System Overhead Management with Virtual Power Centers® for Biometric Applications

Kevin David Howard, Chief Technology Officer, Kevin@massivelyparallel.com
James A. Lupo, Senior Research Scientist, James@massivelyparallel.com
Massively Parallel Technologies, Inc
Louisville, CO

Introduction

Very large-scale biometric applications have, up until now, been confined to systems using field programmable gate array (FPGA) technology. Massively Parallel Technology (MPT) has analyzed the scaling of the data movement methods used to parallelize these applications and has concluded that the standard master-slave models cause enough system level overhead to greatly limit their parallel scaling performance. The low overhead, high-efficiency methods used by MPT Howard Cascades[1] are shown to scale well enough to meet the performance requirements of biometric systems requiring very large databases.

Background

System overhead is defined to be the time cost associated with the parallel environment plus the time cost associated with all required multi-node communication. This overhead can usually be subdivided into two components: the initial overhead time and the overhead growth time. Excessive overhead growth inhibits the ability to scale up the number of nodes in a system to handle required data set sizes in a specified maximum time.

Given the small amount of data required for most characteristic vectors (ie. probe data) submitted to a system for comparison purposes, the primary sources of overhead in biometric systems are time delays in distributing the characteristic vector to the nodes on the compute engine, and the agglomeration and comparison of results from all nodes to determine the best match. Typically done using a master-slave method (Figure 1), the net result is a serialization of the process, resulting in severe suppression of scalability.

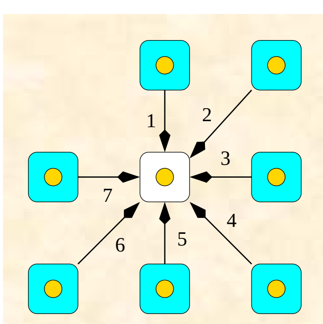

Figure 1. 1-Channel Master-Slave Agglomeration

[1] US Patent 6,857,004 B1, other patents pending.

Cascade Overhead

MPT's approach is to replace the master-slave model with multi-channel Howard Cascade models. The cascade is designed to distribute, agglomerate and compare data in a parallel fashion. This type of processing applies directly to biometric processing requirements, and significantly reduces overhead. The architectures also allow design flexibility which controls overhead while meeting performance requirements.

MPT has observed that the overhead of 1-, 2- and 4-channel Howard Cascades is a function of only the number of time-steps plus a single startup overhead cost.

Equation 1. 2-Channel Overhead

$$T_2 = O_i + O_2\varphi$$

Where: $\varphi \equiv$ the number of expansion steps
$O_i \equiv$ Initialization overhead
$O_2 \equiv$ 2-channel time-step overhead

In fact, the overhead per time step has been found to be essentially independent of the number of channels.

Equation 2. Overhead Independent of Channels

$$O_1 = O_2 = O_3 = O_4 = O$$

Thus data flow is a strong function of the number of channels, while overhead remains constant. The total overhead as a function of the number of channels and nodes is given by:

Equation 3: General Howard Cascade Overhead

$$T_C = O_i + O\log_{(\nu+1)}(N+1)$$

Where: $\nu \equiv$ # of channels per compute node
$N \equiv$ # of compute nodes

Actual results from 1- and 2-channel tests can be used to predict the overhead on 4 channels:

φ	# of Nodes			Total Overhead T_O
	1 Channel	2 Channel	4 Channel	
1	1	2	4	0.4045
2	3	8	24	0.4730
3	7	26	124	0.5415
4	15	80	624	0.6100

Table 1: Calculated Cascade Total System Overhead

The number of channels is thus a very powerful way of controlling the total overhead time in a system.

As an example, the following describes the dynamics of a full finger-print system, including the effects of database size and steps necessary to specify an MPT architecture which meets performance requirements. A complete implementation, using commercially available hardware and software demonstrates the efficacy of the cascade approach.

Figure 2 illustrates agglomeration flow in a 2-channel, depth-3 Howard Cascade. The cascade used for this study was extended to depth-4, producing a cascade containing 80 nodes.

