

Optimizing System Compute Density for Deployed HPEC Applications

Randy Banton, Director, Defense Electronics Engineering
Mercury Computer Systems, Inc.

rbanton@mc.com

Richard Jaenicke, Director, Product Marketing
Mercury Computer Systems, Inc.

rjaenicke@mc.com



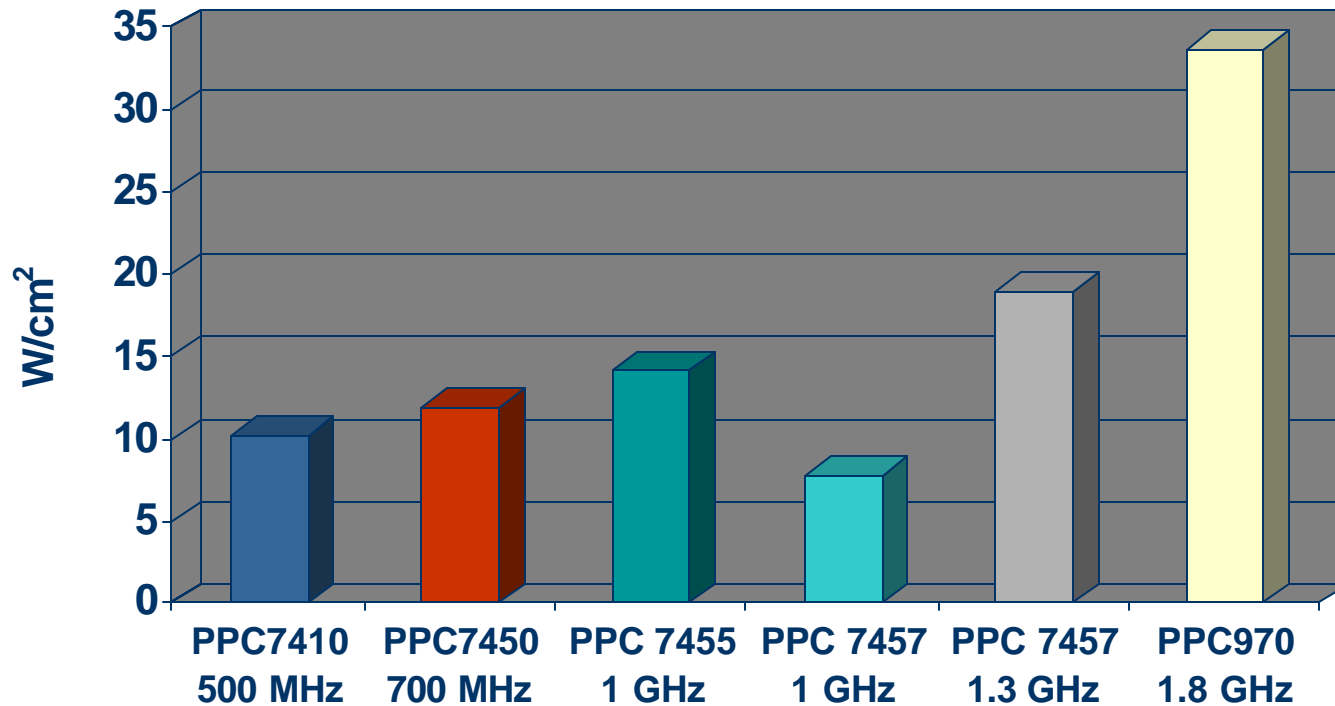
The Ultimate Performance Machine

Objective

- † Many high-end deployed military and commercial HPEC applications share a common need to achieve high to very high compute and bandwidth density in the smallest possible volume.
- † In addition, deployed military applications layer on additional environmental requirements such as higher levels of shock, vibration, endurance vibration, temperature, and condensing humidity. Each of those adds constraints on the solution space for maximizing compute and communication density.
- † Not all HPEC applications can use the same solution due to varying limits on total size or weight and varying levels of ruggedness. For example, the size and weight requirements differ greatly for manned surveillance aircraft, large UAVs, and small UAVs.
- † This poster presentation explores two of the many key dimensions to achieving high functional-density systems:
 - † **Thermal management**
 - † **Board real-estate utilization**

Processor Power Flux Trends

The power flux (W/cm^2) of the PowerPC processor continues to increase.



