

# On-Chip Networks: Do We Need More Research?

José Duato

Dept. of Computer Engineering (DISCA)  
Technical University of Valencia, Spain  
E-mail: [jduato@disca.upv.es](mailto:jduato@disca.upv.es)

NSF Workshop

- 1 Introduction
- 2 New Challenges and Opportunities
- 3 A Case Study
- 4 Conclusions

- 1 Introduction
- 2 New Challenges and Opportunities
- 3 A Case Study
- 4 Conclusions

# Introduction

- The trend toward multi-core processing chips is now a well established one
  - Mass market production of dual-core and quad-core processor chips
  - Trend toward massive multi-core chips based on much simpler cores (e.g. Sun UltraSparc T1, Nvidia 8800 GT/GTX GPUs)
  - Heterogeneous multi-core chips (processors plus accelerators) proposed (mostly for the embedded market)
- Beyond a certain number of cores (say, 8 to 16), an on-chip network with point-to-point links becomes necessary to interconnect cores, cache banks and memory controllers among them (and possibly with on-chip routers for external communication as well) without the constraints imposed by buses.

# Introduction

- The trend toward multi-core processing chips is now a well established one
  - Mass market production of dual-core and quad-core processor chips
  - Trend toward massive multi-core chips based on much simpler cores (e.g. Sun UltraSparc T1, Nvidia 8800 GT/GTX GPUs)
  - Heterogeneous multi-core chips (processors plus accelerators) proposed (mostly for the embedded market)
- Beyond a certain number of cores (say, 8 to 16), an on-chip network with point-to-point links becomes necessary to interconnect cores, cache banks and memory controllers among them (and possibly with on-chip routers for external communication as well) without the constraints imposed by buses.

# Introduction

- Most of the design challenges for current and future on-chip networks are very similar to the ones previously faced by off-chip network designers:
  - Selection of a suitable topology and routing algorithm
  - Definition of efficient flow control and switching techniques
  - Designing a compact and fast router
  - Designing flexible and efficient network interfaces
  - Providing support for fault tolerance
  - Plus some additional support depending on the application area: collective communications, QoS, congestion management

# The Trivial Approach

- Two decades ago, the challenge was to design efficient single-chip routers
  - New switching and flow control techniques (e.g. wormhole and virtual channels) were proposed to drastically reduce buffer sizes and be able to design routers that could fit into a single chip. Also, packets no longer had to be buffered in the host memory
  - Low-dimensional topologies were proposed to use pin bandwidth efficiently: a few wide links are preferred over many narrow links
  - Simple and efficient routing algorithms (e.g. DOR), amenable for hardwired implementations, were proposed to minimize latency
  - Wide links together with pipelined packet transmission (i.e. wormhole) and hardwired routing delivered the lowest latency
  - Virtual channels allowed blocked packets to be bypassed by other packets, achieving a reasonably high utilization of link bandwidth
- Overall, these innovations made single-chip routers feasible, increasing bandwidth and reducing latency by one and three orders of magnitude, respectively.

# The Trivial Approach

- Two decades ago, the challenge was to design efficient single-chip routers
  - New switching and flow control techniques (e.g. wormhole and virtual channels) were proposed to drastically reduce buffer sizes and be able to design routers that could fit into a single chip. Also, packets no longer had to be buffered in the host memory
  - Low-dimensional topologies were proposed to use pin bandwidth efficiently: a few wide links are preferred over many narrow links
  - Simple and efficient routing algorithms (e.g. DOR), amenable for hardwired implementations, were proposed to minimize latency
  - Wide links together with pipelined packet transmission (i.e. wormhole) and hardwired routing delivered the lowest latency
  - Virtual channels allowed blocked packets to be bypassed by other packets, achieving a reasonably high utilization of link bandwidth
- Overall, these innovations made single-chip routers feasible, increasing bandwidth and reducing latency by one and three orders of magnitude, respectively.



