HOME

Relative Sensor with 4π Coverage for Formation Flying Missions

Jeffrey Y. Tien[*], George H. Purcell, Jr.[†],Jeffrey M. Srinivasan[‡], and Lawrence E. Young[§]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109-809

The Terrestrial Planet Finder (TPF) pre-project, an element of NASA's Origins program, is currently developing two architectures for a mission to search for earth-like planets around nearby stars. One of the architectures being developed is the Formation Flying Interferometer (FFI). The FFI is envisioned to consist of up to seven spacecraft (as many as six "collectors" with IR telescopes, and a "combiner") flying in precise formation within ±1 cm of pre-determined trajectories for synchronized observations. The spacecraft-to-spacecraft separations are variable between 20 m and 100 m or more during observations to support various configurations of the interferometer in the planet-finding mode. The challenges involved with TPF autonomous operations, ranging from formation acquisition and formation maneuvering to high precision formation control during science observations, are unprecedented. In this paper we discuss the development of the formation acquisition sensor, which uses novel modulation and duplexing schemes to enable fast signal acquisition, multiple-spacecraft operation, and mitigation of inherent jamming conditions, while providing precise formation sensing and integrated radar capability. This approach performs delay synthesis and carrier cycle ambiguity resolution to improve range measurement, and uses differential carrier cycle ambiguity resolution to make precise bearing angle measurements without calibration maneuvers.

Nomenclature

B	=	baseline (distance between antennas)
BPSK	=	binary phase shift keying
BOC	=	binary offset carrier signal structure
CDMA	=	code division multiple access
f_c	=	carrier frequency
FDMA	=	frequency division multiple access
PR(N)	=	pseudo-random (noise)
TDMA	=	time division multiple access
TPF	=	Terrestrial Planet Finder mission
Δ	=	excess signal path
$\Delta\nu$	=	difference in frequency between two components of the ranging signal
$\Delta\tau$	=	differential delay off arriving signals
θ	=	angular deviation of arriving wave front
λ_{ch}	=	chip length of pseudo-random ranging code
ν	=	generic frequency
σ	=	uncertainty (standard deviation)
σ_ξ	=	uncertainty of the variable ξ
ϕ	=	generic phase

[*] Formation Sensor Testbed Lead, 4800 Oak Grove Drive, M/S 138-212
[†] Senior Member of the Technical Staff, 4800 Oak Grove Drive, M/S 238-600
[‡] Supervisor, Advanced Radio Metric Instrument Development Group, 4800 Oak Grove Drive, M/S 138-212
[§] Supervisor, GPS Systems Group, 4800 Oak Grove Drive, M/S 138-212

I. The TPF Mission

Terrestrial Planet Finder (TPF) is an important mission in the Navigator Program, the planet-finding part of NASA's Astronomical Search for Origins theme. TPF will support the Origins science goals by determining whether habitable or life-bearing planets are orbiting nearby stars. The mission will accomplish its goals by surveying as many as 150 solar type (F-, G-, and K-type) stars in the solar neighborhood, suppressing the radiation from the parent star and looking for the faint optical or infrared reflections from planets in the "habitable zone." If bright planets are found, TPF will make low-resolution spectral observations, looking for evidence of habitability using markers like O_2, CO_2, and H_2O. Finally, TPF will make very sensitive low-resolution spectral observations of the most interesting planets using biomarkers such as O_2, O_3, and CH_4.

TPF expects to use two architectures: (1) a visible/near-infrared coronagraph (TPF-C) and (2) a formation-flying mid-IR nulling interferometer (TPF-I). The coronagraph has the advantages that: (1) the required resolution can be achieved with a smaller filled aperture at the shorter optical wavelengths, and (2) optical telescopes can operate at ambient temperature (about 300 K), while a thermal IR telescope has to be cooled to about 40 K. On the other hand, the contrast between the star and planet is much greater at visible (about 10^9) than at IR (10^6) wavelengths, with the consequence that the required suppression of stellar emission is much easier (but still very hard) to achieve in the IR. Also, the individual apertures can be smaller in an interferometer (3–4 m) than in a coronagraph (6.5–8 m), even though the interferometric baseline is larger. At present, TPF-C is scheduled to be launched in about 2014, and TPF-I will follow before the end of the decade.

In the sections that follow, we focus on one of the technologies required for the TPF-I mission: an "acquisition sensor" that maintains the integrity of the array throughout its lifetime and positions the spacecraft so that a more accurate relative sensor can assume control in preparation for interferometric observations.

II. Requirements on Formation Flying

The acquisition sensor provides to a central processor observations from each spacecraft that enable the processor to determine the relative locations of all the spacecraft in the array. Thus the sensor serves two functions: (1) It enables the array to maintain a known state, preventing both collisions and loss of spacecraft from the array. (2) It helps position the spacecraft for interferometry, so that the lines of sight of a sensor (or a succession of sensors) having a narrower field of view, but greater accuracy, can take over. In order to perform these functions, the acquisition sensor needs to satisfy the following requirements.

(1) The sensor must do its job well regardless of the eventual architecture of the array. At present various architectures are under consideration. Although the dual-chopping Bracewell architecture, with four similar collinear collectors and a combiner (as shown in Fig. 1), is the current baseline, other architectures are under consideration; and the eventual choice may differ from the baseline both in the number of spacecraft and in their arrangement. We want a sensor that will work with any plausible architecture.

(2) The sensor must be able to acquire without a priori knowledge. That is, the sensor on each spacecraft must be able to lock onto the signals from all the other spacecraft without any initial information about their locations or velocities.

(3) Each spacecraft must be able to sense each other spacecraft that is not geometrically blocked. This feature is necessary for acquisition or recovery from a chaotic state.

(4) The sensor must be able to acquire some spacecraft while tracking others. This requirement is necessary for orderly acquisition, given that the procedures for acquisition and tracking are quite different.

(5) The sensor must be able to calibrate bearing angles without maneuvering the spacecraft. Bearing angles (elevation and azimuth, for example) are derived from the phase differences of a received signal at three or more receiving antennas. However, there is initially an n-cycle ambiguity in these differences. The ambiguities can be resolved by a number of schemes that involve translating or rotating the spacecraft. While these procedures may be tractable for two spacecraft, they are intolerable for a large array that could consist of seven or more elements.

(6) The sensor must be able to maintain the array if the sensor fails temporarily on one spacecraft. This requirement assures that there will be no collisions or loss of spacecraft from the array after the sensor resets on one spacecraft, rendering it deaf and dumb for several minutes. A more serious outage might isolate a spacecraft for several days.

(7) The sensor must be compatible with the temperature distribution imposed by the IR telescopes: a warm side of the spacecraft, including most of the electronics, at something like 300 K and a cold side, including all the optics, at about 40 K.