Data taken from publicly available sources such as data sheets, physical measurements, and articles

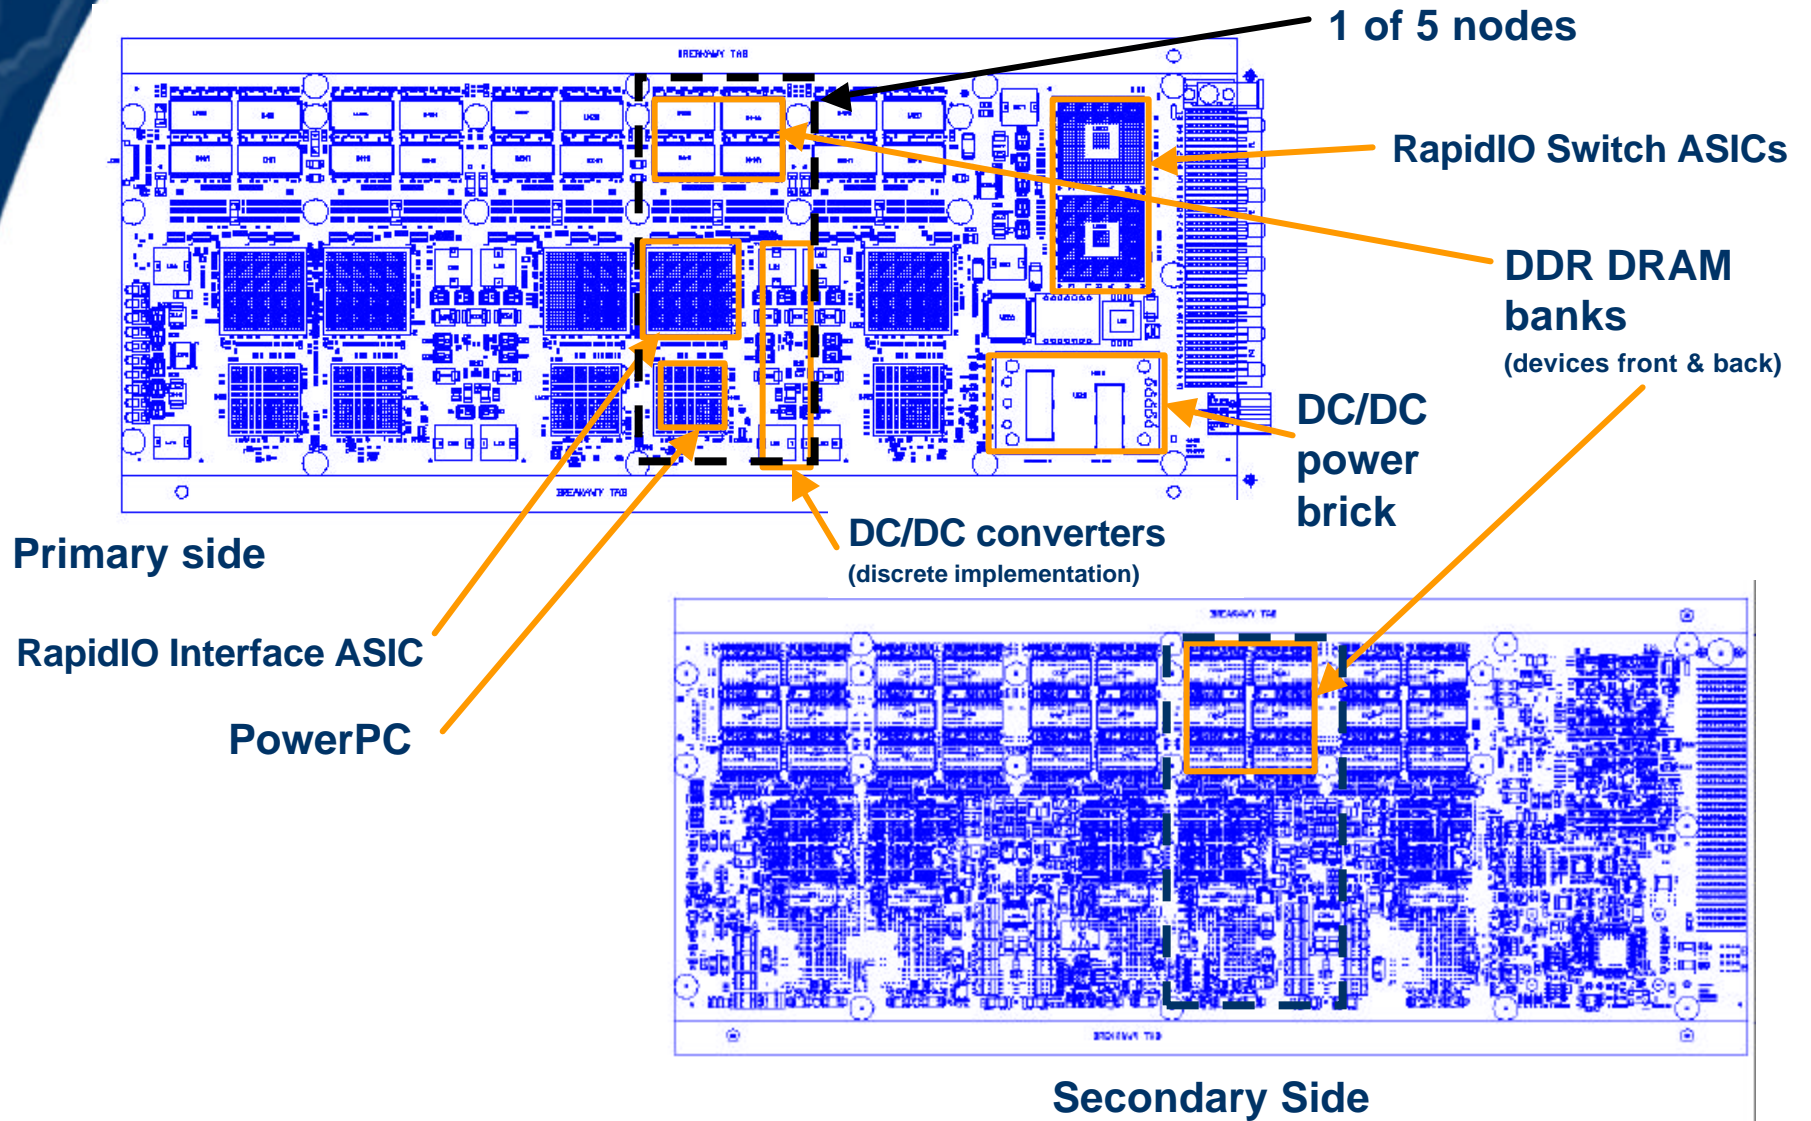
Similar increases are being seen in FPGA compute elements.

