

Technical Overview of Time Synchronized Mesh Protocol (TSMP)

Introduction.....	2
Wireless Sensor Networks.....	3
TSMP Overview.....	5
TSMP Components.....	7
Time synchronized communication.....	7
Frequency hopping.....	8
Automatic node joining and network formation...	11
Fully redundant mesh routing.....	13
Secure message transfer.....	14
Conclusion.....	16
Appendix.....	17

Introduction

TSMP (Time Synchronized Mesh Protocol) is a networking protocol that forms the foundation of reliable, ultra low-power wireless sensor networking. Wireless sensor networks (WSNs) are self-organizing, multi-hop networks of wireless sensor nodes used to monitor and control physical phenomena. Typical WSN applications include industrial process automation, commercial building climate control and security alarming.

TSMP provides redundancy and fail-over in time, frequency and space to ensure very high reliability even in the most challenging radio environments. TSMP also provides the intelligence required for self-organizing, self-healing mesh routing. The result is a network that installs easily with no specialized wireless expertise, automatically adapts to unforeseen challenges, and can be extended as needed without sophisticated planning.

There are five key components of TSMP that contribute to end-to-end network reliability, simple installation and power efficiency.

- *Time synchronized communication*
- *Frequency hopping*
- *Automatic node joining and network formation*
- *Fully-redundant mesh routing*
- *Secure message transfer*

This whitepaper provides a survey of WSN solutions and describes TSMP with enough detail to provide the technical reader with a full picture of protocol capabilities.

Wireless Sensor Networks

Wireless sensor network (WSN) is a term used to describe an emerging class of embedded communication products that provide redundant, fault-tolerant wireless connections between sensors, actuators and controllers. WSNs are deployed to provide access to assets or instruments that were previously deemed unreachable due to physical or economic barriers.

By literal definition, WSN is a term that can be applied to any wirelessly connected instrument (even a garage door opener). In practice, the WSN label is used to describe products that provide performance above and beyond traditional point-to-point solutions, particularly in areas of fault tolerance, power consumption and installation cost.

Wireless Challenges

While wireless provides clear advantages in cost and flexibility, it also brings along a host of challenges. Specifically, point-to-point radio communication links are notoriously variable and unpredictable. A link that is strong today may be weak tomorrow due to environmental conditions, new obstacles, unanticipated interferers and myriad other factors. These factors can be boiled down into three major failure modes: RF interference, changes in the physical environment that block communication links, and loss of individual nodes.

- *RF interference:* The small portion of the electromagnetic spectrum devoted to general-purpose wireless communication devices is crowded with traffic from Wi-Fi networks, cordless telephones, bar-code scanners, and innumerable other devices that can interfere with communications. Because there is no way to predict what interferers will be present in a facility at a given location, frequency, and time, a reliable network must be able to continually sidestep these interferers on an ongoing basis.
- *Blocked Paths:* When a network is first deployed, wireless paths are established between devices based on the immediate RF environment and available neighbors. Unlike wired networks, these variables often change; paths may later be blocked by new equipment, repositioned partitions, delivery trucks, or very small changes in device position. Assuring reliability for the life of the network, not just the first few weeks after installation, requires continually working around these blockages in a transparent, automatic fashion.
- *Node Loss:* Node loss is an important issue to consider with wireless sensor networks. While node failure because of semiconductor or hardware malfunction is rare, nodes may be damaged, destroyed or removed during the life of the network. Additionally, power surges, blackouts, or brownouts can cause nodes to fail unless they have an independent power source. End-to-end reliability requires the networking intelligence that routes around the loss of any single node.

Technical Overview of TSMP

Any of these problems will bring down a point-to-point wireless link. However, with a network architecture designed to protect against these issues, the network can isolate individual points of failure and eliminate or mitigate their impact, allowing the network as a whole to maintain very high end-to-end reliability in spite of local failures. Similarly, a well-designed wireless network architecture will transparently adapt to changing environments, allowing long-term operation with zero-touch maintenance.

