

Center for Hybrid Multicore Productivity Research (CHMPR)

A CISE-funded Center

University of Maryland, Baltimore County, Milton Halem, Director, 410.455.3140, halem@umbc.edu

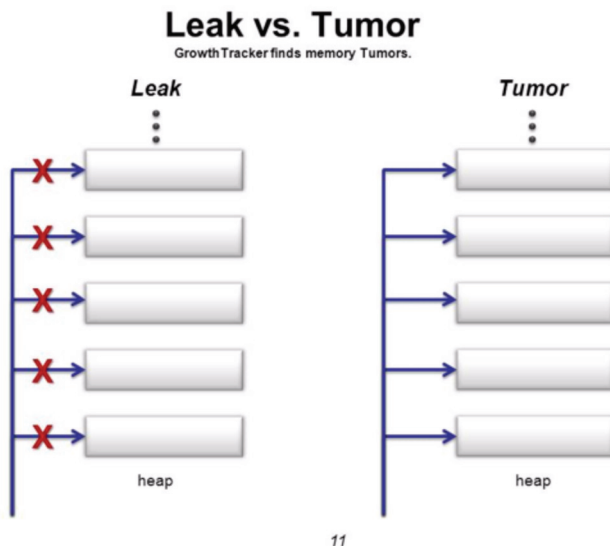
University of California San Diego, Sheldon Brown, Site Director, 858.534.2423, sgbrown@ucsd.edu

Center website: <http://chmpr.ucsd.edu/>

GrowthTracker: Diagnosing Unbounded Heap Growth in C++ Software

Robust, mission-critical software is a fundamental requirement of any nation's cyber infrastructure. This breakthrough by CHMPR researchers at the University of California, San Diego site creates essentially a set of tools that make possible robust, mission-critical execution of software by efficiently tracking use of computer systems' memory resources.

As time to market shrinks and programmers use mash up techniques to rapidly bring software systems live on the Internet, it becomes impossible to test all input conditions of a software system. This breakthrough monitors memory usage during run time executions of programs to avoid program failure and provide a software system with higher reliability and availability.



This graphic shows the difference between a memory leak, in which memory is continuously allocated without being utilized, and a memory tumor, which our tool detects where memory allocations unintentionally continue to grow.

Programs that allocate memory but do not occasionally free it up can be detected using tools known in prior art. C++ Programs where memory growth happens in unbounded fashions are detected in literature using staleness techniques - based on the last time it was accessed. The current techniques are not robust and create many false positives and false negatives. The current breakthrough is more robust and allows

Center for Hybrid Multicore Productivity Research (CHMPR)

memory growth to be tracked at an appropriate granularity to reduce the number of false positives and false negatives. This technique has been peer reviewed and published at the ICST 2013 (IEEE International Conference on Software Testing and Validation).

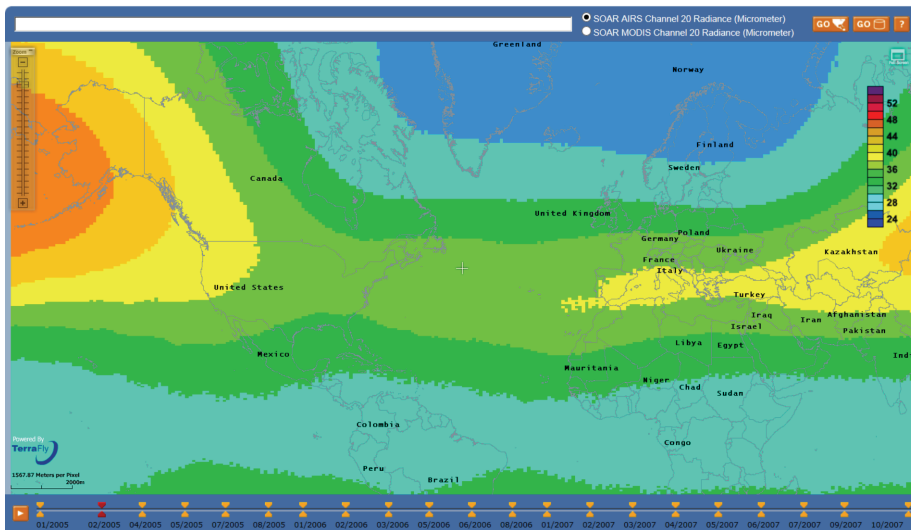
Software systems with extremely large set of C++ objects and threads where the number of objects to be created is not known *a priori*, for example, in large virtual worlds nor in large cloud-based in-memory databases. GrowthTracker can also be applied to embedded systems with very small memory footprints where efficient memory management is critical for system availability, like mobile devices.

