SEPTEMBER 2019

AUTONOMY AT SCALE

Intelligent Machines Advancing Technology to Improve our Future

noblis®

For the best of reasons

NOBLIS.ORG

TABLE OF CONTENTS

FUNDAMENTAL TECHNOLOGY: **A PRIMER ON WIRELESS CONNECTIVITY**

Keith Biesecker

Depending on the application, location or operating environment, the components of an autonomous system may or may not need to communicate—after all, it's autonomous. The system might only need to communicate part of the time (e.g., to upload data), or it might need communications simply for command and control, but not for the exchange of mission data. Most often, however, these systems do require some form of communications, particularly when expected to work at scale.

With the proliferation of micro-sensors and the advent of the Internet of Things (IoT), machine-to-machine (M2M) communication, machine learning, and other smart technologies comes the need for more sophisticated communications—not necessarily faster or more ubiquitous, just different. Some of these new technologies can be used to help define the communication network itself (e.g., self-adaptive cognitive radio networks capable of dynamically re-configuring themselves).

In considering autonomy at scale—variable node densities, dynamic mobile architectures, changing environments—it helps to start with some important fundamentals of wireless communication.

With the proliferation of micro-sensors and the advent of the Internet of Things (IoT), machine-to-machine (M2M) communication, machine learning, and other smart technologies comes the need for more sophisticated communications

Basics of Wireless Communications

One of the more common misperceptions of wireless communication is that a solid link or connection between nodes exists. Sometimes this might be true—and if it's a good link, you might get the performance that's advertised or expected from a service or technology (e.g., 800Mbps WiFi or 10Gbps fifth generation (5G) services)—but these untethered connections are volatile. Wireless communication occurs in a stochastic environment where conditions continuously change. Those depending on the link, such as network engineers, application developers, or users, often assume they have a solid link even when they may not.

Generic Link Budget

$$P_r = P_t + G_t - L_t - L_{fs} - L_p + G_r - L_r \text{ (dB)}$$

- Pr and Pt = receive and transmit power
- Gt and Gr = transmitter and receiver antenna gains
- Lt and Lr = transmitter and receiver losses (connectors, thermal noise, etc.)
- Lfs = free space path loss
- Lp = miscellaneous path losses

Figure 1: Generic Link Budget

A wireless link between any two nodes can be defined by a link budget, (as shown in Figure 1) which is an accounting of all the gains and losses from the transmitter through the operating environment or communications channel and to the receiver. Setting aside system design aspects such as power, antenna gain, and transmitter/receiver efficiencies, consider the more unpredictable and dynamic aspects associated with the communications channel or path. Impediments to radio and microwave communication (300kHz to 300GHz) include: obstructed radio line-of-sight (LOS) and Fresnel Zone clearances, frequency selective fading due to signal multipath, radio frequency (RF) interference from other communications devices/systems, electrical interference from devices such as lighting fixtures and motorized equipment, and attenuation and scattering due to ground clutter (including man-made obstructions), vegetation, atmospheric gases and precipitation (Figure 2). These impairments contribute to path loss (or loss through the channel) and would be included as part of the miscellaneous path losses (Lp) identified in the link budget—miscellaneous, not insignificant.

Wireless communication systems are planned with these challenges in mind. They are designed to meet minimum needs, and margins are built into the link budget to account for uncertainties and challenging situations. If the design is good and the power received sufficient, the link should be capable of exchanging data between the two nodes with the desired performance.

When the communication network must scale to the size, density, and complexity needed to support the autonomous applications discussed in this paper, other aspects need to be considered.

Figure 2: Impediments to Wireless Communications

Requirements & Limitations of Applications

The applications for autonomous systems vary dramatically, and some might not require communication. If they do, the supporting communications systems must account for a variety of interdependent requirements—environmental, architectural, and performance among others.

