}{\tau_l})$ additional sends. For a τ_l of 1 s and a τ_h of 1 h, this is a cost of 11 packets to obtain a 3000-fold decrease in the time it takes to detect an inconsistency (or, from the other perspective, a 3000-fold decrease in maintenance overhead). The simple Trickle policy, "every once in a while, transmit unless you have heard a few other transmissions," can be used both to inexpensively maintain that the network is consistent as well as quickly inform nodes when there is an inconsistency.

Figure 6 shows pseudocode for the complete Trickle algorithm.

3.5. Case study: Maté

Maté is a lightweight bytecode interpreter for wireless sensor networks.¹¹ Programs are tiny sequences of optimized bytecodes. The Maté runtime uses Trickle to install new programs in a network, by making all nodes consistent to the most recent version of a script.

Maté uses Trickle to periodically broadcast version summaries. In all experiments, code routines fit in a single packet (30 bytes). The runtime registers routines with a Trickle propagation service, which then maintains all of the necessary timers and broadcasts, notifying the runtime when it installs new code. Maté uses a very simple consistency resolution mechanism. It broadcasts the missing routines three times: 1, 3, and 7 s after hearing there is an inconsistency.

Figure 7 shows simulation results of Maté's behavior during a reprogramming event. These results come from the TOSSIM simulator,¹² which simulates entire sensor network applications and

models wireless connectivity at the bit level. In these experiments, τ_i is 1 s and τ_h is 1 min.

Each simulation had 400 nodes regularly placed in a square grid with node spacings of 5, 10, 15, and 20 ft. By varying network density, we were able to examine how Trickle's propagation rate scales over different loss rates and physical densities. Density ranged from a 5 ft spacing between nodes up to 20 ft (the networks were $95' \times 95'$ to $380' \times 380'$). Crossing the network in these topologies takes from six to forty hops.^a Time to complete propagation varied from 16 s in the densest network to about 70 s for the sparsest, representing a latency of 2.7 and 1.8 s per hop, respectively. The minimum per-hop Trickle latency is $\frac{\tau_i}{2}$ and the consistency mechanism broadcasts a routine 1 s after discovering an inconsistency, so the best case latency is 1.5 s per hop. Despite an almost complete lack of coordination between nodes, Trickle still is able to cause them to cooperate efficiently.

Figure 8 shows how adjusting τ_h changes the propagation time for the 5 and 20 ft spacings. Increasing τ_h from 1 to 5 min does not significantly affect the propagation time; indeed, in the sparse case, it propagates faster until roughly the 95th percentile. This result indicates that there may be little trade-off between the maintenance overhead of Trickle and its effectiveness in the face of a propagation event.

A very large τ_h can increase the time to discover inconsistencies to be approximately $\frac{\tau_h}{2}$. However, this is only true when two stable subnets ($\tau = \tau_h$) with different code reconnect. If new code is introduced, it immediately triggers nodes to reset τ to τ_p , bringing them quickly to a consistent state.

The Maté implementation of Trickle requires few system resources. It requires approximately 70 bytes of RAM; half of this is a message buffer for transmissions, a quarter is pointers to code routines. Trickle itself requires only 11 bytes for its counters; the remaining RAM is for internal coordination (e.g., pending and initialization flags). The executable code is 1.8 K (90 lines of code). Other implementations have similar costs. The algorithm requires few CPU cycles, and can operate at a very low duty cycle.

3.6. Uses and improvements

Trickle is not just used by Maté; it and its derivatives are used in almost every dissemination protocol today. The Deluge binary dissemination protocol for installing new sensor node firmware uses Trickle to detect when nodes have different firmware versions⁹ ($\tau_i = 500$ ms, $\tau_h = 1.1$ h). The MNP binary dissemination protocol ($\tau_i = 16$ s, $\tau_h = 512$ s) adjusts Trickle so that nodes with more neighbors are more likely to send updates by preventing low degree nodes from suppressing high degree ones.²³ The Drip command layer of the Sensornet Management System uses Trickle ($\tau_i = 100$ ms, $\tau_h = 32$ s) to install commands.²² The Tenet programming ar-

^a These hop count values come from computing the minimum cost path across the network loss topology, where each link has a weight of $\frac{1}{1-\text{loss}}$, i.e. the expected number of transmissions to successfully traverse that link.

