

CORONET Proposers' Day
August 8, 2006
Questions and Answers

All questions that were submitted during the CORONET Proposers' Day meeting are answered below. Some questions may have been slightly edited for clarity. If there are multiple questions regarding the same issue, the questions are listed together under a single question, along with a single answer.

If any statements below conflict with the BAA or PIP, the BAA and PIP take precedence.

Services

1. *Q: What percentage of the aggregate network demand is made up of multicasting traffic?*
A: For the demand scenarios of Table 3, there is no multicast traffic (see page 11 of the PIP). However, support for multicast traffic must be included as part of the proposed solution. A discussion of how multicast is supported is one of the technical deliveries for Phase 1 (see page 30 of the PIP).
2. *Q: For 1:N multicasting, how large is N expected to be?*
A: N should be on the order of 5 to 10, but higher values of N also can be considered.
3. *Q: What are the data rates associated with individual multicasting streams? Does the user care whether the multicasting is IP-based or λ -based?*
A: Support for multicast should cover the entire range of data rates specified in Table 2. The proposed solution should provide support for multicast of both IP and wavelength services (see page 30 of the PIP).
4. *Q: What defines a "stream" or "connection"? End-to-end, WDM path, etc. All one wavelength or multiple/changing?*
A: A 'connection' or a 'flow' is defined end-to-end. It may be routed over multiple lightpaths, where, in general, the lightpaths can have different wavelengths.
5. *Q: Is there anything in common within a stream? source IP, transport connection, ingress subnet, or source add/drop mux?*
A. (We assume this question was regarding the use of the term 'flow'.) Assume IP flows have a common source IP address and port and destination IP address and port.
6. *Q: Will SLAs established in advance for on-demand services place any restriction on service instances (such as ingress and egress points)? It is likely that real SLAs will include such restrictions.*
A: No. For the purposes of Phase 1 of this BAA, QoS metrics for guaranteed bandwidth IP services are on a per-demand basis, and for best-effort IP services are on a long-term average basis. (See Table 7 of the PIP.) In Phase 2, the suggested 'restriction' may be considered if deemed necessary.

7. *Q: Are the traffic distributions of Table 2 based on the raw end-to-end demands, or are they based on the amount of capacity allocated to them. For example, if a 40 Gb/s wavelength is inefficiently packed and carries just one 10 Gb/s guaranteed bandwidth demand, does this count as 10 Gb/s or 40 Gb/s?*

A: The traffic distributions are based on the raw end-to-end demands. In the example given, this would count as 10 Gb/s.

8. *Q: The traffic demand distribution described on pages 8 (bottom of chart) and 29: can we treat that as demand after aggregation and grooming?*

A: This should be the demand distribution prior to aggregation and grooming; i.e., it is the raw end-to-end demands.

9. *Q: Aggregate network demands are symmetric and bi-directional. How should the aggregate demand be counted?*

For example, consider the following bi-directional traffic matrix:

	<i>A</i>	<i>B</i>	<i>C</i>
<i>A</i>		<i>10</i>	<i>10</i>
<i>B</i>	<i>10</i>		<i>10</i>
<i>C</i>	<i>10</i>	<i>10</i>	

Is this 30 units of aggregate demand? Or, should it be counted as 60 units.

A: In the example shown, this counts as 30 units of aggregate demand, not 60.

10. *Q: Does the best-effort bandwidth shown in Table 2 (40% of S) represent the aggregate peak bandwidth to be allocated to this service type, or just the average bandwidth? We assume the former based on the fact that average bandwidths for this service type are difficult to pin down in practice.*

A: The levels of best-effort IP services specified in Table 2 represent long-term averages (see page 11 in the PIP). Reasonable assumptions should be made by the proposers regarding the burstiness of this type of traffic, and these assumptions should be explicitly stated.

11. *Q: Can we assume that in multiple-wavelength demands that the λ 's can be routed separately?*

Must bundled wavelengths traverse the same links? The same fibers? What are the timing requirements on the signals' arrivals (how synchronized)?

A: For services comprised of multiple wavelengths, each wavelength of the service must be routed over the same path. However, it is not necessary that the wavelengths carrying the service be contiguous. Furthermore, if the links of the path have multiple fiber pairs, it is not necessary that the wavelengths carrying the service be on the same fiber pair. No additional timing requirements are specified.

12. *Q: QoS metrics for IP services: Dr. Saleh mentioned these metrics were not averages but rather limits. What statistical probability is to be associated with these limits? (e.g., 2σ , 3σ , 4σ , 5σ ...)*

A: For the Guaranteed-bandwidth IP services, the QoS metrics must be met for each individual flow. For the best-effort IP services, the QoS metrics represent long-term averages. (See page 27 of the PIP.) In addition, as was stated during the Proposers' Day meeting, the QoS metrics of Table 7 only apply under the no-failure condition.

13. *Q: Elaborate on how DARPA wants best-effort IP services, which account for 40%S and are bursty, to be supported?*

A: The best-effort services can be considered as 'always-present, highly variable traffic'. For example, they are not included in the service setup table (Table 5). Research on how best to carry this traffic is part of the CORONET program, including possible edge-aggregation techniques.

