INTERNATIONAL STANDARD

**ISO
20205**

First edition
2015-08-15

Space data and information transfer systems — Spacecraft Onboard Interface Systems — Low Data-Rate Wireless Communications for Spacecraft Monitoring and Control

Systèmes de transfert des informations et données spatiales — Services d'interfaces à bord des véhicules spatiaux — Communication sans fil à faible débit de données pour la surveillance et le contrôle des véhicules spatiaux

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of technical committees is to prepare International Standards. Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75 % of the member bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights.

ISO 20205 was prepared by the Consultative Committee for Space Data Systems (CCSDS) (as CCSDS 882.0-M-1, May 2013) and was adopted (without modifications except those stated in clause 2 of this International Standard) by Technical Committee ISO/TC 20, *Aircraft and space vehicles*, Subcommittee SC 13, *Space data and information transfer systems*.

![CCSDS logo]

The Consultative Committee for Space Data Systems

Recommendation for Space Data System Practices

SPACECRAFT ONBOARD INTERFACE SYSTEMS—LOW DATA-RATE WIRELESS COMMUNICATIONS FOR SPACECRAFT MONITORING AND CONTROL

RECOMMENDED PRACTICE

CCSDS 882.0-M-1

MAGENTA BOOK
May 2013

AUTHORITY

Issue:	Recommended Practice, Issue 1
Date:	May 2013
Location:	Washington, DC, USA

This document has been approved for publication by the Management Council of the Consultative Committee for Space Data Systems (CCSDS) and represents the consensus technical agreement of the participating CCSDS Member Agencies. The procedure for review and authorization of CCSDS documents is detailed in *Organization and Processes for the Consultative Committee for Space Data Systems* (CCSDS A02.1-Y-3), and the record of Agency participation in the authorization of this document can be obtained from the CCSDS Secretariat at the address below.

This document is published and maintained by:

CCSDS Secretariat
Space Communications and Navigation Office, 7L70
Space Operations Mission Directorate
NASA Headquarters
Washington, DC 20546-0001, USA

STATEMENT OF INTENT

The Consultative Committee for Space Data Systems (CCSDS) is an organization officially established by the management of its members. The Committee meets periodically to address data systems problems that are common to all participants, and to formulate sound technical solutions to these problems. Inasmuch as participation in the CCSDS is completely voluntary, the results of Committee actions are termed **Recommendations** and are not in themselves considered binding on any Agency.

CCSDS Recommendations take two forms: **Recommended Standards** that are prescriptive and are the formal vehicles by which CCSDS Agencies create the standards that specify how elements of their space mission support infrastructure shall operate and interoperate with others; and **Recommended Practices** that are more descriptive in nature and are intended to provide general guidance about how to approach a particular problem associated with space mission support. This **Recommended Practice** is issued by, and represents the consensus of, the CCSDS members. Endorsement of this **Recommended Practice** is entirely voluntary and does not imply a commitment by any Agency or organization to implement its recommendations in a prescriptive sense.

No later than three years from its date of issuance, this **Recommended Practice** will be reviewed by the CCSDS to determine whether it should: (1) remain in effect without change; (2) be changed to reflect the impact of new technologies, new requirements, or new directions; or (3) be retired or canceled.

In those instances when a new version of a **Recommended Practice** is issued, existing CCSDS-related member Practices and implementations are not negated or deemed to be non-CCSDS compatible. It is the responsibility of each member to determine when such Practices or implementations are to be modified. Each member is, however, strongly encouraged to direct planning for its new Practices and implementations towards the later version of the Recommended Practice.

FOREWORD

This document is a CCSDS Recommended Practice, which is the consensus result as of the date of publication of the Best Practices for low data-rate communication systems for spacecraft monitor and control in support of space missions.

Through the process of normal evolution, it is expected that expansion, deletion, or modification of this document may occur. This Recommended Practice is therefore subject to CCSDS document management and change control procedures, which are defined in the *Organization and Processes for the Consultative Committee for Space Data Systems* (CCSDS A02.1-Y-3). Current versions of CCSDS documents are maintained at the CCSDS Web site:

http://www.ccsds.org/

Questions relating to the contents or status of this document should be addressed to the CCSDS Secretariat at the address indicated on page i.

At time of publication, the active Member and Observer Agencies of the CCSDS were:

Member Agencies

– Agenzia Spaziale Italiana (ASI)/Italy.
– Canadian Space Agency (CSA)/Canada.
– Centre National d'Etudes Spatiales (CNES)/France.
– China National Space Administration (CNSA)/People's Republic of China.
– Deutsches Zentrum für Luft- und Raumfahrt e.V. (DLR)/Germany.
– European Space Agency (ESA)/Europe.
– Federal Space Agency (FSA)/Russian Federation.
– Instituto Nacional de Pesquisas Espaciais (INPE)/Brazil.
– Japan Aerospace Exploration Agency (JAXA)/Japan.
– National Aeronautics and Space Administration (NASA)/USA.
– UK Space Agency/United Kingdom.

Observer Agencies

– Austrian Space Agency (ASA)/Austria.
– Belgian Federal Science Policy Office (BFSPO)/Belgium.
– Central Research Institute of Machine Building (TsNIIMash)/Russian Federation.
– China Satellite Launch and Tracking Control General, Beijing Institute of Tracking and Telecommunications Technology (CLTC/BITTT)/China.
– Chinese Academy of Sciences (CAS)/China.
– Chinese Academy of Space Technology (CAST)/China.
– Commonwealth Scientific and Industrial Research Organization (CSIRO)/Australia.
– CSIR Satellite Applications Centre (CSIR)/Republic of South Africa.
– Danish National Space Center (DNSC)/Denmark.
– Departamento de Ciência e Tecnologia Aeroespacial (DCTA)/Brazil.
– European Organization for the Exploitation of Meteorological Satellites (EUMETSAT)/Europe.
– European Telecommunications Satellite Organization (EUTELSAT)/Europe.
– Geo-Informatics and Space Technology Development Agency (GISTDA)/Thailand.
– Hellenic National Space Committee (HNSC)/Greece.
– Indian Space Research Organization (ISRO)/India.
– Institute of Space Research (IKI)/Russian Federation.
– KFKI Research Institute for Particle & Nuclear Physics (KFKI)/Hungary.
– Korea Aerospace Research Institute (KARI)/Korea.
– Ministry of Communications (MOC)/Israel.
– National Institute of Information and Communications Technology (NICT)/Japan.
– National Oceanic and Atmospheric Administration (NOAA)/USA.
– National Space Agency of the Republic of Kazakhstan (NSARK)/Kazakhstan.
– National Space Organization (NSPO)/Chinese Taipei.
– Naval Center for Space Technology (NCST)/USA.
– Scientific and Technological Research Council of Turkey (TUBITAK)/Turkey.
– Space and Upper Atmosphere Research Commission (SUPARCO)/Pakistan.
– Swedish Space Corporation (SSC)/Sweden.
– United States Geological Survey (USGS)/USA.

