



**Reducing Server Power Consumption by 20%  
With Pulsed Air Jet Cooling**

**Influent Corp**

**10201 Maple Leaf Court  
Ashland, VA 23005  
804.550.3490**

**[www.influentmotion.com](http://www.influentmotion.com)**

## Executive Summary

Air has been the traditional workhorse for high power electronics cooling, but ongoing escalations in power density have made it more and more difficult for fan-driven air to satisfy cooling requirements. Simply driving more air end-to-end through a product is not viable from an energy standpoint since fan power increases with the cube of air flow rate. Also, increased fan flow requires higher fan speeds which result in higher noise levels and reduced fan reliability. These disincentives appear to be imposing the practical limits of fan-driven air cooling for servers and other high power electronics products.

Influent has developed an air jet thermal management solution that significantly extends the performance horizon of air cooling for high density electronics products. As an example, this paper shows how an Influent air jet system can be retrofitted within an existing server platform to provide additional high precision smart cooling of a 16 DIMM array. The result? Influent's air jets consume less than 1/10 of the power that the server's on-board fans would require to provide the same DIMM temperature.

Providing air jet cooling to a large 16 DIMM array in a space-constrained 1-U server is a particularly challenging thermal problem and showcases the power and unique capabilities of Influent's new air jet system. This relatively simple system solves what is likely the most difficult thermal problem in this server platform in terms of the number of DRAM chips being cooled and the difficulty of delivering an air jet to each chip while not impeding chassis fan flow. In this example, a single Influent pulsed air jet actuator drives a manifold design having 65 jet ports.

Influent's pulsed air jet systems provide unique "smart cooling" solutions such as proportional air delivery and load following (variable flow rate). This new air cooling system has the potential to:

- Advance server performance and reduce total server energy consumption by up to 20%,
- Enable related energy reductions in data center facilities,
- Free up PCB space by reducing the size and weight of heat sinks,
- Enable fan downsizing and significantly reduce fan prices for OEMs,
- Reduce fan related noise levels, and
- Provide these advantages for many other high power products, such as embedded systems, routers, industrial controls, laptops and high power sealed enclosures.

## Pulsed Air Jets

Pulsed air jets first appeared in the published literature over a century ago. Influent's innovation was to make air jet actuators small and powerful enough to enable their commercial viability for high power electronics cooling.

Also referred to as synthetic jets, pulsed air jets have been well known since at least 1878 when Lord Rayleigh<sup>1</sup> demonstrated a small pulsed air jet powerful enough to blow out a candle. In 1957 Dauphinee<sup>2</sup> demonstrated a thermal management application by using a pulsed air jet to maintain a constant temperature in an enclosure. In 1973 Mednikov and Novitskii<sup>3</sup> reported that a pulsed air jet actuator could deliver pulsed air jets to a remote location by driving a conduit. During the 1970s and 1980s Ephraim Gutmark<sup>4,5,6,7</sup> and other researchers began to study the use of pulsed air jets for cooling electronics components including high power computer components. These developments laid the groundwork for distributed pulsed air jet systems in thermal management.

The manner in which pulsed air jets generate flow and pressure is distinctly different from fans and blowers that depend on rotating blades to deliver air flow. A large part of the flow provided by pulsed air jets is created external to the air jet actuator via a process called “entrainment,” which allows the actuator to provide much more flow for its size compared to fans or blowers.

Pulsed air jets deliver pulsating air flow in the form of high velocity “air bullets.” As shown in Figure 1, air bullets are ejected from the actuator through a port when a piston or diaphragm creates an oscillating volumetric change. Air is pushed out of the port on a positive piston stroke and air is drawn back in through the port on a negative piston stroke, with the port being always open. If the piston stroke is large enough to push the air bullet more than one port diameter away, then it escapes and travels away from the port. Each new stroke of the piston fires a new air bullet.




Figure 1 - Pulsed Air Jet Operation

Delivering air in the form of a pulsed jet enables the actuator’s mechanically delivered air flow to be significantly amplified by entrainment. As illustrated in Figure 1, entrainment occurs as a result of friction between the air bullets and the surrounding ambient air. This friction creates rotating vortices (circular blue arrows) that follow each bullet, causing stationary ambient air to be entrained, or pulled into, the pulsating air jet. The size of the bullets decreases with propagation distance as the vortices transform the bullet’s energy into the increasing volumetric flow, illustrated by the growing diameter of the cone.


