

AD-A285 508




1

IDA DOCUMENT D-1451

SCALEABILITY PANEL REVIEW

David R. Cheriton
Dale B. Henderson
Duncan C. Miller
David L. Mills
Stuart D. Cheshire, *Adjunct Panel Member*
Randy L. Garrett, *Task Leader*
Julia A. Loughran

October 1993



1482

Prepared for
Advanced Research Projects Agency

DTIC QUALITY INSPECTED

Approved for public release, unlimited distribution: May 6, 1994.

94-31946



INSTITUTE FOR DEFENSE ANALYSES
1801 N. Beauregard Street, Alexandria, Virginia 22311-1772

DEFINITIONS

IDA publishes the following documents to report the results of its work.

Reports

Reports are the most authoritative and most carefully considered products IDA publishes. They normally embody results of major projects which (a) have a direct bearing on decisions affecting major programs, (b) address issues of significant concern to the Executive Branch, the Congress and/or the public, or (c) address issues that have significant economic implications. IDA Reports are reviewed by outside panels of experts to ensure their high quality and relevance to the problems studied, and they are released by the President of IDA.

Group Reports

Group Reports record the findings and results of IDA established working groups and panels composed of senior individuals addressing major issues which otherwise would be the subject of an IDA Report. IDA Group Reports are reviewed by the senior individuals responsible for the project and others as selected by IDA to ensure their high quality and relevance to the problems studied, and are released by the President of IDA.

Papers

Papers, also authoritative and carefully considered products of IDA, address studies that are narrower in scope than those covered in Reports. IDA Papers are reviewed to ensure that they meet the high standards expected of refereed papers in professional journals or formal Agency reports.

Documents

IDA Documents are used for the convenience of the sponsors or the analysts (a) to record substantive work done in quick reaction studies, (b) to record the proceedings of conferences and meetings, (c) to make available preliminary and tentative results of analyses, (d) to record data developed in the course of an investigation, or (e) to forward information that is essentially unanalyzed and unevaluated. The review of IDA Documents is suited to their content and intended use.

The work reported in this document was conducted under contract MDA 963 89 C 0003 for the Department of Defense. The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.

© 1993 Institute for Defense Analyses

The Government of the United States is granted an unlimited license to reproduce this document.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE October 1993		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE Scaleability Panel Review			5. FUNDING NUMBERS MDA 903 89 C 0003 ARPA Assignment A-132	
6. AUTHOR(S) David R. Cheriton, Dale B. Henderson, Duncan C. Miller, David L. Mills, Stuart D. Cheshire, Randy L. Garrett, Julia A. Loughran				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses (IDA) 1801 N. Beauregard St. Alexandria, VA 22311-1772			8. PERFORMING ORGANIZATION REPORT NUMBER IDA Document D-1451	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Advanced Research Projects Agency (ARPA) 3701 North Fairfax Dr. Arlington, VA 22204-1714			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, unlimited distribution; May 6, 1984			12b. DISTRIBUTION CODE A	
13. ABSTRACT (Maximum 200 words) A peer review panel was held at IDA on August 19 and 20, 1993, to address the challenges of network scaleability as it pertains to the Synthetic Theater of War (STOW) program. The goal of STOW is to conduct a Distributed Interactive Simulation (DIS) demonstration in 1994 with 10,000 entities (live, virtual, and constructive), with 100,000 entities by 1997, over the DSI network. Scaleability addresses techniques for maximizing the number of DIS entities on the network by minimizing the network load, both in terms of bandwidth and packets per second. The panel, which consisted of five network scaleability experts, was presented with the current research being conducted by three separate contractors. The experts provided insights on how STOW's goals could be met. They expressed concern over meeting the requirement for a secure network, given the current network encryption requirements. They recommended T3 circuits as a backbone to handle the network traffic. They also indicated that Multicast Groups (MSGs) appear essential along with other traffic reduction techniques and they stressed the importance of the scenario setup.				
14. SUBJECT TERMS STOW, Distributed Interactive Simulation (DIS), Network Scaleability, Scaleability, Network, Computer Security, T3 Circuits, Multicast Groups (MSGs).			15. NUMBER OF PAGES 162	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR	

UNCLASSIFIED

IDA DOCUMENT D-1451

SCALEABILITY PANEL REVIEW

David R. Cheriton
Dale B. Henderson
Duncan C. Miller
David L. Mills
Stuart D. Cheshire, *Adjunct Panel Member*
Randy L. Garrett, *Task Leader*
Julia A. Loughran

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

October 1993

Approved for public release, unlimited distribution: May 6, 1994.



INSTITUTE FOR DEFENSE ANALYSES

Contract MDA 903 89 C 0003
ARPA Assignment A-132

UNCLASSIFIED

PREFACE

This document was prepared by the Institute for Defense Analyses for the Advanced Research Projects Agency (ARPA), under the ARPA Assignment, Synthetic Theatre of War (STOW), and fulfills the objective of providing "technical support focused to the needs of STOW. Areas to be addressed include . . . technology assessments relevant to the successful execution of STOW (e.g., pertaining to scalability . . .)."

The document was reviewed by Dr. Richard J. Ivanetich, Director of the Computer and Software Engineering Division, Institute for Defense Analyses.

Table of Contents

INTRODUCTION AND SUMMARY	1
APPENDIX A - PANEL MEMBER'S COMMENTS	A-1
Comments on STOW '94 and '97 Scaleability Issues	A-3
<i>David R. Cheriton, Stanford University</i>	A-3
Scaleability for STOW '94 & '97	A-9
<i>Dale B. Henderson, Los Alamos National Laboratory</i>	A-9
Issues of Scaleability for STOW	A-13
<i>Duncan C. Miller, MIT Lincoln Laboratory</i>	A-13
Summary of STOW Panel Review	A-17
<i>David L. Mills, University of Delaware</i>	A-17
Scaleability Issues for STOW	A-23
<i>Stuart D. Cheshire, Stanford University</i>	A-23
APPENDIX B - BIOGRAPHICAL SKETCHES OF PANEL MEMBERS	B-1
APPENDIX C - PRESENTATIONS	C-1
APPENDIX D - READING LIST	D-1
APPENDIX E - LIST OF ATTENDEES	E-1
LIST OF ACRONYMS	Acronyms-1

INTRODUCTION AND SUMMARY

Introduction

At the request of Colonel Robert Reddy, Assistant Director of the Advanced Research Projects Agency/Advanced Systems Technology Office (ARPA/ASTO) and Program Manager for Synthetic Theater of War (STOW) as well as Advanced Distributed Simulation (ADS), a Peer Review Panel was convened to review the work being done on network scalability. The term *scalability*, as used in this document, refers to techniques to maximize the number of distributed interactive simulation (DIS) entities on a network by minimizing the load presented to the network, both in terms of bandwidth and packets per second.

The Panel was asked to consider the following questions:

- Are these approaches the most promising? If not, what would be better?
- Do the separate programs fit together to form a coherent approach?
- Are the programs practicing "good science"? Are good experimental methodologies being used?
- Are the programs incorporating leading edge technology?
- What is the Panel's risk assessment of these approaches, and what can be done to manage risk?

The Panel met at the Institute for Defense Analyses, Alexandria, Virginia, on August 19 and 20, 1993. The Panel members were as follows:

- Dr. David R. Cheriton, Stanford University, Palo Alto, California
- Dr. Dale B. Henderson, Los Alamos National Laboratory, Los Alamos, New Mexico
- Dr. Duncan C. Miller, Massachusetts Institute of Technology Lincoln Laboratory, Lexington, Massachusetts

- Dr. David L. Mills, University of Delaware, Newark, Delaware
- Mr. Stuart D. Cheshire, *Adjunct Panel Member*, Stanford University, Palo Alto, California

The Navy Research, Development, Test and Evaluation Division (NRaD) of the Naval Command, Control, and Ocean Surveillance Center, San Diego, California, provided a "read-ahead" package.

Presentations were given by the following:

- CDR Dennis McBride, ARPA/ASTO Program Manager for Scaleability
- Ms. Linn Flynn, ARPA/ASTO Program Manager for the Distributed Simulation Internet
- Mr. Dan van Hook, Loral Advanced Distributed Simulation (LADS)
- Mr. William Miller, American Telephone and Telegraph (AT&T)
- Dr. Joshua Seeger and Mr. Lou Berger, Bolt Beranck and Newman (BBN)

Summary

Conclusions

The panel provided the following overall conclusions:

- 10,000 entities on a secure network in 1994 seems impossible, given the restrictions of the Motorola Network Encryption System (NES).
- T-1 tail circuits (about 1.5 megabits per second throughput) will require approximately a 90% traffic reduction from the full exercise load that might be presented to the backbone. If a T3 (about 45 megabits per second) backbone is available, then 10,000 entities should not present a problem to the backbone. Therefore, a 90% reduction in traffic down the tail circuit is required.
- MCGs (multicast groups) seem absolutely required for large numbers of entities on the wide area network, in conjunction with other reduction techniques. Much more study is needed on the effects of MCGs versus traffic reduction. If a secure network is essential, then STOW needs to prepare to fall back to less than 10,000 entities for 1994.
- Achieving a 90% traffic reduction will require at least 100 MCGs and perhaps many more. If an upgraded NES will not do more than 16 MCGs each, then

there may be a possibility to trade off aggregating traffic to provide increased shared bandwidth versus using parallel NES to obtain an increased number of MCGs. If about 6 to 10 NES are needed to get 100 MCGs, then 130 NESs are needed to support 13 sites. These 130 NESs would cost about \$4 million.

- The Panel emphasized the importance of a judicious scenario setup to match the virtual world with the physical (network) sites in order to achieve traffic reduction. This must always be considered, and the interactions and logistics are complex. Pathological conditions that would require broadcasting all traffic to all sites must be avoided. Any possible, arbitrary scenario is not feasible. Knowledge of the intricate details of DSI configuration is vital to make the most efficient utilization of the network.
- For 100,000 entities, MCGs will be essential. A single site seems unlikely to need to receive more than 10,000 entities even in future, which should help bound the required size of the tail circuits and capabilities of the site network equipment and computers.

Comments on Specific Presentations

LADS

The LADS presentation was highly impressive. Its work on a network simulation using real data and SAF (semi-automatic forces) is most valuable and needs to continue. The impact of the proposed algorithms on simulation fidelity must be addressed. Also, the effectiveness of the algorithms should be confirmed using real-life field tests. The network simulation needs to be tested against real network experiments. There appears to be good cooperation and division of labor between LADS and BBN. The On-Demand Forwarding algorithm may be more promising than geographical. There is no need to have both LADS and AT&T look at TimeOuts. More detail should be incorporated into the network traffic simulation. Accurate modeling of MCG joins and leaves may not be a 1994 issue, given the static nature of NES MCGs. The effect of compression should be quantified.

AT&T

An Advanced Interface Unit (AIU) or "intelligent gateway" is needed to reduce traffic, but the implementation is very difficult. It was not clear what AT&T was proposing to put in the AIU. There appeared to be some risk in having AT&T develop the AIU due to the lack of previous experience in this particular area. Due to changes in the networking requirements, the AT&T work is no longer as relevant to ASTO programs as when the orig-

inal contract was awarded; hence, some of the comments may seem more negative than otherwise would have been expected.

The AT&T recommendation to use delta state updates and situational information was similar to LADS geographical filtering, except AT&T proposed to reduce the number of bits in the protocol. Using delta PDUs may introduce simulation compromises. This presentation had too much emphasis on speculation and too little experimental science.

General Comments

- There needs to be another component in the Scaleability program analyses than just bottom-up, although there is the danger in getting only fluff with a top-down approach. STOW will need to use multiple techniques in parallel to achieve the required network traffic reduction, but the use of multiple techniques vastly complicates the simulators and networks.
- Some things are network-wide issues that cannot be solved locally.
- STOW should look carefully at DSI transition.
- STOW has no need for network reservation since the projected number of T-1 tail circuits cannot overload a T-3 backbone.
- STOW *does need to be able to set priorities for traffic coming on/off tail circuits.*
- NES is a big concern. STOW needs a more detailed network simulation incorporating NESs to determine real risks and throughput.
- Surprisingly and disconcertingly, many questions that ought to have clear answers did not. Too often there is a need for different people from different organizations to get together to answer questions. The Scaleability program and the DSI have a problem with integration—determining what topics to be investigated and what needs to be coordinated.
- The AIU is a complicated, sophisticated box. Multiple people are working on building it. Functionality continues to be added in, but STOW needs to start building AIUs very soon. Some proposed bandwidth reductions are really DIS protocol changes that could possibly be added into AIUs.
- Using satellites with delay compensation (through the use of dead reckoning with some adjustments for propagation time delay) to minimize the impact of latency should be considered as a replacement for always using terrestrial network links.

APPENDIX A
PANEL MEMBER'S COMMENTS

Comments on STOW '94 and '97 Scalability Issues

David R. Cheriton
Stanford University

A major constraint on STOW '94 is the security requirement, given the apparent necessity to use the Motorola NES to satisfy this requirement. The panel only really came to grips with this limitation on the second day of presentations, and this constraint dominated some issues that were presented and discussed earlier.

The NES introduces some major uncertainties into the program. There were proposals to gang a number of NESs together to increase data rates, but there has not been any experience with this approach to date. The proposed approach did not increase the number of multicast groups beyond the hoped-for 16 per NES because each is configured identically to provide flexible load-sharing. An alternative approach, in which each NES at a site provides a separate set of 16 multicast groups, reduces this flexibility of load sharing but increases the number of multicast groups. Little is known about the behavior of traffic in this configuration either. If ganging NESs at a site to support secure traffic close to the T1 speeds of the line is cost justified (see below), study is required to better understand the behavior of the multicast-based traffic reduction techniques within this network configuration. That is, how well do $16 \cdot k^1$ multicast groups or so per site reduce traffic (assuming we have k NES nodes per site), and can this traffic be expected to be distributed fairly uniformly across the groups so there are not bottleneck NES nodes.

The value of k (NESs per site) was estimated as having to be around 6 to 10 to achieve secure bandwidth that matched the speeds of the T1 site tail circuits. This number might be lowered somewhat, based on further experiments, but one possible merit in using a larger value of k is providing more multicast groups, namely $16 \cdot k$.

¹ [where k is any number.]

At \$17,000 per node, this cost would represent a significant, possibly prohibitive, and likely throw-away (after STOW '94) cost. Given the major cost of this model, one might consider non-trivial alternatives. One further alternative traffic reduction technique that was raised but not explored in detail was to replicate the SAF execution at each site so that this traffic is largely removed from the tail-circuits. It is estimated that 90 percent of the traffic will be SAF generated. Buying extra computer server hardware to run redundant SAF at each site and reducing the capital expenditures on NESs seem preferable based on the reusability of some computer servers after STOW '94. For example, these could be additional SGI workstations or similar. However, this direction would require some further study and development to coordinate SAF between different sites, and may in fact involve modifications to the SAF to increase the determinism of their execution. Work in this area may have long term pay-off in providing another opportunity for trading off between computation and communication, an issue that can be expected to be relevant in '97, given the expected continued dependence on satellite communication. However, it should be emphasized that this is an area that would require effort to investigate and develop as an option; not all of the panel even convinced that this was clearly feasible. Perhaps an effort in this area could be considered in conjunction with Commander McBride's program.

Considering STOW '97 briefly, there seems to be a potentially excessive trust in ATM technology adequately solving all the bandwidth problems. Multicast is not currently well-supported in the ATM standard, ATM has not been tariffed in general, and other competing technologies may come along or be necessary to use, an obvious example being satellite. As one aspect of accommodating these uncertainties with the network technology, there should be greater emphasis on NSA developing a network-technology independent encryption device (N-TIED). I plan to submit a sketch proposal for such a device as a separate document.

In this same vein, I think it would be worth further investigation of delay compensation techniques so that satellite links could be confidently incorporated in the future. They appear necessary for both '94 and '97. The dead reckoning mechanism provides a basis for delay compensation so one may be able to build this as an extension of this existing mechanism.

Further on STOW '97, there is relatively little understood about the scalability problems for STOW '94 and the actual behavior of 10,000 entities. We should try to keep our options as open as possible for '97, given that a success in '94 and further progress with DIS (such as dynamic terrain) may significantly drive up the requirements for '97 beyond

that expected from purely increasing the number of entities. The key point is not that there is a problem but there is definitely risk, and exploring lots of options and keeping them open seem like the best approach to minimizing these risks. As one example, relatively little attention has been paid in the performance studies to the effects of wide-area perception units, such as stealth vehicles and JSTARS. These units may create significant load demands on the network beyond that of normal entities, and their importance may increase by '97.

Considering the presentations overall, I was disappointed that the presenters did not explicitly address fidelity maintenance in most instances of discussing various optimizations. Fidelity must be considered in each and every optimization, not because it is necessarily a problem, but we cannot afford to lose fidelity beyond some tolerance level. I also felt that the presenters did not in general distill out the key issues and solutions. For example, we were dragged through grid-based multicast filtering only to discover that so-called on-demand forwarding (ODF) was better, and then only to discover that ODF required far more multicast groups that would be available in a secure network in '94. The latter suggests a lack of communication between groups. (I pick that particular example because van Hook, as one of the strongest presenters, seems more diplomatic to pick on.)

Comments on Specific Presentations

van Hook (Loral)

Overall, a fine presentation. However, this work needs to consider the problem slightly differently, given the NES. We need to know how well one can do with multicast-based filtering using specific numbers of groups like 16, 32, 48, 64, etc., corresponding to that expected to be supported by the NES. The study presented seemed aimed instead at picking the best approach, viewing multicast group as basically free.

Miller (AT&T)

It was not clear what contribution AT&T was going to make, and I was left with the impression that this group was not going to be able to deliver anything that was a real contribution to STOW '94 or '97.

Seeger (BBN)

I failed to find much of merit in this presentation for several reasons. Firstly, there seemed to be a rather poorly defined focus to this work. I could not really understand what they were tasked to do. Secondly, the recommendations on congestion control seemed like

rather weak opinions with no scientific or engineering backing. I even reviewed the provided Castineya and Escobar paper, and found no more content. For example, they glibly discard FIFO queuing as a viable approach while, at the present, Internet-working luminaries like Dave Clark of MIT are discovering in their research that FIFO queuing has some very nice properties for minimizing average delay in real-time communications settings. It is not clear who is right or wrong, but I didn't get a lot of confidence in the approach. They need data before throwing about raw opinions, especially in an area as difficult as congestion control.

The measurements of multicast join/leave behavior seemed to have a number of strange unexplained characteristics, such as some very short join times, and the traces being used to gather the data seem very weak. It would seem better to integrate this work with the performance evaluations being done by van Hook so we get the complete picture, rather than two independent efforts that both look at only part of the overall load. And van Hook seemed better positioned to carry this work forward.

Finally, the BBN work looking at protocol options for DSInet seemed extremely biased in favor of ST II which BBN developed. It is clear at the present time that IP multicast, the clear competitor, is enjoying far greater success in the commercial marketplace while it was listed as "experimental" versus ST II being listed as "available." I also felt that the amount of further development required for ST II implementations to actually support STOW '94 was not being presented accurately, with considerable confusion between what the protocol was supposed to do, and what implementations actually did. Personally, I feel that IP multicast would be a far better COTS-oriented solution for DSInet, and that it would be feasible to acquire COTS IP multicast routers for STOW '94 which would provide a longer-term investment than an ST II solution. However, given my involvement with IP multicast, this opinion is likely biased as well. It is good to hear that Houston Associates is going to consider this issue as an independent agent, but it does appear unlikely that they will be able to evaluate the options adequately within the short time frame of STOW '94 deployment. The only real benefit that ST II appears to provide over IP is some claimed ability to insulate ST II traffic from IP traffic. I think that this level of resource management can be accomplished by regulating regular IP (multicast) routers. It seems unfortunate to invest in non-standard approaches like ST II for one exercise when a COTS solution is available. We note further that IP multicast is being actively developed and refined by the Internet community, including router vendors, to satisfy the demands of real-time video and

audio. These demands largely subsume those for DIS, so DIS can automatically benefit from the commercial directions here if only it selects IP multicast.