DOCUMENT CONTROL

Document	Title	Date	Status
CCSDS 882.0-M-1	Spacecraft Onboard Interface Systems—Low Data-Rate Wireless Communications for Spacecraft Monitoring and Control, Recommended Practice, Issue 1	May 2013	Current issue

CONTENTS

CONTENTS (continued)

1 INTRODUCTION

1.1 PURPOSE

This document presents the recommended practices for the utilization of low data-rate wireless communication technologies in support of spacecraft ground testing and flight monitoring and control applications. Relevant technical background information can be found in reference [3].

The recommended practices contained in this document enable member agencies to select the best option(s) available for interoperable wireless communications in the support of spacecraft monitoring and control applications. The specification of a Recommended Practice facilitates interoperable communications and forms the foundation for cross-support of communication systems between separate member space agencies.

1.2 SCOPE

This Recommended Practice is targeted towards monitoring and control systems, typically low data-rate and low-power wireless-based applications.

1.3 APPLICABILITY

This Recommended Practice specifies protocols (including at least the Physical [PHY] layer and Medium Access Control [MAC] sublayer of the Open Systems Interconnection [OSI] Model—see reference [F1]) that enable a basic interoperable wireless communication system to support low data-rate spacecraft monitoring and control applications.

1.4 RATIONALE

From an engineering standpoint, mission managers, along with engineers and developers, are faced with a plethora of wireless communication choices, both standards-based and proprietary. This Recommended Practice provides guidance in the selection of systems necessary to achieve interoperable communications in support of wireless, low data-rate monitoring and control.

1.5 DOCUMENT STRUCTURE

This document is composed from a top-down (technology) perspective, first defining the technology as a recommended practice, then providing informative material supporting specific application profiles. (For more information on space mission use cases addressed by wireless technologies, see reference [3]).

Section 2 provides an informational overview of the rationale and benefits of spacecraft onboard wireless technologies for use in spacecraft monitoring and control operations.

Section 3 provides recommended practices and applicable standards relating to low data-rate wireless communication systems.

Annex A justifies the choice of an alternative, scheduled medium access scheme.

Annex B discusses security considerations related to the specifications in this document.

Annex C provides an informative description of the recommended practices, through an overview of the technologies, and a set of application profiles where the recommendations are applicable.

Annex D provides justification for selection of the 2.4 GHz band.

Annex E lists abbreviations used in this document along with their expanded forms.

Annex F provides a list of informative references.

1.6   DEFINITIONS

low data-rate: 250 kbps or less.

NOTE – In general the definition of low data-rate is somewhat ambiguous; for this Recommended Practice it is specified as 250 kbps.

low power: 10 mW or less (typical).

quality of service, QoS: The ability to provide different priority to different applications, users, or data flows, or to guarantee a certain level of performance to a data flow.

1.7   CONVENTIONS

1.7.1   NOMENCLATURE

The following conventions apply for the normative specifications in this Recommended Practice:

a) the words 'shall' and 'must' imply a binding and verifiable specification;

b) the word 'should' implies an optional, but desirable, specification;

c) the word 'may' implies an optional specification;

d) the words 'is', 'are', and 'will' imply statements of fact.

NOTE – These conventions do not imply constraints on diction in text that is clearly informative in nature.

1.7.2 INFORMATIVE TEXT

In the normative section of this document, informative text is set off from the normative specifications either in notes or under one of the following subsection headings:

– Overview;

– Background;

– Rationale;

– Discussion.

1.8 REFERENCES

The following publications contain provisions, which through reference in this text, constitute provisions of this document. At the time of publication, the editions indicated were valid. All publications are subject to revision, and users of this document are encouraged to investigate the possibility of applying the most recent editions of the publications indicated below. The CCSDS Secretariat maintains a register of currently valid CCSDS publications.

[1] *IEEE Standard for Local and Metropolitan Area Networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)*. IEEE Std 802.15.4a™-2011. New York: IEEE, 2011.

[2] *Wireless Systems for Industrial Automation: Process Control and Related Applications*. ISA-100.11a-2011. Durham, North Carolina: ISA, 2011.

[3] *Wireless Network Communications Overview for Space Mission Operations*. Report Concerning Space Data System Standards, CCSDS 880.0-G-1. Green Book. Issue 1. Washington, D.C.: CCSDS, December 2010.

2 OVERVIEW

2.1 RATIONALE AND BENEFITS

Monitoring and controlling the behavior of a spacecraft and launch systems, during testing phases on ground or during nominal operations in orbit, is the key to ensuring the correct functioning of various onboard systems and structures, the responses of these systems in their operational working environments, and the long-term reliability of the spacecraft. These data are also highly significant when compiling lessons learned that will be applied to building better space systems and increasing the reliability of future space components. (Refer to reference [3] for a comprehensive overview of application domains and for a detailed summary of RF communications and restrictions in differing operational environments.)

The quantity of acquired spacecraft functional data depends on the ability to monitor required parameters at precise locations within a given project time and cost envelope. Hundreds and often thousands of data measurement locations are required, steadily increasing the mass (acquisition systems, cables, and harnesses) and the project costs and time (installation and verification of each new sensor).

