RTN-101:

Reference Station Communications (Part 5)

The new wave of reference station receivers are RTN-ready, accommodating multiple communications options

Aut viam inveniam aut faciam.

(Either find a way or make one.)

Raw Observation Data, in Real Time

The role of a reference station in an RTN is to be quite stable, to collect good quality raw observations and to transmit these to the Central Processing Center (CPC) in real-time.

As discussed in the previous installments of this series, the fundamental design of an RTN consists of several key components, all of equal importance: a network of reference stations meeting rigorous metrics with respects to positional integrity and data quality, a central processing center to gather the data and produce correctors, and real-time communications. In Part 4 (October 2006) we discussed the real-time communications between the CPC and the field observers (rovers). Here we discuss the communications between the reference stations and the CPC (**Figure 1**).

Acro-speak

CORS - Continuously Operating Reference Station. When people see the term "CORS" perhaps they immediately think of the network of reference stations stewarded by the National Geodetic Survey (CORS, National CORS, CORS & Cooperative CORS); for many, the NGS CORS represents their first opportunity to utilize a continuously operating station and the products thereof. Any continuously operating reference station could rightfully be called a CORS, if it operates continuously, is of a measurable quality, and can act as a reference resource. So for the purposes of this article, the term CORS will be used to describe the reference stations of an RTN (plus, it is quicker to type).

>> By Gavin Schrock, LS

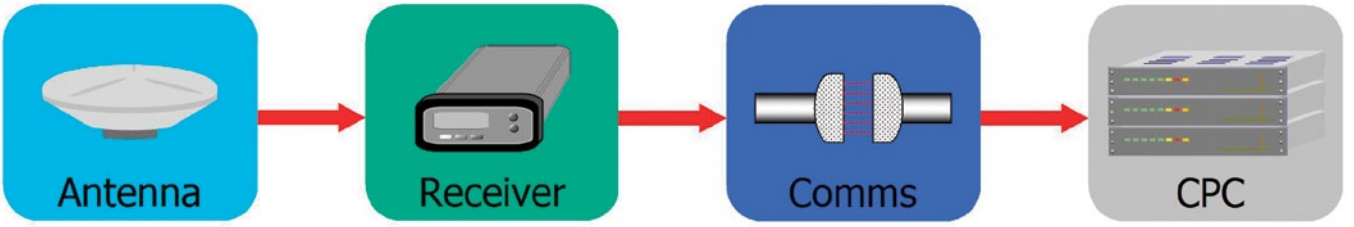

Figure 1 Data flow from rReference station to the CPC

Real-Time

For true network corrections, the CPC must synchronize raw observations from multiple stations. To achieve rapid results and to minimize the observation times for the field user, these observations are sought at a high rate; 1Hz (1-second epoch). Substantial latency in the transmissions would compromise (or likely prohibit) such synchronization. Optimal latencies should not exceed one second from the time of observation through the entire communications labyrinth to the synchronizing software of the CPC.

Though there may be issues of internal latency in some older receivers, communications components, and the circuitous route the data may follow through communications networks (and combinations thereof) you can be fairly confident that if the data is passing though the Internet (with the exception of some Internet satellite links, although there have been great improvements of late) the average total latency will be far less than the one second desired.

Continuous? Some have made the argument that the stations might only need to provide the data during normal hours of field operations, but that is looking at a narrow set of needs. An RTN can (and should) benefit a wider range of users; scientific, academic, engineering, construction, structural monitoring… many round-the-clock. In addition, the function of network integrity monitoring only achieves optimal results over periods of days and weeks.

Antenna to Receiver

The signal from the antenna to the receiver is R/F, and this link is typically well provided for by the equipment manufacturers. The general rule of thumb is that this link is made via a coaxial cable of up to 100m length (**Figure 2**), although greater lengths can be attempted with the aid of in-line signal boosters.

While it is almost always better to house the antenna cable in conduit, some opt simply for heavier gauge coaxial cable that may even be rated for direct burial or exterior exposure. (One network administrator confessed to bribing a cable-TV installer with a few beers to advise on some of the details. I don't recommend this, but it worked).

Many of the antennae paired with specific families of receivers need small amounts of (low voltage) power, but typically receive this power from the respective receiver through the coaxial cable. In choosing lightning surge protectors (and by all means protect your investment!) be sure that the surge protector can allow the low power to pass through it (gas cartridge styles usually do). There are surge protectors designed to work in-line as N-Type connectors adaptable for most cable sizes, and typically one will be placed at the antenna end, and another at the receiver end of the cable.

