Reticulum Primer

v0.04β

Michael Faragher

17JAN2024

Localization note: Commas (,) are used as thousand separators and periods (.) are used for decimal places. Example: 1,234.56

Document targets RNS manual 0.5.5β and may not align with chapter/section numbers from other versions.

Contents

List of Figures

List of Tables

1 Foreword

Communication is fundamentally human. We celebrate a child's first words, long for the sound of a voice, and exchange letters for both business and pleasure. It is the sharing of information, from the mundane to the profound, from education to entertainment. It is what changes isolation to community.

As the distance grows, so does the infrastructure. Postal systems are, at their core, generalized courier services, which are both relatively slow and labor intensive. Point to point electric communication, such as the telegraph, reduce travel time, but still require transport from the hub to the end user. The revolution of the Internet wasn't just electronic communication, but that each endpoint could communicate directly in near real-time. Wi-Fi and mobile data have allowed small, portable devices to have always-on network access, and finally communication is no longer tethered to a single time or point. So long as these services are available.

Amateur radio has long been used both for communication and emergency response, but is rightfully restricted in its content. As an international and potentially global communication system, it needs to continue its mission without interference. Other unrestricted bands exist, and digital modulation can increase the number of signals in a given bandwidth, but high speed radio networks, like Wi-Fi, are short range and long range networks, like LoRa, are slow. Even cellular data networks require many expensive radios attached to a wired backbone. There are few fault-tolerant, efficient, digital radio network technologies available to the individual.

Reticulum is designed to fill exactly this role. It is not a network, but a tool to make networks, ranging from global to personal. Communication over radio, Internet, LAN, or even serial is supported, and all communication is end to end encrypted. It brings people together over any interface required.

Modern communication systems do an excellent job. Cell service is widespread, the Internet covers 65% of the world population, and the problem isn't finding a communication app, it's choosing one. However, 35% of the population does not have this access, and far more are at the mercy of a failure-intolerant system.

Redundant communication, both in network topology and diversity of transmission media, is key to both safety and community. As the world becomes more interconnected, it's more important than ever to have a good relationship with your neighbors, even if they're a continent away.

Reticulum allows that communication from person to person, over networks anyone can build, from individuals to municipalities. Equally useful for a private chat or search and rescue in a woodland, it puts that fundamental human need for communication in our hand, linking us without the need for an intermediary.

Communication is what makes communities, and Reticulum can both form and serve communities as surely as any language or written word.

1.1 Intent

This document is intended to be a primer and first order reference for a Reticulum network, particularly using an RNode. It is not intended to replace previously published information or provide finalized designs for networks or support systems. The information here is designed to provide both hobbyists and subject matter experts the background and terminology for future research and relating to existing knowledge.

Initially intended as a cookbook, it has grown from describing potential applications to explaining the decisions required to design a system. Simple systems, such as an RNode attached to a single board computer to create a local propagation node or gateway, can be described in their entirety, while more complicated systems are generalized to be purpose built for the end user's specific application.

1.2 Reticulum

Reticulum, is a network protocol designed to create cryptographically secure mesh networks across heterogeneous interfaces. It is specifically designed for low bandwidth, high latency links but is capable of running on high-speed infrastructure, making it possible to both run over packet radio or high speed Ethernet connections. Depending on configuration, these networks can either be end to end encrypted public access networks, or secured, closed networks requiring an access key; there is no provision for unencrypted traffic.

Identities can be generated at will and use an asymmetric encryption to establish an initial link, which then uses symmetric encryption with an ephemeral key to provide both efficient and forward secure communications. Each identity can define endpoints specific to an application and use case, and while these addresses are deterministic, they are the product of one-way functions; a known identity can lead to known addresses, but a known address cannot be used to determine the base identity. By querying the network, it is possible to obtain the identity's public key, both verifying the identity and allowing encrypted traffic.

The *Reticulum Network Stack* (RNS) is a full implementation of the protocol and its support elements. Currently there is a fully functional Python implementation, but no stable releases for other languages. Several implementations are in development but no ETA is yet available.

Further details can be found below and in the comprehensive Reticulum documentation.

1.3 NIMS/ICS

The National Incident Management System / Incident Command System (NIMS/ICS) is the US government's emergency management framework for ad-hoc and lean management systems. The key tenets will be reproduced here to explain the intent of many of the following decisions, but the sprawling standards fill multiple volumes and training courses. Further information can be found in the ICS-200 training course and ICS-219 T-Cards, both of which are available from the Federal Emergency Management Agency (FEMA).

This document is not intended to train for life-saving operations and simply uses the NIMS/ICS framework to propose useful network models based on real-world applications.

1.4 Nomenclature

Much of the terminology is based on NIMS/ICS, meaning such absurd terms as "A strike team of dump trucks" may be used. This is for general interoperability and adherence to printed standards and is not intended to propose any sort of suitability or intent. It is worth noting that this system grew from civil firefighting organizations and is currently used for disaster response and any oddities must be viewed through that framing.

1.5 License

This page intentionally left blank.

2 RNode

An RNode is a term for a category of devices designed to act as LoRa modems for Reticulum connections. Known configurations can be configured using the rnodeconf utility, but compliant custom hardware can be used instead.[1] The core of the system is currently the SX1276 or SX1278, with the intention to expand to the SX1262/8 and SX1280. HopeRF modules are currently untested.

2.1 Coding Rate

Coding Rate (CR) defines the ratio of error correction data versus payload data used for forward error correction. This is a per-message setting, and a radio can decode a message regardless of its CR. This is intended to be used dynamically based on local interference and data corruption.

Annoyingly, it is described in two ways, but there is no overlap in the numbers used, making it less confusing than it first seems. The CR is always defined per four units of payload data. A value of one through four is used to signify the additional data transmitted, while five through eight signify the total sent data per four units of payload data. This provides an overhead of between 25 - 100% due to forward error correction.

Sadly, these numbers are not interchangeable, and are application specific.

CR		Overhead
1	5	25%
2	6	50%
3	7	75%
4	8	100%

Table 2.1: Coding Rate Equivalence

2.2 Spreading Factor

A *Symbol* or *Chirp* is a a single piece of information made up of many *Chips*, and the *Spreading Factor* (SF) is a definition of how many Chips are in a Symbol. A single symbol contains 2^{SF} Chips, providing data spread over a greater portion of the spectrum, improving both receiver sensitivity and resistance to interference. A single Symbol can transmit SF bits of information, meaning a single byte of payload is almost always split between two or more symbols.

[1] http://unsigned.io/guides/2022_01_26_how-to-make-your-own-rnodes.html

For example, at SF8, a single Symbol is capable of sending a byte of data per symbol, but at a Coding Rate of 5, it takes 10 bits to transfer one byte of payload. There is no practical reason to use SF 10 to transfer one byte of payload at CR 5 or SF 12 to transfer one byte at CR 6, but it's mentioned for the sake of curiosity.

2.3 Bandwidth

Bandwidth describes the width of the signal used by the radio. LoRa uses bandwidth as a mechanism to transmit in parallel. The wider the bandwidth, the more data that can be transmitted at once. Since frequency tolerance is proportional to bandwidth, a wider bandwidth also allows the use of lower precision crystals.[2]

2.4 Sensitivity

Sensitivity is the minimum signal strength required to properly detect and process a signal. The raw sensitivity is an important metric for detecting weak signals, but the *signal to noise ratio* (SNR) is more important in noisy environments. SNR, as written, is how powerful a signal is compared to the noise floor.

It's important to note that lower numbers imply better sensitivity and that negative values for SNR indicate that a signal can be detected below the noise floor. LoRa is capable of sensitivities between -112 and -146 dB, depending on the configuration, better than WiFi (-80) or Cellular (-120), but less than GPS-style navigation satellites (-165). See SX126X datasheet for specifics.

2.5 Bitrate

The effective data transfer rate of an RNode is measured in bits per second, or *baud.* The effective data rate is a straightforward calculation given a thorough understanding of the LoRa standard, but is more opaque at first glance. For a simple equation without context, you can skip to equation 2.3.

LoRa transmits one chip per second per Hertz of bandwidth. This means at BW 125, the transmitter sends 125,000 chips per second. Since there are $2^S F$ chips per symbol, and each symbol transmits SF bits of raw data, the raw data rate is represented by equation 2.1. To compensate for error correction overhead, this data rate needs to be adjusted using the percentage of payload data to raw data, as calculated in equation 2.2. Multiplying these two equations results in a base data rate calculation shown in 2.3 which is mathematically equivalent to equation 2.4 which is the form presented in Semtech AN1200.22.

$$baud_{raw} = \frac{BW \cdot SF}{2^{SF}} \tag{2.1}$$

$$PayloadPercent = \frac{4}{4 + CR} \tag{2.2}$$

[2]SF 10, 11, and 12 have additional precision requirements as described in Section 2.6

$$baud = \frac{BW * SF}{2^{SF}} \cdot \frac{4}{4 + CR} \tag{2.3}$$

$$baud = SF \cdot \frac{\frac{4}{4+CR}}{\frac{2^{SF}}{BW}} \tag{2.4}$$

Reticulum generally requires at least 500 baud to function correctly, which provides a floor to acceptable settings. Table 2.2 shows bitrates calculated for CR5, with out of specification bitrates and TCXO only bandwidths marked. See section 2.6 for more information on bandwidths below 62,500 Hz.

