

Questions for vIC

Introduction

What kinds of problems are caused by devices with high I/O rates?
What is interrupt coalescing?

Background

What is the VMkernel vs. the VMM in this paper's terminology?
How does interrupt processing work?
Without vIC, can any coalescing of interrupts occur?
Why is IPI useful? What performance or correctness issue(s) can it help with?
What are problems with this basic approach?

vIC Design

Both MIDL and MCC are introduced: what are they?
- Why are both useful in interrupt coalescing?
- Why can't they be used in a hypervisor setting?
The Ratio R is (num virtual interrupts in guest) / (num actual IO completions)
- What range of values can R take on?
CIF is the number of commands in flight
- Why is it useful to have a high CIF?
- What determines CIF?
What is the intuition of increasing R when CIF is low, and decreasing it when CIF is high?
There are three modes in which vIC operates:
- Low IOPS
- Low CIF
- Variable R
For which of these does coalescing occur?
Can coalescing without timeouts lead to correctness problems?
A simple approach to build vIC based on countUp and skipUp is developed; why?
What does Algorithm 1 calculate?
How is it used by Algorithm 2?
In Alg 2, for what values of countUp and skipUp are interrupts not delivered?
What does the IPI optimization do?
HOW MUCH CODE DID THEY ACTUALLY WRITE FOR THIS PAPER?

Evaluation

What is Table 2 showing?
- How much coalescing is occurring as the # of outstanding I/Os increases?
- How does the CPU utilization (per IO) change as OIO goes up?
- How is the rightmost column of Table 3 calculated?
What are Figures 3 and 4 showing us?
How much impact does vIC have for TPC-C*?
What do we learn from Figure 5?

Related Work

What is unusual about this related work section?

* not TPC-C™ compliant