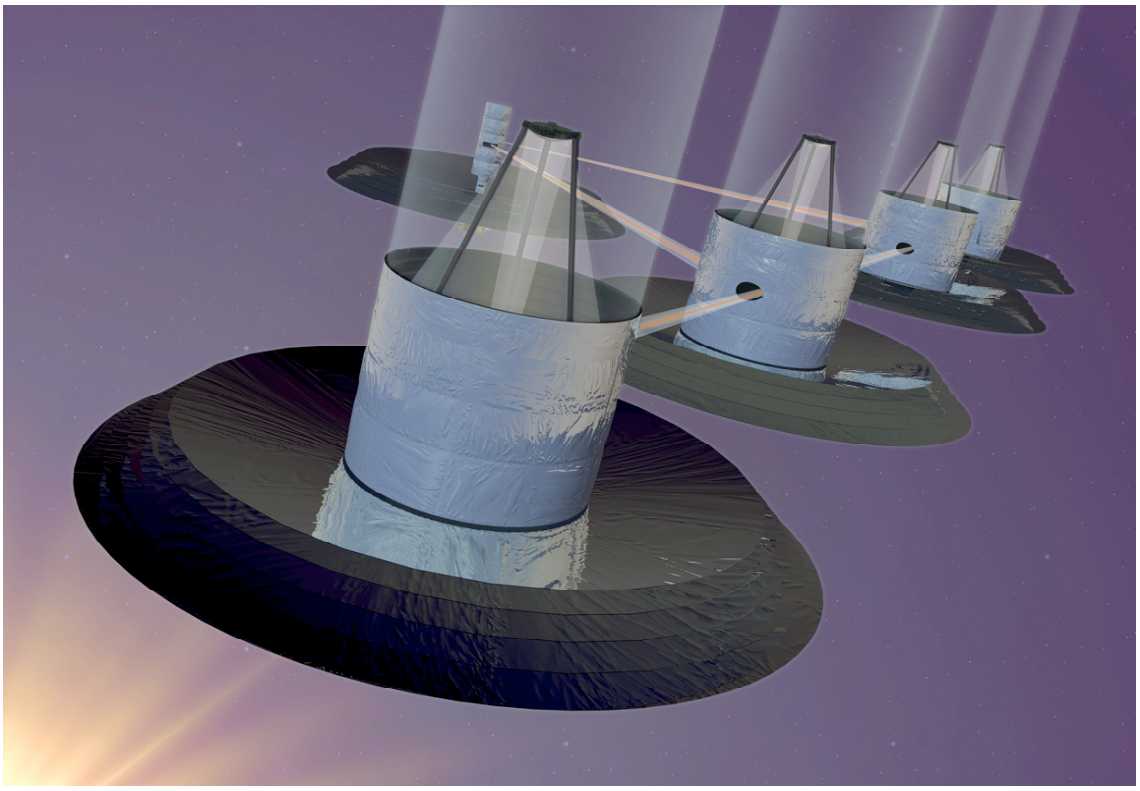

Figure 1. Formation-Flying Interferometer.

Once these general requirements are met, the sensor must provide knowledge to align all the spacecraft well enough so that a more accurate sensor with a smaller field of view can initialize itself. We therefore have the following performance requirements.

In Table 1, note that the nominal operating range for the interferometer is 16–100 m between centers. The envelope value allows for operation with reduced accuracy to 10 km for recovery from faults. For the radar mode, which senses reflected signals when a transmitter is disabled, operation extends only to 200 m.

Table 1. Performance Requirements for Acquisition Sensor.

Requirement	Normal Operation	"Radar"
Operating envelope		
Inter-s/c range (m)	16–10,000	16–200
Inter-s/c bearing (sr)	4π	4π
Inter-s/c range rate (m/s)	2	2
Inter-s/c bearing rate (deg./s)	1	1
Range (m, 1σ)	0.5	1.0
Range rate (m/s, 1σ)	0.001	0.01
Bearing (degrees, 1σ)	1	30
Bearing rate (arcmin/s, 1σ)	1	--

The following sections describes a sensor that meets these requirements.

III. Design Trades

Because of the constraints placed on the acquisition sensor by the spacecraft, and the performance requirements it must meet, it faces formidable design challenges. Below we discuss some of these challenges and how they have been met.

A. Sensing Technology

The most fundamental choice to be made concerning the acquisition sensor is the sensing technology. Radio and optical are the obvious choices; no other possibilities have been seriously considered.

Optical sensors tend to have a relatively small field of view. Hence, to obtain global coverage one must either have a large number of them or extend their field of regard by scanning them. However, the number of individual sensors required for instantaneous coverage may be excessive, and scanning requires moving parts, which pose various problems in the flight environment. Perhaps a stronger objection to scanning is that it requires motors, which generate heat. The problem with heat is that in order to get global coverage, some of the sensors must be placed on the cold side of the spacecraft. Since the spacecraft are passively cooled to 40 K, heat sources must be strictly limited to the milliwatt level.

Another problem with optical sensors is the sun. When the sensing direction approaches the sun, as it may, it is hard to prevent sunlight from saturating or otherwise confusing the sensor. We might try to overcome this problem by using a spectrally narrow (laser) signal and a correspondingly narrow filter on the input, but there is difficulty in making the filter narrow enough to discriminate effectively in favor of the signal and at the same time wide enough to allow for the width of the signal and the variability of the signal and filter.

RF sensors have problems of their own, but they appear more manageable. Multiple antennas are needed to achieve complete directional coverage, but they can have a wide field of view, and there are no moving parts. Nevertheless, heat transfer to the cold side of the spacecraft remains a problem, because transmission lines must connect the cold antennas to the warm electronics. Heat is conducted along these lines, and for the transmitters, signal power is dissipated as well. On balance, an RF sensor is preferable.

B. RF Frequency

Having chosen an RF sensor, we have to pick the frequency at which that sensor will operate. The discriminators are system temperature, multipath, spacecraft accommodation, and regulatory allocation. We can't simply select a convenient frequency based on the needs of the mission: we have to fit into an allocation of the National Telecommunications and Information Administration (NTIA) and the International Telecommunications Union (ITU). Table 2 shows the bands that are available for spacecraft-to-spacecraft communication between 1 and 100 GHz.[1]

Table 2. Frequency Allocations

BAND	FREQUENCY
S	2025–2110 MHz
	2200–2290 MHz
Ku	13.75–14.30 GHz
	14.50–15.35 GHz
Ka	22.55–23.55 GHz
	25.50–27.00 GHz
	32.30–33.40 GHz
W	59–64 GHz
	65–71 GHz

Among these frequencies, phase multipath, proportional to wavelength, is smaller at the higher frequencies. However, this factor is not crucial, because the requirement on bearing (whose accuracy depends mostly on the phase observable) is fairly relaxed. The lowest frequencies (at S band) have the great advantage that we can use coax rather than waveguide for the RF transmission lines. Furthermore, system temperatures are smaller at the lower frequencies. As a bonus, S-band hardware is generally less expensive than it is at the higher frequencies. On the whole, S band meets the requirements most easily and is our choice.