The use of wireless technologies is foreseen to reduce the integration effort, cost, and time typically required to instrument a high number of physical measurement points on a space structure. Technicians should need less time to integrate and verify their installations, while the risk of mechanically damaging interfaces during the process should be reduced. Large structures should see health monitoring equipment mass reduced, while last-minute changes in the instrumentation (e.g., addition/removal of sensing nodes at measurement points) should be easier to accept at project level. One of the byproducts of using wireless technologies in space systems is the extra flexibility introduced when implementing wireless fault-tolerance and redundancy schemes.

An overriding consideration in this document is the desire to provide recommendations that utilize wireless technology to augment the *overall networking infrastructure* in a spacecraft rather than to provide dedicated data transport to particular end-to-end application-specific subsystems. That is, although the recommendations specified in this document are related to relatively small-scale Personal Area Networks (PANs) rather the more familiar Local Area Networks (LANs) such as Ethernet, the desire is for wireless PANs to function as natural extensions of the backbone LAN. This implies in particular that the recommendations specified herein focus on providing wireless data transport across the lower levels of the OSI model (PHY and MAC) and not on achieving higher-level application-specific behavior.

2.2 SCOPE OF INTEROPERABILITY

The intent of the recommended practices promulgated in this book is to provide a framework for establishing a scalable wireless infrastructure for low-rate data transport that will (1) support traffic generated by diverse sensor types, multiple application-specific devices, and devices supplied by multiple different vendors and (2) facilitate operation of multiple wireless networks in the same bandwidth with minimal interference. The recommended

practices will ensure interoperability of low data-rate wireless devices on a common network at the PHY layer and MAC sublayer so that data packets generated by new devices entering the network will be transported by the existing network devices without regard to the sensor or application that generated the data in the packet payload. In its current form, the book's recommendations should allow new nodes to enter a star topology network and begin communicating with a gateway. Should future revisions augment the current recommendations to allow for transport mechanisms such as peer-to-peer communication and multi-hop relaying, new nodes entering the network will not only be able to transmit their own data to a gateway, but they may also be able to communicate with other nodes and to transport data for other network devices.

Adherence to these recommended practices will promote interoperability of the low data-rate wireless networks addressed in this document with other wireless networks using the same bandwidth via the interference mitigation techniques encompassed by the recommendations.

2.3   EVOLUTION OF THE BOOK

The current version of this document specifies two recommended practices for low data-rate spacecraft monitoring and control. Functionally, the current recommendations can be regarded as pertaining only to the behavior of the network at the PHY layer and MAC sublayer of the OSI network stack, not at the Logical Link sublayer or higher. This level of detail in the recommendations is in line with the philosophy discussed in 2.1 above, that the recommended behavior of wireless networks should be specified only at the lower layers of the network stack (similar to the behavior specified for the backbone network in the spacecraft), leaving higher-layer behavior at the discretion of system designers.

Furthermore, the two recommended practices specified in the current version of the document are restricted to a subset of the network functionality generally supported by the PHY and MAC layers of the OSI stack: one for single-hop contention-based access within a star topology and one for single-hop scheduled access within a star topology. Hence, both recommendations provide a mechanism for data packets to be exchanged between a network coordinator or gateway and individual nodes on the wireless network, but they do not address a mechanism for data packets to be exchanged between two non-coordinator nodes in the network or for communication between any two nodes via intermediary nodes in a multi-hop fashion. The evolution of this document is foreseen to propose additional recommended practices for anticipated application profiles, such as recommended practices for peer-to-peer communication in both mesh and star topologies and for multi-hop data transport in mesh topologies.

The current recommendations also do not address a mechanism for exchanging data packets between a node on the network and a device outside of the wireless network. It is assumed that the network coordinator or gateway will somehow be able to communicate with the backbone network of the spacecraft, but the mechanisms for that, which are typically implemented at the Network (NWK) Layer of the stack, are beyond the scope of the current document and are not discussed. Similarly, the recommendations do not discuss or provide mechanisms for end-to-end acknowledgement or re-transmission of data packets sent

between user applications. The mechanisms for that behavior are typically implemented at the Transport or the Application (APP) Layer of the stack and once again are beyond the scope of the current document. While it is anticipated that future recommendations may address some functionality at the NWK layer, such as routing of Internet Protocol (IP) packets within the wireless network, it is not anticipated that protocol behavior above the NWK layer (such as any APP-layer functionality) will be addressed by future recommendations.

2.4 DIFFERENTIATING CONTENTION-BASED AND SCHEDULED CHANNEL ACCESS

There are two predominant types of medium-access schemes currently utilized in wireless sensor networks: *random* or *contention-based* access and *scheduled* access (see reference [F2]). Contention-based schemes require no centralized control of network access and are thus well suited for ad-hoc network architectures as well as other situations where it is desirable to minimize network administration overhead and operational complexity. Nodes are allowed to attempt channel access at arbitrary times in an ad-hoc fashion as dictated by local data traffic flow and must therefore contend with one another for access in a fairly random manner. The most common contention-based access technique utilized in sensor networks is Carrier-Sense Multiple Access (CSMA) with Collision Avoidance (CA), generally abbreviated as CSMA-CA or simply CSMA. In contrast, scheduled access schemes require some type of (generally centralized) control mechanism for coordinating network access for all nodes in the network in a synchronized fashion. Typically, this will be based on predetermined or anticipated traffic flow so that bandwidth is available in a predictable manner that precludes contention among the nodes. This approach increases network administrative overhead and operational complexity but facilitates QoS guarantees and deterministic network behavior. The most common scheduled access technique utilized in sensor networks is Time-Division Multiple Access (TDMA).