Comms-Ready Receivers

The manufacturers of dual-frequency GPS (and/or) GNSS gear offer lines of reference station receivers; usually with multiple options for connecting to communications networks or transmission equipment (**Figure 3**). Whether a CORS-style receiver is going to act as single-base broadcast style (via base-radio, dial-up connection, Internet IP, or NTRIP), connect to an RTN, or simply log data for access via web or FTP; these receivers are set up for any (or all) of these connection options. Furthermore, these CORS are capable of multi-tasking and can perform all of these duties and even accommodate redundant comms where desired.

Where bi-directional communications can be established between a receiver and the CPC, these receivers are designed for remote operation, often via a user-friendly web interface. In fact, some of these units are self-contained web servers, complete with their own cyber-security measures (firewalls, authentication, and redundant data logging).

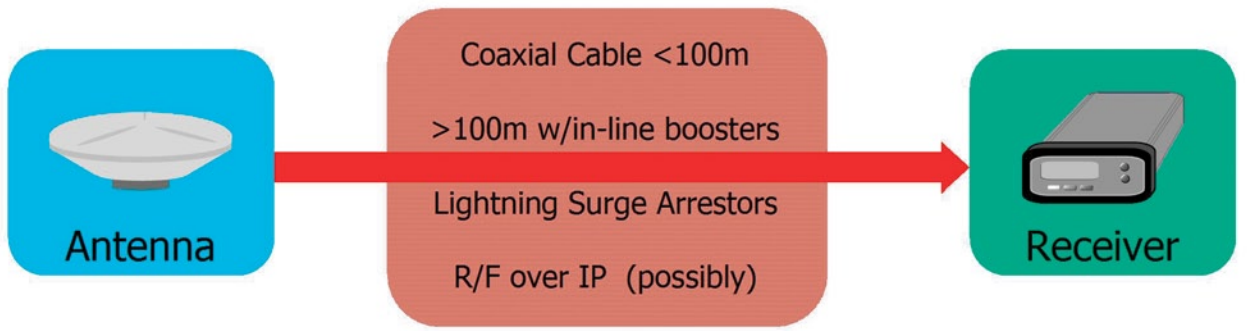

Figure 2 R/F signal from antenna to receiver

Legacy Receivers

If a dual frequency receiver can output 1-sec raw observations via a serial port (without substantial internal latency) then it can act as an RTN CORS. Getting the data to the CPC can follow the same route as the Comms-Ready receivers; with the addition of low-cost communications interfaces like serial-to-IP devices. The disadvantages of using legacy receivers go beyond the limited comms options; limitations on remote operation, limited (or no) multi-constellation support, and possibly poorer quality.

Connection Via Internet

The most commonly utilized communications options, and typically those that can be achieved with the lowest costs are Internet based. Arguably the rapid growth of RTN and the ability to connect wide regional networks of CORS in real-time has mostly to do with the ability to exchange data via the Internet protocols of TCP/IP and UDP, simply and inexpensively.

Unless the developer of an RTN has access to an existing wide regional (perhaps proprietary) communications network (*e.g.*, radio, microwave, WAN) then the best option is the public Internet, and interfacing local comms (*e.g.*, LAN, wireless, satellite) to the Internet. It sounds (and should be) as simple as getting the observation data from the CORS to the Internet, and then from the Internet to the CPC.

The devil, of course, is in the details. These details are governed by what comms options are available, but more challenging so are aspects of what comms options are "allowed". If a CORS is to be installed at a location where commercial Internet access is available (*e.g.*, DSL, broadband wireless, etc.) then the security risks are at the CPC and receiver level (both have accommodations for IP filtering and localized firewalls). But more often than not, the desire to keep comms costs low usually means that some other local system is leveraged to provide the connection to the Internet (*e.g.,* corporate LAN).

When a local system (Information Technology network, LAN, or WAN) is being utilized, and unless all communications between each CORS and the CPC are wholly contained within that local network, the interface with the public Internet causes the most concern. Flow, Firewalls, and Familiarity (or lack thereof) are the 3 Fs that can turn a simple connection into excruciating and frustrating negotiations that can take weeks, months (or, in some cases, years).