BW[3]	Spreading Factor[4]					
(kHz)	7	8	9	10	11	12
7.8	*341*	*195*	*110*	*61*	*34*	*18*
10.4	*455*	*260*	*146*	*81*	*45*	*24*
15.6	682	*390*	*219*	*122*	*67*	*37*
20.8	910	520	*293*	*163*	*89*	*49*
31.25	1370	781	*439*	*244*	*134*	*73*
41.7	1820	1040	586	*326*	*179*	*98*
62.5	2730	1560	879	*488*	*269*	*146*
125.0	5470	3130	1760	977	537	*293*
250.0	10940	6250	3520	1950	1070	586
500.0	21880	12500	7030	3910	2150	1170

Table 2.2: Baud rate at Coding Rate 5

2.6 Frequency Tolerance

The SX127X chip has a maximum allowable frequency difference between sender and receiver. The Datasheet recommends a minimum BW of 62.5 without temperature controlled oscillators (TCXOs) and for maximum compatibility, BW of 125 is recommended to avoid edge cases, especially when a wide range of temperatures is expected.

Onboard crystal tolerance is alluded to as 20 parts per million (ppm) and many devices, including the LilyGO v1.6, lack the ability to use an external oscillator, and will rely on this crystal. A 20 ppm crystal is classified as one with ±20 ppm variance around the central frequency. Using a 915 MHz central frequency, this would be 915,000,000±18,300 Hz.

The required tolerances are ±25% of the bandwidth, except SF 10, 11, and 12 which have an additional limit of 200, 100, and 50 ppm respectively. Table 2.3 shows the required tolerances listed in both Hertz and ppm for convenience's sake.

[3] Boxed bandwidths require TCXO to meet tolerances

[4] SF12 BW125 may require TCXO to operate reliably, boxed baud rates are not within Reticulum specifications

It's worth noting that an individual crystal may be more accurate than specification, but two 20 ppm crystals at the same temperature can have up to 40 ppm deviation from each other, and temperature differences as low as 10-20°C can push a crystal out of specification.

Similarly, since these numbers are based on ratios, a 433 MHz radio with the same ppm deviation would have a smaller absolute deviation, making the special case ppm limits for SF 10 - 12 more relevant. To compare with the example above, a ±20 ppm 433 MHz radio would be 433,000,000±8,660Hz, perhaps allowing BW 41.7 to be used reliably.

BW	Spreading Factor											
	7		8		9		10		11		12	
(kHz)	kHz	ppm	kHz	ppm	kHz	ppm	kHz	ppm	kHz	ppm	kHz	ppm
7.8	1.950	2.1	1.950	2.1	1.950	2.1	1.950	2.1	1.950	2.1	1.950	2.1
10.4	2.600	2.8	2.600	2.8	2.600	2.8	2.600	2.8	2.600	2.8	2.600	2.8
15.6	3.900	4.3	3.900	4.3	3.900	4.3	3.900	4.3	3.900	4.3	3.900	4.3
20.8	5.200	5.7	5.200	5.7	5.200	5.7	5.200	5.7	5.200	5.7	5.200	5.7
31.25	7.813	8.5	7.813	8.5	7.813	8.5	7.813	8.5	7.813	8.5	7.813	8.5
41.7	10.425	11.4	10.425	11.4	10.425	11.4	10.425	11.4	10.425	11.4	10.425	11.4
62.5	15.625	17.1	15.625	17.1	15.625	17.1	15.625	17.1	15.625	17.1	15.625	17.1
125.0	31.250	34.2	31.250	34.2	31.250	34.2	31.250	34.2	31.250	34.2	31.250	34.2
250.0	62.500	68.3	62.500	68.3	62.500	68.3	62.500	68.3	62.500	68.3	45.750	50
500.0	125.000	136.6	125.000	136.6	125.000	136.6	125.000	136.6	91.500	100	45.750	50

Table 2.3: Frequency tolerance, 915 MHz

2.7 Antennas

Antennas are a complicated topic, even compared to the rest of this document, and only general reference will be included. The intent of this section is to understand and select off the shelf antennas, not to design a custom antenna. Quick searches for either antenna design, either for external or PCB work, will result in substantial reference information and ready-built designs.

The key elements of an antenna are that they're tuned to a frequency[5], they have a certain amount of gain, a radiation pattern, have a certain impedance, and interfere with each other.

An antenna is tuned to a frequency based mostly on their length being a clean fraction of a wavelength. The high frequencies of RNodes allow for relatively small antennas, although 433 MHz will be roughly twice the size of 868 MHz, and even unshielded feed line length has tight tolerances due to the small sizes. Wavelengths are determined by diving the speed of light by the frequency of the signal.

Frequency	Wavelength	
(MHz)	Inch	cm
433	27.26	69.24
868	13.60	34.54
915	12.90	32.76

Table 2.4: Common LoRa Wavelengths

Gain is measured in decibels (dB)[6] and is logarithmic. For estimation purposes, this means a gain or loss of 3 dB is equivalent to doubling or halving the signal strength respectively.

Antennas are not amplifiers, however, and this gain isn't free. All antenna gain is based on reduction of radiation. Thus radiation patterns are used to match the antenna to the use case, providing maximum power where it's needed and reducing interference from unwanted directions. By way of example, two very common antenna types are the dipole and the Yagi.

The dipole is typically a half wavelength in size, with both a driven and a ground antenna a quarter wavelength long and pointing away from each other. A reference gain of 2.15 dB is used, and it provides a pattern with greater gain perpendicular to the antenna and lower gain along the axis of the antenna. Thus placing the antenna pointing upwards (with the ground facing downwards by tradition) provides much greater gain closer to the ground and less gain upwards and downwards, directions with far less utility.

a Yagi-Uda antenna has a single driven element (typically a dipole) and a number of passive elements, either to reflect or enhance the signal. The number

[5]Many receive-only antennas operate over multiple frequencies, but a poorly matched transmit antenna, either to wavelength or impedance, will generate waste heat, transmit less power, and potentially damage or destroy the transmitter

[6]The capitalization is based on the base unit being the Bel, named by Bell Labs for Alexander Graham Bell and the useful unit being one tenth of a Bel, or deciBel

and configuration of elements create a varying gain along the axis of the antenna, allowing a dozen or more decibels in gain. The gain to the rear, sides, above, and below is very strongly reduced by consequence. This is very valuable for point to point communication, targeting a certain area, or otherwise focusing all attention in a certain direction. A major downfall is their size, with a 913 MHz 13 dB Yagi being almost one yard/meter in length. Doubling the length only increases the gain by around 2 dB.

Impedance is equivalent to resistance, but for an alternating signal instead of a steady current. There is an additional requirement to match impedance between elements. Without getting into physics or electronic theory, when a wave reaches a discontinuity in its medium, it can reflect instead of transmitting the entirely of its energy, both reducing output and feeding back into the transmitter. This is similar to how light both reflects and refracts in a transparent surface, and is a complicated field of study. Instead of trying to compensate on a case by case basis, the industry standard is to use known impedances, such as 50Ω, and impedance matching circuits to link networks of different impedance together. This is also a critical specification on cables.

Interference is a critical issue facing multiple antenna elements or any nearby conductive surface. Many monopoles rely on a reflective ground plane, such as the Earth or a vehicle chassis, to function effectively. The directionality or lack thereof depends strongly on this, as well as preventing signal loss or radio shadows. Carefully spaced elements, such as those in a Yagi or phased array, can increase directionality and focus the signal[7] but conventional wisdom is to place antennas at least one wavelength away from each other, preferably with three or more wavelengths of separation.[8]

For completeness sake, antennas also have a polarization, but apart from keeping dipoles vertical, there's little effect on RNodes, since there's no reason to use things such as circular polarization in any typical application.

Be aware using a high-gain antenna may change the legal status of a radio due to increasing effective broadcast power. Please check your local laws before changing antennas on a radio.

Generally speaking, a monopole with a good ground plane, a dipole in basically any situation, or a Yagi or parabolic antenna for point to point communications are great starting points. Specific implementations, such as the quarter wave ground plane antenna, are best left to application-specific research.

[7] 1940s era phased array radar, made of precisely spaced dipoles, is an excellent place to start research

[8] This is a minimum and may not remove interference entirely, however, highly directional antennas may be placed with smaller separation in their low-power directions; see also log-periodic antennas

This page intentionally left blank.

3 Interfaces

An Interface is a method the RNS uses to communicate with another node. These components allow the wide interoperability Reticulum enjoys. While individual interface types both require and accept unique options, they all perform the same function and have the same modes.