WSNs aim to overcome these challenges by applying self-organizing and self-healing intelligence to continuously adapt to unpredictable conditions. The goal of WSN technology is to provide extremely high reliability and predictability for years at a time without constant tuning by wireless experts.

Time Synchronized Mesh Protocol (TSMP) provides a mechanism for WSN intelligence. By defining how a wireless node utilizes radio spectra, joins a network, establishes redundancy and communicates with neighbors, TSMP forms a solid foundation for WSN applications.

TSMP Overview

TSMP is a media access and networking protocol that is designed specifically for low-power, low-bandwidth reliable networking. Current TSMP implementations operate in the 2.4 GHz ISM band on IEEE 802.15.4 radios and in the 900 MHz ISM band on proprietary radios. Figure 1 shows the elements of TSMP in the standard wireless network stack and the OSI network stack.

TSMP Stack	Standard Wireless Stack	OSI Stack
Application	Application	Application
Presentation	Presentation	Presentation
Session	Session	Session
TSMP	Network	Transport
	Media Access	Network
Physical	Physical	Data Link
		Physical

Figure 1. Mapping TSMP to Common Protocol Stack Models

TSMP is a packet-based protocol where each transmission contains a single packet and acknowledgements (ACKs) are generated when a packet has been received unaltered and complete. Mechanisms are in place to transport packets across a multi-hop network as efficiently and reliably as possible. All measures of reliability and efficiency are done on a per-packet basis.

Packet Structure

TSMP packets consist of a header, a payload and a trailer. Packets contain fields that identify the sending node, define the destination, ensure secure message transfer and provide reliability and quality of service information. For the purposes of this paper we will discuss the implementation of TSMP on IEEE 802.15.4 radios. The 802.15.4 standard specifies a maximum packet size of 127 B, TSMP reserves 47 B for operation, which leaves 80 B for payload. For a full description of TSMP packet structure please see Appendix A.

MAC Header	NET Header	Payload	APP MIC	MAC MIC	FCS
------------	------------	---------	---------	---------	-----

Figure 2. TSMP Packet Structure

TSMP also defines several packet types. These packet types enable specific functions within the network. Some packet types take priority over others; some allow transparent tunneling while others require packet parsing at each hop in the route.

Technical Overview of TSMP

Definitions

Several terms are used throughout the following sections that are not common and may not be familiar to the reader.

TSMP Node: a wireless device running TSMP

TSMP Network: a network of TSMP nodes

Path: a bidirectional single-hop connection between any two TSMP nodes. Think of this as a line drawn between two nodes to connote connectivity.

Link: a directed communication channel between two TSMP nodes. There are multiple links per path. Links are directional and may be added/removed from a path to increase/decrease available bandwidth.

Route: A series of paths that connect a source node to a destination node. In a mesh network, a route often consists of multiple hops.

Parent Node: a node that is one hop closer to the destination than the reference node. Parent nodes route data for child nodes.

Child Node: a node that is one hop further away from the destination than the reference node. Child nodes pass data off to parent nodes.

Mesh: a network with fully redundant routing for all nodes

Star: a network with non-redundant routes between end nodes and a central router

A wireless device with an embedded microprocessor running TSMP is referred to as a *TSMP node*. A network of TSMP nodes connected by paths is a *TSMP network*. A TSMP network forms a mesh topology where data travels on routes from a source (typically a sensor) to a destination (typically a gateway).

TSMP Components

In the following pages each key component of TSMP is broken out in some detail. After reading this section a technical reader should have a good picture of how a TSMP node works and how a TSMP network behaves.

TSMP consists of five key components:

- Time synchronized communication
- Frequency hopping
- Automatic node joining and network formation
- Fully redundant mesh routing
- Secure message transfer

Time Synchronized Communication

All node-to-node communication in a TSMP network is transacted in a specific time window. Commonly referred to as Time Division Multiple Access (TDMA), synchronized communication is a proven technique that provides reliable and efficient transport of wireless data. Unlike wired systems where nodes can be directly connected by a dedicated wire (media), to the exclusion of neighbors, in a wireless system all devices within range of each other must share the same media. Several other Media Access Control (MAC) mechanisms are available including CSMA, CDMA and TDMA. TSMP is based on TDMA.

