

not be transmitted because of a large drop in signal power at the band edges. The large drop was the result of significant phase imbalance at the input ports of the magic-T hybrid junction resulting from a large difference between the rates of change of phase with frequency, which in turn was caused by a large difference between the electrical lengths of the two TWT signal paths. To correct for this disparity in electrical lengths, it was necessary to add a dispersive circuit element to one of the paths, thereby reduc-

ing the difference between the rates of change of phase by more than an order of magnitude. This correction made it possible to transmit the 622-Mb/s signal at a very low bit error rate ( $10^{-8}$ ).

The helical TWTs used in the above demonstration have bandwidths of at least 9 GHz. By maintaining a balance of phase with changes in frequency at the input ports, it is thus possible to extend the operational bandwidth of the magic-T hybrid junction, which was observed to be at least 3 GHz, to that offered by the

inherently wide band individual TWTs.

*This work was done by Edwin G. Wintucky, Rainee Simons, and Karl R. Vaden of Glenn Research Center; Gary G. Lesny of Alphaport Inc.; and Jeffrey L. Glass of ZIN. Further information is contained in a TSP (see page 1).*

*Inquiries concerning rights for the commercial use of this invention should be addressed to NASA Glenn Research Center, Innovative Partnerships Office, Attn: Steve Fedor, Mail Stop 4-8, 21000 Brookpark Road, Cleveland, Ohio 44135. Refer to LEW-18092-1.*

## Reusable, Extensible High-Level Data-Distribution Concept

Users can optimize distributions for parallel computing, without concern for tedious details.

NASA's Jet Propulsion Laboratory, Pasadena, California

A framework for high-level specification of data distributions in data-parallel application programs has been conceived. [As used here, "distributions" signifies means to express locality (more specifically, locations of specified pieces of data) in a computing system composed of many processor and memory components connected by a network.] Inasmuch as distributions exert a great effect on the performances of application programs, it is important that a distribution strategy be flexible, so that distributions can be adapted to the requirements of those programs. At the same time, for the sake of productivity in

programming and execution, it is desirable that users be shielded from such error-prone, tedious details as those of communication and synchronization.

As desired, the present framework enables a user to refine a distribution type and adjust it to optimize the performance of an application program and conceals, from the user, the low-level details of communication and synchronization. The framework provides for a reusable, extensible, data-distribution design, denoted the design pattern, that is independent of a concrete implementation. The design pattern abstracts over coding patterns that have been found to


be commonly encountered in both manually and automatically generated distributed parallel programs.

The following description of the present framework is necessarily oversimplified to fit within the space available for this article. Distributions are among the elements of a conceptual data-distribution machinery, some of the other elements being denoted domains, index sets, and data collections (see figure). Associated with each domain is one index set and one distribution. A distribution class interface (where "class" is used in the object-oriented-programming sense) includes operations that enable specification of the mapping of an index to a unit of locality. Thus, "Map(Index)" specifies a unit, while "LocalLayout(Index)" specifies the local address within that unit. The distribution class can be extended to enable specification of commonly used distributions or novel user-defined distributions.

A data collection can be defined over a domain. The term "data collection" in this context signifies, more specifically, an abstraction of mappings from index sets to variables. Since the index set is distributed, the addresses of the variables are also distributed.

*This work was done by Mark James, Hans Zima, and Roxana Diaconescu of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP (see page 1).*

*The software used in this innovation is available for commercial licensing. Please contact Karina Edmonds of the California Institute of Technology at (626) 395-2322. Refer to NPO-42538.*



The Items Named in the Boxes are the elements involved in defining a distribution.