Restoration

14. *Q: Table 5: Is the time given for "per connection" or for "all affected demands" to be restored to working state?*

A: All affected demands must be restored to the working state within the times specified.

15. *Q: Given restoration time scales we are discussing, (100-250 msec), why are we concerned with multiple independent failures? It seems that the probability of more than one catastrophic failure within this timescale is essentially zero.*

A: While a demand may be restored within 100 to 250 msec using protection capacity, the failure itself (e.g., a cable cut) may not be repaired for several hours. During the time it takes to physically repair the actual failure (e.g., repair the cable), a second, or even a third failure may occur.

16. *Q: Services may be classified as required to be dependable or non-critical?*

A: None of the services should be considered pre-emptible (or 'non-critical'). As stated on page 16 of the PIP, if a failure occurs that does not bring down a particular demand, then that demand should not suffer a hit due to that failure.

17. *Q: Restoration can be partial or has to be for full service (in other words, is graceful degradation allowed)?*

A: Restoration has to be for the full service. Graceful degradation is not a part of the restoration strategy considered in this BAA; the full service must be restored.

18. *Q: For $1\lambda/2\lambda/4\lambda/8\lambda$ services, if one of the group fails, failure of the whole group?*

A: Yes, if one wavelength of the multiwavelength service fails, the whole demand is considered as failed.

19. *Q: Consider protection times:*

Dedicated : < 50ms; Shared : < 100ms or 250ms

But, a flow can be 1Tb/s (max), so within 50ms up to 6 GB can be lost

→ Is there a need for smaller times?

→ Is there a need for better protection techniques, or can “shared” schemes be acceptable?

A: The requirements for protection times are as specified in Table 6. For purposes of this program, there are no further requirements on restoration time, other than what is specified in this table. The goal is to have restoration schemes that are both very fast and capacity-efficient. While there is no requirement that the restoration be shared, it is unlikely that the capacity requirements can be met through the use of pure dedicated protection.

20. *Q: Is P% protected traffic assumed to be fully restored after 3 failures?*

A: P% of the traffic needs to be arbitrarily selected to be restorable from three failures. In order for a demand to be designated as ‘three-restorable’, it must be possible to find four disjoint paths between the endpoints of the demand. For any demand that has been designated as ‘three-restorable’, there must be enough protection capacity such that the demand can be restored from all combinations of three failures (where at most one of the failures is a node failure).

21. *Q: For purposes of meeting the metrics of Table 6, do we need to consider the case where an IP router fails, but the optical elements at the node are up?*

A: No. Table 6 addresses only two failure types: link failures and total node failures. Node failure means that all equipment at that node (optical network elements, IP routers, and other equipment) fails (see page 24 of the PIP). While proposed solutions must be capable of recovering from other types of failures (e.g., transmitter failure), only link and total node failures are covered by the metrics.

22. *Q: Since we must consider in our analysis non-catastrophic failures, what are the restoration standards for these non-catastrophic failures? (re:Table 6)*

A: There are no metrics for the non-catastrophic failures (i.e., failures other than link or node failures). However, proposed solutions should support recovery from these failures. The delivery for Phase 1 should include a discussion on the mechanisms used to recover from these other failure types, and how the recovery mechanisms perform.

23. *Q: Are all degree-N nodes (N=2, 3, 4) to be treated equally for restoration? Or are some more equal than others?*

A: With regard to node failure, any degree node can fail equally likely. When Table 6 refers to a node failure, it is any node failure, regardless of the nodal degree. With regard to restoration schemes, it is possible to have certain nodes play a more major role in restoration; this is up to the proposers to determine as they design their restoration schemes.

24. *Q: Since we are modeling IP/WDM architecture, we need to count the traffic carefully. Typically in commercial networks, IP links run at 50% to 80% utilization over a peak 5-minute interval. Therefore, “overbuild” for restoration does occur to some natural extent. For wavelength services, this is not an issue. Therefore, when the PIP refers to the “spare capacity ratio” (overbuild) I assume it counts the 40Gig (or larger) pipes of link capacity required to meet the IP restoration objective and not traffic. Please clarify.*
A: The spare capacity ratio is defined based on wavelength-km measurements. (See page 24 of the PIP.) Thus, this ratio is counting the ‘pipes of link capacity’, not the actual traffic that is carried in the pipe.
25. *Q: Should the network topology and demands be constructed so that fault isolation is not required to provide restoration for triple-failure services?*
A: The topology and demands should be realistic. Do not specifically construct the topology or demands so that restoration is more easily achieved. Research on fault isolation is in-scope for the CORONET program.
26. *Q: Unless I missed a more precise definition in the BAA, I believe the definition of B/W on page 24 needs to be corrected. As it stands, a design that routes working traffic on longer-than-necessary paths can result in much lower B/W ratio than if all working traffic were routed shortest path. As IP traffic is customarily routed shortest path (to achieve least latency), it is important to have designs that are faithful to this practice. (Several people asked a similar question during Proposers’ day)*
A: Clearly, performers should not route working traffic over circuitous paths for the purpose of decreasing the spare capacity ratio. Note that performers will be providing plots of W (i.e., the working capacity), as well as their demand set, as part of the proposal. These will be checked to a first order for consistency. We do not want to place explicit constraints on routing because this could limit opportunities for load balancing or grooming.

