

DARPA-PS-26-137 Firefox
Frequently Asked Questions (FAQs)
as of 10/13/2026

101Q: If the intended time interval is $\sim 100\ \mu\text{s}$, would TA2 be expected to retain the previous complex speckle-field image for each ROI in the analog/optical domain to perform temporal correlation before digitization?

101A: The current approach requires that one retain the previous speckle field image (or information) to extract velocity information. If your approach does not require aligning in velocity space and can effectively work in instantaneous phase space, please provide details on your proposed technical approach.

100Q: For a given ROI, are the two consecutive complex-field measurements expected to be acquired during a single ROI dwell (with a short time interval, e.g., $\sim 1\ \mu\text{s}$), or during successive scan cycles (nominally $100\ \mu\text{s}$ apart in Phase 3)?

100A: It is expected that they will be acquired during successive scan cycles so that a proper phase difference (velocity) can be determined

99Q: TA2-1.3 specifies phase alignment within 10 ms, while Section 1.6.2 requires processing faster than tissue speckle decorrelation ($< 1\ \text{ms}$). How should these requirements be applied together in Phase 1 verification?

99A: The phase alignment requirements were softened for Phase 1 based on an understanding that it might be too challenging to meet the final goal requirements this early in the program. You will note that both phase alignment rate and number of speckles increases across the 3 phases of the program.

98Q: What phase accuracy and output bandwidth will be used to verify “phase align N speckle within” the stated time? Please define the start and end events of the alignment latency measurement, including any required sensing, estimation, feedback, and settling. What output should TA2 provide (phase, velocity, or displacement per channel)? May calibration or alignment state be retained across revisits of the same point, and is pre-training permitted?

98A: Performance will ultimately be quantified by evaluating the velocity PSD of the output of TA2 when operated in a low photon count environment. We will introduce a known target velocity with $< 1\ \text{ms}$ decorrelation time for this analysis.

97Q: Under Sections 1.4 and 1.6.2, may digital calibration and feedback control use reduced detector-monitor signals to calculate settings for optical/analog phase alignment and combining, provided the speckle data are aligned, weighted, and combined before digitization? Separately, is digital ROI selection permitted?

97A: Yes. Please clearly explain your technical approach.

96Q: Section 1.4 represents speckles as receiver pixels. For non-camera architectures, how will the required number of individual speckles be counted and verified? Please also clarify “sample” and “channel” in Table 4, including their relationship to scan points and dwell periods.

96A: The exact number of speckles can be determined using a separate camera (not needed for integration into your approach). A single “look” at the target represents a sample for a single channel (think single camera frame from current approach). If you stay at that same scan point, you could acquire multiple samples and extract an instantaneous phase measure. This process starts over if the system scans the sensing point to a new location.

95Q: Will TA2 hardware have access to the system’s reference/local-oscillator light? If so, please indicate its available power, polarization, frequency offset (f_{mod}), and the coherence-gate width and reference-path matching requirements. If not, should TA2 derive or provide its own reference?

93A: Yes, it is expected that TA2 will have access to TA1’s reference/LO light. Please specify what additional parameters/information you require.

94Q: Submission against both Technical Areas. Section 1.5 directs that a team responding to both TA1 and TA2 submit separate proposals, each addressing only one TA. We intend to do so under a single prime. Please confirm this is permissible, and clarify whether any OCI disclosure or mitigation plan is expected at the abstract stage or only in the OPP package.

94A: Yes, this is permissible. The Government will require an OCI disclosure or mitigation plan at OPP submission.

93Q: Phase 3 timing relationship between TAs. TA1-3.2 requires a full-volume scan of 100 points within 100 μs , and TA2-3.2 requires phase alignment of 100 speckle within 1 μs . Should the 1 μs TA2 alignment window be understood as coincident with the 1 μs per-point dwell implied by TA1-3.2?

93A: Yes.

92Q: TA1/TA2 interface boundary for collection efficiency (TA2-1.4, TA2-2.3, TA2-3.3). The metric is defined as “light into the TA2 hardware vs. light coming out of TA1.” At what physical plane is this ratio measured, and on which side of that boundary does mode-sorting and fiber-coupling hardware reside? Is the Government specifying a standard interface between TA1 and TA2 deliverables, or is each performer expected to propose one?

92A: Expectation is that the light coming out of TA1 will be contained within a multi-mode optical fiber. The TA2 performers are then free to re-format that light in whatever approach is best suited to integration into their technology. For instance, one could envision having to re-image the distal end of the multi-mode fiber with a specific magnification to match TA2 collection aperture.

91Q: Optical transmit power basis (TA1-1.3, TA1-2.3). Are the 10 mW (Phase 1) and 100 mW (Phase 2) average power figures specified per illumination channel, or as a total delivered to the subject across all channels simultaneously?

