Why be Normal?

Single Crystal Growth and X-ray Spectroscopy Reveal the Startlingly Unremarkable Electronic Structure of Tl-2201

by

Darren Peets

B.Sc., The University of British Columbia, 2000
M.Sc., The University of British Columbia, 2002

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

The Faculty of Graduate Studies

(Physics)

The University of British Columbia

(Vancouver)

September 2008

Abstract

High-quality platelet single crystals of $Tl_2Ba_2CuO_{6\pm\delta}$ (Tl-2201) have been grown using a novel time-varying encapsulation scheme, minimizing the thallium oxide loss that has plagued other attempts and reducing cation substitution. This encapsulation scheme allows the melt to be decanted from the crystals, a step previously impossible, and the remaining cation substitution is homogenized via a high-temperature anneal. Oxygen annealing schemes were developed to produce sharp superconducting transitions from 5 to 85 K without damaging the crystals. The crystals' high homogeneity and high degree of crystalline perfection are further evidenced by narrow rocking curves; the crystals are comparable to YSZ-grown $YBa_2Cu_3O_{6+\delta}$ by both metrics.

Electron probe microanalysis (EPMA) ascertained the crystals' composition to be $Tl_{1.920(2)}Ba_{1.96(2)}Cu_{1.080(2)}O_{6+\delta}$; X-ray diffraction found the composition of a $T_c = 75$ K crystal to be $Tl_{1.914(14)}Ba_2Cu_{1.086(14)}O_{6.07(5)}$, in excellent agreement.

X-ray refinement of the crystal structure found the crystals orthorhombic at most dopings, and their structure to be in general agreement with previous powder data. Cation-substituted Tl-2201 can be orthorhombic, orthorhombic crystals can be prepared, and these superconduct, all new results. X-ray diffraction also found evidence of an as yet unidentified commensurate superlattice modulation.

The Tl-2201 crystals' electronic structure were studied by X-ray absorption and emission spectroscopies (XAS/XES). The Zhang-Rice singlet band gains less intensity on overdoping than expected, suggesting a breakdown of the Zhang-Rice singlet approximation, and one thallium oxide band does not disperse as expected. The spectra correspond very closely with LDA band structure calculations, and do not exhibit the upper Hubbard bands arising from strong correlations seen in other cuprates. The spectra are noteworthy for their unprecedented (in the high-T_c cuprates) simplicity.

The startling degree to which the electronic structure can be explained bodes well for future research in the cuprates. The overdoped cuprates, and Tl-2201 in particular, may offer a unique opportunity for understanding in an otherwise highly confusing family of materials.

Table of Contents

Appendices

List of Tables

List of Figures

Acknowledgements

I'd like to thank Dr. Ruixing Liang, who offered rather a lot of help over the course of this project, and loves handling things that are small, black and shiny. He performed the vast majority of the crystal harvesting, SQUID magnetometry and annealing, and was an indispensible resource for anything related to crystal growth or characterization. My long-suffering supervisor, Dr. Doug Bonn, deserves special mention for tolerating my jaunts into administration, politics, rule-bending, and random yet protracted acts of sarcasm. While he grumbled mightily about the Fire Hydrant's political career, I know he secretly enjoyed watching its antics. His wife, Colleen, also needs to be thanked at this jucture, for instilling patience in Doug in this regard, and encouraging me to get involved over Doug's protestations. Doug also spent more than his fair share of time turning platelet single crystals into spheroids, a task I feel fortunate not to have had to do myself. Dr. Walter Hardy did most of the design work on my furnace, and was extremely helpful whenever things broke or didn't quite work right. He's very good at both solving experimental problems and convincing things to work. Jess Brewer and Phil Stamp deserve thanks as well, for the role they played in keeping me on track, in line (to the extent possible), on my toes, and, in Jess' case, in gold. Dave Hawthorn, Kyle Shen and Jonathan Denlinger helped a great deal by explaining the X-ray spectroscopies and the physics and instrumentation involved. The physics shops also need to be thanked, because the vast majority of my crystal growth apparatus was built by them, and this research would not have been possible without them.

The X-ray diffraction data were taken by Dr. Brian Patrick in the Chemistry Department's X-ray crystallography facility, the XAS and XES data were taken with help from Drs. David Hawthorn and Kyle Shen, now at Waterloo and Cornell respectively, and the EPMA data were collected by Dr. Mati Raudsepp in the Earth & Ocean Sciences Department. The LDA calculations were performed by Dave Hawthorn.

NSERC, CIFAR and the BC Synchrotron Institute helped out in their usual ways, and Graduate Studies finally saw fit to offer me some nice scholarships once I'd been here a while. The Advanced Light Source is supported by the Director, Office of

Words to Live By

"There were many occasions I wished that Hydrant was there so I could pee on it."

— Dennis Pavlich, UBC's then-VP External & Legal Affairs

"I'm incompetent."

— Dennis Pavlich, UBC's then-VP External & Legal Affairs

"Of course you can quote me, I am incompetent."

— Dennis Pavlich, UBC's then-VP External & Legal Affairs

"I kind of enjoy arguing with you." — Darren Peets
"It's not an argument." — Hon. Rick Thorpe
"Yes it is." — Darren Peets
"No it's not." — Hon. Rick Thorpe

"It will work, because it has to work."

— UBC's transportation consultant Richard Drdul, explaining why UBC's
proposed underground bus terminal would work as designed

"However, it has been pointed out that users of the UBC terminal will be predominately students, who may be more tolerant of crowded conditions than the general public."

— TransLink consultants' report on the proposed bus terminal

1. Introduction

In condensed matter physics, a great deal of effort is presently being applied to the problem of correlated electron systems, where the effects of multi-particle interactions are fundamental to the materials' behaviour, and cannot be scaled out through standard Fermi liquid rescaling of parameters, such as effective masses. The greatest unsolved problem in correlated electron physics, and one that has been the subject of intense study for over two decades, is that of high-temperature superconductivity (HTSC). The work reported in this thesis represents several key advances in the concerted effort to characterize and ultimately make sense of these materials.

1.1 History

The first superconductor was discovered by Heike Kamerlingh-Onnes in 1911 [1], when he observed the resistivity of mercury to suddenly vanish when cooled below 4.2 K. This state was proven unequivocally to be a new phase of matter when, in 1933, Walther Meißner and Robert Ochsenfeld showed that superconductors were also perfect diamagnets [2]. A successful microscopic theory explaining conventional superconductors has been in place since 1957, when Bardeen, Cooper and Schreiffer posited that a phonon-mediated attractive interaction between electrons would lead to this behaviour [3]. This is known as the BCS theory. By the 1980s, superconductivity was widely considered a mature field.

The first high-temperature superconductor, $La_{2-x}Ba_xCuO_{4+y}$, was discovered in 1986 by Georg Bednorz and K. Alex Müller [4]. A multitude of these cuprate superconductors were discovered in the following years, with superconducting transition temperatures (T_cs) as high as 138 K confirmed at ambient pressure [5] (for $Hg_{0.8}Tl_{0.2}Ba_2Ca_2Cu_3O_{8+\delta}$). These high critical temperatures are found in copper oxides with a perovskite-like structure. The only feature common to all HTSCs is a planar CuO_2 layer (sometimes found in multilayers), which is believed to be responsible for the superconductivity. While conventional superconductors are typically ductile metals, the cuprates' crystal structures resemble rock salt, making them brittle, and their normal state is poorly conducting or even insulating (depending on

carrier doping).

Twenty-one years' monumental effort has been applied to characterizing and explaining these materials, resulting in an impressive collection of often contradictory data and a wide variety of theoretical explanations. To date, the HTSCs do not have the benefit any agreed-upon microscopic theory. This is certainly not for lack of effort — thousands of physicists have worked on these compounds over the past two decades — but issues such as sensitivity to impurities, inconsistencies in results between experiments probing similar properties, the absence of a single material to which all techniques may applied, and the extreme difficulty of the theoretical problems have prevented both a coherent experimental story and a detailed theoretical understanding of the underlying mechanism.

1.2   High-Temperature Superconductivity

The high-temperature superconductors can be tuned through a remarkable series of states of matter by adjusting the concentration of hole carriers [6, 7]. Their phase diagram is remarkably rich in new physics, which, while exciting, also makes a detailed understanding even more difficult — it's especially challenging to explain how the superconducting state arises when the state from which it arises is itself unexplained.

Figure 1.1 shows this doping phase diagram schematically, for the hole-doped cuprates. At zero hole doping ($p = 0$), defined conventionally to have exactly one hole per plane copper atom (i.e. all plane Cu atoms are $3d^9$), the system is an antiferromagnetic Mott insulator (AFI). It is worth emphasizing that this $p = 0$ line is the *only* part of the phase diagram for which there is an accepted, solved microscopic theory. As holes are doped into this phase, the antiferromagnetism's Néel temperature (T_N) is quickly suppressed. As more holes are added, superconductivity appears as the ground state of the system.

The nature of the low-temperature region between these phases, whether they meet, whether any overlap exists, or to what extent this varies between materials would make interesting topics for one or more theses, but are well beyond the scope of the present one.

As more holes are introduced, the superconducting transition temperature (T_c) rises, peaks, and falls back to zero. These constitute the underdoped, optimally doped and overdoped regimes respectively. Some properties of the system change dramatically across this dome, while others do not [8] — for instance, the nodes described in section 1.5 below are believed to exist at all dopings probed [9, 10, 11, 12, 13],

Figure 1.1: Schematic depiction of the doping phase diagram of the hole-doped
cuprate superconductors — phases observed as a function of tempera-
ture (T) and hole doping (p). Nothing should be read into the location
of either dashed line.

but the Fermi surface seen in the overdoped regime by angular magnetoresistance
oscillations (AMRO) [14] and Angle-Resolved Photoelectron spectroscopy (ARPES)
[15, 16] is a single hole pocket consuming over 60% of the Brillouin zone, while there
are indications on underdoped $YBa_2Cu_3O_{6.5}$ and $YBa_2Cu_4O_8$ from quantum oscilla-
tions in the Hall resistance R_{xy} [17, 18, 19], and the de Haas-van Alphen technique
[20] that the Fermi surface consists of pockets each enclosing closer to 2% of the Bril-
louin zone. ARPES measurements of underdoped cuprates show the Fermi surface
reduced to arcs around the nodal directions [21, 22, 23]. It should be noted that while
these studies found the behaviour of particles at the Fermi surface to be consistent
with some aspects of Fermi liquid theory, many other properties of the phase remain
unexplained.

Above these low-temperature phases exist several other phases. Above the under-
doped regime, the normal state is referred to as the pseudogap state, as it appears
to exhibit a suppression of the density of states near the Fermi level, but not a full

gap as in the superconducting state. The location of the crossover that marks this phase's upper-temperature boundary, and whether there is only one such crossover, are matters of some debate, while the nature of this state and its relationship to the superconducting dome are themselves contested. It is unclear whether this phase extends down in doping to the insulating phase, is interrupted by another phase, or becomes less exotic.

Above the superconducting dome at optimal doping and above most, if not all, other non-superconducting phases, is a phase that can be successfully described phenomenologically by marginal Fermi liquid theory. This state has a T-linear resistivity and several other unusual power laws in the transport properties, all of which can be described by this model, although not yet explained microscopically. Similar behaviour in the resistivity has been observed above quantum critical points, where it arises from quantum fluctuations and the lack of a temperature scale, but such an explanation would raise more questions than it would answer, as it would imply a mysterious quantum phase transition buried under the superconducting dome where it can't easily be accessed. Such a transition could be linked to Fermi surface reconstruction, changes to the details of the band structure or correlated electron physics, the demise of some hidden order, or any number of other possible explanations, and an actual quantum critical point may be located near or under the superconducting dome along another axis, such as magnetic field or pressure.

On the less-studied overdoped side, there is some limited evidence suggesting a crossover to another phase [24, 25, 26, 27], which is, perhaps optimistically, labelled a Fermi liquid on most schematic phase diagrams. This labelling arises, in large part, due to a lack of evidence to the contrary. If this phase were shown to obey Fermi liquid theory without exhibiting bizarre, unexplained properties, it would constitute a second phase for which we could claim a microscopic understanding. Of course, given the other states present in this phase diagram, perhaps the overdoped end is more likely to be yet another exciting new phase of matter.

Due largely to technology-transfer and materials issues, the vast majority of work has been conducted in the optimal and underdoped regimes. Compounds native to the overdoped regime are scarce, difficult to grow, and typically quite toxic, while any commercial applications of HTSCs are likely to require high transition temperatures and critical current densities. This thesis deals with $Tl_2Ba_2CuO_{6\pm\delta}$, a compound which is intrinsically overdoped.

1.3 Materials

In the struggle to understand the cuprates, the quest for better materials to study remains key. Researchers want samples that can be grown very cleanly, can be doped over a very wide range of the phase diagram, and are suitable for a variety of bulk *and* surface measurements.

Toward the insulating end of the phase diagram, Na-doped $Sr_2CuO_2Cl_2$ (SCOC) yields excellent cleaved surfaces that have been used in surface-sensitive techniques such as ARPES [22] and Scanning Tunnelling Microscopy and Spectroscopy (STM and STS) [28], recently revealing the existence of a new charge-ordered phase between the insulator and superconductor.

Near optimal doping, $YBa_2Cu_3O_{6+\delta}$ (YBCO) stands out as the cleanest material and has been used extensively in studies of the d-wave superconducting state using bulk probes such as thermal conductivity [29], superfluid density [9] and microwave conductivity [30]. These transport measurements indicate that the quasiparticles in this material are well defined at low temperatures and can have very long mean free paths. Unfortunately, this material does not have a cleavage plane that is well-suited to surface-sensitive techniques, limiting the available single particle spectroscopy from ARPES [31] or STS [32]. Instead, the rich spectroscopic information from STS and ARPES on superconducting samples has come mainly from the family of Bi-containing cuprates and in particular $Bi_2Sr_2CuO_{6+x}$ (Bi-2201) and $Bi_2Sr_2CaCu_2O_{8+x}$ (Bi-2212 or BSCCO) [33, 34, 35, 36]. However, this group of materials has greater problems with disorder than the $YBa_2Cu_3O_{6+\delta}$ family, leaving a gulf between the materials for which we have detailed single particle spectroscopy and the materials for which we have bulk measurements indicating well-defined quasiparticles with long mean free paths. This dilemma has precluded strict tests of the theory of the d-wave superconducting state, in which one must connect the single particle spectroscopies to the rest of the properties of the material [37].

A further problem near optimal doping is that the superconducting critical temperature is so high as to effectively prevent suppression of the superconductivity to study the underlying normal state. Extreme overdoping is one means of reducing T_c so that one can study both the superconducting and normal states, but here most samples tend to be less clean because of the need to dope with large concentrations of cation impurities such as Sr in La_2CuO_{4+y} and Pb in $Bi_2Sr_2CuO_{6+x}$. The former is particularly problematic, being prone to phase separation at high dopings.

The structure of Bi-2212, the material most commonly studied by surface probes,
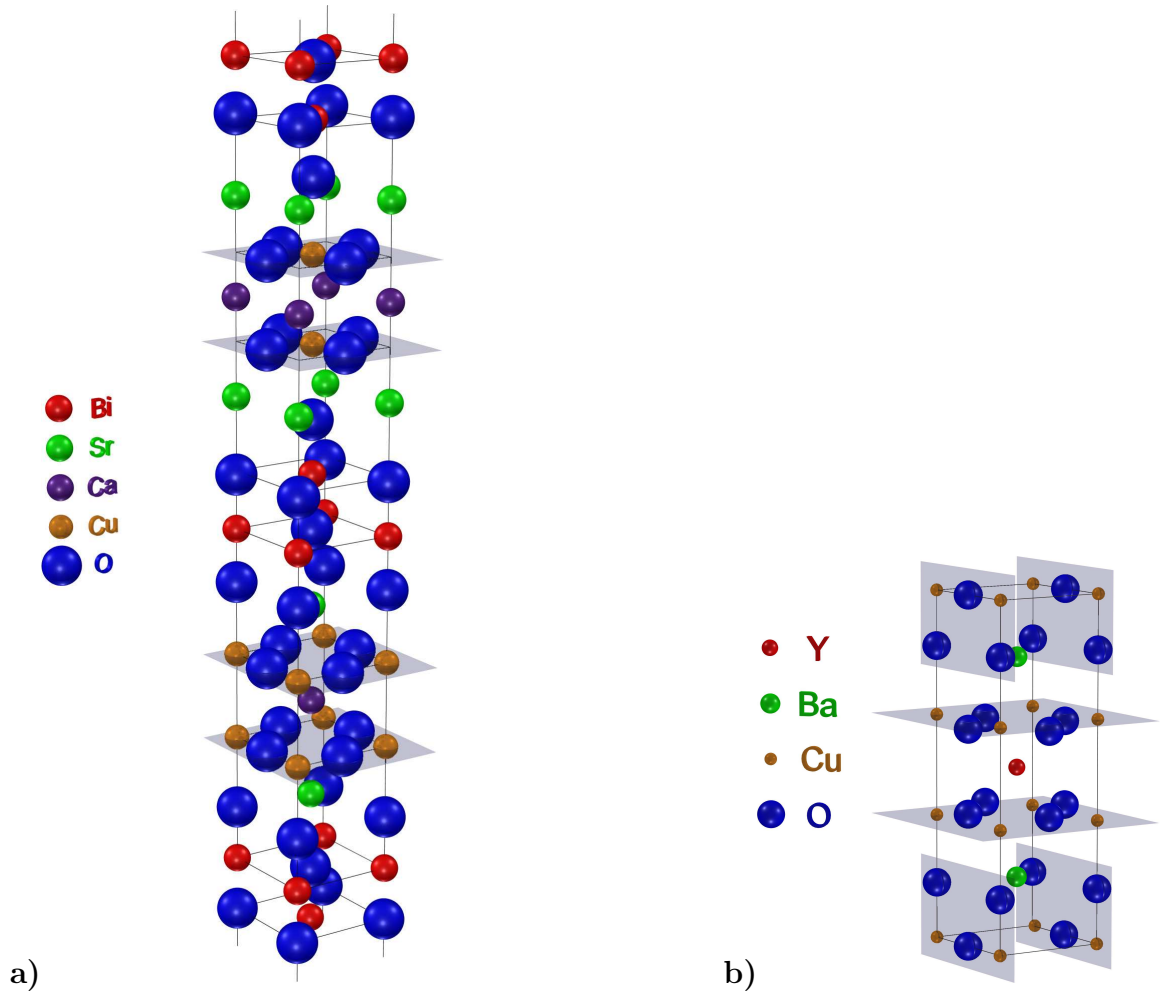
Figure 1.2: Unit cells of the high-temperature superconductors **a)** Bi-2212 and **b)** YBCO (fully doped). for both materials, the c-axis is vertical; YBCO's chains run along the b axis. One-dimensional (chain) and two-dimensional (plane) copper-oxygen features are emphasized with shading. Bi-2212's crystal structure is body-centred, so every unit cell has two identical bilayers.

is shown in Figure 1.2a. Bi-2212 has a BiO double layer held together only by Van der Waals bonds, yielding atomically-flat cleaved BiO surfaces when delaminated using scotch tape. While the ease of cleaving, the excellent likelihood of obtaining a prefect surface and the lack of surface reconstruction or surface states make this material highly prized for surface-sensitive techniques, the extremely two-dimensional nature of its transport properties makes some bulk probes difficult to apply to it, particularly where currents are required to run along the c-axis, perpendicular to the CuO_2 planes.

As regards disorder, normally around 5% of Bi-2212's Sr sites are occupied by Bi

atoms [38]. Since these Sr atoms constitute a charged defect immediately adjacent the CuO_2 planes, they are poorly screened from the charge carriers there. This can produce strong scattering, which tends to mask intrinsic behaviour in transport measurements. The oxygen nonstoichiometry utilized for doping in this compound is in the BiO double layer, however, far from the planes.

Additionally, Bi-2212 is a bilayer material, meaning that it has two adjacent CuO_2 planes separated only by a layer of Ca^{2+} cations. This leads to bonding and antibonding combinations of the bands associated with the CuO_2 plane, and can also lead to concerns as to whether the effects observed are intrinsic to the cuprates or merely an effect of the bilayer.

Bi-2212 may be doped as far down the overdoped side of the phase diagram as $T_c \sim 65$ K (from an optimal T_c of 96 K) when heavily Pb-doped [39], and may be underdoped somewhat as well, but it is difficult to produce homogeneous crystals in that doping range.

For transport and other bulk measurements, the gold standard is YBCO, whose unit cell is shown in Figure 1.2b. Very high purity crystals of this material may be grown [40, 41], and its carrier doping can be made quite homogeneous, notably through the use of oxygen-ordered phases [42, 43]. YBCO is an especially three-dimensional cuprate, does not suffer from intrinsic cation disorder, and has oxygen dopants farther from the planes. However, like Bi-2212 its planes exist in bilayers, plus it has a one-dimensional CuO "chain" layer, leading to substantial differences in the transport properties between the a and b directions.

YBCO 'cleaves' (or, more accurately, breaks) between the BaO and CuO (chain) layers, the latter being responsible for the doping. Since the doping layer picks a side, this leaves two surfaces with dopings very different from the bulk and each other. Compounding the problem, there are no non-polar c-axis surfaces available in the YBCO unit cell (all mirror planes coincide with a layer of atoms), so every atomically-flat surface must undergo electronic reconstruction to avoid a 'polarization catastrophe'. Electronically, the surfaces are fundamentally different from the bulk. This problem has only very recently been addressed, by decorating the surface with potassium atoms to not only stabilize it but to continuously change its doping [44].

$YBa_2Cu_3O_{6+\delta}$'s accessible oxygen doping range extends from $\delta = 0$, the undoped insulator, to $\delta = 1$, just past optimal doping. Adding calcium (which substitutes for yttrium) can overdope it to lower T_cs, but the yttrium site is between the two CuO_2 plane layers, an effective site for scattering due to the lack of any anions to help screen the effect, and the observed suppression of T_c may arise in part from this

disruption, not just from the added doping.

In both YBCO and Bi-2212, the relationship between the oxygen doping x or δ and hole doping (carrier concentration) p is non-obvious, and can be extremely difficult to sort out exactly [45]. For underdoped YBCO, one particular oxygen doping can provide a wide range of p and T_c values, since the number of holes promoted onto the planes depends crucially on the degree of oxygen order.
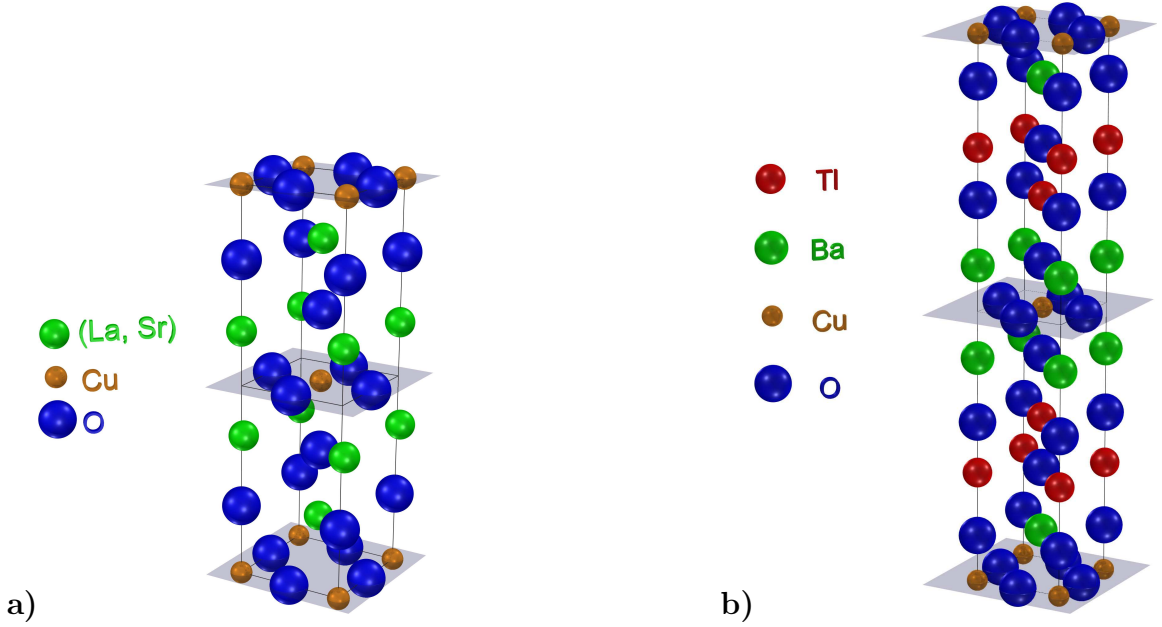
Figure 1.3: Unit cells of the high-temperature superconductors **a)**LSCO and **b)**Tl-2201. Again, the c axis is vertical and copper oxide planes are emphasized through the use of shading. Each compound is body-centred, and thus has two copper oxide planes per unit cell in identical chemical environments.

Another commonly-studied material is $La_{2-x}Sr_xCuO_{4+y}$ (LSCO or 214), shown in Figure 1.3a. This compound does not have complicating multilayer structures, chains, or similar copper features, and can be overdoped to $T_c= 0$. However, doping is primarily accomplished by cation substitution on the lanthanum site, which is located immediately adjacent to the planes. This is the second most destructive possible location for cation substitution after the plane copper sites. The substitution levels required to traverse the superconducting dome are substantial — T_c appears on the underdoped side at $x = 0.05$ and returns to zero on the overdoped side around $x = 0.27$ [46] — so disorder varies wildly with doping. It is important to note that cation doping, unlike oxygen doping, is not reversible. Dopant oxygen atoms are mobile in these compounds at elevated temperatures, and may be added or removed

from the compound, but the cation stoichiometry is set permanently during crystal growth.

LSCO offers the considerable benefit of an especially straightforward relation between its chemical doping x and the hole doping p, making some parts of the data analysis and comparisons to theory comparatively straightforward: $p = x$. Unfortunately, LSCO also exhibits some unusual magnetic behaviour not seen in other systems and is abnormally susceptible to a charge-ordered phase at $1/8$ doping. Its maximum T_c is suppressed by $\sim$50 K relative to other single-layer cuprates, possibly because of these competing phases or the cation disorder.

Tl-2201, like LSCO, is a single-layer compound. While there are two CuO_2 planes per unit cell (see Figure 1.3b), these are in identical environments — as is the case in Bi-2212 and LSCO, the crystal structure is body-centred. Like LSCO, Tl-2201 exhibits cation substitution, but in this case far from the CuO_2 plane, in the TlO double layer. Here, copper substitutes for thallium at the $\sim$5% level. Oxygen doping, which is used as the hole doping control parameter, is accomplished via interstitials in this same double-layer [47, 48], shown in Figure 1.4. Particularly low dopings (i.e. approaching optimal doping) may require the removal of a few oxygen atoms from the TlO layer.

Figure 1.4: Location of the oxygen interstitial within the TlO double layer. The interstitial oxygen dopant is located mid-edge, inside both a tetrahedron of thallium atoms and a tetrahedron of oxygen atoms. Here, the c axis is again vertical, and the Cu–O bonds fall below the lines marking the edges of the layer shown.

Tl-2201 is less anisotropic than some other cuprates, and has been successfully studied by such bulk probes as thermal transport [24, 49], resistivity [25, 26], and microwave conductivity [10, 13]. It is believed to cleave between the TlO layers (albeit not as easily as Bi-2212), so surface probes may be employed. Indeed, as a result of this work, Tl-2201 became the first superconducting cuprate on which bulk and surface measurements agreed quantitatively on the same physical property —

the Fermi surfaces measured via AMRO (a bulk probe) [14] and ARPES [15, 16]. In the $3d$ and $4d$ transition metal oxides, this was only the second such quantitative agreement, after measurements of Sr_2RuO_4 by de Haas–van Alphen [50, 51] and ARPES [52, 53, 54]. The excellent agreement between these results indicates that $Tl_2Ba_2CuO_{6\pm\delta}$ may be the ideal testing ground for finally joining the modern single-particle spectroscopies with a host of well-established bulk probes.

1.4 Crystal Growth

Fundamentally, crystal growth involves condensing a compound into its solid phase from higher-entropy phases (e.g. solution, liquid, vapour) in a controlled manner, such that the resulting solid phase exhibits the best possible long-range order. In most everyday materials (such as water or salt), the compound in the liquid or vapour phase is chemically identical to the solid phase. This is not true in general.

1.4.1 Phase Diagrams

The Gibbs Phase Rule,

$$D = C - P + N \tag{1.1}$$

indicates how many degrees of freedom (D) there are when P phases coexist in a C-component system with N non-compositional variables. For instance, a one-component system wherein pressure and temperature can change can have three phases coexisting, but with no degrees of freedom — they can only coexist at a single point (the triple point of water is such a point). When three components are present but pressure is irrelevant, many types of region are possible, from points where four phases coexist to regions where a single phase exists for a range of temperatures and compositions. This is the case for the Tl-Ba-Cu oxide system.

Like many of the cuprate superconductors, $Tl_2Ba_2CuO_{6\pm\delta}$ melts incongruently, at a peritectic point. Instead of melting into a liquid with identical composition to the solid, Tl-2201 melts into a solid phase and a liquid phase with compositions different from the original solid phase. This makes carefully controlling the inverse process, employed in crystal growth, far more challenging. A peritectic point may be readily observed in Figure 1.5, by following the $Tl_2Ba_2CuO_{6\pm\delta}$ composition up to 930°C, where it decomposes. This figure shows a quasi-two-component slice through the full three-component phase diagram — one degree of freedom has been neglected.

When a liquid with the same composition as Tl-2201 is cooled, the first feature

Figure 1.5: A slice through the Tl-Ba-Cu oxide phase diagram, showing a region where $Tl_2Ba_2CuO_{6\pm\delta}$ is the primary phase, adapted from [55].

it encounters in the phase diagram is a liquidus curve, below which the solid phase $Tl_2Ba_2O_5$ begins to coexist with the liquid. If a more copper-rich composition is used, another liquidus curve exists where $Tl_2Ba_2CuO_{6\pm\delta}$ is the first solid phase to form. Figure 1.5, based on the only published phase diagram for the Tl-Ba-Cu oxide system [55], shows this particular liquidus curve existing between 55 and 72 mol% CuO. When a solid phase begins forming, the fact that its composition is different from that of the liquid implies that the composition of the liquid must also change. The compositions of the two phases follow the boundaries of a forbidden dome. Eventually, a eutectic point may be reached, where the liquidus curve meets another in a minimum. At this lowest-melting composition, the liquid phase freezes directly into two intimately mixed solid phases. The example is at 72 mol% CuO and 896°C. The constant-temperature line passing through this point has only one degree of freedom — composition — and marks the coexistence of two solid phases with the liquid.

It should be emphasized that the phase diagram shown is merely an approxima-

tion. If the cation substitution level of Tl-2201 (assumed to be zero in this slice) can vary, the neglected degree of freedom may be important, and the composition of the crystal may vary as the composition of the melt drifts.

While peritectic points are common complicating factors in the cuprates, Tl-2201 has a further complication. The thallium oxide stable at ambient conditions is Tl_2O_3, which is also the only thallium oxide available commercially. However, the thallium oxide stable at growth temperature is monovalent Tl_2O [56], requiring the removal of oxygen:

$$Tl_2O_3 \rightleftharpoons Tl_2O + O_2 \uparrow$$

Normally, this would not be a concern — barium precursors in particular are typically carbonates, requiring that CO_2 be evolved and removed from the system. However, Tl_2O is itself volatile at growth temperatures, having a vapour pressure on the order of a tenth of an atmosphere [56]. The thallium oxide vapour pressure is non-negligible as low as 500°C, at which temperature it's $\sim 10^{-7}$ atm. If the thallium oxide vapours are not contained, they will escape, leading to crystals grown at best at a time-varying location in the phase diagram, or at worst in a thallium-free environment. Thallium is also toxic (it was once commonly used as a rat poison), adding one more incentive to keep it contained. A further issue is that Tl_2O is quite reactive, most notably reacting with the quartz often used for encapsulation and furnace tubes.

Another difficulty in working on $Tl_2Ba_2CuO_{6\pm\delta}$ is the dearth of phase diagrams. The almost total lack of phase diagrams is both an artifact and a cause of the scarcity of chemical and materials work on this system. The determination of a phase diagram is made more difficult by the thallium vapour pressure.

The only two (even partial) phase diagrams for the Tl–Ba–Cu–O system were published in reference [55] and are reproduced in Figures 1.5 and 1.6. The former is a quasi-binary phase diagram between $Tl_2Ba_2O_5$ and CuO, while the latter is a detail of the area around Tl-2201, as inferred from the known or suspected solid phases. It is not known whether the two Tl-2201 phases shown are each chemically distinct, whether solid phases exist continuously along what is depicted as the tie line between them, or whether each exists over some range in composition, but with a disallowed region between them. It is also not known whether crystals of a strictly stoichiometric phase may be grown. It should be noted at this juncture that phase diagrams drawn this way assume that the behaviour of the solid and liquid phases is completely independent of the surrounding atmosphere. This is normally accurate, but in this system the partial pressure of oxygen affects the oxidation state of thallium, and

Figure 1.6: A detail of the Tl-Ba-Cu oxide phase diagram in the vicinity of Tl-2201, under 1 atm O_2, based on a figure in [55]. There may be a continuum of phases connecting the two Tl-2201 compositions.

could have a material effect on the phase diagram.