Once full entrainment is reached, the volumetric flow  $nQ$  is  $n$  times larger than the flow  $Q$  created by the actuator. Depending on the internal dynamic pressure of the actuator, the flow amplification factor  $n$  can range from 10 to 30. Influent's high power actuators enable these large flow amplification values and allow high flow rates to be generated by actuators that are physically much smaller than fans or blowers providing the same flow rate.

### Influent's Pulsed Air Jets

Reaction-Drive™ technology is the breakthrough solution that powers Influent's pulsed air jets. The essence of the Reaction-Drive™ technology is a new dynamic drive system synergistically combined with the power of resonance, resulting in the largest volumetric displacement in the smallest package. A functional description of Reaction-Drive™, as well as an animation of an actual design, can be seen on Influent's web site at <http://www.influentmotion.com/technology-overview.html>.

Competing pulsed air jet actuators simply cannot provide the significant pressure-flow performance and design advantages described in the following sections. These new air systems are enabled by the high power densities provided by Influent new Reaction-Drive™ pulsed air jets actuators.

### Manifolds



In addition to entrainment, Influent's air jets are powerful enough to drive manifolds and this provides a new set of design dimensions for air cooling. These new design dimensions and their resulting features are expected to extend the performance horizon of air cooling in current air cooled systems and also allow air cooling to be used in applications that were previously out of reach due to internal product space constraints.

Figure 2 – Pulsed Air Jet Manifold

The previous section explained how an actuator can create a pulsed air jet through a single port. A significant advantage of Influent's actuators is that they can drive multiple ports thereby creating multiple pulsed air jets to cool many hot spots. This capability can be extended significantly by attaching a manifold to the jet actuator and adding ports to the manifold. Figure 2 shows an example where an actuator is driving a gold manifold having four ports, with each pair of ports cooling separate heat sinks.

Manifolded air jets provide a powerful value proposition in air-cooled products by delivering high-velocity jets to hot spots within space-constrained areas, thereby reducing the load on the system fans resulting in reduced fan power.

To grasp the real power of these systems it's important to understand the synergistic results of combining Influent's high-power jets with multi-port manifolds.

1. Influent's air jet actuators generate oscillating pressure in the manifold and manifold ports convert this oscillating pressure into pulsating air jets at the ports. Since energy is conveyed through the manifold primarily in the form of pressure, there is minimal flow loss in the manifold. Fans are unable to drive manifolds since they lack the pressure required to push 100% of their flow through the thin manifolds required in space constrained products.
2. The total entrained flow of all manifold ports increases with the square root of the number of ports (this assumes that total port area is kept constant). For example, an air jet actuator that can provide 22 CFM of entrained flow through a single port would deliver a total of 70 CFM of flow through a 10 port manifold! This enables even further miniaturization of the air cooling system which is vital for fitting into space-constrained products.

Therefore Influent's air jet manifolds provide a disruptive capability that fans cannot provide: low loss transmission of high velocity air jets to difficult hot spots throughout a product chassis.

## Meeting Thermal Management Challenges in Servers

Influent's pulsed air jet manifolds enable a new paradigm of air cooling referred to as "*proportional air delivery*" which has the potential to significantly reduce the power that fans consume to cool servers. Fans used for heat removal in servers must be able to provide both flow and pressure. The fan must provide enough pressure to overcome the flow resistance of all the hot components as the fans drive the air from one end of the chassis to the other. The air flow rate required is determined by providing enough air velocity at each hot component to maintain target temperatures. The power consumed by these fans is the product of air volume flow rate and total pressure drop.

Higher server processing performance results in more heat which requires higher air velocities to remove that heat at the component level. To provide the extra cooling, fans must run faster and will consume more power, generate higher noise levels and experience reduced life and reliability. Increased power consumption is especially challenging since fan power increases with the cube of the flow rate. As such, server OEMs are beginning to experience the practical limitations of air cooling.

Pulsed air jet manifolds deliver air directly to the hot spots, thereby avoiding the large pressures required to pull air across the entire chassis. This substantially reduces the air flow rate and pressure that the system fans must provide, which significantly reduces system fan power and fan prices for OEMs. Manifolds driven by pulsed air jets are the only way to provide proportional air delivery with the resulting power savings sought in high density electronics products like servers.


As server designs change, the thermal bottleneck can shift from one component to another. The flexibility of Influent's air jet manifolds allow air jets to be delivered wherever the thermal problem exists in the chassis be it CPUs, CPU chip sets, power supplies or voltage regulators for example. Some OEMs have reported that cooling memory is now more challenging than CPUs in their architecture.