91A: Illumination channel

90Q: In the reference processing, is the weighted phase formed per sample and then differenced, or is the per-pixel phase change formed first and then weighted? Is the weighting by amplitude or by power (amplitude squared)?

90A: The current approach weights phase accuracy on a per sample basis based on optical power (digitized value). Other approaches are welcome if they improve system performance.

89Q: How is collection efficiency (TA2-3.3) defined and measured: optical, electrical, or end to end?

89A: Collection efficiency is quantified in terms of optical end-to-end performance (ratio of photons/optical at output to photons/power at input). Goal is to maximize this efficiency to ultimately minimize quantization and shot noise.

88Q: What range of velocities (um/s) and what bandwidth must be resolved, and over what averaging window is the metric scored?

88A: Current velocities for neural activity are estimated at <300um/sec, but the challenge is that the clutter (heart rate, respiratory rate) can sometime produce much higher velocities which result in phase wrapping. These are instantaneous velocities, so not sure what you mean by "window"

87Q: Beyond "under 1 ms," how fast do speckle amplitude and phase decorrelate over one dwell and over one revisit? Is the field correlation $g_1(\tau)$ available, and is the 1 ms figure for field or intensity?

87A: The current measurements show that the speckle intensity decorrelates in less than 1 ms, which drove the need to sample as fast as possible to minimize target decorrelation-based phase noise. The current system is limited to a 10kHz frame rate, so exact decorrelation times are difficult to determine.

86Q: Velocity needs two visits to the same point minimum (revisit up to 100 us), so we assume the earliest velocity output is about 101 us after the first visit begins. Is that the intended meaning of "under 1 us" per sample in Phase 3, and what limit applies to pipeline delay beyond it?

86A: The current system makes 4 sequential measurements to extract an accurate single-phase measurement. The system then makes four more sequential measurements to get the second phase measurement. These are then differenced to derive the velocity (so 800us for

the first velocity data point). This just beats the 1ms decorrelation time and it is expected that sampling faster would reduce phase noise due to target decorrelation error.

85Q: Will TA2 receive TA1's raster clock with a point index and dwell-start and frame-start markers? Is the scan order fixed?

85A: Yes. It is expected that all timing information will be provided as it is understood that this is required for data synchronization. Please include what information you need from TA1 just for clarification.

84Q: Is 1310 nm the working wavelength for every dataset, and what polarization state reaches TA2?

84A: Yes, 1310nm is the working wavelength for every dataset. Any polarization state is allowable with the only requirement being that the reference beam present the same polarization.

83Q: For the reference provided to TA2, what phase stability should we expect/assume over 100 us, and is it common-path with the signal? We understand TA1 is in development as well, so any information in this regard is appreciated.

83A: It is expected that there is excellent phase stability as we require this phase stability to accurately record phase shifts. The current system path-length matches the object/reference beams to help mitigate impact of any phase drift.

82Q: Is the transmit or reference beam frequency-shifted at all, and by how much? Will an RF reference for that shift be provided?

82A: The current system frequency shifts (phase-steps) both the reference and object beam to support heterodyne processing. The total shift is currently ~2.5kHz, which is driven by the maximum frame rate of the camera (10kHz)

81Q: What reference light will TA2 receive from TA1 (source, power, fiber type, polarization), and is it continuously available or just when the raster points are active?

81A: The expectation is that the reference light will be continuously available.

80Q: For TA2 the stated metrics are the devices' ability to phase align a given number of speckle within a given time. Our proposed solution can passively decompose a high order speckle mode into a given orthogonal basis to then be processed on a photonic integrated circuit. Would it be acceptable that our solution phase aligns an equivalent number of orthogonal spatial modes rather than individual speckle? (i.e. phase alignment of 100 independent modes → phase alignment of 100 speckle)

80A: The current system requires that multiple speckles must be phase aligned to achieve the desired velocity measurement accuracy due to CNR limitations. Please show in your

proposal how you can achieve the same performance using a single speckle under the realization that you will have <10 photons per speckle. This analysis should also account for the expected speckle drop-out.

79Q: In TA1, is it acceptable if the developed device requires pulsed laser illumination, or must it work under continuous-wave laser illumination?

79A: TA1 must operate under CW laser illumination.

78Q: Section 5.3 indicates that proposers should be prepared to grant the Government rights to “background IP” as part of the proposed licensing structure for TA1/TA2 technologies. Can DARPA confirm that “background IP” in this context refers only to intellectual property that is actually used in the performance of the funded project — and not to a performer’s broader intellectual property holdings?

78A: Background IPs refer to the intellectual property that will be used in performance of Firefox program that are requested to be delivered to the Government.