1.4.2 Growth Techniques

A wide variety of crystal growth techniques are available for use with the cuprate superconductors, with the most suitable technique generally determined by the desired properties of the crystals (e.g. purity, size) and system-specific technical considerations (e.g. a corrosive melt).

One method commonly used to grow non-cuprate oxide superconductors, most notably Sr_2RuO_4 [57, 58], and a few cuprates is travelling-solvent floating-zone (TSFZ, see Figure 1.7). In this technique, two rods of the material to be grown (in ceramic or pressed powder form) are melted where they meet, and rotated counter to one other to stir the liquid there. The molten zone is moved along the pair of rods by moving either them or the heat source, and the rod left behind is crystalline. After some length, for crystals that grow anisotropically, one crystal domain with a preferred

Figure 1.7: Travelling solvent floating zone (TSFZ) crystal growth method: two coun-
terrotating rods of the desired material (brown) are heated where they
meet, and as this molten zone (yellow) or the sources heating it move,
crystalline material (black) is left behind.

orientation will squeeze out all competing domains, and the remainder of the rod will
be a single crystal. This technique works best with materials that grow comparatively
quickly. In systems with peritectic points, TSFZ can work, but it is limited by the
solubility of the solid phase in the liquid phase. TSFZ can grow Bi-2212 or LSCO
crystals on the order of 10 cm^3 in size, but fails for YBCO because its liquid phase has
strong surface energy effects that cause it to spread out, particularly in temperature
gradients. TSFZ is inappropriate for Tl-2201 because it involves a small heated zone
within a furnace that is otherwise quite cool. Thallium oxide would evaporate from
the heated zone and condense out in cooler parts of the furnace.

Figure 1.8: Czochralski growth: a rotating rod dips a seed crystal into the melt, then
slowly pulls it out as a crystal grows on the seed.

A related method is Czochralski growth (Figure 1.8), in which a seed crystal
is dipped into a crucible of the melt, then slowly pulled out as the desired phase
crystallizes onto the seed. Here too, the seed is typically rotated to stir the melt, and

the resulting crystal can be tens of cubic centimetres. This method is commonly used to grow crystals of congruently-melting compounds such as silicon and sapphire. The high-T_c cuprates melt incongruently, however, and only YBCO crystals have been grown by an adaptation of method — solute-rich liquid crystal pulling (SRL-CP), wherein the sparingly-soluble peritectic solid phase (Y_2BaCuO_5) is transported to the growth front by carefully controlled convection. Applying SRL-CP to Tl-2201 would require a seal at 950°C between an airtight furnace and at least one rotating rod in order to avoid the problems associated with thallium oxide vapour.

There are various forms of liquid-phase and vapour-phase epitaxy, typically for growing single-crystalline films on substrates. While this is quite useful for applications, it is of limited value for basic research into the fundamental behaviour of the material, largely thanks to disorder (impurities and particularly strain and dislocations) — the electronic structure of the cuprates is particularly sensitive to disorder.

Figure 1.9: Flux growth: a melt containing the constituents to be crystallized is cooled in one area. Crystals nucleate and grow there, then must be separated from the flux.

Smaller single crystals may be grown by cooling a high-temperature solution of the compound (the flux) through the liquidus line, much as salt or sugar crystals may be grown from aqueous solution (see Figure 1.9). Typically, a cold spot is provided to generate convection and control the crystals' nucleation, by blowing air on the crucible, bringing a quartz rod up to it to act as a lightpipe, or designing the furnace to have a temperature gradient. This method can sometimes reduce the temperatures required (which would be useful for controlling the loss of thallium in the case of Tl-2201). Flux growth produces smaller (~ 1 mm^2 $\times$ 10 μm) crystals than do the above methods, but with a high degree of crystalline perfection. Tl-2201 has been successfully grown in a KCl flux [59], but without any substantial reduction in the temperature required. The crystals may be separated from the solidified KCl by dissolving it in water, although this can damage the crystals if any CO_2 is dissolved in

the water. However, a chief drawback with this technique is a tendency to introduce unwanted impurity atoms (such as K or Cl in this example) into the crystals.

A variant of flux growth is self-flux growth, in which the melt is an excess of one or more of the components in the crystal. YBCO, for instance, is typically grown from a melt that's extremely rich in CuO [40, 41]. This can produce extremely high-quality crystals, free from impurities. The limitation, alluded to above, is that the crystals produced are far smaller than those produced by TSFZ or similar methods. For many techniques, particularly transport and surface probes, small crystals may be better-suited to the measurement. However, it is far easier to make one large crystal into many tiny crystals than to do the reverse. Techniques such as μSR or especially neutron scattering, which require larger sample volumes for adequate signal-to-noise, require large mosaics of crystals this size, where up to several hundred crystals must be individually annealed, characterized and aligned [60]. It is important in the Tl-2201 system to have the crucible either itself be sealed, or be contained in a small, airtight vessel, to minimize the thallium oxide losses. Encapsulation, however, complicates the matter of separating crystals from flux. When the solidified flux has a very similar composition to the phase being grown, they cannot be separated by dissolving the flux. In the case of YBCO, the crucible may be tipped over at the conclusion of crystal growth, to decant the melt. This step is extremely difficult in the case of Tl-2201.

No one crystal growth method is perfect for all purposes. For research on the fundamental physics of the cuprates and particularly their electronic structure, minimizing disorder is key. Self-flux growth was employed in this work.

1.4.3   Annealing and the Pursuit of Order

To access the intrinsic transport and electronic properties of a d-wave superconductor, a high degree of crystalline perfection is required. This means the lowest possible concentration of impurity atoms, a minimum of dislocations, and dopants that are homogeneously distributed (ideally, ordered).

One obvious and important step is to use the purest precursors possible. However, in the Tl-2201 system, there are some unusual considersations. For instance, $BaCO_3$ would be used as a barium source for growth of most cuprate crystals, since it is a stable stoichiometric compound, allows accurate weighing by absorbing little water from air, and is inexpensively available in high purity. However, the Tl-2201 system is comparatively low-melting, so the evolution of carbon dioxide is significantly slower

than for other systems such as YBCO. Barium carbonate is quite stable, so it does not react to form Tl-2201 until the CO_2 has evolved, potentially reducing the effective barium content of the melt in a time-varying manner. Additionally, there are a number of known superconducting thallium copper oxycarbonates [61] which should be avoided.

Some cuprates, notably YBCO, grow from extremely corrosive BaO-CuO melts, which can dissolve common crucible materials. Aside from weakening the crucible, which can be catastrophic in its own right, this contaminates the melt with cations from the crucible material. If these cations can find an amenable site in the crystal, they will contaminate it. Crucible corrosion is less of a problem for Tl-2201 than for YBCO because of the lower growth temperature.

1.4.3.1 Doping

Most of the superconducting cuprates can be chemically doped to traverse at least part of the superconducting dome, either via cation substitution or variation of the oxygen content.

In many cuprates, most notably LSCO ($La_{2-x}Sr_xCuO_{4+y}$) and the electron-doped materials (e.g. $Pr_{1-x}Ce_xCuO_{4-y}$), carrier doping is accomplished via cation doping. In LSCO, some La^{3+} cations are replaced with Sr^{2+} — a less-positive cation layer requires a less-negative CuO_2 layer for overall charge neutrality, so this adds holes to the copper-oxygen plane. Reaching the overdoped end of the superconducting dome requires that nearly 15% of the lanthanum atoms be replaced[1], which cannot be considered a small perturbation.

Bismuth-based cuprates exhibit excess Bi on the Sr site, replacing Sr^{2+} with Bi^{3+}, while Tl-2201 has substitution of excess copper onto the thallium site, replacing Tl^{3+} with Cu^{2+} or possibly Cu^+. However, unlike in LSCO, where the cation substitution is deliberately varied to dope the material, these compounds' cation substitution is intrinsic, difficult to vary, and is not utilized to adjust the doping. While the substitution is undesirable, it is at least consistent. Deliberate cation substitution often increases strain in the crystal, but can be used to relax intrinsic strain; in intrinsically cation substituted systems such as Tl-2201, the substitution is typically driven by a need to reduce strain.

Having a uniform copper substitution level requires that the composition of the crucible contents not drift significantly over the course of the growth, for instance

---

[1]Because there are two lanthanum atoms per unit cell, $p = 0.27$ means 13.5% of lanthanum atoms are replaced.

through loss of thallium. This is a problem for flux growth of all materials with variable cation doping, as the composition of the melt will move along the liquidus line on cooling, potentially changing the cation ratio in the crystal.

In oxygen-doped compounds, the introduction of an oxygen atom as O^{2-} to an otherwise vacant site requires that two electrons be captured from the crystal, so two holes are introduced. In some compounds, notably LSCO, the only band near the Fermi level is that for the CuO_2 planes, so every dopant oxygen atom dopes two holes into the plane. In other cases, such as YBCO, other electronic states exist near the Fermi level, and the correlation between chemical doping and hole doping, while generally monotonic, may be difficult to determine. YBCO's Fermi surface consists of chain bands and bonding and antibonding combinations of the CuO_2 planes' bands, with BaO bands nearby. The relationship between oxygen and hole doping in YBCO has only recently been determined, through measurements of the c-axis lattice parameter [45]; the relationship in $Tl_2Ba_2CuO_{6\pm\delta}$ is not known exactly and will depend on the precise level of cation substitution, which will effectively give the crystal's carrier doping level an offset. Band structure calculations suggest that Tl-2201 should have a small TlO electron pocket around (0,0) which would complicate matters, but neither this nor a BiO analogue in the bismuth-based cuprates has ever been observed.

At temperatures where the oxygen atoms may enter and leave the crystal, a chemical equilibrium is established between the oxygen content of the crystal and the oxygen partial pressure around it, the location of the equilibrium being dependant upon the temperature of the system. In the case of $Tl_2Ba_2CuO_{6\pm\delta}$, the equilibrium is

$$Tl_2Ba_2CuO_6 + \frac{\delta}{2}O_2 \rightleftharpoons Tl_2Ba_2CuO_{6+\delta}$$

with equilibrium constant

$$K = \frac{\delta}{P_{O_2}^{\frac{\delta}{2}}} \tag{1.2}$$

and expected temperature dependence

$$\ln K = -\frac{\Delta G^{\circ}_{rxn}}{RT} \tag{1.3}$$

where P_{O_2} is the partial pressure of O_2 gas around the crystal, R is the ideal gas

constant, and ΔG_{rxn}° the product's standard Gibbs free energy minus the reactants' total standard Gibbs free energy. Thus an oxygen partial pressure and anneal temperature uniquely determine an oxygen content, assuming enough time is allowed for the oxygen content to reach equilibrium.

In materials that may be chemically doped, it is extremely desirable to have a homogeneous distribution of dopant atoms. Regions of the crystal having different concentrations of dopant atoms will have different T_cs, blurring and obscuring experimental features. Cation dopants may have a limited mobility at temperatures close to the growth temperature, but this can only be used to homogenize the dopants, not change the doping. In oxygen-doped systems, oxygen atoms may enter and leave the material at elevated temperatures, and may remain mobile to just above room temperature, allowing the doping to be both adjusted and homogenized.

$YBa_2Cu_3O_{6+\delta}$ and most other cuprates grow intrinsically strongly underdoped, and in YBCO's case oxygen must be added to reach doping levels at which superconductivity may be observed. As-grown $Tl_2Ba_2CuO_{6\pm\delta}$, however, is intrinsically overdoped, and the addition of oxygen suppresses superconductivity — lower dopings require oxygen removal. Regardless of what doping range is desired, however, the crystals must be annealed long enough to ensure homogeneity.

In some systems, the oxygen dopants can be ordered to further reduce scattering. For instance, YBCO exhibits several oxygen-ordered phases, including Ortho-II, Ortho-III, Ortho-V and Ortho-VIII [43, 62, 63], and LSCO can exhibit oxygen dopant ordering, including Stage 6, Stage 4, Stage 3, and Stage 2 [64, 65], although the staging in LSCO is weak and runs along the c-axis, making it less helpful for most transport properties. No ordering has been observed (or indeed sought) in Tl-2201.

1.4.4   Symmetry and Cation Substitution

One potential drawback to Tl-2201 is that it can exist in either of two symmetries, orthorhombic and tetragonal, which are commonly believed to be chemically distinct, the former with the stoichiometric formula $Tl_2Ba_2CuO_{6\pm\delta}$, and the latter a defected phase, with roughly $\sim 8\%$ of Tl atoms replaced by Cu[25, 66]. Prior to this work, all crystals measured have been tetragonal. Orthorhombic symmetry, or (worse) a change from one symmetry to another, may make the interpretation of some data less straightforward.

The orthorhombic distortion is thought to arise from a lattice mismatch between the CuO_2 plane layer and the naturally larger $Tl_2O_2^{2+}$ double layer. Powder X-ray

Figure 1.10: The CuO_2 plane and its slice through the unit cell (in green) are shown for the **a)** tetragonal phase and the **b)** orthorhombic phase. The orthorhombic distortion is exaggerated for clarity.

diffraction has shown that the orthorhombic distortion is an elongation along one plaquette diagonal, leading to a rhombic CuO_2 plaquette instead of square; the Cu–O bond lengths remain the same in both directions. The distortion is minor: instead of being at right angles, the Cu–O bonds meet at $\approx 89.7°$ in orthorhombic samples. In this circumstance, with the plaquette diagonals (coinciding with the directions of the nodes in the superconducting gap) remaining orthogonal, no mixing of order parameter symmetries should be required. The tetragonal and orthorhombic CuO_2 layers are depicted in Figure 1.10.

There have been doubts expressed as to whether the orthorhombic version of the compound superconducts at all. Orthorhombic ceramic samples have been observed to superconduct, but it hasn't been clear to all in the field whether this superconductivity was due to the orthorhombic phase itself or due to filamentary superconductivity of tetragonal phase impurity regions.

In fact, Tl-2201's crystal symmetry has been shown to depend on the level of cation substitution and the oxygen content — interstitial oxygens increase the distortion, while cation disorder seems to suppress it. Stoichiometric and near-stoichiometric ceramics are orthorhombic for all oxygen contents, while heavily-substituted ones are strictly tetragonal [66]. However, for intermediate levels of substitution, the orthorhombicity depends on the oxygen content: samples with higher oxygen contents
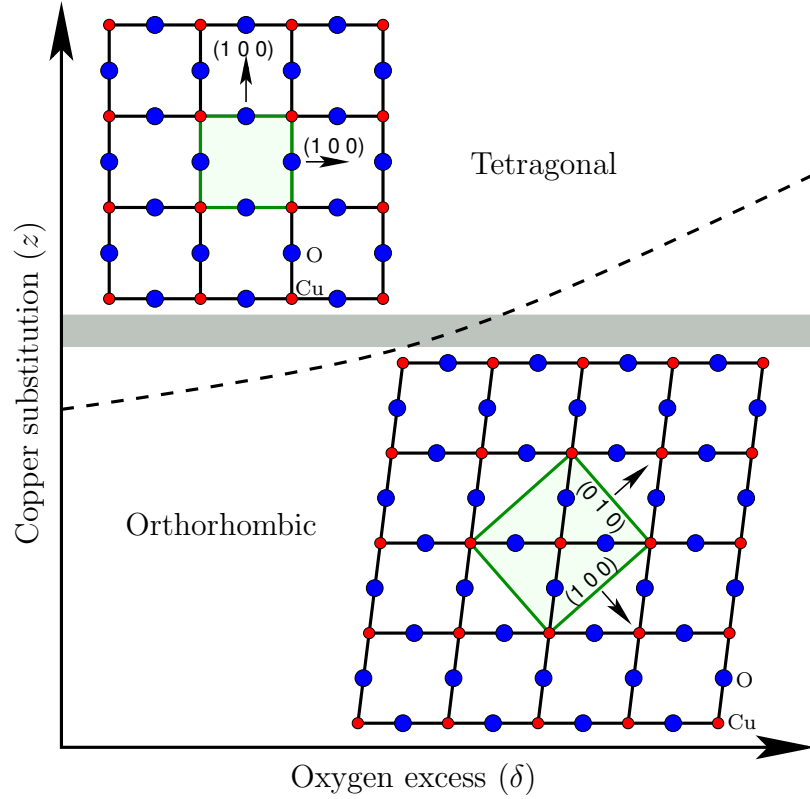
Figure 1.11: Schematic depiction of the dependence of the orthorhombic/tetragonal structural transition on z and δ in $\text{Tl}_{2-z}\text{Ba}_2\text{Cu}_{1+z}\text{O}_{6+\delta}$. This scenario is suggested by the data from polycrystalline samples of Wagner et al. [47] and our own single crystal diffraction results as reported in Section 2.5.4 and reference [16]. The gray bar indicates where our crystals fall on this plot: z=0.080. More heavily overdoped materials tend to be farther to the right or down in this phase diagram.

are more orthorhombic [47]. Figure 1.11 shows the inferred schematic phase diagram relating the symmetry to the oxygen content δ and the copper substitution z.

1.5   Electronic Structure

Fundamentally, superconductivity is an effect of the interactions between charge carriers. Electrons (or, in the case of most cuprates, holes) near the Fermi level are paired via some boson, and these Cooper pairs (now bosons themselves) behave collectively, condensing to form a superconducting phase [67] (the language of Fermi liquid theory is used here for simplicity — there may not be distinct Fermions or a Fermi surface in these materials).

In conventional superconductors, those explainable through BCS theory, the energy saved due to pairing of carriers results in an energy gap at the Fermi surface — to excite an electron or hole, one must first supply the energy required to break the pair in which it resides, which is Δ per particle. The pairing strength, which is the order parameter for superconductivity, generally does not depend substantially on the carriers' direction in k-space and does not break the symmetry of the lattice, so these materials are denoted s-wave. Many elemental metals, including mercury, aluminum, lead, tin, and niobium are conventional superconductors at low temperatures, as are a number of simple compounds such as V_3Si, Nb_3Sn, or MgB_2. An s-wave pairing gap is shown schematically in Figure 1.12a.

Figure 1.12: Schematic depictions of an **a)** s-Wave and a **b)** d-Wave pairing gap shown at the Fermi surface in k-space. The zone centre (Γ point) is at the centre of each panel, k_x and k_y are horizontal and vertical respectively, states in the centre of each zone are occupied, and the dashed line represents the Fermi surface.

In the cuprates, by contrast, the order parameter has the symmetry of a $d_{x^2-y^2}$ orbital, changing sign (quantum mechanical phase) at $\pi/4 \pm n\pi/2$ [9, 11, 12] (zero is taken along the Cu–O bond direction). These materials are denoted d-wave; the order parameter's phase change breaks underlying symmetries of the lattice. The nodal points (or lines, taking into account the c direction) constitute locations in k-space where the order parameter passes linearly through zero, changing phase as it does so. A schematic depiction of a d-wave pairing gap is shown in Figure 1.12b. At low temperatures, where transport properties are dominated by low-lying excitations,

these nodes dominate all observed transport behaviour. However, the addition of impurities or strong scattering sites alters these nodal areas substantially, leading to the observation of different power laws, or even exponentially activated behaviour, which would not otherwise be seen [68, 69, 70]. For the purposes of fundamental research on the electronic structure of the cuprates or any other superconductor with nodes in the pairing potential, the availability of pure, well-ordered samples is crucial.

In this discussion, dispersion along the c direction is neglected, because it is small [14]. The cuprates are essentially two-dimensional both in terms of the crystal structure (CuO_2 plane layers) and electronic structure. They do superconduct along the c direction, albeit weakly, through the hopping or tunnelling of Cooper pairs. The crystal structure is also assumed to be tetragonal, rather than orthorhombic, for simplicity — some orthorhombic crystal structures require a small mixing of order parameter symmetries.

Once the abovementioned cleanliness criterion is satisfied, transport properties can yield important information on the intrinsic electronic structure and the pairing in these compounds. Characterizing the cuprates' electronic structure is a key, but not necessarily a sufficient, condition for understanding their behaviour.

1.5.1   Correlated Electron Physics

Adding to the challenge of understanding the cuprates, their electronic structure is dominated by strong electron-electron correlations, leading to a wide variety of consequences in the electronic structure, many of them counterintuitive [71, 72]. Copper atoms in an undoped CuO_2^{2-} plane are Cu^{2+} (electronic configuration $[Ar]3d^9$) — with a half-filled $3d_{x^2-y^2}$ band, this material would be expected to conduct. However, the $3d$ orbitals have the same radial probability distribution as a $1s$ orbital, but with the capacity for holding up to ten electrons, making electron-electron interactions crucial. Indeed, the on-site repulsion all but precludes two carriers sharing a copper atom's $d_{x^2-y^2}$ orbital, localizing the carriers, and the undoped cuprates are antiferromagnetic Mott insulators to several hundred Kelvin.

One theoretical model commonly used as a starting point for explaining this behaviour is the Hubbard Hamiltonian [73, 74]

$$H = U \sum_i c_{i,\downarrow}^\dagger c_{i,\downarrow} c_{i,\uparrow}^\dagger c_{i,\uparrow} - t \sum_{<i,j>,\sigma} \left(c_{i,\sigma}^\dagger c_{j,\sigma} + c_{j,\sigma}^\dagger c_{i,\sigma} \right) \tag{1.4}$$

where U is the on-site Coulomb repulsion, t is a kinetic energy term, the energy

associated with a hop to the nearest-neighbour copper site, while $c_{i,\sigma}^{\dagger}$ and $c_{i,\sigma}$ respectively create and annihilate a hole with spin σ at site i. The second sum is over nearest neighbours only. The onsite repulsion U in the cuprates is typically a few eV, effectively preventing double occupancy, while the hopping term promotes antiferromagnetism — anti-aligned nearest neighbour holes can hop onto each others' sites and back without violating the Pauli exclusion principle, lowering their energy. A common simplification is the $t - J$ model [75, 76], where double occupancy is projected out and this virtual hopping process explicitly considered. The simplest version of this, with only nearest-neighbour exchange, has the effective Hamiltonian [77]

$$H = -t \sum_{<i,j>,\sigma} \left(\hat{c}_{i,\sigma}^{\dagger} \hat{c}_{j,\sigma} + \hat{c}_{j,\sigma}^{\dagger} \hat{c}_{i,\sigma} \right) + J \sum_{<i,j>} \left(\vec{S}_i \cdot \vec{S}_j - \frac{1}{4} n_i n_j \right) \qquad (1.5)$$

where $n_{i,\sigma} \equiv c_{i,\sigma}^{\dagger} c_{i,\sigma}$ is the hole number operator, $\hat{c}_{i,\sigma} \equiv c_{i,\sigma}(1 - n_{i,-\sigma})$ is an adapted form of $c_{i,\sigma}$ that projects out doubly-occupied states, $J = 4t^2/U$ is the exchange coupling constant, and $\vec{S}_i = c_{i\alpha}^{\dagger} \vec{\sigma}_{\alpha\beta} c_{i\beta}$ is the spin operator at site i ($\vec{\sigma}_{\alpha\beta}$ refers to the Pauli spin matrices).

In a material described by the above Hamiltonians at half-filling, the band that would otherwise be at the Fermi level splits, and the electronic structure is best described by electron-addition and electron-removal states of the system, the upper and lower Hubbard bands (UHB and LHB) respectively, which are separated in energy by U. In the cuprates, a further modification to this picture is required — the charge transfer energy from copper to oxygen is less than U, making the material a charge-transfer insulator in the Zaanen-Sawatzky-Allen scheme [72, 78]. When holes are added to the plane, the lowest-energy states they can occupy are not on the copper atoms themselves, but primarily on the oxygen ligands.

Although the electronic properties of the CuO_2 plane can only be properly determined when all n electrons are considered simultaneously, the n-particle basis states required in the infinite systems preferred by theorists, or even in real systems where n is merely $\sim 10^{26}$, are not amenable to most theoretical approaches. It is therefore common to create single-particle basis states that have some of the correlation effects accounted for internally. One common such basis state is the Zhang-Rice singlet [77], which consists of an appropriate combination of oxygen $2p_x$ and $2p_y$ orbitals with a copper $3d_{x^2-y^2}$ orbital. The Zhang-Rice singlet state is depicted in Figure 1.13. By sharing the hole between the copper atom's $3d_{x^2-y^2}$ orbital and its neighbouring oxygen ligands' $2p_x$ or $2p_y$ orbitals, this state has a lower energy cost than simply adding the hole to the copper atom. There is another small energy savings from hopping if

the spin is anti-aligned with that of the copper hole, hence its singlet nature.

Figure 1.13: Zhang-Rice singlet: a hole is delocalized on the four oxygen $2p$ orbitals to avoid the onsite repulsion at the copper site.

If Zhang-Rice singlets are created to accommodate holes as they're added to the system, they would manifestly exist at the Fermi level. Promoting a core electron to the Fermi level would cost the least energy if it annihilated a Zhang-Rice singlet, which is constructed to accommodate holes that can't otherwise easily be accommodated. Their destruction represents the lowest-energy excitations of the system, at least at the eV scale (note that the energy scale of the superconductivity that ultimately arises is tens of meV). It should be noted that the validity of this basis state requires that the added holes be dilute — holes around adjacent copper atoms would need to share an oxygen orbital, requiring a modified basis state. Near or above a doping level of $p = 0.25$, where every copper with an extra hole would expect to have one nearest-neighbour copper atom with an excess hole, there's no reason to imagine that this would continue to hold.

Parts of the crystal structure that do not strongly interact with Cu atoms (such as the $Tl_2O_2^{2+}$ layer) and contain no other $3d$ or $4f$ elements are not expected to exhibit strong correlation effects, and it's thought that their band structure may be accurately calculated by conventional means.

X-ray absorption and emission spectroscopies directly access electronic structure

on the atomic scale, showing these correlation effects, effects which may be partially responsible for the high-temperature superconductivity.

2. Crystal Growth and Preparation

As mentioned above, experimental studies of the HTSCs rely critically upon the availability of high-quality single crystals. This chapter details some of the considerations involved in the growth of single-crystalline Tl-2201, and describes how it was performed in the present work, along with the annealling process. Finally, the various methods used to characterize the crystals and their respective results are summarized.

2.1 Tl-2201: Starting Point

Despite Tl-2201's significant advantages as an overdoped single-layer cuprate, very few experimental techniques have been applied to it, due largely to a dearth of high-quality single crystals. The toxicity and volatility of thallium oxides make preparation of the material especially challenging, and only a handful of groups have succeeded in preparing single crystals.

A. P. Mackenzie's group, then at Cambridge, developed a technique involving an alumina (Al_2O_3) crucible with a tight-fitting gold lid as a barrier to the diffusion of thallium oxide vapours [25, 79]. They found a two-degree by two-minute window for successful crystal growth — any longer or higher, and the thallium would be lost, any cooler and Tl-2201 would not form. This crystal growth effort produced a significant number of crystals up to $\sim 1 \times 1$ mm^2 that have been studied extensively.

N. N. Kolesnikov and collaborators, in Chernogolovka, Russia, grew crystals using alumina crucibles [80, 81, 82]. Few details of the technique were published, but the crystals produced were typically larger, up to a few millimetres across. They were also reported to have transition temperatures as high as 115 K, while no other group has observed a T_c above 94 K on either a ceramic or single crystalline sample. One early paper [80] from this group mentions that the crystals grew far better in the presence of calcium — if that practice continued, it is possible that some of these crystals may have had intergrowths of $Tl_2Ba_2CaCu_2O_8$ (Tl-2212) or $Tl_2Ba_2Ca_2Cu_3O_{10}$ (Tl-2223) which could have raised the transition temperature of the remainder of the crystal by proximity effect [83]. They do not report any special precautions for controlling the loss of thallium. This group produced a substantial number of crystals that were

used for many techniques, but have not been producing crystals of this material in the past year or two.

M. Hasegawa *et al.*, then at the University of Tokyo, grew ~ 1 mm^2 Tl-2201 crystals using an encapsulation scheme [84, 85, 86]. Their mixed powders were pre-reacted, then sealed in a gold crucible within an alumina or steel bomb for growth. No research has been published on these crystals outside the originating group.

Finally, in the prelude to this work, the author's M.Sc. work produced crystals using a weak encapsulation scheme, with a 1 kg mass compressing a gold seal between an alumina crucible and lid [16, 87, 88]. Crystals obtained were roughly 1 mm^2 in size, but the technique did not have the desired reproducibility. The present work set out to improve upon that technique. The growth parameters reported in the literature are summarized in Table 2.1.

Table 2.1: Previously reported growth parameters. It is not clear from the Kolesnikov group's papers whether they used inverted crucibles as lids or no lids at all. This difference would likely have had little impact on their results.

Group	Tl : Ba : Cu	Crucible	Lid	Temperature Program
Torardi [89]	1 : 1 : 1	Au	Au	"Slow cooling"
Liu [79]	1 : 1 : $\frac{3}{2}\sim 2$	Al$_2$O$_3$	Au	910–870°C in 40 h
Tyler [25]	1.34 : 1 : $\frac{3}{2}$	Al$_2$O$_3$	Au	892$\sim$895°C, 3 min
Kolesnikov [82]	2 : 2 : 1	Al$_2$O$_3$	Unclear	950°C, -5°C/h
Kolesnikov [81]	1 : 1 : 1	YSZ	Unclear	950°C, 30 min
Hasegawa [86]	1 : 1 : 1	Au	Au	10 h@915°C, 10 h to 895°C
Hasegawa [85]	1 : $\frac{5}{4}$: 1	Au	Au	10 h@920°C, 57 h to 880°C
Peets [16, 87]	1.05 : 1 : 1	Al$_2$O$_3$	Au	935–890°C in 90 h

It is worth noting that none of these techniques allowed decanting. In all cases, the product was a solidified ingot of flux, with the Tl-2201 crystals enveloped in the flux or occasionally occupying voids within the flux ingot. Crystals were harvested by breaking the crucible and flux ingot open, then mechanically separating the crystals and flux to the extent possible.

Annealing of Tl-2201 crystals has proven even less successful. Powders and ceramics have been annealed successfully [66, 90], but in the cases of ceramics or powders surface decomposition due to loss of thallium oxide may go unnoticed. Crystals, with their shiny surfaces, visibly tarnish if any decomposition occurs. The most comprehensive annealing study to date is [90].

In the case of crystals, Mackenzie's group was able to anneal them, without sig-

nificant damage, to T_cs ranging from 4 K to 85 K. Hasegawa reports attempts at annealing by holding the ingot at an elevated temperature for 24 hours during cooling [85], while Kolesnikov's group reported [82]:

> Only as-grown crystals have been used in the measurements. Some attempts at adjusting the transition temperature of crystals by means of annealing them in an oxygen atmosphere have been done which have revealed that the crystal quality deteriorates strongly after such a procedure. And though their transition temperature may be indeed reduced in such a way down to zero, they have been found to be unsuitable for measurements.

Annealing was not within the scope of the earlier work at UBC.

It must be noted that because much of the annealing work performed to date has concentrated on reaching optimal doping, which requires removal of excess oxygen from the material, it has frequently been performed in atmospheres such as argon or hydrogen, or in "vacuum". In flowing argon, the oxygen partial pressure around the crystal, which sets the oxygen content therein, may vary by several orders of magnitude over the course of an anneal, or be determined by the concentration of oxygen impurities in the flowing argon. Annealing under "vacuum" (i.e. with a vacuum pump attached) will result in a similar time-varying oxygen partial pressure, determined by the amount of material being annealed, the pumping speed, and any other sources of oxygen (e.g. outgassing). In the case of hydrogen annealing, oxygen molecules react with hydrogen to form water; however, the annealing temperature must be kept quite low to avoid decomposition of Tl-2201, so the equilibrium between water, oxygen and hydrogen vapours is not established, and the oxygen partial pressure is not well defined. None of these processes can be considered controlled, nor can they be expected to produce the stable chemical equilibrium required for a homogeneous oxygen content.

Because of the abovementioned issues with annealing, T_c widths reported in the literature are typically $\sim 10 - 20$ K. This level of inhomogeneity is sufficient to blur many features, obscure transport power laws, and make the interpretation of doping-dependent studies extremely difficult.

2.2 Growth of Single-Crystalline Tl-2201

For crystal growth a copper-rich self-flux technique with novel variable encapsulation was employed, which is described below. Determining the optimal parameters for

crystal growth in a system with complications such as a lack of phase diagrams, three independent components, toxic ingredients, and a volatile component can be long and tedious — the parameter space has a very large number of dimensions, which are frequently not orthogonal. The effects of small variations in a handful of parameters will be discussed, but most parameters will simply see their optimal values reported.

Figure 2.1: The furnace used for crystal growth. The assembly surrounding it permits pressure to be applied to rams running through the centre of the furnace, and enables their controlled rotation.