CPU thermal mitigation techniques are most often the top priority of packaging engineers, but memory module cooling skills are increasingly becoming almost as important. Thermal management problems present significant design challenges within embedded environments, requiring knowledge, precision, and creativity to satisfy memory subsystem thermal requirements.

In the past, memory was not as complex and did not warrant the same level of thermal attention that designers paid to the CPU. Processors and chipsets always required cooling and so came equipped with heat sinks as a production standard. Specialized fan-sink combinations, aftermarket coolers, liquid loops, and Peltier device usage have become more common. By comparison, memory modules only needed minimal airflow and on occasion minor adjustments to keep temperatures within the manufacturer's recommended range. However, with the advent of DDR3 and DDR4 technology, increasing speeds and higher heat dissipations demand solutions beyond simple fan-driven airflow.

Clock speed, higher voltage, the addition of fully buffered (FB) chips, hotter inlet air, and board location are just a few of the problems that are facing today's memory packaging engineers, and the problem appears most acute in 1-U servers. In response, DRAM (Dynamic Random Access Memory) manufacturers have increased the allowable case temperature of DRAMs from 85°C to 95°C. Aftermarket DIMMs often have heat spreaders and large aluminum heat sinks attached to the chips to minimize the temperature.

Although these techniques have been somewhat successful, they usually only reduce the chip case temperature by less than 5°C. Increasing the system fan speed is another solution, but many server fans are already running at over 15,000 RPM and suffer from shortened life, high power consumption, and undesirable acoustics. In extreme over-clocked systems, DIMMs may use additional fans, heat sink/heat pipe, and vapor chamber arrangements that are many times larger than the DIMM itself. Although



these cooling arrays can be effective, they are prohibitively expensive and too large for systems such as 1-U servers. In some applications, where air flows past the processors first and then the DIMMs, the upstream processor temperature is increased because the large downstream DIMM heat sinks increase the pressure drop, thereby reducing overall air flow volume. For example, the heat sink shown in Figure 3 uses approximately 22 times the volume of the DIMM and like most large aftermarket DIMM heat sinks, is only beneficial if the air is flowing perpendicular to the DIMMs. For a server having 16 DIMMs, this solution would add 2.48kg (5.47lbs) to the chassis weight and is far too large to fit within a 1-U server.

**Figure 3 - Thermaltake V1R**

The following sections describes a design concept for cooling a 16 DIMM array in an Intel server and provides an example of how the unique properties of Influent's pulsed air jets can be used to provide additional cooling in high density products at greatly reduced power consumption.

## CFD Study 1: Jet Cooling of Large Server DIMM Arrays

Influent Corporation has developed small high-power pulsed air jet actuators that provide the flows and pressures needed to drive large multi-port manifolds. A manifold and pulsed air jet design was recently modeled with commercially available CFD software in an Intel server to demonstrate the additional cooling that can be provided to the server's 16 DIMM array.

In this design the high velocity transient air jet pulses are used to disrupt the thermal and velocity boundary layers that form on the surface of the DRAMs. These boundary layers act to insulate the object from the beneficial heat transfer caused by convection. Disrupting the boundary layer introduces highly turbulent air flow much closer to the DRAM surface which significantly improves heat transfer from the chip surface and results in much lower case temperatures.

Figure 4 shows the solid model of an Intel SR1560SFHS 1-U server used in the CFD study. This model has two 130W processors and sixteen 10W DIMMs. Five Sanyo Denki 9CR0412S501 dual rotor fans deliver air across the main board. In this model, the storage units are Western Digital Velociraptor 300GB drives which dissipate about 6.1W each. There is also a 600W power supply shown.




Figure 4 - Intel 1560 1U Server CFD Model

The server CFD model was first run in its stock condition to establish the server's baseline thermal performance when the fans are running at full speed and the air jets are turned off. Figure 5 shows the resulting velocity contours on a data plane cut through the chassis at 0.5U. Roughly  $0.0392\text{m}^3/\text{s}$  (83.1CFM) of  $35^\circ\text{C}$  air is drawn into the server through the front bezel openings and crosses three storage drives. A divider separates the air into two paths. One air path of about  $0.00673\text{m}^3/\text{s}$  (14.3CFM) cools the 600W power supply, and a larger volume of air is diverted to cool the main board components.