77Q: Will an abstract or OPP be considered non-conforming or non-responsive if it does not include a proposed Section 5.3 licensing structure, or if it proposes an alternative structure that preserves the proposer’s standard practice of commercializing results through its own technology transfer or licensing office?

77A: No. However, if the proposed effort is heavily dependent on intellectual properties yet proposer did not provide an IP structure, evaluation team may find the technical approach risky or not feasible. As such, it is recommended that all proposers provide a complete, comprehensive IP structure being offered for Firefox.

76Q: Section 5.3 directs proposers to “propose a structure via unique license” and to review the Model Other Transaction Agreement referenced in Attachment F (OPP Guidelines). Will Attachment F, or the forthcoming invitation to submit an OPP, include a template, minimum required terms, or model license language for this structure? If not, please confirm if proposers are expected to draft this licensing structure without a DARPA-provided template and where that should be included in the proposal package.

76A: Proposers may revise the Data Rights article within the Model Other Transaction Agreement to reflect the IP structure being proposed.

75Q: Is the “unique license” structure described in Section 5.3, under which the Government would retain Unlimited Rights to TA1/TA2 background and foreground IP, a required element of every abstract and/or Oral Presentation Proposal (OPP), or is it intended only for proposers who do not intend to commercialize the technology developed under the program?

75A: DARPA wants Government Purpose Rights for all data and IPs generated from Firefox Program. The request for Unlimited Rights per Section 5.3 is intended toward proposers who do not intend to commercialize or pursue the technology in a future state.

74Q: Could you clarify the coherence length of the 1310nm laser? Is this a property of the current optical system that TA1 and TA2 solutions must align with, like the other specifications in table 1? Are there any additional details about the laser source you can provide?

74A: The current system utilizes a laser from Thorlabs (SLD1018P). It is required that the properties of this laser be maintained to allow the system to maintain its sensing capabilities.

73Q: Can a TA2 solution introduce additional elements prior to the TA1 scanner to support processing? If these provide loss, can the upstream source power be increased to compensate, provided the delivered illumination remains unchanged and ANSI-compliant?

73A: Yes, that should be acceptable. Please make sure to include estimates of anticipated optical loss as additional laser power amplification can sometimes result in added optical phase noise.

72Q: Could you clarify what is the required output of the TA2 analog computation? The metrics state “phase align and combine S speckles in T time”. Should the analog computation then estimate the combined phase for the region of interest, or a temporal difference/rate as described on the second equation of page 7? Or is simply any phase-aligned and combined signal sufficient, with phase or velocity estimation performed downstream after digitization?

72A: It is expected that TA2 will perform the phase alignment of the measured speckle, which would represent a combined phase for the region of interest. Proper alignment of the optical phase will result in a significant improvement in the recorded velocity waveform, which can be quantified against the measured velocity achieved using information from a single speckle return. We currently quantify performance enhancement through calculation of the PSD for a known target vibration.

71Q: Are the existing Rx collection optics expected to remain fixed during scanning, or should TA1 include corresponding Rx steering optics to track the same ROI as the Tx beam?

71A: Performers are encouraged to explore both approaches, especially as how they relate to optical collection efficiency and bi-static transmit/receive offset angle.

70Q: How is the specified 1 sr FOV defined? Does it refer to the angular scan range entering the focusing optics (e.g., a scan lens), or the target coverage of the complete optical system as viewed from a specified reference point? If the latter, could you clarify the reference point and working distance?

70A: The realized scan angle should meet or exceed the effective area covered at depth assuming the scan reference point is located at the surface of the scalp (outermost portion of the device).

69Q: TA1 specifies that the developed device should be in "wearable form factor". Do you have any specifications for what that entails exactly (e.g., target dimensions, weight, power consumption, as well as whether it is allowed to have control and power electronics outside of the wearable device)?

69A: It is incumbent upon proposers to define and justify their own specifications for what constitutes a practical wearable device within the context of their specific technical approach. Proposers should include specifications for a form factor that is suitable for the intended use-case (e.g. tethered clinical validation and human subject testing) and explain how their design choices address relevant constraints such as those mentioned.

68Q: Does the program have any restrictions on how the requested budget should be spread across Phases 1-3?

68A: The budget should reflect the level of effort and tasks for each phase. The Government will conduct cost realism to analyze whether the proposers understand the goals and objectives of the program as reflected within the Statement of Work/Task Description Document and that the costs are realistic to the proposed tasks for each phase.

67Q: The solicitation instructions say "Embedded text in figures and tables may be no smaller than 10pt font." Attachment B (Abstract Template) instructions say "Font sizes of 8- or 10-point may be used for figures, tables, and charts." Could you please confirm the allowable font size for figure captions/tables/figures/etc?