C. Signal Structure

Recall that the acquisition sensor has four functions:

(1) It acquires the signals transmitted by other spacecraft.

(2) Having acquired the signals, it makes measurements that allow the determination of the ranges to the other spacecraft.

(3) It also makes measurements that allow the determination of the directions of the other spacecraft. (These directions are in a local reference frame that moves with the receiving spacecraft, and they might be expressed as azimuth and elevation, for example.)

(4) If the sensor on one of the spacecraft fails temporarily (because of a power interruption, for instance), it enables the other spacecraft to maintain the array indefinitely, preventing their disabled comrade from drifting out of contact or colliding with one of the others.

Designing a signal structure that fulfills these requirements expediently requires some subtlety, and we draw heavily on our experience designing GPS receivers. Range is determined in the first instance by a pseudo-random (PR) ranging code modulated onto a carrier, and bearing angles are calculated from the phase differences of carriers arriving at different antennas, as shown in Fig. 2. For simplicity, the figure shows only one bearing angle and two antennas; to determine both angles, the signal must be received at three antennas.

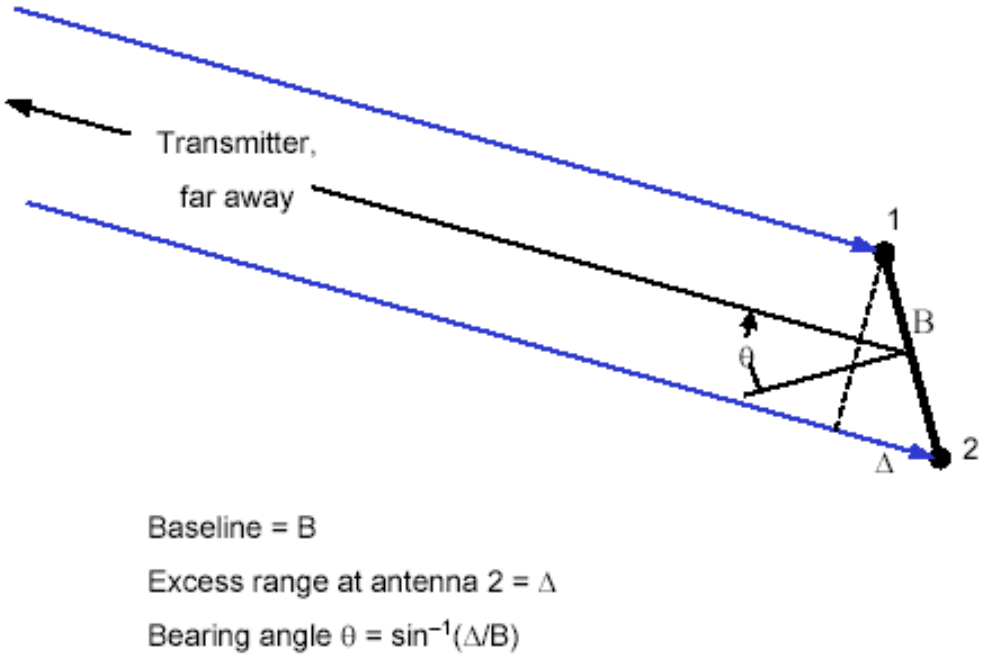

Figure 2. Bearing Angle of Transmitter

A complication of the bearing-angle determination is that there is an integer-cycle ambiguity in measuring the phase difference between signals received at two antennas, and generally the range determinations are not good enough to resolve the ambiguity. We can resolve the ambiguity in several ways by maneuvering the spacecraft, using either rotation or translation. However, since the maneuvers consume precious time (not to mention propellant) during acquisition, have to be repeated whenever the receivers "lose lock" on a signal, and involve

several spacecraft, this solution was rejected as impractical. It would be better to use a signal structure that permits resolution of the ambiguity internally, eliminating the need for maneuvers.

Such a signal structure has been devised, and as explained below, it also speeds acquisition. As shown in Fig. 3, this signal is a generalization of binary offset carrier (BOC) modulation that we informally call "ultra-BOC." It consists of a central carrier modulated by a pseudo-random noise (PRN) code at 10 Mchips[**]/s, two inner tones modulated slowly with data at about 100 bits/s, and two outer tones that are unmodulated. (The central component may also be modulated with data at perhaps 10 kbit/s, but that modulation does not bear on the current discussion.)

With this signal, acquisition follows the steps shown in red in Fig. 3. The search starts with an outer tone (step 1), which is easy to find because it requires a search over frequency but not time. Once this tone is found, we know from the frequency which transmitter sent it and can do a relatively short search for the inner tone (step 2), whose modulation contains some fixed bits and a time code, perhaps at intervals of a second. Knowing the time, we then do another short search over delay to acquire the rapid PRN modulation on the central carrier (step 3).

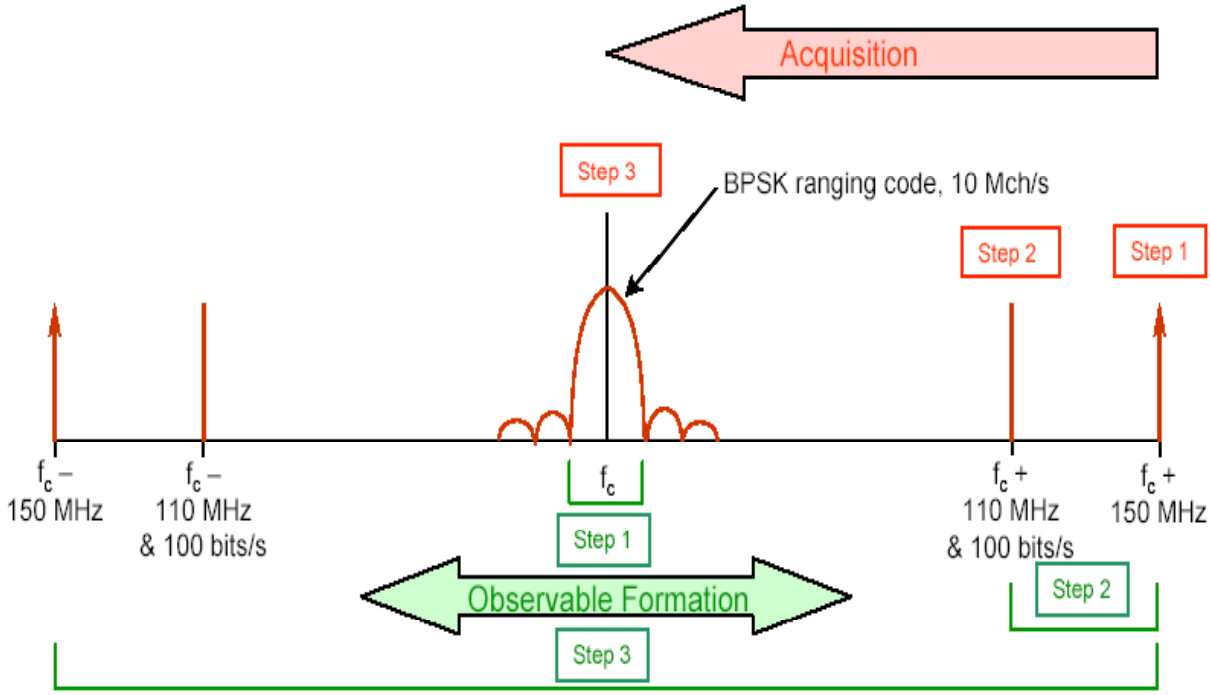

Figure 3. Ultra-BOC Signal Structure and Processing Steps.