This larger flow volume is split again. Four fans are used to channel air into a duct which contains two 130W processors, sixteen 10W DIMMs (Dual In-Line Memory Module), the 25W Northbridge chipset, and 25W of other support components. An additional fan cools the motherboard components that are not within the ducted flow region. The ducted fan flow volume is approximately  $0.026\text{m}^3/\text{s}$  (55.1CFM).




Figure 5 - Server showing velocity contours

The four fans consume 52.8W total. The  $35^\circ\text{C}$  inlet air temperature rises to  $37.6^\circ\text{C}$  after absorbing heat from the storage drives and the four fans. Each of the two 130W processors has a copper heat sink resulting in a maximum case temperature of  $80.7^\circ\text{C}$  for processor 1 and  $77.6^\circ\text{C}$  for processor 2. The manufacturer's recommended maximum CPU case temperature is  $85^\circ\text{C}$ . After absorbing heat from these components the ducted air entering the DIMM stack has reached  $47.5^\circ\text{C}$ . After the air has absorbed 160W of heat from the DIMMs, the ducted air exiting the chassis is  $54^\circ\text{C}$ . The average cooling air in the DIMM region is therefore about  $51^\circ\text{C}$ . The DRAMs in this model are rated for a case temperature of  $95^\circ\text{C}$ , which means that the maximum allowed DRAM temperature rise is  $47.5^\circ\text{C}$ . Figure 6 shows the temperature contours in the ducted flow area.





Figure 6 - Server Temperature Contours

In this model each DIMM has eight 0.84W DRAMs on one side along with a 2.3W fully buffered (FB) chip. On the opposite side are 9 chips dissipating a total of 1W. The model indicates that although the average case temperature is 87.2°C, which is below the 95°C maximum case temperature, one DRAM reaches 101.1°C, which is over the manufacturer's recommended maximum case temperature.


Influent designed a small high performance pulsed air jet system to cool the DIMM array. The overall dimensions of the air jet actuator are 60x60x10mm. The piston drives two compression chambers and each compression chamber is connected to one of the two manifolds. The manifold allows the actuator to be located remotely from the DIMM array. Each manifold runs perpendicular to the DIMMs and has 8 smaller runners that extend between every other DIMM. One manifold runs along the leading edge of the DIMMs while the other manifold is located at the trailing edge of the DIMMs. Each runner has 5 ports that match up with selected DRAM chips. Restriction of the chassis fan airflow is minimized by locating the manifold and runners along the surface of the motherboard. The actuator and manifold assembly are shown in Figure 7. In the final model, the runners for DIMM slots 1, 9 and 16 were removed since these slots receive fan air that has much lower preheat and higher air velocities, as suggested in Figures 5 and 6. Removing these runners reduces the pulsed jet actuator power required by providing jets only to the hottest DIMMs.



**Figure 7 - Pulsed Jet Actuator mounted to DIMM Assembly**

The actuator piston provides 5mm strokes at a 50Hz operating frequency. Jet flow rate increases with operating frequency and higher frequencies can provide much better thermal performance. However, 50 Hz was chosen for the present air jet design to provide low noise operation. Piston oscillations alternately create an oscillating pressure in the leading edge manifold. When the overpressure condition occurs, air is forced out of the runner ports at about 70m/s and impinges against the DRAM chip surfaces, disrupting the boundary layer and increasing the heat transfer coefficient. When the under pressure condition occurs, air is pulled back into the manifold through the ports. The trailing edge manifold works in the same way, but its pulses are 180° out of phase with those of the leading

edge manifold. Figure 8 shows the air jet velocity vectors during a positive pulse. The actuator provides continuously variable capacity. Changing the actuator drive voltage will cause the piston displacement to change, which in turn varies the flow rate of the jets. Variable flow rates offer further efficiency gains wherein the drive voltage is used to vary the air flow rate in response to changing heat loads.



**Figure 8 - Pulsed Port Velocity Vectors**

Manifolds can be highly customized for optimization of specific applications. For example, in this configuration only the hottest DRAMs receive jet impingement cooling. Also note that the port delivering an air jet to the center FB chip is larger than the other ports. Further optimization of this particular manifold design is also possible. In this CFD model the port sizes are identical on each runner. In actual practice, because of the heat sinks upstream of the DIMMs, each channel between DIMMs receives a different volume and temperature of inlet fan air flow. Some DIMMs receive less air volume at a corresponding higher temperature and lower velocity than other DIMMs, as shown in Figure 6 and as suggested in Figure 5. Consequently, many ports could be reduced in size since they do not need to deliver as much flow as the ports cooling the hotter DRAMs. This re-sizing of the ports would provide further reductions in jet actuator power without reducing thermal performance. CFD provides an excellent tool for manifold optimization as it is relatively easy to change port size, port placement, manifold cross section and other parameters to optimize proportional air delivery for a given platform.