The furnace, shown in Figure 2.1 is a Thermolyne 79300 tube furnace, with an apparatus attached to allow the application of pressure to rotatable rams in the heated zone. The furnace includes a Platinel-II thermocouple and a controller capable of several different programs, each of up to 8 ramp rate–level–dwell steps. The crucible is located in the centre of the heated zone, within a large quartz furnace tube. This tube is sealed, and filled with flowing oxygen gas. The oxygen gas is introduced at the upper (right-hand) end of the furnace tube, then exits via two small quartz exhaust tubes that start near the centre of the heated zone. Removing the exhaust gases from the centre of the heated zone instead of from the cooler end helps ensure that any thallium oxide vapours that condense out do so inside the small, easily replaced quartz tube, rather than on the larger, more expensive, difficult to replace furnace tube or, worse, on another piece of apparatus.

A cutaway schematic of the furnace and some of the attached apparatus is shown in Figure 2.2. Pressure is applied via a pneumatic ram located to the right and fed with building compressed air and a pressure regulator, and transferred to the crucible's

Figure 2.2: Cutaway view of furnace, showing the key components. The function of each is discussed in more detail in the text. The unlabelled channels through the brass are for exhausting oxygen, and are plugged at one end of the furnace. Cast alumina plugs inside the furnace are used to prevent radiation from damaging the brass.

gold seal via an alumina ram. Gears and a rod through the furnace's casing allow that alumina ram and another inserted from the other end of the furnace to be rotated together without introducing strains from twisting. The alumina rams are aligned and held in place with pieces of brass as shown in the figure, equipped with sliding seals to allow the translation and rotation required.

The right-hand end of the furnace is placed on a car jack to allow the angle of the furnace to be varied. Its lower element is connected in series with a space heater via a dimmer switch, to reduce its share of the current and create a vertical temperature gradient through the furnace. The gradient was set to be 25°C across the 74 mm quartz furnace tube, or about 3.4°C/cm. The actual gradient across the crucible when filled with melt is not known as it cannot be readily measured without compromising a seal.

Oxygen gas at one atmosphere is passed through the furnace tube at the rate of ∼1 mL/s, or enough to flush out the unoccupied volume in the furnace tube in roughly twenty minutes at room temperature. The oxygen gas is not intended to mix

with the contents of the crucible, which should be at a higher pressure than the rest of the furnace at all times — it is intended to sweep away evolved gases, particularly carbon dioxide and thallium oxides. Accordingly, high-purity oxygen has not been used. Upon being removed from the furnace, the oxygen gas is bubbled through sulphuric acid and water, to capture any thallium oxides it may be carrying. The scrubbed oxygen is then released into a fumehood. The furnace itself is kept in the fumehood in case of a failure of the thallium oxide containment systems.

The crucible used is a (nominally) 20 mL CoorsTek CN-20 crucible, made from 99.8% pure alumina, although most testing and earlier growth runs were in smaller CN-10 (10mL) crucibles. Alumina cannot be used for many cuprates because of their corrosive melts or sites that can accept Al^{3+} as an impurity ion, but Tl-2201 is not believed to have any such site, and no corrosion of the crucible has been observed. The crucible does get impregnated by thallium oxide vapour to some degree, but this generally does not make it through the crucible wall. Before the crucible can be used, the rim must be flattened and made parallel to the base, to ensure proper sealing. Some degree of flexibility has been designed into the system through use of gold pads, but best results are obtained when the height of the crucible does not vary by more than 0.15 mm and the rim has been well-flattened for a uniform, controlled seal. A gold sieve is inserted into the crucible for decanting (a process described below), and its position marked on the exterior of the crucible for alignment purposes.

To make the powder mixture, precursor powders of CuO (Alfa Aesar Puratronic, 99.995% pure, ground fine), BaO_2 (Aldrich, 95% pure) and Tl_2O_3 (Alfa Aesar 99.99%) are weighed out into a glass vial, and shaken (not stirred) for several minutes to ensure they are intimately mixed. They are then further mixed with a mortar and pestle to eliminate any clumps of Tl_2O_3. Barium peroxide was chosen for its low (tolerable, but non-zero) carbonate concentration, to minimize the production of oxycarbonates [61].

The cation ratio that produced the best crystals was Tl : Ba : Cu = 1.1 : 1 : 0.9, but this can be varied to some degree. Higher thallium contents were challenging to seal and typically produced fewer, smaller crystals with poorer-quality surfaces. Slightly lower thallium contents (e.g. Tl : Ba : Cu = 1.05 : 1 : 0.9) typically produced more and slightly larger crystals, but significantly lower thallium contents produced no particular benefit. Far greater improvements to crystal size and yield were realized by increasing the volume of the melt beneath the gold sieve, providing better opportunities for convection. It was desired that the cation substitution be minimized to the extent possible, so high thallium contents were favoured. Changing the barium

content led to deleterious effects on crystal quality and yield. The composition window within which Tl-2201 will grow is small, limiting the parameter space that can be productively explored.

All steps involving the mixing or dispensing of Tl_2O_3 or mixed powders containing it are performed in a glovebox under negative pressure for safety, and as many steps as possible are done without thallium present, for convenience.

The mixed powder is loaded into the crucible and tamped down with a plastic rod, the crucible rim is cleaned, and a gold disc and Al_2O_3 lid are glued on top with DevCon 5-minute epoxy[2]. The epoxy is only applied to the outside of the crucible, not to the rim itself or the interior of the crucible, to minimize carbon contamination; 5-minute epoxy's high viscosity ensures that the epoxy stays where it is applied. Once this has cured, the crucible is epoxied to the end of an alumina ram, which is then inserted into the furnace. Gentle pressure ($\sim$2 atm) is applied to the crucible by the pneumatic ram, to hold it in place once the epoxy burns off, and the furnace is tilted to 23°.

The epoxy is expected to burn off to CO_2 and H_2O around 200°C — by this time, the atmosphere in the furnace will have been fully purged several times, and will essentially be pure oxygen. If the crucible is sufficiently open to the furnace to allow some CO_2 in, it is open enough to allow the flowing oxygen gas to purge it out. The carbon left in the crucible should be only that in the form of $BaCO_3$ — the main impurity in BaO_2 — and the roughly 400 ppm naturally occurring in the $\sim$10-15 mL of air in and above the powder.

Figure 2.3 depicts the growth process, including initial assembly, melting once the furnace is tilted, and decanting, and Figure 2.4 shows the temperature and pressure programs used. Referring to the temperature at the centre of the furnace's heated zone, the furnace is heated at 100°C/h to 400°C, then at 50°/h to 600°C, where the barium peroxide is allowed to fully decompose for four hours. The furnace is then brought to 935°C at 100°C/h with a one-hour pause at 820°C to let the thallium oxide equilibrate. The melt is left at 935°C for just over two hours to ensure complete reaction of the precursors and to allow all O_2 and CO_2 to evolve and escape. Then, the pressure applied to the rim is increased to 22.0 atm (after accounting for 1/3 atm for balancing thallium oxide's partial pressure), sealing the crucible for the remainder of the growth. This can require rotating the rams slightly to ensure the transfer of pressure to the crucible lid, since the seals around the rams can sometimes bind.

From this point, the temperature is reduced at -2°C/h to 905°C, closer to the

---

[2]5-minute epoxy takes a minimum of 20 minutes to set.

PSfrag replacements

Figure 2.3: Schematic depiction of the decanting scheme. **a)** With the crucible filled with the mixed powders, it is inserted between two rams, with a gold disc to seal it against an Al_2O_3 lid. **b)** Once the furnace is tilted and the powder melted, slow cooling allows crystals to form within the melt. **c)** The rams are rotated, separating the crystals from the excess flux with the aid of a gold sieve.

Figure 2.4: Temperature and pressure programs for crystal growth. Several minutes before the temperature is reduced for the cation anneal, the rams and crucible are rotated by 180°. Pressures are calculated for the rim, assuming that 1/3 atm is consumed balancing the thallium oxide's vapour pressure. An oxygen overpressure inside the crucible would significantly reduce these numbers.

liquidus. Crystals are then grown by reducing the temperature at $-0.3°C/h$ to $890°C$, just above the eutectic point, which we found to be slighly lower than that shown in Figure 1.5. The furnace program pauses for up to 24 h at this temperature to allow the user to decant the melt.

Decanting (shown schematically in Figure 2.3c) involves rotating the rams by 180° over the course of ~15-20 minutes. This rotates the crystals up out of the melt while leaving enough time for the melt to flow out, and being gentle enough that the seal is not broken in the event of any misalignments, for instance due to a crucible rim not being perfectly parallel to its base. The gold sieve across the crucible ensures that the crystals do not fall back down into the melt once raised clear of it. Once the crystals have had 5-10 more minutes to fully drain, the temperature is reduced, quickly solidifying the excess melt. After a final anneal step, power is cut to the

furnace elements, and the system cools freely to room temperature.

If there is a significant oxygen overpressure inside the crucible in addition to the thallium oxide vapour pressure, the actual pressure applied to the gold seal will be substantially lower. For instance, in the early, lower-pressure part of the process, if the pressure is equal on the rim and over the crucible (due to an oxygen overpressure and burping of the crucible), the pressure will be 0.85 atm, whereas if there is only a 1/3 atm overpressure, the pressure on the rim will be 2.4 atm. The former scenario is more likely at higher temperatures, and this scheme relies upon the ability to release pressure gently this way.

In general, annealing at a temperature slightly below the growth temperature can relax stress and reduce defect concentrations in the crystals. After other growth conditions were optimized, a 48-hour annealing step was appended at a temperature $\sim$50-100°C cooler than the eutectic (which the furnace reaches via a -100°C/h ramp), followed by free cooling. At these temperatures, cations are still slightly mobile within the crystals, and the atmosphere of oxygen and thallium oxide vapours endures. This allows the substituted copper atoms to distribute themselves more homogeneously within the crystal, but is unlikely to permit any substantial net change in composition, since the crystal is no longer in contact with the melt. Such a high-temperature anneal would be exceedingly difficult to perform once the crucible's seal has been broken, due to thallium loss.

The crucible, now containing crystals on one side and solidified flux on the other, is broken open parallel to the gold sieve, and crystals are picked out from the large mass of crystals typically present. Single crystals harvested in this manner are typically 1$\sim$2 mm^2 and 10$\sim$100 μm thick. As shown in Figure 2.5, these crystals are irregularly shaped and exhibit mirror surfaces, frequently with fine curved growth steps. All work reported here was performed on flux-free single crystals.

The mass of the crucible and its contents is measured before and after growth, to monitor the amount of thallium lost. Typically, the mass lost is less than that expected from the loss of one oxygen atom from BaO_2 and two from Tl_2O_3 (it would require the unphysical *gain* of $\sim$2% of the remaining TlO). It should be noted, however, that while Tl_2O is stable at growth temperatures, Tl-2201 contains strictly trivalent thallium, and the majority of the thallium in the crucible may take the form $Tl_2Ba_2O_5$, which is also stable.

If we instead assume that no oxygen evolved by thallium is lost, the mass loss unaccounted-for by barium deoxygenation would constitute $\sim$5% of the thallium atoms. The truth should lie somewhere in between — some of the oxygen evolved

Figure 2.5: Optical micrographs of several resulting Tl-2201 crystals. The crystals in pane **a)** are on millimetre graph paper for scale, but the scale in pane **b)** is not the same.

by the Tl_2O_3 must escape, and it's likely that some thallium oxide does as well, but not more than a few percent. If CO_2 lost from barium carbonate in the BaO_2 were accounted-for, the thallium loss percentages calculated would be marginally lower, but as the extra carbon and oxygen atoms lost are light (compared to barium) and carbonates only make up a small fraction of one constituent, that should be a small effect. A further small effect is that the crucible appears to become impregnated with thallium oxide to some extent, which would result in a reduced quantity of thallium available for crystal growth.

Since the thallium, barium and copper oxides dealt with in this work are acid-soluble, gold and unbroken ceramic components that have come in contact with the melt or the vapours inside the crucible may be reused after cleaning in acid. The gold lid is cleaned, melted, and flattened into a new disc between growths. Most metal impurities are drawn to the surface through this process by the chemical potential, because they oxidize there; most non-metallic impurities do not mix with gold and float on its surface because of their lower density. This allows impurities to be readily dissolved off or removed mechanically, and the gold may be reused many times. The chief exceptions would be noble metal impurities (i.e. silver), but these should not be present in this system.

Over more than 100 growth runs, this growth technique was optimized to the point where crystals could be grown successfully, with minimal loss of thallium, roughly 75%

of the time. Yields can fluctuate wildly, but typically range from 10 to 50 crystals, limited by patience — smaller crystals are far more plentiful, but crystals smaller than $\sim 1/10$ mm^2 are frequently not collected.

No group has previously reported either a decanting scheme, nor a mechanism to allow the evolution of gases before full encapsulation. Hasegawa did pre-react his powders to enable encapsulation, but the evolved gases still caused some of his crucibles to rupture. The work reported here constitutes the first controlled, reproducible technique for growing Tl-2201 single crystals with the capacity to decant the melt.

2.3 Hole Doping in Tl-2201

As was mentioned in Section 1.4.3.1, Tl-2201's hole doping may be controlled reversibly by varying the oxygen content, permitting ready access to most of the overdoped regime. The location of the oxygen interstitial was shown in Figure 1.4. In a simple picture, adding one oxygen atom would dope two holes into the CuO$_2^{2-}$ plane, to create an O^{2-} ion. However, this picture is complicated by the cation substitution.

There have been indications that increased copper substitution levels are correlated with lower T_cs and steeper changes of T_c with oxygen dopant concentration [47, 66], but no explanation has been offered and this has not been studied in detail in this or previous work. A possible explanation is offered here, and more light will be shed on this issue in Section 5.5.

Table 2.2: The first four ionization potentials of copper and thallium, in eV. Taken from Reference [91] except for the fourth ionization energy of Tl, which is from [92].

	1^{st}	2^{nd}	3^{rd}	4^{th}
Cu	7.72638	20.2924	36.841	57.38
Tl	6.108194	20.4283	29.83	50.8

Considering the cations naively as non-interacting and neglecting oxygen, the first few ionization potentials for copper and thallium are listed in Table 2.2; the zero for both is a free electron, so they can be directly compared. Given that the thallium exists as Tl^{3+}, copper atoms at the thallium site are expected to be Cu^{2+} (or [Ar]3d^9), but likely not quite Cu^{3+}. The next electrons removed would be from copper. So long as the copper remains Cu^{2+}, each substituted copper atom would be expected

to dope one hole onto the CuO_2 plane. This may help to explain why Tl-2201 is considered intrinsically overdoped — for a Tl-2201 sample with cation composition $Tl_{2-z}Cu_{1+z}$, the starting point for oxygen doping would be $p = z$, not $p = 0$ as for other materials, and z is commonly 0.1–0.15 (optimal doping is 0.16). It must be noted at this juncture that the cation ratio is not used as a control parameter in doping, and crystals grown under the same conditions should, in principle, have very similar compositions.

When oxygen is added, incorporated as O^{2-}, one would normally expect each oxygen atom to dope two holes into the plane. However, it is also possible that some of these holes may end up on substituted copper atoms, making them Cu^{3+}. As alluded to, the same change in T_c was found to require roughly half the change in oxygen content in stoichiometric samples as in samples with $z \sim 0.1$ [66] (the change in oxygen content did not exceed the substitution level in that study). This suggests that each dopant oxygen may collect one hole from a nearby substituted copper atom and one hole from the plane.

As mentioned, this doping discussion, as with arguments purporting to explain carrier doping in all but a handfull of cuprates, is quite simplistic, and independent checks of its validity are required. It suggests the existence of weak Cu_{Tl}-O(4) electronic structure near the Fermi energy that may be visible by ARPES or quantum oscillations, given sufficient sensitivity. No such feature has been seen.

Determining the mechanism by which chemical doping is translated into hole doping is extremely difficult in most cuprates, and Tl-2201 is no exception. Indeed, it is common to work solely in terms of chemical doping, then determine an approximate hole doping via T_c when necessary using formulas determined for other materials. X-ray spectroscopies are among the few techniques that can help elucidate the mechanism, provided the dopant concentrations are high enough to be observed and the changes in the spectra are significant.

2.4   Annealing

As-grown crystals typically have superconducting transitions of either $\sim$90 K or 5 $\sim$ 10 K (onset) and several Kelvins wide, a width comparable to other groups' as-grown samples (note that this was not often measured in the current work). For a benchmark, well-annealed high-purity $YBa_2Cu_3O_{6+\delta}$ crystals typically have T_c widths of less than a Kelvin [41], although this figure of merit is strongly dependent on the superconducting dome's slope dT_c/dx at the doping x being considered. As-grown

T_cs in Tl-2201 are determined by the cation substitution level, crystal dimensions (it takes longer times for oxygen to diffuse throughout a larger crystal), cooling rate, and the atmosphere around the crystals, the last two parameters being coupled through the equation of state for oxygen gas. Since the crystals' as-grown superconducting transitions are typically quite broad, it is crucial that the crystals be annealed before they are studied beyond a simple visual inspection. Annealing is not used only to obtain the desired T_c, but also to make the oxygen content (and carrier doping) homogeneous as described in Section 1.4.3. Controlled post-growth annealing of Tl-2201 is experimental territory largely unexplored in the literature.

The crystals were annealed under controlled oxygen partial pressures at temperatures between 290°C and 500°C. This produces samples whose oxygen contents place them on the overdoped side of the phase diagram, with superconducting transitions that range from 5 to 85 K and are sharp. The measurements of these widths are discussed more thoroughly in Section 2.5.1. We have not attempted to reach optimal doping, which in this system is believed to correspond to a superconducting T_c of at least 93 K [47], because even higher annealing temperatures would be required, increasing the risk of thallium loss.

As mentioned earlier, thallium oxide has a partial pressure dependent on the oxygen partial pressure and temperature [56]. Since obtaining most T_cs requires the removal of oxygen from the Tl-2201, most annealing requires low partial pressures of oxygen. This in turn will promote the evolution of Tl_2O, and thus the annihilation of the crystals' surfaces. It is thus unsurprising that annealed Tl-2201 crystals are scarce.

An additional complication is that quartz and Tl-2201 react with each other. Monovalent thallium oxide is chemically similar to monovalent sodium oxide, which is easily incorporated in quartz, making it brittle and cloudy. Where Tl-2201 is in proximity to quartz, any vapour pressure of thallium will result in erosion of the quartz while decomposing the Tl-2201 — this amounts to putting a thallium sink in contact with a thallium source. Anneals at high temperatures or low partial pressures of oxygen are particularly susceptible to this problem.

Figure 2.6 depicts one annealing scheme developed by Ruixing Liang. The crystals and some Tl-2201 powder are contained in a sealed quartz capsule, one end of which contains pieces of porous ceramic to serve as an oxygen buffer. The buffer may be any system that has a well-established relationship between oxygen partial pressure and temperature, has enough free oxygen atoms to stave off depletion, does not react significantly with quartz, and has no volatile components that may contaminate

Figure 2.6: Oxygen-buffered annealing scheme. The oxygen buffer may be any system
that has a well-characterized relationship between oxygen partial pressure
and temperature, adequate free oxygen atoms to enforce this relationship,
and no other atoms with vapour pressures that could contaminate the
Tl-2201. Tl-2201 crystals are embedded in Tl-2201 powder as a thallium
buffer, and the two ends' temperatures may be controlled independently.

the Tl-2201. Common buffers were $YBa_2Cu_3O_{6+\delta}$, CuO/Cu_2O, and Cu/Cu_2O. For
example,

$$CuO \rightleftharpoons \frac{1}{2}Cu_2O + \frac{1}{4}O_2 \uparrow$$

$$P_{O_2}^{\frac{1}{4}} = e^{-\frac{\Delta G_o}{RT}} \tag{2.1}$$

The Tl-2201 crystals are embedded in Tl-2201 powder (calcined from the same
Tl_2O_3, BaO_2 and CuO used for crystal growth, but in a stoichiometric ratio) at the
opposite end of the quartz tube. The use of Tl-2201 powder minimizes reactions
from contact between the quartz and crystals, while providing a buffer for thallium
oxide vapours. Having an oxygen partial pressure itself, it helps ensure that the
oxygen partial pressure is homogeneous around the crystals and more robust against

fluctuations in temperature.

The ends of the quartz tube may be held at different temperatures in a multiple-zone furnace, permitting independent control over the oxygen partial pressure and temperature applied to the crystals. This requires that the two temperature zones be adequately insulated, such that the quartz tube's ends are at the desired temperatures and not in a temperature gradient. The quartz ampoule is evacuated at room temperature before heating, to ensure purity.

O$_2$ flow

PSfrag replacements

Figure 2.7: Open annealing scheme. If the anneal requires oxygen partial pressures not far below 1 atm and the temperature required is not excessive, the crystals may be embedded in Tl-2201 powder as in Figure 2.6, then annealed in flowing oxygen or an appropriate mixed gas.

To obtain more strongly overdoped samples, higher oxygen partial pressures are required, and the thallium volatility problems are substantially reduced. The crystals are again embedded in powder, but they are typically placed in an open quartz tube, which is heated in flowing gas. This annealing scheme is depicted in Figure 2.7. The process gas need not be pure oxygen — for partial pressures that remain a significant fraction of an atmosphere, oxygen may be mixed with any inert gas and the technique will still work. The process gas is typically purified before reaching the furnace tube, then bubbled through oil or acid to monitor flow and exclude any contaminants entering from the room.

For a particular oxygen partial pressure, the dopant concentration will be a strongly-varying function of temperature. If the crystal is cooled too slowly after annealing, the surface of the sample may attempt to reach equilibrium under different temperature and pressure conditions, leading to shells of varying T_c. To preserve the intended oxygen content, crystals must be removed from the furnace and quenched at the conclusion of annealing. The quartz tube is removed rapidly from the furnace

and plunged into a bath of ice water.

Determining the effectiveness of the annealling schemes detailed in this section requires that the crystals be characterized by several different techniques. The role of substituted copper atoms in the doping will be revisited elsewhere.

2.5 Characterization

Several characterization techniques were applied to the crystals, to check their homogeneity, degree of crystalline perfection, cation substitution level, and other properties.

2.5.1 SQUID Magnetometry

The crystals' transition temperatures and the widths thereof were characterized using a Quantum Design MPMS SQUID magnetometer [93]. The SQUID magnetometer employs a superconducting magnet to apply highly stable and uniform DC fields, then uses a superconducting pickup loop array around the sample space to measure the sample's resulting magnetization. The pickup loop array uses a second-derivative design — it has three loops, two wound in one direction, and a middle one with twice as many turns wound the other way. This arrangement has zero net windings in the field and is insensitive to field gradients. Because the pickup coil is superconducting, it can detect the presence of a stationary magnetized sample within it — bringing such a sample into one of the pickup loops changes the amount of flux passing through the loop, driving a persistent current.

The coil array is connected to a SQUID sensor system, which supplies a nulling current to ensure that the net current through the pickup coils remains zero. The actual magnetization measurement is of a voltage proportional to this nulling current.

The crystal is mounted between two small quartz rods, within a quartz tube, at the end of a long probe. The probe is inserted into a central vacuum space within the magnet and pickup coils, which is filled with a few torr of helium gas for heat conduction. A heater in the vacuum space is used to control the temperature of the quartz assembly and sample for temperatures above 5 K. A separate mode of operation is automatically employed for lower temperatures, whereby liquid helium is pumped on; this can reach temperatures as low as 1.8 K.

The crystal is pulled through the pickup coil array in discrete steps, the SQUID voltage output is recorded at each step, and the resulting curve is fit to obtain a mag-

Figure 2.8: Schematic depiction of Quantum Design's MPMS SQUID magnetometer. The crystal (black) is held in place by quartz rods centred in a quartz tube, which is attached to the end of a rod. The space around the quartz tube is filled with a few torr of helium gas to aid temperature equilibration, the temperature being set using a resistive heater within this space. Around this are the second derivative pickup coils and the large DC magnet, which are in liquid helium. The quartz tube is pulled vertically through the pickup loops for measurement.

netization point for that temperature and field. Typically, a few such magnetizations are averaged.

Magnetic measurements have the advantage of probing the entire sample volume[3], whereas resistivity measurements can measure a sample to be fully superconducting when only a handful of filamentary paths are superconducting, as these paths short-circuit the rest of the sample. However, in Type-II superconductors, which have a 'mixed' or 'vortex' state that allows penetration of magnetic field through the crystal in vortices containing magnetic flux $\Phi_0 = h/2e$, the transition temperature can appear broadened in high magnetic fields — the technique cannot easily distinguish between

---

[3]There are specific exceptions to this statement, notably the "baked Alaska", a crystal whose exterior has a higher transition temperature than its interior. Below the exterior's transition temperature, the non-superconducting interior is shielded from magnetic field, and is detected as superconducting.

a sample that is only half superconducting and a sample that is fully superconducting but has allowed half of its volume to admit magnetic flux.

Even a perfectly homogeneous crystal of a Type-II superconductor will not have a perfectly narrow transition in its magnetization, as part of its volume continues to permit the penetration of magnetic flux between H_{c2}, the field at which the vortex state gives way to the normal state, and H_{c1}, the field at which the first vortices are allowed to penetrate (or the last vortices are excluded). A magnetically-measured transition width is an upper bound, not an exact value. Because H_{c1} has stronger temperature dependence than H_{c2}, the temperature range over which the vortex state is encountered grows in higher fields, making the transition appear broader than it actually is. For Tl-2201, it is particularly important to measure magnetization curves at low fields (a few Oersteds), because H_{c1} is particularly low in this material.

Figure 2.9: SQUID magnetization curves of several Tl-2201 crystals. Transition widths are typically around 0.75–1.5 K. All data were collected in an applied field of $1 \sim 2$ Oe.

SQUID magnetization curves for Tl-2201 crystals at a variety of dopings are shown in Figure 2.9. These curves were measured in applied fields of $1 \sim 2$ Oe, although the exact value of the field depends upon the field history of the magnet. The superconducting transition temperature of each crystal is reported in the figure's legend,

as determined from the midpoint of each transition. The lowest-T_c crystal was measured in a different operating mode of the SQUID, RSO (reciprocating sample option), which is AC but very low frequency. All curves have been normalized. Converting the voltage measured into an absolute magnetization requires detailed knowledge of the sample's dimensions. Given that the location and width of the transition are the primary results sought, such a calculation was not considered justified in general. The Meissner fractions measured (the fraction of magnetic flux excluded when cooling through T_c) were on the order of $40 \sim 60\%$, which compares well with field-cooled results on other cuprate superconductors and is higher than that observed in YBCO due to weaker pinning of vortices.

Table 2.3: Widths of superconducting transitions (ΔT_c) for data shown in Figure 2.9, along with the anneal conditions and the technique used to moderate the oxygen partial pressure in the annealing process. Note that the transition temperatures obtained are sensitive to the amount of cation substitution — the $T_c = 70.5$ K crystals are believed to have a high level of substitution, for instance. Widths of $1 \sim 2$ K may be obtained throughout the doping range studied.

T_c	ΔT_c	Anneal T	P_{O_2}	P_{O_2} maintained by
18.3 K	1.3 K	475°C	1 atm	Flowing gas: 1 atm O_2
36.2 K	2.9 K	480°C	0.01 atm	Flowing gas: 1% O_2 in N_2
43.5 K	3.0 K	480°C	10^{-4} atm	Flowing gas: 0.01% O_2 in N_2
60.5 K	1.5 K	400°C	10^{-5} atm	Cu_2O/CuO buffer, 650°C
61.2 K	1.3 K	400°C	10^{-5} atm	Cu_2O/CuO buffer, 650°C
68.9 K	0.9 K	400°C	3.9×10^{-7} atm	Cu_2O/CuO buffer, 570°C
70.5 K	1.1 K	400°C	6.7×10^{-10} atm	Cu_2O/CuO buffer, 450°C
72.3 K	1.3 K	400°C	3.9×10^{-7} atm	Cu_2O/CuO buffer, 570°C
74.4 K	1.0 K	400°C	6.7×10^{-10} atm	Cu_2O/CuO buffer, 450°C
75.3 K	0.9 K	400°C	6.7×10^{-10} atm	Cu_2O/CuO buffer, 450°C
90.0 K	1.3 K	*As grown*		

The widths of the superconducting transitions, ΔT_c, are determined from the points at which the sample has reached 10% and 90% of its full diamagnetic signal. The transition widths observed for the crystals shown in Figure 2.9 are reported in Table 2.3. Widths on the order of $1 \sim 2$ K can be obtained throughout the doping range from 18 K to 75 K, indicating that the annealing techniques developed, discussed in Section 2.4, have successfully produced homogeneously doped crystals. Widths are typically narrower for lower dopings (higher transition temperatures),

which is expected since dT_c/dp is lower at those dopings — identical doping inhomogeneity produces a much greater spread in T_c where this is steeper. For this reason, the as-grown crystal with a T_c of 90.0 K, close to optimal doping, has a misleadingly narrow transition width. Variations in T_c between seemingly identical annealing conditions are attributed to differences in thermal isolation between the crystals and buffer, and to variations in cation substitution level — most of the crystals reported on in this table were grown from a slightly less thallium-rich melt than was ultimately settled upon.

As discussed above in Section 2.1, annealing of Tl-2201 crystals has been seldom attempted and even more seldom successful. Hawegawa [85] reported on one crystal, which had a T_c of 25 K with a ΔT_c of 3 K. Mackenzie's group [25] produced T_c widths as narrow as $1 \sim 2$ K at some dopings. Both groups measured T_c resistively which, as mentioned above, can produce narrower widths and be less sensitive to inhomogeneity. Kolesnikov's group did not attempt annealing, working instead with as-grown crystals with T_c widths on the order of 10 K, measured magnetically.

The widths of our crystals' superconducting transitions are narrower than those reported previously. The gold standard, ultra-high-purity YBCO, generally has sub-Kelvin T_c widths, narrower than most of those reported here. However, YBCO is the cleanest cuprate, exhibits no cation substitution, and does not decompose under normal annealing conditions, so this should not be surprising. In addition, Tl-2201 is believed to have a very low lower critical field H_{c1}, which can lead to broadening of the magnetic transition. Given that the widths reported in this section compare very well with many other cuprates, and are the narrowest reported on Tl-2201, it can be concluded that the crystals are quite homogeneous, and the annealing processes used work well.

2.5.2   X-ray Rocking Curve

One key measure of crystalline perfection is the crystal's X-ray rocking curve. In this technique, the crystal is mounted in an X-ray beam, and the detector is aligned for a 2θ value corresponding to a known reflection. The width of that reflected peak is measured by rocking the sample's angle with respect to the beam (θ). Widths are reported as the full width at half-maximum (FWHM). A perfect crystal at zero temperature would produce a close approximation to a delta function (limited by the X-ray apparatus), but a real crystal's rocking curve width is limited by spatial variations in the alignment of the atoms. Such variations lead to a distribution in the

reciprocal lattice vectors' direction, which leads to a finite width in the reflection.

What constitutes 'good' alignment (or a 'good' rocking curve width) is dependent on the system, crystal size and growth technique, and is most useful as a benchmark in comparisons among similar materials grown by similar techniques. For instance, $YBa_2Cu_3O_{6+\delta}$ grown by the top-seeded melt-growth technique, where crystals are cubic-centimetre-sized and impurities are deliberately added on the 2-30% level, may have rocking curves several degrees wide [94], $YBa_2Cu_3O_{6+\delta}$ grown by a self-flux technique in yttria-stabilized zirconia (YSZ) crucibles may be a couple hundredths of a degree and $YBa_2Cu_3O_{6+\delta}$ grown in $BaZrO_3$ crucibles may be a few thousandths of a degree [41]. Rocking curves for high-purity single-crystalline silicon, with only one constituent atom instead of four, are narrower still.

Figure 2.10: X-ray rocking curve for a square-millimetre-sized $T_c = 75$ K crystal, $(0\ 0\ \overline{10})$ peak. The fit is to a squared Lorentzian, and the full width at half-maximum is $0.0405°$.

X-ray rocking curves of several crystals were collected on a Philips X'Pert Pro X-ray diffractometer, using a copper $K\alpha_1$ vertical line source excited by 45 kV and 40 mA. The X-rays diffracted off the crystal were detected using a serial detector with no entrance slit. Each crystal's position and orientation were optimized in x, y, z, ϕ, and ψ, and 2θ was also optimized. Crystals were mounted to glass slides using

vacuum grease, to avoid spurious signals from the mount. Each ~ 1 mm^2 crystal was fully illuminated by the line source, allowing all parts of the crystal to contribute to the rocking curve. The $(0\ 0\ \overline{10})$ rocking curve for a $T_c = 75$ K crystal is shown in Figure 2.10 — its width (FWHM) is 0.0405°. Measured widths varied from about 0.025° to 0.050°, comparable to YSZ-grown YBCO.