1) Thermal Management Study

- † **Non-processor power is increasing, too**
 - † In addition to the increasing power flux of processors, non-processor components in these COTS designs must be thermally managed like never before.
 - † While the PowerPC has been ~80% of a processing nodes total power, it is now closer to 50-60% with the balance of the power dissipation coming from other components of the node; memory, control and interconnect ASICs, FPGAs, and DC/DC power converters.
- † **Modeling the future**
 - † The next few slides contrast the past with the present and future in the areas of thermal management. These CFD modeling results are from a customer-driven study with results correlated against actual similar hardware under operational test.
 - † For visual simplicity, these are modeled as straight-thru air-flow. In a rack-mount system with front-to-back flow where air is taking a 90° turn in and a 90° turn out, resulting temperature may be 10-20+C higher without proper chassis level management.

Processor Board Under Study

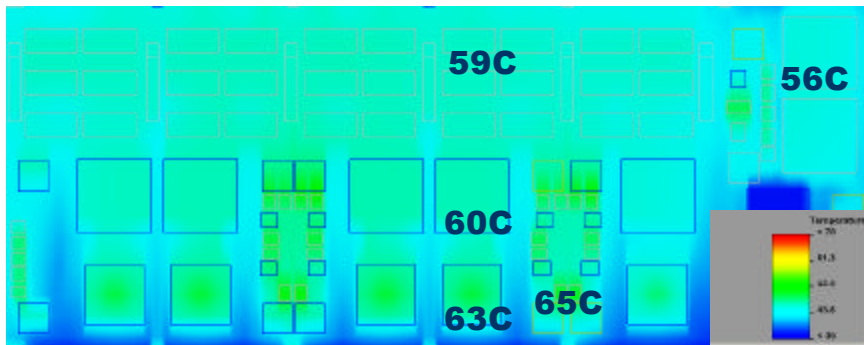
Very high routing and component density




Past Techniques on Past Processors

Past thermal management solutions were fine for commercial designs and had headroom for extending designs to MIL deployed

Models for 8W PowerPC



Commercial: 35C inlet; 5k feet with “coarsely” managed air



Military: 55C inlet; 10k feet with “coarsely” managed air

The past:


- † “Life was good” for commercial environments...
- † And, headroom was available for moving to MIL-deployed environments such as 55C concurrent with 10K altitude

Ideal flow shown in this study – temperatures may increase an additional 10-20+C in a non-optimized chassis.


Past Techniques on New Processors

New processors push even the commercial designs to meet spec and make MIL-deployed derivatives unattainable using past techniques

Models for **20W PowerPC**



Commercial: 35C inlet; 5k feet with “coarsely” managed air



Military: 55C inlet; 10k feet with “coarsely” managed air

© 2003 Mercury Computer Systems, Inc.

Today's problem:

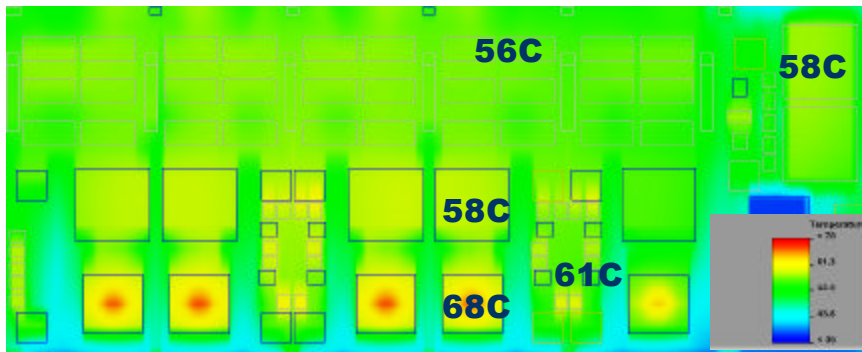
- ↑ Even with each processor getting directed air at the inlet temperature (i.e., no preheating), the new generation of processors have little to no margin against their T_j or T_c max temperatures in MIL deployed environments
- ↑ New solutions are required!

Ideal flow shown in this study – temperatures may increase an additional 10-20+C in a non-optimized chassis.


Solution: Finely Managed Airflow

“Finely Managed” air at the board level explicitly shapes and tunes airflow with precise control over impedance drop and impedance distribution over the boards surface

Models for **20W** PowerPC



Commercial: 35C inlet; 5k feet with “finely” managed air



Military: 55C inlet; 10k feet with “finely” managed air

© 2003 Mercury Computer Systems, Inc.

New solution:

† To allow use of the new generation of processor, memory, ASIC components - new airflow shaping and management techniques are required to achieve similar temperatures to the past commercial designs

† These same techniques enable use of these next-generation components in MIL-deployed environments with only a little headroom

Ideal flow shown in this study – temperatures may increase an additional 10-20+C in a non-optimized chassis.


Finely Managed Air Flow at the Chassis Level

Future:

- † More aggressive use of these techniques throughout the chassis can result in even more benefit
- † *If these improvements weren't achieved, processing and memory frequencies would be reduced by 1/2 to 1/3 – i.e., 400 to 600 GFLOPS (peak) at the chassis level*

Model inputs:


- † Processor: 20W PowerPC
- † Other devices w/ representative power
- † MIL Deployed: 55C inlet, 10K alt.
- ✧ Finely managed airflow, more aggressive throughout the chassis



2) Board Routing and Component Density

- † Aside from thermal management, use of board real estate has also become a huge challenge to functional-density.
- † The mechanical features such as board stiffeners, rugged heat-sink mounting, and such for MIL-rugged, air-cooled boards takes precious inches² of board space. MIL requirements for board- and system-level endurance vibration and shock pulses drive many of the rugged structural requirements.
- † The holes, pads, and keep-outs for these items was now causing significant loss of density – for both internal routing and components.
- † Example: Due to routing and placement restrictions, using conventional techniques yielded one less processor node on the board format under study, or 24 less nodes at the chassis level. Thus approximately 250 GFLOPS (peak) would be left on the table using conventional techniques.
- † Innovation for forced air-cooled boards was required and resulted in putting thermal management and structural rigidity features in the Z-dimension via a cover mechanism ...