The CFD model indicates that the heat transfer coefficient of highest temperature DRAM increases from  $29.6\text{W/m}^2\text{K}$  to  $44.4\text{W/m}^2\text{K}$  when the pulsed air jets are activated. The increase in the heat transfer coefficient lowers the maximum DRAM case temperature  $19.6^\circ\text{C}$  from  $101.1^\circ\text{C}$  down to  $81.5^\circ\text{C}$ .

Figure 9 shows the before and after maximum case temperatures for each DIMM. The minimum temperature drop is  $3.9^\circ\text{C}$  in DIMM slot 16 and the maximum is  $20.1^\circ\text{C}$  in DIMM slot 5. The average temperature drop across all DIMMs is  $11.2^\circ\text{C}$ .

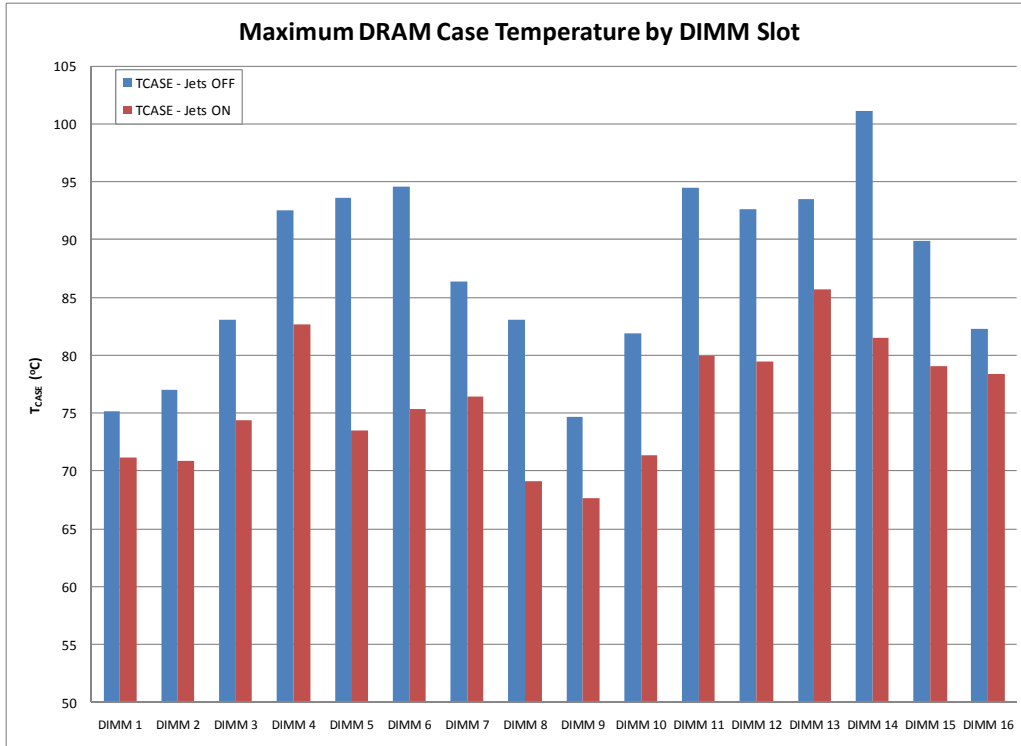


Figure 9 - DIMM Temperature Comparison

### CFD Study 2: Air Jet Power vs. Server Fan Power

A second CFD study was conducted using the same Intel SR1560SFHS 1-U server model. In this study, the pulsed air jets were turned off and the flow rate from the 4 DIMM array fans was increased until the same DIMM temperatures provided by the air jets were achieved. This study provides a head-to-head comparison of power consumption for jets vs. fans. The fans are limited to 55.1 CFM at full power, but the CFD model’s ability to increase the fan air flow rate to any level is used to do this “what-if” study.

### Results for CFD Studies 1 and 2

The tables of Figure 10 and 11 provide summary results for the 2 CFD studies.

