

DARPA Funds Development of New Type of Processor

— R. Colin Johnson, Advanced Technology Editor, EE Times

A completely new kind of non-von-Neumann processor called a HIVE — Hierarchical Identify Verify Exploit — is being funded by the Defense Advanced Research Project Agency (DARPA) to the tune of \$80 million over four-and-a-half years. Chipmakers Intel and Qualcomm are participating in the project, along with a national laboratory, a university and a defense contractor North Grumman.

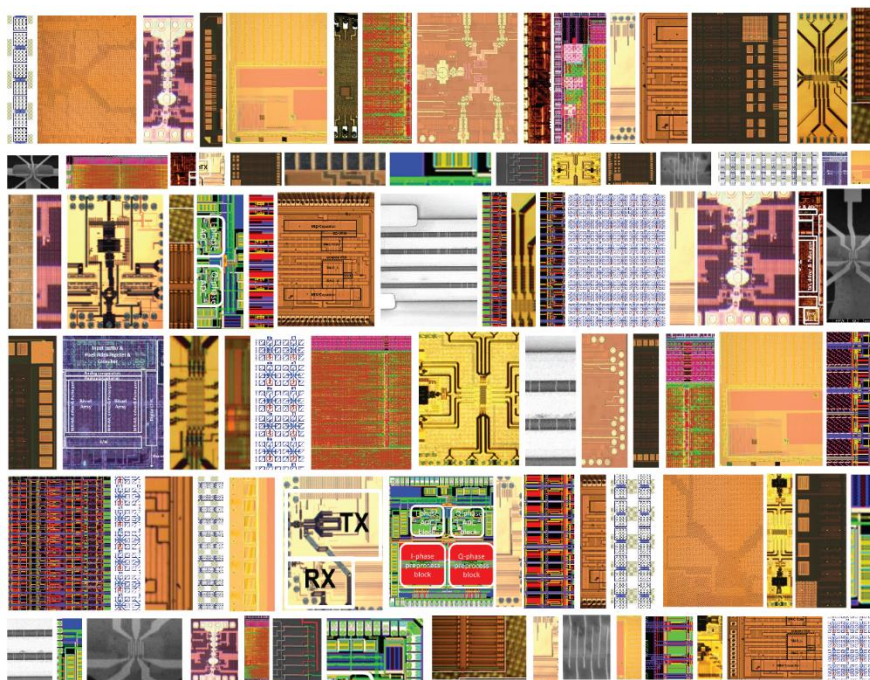
Pacific Northwest National Laboratory (Richland, Washington) and Georgia Tech are involved in creating software tools for the processor while Northrup Grumman will build a Baltimore center that uncovers and transfers the Defense Departments graph analytic needs for the what is being called the world's first graph analytic processor (GAP).



Hierarchical Identify Verify Exploit (HIVE) uses a sequence that begins with the multi-layer graphical representations of data (see figure) that open the way for graph analytic processing to identify relationships between data within and perhaps between the layers. (Source: DARPA)

"When we look at computer architectures today, they use the same [John] von Neumann architecture invented in the 1940s. CPUs and GPUs have gone parallel, but each core is still a von Neumann processor," Trung Tran, a program manager in DARPA's Microsystems Technology Office (MTO), told EE Times in an exclusive interview.

"HIVE is not von Neumann because of the sparseness of its data and its ability to simultaneously perform different processes on different areas of memory simultaneously," Trung said. "This non-von-Neumann approach allows one big map that can be accessed by many processors at the same time, each using its own local scratch-pad memory while simultaneously performing scatter-and-gather operations across global memory."



The patchwork of dies represents work funded by DARPA. "Beyond Scaling: An Electronics Resurgence Initiative" is pushing for a new era of microsystem structures and capabilities. (Source: DARPA)