# The Trivial Approach

- The first microprocessors implemented processor architectures similar to those used two decades back in time. Similarly, on-chip networks can be designed now by using the techniques proposed two decades ago for single-chip routers.
  - Low-dimensional network topologies seem to be very appropriate for on-chip implementations. A 2-D mesh can be implemented by using only short links for neighbor to neighbor communication
  - Wormhole switching allows the design of very compact routers by implementing very small buffers
  - Low-dimensional topologies allow the design of very compact and fast routers, with small internal crossbars and simple arbiters
  - The use of wide links leads to very low packet latency
  - Ironically, the design constraint proposed by Bill Dally two decades ago (i.e. bisection bandwidth) has never become true for off-chip networks but will very likely become true for on-chip networks
- What is new? Do we need further research on on-chip networks?

# The Trivial Approach

- The first microprocessors implemented processor architectures similar to those used two decades back in time. Similarly, on-chip networks can be designed now by using the techniques proposed two decades ago for single-chip routers.
  - Low-dimensional network topologies seem to be very appropriate for on-chip implementations. A 2-D mesh can be implemented by using only short links for neighbor to neighbor communication
  - Wormhole switching allows the design of very compact routers by implementing very small buffers
  - Low-dimensional topologies allow the design of very compact and fast routers, with small internal crossbars and simple arbiters
  - The use of wide links leads to very low packet latency
  - Ironically, the design constraint proposed by Bill Dally two decades ago (i.e. bisection bandwidth) has never become true for off-chip networks but will very likely become true for on-chip networks
- What is new? Do we need further research on on-chip networks?

# The Trivial Approach

- The first microprocessors implemented processor architectures similar to those used two decades back in time. Similarly, on-chip networks can be designed now by using the techniques proposed two decades ago for single-chip routers.
  - Low-dimensional network topologies seem to be very appropriate for on-chip implementations. A 2-D mesh can be implemented by using only short links for neighbor to neighbor communication
  - Wormhole switching allows the design of very compact routers by implementing very small buffers
  - Low-dimensional topologies allow the design of very compact and fast routers, with small internal crossbars and simple arbiters
  - The use of wide links leads to very low packet latency
  - Ironically, the design constraint proposed by Bill Dally two decades ago (i.e. bisection bandwidth) has never become true for off-chip networks but will very likely become true for on-chip networks
- What is new? Do we need further research on on-chip networks?

# The Trivial Approach

- The first microprocessors implemented processor architectures similar to those used two decades back in time. Similarly, on-chip networks can be designed now by using the techniques proposed two decades ago for single-chip routers.
  - Low-dimensional network topologies seem to be very appropriate for on-chip implementations. A 2-D mesh can be implemented by using only short links for neighbor to neighbor communication
  - Wormhole switching allows the design of very compact routers by implementing very small buffers
  - Low-dimensional topologies allow the design of very compact and fast routers, with small internal crossbars and simple arbiters
  - The use of wide links leads to very low packet latency
  - Ironically, the design constraint proposed by Bill Dally two decades ago (i.e. bisection bandwidth) has never become true for off-chip networks but will very likely become true for on-chip networks
- What is new? Do we need further research on on-chip networks?

# The Trivial Approach

- The first microprocessors implemented processor architectures similar to those used two decades back in time. Similarly, on-chip networks can be designed now by using the techniques proposed two decades ago for single-chip routers.
  - Low-dimensional network topologies seem to be very appropriate for on-chip implementations. A 2-D mesh can be implemented by using only short links for neighbor to neighbor communication
  - Wormhole switching allows the design of very compact routers by implementing very small buffers
  - Low-dimensional topologies allow the design of very compact and fast routers, with small internal crossbars and simple arbiters
  - The use of wide links leads to very low packet latency
  - Ironically, the design constraint proposed by Bill Dally two decades ago (i.e. bisection bandwidth) has never become true for off-chip networks but will very likely become true for on-chip networks

- What is new? Do we need further research on on-chip networks?