I'd like to see a clearer definition of BBN's role in the networking component and more careful work in that role than what was evidenced by the presentation.

Lou Berger (BBN)

I got the impression that he was just providing an overview of DSInet, and could not very well identify or articulate the key issues on which we should focus. He did under tight questioning provide some good insight into the DSInet and its problems and limitations. I wish he could have developed his own candidate list, saving us some time in teasing it out of him.

Overall, it would seem beneficial to have someone capable taking an overall view of the project and studies, rather like this review panel as tried to do, but on an on-going basis. Only a few aspects of the current studies seem useful and there are many issues to consider and study, as pointed out above and by the other review members.

Comments on David Mills review

These are a few brief reactions to David Mills review input, which I thought was very good on the whole, but did have a few comments that either in my interpretation or his intent, I have some question on.

1. I feel the description of the "current technology" for network resource management is far too optimistic. Ferrari's work has never really been used and I don't think is really practical or useful for the problems of DIS. Clark's work is far more promising, but is still in the proposal and evaluation stage. I think it is very premature to "insist" that the commercial vendors incorporate this stuff into their products. They have ample motivation to address these problems for other commercial reasons. And I believe they are viewing that current state as premature to use as well. I would also mention that there is a long history of "resource management" mechanisms making situations worse rather than better, so just because we have the need, don't assume that current medicine will make you better.
2. I didn't fully understand the point on group membership and the recommendation. IP multicast provides joins within roundtrip times, which is the best one can do. The leave latency is controlled by a timeout, which defaults to 90 sec-

onds. One could clearly tighten that up substantially. I agree that there may be issues to explore here, but let's make sure we don't fund a whole research program when 90 percent of the benefit would be gained by just tightening this timeout on leaving a group. Perhaps more strongly, the latter parameter adjustment should be tried before we really decide we have a problem. It is not clear to me that there is one. I expect one does have to tolerate some extra level of network traffic because of joining and leaving not being instantaneous. However, I don't see this leading to spikes. We have to remember that the multicast techniques are just an optimization over the broadcast approach currently being used. That is, we are improving on what is there, not departing from some absolute optimum.

3. I didn't understand how to do header compression without changes to otherwise COTS routers. I was assuming, and understood, that the backbone net capacity was going to be improved. I am not convinced that header compression is going to be worth the pain in this light, although I really do not understand what Mills has in mind to use here. Perhaps this should be examined and clarified before the recommendation is taken too seriously.

Scaleability for STOW '94 & '97

Dale B. Henderson
Los Alamos National Laboratory

The scaleability project was very well articulated by Commander Dennis McBride, whom I had not the pleasure of meeting before. He seemed to appreciate both the opportunities and risks to his tasks of 10,000 and 100,000 entity exercises. The solutions which he proposed seem to be the likely ones to consider. That of which I am a little unsure is the overall context in which these solutions are to apply. While perhaps outside of or at the edge of the charge to the present panel, some of these solutions may be taken to compromise the basic genius of SIMNET. Does this matter, or is anything really being sacrificed? I do not know, but I feel that the questions should be addressed. Perhaps Colonel Reddy would have addressed them had he been able to attend. Perhaps you need another panel or a funded (contractor? FFRDC?) study.

To me the basic genius of SIMNET is the idea of a broadcast state message combined with the notion of dead reckoning to reduce the network traffic. The various new ideas which we were asked to consider are compromises on this theme. (I say this while acknowledging having been a member of the previous panel which worked at inventing some of the ideas under discussion.) A simulation employing a network with interest filtering or fidelity channels or interest registration is a lot more complicated and also less elegant than the original SIMNET. The panel had only slight briefings on and no internal discussion of the computational requirements of implementing these algorithms. I worry [about] the added loads, both at the intelligent gateways and at the entities may be underestimated. This should, however, be explored with the tool set discussed by Mr. van Hook.

To address whether or not the network traffic reducing methods compromise the mission of the exercise or simulation requires that we consider what this mission is. Clearly it is not just to put 10,000 or 100,000 entities on a network; that is too much like climbing a mountain because it exists. To justify any of the filtering schemes requires consideration

of the negative effects, if any, on the credibility and applicability of the simulation. With SIMNET the goals were training. Training value should be measurable through after-action reports, testing, and other measures. Now that we are talking seriously about other applications, we need to include a broader assessment of the very efficacy of the simulation, especially as we introduce some of these complicated mechanisms to enable scaling. By the way, some of the problems of using "time warp" schemes in parallel computer-based simulation are analogous.

The role and dominance of semi-automatic forces [SAF] raises similar philosophical questions. As I thought I understood the role of semiautomatic forces in SIMNET, they served to add data richness as an enhanced background to a fundamentally manned system of trainers. Now, with the larger numbers of total entities, they (SAF) are coming to predominate. This brings us to that which I know more about: very-many-instance simulations with much (or all) of the communication within the shared memory of a central supercomputer. Add a few manned instances of objects to one of these through external interfaces, and one comes to the same point from the other direction. But looking at it this way brings along questions (typically addressed in the non-DIS/DSI world) of significance, fidelity, applicability, and accreditation which deserve answers.

Such consideration should include the whole system. An intelligent gateway, for instance, is not just a passive service but is an integral part of the simulation engine. Interest registration and other ideas for tailoring the traffic move important decisions away from the entities which are best prepared to consider their effects.

Although we did not discuss them, there were suggestions of reductions in the arithmetic precision of field quantities, from 64 to only 16 bits. I occasionally worry that any simulation is an iterated remapping of data into its own space. Iterated remaps can be numerically unstable; this is the basis of numerical chaos. The conditions for the numerical instability of the Lanchester equations are known and demonstrated. For more complex simulations, we (I) can't guess, but I would worry that reducing numerical significance can not help.

I support the group-think conclusion of our final session that the 1994 goal of 10,000 entities operating in a secure environment is not possible. The arithmetic is simple and compelling. Thinking about the ARPA presentations, I suspect that they already knew this and were really looking for our confirmation.

While the government seemed to have its act together, the contractors left me with more of a mixed impression. I share the group's impression that AT&T seemed especially weak. Could it be that they are really just getting started?

I would rate van Hook's contractor briefing as best in that the tool set appears to be the right capability at the right time. The results should replace our collective guesses about traffic reduction with real information. That they can replay collected data from past exercises will make the results more convincing. The addition of some kind [of] a loop closure might also help: how would any of the network performance enhancing algorithms affect the simulation? Generally the network measures of performance need to be linked to the scenario bottom line.

Perhaps this tool could be extended to address some of the questions of validity and accreditation for simulations running under the various filtering and other network enhancement schemes. Fidelity channels, on-demand forwarding, and interest registration each raise questions which might be addressed. Even multicast grids should be proven safe before they are utilized in a big, expensive exercise.

The DSI understanding discussed by Lou Berger appears to be essential. He seems to have a more comprehensive knowledge than anyone else. Yet, the panel easily asked relevant questions to which the answers were not immediate or clear; the obvious (hopefully incorrect) conclusion is that nobody quite knows the answers. An example of this was our questions about the addressing through the T-20s and NESs—did we finally get it straight or did we just quit asking about it? I am unsure. The tool set should also help with these issues. Does it include the Network Encryption System? If not, the extension could be very important.

In this regard, why did we get such an impression that the DSI is more of a hostage to the NSA than in partnership with the NSA? I once raised an issue of key policy to the NSA and received a very positive and helpful response. (Furthermore, in comparison with STOW, the government had much less on the table in my case.) Having NSA participation at the review was probably a good step by Lynn Flynn in establishing a feeling of being in it together. Maybe you should have invited someone not directly involved from NSA, perhaps its Chief Scientist or the director of the SRC, to be part of our panel. There is always next time.

I want to repeat a specific remark from the discussion. I suggest that the Magic Carpet or Plan View devices ought to be reprogrammed so that the data are filtered locally (gen-

erally West Coast?) according to commands from the viewer (generally East Coast?) and only images be transmitted back. (Even if these stealth viewers are directly on the T-3, the data must be filtered somewhere.) As best I now understand, these viewing platforms are the most data-demanding entities, and they clearly do not feedback into the scenario's bottom line, so that none of my questions of significance would apply.

Issues of Scaleability for Stow

Duncan C. Miller
MIT Lincoln Laboratory

Overall Conclusions

The briefings and panel discussions focused primarily on the goal of a 10,000 entity exercise for STOW '94. Secondly, we considered the applicability of these approaches for STOW '97.

The consensus of the panel was that a 10,000 entity exercise a year from now is extremely risky if it must be conducted on a secure network. The principal constraint is imposed by the NES devices.

In a worst-case, "broadcast" approach, in which all exercise traffic is sent to every site, the T-1 tail circuits to the sites would be overloaded by a factor of about 10. Therefore, it is clear that an efficient filtering mechanism, probably based on multicast groups, will be necessary. The algorithms used must be clever enough to achieve a 90% traffic reduction on each tail circuit during the peak load periods of the exercise.

The principal point of risk is that the Motorola NES devices, which are currently the only approved end-to-end encryption devices for such applications, will not be able to support the necessary numbers of multicast groups. If clever workarounds cannot be found to the constraints imposed by the NES devices, and the requirement for a secure network cannot be relaxed, there seem to be only two fallback paths to pursue:

1. Reduce the magnitude of the exercise by a factor of three or four.
2. Procure multiple NES devices (and perhaps multiple T-1 lines) for each site. We suspect that as many as 10 encryption devices may need to be ganged together to support the needed number of multicast groups. This approach might require 130 NES devices for the 13 planned sites.

In my opinion, if effective multicast grouping algorithms can be developed, the per-site traffic loads for a 100,000 entity exercise in 1997 will not be substantially greater than for a 10,000 entity exercise. My basis for this opinion is that on a real battlefield, most units reach a saturation point in the number of other units they can consider. A company-size unit will continue to focus on approximately the same number of nearby company-, battalion-, and regimental-size units no matter how large the total battlespace becomes.

Although STOW '97 represented only a secondary focus of our discussions, we felt reasonably comfortable that ATM technology, wideband fiber-optic transmission media, and high-speed interfaces will be sufficiently well developed within the next three years that exercises involving an aggregate traffic load of 100,000 packets per second will be feasible. Whatever lessons are learned during STOW '94 can be readily applied to the implementation of STOW '97.

Comments on Specific Presentations

Dan van Hook (Loral)

Dan van Hook's presentation on simulation tools and algorithm evaluation was fascinating. It has been clear for a long time that filtering algorithms will be critical in determining how far distributed simulation can be scaled up, and many of us have speculated vigorously about what degree of traffic reduction can be achieved in realistic exercise scenarios. Dan's is the first data that has come out that can be used to begin to calibrate our expectations. It is exciting to see that traffic reductions by factors of three to five or more seem possible. The "On-Demand Filtering" plus "Fidelity Channel" plus "Dead Entity Server" combination of algorithms seems particularly promising.

It is essential that these tools be brought as quickly as possible to a state where various organizations can begin to use them and to explore hypothetical scenarios and algorithms. Then an objective dialog can begin regarding the tradeoffs involved in traffic control algorithms, network topology, and exercise scenario design. Probably the most crucial issue that needs to be understood is that of the dynamics of multicast subscription management under alternative assumptions. It may well be that the best approach in this area has not yet been conceived, so it is important to get additional creative people thinking about it.

Bill Miller (AT&T)

This presentation was less impressive. We heard very little by way of innovative approaches for saving communications bandwidth. It was suggested that changes in the DIS protocols could produce more compact PDUs by transmitting incremental changes in vehicle state rather than the vehicle state and its derivatives. This is undoubtedly true, but it would require a fundamental modification to the approach now codified in IEEE 1278-1993. If this idea were implemented, it would be better to intervene at a local area network gateway to retransmit the more compact PDUs over the Wide Area Net and reconstitute them in standard form at the receiving LANs.

It was also suggested that derivatives could be inferred at a receiving node rather than transmitted explicitly. This proposal, while technically possible, reveals a serious misunderstanding of the dead reckoning concept.

When a pilot opens his throttle and initiates an acceleration of his aircraft, his simulator knows immediately that the vehicle's time derivatives have changed. By transmitting the new state vector derivatives to other entities immediately, they can initiate an extrapolation of the future state of his aircraft that will remain within discrepancy threshold tolerances for a significantly longer time than a simple, constant velocity extrapolation. Under the proposed alternative scheme, after a sufficient period of time, the receiving nodes could infer and begin to make use of the aircraft's acceleration, but not until many unnecessary PDUs had been transmitted. This scheme is inferior to a higher-order dead reckoning extrapolation approach in terms of information transmission under almost all circumstances.

The other AT&T responsibility described to us was the development of detailed specifications for the communications interface gear. We heard nothing to suggest that this was more than a straightforward systems engineering effort that could be successfully undertaken by any number of SETA organizations. It was not clear why the capabilities of AT&T Bell Labs were required.

Josh Seeger and Lou Berger (BBN)

Josh Seeger presented a summary of the network modeling being carried out in conjunction with the LADS application-level modeling. This work appeared to be well integrated and quite useful for predicting performance bottlenecks for STOW '94. The most critical technical issue in this area is the dynamic multicast group subscription/description rates that result from various grouping algorithms. The preliminary data presented was

intriguing but rough. There seemed to be some anomalies, especially in the numbers of entities that maintained a subscription to a particular multicast group for less than five seconds. My intuition tells me that the numbers we were shown were much higher than could be accounted for by a few fast-moving aircraft and by entities grazing each other's areas of interest. The algorithms being simulated should be checked for dynamic stability.

Lou Berger's presentation on the state of the Defense Simulation Internet turned out to be the focus of most of the "bad news" reflected in the consensus summarized above. Here was where we discussed security issues and the serious limitations of the NES devices. Here was also where the "religious issues" of ST vs. IP multicasting were debated. My overall perspective on these issues is that we must remember that we are dealing with some near-term problems for STOW '94 that we expect to be resolved by hardware and software developments over the next two or three years. ARPA should therefore resist pouring more resources than are necessary into short-term solutions to these problems. Other factors are at play, both in terms of market forces for commercial communications and in DoD requirements for secure communications, that will dominate the specific requirements of distributed simulation in general and STOW in particular.

The best strategy seems to be to do whatever can be done with the existing hardware and software for STOW '94 while closely monitoring developments in other technical arenas. If, as Dave Cheriton argues, other forces are working to make IP multicasting an industry standard, that's fine. The DIS community should cooperate in this development, making sure that our unique needs for highly dynamic multicast groups are understood and addressed by the larger community. (A modest level of support for ST-II multicasting may be the best guarantee that the larger community will pay attention to these needs.) Similarly, we must make sure that NSA understands our specific needs for higher-bandwidth end-to-end encryption devices that can handle large numbers of small packets being routed to multiple destinations.

Finally, I must add that the prevailing view that ATM will resolve all these problems within the next two or three years leaves me somewhat apprehensive. I don't doubt that they will be resolved, but I suspect the solutions will not be handed to us on a platter. We will have to speak up forcefully to make sure our needs are being adequately considered in a marketplace that is being driven by other forces.

Summary of STOW Panel Review

David L. Mills
University of Delaware

Overall Conclusions

On August 19 and 20, 1993, a Synthetic Theater of War (STOW) program review was held at the Institute of Defense Analysis in Alexandria, Virginia. The review panel, of which I was a member, met to hear technical briefings on the STOW program, specifically on the demonstrations to be held in FY94 and another in FY97. Briefings were presented by Loral Advanced Distributed Simulation (LADS), American Telephone and Telegraph (AT&T), and Bolt, Beranek Newman (BBN). The following are my impressions on these and related issues.

The BBN briefing emphasized network-layer modeling while the LADS briefing emphasized application-layer modeling. These briefings included considerable detail and convincing arguments. The LADS briefing, in particular, demonstrated well-founded design and implementation strategies and understanding of the issues. The AT&T briefing described candidate gateway designs, compression algorithms, and dead-reckoning algorithm improvements, but was generally less effective than the others. The BBN briefing raised painful issues of near-term design and implementation shortfalls which raise real doubt that the planned level of 10,000 entities at 15 sites is sustainable in STOW '94.

My overall assessment on the technical direction of the STOW program is mixed in view of the briefings and related background information. On one hand, it is apparent that the networking technology issues mentioned in previous review panels are being addressed and studied on their own merit. The LADS briefing demonstrated that on-demand forwarding, fidelity channels, and dead-entity servers have the potential to reduce the network traffic by factors of between 3 and 5, at least in certain topologies and traffic scenarios. However, in themselves these schemes do not appear capable of traffic reductions by factors of 10 or more, as may be required by STOW configurations. The AT&T

briefing, while considering some issues not evident in the LADS briefing, did not bring comfort that sufficient coding gain, delta PDUs, or dead-reckoning improvements would make up the shortfall.

I don't think it reasonable in a report such as this to offer specific guidance on technical issues, such as how many NES are required and what available options should be pursued. I can, however, comment on the design philosophy and likely outcome, should certain design approaches be pursued. In anticipation of the commentary to follow, it is clear that no magic bullet is likely to do the job of reducing backbone and tail traffic by a factor of at least 10, as required by a quick analysis of the projected topology of STOW '94. Rather, it is likely that this goal will be achieved, if it is achieved at all, by a systematic assault on minor inefficiencies, each of which may contribute a small fraction of the overall budget.

Network Model and Assumptions

First, I assume there is no avoiding encryption for end-end data transmitted over the WAN. However, if commercial WAN providers are to play and promiscuous broadcast is to be avoided, there has to be some way for the (classified) simulation community to inform the (unclassified) network providers about the protocol data unit (PDU) address and service class, in order to route the traffic along multicast trees and perform whatever traffic mitigation is required. Following the SDNS model proposed by NSA, which I would assume is politically correct, the only way to do that is in the unprotected PDU header which can be IP, ISO, or even ST-II. Furthermore, I assume the SDNS PDU is manufactured at the point of entry from the LAN to the WAN and that the PDU is protected by a message digest, [and] so cannot be modified en route. Therefore, the only thing the WAN can do is forward, replicate, or discard the PDU. Of course, PDUs can be aggregated by encapsulation as long as the original format is restored upon exit. Note that there may be some argument over the degree to which the application security is compromised by the unprotected PDU header, which admits of traffic analysis.

I assume most of the lessons learned in current practice and experiment have been internalized in the network design. The traffic to be handled at the entrance and exit of the WAN is estimated at about 800 2,000-bit packets per second, which is not high by current standards. However, the encryption devices may not support this rate, which is surprising when compared with similar commercial devices which operate at much higher speeds. I believe that admission controls will be necessary to prevent network collapse by denying service or group formation, that leaky-bucket traffic grooming will be necessary to reduce

the occurrence of timeouts due to PDU bunching and timeouts, while multiple-class, round-robin scheduling will be necessary to mitigate delay bounds on the basis of absolute or statistical guarantees or no guarantee at all.

We agreed at the meeting that these features of this type will be necessary in a fully developed DSI, whether or not the network is dedicated to the simulation mission. Current technology based on the work of Ferrari, Clark, and others suggests a design based on these principles is practical and quite likely would be an important factor in ensuring success of the STOW '94 simulation. Of particular current interest is the guaranteed/stochastic/best-effort delay-bound model which maps nicely to the types of data used in STOW simulations.