In terms of application support, CSMA is best suited for situations where tight bounds on packet latency and packet jitter are not required but nodes may sometimes require relatively large amounts of available channel bandwidth for relatively short periods of time in a relatively unpredictable manner. CSMA does not readily support *deterministic* network behavior but does readily support *bursty* and *aperiodic* traffic flow. In contrast, TDMA is well suited for applications requiring much tighter bounds on packet latency and jitter but for which the traffic flow from the nodes is more uniform and predictable. TDMA readily supports deterministic network behavior but is generally better suited for applications with less bursty and more periodic traffic flow. In addition, interference avoidance schemes such as frequency hopping are far more easily implemented in a scheduled TDMA MAC sublayer than in a contention-based CSMA MAC sublayer. The same applies to maintaining connectivity in a mesh network topology that supports multi-hop relay traffic with battery powered nodes on a low duty cycle (long sleep period, short active period), although multi-hop transport is beyond the scope of the current Recommended Practice.

2.5 SECURITY PROVISIONING

Wireless networks suffer the maladies of both active tampering and passive eavesdropping due to the inherent nature of wireless communications where access to the transmission media is not a physical constraint as within wired communications. In addition, wireless sensors have severely limited computational processing power and may have no available onboard data storage. Because of the computational complexity of cryptographic algorithms, coupled with the limited battery-based lifetime of a wireless sensor node, security provisioning in these types of devices is a pragmatic engineering balance.

The cryptographic mechanism in this standard is based on symmetric-key cryptography and uses keys that are provided by higher-layer processes; the mechanism assumes a secure implementation of cryptographic operations and secure and authentic storage of keying material (reference [1]). For the recommended practice contained in this document, the PHY/MAC layer provides services that support data confidentiality, data integrity (authenticity), and replay protection:

Table 2-1: PHY/MAC Security Service Provisioning

Security service	Description
Data confidentiality	Transmitted information is disclosed only to parties for which it is intended
Data integrity	Assurance of the source of transmitted information (and, hereby, that information was not modified in transit)
Replay protection	Assurance that duplicate information is detected

NOTE – Per annex B some of the required security architectural elements may be implemented at higher layers (e.g., key management) in the OSI stack and are not strictly defined, or implemented, at the PHY/MAC layer.

2.6 QUALITY OF SERVICE PROVISIONING

Both of the recommended practices prescribed in 3.2 provide support for implementing QoS provisioning so that system designers can implement their QoS policies over the wireless network.

In the 802.15.4 CSMA-CA operational mode, QoS primitive operations are achieved utilizing Guaranteed Time Slots (GTS) as shown in figure 2-1. Briefly, the active portion of the superframe is composed of a beaconing period, a Contention Access Period (CAP) and a

Contention Free Period (CFP); the slotted CSMA scheme is utilized during the CAP, and the GTS scheme is utilized during the CFP period.
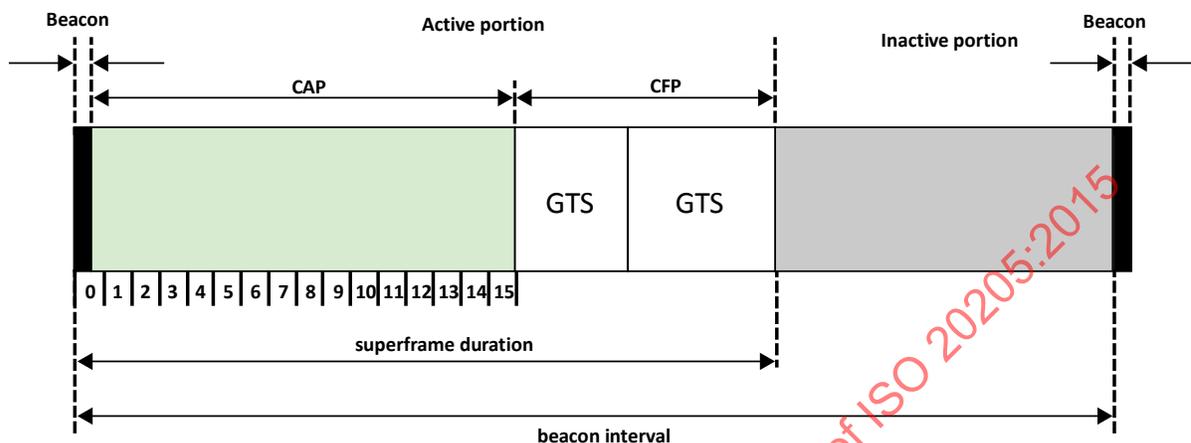
Figure 2-1: IEEE 802.15.4 Superframe

The GTS scheme enables bandwidth reservation between an 802.15.4 PAN coordinator and a PAN device. Notably, more sophisticated QoS schemes that attempt to enforce either some type of fairness for all nodes in the network and/or to handle nodes entering and leaving the network are advanced functionality that is typically implemented at the higher Network (NWK) Layer of the communications stack.

In the 802.15.4 scheduled medium-access operational mode, which is TDMA-based, the available TDMA slots are analogous to CSMA GTS slots during the CFP. Integrated communication stacks based on 802.15.4 (e.g., ZigBee, ISA100, 6LoWPAN, 802.15.4e—see reference [F7]) all enable deployment-wide QoS at the NWK layer. (Refer to annex A for additional QoS provisioning provided in the ISA100.11a recommendation.)

3 RECOMMENDED PRACTICES FOR LOW DATA-RATE WIRELESS COMMUNICATIONS FOR SPACECRAFT MONITORING AND CONTROL

3.1 OVERVIEW

This section presents the recommended practices for *spacecraft monitoring and control applications using low data-rate wireless communication technologies*. (See table C-2 for a non-exhaustive set of example use-cases that may benefit from using low data-rate wireless communications.)

As discussed in section 2, in order to ensure the most basic interoperability between low data-rate wireless communication devices, the current recommendations are focused on specification of functionality at the air interface PHY layer and the MAC sublayer of the OSI model. Following this guideline, two different compliant systems would thus be able to share the medium and potentially join the same wireless network.

3.2 RECOMMENDED PRACTICES

3.2.1 APPLICATIONS SUITED FOR SINGLE-HOP CONTENTION-BASED COMMUNICATIONS

For spacecraft monitoring and control activities employing low data-rate contention-based wireless communications in single-hop configurations, both the air interface PHY layer and the MAC sublayer shall comply with the IEEE 802.15.4-2011 specification (reference [1]).