Multi-Purpose Cover Increases Precious Board Space




† The cover provides:

- † Rugged mounting points for heat sinks
- † Rugged structural members that the board mounts into
- † Facilities for finely managed air features
 - In the cover surface itself
 - On the inlet
 - On the outlet
 - At the interface to the card-cage
- † Finely controlled air plenums are created at a slot-by-slot level at the board and card-cage structures

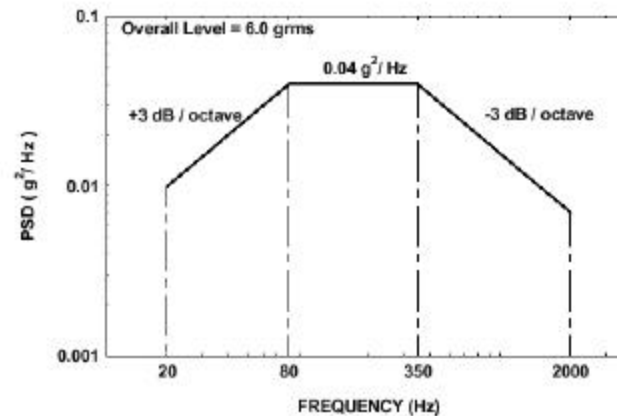
These features can be tuned as platform and application requirements drive needs

Cover Structure Has Improved Performance







Performance of full cover using the Z-dimension is similar to or better than other structural methods that take up significant board space and are often at odds with the direction and amount of airflow required.

Evaluation done using
NAVMAT P-9492 "Willoughby"
—
A typical random vibration
test for components and PWBs



Comparison Among Other Alternatives For This Module Format



Typical goal for this style of module is:

- † $125\text{Hz} < f < 150\text{Hz}$
- † $0.025 < d < 0.040$

A Glimpse at a System Solution

- † For a total system solution, solving these present day board-level thermal and structural management isn't the whole issue, chassis-level functional density must complement these solutions.

- † The *balanced HPEC TFLOP* under study, in a small volume chassis, would need all of this supporting cast in addition to the processors:
 - † >100 GB total memory


 - † >50GB/s aggregate and bisection backplane interconnect
 - † 10-20GB/s aggregate and bisection inter-chassis interconnect

 - † 10-20GB/s concurrent external I/O
(e.g., streaming sensor data over fiber, VITA-17.1)
 - † 2-3 dozen open standard I/O slots or sites
(e.g., IEEE 1386.1 PMC, VITA-42 XMC)

A Chassis-Level View

Two processor modules mounted vertically sharing the same air intake to minimize vertical space.

- † This picture represents the results of one study which determined a means to meet these high functional-density requirements (e.g., “a balanced TFLOP”) in a MIL-deployable HPEC system.
- † This arrangement achieves two leading edges of inlet air to enable high-density without having to have “columns” of the hot components, such as processors, heating each other.
- † From the previous thermal study presented, it is obvious that the heating effects of these new generation components won’t allow column organizations used in the past.



Processor module under study

RapidIO backplane interconnect w/ >50GB/s
aggregate and bisection bandwidth

Approx 21” H x 23” D,
19” rack mountable

Patent pending

Investment in Innovation

- † Mercury has invested in innovations toward solving this new class of problems in the COTS MIL-deployed, high-density HPEC application space.
- † To date, 8 patents related to the IP for these methods described have been filed. Five of these filings have "notice of allowance" which is the last step to patent registration. The remaining 3 filings are in the last stages of the PTO process.
- † The IP represented by these patents will appear in Mercury's future COTS MIL-deployable forced air-cooled HPEC products under development.