I believe it necessary that class-oriented service policies be implemented in the network. These policies provide for a set of priority classes, with each class serviced by deadline and/or round-robin. This may be the single most important feature allowing guaranteed deadlines for such things as impact reports, stochastic deadlines for position reports, and best effort for other traffic. These policies also provide the framework for an engineered scheme to discard excess traffic upon overload.

Recommendation: Make sure the vendors are aware of admission control, leaky-bucket and class-oriented service principles, and insist they be incorporated in the routers, encryption devices, and end systems.

Simulator Simulation

In my opinion the most important development to come from the briefings is a simulator for the simulation system itself. This is a program which emulates the network operations induced by a functioning discrete-time simulation. It embodies the PDU compression, multicasting, and forwarding algorithms of the WAN, and can be implemented either as a synthetic simulator driven by software scripts or as an intelligent local router joining Ethernet segments, for example.

With this tool it is possible to determine the effectiveness of various compression schemes, multicast protocols, and class-oriented scheduling. If my conclusion [is] that there is no magic bullet and that substantial reductions in traffic level are possible only by a series of incremental improvements, this tool will be critical to the mission success.

Recommendation: Give the simulator-simulator area [an] increased emphasis in tasking statements and support the construction of a suitable intelligent local router based

on an ordinary workstation and supporting toolkit. A possible point of departure could be the Sun workstation routers developed for the DARTnet research network for ARPA.

SAFOR Servers

One of the things that caught my interest during the meeting is the observation that SAFOR traffic accounts for 90% of the total. If so, the effort to reduce traffic levels would be most effective if concentrated in this area. Presumably, SAFOR traffic is generated by a suite of replicated algorithms in each entity and orchestrated by an operator. This suggests an approach in which a state machine representing each operator is instantiated in all entities. Its operations could be managed by a low-rate protocol that avoids the transmission of entity-state PDUs on the WAN. This is the same kind of idea that leads to the dead-entity server.

There may be applications of the replicated state-machine approach other than dead entities and SAFOR, such as weather and terrain changes, craters, etc. I believe that the benefits of such techniques as on-demand forwarding, fidelity channels, and compression may well have approached their limits and the replicated state-machine approach may be the single most useful tool remaining.

Recommendation: Pursue the replicated state machine approach with SAFOR, hunt for other applications that yield to this approach, [and] generate a suitable model, protocol, and PDU encoding to generalize the application.

On-Demand Forwarding

I have some concern that on-demand forwarding, which can be described as receiver-directed multicasting, will result in a significant reduction in traffic for the STOW '94 simulation. My reasoning for this is as follows. The total PDU flux emitted by all entities in the STOW '94 simulation is expected in the order of 30 Mbps with 10,000 entities at 15 sites. For the purposes of discussion, assume the emissions of all sites are substantially equal and that coding and aggregation gains permit the use of a 1.544 Mbps T1 tail circuit. It follows that this pipe will probably be fully loaded.

In the topology anticipated for the STOW '94 simulation, the sites will not be richly connected and will probably take the form of a linear network or ring. On the assumption that the backbone will probably be assembled from 1.544-Mbps T1 circuits, there will be traffic bottlenecks where the tails join the backbone. It does not seem likely that every single member of a group will be at a particular site, so there will be at least some degree of

PDU replication in the backbone. This will certainly limit the effectiveness of the spanning-tree pruning feature of on-demand forwarding. Since the backbone is not likely to be richly connected, I conclude a T1 backbone is not likely to support the full STOW '94 mission no matter how the multicast groups are formed or how the spanning trees are pruned.

Recommendation: Increase the capacity of the backbone network either by increased fabric speed or richer connectivity. Make sure admission controls are in place and that class-oriented discard policies are in effect.

Group Membership

Now, assume that some magic multicasting scheme now in use or yet to be discovered is available and does the job; that is, a particular traffic flow can be supported by realizable bucket parameters, routing tables, replication points, and class policies. When an entity moves from one multicast group to another, a group-management protocol must react quickly. To avoid dropouts as the tables are being adjusted, the entity will probably belong to both groups for some interval spanning the reaction time of the protocol. This is likely to induce a traffic spike in the network for at least that period, as well as complicate the hand-over protocol.

Unless the hand-over protocol is incorporated directly in the design of the multicast routing algorithm, there will be additional delay as the routing tables are adjusted throughout the network. Current commercial networking practice is to use a standard link-state routing protocol like ISO IS-IS for the network routing fabric and an encapsulation overlay with a special-purpose multicast routing protocol to determine routing and replication points. In designs known to me, the group-management functions are handled by a dedicated protocol layered on the routing substrate. The required speed of response for a real-time system would seem to require that the link-state routing, multicast routing/replication, and group-management functions must be implemented as a functional unit. While protocols such as ST-II and IP multicast address one or more of these issues, they do not address them all in an integrated way.

Work is needed toward a comprehensive, real-time, multicast group management; admission control; and routing management protocol. While it is unlikely that the development and implementation of such a protocol could be completed by the time of STOW '94, it will be necessary in the long run and should be pursued.

Data/Header Compression

There has been some attention to the issue of data compression, either through the use of variable-length encoding or motion compensations as described by AT&T. If the average PDU size is 2,000 bits, which seems to me on the high side, it would make sense to be somewhat more aggressive in this area; however, the issue seems not to have stirred the imagination of the standards community at best. However, my commentary here is primarily aimed at the WAN and its routers which see the PDU as a protected (encrypted) data area and an unprotected header.

Since the encryption operation effectively randomizes the data area, it is unlikely that any compression algorithm can reduce its size; however, this leaves open the issue of header compression which uses an associative table to map the PDU header upon entry to a minimal header for transmission through the network, then restores it upon exit. The associative table must be constructed either at flow setup or on the fly, with timeout as in other protocols such as ARP. From experience in the Internet, these techniques could reduce the backbone and tail traffic by 5 to 10%.

Recommendation: This is a relatively easy thing to do and can be incorporated at the network or physical layer of the network software. It is an ideal candidate to be explored using the simulator mentioned above.

Scaleability Issues for STOW

Stuart D. Cheshire
Stanford University

Comparison of ODF and Geographical Grid Multicasting

I saw something in Dan van Hook's talk that no one else seemed to stress, perhaps even Dan van Hook himself (but that's okay—scholars find more in Shakespeare than he was probably ever consciously aware of).

I heard several times that “On-Demand Forwarding” would “probably” reduce traffic better than geographical multicast groups would, but it was not stated very forcefully and it did not seem to evoke any conclusive agreement. I would say that ODF would *definitely* reduce traffic better than geographical multicast groups would, almost by definition, and I was not in the least surprised to see Dan van Hook's graph showing that ODF with 170 multicast groups achieved better traffic reduction than a geographical grid with 678 groups.

If I understood ODF correctly, then it allows each receiver to “determine which entities they need to receive state from.” Surely, if each receiver is receiving precisely the packets it needs and no others, then this is the optimum solution? However, I also feel that this could not actually be made to work because each receiver does not have perfect knowledge with which to make the decision, nor unlimited MIPS with which to do the calculation, nor unlimited bandwidth for control information. I view it more as a kind of DIS “Turing Machine”—something one would never build but a model against which to compare other proposals.

One (undesirable) way we might actually achieve this limit could be with an omnipotent gateway, with full knowledge of the DIS application semantics and full models of the internal states of all the participants. It could arbitrarily compress, alter, combine, filter, and discard PDUs to send precisely the information needed by the participants, and nothing more.

Multicast is a powerful general purpose mechanism, but how do we know how well it will work unless we already have some notion of how well is "well" and how well is "not well"? If Dan van Hook can show with his simulator that he can achieve within 10% of our idealized ODF target with a reasonable number of multicast groups using off-the-shelf IP multicast hardware, then I would be prepared to call that a win. If he cannot, then perhaps we will have to consider building our "smart" gateways, distasteful as that may be. If even our idealized ODF model does not produce the required traffic reduction, then we have to rethink the entire approach. In any eventuality, I think it is valuable for us to know which of the above three scenarios we are facing.

Encryption Devices

Discussion was between using Motorola NES (too slow), or alternatively abandoning encryption altogether for STOW '94 (not acceptable to NSA).

A third possibility not mentioned was that of keeping encryption, but lowering the requirement. The NES costs about \$15,000 and has a few hundred kbps throughput at best. Vendors at Interop were hawking boxes which can DES-encrypt full rate ethernet frames. One company I saw was Semaphore who had ethernet to T1 gateways for \$7,000. Its encryption engine can do up to 6,000 frames per second, and up to 9Mbps—easily enough to keep the T1 link completely saturated with data.

Perhaps as proof of concept, we could achieve STOW '94 using DES, and leave that component of the system upgradeable to a Level 1 NSA security device when it becomes available. This could demonstrate that the goals of the STOW project are in fact achievable, and that it is the current quality of encryption devices which are the weakest link rather than any other aspect of the project.

ATM Technology Hopes For STOW '97

I agree with David's and Duncan's reservations about the prevailing blind faith in ATM. I do not think that ATM is going to come riding up like a knight in shining armor and rescue the networking industry. I'm not old enough to have seen a great deal of history in the computer industry, but I can mention three "data points" that I am familiar with (Table A-1, "Apple, Ethernet, and ISDN," on page 25).

You can make your own decision about whether you think ATM falls into the Apple camp, the Xerox PARC camp or the AT&T camp, but personally I'm not expecting to see anything real within a three-year time scale, and I wouldn't bet the farm on it.

Table A-1. Apple, Ethernet, and ISDN

1. **Apple.** Computer company.
 - 1990, Apple IIFX costs about \$8,000.
 - 1993, Apple Newton gives twice the MIPS in a package that weighs under a pound, fits in the palm of your hand, runs for two weeks on alkaline AAA batteries and costs about \$700.
 - Time scale: 3 years.
2. **Ethernet.** Borderline computing/communication technology.
 - 1977, Xerox Ethernet runs at 10Mbps.
 - 1993, first faltering steps towards a 100Mbps Ethernet standard.
 - Time scale: 16 years.
3. **ISDN.** Telecommunications technology.
 - 1963, ISDN announced.
 - 1993, still not widely available even in the USA, never mind in less-developed countries in the world.
 - Time scale: 30 years.

Dead Entity Server

Useful as this is, I would resist too much development in this direction. For DIS to be a success, it *has* to solve the dynamic terrain problem, and when this is accomplished, the facilities promised by the dead entity service will be subsumed into the more general mechanism.

If the dead entity server has to be built, it should be regarded in the same light as the NES boxes, i.e., as a temporary solution for STOW '94, to be discarded for STOW '97.

Comments on Individual Presentations

Bill Miller (AT&T)

His lack of understanding of some basic principles worried me. His proposal to send only incremental changes exposes the fact that he had not even considered that a network of this scale loses packets. If it did not lose a single packet, then that would be evidence of gross over-engineering and waste of money.

He was obviously too used to thinking in terms of telephone calls where you have a very small bandwidth, only two endpoints, and charge a dollar a minute to pay for an over-engineered network. I do not know how much AT&T would charge to set up a 100-way conference call, or even if they are capable of doing it. I am fairly certain that they could not set up a 1,000-way or 10,000-way conference call.

STOW seems so far outside the gamut of AT&T's experience that Bill Miller was not even able to answer questions with plausible sounding guesses.

APPENDIX B
BIOGRAPHICAL SKETCHES OF PANEL MEMBERS

David R. Cheriton

David Cheriton is a Professor of Computer Science at Stanford University, Palo Alto, California. Professor Cheriton's primary focus of research has been on high performance distributed systems with considerable emphasis on real-time applications. Additional research foci include issues in high-speed communication, multiprocessor architecture, and distributed database support. His research group developed a distributed operating system called *V*. The central idea in the *V* system is to provide a fast, reliable, and secure real-time distributed operating system kernel on which real-time applications, distributed databases, and ordinary software development/document production can be implemented with maximal performance and minimal implementation difficulty. The main focus has been on communication performance.

Professor Cheriton's research has been largely sponsored by the Advanced Research Projects Agency (ARPA), and he plays an active role in the directions of this community. In particular, he was a member of the ARPA Gigabit Working Group, a task force chairman for the Distributed Systems Architecture Board, and a member of the End-to-End Protocols Task Force under the Internet Advisory Board. In this role, he and his group developed the IP multicast extension (IEEE Request for Comments (RFCs) 966, 988) which is now a draft standard. He also developed the VMTP protocol (IEEE RFC 1045) which is considered a strong candidate as a next-generation transport protocol for the Internet, addressing issues of real time, security, and fault-tolerance. Finally, Professor Cheriton has been an Associate Editor of the *ACM Transactions on Computer Systems* and the *Distributed Computing* journal, and a referee for all the major publications in the field. Professor Cheriton is a member of the Institute of Electrical and Electronics Engineers (IEEE), Inc., and the Association for Computing Machinery (ACM).

Professor Cheriton received his Ph.D. in computer science from the University of Waterloo, Canada, in 1978, with his thesis work growing out of the development of the "Thoth" portable real-time operating system.¹ This system was used in a number of real-time applications and has served as a basis for several other commercial real-time systems. He subsequently spent three years at the University of British Columbia in Canada. He has been at Stanford University since 1981.

¹ D.R. Cheriton, M.A. Malcolm, L.S. Melen, G.R. Sager, "Thoth: A Portable Real-Time Operating System," *Communications of the ACM*, 22/2 (February 1979), pp. 105-115.

Based on this expertise, Professor Cheriton has worked as a consultant for a wide range of companies, including ESL/TRW, Rockwell International, IBM, Silicon Graphics, Digital Equipment Corporation, Institute for Defense Analyses, Texas Instruments, SRI, and Dynamics Research Corporation. He regularly gives talks and lectures on distributed systems research at the major universities and research laboratories throughout the United States.

Dale Henderson

Dale Henderson is employed at Los Alamos National Laboratory, Los Alamos, New Mexico. He joined Los Alamos in 1966 upon completing his Ph.D. at Cornell University. After four years in experimental plasma physics, he moved to the (then new) laser induced fusion program. In 1975 Dr. Henderson became the leader of the theory group in the laser induced fusion project. In 1979 he moved to the project management of computer code development for the nuclear weapons design program. Soon after President Reagan's "Star Wars" speech, Dr. Henderson recognized the need for a flexible comprehensive simulation model and began the DETEC (Defense Technology Evaluation Code) project at Los Alamos. DETEC was adopted as the major software vehicle at the Strategic Defense Initiative Organization's National Test Bed (NTB) in 1988. Having served the NTB Joint Program Office from its beginning, Dr. Henderson undertook a FY 1990 assignment to the SDIO as Chief Scientist of the NTB.

Duncan C. Miller

Duncan Miller is currently a Group Leader at the Massachusetts Institute of Technology (MIT) Lincoln Laboratory in Lexington, Massachusetts. Dr. Miller has been involved with man-in-the-loop simulation since his master's thesis at MIT's Man-Machine Systems Laboratory in 1964-65. In 1983, while at Bolt, Beranek and Newman (BBN), Inc., he formed and led the group that developed the original Simulator Network (SIMNET) architecture and software. He has been actively involved in the extension of distributed interactive simulation systems and concepts since that time. He is a member of the Steering Committee for Distributed Interactive Simulation Standards (IEEE Std 1278-1993) and the Defense Science Board Task Force on Simulation, Readiness, and Prototyping. He previously served on the Naval Research Advisory Committee Panel on the Impact of Advancing Technology on Exercise Reconstruction and Data Collection.

Dr. Miller received his bachelor's (1965), master's (1965), engineer's (1967) and doctorate (1969) degrees in mechanical engineering from MIT. His principal areas of study

included control theory, human operator performance modeling, human factors, and perceptual psychology.

David L. Mills

David Mills is a Professor of Electrical Engineering at the University of Delaware, Dover. He currently leads projects in high-speed networks and internetworking research sponsored by ARPA and National Science Foundation (NSF). His research activities have been concentrated in the areas of network architecture, protocol engineering, and experimental studies using the Internet system. He is a member of the Internet Research Steering Group and formerly chaired the Internet Architecture Task Force. He is also an advisor to the NSF and was principal architect of the NSFNET Phase-I Backbone network.

Before joining the Delaware faculty in 1986, Dr. Mills was a Director (Networks) at M/A-COM Government Systems Division (Linkabit), and led ARPA-sponsored research and development projects in packet-switching network architectures and application protocols. Previously, he was a Senior Research Scientist at COMSAT Laboratories where he worked in the areas of packet-switching satellite and internetworking technologies, and was an assistant professor of computer science at the University of Maryland, Takoma, Maryland, where he worked on several research projects in distributed computer networks and operating systems.

Dr. Mills earned a doctorate in computer and communication sciences at the University of Michigan, Ann Arbor, Maryland, in 1971, and has held postdoctoral positions at the University of Edinburgh, Scotland, and the U.S. Defense Communications Agency [now the Defense Informations Systems Agency]. He has published and lectured extensively on data communications, computer networks, and operating systems, and has been a consultant to a number of corporations and government agencies. He is a member of Sigma Xi, ACM, and the IEEE Computer Society.

Stuart D. Cheshire

Stuart Cheshire is a fourth-year Ph. D. student at Stanford University, Stanford, California, studying under David Cheriton in networks and distributed systems. Mr. Cheshire's primary focus of research has been in protocols for distributed interactive simulation. The results of some of this work can be seen publicly in the form of the popular Internet game "Bolo," a networked multi-user, real-time tank battle simulation.

His research has been sponsored in part by the U.S. Army Research Institute, Alexandria, Virginia, for its contribution to Army research into team training methodologies. Mr. Cheshire received his First Class Honours B.A. in Computer Science from Sidney Sussex College, Cambridge, England, in 1989, and was awarded an honorary M.A. in July 1993. He has written for numerous computer magazines, and has worked in the computer industry with Madge Networks, a company manufacturing IBM-compatible Token Ring network hardware and software.

Mr. Cheshire is the author of the Stanford print accounting package which authenticates and bills students for network printing from their in-room connections, and the Stanford Software Librarian, a fault-tolerant distributed software license management system.

APPENDIX C
PRESENTATIONS

Simulation Tools for Developing and Evaluating Networks and Algorithms in Support of STOW 94 C-3
Daniel J. Van Hook, Loral Advanced Distributed Simulation


Intelligent Gateway for Defense Simulation Internet C-37
William D. Miller, AT&T Bell Laboratories

STOW 94 Network Requirements & Evaluation of Solutions C-73
Joshua Seeger, BBN Systems & Technologies

Today's DSI C-89
Lou Berger, BBN Systems and Technology

Summary: Key Technical Issues for STOW 94 C-103

Exploring the Limits of DS Scalability...



Simulation Tools for Developing and Evaluating Networks and Algorithms in Support of STOW 94

Presented for:
Scalability Peer Review
19-20 August, 1993


Daniel J. Van Hook
dvanhook@camb-lads.loral.com

LORAL
Advanced Distributed
Simulation



BBN Systems and Technologies

Exploring the Limits of DIS Scalability...



Topics for Today

- Project Overview
- Challenges for STOW
- Technical Approach
- Current Status
- Deliverables
- Conclusions

LORAL
Advanced Distributed
Simulation



BBN Systems and Technologies



Project 10-4 Overview


- ▼ **Phase 1 — complete (7/92 to 11/92)**
 - Develop Scenario and Network Simulation Tool
 - Develop Algorithm Demonstrator
 - Demonstrate 10^4 Entity Scenario

- ▲ **Phase 2 — on-going (1/93 to 10/93)**
 - Refine Tool Set
 - Develop and Research Algorithms
 - User Documentation
 - Added to STOW Team (March, 1993)

- ▲ **STOW Scenario — on-going (5/93 to 11/93)**
 - Simulate Projected STOW 94 Scenario
 - Analyze Scenario and Provide Feedback on Scenario, Network and Algorithm Performance



Exploring the Limits of DIS Scalability...