TCP Clients allow Reticulum to operate over the Internet, and despite being a client, the node can operate as a propagation node and is otherwise a peer to the nodes it connects to. It connects to TCP Server interfaces, which require either a domain name or IP address to allow clients to connect. I2P and UDP Interfaces function in a similar manner.

The other critically important interface is the RNode Interface, which allows for LoRa links between nodes. It is a long range, low bandwidth interface that provides a license-free[1] low cost radio interface.

Please refer to Chapter 6 for further information on Interfaces.

```
[[RNS Testnet BtB]]
  type = TCPClientInterface
  enabled = yes
  target_host = betweentheborders.com
  target_port = 4242
```

Figure 3.1: Sample Client Interface

3.1 Modes

Both the Full and Gateway modes are functionally equivalent except a Gateway will attempt to proactively provide pathing information to any clients that query the interface.

An Access Point (AP) is designed to provide the same pathing functions as the Gateway endpoint, but is designed for short-term connections and low data transfer when not in active use.

Roaming mode is specifically designed for mobile nodes, and both announce propagation and TTL reflect this.

Boundary mode is designed to prevent high speed networks from saturating lower rate networks. As noted in the Reticulum manual, placing an Internet

[1]This statement depends on region; check your local regulations

```
[[BtB Server Interface]]
  type = TCPServerInterface
  interface_enabled = True

  # Listen on all TCP interfaces
  listen_ip = 0.0.0.0
  listen_port = 4242

  # Setting interface mode
  mode = gateway

  # Data rate cap
  bitrate = 112000 #Baud / bits per second

  # Target announce rate (seconds) - Once per hour
  announce_rate_target = 3600
  # Number of violations before penalty
  announce_rate_grace = 6
  # Time announce delayed after violation (sec) - 4 hrs
  announce_rate_penalty = 14400
```

Figure 3.2: Sample TCP Server Interface

connection in boundary mode will prevent excessive data on an attached RNode interface.

See section 6.12 in the Reticulum Manual for more information.

3.1.1 Announce Propagation Rules

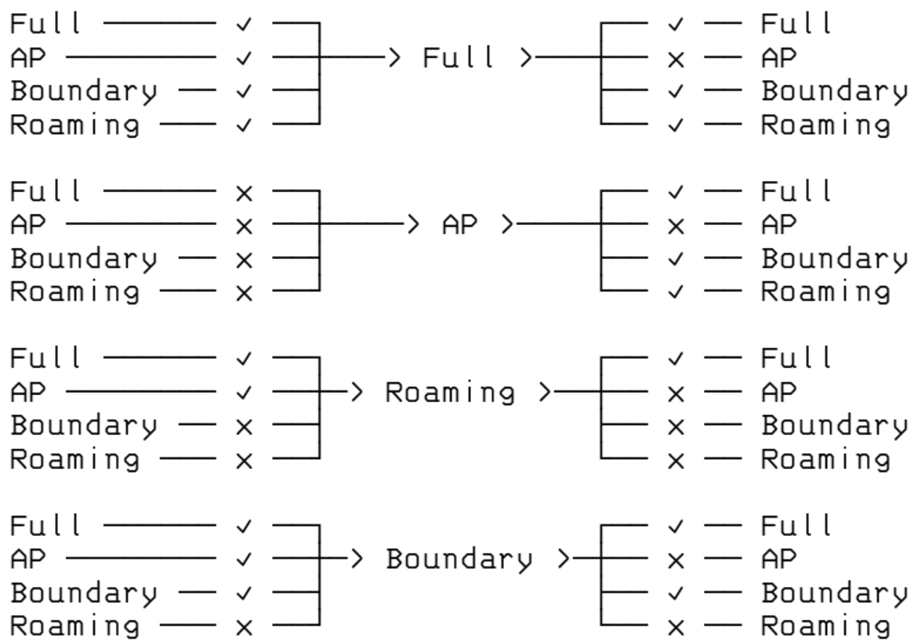

Figure 3.3: Announce Propagation Rules (Qvist)

3.2 Group ID

AutoInterfaces with a set Group ID act similarly to a VLAN, and makes a logically separate network. By grouping interfaces in this manner, a single piece of hardware can run multiple networks with multiple interfaces to save on hardware costs in a single location.

If any other node ends up bridging these interfaces, they will become part of a single network again, just with extra hops, but this is a network architecture concern.

Under most use cases, this is likely not very useful. Multi-NIC servers and major infrastructure may benefit from this separation, bur that is beyond the scope of this document. To run multiple networks over the same medium, it is generally advised to use IFACs (see below).

3.3 Access Codes

By supplying an access code, a network becomes closed. Only packets with the proper authentication are passed and destination addresses are encrypted. This is functionally equivalent to using a password on a Wi-Fi network, even through Reticulum messages are always encrypted even without an access code.

As noted, this creates a defacto network, and networks with differing names/-passcodes will be treated as separate networks.

4 Off-Grid Power Planning

4.1 Basic Electric Power Theory

Caution: Electrical power transmission and storage is dangerous. The generalized guidance provided here is designed to guide people with off-the-shelf part selection or expand the knowledge of people already familiar with electrical system design and construction. Complying with code, design rules, and safety standards is the responsibility of the reader.

Apart from respecting voltage limits and mitigating losses, the most important planning number for off-grid systems is the power consumption in Watts and the storage in Watt-Hours. A one Watt-Hour storage system can, ideally, power a one Watt load for an hour, a two Watt load for a half hour, or a half Watt (500 milliwatt) load for two hours. However, many production and storage systems aren't measured in Watts, requiring a basic understanding of electrical theory.

Watts are the product of Volts and Amps, as described in equation 4.1. Similarly, 4.2 is a simple equation for direct conversion of charge. Using these in a system can be complicated and will be discussed later, but by way of example, three batteries[1] will be compared: a NiMH AA, a 18659 Li-Ion, and a small lead-acid battery. The AA is 1.2 volts at 2400 mAh, the Li-Ion is 3.7 volts at 2800 mAh, and the lead-acid is 6 volts at 224 Ah. The batteries are not comparable in size, as the first two are single cells, and the last is a 68 pound battery, but the math is similar to both types of storage systems.

$$Watts = Volts \cdot Amps \tag{4.1}$$

$$WattHours = Volts \cdot AmpHours \tag{4.2}$$

There are also charge and discharge rate limitations. Batteries have a C rating that describes the percentage a battery can discharge per hour. A 100 Amp-Hour battery with a C of 1 can supply 100 Amps, while a C of 20 would have a maximum discharge rate of 2,000 Amps.[2] For this discussion, the NiMH has a rating of 3 C, the Li-Ion is rated at 35A (effective C of 12.5), and the lead-acid has a C of 0.05.

[1] Technically, the AA and 18659 are cells, whereas a collection of cells is referred to as a "battery." The term "battery" will be used to mean any cell or battery-based storage system.

[2] Discharge rates and properties are *far* more complicated in practice. Check your datasheets for proper design information.

Type	Voltage	Capacity, Ah	Watt Hours	Max Discharge, W	Min at 25W
NiMH	1.2	2.4	2.88	8.64	Overcurrent
4xNiMH	4.8	2.4	11.52	34.56	28
Li-Ion	3.7	2.8	10.36	129.50	25
Lead-Acid	6.0	244	1464	73.20	3500

Table 4.1: Example battery chemistry comparison

Type	Weight	Cost	Fire Risk	Lifespan	Trickle
Lead-Acid	High	Low	Low	Low - Mid[3]	Yes
NiMH	Mid	Mid	Low	Mid - High[4]	Slow
NiCd	Mid - High	Mid	Low	High[5]	Yes
Li-Ion	Low	High	Very High	High	No
Li-Po	Low	High	Very High	Mid - High	No

Table 4.2: Battery general attributes

Wire and cable selection are also important. A given wire can only pass so much current before it gets intolerably hot. The internal resistance causes a flow restriction which generates heat. The heat generated is the product of the voltage drop times the amperage, but it's difficult to determine voltage drop in the planning stages and much simpler to use resistance (which is related to voltage drop). Ohm's law (4.3) states that voltage equals current times resistance, and knowing that power is volts times amps, determining the power without knowing the voltage is possible. (4.4)

$$V = I \cdot R \tag{4.3}$$

$$P = I^2 \cdot R \tag{4.4}$$

Gauge	Diameter (Inch)	R/foot (mOhms)	Heat/foot (mW) .1 A	1 A	3 A	Limit (Amps)
12	.081	1.5	.015	1.5	13.5	9.300
14	.064	2.5	.025	2.5	22.5	5.900
18	.040	6.4	.064	6.4	57.6	2.300
20	.032	10.2	.102	10.2	91.8	1.500
22	.025	16.1	.161	16.1	144.9	0.920
24	.020	25.7	.257	25.7	231.3	0.577
26	.016	40.8	.408	40.8	367.2	0.361
28	.013	64.9	.649	64.9	584.1	0.226

Table 4.3: Copper wire chart

[3] Lifespan increases dramatically at lower depths of discharge

[4] Minor memory effect

[5] Battery memory causes early failure if not regularly fully discharged

4.2 Power Generation Systems

4.2.1 Solar

There are two major systems for harvesting solar power, photovoltaic (PV) and thermal. A PV cell turns solar radiation directly to electricity while thermal cells are used to heat a fluid, often water. Thermal cells, while robust and useful, will not be discussed and should be rolled into Heat Scavenging. (4.2.5)

PV module efficiencies[6] range from 20 - 40% and further efficiency losses from storage and conversion are considerable. Due to this the entire system should be designed to a single standard with the minimum number of conversions. Due to battery construction and off-grid systems, this typically means a 6 or 12 volt nominal system, with a final power conversion from raul voltage to 5 or 3.3 volts as required by end-use hardware.