# The Trivial Approach

- The first microprocessors implemented processor architectures similar to those used two decades back in time. Similarly, on-chip networks can be designed now by using the techniques proposed two decades ago for single-chip routers.
  - Low-dimensional network topologies seem to be very appropriate for on-chip implementations. A 2-D mesh can be implemented by using only short links for neighbor to neighbor communication
  - Wormhole switching allows the design of very compact routers by implementing very small buffers
  - Low-dimensional topologies allow the design of very compact and fast routers, with small internal crossbars and simple arbiters
  - The use of wide links leads to very low packet latency
  - Ironically, the design constraint proposed by Bill Dally two decades ago (i.e. bisection bandwidth) has never become true for off-chip networks but will very likely become true for on-chip networks

● What is new? Do we need further research on on-chip networks?

# The Trivial Approach

- The first microprocessors implemented processor architectures similar to those used two decades back in time. Similarly, on-chip networks can be designed now by using the techniques proposed two decades ago for single-chip routers.
  - Low-dimensional network topologies seem to be very appropriate for on-chip implementations. A 2-D mesh can be implemented by using only short links for neighbor to neighbor communication
  - Wormhole switching allows the design of very compact routers by implementing very small buffers
  - Low-dimensional topologies allow the design of very compact and fast routers, with small internal crossbars and simple arbiters
  - The use of wide links leads to very low packet latency
  - Ironically, the design constraint proposed by Bill Dally two decades ago (i.e. bisection bandwidth) has never become true for off-chip networks but will very likely become true for on-chip networks
- What is new? Do we need further research on on-chip networks?

- 1 Introduction
- 2 New Challenges and Opportunities**
- 3 A Case Study
- 4 Conclusions



# The Microprocessor Analogy

- Although microprocessor microarchitecture was initially based on simple accumulator-based designs and, for several years, new designs incorporated previously proposed ideas, a lot of novel research has been developed for microprocessors
- There have been even disruptive approaches (RISC vs. CISC) that dramatically changed the way processors were architected
- So, we should expect innovation in on-chip networks in the coming years, but:
  - What is the equivalent to the RISC revolution?
  - Was it the design of single-chip routers 20 years ago? Or should we expect dramatic enhancements in the coming years?
- I believe that most of the breakthrough architectural ideas have already been proposed, but there is still room for optimized designs tailored to particular requirements as well as new technologies

# The Microprocessor Analogy

- Although microprocessor microarchitecture was initially based on simple accumulator-based designs and, for several years, new designs incorporated previously proposed ideas, a lot of novel research has been developed for microprocessors
- There have been even disruptive approaches (RISC vs. CISC) that dramatically changed the way processors were architected
- So, we should expect innovation in on-chip networks in the coming years, but:
  - What is the equivalent to the RISC revolution?
  - Was it the design of single-chip routers 20 years ago? Or should we expect dramatic enhancements in the coming years?
- I believe that most of the breakthrough architectural ideas have already been proposed, but there is still room for optimized designs tailored to particular requirements as well as new technologies

# The Microprocessor Analogy

- Although microprocessor microarchitecture was initially based on simple accumulator-based designs and, for several years, new designs incorporated previously proposed ideas, a lot of novel research has been developed for microprocessors
- There have been even disruptive approaches (RISC vs. CISC) that dramatically changed the way processors were architected
- So, we should expect innovation in on-chip networks in the coming years, but:
  - What is the equivalent to the RISC revolution?
  - Was it the design of single-chip routers 20 years ago? Or should we expect dramatic enhancements in the coming years?
- I believe that most of the breakthrough architectural ideas have already been proposed, but there is still room for optimized designs tailored to particular requirements as well as new technologies

# New Design Constraints and Opportunities

- Power consumption is leading to new strategies for both on-chip and off-chip networks (DVFS, on/off links, variable width links)
- Multiple metal layers within the chip allow for optimizations not feasible outside the chip (flit-reservation flow control, link widths tailored to packet types, optimized topologies, etc)
- Process variability requires new transmission techniques to avoid having to lower the clock frequency for the entire chip
- Reducing chip-kill in the presence of permanent faults requires some network reconfiguration technique
- There are several ongoing efforts toward developing tools for design space exploration, both for on-chip and off-chip interconnects