67A: Font sizes should not be smaller than 10pt.

66Q: What are the expected minimum, nominal, and maximum optical powers per speckle at the TA2 interface, and what per-speckle CNR or phase-noise range should TA2 use for design and IV&V?

66A: The measured CNR on a per camera pixel basis is <10dB. This corresponds to a velocity noise of ~200-600um/s, which is why the current system requires phase aligning of multiple pixels (and speckle) to reduce this velocity noise to reasonable values to support accurate recording of the velocities associated with neural activity. These CNR values can be approximated by reducing the object beam power to extremely low values and recording velocities from a vibrating, diffuse mirror. For reference, operating the system in this manner with object beam powers <1nW will result in equivalent CNR/pixel values. It is expected that the TA1 and TA2 technologies will contribute negligible velocity (phase) error to the system. Based on the required velocity sensitivities, the velocity error should be <10um/s.

65Q: Would collaboration with biomedical research teams experienced in studying neuronal activity be encouraged?

65A: Knowledge or expertise in neuroscience and brain-computer interfaces is optional but not required, as Firefox is primarily a coherent optics program.

64Q: If performers are responsible for generating the speckle data, will the program provide specifications or standardized models, samples, phantoms, measurement conditions, or evaluation protocols to ensure consistency across performer teams?

64A: Yes. All required information will be provided to ensure that the performers can test their systems in the same manner that the government team will test their technologies. The IV&V team uses a straightforward experimental approach to generate speckle data that contains a small velocity signal of interest. Specifically, a function generator drives an amplified piezoelectric actuator to achieve 10's of nanometers displacements at 75-100Hz. A diffuse reflector is attached to the actuator, which when illuminated, provides a speckle field at the receiver with a known velocity signal. One can use either spectralon for a Lambertian target reflectance or silver reflective ground glass.

63Q: What types and formats of data will be provided to performers?

63A: Data will be 1-D and 2-D time series data

62Q: Will the program provide ground-truth of neuronal-motion or phase data for the development and validation of performer solutions?

62A: The IV&V team can provide baseline data from calibrated targets and instructions on how this data was collected so that the performers can collect and process data themselves.

61Q: Can a single subcontractor participate in multiple (two) proposals for the same TA?

61A: Yes.

60Q: Must the 100 Phase 3 sensing coordinates be independently programmable anywhere within the > 1.0 -sr addressable FOV, or is a predetermined raster/grid acceptable?

60A: Yes, a predetermined grid is acceptable.

59Q: A Gaussian beam with a 1.0mm $1/e^2$ diameter will produce a < 50 micron $1/e^2$ diameter spot size when focused by a 25 mm focal length lens at a 1310 nm wavelength (scattering effects are not included). The program objective is to achieve 50-micron spatial resolution. Will a 1.0mm $1/e^2$ diameter Gaussian beam deviation from the TA1-1.3 metric be accepted or be determined as non-compliant by the program?

59A: Please explain the justification for requiring this smaller beam diameter in the abstract submission. The beam diameter also impacts allowable optical power based on skin safety requirements, which will ultimately impact system SNR and CNR.

58Q: Seeking clarification of the TA1-1.3 metric definition of the “beam waist at aperture of > 2mm.” If a proposer is using a Gaussian beam ($M^2 \sim 1$) will the > 2mm beam waist correspond to the $1/e^2$ diameter?

58A: Yes – ensure that your design parameters are clearly stated in the abstract submission.

57Q: We see that page 25 of the solicitation (section 5.5.2. Awards - Fundamental Research and FRRBS) says that Biographical Sketches and Current and Pending Support must be submitted by “all submissions proposing to this solicitation.” Could you please further clarify if these 2 documents are required at the Abstracts phase, the OPP phase, or both?

57A: These documents are required at the OPP phase. Only Attachments A and B are needed for the abstract phase.

56Q: Given a ROM budget is required for the Abstract, could you share an anticipated project start date for the 36-month program?

56A: DARPA anticipates program kickoff will occur within the first 3 months of 2027.

55Q: Does the proposed solution need to involve strictly no calibration or fine-tuning for individual users, or can a brief phase of calibration to each new user be allowed? (e. g., showing the user a set of videos and using recordings while they view these videos to inform a set of physical and post-processing parameters)

55A: A brief calibration is allowed. Please explain how this will work and impact final system performance in the abstract submission.

54Q: Is a solution that provides, by optical coherent processing, the field components of the speckle field signal that are correlated with visual stimuli considered responsive, or do responsive solutions strictly need to be about aligning the phase of these different speckles?

54A: The field components that are correlated with the visual stimuli and found in each of the speckle is the goal of the program. Processing each of these speckles independently has been shown to limit system sensitivity due to shot and quantization noise. If you have a different approach, please explain with appropriate details how you will achieve high system performance.