Figure 7: Time to consistency in 20×20 TOSSIM grids (seconds). The hop count values in each legend are the expected number of transmissions necessary to get from corner to corner, considering loss.

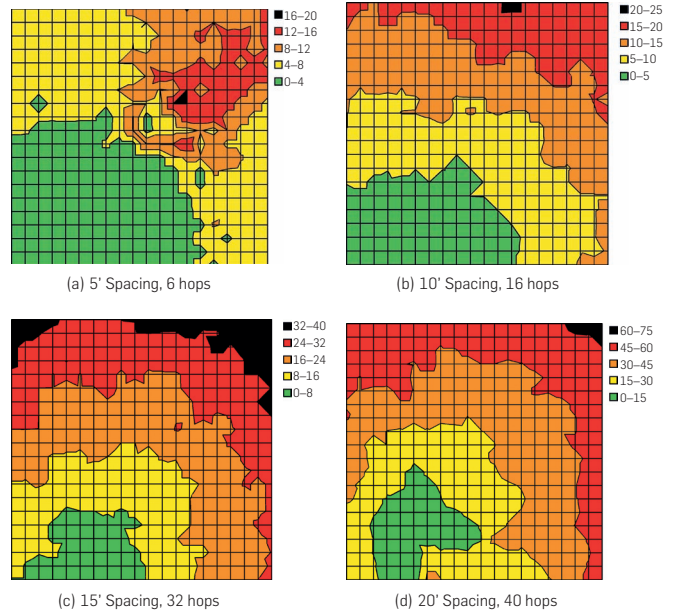
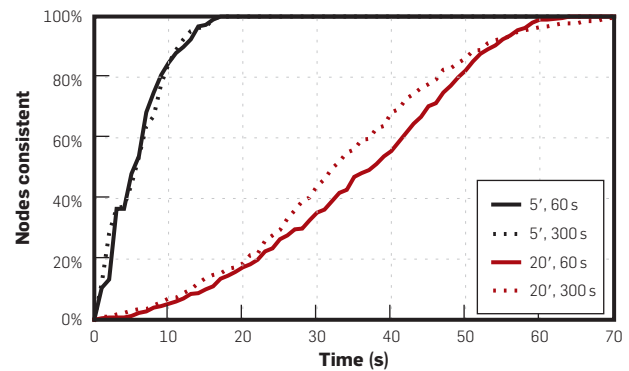


Figure 8: Rate nodes reach consistency for different τ_h s in TOSSIM. A larger τ_h does not slow reaching consistency.



chitecture uses Trickle ($\tau_i = 100$ ms, $\tau_h = 32$ s) to install small dynamic code tasks.⁷

In the past few years, as collection protocols have improved in efficiency, they have also begun to use Trickle. The CTP, the standard collection layer in the TinyOS operating system distribution,²¹ uses Trickle timers ($\tau_i = 64$ ms, $\tau_h = 1$ h) for its routing traffic. The 6LoWPAN IPv6 routing layer in Arch Rock's software uses Trickle to keep IPv6 routing tables and ICMP neighbor lists consistent.¹ As protocols continue to improve, Trickle's efficacy and simplicity will cause it to be used in more protocols and systems.

One limitation with Trickle as described in this paper is that its maintenance cost grows $O(n)$ with the number of data items, as nodes must exchange version numbers. This growth may be a hindering factor as Trickle's use increases. Recent work on the DIP protocol addresses this

concern by using a combination of hash trees and randomized searches, enabling the maintenance cost to remain $O(1)$ while imposing a $O(\log(n))$ discovery cost.¹⁴

4. DISCUSSION

Wireless sensor networks, like other ad hoc networks, do not know the interconnection topology a priori and are typically not static. Nodes must discover it by attempting to communicate and then observing where communication succeeds. In addition, the communication medium is expected to be lossy. Redundancy in such networks is both friend and foe, but Trickle reinforces the positive aspects and suppresses the negative ones.