BBN/LADS Team


- **BBN Systems and Technologies**
 - Network tools and modeling
 - Network algorithms and architectures
 - Debbie Deutsch, Julio Escobar, Isidro Castineyra, Walter Milliken, Josh Seeger, Shawn Smith

- **Loral Advanced Distributed Simulation**
 - Scenario tools and modeling
 - Scaling algorithms and DIS Systems
 - Maureen Saffi, Richard Schaffer, Dan Van Hook, Rob Vrablik, Deborah Wilbert, Chris Williams



BBN Systems and Technologies

Exploring the Limits of DIS Scalability...



STOW Challenges

- **Large number of entities**
 - High Packet and Bit Rates
- **Large region of interest simulators**
 - PVDs, High Altitude Sensors, e.g., AWACS, Satellites
 - Effectiveness of geographic filtering limited due to need for data from a large area
- **Network Limitations (DSI)**
 - Backbone: Links and Switches
 - Tail Circuits
 - Host processing
 - Security Devices (NES)
 - DIS has Tight Delay Requirements
 - Shifting Traffic Patterns are Heavily Scenario Dependent

LORAL
Advanced Distributed
Simulation



BBN Systems and Technologies

Exploring the Limits of DIS Scalability...



STOW Challenges (cont.)

- **After Action Review**
 - Distributed Logging and Replay
- **Integrating Constructive Simulations**
 - Disparate Frame Rates
 - Aggregate to Entity Mapping and Vice Versa
- **Integrating Live Vehicles**
 - Latencies of Range Instrumentation
 - Bandwidth Limitations of Range Instrumentation
- **Coordination and Interaction**
 - Exercise initiation and control
- **Radio Traffic (Not Required for STOW94)**
 - Data rates: 64kbits/sec and 20 packets/sec per transmitter
 - Long propagation distance implies many potential receivers

LORAL
Advanced Distributed
Simulation



BBN Systems and Technologies



Technical Approach

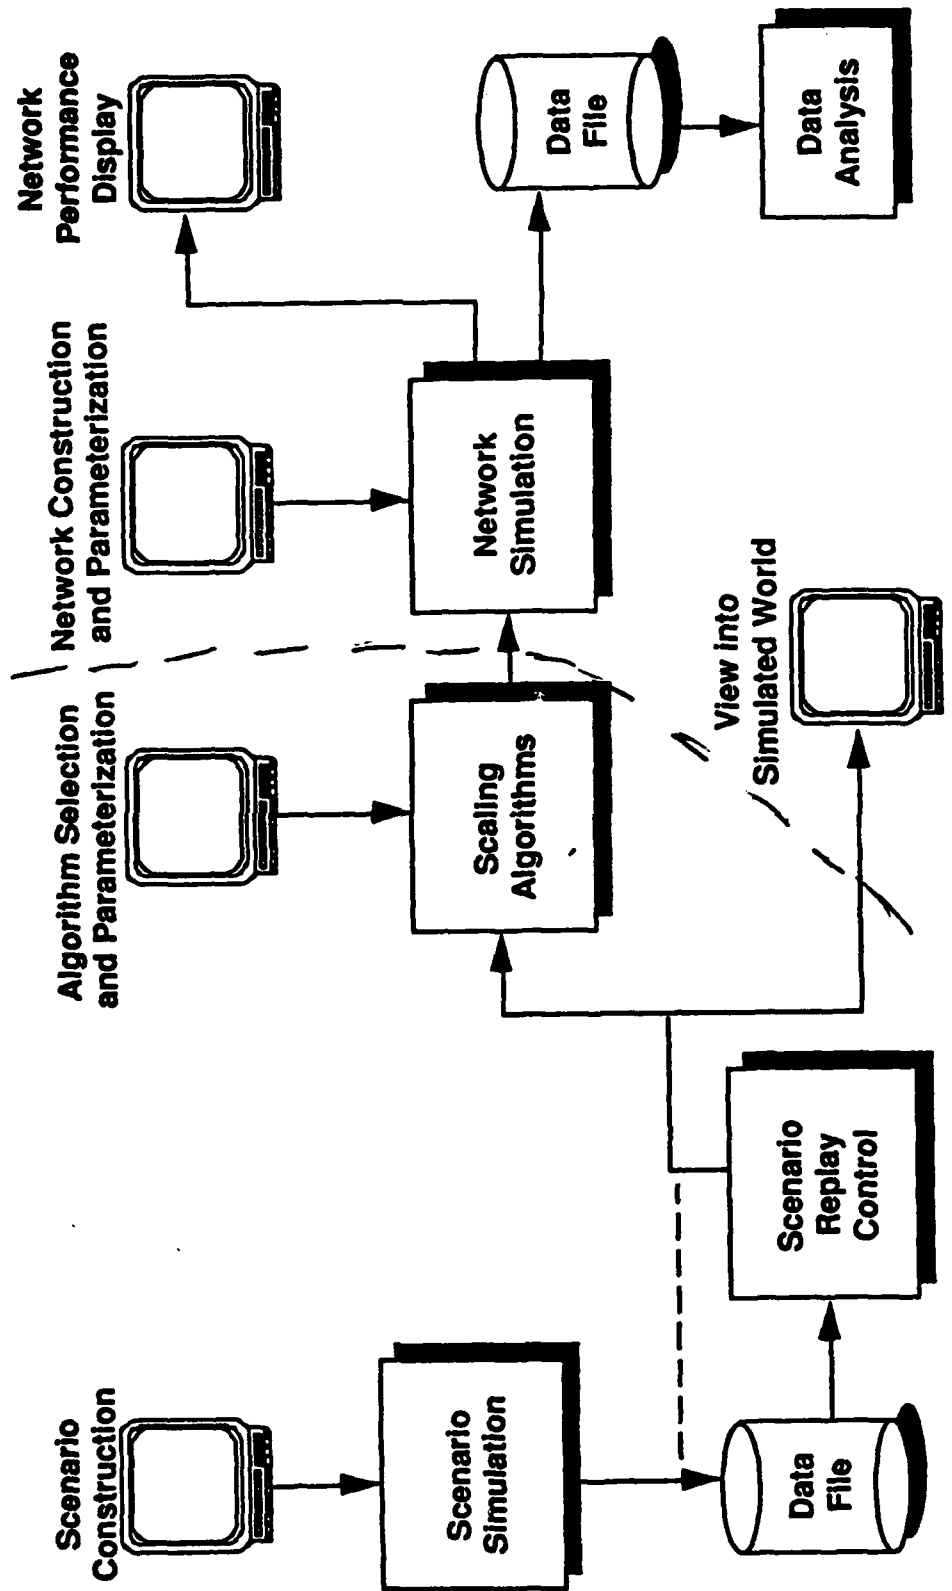
- **Key Concept: "A Simulation of the Simulation"**
 - simulate before you build
- **The Scalability Toolset**
 - Network Topologies and Parameters
 - Network Algorithms for Routing, Congestion Control, Resource Reservation
 - Scaling Algorithms for Bit/Packet Rate Reduction
 - Simulation Scenarios
- **Scaling and Network Algorithms**
 - Techniques to Facilitate "Scaling Up" of DIS Exercises on Existing and Projected System Underpinnings
- **Scenario Analysis**
 - Analyze Performance of Scenarios, Networks, Algorithms



Exploring the Limits of DIS Scalability...




Toolset Architecture



LORAL
Advanced Distributed Simulation



BBN Systems and Technologies



Scaling Techniques

- **No Magic Bullet: Combination of Techniques Required**
- **Where Do the Algorithms Reside?**
 - Network, Hosts, Intelligent Gateway (IG)
 - Flexibility Permits Evolution as Technology Evolves
- **Filtering: Transmit Only the Required Information**
 - Geography: Within Range
 - Spectrum: Within Spectral Band
 - Entity Type and Role: Tactical Significance
- **Coding/Compression: Bit/Package Rate Reduction**
 - Exploit DIS Application Characteristics
 - Application Independent: Bit Encoding



Exploring the Limits of DIS Scalability...



Scaling Techniques (Cont.)

- **Overload Management/Graceful Degradation**
 - Application Level: Relax Fidelity Requirements in Host or IG
 - Network Level: Congestion Control Policies in Switches, Routers
- **Architecture**
 - Distribute Simulation, e.g, Distributed SAF, Dead Entity Service
- **DIS Protocol Changes**
 - Timeout Value
 - Representation of Distributed Information
 - Efficient Data Representation

LORAL
Advanced Distributed
Simulation



BBN Systems and Technologies



Scaling Techniques (Cont.)

- **Network**
 - Choice of Topology and Technologies
 - Routing Algorithms
 - Multicast Support
 - Resource Reservation
 - Congestion Control Policies
- **Scenario Tailoring**
 - Constrain the Scenario to Fit the Available System Resources
 - Limit Unit Interactions and Area Occupied





Geographic Filtering

- **Idea**
 - Receive data only from the *region of interest*, i.e., entities within visual or sensor range
 - Permits construction of simulators that can operate in exercises of any scale - performance dependent only on local entity density
- **Approach**
 - Use multicast capabilities of the network to deliver only necessary data
- **Issues**
 - Effectiveness reduced as entities come into proximity
 - Effectiveness not constant over the course of an exercise
 - Multicast groups and change rate are limited
 - Region of interest must extend beyond sensor range due to transmission timeout, relative velocities, group change latencies




Exploring the Limits of DIS Scalability...


Project
107



Time = t



Time = t + T



- A = Viewing Radius
- R = Region of Interest Radius
- v = Closing Velocity
- T = Transmit Timeout
- A - R = v x T
- For v = 300m/s and T = 5 Seconds A - R = 1500m



BBN Systems and Technologies



Multicast Grid Geographic Filtering

- **Idea**
 - Receive data from a terrain region by subscribing to its associated multicast groups
- **Approach**
 - Divide terrain into grid cells
 - Multicast group associated with each grid cell
 - State updates sent to group for cell each entity lies in
 - Subscribe to groups for surrounding cells to receive traffic from surrounding entities
- **Issues**
 - Cell size tradeoff
 - Small for efficient filtering
 - Large to reduce number of groups and group change rate

Exploring the Limits of DIS Scalability...



Multicast Grid Geographic Filtering (Cont.)


- **Extensions**
 - Dynamically assign group addresses to cells to conserve addresses
 - Static non-uniform grid: dependent on entity density and network topology
 - Dynamic non-uniform grid: adjust grid morphology at run-time based upon entity density and network load

LORAL
Advanced Distributed
Simulation

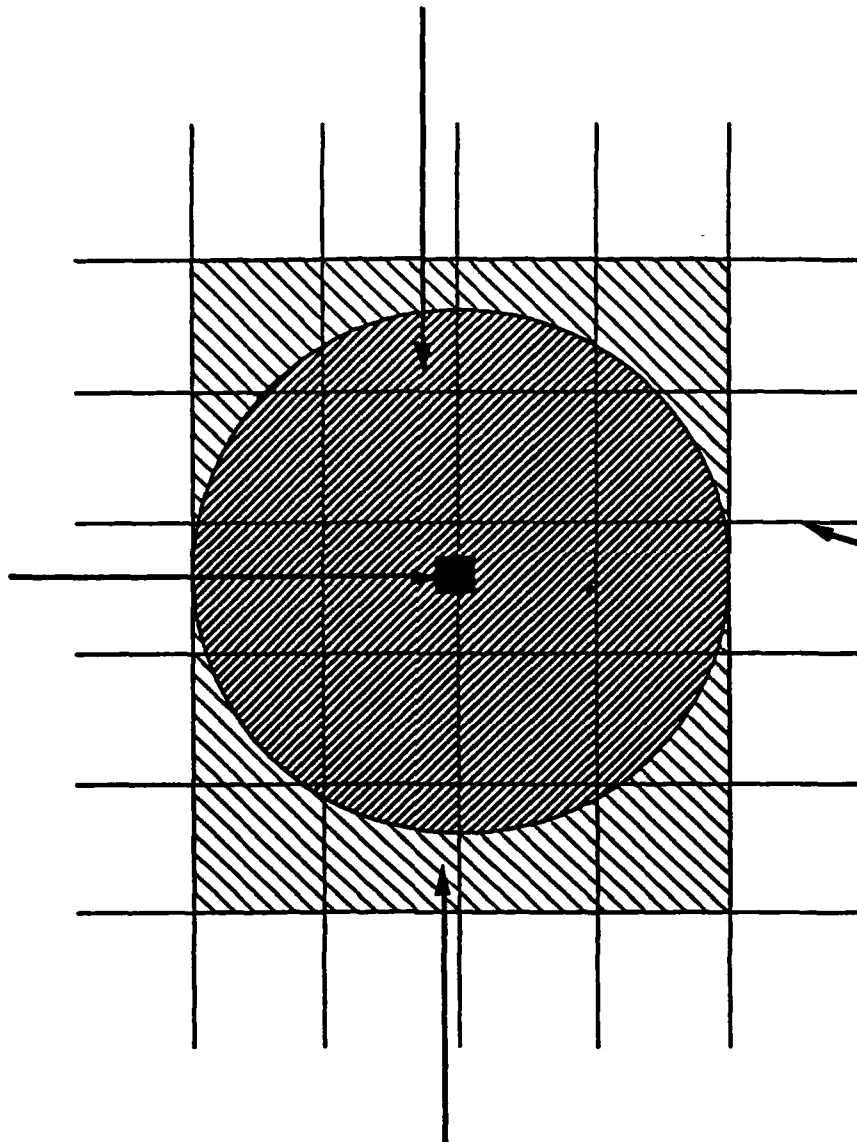


BBN Systems and Technologies

Exploring the Limits of DIS Scalability...



Entity



Terrain Area PDUs Received From

Entity Region of Interest

Grid System

LORAL
Advanced Distributed Simulation




BBN Systems and Technologies



On-Demand Forwarding (ODF) Geographic Filtering

- **Idea**
 - Receivers determine which entities they need to receive state from via a range test
- **Approach**
 - Broadcast low rate summary PDUs that tersely describe entities
 - Range squared test performed by receivers determines entities within range; state update requested
 - Multicast group formed on set of requesting receivers; state updates forwarded to the set of requesting receivers
- **Issues**
 - Visibility computations intensive, but not unreasonable
 - Summary PDU bit rate is reasonable
 - Tradeoff between visibility checking on entities vs. aggregates: finer grain traffic control vs. compute efficiency





ODF Geographic Filtering (Cont.)


- **Extensions**
 - Summary PDU may help solve the wide area viewer problem
 - Combination of multicast, unicast and broadcast if multicast capabilities are limited
 - DIS Aggregate Protocol may be applicable
- **Some Example Resource Requirements**
 - Assume
 - 100 instructions per range test
 - 10 second summary update period
 - 18 sites
 - 64 bits per entity per summary update
 - For 10000 entities: 52 mips and 64 kbps
 - For 5000 entities: 13 mips and 32 kbps
 - An R4000 processor is nominally about 50 mips
 - Range test is parallelizable



Exploring the Limits of DIS Scalability...



ODF/Grid Comparison




LORAL
Advanced Distributed Simulation



BBN Systems and Technologies

Exploring the Limits of DIS Scalability...



Dead Entity Service (DES)

- **Idea**
 - Distribute packet generation for destroyed entities
- **Approach**
 - Simulating host hands off packet generation to dead entity service upon destruction of an entity
 - Dead entity service is distributed - **not a server**
 - Service emits local copy of state PDU
 - Simulating host transmits at a low rate or upon state changes
- **Issues**
 - Synchronization difficulties preclude handing off simulation of dead entities - only packet replication handed off
 - Handoff protocol design

LORAL
Advanced Distributed
Simulation




BBN Systems and Technologies




Status of Our DES Analysis


- Lots of packets from dead entities at the end of an exercise
 - Pre-STOW Scenario: 30%
 - STOW Scenario B: 50%
- Preliminary Result
 - Marginal gain observed when used in conjunction with either ODF or Grid Geographic Filtering
- Current Hypotheses
 - Filtering and DES effects not additive
 - Benefits may not be apparent in aggregate measures
 - Effects are scenario dependent
- Next Steps
 - Examine effect of DES on individual network elements
 - Test with additional scenario(s)



Exploring the Limits of DIS Scalability...



Complementary Algorithms



Scenario Time Course





Fidelity Channels

Idea

- Deliver traffic at a fidelity level appropriate to the receiver
- Reduced fidelity allows reduction in data rate
- Especially applicable to wide region of interest simulators

• Approach

- Synergy with geographic filtering: ODF or Grid
- Receive data of a specified fidelity for a region of interest
- Channels supported by multicast addressing
- Remote entity approximation (dead reckoning) used to achieve fidelity specifications for each channel

• Issues

- Multicast address usage - number and change rate

• Extensions

- Receive lower fidelity data at edges of viewing range
- Possibly useful as an overload management technique

LORAL
Advanced Distributed
Simulation

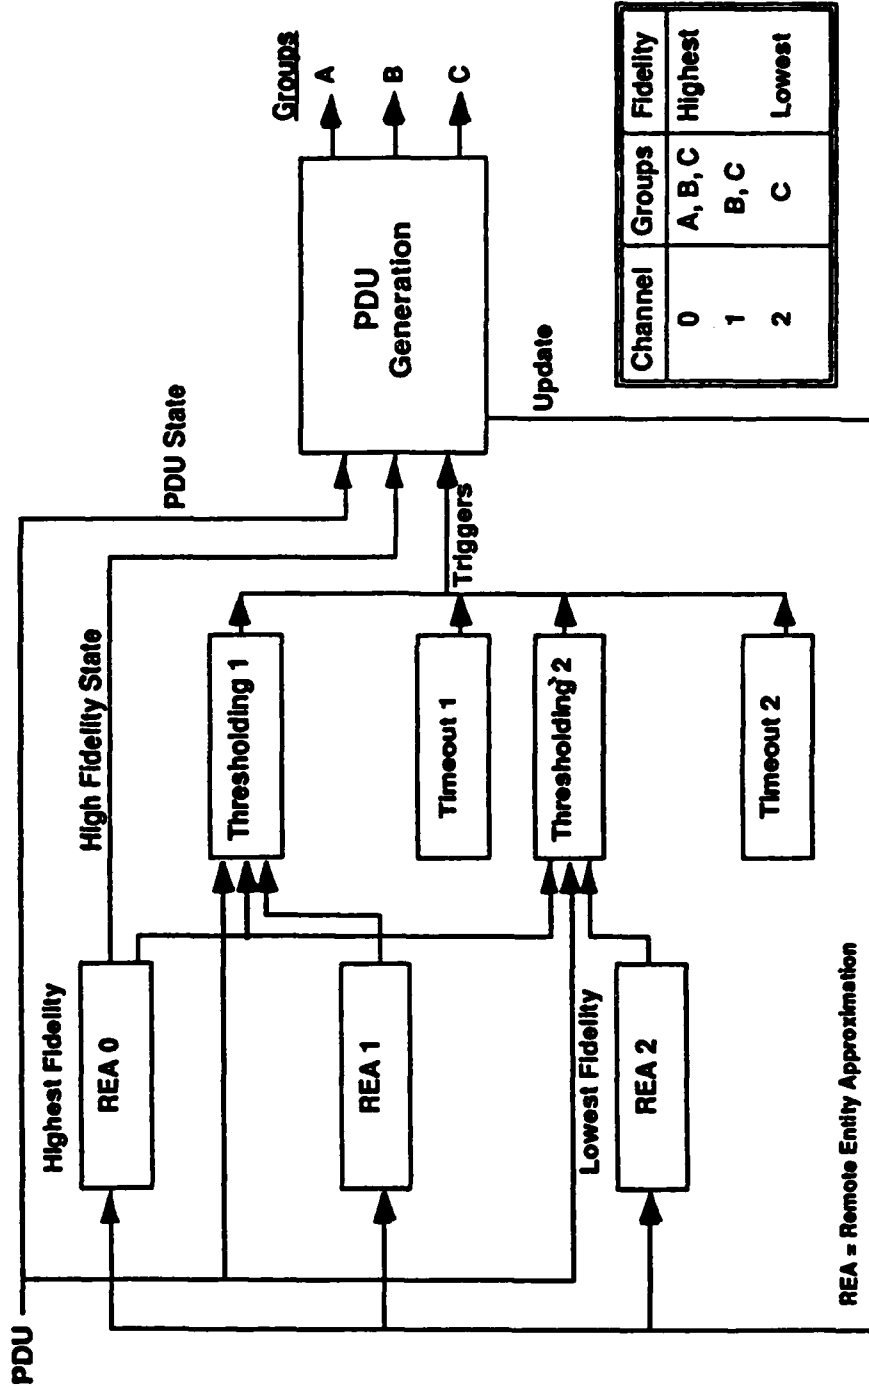


BBN Systems and Technologies

Exploring the Limits of DIS Scalability...