Generation of the observables proceeds in the opposite sense, working from the central component outward, as shown in green in Fig. 3. In the first step, we start from the delay given by the ranging code. Then in step 2 we form a "synthesized delay" by using the closely spaced tones on one side (inner and outer tones from either side) of the center frequency. That is, we compute

$$\Delta\tau = (\phi_1 - \phi_2)/(\nu_1 - \nu_2), \tag{1}$$

where the ϕ's are the observed phases at the two frequencies and the ν's are the frequencies themselves. Because the integral cycles of phase are unknown for each measurement, there is an n-cycle ambiguity in the phase difference. We resolve this ambiguity using the known delay from the ranging code. For the procedure to work, the error in the ranging-code delay must be much less than the error made by offsetting the phase difference by a half cycle. Furthermore, the instrumental components of the phases must have been calibrated to a small fraction of a cycle. Hence the procedure is more likely to succeed for the differential phases (between receiving antennas) used to compute bearing angles than for the undifferenced phases that could be used to compute ranges.

[**] A chip is the shortest time between transitions in the PRN code.

In the third step of observable formation, we continue in the same way, forming a second synthesized delay from the two outer tones. This step helps because the error in the delay determination is proportional to $1/(\nu_1 - \nu_2)$, which is inversely proportional to frequency difference. The same caveats discussed above apply.

In the last step of the ultra-BOC procedure, we expect to resolve the phase ambiguity at the carrier frequency itself. Thus the process has removed the integer-cycle phase ambiguity from the differential phase measurements using the signals alone: no spacecraft maneuvers are necessary.

Before leaving the subject of ultra-BOC, we need to determine what are the optimum spacings of the various components. Without going into the details of the calculations, we can say that Eq. (1), along with the assumptions that (1) the error on the output delay at each step is much smaller than the error on the input delay, and (2) the error on the input delay at each step is less than 1/6 of the delay ambiguity at the output of that step (so that the probability of an error at each step is approximately 0.0027), leads to

$$\Delta\nu_1 \le \frac{c}{6\sqrt{2}\sigma_\rho}, \tag{2}$$

$$\frac{\Delta\nu_2}{\Delta\nu_1} \le \frac{1}{6\sqrt{4}\sigma_\phi}, \tag{3}$$

and

$$\frac{\nu_{rf}}{\Delta\nu_2} \le \frac{1}{6\sqrt{4}\sigma_\phi}, \tag{4}$$

where σ_ρ is the system-noise error on range,

λ_{ch} is the chip length of the PRN code, here equal to c/(10 MHz) = 30 m,

$\Delta\nu_1$ is the frequency difference of the narrow separation (Hz),

$\Delta\nu_2$ is the frequency difference of the wide separation (Hz), and

σ_ϕ is the uncertainty of an undifferenced phase measurement (cycles).

Assuming that σ_ρ = 0.5 m and σ_ϕ = 0.01 cycle, we have tentatively chosen $\Delta\nu_1$ = 40 MHz and $\Delta\nu_2$ = 300 MHz to go along with $\nu_{rf} \approx$ 2200 MHz.

D. Antenna Configuration

Getting true instantaneous 4π-steradian coverage requires many antennas. If each antenna's beam extends to 70° off boresight, then it can see almost a third of the sky. Because the beams overlap (unavoidable if we are to attain complete coverage), it takes four transmitting antennas to cover the entire sphere. These antennas are most simply pointed in the directions of the vertices of a regular tetrahedron, as shown in Fig. 4. In that case, the beam width (cone half-angle) required for complete coverage is $\cos^{-1}(1/3) = 70.53°$.

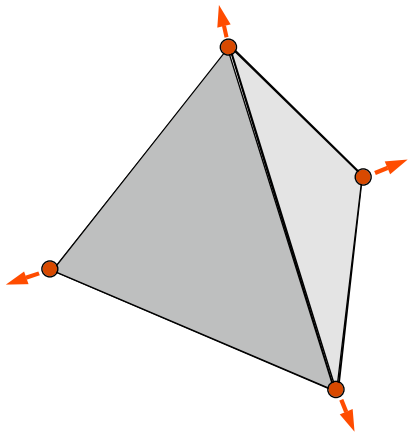

Figure 4. Geometry of Transmitting Directions.

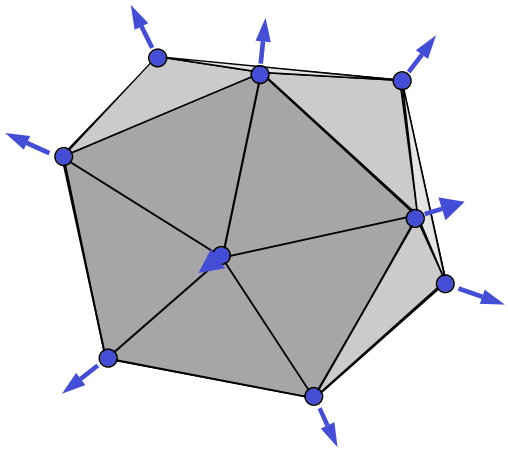

Figure 5. Geometry of Receiving Directions.

Since we need three receiving antennas to be visible from any direction, in order to sense the two bearing angles of the transmitter, we expect to need about three times as many receiving antennas as transmitting antennas, or twelve. These antennas can point toward the vertices of a regular icosahedron, as shown in Fig. 5. The beam angle (again, the cone half-angle) needed to have three antennas looking in every direction is then $\cos^{-1}\left(1/\sqrt{5}\right) = 63.43°$, which is in fact the angular separation of neighboring antennas.

In all, therefore, this scheme involves 16 antennas, which might be placed as shown in Fig. 6. It should be remembered that the issues of (a) the directions of the antenna beams and (b) the placement of the antennas, are distinct but related. One would expect the beam to point outward from the spacecraft at the point of attachment, but the beam size and the spacecraft structure in the neighborhood of the antenna may allow some latitude in direction at a given point.