Timeslots and Frames

In TSMP, each communication window is called a timeslot. A series of timeslots makes up a frame, which repeats for the life of network. Frame length is counted in slots and is a configurable parameter—in this way a particular refresh rate is established for the network. A shorter frame length increases refresh rate, increasing effective bandwidth and increasing power consumption. Conversely, a longer frame length decreases refresh rate, thereby decreasing bandwidth and decreasing power consumption. A TSMP node can participate in multiple frames at once allowing it to effectively have multiple refresh rates for different tasks. The concept of slots and frames is illustrated in Figure 3.




Figure 3. TSMP Slots and Frames

Synchronization

A critical component of any TDMA system is time synchronization—all nodes must share a common sense of time so that they know precisely when to talk, listen, or sleep. This is especially critical in power-constrained applications like WSNs where battery power is often the only option and changing batteries can be costly and cumbersome. In contrast to ‘beaconing’ strategies employed by other WSN implementations, TSMP does not begin each frame with a synchronization beacon. Beaconing strategies can require long listen windows which consume power. Instead, TSMP nodes maintain a precise sense of time and exchange offset information with neighbors to ensure alignment. These offset values ride along in standard ACK messages and cost no extra power or overhead.

A common sense of time enables many network virtues: bandwidth can be pre-allocated to ensure extremely reliable transmission and zero self-interference; transmitting nodes can effectively change frequencies on each transmission and the receiving node can keep in lock-step; bandwidth can be added and removed at will in a very predictable and methodical way to accommodate traffic spikes; and many others.

Duty Cycling

It is important to note that TSMP nodes are only active in three states: 1) sending a message to a neighbor, 2) listening for a neighbor to talk, and 3) interfacing with an embedded sensor or processor. For all other times the node is asleep and consuming very low power. In a wireless device the majority (generally >95%) of the total power budget is consumed by the radio. To achieve low power it is clear that one must minimize radio on time. TDMA is very good at this. Timeslots are measured in milliseconds and in typical WSN applications this leads to a duty cycle of less than 1% for all nodes in the network (including those relaying messages for neighbors). Because all nodes (including those often called ‘routers’) can be aggressively duty cycled, TDMA is the only practical solution for a fully battery-powered network.

Frequency Hopping

In addition to slicing the wireless media across time, TSMP also slices it across frequency. This provides robust fault tolerance in the face of common RF interferers as well as providing a tremendous increase in effective bandwidth. Commonly referred to as Frequency Hopping Spread Spectrum (FHSS), hopping across multiple frequencies is a proven way to sidestep interference and overcome RF challenges with agility rather than brute force.

Another technique to overcome RF challenges is Direct Sequence Spread Spectrum (DSSS). DSSS provides a few dB of coding gain and some improvement in multi-path fading. While beneficial, DSSS is not sufficient in the face of common interferers in the band, including Wi-Fi equipment, two-way radios or even Bluetooth (see figure 4 below). It should be noted that a combination of FHSS and DSSS provides both interference rejection (FHSS) and the coding gain (DSSS).

The other technique for overcoming interference is increasing the radio power—effectively “turning up the volume.” Although often effective, turning up the volume on 802.15.4 radios kills battery life and is not an ideal solution for low-power WSNs.




Figure 4. Frequency Hopping vs. DSSS in 802.15.4 Networks (Source: Dust Networks)

Hopping Sequence

Upon joining a network, a TSMP node (call it node C) will discover available neighbors and establish communication with at least two nodes already in the network, call them parent A and parent B (more on this in later sections). During this process node C will receive synchronization information and a frequency hopping sequence from both parent A and parent B. The 802.15.4 standard specifies 16 distinct frequency channels within the 2.4000-2.4835 MHz ISM band – so let’s use 16 as our number. The hopping sequence is a pseudo-random sequence of all available channels. For example the sequence may be: 4,15,9,7,13,2,16,8,1,etc. Node C receives a distinct start point in the sequence from each parent, and when a new node joins it, it will in turn give a distinct start point to this new child node. In this way each pair-wise connection is ensured to be on a different channel during each timeslot enabling broad use of the available band in any one location.