Service Setup

27. *Q: In BAA PIP, Table 5, regarding very fast service setup, is there any burst assumption on back-to-back service setups?*
A: The service-request arrival process is assumed to be independent.
28. *Q: What is the process relationship between the holding times and the set-up times?*
A: See Table 5 of the PIP for the holding times and setup times for each of the classes of service setup.
29. *Q: Slide 17 and table 5, for each service category, the holding times can vary by a factor of 60 or more. The actual distribution within this range will affect bandwidth requirements. Please specify a distribution (I suggest a uniform distribution).*
A: Proposers should select a reasonable distribution that covers the range of holding times specified in Table 5 of the PIP (e.g., uniform, truncated exponential, bimodal distribution). Proposers should ensure that their solution is robust to other holding time distributions as well.

30. *Q: For scheduled services, when does the “clock start” on the 10 second provisioning time?*

A: The clock starts when the service initiation process begins (even though the service setup request may be known much earlier). Thus, the 10 seconds includes, for example, checking the state of the network, computing routes, and establishing connections.

31. *Q: If very fast set-up traffic is impacted by a network failure, can we assume a change in the assumed blocking probability? (e.g., 10^{-3} blocking becomes 10^{-2} or 10^{-1})*

A: The blocking probability metrics in Table 5 pertain to the case of non-failure conditions. It is expected that the performance will be worse under failure conditions.

32. *Q: For very fast services, can we assume a new set up will take care of the protection requirement?*

A: Very fast services need to be protected against just a single failure. It is possible, though not required, that protection for these services be provided by issuing another service setup request (see page 24 of the PIP).

33. *Q: Usually, provisioning times objectives are stated in tables, e.g. “ n_i connections can be made not to exceed t_i seconds”, for various values of n_i . I assume your service set-up is $n_i=1$?*

A: Yes, the metrics of Table 5 pertain to the establishment of a single connection ($n_i=1$).

34. *Q: Do you expect to use the existing GMPLS control plane or innovations to support very fast set-up, etc.?*

A: It is up to the performers to select the mechanism to use to support very fast set-up. Whatever mechanism is selected must be implementable in Phase 2 of the program.

Security

35. *Q: The BAA is very shallow on security requirements and there are no evaluation criteria focused on security features. Why is this?*

A: While there are no metrics pertaining to it, security is an important part of the CORONET program. As stated on page 20 of the PIP, effects of security systems and measures should be taken into account in architecting the system as well as in assessing the performance of the system. Also, software hooks for security systems should be maintained in the code to be developed in Phase 2.

36. *Q: On slide #5, last bullet – Hooks for security..... What does security hooks refer to? What types of security? (Encryption, authorization, etc.)*

A: Proposers need to be aware of the various security protocols that are already established and that are currently being developed. While the security protocols do not have to be implemented in Phase 2, the system should include the necessary hooks to be compatible with such protocols. Security areas include, but are not necessarily limited to: encryption, confidentiality, authentication and authorization techniques, intrusion detection, protection against unauthorized attempts for traffic pattern detection, and protection against denial of service attacks. (See page 20 of the PIP.)

37. *Q: Does Phase I include a secure control plane? Is this included in the 100ms CONUS setup time?*

A: Phase 1 simulations should include overhead due to security measures. This overhead time should be included in the set up time when meeting the metrics of Table 5.

Hardware

38. *Q: Regarding optical hardware technology, if a different node architecture is used, what kind of metrics needs to be provided (e.g., cost saving)?*

A: It is up to the performers to propose appropriate metrics that adequately demonstrate that their node architecture is preferable to the baseline node architecture of the BAA. Estimated performance with respect to cost, size, and power consumption are likely to be some of the important measures.

39. *Q: Can we buy and use COTS Optical or Edge Hardware for fidelity studies or Phase 2 prototyping?*

A: No.

40. *Q: Can we assume all-optical regeneration provides wavelength conversion?*

A: Yes, all regenerators (OEO or OOO) can be assumed to provide full any-to-any wavelength conversion. (See page 13 of the PIP.)

41. *Q: What is the tunability speed for the Tx (ms, μ s, ns)?*

A: Reasonable assumptions should be made regarding tunability speeds. Any assumptions regarding tunability speeds, if relevant for the proposed solution, should be explicitly stated.

42. *Q: Is switch configuration time included in set-up time?*

A: Yes, switch configuration times are included in service set-up times (and restoration times).

43. *Q: Are network element response times excluded from the restoration time calculations? Deducing these or predicting them based on new designs would be substantial exercises that appear to be out of scope.*

A: Network element response times should be included in the restoration (and service setup) time calculations. Reasonable assumptions should be made regarding the switching times based on the switching architecture that has been assumed. Assumptions regarding switching times should be explicitly stated.

44. *Q: What is the assumption on speed of cross-connecting wavelengths? MEMS-based ROADM switches can switch in a few milliseconds. Can we assume a certain node switching speed?*

A: Reasonable assumptions should be made regarding the switching times based on the switching architecture that has been assumed. Assumptions regarding switching times should be explicitly stated.