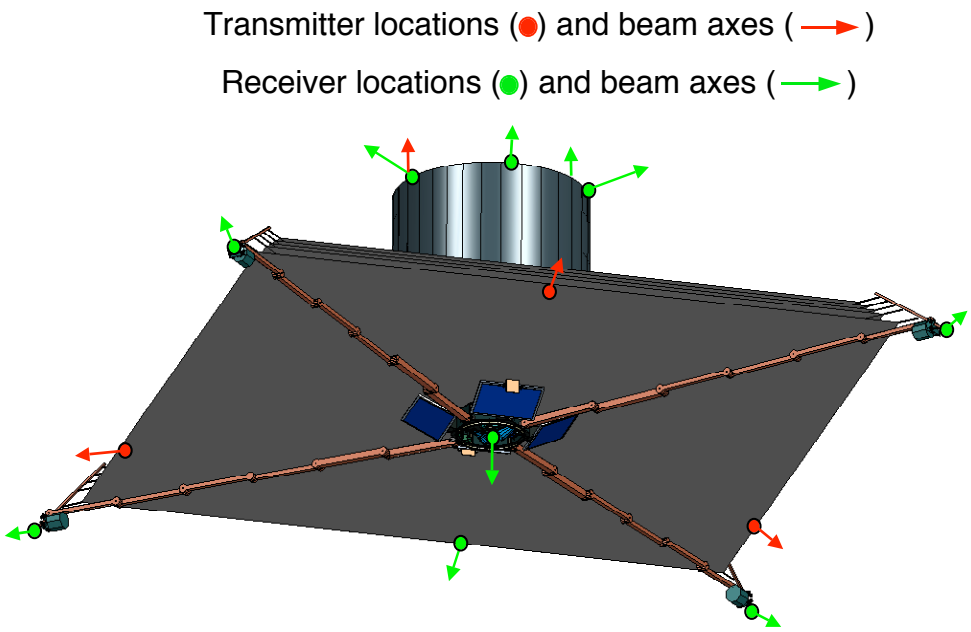

Figure 6. Possible Locations of Antennas on Spacecraft.

Another issue, at least for receiving, is the distribution of antennas over the spacecraft. In principle, we might gain some advantage by clustering receiving antennas together at convenient points. However, the arithmetic of the bearing-angle measurements makes this clustering difficult or impossible. Imagine that we are on a transmitting spacecraft, looking at the projection onto the plane perpendicular to the line of sight of the triangle whose vertices are the three receiving antennas, as shown in Fig. 7. The uncertainty of a bearing angle of the transmitter, as seen from the receiving spacecraft, is inversely proportional to the projection of the triangle onto the associated coordinate direction. Thus, if the triangle is thin, the bearing-angle determination in the thin direction is relatively bad. Consequently neighboring antennas need to be well separated and uniformly distributed in all directions if the bearing angles are to be well determined.

A final problem for antenna distribution, specific to the current TPF spacecraft configuration, is the presence of the thermal shield, because it effectively partitions the observable directions into two distinct hemispheres, those above and those below the shield. Since the shield is by definition thermally opaque, it tends to be opaque to RF energy as well: a transmitter or receiver on one side of the shield can't penetrate to the other side. This circumstance complicates sensing near the plane of the shield, which is crucial, since the shields of all the spacecraft are in the same plane during interferometry. The solution adopted in the placement of Fig. 6 is to place antennas near the edge of the shield, where they can peek around it.

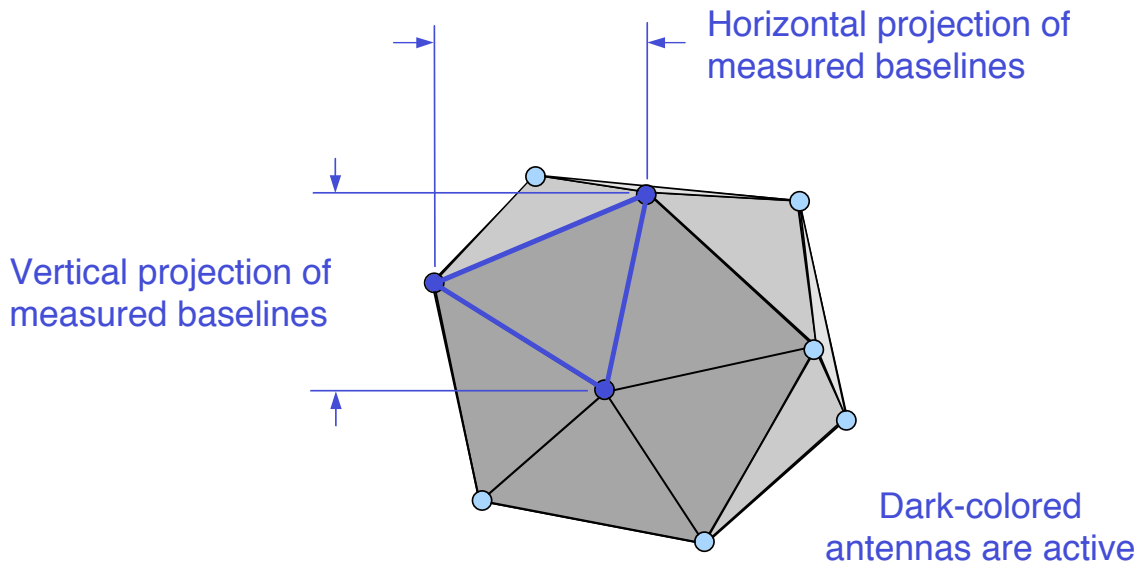

Figure 7. Receiving Antennas As Seen from Transmitter (Schematic)

The current plan is to use a total of 16 antennas on each spacecraft. If this number is unacceptably large, at least two means of reduction are available, alone or in combination:

(1) *Use wider beams.* For, example, if the beams extended to 90° off axis, we could in principle get the required coverage with two transmitting antennas and six receiving antennas per spacecraft. However, the on-axis gain would be lower, multipath would be stronger, and the placement of the antennas would be limited to locations with a hemispherical field of view.

(2) *Relax the requirements.* For instance, we might decide that we need less accuracy in directions not needed for handoff to another relative sensor. In that case we could concentrate our antennas in the handoff directions and reduce the number in directions used only for acquisition. However, we might have to give up considerable performance in order to get a worthwhile reduction in the number of antennas.

At present, all options are open with respect to placement of the antennas, and accommodation with the spacecraft will evolve as the spacecraft itself develops.

E. Interference between Signals

In any particular receiving channel, we may have the wanted signal from a remote spacecraft, unwanted (because they contribute to receiver "noise") signals from other remote spacecraft (or possibly from another transmitter on the same remote spacecraft), and unwanted signals from transmitters on the local spacecraft. The local signals, in particular, may saturate the front end of the receiver and cause severe loss of SNR in the wanted signal. In addition, we need to be able to identify the transmitted signals with their spacecraft and discriminate among them in signal processing.

We can deal with these problems in several ways. One way is code-division multiple access (CDMA), which we intend to use to distinguish among the BPSK-modulated central components of the ultra-BOC signal structure. Like GPS signals, each transmitter has its own code that enables the receivers to distinguish and identify them. Another way is frequency-division multiple access (FDMA): We expect to offset the ultra-BOC tone frequencies on each transmitter, so that the tones (and data sidebands) are distinguishable and don't interfere with one another.