Single-hop contention-based communication networks and devices **should** utilize the 2.4 GHz frequency band. (See annex D for rationale pertaining to 2.4 GHz band preferences; see reference [3] for Electromagnetic Interference (EMI) considerations of the 2.4 GHz frequency band.)

3.2.2 APPLICATIONS SUITED FOR SINGLE-HOP SCHEDULED MEDIUM-ACCESS COMMUNICATIONS

For spacecraft monitoring and control activities employing low data-rate communications utilizing a scheduled medium-access scheme in a single-hop configuration, both the air interface PHY layer and the MAC sublayer shall comply with the ISA100.11a-2011 PHY-layer and MAC-sublayer specifications (reference [2]).

3.2.3 RESTRICTIONS/HAZARDS

When selecting a wireless technology for application in a spacecraft environment, the risks associated with the selected radio frequency band, transmission power level, and physical location should be taken into account for the following governing environmental factors:

a) Operation in explosive environments;

b) RF exposure levels in excess of governmental limits (see annex D);

c) Electromagnetic Compatibility (EMC).

ANNEX A

JUSTIFYING THE SCHEDULED MEDIUM ACCESS RECOMMENDATION

(INFORMATIVE)

A1   BACKGROUND

From its introduction 2003, application of IEEE 802.15.4 to embedded sensing tasks has been steadily increasing. Use has been largely limited to home and office automation, however, since it has been found that 802.15.4 reliability suffers as the RF complexity of the environment in which it is deployed increases. Specifically, industrial deployments of 802.15.4 are often observed to exhibit unacceptably low reliability and high latencies. Amendments incorporated in the IEEE 802.15.4-2006 revision recommended, and even those subsequently incorporated in the 802.15.4-2011 revision, have failed to address these concerns adequately, leading to native IEEE 802.15.4's being widely considered a poor solution for process monitoring and control in harsh industrial environments.

This discrepancy is documented in the IEEE 802.15.4e-2012 amendment, which incorporates a scheduled MAC layer very similar to the ISA100.11a MAC recommended in this Magenta Book. In justifying the update to 802.15.4-2011 provided by the 802.15.4e amendment, the IEEE states that "this amendment to IEEE Std 802.15.4-2011 specifies additional media access control (MAC) behaviors and frame formats that allow IEEE 802.15.4 devices to support a wide range of industrial and commercial applications that were not adequately supported prior to the release of this amendment."   It goes on to observe "industrial applications (and some commercial applications) have critical requirements such as low latency, robustness in the harsh industrial RF environment, and determinism that are not adequately addressed by IEEE Std 802.15.4-2011" (reference [F7]).

Given that many spaceflight applications have constraints on reliability and latency similar to those in industrial process control, it was determined that this Magenta Book required a recommendation that rectifies many of the shortcomings in IEEE 802.15.4-2006 (and later, 802.15.4-2011) that the IEEE itself recognizes. Unfortunately, IEEE 802.14.4e-2012 is new enough that there are not sufficient commercial parts available for testing to justify its inclusion in the present edition of this Magenta book, although future editions may adopt it as the scheduled MAC recommendation.

Instead, the ISA100.11a MAC, which along with the WirelessHART standard inspired the IEEE 802.15.4e-2012 recommendation, is adopted here due to the availability of radios for testing. Indeed, testing by the authors of this Magenta Book has confirmed the relative robustness of ISA100.11a and the relative weakness of IEEE 802.15.4 in the presence of Wi-Fi interference (reference [F8]). ISA100.11a is chosen over WirelessHART since it is capable of supporting a greater variety of application layers and is in general more customizable. (For a detailed comparison of the two, see reference [F9].)

A2   MECHANISMS FOR INCREASED ROBUSTNESS

ISA100.11a provides a number of mechanisms for increasing the overall quality of service. As mentioned in 2.4, the scheduled ISA100.11a MAC provides greater determinism for channel access. Time in an ISA100.11a network is divided into slots, and a time distribution mechanism embedded in packet acknowledgements ensures that radios keep slot boundaries synchronized with respect to neighboring radios to ensure coordinated transmission/reception between communicating pairs.

A Network Manager overseeing operation of the ISA100.11a network allocates communication opportunities to radios requesting bandwidth, thereby ensuring time diversity within the ISA100.11a network. That is, an individual radio within the network will only attempt to use the wireless medium in a time slot assigned to that radio for transmission. If the attempt fails for any reason (e.g., excessive RF interference), the transmission will be retried at the radio's next scheduled opportunity. Furthermore, frequency diversity is added by the Network Manager, assigning one of the up to 16 channels available under the 802.15.4 2.4 GHz DSSS PHY employed by ISA100.11a to the communication attempt. Should a retransmission be required, the next scheduled attempt will be on a different channel, drawn from a predetermined channel-hopping sequence. Time synchronization between radios allows the Network Manager to configure the receiving radio to have its receiver tuned to the channel of the transmitter for the scheduled transaction. ISA100.11a supports adaptive blacklisting, so that channels on which communication attempts repeatedly fail can be removed from sending and receiving radios' channel hopping sequences. For multi-hop topologies, spatial diversity is also added through the use of routing graphs with redundant next-hop paths, although that is outside the scope of this Magenta Book's current recommendation for single-hop communication.

These diversity features, taken together, enhance the ability of ISA100.11a to coexist with other RF systems that are acting as interferers. It should also be noted that, since ISA100.11a uses the 2.4 GHz DSSS PHY specified in IEEE 802.15.4-2006, it inherits the benefits of the Clear Channel Assessment (CCA) service used by the CSMA MAC of 802.15.4, which in particular promotes non-interference of the ISA100.11a radios with other systems operating in the same RF band.

ANNEX B

SECURITY CONCERNS FOR WIRELESS SYSTEMS

(INFORMATIVE)

B1 INTRODUCTION

The 802.15.4 and ISA100.11a specifications recommended in this book describes RF wireless PHY-layer and MAC-sublayer protocols for low-power and relatively low data-rate networked communications. These specifications support a diverse application domain; wireless applications for space operations can benefit from the security features provided in these PHY-layer/MAC-sublayer protocol specifications.