CFD Study	Temp. of Hottest DRAM	Fan Flow Rate	Fan Power (20% efficiency)	Jet Power Draw	Total Power Draw
	(°C)	(CFM)	(W)	(W)	(W)
fans only (baseline)	101.6	55.1	19.2	0	19.2
fan + jets	81.5	55.1	19.2	8	27.2
fans only (elevated flow)	81.2	127.1	202.2	0	202.2

Figure 10 - Summary Results of CFD Study 1

Flow Rate	Pressure	Pumping Power	Fan Power (at 20% efficiency)*	DRAM Temp.	DRAM Temp. Drop	DIMMSlot Air Velocity
(CFM)	(Pa)	(W)	(W)	(°C)	(°C)	(m/s)
55.1	148	3.8	19.2	101.6	0.0	5.5
77.1	269	9.8	49.0	92.1	9.5	7.7
105.9	481	24.1	120.3	85.4	16.2	10.6
127.1	674	40.4	202.2	81.2	20.4	12.8

**Figure 11 - Summary Results of CFD Study 2**

*\*Fan power is calculated only as a point of interest. At the nominal condition, each of the four fans delivers 13.78 CFM (55.1 CFM for all 4 fans) for which the manufacturer’s PQ and efficiency chart (see appendix A) gives an efficiency of 20%. The physics of interest can be appreciated by referring only to the values in the pumping power column, which shows just the hydraulic power calculated as the product of pressure and flow. For consistency in the comparison, the 20% efficiency number is assumed for the higher fan flow rates even though the fans are not capable of providing the higher flows.*

Study 1 (1<sup>st</sup> row) in Figure 10 shows the baseline thermal performance of the server’s DIMM cooling with the jets turned off and the four server fans providing their full flow rate of 55.1 CFM to the DIMM array. Temperatures are referenced to the hottest DRAM.

In study 2 (2<sup>nd</sup> row), the jets are turned on resulting in a DIMM temperature drop of 20.1°C. For this configuration the server fans continue to provide their full flow rate of 55.1 CFM.

Study 3 (3<sup>rd</sup> row) demonstrates that with the jets off, the server fans would have to deliver a total of 127 CFM of air flow to reach the same 20.1°C DIMM temperature drop provided by the jets. In Figure 10 the fan power is based on the actual pumping power (i.e pressure x flow) and then adjusted to account for the manufacturer’s published efficiency of the Sanyo Denki 9CR0412S501 fans. As shown in Figure 11, the fans would have to produce 127 CFM at a pressure lift of 674 Pa resulting in a pumping power of 40.4 W. The Sanyo fans would have to run at a minimum speed of 36,461 rpm (or 607 Hz) to provide these flows, which would generate high noise levels at a very objectionable frequency.

The 1<sup>st</sup> CFD study makes the following important points:

- Air jets provide significant cooling in space constrained products for far less power than fans,
- To provide performance equal to the 8W jets, 202 Watts of fan power would be required,
- No fans on the market can provide this performance and fit in the 1-U server chassis.

CFD study 2 illustrates several reasons why Influent’s pulsed air jets can provide much higher thermal performance than fans for comparatively little power. The reward for delivering higher air velocities to the DRAMs is a higher DRAM heat transfer coefficient. A primary advantage of pulsed air jets is that they consume only the power required to deliver high air velocities directly to the DRAMs (or other hot

targets). In contrast, for the fans to deliver high velocity air to the DRAMs they must consume the power required to deliver that same air velocity throughout the entire DIMM air flow duct. Fans are further disadvantaged in that their power consumption is increasing with the cube of the air flow rate.

The data of Figure 11 illustrates the comparative disadvantages of fans. While the pulsed air jets are delivering 70 m/s (13,700 LFM) to the DRAM chips, the fans are delivering only 5.5 m/s of air velocity at their full design speed. Fan flow rates would have to be increased from 55.1 CFM to 127.1 CFM to provide the same DIMM temperature drops and the fan power can be seen to increase cubically with flow rate as shown in Figure 12 (“air pumping power” is the actual pressure x flow power without accounting for fan efficiency).