Still another way is time-division multiple access (TDMA) to reduce interference in a receiver from local transmitters. This scheme involves transmitting and receiving alternately, so that the transmitted signal can't interfere with reception.

One way to do this is to have the spacecraft take turns transmitting and receive the rest of the time, as shown for two spacecraft in Fig. 8a. This approach minimizes transmissions and prevents interference between spacecraft completely. On the other hand, it requires synchronization among the spacecraft and would be awkward to

implement for all possible situations in a multiple-spacecraft array involving acquisition, non-functioning spacecraft, and so on.

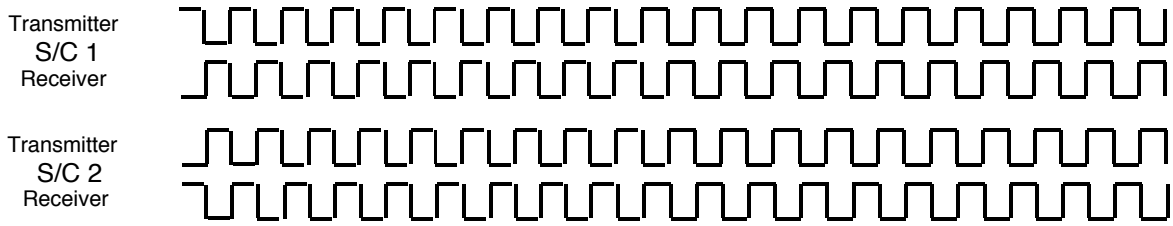

Figure 8a. Synchronous Transmission and Reception.

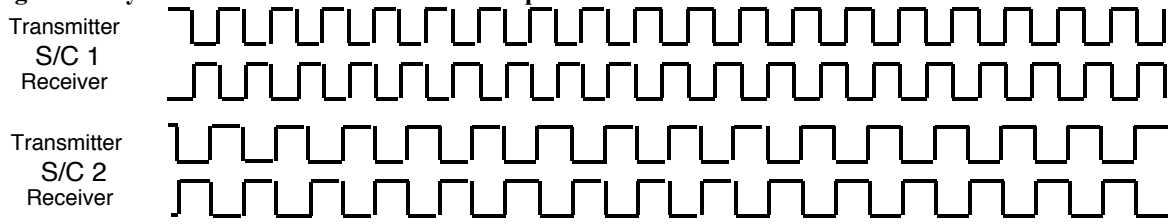

Figure 8b. Asynchronous Transmission and Reception.

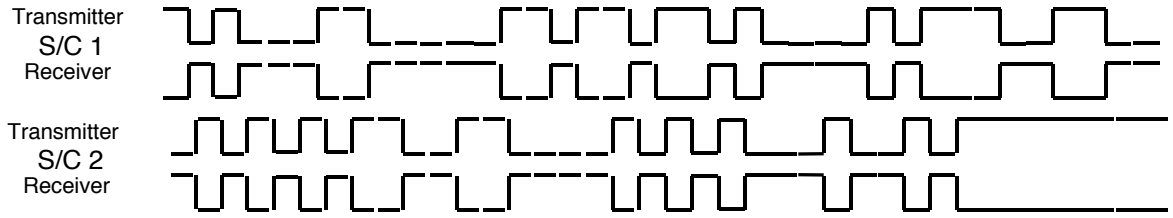

Figure 8c. Pseudo-Random Transmission and Reception.

Figure 8. Coordination of Transmission and Reception to Reduce Interference

An easier procedure to implement is to let each spacecraft alternate between transmitting and receiving without consideration of what the other spacecraft are doing. The most straightforward approach is to follow a fixed pattern, for example transmitting for a certain interval and then receiving for an equal interval and repeating the pattern, as shown in Fig. 8b. In this case the spacecraft should use different intervals to assure that over a period of many intervals, each spacecraft receives each of the others equally. A typical interval for transmission or reception might be 100 μs. Of course, this approach has the disadvantage that the remote spacecraft will interfere with one another's signals, but it does avoid interference from the local spacecraft.

A more exotic approach is to switch the transmitters on and off according to a pseudo-random code that is different for each spacecraft, as in Fig. 8c. The reason for this apparently gratuitous complexity is that for any two-way travel time longer than a chip of the code, a transmitted signal can bounce off a remote spacecraft and have a 50% chance of arriving back at the local spacecraft when the transmitter is off and the receiver is on. Thus, the signal can be used for radar detection of a spacecraft whose coarse sensor is temporarily inoperative. Obviously the chipping rate has to be high for this technique to work for closely spaced spacecraft. If we insist that half the return take place while the receiver is on, then for a minimum distance of 15 m between two spacecraft, the chipping rate has to be at least 10 MHz.

F. Heat Transfer

Heat transfer from the warm side (~300 K) of the spacecraft to the cold side (~40 K) containing the infrared optics must be absolutely minimized. Quantitative limits are being developed, but it's clear that milliwatts are significant, since the ability of the cold side to lose heat passively by radiation is limited. The acquisition sensor is implicated in this problem because of the coax cables connecting the transceiver, on the hot side, with antennas on the cold side. These cables can heat the cold side in two ways: by conduction through the metal conductors and by dissipation of power being carried from the transmitters.

To get an idea of the magnitude of the problem, consider 50 Ω 32081 coaxial cable. Analysis shows that the heat loss to be expected from one of these cables, of the length necessary to connect the transceiver to an antenna on the

cold side of the spacecraft (~5 m), through the temperature difference indicated above, is on the order of 30 mW. Furthermore, the dissipation of energy carried along the cable at S band is about 1dB/m, so we can expect to lose about 68% of the original power in transit. If 10 mW reaches the antenna, then we lose about 22 mW along the way, and some of that will heat the cold side.

There are several ways in which this heat leakage to the cold side can be reduced. One is to minimize the number of antennas there, particularly transmitters. We can get by with a single transmitting antenna on the cold side if the others are allowed to peek upward from the warm side around the edge of the heat shield. Another way is to turn off transmissions to antennas on the cold side when they are not being used. For example, we could certainly arrange for warm-side transmitters to do the sensor's work while the array is in the normal observing configuration for interferometry. A third way would be to make the transmission lines from materials that have high electrical conductivity but low thermal conductivity; unfortunately, according to the Wiedemann-Franz Law[2] the two are proportional for metals (both pure and alloyed), so that avenue appears unrewarding. However, since the high-conductivity region of the cable needs to be only a few skin depths thick (and the skin depth for copper at S band is only about 1.4 μm), we could, for example, use silver- or copper-plated stainless steel inner and outer conductors to get good electrical conduction while obstructing the flow of heat.[3]

The true magnitude of the problem is still unknown, since thermal models are still being developed, and coarse-sensor transmission lines are just one of many sources of heat leakage. Furthermore, the radiating area on the cold side can be made large, so that if the heat conduction and emissivity are made high enough, a leakage rate of several watts can be accommodated at a temperature of 40 K. In view of practical limitations, however, the spacecraft may not be able to support a leakage rate of more than a few tenths of a watt.

