Chapter 8

Conclusions

Still round the corner there may wait,
A new road or a secret gate.

- J. R. R. Tolkien

The primary contribution of this thesis is the development of flit-reservation flow control, a flow control protocol where data flits are led by control flits, and the control flits traverse the network ahead of data flits, reserving channel and buffer resources in advance dynamically. By reserving resources in advance, flit-reservation flow control is able to eliminate routing and arbitration decision latency, reducing network latency. Reserving buffers in advance also allows buffers to be reused immediately after they are released, resulting in high buffer utilization and thus high network throughput.

This thesis also develops a delay model which estimates router pipeline delay using parameterized equations derived through detailed gate-level design and analysis. With this model, router architects and designers can more realistically evaluate and compare the performance of different router microarchitectures. The model is used to demonstrate the effectiveness of a micro-architectural technique, speculation, in reducing the pipeline delay of both virtual-channel and flit-reservation routers.

8.1 Thesis summary

In this thesis, we explore flow control and micro-architectural mechanisms for extending the performance of interconnection networks. We develop flit-reservation flow control, where network buffers and channels are reserved in advance by control flits, on behalf of the data flits they lead. This prior reservation determines if data flits should be buffered or forwarded when they arrive. Reserving resources ahead of time yields two advantages.

First, control flits can be injected in advance when the destination is known before data is ready. Routing and arbitration can then be performed in advance, reducing network latency to the minimal wire delay. Second, reserving buffers in advance allows buffers to be immediately reused, unlike previous flow control methods which leave buffers idle in between usage. Immediate buffer reuse improves buffer utilization and enables flit-reservation flow control to extend network throughput.

Control flits lead data flits using timestamps that indicate the relative arrival time of each data flit. With these timestamps, control flits schedule the earliest possible departure time for their data flits. They accomplish this with the help of reservation tables, which store the status of output channels and buffers from the present to the scheduling horizon. With these tables, the control flits can identify when a channel will be available for its data flit, and when a buffer will be available at the next hop to house this data flit when it arrives. Thus, the movements of the data flits are completely orchestrated by their control flits.

By applying the router delay model we proposed to flit-reservation flow control, we derived a detailed router micro-architecture, and verified that the protocol can be implemented without excessive overhead. Simulations based on realistic pipelines affirm that flit-reservation flow control is able to hide routing and arbitration latency, reducing network latency by up to 33%, to the minimal wire delay. Also, our experiments show flit-reservation flow control achieving high buffer utilization, keeping buffers 93% full on average, as compared to virtual-channel flow control, which keeps buffers 40% full. This results in flit-reservation flow control being able to match or better the throughput of virtual-channel flow control with half the number of buffers.

The simulations are based on pipelines prescribed by our proposed delay model. Our proposed model takes as input the clock cycle time and router parameters and outputs a suitable router pipeline. Router delay is estimated using parameterized delay equations derived through detailed gate level design and analysis. We contrast our model with previous models, highlighting where previous models' assumptions differ from reality.

With this model, we show how speculation can be used to shorten the critical path of routers. By arbitrating for physical and virtual channels simultaneously, speculation enables the delay of flit-reservation and virtual-channel routers to be the same as that of wormhole routers. Thus, contrary to the findings of previous models, our model shows that virtual-channel flow control is able to improve network throughput by up to 40% over wormhole flow control, while maintaining the same zero-load latency.

Also, while the unit-router-latency assumption is frequently adopted in literature, our experiments show that the assumption results in huge inaccuracies. When compared to the realistic pipelines prescribed by our model, experiments assuming a single cycle of router latency quote 125% lower network latency and 33% higher network throughput. We hence conclude that realistic router pipelines should be used when comparing different router designs.

8.2 Future directions

As technology scales, routers will operate with higher clock rates and shorter cycle times. Existing router designers have been relying on pipelining to alleviate this design constraint, leading to routers with deeper and deeper pipelines. This thesis demonstrates the adverse impact that deep pipelines can bring -- longer pipelines result in longer network delay and lower network throughput due to slower buffer turnaround. Thus, as clocks get tighter, router architects cannot continue leveraging pipelining, or performance will suffer.

Future flow control protocols have to overcome the impact of long pipeline delays on network performance. This thesis has shown how advance reservation of network resources is able to alleviate the impact of pipeline delays on network latency and throughput. Prior knowledge allows a router to prepare and reserve passage for a message ahead of time, eliminating decision latency. Prior knowledge also allows a router to know when a buffer will be released, and reuse it immediately after it is vacated. Thus, advance-reservation architectures are required to mitigate the adverse impact of deep pipelining on network performance.

While flit-reservation flow control demonstrates how advance knowledge can be leveraged for efficient buffer and channel utilization, it leaves unanswered how much further performance can be raised. If we have complete advance knowledge of the network traffic, what is the best possible performance we can achieve? Establishing such a yardstick will allow us to measure the potential of advance-reservation router architectures.

The ever-tightening cycle time constraint faced by router architects demands new micro-architectures too. Future micro-architectures have to explore new avenues beyond simple pipelining. In this thesis, we demonstrate how speculation can be used to shorten router pipelines. Speculation enables routers to optimistically execute router functions in parallel, squeezing more functionality into a single clock cycle. Future micro-architectures will have to carry this further as cycle times get tighter.

Throughout this thesis, our focus is on network performance. While this is the predominant concern for interconnection network designers, other factors such as area and power dissipation may pose severe constraints in some applications. Thus, a realistic router model should not just provide delay estimates, but also area and power estimates. Similarly, though this thesis demonstrated the strengths of flit-reservation flow control taking into account router delay, its cost in terms of area and power remains to be evaluated. Building a prototype of a flit-reservation router will allow these to be culled accurately.

Digital systems are placing ever more demanding requirements on their network fabrics -- multiprocessors are gearing towards petaflops processing; Internet routers have to handle hundreds of terabits per second in the future; and I/O interconnects have to handle

Gb/s and even Tb/s bandwidth. This pushes interconnection network designers to continue pushing the latency-throughput envelope, amidst tighter design constraints, developing new mechanisms for extending network performance, so as to improve overall system performance.