

Exploring Garbage Collection
with
Haswell Transactional Memory Hardware

University of
Kent



KOCHI UNIVERSITY OF TECHNOLOGY

Carl Ritson
Tomoharu Ugawa
Richard Jones

2014/06/12

Intel's Haswell

Transactional Memory Extensions (TSX-NI)

Transactional Synchronisation Extensions

(Restricted) Transactional Memory

... with limited processor support

XBEGIN ... XEND

Up to ~16KiB of read and writes

Complexities

Setup of transaction is expensive (3x CAS)

Fallback required if transaction fails

Aborted transactions expensive

Uses for TSX-NI

Simplify parallel collector activities

Facilitate concurrent bitmap marking

Accelerating concurrent copying collection


Concurrent Copying Collection

Sapphire on-the-fly concurrent copying collector

[Hudson and Moss, 2001 & 2003]

Start


from-space




**Mark
Phase**

from-space


to-space




**Copy
Phase**



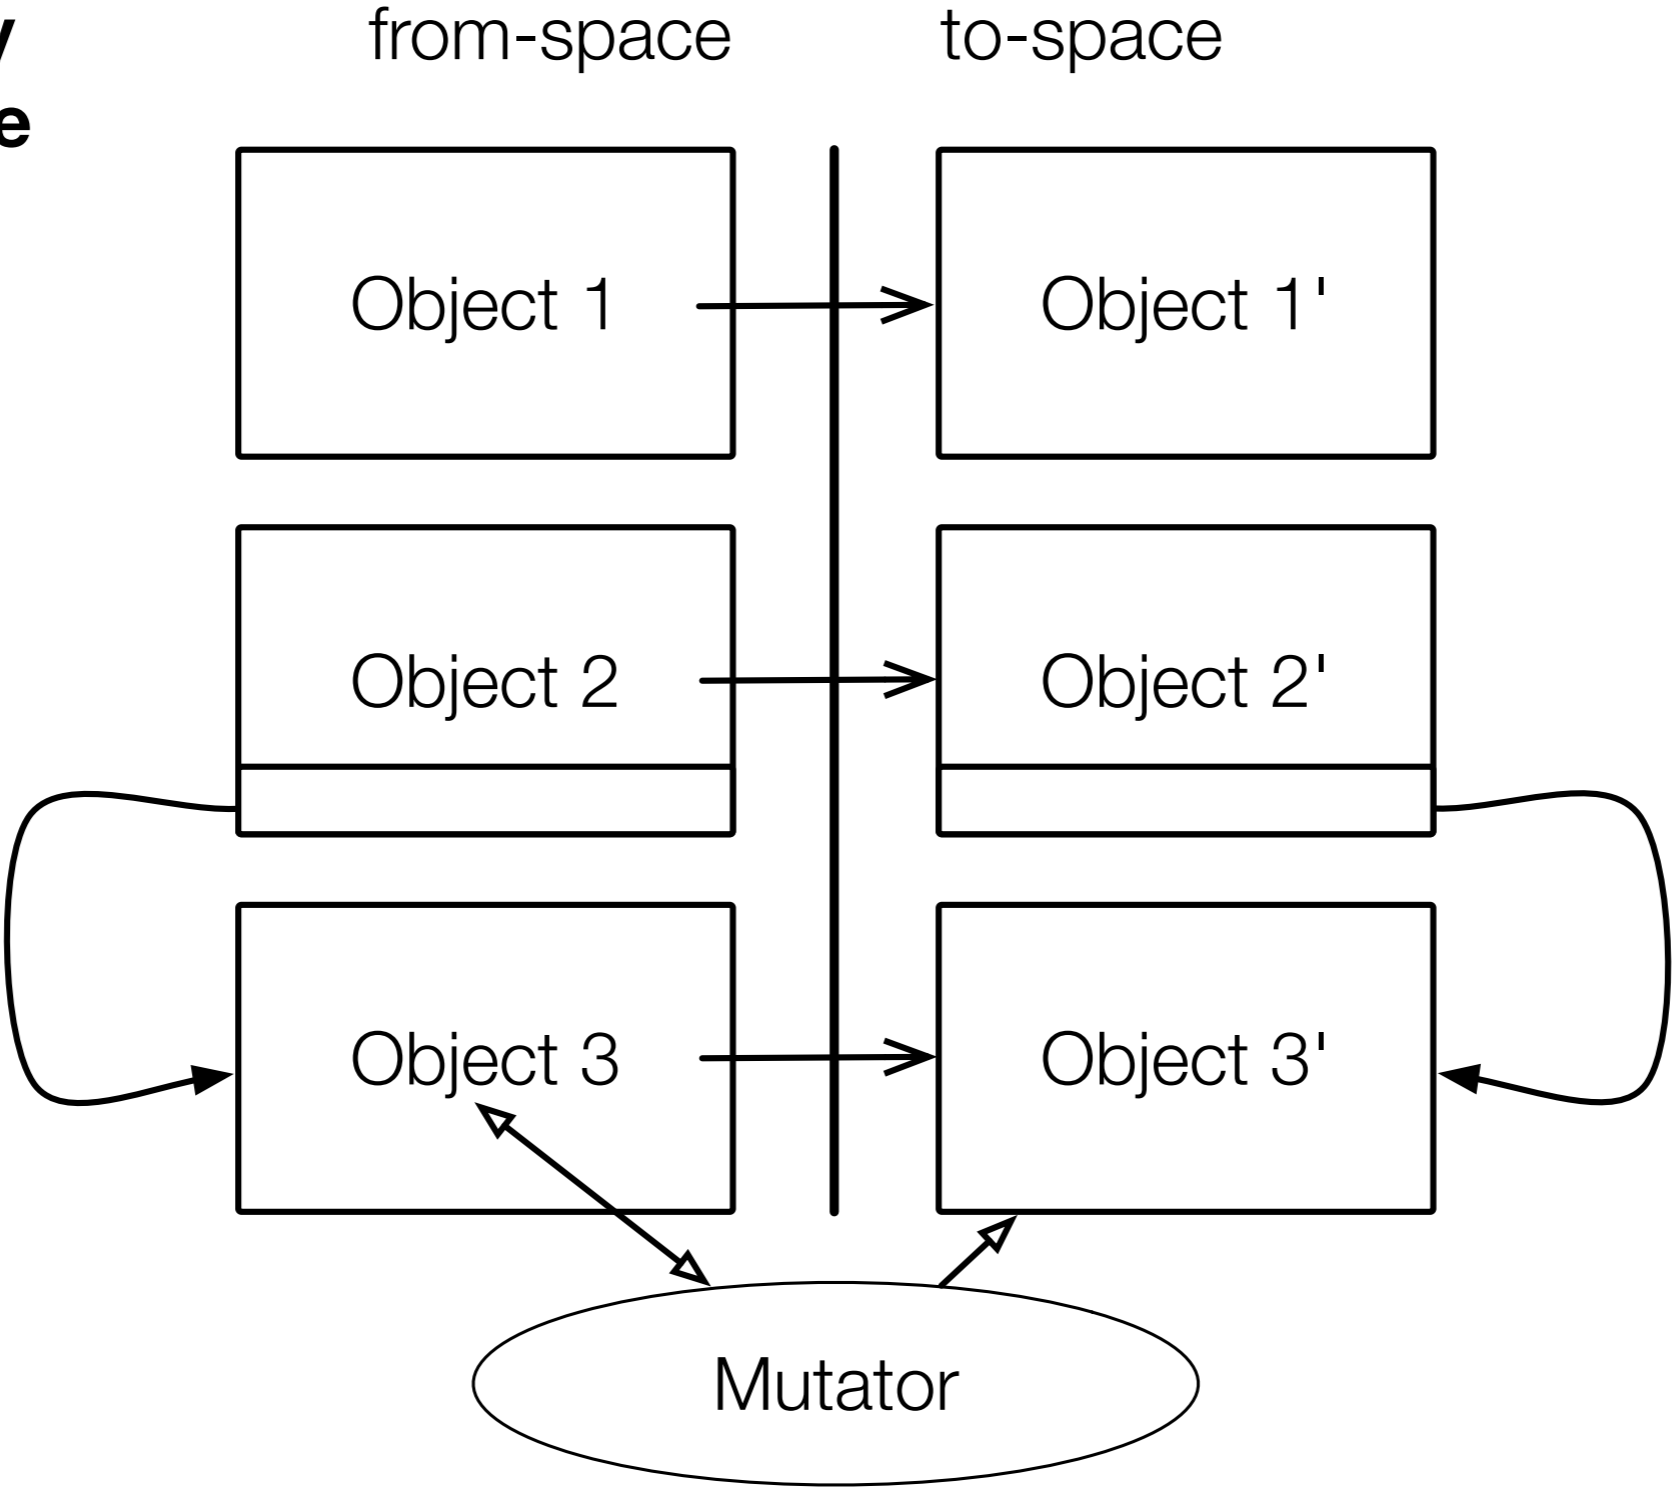
