

entrance pupil points of interest. These are chosen to be the edge of each space, so that these rays produce the bounding volume for the beam. The x and y global coordinate data is collected on the surface planes of interest, typically an image of the field and entrance pupil internal of the optical system. This x and y coordinate data is then evaluated using a convex hull algorithm, which removes any internal points, which are unnecessary to produce the bounding volume of interest. At this point, tolerances can be applied to expand the size of either the

field or aperture, depending on the allocations. Once this minimum set of coordinates on the pupil and field is obtained, a new set of rays is generated between the field plane and aperture plane (or vice-versa).

These rays are then evaluated at planes between the aperture and field, at a desired number of steps perceived necessary to build up the bounding volume or cone shape. At each plane, the ray coordinates are again evaluated using the convex hull algorithm to reduce the data to a minimal set. When all of the coordi-

nates of interest are obtained for every plane of the propagation, the data is formatted into an xyz file suitable for FRED optical analysis software to import and create a STEP file of the data. This results in a spiral-like structure that is easily imported by mechanical CAD users who can then use an automated algorithm to wrap a skin around it and create a solid that represents the beam.

This work was done by Joseph Howard and Lenward Seals of Goddard Space Flight Center. Further information is contained in a TSP (see page 1). GSC-16176-1

Σ High-Performance, Multi-Node File Copies and Checksums for Clustered File Systems

Ames Research Center, Moffett Field, California

Modern parallel file systems achieve high performance using a variety of techniques, such as striping files across multiple disks to increase aggregate I/O bandwidth and spreading disks across multiple servers to increase aggregate interconnect bandwidth. To achieve peak performance from such systems, it is typically necessary to utilize multiple concurrent readers/writers from multiple systems to overcome various single-system limitations, such as number of processors and network bandwidth. The standard cp and md5sum tools of GNU coreutils found on every modern Unix/Linux system, however, utilize a single execution thread on a single CPU core of a single system, and hence cannot take full advantage of the increased

performance of clustered file systems.

Mcp and msum are drop-in replacements for the standard cp and md5sum programs that utilize multiple types of parallelism and other optimizations to achieve maximum copy and checksum performance on clustered file systems. Multi-threading is used to ensure that nodes are kept as busy as possible. Read/write parallelism allows individual operations of a single copy to be overlapped using asynchronous I/O. Multi-node cooperation allows different nodes to take part in the same copy/checksum. Split-file processing allows multiple threads to operate concurrently on the same file. Finally, hash trees allow inherently serial checksums to be performed in parallel.

Mcp and msum provide significant performance improvements over standard cp and md5sum using multiple types of parallelism and other optimizations. The total speed-ups from all improvements are significant. Mcp improves cp performance over 27x, msum improves md5sum performance almost 19x, and the combination of mcp and msum improves verified copies via cp and md5sum by almost 22x. These improvements come in the form of drop-in replacements for cp and md5sum, so are easily used and are available for download as open source software at <http://mutil.sourceforge.net>.

This work was done by Paul Z. Kolano and Robert B. Ciotti of Ames Research Center. Further information is contained in a TSP (see page 1). ARC-16494-1

Σ Stiffness and Damping Coefficient Estimation of Compliant Surface Gas Bearings for Oil-Free Turbomachinery

Initial applications include design of turbochargers, blowers, compressors, pumps, and turbine engines.

John H. Glenn Research Center, Cleveland, Ohio

Foil gas bearings are a key technology in many commercial and emerging oil-free turbomachinery systems. These bearings are nonlinear and have been difficult to analytically model in terms of performance characteristics such as load capacity, power loss, stiffness, and damping. Previous investigations led to an empirically derived method to estimate load capacity. This method has been a

valuable tool in system development. The current work extends this tool concept to include rules for stiffness and damping coefficient estimation. It is expected that these rules will further accelerate the development and deployment of advanced oil-free machines operating on foil gas bearings.

Foil gas bearings are self-acting hydrodynamic bearings comprised of a series

of sheet-metal foil layers from which they derive their name. They are compliant bearings that offer high-speed rotor support while accommodating shaft misalignment and distortion often encountered in turbomachinery. Lightly loaded, low-temperature foil gas bearings are commodities that predominate in the rotor support for aircraft air cycle machines (ACMs). More highly loaded foil