Environmental Domains (Figure 3)

Terrestrial/Land — The difference in operational environments on the Earth's surface can be dramatic, ranging from densely populated urban landscapes to open rural spaces. Urban environments generally have more networking options available (e.g., commercial providers, architectures, infrastructure), more autonomous system nodes to support the communications network (i.e., node density), and less distance between nodes. Disadvantages of the urban landscape include increased interference and greater losses due to diffraction, fading, and a variety of attenuation factors. In a rural environment, the advantages and disadvantages are generally opposite those of its urban counterpart – fewer causes of interference, diffraction, and attenuation, but also fewer networking options, fewer autonomous system nodes to support the network, and greater distances between nodes.

Air — Low altitudes present similar environmental challenges and benefits as those on land, though usually with less pronounced obstacles. At higher altitudes, fewer impediments exist.

Space — More than 100km above sea level, space has still fewer challenges associated with signal propagation. There may be some atmospheric losses in ground/air and space exchanges, some cross polarization, or some doppler shifting in high-speed deep-space communications.

Water — Under water, the communication channel often exhibits severe attenuation, multipath, frequency distortion, and other impediments—making this environment one of the most complex and difficult wireless channels in nature. Underwater radio is typically limited to lower frequency communication (below 10MHz), achieving rates of a few hundred kbps at distances on the order of tens of meters. Higher radio frequencies (2.4GHz) can be used to achieve higher data rates, but only at distances less than a meter. Acoustic and optical communications provide slightly higher data rates and at longer distances (30kbps at 2500m to 10Mbps at 11m). Communication along the surface of the water can also be challenging, but radio can be used and expected to function if antennas remain far enough above the surface.

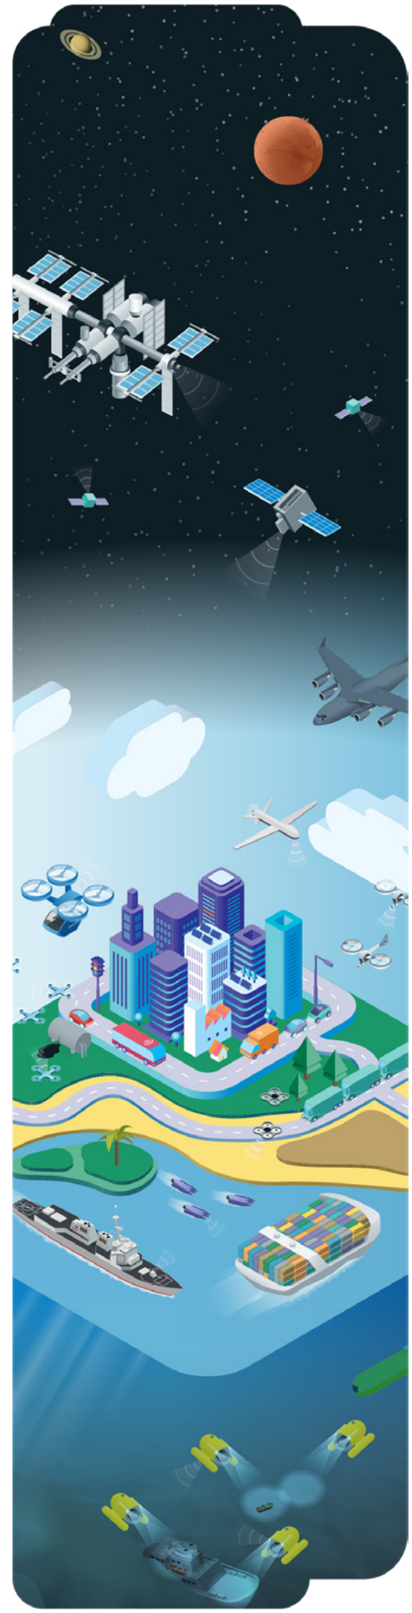

Figure 3: Environmental Domains for Wireless Communications

Architectural Factors

Wireless network architectures generally comprise point-to-point, point-to-multipoint (e.g., typical Wi-Fi with access points), and multipoint-to-multipoint (e.g., Wi-Fi ad-hoc, meshed networks) connections. When taking autonomous systems to scale, the supporting communications networks must account for various architectural factors, as depicted in Figure 4:

- **Node Density** — How many nodes are in the system? Does this number change? Is there a limit?
- **Node Type** — Do only some nodes require communications? Are communications different between different node types?
- **Node Mobility** — Are any nodes moving? All of them? Are they moving at different speeds or different vectors? Are they moving in or out of different networks?
- **Network Infrastructure** — How is the network arranged? Communications for many autonomous systems involve mobile ad-hoc networks (MANET)—continuously self-configuring infrastructure-less networks of nodes connected wirelessly. Some autonomous systems are completely independent and rely only on an isolated MANET, but most usually have at least one node (or gateway) also providing access to a larger communications infrastructure that can be used as backhaul to central command or to help bridge different MANETs.
- **Network Composition** — Autonomous systems might employ homogeneous or heterogenous communications network elements. If homogeneous, can the communications network scale with a dynamic autonomous system? If heterogenous, can the autonomous system successfully adapt to different communication network performance (e.g., data rate) or operational parameters (e.g., security)? Can operational parameters negotiate as nodes move through different network segments?

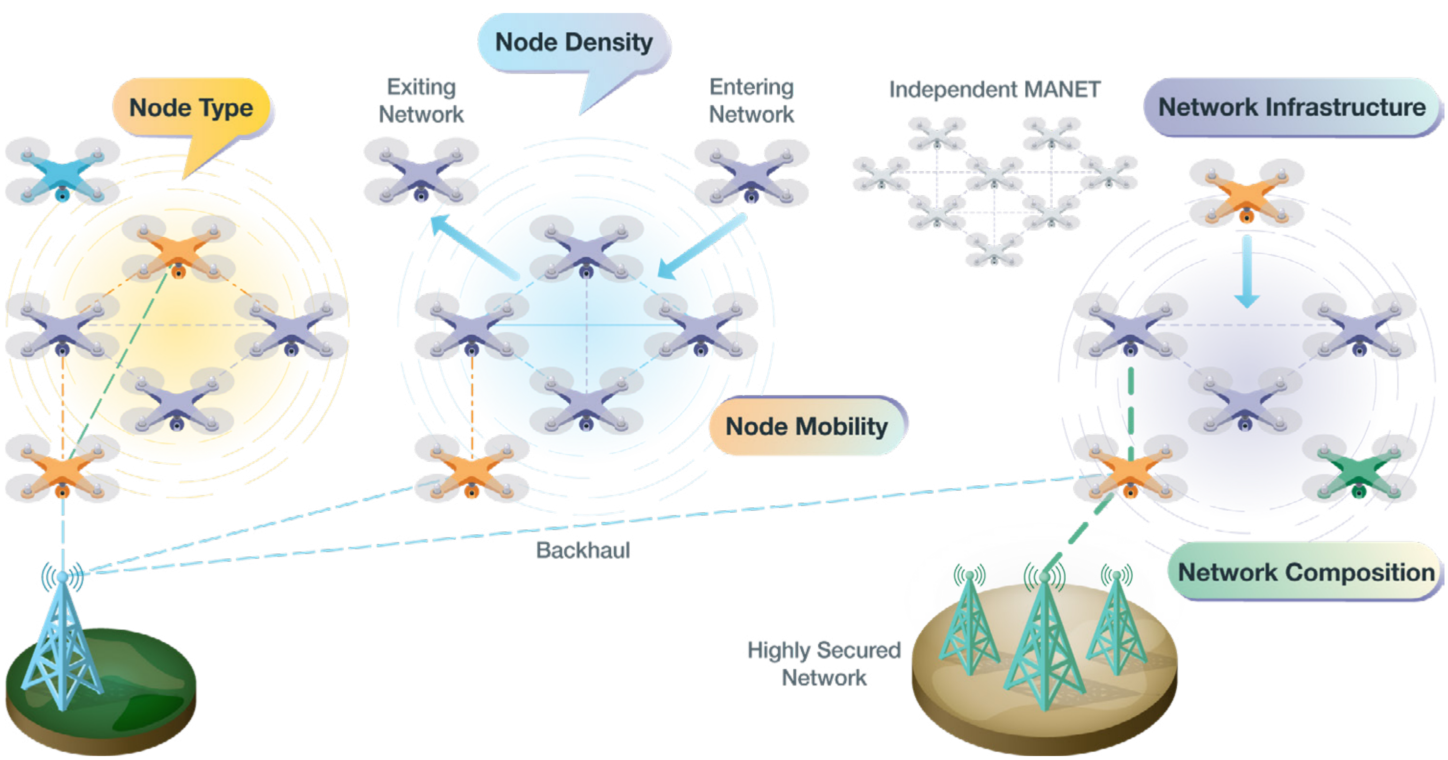

Figure 4: Architectural Factors for Wireless Communications

Performance Metrics

While not always the most important, performance metrics are often the most debated communications system requirements. Some notable metrics include:

- **Data Rate and Bandwidth** — Data rate (bps) indicates the rate at which data can be moved through a channel. Bandwidth (Hz) is the size of the pipe. Recent developments in modulation and channel coding techniques have improved spectral efficiencies (bits/Hz), and increased the amount of data movable through the same size pipe.
- **Availability and Reliability** — Communication availability between any two nodes considers the availability of the link, and the availability of the equipment comprising that link. Link availability accounts for environmental conditions such as rain, snow and fog. Some manufacturers and providers offer a projected annual link availability for their equipment or system when functioning in a specified region or environment. Equipment availability is a statistical estimate based on the reliabilities and repair times of all equipment between the two nodes. The total communication availability from source to destination is the sequence product of all component availabilities through all nodes in the path. Advancements in equipment, protocols, and architectures have vastly improved the communications availability in wireless networks.