**Copy
Phase**



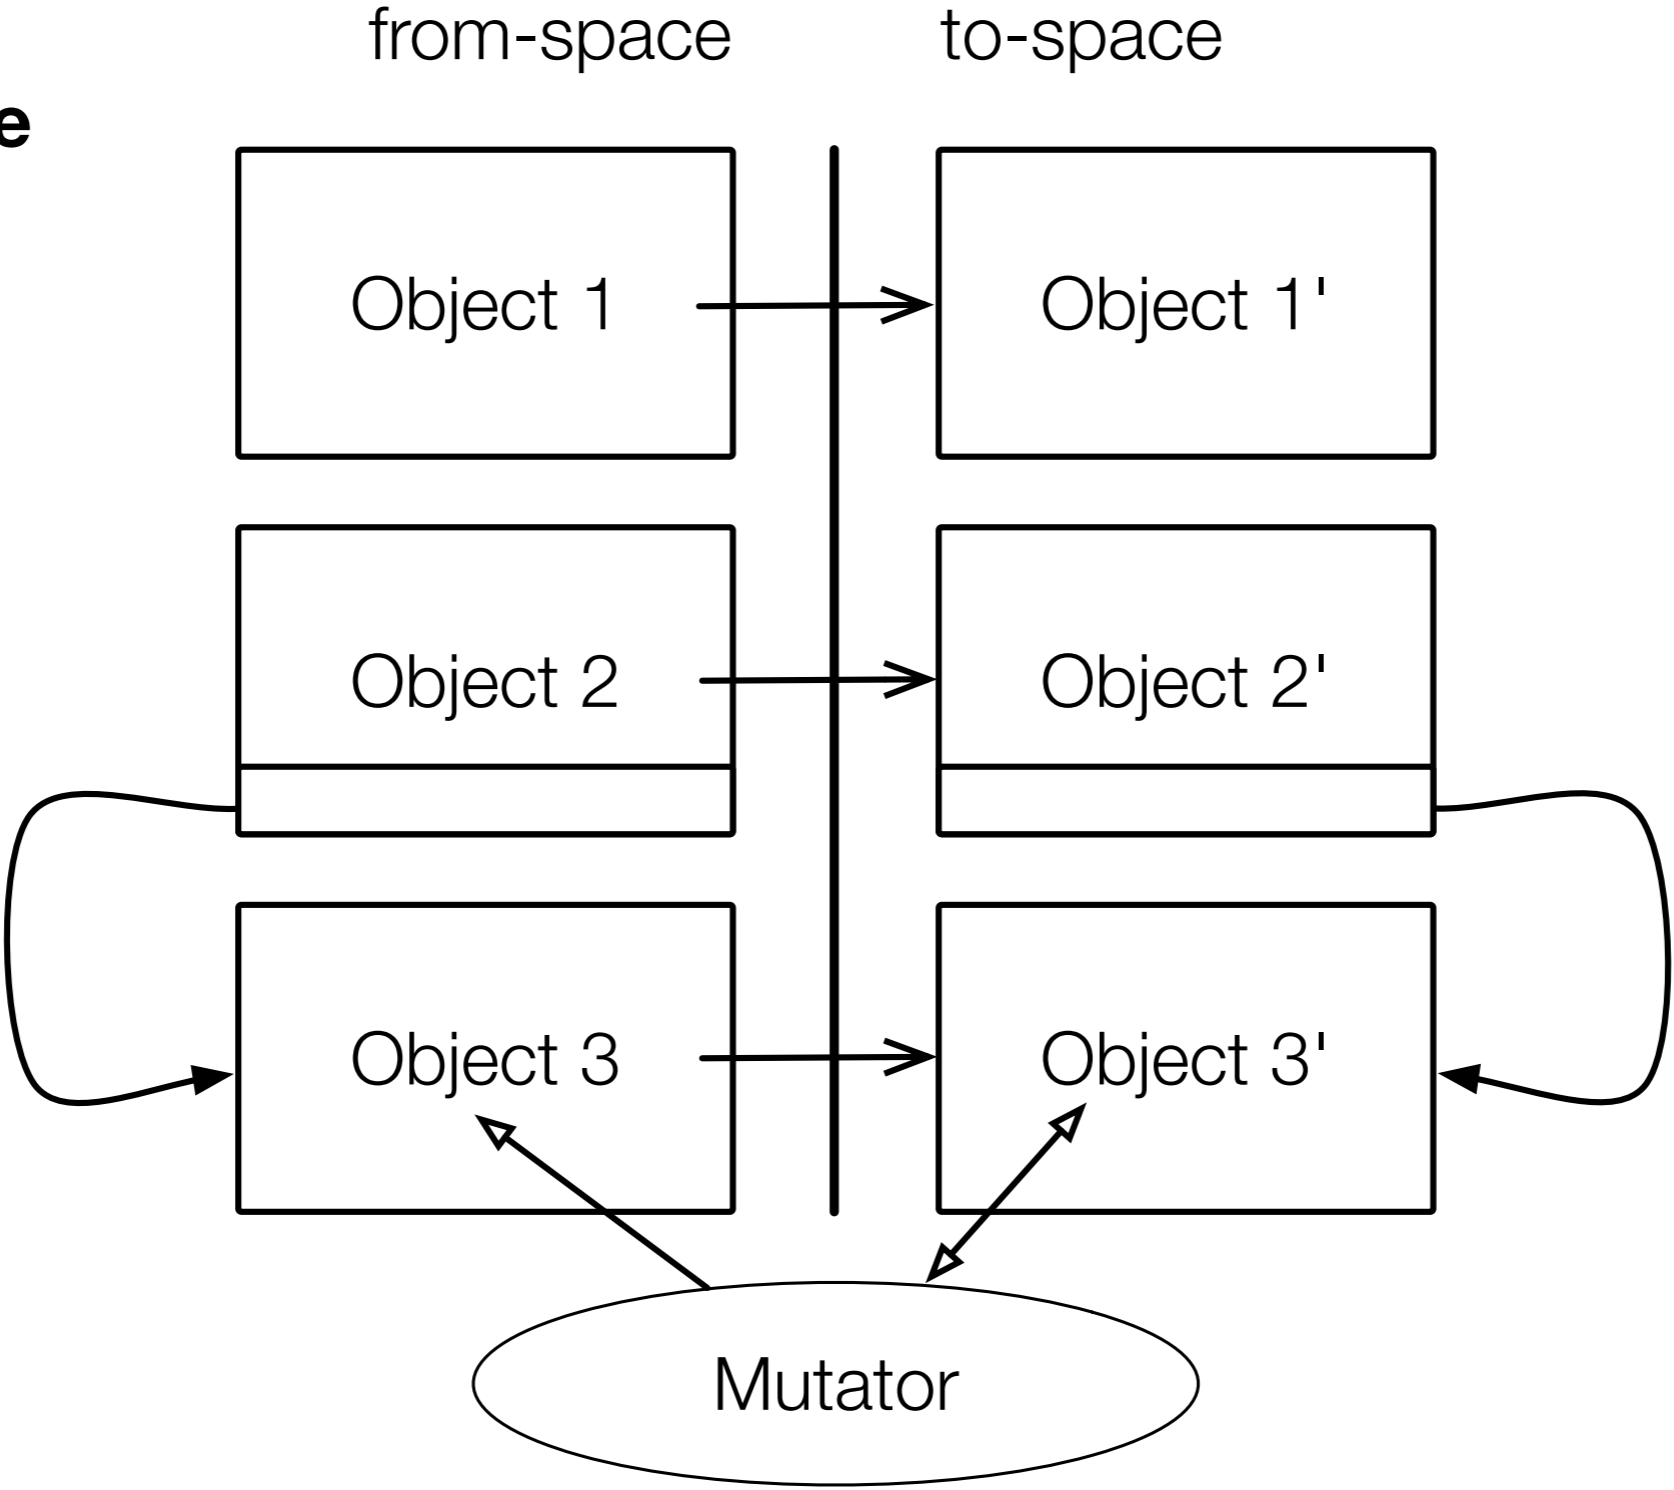
**Copy
Phase**



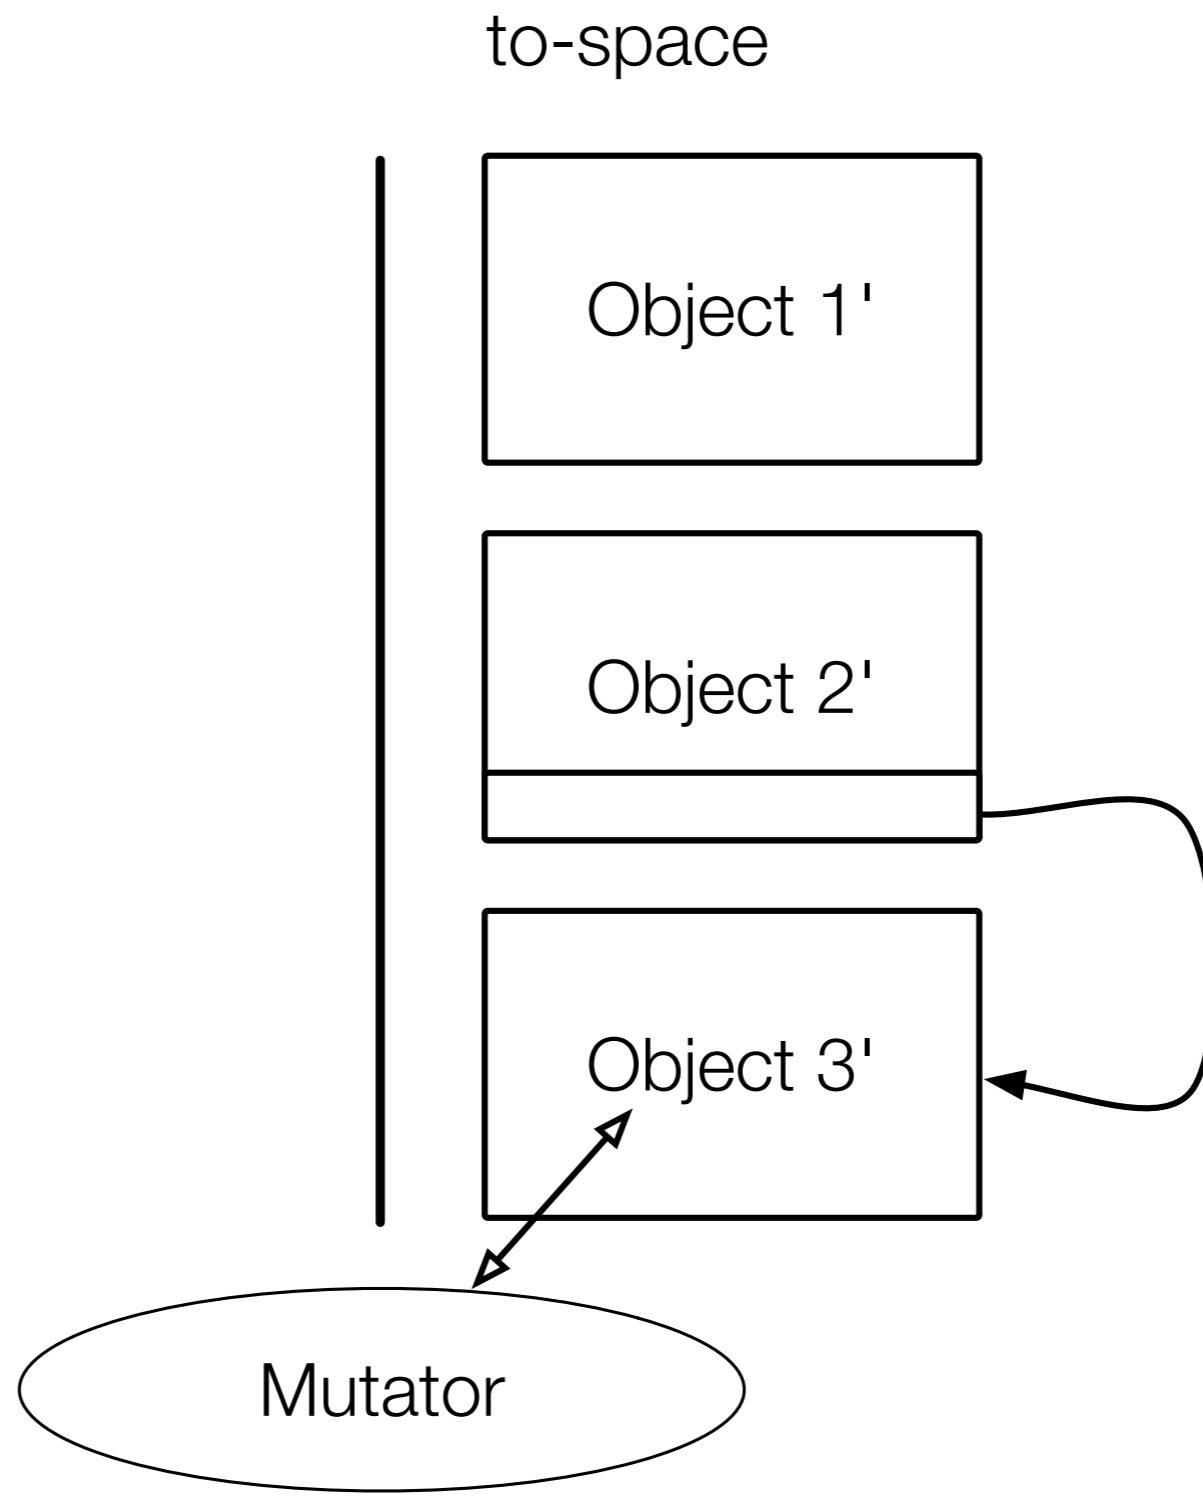
**Copy
Phase**



**Flip
Phase**



Finish



Synchronised Replication (Copying Phase)

Mutator updates from-space and to-space

Collector copies data to to-space


```
u = load (to-space)
v = load (from-space)
if (u != v)
    compare-and-swap (to-space, u, v)
    restart
else
    break
```

unsafe

`v = load (from-space)`
`store (to-space, v)`

XBEGIN

$v_1 = \text{load (from-space)}$

$\text{store (to-space, } v_1)$

$v_2 = \text{load (from-space)}$

$\text{store (to-space, } v_2)$

...