Trickle draws on two major areas of prior research. The first area is controlled, density-aware flooding algorithms for wireless and multicast networks.^{5,17} The second is epidemic and gossiping algorithms for maintaining data consistency in distributed systems.⁴ Although both techniques—broadcasts and epidemics—have assumptions that make them inappropriate in their pure form to eventual consistency in sensor networks, they are powerful techniques that Trickle draws from. Trickle's suppression mechanism is inspired by the request/repair algorithm used in Scalable and Reliable Multicast (SRM).⁵ Trickle adapts to local network density as controlled floods do, but continually maintains consistency in a manner similar to epidemic algorithms. Trickle also takes advantage of the broadcast nature of the wireless channel, employing SRM-like duplicate suppression to conserve precious transmission energy and scale to dense networks.

Exponential timers are a common protocol mechanism. Ethernet, for example, uses an exponential backoff to prevent collisions. While Trickle also has an exponential timer, its use is reversed. Where Ethernet defaults to the smallest time window and increases it only in the case of collisions, Trickle defaults to the largest time window and decreases it only in the case of an inconsistency. This reversal is indicative of the different priorities in ultra-low-power networks: minimizing energy consumption, rather than increasing performance, is typically the more important goal.

In the case of dissemination, Trickle timers spread out packet responses across nodes while allowing nodes to estimate their degree and set their communication interval. Trickle leads to energy efficient, density-aware dissemination not only by avoiding collisions through making collisions rare, but also by suppressing unnecessary retransmissions.

Instead of trying to enforce suppression on an abstraction of a logical group, which can become difficult in multihop networks with dynamically changing connectivity, Trickle suppresses in terms of an implicit group: nearby nodes that hear a broadcast. Correspondingly, Trickle does not impose the overhead of discovering and maintaining logical groups, and effortlessly deals with transient and lossy wireless links. By relying on this implicit naming, the Trickle algorithm remains very simple: implementations can fit in under 2K of code, and require a mere 11 bytes of state.

Routing protocols discover other routers, exchange routing information, issue probes, and establish as well as tear

down links. All of these operations can be rate-controlled by Trickle. For example, in our experiences exploring how wireless sensor networks can adopt more of the IPv6 stack in 6LoWPAN, Trickle provides a way to support established ICMP-v6 mechanisms for neighbor discovery, duplicate address detection, router discovery, and DHCP in wireless networks. Each of these involves advertisement and response. Trickle mechanisms are a natural fit: they avoid loss where density is large, allow prompt notifications of change and adapt to low energy consumption when the configuration stabilizes. By adopting a model of eventual consistency, nodes can locally settle on a consistent state without requiring any actions from an administrator.

Trickle was initially developed for distributing new programs into a wireless sensor network: the title of the original paper is "Trickle: A Self-Regulating Algorithm for Code Propagation and Maintenance in Wireless Sensor Networks."¹³ Experience has shown it to have much broader uses. Trickle-based communication, rather than flooding, has emerged as the central paradigm for the basic multihop network operations of discovering connectivity, data dissemination, and route maintenance.

Looking forward, we expect the use of these kinds of techniques to be increasingly common throughout the upper layers of the wireless network stack. Such progress will not only make existing protocols more efficient, it will enable sensor networks to support layers originally thought infeasible. Viewing protocols as a continuous process of establishing and adjusting a consistent view of distributed data is an attractive way to build robust distributed systems.

Acknowledgments

This work was supported, in part, by the Defense Department Advanced Research Projects Agency (grants F33615-01-C-1895 and N6601-99-2-8913), the National Science Foundation (grants No. 0122599, IIS-033017, 0615308, and 0627126), by the California MICRO program, Intel Corporation, DoCoMo Capital, Foundation Capital, and a by Stanford Terman Fellowship. Research infrastructure was provided by the National Science Foundation (grants No. 9802069). We would also like to thank Sylvia Ratnasamy for her valuable insights into the early stages of this research, as well as Jonathan Hui, for his ideas on applying Trickle to new problems. C