Fidelity Channels Processing



REA = Remote Entity Approximation



BBN Systems and Technologies



Transmit Timeout

- **Idea**
 - Increase maximum allowable time between state updates
 - Packets due to timeout
 - Pre-STOW Scenario: 50 to 70%
 - STOW Scenario B: 50 to 70%
- **Approach(es)**
 - Protocol parameter change: DIS standard default is 5 seconds
 - Distributed service ala the dead entity service
- **Issues**
 - Existing simulators, e.g., SIMNET
 - Time for hosts (re)joining an exercise to (re)acquire current state
 - Possible approach: query protocol and distributed database
 - Interaction with region of interest and viewing range for geographic filtering algorithms



Exploring the Limits of DS Scalability...



Scenario Analysis

Units
Terrain
Missions
Force Laydown
Network Sites

Scenario
Creation

Simulation

Analysis


Decisions
and
Feedback

LORAL
Advanced Distributed
Simulation



BBN Systems and Technologies

Exploring the Limits of DIS Scalability...




Measures of Effectiveness

- **Delay histogram**
 - Site to site delays
- **Congestion time**
 - Percent of time that any network element is heavily loaded
- **Drop time**
 - Percent of time that any network element is dropping packets
- **Reduction in aggregate offered output load**
 - Summed over all the tail circuits
- **Reduction in aggregate offered input load**
 - Summed over all the tail circuits
- **Dropped packets**
 - Percentage of packets dropped for each network element
- **Multicast usage**
 - Number of groups in use
 - Average and peak aggregate membership change rate

LORAL
Advanced Distributed
Simulation



BDN Systems and Technologies



Wednesday, August 18, 1993

Exploring the Limits of DIS Scalability...



Algorithm Research Status

	Design	Implement	Analyze	Report
Grid	100	100	50	0
ODF	100	100	50	0
ODF (constrained)	20	0	0	0
DES	100	100	50	0
Timeouts			80	0
Congestion Control			100	100
Fidelity Channels	40	20	0	0
Grid + DES			20	0
Composite	0	100	0	0



BBN Systems and Technologies

Exploring the Limits of DIS Scalability...



STOW Scenario Simulation Status


- *Pre-STOW Scenario - complete Fall 1992, analysis foregone*
- *Scenario A - complete, analysis foregone*
- *Scenario B - preliminary analysis underway, needs revision*
- **Preliminary Analysis**
 - Sites (with entities) 13
 - Terrain SAKI
 - Duration 4905 4:21
 - DIS Entities ~~4095~~ to 5654
 - DIS Entities (missiles) 749
 - Offered Load (pps) 1850
 - PDUs (appearance) 99.5%
 - PDUs (dead entities) up to 50%
 - PDUs (timeout) 50 to 70%

LORAL
Advanced Distributed Simulation



BBN Systems and Technologies

Exploring the Limits of DIS Scalability...




STOW Scenario B

- Need definition of projected STOW 94 network
- Need definition of wide area viewers
 - Which sites?
 - Characteristics of use?
- Review number of entities destroyed



Exploring the Limits of DIS Scalability...



Deliverables

- **Toolset Software**
 - SAF, PVD, Network Simulator, Scaler, Analysis Tools
- **Toolset Users' Guide**
- **Algorithm Emulation Software**
- **Reports on Algorithm Research**
 - Theory
 - Performance measurements
 - Fielding issues
- **Scenario Files**
- **Report on Scenario Simulation**
 - Scenario description
 - Performance measurements
 - Recommendations for scenario, network, algorithms

LORAL
Advanced Distributed
Simulation



BBN Systems and Technologies

Exploring the Limits of DIS Scalability...



Conclusions

- A combination of techniques and approaches is required
- Multiple, complementary scaling algorithms are necessary
- Network enhancements will be required
- Scenario tailoring will be required

LORAL
Advanced Distributed
Simulation



BBN Systems and Technologies

STOW Scaleability Peer Review
Intelligent Gateway
for Defense Simulation Internet

William D. Miller
AT&T Bell Laboratories
Thursday, August 19, 1993

Institute for Defense Analyses, Alexandria, VA

Topics

- **Problem Definition**
- **Technical Approach**
- **Status**
- **Deliverables**
- **Summary**

Problem Definition

- **User Needs**
- **Environmental and Design Constraints**
- **Functional Schematic**

User Needs

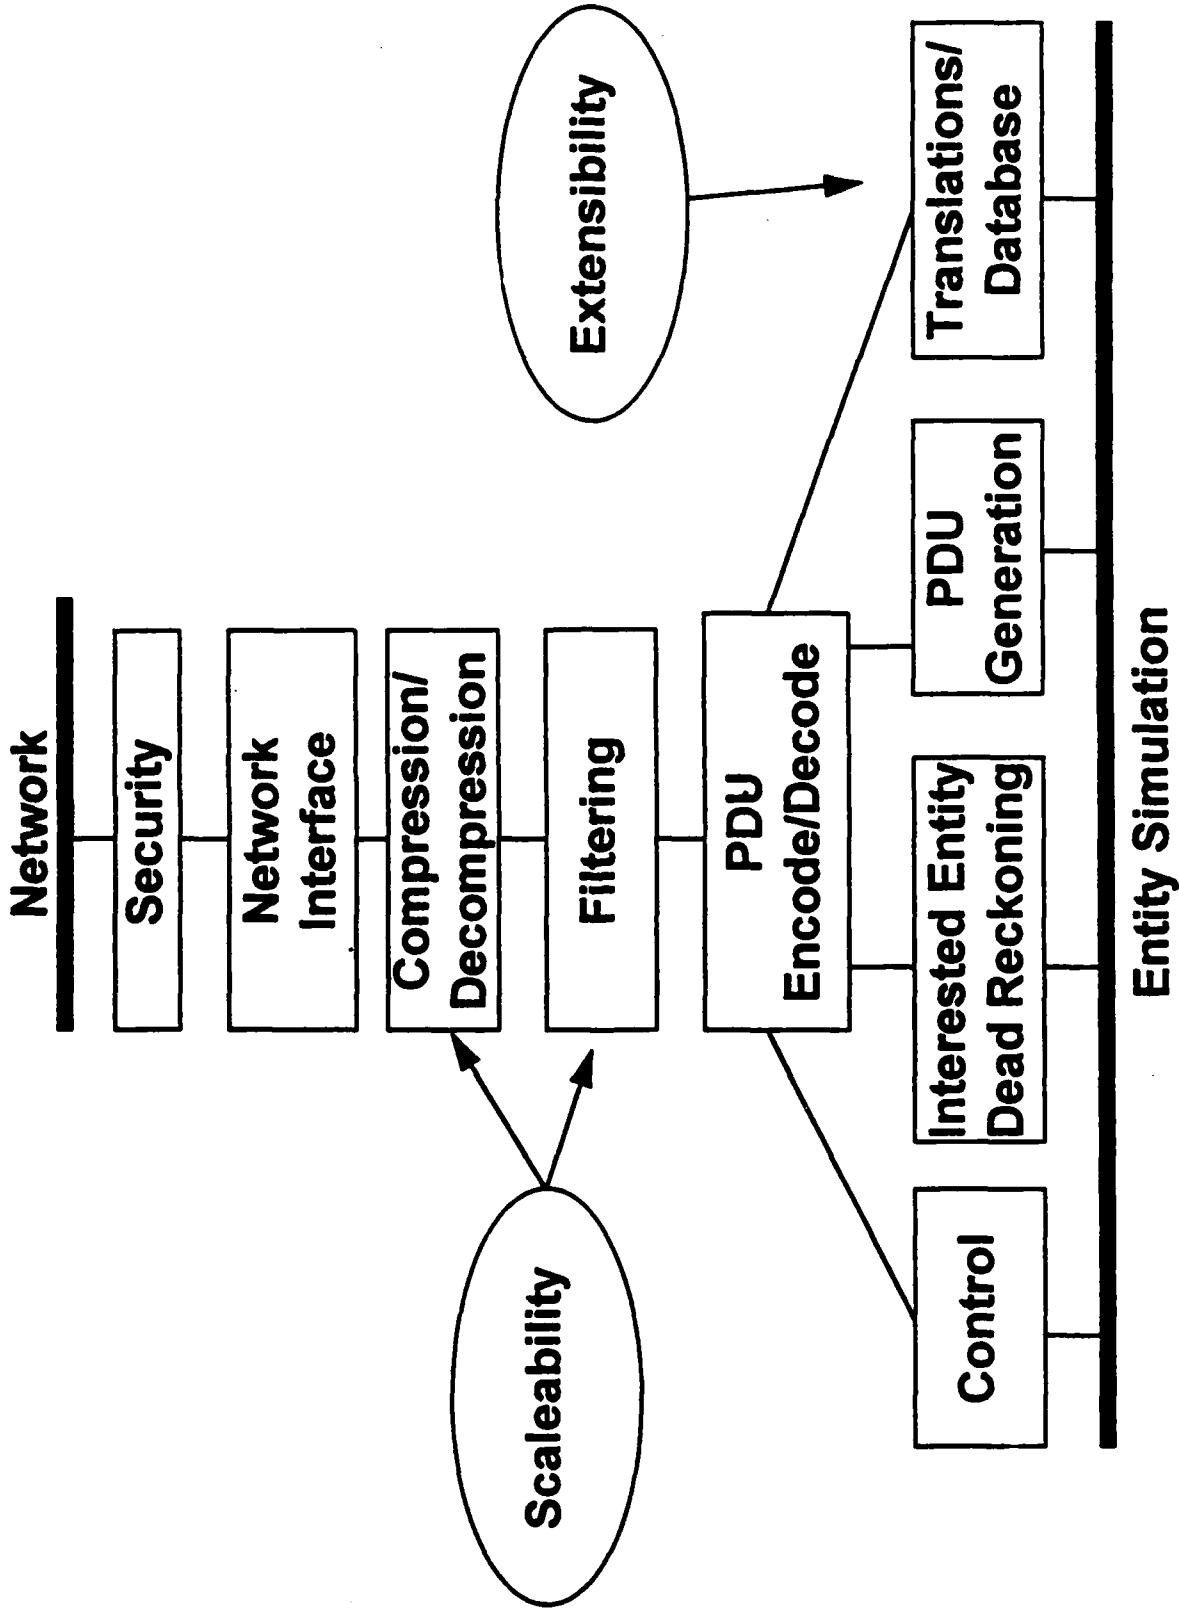
- **Simulation Scaleability**
 - 100,000 simulated entities
 - Data Culling
 - Data Compression
- **Real-Time**
- **High Reliability**
 - Low Error Rate
 - Low Overflow Rate
 - High Availability
- **Security**
- **Flexibility**
 - Programmability
 - Reconfigurability
 - Dynamic Simulation Management
- **Available for STOW 94 (10,000 entities)**

Environmental and Design Constraints

- Commercial Office/Laboratory Environment
- Open Architecture
- Physical Weight Volume Power
- Life Cycle Cost Development Production Support Maintenance
- Simulator Sites
- Network Traffic Profiles
- Processor Budget Allocation
- Latency Budget Allocation
- Intelligent Gateway Platform Selection
- Local Area Network (LAN) Interfaces
- Network Management Interfaces
- Defense Simulation Internet (DSI)
- Security Standards
- Nonuniform Simulation Fidelities

WDM 8/19/93

Functional Schematic



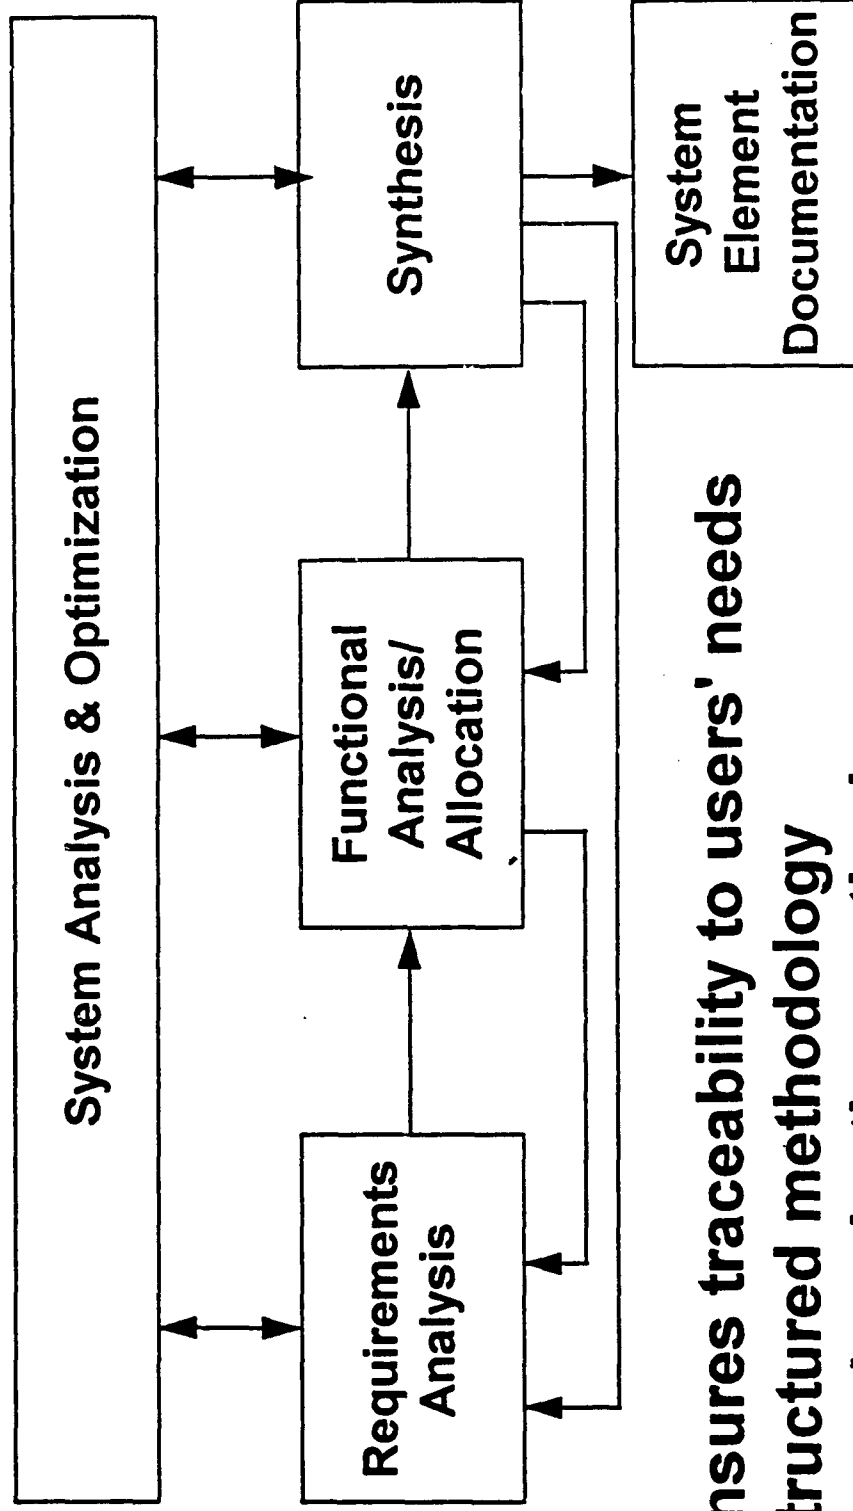
Technical Approach

- **Systems Engineering Standard Process**
- **Candidate Algorithms**
- **Analysis, Modeling, Simulation, Demonstration**
- **Candidate Platforms**

WDM 8/19/93

Systems Engineering Standard Process

System Requirements Definition:



Ensures traceability to users' needs
Structured methodology
Formal evaluation methods
Reduces false starts
Tailor processes to needs of program

Candidate Algorithms

- **Compression/Decompression**
 - Transmit only changed PDU fields
 - Transmit delta's of PDU fields
 - Higher order dead reckoning
 - Cryptic PDU
 - Lossless commercial algorithms
- **Filtering (Situational Awareness)**
 - Incoming filtering from network
(Minimal Intelligent Gateway coordination)
 - Outgoing filtering to network
(Complex Intelligent Gateway coordination)
 - Both outgoing/incoming filtering
- **Increase Entity State PDU "Timeout" Value
from 5 second default**

Candidate Algorithms Situational Awareness Filtering

CONSIDERATIONS

- **Geographic Position**
- **Line of Sight**
- **Emissions**
- **Communications Links**
- **Decision Support System Knowledge of Truth**

Analysis, Modeling, Simulation and Demonstration

- **Analyze Actual PDU Data**
 - **Plenary Scenario**
 - **Pre-STOW Mini-Demonstrations**
- **Derive Empirical Behavior Models from PDU Data Analysis**
- **Simulate Algorithms**
 - **RDD-100 CASE Tool Discrete Event Simulation**
- **Prototype Algorithms in Mini-Demos**

Candidate Platforms

- **AIU/AIU++ (GID)**
 - **Advantage of Existing Software Base**
 - **Low Risk for Mini-Demos**

- **Sun/SGI Alternatives**
 - **Potential Risk to Porting AIU Software**
 - **Multiprocessor SGI for 10-5 Solution**

- **AT&T DSP-3 Massively Parallel Processor**
 - **Potential Risk to Porting AIU Software**
 - **Viewed in Context of 10-5 Solution**

Status

- **Schedule**
- **Traffic Reduction Analysis**
- **PDU Data Analysis**
- **December Mini-Demo Plan**

Schedule

ID	Name	1993												1994										
		O	N	D	J	F	M	A	M	J	J	A	S	O	N	D	J	F	M	A	M	J	J	1994
1	Systems Engineering Planning																							
2	Define Requirements																							
3	In-Process Review																							
4	CDRL A002 (Draft)																							
5	CDRL A002 (Final)																							
6	Design Requirements																							
7	In-Process Review																							
8	In-Process Review																							
9	CDRL A005 (Draft)																							
10	CDRL A005 (Final)																							
11	Integration Requirements																							
12	Pre-STOW Mini-Demo																							
13	In-Process Review																							
14	In-Process Review																							
15	CDRL A001 (Draft)																							
16	CDRL A001 (Final)																							
17	Technology Transition Plan																							
18	CDRL A004 (Draft)																							
19	Final Review																							
20	CDRL A004 (Final)																							

WDM 8/19/93

Traffic Reduction Analysis

- Estimate of 40 Mbps Aggregate STOW Traffic
 - 24 Mbps "Low Rate" Traffic
 - 16 Mbps "High Rate" Traffic

- Reduce to 6 Mbps Aggregate Traffic
 - 2 Mbps "Low Rate" Traffic
 - using 60 second Entity State PDU "Timeout"
 - 4 Mbps "High Rate" Traffic using "Delta" PDUs