Widths were extracted from nonlinear least squares fits to the data, using a squared Lorentzian as a fitting function:

$$N(\omega) = \left(\frac{A}{\pi} \frac{\left(\frac{\Gamma}{2\sqrt{\sqrt{2}-1}} \right)^2}{(\omega - \omega_0)^2 + \left(\frac{\Gamma}{2\sqrt{\sqrt{2}-1}} \right)^2} \right)^2 \tag{2.2}$$

where $N(\omega)$ represents the count rate as a function of angle ω, A is the amplitude of the function, Γ is the FWHM, and ω_0 is the angle corresponding to the centre of the peak. A squared Lorentzian is used because the distribution of dislocations leading to broadening is expected to produce a Lorentzian distribution in each of the three directions, and the use of a line X-ray source effectively integrates over one dimension, that parallel to the line. In rocking curves, count rates are by convention and for clarity graphed against $\Delta\omega \equiv (\omega - \omega_0)$, rather than ω, which is dependent upon the reflection being investigated.

The rocking curve widths measured indicate that the crystals are high-quality, comparable with good YSZ-grown YBa$_2$Cu$_3$O$_{6+\delta}$, but not as good as BaZrO$_3$-grown YBCO. Given that YBa$_2$Cu$_3$O$_{6+\delta}$ is an extremely clean system and free from cation cross-substitution, falling short of the BaZrO$_3$-grown gold standard should be neither alarming nor surprising. The widths measured are considered very good nonetheless — the best HTSC crystals ever grown may not be a completely fair benchmark. Unfortunately, the only Tl-2201 rocking curve width which has previously been reported is from an earlier growth method by the author [16]. No significant change is observed.

2.5.3 Electron-Probe MicroAnalysis

To determine the levels of cation substitution in our crystals, electron-probe micro-analyses (EPMA) of several crystals were performed. In EPMA, a small region of the sample (in this case, ten microns in diameter) is illuminated with an electron beam, and the characteristic X-rays produced are collected. Similar materials are employed

as standards and the software takes account of the sample's self-absorption, so the data produced are weight percents of the atoms believed to be present, typically accurate to a fraction of a percent. Extracting the exact chemical formula only requires knowledge of the atomic weights of the atoms in the compound.

In Tl-2201, previous experience has found oxygen results to be unreliable, so the cations are considered alone. Usually, our normalization technique would assume that barium was exactly stoichiometric (i.e. $Tl_xBa_2Cu_yO_z$), as nothing is known to substitute for it in this system. However, in this case, the barium results were found to be highly variable (at the 2.5% level), likely due to topography of the surface, while the thallium and copper contents were found to sum consistently to about three ($x + y = 3$). No vacancies on the thallium site have ever been observed, and no other elements are expected to be present on that site in any significant quantity, so the normalization was to $(Tl + Cu) = 3$.

Crystals to be measured by EPMA were epoxied to the face of a plexiglas disc, with efforts made to keep their faces parallel to that of the disc to the extent possible. It is important to note that this is unusual — normally, crystals would be encased in epoxy and polished to expose a flat surface, smooth on the submicron scale. In this case, because the crystals are thin, flat and generate hazardous waste when polished, it was decided to proceed without polishing. The software assumes an angle of exactly $40°$ between the collected X-rays and the normal to the crystal surface, and surface topography can cause absorption of X-rays in unpredictable ways, making the data less reliable. To account for this to the extent possible, EPMA data were collected on regions of the sample that were especially flat and uniform. An additional potential effect of this mounting scheme is poorer grounding. A thin layer of graphite is evaporated onto the surface of the disc to allow the charge deposited on the disc by the electron beam to be dissipated. Areas that are improperly coated can charge up, broadening and deflecting the electron beam. With the crystals protruding rather than flush with the surface, there exists the possibility that the crystal's top surface could shield areas around it from the graphite. However, a check for charging effects using an attached scanning electron microscope failed to observe any.

The EPMA measurements were performed on a fully-automated CAMECA SX-50 instrument in wavelength-dispersion mode, using an excitation voltage of 15 kV, a beam current of 20 nA, a peak count time of 80 s, and a background count time of 40 s. Initial data reduction was performed by the instrument's software using the "PAP" $\varphi(\rho Z)$ method [95].

To measure the thallium content, thallium metal was used as a standard, and the

Tl$M\alpha$ line was detected using a PET (pentaerythritol) crystal as a monochromator. Because elemental thallium oxidizes very rapidly in air, this reference sample was sheared off flush with its plexiglas mount, then immediately coated with graphite and transferred into the instrument's vacuum space. To correct for oxide on the surface, the thallium metal was calibrated assuming it to be pure thallium, then analyzed for thallium and oxygen. This technique detected roughly 3% oxygen by mass. This composition of slightly oxidized thallium was then entered as revised calibration data and used as the standard for all later measurements.

For barium and copper, $YBa_2Cu_3O_{6.920}$ served as the standard. Barium employed the Ba$L\alpha$ line and a PET monochromator, while copper used Cu$K\alpha$ and LiF. Oxygen was also measured, using the same YBCO crystal as a standard, the oxygen K line and the synthetic crystal PC1 as a monochromator, but the resulting oxygen contents were unphysically high and wildly varying, and are not reported. Tight Pulse Height Analysis (PHA) control was used to eliminate to the degree possible any interference from higher-order lines.

Figure 2.11: EPMA results on a Tl-2201 crystal grown from a melt of cation ratio Tl:Ba:Cu = 2.2:2:1.8, normalized to Tl + Cu = 3. The red, blue and green points are the thallium, barium and copper contents respectively measured at various points on the crystal, while the lines represent the averages: $Tl_{1.920(2)}Ba_{1.96(2)}Cu_{1.080(2)}O_{6+\delta}$.
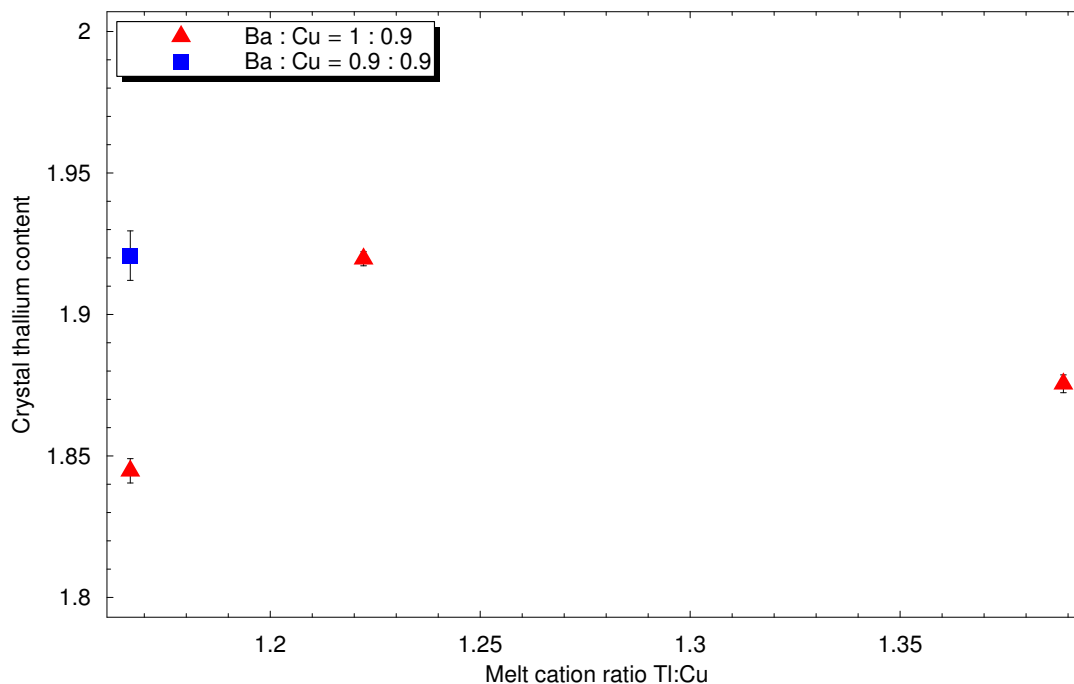
Figure 2.12: EPMA composition results vs. cation ratio in melt used for growth. The thallium-to-copper ratio in the mixed powders constitutes the x-axis, EPMA compositions are shown vertically, and an additional point is shown for a mixture with decreased barium concentration.

The cation composition (normalized as described above) was determined to be $Tl_{1.920(2)}Ba_{1.96(2)}Cu_{1.080(2)}O_{6+\delta}$ for a crystal grown as described in Section 2.2 from the initial composition Tl:Ba:Cu = 2.2:2:1.8. Figure 2.11 shows how the thallium, barium and copper contents varied from point to point on the surface of the crystal — the cation substitution level is not position dependent. The composition measured corresponds to a level of copper substitution at the thallium site slightly lower than that reported previously [25, 48, 79, 80] for the cation-substituted phase of Tl-2201. A similar study on the author's earlier-generation crystals also found a higher cation substitution ratio, $Tl_{1.884(6)}Ba_2Cu_{1.11(1)}O_{6+\delta}$ (normalized to barium) [16, 87].

The high (and unphysical) variability in barium results was seen on all Tl-2201 crystals studied; as mentioned above, this is attributed to effects of the non-standard topography. Similar variations were seen when the locations probed were shifted slightly, but were not observed on a barium sulphate standard.

Crystals were tested that were grown from several different melt compositions, and the results shown in Figure 2.12. Three of the batches had varying thallium

contents — Tl:Ba:Cu ratios of 2.1:2:1.8, 2.2:2:1.8 and 2.5:2:1.8 (although the latter batch leaked slightly). The fourth had a reduced barium content, 2.1:1.8:1.8.

Aluminum contamination was checked separately and found to be below the 50 ppm detection limit (by mass), corroborating to the extent possible our observation that the crucibles were not corroded by the melt. To test for this aluminum contamination, the beam current was increased to 40 nA, the peak count time to 500 s and the background count time to 250 s. The matrix composition, for absorption calculations, was fixed at that measured for the crystal studied. The $AlK\alpha$ was detected using a TAP (thallium acid phthalate) crystal monochromator, and standardized on α-Al_2O_3. Due to interference from a higher-order barium line, this technique involved measuring the apparent aluminum content of our crystal, followed by the apparent aluminum content of a barium sulphate crystal, which should have contained no aluminum. The aluminum ghost signal from the barium sulphate exceeded that from our crystal, indicating the absence of aluminum. A similar check for gold contamination was not feasible due to a much more serious interfering thallium line.

The error bars used and uncertainties quoted in this section are 2σ, and beyond the data in Figure 2.11, are based on scatter in the data alone. An overall shift would not be captured by this, and any issue with the standards could cause additional shifts. Finally, the T_c measured on multiple crystals from the same batch, annealed together, may vary by as much as a Kelvin, so crystals from the same batch may have slightly different cation ratios. Despite the non-standard surface topography, given the extremely high precision and uniformity found in the cation substitution ratios for each crystal studied, it is possible to have a high degree of confidence in the results obtained.

2.5.4   Single Crystal X-ray Diffraction

X-ray diffraction is an established and sensitive technique for determining the crystal structure of a material and the positions of the atoms within the unit cell. X-rays diffract off the material's reciprocal lattice vectors in a manner that depends on the space group; the details depend sensitively on how electron density is distributed within the unit cell. Given basic information about the material and a set of reflections covering a significant region in reciprocal space, standard software packages can quickly and easily compute the lattice parameters and atomic positions.

Several X-ray and neutron diffraction studies have been performed previously on $Tl_2Ba_2CuO_{6\pm\delta}$, finding both orthorhombic and tetragonal symmetries [66, 96], and

indicating the existence of oxygen interstitials [89], oxygen vacancies in the $Tl_2O_2^{2+}$ layer near optimal doping [47], and the substitution of copper onto the thallium site [79, 80, 97]. Almost all of this work was on powder or ceramic samples, which may allow compositions not possible in crystals and may be susceptible to incomplete reaction, surface decomposition in air, or inhomogeneity.

Single crystal X-ray diffraction has been previously performed on Tl-2201 [79, 80], but on first-generation crystals and only after the structure had already been established on powder samples. Both studies found tetragonal symmetry, and were refined in the body-centred space group I4/mmm; Liu also tried several reduced-symmetry space groups, including I422, I4mm, I$\bar{4}$m2 and I$\bar{4}$2m [79]. There is no evidence that Kolesnikov tried other space groups once the crystal was found to be tetragonal. Kolesnikov's group reported the T_c of the crystal they studied to be 110 K, 16 K higher than any other group has ever found for this material, likely near optimal doping. Liu did not report a T_c for their crystal. There is no indication that either group's crystal was annealed, and O(4) oxygen interstitials were only considered by Liu, who found no evidence for them. The O(4) site, sparsely populated at any doping, is expected to have very low, if any, occupancy near optimal doping, so it should not be surprising that no occupancy of the site was observed.

Measurements of the lattice parameters of the author's earlier crystals showed evidence of orthorhombic symmetry [16]. This symmetry has been previously observed in powders but never in other crystals; it is possible, however, that other groups prepared orthorhombic crystals and simply didn't look at high enough dopings to observe the orthorhombicity.

The most thorough investigation on the structural effects of doping and doping-related defects was a powder study by J. L. Wagner [47]. Through a detailed examination of the interplay between oxygen interstitials, oxygen vacancies, T_c, cation substitution, and orthorhombicity, this paper suggests a method of estimating the cation and oxygen doping levels if T_c and the lattice parameters are known. If Tl-2201's lattice constants correlate with the oxygen doping and the level of cation substitution, and both defects may be expected to dope the crystal, their interplay is of interest. In particular, the ability to determine a crystal's cation substitution level from comparatively easily measured quantities such as the lattice parameters and T_c would prove useful.

Determining the lattice constants does not require a full set of X-ray diffraction data — from one or two peaks along each crystallographic direction, a good approximation to the lattice parameters can be arrived at. Correlating this with T_c, EPMA

results and possibly titrations for the oxygen content would complete the picture. However, early attempts at finding the lattice parameters turned up weak peaks that should have been forbidden by symmetry, suggesting that previous investigations had failed to fully capture the crystal structure. A full X-ray structure refinement was performed on crystals at several dopings, to attempt to elucidate Tl-2201's crystal structure.

X-ray diffraction data were collected on a Bruker X8 Apex diffractometer operated by the UBC Chemistry Department's Structural Chemistry Facility. Peaks were captured on a CCD area detector using a molybdenum X-ray source (Mo $K\alpha$, 0.71073Å). The crystal was measured while cooled to -100°C in flowing nitrogen gas, to improve data quality. It must be noted that the unit cell may be expected to contract slightly on cooling, so these data cannot be compared directly against room temperature lattice parameters without caveats.

Tl-2201 single crystals with T_cs of 55 K and 75 K were measured by X-ray diffraction. The orthorhombic lattice parameters (in space group Fmmm) for the $T_c = 75$ K crystal were $a = 5.4477(9)$ Å, $b = 5.4484(9)$ Å, and $c = 23.1711(35)$ Å, while the $T_c = 55$ K crystal's lattice parameters were $a = 5.4424(12)$ Å, $b = 5.4550(12)$ Å, and $c = 23.1428(51)$ Å (all uncertainties in this section are quoted as 1σ). As can be seen from these values, the former was tetragonal within the instrumental resolution, while the latter's a and b lattice parameters differed by 10.5σ, indicating it to be orthorhombic. The figure of merit used by Wagner [47] to compare the strength of the orthorhombic distortion among different samples is the orthorhombic strain

$$\eta = \frac{2(b-a)}{(b+a)} \tag{2.3}$$

The orthorhombic strain for the $T_c = 55$ K crystal is 2.31‰.

The doping dependence of the c-axis lattice parameter is shown in Figure 2.13, compared against published data from Wagner [47], Shimakawa [66] and the author's earlier crystals [16]. On the latter crystals, only the lattice parameters were measured, and using a different instrument. Wagner's data are indexed by the average thallium contents of the batches. The general trend of a decrease in the c-axis with overdoping (or the addition of oxygen interstitials) holds, but the present data do not agree quantitatively with earlier numbers. All earlier data were collected at room temperature, while these data were collected at -100°C, so the lattice parameters might be expected to be shorter, as is observed.

The orthorhombic strain's doping dependence is shown in Figure 2.14. The trends

Figure 2.13: Doping dependence of the *c*-axis lattice parameter. The present data are compared against published data from Wagner [47], Shimakawa [66] and two of the author's earlier crystals [16]. The present data were collected at -100°C while all previous data were collected at room temperature, leading to a relative shift.

observed by Wagner, that the orthorhombicity is suppressed as optimal doping is approached and as cation substitution is increased, are reproduced. Our crystals' orthorhombicities most closely track those for a set of Wagner's samples that had an average thallium content $Tl_{1.881}$. A lower cation substitution level would have been expected given the EPMA results reported in Section 2.5.3, but this technique for determining cation substitution is highly inexact, and there is known to be some variation from crystal to crystal even within a single growth run, so this result is not alarming.

The reduced temperature of our measurement could conceivably affect the orthorhombicity results, but if the orthorhombic distortion is due to static distortions of the lattice, vibrations are close to isotropic in the *ab*-plane, and no atoms have significant mobility, the impact of temperature should be minimal. However, it should be pointed out that there is not yet a detailed understanding of the mechanism behind this distortion.
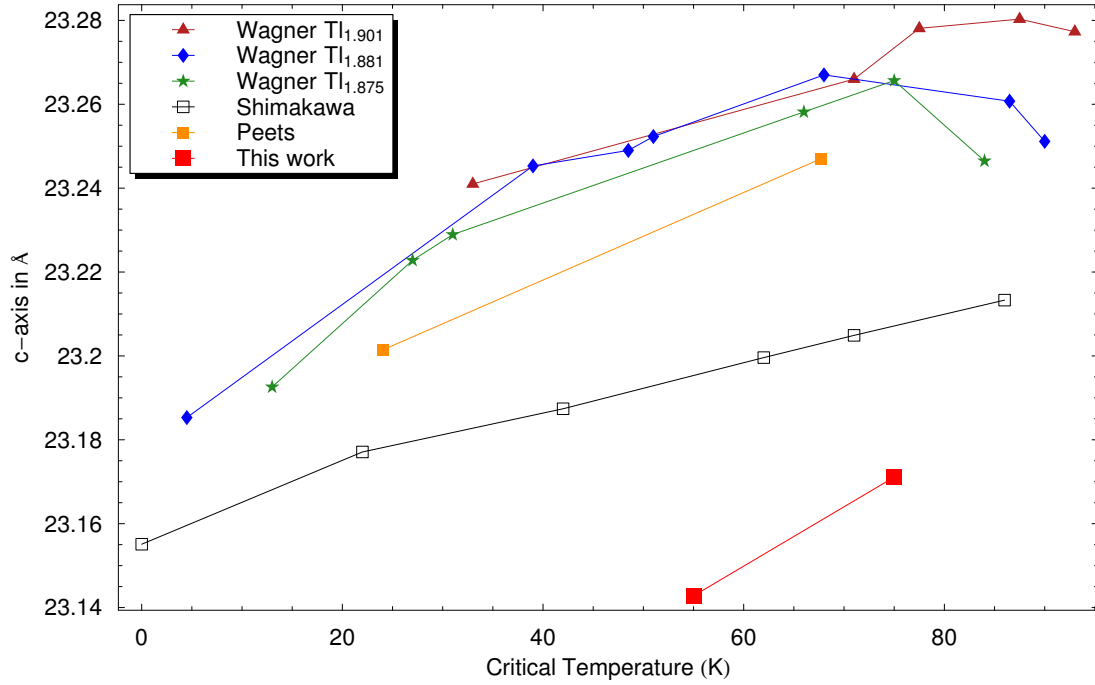
Figure 2.14: Doping dependence of the orthorhombic distortion. The present data are compared against published data from Wagner [47], Shimakawa [66] and two of the author's earlier crystals [16]. Differences in temperature between the datasets should have little effect on this plot.

Gauging the cation substitution ratio by the relationship between the lattice parameters and T_c may well work, but it would need to be better calibrated, by measuring both EPMA and X-ray diffraction on several crystals. The temperature-dependence of the lattice parameters would also need to be characterized. For now, it serves as a useful tool for approximating the substitution ratio, but it may not be useful for distinguishing the slightly different defect levels among our crystals.

While the crystals studied (and those studied earlier in this work and reported in reference [16]) behaved as expected given Wagner's results, obeying the qualitative relationship between the tetragonal and orthorhombic phases depicted in Figure 1.11, it must be emphasized that no crystals of Tl-2201 had ever previously been found to be orthorhombic, although possibly due to X-ray diffraction having been reported on only two crystals previously, both unannealed and likely near optimal T_c. Wagner's results are not well known, and some in the field still believe that the orthorhombic phase exists only in powders or ceramics and is non-superconducting. In fact, the orthorhombic and tetragonal phases are almost universally identified with stoi-

chiometric and defected Tl-2201 respectively. Given the results outlined here, these are clearly no longer tenable positions. These results, that crystals of orthorhombic Tl-2201 may be grown, that the orthorhombic phase superconducts, and that the orthorhombic phase need not be stoichiometric, eliminate any remaining doubt from the material's symmetry phase diagram, leaving only a few specific details.

A structure refinement was performed on the T_c = 75 K crystal, using SHELXL-97 [98, 99], and assuming orthorhombic (Fmmm) symmetry for generality. A total of 3319 reflections were analyzed, 312 of them unique, in the k-space region $-7 \leq h \leq 7$, $-7 \leq k \leq 7$, $-31 \leq l \leq 32$. All atoms were constrained to their symmetric positions (relaxing this constraint led to only minor improvements), and occupancies of all sites except O(4) were fixed at full occupancy; atoms other than O(4) were allowed anisotropic thermal parameters. Besides refining for occupancy of the O(4) site, a refinement was performed for the thallium-copper substitution level. The R_1 factor, indicating the quality of the refinement, was 2.35% — values below 2.5 $\sim$ 3% are considered very good; wR_2, a similar factor more closely analogous to the χ^2 of least squares techniques, was 5.68% on all data.

Table 2.4: Refined atomic co-ordinates for the T_c = 75 K crystal, along with an equivalent isotropic thermal displacement parameter. The latter, U_{eq}, is defined as $\frac{1}{3}$ of the trace of the orthogonalized U_{ij} thermal displacement tensor (reported in Table 2.5). Also see Figure 2.15 for the oxygen sites' identities.

Atom	x/a	y/b	z/c	U_{eq}
Tl	0.0000	0.0000	0.29732(2)	0.0108(2)
Ba	0.0000	0.0000	0.08306(3)	0.0028(3)
Cu	0.0000	0.0000	0.5000	0.0023(4)
Cu_{Tl}	0.0000	0.0000	0.29732(2)	0.0108(2)
O(1)	0.2500	0.2500	0.0000	0.0043(13)
O(2)	0.0000	0.0000	0.3833(4)	0.0087(16)
O(3)	0.0000	0.0000	0.2112(5)	0.055(6)
O(4)	0.2500	0.2500	0.283(14)	0.14(15)

The refined atomic positions are presented in Table 2.4. Cu_{Tl} refers to copper atoms on the thallium site. With the total occupancy of this site fixed (full), its composition was refined to $Tl_{1.914(14)}Cu_{0.086(14)}$, in excellent agreement with the result found by EPMA in Section 2.5.3 of $Tl_{1.920(2)}Cu_{0.080(2)}$. The occupancy of the O(4) site was also refined, but the low sensitivity of this technique to light atoms and the low

occupancy of the site led to a value which, while consistent with the low occupancy expected, is also consistent with zero to within 1.5σ — $O_{0.07(5)}$. To produce a T_c of 75 K, a very low occupancy of this site is expected.

Table 2.5: Anisotropic displacement parameters for the atomic positions in Table 2.4 which were refined assuming an anisotropic electron density. The U_{ij}s effectively describe an ellipsoid occupied by the electron density.

Atom	U_{11}	U_{22}	U_{33}	U_{23}	U_{13}	U_{12}
Tl	0.0130(3)	0.0190(4)	0.0006(3)	0	0	0
Ba	0.0000(4)	0.0052(4)	0.0031(4)	0	0	0
Cu	0.0000(8)	0.0030(8)	0.0040(9)	0	0	0
Cu_{Tl}	0.0130(3)	0.0190(4)	0.0006(3)	0	0	0
O(1)	0.000(3)	0.006(3)	0.007(3)	0	0	0.004(2)
O(2)	0.006(4)	0.017(4)	0.003(3)	0	0	0
O(3)	0.086(15)	0.080(14)	0.000(4)	0	0	0

The thermal parameters U_{ij} extracted for these atomic positions are reproduced in Table 2.5. This tensor represents the spread in electron density or atomic position due to temperature or disorder. The anisotropic displacement factor takes the form

$$T = e^{-2\pi^2\left[h^2 a^{*2} U_{11} + k^2 b^{*2} U_{22} + l^2 c^{*2} U_{33} + klb^* c^* U_{23} + hla^* c^* U_{13} + hka^* b^* U_{12}\right]} \tag{2.4}$$

where a^*, b^* and c^* are the reciprocal lattice parameters. The resulting orthorhombic crystal structure is shown in Figure 2.15. Note that the conventional tetragonal axes are at $45°$ to those shown.

The thallium layers' and O(2) (barium oxide layer — see Figure 2.15 for the oxygen sites' identities) site's ellipsoids are deformed, particularly the former, in the ab-plane. This has been observed before [47, 48, 80, 100], and dealt with by splitting these sites for a better refinement. The disc-like nature of our anisotropic displacement parameters indeed suggests that this layer's atoms are displaced from their ideal positions, but it is not clear that splitting the site and allowing the atoms to be displaced would add any information. More is happening here than splitting the site can capture; this issue will be revisited later in this section.

The O(4) site, refined with isotropic thermal parameters, corresponds to an extremely large ellipsoid, indicative of the very low occupancy of that site — the electron density due to an extra oxygen anion (10 electrons) every tenth or hundredth unit cell is difficult to distinguish in a $Tl_2O_2^{2+}$ layer (176 electrons). The large apparent

Figure 2.15: Refined crystal structure with 98% probability ellipsoids for each atom's nuclear position. The plane oxygen has been re-refined isotropically to eliminate an issue with the U_{12} parameter, and the O(4) interstitial has been excluded for clarity.

size of this atom suggests that its already-low occupancy might be inflated, padded by electron density from other nearby atoms. Furthermore, the O(4) site refined to a location different from where it has been found before. Having the refinement wander off suggests that the occupancy is very low indeed. This refinement does not provide strong evidence for occupation of the O(4) site for $T_c = 75$ K crystals.

It should also be noted at this juncture that the Cu and O(1) atoms are quite isotropic and have small displacement parameters despite being constrained to symmetric locations within the unit cell. This confirms that the CuO_2 planes are extremely flat. While this is unsurprising, it is useful to confirm for the purposes of electronic structure, as any deviations would require admixtures of orbitals. The

analysis of X-ray spectroscopic data, for instance, becomes less straightforward when some mixing of symmetries is present in every band.

The X-ray diffraction results are compared in Table 2.6 against the closest comparators of Wagner's samples. Aside from the drift of O(4) mentioned earlier and a 0.4% contraction in the axes, most likely associated with the measurement temperature, the agreement is excellent, particularly with Wagner's "E2" sample. Note, however, that Wagner left more free parameters — for instance, the O(3) and Tl sites were permitted to depart from their symmetric positions.

Table 2.6: Results of the structure refinement for the $T_c = 75$ K crystal, compared against Wagner's results [47]. Wagner's samples E2 and E3 are used for comparison — the E samples had an average thallium content of $Tl_{1.881}$, and E2 and E3 had T_cs of 86.5 K and 68 K respectively. (Note: directly comparing R and wR values between powder and crystal data is not reasonable, but all similarly indicate how well the respective models reproduce the data.)

	$T_c = 75$ K	Wagner E2	Wagner E3
T_c (K)	75.0	86.5	68.0
a (Å)	5.4477(9)	5.47031(22)	5.46662(9)
b (Å)	5.4484(9)	5.47242(22)	5.48238(9)
c (Å)	23.1711(35)	23.26073(33)	23.2670(4)
$\eta = 2(b-a)/(b+a)$	0.00013	0.000386	0.00288
Tl z/c	0.29732(2)	0.29734(5)	0.29728(5)
Tl n	1.914(14)	1.912(8)	1.9864(8)
Ba z/c	0.08306(3)	0.08329(8)	0.8370(8)
O(2) z/c	0.3833(4)	0.38313(7)	0.38331(7)
O(3) z/c	0.2112(5)	0.21115(10)	0.21078(11)
O(4) z/c	0.283(14)	0.277(3)	0.274(3)
O(4) n	0.07(5)	0.026(6)	0.028(3)
R_1 (R_p)	2.35%	3.34%	3.67%
wR_2 (R_{wp})	5.68%	4.65%	5.24%

As mentioned above, an earlier investigation into our crystals' lattice parameters turned up weak peaks that had not previously been observed. The above structure refinement found anomalies in the Tl, O(3) and O(2) sites, but refined very well. However, a closer inspection of the data indicates that the above structure refinement is incomplete.

The $(hk0)$ plane for the $T_c = 75$ K crystal is shown in Figure 2.16, a synthesized precession photograph. The brightest spots are those expected and those most influ-

ential in the structure refinement, while the weaker ones correspond to a superlattice modulation that doubles the in-plane lattice constants. As can be seen from the $(h0l)$ plane (Figure 2.17), these more closely resemble rods than points, although they do exhibit some point-like character. This indicates that the superlattice modulation is largely, but not entirely, uncorrelated along the c-axis. Given that the Tl, O(2) and O(3) sites were especially anisotropic, the superlattice modulation most likely involves these sites. It should be noted that the temperature at which these data were taken, -100°C or 173 K, make it unlikely that this would originate from charge order in the CuO_2 planes.

Figure 2.16: X-ray diffraction pattern for a $T_c = 75$ K Tl-2201 crystal, $(hk0)$ plane. The weak spots are superlattice peaks which have not been previously observed.

One important characteristic of the superlattice modulation is that it is commensurate — the extra features appearing are exactly consistent with a doubling of the unit cell, and are not split. An incommensurate modulation would make the atomic sites in the CuO_2 plane inequivalent, which be extremely disruptive for single-particle spectroscopies, particularly STM/STS. Incommensurate superlattice

Figure 2.17: X-ray diffraction pattern for a $T_{\mathrm{c}} = 75$ K Tl-2201 crystal, $(h0l)$ plane. The weak superlattice peaks now appear as vertical rods.

modulations in Tl-2201 have been previously suggested based on weak peaks in powder data [66, 97, 101], but this is inconsistent with our data. The earlier data were chiefly collected on more heavily overdoped samples, and Reference [66] found the incommensurate modulation's peaks to weaken substantially as the doping was reduced, so there remains the possibility that the commensurate superlattice modulation observed here may become incommensurate on further overdoping. A full doping study has not yet been completed.

A full investigation of the diffuse scattering is beyond the scope of this thesis, and firm conclusions as to which layer exhibits superlattice structure would require resonant X-ray scattering or neutron scattering, so the exact nature of the superlattice modulation remains unresolved. This would make an interesting topic for further research by collaborators in neutron, diffuse X-ray or resonant X-ray diffraction. In fact, determining the nature of this structure may prove quite important — to minimize scattering, it may become necessary to gain some degree of control over the superlattice.

In summary, high-quality platelet single crystals of Tl-2201 have been grown using a novel variable-encapsulation scheme with the ability to decant, and optimum growth parameters have been determined. Loss of thallium has been minimized, allowing highly homogeneous crystals with a low level of cation substitution. Annealing schemes have been developed which allow the crystal's oxygen content to be set and homogenized without damage to the crystal, leading to homogenous T_cs (and narrow superconducting transitions) from 5 to 85 K. The chemical formula as determined by EPMA for the growth technique settled upon is $Tl_{1.920(2)}Ba_{1.96(2)}Cu_{1.080(2)}O_{6+\delta}$, although the barium content is believed to be exactly two. The crystals are high-quality and homogeneous as determined by SQUID magnetometry and X-ray rocking curves, being on par with $YBa_2Cu_3O_{6+\delta}$ grown in YSZ crucibles.

The X-ray diffraction study found a crystal structure and lattice parameter behaviour very much like those published by Wagner [47] for cation defect levels similar to those determined on our crystals by EPMA. Orthorhombicity matching Wagner's results is also observed. Our crystals become orthorhombic on overdoping, ceasing to be tetragonal at a doping leading to a T_c between 55 K and 75 K. The results that crystals may be orthorhombic, the cation-substituted phase may be orthorhombic, and the orthorhombic phase superconducts all defy conventional wisdom, although that may be more indicative of a shortcoming of conventional wisdom than of any new result.

The full X-ray structure refinement for the 75 K crystal found the thallium site's cation composition to be $Tl_{1.914(14)}Cu_{0.86(14)}$ and an O(4) site occupancy of 0.07(5), leading to a final composition of $Tl_{1.914(14)}Ba_2Cu_{1.086(14)}O_{6.07(5)}$, in excellent agreement with the EPMA result, which is less sensitive to oxygen.

However, there was evidence for a commensurate superlattice modulation, most likely in the blocking layers, a significant departure from the existing literature. X-ray diffraction was not able to determine the exact nature of this modulation, so other techniques such as neutron diffraction or resonant X-ray diffraction will need to be applied.