References

1. Arch Rock Corporation. An IPv6 Network Stack for Wireless Sensor Networks. <http://www.archrock.com>.
2. Couto, D.D., Aguayo, D., Bicket, J., and Morris, R. A high-throughput path mMetric for multi-hop wireless routing. *Proceedings of the Ninth Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2003.
3. Crossbow, Inc. Mote in Network Programming User Reference. <http://webs.cs.berkeley.edu/tos/tinyos-1.x/doc/Xnp.pdf>.
4. Demers, A., Greene, D., Hauser, C., Irish, W., and Larson, J. Epidemic Algorithms for Replicated Database Maintenance. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 1987.
5. Floyd, S., Jacobson, V., McCanne, S., Liu, C.-G., and Zhang, L. A reliable multicast framework for light-weight sessions and application level framing. *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, 1995.
6. Fonseca, R., Gnanawali, O., Jamieson, K., and Levis, P. Four bit wireless link estimation. *Proceedings of the Sixth Workshop on Hot Topics in Networks (HotNets VI)*, 2007.
7. Gnanawali, O., Greenstein, B., Jang, K.-Y., Joki, A., Paek, J., Vieira, M., Estrin, D., Govindan, R., and Kohler,



Association for
Computing Machinery

Advancing Computing as a Science & Profession

ACM Transactions on Asian Language Information Processing

March 2008, Volume 1, Number 1

K. F. WONG J. T. TSIEN	1	Prologue
J. GAO J. GOODMAN M. LI R. F. LEE	3	Toward a Unified Approach to Statistical Language Modeling for Chinese
Y. JIA C. H. WU	34	Meaningful Term Extraction and Discriminative Term Selection in Text Categorization via Unknown-Word Methodology
B. XIA D. G. LEE J. H. LEE	65	Morpheme-Based Graphs to Phoneme Conversion Using Phonetic Features and Morphosyntactic Connectivity Information
T. LEE W. LAU Y. W. NGONG F. G. CHENG	83	Using Tone Information in Cantonese Continuous Speech Recognition

ACM Transactions On Asian Language Information Processing

The Asian Language Information Processing Transaction (TALIP) publishes high quality original archival papers and technical notes in the areas of computation and processing of information in Asian languages and related disciplines. Some of the subjects to be covered by this quarterly publication are: Computational Linguistics; Linguistic Resources; Hardware and Software Algorithms and Tools for Asian Language Processing; Machine Translation; and Multimedia Asian Information Processing. Emphasis will be placed on the originality and the practical significance of the reported research.

To learn more about TALIP, please visit
www.acm.org/pubs/talip/

SUBSCRIBE TODAY!

PRODUCT INFORMATION

ISSN: 1530-0226

Order Code: 138

Price: \$38 Professional Member

\$33 Student Member

\$160 Non-Member

\$16 Air Service (for residents
outside North America only)

TO PLACE AN ORDER

Contact ACM Member Services

Phone: 1.800.342.6626 (U.S. and Canada)
+1.212.626.0500 (Global)

Fax: +1.212.944.1318
(Hours: 8:30am—4:30pm, Eastern Time)