Physical Layer

45. *Q: Are you placing any restriction on assumed fiber type or variation of fiber types?*

A: The assumptions for optical reach, number of supported wavelengths, and amount of supported capacity per fiber have been specified as part of the program (see Tables 3 and 4 of the PIP). Proposers should assume that the fiber types that will be used in actual deployments will meet these specifications. Research on the capabilities of various fiber types is not in-scope for this program.

46. *Q: There are a number of assumptions at the physical layer. Don't we need to do simulations at the physical layer and consider cross layer issues? For instance, SNR and dispersion.*

A: The assumptions for optical reach, number of supported wavelengths, and amount of supported capacity per fiber have been specified as part of the program (see Tables 3 and 4 of the PIP). It is not necessary to do simulations to derive the physical layer performance; this is not in-scope.

47. *Q: With bit rates increasing at a rate of 4X from 2.5Gbps to 10Gbps to 40 Gbps, why did you specify 100Gbps for scenarios 3 & 4?*

A: Due to the many challenges of implementing 160 Gb/s, the assumption is that 100 Gb/s will be the next likely transport rate after 40 Gb/s. Implementing 100 Gb/s Ethernet transport is currently being considered in the industry.

Phase 2 and Commercialization Issues

48. *Q: Can you clarify how much specificity you would like to see in the technology transfer/commercialization plans?*

A: For purposes of the Phase 1 proposal, the plans for technology transfer and commercialization should be presented at a high level. However, there should be enough detail to indicate that the proposers understand the technical and programmatic challenges that must be undertaken.

49. *Q: Can you specify level of maturity for NC&M software (e.g., Technology Readiness Level (TRL)) needed for Phase II?*

A: The software must be developed to the point that a commercial carrier will be confident of the concepts' utility in a real network, and confident of the developers' capability to translate the prototype into deployable software. However, this point will be revisited during the execution of Phase 1, when more details regarding Phase 2 are provided.

50. *Q: PIP, page 20: "credible commercialization plan needs to be included..."; to what extent should the plan go beyond technology transfer to include creating the services enabled by the technologies and the market demand for those services?*

A: Proposers can propose whatever they think is necessary as part of their high-level proposal for Phase 2. However, this point will be revisited during the execution of Phase 1, when more details regarding Phase 2 are provided.

51. *Q: Can options in Phase 2 be assumed to support hardware development if required to transition?*

A: No, hardware development is not in-scope for either phase of this program.

52. *Q: At the end of CORONET (last phase) what is envisioned as the embodiment of the “system” or “network” developed to test and validate functional and performance claims:*

existing contractor network?

existing government network (including leased)?

new network (not currently developed)?

A: Given that the ultimate goal of the program is transition to commercial carriers, validation testing during Phase 2 could be done on an existing commercial or government network, an experimental network, a network emulation testbed of sufficient fidelity, and/or any other means that, when combined, would be convincing to commercial carriers.

53. *Q: Section II.3 and Section IV of the Phase 1 Proposal Format are both asking for descriptions of technology transition plans. Can you clarify what aspects of transition should be covered in these two places?*

A: Section IV should serve as the main focus of the technology transition plan. Sections II.3 and IV can cross-reference each other to minimize repetitions.

Testing and OPNET-related Issues

54. *Q: If not explicitly required, OPNET is at least effectively required. Has DARPA arranged for pricing for OPNET within CORONET?*

A: No. Note that OPNET is not required in the proposal stage.

55. *Q: Is OPNET tool/license being given free for all interested parties?*

A: No. Note that OPNET is not required in the proposal stage.

56. *Q: OPNET scalability issue (in Phase I) (simulation time for 75 nodes in CONUS is going to be extremely long) especially when there are many sub-wavelength connections for guaranteed-bandwidth IP services and best-effort IP services.*

A: Simulations should be run at an appropriate level of fidelity based on what is being investigated. For example, protocol verification may use a higher level of fidelity than statistical simulations aimed at measuring blocking probabilities. We anticipate that scalability with OPNET will not be an issue, assuming appropriate level models are run.

57. *Q: Can large-scale simulations (with 75 nodes and sub wavelength service traffic) be done at MIT/LL (where a supercomputer may be available)?*

A: No.

58. *Q: Can the independent test and evaluation lab be involved with the teams as they prepare their proposals? In the proposal process, can questions be asked directly of the independent testing lab?*

A: No. Interactions with the independent test and evaluation lab should not occur until after contract award. During the proposal stage, all questions must be directly sent to DARPA at: BAA06-29@darpa.mil.

59. *Q: To achieve proof-of-concept verification at end of Phase I, are all research innovations to be software based since hardware development, test & evaluation is not permitted in CORONET? If yes, what is/are assumption(s) to make for software processing platform limits?*

Processing platform=proc type(s) + proc speed(s) + memory + bus specifications

A: The verification at the end of Phase 1 is all software-based. In the simulations, state-of-the-art CPUs should be assumed to be deployed at the network nodes to perform network control and management duties.

Standards

60. *Q: Should research proposals address industry standards and inter- operability issues with disparate/immature standards? If so, which standards are of pertinence/interest to CORONET?*

A: Awareness of standards activity is encouraged. However, there are no explicit standards that are required to be included as part of the program. Clearly, it would be advantageous if the solution is compatible with established and emerging standards.