53Q: The Phase 1 hand-off reads "Deliver software models/processing software to IV&V" (PS Table 6, MS5, Month 11, p. 14), while the Phase 2 and 3 hand-offs read "software

models/processing hardware" (MS9, Month 22, p. 14; MS14, Month 34, p. 15). Is "processing software" at Month 11 intended, i.e., is no processing hardware delivered to IV&V in Phase 1?

53A: The primary objective of structuring this hand-off around software, simulation, and modeling rather than hardware is risk mitigation under fixed program timelines. Hardware implementations that may involve custom microelectronics involve extended fabrication schedules that may exceed the operational calendar of Phase 1. Establishing rigorous, validated models and simulation in Phase 1 ensures that performers will have sufficient time for successful fabrication of hardware components in subsequent phases.

52Q: The receiver images both the shallow gate and the deep gate onto the same FPA, but the example speckle image and phase calculation appear to use only the shallow gate. How is the deep-gate signal used in the measurement or the phase calculation?

52A: Speckle images are both from the same depth (surface of the cortex).

51Q: The PS and slides treat each pixel as one speckle, i.e. one phase sample (PS Sec. 1.4, p. 6; slide 18). In the existing setup, how large a patch of pixels shares a uniform phase, and does it roughly correspond to one speckle grain? Could you also share speckle images from consecutive frames, so we can estimate the decorrelation time separately from the heterodyne modulation? We have assumed the tissue decorrelation time < 1 ms, PS Sec. 1.6.2, p. 10) is longer than the heterodyne modulation period, but could not confirm this.

51A: The system currently operates with $\sim 3 \times 3$ pixels/speckle and we could reduce that to be closer to 4 pixels/speckle. Yes, the decorrelation time of tissue is < 1 ms and the current system operates at a heterodyne frequency of ~ 2.5 kHz. This is probably marginal, but we were limited by the 10 kHz frame rate of the camera and would like to have the ability to operate at even higher heterodyne frequencies.

50Q: There are a few places where values are different within the PS. Can you confirm which values proposers should use for the following? TA1 Phase 1 point count – should this be 10 or 20 independent points? TA1 Phase 2 scanning angle and range – 50 points within 5 ms over 0.75 sr, or full-volume scan within 10 ms over > 0.5 sr? Receiver fiber core size – should this be the 62 micron core MMF or 50 micron MMF?

50A: Point count for Phase 1: 10 points. Scan angle and range for Phase 2: full-volume scan within 10 ms over > 0.5 sr. Receiver fiber core size: 62 micron core MMF.

49Q: For signal acquisition, do we need to use four-step phase shifting or an off-axis holographic configuration? If four-step phase shifting is required, what heterodyne frequency would you recommend? (Maybe around 10 MHz or higher to meet TA2-3.2)

49A: 4-phase step phase shifting or some other form of heterodyne approach is desired as it provides the ability to use CNR weighting at the receiver since that CNR requires that the signal and local oscillator can coherently mix. We require a minimum of 3 kHz to account for

decorrelation time of the tissue. There is no maximum value and the exact value can be driven by scan parameters (i.e. do you stay at one scan point for the full 3kHz period or do you can across all of the points once at one phase value and then phase shift by $\pi/2$ and scan them all again). There are pros and cons to both approaches in terms of system complexity and SNR.

48Q: (TA2-1.4) Optical collection efficiency is defined as “light into the TA2 hardware vs. light coming out of TA1”. Does this include the collection efficiency of the receiver MMF, or does it refer only to the coupling efficiency from the MMF into the TA2 hardware?

48A: Optical efficiency is defined as full system transmission efficiency, so the ratio of the light coming out of the single mode fiber illumination coupler to the light exiting the MMF. This is obviously also dependent on the target reflectivity, so could assume a 10% Lambertian target for simplicity. This analysis will allow us to do a relative comparison between different approaches.

47Q: (TA2-1.2) How will the "24-bit detection-equivalent bit depth" be defined and tested? Does this correspond to an after-photodetection CNR of approximately $6.02 \times 24 + 1.76 \approx 146.2\text{dB}$?

47A: 24-bit detection-equivalent bit depth is simply the bit depth of the digitizer utilized in the system. This was expected to be an easy spec to meet, but is significantly better than what could be achieved if using a FPA for detection and processing.

46Q: Table 1 lists a 62 μm , NA 0.275 receiver fiber, while TA1-2.4 specifies a 50 μm core. What NA, index profile (graded or step) and tether length should we assume? A standard 50 μm , NA 0.2 graded-index fiber carries only about 70 modes per polarization at 1310 nm, which is below the 100-speckle Phase 3 target.

46A: The system utilized the 62 μm , NA0.275 receiver fiber, so please design based on that.