Economic Impact: As the number of cores increases exponentially in datacenters and mobile devices, efficient techniques for memory management are becoming increasingly important. In today's increasingly complex computing systems, the reliability and scalability of systems becomes critical. This work greatly improves the robustness and reliability of software. We have deployed this approach to find heap management errors in popular software such as Google's Chrome web browser, Apple's WebKit (Safari Browser), the Ogre3D rendering engine, and the Bullet physics simulation toolkit. We submitted fixes to all of these software systems, greatly improving their stability. The economic impact on widespread adoption of GrowthTracker will likely have significant impacts on the stability of software systems, ranging from increases in productivity gained by software not crashing, to preventing catastrophic failures of software systems whose memory usage aberrantly grows over time in manners that were heretofore very difficult to detect. Not detecting aberrant growth eventually causes systems to crash. We call this growth a "tumor" as it is analogous to the biological phenomenon. We disseminated our approach to identifying and solving these software engineering problems in publications and conference presentations. This allows industry to learn from and be influenced by the methodology. It also helps industry to prepare for new challenges in scale and efficiency.

For more information contact Sheldon Brown at the University of California, San Diego, 858.534.2854, sgbrown@ucsd.edu.

Distributed Cloud Computing: 3-D Visualization Services for Climate Data on Demand

This breakthrough results from collaborations between two IUCRCs, the Center for Hybrid Multicore Productivity Research (CHMPR) at the University of Maryland, Baltimore County (UMBC) and the Center for Advanced Knowledge Enablement (CAKE) at Florida International University (FIU) and Florida Atlantic University (FAU). See the “Center for Advanced Knowledge Enablement (CAKE)” on page 27 for more information.




Satellite imagery enables precise measurement of global temperatures. This image presents the 8-year global average surface temperature (as Brightness Temperature colors higher on the scale are representative of warmer temperatures); by comparing successive average surface temperatures, global temperature changes can be detected.

Measuring the surface temperature of the entire earth on a daily basis is a difficult challenge because 75% of the planet is covered with oceans and ice. Continuously determining, for several days to weeks, the vertical thermal (i.e., temperature) field around a hurricane surrounded by dynamically rotating clouds is needed for more accurate landfall predictions. Thus, for applications ranging from climate change to hurricanes, satellites measure the earth's emitted infrared radiation twice daily with sufficiently high spatial and spectral (related to the spectrum) resolution to provide an estimate of vertical profiles of regional or global surface brightness temperature (BT). However, in order to assess global warming, these temperatures need to be measured to within an accuracy of 0.10 °C per year since models indicate CO₂ warming of ~20-30 °C over 100 years. Moreover, to resolve the structure around hurricanes, infrared data at resolutions of 1-5 km are needed. Not until 2002, when the Aqua (Latin for water) satellite was launched, has there been a single satellite with instruments that can meet both the accuracy and the precision required.

Center for Hybrid Multicore Productivity Research (CHMPR)

This breakthrough work makes it possible to deliver a decade of 3-D animated visualizations of spectral infrared (IR) satellite radiance data from instruments on Aqua. These animations use 3-D to show the vertical structure of a decade of global and regional temperature trends occurring at the surface and lower troposphere. In addition, the algorithms developed by CHMPR have been providing CAKE with 3-D temperature profiles that specify the thermal structure around hurricanes in order to improve their landfall prediction.



Atmospheric temperature layers up to 20,000 meters (65,619 feet). The vertical axis shows the height above sea level. The coldest (blue) and hottest (red) points in the eye of the hurricane are shown. The horizontal axes show the location of the hurricane (latitude and longitude).

CAKE and CHMPR have implemented a distributed cloud computing web-based service, called SOAR. This service incorporates visualization as a public service available on a multi-core IBM-based server cluster. This system provides researchers and students with the ability to select regional and chronological periods and automatically transform IR orbital satellite data into spherical grid arrays of 3-D temperature profiles for viewing the continuous changing thermal structure of the atmosphere. The FIU site at CAKE added value to the satellite data visualization by providing spatiotemporal (i.e., space-time) visualization and animation of the data (<http://cake.fiu.edu/SOAR>) using the FIU TerraFly Geospatial Data Management Service (<http://terrafly.com>). The FAU site at CAKE developed tools for 3-D visualization of the vertical temperature profiles; when coupled with CHMPR's data-gridding techniques, this partnership has created the first integrated, scientifically-validated, multi-year infrared brightness temperature record.

Economic Impact: Fundamental Decadal Data

Records are highly desired products recommended by the National Academy of Science/National Research Council. The SOAR distributed cloud computing web-based service enhances NASA's ACCESS program by providing fundamental brightness temperature records. This can go a long way towards improving scientific and public understanding of the nature of global and regional climate change. As a result, everyone can be better positioned to design policies and actions for mitigating negative climate impacts on the economy, which could include billions of dollars of property value lost to sea-level rise and billions of dollars of insurable losses due to increases in extreme weather-related disasters.