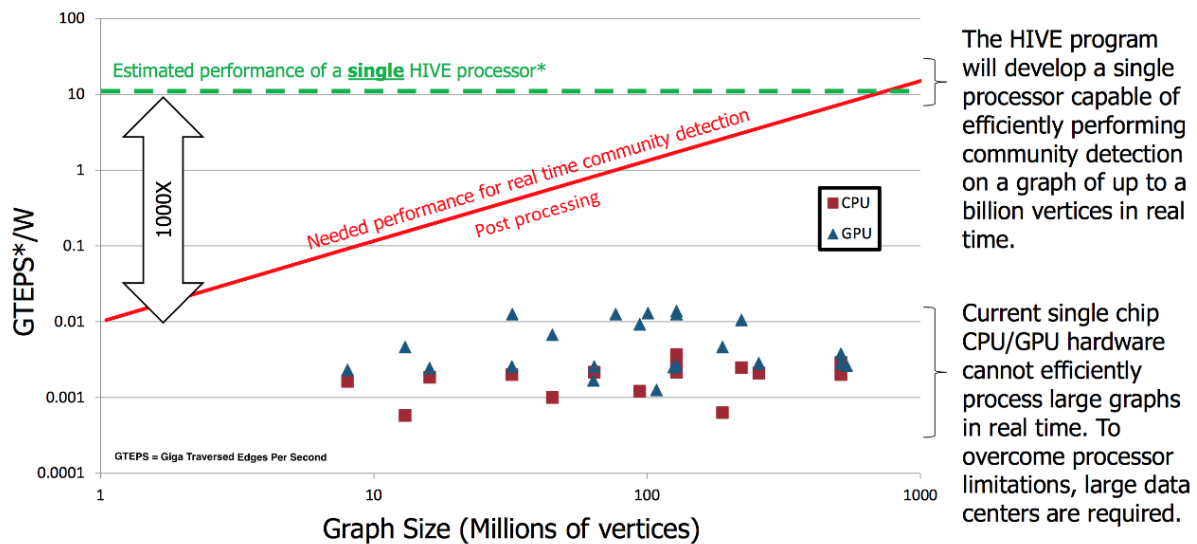
Graph analytic processors do not exist today, but they theoretically differ from CPUs and GPUs in key ways. First of all, they are optimized for processing sparse graph primitives. Because the items they process are sparsely located in global memory, they also involve a new memory architecture that can access randomly placed memory locations at ultra-high speeds (up to terabytes per second).

Today's memory chips are optimized to access long sequential locations (to fill their caches) at their highest speeds, which are in the much slower gigabytes per second range. HIVEs, on the other hand, will access random eight-byte data points from global memory at its highest speed, then process them independently using their private scratch-pad memory. The architecture is also specified to be scalable to up to however many HIVE processors are needed to perform a specific graph algorithm.

"Of all the data collected today, only about 20 percent is useful — that's why its sparse —making our eight-byte granularity much more efficient for Big Data problems," said Tran.

Together, the new arithmetic-processing-unit (APU) optimized for graph analytics plus the new memory architecture chips are specified by DARPA to use 1,000-times less power than using today's supercomputers. The participants, especially Intel and Qualcomm, will also have the rights to commercialize the processor and memory architectures they invent to create a HIVE.

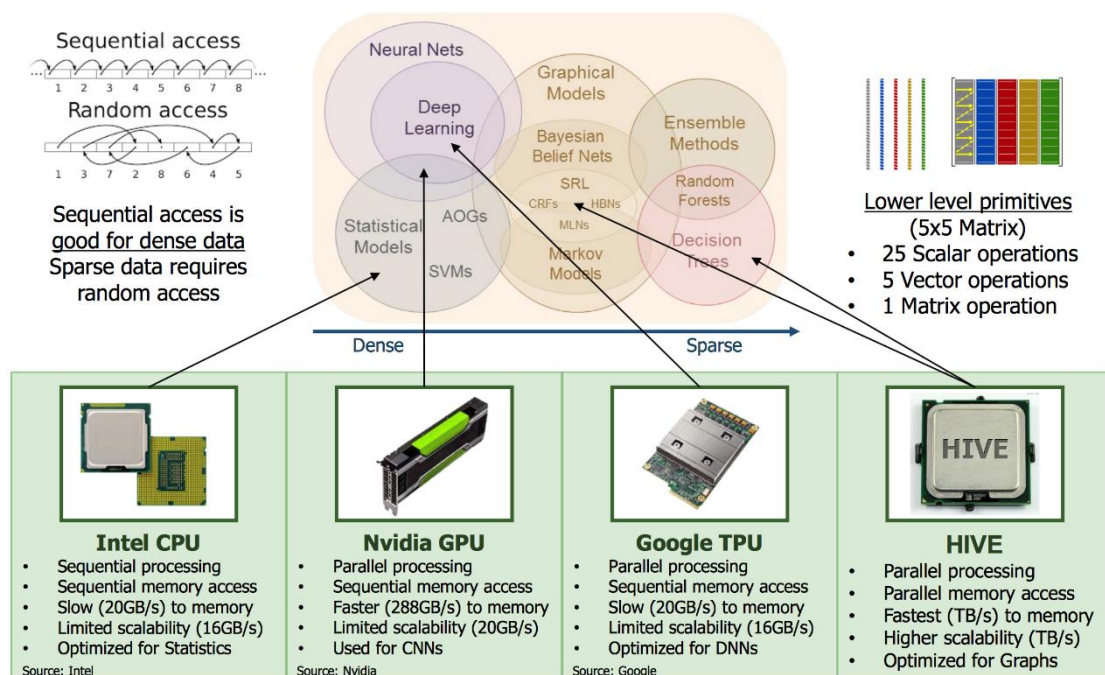
The graph analytics processor is needed, according to DARPA, for Big Data problems, which typically involve many-to-many rather than many-to-one or one-to-one relationships for which today's processors are optimized. A military example, according to DARPA, might be the the first digital missives of a cyberattack. A civilian example, according to Intel, might be all the people buying from Amazon mapped to all the items each of them bought (clearly delineating the many-to-many relationships as people-to-products).