# New Design Constraints and Opportunities

- Power consumption is leading to new strategies for both on-chip and off-chip networks (DVFS, on/off links, variable width links)
- Multiple metal layers within the chip allow for optimizations not feasible outside the chip (flit-reservation flow control, link widths tailored to packet types, optimized topologies, etc)
- Process variability requires new transmission techniques to avoid having to lower the clock frequency for the entire chip
- Reducing chip-kill in the presence of permanent faults requires some network reconfiguration technique
- There are several ongoing efforts toward developing tools for design space exploration, both for on-chip and off-chip interconnects

# New Design Constraints and Opportunities

- Power consumption is leading to new strategies for both on-chip and off-chip networks (DVFS, on/off links, variable width links)
- Multiple metal layers within the chip allow for optimizations not feasible outside the chip (flit-reservation flow control, link widths tailored to packet types, optimized topologies, etc)
- Process variability requires new transmission techniques to avoid having to lower the clock frequency for the entire chip
- Reducing chip-kill in the presence of permanent faults requires some network reconfiguration technique
- There are several ongoing efforts toward developing tools for design space exploration, both for on-chip and off-chip interconnects

# New Design Constraints and Opportunities

- Power consumption is leading to new strategies for both on-chip and off-chip networks (DVFS, on/off links, variable width links)
- Multiple metal layers within the chip allow for optimizations not feasible outside the chip (flit-reservation flow control, link widths tailored to packet types, optimized topologies, etc)
- Process variability requires new transmission techniques to avoid having to lower the clock frequency for the entire chip
- Reducing chip-kill in the presence of permanent faults requires some network reconfiguration technique
- There are several ongoing efforts toward developing tools for design space exploration, both for on-chip and off-chip interconnects

# New Design Constraints and Opportunities

- Power consumption is leading to new strategies for both on-chip and off-chip networks (DVFS, on/off links, variable width links)
- Multiple metal layers within the chip allow for optimizations not feasible outside the chip (flit-reservation flow control, link widths tailored to packet types, optimized topologies, etc)
- Process variability requires new transmission techniques to avoid having to lower the clock frequency for the entire chip
- Reducing chip-kill in the presence of permanent faults requires some network reconfiguration technique
- There are several ongoing efforts toward developing tools for design space exploration, both for on-chip and off-chip interconnects



# Tightly-Coupled Network Interfaces

One of the areas where significant progress can be made is the network interface:

- For off-chip networks, most current network interfaces require moving data through the PCI-Express interface, thus increasing latency dramatically
- For on-chip networks, extremely low latency is required
- Interprocess communication (either shared memory or message passing) should go from cache to cache, without accessing external memory
- Many processors/cores concurrently accessing the network interface: need for scalable virtualization

*We need low-latency tightly-coupled virtualized network interfaces*

# Standardization?

- At first glance, standardization of on-chip networks seems to be a stupid idea: No external interfaces; each manufacturer will use its own proprietary interconnect
- VLSI design based on reusing existing IP cores is becoming common practice
- Defining a standard interface for on-chip networks would allow a much simpler and efficient reuse of existing IP cores in future homogeneous and heterogeneous multi-core designs

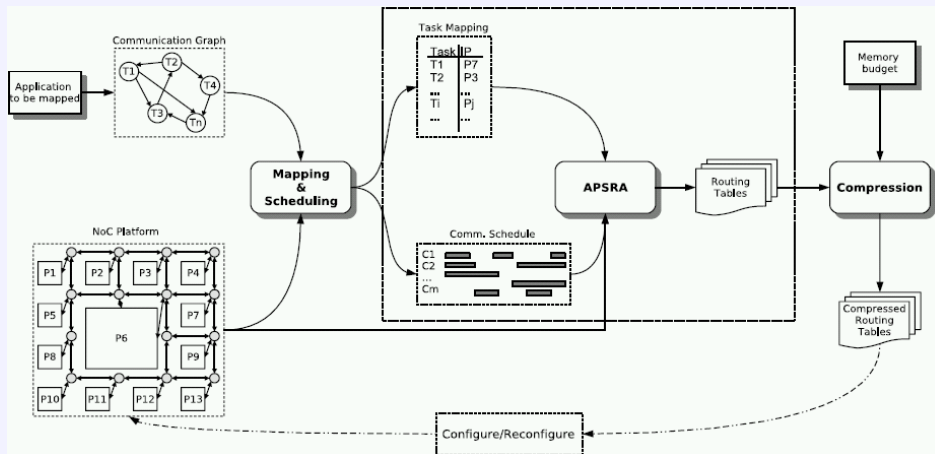
- 1 Introduction
- 2 New Challenges and Opportunities
- 3 A Case Study**
- 4 Conclusions