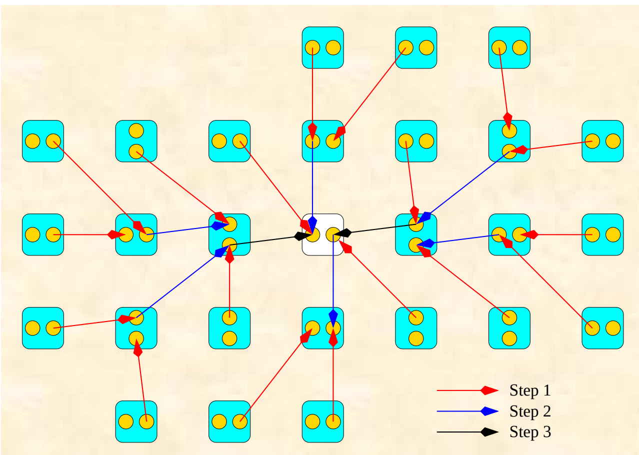

Figure 2. 2-Channel Howard Cascade Agglomeration

Each of the 80 nodes was a Dell PowerEdge 850 server with the following specifications:

1 Pentium-4 521, 2.8 GHz, 1 MB Cache
800 MHz FSB
250GB, 7,200 rpm SATA HDD
Windows 2003 Server Standard Edition
Dual On-Board NICS (10/100/GigE)
4GB of memory.
Extreme Networks Black Diamond switch.

Each node ran a dongle-protected binary-only version of the Verifinger® finger print identification product with only 3 small vendor provided modifications to its API. The binary was embedded in the MPT parallelization framework which handles all communication and work allocation. The algorithm performs a full 180° rotational comparison of the probe fingerprint against each target record. A 1.1×10^6 record fingerprint database was provided by Raytheon Corporation. To achieve worst-case times, the system was configured to do an exhaustive search of the database for every probe. The front end was a Web-based program running on a laptop computer with a USB connected commercial fingerprint reader. The web server performed a normal job submission to the cascade, and returned results to the user.

Results and Conclusions

For comparison purposes, a single node was removed from the cascade environment and run in stand-alone mode with all input and output performed local to the machine. Each record in a standard probe set, also provided by Raytheon, was applied to the system. The probe set includes known miss and match records, with matches occurring with varying, but also known, confidence levels. The location in the database of matching records is also known, though any advantage this might provide is eliminated by the exhaustive search requirement. The worst-case stand-alone end-to-end time recorded with this system was 332.24 seconds.

Using the same probe dataset, the 80-node cascade achieved a worst-case end-to-end time of just 4.5 seconds, equivalent to a parallel scaling efficiency of 92% as compared to the stand-alone result (Figure 3). Given that other MPT architectures offer even greater advantage and the ease of tuning performance to match specifications, it is clear COTS-based clusters can serve as the compute engine for biometric matching systems containing vastly larger databases.

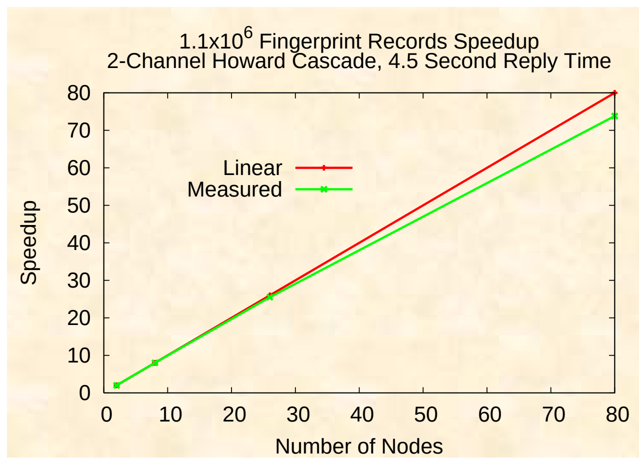

Figure 3. Worst-Case Finger-print Speedup Chart