By **flow**, we mean **bandwidth**. This is the easy one. Assure the respective Information Technology (IT) folks that the 'flow' of raw data is often less than 500 bytes per second (tiny) and will not "hog the pipe". By **familiarity** we mean **knowledge of RTN and GPS/GNSS**. A typical corporate IT network does not have to deal with a lot of remote sensors (which is essentially what a CORS is). Make the analogy of weather stations, SCADA, or other command/control/monitor systems.

The biggest heartbreak is the connection to the public Internet; most IT networks do not want (or allow) external sources to "request" a connection to their network. With respects to an RTN, the optimal type of connection is one where the CPC can initiate the contact with the CORS (*i.e.*, a port on the CPC server is a "socket client" to a "socket server" port on the CORS) to request the data stream. Another advantage of this relationship is that it can accommodate remote operation of the CORS from the CPC and ad-hoc requests for other station quality data, meteorological info, ephemeris, or data from other sensors. This requires a "static IP" for the CORS (as opposed to a Dynamic IP or DHCP like your home Internet connection that can change as often as each new login). Though **firewall** rules can be established to filter only for requests from the CPC, most IT concerns are (arguably) justified in steering clear of this configuration.

Where the CPC is not allowed to initiate communications, then the next best option is to configure the CORS to act as a "socket client" to a dedicated "socket server" port on a CPC server (this can also be filtered to only accept connections from the IP of the receiver). In this scenario, the local IT "call the shots" and initiate all communications sessions; something much more easily managed with firewall rules. The downside is that many of the features of a comms-ready receiver are not available: remote operation, redundant logging, and web interface. And some comms-ready receivers cannot be configured as a device client, and must be connected to the Internet via a "serial-to-IP" device. This "device

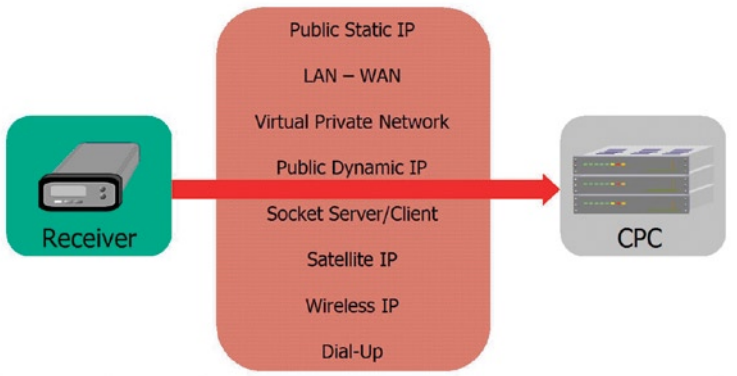

Figure 3 Connection via the Internet

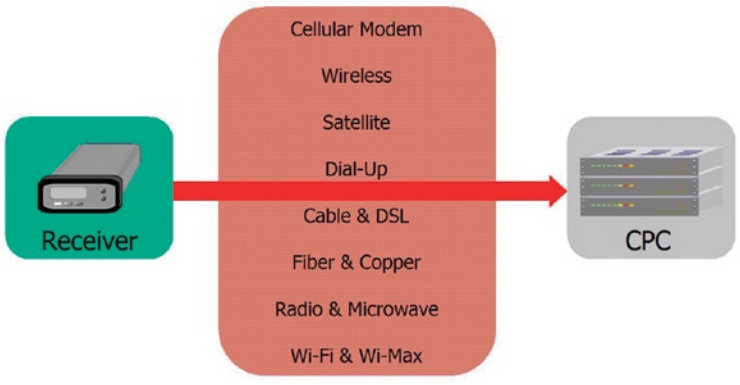

Figure 4 Commercial Internet connection options

client" style may be the only option if a "dynamic IP" Internet connection is all that is available.

The preferred method of connection for legacy receivers to the Internet is also a static IP, and the ability of the CPC to initiate communications. This can be configured on a simple and inexpensive (<$150) serial-to-IP device (usually the size of a deck of cards with one or more serial ports on one end and one or more Ethernet ports on the other). If the "device client" option is all that is allowed, or if a dynamic IP is all that is available, then the serial-to-IP devices can be configured as such.

Another option is to extend the IT network safely through a "tunnel" in the public Internet via VPN (Virtual Private Network). VPN utilizes dedicated servers and authentication clients to provide this very secure pipe. The downside is more potential points of failure and the possible need for a lot more equipment at each remote site.