- Investigate Situational Awareness Filtering to Reduce Aggregate STOW Traffic Below 1.5 Mbps (T-1 Rate)

PLENARY SCENARIO

- **Analysis In-Process**
 - **Distribution of change in position, orientation, velocity, and acceleration**
 - **Distribution of PDU types**
 - **Interarrival time distribution for PDUs from each entity ID**
 - » **aggregated by entity class**
 - » **aggregated by PDU type**
 - **Distribution of PDU latency**
 - **Network Loading**

PLENARY SCENARIO

• Scenario Entities

- 2 F/A-18s
- Sparrow Missile
- AH-64 Attack Helo
- 2 Guided Missile Frigates - O. H. Perry
- LHD-1 Amphibious Assault Ship
- 4 SU-25 Frogfoots
- F-16C
- E-2
- Miscellaneous Weapons System (RADAR)
- 4 M1A1 Abrahms Tanks
- 6 BMP-1 Armored Fighting Vehicles
- 5 T-72 Tanks
- 4 Towed Artilleries
- M2 Bradley Fighting Vehicle
- Lifeform

PLENARY SCENARIO

- **Bounding Box**
 - X = 90 kilometers
 - Y = 78 kilometers
 - Z = 144 kilometers
- **Location**
 - latitude = 36 ° N
 - longitude = 122° W
 - where = California

PLENARY SCENARIO

- **8654 PDUs over 15 minutes for 35 entities**
 - **F-16C generates 3219 PDUs (37%)**
 - **E-2 generates 1330 PDUs (15 %)**
- **Apparent Timeout ~10 seconds**

PLENARY SCENARIO

- **INITIAL OBSERVATIONS**
 - Two-thirds of the entities do not move at all
 - 99% of the ESPDUs represent movement by less than 500 meters
 - » Suggests that changes in position (X,Y,Z) can be represented by 16 bits instead of 64 bits
 - » 16 bits can represent centimeters over a 500 meter range
 - 2 out of 35 entities generate half the PDU traffic

```

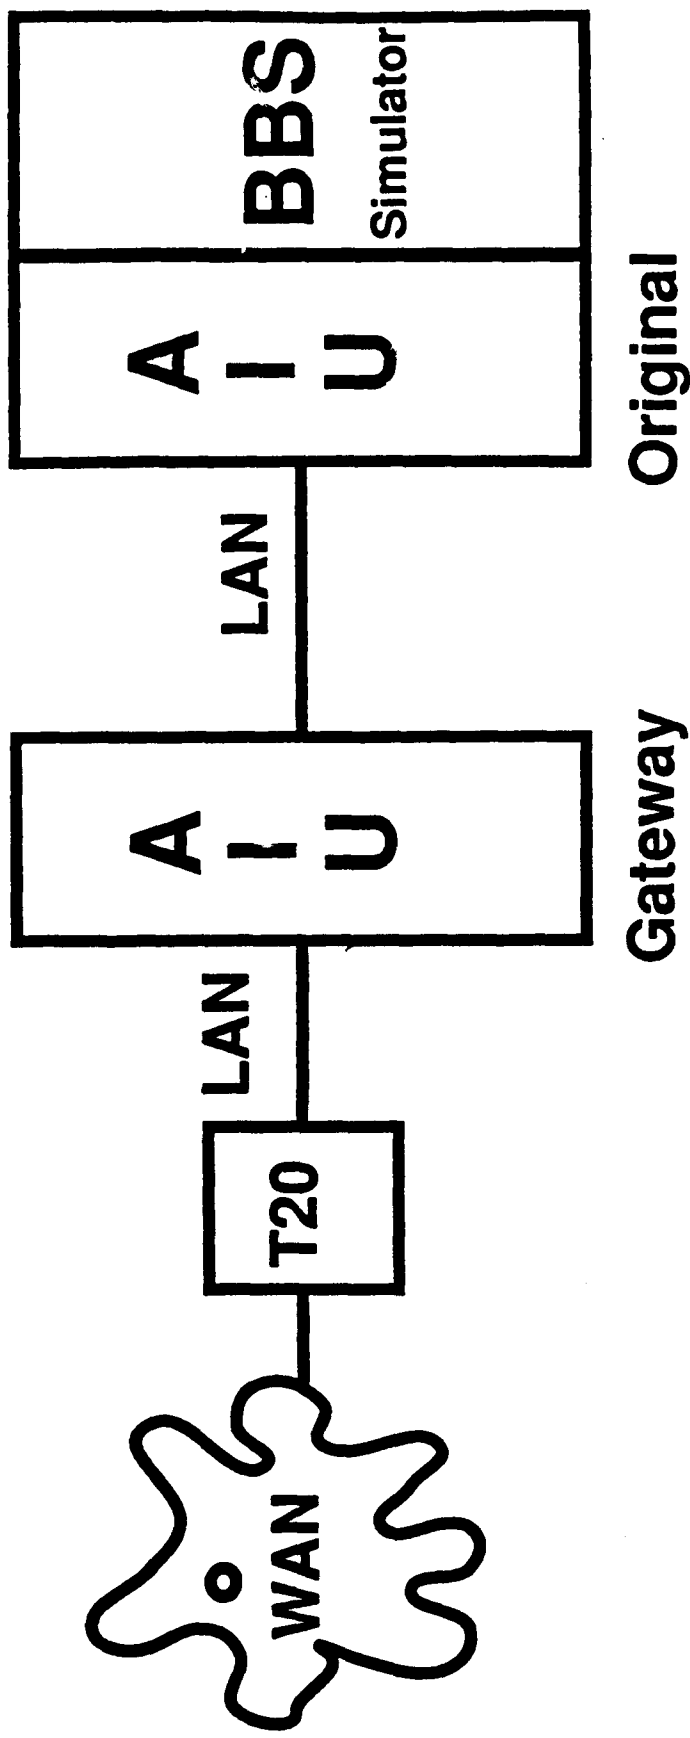
#S (LOG-RECORD
:LOG-TIMESTAMP 0
:PDU-SIZE 160
:PDU #S (ENTITY-STATE
:HEADER #S (PDU-HEADER
:VERSION 1
:EXERCISE 1
:KIND 1
:UNUSED_8 0)
:ENTITY-ID #S (ENTITY-ID
:SIMULATOR #S (SIMULATOR-ADDRESS
:SITE 1005
:HOST 1)
:ENTITY 26)
:UNUSED_8B 0
:FORCE-ID 2
:ENTITY-TYPE #S (ENTITY-TYPE
:ENTITY-KIND 1
:DOMAIN 2
:COUNTRY 164
:CATEGORY 2
:SUBCATEGORY 8
:SPECIFIC 0
:EXTRA 0)
:GUISE #S (ENTITY-TYPE
:ENTITY-KIND 0
:DOMAIN 0
:COUNTRY 0
:CATEGORY 0
:SUBCATEGORY 0
:SPECIFIC 0
:EXTRA 0)
:TIMESTAMP 2378810926
:LOCATION #S (WORLD-COORDINATES
:X -2635379.0267692064d0
:Y -4397589.30472663d0
:Z 3791781.967826985d0)
:VELOCITY #S (LINEAR-VECTOR :X 0.0 :Y 0.0 :Z 0.0)
:ORIENTATION #S (EULER-ANGLES
:PSI 3559208823
:THETA 413347928
:PHI 2669003701)
:DEAD-RECKON-PARMS #S (DEAD-RECKON-PARMS
:ALGORITHM 2
:UNUSED-8 0
:UNUSED-16 0
:UNUSED-32 0
:UNUSED-32-2 0
:UNUSED-32-3 0
:ACCELERATION #S (LINEAR-VECTOR
:X 0.0
:Y 0.0
:Z 0.0)
:ANGULAR-VELOCITY #S (ANGULAR-VELOCITY-VECTOR
:ROLL 0
:PITCH 0
:YAW 0))
:APPEARANCE 0
:MARKING #S (ENTITY-MARKING :CHARACTER-SET 0 :TEXT "MDTS")
:CAPABILITIES #S (ENTITY-CAPABILITIES
:AMMUNITION-SUPPLY 0
:FUEL-SUPPLY 0
:MISC-SUPPLY 0
:REPAIR 0
:UNUSED 0)

```

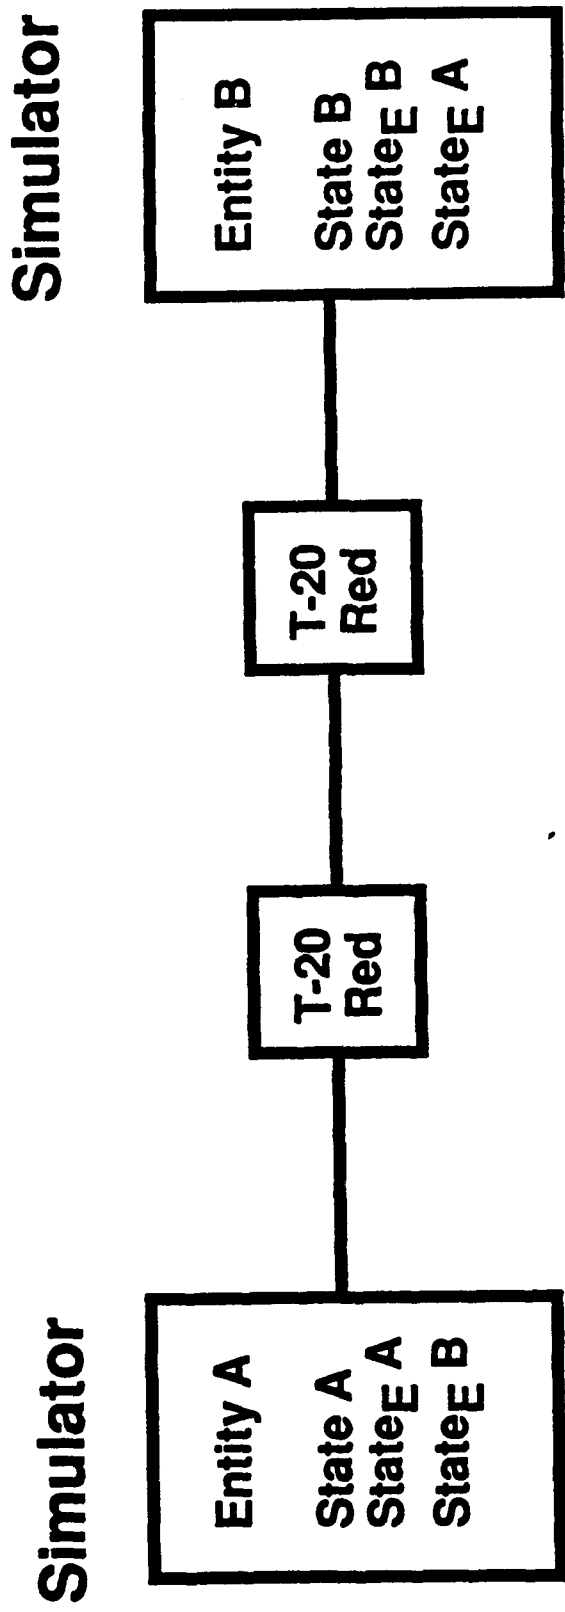

:UNUSED-16-2 0
:UNUSED-8-2 0
:NUM-PARTS 0
:PARTS #(#S(ARTICULATED-PART
:CHANGE 1081
:PART-ID 31232
:TYPE 4011
:VALUE #S(SIXTY-FOUR-BITS
:WORDS #(1357661873 0))))))

Pre-Stow Demonstration

November 1993



DSI Network



Timeout Algorithm Variables

t	clock time of entity A as maintained by simulator A and placed in timestamp of PDU sent out for entity A. Also time that B uses for modeling A.
$t+1$	next clock time of entity A as maintained by simulator A and placed in timestamp of next PDU sent out for A. Also used by B for modeling A.
$S(t)$	State of entity A at time t
$S_E(t)$	Estimated State of entity A at time t using dead reckoning model
t_T	time when last PDU sent out for entity A
T	Timeout parameter
T'	Extend timeout parameter
Δ_{MAX}	Maximum allowed deviation between $S(t)$ and $S_E(t)$ before PDU sent out.

Timeout Demonstration Assumptions

- 1 Simulators uses timeout of T (5 seconds)**
- 2 Gateway enforces timeout of T' (60 seconds) for entity state PDU's only**
- 3 "A" represents a generic entity on the WAN side.**
- 4 Gateway AIU will maintain dead reckoning models of all A entities using timestamp stored in EntityState PDU. For each A, the gateway will decidewhether the PDU should be sent out under the new timeout period.**
- 5 Demonstration will show effect of new timeout period for A entities only.**

Timeout Algorithm for Entity A

(current algorithm)

- 1 $S(t) \rightarrow S(t+1)$
- 2 $S_E(t) \rightarrow S_E(t+1)$
- 3 if $(t+1-t_T) > T$
 - Send PDU with entity A state information
 - $S_E(t+1) = S(t+1)$
 - $t_T = t+1$
 - goto 1
- 4 if $(S_E(t+1) - S(t+1)) > \Delta_{MAX}$
 - Send PDU with entity A state information
 - $S_E(t+1) = S(t+1)$
 - $t_T = t+1$
 - goto 1
- 5 goto 1

Timeout Algorithm for Simulator B (current algorithm)

- 1 if receive PDU from A
 $S_E(t)$ = State information from PDU**

Timeout Demonstration

Algorithm for IG

(for each A)

- 1 Input PDU from entity A
- 2 $t = \text{time (PDU)}$
- 3 if $(t - t_T) > T$:
 - $S_E(t) = \text{State (PDU)}$
 - $t_T = t$
 - Send PDU onto B
 - goto 1
- 4 $S_E(t_T) \rightarrow S_E(t)$
- 5 if $|\text{State (PDU)} - S_E(t)| > \Delta_{\text{MAX}}$:
 - $S_E(t) = \text{State (PDU)}$
 - $t_T = t$
 - Send PDU onto B
 - goto 1
- 4 goto 1

Deliverables

Requirements Document

CDRL A002

(System Specification)

- 1. SCOPE
 - 2. APPLICABLE DOCUMENTS
 - 3. REQUIREMENTS
 - 3.1 Definition
 - 3.2 Characteristics
 - 3.2.1 Performance Characteristics
 - 3.2.2 System Capability Relationships
 - 3.2.3 External Interface Requirements
 - 3.2.4 Physical Characteristics
 - 3.2.5 System Quality Factors
 - 3.2.6 Environmental Conditions
 - 3.2.7 Transportability
 - 3.2.8 Flexibility and Expansion
 - 3.2.9 Portability
 - 3.3 Design and Construction
 - 3.3.1 Materials
 - 3.3.2 Electromagnetic Radiation
 - 3.3.3 Nameplates and Product Marking
 - 3.3.4 Workmanship
 - 3.3.5 Interchangeability
 - 3.3.6 Safety
 - 3.3.7 Human Engineering
 - 3.3.8 Reserved
 - 3.3.9 System Security
 - 3.3.10 Government Furnished Property Usage
 - 3.3.11 Computer Resource Reserve Capacity
 - 3.4 Documentation
 - 3.5 Logistics
 - 3.5.1 Support Concept
 - 3.5.2 Support Facilities
 - 3.6 Personnel and Training
 - 3.6.1 Personnel
 - 3.6.2 Training
 - 3.7 Characteristics of Subordinate Elements
 - 3.8 Precedence
 - 3.9 Qualification
 - 3.10 Standard Sample
 - 3.11 Preproduction Sample
- 4. QUALIFICATION ASSURANCE PROVISIONS
 - 4.1 Responsibility for Inspections
 - 4.2 Special Tests and Examinations
 - 4.3 Requirements Cross Reference
 - 5. PREPARATION FOR DELIVERY
 - 6. NOTES
- APPENDICES

Design Requirements/Methodologies

CDRL A005

1. SCOPE
 2. APPLICABLE DOCUMENTS
 3. REQUIREMENTS
 - 3.1 System Definition
 - 3.1.1 Missions
 - 3.1.2 Threat
 - 3.1.3 System Modes and States
 - 3.1.4 System Functions
 - 3.1.5 System Functional Relationships
 - 3.1.6 Configuration Allocation
 - 3.1.7 Interface Requirements
 - 3.1.8 Furnished Property List
 - 3.2 System Characteristics
 - 3.2.1 Physical Requirements
 - 3.2.2 Environmental Conditions
 - 3.2.3 Reserved
 - 3.2.4 Materials, Processes, and Parts
 - 3.2.5 Electromagnetic Radiation
 - 3.2.6 Workmanship
 - 3.2.7 Interchangeability
 - 3.2.8 Safety
 - 3.2.9 Human Performance/Human Engineering
 - 3.2.10 Deployment Requirements
 - 3.2.11 System Effectiveness Models
 - 3.3 Processing Resources
 - 3.4 Quality Factors
 - 3.4.1 Reliability
 - 3.4.2 Modifiability
 - 3.4.3 Availability
 - 3.4.4 Portability
 - 3.4.5 Additional Quality Factors
 - 3.5 Logistics
 - 3.5.1 Support Concept
 - 3.5.2 Support Facilities
 - 3.5.3 Supply
 - 3.5.4 Personnel
 - 3.5.5 Training
 - 3.6 Precedence
4. QUALIFICATION REQUIREMENTS
 - 4.1 General
 - 4.1.1 Philosophy of Testing
 - 4.1.2 Location of Testing
 - 4.1.3 Responsibility for Tests
 - 4.1.4 Qualification Methods
 - 4.1.5 Test Levels
 - 4.2 Formal Tests
 - 4.3 Formal Test Constraints
 - 4.4 Qualification Cross Reference
5. PREPARATION FOR DELIVERY
6. NOTES
- APPENDICES
- A. CANDIDATE COMPRESSION/FILTERING PARADIGM TRADE STUDY REPORT
 - B. DISCRETE EVENT SIMULATION MODELS AND ANALYSES
 - C. SOFT COPY DISKETTE

Final Report Outline

CDRL A001

(Product Spec)

- 1. SCOPE
- 2. APPLICABLE DOCUMENTS
- 3. REQUIREMENTS
 - 3.1 System Definition
 - 3.1.1 Missions
 - 3.1.2 Threat
 - 3.1.3 System Modes and States
 - 3.1.4 System Functions
 - 3.1.5 System Functional Relationships
 - 3.1.6 Configuration Allocation
 - 3.1.7 Interface Requirements
 - 3.1.8 Furnished Property List
 - 3.2 System Characteristics
 - 3.2.1 Physical Requirements
 - 3.2.2 Environmental Conditions
 - 3.2.3 Reserved
 - 3.2.4 Materials, Processes, and Parts
 - 3.2.5 Electromagnetic Radiation
 - 3.2.6 Workmanship
 - 3.2.7 Interchangeability
 - 3.2.8 Safety
 - 3.2.9 Human Performance/Human Engineering
 - 3.2.10 Deployment Requirements
 - 3.2.11 System Effectiveness Models
 - 3.3 Processing Resources
 - 3.4 Quality Factors
 - 3.4.1 Reliability
 - 3.4.2 Modifiability
 - 3.4.3 Availability
 - 3.4.4 Portability
 - 3.4.5 Additional Quality Factors
 - 3.5 Logistics
 - 3.5.1 Support Concept
 - 3.5.2 Support Facilities
 - 3.5.3 Supply
 - 3.5.4 Personnel
 - 3.5.5 Training
 - 3.6 Precedence
- 4. QUALIFICATION REQUIREMENTS
 - 4.1 General
 - 4.1.1 Philosophy of Testing
 - 4.1.2 Location of Testing
 - 4.1.3 Responsibility for Tests
 - 4.1.4 Qualification Methods
 - 4.1.5 Test Levels
 - 4.2 Formal Tests
 - 4.2.1 Formal Test Constraints
 - 4.2.2 Qualification Cross Reference
- 5. PREPARATION FOR DELIVERY
- 6. NOTES
- APPENDICES
 - A. SOFT COPY DISKETTE