The Giga Traversed Edges Per Second Per Watt needed for realtime graphic analysis that identify relationships as they unfold in the field is 1000-times faster (dotted line) than the fastest GPU (triangle mark) or CPU (square mark) today. (Source: DARPA)

"From my standpoint, the next big problem to solve is Big Data, that today is analyzed by regression which is inefficient for relations between data points that are very sparse," said Tran. "We found that the CPU and GPU leave a big gap between the size of problems and the richness of results, whereas graph theory is a perfect fit for which we see an emerging commercial market too."

Besides the HIVE chip, the DARPA mandate calls for the development of software tools to help programming the new architecture, which goes beyond today's parallel processing paradigm by also allowing simultaneous parallel access to random memory locations. If successful, DARPA claims that the graph analytics processor will be able to recognize and identify many types of situations that are intractable for conventional CPUs and GPUs.



Applications (top) and performance (bottom) comparisons between Intel CPUs, Nvidia GPUs, Google TPUs, and DARPA's proposed HIVE processor. (Source: DARPA)

DARPA describes its Big Data as sensor feeds, economic indicators, scientific- and environmental-measurements as the nodes of a graph, and the edges of the graph as the relationships between the nodes, such as "bought" in the Amazon example.

The basis of graph theory analytics can be traced back to the famous philosopher Gottfried Wilhelm Leibniz, but is usually attributed to the first paper on the subject, the "Seven Bridges of Königsberg" published in 1736 by Leonhard Euler. Since then it has been developed into a host of algorithms and mathematical structures that model the relationships between random data points. The HIVE architecture is designed to use these graph analytics to identify threats, track disease outbreaks, and otherwise answer Big Data questions that today are intractable for conventional CPUs and GPUs.

The four-and-a-half-year DARPA program will spend the first year with Intel and Qualcomm designing rival architectures, while Georgia Tech and PNNL design rival software tools. After the first year, one hardware design and one software design will be chosen. DARPA will provide the company with the winning hardware design with \$50 million in funding, on the condition that the company kick in \$50 million of its own. DARPA will also provide \$7 million to the organization that provides the winning software design.

Meanwhile, Northrup will be given \$11 million in non-matching funds to set up the Baltimore center to survey all of the Defense Department needs in graph analytics and make sure that the hardware and software builders meet those needs.

"HIVE is a team effort to collaborate on data handling that leverages machine learning and other AI using graph analytic processors," Dhiraj Mallick, vice president of the Intel's Data Center Group, told EE Times.

Confident that Intel will beat out Qualcomm with the winning chip design, Mallick continued: "Intel has been asked to provide a 16-node platform at the end of the program using 16 HIVE processors on a single printed circuit board. Intel will also have the rights to productize versions for the worldwide market."

The resulting HIVE processor will enable realtime identification and awareness of strategic assets as situations unfold. Whereas today we have to depend on after-the-fact analysis "closing the barn door after the horse has been stolen," Mallick said.