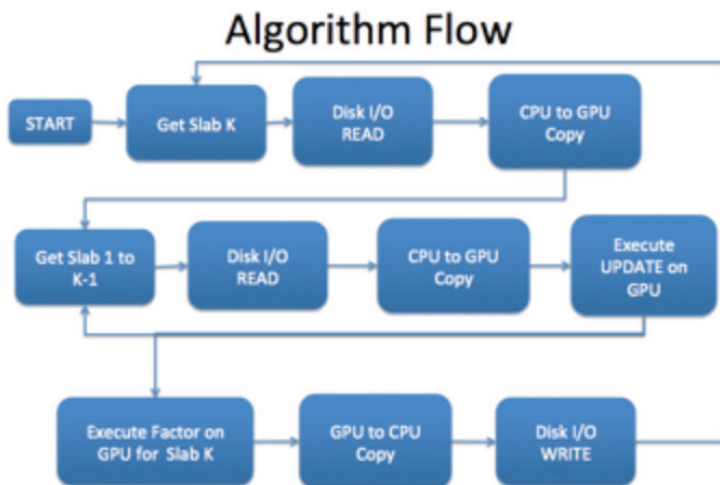
For more information, contact Milton Halem, 410.455.2862, halem@umbc.edu, or Naphtali Rishe, <http://cake.fiu.edu/Rishe>, or Borko Furht, 561.297.3180, borko@cse.fau.edu.

Specialized Graphic Processors

Until this year, supercomputers were based on tens of thousands of commodity processors like Intel and AMD multicore chips with 2 to 8 processors found in ordinary personal computers (PCs). These PCs contain specialized graphics cards that use hundreds of processors on their chips to render animations for games, simulations, and videos that are very fast and cheap. The graphics chips (GPUs) have evolved software and hardware that can not only do more than graphic renderings but can also perform complex floating point arithmetic.

Lockheed Martin, a CHMPR member, supported a project at UMBC to study and test the performance of these GPUs when added to commodity based clusters. The company wanted to know whether such GPUs can accelerate the performance of the solution of a system of equations with more than a million unknowns.

Such problems lead to enormous matrices of 1 million by 1 million terms or more than 30 Terrabytes (32×10^{12} or 32 million million), well beyond the capability of any computer to hold all these data internally in memory. Thus, this data intensive problem requires continuous moving of data from disks in and out of memory so that the processors can compute on them and then store them back on disks for future operations. It requires that all of the operations work in parallel. The method chosen for solving such equations is known as Gauss elimination and for implementation uses a transformation of the matrix into lower and upper triangular forms for direct and very fast solutions. These problems are commonly used in economics, chemistry, computer science, physics, and engineering.



The algorithm performing operations in disk IO, CPU, and GPU.

Even with high speed interconnects, disks, and CPUs, the solution time for 1 million unknowns exceeds 25 days on a single multicore commodity chip. As a test case for Lockheed Martin, this project used two sys-

Center for Hybrid Multicore Productivity Research (CHMPR)

tems to perform timing tests. One system was based on their Cray computing node with an AMD chip and an Nvidia GPU. The other system used the CHMPR computing node with an Intel chip and also an Nvidia GPU processor. A key result obtained with the additional graphic co-processor added to the system (the Nvidia GPU) was a reduction in clock time for solving a problem of 40,000 unknowns from 5 hours of solid computing to 40 minutes. Further studies indicate that the potential exists for reducing this time to less than 2 minutes when more recent available GPUs are used, combined with solid-state disks.

Other government sponsors such as the NOAA/National Center for Environmental Prediction, which is responsible for operational weather and climate forecasting, and The NSA/Laboratory for Physical Sciences are supporting research into the resiliency of such hardware configurations when scaling to hundreds of millions of such processors.

Economic Impact: This work exploited the extraordinary computational power of GPUs to accelerate data- and compute-intensive applications, which had not been investigated previously. Findings are being used to help improve the efficiency of computing systems. When using the capabilities developed at CHMPR for capitalizing on the parallel nature of the architecture, significant cost benefits, savings, and new performance studies are possible for many critically important real world applications. This work has made general accelerator technologies more feasible for solving large 64-bit complex valued matrices that exceed 1 million unknowns. In addition, more efficient use of accelerators such as GPUs will make possible significant reductions in cooling costs. For large production-quality computer systems the annual saving can be expected to approach a million dollars.

For more information, contact Shujia Zhou, szhou@umbc.edu or Milton Halem, 410.455.3140, halem@umbc.edu.