- **Latency** — For communications systems, network latency is an expression of how much time it takes for a packet of data to get from one point to another. Contributing factors include simple propagation delay through the medium itself, packet size, packet transmission time, the number of hops a packet must make through the network path, and routing or switching delays. Low latency communication is crucial in applications such as automated industrial control, financial trading, transportation, and applications of augmented and virtual reality (AR/VR). where the requirements can be on the order of 1 millisecond (ms).

Many of the recent advances in wireless technologies and services focus on improving these particular metrics. Consequently, the International Telecommunication Union (ITU) recently defined a set of next generation service categories based on data rate, latency, and availability/reliability.

Other Requirements

Some additional interdependent requirements affecting the communications used to support autonomous networks include:

- **Physical Design** — What power, antenna systems and radios must be selected to compensate for the anticipated path losses identified in the link budget?
- **Spectrum** — With what frequencies will the communications systems operate? Are the frequencies licensed, unlicensed, shared?
- **Security**— Protecting communications often comes at the expense of lower latency and higher data rate. Can the application function properly and be adequately secured?

Development and Innovation

A variety of applications drive the high data rate, low latency, and high availability requirements for the next generation of wireless networking. These include industrial control systems, autonomous vehicles, AR/VR, high-frequency financial trading, and electrical smart grids. Next generation wireless networks are also key to mission-critical IoT, M2M communication, and the Tactile Internet—the evolution of IoT that will add a new dimension to human-to-machine interaction by enabling tactile and haptic sensations that allow people to interact with their environment in real-time.

The initial 5G specifications simply enable new radio technologies to work with the existing 4G infrastructure; however, they include additional provisions for three new service categories that specifically address the strict data rate, latency, and availability/reliability needs of next generation wireless technologies.