61. *Q: Is concurrent standards development activities in CORONET Phase I desired/supportable?*

A: No, the CORONET Phase 1 funding does not include standards activities.

62. *Q: Are layer 3 transport protocols to be redesigned? Clearly TCP is not consistent with some of the requirements (e.g., 1 sec download of 10 Gb).*

A: Redesigning protocols such as TCP is not in-scope for this project. However, proposers may consider cross-layer IP and optical layer optimizations to achieve the goals of the program (e.g., rapid service setup, rapid and efficient restoration, efficient demand aggregation).

Miscellaneous (related to the network design exercises)

63. *Q: Geographic distribution of traffic generating nodes?*

A: The topology and demand matrix used must be realistic. Customizing the topology and demands for the sole purpose of meeting the metrics is not acceptable. For example, having all traffic be between neighboring nodes in order to meet the capacity requirements is not acceptable.

64. *Q: In Table 3, for the scenarios where multiple fiber pairs are allowed, are links required to have the multiple fiber pairs, or just the links that need them?*

A: Only the links that need this much capacity need to have multiple fiber pairs.

65. *Q: Table 4: when it states that a maximum of 10 (or 7) optical network elements can be traversed before requiring regeneration, does this include amplifier sites?*

A: No.

66. *Q: How do you define a node?*

A: As stated on page 8 of the PIP, nodes are the sites in the network where terminals and switches are located, where traffic is added, dropped, switched, and/or groomed. For example, sites that simply have a line optical amplifier are not considered nodes.

67. *Q: Are the scenarios in table 3 stand-alone (treat each as desert-start) or do designs build one each other?*

A: Each design scenario should be treated as a 'greenfield' (desert-start) design; i.e., the designs do not have to build on the previous designs. However, note that practical issues related to evolving a network (e.g., supporting a mix of line rates) should be investigated as part of Phase 2 of the program.

68. *Q: To engineer the capacity to support the IP traffic, the cost ratio between IP and optical elements (e.g. line card) is a very important parameter. Are we allowed to make assumptions on relative cost numbers?*

A: Performers should make realistic assumptions with regard to factors such as the cost ratio between IP and optical elements. Any such assumptions should be stated, along with the justification for the assumed value.

Programmatics

69. *Q: What are the expected number and amount of awards in Phase I and phase II, respectively?*

Will multiple awards be made for Phase I?

How many Phase I awards do you anticipate?

A: The government reserves the right to select for award all, some, or none of the proposals. The expected number and amount of awards in each phase is not specified. However, it is possible that multiple awards will be made in Phase 1.

70. *Q: Can you comment on the level of funding available for this program and the allocation between Phase I and II?*

What is the value of the overall program?

What is the projected DARPA funding for the Phase I and Phase II efforts?

How much money is expected to be available in Phase I and how many projects are expected to be funded?

A: The level of funding for the program as a whole and its allocation between the phases has not been determined.

71. *Q: When do you expect to announce the winner(s)?*

A: The earliest anticipated award is planned to occur in December 2006. (See page 5 of the PIP.)

72. *Q: Will an attendance list of Proposers' Day be made available to assist in teaming discussions?*

A: The attendee list from Proposers' Day will be made available. However, individuals may request that their names and contact information not be included on the list.

73. *Q: Must primes be U.S. companies excluding universities or does "companies" include universities?*

A: While there is no restriction on universities serving as a prime, it is anticipated that, as a practical matter, US companies will serve as the primes.

74. *Q: Do you expect the results of the proposal-stage network study to be provided in OPNET format?*

A: No. OPNET is not required for the proposal.

75. *Q: Can a person or an organization be a subcontractor in multiple teams?*

A: Yes, assuming the primes are aware of, and approve of, this arrangement.

76. *Q: Can a company be a prime of multiple teams?*

A: Yes

Corrections to PIP

Correction: Table 5, Footnote (2). This should read:

"The required very fast service-setup times can be defined more generally as: 50 msec + The round trip fiber transmission delay between the source and destination nodes"

Correction: Table 6, Footnote (2). This should read:

"*Single-Failure Restoration Time (SFRT)* can be defined more generally as: 50 msec + The round trip fiber transmission delay between the source and destination nodes"

CORONET BAA/PIP
Round 2 of Questions and Answers
August 30, 2006

Some questions may have been edited slightly for clarity.
Note that the question numbers are a continuation of the list from the first round of Q&A.

If any statements below conflict with the BAA or PIP, the BAA and PIP take precedence.

Traffic Profile

77. *Q: Page 9 of the PIP states: "It is also important that the performers demonstrate realism in generating a traffic distribution to meet the specifications of the CORONET program. ... Traffic distributions based on 'all-to-all' traffic are not admissible." Non-uniform node sizes and corresponding traffic patterns are described. Is a distance-independent traffic profile considered a realistic and/or desirable evaluation profile?*

A: A totally distance-independent traffic profile is not very realistic or desirable for evaluation purposes. For example, traffic between small nodes typically is somewhat distance-dependent. Also, there are generally communities of interest among some of the large nodes that generate a greater amount of traffic; for example, traffic between some of the large nodes along the East coast, or traffic between the East and West coasts.