Communications security attempts to ensure the confidentiality, integrity, and/or authenticity of transmitted data, as required depending on the threat, the mission security policy(s), and the desire of the mission planners. It is possible, and often likely, to require all three of these security attributes to ensure that the communications data payload is not disclosed, not altered, and not spoofed.

Specific potential threats and attack scenarios relevant to the space flight operations domain are addressed in more detail in reference [F5]. A summary of typical space agency use cases is provided in reference [3]. Threats and attack scenarios for ground segment operations, i.e., AIT activities, are typical in scope to general terrestrial security concerns.

B2 GENERAL RISKS

A MAC-sublayer security protocol provides four basic security services: access control, message integrity, message confidentiality, and replay protection (reference [F6]).

Access control and message integrity. Access control means the MAC-sublayer protocol should prevent unauthorized parties from participating in the network. Legitimate nodes should be able to detect messages from unauthorized nodes and reject them. Also, a secure network should provide message integrity protection: if an adversary modifies a message from an authorized sender while the message is in transit, the receiver should be able to detect this tampering. Including a message authentication code with each packet provides message authentication and integrity (reference [F6]).

Confidentiality. Confidentiality means keeping information secret from unauthorized parties. It is typically achieved with encryption. Preferably, an encryption scheme should not only prevent message recovery, but also prevent adversaries from learning even partial information about the messages that have been encrypted (reference [F6]).

Replay Protection. An adversary that eavesdrops on a legitimate message sent between two authorized nodes and replays it at some later time engages in a replay attack. Since the

message originated from an authorized sender it will have a valid message authentication code, so the receiver will accept it again. Replay protection prevents these types of attacks (reference [F6]).

The 802.15.4 security implementation is handled at the MAC sublayer, below application control. The application specifies its security requirements by setting the appropriate control parameters in the radio stack. If an application does not set any parameters, then security is not enabled by default. That is, when using 802.15.4, an application must explicitly enable security. As discussed in C1.3, security is implemented automatically in ISA100.11a.

B3   SECURITY CONCERNS

Several security concerns specific to the IEEE 802.15.4 design are addressed in more detail in reference [F6]; the reader is strongly encouraged to review the identified shortcomings. It may be necessary to apply security services at multiple layers within the protocol stack, to account for distributed processing and cross-support, to account for different classes of data or end users, or to account for protection of data during unprotected portions of the complete end-to-end transmission (e.g., across ground networks). The specification of security services at other layers is outside the scope of this document.

B4   CONSEQUENCES OF NOT APPLYING SECURITY

Without authentication, unauthorized commands or software might be uploaded to a spacecraft or data received from a source masquerading as the spacecraft. Without data integrity, corrupted commands or software might be uploaded to a spacecraft potentially resulting in the loss of the mission. Without data integrity, corrupted telemetry might be retrieved from a spacecraft that could result in an incorrect course of action being taken. If confidentiality is not implemented, data flowing to or from a spacecraft might be visible to unauthorized entities resulting in disclosure of sensitive or private information.

ANNEX C

DISCUSSION ON LOW DATA-RATE
WIRELESS COMMUNICATIONS FOR
SPACECRAFT MONITORING AND CONTROL

(INFORMATIVE)

C1   GENERAL

C1.1   OVERVIEW

The following subsections contain engineering discussions applicable to the recommended practices in section 3.

C1.2   CONTENTION-BASED CHANNEL-ACCESS MECHANISM

As discussed in 2.4, the operation of a contention-based channel-access mechanism cannot readily support packet delivery with reliably low and predictable latency in many situations, particularly when the number of active nodes in the network grows to even moderate levels. As such, it is generally not appropriate for use in situations requiring deterministic or 'real-time' behavior, such as spacecraft control guidance and navigation loops or life-critical applications.

Similarly, the 802.15.4 MAC sublayer specified in Recommended Practice 3.2.1 provides no specific mechanisms for adaptive channel selection or interference avoidance. The recommendation as stated presumes operation on a single, predetermined subchannel of the 2.4 GHz ISM band and persistent interference on the selected channel will lead to substantial performance degradation. Mechanisms for detecting and avoiding such interference, if necessary, must be implemented at higher layers of the protocol stack. As such, the current recommended practice may not be well suited for operation in a very cluttered spectral environment with many different wireless systems contending for the same bandwidth. Additionally, the environment may induce interference effects such as multi-path fading. When these effects are time-varying and not well characterized a priori, the current recommended practice may not be well suited. Conversely, the current recommended practice can be expected to work very well in environments for which the available spectrum is well understood over time and carefully managed..

Although Recommended Practice 3.2.1 applies specifically only to single-hop communication between a client node and the network coordinator in a star network topology, the recommended standard protocol, IEEE 802.15.4, will also support communication modes and services not specifically addressed in the recommendation if appropriate functionality is provided by higher layers of the protocol stack. For example, while considerations regarding communication security are beyond the scope of the current

recommendation, the 802.15.4 MAC sublayer specification defines encryption and decryption services for symmetric-key cryptographic techniques that will support secure communication if procedures for establishing and maintaining the necessary keys are provided by higher layers. Similarly, while peer-to-peer and multi-hop communication within an arbitrary mesh topology are beyond the scope of the current recommendation, the 802.15.4 MAC sublayer specification does provide support for these communication modes if the necessary routing and synchronization mechanisms are implemented in higher layers of the protocol stack.

These limitations and capabilities must be understood and considered carefully when making an engineering decision regarding the applicability of a contention-based access mechanism in general or recommendation 3.2.1 in particular.

C1.3   SCHEDULED CHANNEL-ACCESS MECHANISM

A scheduled channel-access mechanism requires a method for synchronizing transmissions/receptions among the nodes in the network. Furthermore, the ISA100.11a recommendation allows nodes to switch among the 16 available channels in the 802.15.4 2.4 GHz PHY with each subsequent transmission attempt, coordinating transmitters and receivers so that they both use the same channel at the same time. As discussed in annex A, a centralized Network Manager entity is required to establish this 'channel hopping' mechanism for each node in the network and mediate bandwidth usage through granting communication 'contracts' to nodes.