3. X-Ray Absorption and Emission Spectroscopies

This chapter concentrates on the concepts behind and details of X-ray absorption and emission spectroscopies, data reduction and analysis. Because the apparatus was designed and constructed by groups at other universities over a decade ago and is now operated, upgraded and maintained as a user facility, technical details will be reported only where relevent. It should be noted that the excluded details are highly non-trivial, and many technical issues did affect the design fundamentally — for instance, soft X-rays have a short penetration length in virtually all materials, including air [102], so not only must the entire experiment and beamline be operated under ultra-high vacuum (UHV), but the beamline and its endstation must be open to the synchrotron storage ring and all other active beamlines. This requires a series of interlocks such that an elevated pressure in any part of the system immediately isolates it from other parts of the system, and scans are interrupted periodically as a result of such events.

3.1 Introduction to X-ray Absorption and Emission Spectroscopies

X-ray absorption spectroscopy (XAS) and X-ray emission spectroscopy (XES) are powerful techniques for investigating materials' electronic structure at the eV scale. Detailed information about the bands involved in bonding may be obtained, including the symmetry of the orbitals from which they're constructed or their evolution with doping or temperature. In a correlated electron system, features such as the upper and lower Hubbard bands or the cuprates' Zhang-Rice singlet states may be observed.

A schematic representation of these techniques is shown in Figure 3.1; this will be used to present a schematic, conceptual introduction to the techniques. This approximation should be valid in the absence of strong correlation effects, which makes it inherently inapplicable in the cuprates. Nonetheless, it's useful to see how the technique works before complication is added.

ε_F

PSfrag replacements

Figure 3.1: Schematic depiction of the processes behind X-Ray Absorption Spectroscopy and X-Ray Emission Spectroscopy (in the absence of correlations). An incident photon is absorbed, exciting a core electron into unoccupied states. An electron in an occupied state falls into the core hole, emitting an X-ray photon.

In the case of either XAS or XES, monochromatic, polarized sychrotron X-rays are shone on the sample. Their energy is chosen to be near an atomic absorption edge, so electons may be excited from core levels into the unoccupied states. By changing the parameters of the insertion device used to generate the X-rays in the synchrotron and of the monochromator, the X-rays' incident energy may be scanned.

When the incident energy is in resonance with a transition to an unoccupied state and selection rules permit, there is a collosal enhancement in absorption. This may be detected either via the core electrons liberated (total electron yield, or TEY) or the X-rays emitted as other electrons fall down to fill holes in the core level (total fluorescence yield or TFY). An XAS spectrum, then, shows at what energies unoccupied states exist and, through the intensity, the density of states at that energy.

Selection rules can be used to extract further information. For an incident x-ray photon to be absorbed by an electron, that electron must make a transition in which its parity changes along the photon's polarization direction (it is further constrained to changes in quantum number of $\Delta\ell = \pm 1$ and $\Delta m_\ell = 0, \pm 1$). If we assume that the core electron being excited is localized and thus not involved in bonding, this electron starts out in a spherically symmetric state — in a $p \to d$ transition, there is no distinction between p_x and p_z orbitals, for instance, as they are degenerate and exist in superpositions. Under these selection rules, for the polarization $\vec{E} \parallel \hat{c}$, an $s \to p$ transition can only occur if there's a vacant p_z orbital. Similarly, a vacant $d_{3z^2-r^2}$ orbital can only receive electrons excited from a p_z orbital. By probing the

sample with X-rays parallel to each crystal axis in turn, it is possible to determine the symmetry of the orbitals from which each band is formed.

In XES, the x-rays emitted through the fluorescence process are collected and their energies analyzed. The energy spectrum of these X-rays provides much the same information as an XAS spectrum, but for occupied states. In some cases, XES may also provide more information about the unoccupied state being probed — for instance, if the state into which the core electron was excited is stable and localized, it will have a long lifetime and will appear in the emission spectrum, but only when excited at that specific energy.

In the presence of strong correlation effects, these techniques probe states of the system, rather than electronic states. The liberation of a core electron would not normally be a substantial perturbation on the system, as the core levels are tightly-bound, localized and full, and thus largely uninvolved in bonding or the correlations. However, this electron is not removed from the system — it is added to what would be, in the absence of correlations, an unoccupied state near the Fermi level. This makes XAS an electron-addition spectroscopy. Similarly, when the state relaxes, an electron is removed from what would be an occupied valence state, and deposited in the non-interacting core state, making X-ray emission an electron-removal spectroscopy.

XAS studies of several cuprates have been performed previously, but with lower resolution and signal-to-noise, and on only a few compounds at a few dopings. It has been applied to LSCO [103, 104, 105], YBCO [105, 106, 107, 108] and near-optimally doped multilayer Tl-based cuprates [109, 110], along with a handful of other cuprates, but not to Tl-2201, nor to ordered phases of YBCO. The sum total of doping-dependent studies of the overdoped regime previously performed on any compound is reproduced in Figure 3.2.

3.2   Apparatus and Data Collection

X-ray emission and absorption spectroscopies were performed at beamline 8.0.1 of the Advanced Light Source (ALS) at Lawrence Berkeley National Laboratory (LBNL) in Berkeley, California [111]. This beamline's X-rays are produced by bunches of 1.9 GeV electrons passing through an undulator whose gap (and thus the peak energies in its spectrum) may be controlled remotely by computer, using a LabVIEW-based interface. The undulator's third harmonic was the most intense accessible in the energy range of interest, and was used for all work reported here.

Amongst other steps, the X-rays are diffracted off a monochromator grating to
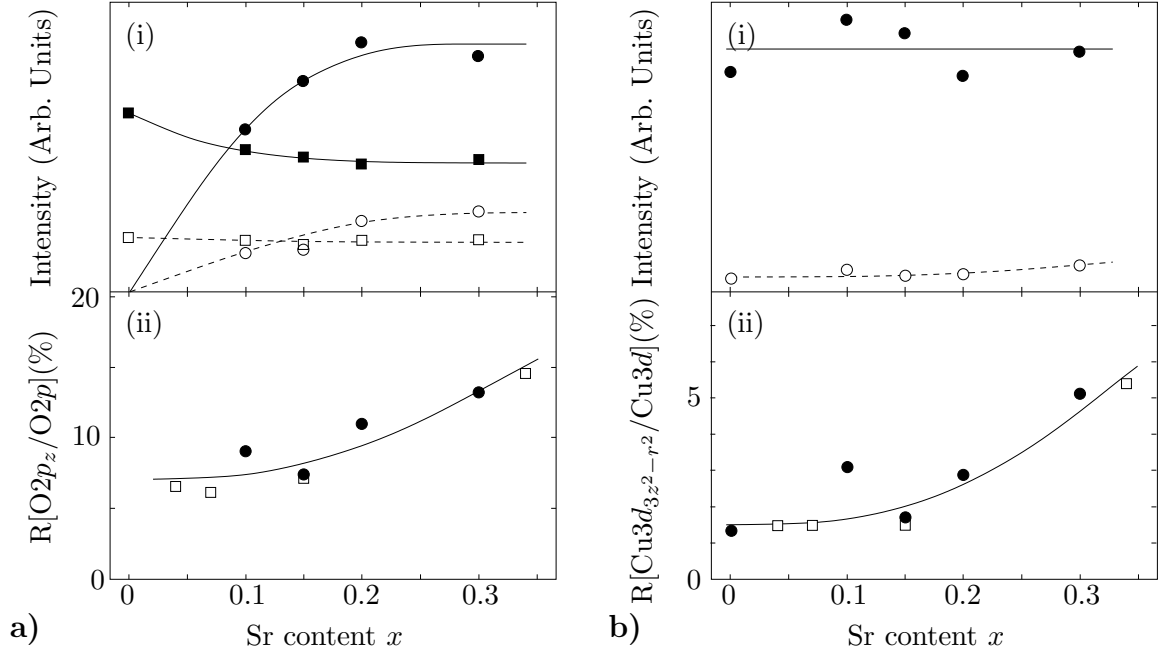
Figure 3.2: Prior to the work reported in this thesis, this constituted all doping-dependent XAS data taken on the overdoped side of the phase diagram (optimal doping is around $x=0.18$, but $x=0.2$ is still quite close to optimal). Adapted from [103]; the material is LSCO, x is the Sr content, and intensity refers to integrated spectral weight (in some cases, this required subtraction of an $x = 0$ background). **a)** Oxygen K edge: i) squares are the upper Hubbard band, circles are the Zhang-Rice singlet or valence band, open symbols are $2p_z$ symmetry, closed symbols are $2p_{x,y}$ symmetry. ii) Fraction of $2p_z$ character — open symbols are taken from [104]. **b)** Copper $L\alpha$ edge: i) closed symbols are $3d_{x^2-y^2}$ while open symbols are $3d_{3z^2-r^2}$ symmetry. ii) Fraction of $3d_{3z^2-r^2}$ character — open symbols are again taken from [104]. The Zhang-Rice singlet band moved to lower energy as holes were added, while the upper Hubbard band moved to higher energy.

select an energy, then focussed to a vertical cigar-shaped spot on the sample on the order of 100 μm wide by 0.5 mm high, with the polarization horizontal. The peak energy of the undulator and the energy of the monochromator are controlled together to maximize flux, and a LabVIEW-based program continually optimizes the undulator gap.

Three monochromator gratings are available, to allow the beamline to operate throughout the energy range 70 to 1400 eV (the gratings must be changed manually). Of particular interest here are copper and oxygen spectra, specifically the Cu $L\alpha$ and $L\beta$ edges at 929.7 and 949.8 eV respectively, and the oxygen K edge at 524.9 eV[112].

Copper spectra were taken on the high-energy grating, while oxygen could be taken on either the medium-energy or high-energy grating — the oxygen K edge is in an energy range where either grating may be used. The high-energy grating was used for several experiments (notably most of the YBCO data), but was ultimately found to provide lower intensity. There is no significant difference between spectra collected on the same crystal on different gratings aside from an overall shift due to differences in energy calibration, as shown in Figure 3.3.


Figure 3.3: Comparison of oxygen K edge scans taken on the same crystal using the high- and medium-energy gratings. The sample is Ortho-II ordered $YBa_2Cu_3O_{6.5}$, with beam polarization along the b axis. The October 2005 data were collected using the high-energy grating, while the June 2006 data used the medium-energy grating. An overall shift in energy has been performed ($+6.25$ eV to the October data, $+0.70$ eV to the June data), along with the normalization procedure described in detail later. There are minor differences between the scans, possibly attributable to aging or ordering of the sample, exposure to air in the intervening 8 months, or a different fraction of the beam illuminating the sample.

The monochromator's exit slit width may be adjusted manually, with a wider slit permitting higher intensities with reduced energy resolution. The beamline was designed for a resolving power of 10000 with both the entrance and exit slits set at

10 μm, but in practice the resolving power is closer to 6600 [113]. Entrance and exit slit widths employed in this work were 30 and 30 μm respectively ("30/30"), 50/50, 70/100, 80/300, and 80/450. The last has a non-Gaussian energy distribution and was used only once, while the others provided energy resolutions for oxygen K edge spectra of 0.24 eV, 0.40 eV, 0.8 eV and 2 eV respectively, and energy resolutions for copper L edge spectra of 0.42, 0.70, 1.4, and 4 eV respectively. The typical slit widths used were 50/50 for XAS and 70/100 for XES, and any spectra collected with slit widths other than these are identified. Changes to the entrance slit width above 30 μm have a negligible effect on the resolving power.


Figure 3.4: XAS manipulator and sample plates. **a)** Here, samples are mounted on two bolts on a copper puck, which is loaded onto the end of the manipulator. The X-rays' paths from synchrotron to sample and from sample to channeltron (for those detected) are indicated with arrows. XES spectra are collected through the metallic cone on the far side of the sample. **b)** View from the other side of the chamber. In this case, a stainless steel sample plate with gold evaporated onto it has been loaded on a copper sample puck.

Figure 3.4 shows the geometry of the measurements. The sample plate is screwed onto a copper sample puck, which is then loaded onto a larger block of copper at the centre of the vacuum chamber. This copper block is attached to a manipulator capable of x, y and z translations as well as rotation about the vertical (z) axis,

and is at the end of a helium flow cryostat (the vacuum space of which is the main chamber itself). This copper block, which will be referred to as the cryostat, contains a thermocouple and heater for controlling and monitoring the sample's temperature. Sample preparation, handling, mounting, insertion and alignment with the beam are described in Section 3.3.

The vacuum chamber is grounded, and the cryostat is connected to the input of a current amplifier but otherwise electrically isolated. When the sample is illuminated by the X-ray beam, electrons may be liberated from the sample. They drift to the walls of the chamber, so a current is required to maintain the sample's charge neutrality. This current is measured as total electron yield (TEY), using a Keithley current amplifier.

Simultaneously, a channeltron in front of the cryostat detects fluorescence X-rays for total fluorescence yield (TFY), in much the same way a photomultiplier tube detects visible light. When the channeltron is active, the TEY signal from the cryostat is greatly reduced, because the high voltage on the channeltron interferes, repelling the electrons back toward the sample; the TEY signal, however, is still quite strong. TEY spectra measured without the channeltron active have far better signal-to-noise than TFY.

To record a spectrum, a LabVIEW-based program adjusts the monochromator to sweep the incident photons' energy and the undulator gap to ensure maximum flux. The monochromator's exit slit position may also be adjusted during a sweep, since the diffraction geometry changes somewhat with energy, but this was not done because of the narrow energy ranges studied and the considerable delays incurred from this extra step at each energy. The monochromator, undulator, exit slit, and sample position are optimized together for the peak or prepeak of greatest interest, to ensure the optimum geometry and intensity at that energy.

While the raw TEY and TFY data are measured, the photoelectron current on the last focussing mirror is recorded so variations in the incident beam intensity may be normalized out.

The software allows the user to set the spacing in energy of the data points and the dwell time per point in separate energy ranges, allowing higher quality, denser points in areas of interest without interrupting the scan. Typical energy steps and dwell times for 0° and 70° sample angles are shown in Tables 3.1 and 3.2 for Cu L edge and O K edge scans respectively. Dwell times are longer for the 70° sample angles, because higher angles require longer dwells for similar data quality. The angle refers to the angle between the beam polarization and the sample's ab-plane — illuminating the

crystal with a beam at normal incidence is $0°$. Scans of c-axis samples also typically required longer dwell times, due to smaller sample area.

Table 3.1: Typical parameters for a Cu L edge scan. Points are denser in regions of interest, and higher dwell times are required at higher angles as less of the beam can illuminate the sample (the crystal becomes narrower than the beam). The main peak visible in Cu L edges is strongly suppressed in the c-axis polarization, and the extraction of the c-axis component from angle-dependence is sensitive to noise, providing an impetus for longer dwell times. c-Axis samples are also typically narrow, requiring similarly longer dwell times.

Energy Range	Spacing	Dwell ($0°$)	Dwell ($70°$)
910 – 925 eV	0.25 eV	2 s	6 s
925 – 936 eV	0.1 eV	2 s	6 s
936 – 945 eV	0.25 eV	2 s	6 s
945 – 957 eV	0.1 eV	2 s	6 s
957 – 980 eV	0.1 eV	2 s	6 s

Table 3.2: Typical parameters for an O K edge scan. As for copper, points are denser in regions of interest, and higher dwell times are required at higher angles and for c-axis samples as less of the beam is incident on the crystal.

Energy Range	Spacing	Dwell ($0°$)	Dwell ($70°$)
510 – 525 eV	0.2 eV	0.5 s	1 s
525 – 535 eV	0.05 eV	0.5 s	1 s
535 – 545 eV	0.1 eV	0.5 s	1 s
545 – 560 eV	0.25 eV	0.5 s	1 s
560 – 600 eV	0.5 eV	0.5 s	1 s

As internal consistency checks, spectra were taken at multiple locations on most crystals; at the original location for an angle of $0°$ after measuring either at other angles or on other crystals; and on gold, stainless steel, silver paint and phosphor. There were generally no significant variations from point to point on a sample, or over the course of a week. A much faster, coarser spectrum is often collected before each final spectrum, to ensure that the system's operating parameters are correct, the beam is on the sample, and so forth.

To collect emission spectra, the sample plate is rotated to $\sim 45°$, and emitted X-rays are collected at $90°$ to the incident beam. The entrance slit to this part of the

Figure 3.5: Schematic depiction of XES geometry as viewed from above. The synchrotron X-rays incident on the sample are absorbed, then X-rays emitted toward the XES setup are collected through a slit, diffracted off a grating, and detected on a channelplate, which may be translated horizontally to detect different energies.

apparatus, visible as a cone in the background in Figure 3.4a or on the left in Figure 3.4b, is controlled by piezos. A vertical diffraction grating appropriate to the energy range being studied is placed in the path of the X-rays, and they're detected on a channelplate (see Figure 3.5 for the geometry). The emitted X-rays arrive spread out by the diffraction grating, so the horizontal axis on the channelplate corresponds to energy. In this case, instead of rocking the grating to select the energies incident on the channelplate, the channelplate can be translated horizontally to allow the collection of different energy ranges.

The energy of the X-rays incident on the Tl2201 crystal is fixed by setting the beamline's monochromator and undulator, so recording the emission spectrum for each energy merely involves patience — the intensity of the emitted photons is low, particularly since only a narrow slice of solid angle is being examined and we're seeking to discern their energy spectrum (in detail) from a diffracted beam. Even with the beamline's entrance slit set fairly open to improve intensity, count times can be a substantial fraction of an hour for some incident energies, and spectra are typically left noisy by necessity. Data are collected for different incident energies to examine how the emitted spectrum changes as a function of excitation energy. Data for the Cu L edge and O K edge were collected using the same grating, but copper data were from the second-order diffraction spectrum.

The channelplate comprises a 1024×1024 grid of detectors. Since vertical variations are minor, largely uninteresting, and slowly-varying, sets of adjacent rows are binned, resulting in a user-selectable effective resolution, typically 1024×32. The
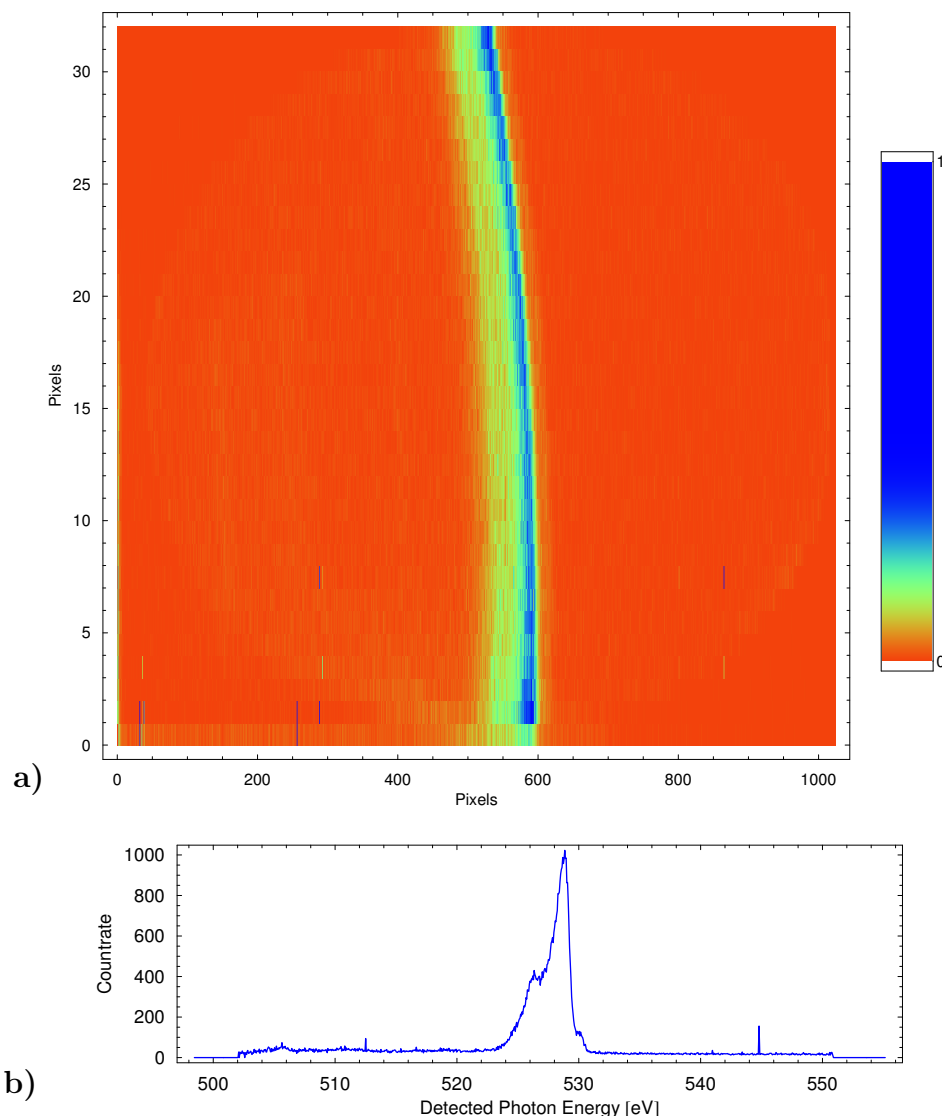
a)

b)

Figure 3.6: **a)** Raw XES data for the oxygen K edge with incident photons of 530.0 eV. It should be possible to discern side peaks to the left and right, separated from the main peak by troughs. Photons can strike the grating at angles other than normal to the grating lines, leading to the detection of an arc by the channelplate. The channelplate is not mounted at precisely the correct angle, hence the rotation. The top and bottom ~ 3 rows are typically discarded when the rows are combined. **b)** The spectrum resulting from this, once the rows are averaged and position converted to energy.

standard grating equation is

$$d\left(\sin\theta_m(\lambda) + \sin\theta_i\right) = m\lambda \tag{3.1}$$

for grating spacing d, diffraction order m, incident angle θ_i and diffracted angle $\theta_m(\lambda)$. Since photons striking the grating at a slight upward angle, for instance, encounter fewer grating lines over the same distance travelled (a larger d spacing), they're diffracted at a lower angle. A single-energy peak therefore appears on the channelplate as an arc rather than as a line (see Figure 3.6). LabVIEW-based software is available at the beamline that corrects for this, such that all rows may be averaged without smearing, then converts position on the detector into energy.

3.3   Sample Handling and Mounting

Like most other barium-containing oxides, YBCO and Tl-2201 react with the CO_2 in air, particularly in the presence of water vapour, to decompose into $BaCO_3$ and oxides of the other cations. Tl-2201 decomposes more rapidly than many of the other cuprates, and particularly YBCO, so special precautions must be taken during storing and handling — the crystals are stored in a special dessicator complete with NaOH pellets to absorb CO_2. Still, because even a thin surface layer of carbonate could potentially affect the results due to the limited penetration lengths of the X-rays, the crystals are etched before being mounted on the sample plate. The mounting is done shortly before departure, and the crystals and sample plates travel in a small airtight container with dessicant and NaOH. NaOH is not considered necessary for YBCO, but dessicant is still used.

The etching was performed in a dilute solution of bromine (Br_2) in anhydrous ethanol ($\sim 1\%$ by volume for YBCO, $< \frac{1}{2}\%$ for Tl-2201). Tl-2201 etches significantly faster than YBCO, and the first etchant used for Tl-2201 was too concentrated and left etch pits. Figure 3.7a shows one of the crystals used in the XAS study after such an etching treatment — etch pits are clearly visible. Provided these are shallow compared to their width, they're irrelevant for XAS. The Tl-2201 etch pits were square, as expected, and are believed to be aligned with the tetragonal ab axes of the crystal, as for other cuprates. Without etch pits, the easiest way to locate these axes is by X-ray diffraction — unlike YBCO and most other cuprates, Tl-2201 grows in irregular-shaped platelets, with few hints as to the orientation of the a and b axes. YBCO crystals, being cuboidal, are straightforward to align by eye.

Given that the crystals are stored in dessicators or under vacuum except during handling, it is unclear that etching is actually required for fluorescence-detected XAS measurements. It is mentioned here because the step was performed, impacts surface topography, and provides a means to align the crystals' a axes to within a couple degrees. While the electronic structure should be identical for all in-plane polarizations due to the symmetry of the lattice, this alignment is nonetheless desirable.


Figure 3.7: A $T_{\mathrm{c}} = 9.5$ K Tl-2201 crystal and XAS sample plate. **a)** The $T_{\mathrm{c}} = 9.5$ K Tl-2201 crystal most thoroughly studied by XAS, after etching but before mounting. Square etch pits are visible, indicating the orientation of the crystal axes. **b)** The same crystal is shown, attached to the sample plate with silver paint and masked off with gold foil. **c)** An XAS sample plate with five Tl-2201 crystals mounted on it, including the one shown in the other two panes. From top to bottom, their T_{c}s are 46 K, 69 K, 69 K, 9.5 K, and 9.5 K. For scale, the sample plate measures $\frac{5}{8}$"$\times$1".

Crystals are mounted onto a stainless steel sample plate using silver paint, chosen because it is conducting and does not require elevated temperatures to cure — such

temperatures could make oxygen atoms mobile and change the doping near the surface of the crystal. For crystals less than about 2 mm in length along the long direction of the sample plate, as is the case for all Tl-2201 and almost all YBCO crystals studied, gold foil is attached (again with silver paint) around the crystal. Recall that the beam spot is longer in this direction; this helps ensure that stainless steel or excess silver paint cannot contribute to the detected signal. Gold does not oxidize and has no absorption edges near the O K or Cu L edges, so it will not add spurious signals if struck by the beam. It will contribute, at most, a smooth, slowly-varying background, which is easily dealt with. For compounds in which the oxygen atoms are not mobile at low temperatures (such as those shown in Figure 3.4b), gold may be evaporated onto the sample plate, but this technique is not suitable for most oxygen-doped cuprates because of the temperatures associated with the deposition.

An additional complication arises for some underdoped YBCO crystals. In these crystals, dopant oxygen atoms are mobile at or just above room temperature, and the degree of order affects the carrier doping. Since studying this effect was one of the key reasons for performing the experiment, it was necessary to transport the crystals to the synchrotron cold, and ultimately insert them into the vacuum chamber cold. Such crystals, once mounted on the stainless steel plate, were taken to the synchrotron sealed in a bag of sealing film, either on dry ice or in a cooler bag filled with ice packs. Once there, they were stored in a freezer.

Once at the synchrotron beamline, every sample plate must be attached with #0-80 screws to a copper puck mentioned earlier. A load-lock chamber attached to the main vacuum chamber can accommodate up to five such pucks, which are then selected and transferred to the main chamber using a transfer arm. When temperature was important, the sample garage was replaced with a cooled version which had a lower capacity, but could monitor the temperature of the sample pucks. In addition, a dry box made of plastic film was assembled around the load-lock, flushed with nitrogen, then equipped with a copper plate cooled by a dry ice and acetone bath, allowing the sample plate to be attached to the copper puck without warming it.

In cases where the sample's temperature is to be controlled during a measurement, a set screw on the back of the cryostat is tightened to more snugly attach the copper puck, using a screwdriver attachment which may be affixed to the load-lock arm in place of a sample puck.

To align the sample in the beam, several techniques are used, often in combination. A phosphor is dusted or painted onto the sample plate, to make the X-ray beam visible. This is extremely useful for coarse alignment, and is normally adequate

for crystals larger than those studied here. To make this technique less coarse, a rudimentary telescope is available. To more accurately align the sample in the beam, the undulator and monochromator are adjusted to a prepeak energy in the case of the oxygen edge, or to the copper edge. The TEY or TFY signal may then be monitored on a multimeter as the sample is moved across the beam. Gold, silver paint or stainless steel will generally give a very different signal than will the crystal, allowing the edges of the crystal to be readily located. Finding the centre of the crystal (or avoiding blemishes on the crystal surface) is then straightforward. Having the system tuned and aligned at the edge or prepeak of interest ensures that the energy range of highest importance has the best signal-to-noise and ensures the best alignment should the position of the beam vary with energy due to misalignment of the monochromator.

3.4   Data Reduction and Analysis

Once the raw data have been collected, several data handling steps are required before spectra can be meaningfully compared against each other.

When a dataset is collected, for every incident energy there is recorded a measurement of the photoelectron current on one focussing mirror, the unnormalized total electron yield (a current) and the unnormalized total fluorescence yield (a count rate). The first step in analyzing these data is to normalize the electron and fluorescence yields by the incident intensity (I_0) as measured on the focussing mirror, to normalize out changes in the beam intensity. The intensity of X-rays incident on the sample is, in general, a strong function of energy, and is also a function of time since the current in the ALS storage ring decayed slowly over each 8-hour period between refills. Figure 3.8 shows raw TEY, TFY and I_0 datasets before any normalization or other manipulation.

In principle, TEY gives a much stronger signal than TFY, requiring far shorter count times for the same quality of data. However, electrons liberated within the bulk of the crystal cannot reach the surface as readily as can fluorescence X-rays — the penetration length for liberated electrons is on the order of 10-100Å, whereas the penetration length for X-rays is generally on the order of tenths of microns or microns, making TFY more of a bulk probe and TEY quite surface-sensitive [102, 114]. Due to the small size of the crystals, difficulties in cleaving them, background signal issues, and differences between TEY and TFY spectra, TFY was selected. Figure 3.9 compares simultaneously-recorded TEY and TFY datasets — the TFY version

Figure 3.8: Raw TEY, TFY and I_0 traces (taken simultaneously) for the most thoroughly studied $T_c = 9.5$ K Tl-2201 crystal, $\vec{E} \perp \hat{c}$, O K edge.

has less of a background problem, and while the two appear to have similar features, their relative intensities do not agree quantitatively. Examples of such discrepancies are discussed in the figure's caption. Recall that while TFY is being measured, high voltages on the channeltron interfere with TEY, greatly reducing its signal strength and making its noise levels similar to those of the simultaneously-collected TFY.

It should be noted that since TFY arises from a count rate related to the fraction of the beam illuminating the sample while I_0 from a photoelectron current elsewhere in the beamline, detected via an adjustable-range current-to-voltage converter, the absolute numbers on the vertical scale have little physical meaning, and further renormalization steps should not result in any loss of information. Cross-sections could be computed and used for the vertical scale, given an appropriate calibration sample, but this was not done, as it would not result in any additional useful information.

Next, to compare datasets at different dopings, temperatures, locations on the same crystal, or orientation, the datasets must be normalized against each other to eliminate background effects. This is done by selecting one energy below the first features, where no X-rays should be absorbed, and another in the continuum well above the absorption edge, at an energy where the intensity observed is slowly-

Figure 3.9: TEY compared with TFY, same raw data as Figure 3.8. Both traces have been normalized against I_0, then TEY scaled to a similar intensity to TFY and shifted vertically to aid comparison. For examples of discrepancies between the two methods, the prepeaks' intensity ratio differs, and several features above the main edge have different strengths compared to the prepeak and main edge. All traces in all following graphs are TFY normalized by I_0 unless otherwise indicated.

varying and should be invariant under changes in orientation, doping, temperature, etc. The intensities at these points are arbitrarily set to zero and one respectively for all datasets, allowing comparisons amongst them. Normalization points used for the oxygen edge in Tl-2201 were 526.75 eV and 576 eV. Figure 3.10 shows the effect of this process, for a series of spectra as a crystal is rotated from $\vec{E} \perp \hat{c}$ toward $\vec{E} \parallel \hat{c}$. Changes in the background and overall signal strength are attributed to how much of the beam is illuminating the sample surface versus the surrounding surfaces — at high angles, the area of the crystal exposed to the beam becomes quite narrow, and the data eventually become unreliable (typically between 70° and 80° for the Tl-2201 crystals studied).

One further step is energy calibration. The monochromator we were using had energy offsets as high as 10 eV for some datasets, and the energy calibration would vary slightly from run to run, particularly if a different grating were used to study

a different edge between runs. Variations in energy calibration from run to run can be dealt with by shifting the data to align peaks that are known to be at the same energy. Larger offsets, due to switching gratings or from month to month, could also be accounted for by measuring standards periodically, which would provide confidence in our energy calibrations. Additionally, comparison to an electron energy loss spectroscopy (EELS) measurement could be used, since that technique is not subject to the same variability. However, the absolute energy scale is generally not crucial to our measurements, as our analysis is limited to doping- and polarization-dependence for peak identification and relative differences for peak position. The band structure calculations to which our data are compared are referenced to the Fermi energy in any case. Some of the data presented in this paper have been shifted in energy to agree with the earlier results of Pellegrin [109] on other thallium cuprates or Nucker [106] on YBCO as appropriate.

If crystals of the material being studied are plentiful and thick enough, the straightforward way to measure c-axis polarization is by having a c-axis sample. In the cases of cuprates that may be grown by TSFZ, such as LSCO or Bi-2212, these conditions are met, c-axis faces are large and plentiful, and the collection of high-quality data is rapid. YBCO cannot be grown by TSFZ, but crystals up to $\sim$200 μm thick exist, so ac or bc faces are prepared by cutting a thin slice off the crystal and polishing that slice. As mentioned above, these faces are smaller than ab faces, requiring longer count times for the same quality of data. Tl-2201 cannot be grown by TSFZ, and our crystals are typically 10 μm thick (occasionally up to 100 μm), making c-axis faces impractical.