45Q: TA2-1.1) Is a “speckle” an independent speckle grain, with the number of speckles related to the number of supported modes in a given multimode fiber (MMF)?

45A: Yes, a speckle is an independent speckle grain. The number of speckles is driven by both the number of speckles at the input to the MMF as well as propagation through the MMF itself.

44Q: Regarding Technical Area 1, would an architecture deploying an array of approximately 100 such fiber-triplet units—rapidly switched/multiplexed on-chip to achieve the required scan rate and coordinate count—be considered fully responsive to the TA1 scanning and field-of-view metrics?

44A: Yes, this approach is considered responsive.

43Q: Do we need to consider or implement a method to synchronize the TA1 raster scan with the TA2 processing?

43A: Timing information must be made available from TA1 so that TA2 performers can synchronize data collection

42Q: On page 12 of the Proposers Day slides, there are two MMFs and one SMF. Do we need two MMFs? If they are used to increase the light-collection efficiency, could we use an MMF bundle instead?

42A: Yes, the two MMFs provided increased collection efficiency. MMF bundles are an option and please provide estimate of relative collection efficiency. Ideally would like to achieve increased collection efficiency over current approach.

41Q: (TA1-1.5) Do TA1 deliverables include the collection optics and the coupling into the receive fiber?

41A: Yes.

40Q: Table 2 lists 10 points for Phase 1, but Table 3 TA1-1.7 lists 20 points. Which applies? Must points be randomly addressable, or is a fixed raster acceptable?

40A: 10 points is the correct value for Table 3. A fixed raster is acceptable.

39Q: Would illuminating several points simultaneously (with longer averaging times and requiring multiple receiver MMFs or MMF bundle), instead of raster-scanning a single focused beam, be within the scope?

39A: Illumination of several points is allowed; however, details must be provided as to how you will distinguish between the return photons from the different points. The medium is highly scattering, so it is not expected that simple geometric filtering will be adequate.

38Q: Must the beam stay stationary on each point during its dwell, or is continuous scanning acceptable? (I guess it should stay stationary for the TA2 processing) If the beam should stay stationary on each point during its dwell, is there a requirement for the dwell time?

38A: Beam should stay stationary at each point with the goal of maximizing dwell time. Please provide estimates of dwell time.

37Q: (TA1-1.6) How will "pass-through of 1310 nm without affecting coherence" be measured? Is there a limit on the phase noise or jitter the steering optics may add in the neural signal band (roughly 10 Hz-10 kHz)?

37A: The plan was to measure the intrinsic coherence length of the laser (or linewidth) prior to passing through TA1 architecture and then again after detection/collection. Phase noise in the 10Hz-10kHz band is also an important consideration, so please quantify how you plan to

quantify performance. It is expected that the approach developed for TA1 will not appreciably change the phase noise characteristics (for example, quantified with a PSD).

36Q: Since the source coherence length is about 50 μm and sets the depth resolution, do we need to keep the signal and reference path lengths matched as the beam scans over the 1-steradian field of view? If not compensated, the constant-delay image plane would become curved. Is imaging on this curved surface acceptable, or do we need dynamic path-length compensation?

36A: You do not need to provide dynamic path-length compensation and it is understood that this will result in some sensing curvature

35Q: What is the nominal and allowable range of the heterodyne modulation frequency f_{mod} at the TA2 input?

35A: $F_{\text{mod}} > 3 \text{ kHz}$, which is the minimum to account for decorrelation time of the tissue. There is no maximum value and the exact value can be driven by scan parameters.

----------

34Q: Is there a position repeatability requirement for the scanned points from full volume scan to full volume scan?

34A: Repeatability should be better than 10 mrad (1/10 of the acquisition spot separation).

33Q: Are the scanned points required to be on a rectangular grid or can they be distributed in a pattern that mostly fills the scan volume uniformly?

33A: They can be distributed in a pattern that mostly fills the scan volume uniformly.

32Q: Does the TA1 system require a component that physically scans $> 1.0 \text{ sr}$ or does it need to scan an area formed by a $> 1.0 \text{ sr}$ FOV with a 15 mm to 25 mm focal length lens?

32A: The system needs to physically scan $> 1.0 \text{ sr}$ with a diffraction limited spot (in air) at depth as defined by a 2mm beam waist and a 25mm focal length lens.

31Q: Are the TA1 prototype scalable components, or system of components, expected to work with low temporal-coherence lasers or light sources? Or can they be designed to work only with narrow linewidth lasers

31A: No, the system must work with short coherence length lasers so that it can range gate and achieve sub-wavelength ($\lambda/1000$) displacement sensitivity for recording of neural activity.