Technology Transfer Plan

CDRL A004

1. SCOPE
 2. APPLICABLE DOCUMENTS
 3. DEFINITION REQUIREMENTS
REPORT (RDD-100 FORMAT)
 4. DESIGN REQUIREMENTS REPORT
(RDD-100 FORMAT)
 - 4.1 Functional Flow Block Diagrams
 - 4.2 Requirements Allocation
 - 4.3 Schematic Block Diagrams
 5. DISCRETE SIMULATION MODELING
AND ANALYSIS REPORT (RDD-100 FORMAT)
 6. EMULATION SOFTWARE LISTINGS
 7. STOW 94 TRANSITION
 - 7.1 STOW 94 Systems Engineering Master Schedule (SEMS)
 - 7.2 Systems Engineering Detailed Schedule (SEDS)
 - 7.2.1 Intelligent Gateway Platform Selection
 - 7.2.2 Intelligent Gateway Platform Integration
 - 7.2.3 Intelligent Gateway for Defense Simulation Internet
Interface Control Document (ICD)
 - 7.4 Intelligent Gateway Internal Interface Agreements (IA)
 - 7.5 Intelligent Gateway Master Test and Evaluation Plan
(MTEP)
 - 7.6 Intelligent Gateway for Defense Simulation Internet
Development Plan
 - 7.7 Intelligent Gateway for Defense Simulation Internet
Controlled Introduction Plan
 - 7.8 Intelligent Gateway for Defense Simulation Internet
Training Plan
 - 7.9 Intelligent Gateway for Defense
Simulation Internet Operations Plan
 - 7.9.1 Installation Plan
 - 7.9.2 Administration Plan
 - 7.9.3 Maintenance Plan
 - 7.9.4 Pre-Operations Plan
 - 7.9.5 Operations Plan
 - 7.9.6 Post-Operations Plan
 8. DISTRIBUTED INTERACTIVE SIMULATION (DIS)
APPLICATIONS
 9. COMMERCIAL APPLICATIONS
 10. NOTES
- APPENDICES
- A. DEFINITION REQUIREMENTS DATABASE
(RDD-100 FILE ON DISK MEDIA)
 - B. DESIGN REQUIREMENTS DATABASE
(RDD-100 FILE ON DISK MEDIA)
 - C. DISCRETE TIME SIMULATION MODELS
(RDD-100 FILE ON DISK MEDIA)
 - D. EMULATION SOFTWARE
(SOURCE CODE ON DISK MEDIA)
 - E. EMULATION TEST SOFTWARE
(SOURCE CODE ON DISK MEDIA)

WDM 8/19/93

Technology Transfer Plan Contents

- **Technical Data**
- **RDD-100 Database Soft Copy**
- **Compression/Filtering Emulation
Source Code**
- **Compression/Filtering Test Case
Source Code**
- **STOW 94 Transition Plan**
- **Distributed Interactive Simulation (DIS)
Applications**
- **Commercial Applications**

WDM 8/19/93

Summary

- **Emphasis on Structured, Documented Analysis**
- **Seek a Common Intelligent Gateway Requirements Document**
- **Trade Studies Crucial to Identify Algorithms, Architectures**
 - **Analyze Actual PDU Data**
 - **Prototype Algorithms in Demonstrations**
- **Development of Integration Vehicle Driven by Requirements**
- **Identify Broader Applications in Technology Transfer Plan**



STOW 94 Network Requirements & Evaluation of Solutions

**Dr. Joshua Seeger
Manager, Gigabit Networks Research
BBN Systems & Technologies**

**BBN Systems & Technologies
10 Moulton St.
Cambridge, MA 02138**

**Phone: (617) 873 2456
FAX: (617) 873 3776
email: jseeger@bbn.com**

Agenda

- **Project: Scalability architecture for the DSI**
 - **Project description & objectives**
- **Approach**
 - **BBN/LADS unique collaboration**
 - **Requirements definition**
 - **Identifying & evaluating solutions**
- **Accomplishments to date**
 - **Requirements defined**
 - **Solutions considered**
- **Work remaining**

Project Description

- **To define a scalability architecture for the Distributed Simulation Internet to support STOW (and more).**
 - **Current DSI does not support multicast-dependent scaling algorithms under study by simulator developers.**
 - **Current DSI lacks sufficient backbone bandwidth for expected traffic loading of STOW-94**
 - **Current suite of site equipment represents large investment & new sites are currently being installed.**
- **Objectives:**
 - **To determine requirements for network bandwidth and functionality imposed by STOW (and other DSI users).**
 - **To apply requirements and real-world constraints to evaluate alternate approaches to DSI improvement.**
 - **To provide the Gov't. with near- and long-term recommendations for DSI evolution.**

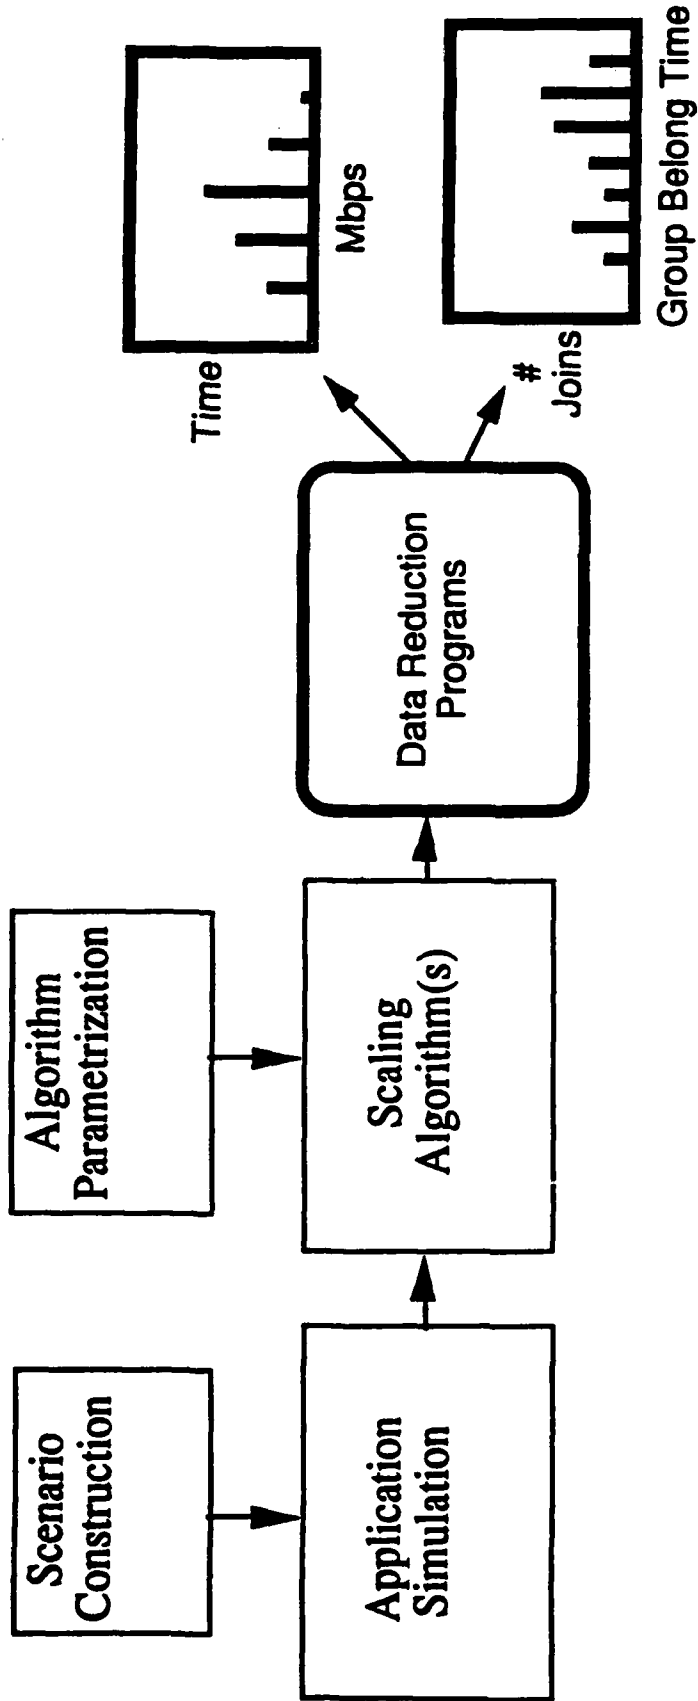
Approach

BBN & LADS Collaboration

- **Unique combination of networking and distributed simulation expertise.**
- **LADS provides information about simulator behavior and attainable traffic reduction.**
- **BBN provides information on capabilities of present and emerging networking technologies.**
- **LADS assembles sample scenarios; currently the best estimate of network subscription and traffic loading.**
- **LADS and BBN translate scenarios into expected network loading (traffic and switch processing).**
- **Process involves frequent feedback and iteration.**

Approach

- **STOW Requirements Definition:**



Approach

- **Key Requirements**
 - **Minimum number of multicast groups**
 - **Peak & average frequency of joins/leaves**
 - **Maximum join/leave latency**
 - **Peak backbone & tail circuit loading vs. multicast capability**

Approach

- **Identifying solutions:**
 - **By participation in IETF WGs and ARPA-sponsored networking research**
 - **Awareness of standards, products, ongoing research**
- **Evaluating solutions:**
 - **Requirements fulfillment & expected performance**
 - **Other evaluation criteria, e.g:**
 - Off-the-shelf availability**
 - How much it capitalizes on installed base**
 - Amount of additional development required**
 - Cost, complexity, risk, lead time**
 - Adherence to current standards**
 - Ease of migration to emerging standards**

Accomplishments

- Studied the benefits of network congestion management for distributed simulation
- Analyzed LADS' original 10,000-entity scenario
 - Histograms:
 - multicast group membership "hold time"
 - multicast group joins/leaves per unit time
 - backbone traffic load

Congestion Management for the DSI

- **Congestion management cannot rescue an overloaded network.**
- **Congestion management can avoid catastrophic degradation of simulation quality by filtering traffic fairly in case of excessive load.**

Recommendations specific to distributed simulation:

- **Congestion mgt. mechanisms should be based on local node info. at the exits from the backbone.**
- **Main control action is deciding which packets to drop.**
- **Fairness should be enforced on a per-flow basis (origin-destination).**
- **Packet priorities and time stamps would be effective dropping criteria.**
- **Resource reservation could confine adverse effects of congestion within specific classes of traffic.**

Solutions Under Study

- Available* today - multicast:
 - MOSPF
- Available* today - multicast & resource reservation:
 - ST-II
- Available today for experimental use - multicast:
 - IP Multicast (Deering's current mrouter-based version)
- Proposed - multicast:
 - Core Based Trees
 - IP Multicast (Deering's proposal(s) for improvement)
- Proposed - resource reservation:
 - RSVP


(* - available in vendor-supported products)

Work Remaining


- **Analysis of LADS' STOW-94 scenario**
- **Iteration with LADS on requirements arising from different combinations of scalability algorithms**
- **Final definition of evaluation criteria (w/ Gov't.)**
- **Elaboration of candidate solutions**
- **Characterization of desired long-term solution**
- **Evaluation of candidates via analysis & modeling**
- **Near-term recommendation**

Multicast Group Subscription Hold-Time Cumulative Distribution


grid = 2.5 Km —



Total Bandwidth Histogram




Frequency of Join/Leave Signalling



Multicast Group Subscription Hold-Time Distribution

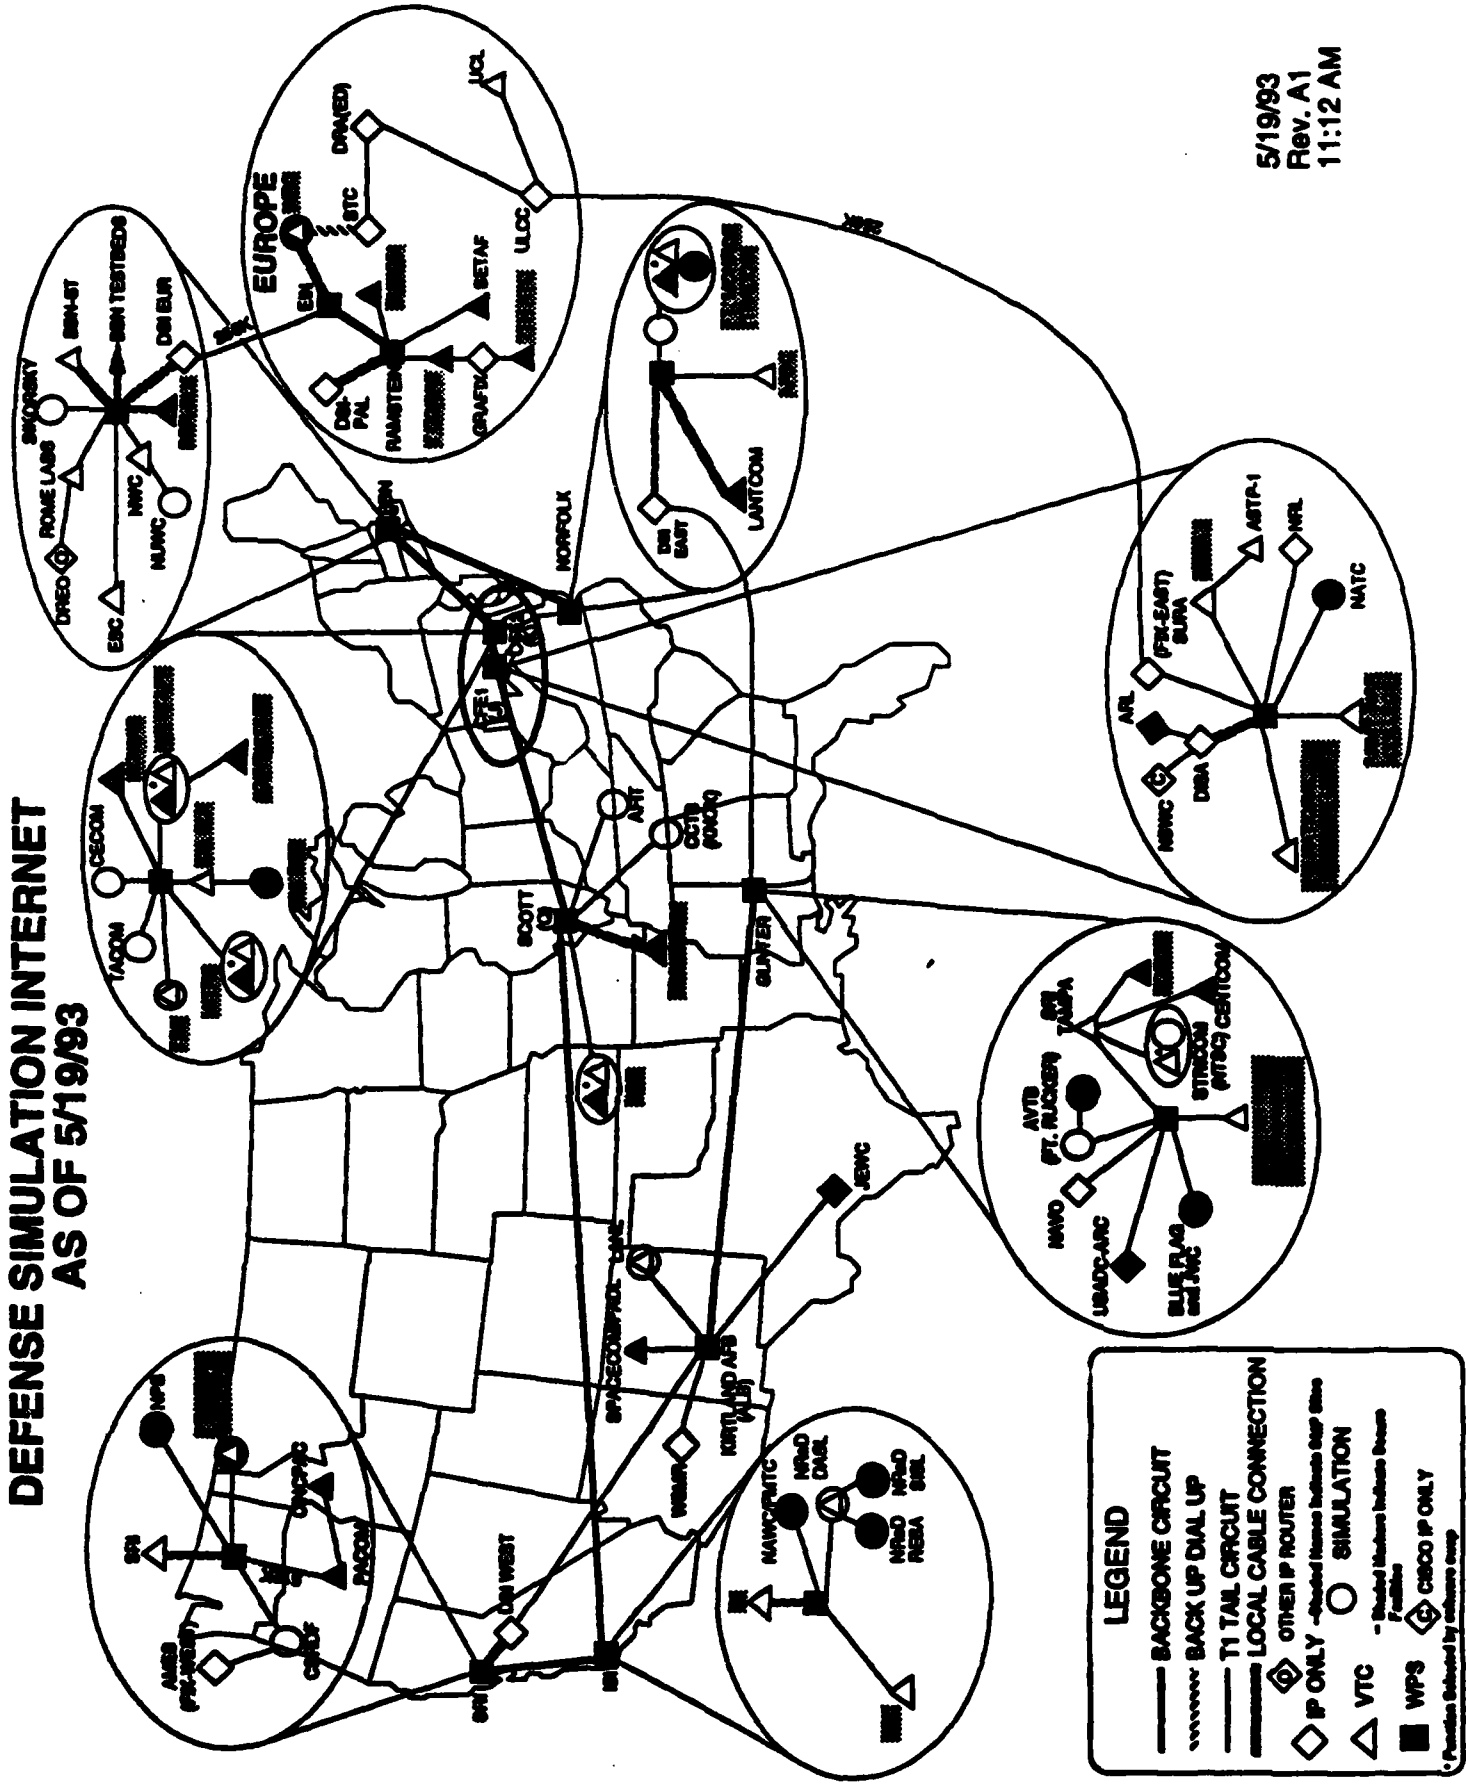
grid = 2.5 Km



Today's DSI

Lou Berger
August 19, 1993
BBN Systems and Technology

DEFENSE SIMULATION INTERNET AS OF 5/19/93




5/19/93
Rev. A1
11:12 AM

LEGEND

- BACKBONE CIRCUIT
- - - - BACK UP DIAL UP
- T1 TAIL CIRCUIT
- LOCAL CABLE CONNECTION
- ◇ IP ONLY - Shaded Nodes Indicate SAMP Sites
- △ VTC - Shaded Nodes Indicate Power Facilities
- WPS - CISCO IP ONLY
- SIMULATION