In the case of Tl-2201 and other systems where the thickness of available crystals precludes c-axis samples, angle-dependent scans must be taken (see Figure 3.10), from which the c-axis component may be extracted. The X-rays' electric field amplitude incident on the crystal as a function of angle, $\vec{E}(\theta)$, constitutes a combination of ab and c polarizations:

$$\vec{E}^2(\theta) = \vec{E}_{ab}^2 \cos^2 \theta + \vec{E}_c^2 \sin^2 \theta \tag{3.2}$$

Now, recalling that the intensity $I \propto \vec{E}^2$ and solving for the c-axis component:

$$I_c = \frac{I(\theta) - I_{ab} \cos^2 \theta}{\sin^2 \theta} \tag{3.3}$$

which allows calculation of the c-axis component from a scan at normal incidence
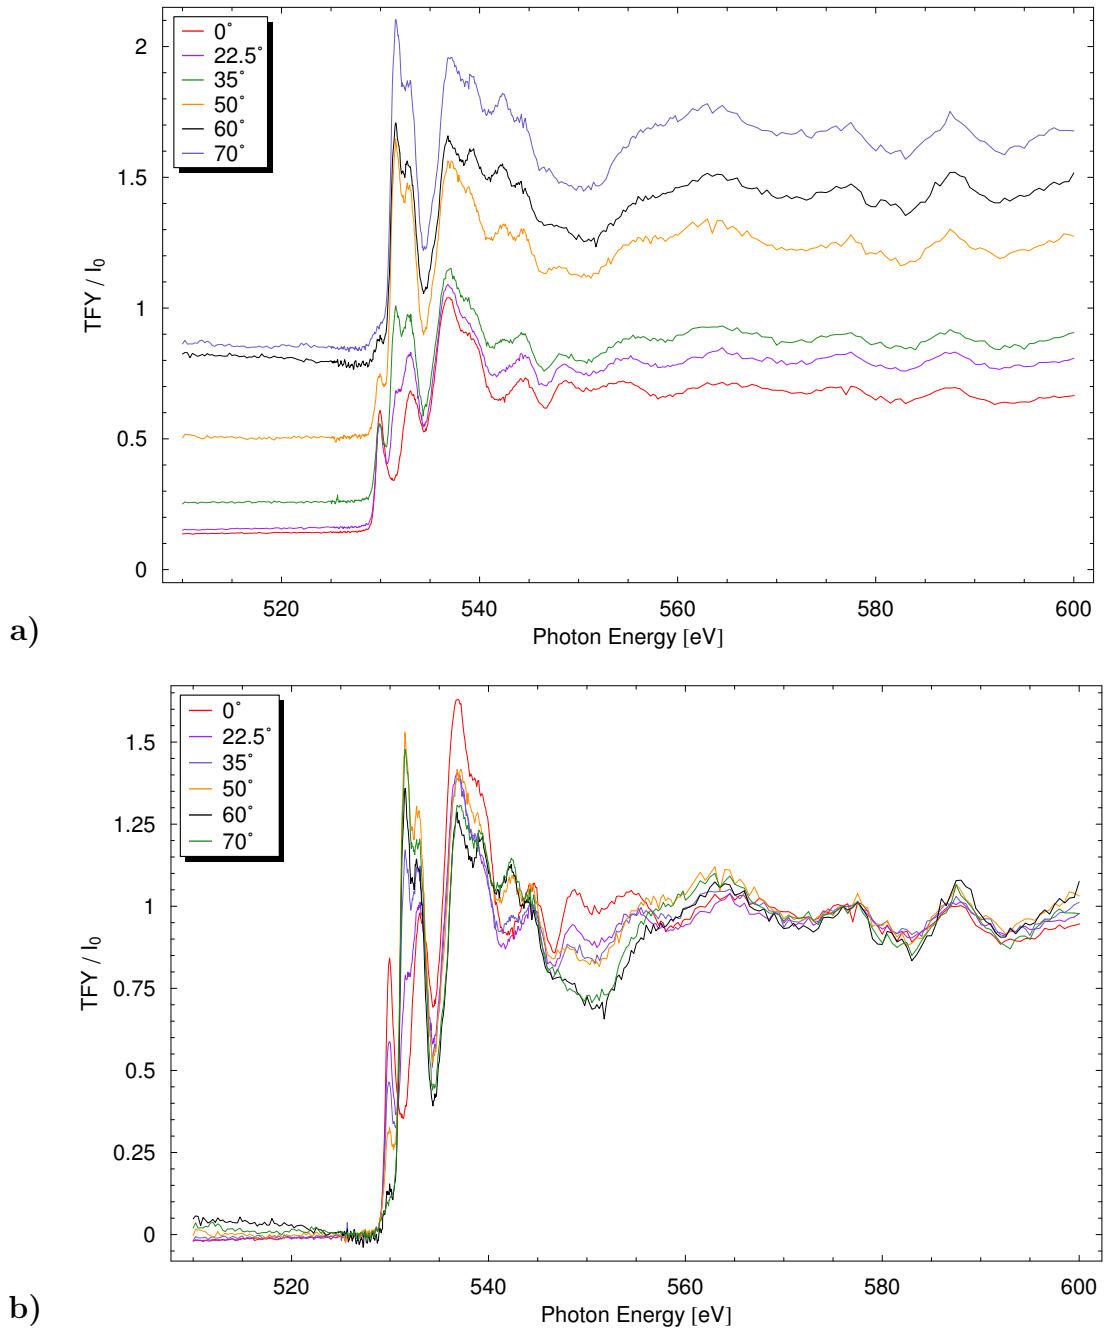
a)

b)

Figure 3.10: TFY datasets as the sample's angle relative to the photons' polarization is changed from $\vec{E} \perp \hat{c}$ (0°) toward $\vec{E} \parallel \hat{c}$ (90°), same crystal as Figure 3.8. **a)** The traces have been normalized by I_0 but have not yet been scaled or shifted relative to each other. **b)** Points below the first prepeak and well above the main oxygen K edge have been set to zero and one respectively to scale the traces against each other.

and a second scan at any other angle. The closer the latter angle is to 90°, the more reliable the results.

While this calculation is performed for all angles, the *c*-axis contribution is not large until higher angles, so the *c*-axis component extracted from small-angle spectra is generally quite unreliable and frequently useless, possibly due to a high sensitivity to slight misalignments between the crystal, the sample plate, cryostat or manipulator, and the beam. While the extracted *c*-axis component tends to converge at higher angles (i.e. beyond about 50°), this method should still not be considered quantitative. Frequently, glitches may be seen where features have been suppressed with angle, even in spectra obtained from high-angle data. Figure 3.11 demonstrates the improvement in the quality of extracted *c*-axis components for higher angles.

The data have not been corrected for effects like saturation or self-absorption. Fluorescence X-rays emitted deep within the sample must reach its surface unmolested to be counted, and a line that is too strongly absorbing can absorb all photons at an abnormally shallow depth, measuring less of the material. The latter effect, saturation, could reduce the height of the main copper-edge peak. If a point is reached before the maximum absorption where nearly all photons are already being absorbed by this edge at a shallow depth, measuring an absorption that's stronger than 'nearly all' may be difficult.

As for self-absorption, since XAS depends upon the material absorbing X-rays in a strongly energy-dependent manner, some features, particularly very strong ones, may suppress themselves by reabsorbing the emitted photons. For less strongly-absorbing features, the incident beam can again penetrate further and the emitted X-rays can traverse more of the sample to escape, allowing more of the sample's volume to contribute effectively to the signal.

While the data can be corrected for these effects, a qualitative anaylsis such as is done here does not require it, and changes in prepeaks and shoulders of main peaks should not be significantly affected.

X-ray absorption data were collected on several compounds, including Tl-2201, YBCO and LSCO, and emission data were collected on LSCO and Tl-2201. This thesis concentrates on Tl-2201, with some limited YBCO data included for contrast.
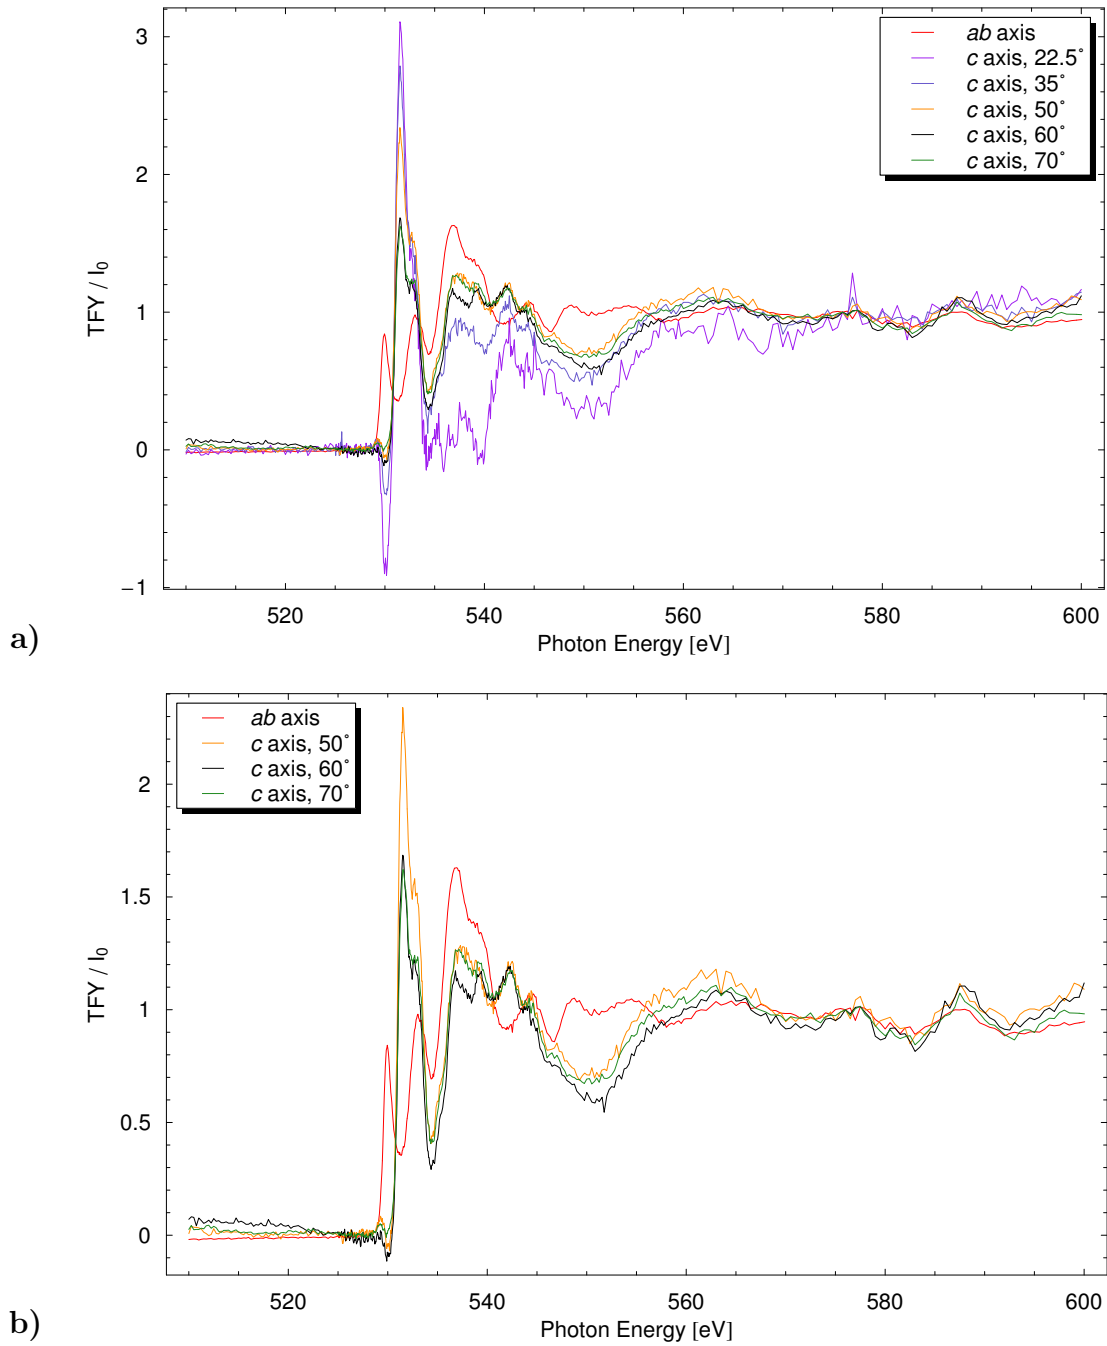
a)

b)

Figure 3.11: The *ab*-plane (0°) data compared against the *c*-axis contribution extracted from the oxygen K edge traces in Figure 3.10. **a)** All angles are shown. For lower angles, the small *c*-axis contribution leads to unreliability in the extracted *c*-axis component. **b)** For higher angles, the traces converge (note especially the *c*-axis components from 60° and 70°). The *c*-axis result still isn't quantitative — note the glitch at 530 eV — but is useful for qualitative analyses.

4. X-ray Absorption Results on YBCO

XAS data were collected for electric field polarizations along the a, b and c-axes of $YBa_2Cu_3O_{6+\delta}$ crystals with oxygen contents $6+\delta$ of 6.0, 6.35, 6.5, and 6.99. Due to the availability of thick crystals, separate samples were prepared for each axis, obviating the need for angle-dependent scans. However, YBCO is orthorhombic for oxygen contents above about 6.3, with significantly different electronic structure expected along each axis, requiring three samples for all oxygen dopings but 6.0.

YBCO's chain layer, shown shared between the top and bottom of its unit cell in Figure 1.2b, consists of a nearly-square two-dimensional array of copper atoms with oxygens above and below each (the apical oxygens). As oxygen dopants are added, they occupy sites along the b axis edges, forming one-dimensional structures and differentiating the b axis from the shorter a axis. Oxygen dopant atoms are incorporated neutral (from O_2 gas), but exist in the crystal structure as O^{2-}, requiring that they remove electrons from (introduce holes into) the lattice. How this can lead to carrier doping of the planes in a manner dependent on the length of a chain segment ("chainlet") is depicted in Figure 4.1, which shows a chain layer where, proceeding from a) to d), each chain has one more oxygen atom than the one before it.

Chain copper atoms in $YBa_2Cu_3O_{6.00}$ are generally taken to be Cu^{1+}, $[Ar]3d^{10}$ (neutral copper atoms have the electronic configuration $[Ar]4s^13d^{10}$). Adding an isolated oxygen atom to this layer will take an electron from each of the two neighbouring copper sites, ionizing them to Cu^{2+} ($[Ar]3d^9$), leaving both depleted (triply ionizing Cu is energetically quite unfavourable). Adding an oxgyen atom in a site adjacent one of these coppers will thus require that an electron be taken from somewhere else in the crystal. The number of holes introduced to the rest of the crystal by a chainlet of n oxygen atoms will be $n-1$. This is a justification rather than a full explanation, and should not be taken too literally — it will be tempered in reality by correlated electron physics, any delocalization of Cu $3d$ core orbitals, and details of the band structure.

YBCO's chain layer is fully deoxygenated for $6+\delta = 6.0$ and very close to fully oxygenated for $6+\delta = 6.99$, leaving little opportunity for mobile oxygen atoms to diffuse, order, cluster or in any way vary the carrier doping. Between these dopings,

Figure 4.1: Depiction of how carrier doping in YBCO can depend on chain length. **a)** An isolated oxygen atom, as shown in the first chain, will pull one $3d$ electron from each of its neighbouring copper atoms. **b)** When two oxygen atoms are added adjacent each other, the shared copper atom can only donate an electron to one of the oxygens, requiring that an electron be taken from elsewhere in the crystal. **c)** and **d)** show the continuation of this as more oxygens are added, allowing chainlets to grow and merge.

however, the degree of order in the chain layer (via the number of chainlets) plays a crucial role in setting the carrier doping. As mentioned earlier, a variety of oxygen-ordered phases exist, including Ortho-II, Ortho-III, Ortho-V, and Ortho-VIII [43, 62, 63], where the roman numerals refer to the periodicity along the a axis of the ordering. For instance, Ortho-III has two full chains followed by an empty chain. The chain layer in Ortho-II is shown in Figure 4.2. Dopant order makes for more defect-free crystals, with a good periodic crystal structure, properties which are particularly important for techniques that rely upon long mean free paths, such as microwave spectroscopy [115, 116] or quantum oscillations [17, 19].

Chain length and the degree of order can be controlled by aging the crystals near room temperature to promote lower-entropy arrangements, or raising the temperature to enter a less-ordered phase and more randomly distribute the dopants. The application of pressure also promotes lower-entropy, more space-efficient phases. Near room temperature, oxygen atoms can slowly diffuse through the crystal, gradually forming whichever phase is favoured in that temperature and pressure regime,

but the oxygen atoms are not believed to enter or leave the crystal to any significant degree at these temperatures. One component of the YBCO work was to examine changes in YBCO's electronic structure due to the ordering of dopant oxygen atoms.


Figure 4.2: Oxygen chain order in Ortho-II YBa$_2$Cu$_3$O$_{6.50}$ — full and empty chains alternate.

One set of samples at $6+\delta = 6.35$ was transported to the ALS fully disordered, after having been quenched from a temperature just below the tetragonal-orthorhombic transition. These were allowed to order at room temperature for several hours at the beamline, a length of time that would correspond to a several-Kelvin change in T_c but not necessarily a large change in doping — the slope of T_c with doping is thought to be quite steep around 6.35. Additionally, a set of Ortho-II crystals were measured, then disordered by heating to 200°C, well above the disordering temperature, for several hours, after which they were rapidly reinserted onto the cryostat and remeasured. This produces crystals with every chain half-occupied and the oxygen atoms distributed randomly. The data taken on 6.35 crystals, along with data on several other dopings, are presented in Appendix A.

4.1   XAS Results on YBa$_2$Cu$_3$O$_{6.99}$ Crystals

XAS results on fully-oxygenated YBa$_2$Cu$_3$O$_{6.99}$ crystals are presented in this section, to demonstrate some capabilities of the technique and to show the significance of the differences that may be encountered between different polarizations.


Figure 4.3:   Copper $L\alpha$ edge for YBa$_2$Cu$_3$O$_{6.99}$ crystals as a function of the incident X-ray beam's polarization.

Figure 4.3 shows the copper $L\alpha$ edge for YBa$_2$Cu$_3$O$_{6.99}$ for X-rays with $\vec{E}$ along each crystallographic axis. A very clear polarization-dependence is visible: the main peak, around 931 eV, has no out-of-plane (c) component, while features around 932 and 933 eV exist in b and c polarizations, but have little to no weight in the a polarization. The lowest-energy peak, representing the lowest-level excitations involving copper orbitals, is thus identified with the CuO$_2$ planes, while the other features are identified with the chains.

The corresponding oxygen K edge scans are shown in Figure 4.4. Again, a clear polarization-dependence is visible. It is not productive to enumerate all the oxygen K edge peaks and their polarization dependences, other than to note that the features exhibit tremendous variety in this regard. The oxygen K edge shows features associated with almost every layer in the crystal structure, with copper oxide chains

Figure 4.4: Oxygen K edge for $YBa_2Cu_3O_{6.99}$ as a function of polarization.

and planes, chainlet ends, and correlated electron features such as the upper Hubbard band. There are at least eight distinct features visible below the edge itself around 537 eV. This level of complexity is common in the cuprates, and makes interpretation quite challenging.

The dependence of the spectra on polarization allows us to determine which specific orbitals (e.g. Cu $3d_{3z^2-r^2}$) are contributing to each feature.

4.2 Doping Dependence

Since several dopings were examined spanning $YBa_2Cu_3O_{6+\delta}$'s entire doping range, the doping dependence of the spectra can be examined. This provides more information as to the nature of the features observed — for instance, oxygen K edge features associated with chains will be absent for $6 + \delta = 6.0$.

The doping dependence of the Cu L and O K edges for polarization along the a axis are shown in Figure 4.5. The copper edge shows some broadening on the high-energy side of the main peak as doping is increased, along with complete suppression of a second peak at 934 eV. Data taken with $\vec{E} \parallel b$, not shown in this chapter, exhibit a much more pronounced hump on the high-energy side of the edge. The oxygen edge
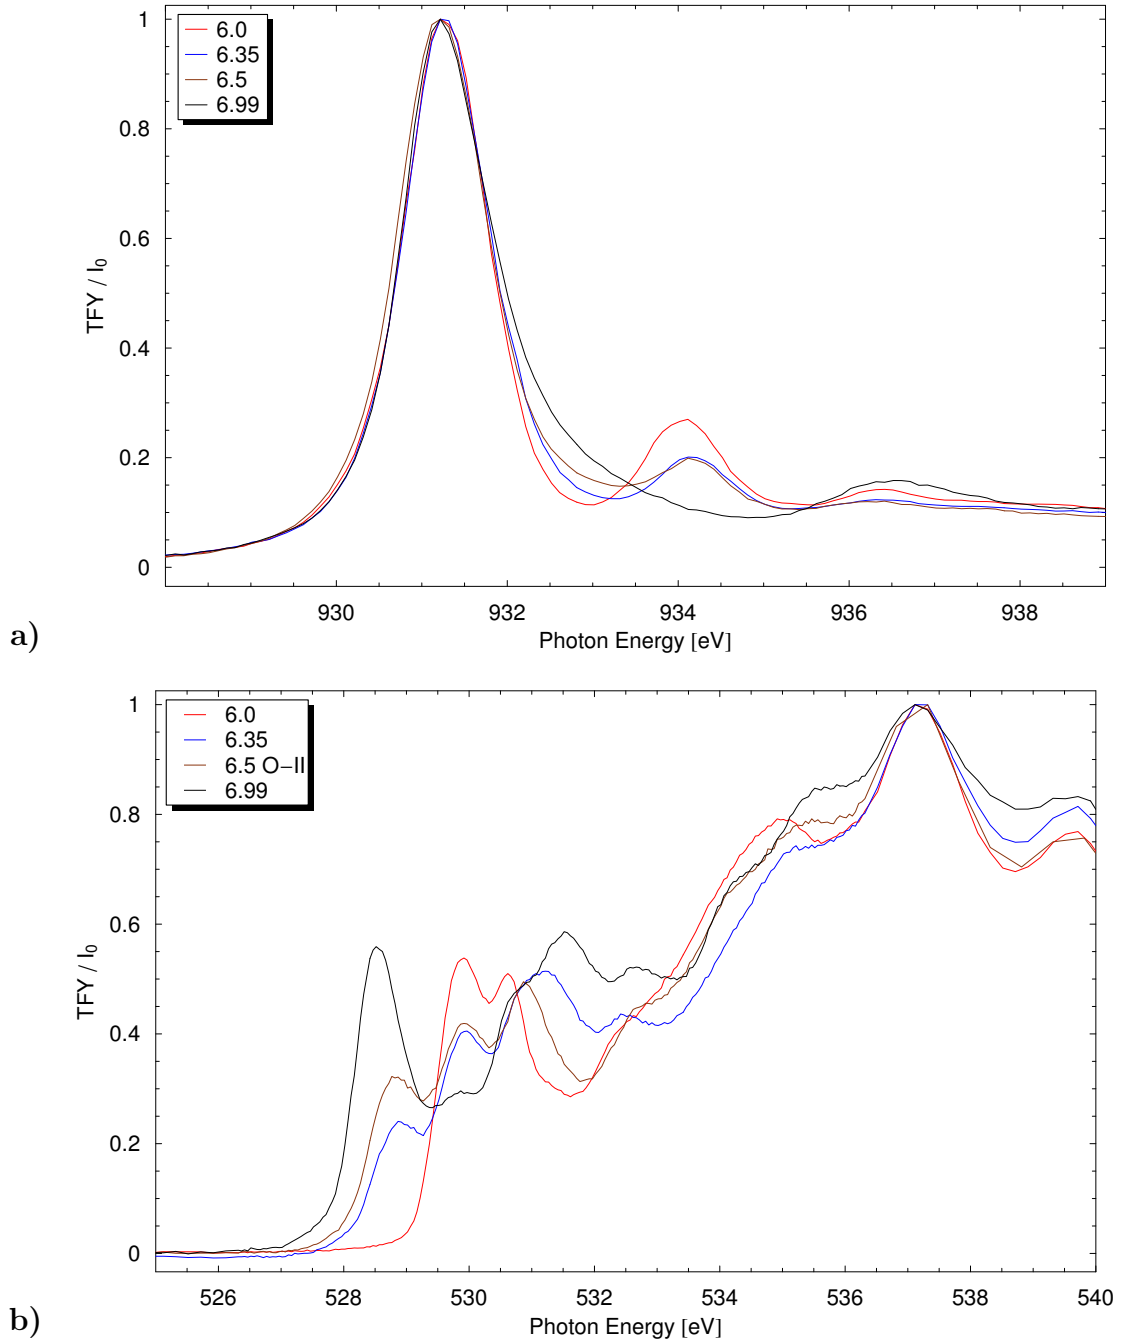
Figure 4.5: Doping dependence of YBCO's **a)** copper L and **b)** oxygen K edges for X-ray polarization $\vec{E} \parallel a$.

shows the appearance (and shift to lower energy) of a significant prepeak at 528.5 eV that doesn't exist in the undoped compound, partial suppression of a feature at 530 eV, and the appearance of a feature around 531.5 eV.

Figure 4.6 shows the doping dependence of the Cu L and O K edges for polarization along the c axis. Here, the copper edge changes substantially — the highest peak for $6 + \delta = 6.0$ is suppressed entirely, while the small lower-energy peak grows into a strong double peak. The oxygen edge similarly undergoes massive changes, with a double peak appearing at the low-energy end and most other features being largely suppressed as doping is increased.

4.3   Ortho-II Order in YBa$_2$Cu$_3$O$_{6.5}$ Crystals

For the oxygen dopings $6 + \delta = 6.35$ and 6.5, crystals were measured with differing degrees of oxygen chain order. As chain length increases, any features due to the ends of chainlets will weaken and the carrier doping will change. Figure 4.7 shows this effect for YBa$_2$Cu$_3$O$_{6.5}$. After this set of crystals was measured Ortho-II ordered, their oxygen order was scrambled at 200°C, and they were remeasured in the disordered (disordered Ortho-I) state. This process was repeated eight months later, in June 2006; the changes on disordering were consistent, but less substantial. This is attributed to incomplete ordering or disordering.

Changes in the copper L edge c-axis spectra were not sufficiently reproducible to permit firm conclusions to be drawn; the Appendix, section A.4 shows the June data for that edge and polarization. The oxygen K edge c-axis scan after disordering the first time was weak and exhibited an unusual background, so data from the June beamtime are used in their place. The a- and b-axis data for both edges, both here and in the Appendix, are otherwise from October, 2005.

The prepeak around 528.8 eV for polarizations along a and b is suppressed on disordering, as is the predominantly c-axis feature (also weakly visible along the b-axis) at 528 eV. A prepeak at 530 eV in the a and b polarizations is enhanced, while changes to other features were not reproducible enough to draw conclusions. The changes observed are in almost all cases consistent with a reduction in hole doping on disordering. One notable exception is the first prepeak on the c-axis spectrum, which is suppressed slightly by disorder but enhanced by increased oxygen doping. This feature does not exist in scans with $\vec{E} \parallel a$, suggesting that it may be chain-related. If so, it would not be surprising that increased hole doping on the planes would result in reduced hole doping on the chains for the same oxygen content.
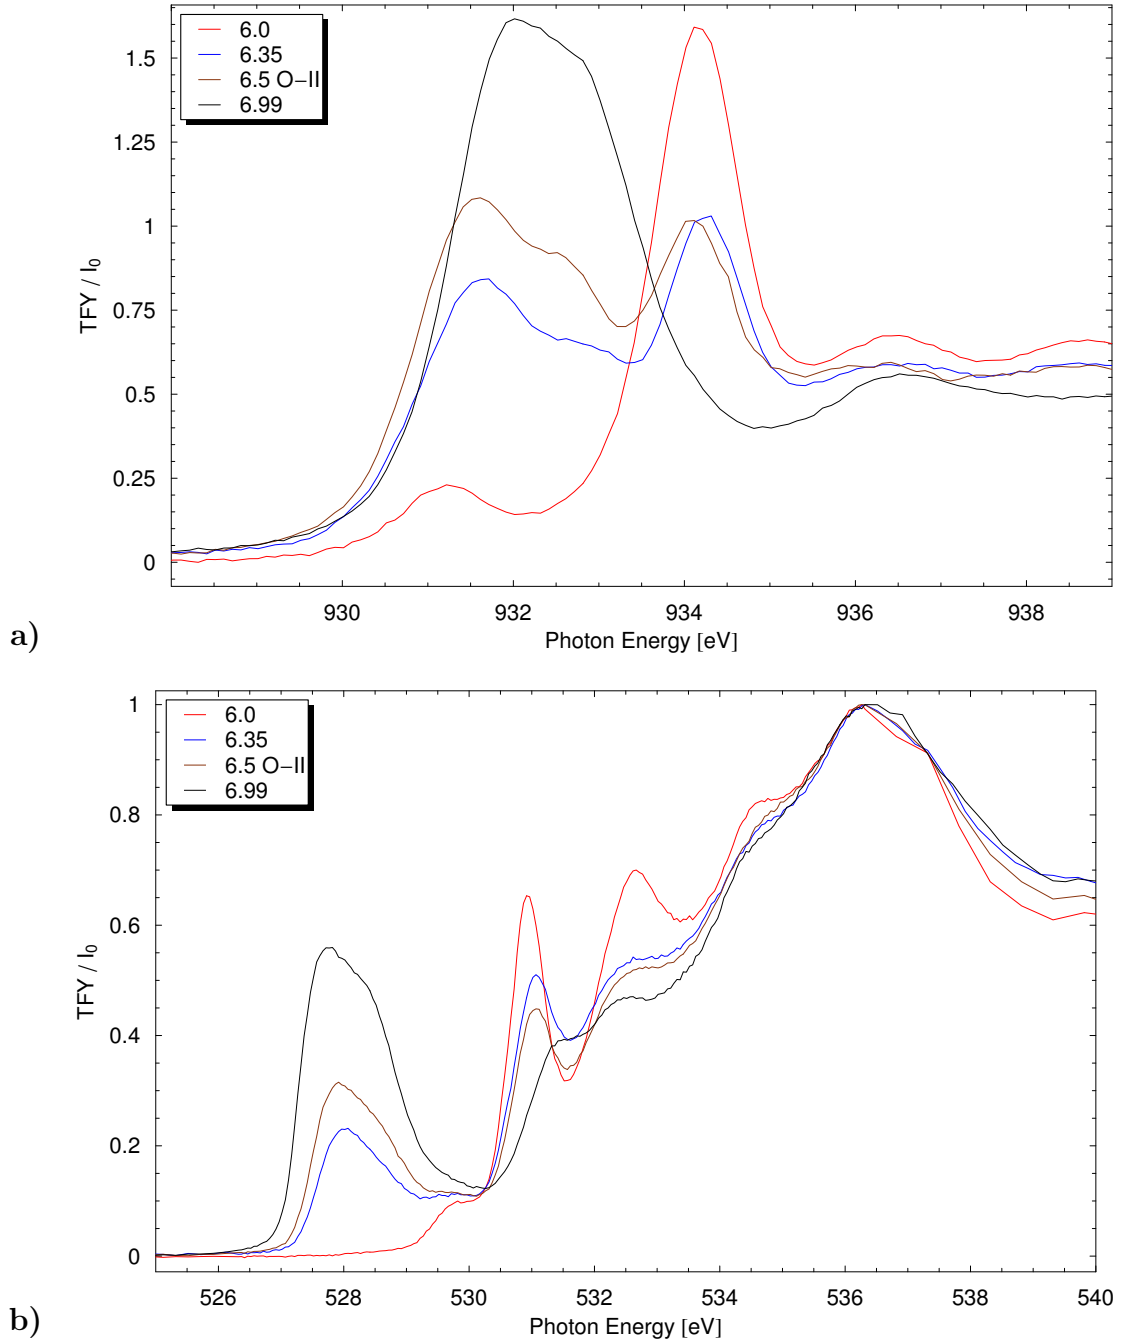
Figure 4.6: Doping dependence of YBCO's **a)** copper L and **b)** oxygen K edges for X-ray polarization $\vec{E} \parallel c$.

Figure 4.7:   Oxygen K edge for Ortho-II ordered and for disordered (Ortho-I) $YBa_2Cu_3O_{6.5}$ crystals. Changes to the prepeaks are clearly visible.