Architecture

78. *Q: What should be done if performers wish to propose an architecture that is different from what is specified in the PIP?*

A: At a minimum, performers must evaluate their proposed scheme using the model specified in the PIP. Performers also may propose alternative models if these models achieve significantly higher performance with comparable or lower cost, or comparable or higher performance with significantly lower cost. If performers believe that the development and analysis of an alternative model will represent a significant cost, then performers may include this task as a *fully costed option* in the proposal. These program options may or may not be exercised by the Government at the time of selection of the performers for Phase 1.

Control Plane

79. *Q: For the control plane implementation, is there an assumption on the control channel capacity?*

A: It is up to performers to determine the control channel capacity that is required by their proposed system. The required capacity should be included as part of the description of the control plane architecture in the Phase 1 deliverable.

80. *Q: A follow-up question along the lines of Question 37 in the Proposers' Day Q&A document (regarding a secure control plane): What assumptions can/are to be made about control/management plane message passing:*
- a) Is the control plane to be realized over a public or dedicated network?*
 - b) If dedicated, does the control plane support access, e.g., craft interface, via the public Internet?*
 - c) Does the control plane follow the same link topology as the data plane?*
 - d) If the control plane has the same link topology as the data plane, does a link failure also fail the control plane as well as the data plane, such as would be the case if the control plane were carried within a data-plane wavelength or on a separate data-plane wavelength?*

A:

- a) The control plane is to be realized over a dedicated network.
- b) Access via the public Internet could be a useful feature. It is up to performers to decide if and how to implement this feature.
- c) One can assume that the control plane follows the same link topology as the data plane.
- d) A link failure causes both the control and data planes to fail on that link. The same is also true in the case of a node failure (see page 24 of the PIP).

OPNET-related Issues

81. *Q: Is there a preferred OPNET baseline for CORONET? Modeler version or Kernel release?*

A: OPNET Modeler version 11.5 or later.

82. *Q: Is there a preferred OPNET Hardware? (E.g., Intel?) Uniprocessor? multiprocessor?*

A: No specific hardware is prescribed. A faster/dual processor system will reduce the run time of the simulation over a slower/single processor system, but will not affect the simulation results.

83. *Q: Is there a preferred OPNET OS? (UNIX? WIN XP? etc.)*

A: The independent test and evaluation lab can run OPNET in either a Windows or a UNIX environment.

Standards

84. *Q: A follow-up question to Question 62 in the Proposers' Day Q&A document: New IP/optical architectures may require changes in the IP layer such as new routing protocols. We understand that TCP is an end-system protocol and is out-of-scope. However, some of the cross-layer aspects of radical new IP/optical architectures will require changes to IP protocols. How much of layer 2/3 IP protocol layers are we allowed to or precluded from modifying?*

A: While it is desirable to use existing standards as much as possible, performers are permitted to propose modifications to the IP protocol if such changes are required to achieve the goals of the CORONET program. Performers should describe how their proposed solution would be impacted if the required changes to IP were not adopted as part of the standard. Note that funding of standards activities is not included in the program.

Miscellaneous (related to the network design exercises)

85. *Q: In order to assess the packet loss performance listed in Table 7 using OPNET simulations, is packet loss measured from ingress router to egress router including any losses due to router buffer overflows (in which case the queuing structure of the routers would be modeled in OPNET as well)?*

A: Packet loss is measured from ingress router to egress router and includes losses due to buffer overflow. Performers need to determine the appropriate level of modeling that will demonstrate that the metrics of Table 7 are being met. Ensuring that the metrics of Table 7 are met is clearly one of the technical challenges of the CORONET program. Note that a description of the methodology used to ensure and measure the IP QoS metrics is one of the technical deliverables in Phase 1 of the program (see page 29 of the PIP). Also note that the metrics of Table 7 are not a part of the proposal-stage network study (see page 41 of the PIP).

86. *Q: When delivering the Excel file with the network topology information, should the link distances be in kilometers or miles.*

A: All distances should be in kilometers.

CORONET BAA/PIP
Round 3 of Questions and Answers
October 2, 2006

Some questions may have been edited for clarity.

Note that the question numbers are a continuation of the list from the previous two rounds of Q&A.

If any statements below conflict with the BAA or PIP, the BAA and PIP take precedence.

Protection

87. *Q: Is there a minimum amount of traffic that must be 1+1 protected in the proposal-stage network study?*

A: There is no requirement that any of the traffic be 1+1 protected in the proposal-stage network study. As stated on page 16 of the PIP, it is likely that some type of shared mesh protection will be needed in order to meet the capacity metrics.

OPNET-related Issues

88. *Q: Will the BAA accept OPNET models that incorporate processes running outside of OPNET through HLA (High Level Architecture)?*

A: The BAA will accept OPNET models that incorporate processes running outside of the OPNET environment subject to several conditions. First, proposers must specify why the processes are needed to meet the program requirements and must substantiate that the processes cannot be run using standard OPNET modules. Second, the processes that are run outside of OPNET must be written in a standard programming language (e.g., C) and must be sufficiently well documented and validated such that they can be verified by the independent test and evaluation lab.

89. *Q: For the purposes of budgeting for OPNET modeling resources in Phase 1, is it possible to describe the level of OPNET model fidelity required by BAA?*

A: It is up to the performers to determine the level of modeling fidelity that is required to demonstrate that the program metrics have been met.