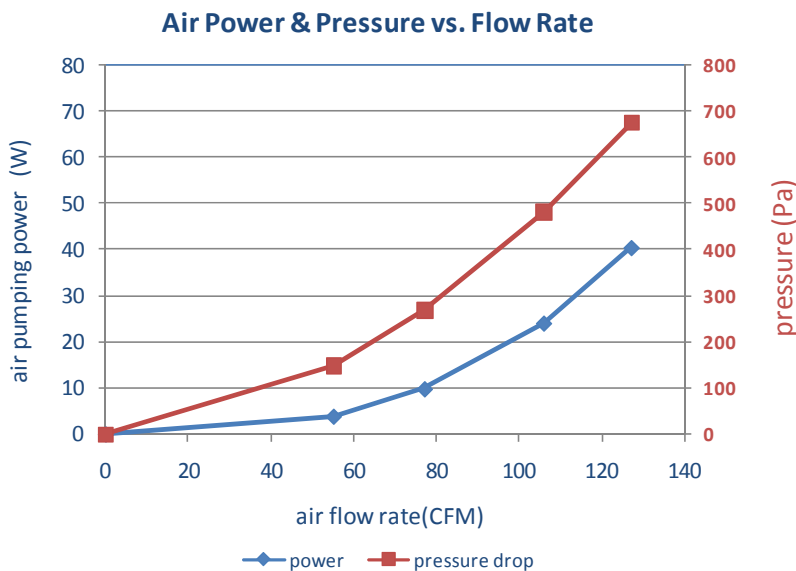



Figure 12 - Fan Power Issues

From Figure 11, the fan air velocity required to match the air jet thermal performance is 12.8 m/s. This air velocity is lower than the jet velocity but provides the same thermal performance since it is applied to the entire DIMM surface area (i.e. DRAMs and pcb surface). However, the fan power required to drive a 12.8 m/s air velocity throughout the DIMM duct makes this approach practically impossible.

In summary, a pulsed jet’s ability to cool with such comparatively low power consumption results from:

- (1) delivering high-velocity direct-impingement jets to a large number of hot spots (direct impingement significantly reduces the pressure-related power required to pull air across the entire platform),
- (2) using low-profile manifolds that do not interfere with the server’s fan air flow,
- (3) creating higher heat transfer coefficients due to high-velocity pulsating turbulent jet flow,
- (4) flow amplification due to jet entrainment, and
- (5) introducing a source of cooler air in the DIMM channels via the jets.

The synergistic combination of these features significantly extends the practical horizon of air cooling. However, the key to driving these systems is a high-power jet actuator that can deliver the required flow and pressure, is small enough to fit inside the product and provides the required life and reliability. Influent can develop the air jets actuators required to enable these new air systems. Additional air jet performance gains are certainly possible for the modeled CFD design. Higher power actuators, further manifold optimization and higher operating frequencies can all contribute to higher jet velocities and flow rates with commensurate improvements in thermal performance.



Actual Size Comparisons

	Jet Actuator	Blower
<b>Size (smallest dimension)</b>	0.60 inches	5.8 inches
<b>Weight</b>	0.22 lbs	5 lbs
<b>Power</b>	8 W	400 W (0.53 HP)
<b>Cost</b>	\$30 low volume*	\$209 retail

\*does not include cost reduction of system fans which offsets jet system price

Figure 13 – Blower size required to match jet performance

Figure 13 shows the size of a blower required to provide the same DIMM temperature drop as Influent’s pulsed air jet system. The blower is an Ametek 150446E and will deliver the required flow and pressure of 127 CFM at 674 Pa (table from Figure 1, row 4). This blower was the smallest commercially available blower found by the authors that could deliver the required performance.

## Advantages of Pulsed Air Jet Cooling

### (1) Server Performance Gains and Energy Reduction

Upgrading servers to higher-performance higher-power components can be prohibited by the flow rate, power consumption, noise and life limitations of chassis fans. If additional air cooling is needed for higher performance components within a server, pulsed air jets can be used to achieve target component temperatures for a fraction of the power required by increasing the fan air flow. This approach minimizes any increase in server-level cooling power in order to provide the desired thermal performance.

Air jets can be used for all of the high power components within a server to reduce the server's total internal thermal management power consumption. A single jet actuator can be used to provide high-velocity direct impingement air jets to the high power components across the platform. The resulting reduction in system fan flow and pressure requirements would enable the use of lower power fans which would also reduce fan prices for OEMs. The above study suggests that up to 90% reductions in a server's internal fan power draw may be possible, which includes the air jet power consumption (see the change in "pumping power" in Figure 11).

Server fans can account for 10-25%<sup>8</sup> of total server power consumption. Therefore a 90% reduction in server fan power could result in total server energy reductions of up to 20%.

While performance gains and energy reduction were demonstrated separately in the CFD models, it's important to note that both performance gains and energy reduction can be achieved together in a single server system.