IV. The Sensor

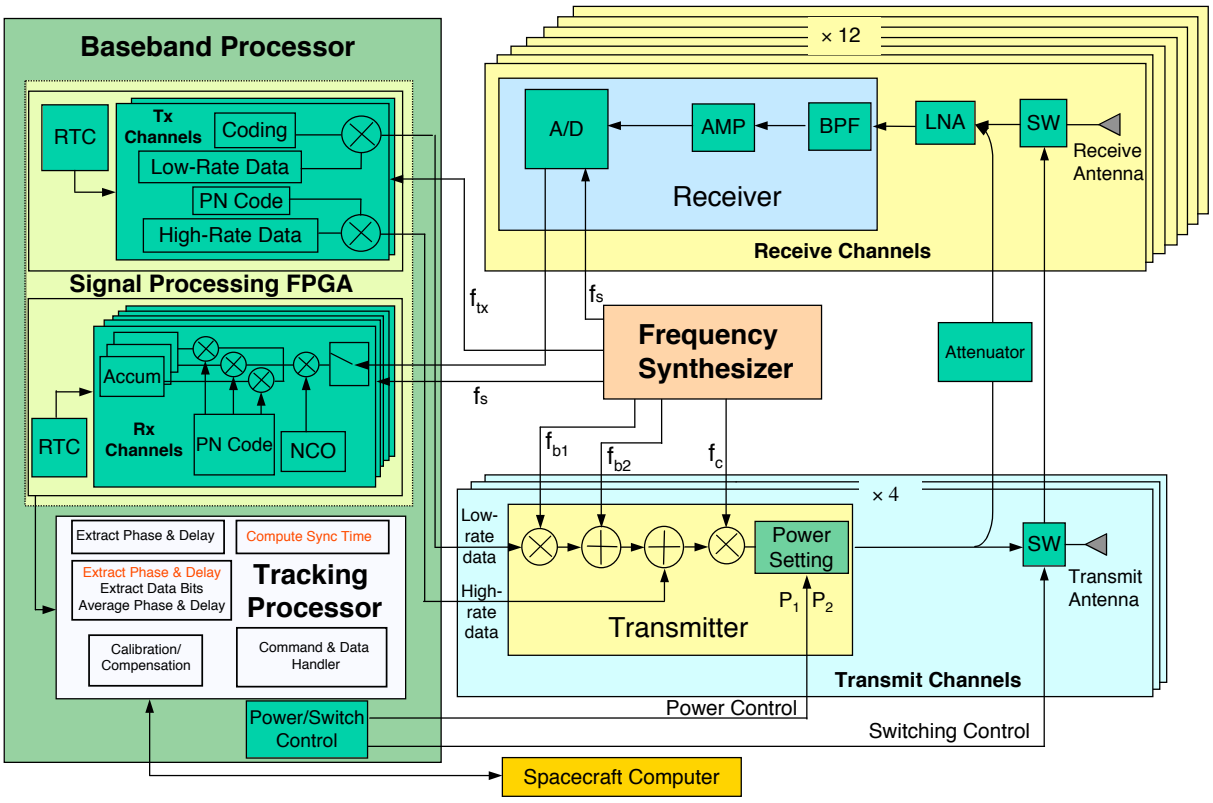

Figure 9. Block Diagram of Acquisition Sensor

Guided by the issues discussed in the preceding section, and others, we have designed a sensor, shown schematically in Fig. 9, that satisfies the requirements for an acquisition sensor for TPF. The generation of the

transmitted signals begins at the upper left in the baseband processor. There is one transmission channel for each transmitting antenna, or four in all. In each channel, the low- and high-rate data of the ultra-BOC signal are combined with codes at the rates specified by frequency f_{tx} and a real-time clock (RTC). These signals then go to the RF section, where they are combined (mixed or added, as shown at the lower right) first with the offset of the inner ultra-BOC tone from the center frequency (see Fig. 3), f_{b1}, then with the offset of the outer tone, f_{b2}, and finally with the carrier, f_c. The power level is set either to P_1 or P_2, depending on the proximity of the remote spacecraft. A very small part of the transmitted signal is routed to appropriate receivers on the same spacecraft to be used as a calibration signal. The rest passes through a switch (SW) that turns the signal off and on to prevent interference with reception.

In the receivers, after being picked up by an antenna, the remote signals pass through a switch that allows reception only when the local transmitters are off. Then the local calibration signals are injected, as needed, and the composite signal goes through a low-noise amplifier (LNA) and a band pass filter (BPF)/amplifier chain before it is down-converted and digitized by the A/D converter. The digitized data are then output to the baseband processor. In the baseband processor the ultra-BOC signals are each phase counter-rotated to DC and then correlated with the appropriate PRN code, if any, to determine the range from the central ultra-BOC signal and the phases of all the signals. This processing resembles the GPS paradigm, with early, prompt, and late integrations to find the peak of the correlation. Each component of the ultra-BOC signal is processed independently.

The tracking processor controls the entire process, accepting input from the spacecraft computer, extracting phase and delay observables, recovering the data bits (both fast, on the central component, and slow, on the inner ultra-BOC tones), averaging and calibrating the observables, assigning time tags to them, and finally sending the calibrated observables to the spacecraft at one-second intervals. This processor also controls the power output of each of the transmitters and switches the transmitters and receivers off and on according to one of the schemes discussed above.

Central to the whole process are a frequency standard and synthesizers that provide coherent signals to the various parts of the sensor for generating and demodulating the components of the ultra-BOC signal, controlling sampling, and keeping time.

V. Testbeds

The acquisition sensor is complicated and contains many pieces of unproven technology. The TPF Project can ask, for example: Can the receiving and transmitting antennas be accommodated on the spacecraft? Can the sensor really attain 4π-steradian coverage? Can multipath be managed adequately? Can the data be calibrated well enough so that ultra-BOC will work as intended? Can we track a combination of close and distant spacecraft? Can we hand off gracefully from one antenna to another as the bearing angles change? Can the computational machinery keep up with all the model calculations, correlations, calibrations, I/O, and other functions required for many complex signals? All these technologies must be shown to work on the ground before they can fly.

To validate the various new technologies, two ground testbeds have been devised. The first is an indoor testbed, shown schematically in Fig. 10. Here three "spacecraft" will be connected by cables with attenuators adjusted to simulate space loss. This testbed will demonstrate:

1. That we can generate and interpret the ultra-BOC signal with the needed accuracy,
2. That we have the dynamic range to cope with a mixture of close and distant spacecraft,
3. That internal calibration works well enough to make ultra-BOC feasible,
4. The operation of the software being written to manage the sensor, and its interaction with the hardware, and
5. The ability of the software to keep up with the required computations.