30Q: Is it possible to include the development of a light source (functionally identical to the one used in the single-channel proof-of-concept by APL) with emission wavelength centered at 1310 nm in the abstract submission if this is important to the proposed TA1 or TA2 solution?

30A: Yes. The source must be made available to both TA1 and TA2 as it would be required for coherent processing. The current system operates with a center wavelength of 1310nm and a coherence length that can be tuned from 50microns to 3mm. At the same time, justification for this component development should be included in the abstract.

29Q: How is the “24-bit detection-equivalent depth” defined and measured? Does it refer to SNR or dynamic range of the recovered phase/displacement, an equivalent ADC resolution, or another system-level metric?

29A: This is defined as the bit depth of the optical detector digitizer. The current system uses a camera to digitize the optical data and the bit depth of the camera is 14-bit. Increasing bit depth will ensure that the system operates in the shot noise limit and not the quantization noise limit.

28Q: Is the sensing illumination CW, pulsed, or otherwise temporally modulated? For the current single-channel system at the stated 10 mW/100 mW transmit-power levels, what is the typical optical power at the output of the receive MMF, both in total and approximately per speckle or supported fiber mode?

28A: The system uses CW illumination that is frequency modulated to support heterodyne operation. Heterodyne modulation currently occurs at ~2.5kHz. The current system operates can successfully measure target velocities with <10 signal photons per pixel in the camera. This is specified as signal photons as this is a coherent optical receiver and therefore utilizes a local oscillator or reference beam to provide optical gain and ensure that the system operates in the shot-noise limit.

27Q: The solicitation describes a weighted CNR calculation. Is this weighting derived under the assumption that the noise level is the same across all channels, or is it based on a different noise model or statistical assumption?

27A: The current weighting approach is based on a real-time measurement of CNR for each pixel and weighting is then applied across all of the pixels that normalizes these weights such that velocity calibration is maintained. There is no assumption that the noise is constant across all channels (pixels).

26Q: Are TA2 performers expected to follow the signal processing or reconstruction algorithm described in the solicitation, or are alternative approaches permitted if the required outputs and performance metrics are achieved? For example, would a neural-network-based approach for velocity estimation or prediction be considered compliant within the specified TA2 hardware and processing constraints?

26A: Yes, performers are encouraged to explore novel approaches that can provide enhanced performance over what has currently been implemented. The effectiveness of their approach will be baselined against the identified reconstruction/processing algorithms.

25Q: What are the boundaries on the use of digital electronics in TA2? For example, may and FPGA or other digital processor be used for adaptive control, calibration, or low-rate monitoring, provided that the high-bandwidth sensing signals are not digitized or digitally processed?

25A: There are no boundaries on the use of digital electronics for TA2, provided that optical signals are not digitally processed.

24Q: How will TA2 processing results be evaluated during each phase? In particular, what will serve as the reference or ground truth for evaluation the recovered neural signal or predicted velocity, and which quantitative metrics will be used to determine performance? Will a common dataset, simulation model, or experimental measurement be used for evaluation?

24A: The best way to test the performance of TA2 will be to quantify the performance using experimentally derived data. The IV&V team currently simulates neural activity by using a PZT driven with 10's of nanometers displacements at ~100Hz to simulate the expected neural activity displacements and frequencies. To account for signal degradation, the IV&V team drops the object beam power to extremely low values. Finally, target decorrelation can be re-created by spinning the vibrating target. The exact parameters for approximating neural activity for TA2 will be provided by DARPA so that the performers can test at their own facilities.

23Q: What is the expected input to the TA2 processing architecture? Should TA2 assume a camera-acquired intensity/speckle image or may the architecture directly process an optical field or holographic signal that retains phase information?

23A: It is expected that TA2 will directly process an optical field or holographic signal. The camera-acquired approach operates as a baseline for expected performance and does not need to be implemented.

22Q: In clarifying the definition of the phase-alignment time, must the system establish phase alignment for a previously unknown speckle field within 1 microsecond, including measurement, calculation and control? Or does the metric refer to the update rate or latency of the optical processor after initialization or calibration?

22A: The decorrelation time of tissue should drive the phase alignment time requirement, which is expected to be < 1ms. Therefore, the phase alignment rate should be 1-10kHz.

21Q: In scattering tissue, some object light should contribute only background. Should TA2 solutions explicitly account for this non-interfering background when estimating CNR and weighting speckles?

21A: In the proposed data collection approach, a carrier should only be present with constructive interference between the object and reference beam. This is why we weight on CNR and not simply intensity on a pixel, because a pixel's intensity will contain both background (noninterfering) and carrier components.

20Q: In the detector-power equation in the program solicitation, does P_O represent the total detected object power or only the mutually coherent, coherence-gated portion?

20A: Total detected object power.