The Network Manager is the key to an ISA100.11a network's operation and is its most complicated component. A Network Manager is constantly optimizing the channel-hopping scheme in response both to nodes' requests for communication bandwidth and nodes' reports of the channel qualities in their individual locations. Implementing this functionality from scratch, while possible, may prove time-consuming and it may be more feasible to employ a pre-certified ISA100.11a Network Manager. This, however, comes with a caveat: ISA100.11a is designed as a complete networking solution for high-reliability industrial process monitoring and control. As a result, an ISA100.11a-compliant Network Manager functions on all levels (PHY through APP) of the OSI model. To achieve the PHY- and MAC-sublayer behavior specified in this Recommended Practice, the use of a complete ISA100.11a stack configured so that behavior at layers above the MAC sublayer is either disabled or transparent to the user is advised. Specifically, the following configuration is recommended:

   a) All nodes, except for the network gateway, should be configured as non-routing devices.

   b) APP layer tunneling should be used to bypass the object-oriented APP layer scheme recommended by ISA100.11a.

Configuration a) results in a star network topology, giving the single-hop behavior mandated in this recommended practice. It reduces functionality at each of the upper Data Link and

NWK layer to a pass-through, since the upper Data Link layer is responsible for multi-hop routing within an ISA100.11a mesh network and the NWK layer is responsible for routing outside of the gateway on the backbone network (a recommendation for which is not covered in this document). Configuration b) reduces functionality at each of the Transport and APP layers to a pass-through as well.

It is worth noting that over-the-air transmissions must be secured in an ISA100.11a network. While security is optional in the 802.15.4 PHY/MAC recommendation, some level of security is required in the ISA100.11a PHY/MAC recommendation implicitly through the use of an ISA100.11a stack configured as directed above. A Security Manager entity joins the Network Manager in a proper ISA100.11a implementation, and its inclusion is non-optional. Messages are encrypted on both a hop-by-hop and end-to-end basis, and distribution and maintenance of encryption keys is handled automatically by the Security Manager.

As such, this Recommended Practice covers secure, single-hop communications. Should a user wish to extend this functionality to multi-hop communication, configuration a) can of course be ignored, but such functionality is outside the scope of the current recommended practice.

It is also worth cautioning the user that ISA100.11a is a relatively resource-heavy protocol with regards to computational complexity at the Network Manager. Network formation will generally take longer compared to the 802.15.4 PHY/MAC recommendation, and support for node mobility will be more limited. The same caveat applies to administrative messages to the nodes from the Network Manager (and vice versa). A greater percentage of available bandwidth will be used to maintain the ISA100.11a network to achieve more efficient use of the remaining bandwidth in contention-based environments. Thus the current Recommended Practice can be expected to work quite well in an environment in which contention for bandwidth from other systems and interference effects are significantly present but not well modeled. Conversely, when the available spectrum is well understood over time and carefully managed, the current Recommended Practice may not be well suited.

C2   APPLICATION PROFILES

C2.1   OVERVIEW

An application profile is an explicit listing of the configuration settings of a typical implementation that may be suitable for multiple use cases or applications. Table C-1 is a quick-look table, which lists the most common application profiles targeted by the two recommendations specified in this document. It should be noted that all of these application profiles are based on a star network topology in which the individual nodes in the network all communicate directly with a central gateway node that aggregates data, disseminates commands, or both. Both the 802.15.4 standard, which is specified in 3.2.1 and the ISA100.11a standard, which is specified in 3.2.2, are well suited for applications based on such a topology and can be expected to work well for both periodic, fixed-length, block data transfer as well as aperiodic, variable-length, bursty data transfer.

Table C-1: Application Profile Quick Look-Up Table

List of application profiles falling under the recommended practice
1. Single-hop periodic data aggregation
2. Single-hop triggered (event-driven) data aggregation
3. Single-hop, latency tolerant command and control or command-driven data aggregation (polling)

Table C-2 presents a set of use-cases that may benefit from using low data-rate wireless communications.

Table C-2: Quick-Look Table for Scenarios That Can Utilize Low Data-Rate Wireless Communications

Use-case	Typical examples
Assembly, Integration and Testing (AIT) / Ground Support Equipment (GSE) / Developmental Flight Instrumentation (DFI) activities	*Thermal chamber testing, vibration testing, data bus monitoring...*
Spacecraft onboard health monitoring	*Temperature and radiation level monitoring, impact detection...*
Scalability / extensibility / retro-fit of instrumentation capabilities	*Instrument replacement, adding capability to existing vehicles...*
Habitat environmental monitoring and control	*Temperature, humidity, pressure monitoring...*
Crew (physiological) monitoring	*Heartbeat, temperature, location...*
Scientific monitoring and control	*Periodic observation of experimental variables...*
Intra-spacecraft robotic activities	*low data-rate positioning telemetry, health data...*

C2.2 SINGLE-HOP PERIODIC DATA AGGREGATION

The single-hop periodic data-aggregation profile covers the most common implementation of a wireless sensor network, one that consists of a central data sink (i.e., a gateway or network coordinator) and a number of child nodes that perform periodic data acquisition. The network is configured in a star topology, with each child node having a direct link to the coordinator. Typically, a child node wakes up from a very low-power (sleep) mode on a predetermined periodic schedule, executes a data acquisition task, formats the acquired data, transmits a data packet to the network coordinator, and then goes back into sleep mode. Alternatively, the acquisition node may sample data during each wake cycle but only transmit data to the coordinator when a full packet's worth of data has been accumulated. The coordinator node, which either never sleeps or sleeps only infrequently, aggregates the data from all of the child nodes and relays it over a backbone network to user applications that consume the data. Generally, the duty cycle of the child nodes is quite low, with data acquired at rates from one observation per second down to one observation every several minutes and children often spending 99 percent or more of their lifetimes in sleep mode. For this profile, the data payload transmitted in each packet is generally small and fixed in size.