For insolation and other data, the National Renewable Energy Laboratory has general reference maps. (https://www.nrel.gov/gis/solar-resource-maps.html) While more specific maps are required for planning, they're useful for getting real-world numbers for feasibility purposes. More detailed information used for the following examples can be found at https://www.nrel.gov/docs/legosti/old/789.pdf. These examples use station 14898 in Green Bay, Wisconsin.

Month	Insolation kWh/m^2	Elevation (degrees)	Effective Area (m)	Ideal Insolation
Jan	1.42	23.2	0.39	3.64
Feb	2.29	31.7	0.53	4.32
Mar	3.48	41.7	0.74	4.70
Apr	4.53	52.5	0.79	5.73
May	5.42	60.5	0.87	6.23
Jun	6.01	64.8	0.90	6.68
Jul	5.95	63.7	0.90	6.61
Aug	5.11	56.7	0.84	6.08
Sep	3.84	45.5	0.71	5.41
Oct	2.59	33.8	0.56	4.62
Nov	1.47	24.2	0.41	3.59
Dec	1.10	20.3	0.35	3.14

Table 4.4: Sample Insolation, station 14898, Green Bay, WI

The data provided shows an average insolation of a flat one square meter area on the ground, and is not representative of a panel's total exposure. As shown in Table 4.4 there is a wide variance in the elevation of the sun during the year, and this changes the areal insolation more than the distance to the sun or the atmospheric attenuation. By tilting the panel to 40°for Green Bay, the effective area changes from between 35 - 95% to 90 - 100% as shown in Figure 4.1. The lateral movement of the sun is not compensated for, and a sun-tracking

[6] Not cell efficiencies; beware marketing material discussing cell efficiency, which is 5-10% lower than actual output

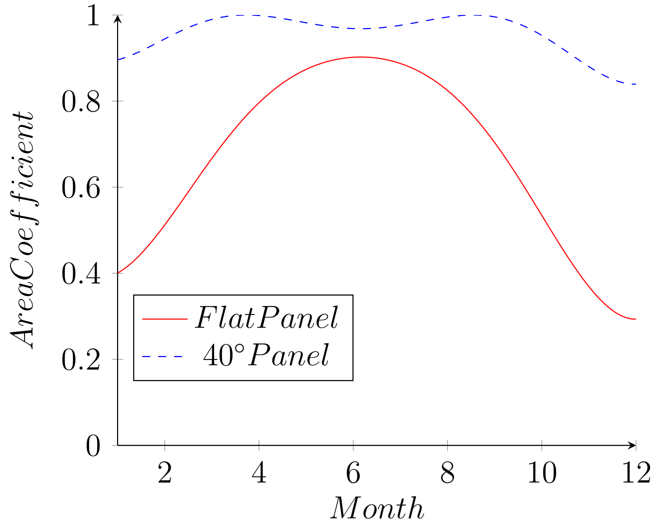

Figure 4.1: Effective area, station 14898, Green Bay, WI

system may produce more power at the cost of additional complexity and energy consumption.

Given this information, the planning numbers for available solar energy in Green Bay is 2.5 - 6 kWh per day per square meter. This assumes clear sky free from terrain obstruction. Given a cheap solar panel with an efficiency of 20% this gives a power output of 500 Wh to 1.2 kWh per square meter of panel. Regardless of the claimed rate of the solar panel, more than 50% of available energy is simply not possible. Solar panels must always be rated against the available energy, not the claimed capacity of the panel.

There are ways to increase the output of a panel, mostly revolving around using reflectors or guides to increase the sunlight reaching the panel. However, these also increase the heating of the panel, and increased heating will drastically reduce the life of the solar panel. Actively cooled panels or clever geometry that only reflects light during certain times can mitigate overheating, but is beyond the scope of this document.

4.2.2 Wind

Wind power is one of the most temperamental and size-intensive renewable energy solutions. One of the shortest effective mast heights is 10 meters / 30 feet. Obstructions and ground effects significantly reduce wind speed and increase turbulence below this level. Even at this height, the wind speed is several meters per second slower than 80 meter towers. While it's tempting to double up by placing an antenna on the turbine, RF leakage from cheap generators and distortion from the blades may cause issues with reception.

The design of a wind turbine is a complicated affair and well beyond the scope of this document; it is assumed an off the shelf turbine with an applicable mount and overspeed protection is used for this exercise.

A final caution, an exclusion zone around the mast equal to its height is required for safety reasons. A 30 foot mast would thus require a circle 60 feet across to be clear of obstructions and people for optimal performance and safety.

The amount of energy that can be extracted from wind is based on the square of the wind speed, meaning doubling the speed means eight times the available energy, making high speed and consistent winds the most important factor for wind power. Finding average local wind power is critical for any planning, but NREL data shows the 10 meter wind speed over most of the continental United States to be 3 - 5 meters/second, and any reputable vendor will provide energy output at a given wind speed.

An example of a currently offered unit, a 400 Watt 50-inch turbine, is rated for 3 - 25 meters/second, with a startup speed of 2.5 meters per second and a survival speed of 40 meters per second. Looking at its published power curve, you could only expect 50 - 100 Watts at 3 - 5 m/s and some days it may not start at all. This is why it's important to match site conditions to equipment instead of looking at the peak ratings.

Failing this, many calculators exist to determine the maximum power extracted from field conditions. The most important consideration is that wind power requires even more battery capacity than solar, as it can experience long periods of high or low average speed.

In certain regions wind power can make a tremendous amount of sense, but the wind variability, power storage, and maintenance requirements must be considered, and while much of the area can be empty, it does require the largest land area to be dedicated to its use, and it poses the second largest risk to people and animals in the area (after steam power).

4.2.3 Internal Combustion

Internal combustion engines burn fuel in a combustion chamber, and while large scale turbines exist, small turbines are very inefficient and have tremendous exhaust temperatures, even with regenerators. Gasoline and diesel engines attached to a generator are very common, power dense, and quick to refuel. However, they do require a constant supply of fuel. An electronics-rated generator will output much cleaner power, and some can throttle or even self-start. It's important to determine the proper system, as the differences can completely change the performance of the system. A battery system attached to a generator that runs at peak efficiency only when the charge is low may be more efficient than a throttling constantly running system, but it also has a greater risk of failure due to the starter. An inverter equipped system will also have lower efficiency than a system that outputs power directly from the generator.

These systems can also run on wood gas, liquid propane, or natural gas. These usually have the same general properties as gasoline or diesel systems, but natural gas systems can be attached to municipal supplies; this is not off-grid, but can be helpful for backup power during outages.

Internal combustion engines may have more application for off-grid but camp based applications than autonomous stations, but it's a great resource when combined with other requirements, and when combined with vehicle power take-offs and large scale manned applications it's a good high-power, environment-tolerant system.

4.2.4 External Combustion

External combustion engines are devices that use a fuel to heat a working fluid (or solid state device) externally, then extracting power from this temperature change. It is possible to do this with Peltier elements, it is an expensive and wasteful process, and the primary engines of this type are steam engines and Stirling engines.

DANGER

Steam engine boilers are pressure vessels and are extraordinarily dangerous. They are highly regulated and require specialty construction skills. A boiler failure may cause an explosion with a substantial lethal radius. Amateur construction is highly, highly discouraged and a boiler cannot be left unattended, severely limiting usefulness.

Only use professionally manufactured purpose built boilers.

These kinds of devices are mostly useful when there's already a working fire for other purposes (heating, cooking, etc.) or there is a constant supply of fuel and water, but a lack of electrical power.

Using a steam engine to feed a turbine can generate power using combustible fuels, but the efficiencies of scale of a turbine dictates a rather low efficiency and considerable fuel consumption for limited electrical output. Multi-stage piston engines are the best way to maximize the efficiency of low quality steam, but they are large, expensive, and only useful in static applications. A steam-based system is really only useful when a fire or steam is already available and efficiency is not an issue.