On these same lines, NTRIP (the internationally accepted protocol for Internet transmission of GNSS data and corrections) is solving these Internet security issues for many. The same protocol and authentication tool as used by the majority of rovers to connect to the CPC for correction data can be used in a similar manner to connect CORS to the CPC. This can utilize the same port as the rovers and send data directly to mount points on an NTRIP caster residing within the CPC (or an external location). The CPC can simply "pick-up" the data from the NTRIP caster. (Watch for an upcoming article on NTRIP.) The preceding discussion has focused on Internet connections via a local LAN (which can pose the most security concerns, both real and imagined), but what of the other methods to directly connect to the Internet?

Other Internet Connections

Essentially, all of the devices and methods listed in **Figure 4** do the same thing; they provide the link from the receiver to the Internet, and the less steps the better. If your antenna is mounted on or next to a facility that already has Internet connectivity (be that via the LAN/Intranet, or say a commercial source like DSL) then plug in and deal with whatever security concerns there may be. But if there is no Internet connectivity nearby, or if trenching a line over to your CORS is costly, then you might want to consider some of the other options.

If it is a matter of jumping those last few hundred yards (or mile) to the nearest Internet connection point, then perhaps a few hundred dollar pair of radios that can transmit the serial data (with those little YAGI-style antennas) can do the trick; other folks set up a small Wi-Fi network; but remember, the more "moving parts" the more potential for failure.

While something like a direct hard-wired T-1 span (fiber *et al*) to an Internet node is wonderful, it is a very costly option, and complete overkill with respects to the bandwidth needs of a CORS. This does provide you with static IP capability, but you can request static IP through most commercial DSL providers and an increasing number of cable Internet providers (though at an additional cost). These are all good options with surprisingly good track records. The downside is a monthly fee.

Dial-up via conventional copper should be viewed as a last resort; there are few guarantees of service for data via analog, and then you have the added headache of keeping those finicky modems happy and connected (folks end up resorting to dial-up power strips and lighting timers to cycle the modems). A dedicated phone line can help a bit, or perhaps pair of modems (one at the CORS and another at a location with a good Internet connection) are other configurations that have been tried with varying levels of success. Dial-up via cellular can be costly (especially if per minute) and terribly flaky, but using cellular (phone or modem) to connect to the Internet via IP (which is typically how they are used for rover operations) have also been utilized with varying levels of success. How to deal with dropped connections and getting the devices reconnected is the headache.

Broadband wireless in most instances utilizes the same cellular network and you are provided with a cellular card-modem or modem box. There is usually no option for static IP (although there are some intriguing Internet services by third parties that can offer to keep rerouting the dynamic IP the broadband provider keeps generating for you to a virtual static IP…but for a fee). The coverage for these broadband cellular services is increasing rapidly, and is a viable option if you can figure out contingencies for resetting the components after outages.

Wi-Fi, Wi-Max, and other options in the 80211 family are becoming readily familiar at the consumer level with the advent of home wireless networks. The drawback is limited range – 1,500m at best for Wi-Fi . Wi-Max has been heralded by the recent wave of wide-area wireless Internet providers, basically DSL speed or near-cable-speed portable modem. An excellent choice where available, but monthly fees apply, and static IP is rarely available.

Satellite has the disadvantage of often having high latencies and there can be a lot of variation in these offerings. The best bet is to test the latency of a solution before committing. There are many examples of networks that have tapped satellite, even through inexpensive satellite-Internet providers, where no other options existed. While sub-second latencies can be maintained to a manageable level (90%+) such systems can be susceptible to outages and spikes in latency.

I have only listed some of the more common solutions. One could go on for endless pages on the variations and nuances of the options (but the editor would not be very happy).

The most important thing you have to do is keep your eye on the underlying simplicity of the need: get the observation data to the CPC in real-time, and that usually means simply getting it to the Internet. This should not be a painful experience, and if someone stands in the way and gives you too many reasons why it can't be done, go ask somebody else.

Nolite id cogere, cape malleum majorem. (Get a bigger hammer.)

Gavin Schrock is a surveyor in Washington State where he is the administrator of the regional cooperative real-time network, the Washington State Reference Station Network. He has been in surveying and mapping for more than 25 years and is a regular contributor to this publication.