Vehicle ground test applications require flexibility in the implementation of the tests and the location and orientation of the nodes and antennas. Hence, it is often the case that all nodes will have omnidirectional antennas rather than directional higher-gain antennas.

The RF transmit power is a very application-specific parameter and heavily depends on the operational environment and on EMI/EMC constraints. Some spacecraft will not allow transmission powers higher than perhaps –15 dBm, while others may permit powers up to 10 dBm. In contrast, for other applications such as structural testing of small components in a laboratory thermal-vacuum chamber, relaxed transmit power constraints are often seen. The permissible transmit power is thus one of the first parameters/constraints to be identified before setting up a wireless sensor network.

The number of acquisition nodes in the wireless network is also very application-dependent. In a typical laboratory testing activity, a few nodes, each with several sensors, may well prove to be enough for the task at hand. Spacecraft testing and monitoring on the other hand may require the utilization of hundreds of wireless nodes.

Table C-3 summarizes the high-level implementation parameters and operational configurations for the periodic data aggregation application profile.

Table C-3: **Typical Operating Parameters for the Single-Hop, Periodic Data Aggregation Application Profile**

Implementation parameter / operational configuration	Typical value
Topology	Star
Antenna type	Typically omnidirectional
Transmit power	Typically -15 dBm to +10 dBm
Typical number of nodes	10 – 100
Antenna Polarization (master/slave)	Linear/linear; circular/linear
Spectrum/Channel utilization	Per IEEE 802.15.4 specifications; spectrum and channel management
Typical communication range	0 – 10 m
Typical transmit periodicity	Seconds to minutes
Expected battery life	Months to years
Typical receiver periodicity	Low
Latency constraints	Typically relaxed
Routing	None
Data payload characteristics	Periodic, fixed-length, uniform rate

C2.3   SINGLE-HOP TRIGGERED EVENT-DRIVEN DATA ACQUISITION

The single-hop triggered-event-driven data-acquisition profile covers an implementation of a wireless sensor network that consists of a central data sink and a number of child nodes that perform non-periodic data acquisition. The network is configured in a star topology, with each child node having a direct link to the coordinator. For this profile, however, a child node wakes up to acquire data only when triggered by the occurrence of some local event rather than on a predetermined periodic schedule. The triggering event is sensed by the child node using a low-power circuit that remains active even in sleep mode. When data collection is triggered, the acquisition node collects some amount of data, which may be either predetermined or based on the length or intensity of the triggering event. The collected data may be transmitted back to the sink in raw form or may be processed locally to reduce the data in some fashion. In either case, the resulting data payload is formatted and transmitted back to the sink via a single packet or subdivided into several sequential packets, as necessary. The coordinator node, which either never sleeps or sleeps only infrequently, aggregates the data from all of the child nodes and relays it over a backbone network to user applications that consume the data. For this profile, the duty cycle of the child nodes is obviously determined by the frequency of triggering events, but is generally extremely low.

RECOMMENDED PRACTICE FOR LOW DATA-RATE WIRELESS COMMUNICATIONS

General considerations regarding antenna configuration, power level, and network size are identical to those discussed in table C-3. Table C-4 summarizes the high-level implementation parameters and operational configurations for the event-driven data aggregation application profile.

Table C-4: Typical Operating Parameters for the Single-Hop Triggered, Event-Driven Data Acquisition Application Profile

Implementation parameter / operational configuration	Typical value
Topology	Star
Antenna type	Typically omnidirectional
Transmit power	Typically -15 dBm to +10 dBm
Typical number of nodes	10 – 100
Antenna Polarization (master/slave)	Linear/linear; circular/linear
Spectrum/Channel utilization	Per IEEE 802.15.4 specifications; spectrum and channel management
Typical communication range	0 – 10 m
Typical transmit periodicity	Event driven
Expected battery life	Months to years
Typical receiver periodicity	Low, depends on beacon and acknowledgement mode
Latency constraints	Typically relaxed
Routing	None
Data payload characteristics	Non-periodic, variable-length, bursty

C2.4   SINGLE-HOP COMMAND AND CONTROL OR COMMAND-DRIVEN DATA AGGREGATION

The single-hop command-and-control or command-driven data-aggregation profile again covers an implementation of a wireless sensor network that consists of a central coordinator and a number of child nodes. In this case, however, the child nodes may acquire data from a sensor, control an actuator, or both. Further, in this profile, data may flow not only from the child node to the coordinator in the form of telemetry or command status, but also from the coordinator to the child node in the form of commands. The network is configured in a star topology, with each child node having a direct, bi-directional link to the coordinator.

For the command-driven data aggregation application, a child node wakes up on a periodic schedule and communicates with the coordinator for a possible command to acquire data. If there is no command waiting, the node goes back into sleep mode. If there is a data acquisition command waiting, the node decodes the command, acquires and formats the requested amount of data, transmits the data back to the coordinator in as many packets as necessary, and goes back into sleep mode. For the command and control application, the child node wakes up on a periodic schedule and polls the coordinator for a possible command to change an actuator setting. If there is no command waiting, the node goes back into sleep mode. If there is an actuation command waiting, the node retrieves the command, decodes it, activates an appropriate control signal for the actuator, optionally transmits a command status to the coordinator (e.g., success/failure), and goes back into sleep mode. One could again envision such an operation being conducted in conjunction with the periodic or event-triggered transmissions in C2.2 and C2.3. Should a control algorithm interfacing with the gateway decide a local actuation (e.g., turning on a heater or a fan) is necessary based on measured data (e.g., a temperature reading), a command for that actuation would be sent to the node which measured the data and is capable of actuating the control device.

For either application, the coordinator aggregates the data from all of the child nodes (either telemetry or command status data) and relays it over a backbone network to user applications that consume the data. The coordinator once again sleeps only infrequently. The duty cycle of the child nodes is command-driven, but is generally extremely low.

General considerations regarding antenna configuration, power level, and network size are again identical to those discussed in C2.2. Table C-5 summarizes the high-level implementation parameters and operational configurations for both the command and control and command-driven data aggregation application profiles.