A Stirling engine uses heat differential to produce mechanical energy, which can then be fed into a generator. Power is based on the temperature differential between hot and cold ends, and it's important to keep these ends from coming to equilibrium. By heating the hot end and attaching the cold end to a reservoir, such as the ground (with good surface area and thermal conductivity) the engine can convert this temperature differential into power. Given the requirement for large thermal masses, this requires either tremendous air cooling or a semi-static emplacement. Once again, this is most useful when there is already a fire for other purposes, but it is a much more feasible and safe solution than small-scale steam power.

4.2.5 Heat Scavenging

The Stirling and Peltier production methods above don't require fire, but rather require a temperature differential. Similar to Combined Heat and Power plants (CHP) this allows waste heat to be used in a productive manner. This is not free energy, there are backpressure and cold end considerations, however it is using waste energy in a positive way. Heat from flue gases from a wood-fired stove can be used to heat a Sterling's hot end while the cold end can be attached to a cooling system, and small amounts of power can be generated. Specifications for both Peltier elements and Stirling engines vary; check their documentation for temperature requirements and efficiency. See Chapter 5 for general information on available energy across a temperature gradient.

4.3 Battery Chemistry

A chemical battery is, at its most basic, two materials with differing electronegativities producing a voltage. A wet cell uses a liquid electrolyte while a dry cell uses a paste or similar technology. Rechargeable batteries use a reversible process where chemical energy is released when there is a load and the chemical reaction is reversed when power is applied. The charging mechanisms and limits when the process becomes irreversible change by chemistry; check your specific battery's datasheet.

4.3.1 Nickel-Cadmium

Nickel-Cadmium (Ni-Cd, trade name NiCad) batteries have been losing popularity against NiMH cells, but they have certain advantages, such as power density, high cycle life, and tolerance of deep discharge. However, their greatest flaw, aside from the toxicity of cadmium, is their memory; a Ni-Cd battery that is repeatedly partially discharged will have a reduction in effective capacity. The need to have well managed charge controllers, including periodic deep discharge, makes their use as renewable energy storage system difficult. However, they are still useful in their traditional role: to be replaced with fresh cells when charging is required.

4.3.2 Nickel-Metal Hydride

Nickel-Metal Hydride (NiMH) has become one of the major standards of batteries, and the de-facto rechargeable AA and AAA cell. They are capable of trickle charging, high discharge rates, and good power per volume and price. However, they are not straightforward to charge, and possess a risk of individual cells being overdischarged if a battery is drained too far. These issues can be mitigated with proper usage and charge controllers, making them a reasonable option for small, portable systems or complicated systems with very low draw.

4.3.3 Lithium Ion

Lithium Ion (Li-Ion, liion) batteries are expensive and heavy, but have extraordinary power density. They require extensive support hardware to perform at their peak, as elevated temperatures or state of charge will quickly reduce the lifespan of the battery, even with no load applied. A properly supported Lithium Ion battery does not have the cycle life of the more durable battery types. An improperly charged, overheated, or physically damaged Li-Ion battery can ignite or explode.

4.3.4 Lithium Polymer

Lithium Polymer (Li-Po) batteries are similar to Lithium Ion in many regards, including safety concerns and high power density, but are considerably more expensive, lighter, and have a lower charging efficiency. They are widely available for drone usage, but do recommend being charged in a protective enclosure in case of damage. They are widely used for mobile applications but absolutely must have proper charging circuits and mechanical protection. Use extreme caution when ordering components, as many vendors wire their plugs in the reverse of expected polarity, putting both battery and device at risk.

4.3.5 Lead-Acid

Prior to any detailed discussion, it's important to note that there are two major types of lead-acid batteries. SLI batteries, made for Starting, Lights, and Ignition, are the traditional automotive starter battery, are designed for maximum short-term current, and are rapidly destroyed by deep discharge. The more useful batteries for off-grid storage are deep cycle (often marine) batteries. These use a more durable plate design and while they produce much lower peak current, they are designed to be used for power storage and are often found in static power storage systems and in vehicles like forklifts, where their weight is an advantage.

Lead-acid batteries are a heavy, wet cell battery comprised of lead and lead oxide plates submerged in acid. Sealed batteries are virtually maintenance free and have a long lifespan over a wide range of operating conditions. Colder temperatures do reduce current, but this will likely not affect the kinds of drain experienced by a Reticulum system. They are available in multiple voltages and form factors from small 6 volt units to batteries well over a thousand pounds.

The largest advantage to lead-acid batteries, particularly the 12 volt models, is a wide range of support from both the renewable energy and automotive industries. While a 12 volt bus is less useful for small electronics, it's also a bus voltage for modern computing and producing 5 and 3.3 volts from 12v is easy using off the shelf components. Given the ease of charging a lead-acid battery and the widespread availability of charge controllers, it's a great option for any static or vehicle-based power storage system.

Be aware "12 volt" automotive systems are actually 12.6 - 15 volts when running, and a true 12v rail still needs a voltage regulator if devices are not rated for these higher voltages from lead-acid batteries and their charging equipment.

4.4 System Design

The following is a single process that can be used to produce an off-grid system, and is not intended to be taken as the only way or applicable for all situations. Assumptions and goals for this design are not suitable for all systems.

First, the load must be determined. A 5v 3A single board computer powering an RNode provides a maximum ideal draw of 15 Watts. Given a 100% duty cycle, this means 15 Watts for 24 hours per day, or 360 Watt-hours. This is our core power consumption.

Before specifying individual components, a solar power supply connected to a 12v lead-acid battery is selected. An automotive rated 12v - 5v DC/DC converter with a rated 95

Second, the charging system efficiency needs to be considered. Assuming 100% of the power is from the battery and a charging efficiency of 70%, the 380 Wh will require 450 Wh to recharge. An inexpensive solar charge controller is around 80% efficient, meaning a total solar input of around 560 Wh.

Given the Green Bay examples from Section 4.2.1 we find a 500 - 1,200 Wh/m^2-day power availability, and would select solar panels equaling 1.5 square meters, providing an excess of energy in the summer, but still allowing for poor battery performance in the cold months.

Finally, we determine that three days without recharge is as long as we can expect the unit to run, and even during heavy rain there's at least some incoming power. This does, however, require design such that the panels will passively de-snow themselves in that period or maintenance would be performed prior to this limit. This leaves a total working capacity requirement of 1680 Wh, or 140 amp hours at 12v, nominal. Using a battery this size would quickly destroy it, and a 50% depth of discharge system should last 700+ cycles, giving a working life span of two to three years. Thus a minimum capacity of 280 amp-hours is selected. A battery this size is generally difficult to find, but two 140 amp-hour batteries in parallel can provide this capacity.

These batteries are considered end of life at 80% of their rated capacity remaining, so even after this time period, they can still last a day or two without charging. If this is acceptable, they can remain in service, or be put in lower-demand applications or even sent for recycling.

It's important to note that while this was designed for worst-case draw, and the entire system could be scaled back with a lower demand SBC, it contains no redundancy and is not resilient. It could be suitable for placing at a cabin or in a barn, but it is not fault-tolerant enough to be used for any critical function. Similarly, this is purely for the Reticulum node, and includes no "balance of plant," meaning all the support systems, environment sensors, housings, and control hardware are entirely missing. This is provided as an example to guide purpose-built designs, not as a reference design.

This page intentionally left blank.

5 Thermal Management

Thermodynamics begins with very simple rules that, while intuitive, need to be understood before any thermal management can be planned. Generally speaking, the First Law of Thermodynamics states that heat cannot be created or destroyed, but more specifically it states that the increment of heat in the system is equal to the sums of all heat transferred to the system, less the work done by the system. This is described in equations 5.1 and 5.2 which are both equivalent to the generalized and specific definitions above. In the case where the system is doing no appreciable work, such as an electronics enclosure, the statement can be simplified to say that the increment of heat in the system is equal to the sum of all heat transfers in or out of the system. (5.3)

$$HeatEnergy = NetHeatAdded - WorkDone \tag{5.1}$$

$$\Delta U_{system} = Q - W \tag{5.2}$$

$$\Delta U_{system} = \sum Q_n \tag{5.3}$$

The Second Law of Thermodynamics states that heat energy only flows from regions of higher energy to regions of lower energy. Under all practical circumstances, all connected regions reach equilibrium, they do not increase their temperature differential.[1]

Given these laws, a design rule emerges: for every Watt of heat energy in, one Watt must be released. For open systems, this is a matter of adequate air flow, but many use cases require a sealed enclosure, creating a closed system that needs proper design.

From an engineering standpoint, there are two ways to approach this issue, statically or dynamically. Dynamic systems can be more accurate and account for more variables, but a static analysis is far simpler and is accurate enough once steady-state is reached.

Given our design rule and equation 5.3 we can see the problem with an example 25 Watt system. Since 25 Watts are added to the system, 25 Watts must be removed, and this can only occur when flowing from a higher temperature to a lower temperature. Thus, the enclosure temperature will rise until thermal equilibrium is achieved. The design must be tested to see if this is acceptable, as

[1]Endothermic chemical reactions, refrigeration cycles, and Peltier Elements obey these laws by using energy from work to transport heat energy in a process that converts work into heat, but in a way that increases the ΔT between regions.

a device in a shaded region in northern Canada will have a different equilibrium point than a box in direct sunlight in Sierra Leone.