Miscellaneous

90. *Q: If we think a different set of conditions may be more appropriate in the metrics tables (e.g., for commercial applications), how should we proceed?*

A: Note that the CORONET BAA is trying to push the envelope of this technology. Thus, at a minimum, proposers need to satisfy the conditions that have been specified in the PIP. In addition, proposers may also present results for scenarios that they feel are more appropriate.

CORONET BAA/PIP
Round 4 of Questions and Answers
October 20, 2006

Some questions may have been edited for clarity.

Note that the question numbers are a continuation of the list from the previous three rounds of Q&A.

If any statements below conflict with the BAA or PIP, the BAA and PIP take precedence.

Protection

91. *Q: There is a requirement for certain types of traffic to be resilient to 1, 2, and 3 network element failures. Does this mean that for a traffic type that is resilient to n failures, that it should be resilient to only n failures in the entire network, or that it should be resilient to n failures that afflict its specific data paths? Thus, if we have k types of traffic that are n -resilient, then all of these traffics are required to have service despite kn failures?*

A: There are three restoration classes specified in Table 6: single-failure, double-failure, and triple-failure. These classes are referring to link and node failures only, where at most one of the failures is a node failure. For purposes of meeting the metrics, when there are multiple failures, assume there is enough time between failures such that recovery from one failure occurs before the onset of the next failure. These specifications mean:

- For any single failure in the entire network (either one link failure or one node failure), all single-protected, all double-protected and all triple-protected traffic in the network affected by the failure must be restorable.
- For any combination of a first failure and a second failure in the entire network (either two link failures or one link failure and one node failure), all double-protected and all triple-protected traffic affected by either or both of the failures must be restorable (of course, all single-protected traffic affected by just the first of the two failures must be restorable).
- For any combination of a first failure, a second failure, and a third failure in the entire network (either three link failures or two link failures and one node failure), all triple-protected traffic affected by any one or any two or all of the failures must be restorable (of course, all double-protected traffic affected by just one or both of the first two failures must be restorable; all single-protected traffic affected by just the first of the three failures must be restorable).

Miscellaneous (related to the network design exercises)

92. *Q: On page 41 you state (concerning the network study): "The traffic should be assumed to be 100% fixed (i.e., no churn). The proposal must include a discussion on how the results will be affected by traffic churn."*

To me, no churn implies a holding time of infinity. This obviates the dynamic traffic classes in Table 5 ("holding time") row.

If above is indeed the correct interpretation, then please clarify your IP-layer and WDM-layer link utilization assumptions for carrying the working traffic in the network study.

A: In order to simplify the proposal-stage network study, the PIP specifies that the traffic should be considered as fixed. Thus, proposers should generate an appropriate traffic matrix that meets the service and protection requirements of Table 2 and Table 6, respectively, and produce a network design where all the traffic is added at once to the network. The holding times are infinite. As stated on page 41 of the PIP, the metrics of Table 5 are not a part of the proposal-stage network study. (Also, as stated on page 21 of the PIP, any results obtained via simulation should be based on a sufficient number of runs to produce reasonable confidence intervals. Thus, it is expected that performers would generate several traffic matrices and generate designs for each one, and present results that are, for example, the average of all designs.)

While there is no dynamic component to the proposal-stage network design, proposers should include a discussion of how their algorithms will likely be impacted by churn. For example, proposers may want to include estimates along the lines of: dynamic traffic will result in X% less efficient utilization of the bandwidth; or, dynamic traffic will result in Y% more required shared protection resources.

Dynamic traffic and the metrics of Table 5 are applicable in the Phase 1 network studies.

93. *Q: What is the maximum utilization of a wavelength that can be assumed for IP traffic; i.e., how much overhead should be reserved to account for the burstiness of the best-effort IP traffic?*

A: Proposers should make reasonable assumptions as to the overhead required for the bursty IP traffic, and should justify these assumptions. As was answered in Question 10 of the first round of Q&A, reasonable assumptions should be made by the proposers regarding the burstiness of the best-effort IP traffic, and these assumptions should be explicitly stated.

Programmatics

94. *Q: Will DARPA accept and fully review proposals that address BAA 06-29 Phase 1 requirements if the proposals are received by DARPA after 11-17-06? For example, if DARPA receives a proposal addressing Phase 1 requirements and requesting Phase 1 funding on 12-17-06 (i.e., 1 month after the initial deadline), is it possible that the proposal response will be accepted, reviewed, awarded, and funded? If the answer is yes, then up to what time period will Phase 1 proposal responses be considered for funding?*
A: Proposers may submit a proposal under BAA06-29 at any time during the open period for the announcement, which closes 4 August 2009. However, in order to be considered for the first round of selections, proposals must be received no later than 12:00 Noon EST, 17 November 2006.

95. *Q: We would like to bid the CORONET program as a fixed support TIA (see DoD Grant and Agreement Regulation DoD3210.6-R, March 26, 2003, and Interim Guidance for Technology Investment Agreements Revision 2, February 3rd, 1999). Our motivation to do this is to minimize issues related to financial management systems, time card and cost reporting, and other factors with our consortia and particularly our subcontractors. Section 37.215 of this document states that the granting agency must seek cost share for TIA's in general, but section 37.205 states that a fixed-support TIA does not require cost share.*

Question 1: Will DARPA accept a fixed-support TIA as a proposed procurement vehicle for CORONET?