19Q: The solicitation requires data compression, phase alignment, and speckle processing to occur in the analog or optical domain before digitization. Would a brief initial or periodic digitized calibration measurement be permitted solely to determine the phase-alignment settings, provided that the continuous neural signal is subsequently phase-aligned, combined, and compressed in the analog or optical domain before final digitization?

19A: Some initial and periodic digitized calibration is allowable. For clarification, it is expected that the speckle will decorrelate in under a millisecond; thus, phase alignment will need to occur at > 1 kHz rate during data acquisition.

18Q: Do TA2 performers need to develop and integrate the photodetector array with the analog processor for an in-sensor approach?

18A: No. The photodetector array was provided as a baseline design and is what performer solutions will be tested against.

17Q: There is a requirement on the solid angle, but do solutions that achieve the same effective interrogation area on the brain achieve the desired goal?

17A: Yes.

16Q: Do the power and $M^2 < 2$ requirements need to be maintained at “every scan position” across the full FOV, including the edge?

16A: It is expected that some degradation will occur based on scan angle. The goal is to come as close to this value as possible across all scan angles. Proposers should provide details on predicted degradations.

15Q: Does “full volume scan rate” include beam transition, settling, and acquisition time at every coordinate, or only the response time of the steering device?

15A: This should include beam transition, settling, and acquisition time.

14Q: For the 10 mW Phase 1 and 100 mW Phase 2 requirements, where is the optical power measured? Is it at the output of the TA1 beam steering system, or at the tissue surface?

14A: This optical power should be measured at where the tissue surface would be located. These optical powers are based on skin safety limits.

13Q: Can Phase 1 leverage COTS components and transition to a custom, integrated solution as the program progresses to meet final Phase 3 metrics?

13A: Yes, proposers must clearly lay out the technical path that will enable transition to custom design that meets all performance metrics.

12Q: In the current optical system, can proposers remove the transit SMF fiber collimator and propose a custom replacement?

12A: Yes, the expectation is that relative performance should be maintained.

11Q: Does the beam steering system need to integrate with the current head-mounted prototype without modifying any components of the current prototype? In other words, does the beam-steering system proposed need to fit between the current transmit SMF fiber collimator and lens? Or can the input fiber be modified (i.e. to enable fiber packaging to the transmitter) or can the current lens be replaced with a different lens/focusing system?

11A: It is expected that the performer will modify the existing design to incorporate their scanning technology. The provided optical design is provided as design guidance and reference.

10Q: Is the 2 mm beam waist defined for the normal-incidence on-axis situation only, or is this required for all steering angles?

10A: This beam waist should ideally be for all steering angles. Proposers should provide estimated deviation from this beam diameter if it is scan angle dependent.

9Q: What assumptions can be made by TA2 about TA1 scanning? Can it be assumed that TA2 will receive optical input?

9A: Yes, TA2 should assume that it will receive optical input. This input can be simulated using simple optical hardware for proper testing of TA2 performance. Specifics of the optical hardware will be provided to performers once the program starts.

8Q: For TA1, are there any requirements on the receiver?

8A: The receiver must support use of a multi-mode optical fiber to transmit collected photons to TA2 technologies. Estimates for collection efficiency are expected to be provided in the proposal, as high-collection efficiency is required for successful neural recording.

7Q: Are international collaborators allowed to participate in the program?

7A: Yes.

6Q: For scanning across 100 points, do program goals require the optical beam to scan continuously between individual points, or just the 100 discrete points at specific locations across the desired area?

6A: To meet program goals, scanning at discrete points at specific locations is ideal.

5Q: The PS uses optical and analog processing. To be clear, are approaches that propose all-analog processing considered responsive (as opposed to all-optical)?

5A: Both analog and optical (or some combination) approaches are considered responsive.

4Q: Is the system tested against human brain activity specifically? Or should the system meet the metrics in general regardless of human or animal brain activity or scattering media?

4A: The system should meet the metrics in general. These metrics have been identified as representative of what is required to successfully record human brain activity.

3Q: What data will be made available to performers over the course of the program?

3A: Information will be provided on experimental configurations for testing TA1 and TA2. In addition, specific data will be made available from the IV&V team upon request from performer teams.

2Q: The program solicitation mentions that “multiple awards are anticipated between \$4,000,000 and \$7,000,000 across all Phases 1 – 3 of the program...” Is the \$4M - \$7M range for a single awarded project, or total for all Firefox projects?

2A: This range is for a single award.

1Q: Are non-U.S. universities eligible to participate as a prime or sub-performer for Firefox? If non-U.S. participation is permitted, are there particular contracting or registration requirements we should address before submitting an abstract?

1A: Yes, non-U.S. universities are eligible to participate as prime or sub-performers for the Firefox program.