Consider a theoretical project box in the shade. The design limits are considered to be 70°C, the thermal load 26 Watts, and a Pelican M50 case used as a project enclosure. Given a surface area of approximately 700 cm^2, a wall of 5mm polycarbonate, and enough air flow to carry away all rejected heat, we can determine the maximum operating temperature.

Using Fourier's Law (5.5) we can determine the heat transfer using this information. With a k of 0.19, equations 5.7 to 5.9 result. Under these idealized conditions, maximum operating temperature is up to 60°C atmospheric temperature.

$$Flux = \frac{C_{ThermalConductivity}}{Thickness} \cdot Area \cdot TempDifferential \tag{5.4}$$

$$q = \frac{k}{s} A \Delta T \tag{5.5}$$

$$q = UA\Delta T \tag{5.6}$$

$$25 = \frac{0.19}{0.005} 0.07 \Delta T \tag{5.7}$$

$$\Delta T = \frac{25 \cdot 0.005}{0.19 \cdot 0.07} \tag{5.8}$$

$$\Delta T = 9.40°C \tag{5.9}$$

While it doesn't seem the case is holding a considerable amount of heat, it's providing substantial insulation and increasing internal temperatures tremendously. By way of comparison, a 1mm aluminum enclosure would provide the following results:

$$\Delta T = \frac{25 \cdot 0.001}{240 \cdot 0.07} \tag{5.10}$$

$$\Delta T = .00015°C \tag{5.11}$$

The assumptions made above include an infinite thermal mass on the outside of the box. It's important to also consider the limits of convection. Newton's Law of Cooling (5.12) is fundamentally similar to Fourier's Law, but uses a different, variable, value for heat transfer and, for obvious reasons, lacks a thickness component. The value of h_c can be estimated using equation 5.13, where v is windspeed in meters per second.

$$q = h_c A \Delta T \tag{5.12}$$

$$h_c = 12.12 - 1.16v + 11.6v^{(\frac{1}{2})} \tag{5.13}$$

Assuming 50% of the enclosure is experiencing a fairly smooth flow of air, a worst case scenario can be determined. With an effective surface area of 0.035

square meters and a wind speed of 0 (making h_c 12.12) this leads to equations 5.14 to 5.16 and shows that the limiting factor in this cooling system is, in fact, the amount of convective heating to open air, generating a system which is only within design limits when the atmosphere is under 10 °C.

$$25 = 12.12 \cdot 0.035 \Delta T \tag{5.14}$$

$$\Delta T = \frac{25}{12.12 \cdot 0.035} \tag{5.15}$$

$$\Delta T = 59°C \tag{5.16}$$

The ways to solve this issue are common and widely understood. Increasing the surface area, as in a heat sink, or increasing air flow, such as with fans, or a combination of the two are the simplest solutions. The ability to use conductive compounds to provide a direct path from chips to cooled housing is a huge advantage of a thermally conductive enclosure.

As an example, we can consider an objective of a 20°C convection system and determine the requirements of the enclosure. We can maintain a purely passive design[2] and see the required surface area, or we can introduce forced draft and determine the feasibility of such a design. The passive system is scaled up with a constant 50% effective surface area, as shown in 5.17 to 5.19, resulting in a required surface area of 0.2 square meters (total enclosure surface).

$$25 = 12.12 \cdot A \cdot 20 \tag{5.17}$$

$$A = \frac{25}{20 * 12.12} \tag{5.18}$$

$$A = 0.103 \tag{5.19}$$

This case geometry cannot be actively cooled within the estimations provided above, and it's worth noting that even high volume 120mm fans only provide around one meter per second of velocity. It's likely more effective to use an open cooling loop, compensating for dust and insects, or a closed loop with an external heat sink. A combination of the two methods, however, can be effective, with a one meter per second forced draft almost doubling the effective surface area of a heat sink/high surface area case.

[2] We will assume any sort of cooling fins will act within the Newtonian estimations, although fin arrangement for passive cooling is a non-trivial design question

This page intentionally left blank.

6 Propagation Nodes

Reticulum has a number of closely related utilities that are usually lumped in with, but are distinct from, the protocol and network stacks themselves. This is similar to grouping e-mail and web pages into TCP/IP. While people may describe it all as "the Internet" it's important to know what is and isn't part of the protocol.

The *Lightweight Extensible Messaging Format* (LXMF) and its associated tools (LXMessage(s), LXMRouter, etc.) are not a core part of Reticulum, but are designed to work within its strengths. They are typically used by Sideband and NomadNet to act as a messaging standard. While generally designed to be routed and delivered immediately, the propagation node allows for asynchronous delivery.

A propagation node, for example one within NomadNet, acts as a distributed message store. If a destination is not available, the sender has the option to send via the propagation network. This sends the encrypted message to a propagation node, which stores it until the destination is seen. Nodes automatically synchronize with each other and create a secure, failure tolerant, and distributed data store.

6.1 Offline Nodes

The automatic synchronization of nodes provides a secondary function: the transport of messages between networks. Message stores can be used to store messages intended for recipients of other networks as well as offline members of the same network. For example, a vehicle could store all messages from the users in the vehicle until it comes in range of an uplink. This is conceptually similar to mail being taken from/delivered to a facility or ship. It is also similar to old BBS functionality, using bulk data transfers over intermittent network connections to transfer messages to local users.

This page intentionally left blank.

7 Examples

These generalized examples are designed to guide creation of bespoke systems. They are deliberately vague and not matched to any particular application, but act as examples of both applications and useful arrangements of interface modes.

7.1 RNode/WiFi Access Point

A simple Reticulum access point for both LoRa and Wi-Fi. Given a single board computer with an Ethernet port and a Wi-Fi radio, this could function as both a Reticulum and TCP/IP access point, providing the widest variety of access. This is also possible with a Wi-Fi dongle or without a Wi-Fi downlink. With two Wi-Fi devices, or a single radio in bridge mode, it may also be used as an extender for an existing wireless network.

The use of AP and Boundary modes prevents all announces from the high-speed links from propagating over the radio, but both the AP and Gateway interface modes act as a path provider for their respective downlinks.

UPLINK			
Destination	Ethernet / WiFi uplink		
Intent	Connects subnet to larger networks		
Mode	Boundary	**IFAC**	Disabled
Hardware	Ethernet NIC or Wi-Fi radio/dongle		
Rate			
Settings			

DOWNLINK			
Destination	RNode		
Intent	Provides network connection over LoRa		
Mode	Access Point	**IFAC**	Disabled
Hardware	RNode		
Rate			
Settings			

DOWNLINK			
Destination	Wi-Fi User/Router		
Intent	Provides network connection for Wi-Fi systems		
Mode	Gateway	**IFAC**	Disabled
Hardware	Wi-Fi radio / dongle		
Rate			
Settings			

7.2 RNode Solar Propagation Node

This node is designed as a stand alone unit. Rather than linking to a larger network, this accepts connections from local radios and stores messages for offline nodes. It is not capable of establishing its own uplinks, but will transfer over any network that connects to it. The system must be configured as a propagation node.

Limited data transmission reduces overall power consumption, reducing power requirements. See Chapter 4 for guidance on power planning.

DOWNLINK			
Destination	RNode		
Intent	Provides network connection over LoRa		
Mode	Access Point	**IFAC**	Disabled
Hardware	RNode		
Rate			
Settings			

7.3 RNode/Cellular Solar Node

A node that functions similarly to the above systems, but includes a cellular data uplink. This is a good option for off-grid locations that have good cellular coverage, such as rural areas near cities or areas such as parks or farm fields. Due to a wide variety of hardware, uplink methods, and subscriptions, see your vendor's technical documents. TCP Client interfaces are the most likely interface type to use with this type of uplink.

Cellular uplinks require considerable power compared to an RNode on standby, and will require a much more robust power system than the above solar node.

UPLINK			
Destination	Cellular Data Network		
Intent	Connects subnet to larger networks		
Mode	Boundary	**IFAC**	Disabled
Hardware	Commercial cellular uplink modem		
Rate			
Settings			

DOWNLINK			
Destination	RNode		
Intent	Provides network connection over LoRa		
Mode	Access Point	**IFAC**	Disabled
Hardware	RNode		
Rate			
Settings			

DOWNLINK			
Destination	Wi-Fi User/Router		
Intent	Provides network connection for Wi-Fi systems		
Mode	Gateway	**IFAC**	Disabled
Hardware	Wi-Fi radio / dongle		
Rate			
Settings			

7.4 RNode Solar Repeater

Similar to the Solar Propagation Node, this is a stand-alone system designed to increase network coverage. Unlike the Propagation Node, this is not designed as an access point, but is a piece of fixed infrastructure. It transmits all announces and is used to link nodes together. With a single interface and no storage potential, it is the simplest piece of infrastructure possible.