Technical Overview of TSMP




Figure 5. Nodes A and B are Parents of Node C

In operation, each node-to-node transmission (say $C \rightarrow A$) is on a different frequency than the previous transmission. And should a transmission be blocked, the next transmission will be to an alternate parent ($C \rightarrow B$) on a different frequency. The result is simple but extremely resilient in the face of typical RF interferers.

Bandwidth and Scalability Effects

As with most communications mechanisms, increasing the number of distinct channels proportionally increases the throughput of the system. In the case of TSMP, employing FHSS on top of the 802.15.4 radio effectively increases bandwidth by 16 times. This is because two pairs of nodes communicating on different frequencies will not interfere with each other even if they are within range. Conversely, for low data rate applications this means that even if the majority of the band is blocked by RF interference, the messages will still find a clear channel and get through. In either case the effect of FHSS is to greatly increase the reliability of the system.

Combining the frequency and time division into one map provides the following matrix. Each vertical column is a timeslot and each horizontal row is a frequency. Every cell (box) is a unique communication opportunity for a pair of TSMP nodes.

Technical Overview of TSMP




Figure 6. Frequency/Time Matrix

For example, a TSMP implementation on 802.15.4 radios with 60 timeslots per second provides:

$$16 \text{ channels} \times 60 \text{ slots/second} = 960 \text{ transmissions/second}$$

Assuming an 80 B payload the effective total bandwidth is:

$$960 \text{ transmissions/second} \times 80 \text{ B/transmission} = 76.8 \text{ KB/second}$$

Given that scalability in wireless systems is primarily governed by access to the media, the more efficient the media access protocol the more scalable the network. A frequency hopping TDMA protocol is a very efficient means to coordinate node communications. It has been demonstrated that over 1,000 TSMP nodes can operate in the same radio space with each other without effecting end-to-end reliability. In contrast, dense networks of nodes using collision-based protocols like CSMA (Carrier Sense Multiple Access) often experience cascading collisions and network failure.

Automatic Node Joining and Network Formation

A key attribute of a TSMP network is its ability to self-organize. Indeed, this is one of the key reasons for mesh networking in the first place. Every TSMP node has the intelligence to discover neighbors, measure RF signal strength, acquire synchronization and frequency hopping information, and then establish paths and links with neighbors.

Technical Overview of TSMP

For the purposes of this discussion, it is important to note that all TSMP nodes are fully capable mesh networking nodes. In TSMP there is no concept of reduced function, non-routing sensor nodes or end nodes. Every TSMP node has the ability to route traffic from neighbors as dictated by RF connectivity and/or network performance requirements. During the life of an installation it may be the case that a node joins as an end node, becomes a routing node due to changing RF conditions and then reverts back to an end node. This type of behavior is not uncommon in mesh networks and must happen automatically to provide long-term network reliability.

Node Joining

In this section we will describe how a TSMP node joins an established network. An established network is simply a set of nodes that share a network ID and password and are synchronized with each other. A network is typically seeded by a gateway node that serves as the timing master and relays configuration information to all other network nodes.

In addition to the timeslots that carry application messages across the network, there are other timeslots dedicated to network configuration, neighbor discovery and listening for new join requests. Just like all other timeslots, these cycle in time on a refresh rate defined by the frame length. Additionally, as network nodes communicate with each other they include special codes in the messages that advertise key network settings like frame length, open listen slots and frequency, network ID and current time. When a TSMP node is powered-up or reset, it will begin listening for these codes.

Here is a simplified state machine of a joining node:

- Listens on frequency A for a period, listens on B, listens on C,...
- Hears a neighbor and locks in on timing information and then only listens at the beginning of each slot to determine if there is a message to receive, reducing power consumption.
- Listens to this frequency for a period. During this period the node is building a neighbor list. This includes nodes within radio range that have transmitted during the period on this frequency.
- Report neighbor list including RSSI.
- Choose a neighbor and send a join request.
- Receive an activate command from the neighbor node and establish links.