However, the full functionality of the sensor can be demonstrated only in a spacecraft-like environment. For this purpose we are designing a second testbed outdoors, as shown conceptually in Fig. 11. It will be implemented after the indoor testbed has done most of its work and will build on the validation achieved there. In the outdoor testbed the RF signals will be broadcast through the air between actual antennas mounted on three realistic models of the spacecraft that can be translated and rotated. The outdoor testbed will demonstrate:

1. 2π-steradian coverage with two transmitting antennas and six receiving antennas on one modeled spacecraft, including handoff from one antenna to another as the bearing angles of the remote spacecraft change,
2. Operation of ultra-BOC in a flight-like environment, giving bearing angles without calibration maneuvers,
3. Satisfaction of the performance requirements, confirmed by truth sensors, between 16 and 1000 m range,
4. Operation of the radar mode, with the sensor disabled on one spacecraft,
5. Management of multipath adequate to meet the performance requirements, and

6. Ability of the sensor to track a distant spacecraft in the presence of one nearby.

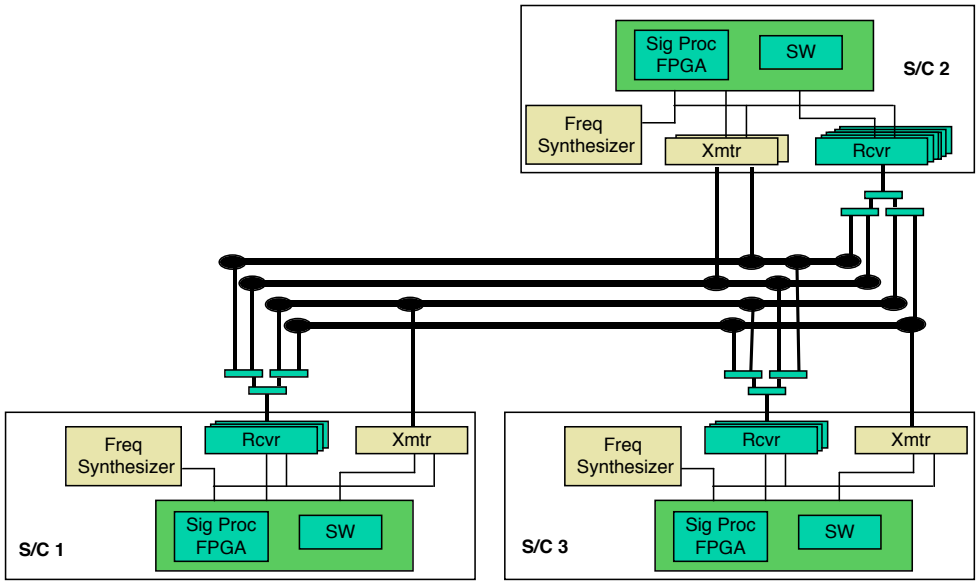

Figure 10. Indoor Testbed (Schematic)

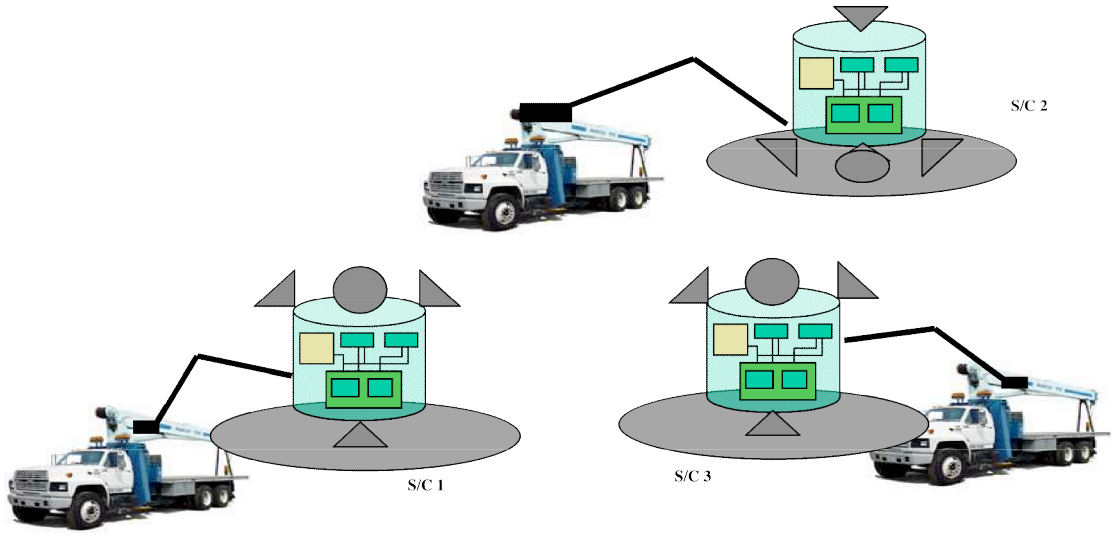

Figure 11. Outdoor Testbed (Concept)

Together, these two testbeds will provide the means to advance the technology level of the sensor to flight readiness. They will show that the sensor can do its part in realizing the TPF-I Mission.

VI. Summary

Formation-flying interferometry has been selected as the architecture for TPF-I, the second half of a mission to detect Earth-like planets orbiting nearby stars. This architecture requires a reliable "acquisition sensor" that can maintain the integrity of an array of at least five spacecraft and position them for handoff to another sensor for more

accurate control. According to current requirements, the acquisition sensor would need to have a 4π-steradian field of view and provide relative positioning between spacecraft with an uncertainty of 50 cm in range and 1° in the bearing angles, at separations of 16 to 10,000 m.

Such a sensor does not exist and faces formidable technical challenges. First, the technology must meet the requirements while co-existing with the spacecraft and interferometer. At the same time, the signal structure must facilitate acquisition of the signal, expedite calibration, meet the performance requirements, and tolerate failures. Finally, the sensor must operate reliably for at least five years and scale gracefully from five to seven or more spacecraft.

We have designed a sensor to meet these requirements. To develop the new technology we are building two testbeds, one indoors and one outdoors, that will validate the various elements singly and in combination. The indoor testbed will provide maximum control for debugging, and the outdoor testbed will provide maximum fidelity to the actual operating environment for a subset of the flight sensor.

Acknowledgments

The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1]National Telecommunications and Information Administration, *Manual of Regulations & Procedures for Federal Radio Frequency Management*, U.S. Government Printing Office, Washington, D.C., May, 2003.

[2]Rosenberg, H.M., "The Behavior of Materials at Low Temperatures," *Advanced Cryogenics*, Bailey, C.A. (ed.), New York, Plenum Press, 1971, pp. 77–101.

[3]Weinreb, S., "Cryogenic Performance of Microwave Terminations, Attenuators, Absorbers, and Coaxial Cable," National Radio Astronomy Observatory Electronics Division Internal Report 223, Charlottesville, Virginia, January, 1982.