NOTE: This is a very bandwidth-intensive mode, and may require RNode speeds more than twice what is typically expected. It will function adequately on higher speed interfaces.

DOWNLINK			
Destination	RNode		
Intent	Provides network connection over LoRa		
Mode	Full	**IFAC**	Disabled
Hardware	RNode		
Rate			
Settings			

7.5 Vehicle Crew

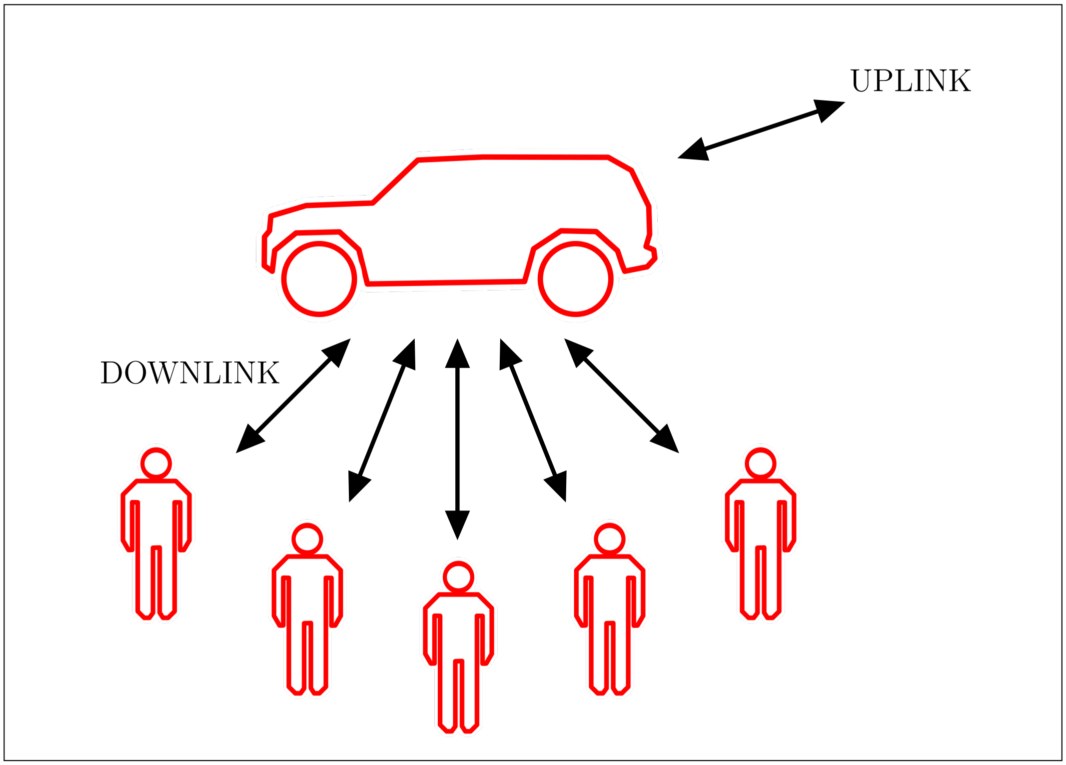

Type: Mobile Node

Use case: A vehicle or other mobile node servicing a small team and using a stable uplink. This is not appropriate for meshing between peer vehicles, and requires communication to other mobile nodes to path through a stable uplink. Vehicle will act as a gateway to any station on the frequency and provide path information for its attached clients. If using an LXMF-type system, it can be used as a propagation node to store messages while out of contact to be uploaded when uplink is re-established.

UPLINK			
Destination	Base station uplink		
Intent	Telemetry digests and uplink router		
Mode	Roaming	**IFAC**	Disabled
Hardware	RNode		
Rate			
Settings			

Explanation: Vehicle uplink is designed to pass text messages and automated telemetry at a relatively low data rate. By running in roaming mode, it signals to the network that it's not expected to be in constant service and it will only pass announces to both full and gateway interfaces. This prevents vehicles on the same network from treating each other as valid paths for known paths.

DOWNLINK			
Destination	Individual Radios		
Intent	Collecting GPS data, broadcasting telemetry, text messages		
Mode	Gateway	**IFAC**	Disabled
Hardware	RNode		
Rate			
Settings			

Explanation: Vehicle acts as a gateway for the individual radios in the unit. If other radios with similar settings are in range, they will also use the vehicle as a node.

7.6 Medical Vehicle Crew

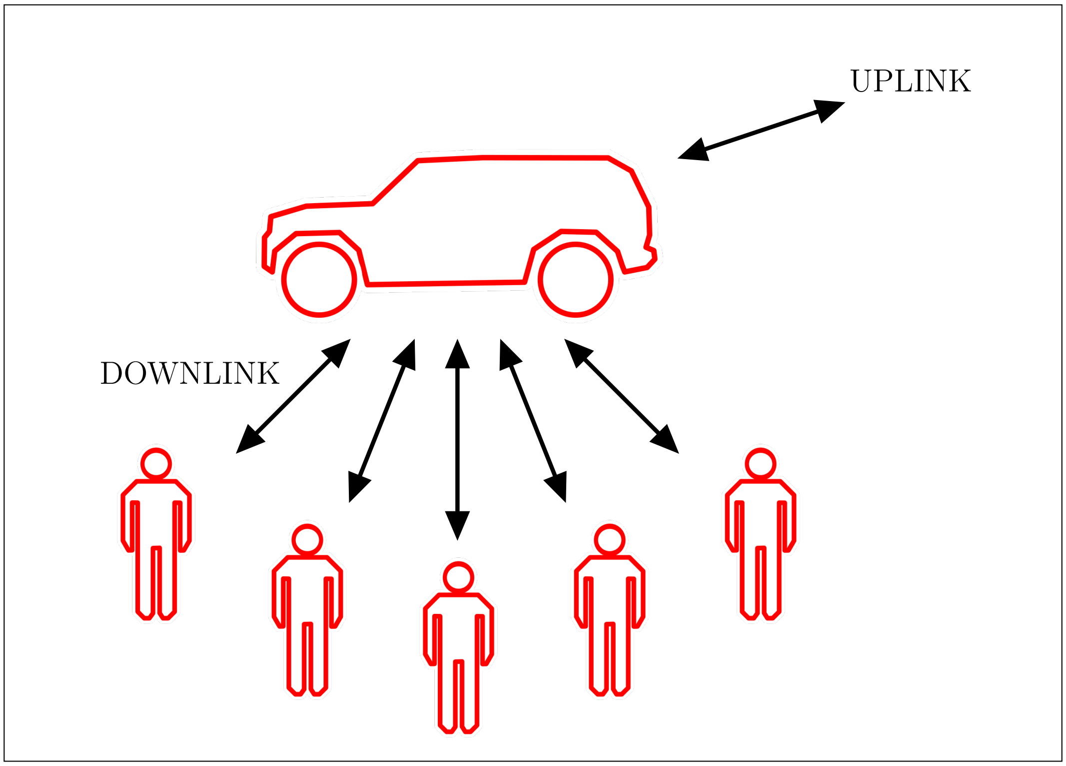

Type: Critical Mobile Node

Use case: A vehicle or other mobile node servicing a small team and using a secure mesh network for transport. Interfaces protected with IFAC are designed not to protect the data (all data is encrypted in all cases) but rather to prevent radios from using their bandwidth for critical operations, such as medical or emergency communications. Vehicle will act as a gateway to any station on the frequency and provide path information for its attached clients. If using an LXMF-type system, it can be used as a propagation node to store messages while out of contact to be uploaded when uplink is re-established.

UPLINK			
Destination	Base station uplink, peer nodes		
Intent	Telemetry digests and critical text communications		
Mode	Full	**IFAC**	Enabled
Hardware	RNode		
Rate			
Settings			

Explanation: Vehicle uplink is designed to pass text messages and automated telemetry at a relatively low data rate. In contrast to the previous example, this system will pass announces through other Critical Mobile Nodes, creating a situation with greater availability in poor reception conditions at the cost of sub-optimal overhead.

DOWNLINK			
Destination	Individual Radios		
Intent	Collecting GPS data, broadcasting telemetry, text messages		
Mode	Gateway	**IFAC**	Enabled
Hardware	RNode		
Rate			
Settings			

Explanation: Vehicle acts as a gateway for the individual radios in the unit. Only radios with the proper IFAC code will be able to connect to this network.