Question 2: If we bid a fixed-support TIA, will we be automatically excluded from evaluation?

Question 3: Will DARPA accept a fixed-support TIA with no cost share as a proposed procurement vehicle for CORONET?

Question 4: If we bid a fixed-support TIA with no cost share, will we be automatically excluded from evaluation?

Question 5: Can DARPA provide any templates, forms, or guidance for submission of a cost proposal that conforms to a fixed support TIA with no cost share to enable a response to CORONET?

Question 6: Can DARPA provide the URL's for Federal Government Web sites that have specific information related to submission of a cost proposal that conforms to a fixed support TIA with no cost share?

A1: Yes.

A2: No you will not be excluded from evaluation.

A3: Yes, if sufficient rationale is provided to warrant a no cost share arrangement.

A4: No you will not be excluded from evaluation.

A5: None are available.

A6: No known URL's for the requested information.

96. *Q: As it is the ultimate goal of the program to transition the CORONET technology to commercial telecommunications carriers, does DARPA anticipate that any Organizational Conflicts of Interest would arise under the CORONET program that would eliminate the CORONET contractor from the competition for a procurement contract to transition the technology or require a mitigation plan due to Organizational Conflicts of Interest?*

A: There is not enough information provided to specifically address this question. However, in general, if a proposer believes there is a potential organizational conflict of interest, as discussed in Paragraph 3.3 of the Proposer Information Pamphlet, "...All facts relevant to the existence or potential existence of organizational conflicts of interest (Federal Acquisition Regulation 9.5) must be disclosed."

CORONET BAA/PIP
Round 5 of Questions and Answers
October 27, 2006

The questions may have been edited for clarity.

Note that the question numbers are a continuation of the list from the previous four rounds of Q&A.

If any statements below conflict with the BAA or PIP, the BAA and PIP take precedence.

Miscellaneous (related to the network design exercises)

97. *Q: The answer to Question 92 included the statement: "Also, as stated on page 21 of the PIP, any results obtained via simulation should be based on a sufficient number of runs to produce reasonable confidence intervals. Thus, it is expected that performers would generate several traffic matrices and generate designs for each one, and present results that are, for example, the average of all designs." Does this comment apply to the proposal-stage network study?*

A: The main intent of this comment was to ensure that the results presented for the proposal-stage network study are meaningful and valid. However, performers could do what they feel is necessary to produce valid results; they should also state what they did.

Programmatics

98. *Q: Where in the proposal should the performers include any "options" for tasks that go beyond the core requirements specified in the PIP? Is there any page limitation associated with these options? How should the associated deliverables, rationale, approach and plan for these "options"-related tasks be organized in the proposal, given the current limitation on the number of pages?*

A: See CORONET PIP Modification 2.

CORONET BAA/PIP
Round 6 of Questions and Answers
November 7, 2006

The questions may have been edited for clarity.

Note that the question numbers are a continuation of the list from the previous five rounds of Q&A.

Programmatics

99. *Q: Will the CORONET project fall under the ITAR (International Traffic in Arms Regulations)?*

A: Note that the International Traffic in Arms Regulations (ITAR) is only one of the U.S. Export Control Laws that apply to DARPA projects, though it is one of the two "biggest" applicable laws. The other is the Export Administration Regulation (EAR), administered by the Commerce Department. The ITAR contains a list of defense-related articles and associated information (the U.S. Munitions List) that are subject to U.S. Export Control Laws.

With this in mind, the answer to your question is addressed in the CORONET PIP in the following sections:

- Page 4, End of Sec. 1.2.
- Page 20, End of Sec. 2.4.1.
- Page 36, Sec. 3.8.

To summarize, a direct answer to the question is that DARPA does not anticipate that the research performed as part of Phase 1 of the CORONET Program will be ITAR controlled. However, this could change as the program evolves. Also, it is possible that the development stage of Phase 2 will fall under ITAR control. Note that in general the contractual responsibility for determining the requirements of US Export Control Laws (not just ITAR) and continuously complying with them is placed on the performers.

Proposal Formatting

100. *Q: In Vol. I, Sect III (Network Study) the network topology, including cities, links, and link distances, for a 100-node network is requested. It does say that you want this information in "MS-Excel-readable format", but does not state whether it must also appear printed in the proposal. Does the list of links in the network topology need to be included in the printed proposal or can it just be provided on the CD-ROM?*

A: The description of the network topology used in the proposal-stage network study (i.e., the list of network links) is to be included in an MS-Excel-readable file on the CD-ROM. Also, as stated on page 41 of the PIP, an illustration of the topology must be included as part of Section III of the proposal. There is no requirement to include the list of network links as part of the proposal text.

101. Q: The CORONET PIP explicitly states that 12pt font is the smallest allowed in Volume I. May we use a font smaller than 12pt within tables and figures in Volume I, as is allowed in Volume II (Cost proposal)?

A: If needed, you may use a font type smaller than 12 point within any tables and figures included in the proposal, however, the font type in the tables and figures can be no smaller than 8 point.