Next generation wireless will need to support applications with unique and extremely stringent requirements—end-to-end latency on the order of a few milliseconds, availability greater than 99.9999%, and novel traffic types that use short data packets.

Previous developments in wireless networking focused on improvements to throughput, mobility, and coverage, mostly catering to the human-centric and delay-tolerant content (e.g., streaming media). Next generation wireless (e.g., 5G) will need to support applications with unique and extremely stringent requirements—end-to-end latency on the order of a few milliseconds, availability greater than 99.9999%, and novel traffic types that use short data packets. The degree to which these needs can be met remains to be seen, but it's likely that only one or two of these requirements can be achieved at any one time, particularly at large scale.

Some of the technologies that show promise in meeting these lofty goals are based on 5G mobile technology standards. Unlike previous generations, 5G technology allows for multiple connectivity schemes, heterogeneous networks, and the use of new and existing technologies. The ITU and the 3rd Generation Partnership Project (3GPP), which authorize, create and maintain technical standards for global mobile communication technologies, recently approved an interim set of specifications for 5G communications.

- **Enhanced Mobile Broadband (EMBB)** for supporting stable connections with very high peak data rates.
- **Massive Machine-Type Communication (MMTC)** for supporting the extremely large number of IoT devices, which are only occasionally active and send small data payloads.
- **Ultra-Reliable Low-Latency Communication (URLLC)** for supporting low-latency transmission of small payloads with very high reliability from a limited set of terminals, such as alarms.

Engineering next generation wireless must account for both the stringent requirements of new applications and the traditional performance of today's networks. Current research involves new theoretical principles and modeling techniques, changes to physical design and the communications protocol stack, adaptations to architecture and infrastructure, and other studies. Some more specific areas of research and development include:

- **Physical:**
  - New modulation schemes, multiple access techniques, error coding, and antenna design
  - Distributed signal processing, interference detection, and radio sensing
  - Cognitive Design (e.g., new cognitive radios that can listen to the surrounding environment and select appropriate frequency bands, modulation schemes, or specific power levels)
- **Upper-Layer Communication Protocols:**
  - Low-latency multipath routing schemes based on multipath link availability
  - Opportunistic routing to increase forward node probabilities and improve link availability
  - Machine learning based on smart steering, which is another advanced routing technology
- **Information Theory:**
  - Work on fundamental limits, performance analysis, and network theoretic approaches (e.g., stochastic network calculus)
  - Combining queuing theory and communication theory
  - New communication channel models with adaptations for more dynamic environments
- **Architecture and Infrastructure:**
  - Backhaul and core network adaptations for MMTC and URLLC edge networks
  - Integration of high-performance wired infrastructures

The tremendous amount of research in these areas is evidenced by the number of topics addressed in most current IT/communications organizational journals, conferences, and solicitations.

Conclusion

By their very nature, autonomous systems employ some aspect of self-governance. As such, they might be able to sustain themselves without communication for various amounts of time. Most of these systems, though, will require some form of communication, particularly at scale.

In supporting large-scale autonomous systems, the implementation of a wireless solution is often more important than the selection of any particular wireless technology or service.

When considering wireless communications, remember the basics, understand the requirements and limitations the application imposes, and stay cognizant of developments and innovation. Wireless technologies are volatile, lifecycles are short, and upgrades happen frequently. In supporting large-scale autonomous systems, the implementation of a wireless solution is often more important than the selection of any particular wireless technology or service.

ABOUT NOBLIS

Noblis is a nonprofit science, technology, and strategy organization that brings the best of scientific thought, management, and engineering expertise with a reputation for independence and objectivity. We work with a wide range of government and industry clients in the areas of national security, intelligence, transportation, healthcare, environmental sustainability, and enterprise engineering. Together with our wholly owned subsidiary, Noblis ESI, we solve difficult problems of national significance and support our clients' most critical missions.

NOBLIS.ORG

noblis®

For the best of reasons

703.610.2000 answers@noblis.org @NoblisInc