7.7 Strike Team

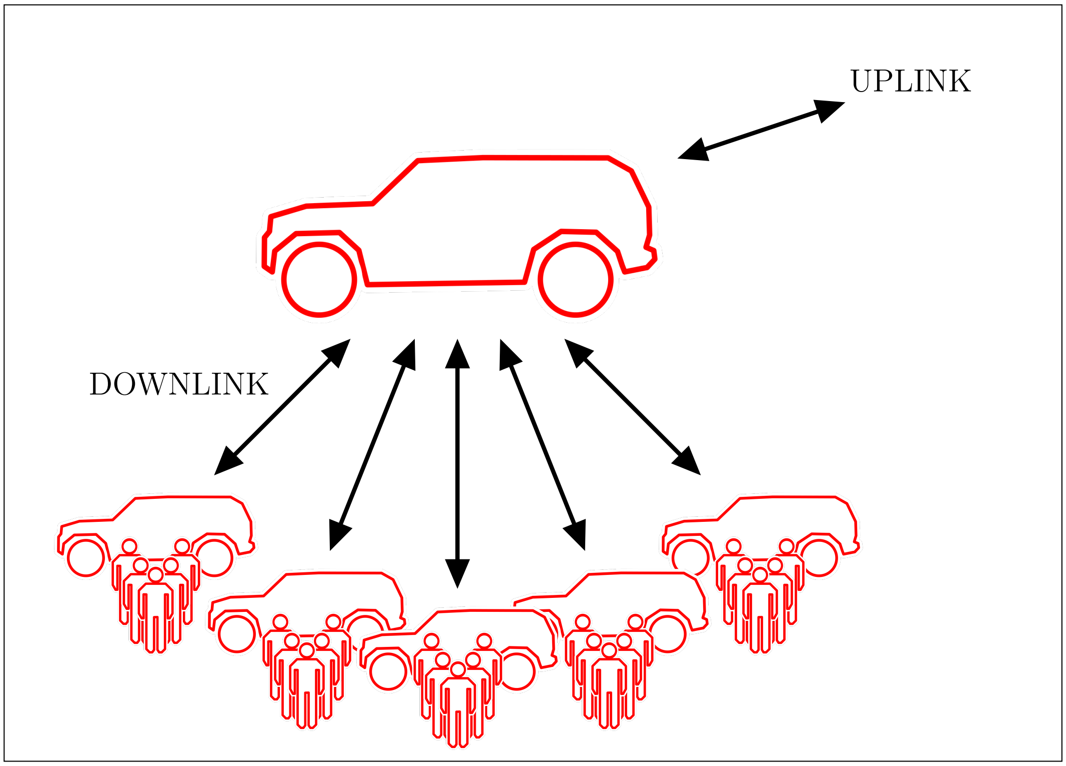

7.8 Incident Command

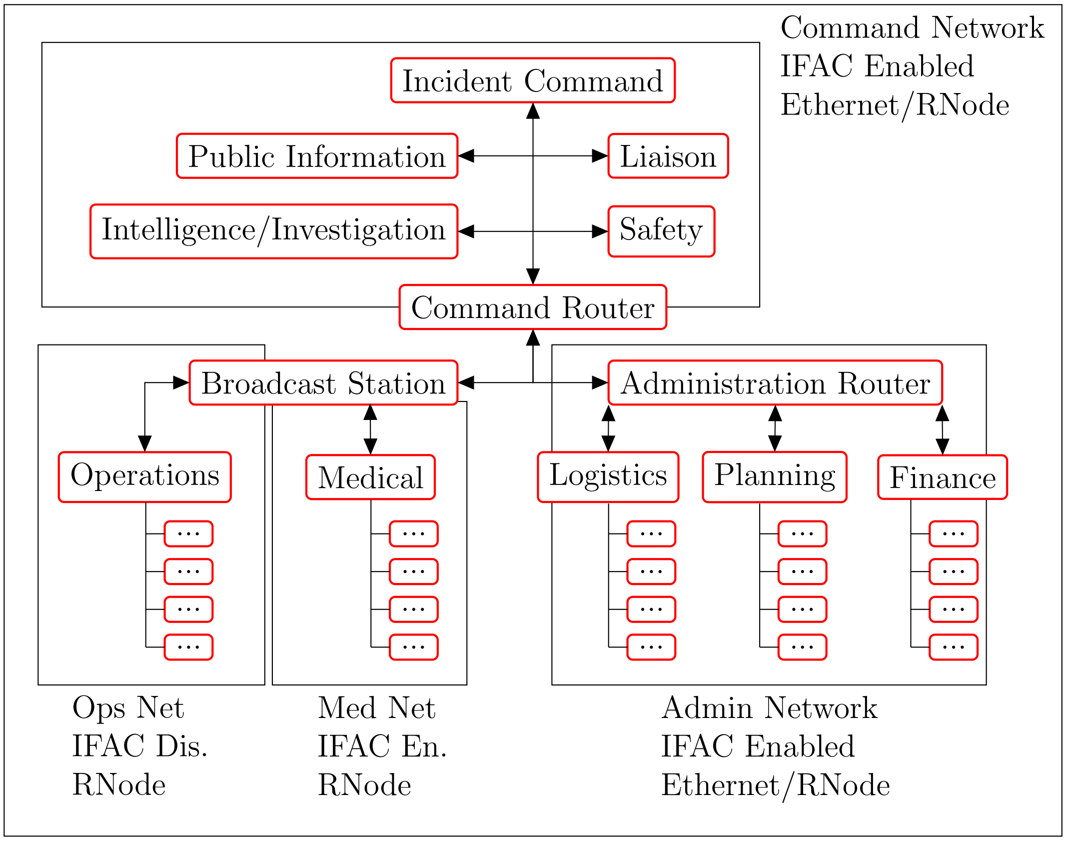

Overall: A basic Incident Command structure for general reference purposes. Using IFAC on all interfaces other than the operations broadcast radio is to keep the network free of non-critical traffic. The Operations Network is open to allow for drop in/out of the network without pre-arranged access codes. While it could also use IFAC with no negative effects other than access, it allows for the widest variety of devices to access the network, be it personal radios or telemetry units. The Medical Network could be synonymous with any closed network, and is access controlled for the same safety reasons as the Command Network.

The Administration network is a fixed network to pass information mostly amongst itself. Due to fixed infrastructure, the Admin and Command networks are likely to have a high speed connection, but technically any node could contact any other, although Finance would have little need to contact someone in Operations below the officer level.

The primary advantage of Reticulum on the fixed nodes is the multi-path capability. If the Broadcast Station and Administrative Router were connected by Ethernet and WiFi, a failure in the fixed network would automatically fail over into the wireless. Given the proliferation of Cellular uplinks, RNodes, and systems such as Starlink, these options provide a fault tolerant network in even the worst conditions while leveraging the fastest available pathing.

Additionally, any station can be made up of multiple nodes in parallel, even in different locations. This provides High Availability across machines and communication layers. Combined with multiple paths provided by mesh networking, Reticulum provides a secure, fault tolerant, and multi-band communications network for both human-human and machine-machine communication, complete with offline message storage and caching.

COMMAND ROUTER DOWNLINK			
Destination	Command Network		
Intent	Transferring mission critical information		
Mode	Gateway	**IFAC**	Enabled
Hardware	Ethernet / WiFi/ RNode / Cellular		
Rate	N/A		
Settings	N/A		

COMMAND ROUTER UPLINK			
Destination	Broadcast Station / Admin Router		
Intent	Routing message between offices		
Mode	Full	**IFAC**	Enabled
Hardware	Ethernet / WiFi / RNode / Cellular		
Rate	N/A		
Settings	N/A		

8 NIMS/ICS Terms

From ICS 200 training documents.[1]

Incident Commander: The Incident Commander is the individual responsible for overall management of the incident.

Command Staff: The Command Staff consists of the Public Information Officer, Safety Officer, and Liaison Officer. They report directly to the Incident Commander.

Officer: Officer is the ICS title for the personnel responsible for the Command Staff positions of Safety, Liaison, and Information.

General Staff: The General Staff are assigned functional authority for Operations, Planning, Logistics, and Finance/Administration. The General Staff also report directly to the Incident Commander.

Section: A Section is the organizational level with responsibility for a major functional area of the incident (e.g., Operations, Planning, Logistics, Finance/Administration).

Section Chief: Chief is the ICS title for individuals responsible for functional sections: Operations, Planning, Logistics, and Finance/Administration

Branch: A Branch is the organizational level having functional or geographic responsibility for major parts of the Operations or Logistics functions.

Branch Director: Branch Director is the ICS title for individuals responsible for supervision of a Branch.

Division/Group: Divisions are used to divide an incident geographically. Groups are used to divide an incident functionally.

Division/Group Supervisor: Supervisor is the ICS title for individuals responsible for a Division or Group.

Strike Team: A Strike Team is a specified combination of the same kind and type of resources with common communications and a Leader.

---

[1] U.S. government works created by a Federal employee as part of their official duties are in the public domain. Contact author if this status is disputed.

Task Force: A Task Force is a combination of single resources assembled for a particular tactical need with common communications and a Leader.

Unit: A Unit is the organizational element having functional responsibility for a specific incident planning, logistical, or financial activity.

Task Force/Strike Team/Unit Leader: Leader is the ICS title for an individual responsible for a Task Force, Strike Team, or functional Unit.