Two features common to the cuprates that will be looked for in Tl-2201 warrant mention here. The Zhang-Rice singlet band is expected to correspond to the lowest-energy excitations in the system (possibly tied with the chain states that provide the holes), as these were the holes that were added on doping and had to be accommodated in the plane oxygens. This band, composed of oxygen $2p_x$ and $2p_y$ orbitals as shown in Figure 1.13, should by definition not exist at zero doping, and increase in weight more or less linearly as holes are added. Such a feature is seen in Figure 4.5b, the doping evolution of the a-axis spectra, around 528.5 eV. Its energy (the binding energy) is reduced on doping. The second feature worth mentioning is the upper Hubbard band. This band is expected to be the second-lowest in energy, chains notwithstanding, be composed primarily of in-plane orbitals, and decrease with doping as the Zhang-Rice singlet band bleeds states out of it. Such a feature is observed in Figure 4.5b at 530 eV. Similar identifications of these peaks were made by Nücker [106, 117].

The key points here are the high sensitivity of this technique, the complexity of the spectra, and the magnitude of the changes as hole carriers are added to the material. In YBCO, as in other underdoped cuprates, massive changes are observed

with doping, and there are a plethora of features to identify, so a great deal of information is available. Significant changes can even be observed reproducibly when the dopant oxygen atoms in a crystal are disordered. Overdoped cuprates have not been studied to any significant degree by this technique.

5. X-ray Absorption Results on Tl-2201

As mentioned in Section 1.2, it is commonly believed that the overdoped cuprates may exhibit much more Fermi-liquid-like properties than other regimes in the phase diagram. While the overdoped regime has seen far less research than the underdoped side, tantalizing hints of normalcy have been seen by the handful of transport and electronic structure techniques that have been applied there. In Tl-2201, the hole states are responsible for the unconventional superconductivity, so there is keen interest in this material's electronic structure, and X-ray absorption spectroscopy, which accesses the unoccupied (equivalently, hole) states in question, can provide crucial pieces of the puzzle.

Using the procedures outlined in Section 3.2, XAS data were obtained on Tl-2201 crystals with T_cs of 9.5 K, 60 K and 69 K at incident angles from $0°$ to $80°$. An additional crystal with a T_c of 46 K was also measured, but the O K edge results did not resemble those at other dopings, and it has been omitted from this discussion as it is considered unreliable. This crystal appeared visibly discoloured (see Figure 3.7c), suggesting a layer of contamination on the surface.

5.1   XAS Results on Heavily Overdoped Tl-2201 Crystals

The copper L edge c-axis components extracted for a $T_c = 9.5$ K crystal are shown in Figure 5.1. The c-axis spectra extracted at angles $\geq 40°$ are in near-quantitative agreement. While there are clear differences from data with polarization in the ab-plane, it should be noted that the main edge appears not to be completely suppressed. Indeed, it seems to retain roughly 35% of the in-plane weight. However, the peak has shifted to lower energy, strongly suggesting that this may be a distinct feature. A full discussion of this feature may be found in Section 5.4.

At this doping, for the oxygen K edge, there are three energy ranges at which setting the intensity to unity for normalization purposes was contemplated: well above the edge, at the edge itself, or in the notch between the edge and its prepeaks. The first should, in principle, be more reliable if the background is not changing, because the edge or notch could potentially change intensity with angle. If the background

Figure 5.1: A comparison of the *ab*-plane and *c*-axis copper L edge spectra for a $T_c = 9.5$ K crystal The *c*-axis component extracted for higher angles converges, and is markedly different from the *ab* component.

changes with angle, however, a normalization point nearer the prepeaks may be best. Surprisingly, the most consistent results were obtained when the spectra were normalized by the third method. Graphs obtained by all three normalization methods for this and other dopings are included in Appendix B for reference.

Figure 5.2 compares the in-plane spectra with *c*-axis spectra extracted from scans taken at several high angles. The prepeak at 530 eV is entirely absent for $\vec{E} \parallel c$ (the glitch there suggests a slight misalignment of the sample), while a strong prepeak at 531.5 eV, absent in the *ab*-plane data, appears. The intensity of the third prepeak, around 533 eV, does not appear to change significantly. Modelling this peak to determine whether it was polarization-dependent was unsuccessful, due primarily to the change in the background level across this peak.

5.2   XAS Results on a $T_c = 60$ K Tl-2201 Crystal

The *ab*-plane and *c*-axis spectra for a $T_c = 60$ K Tl-2201 crystal are displayed in Figure 5.3. As for the $T_c = 9.5$ K crystal described in Section 5.1), the reliability

Figure 5.2: A comparison of the *ab*-plane and *c*-axis O K edge spectra for a $T_c = 9.5$ K crystal. The *c*-axis component extracted for higher angles converges, and differs significantly from the *ab* component.

is expected to be higher for higher angles, although in this case measurements were taken at fewer angles. The spectrum extracted from the scan at 60° will be used for future analysis. Features' polarization dependence is similar to that observed on the $T_c = 9.5$ K sample; this comparison will be discussed further in Section 5.4.

The oxygen edge *ab*-plane and *c*-axis polarizations for the $T_c = 60$ K crystal are compared in Figure 5.4. Again, the extracted *c*-axis spectra converge. As before, the prepeak at 530 eV is entirely absent for *c*-axis polarization, while the strong *c*-axis polarization prepeak at 531.5 eV remains, and the third prepeak, around 533 eV, does not appear to exhibit polarization-dependence.

5.3   XAS Results on T_c = 69 K Tl-2201 Crystals

The $T_c = 69$ K crystal's *ab*-plane and *c*-axis copper L edge components are compared in Figure 5.5. As for the $T_c = 60$ K crystal, data were taken at fewer angles, and the spectrum extracted from the 60° data will be used for all further analysis. The spectra behave much like those at the other two dopings.

Figure 5.3: A comparison of the ab-plane and c-axis copper $L\alpha$ edge spectra for a $T_c = 60$ K crystal.

The oxygen K edge ab-plane and c-axis polarization spectra for the $T_c = 69$ K crystal are shown in Figure 5.6. The data taken at 70° had a strong background that could not be corrected for, and is thus considered unreliable, so the c-axis component extracted from the 60° spectrum is used in all following analysis. The 70° spectrum shown has had a substantial parabolic background subtracted from its raw data.

As for the other two dopings, the prepeak at 530 eV is absent for $\vec{E} \parallel c$, a strong prepeak at 531.5 eV is only observed in the c-axis polarization, and the 533 eV prepeak does not exhibit clear polarization-dependence.

5.4   Doping Dependence

As can already be seen by comparing spectra shown above for different dopings, very little doping dependence is observed. This stands in stark contrast to doping dependence in the underdoped regime presented in Section 4.2 and reported in previous work, for instance on YBCO [106] and LSCO [103]. There, there were substantial changes between zero doping and optimal doping, with significant features appearing, shifting in energy and disappearing altogether with doping. On Tl-2201, the changes

Figure 5.4: A comparison of the *ab*-plane and *c*-axis oxygen K edge spectra for a T_c =60 K crystal. The *c*-axis component extracted for higher angles again converges and is quite different from the *ab* component.

for most polarizations and edges are much less significant.

5.4.1 Copper *L* Edge

In the underdoped cuprates, significant changes with doping have been observed on the copper L edge, particularly on the high-energy side for in-plane polarizations (see, for example, Figure 4.5a) and throughout the *c*-axis spectrum (see Figure 4.6a). It should be noted that the YBCO examples offered have chain states, and YBCO's CuO_2 planes are not perfectly planar, leading to some mixing of orbital symmetries. In Tl-2201, where the CuO_2 planes are thought to be genuinely planar, and chains and bilayers are absent, the interpretation should be more straightforward. As will be shown, the changes with doping are substantially less pronounced in this material.

Figure 5.7 shows the doping dependence of the copper L edge for in-plane polarization. There may be slight changes with doping on the shoulder around 930.5 eV — it appears slightly higher in overdoped crystals — but there is otherwise very little evidence for doping dependence. A spectrum taken by Pellegrin on the trilayer

Figure 5.5: A comparison of the *ab*-plane and *c*-axis Cu *L* edge spectra for a $T_c = 69$ K crystal.

cuprate Tl-2223 at optimal doping [109] is included for comparison, showing the same asymmetric lineshape. These points do not, however, allow a meaningful comparison with Tl-2201's possibly doping dependent shoulder.

To check the shoulder for a systematic doping dependence, three schemes were used. First, the $T_c = 9.5$ K crystal's spectrum was subtracted from each spectrum, as shown in Figure 5.8a, so that any absolute increase could be observed. Data are included for a second $T_c = 9.5$ K crystal, a second $T_c = 69$ K crystal, and a second spectrum on the original crystal. A scaled version of the $T_c = 9.5$ K crystal's spectrum that was subtracted from the others is included for reference. The second scheme involved instead dividing each spectrum by that for the $T_c = 9.5$ K crystal (again scaled and included for reference), to check for relative changes. The results obtained by this method, shown in Figure 5.8b, are similar, but more clear.

Because the spectrum has a non-zero signal above the edge, with an unknown onset, fine details such as the shoulder at 933.5 eV cannot be easily modelled for quantitative comparisons between dopings or angles. This analysis should be considered qualitative only.

By both of these two methods, for all data at all dopings, there does appear to be

Figure 5.6: A comparison of the *ab*-plane and *c*-axis O K edge spectra for a $T_\mathrm{c} = 69$ K crystal.

a systematic change with doping, although it is difficult to conclude this definitively due to the scale of the change in comparison to the noise level. The $T_\mathrm{c} = 9.5$ K scans are not depressed noticably around 930.5 eV, while those on $T_\mathrm{c} = 60$ K and 69 K crystals are, the 69 K ones more so than the 60 K.

The third method assumes that the *c*-axis peak arises entirely from copper atoms on the thallium site. Band structure calculations shown in Figure 5.16a and 5.17b indicate that the lowest-energy unoccupied band in this layer contains mainly thallium $5d_{3z^2-r^2}$ character, with some $6s$ (ignoring other species). If a copper atom in that site is assumed to contribute mainly $3d_{3z^2-r^2}$ weight, subtracting $\frac{1}{4}$ of the *c*-axis spectrum from the *ab*-plane data would leave only the CuO_2 plane component of the spectrum[4]. The results of this method are shown in Figure 5.9. The traces have been scaled by

---

[4]The final state, $d_{3z^2-r^2}$, is $d_{2z^2-x^2-y^2}$, with twice as much weight along z as along x or y. The initial state is a p orbital, where the set of occupied p orbitals is spherically symmetric. The transition probability from Fermi's golden rule is proportional to $|\langle \psi_f | T | \psi_i \rangle|^2 \delta (E_f - E_i - \hbar\omega)$, where the transition operator T is proportional to $e^{i\vec{k}\cdot\vec{r}}$ from the electric field, with direction determined by the photon's polarization vector. After Taylor expanding, $T \sim 1 + ikr + ...$, the first term of which integrates to zero. The second term, combined with the polarization direction, ends up $\propto \hat{\varepsilon}\cdot\vec{r}$, which serves to select out components along the electric field. The resulting integral has twice as much weight along z as along the other axes, an anisotropy which squares to 4 [118].
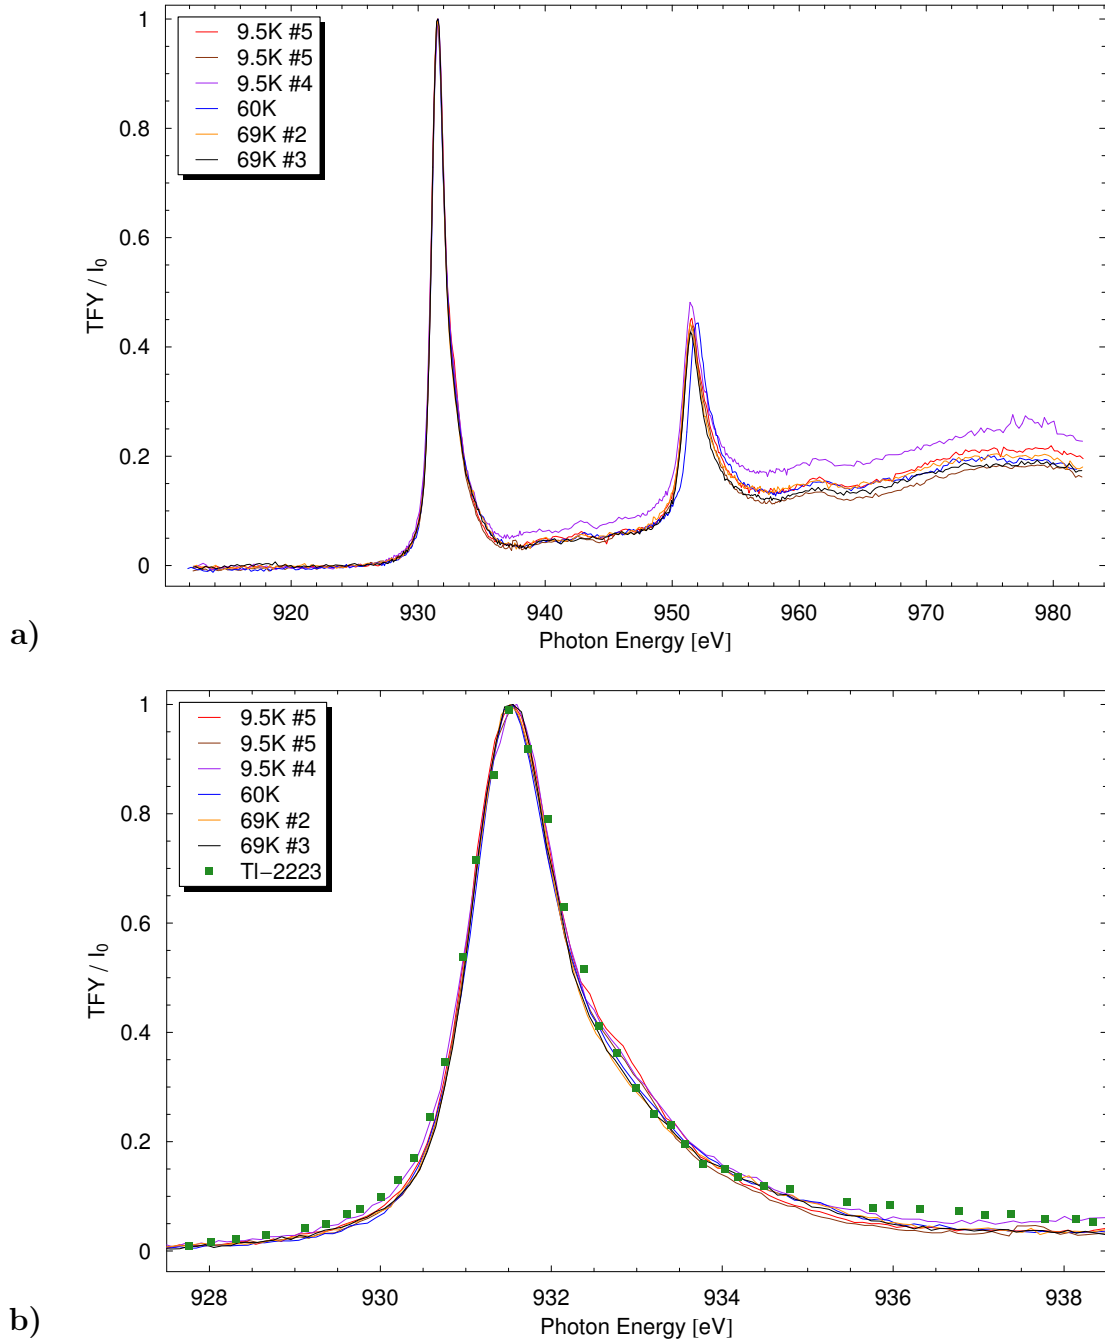
Figure 5.7: Doping dependence of the copper L edge in-plane spectra. **a)** The *ab*-plane spectra, normalized by peak. **b)** Detail of the $L\alpha$ edge and comparison to Pellegrin's results on the trilayer compound Tl-2223 [109] — the shoulder around 933 eV may evolve slightly with doping.
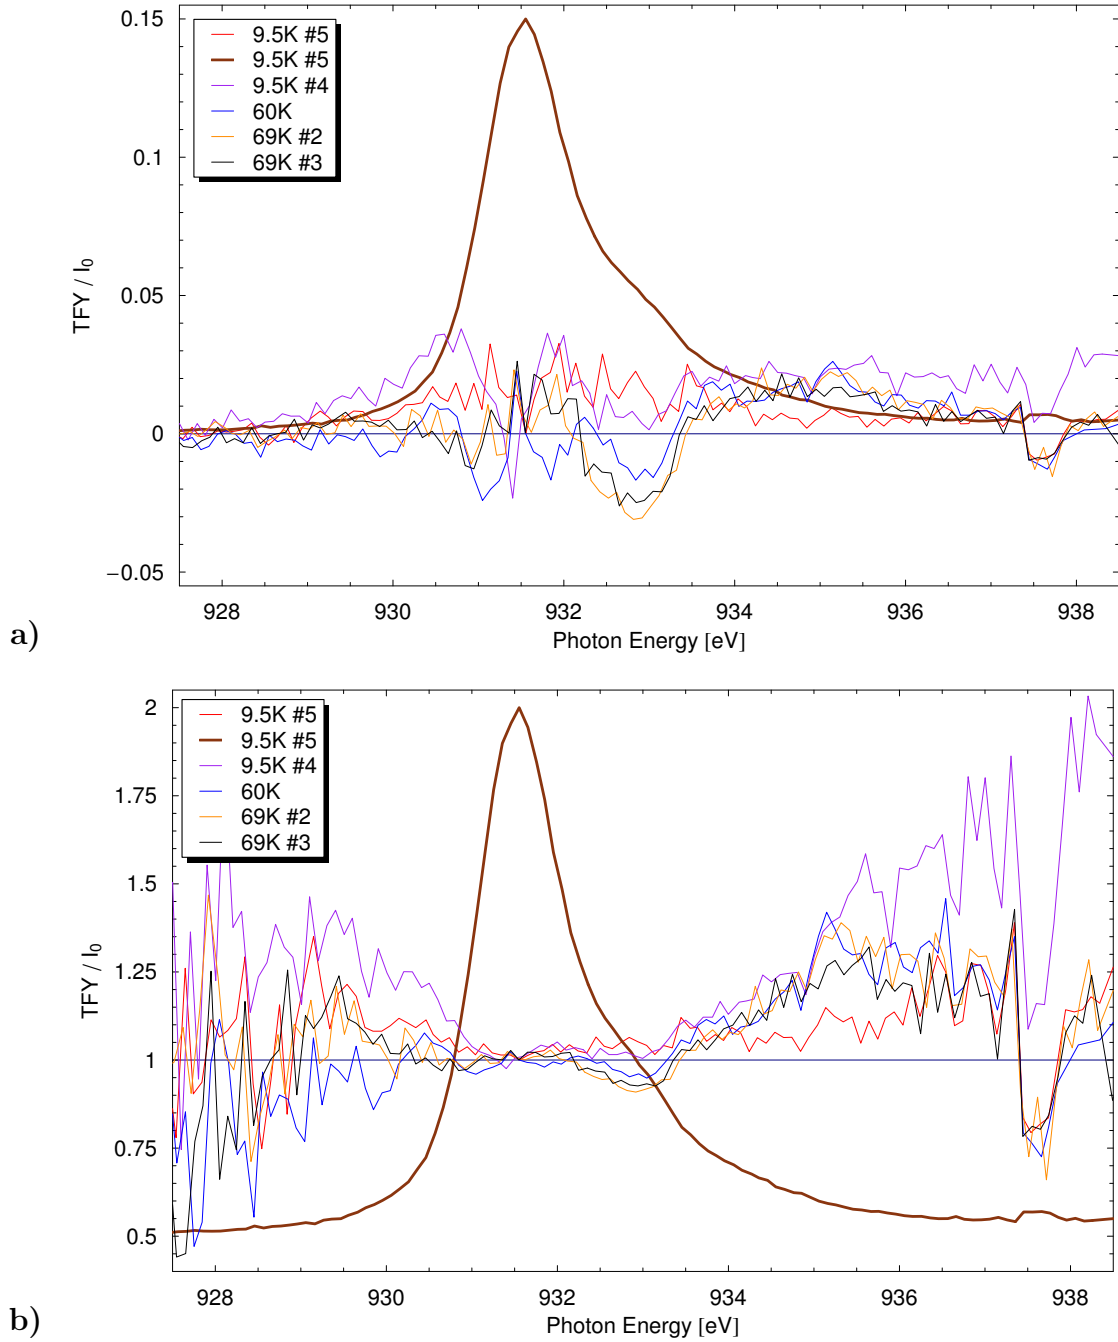
Figure 5.8: Doping dependence of the shoulder at 933 eV. **a)** Subtraction of $T_c = 9.5$ K crystal's spectrum from other spectra: traces from other dopings exhibit a dip at 933 eV. **b)** Division by $T_c = 9.5$ K crystal's spectrum: traces from other dopings exhibit a $5 - 10\%$ dip. The dip is less pronounced for the $T_c = 60$ K crystal than for 69 K.

Figure 5.9: In-plane copper L edge spectra corrected by attributing the c-axis weight to $3d_{3z^2-r^2}$ orbitals of copper atoms on the thallium site.

maximum again — no systematic dependence of peak height on doping was observed.

Figure 5.10 shows the doping dependence of the copper L edge for out-of-plane polarization. Here, substantial increases may be observed in the strength of the $L\alpha$ edge as doping is increased (T_c is reduced). Data points from optimally doped Tl-2223, a trilayer material, are included for comparison, and show a weak peak consistent with the less-doped 60 K and 69 K T_c Tl-2201 samples.

To characterize the magnitude of these changes, the spectra were divided by each other to obtain the relative increase associated with the peak. The shoulder seen in the ab-plane data around 933.5 eV is not strong enough here to justify a comment on any changes with doping. It is worth reiterating that because these spectra are calculated by weighing and subtracting datasets, they contain significantly greater uncertainty than the raw ab-plane data, and cannot readily be used for quantitative analysis. As an example, spectra for a $T_c = 9.5$ K crystal extracted from data taken at 60° and 70° are included in Figure 5.10, and they differ by ∼20%.

The $T_c = 60$ K crystal was used as a baseline for Figure 5.10b because its spectrum was cleaner than that of the 69 K sample, and to allow easier comparisons amongst the spectra than would be possible using a stronger 9.5 K spectrum as a baseline.
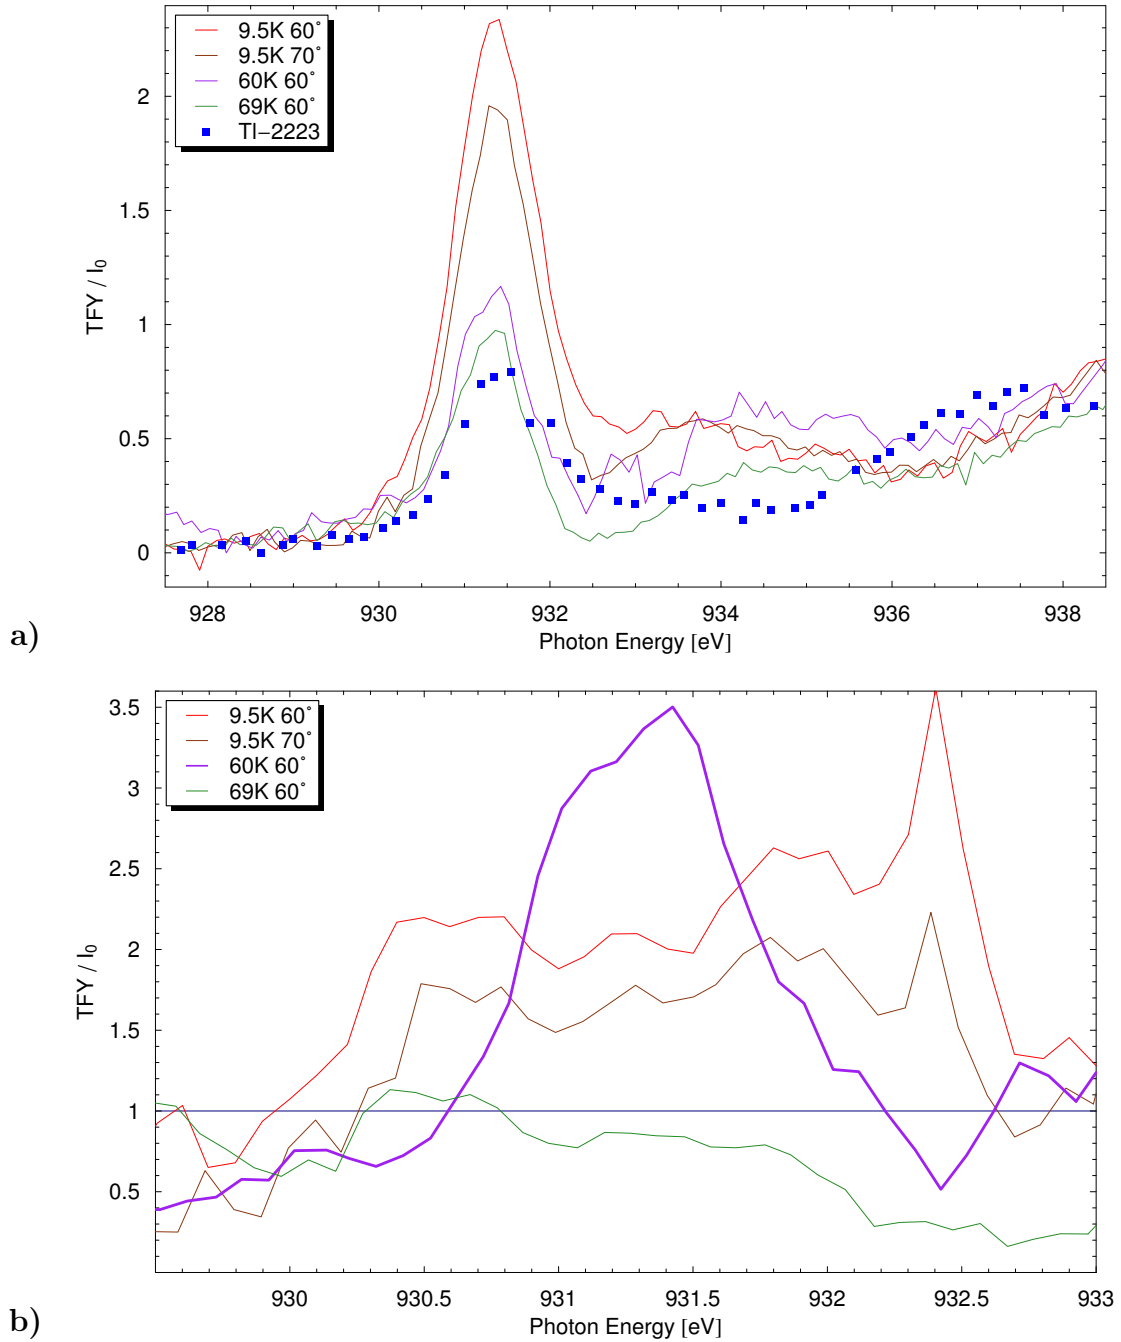
Figure 5.10: Doping dependence of the copper L edge c-axis spectra. **a)** The $L\alpha$ edge — the intensity of the edge itself evolves with doping in this polarization. **b)** Examining doping dependence: division by $T_c = 60$ K crystal's spectrum. As in Figure 5.7, the data points are optimally doped trilayer Tl-2223, from reference [109].

The peak extracted for the $T_c = 9.5$ K crystal is $\sim 1.7 - 2.1$ times stronger and the 69 K crystal's peak about 20% weaker than that for the 60 K crystal. There is no other unambiguous doping evolution visible in these spectra.

As described above in Section 5.1, the weight in the Cu L main edge peak has been calculated as an approximate proportion of the equivalent peak in the respective ab-plane spectrum. For the $T_c = 60$ K crystal, c-axis data extracted from a 60° scan retained 16% of their original weight. For the $T_c = 69$ K crystal and 60°, 13% of the weight survived. As mentioned above, the $T_c = 9.5$ K crystal was found to keep 27% of its ab-plane weight as extracted from a 70° scan. These are all roughly consistent with the results shown in Figure 5.10, and are significantly higher than would be expected.

Table 5.1: Amplitude and integral of the first peak in the copper L edge c-axis spectrum, normalized to the $T_c = 69$ K crystal.

T_c	Lorentzian amplitude /69 K	Lorentzian integral /69 K
9.5 K	4.56	2.46
60 K	1.37	1.24
69 K	1	1

The c-axis peaks were fit to Lorentzians to allow better comparisons of their relative weights and better modelling of the ab-plane shoulder. Table 5.1 and Figure 5.11 show the amplitude of this peak and its integral, as extracted from the Lorentzian fits. A significant change is observed through this doping range, although it is important to realize that the nature of the calculations involved means error bars on these points could easily be $\pm 30\%$.

Since the crystals studied were grown under the same conditions, cation substitution levels at all three dopings should be nearly identical, thus the changes in c-axis weight with doping can only be attributed to hole doping or to the concentration of oxygen dopants. As mentioned above, that the position of this peak is lower in energy than for the ab-plane spectra strongly suggests that this peak is associated with a different copper site, namely copper atoms in the thallium oxide layer. In that case, the most obvious explanation is that proximity to an oxygen dopant enhances the hole doping of these copper atoms. Since about 5% of the thallium sites are occupied by copper atoms and there are at most 0.1 to 0.15 excess oxygen atoms per formula unit [47], the odds of a copper atom having an adjacent interstitial oxygen are still quite low even at the highest dopings. It is conceivable, however, that the dopant

Figure 5.11: Doping dependence of integrated weight in the copper L edge c-axis peak at 931.3 eV, after fitting to a Lorentzian lineshape and subtracting the resulting offset. Error bars of $\sim \pm 30\%$ would not be unreasonable. Hole doping values were calculated from T_{c} using Equation 5.1.

oxygens may preferentially fill sites adjacent copper atoms. It should be noted that Figure 5.11 suggests that the intensity in this peak tends toward zero near or above optimal doping — Wagner [47] found that occupation of the O(4) site went to zero around optimal doping.

The comparison with Tl-2223 in Figure 5.10a is not particularly helpful in confirming this peak's identification. In the trilayer compound, there are three CuO_2 planes instead of one, separated by layers of calcium atoms, while the blocking layers are identical. It is likely that the Tl-2223 had more copper substituted onto the thallium site than did Tl-2201, because of our efforts to reduce this defect. Regardless of which site is responsible, there should be more such copper atoms per unit cell in Tl-2223. In addition, Tl-2223 does not have the same relationship between dopant concentration and T_{c}, because the doped charge carriers are spread over three CuO_2 planes instead of one (and they may not be distributed evenly amongst those three). The comparisons to Tl-2223 are, however, useful in highlighting the changes between the optimally doped and highly-overdoped regimes: on the copper L edge, Tl-2201's

only substantial departure from its trilayer analogue is an increase in weight in the
c-axis peak attributed to substituted copper, and this is consistent with the optimally
doped trilayer compound.

The significant departures from underdoped cuprates bear reiterating here — for
in-plane polarizations, almost no doping dependence is observed, while significant
doping dependence is seen for this polarization in YBCO (Figure 4.5a) and LSCO
[103]. For out-of-plane polarization, some doping dependence is observed, but less
than for YBCO (Figure 4.6a). Further discussion of the features observed and their
doping dependence is provided in Section 5.5.

5.4.2   Oxygen *K* Edge

In $YBa_2Cu_3O_{6+\delta}$, many substantial changes were observed to the oxygen K edge, for
all polarizations, on hole doping. The Zhang-Rice singlet band appeared and grew to
be quite strong, while moving to lower energy, the upper Hubbard band weakened, and
several other features changed substantially as well, although these may be associated
with chain states. For reference, the a-axis and c-axis spectra for YBCO were shown
in Figures 4.5b and 4.6b respectively. Similarly substantial doping dependence was
observed in LSCO [103], although LSCO showed little change between a lightly and a
heavily overdoped sample. This section establishes that Tl-2201 more closely follows
the trends (or lack thereof) observed in overdoped LSCO than those in underdoped
materials.

Figure 5.12 shows the doping-dependence of the oxygen K edge for in-plane polar-
ization. The most astonishing result here is the almost complete lack of any doping
dependence over a wide range in doping. As was seen in Section 4.2, YBCO's oxygen
prepeaks exhibit massive changes with doping; this is common on the underdoped
side of the superconducting dome.

There are slight changes with doping on the first prepeak — it is stronger and
at slightly lower energy in the more strongly overdoped $T_c = 9.5$ K crystal. To
characterize this change, the prepeak was integrated (from 528 eV to 531 eV) for
each doping. The $T_c = 9.5$ K crystal's prepeak had 11.4% more weight than the
60 K crystal's and 10.9% more than the 69 K crystal's. Other normalization schemes
produced different values here, but all agreed that the increase was in the vicinity of
10%.