All TSMP messages are encrypted and include a network ID. The network ID is used to bind nodes together into a network, allowing multiple TSMP networks to operate in the same radio space without the risk of sharing data or misrouting messages. If a mote hears a node with a network ID that does not match its own, then it will not initiate joining but will continue unsynchronized listening until it hears the right ID. There is also a join key used to encrypt messages. If the mote has the wrong join key, then its join request will not be accepted by the parent node, the mote will time out, and revert back to unsynchronized listening.

Fully Redundant Mesh Routing

Redundant routing is a must have in real world RF environments. Conditions change dramatically over time due to weather, new or unknown RF systems, moving equipment and population density. Combine this with the utter unpredictability of node placement, installer practices, and future network expansions or repurposing, and one gets a clear picture of the challenges facing RF reliability. A full mesh topology with automatic node joining and healing lets the network maintain long-term reliability and predictability in spite of these challenges. As with water flowing downhill, only self-organizing full mesh networks can find and utilize the most stable routes through the available node topology.

Fully redundant routing requires both spatial diversity (try a different route) and temporal diversity (try again later). TSMP covers spatial diversity by enabling each node to discover multiple possible parent nodes and then establish links with two or more. Temporal diversity is handled by retry and failover mechanisms.

Spatial Diversity – Redundant Routing

As mentioned previously, all TSMP nodes are router nodes. This is a fundamental advancement over star or hybrid star-mesh architectures. A full mesh topology is the only way to accommodate changing conditions. A full mesh or ‘flat’ network (no concept of higher or lower functioning nodes) does not rely on special-purpose routers, base stations, or aggregators, and does not require nearly the wireless expertise and installation skill of other solutions. There is no need to survey, engineer and then ultimately overbuild point-to-point connections. As a full mesh is installed, all connected nodes form one giant antenna for other joining nodes. This allows for extremely quick and robust installations. Additionally, should an installed network need to be expanded, only a full mesh network can gracefully accommodate new nodes by relying on edge nodes to automatically assume routing duties. Note that in some applications (where power is at an extreme premium) it may be desirable to have an end node remain an end node, and selectively refuse to assume routing duties. TSMP supports this type of customization through configurable settings.




Figure 7: Network Topologies

Each TSMP node maintains its own neighbor list. This neighbor list includes parent nodes and child nodes. A node may have as many parents as required (this is a configurable

parameter). For example, a particular high-value node may have four parents to ensure utmost reliability. Conversely, a node with little value may be configured to only acquire one parent to reserve bandwidth for other traffic.

A key enabler of TSMP's full mesh capability is the effective duty cycling of router nodes. Because a router can maintain less than a 1% duty cycle, it may be powered just as an end node. This device parity means that installation and commissioning need not consider device type, power source, etc.

Temporal Diversity – Retries and Failover

Once a link is established, TSMP provides communication mechanisms to ensure reliable operation. As discussed above, a node-to-node message transmission occurs in one timeslot on one frequency. Within one timeslot a message is sent, the sending node switches to receive mode and awaits an acknowledgement (ACK). Should an ACK not arrive within the timeslot, the sending node will retry in the next available slot. This will generally be to an alternate parent and will always be on a different frequency. Similarly, if a NACK (a message indicating the expected packet was not properly received) is received, then the sending node will retry on the next available slot. NACKs are generated a number of ways: invalid checksum (FCS), invalid message integrity code (MIC), or the receiving node has a full message queue.

Each sending node keeps track of missing ACKs and NACKs. Should a number of transmissions go unacknowledged, the sender will consider the path invalid and initiate communication with the next available node on its neighbor list.

Secure Message Transfer

There are three pillars of secure message transfer: encryption, authentication and integrity. Encryption keeps the information carried by the message from being read by other parties. Authentication ensures that the sender is actually the sender. Integrity ensures that the message was delivered unaltered. TSMP provides mechanisms for each of these functions. It is worth noting that frequency hopping provides some level of security in its own right. Because of the pseudo random hopping sequence maintained by each pair of nodes, if a snooping receiver did manage to hear one transmission, then it only has a 1 in 16 chance (for 802.15.4 radios) of hearing the next transmission.