# Multi-core Embedded Devices

- Embedded devices usually run just one or a few applications
- General-purpose processors are not suitable for running those applications because they are either too slow or too power hungry
- Ultra low-power processors combined with hardware accelerators are the preferred choice for most designers
- Recent designs for the embedded market use multi-core processors to reduce power consumption: more efficient than just DVFS
- Future embedded systems may use a large number of heterogeneous cores to deliver the best tradeoff between performance and power consumption

# Opportunities for Optimization

- For a given application running on a multi-core chip, each process may not communicate with every other process



# Opportunities for Optimization

- Task mapping can be optimized so as to minimize communication among cores: less bandwidth requirements and power consumption
- Routing algorithms do not need to provide paths among every pair of nodes
- Routing algorithms may dynamically change according to application needs

# Required Support

- Task mapping strategies for heterogenous multi-core systems that take the communication graph into account
- Optimized routing algorithms that minimize communication latency in partially connected networks
- Network reconfiguration techniques that dynamically adapt routing algorithms, and even network topology, to application communication requirements

# What is Really New?

- Most task mapping strategies proposed up to now did not consider communication costs
- Routing algorithms that do not connect all the nodes in the network:
  - Do we allow routing through nodes that will never inject/receive a packet?
  - How much power is saved by not routing among all the nodes?
- Network reconfiguration techniques for on-chip networks:
  - How do we implement reconfigurable routing algorithms without having to implement large and slow routing tables?
  - How can we implement dynamic topology reconfiguration without introducing too much overhead?



# What is Really New?

- Most task mapping strategies proposed up to now did not consider communication costs
- Routing algorithms that do not connect all the nodes in the network:
  - Do we allow routing through nodes that will never inject/receive a packet?
  - How much power is saved by not routing among all the nodes?
- Network reconfiguration techniques for on-chip networks:
  - How do we implement reconfigurable routing algorithms without having to implement large and slow routing tables?
  - How can we implement dynamic topology reconfiguration without introducing too much overhead?

# What is Really New?

- Most task mapping strategies proposed up to now did not consider communication costs
- Routing algorithms that do not connect all the nodes in the network:
  - Do we allow routing through nodes that will never inject/receive a packet?
  - How much power is saved by not routing among all the nodes?
- Network reconfiguration techniques for on-chip networks:
  - How do we implement reconfigurable routing algorithms without having to implement large and slow routing tables?
  - How can we implement dynamic topology reconfiguration without introducing too much overhead?

- 1 Introduction
- 2 New Challenges and Opportunities
- 3 A Case Study
- 4 Conclusions**

# Conclusions

- On-chip networks are the result of increasing interest on multi-core processor chips (CMPs). Many systems can be designed just by reusing results previously proposed for off-chip networks
- Many people are reinventing the wheel. It is necessary to analyze what is really new
- Most breakthrough ideas have already been invented, but there are still many opportunities for innovation
- New application areas will impose new sets of design constraints:
  - Previously proposed generic techniques need to be adapted to the new constraints
  - Engineering process: Highly optimized designs will be required

# Conclusions

- On-chip networks are the result of increasing interest on multi-core processor chips (CMPs). Many systems can be designed just by reusing results previously proposed for off-chip networks
- Many people are reinventing the wheel. It is necessary to analyze what is really new
- Most breakthrough ideas have already been invented, but there are still many opportunities for innovation
- New application areas will impose new sets of design constraints:
  - Previously proposed generic techniques need to be adapted to the new constraints
  - Engineering process: Highly optimized designs will be required