XEND

XBEGIN

copy object 1

copy object 2

...

copy object n

XEND

scan and record object 1 .. n

XBEGIN

copy object 1

copy object 2

...

copy object n

XEND

$v_1 = \text{load}(\text{from-space})$

$\text{store}(\text{to-space}, v_1)$

$v_2 = \text{load}(\text{from-space})$

$\text{store}(\text{to-space}, v_2)$

...

MFENCE

verify to-space with from-space

Test Setup

Intel Core i7-4770 3.4Ghz, 16GiB RAM


Ubuntu 12.04.3 LTS


DaCapo 2006 (sunflow from 2009)

Fixed 350m heap and GC trigger

Copy speed

bytes copied / time in copy phase





Transactional Penalty

Normal reads = 11 bytes/ns

Normal writes = 6 bytes/ns

Transactional reads = 9 bytes/ns (~80%)

Transactional writes = 6 bytes/ns (100%)

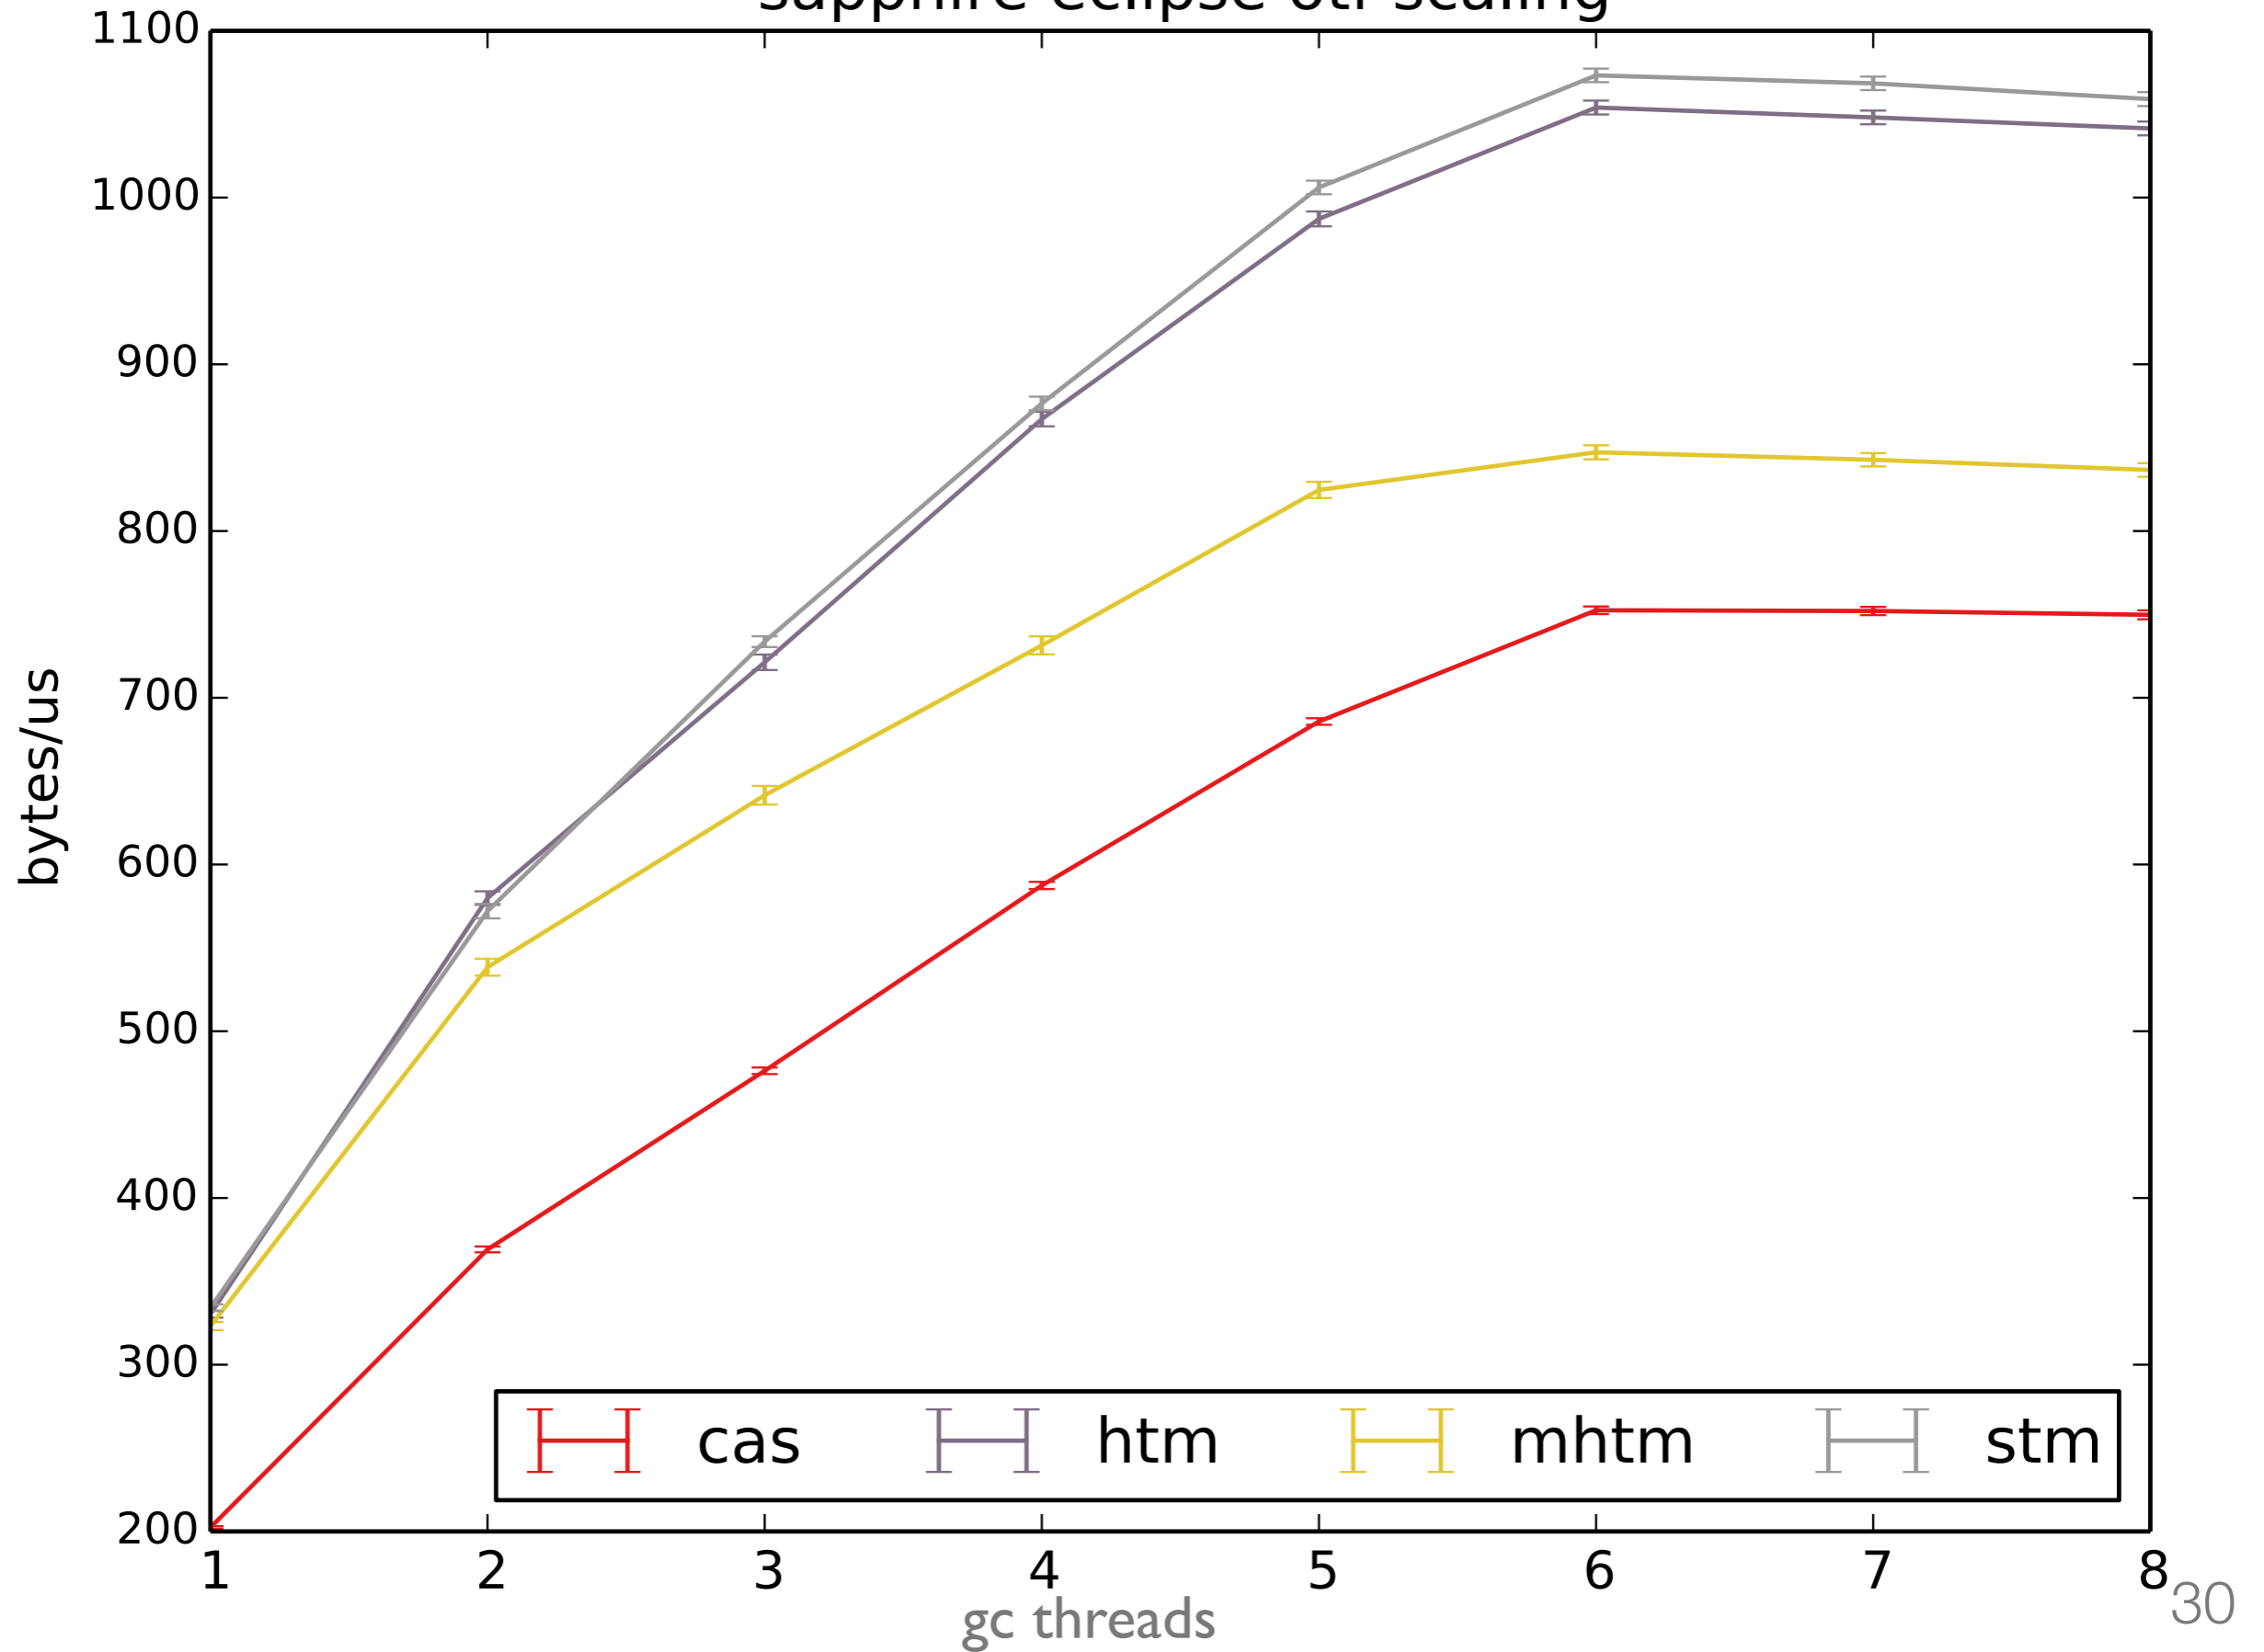
Peak Performance

Unsafe = 440 bytes/us

$440 * 0.8 = 352$ bytes/us (exactly as observed)

Scaling

sapphire eclipse of scaling



Conclusions (Specific)

Speed up of 48-101% for concurrent copying collection

Weak consistency requirements make STM possible

No improvement for bitmap marking and parallel copying

Conclusions (General)

Transactional work must be sufficient and concise

Engineering this can add overheads which outweigh benefits

Future Work

New collector designs which exploit strong consistency

Ground up batching of work for transactions

... mostly barrier free on-the-fly collection?

Questions?

Code: <http://github.com/perlfu/sapphire>