Email: acmhlp@acm.org

Mail: ACM Member Services

General Post Office

PO Box 30777

New York, NY 10087-0777 USA

www.acm.org/pubs/talip/

AD28

References

- E. The TENET architecture for tiered sensor networks. *Proceedings of the Fourth International Conference on Embedded Networked Sensor Systems (Sensys)*, 2006.
- Hill, J., Szewczyk, R., Woo, A., Hollar, S., Culler, D.E., and Pister, K.S.J. System architecture directions for networked sensors. *Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLoS)*, 2000.
- Hui, J.W. and Culler, D. The dynamic behavior of a data dissemination protocol for network programming at scale. *Proceedings of the Second International Conference on Embedded Networked Sensor Systems (SenSys)*, 2004.
- Intanagonwiwat, C., Govindan, R., and Estrin, D. Directed diffusion: a scalable and robust communication paradigm for sensor networks. *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2000.
- Levis, P., Gay, D., and Culler, D. Active sensor networks. *Proceedings of the Second USENIX/ACM Symposium on Network Systems Design and Implementation (NSDI)*, 2005.
- Levis, P., Lee, N., Welsh, M., and Culler, D. TOSSIM: accurate and scalable simulation of entire TinyOS applications. *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2003.
- Levis, P., Patel, N., Culler, D., and Shenker, S. Trickle: a self-regulating algorithm for code maintenance and propagation in wireless sensor networks. *Proceedings of the First USENIX/ACM Symposium on Network Systems Design and Implementation (NSDI)*, 2004.
- Lin, K. and Levis, P. Data discovery and dissemination with DIP. *Proceedings of the Seventh International Symposium on Information Processing in Sensor Networks (IPSN)*, 2008.
- Madden, S., Franklin, M.J., Hellerstein, J.M., and Hong, W. TinyDB: an acquisitional query processing system for sensor networks. *Transactions on Database Systems (TODS)*, 2005.
- Mao, Y., Wang, F., Qiu, L., Lam, S., and Smith, J. S4: small state and small stretch routing protocol for large wireless sensor networks. *Proceedings of the Fourth USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2007.
- Ni, S.-Y., Tseng, Y.-C., Chen, Y.-S., and Sheu, J.-P. The broadcast storm problem in a mobile ad hoc network. *Proceedings of the Fifth Annual International Conference on Mobile Computing and Networking (MobiCom)*, 1999.
- Paek, J. and Govindan, R. RCRT: rate-controlled reliable transport for wireless sensor networks. *Proceedings of the Fifth International Conference on Embedded Networked Sensor Systems (SenSys)*, 2007.
- Rangwala, S., Gummadi, R., Govindan, R., and Psounis, K. Interference-aware fair rate control in wireless sensor networks. *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, 2006.
- Sun Microsystems Laboratories. Project Sun SPOT: Small Programmable Object Technology. <http://www.sunspotworld.com/>.
- TinyOS Network Protocol Working Group. TEP 123: The Collection Tree Protocol. <http://www.tinyos.net/tinyos-2.x/doc/txt/tep123.txt>, 2007.
- Tolle, G. and Culler, D. Design of an application-cooperative management system for wireless sensor networks. *Proceedings of the Second European Workshop of Wireless Sensor Networks (EWSN)*, 2005.
- Wang, L. MNP: Multihop network reprogramming service for sensor networks. *Proceedings of the Second International Conference on Embedded Networked Sensor Systems (SenSys)*, 2004.
- Woo, A., Tong, T., and Culler, D. Taming the underlying challenges of multihop routing in sensor networks. *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2003.
- Yang, J., Soffa, M.L., Selavo, L., and Whitehouse, K. Clairvoyant: a comprehensive source-level debugger for wireless sensor networks. *Proceedings of the Fifth International Conference on Embedded Networked Sensor Systems (SenSys)*, 2007.

Philip Levis (pal@cs.stanford.edu)
Assistant Professor, Stanford University,
Stanford, CA, USA

Eric Brewer (brewer@cs.berkeley.edu)
Professor, U.C. Berkeley, Berkeley,
CA, USA

David Culler (culler@cs.berkeley.edu)
Professor, U.C. Berkeley, Berkeley,
CA, USA

David Gay (david.e.gay@intel.com) Senior
Researcher, Intel Research Berkeley,
Berkeley, CA, USA

Samuel Madden (madden@csail.mit.edu) Associate Professor, MIT CSAIL,
Cambridge, MA, USA

Neil Patel (neilp@cs.stanford.edu) Ph.D.
Student, Stanford University, Stanford,
CA, USA

Joe Polastre (joe@sentilla.com) CTO,
Sentilla Corporation, Redwood City,
CA, USA

Scott Shenker (shenker@icsi.berkeley.edu) Professor, U.C. Berkeley, Berkeley,
CA, USA

Robert Szewczyk (rob@sentilla.com)
Principal Engineer, Sentilla Corporation,
Redwood City, CA, USA

Alec Woo (awoo@archrock.com)
Technical Staff, Arch Rock Corporation,
San Francisco, CA, USA

© 2008 ACM 001-0782/08/0700 \$5.00