### (2) Data Center Energy Reduction

Influent is also exploring the implications of pulsed air jet systems for energy savings at the data center level. Current studies<sup>9</sup> report that as much as 50% of the power consumed by today's data centers is spent on cooling. With new industry standards such as ASHRAE TC-9.9, data centers are considering raising the server inlet air temperature in order to reduce the cooling energy load of the facility. However, as inlet air temperatures are increased air velocities through the server must be increased to maintain required component temperatures. Consequently, the facility energy savings associated with higher inlet air temperatures can easily be offset by the increased fan power needed to deliver the higher air velocities through the server.

As discussed above, pulsed air jets save energy by delivering high velocity air only to the hot components where it's needed. For this reason, Influent expects that pulsed air jets in servers would allow data centers to harvest the energy savings of higher inlet air temperatures, without the offset of higher chassis fan power. The resulting savings could be significant since annual energy costs alone exceeded the cost of the servers in 2008 and infrastructure and energy cost (I&E) is projected to be 75% of the server cost in 2014<sup>9</sup>.

### **(3) Advantages for all High Power Products**

All high power products face similar thermal design challenges. The same pulsed air jets advantages described above to solve thermal problems in servers can also be used in embedded systems, routers, industrial controls, laptops and providing air cooling inside of sealed enclosures. A further advantage of reducing system air flow is that lower fan speeds will reduce noise levels.

Pulsed air jets can also reduce the size and weight of heat sinks. The high jet velocities and significant turbulence increases the heat transfer coefficient of hot targets. When pulsed air jets are used to cool heat sinks, the higher heat transfer coefficients allow heat sinks to be downsized, which frees up PCB space and reduces heat sink weight. Influent's CFD studies have shown that pulsed air jets can enable CPU heat sink volume and weight reductions of 41% and 46%.

### **Developing Air Jet Systems**

Manifolds are the interface between the air jet actuator and the hot targets. As such, designing an air jet system for a given platform is primarily an exercise of customizing a manifold for that specific product. Influent plans to develop a limited number of air jet actuators having different power levels that could then be used to drive customized manifolds for different end products.

The CFD process used to design the Intel 1-U server manifold was simple and straightforward. First the routing of the manifold is designed taking care not to block existing chassis fan air flow. Next the model is run at the stock air flow condition with jets off and the temperature distribution across the DIMMs (or other targets) is examined. The temperature distribution provides a map of where jet cooling is required and this information is used to choose the manifold's number of ports, port location, port size and port direction. A model run is then completed with the manifold design and further manifold adjustments are made as needed in response to the new thermal map and resulting temperatures.

The Intel server model shown in this paper was designed with only 3 of the above iterations. While further optimizations of the Intel server manifold are possible, the intent was to also demonstrate a rapid and repeatable design process that can provide high-impact performance and is easily accessible to Influent's customers.

## References:

- [1] Theory of Sound, Lord Raleigh, Volume II, p. 217 & 333, 1878
- [2] Acoustic Air Pump, T. M. Dauphinee, Review of Scientific Instruments, Vol. 28, No. 6, 1957
- [3] Experimental Study of Intense Acoustic Streaming, E.P. Mednikov, B.G. Novitskii, Sov. Phys. Acoust., Vol. 21, No. 2, 1973
- [4] Heat Transfer Augmentation in Turbulent Impinging Jets, E. Gutmark, et al, Proceedings of the Second Symposium on Turbulent Shear Flows, Imperial College London, July 1979, pp. 9.1-9.4
- [5] The Flow of a Row of Turbulent Vortex Rings Impinging Onto A Plate, E. Gutmark, Y. Yassour and M. Wolfshtein, Proc. Symp. Turbulent Shear Flows, Univ. CA Davis, (1981)
- [6] Acoustic Enhancement of Heat Transfer in Plane Channels, E. Gutmark, Y. Yassour and M. Wolfshtein, 7<sup>th</sup> International Heat Transfer Conference, Munchen, Vol. 3, pp. 441 (1982)
- [7] Pulsating Flow Generator, Israeli Patent Number 77,461 Issued 12-27-1985, Inventors: M. Wolfshtein, E. Gutmark, Y. Yassour
- [8] Delivering Energy Proportionality with Non Energy-Proportional Systems—Optimizing the Ensemble, Niraj Tolia, Zhikui Wang, Manish Marwah, Cullen Bash, Parthasarathy Ranganathan, and Xiaoyun Zhu, *HP Labs, Palo Alto*, Proceedings of HotPower 08, San Diego, CA December 08
- [9] In the Data Center, Power and Cooling Costs More Than the IT Equipment It Supports, C. Belay, Electronics Cooling, Vol. 13, No. 1, 2007

# APPENDIX A