© Provided by various users


DSI Simulation Node and C²



Red Side

Black Side

Motorola NES Profile



Protocols: IP, CLNP, SDNS,
 Ethernet/ 802.3, X.25 (Black side only)
 Can support others via encapsulation

Performance: ~60 pps @ 1440 byte packets = ~700 kbps
 ~125 pps @ 512 byte packets = ~512 kbps
 ~200 pps @ 64 byte packets = ~102 kbps

Black Interface: Standard IP Host

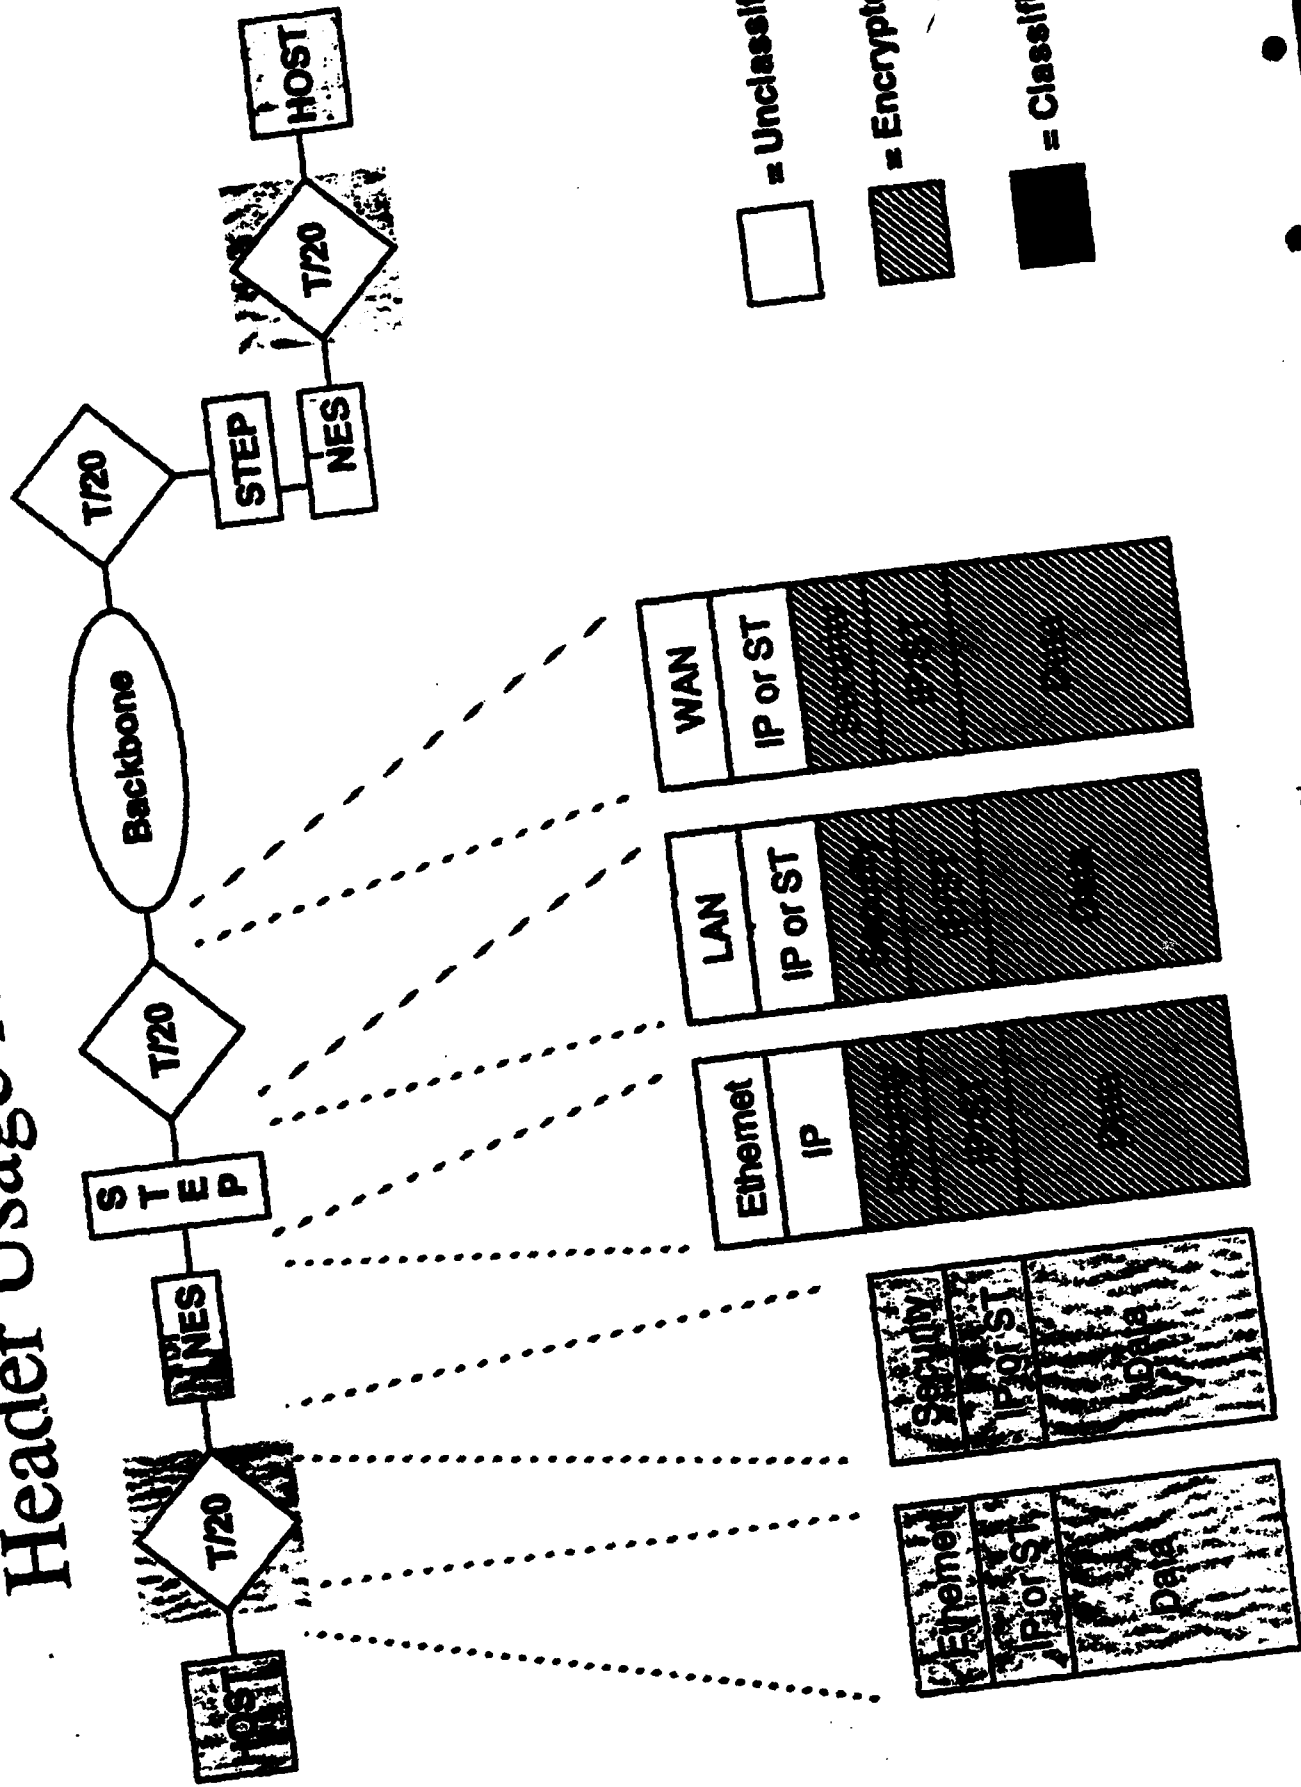
Red Interface: IP Transparent Bridge/Router
 Transparent Bridge

Number of multicast groups: 1 = Broadcast

Target E3 Performance Using NESEs

- Goal:** To support full T1 capacity
- Approach:** Parallel NESEs
Aggregation in applications
- Number of NESEs needed:** 5
- Why:** 5 NESEs @ ~700 kbps = ~3.5 Mbps
T1 = 1.544 Mbps in each direction
= ~3 Mbps aggregate (in+out)

Header Usage Across NES E3 DSI



Security Header Definition

Standard Ethernet Header

Destination Address	(48 bits)
Source Address	(48 bits)
Type/Length	(16 bits)

Security Header (Unicast)

Unused	(8 bits)
Destination E3D	(8 bits)
Destination IP Address	(32 bits)
Unused	(8bits)
Source E3D	(8 bits)
Source IP Address	(32 bits)
Type/Length	(16 bits)

Security Header (multicast)

Destination Multicast Address (ff:ff:ff:ff:ff:ff)	(48 bits)
Unused	(8 bits)
Source E3D	(8 bits)
Source IP Address	(32 bits)
Type/Length	(16 bits)

DSI Attributes and Limitations

- ST - II
- Backbone Related
- T/20
- Application Interface

DSI Attributions and Limitations: ST-II

ST2 Attributes

- IETF experimental protocol - RFC 1190
 - IETF Working Group being formed to simplify ST-II spec
- Provides:
 - Circuit oriented connections
 - Resource reservation
 - Dynamic multicast
- Does not require custom routing protocol
- Can scale to support large number of simultaneous connections

DSI Attributes and Limitations: ST-II (continued)

ST-II Limitations

- Protocol specification needs:
 - Clarification
 - Formal definition of states and state transitions
 - Minimum subset definition
 - FlowSpec refinement
 - Guidance on use of standard routing protocols
- Scalability
 - A single connection can not handle a large number of targets
- Not widely available
 - There are only 4 operational and about 9 in progress implementations

DSI Attributes and Limitations

Backbone Related

Backbone Attributes

- Supports up to T1/E1 trunks and tail circuits
- Provides resource allocation
- Provides static and dynamic multicast groups
- Network level switching

Backbone Limitations

- Cannot support needed T3 backbone trunks
- Current T/20 implementation uses static multicast groups
 - Number of static backbone groups is limited
 - Membership in groups is controlled manually ✓

DSI Attributes and Limitations:

T/20

T/20 Attributes:

- Supports ST, standard IP, and dynamic ST and IP routing Protocols: IP, ST-II, Ethernet/802.3, Token Ring 802.5, PPP, Frame Relay DTE, BBN Trunk, X.25, EGP, BGP, SNMP, Telnet, TFTP
- Provides resource management services to ST-II
- Provides platform for ST-II application interfaces DIS, VTC, E3, and Voice

T/20 Limitation:

- Provides limited performance - ~6000 pps (max)
But it is adequate for T1 line limited sites
- Supports a compiled in maximum of ~128 simultaneous ST connections
- Does not implement full ST-II specification
Prioritized streams and preemption are not supported
- Does not implement full resource allocation
T/20 implements absolute, not managed, ST-II priority over IP

DSI Attributes and Limitations: DIS Application Interface

DIS Application Interface Attributes

- T/20 based
- Bridges traffic from multiple sites
- Uses ST-II multicast and resource allocation services

DIS Application Interface Limitations

- Requires manual (remote) control
- Support single output connection
One interface maps into one multicast group
- Not DIS knowledgeable

SUMMARY:

Key Technical Issues for STOW 94

- o NES currently supports 1 multicast group, may support 16 in 1994
- o ST-II limits number of end points per ST connection
- o T/20 limits number of simultaneous connections per T/20 to 128
This can be changed through recompilation and memory upgrades
- o Backbone is now T1 will be upgraded completely to T3 by mid 1994.
- o Current DIS application gateway ^{for T20} needs to be replaced or enhanced to support STOW 94

SUMMARY:
Key Technical Issues for STOW 94

- o DSI will support T1 site tails and T3 backbone trunks
- o DSI will support IP and ST-II, will consider recommendations for alternatives
- o NES will be used in parallel to provide T1 rate access
Number of needed NESes will depend on average packet size

APPENDIX D
READING LIST

NRaD: Scaleability Peer Review: STOW 94. San Diego, CA: Naval Command, Control and Ocean Surveillance Center, RDT&E Division.

Notebook containing the following read-ahead material and material distributed at the review. Presentation is in order of the notebook.

- Purpose handout
- General Background Information.
- Scaleability Vision and Strategy
- Scaling Algorithms and Multicast Requirements
- Loral Advanced Distributed Simulation (LADS) Algorithm Description
- Viewgraphs from the June 11, 1993, STOW IPR, Washington, D.C.
- Plans
- Tool Kit Data
- Defense Simulation Internet (DSI)
- Distributed Interactive Simulation (DIS). Proposed IEEE standard draft. *Standard for Information Technology—Protocols for Distributed Interactive Simulation Applications*, Version 2.0, Third Draft. Orlando, FL: Institute for Simulation and Training, May 28, 1993.
- Selections from IDA Document D-780, *Panel Review of Long-Haul Networking in Distributed Simulation*, June 1990: abstract, table of contents, abbreviations, summary.
- *Summary Report: The Eighth Workshop on Standards for the Interoperability of Defense Simulations. Volume II: Minutes from the Working Sessions*. Orlando, FL: Institute for Simulation and Training, March 22-26, 1993.

APPENDIX E
LIST OF ATTENDEES

Panel Members

Dr. David R. Cheriton
Computer Science Dept.
Stanford University
Stanford, CA 94305-2140
(415) 723-1131
cheriton@cs.stanford.edu

Dr. Dale B. Henderson
Mail Stop F606
Los Alamos National Laboratory
Los Alamos, NM 87545
(505) 665-2151
dbh@lanl.gov

Dr. Duncan C. Miller
Massachusetts Institute of Technology
Lincoln Laboratory
Room B-137
244 Wood St.
Lexington, MA 02173-9108
(617) 981-0158
fax (617) 981-7239
fax (617) 981-0202
dmiller@ll.mit.edu

Dr. David L. Mills
Electrical Engineering Department
Evans Hall, Room 140
University of Delaware
Newark, DE 19716
(302) 831-8247
mills@udel.edu

Mr. Stuart D. Cheshire
Adjunct Panel Member
29C Escondido Village
Stanford University
Stanford, CA 94305
(415) 497-2399
cheshire@cs.stanford.edu

Government

Mr. Tim Bailey
National Security Agency
Attention: V33
9800 Savage Road
Ft. Meade, MD 20755
(410) 859-6670

Lin Flynn
Advanced Research Projects Agency/ASTO
3701 N. Fairfax Drive
Arlington, VA 22203-1714
(703) 437-3223
lflynn@arpa.mil

RAdm Lee Kollmorgen (ret)
P.O. Box 646
1 Lost Pond
Nellysford, VA 22958
Advanced Research Projects Agency
Consultant
(804) 361-1580
kollmorg@charm.isi.edu

Mr. Ken Neubecker
National Security Agency
Attention: V51
9800 Savage Road
Ft. Meade, MD 20755
(410) 859-6281

Mr. Kevin Rainville
National Security Agency
Attention: V51
9800 Savage Road
Ft. Meade, MD 20755
(410) 859-6679

Industry

Mr. Lou Berger
BBN Systems and Technologies
1300 N. 17th Street, Suite 1200
Arlington, VA 22209
(703) 284-4651
lberger@bbn.com

Mr. William D. Miller
AT&T Bell Labs
Room 15H408
67 Whippany Road
P.O. Box 903
Whippany, NJ 07981
(201) 386-5339
william.d.miller@att.com

Ms. Barbara Perry
Loral ADS
Manager of R&D Programs
50 Moulton Street
Cambridge, MASS 02138
(617) 873-3244
bperry@camb-lads.loral.com

Dr. Joshua Seeger
BBN Systems and Technologies
Mail Stop 6/5a
10 Moulton St.
Cambridge, MA 02138
(617) 873-2456
jseeger@bbn.com

Ms. Perri-Anne Stiffler
Houston Associates, Inc.
4601 N. Fairfax Drive
Suite 708
Arlington, VA 22203
(703) 812-8730
pstiffle@charm.isi.edu

Mr. Daniel J. van Hook
Massachusetts Institute of Technology
Lincoln Laboratory
Office B-139
244 Wood Street
Lexington, MA 02173
(617) 981-4153
Fax (617) 981-0202
vanhook@ll.mit.edu

Institute for Defense Analyses

Dr. Peter S. Brooks
Strategy, Forces & Resources Division
(SF&RD)
1801 N. Beauregard St.
Alexandria, VA 22311-1772
(703) 845-2170
pbrooks@ida.org

MG Larry Budge (ret)
Simulation Center
1801 N. Beauregard St.
Alexandria, VA 22311-1772
(703) 845-6825

Dr. Randy L. Garrett
Computer Software and Engineering Division (CSED)
1801 N. Beauregard St.
Alexandria, VA 22311-1772
(703) 845-6836
rlg@ida.org

Ms. Julia A. Loughran
Computer Software and Engineering Division (CSED)
1801 N. Beauregard St.
Alexandria, VA 22311-1772
(703) 845-6681
loughran@ida.org

LIST OF ACRONYMS

ACM	Association for Computing Machinery
ARPA	Advanced Research Projects Agency
AT&T	American Telephone and Telegraph
ADS	Advanced Distributed Simulation
AIU	Advanced Interface Unit
ARP	Address Resolution Protocol
ARPA	Advanced Research Projects Agency
ASTO	Advanced Systems Technology Office
ATM	Asynchronous Transfer Mode
BBN	Bolt, Beranek and Newman, Inc.
CDR	Commander, U.S. Navy
COL	Colonel, U.S. Army
COTS	Commercial off the Shelf
DETEC	Defense Technology Evaluation Code
DIS	Distributed Interactive Simulation
DSI	Defense Simulation Internet
FFRDC	Federally Funded Research and Development Center
FIFO	First In, First Out
IEEE	Institute of Electrical and Electronics Engineers, Inc.
IP	Internet Protocol
ISDN	Integrated Services Digital Network
IS-IS	Intermediate System – Intermediate System (ISO)
ISO	International Organization for Standardization
JSTARS	Joint Surveillance Target Attack Radar System
kbps	Kilobits per second
LADS	Loral Advanced Distributed Simulation
LAN	Local Area Network
Mbps	Megabits per second
MCG	Multicast Groups

MIPS	Machine Instructions Per Second
MIT	Massachusetts Institute of Technology
N-TIED	Network-Technology Independent Encryption Device
NES	Network Encryption System
NRaD	Navy Research, Development, Test and Evaluation Division
NSA	National Security Agency
NSF	National Science Foundation
NTB	National Test Bed
ODF	On-Demand Forwarding
PARC	(Xerox) Palo Alto Research Center
PDU	Protocol Data Unit
RFC	[IEEE] Request for Comment
SAF	Semi-Automatic Forces; also referred to as SAFOR
SDIO	Strategic Defense Initiative Organization
SGI	Silicon Graphics, Inc.
SIMNET	Simulator Network
SRC	Supercomputing Research Center (IDA)
ST	Stream Protocol
STOW	Synthetic Theater of War
VMTP	Versatile Message Transaction Protocol
WAN	Wide Area Network