Encryption

TSMP uses industry-standard 128-bit symmetric key encryption for end-to-end confidentiality of packet payload. Nodes that share keys communicate by encrypting messages with CTR-mode cipher. Since all nodes are time-synchronized, unique timestamps are used to generate non-repeating nonce (numbers used once) as encryption vectors.

Authentication

Technical Overview of TSMP

While encryption provides confidentiality of messages, authentication is needed to ensure source identity. To make sure that every packet in a TSMP network is generated by an authorized node, TSMP uses packet source addresses protected by 32-bit Message Integrity Codes (MIC). Every packet carries two MIC codes to provide authentication: end-to-end source address authentication guaranteed by the network layer MIC, and node-to-node source address authentication, guaranteed by the MAC layer MIC. The MAC layer authentication is particularly important in protecting ACKs.

Integrity

The same 32-bit Message Integrity Codes (MICs) that authenticate the sending node's address also serve to ensure content integrity. Any message tampering would invalidate the MICs and be immediately recognizable by the receiving node.

Conclusion

The reliability of the TSMP protocol has been proven over the past three years in challenging network deployments. The simple but powerful concepts of temporal, frequency and spatial diversity provide an extremely robust networking protocol that stands up to the challenges of real-world commercial and industrial environments. Embedded self-organizing and self-healing intelligence radically reduces the installation complexity and ensures long-term predictable behavior. All Dust Networks products are built on top of TSMP. Dust Networks is currently working with leading organizations to standardize the core components of TSMP.

Trademarks


Dust, Dust Networks, the Dust Networks logo, and TSMP are registered trademarks of Dust Networks, Inc. All third-party brand and product names are the trademarks of their respective owners and are used solely for informational purposes.

Copyright

This documentation is protected by United States and international copyright and other intellectual and industrial property laws. It is solely owned by Dust Networks, Inc. and its licensors. This product, or any portion thereof, may not be used, copied, modified, reverse assembled, reverse compiled, reverse engineered, distributed, or redistributed in any form by any means without the prior written authorization of Dust Networks, Inc.

Document Number: 025-0003-01
Last Revised: June 20, 2006.

Appendix: TSMP Packet Structure



PHY preamble	MAC header	NET header	Encrypted application payload	APP MIC-32	MAC MIC-32	FCS-16
--------------	------------	------------	-------------------------------	------------	------------	--------

Packet section	Description
PHY preamble	Preamble, start of frame delimiter, and length
MAC header	Per-hop addressing and timing information
NET header	End-to-end addressing and routing information
App payload	Encrypted application payload, packet-type dependent
APP MIC-32	End-to-end message integrity code for application data and nonce (32 bit).
MAC MIC-32	Per-hop message integrity code for the entire packet (32 bit)
FCS-16	Frame checksum for the entire packet (16 bit) per 802.15.4

Packet details

PHY preamble

The PHY preamble is used to achieve RF sync between radios and define the length of the packet as specified in IEEE 802.15.4.

MAC header

MAC header contains the fields necessary to deliver packets on a per-hop basis, as well as, timing information for mote synchronization. The MAC header includes:

- Source and destination addresses for current hop
- Network identifier
- Synchronization and joining information

NET header

The network header contains the fields necessary for end-to-end delivery of packets in TSMP network. The NET header includes:

- Source and destination addresses of communicating nodes
- Packet priority

- Routing information

App payload

Application payload is a variable-sized portion of the packet containing the actual payload of commands and/or sensing data. The payload is always sent end-to-end encrypted with 128-bit key.

APP MIC-32

APP MIC-32 is a Message Integrity Code used for end-to-end authentication of payload. This APP MIC ensures that the packet is not tampered with as it is routed from node to node.

MAC MIC-32

MAC MIC-32 is a Message Integrity Code used on a per-hop basis. This field is recalculated by sender prior to being sent on the radio. The MAC MIC ensures that illegal packets can not be injected into the network and packets can not be falsely acknowledged.

FCS-16

Every packet contains FCS-16 checksum field as specified in 802.15.4. This field ensures that corrupt packets are not processed by communicating nodes.