The lowest-energy oxygen prepeak is expected to be the planes' Zhang-Rice singlet
excitations discussed earlier, rather than interstitial dopants (it is too strong to be

Figure 5.12: Doping dependence of the oxygen K edge in-plane spectra.

dopants in any case — the dopants are dilute). A useful benchmark, then, is p, the hole concentration. An empirical formula relating p to T_c has been developed based on LSCO [46] which is commonly assumed to be a good approximation for other cuprates:

$$1 - \frac{T_c}{T_c^{max}} = 82.6(p - 0.16)^2 \tag{5.1}$$

This formula neglected an anomaly around $p = 1/8$ on the underdoped side, and requires corrections around this doping in several systems, including YBCO (around an oxygen content of 6.5) [45], but has been roughly confirmed on that system over much of the underdoped regime. It has seldom been applied in the overdoped regime, and never verified there on a system other than LSCO, so its reliability there is unclear.

Using this formula and a maximum T_c of 94 K, the hole doping p would be 0.264 for the 9.5 K crystal, 0.226 for the $T_c = 60$ K crystal, and 0.217 for the $T_c = 69$ K crystal, meaning the $T_c = 9.5$ K crystal should have 22% more holes than the $T_c = 69$ K crystal and 17% more holes than the $T_c = 60$ K crystal. The small enhancement seen in this prepeak is inconsistent with such a large change in doping.

One explanation for this would be that T_c declines significantly faster than Equa-

tion 5.1 states. There is certainly no evidence linking T_c to doping in this compound and very little such data on any other overdoped compound. However, the Fermi surface areas measured by AMRO [14] and ARPES [15, 16] for Tl-2201 crystals with T_cs of 20 K and 30 K respectively correspond to hole dopings of $p = 0.24$ and 0.26 respectively, when the Presland formula (Equation 5.1) would call for 0.26 and 0.25. This level of agreement with the formula would be extremely difficult to square with a $T_c = 9.5$ K crystal having $p = 0.24 \sim 0.25$, the result obtained from the enhanced weight assuming the higher-T_c crystals to still obey the formula. Similarly, if the highly overdoped crystal were assumed to obey Equation 5.1, the other two dopings would need to be around $p = 0.24$.

Another explanation is a breakdown in the Zhang-Rice singlet approximation. As mentioned in Section 1.5, Zhang-Rice singlets cannot be formed where two adjacent copper atoms have holes. At dopings of $p \sim 0.20 - 0.30$, where the Tl-2201 crystals studied are expected to be, the likelihood of holes on adjacent copper atoms reaches unity. It should not be surprising if the Zhang-Rice singlet approximation is breaking down. However, the holes should be going somewhere else near the Fermi energy, and there is no further evidence for doping-dependence on the oxygen K edge for this polarization.


Figure 5.13: Doping-dependence of the oxygen K edge c-axis spectra.

Figure 5.13 shows the doping-dependence of the oxygen K edge for out-of-plane polarization. The prepeaks' general shape is similar, and variations in overall intensity may be an artifact of the extraction of the c-axis component. There is no evidence for systematic doping dependence in these spectra.

Figure 5.14 compares the Tl-2201 oxygen K edge spectra with those of optimally doped Tl-2212 and Tl-2223 reported by Pellegrin [109]. The upper (isotropic) thallium oxide band is about 1 eV lower in these materials than in Tl-2201, and lower in the trilayer than the bilayer material. In the c-axis spectra, it merges with the lower TlO band. There is no clear evidence of an upper Hubbard band, but it should be weak at optimal doping, and the data quality and proximity of the thallium oxide band may mask it. For instance, both datasets show a possible small bump around 531-531.5 eV, but the point spacing and the scatter in data make this consistent with noise. The Zhang-Rice singlet band is stronger in these materials because they have multiple CuO_2 layers contributing to it.

While the doping dependence on both edges might appear uninteresting at first, due to the near-absence of any doping dependence, this is unexpected, and represents a marked departure from the bahaviour observed in the underdoped regime. As mentioned previously, spectra taken on LSCO and YBCO on the underdoped side of the phase diagram (see Chapter 4) undergo collosal changes from underdoping to optimal doping. That this should abruptly all but stop near optimal doping, giving way to barely-perceptible changes, is truly surprising. In the cuprates, normalcy is not normal.

The fact that the doping dependence (on all but the copper L edge, where it is attributed to interactions of dopants) appears to essentially vanish in the vicinity of optimal doping is particularly interesting, since many theoretical models predict quantum critical points in this regime, and the T^* crossover (marking the elimination of the pseudogap) vanishes or merges with T_c near optimal doping [8]. Several other techniques have also seen hints of changes in behaviour above optimal doping, including ARPES lineshapes [15], reconstruction of the Fermi surface as mentioned in Section 1.2 and references [21, 22, 23], and the μSR relaxation rate [119]. How the overdoped side of the superconducting dome can behave so differently from its underdoped counterpart will no doubt remain a topic of intense study in the years to come.
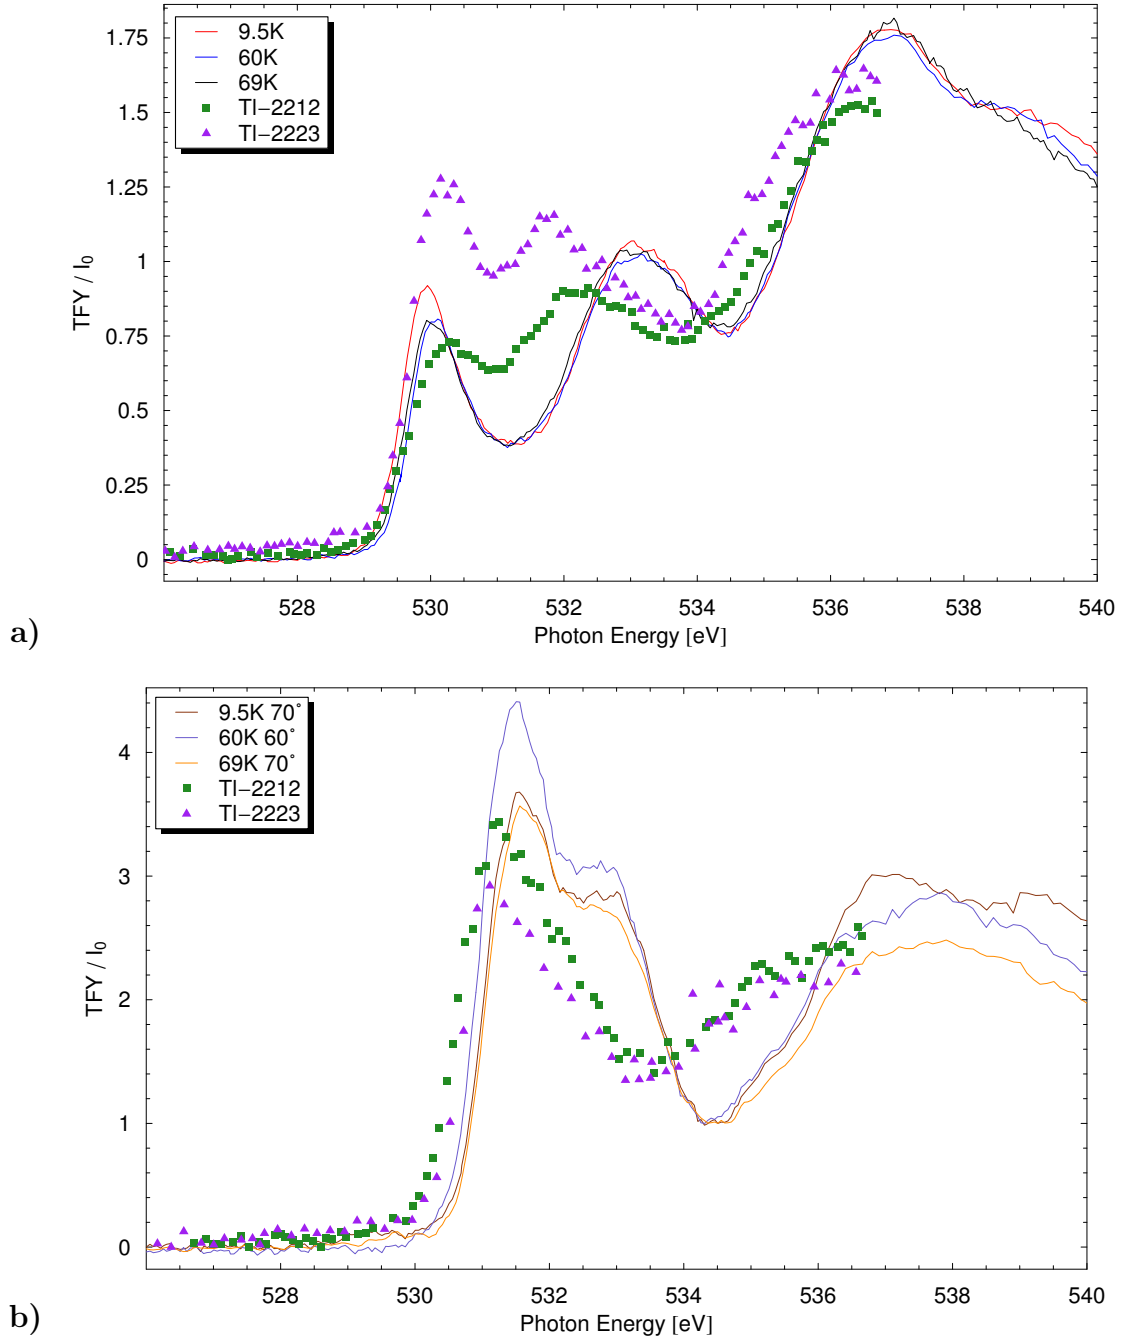
Figure 5.14: Comparison of the oxygen K edge spectra in Figures 5.12 and 5.13 with those of Pellegrin for optimally doped Tl-2212 and Tl-2223 [109]: **a)** in-plane and **b)** out-of-plane polarization.

5.5   Band Structure Calculations

The electronic band structure and resulting densities of states near the Fermi energy
have been calculated for stoichiometric $Tl_2Ba_2CuO_{6\pm\delta}$ by Hawthorn [120], using the
Local Density Approximation (LDA) in the tight-binding linear muffin-tin orbital
(TB-LMTO) approximation and using the atomic positions from Ref. [79] as input
parameters. The contributions to the density of states from different source orbitals
were calculated, and the results are reproduced here. This is invaluable in identifying
features in the XAS spectra.

Figure 5.15: Densities of states at and above the Fermi level for the various atomic
species and sites, from LDA band structure calculations.

Figure 5.15 shows the contribution to the density of states from each atomic
species and crystallographic site in Tl-2201, and Figure 5.16 shows the weight arising
from thallium, barium and copper orbitals. Figure 5.17 shows the thallium and copper
contributions by orbital, and 5.18 breaks down the oxygen contribution by site and
symmetry. These figures will not be discussed further, instead serving as a reference
for a discussion on the features observed in the spectra.

On the copper edge, it is fairly straightforward to identify several of the main
features. The main peak visible for in-plane polarization (see Figure 5.9) is seen in all
cuprates, and is identified with excitations into a $3d_{x^2-y^2}$ orbital. Figure 5.17b, which
shows the contributions of the various Cu $3d$ orbitals, concurs with this evaluation
— the weight around the Fermi energy has predominantly $3d_{x^2-y^2}$ character. This
is commonly referred to as the exciton peak, to reinforce that interactions with the
core hole are strong. The strength of these interactions prevents any quantitative
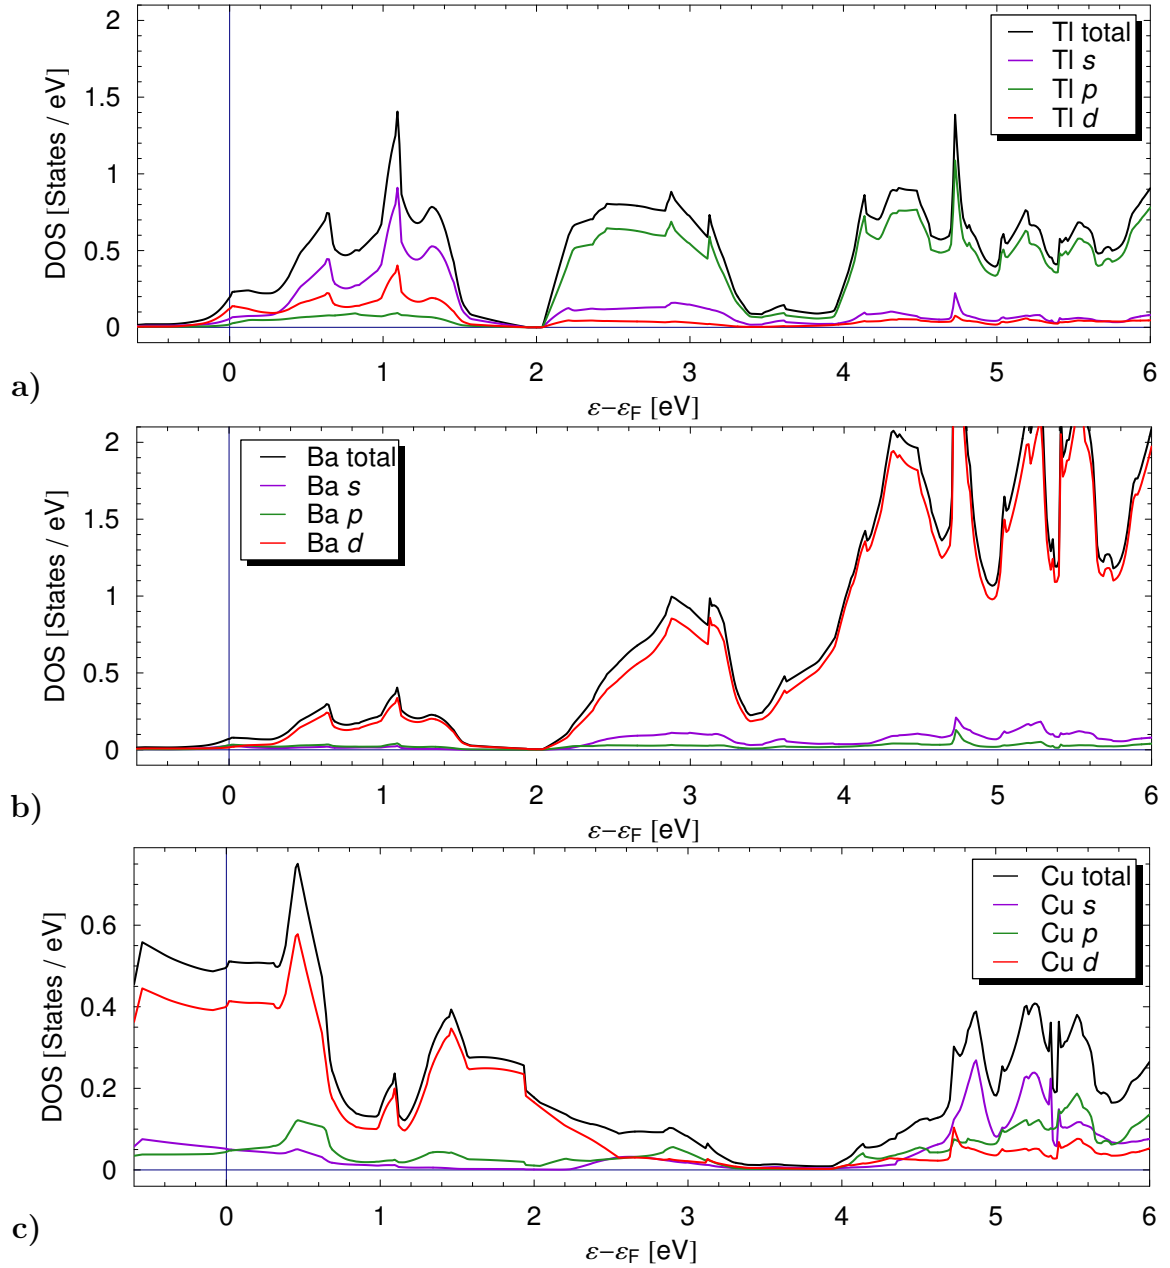extraction of the density of states.

Figure 5.16: Densities of states at and above the Fermi level for **a)** thallium, **b)** barium and **c)** copper, from LDA band structure calculations.

Figure 5.17: Densities of states at and above the Fermi level for different symmetry **a)** thallium and **b)** copper orbitals, from LDA band structure calculations.

About 2.5-3 eV above the main peak, a second peak is visible in the *c*-axis polarization (see Figure 5.10a). It may also be present in the in-plane spectra, but the strength of the adjacent exciton peak makes that difficult to determine. A similar feature in YBCO has been ascribed to excitations into hole states drawn from chain copper and apical oxygen orbitals [106], but the absence of chains in Tl-2201 makes that explanation implausible here. Thallium atoms occupy a site directly above the apical oxygen, so copper substitution could explain this feature. This peak is weak, as would be expected for a feature arising from a dilute defect. It is revisited below, where the oxygen edge is discussed.

The asymmetric lineshape for the main copper *L* edge in-plane peak has been explained as a consequence of hole doping [103, 121, 122]: copper atoms harbouring extra holes (delocalized on the oxygen ligands) will have a higher energy. Figures 5.8 and 5.9, showing the doping-dependence of the in-plane spectra, show a slight enhancement around 933 eV, roughly 1.2 eV above the peak, as doping is increased.

This should correspond to these atoms where the initial state is $3d^9\underline{L}$ ($\underline{L}$ denoting a ligand hole) instead of $3d^9$, and the small change observed on this shoulder would be expected to correspond to a change on the order of 20% to the underlying feature. Note that selection rules forbid the filling of the p-symmetry orbitals holding the ligand hole from a copper $2p$ initial state, so the final state has a full $3d$ shell, with a $2p$ core hole and the original $\underline{L}$ ligand hole.

That leaves the lowest-energy c-axis peak, which increases in weight substantially with overdoping. This was discussed in Section 5.4, where it was tentatively attributed to copper atoms in the thallium site located next to dopant oxygens.

Traversing the overdoped side of the phase diagram with cation-substituted Tl-2201 requires changing the oxygen content by 0.08 to 0.12 per formula unit [47, 66], by which p is changed by 0.11 according to the Presland formula (Equation 5.1). Adding this much oxygen ought to introduce 0.16 to 0.24 holes per formula unit, leaving 0.05 to 0.13 holes unaccounted for. In the crystals studied, which were from an earlier batch grown with a slightly lower thallium content, there should be roughly 0.10 excess copper atoms per formula unit, which could accommodate these holes by becoming Cu^{3+}. No bands besides those from the CuO_2 plane are observed at the Fermi surface to accommodate the extra holes.

One alternative explanation for this feature is that the crystals studied had vastly different copper substitution levels. The crystals with T_cs of 69 and 9.5 K were grown in identical conditions two weeks apart; both batches leaked slightly, the former more than the latter. This suggests a slighly reduced thallium content in the melt and a slightly higher copper substitution level in the $T_c = 69$ K crystal. With higher substitution producing a greater offset in hole doping, this crystal should exhibit a lower T_c for the same oxygen content (or a lower oxygen content for the same T_c), and a stronger c-axis peak than the 9.5 K sample, if this interpretation is correct. The opposite is observed.

If the oxygen dopants are illuminating substituted copper atoms by changing their oxidation state, the 69 K crystal's already lower oxygen content, potentially reduced further once the last paragraph's argument is taken into account, would produce a substantially weaker peak, consistent with the spectra. Taken in combination with the literature, there is strong circumstantial evidence here that the oxygen dopants each dope one hole onto a substituted copper atom and one hole onto the plane, suggesting a non-random proximity of oxygen dopants to substituted copper atoms.

Figure 5.17 shows the contribution to the density of states by each symmetry of d orbital, for thallium and copper. The only thallium orbital with any weight near the

Fermi energy is the $5d_{3z^2-r^2}$ orbital. In the case of the two features identified with copper on the thallium site, the copper atom can be expected to contribute an orbital of the same $d_{3z^2-r^2}$ symmetry as does the thallium atom, making it visible primarily in the c-axis polarization.

Figure 5.18 shows the contribution to the density of states by each symmetry of p orbital, for the three primary oxygen sites. Plane oxygen atoms largely contribute in-plane orbitals, those in the thallium oxide layer contribute mainly p_z character, in agreement with the z-orientation of the Tl d orbitals they're bonding with. Apical oxygens have almost exclusively p_z character, and appear to also contribute most strongly to the thallium oxide band.

The oxygen K edge has three pre-edge features — a peak for in-plane polarization around 530 eV (refer to Figure 5.12), a strong peak at 531.5 eV for c-axis polarization (see Figure 5.13), and a peak around 533 eV with no clear polarization dependence.

The lowest-energy excitations are expected to be into Zhang-Rice singlets. These are comprised of oxygen $2p_x$ and $2p_y$ orbitals at the O(1) (CuO$_2$ plane) site, and Figure 5.18a shows such a band at the Fermi energy for in-plane polarization. Note that, where the O(1) site is formally located, $2p_x$ orbitals point toward the adjacent copper atoms, so there is little $2p_y$ weight; a matching site with $2p_y$ character can be generated by symmetry operations on the unit cell. The identities of the oxygen sites are shown in Figure 2.15. The peak at 530 eV has strictly in-plane character as would be expected for a Zhang-Rice singlet band, and its intensity increases with hole doping, as would be expected for a band of states created to accomodate holes. As mentioned above, the increase in intensity is significantly lower than expected, possibly because of a breakdown in the approximation that the Zhang-Rice singlets be dilute. The lack of clear doping dependence on the low-energy side of the copper L edge indicates that holes are still being added to oxygen atoms, rather than copper atoms.

The Zhang-Rice singlet band in Figure 5.12 also shows a possible slight shift to lower energy as doping is increased. This has been observed in other systems, such as YBCO (see reference [106] or Section 4.3), and is considered unsurprising — as holes are doped in and the Fermi sea's tide retreats, more of the underlying band structure is exposed. Whether the Zhang-Rice singlet states' energy is referenced to the Fermi level or they are considered part of the band structure, introducing more holes creates more such states at the Fermi level, which is retreating, and the energy of the onset should decline. The feature retains its shape, so the peak shifts to lower energy as well. The shift observed over the doping range probed was roughly 0.1 eV,
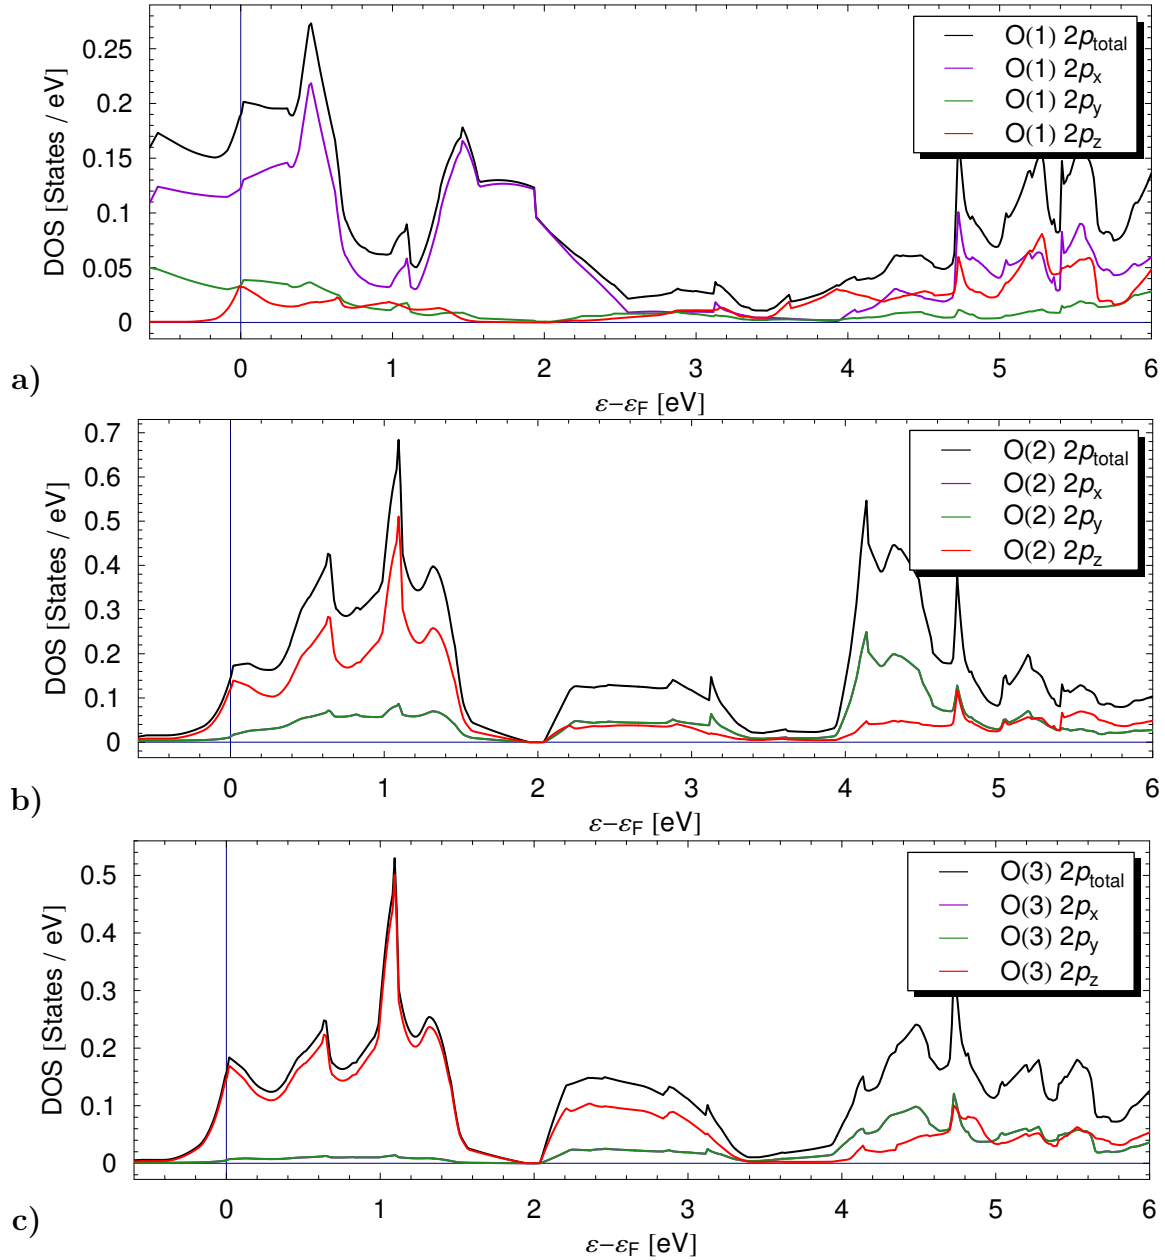
Figure 5.18: Densities of states at and above the Fermi level for different symmetry oxygen orbitals at each major oxygen site, from LDA band structure calculations. As shown in Figure 2.15, the **a)** O(1), **b)** O(2) and **c)** O(3) sites refer to the plane, TlO-layer and apical oxygens respectively. The O(2) and O(3) sites have four equivalent cation nearest neighbours, along the orthorhombic a and b directions with some c component, so their p_x and p_y spectra are identical.

comparable to that seen in YBCO for $\vec{E} \parallel \hat{a}$ for a similar change in doping (see the fourth figure in [106] or Figure A.14).

Centred about 1 eV above the Fermi level, LDA produces a band formed from the thallium oxide and barium oxide (apical oxygen) layers, with a predominantly $2p_z$ contribution from oxygen atoms (see Figure 5.15 or Figures 5.18b and 5.18c). This band also shows a small contribution from plane copper atoms' $3d_{3z^2-r^2}$ orbitals (Figure 5.17b). The c-axis peak at 531.5 eV matches this description, although its energy is higher than expected. Note that, since "zero" for XAS corresponds to the energy of a core orbital and can vary by crystallographic site, the energies of bands derived from different atomic sites cannot be directly compared quantitatively without also knowing the core level shifts. In this material, the LDA band structure calculations indicate similar core hole potentials for all three oxygen sites, so the shifts mentioned should be close to correct.

Around 3 eV above the Fermi energy, LDA shows another band shared by the TlO and BaO layers (see Figure 5.15). In this case, the cation orbitals are thallium and barium d orbitals, while the apical oxygens contribute mainly $2p_z$ weight and those in the TlO layer contribute roughly equal contributions from all three p orbitals (with slightly less coming from $2p_z$). This would suggest an oxygen peak slightly stronger in c-axis polarization than in-plane, but with little anisotropy. As mentioned, such a feature is observed 3 eV above the Fermi level, with no significant polarization dependence visible.

In Section 5.4.1, a weak copper L edge feature found 2.5-3 eV above the exciton peak was tentatively assigned to copper atoms in the thallium oxide layer. It was visible in c-axis polarization but in-plane data were inconclusive. That energy range corresponds with that of this final oxygen feature. The thallium orbitals expected to contribute to this band are of $d_{3z^2-r^2}$ symmetry, so copper would be expected to contribute similarly, and this peak should be strongest for c-axis polarized X-rays.

One further issue warrants mention — a thallium oxide electron pocket has long been predicted at the Γ point [15, 123, 124] (unlike almost any other band, this one is thought to show dispersion along the c-axis). Similar pockets have been predicted in many cuprates, including the bismuth-based systems, but have never been observed. The Fermi surfaces measured by ARPES and AMRO on Tl-2201 are notably devoid of electron pockets [14, 15].

Sahrakorpi and Bansil recently examined the thallium oxide hole pocket in detail, finding the TlO band to be 0.7 eV above the Fermi level for $p = 0.20$ [125] and 0.9 eV above for $p = 0.24$ [126], implying the pocket's absence at the dopings studied in this

work — using Presland's formula as reproduced in Equation 5.1, these hole dopings would correspond to T_cs of 82 K and 44 K respectively. As mentioned, the LDA calculations presented here show the band 1 eV above the Fermi level, and our data show it at 1.5 eV (modulo core level shifts), confirming the absence of the electron pocket.

Curiously, there was absolutely no evidence that this band's position varies with doping, when band structure calculations indicate that it should. One possible explanation for this sole discrepancy is that while Bansil's calculations distributed dopants randomly, the copper L edge absorption spectra suggest that the oxygen dopants and substituted copper may colocate.

The upper Hubbard band, which was seen in Figure 4.5b to weaken as YBCO's doping was increased, was not observed at any doping studied in Tl-2201. It could be absent entirely, or masked by other features and doping-independent. A doping-independent upper Hubbard band would be surprising, since even the Zhang-Rice singlet band still exhibits some limited doping dependence; the upper Hubbard band does, however, weaken substantially on overdoping, so such small relative changes to a small feature might simply not be visible. Given the trends observed in YBCO, this feature's outright absence might not be shocking — it was already weak for mildly overdoped $YBa_2Cu_3O_{6.99}$ — but its disappearance above optimal doping would nevertheless represent a fundamental change in the electronic structure. The only possible remnant of strong correlation effects observed is the Zhang-Rice singlet band, for which Zhang-Rice singlets may no longer be appropriate basis states — the Zhang-Rice singlet approximation requires that the holes be dilute, and the hole dopings studied here cannot be considered dilute. LDA shows oxygen-derived bands at the Fermi surface, so this feature may correspond to perfectly normal bands. The change in doping dependence may indicate a fundamental change in the bands' character.

What must be emphasized in this otherwise seemingly clearcut and unsurprising section is that the degree of agreement with band structure calculations, the near-lack of doping dependence, and the lack of features associated with strong correlations, most notably the upper Hubbard band, are unprecedented in the cuprates. This strongly suggests a fundamental change in the electronic structure and its behaviour from the underdoped to the overdoped side of the phase diagram. Determining whether this change is gradual or abrupt and related to a phase transition will require further work.

6. X-ray Emission Results on Tl-2201

While the unoccupied states probed by X-ray absorption (reported in Chapter 5) are of primary interest in the hole-doped cuprates, the occupied states can also show important signatures of strong correlation effects. Agreement with band structure calculations on one side of the Fermi energy does not imply agreement on the other, so the indications of comparative normalcy seen by XAS require confirmation by XES. While the lack of previous emission measurements on Tl-2201 makes this study interesting regardless, it will serve primarily to verify the conclusions from the previous chapter.

Emission work performed on the cuprates to date has mainly taken the form of resonant inelastic X-ray scattering (RIXS) [127], a near-resonant special case of XES. Among other studies, it has been performed on LSCO at several dopings [128, 129, 130, 131]; on $Sr_2CuO_2Cl_2$, Bi-2212 and $Nd_{1.85}Ce_{0.15}CuO_4$ [132]; on $HgBa_2CuO_{4+\delta}$ [133], and on a Tl-2212 film [134]. This work has focused primarily on optimal doping and undoped parent compounds, and on copper edges.

6.1   XES Results on Tl-2201

XES measurements were performed on a Tl-2201 crystal with a T_c of 60 K by procedures described in Section 3.2. Since the absorption spectra displayed an almost complete lack of doping-dependence, the benefits of performing XES measurements at the other dopings were not considered likely to justify the time required, and such measurements were not performed. The data were likewise not collected at multiple angles or for incident X-ray polarizations along different axes (in Tl-2201, these statements are equivalent).

Emission measurements are quite time-consuming and the data noisy, making quantitative interpretation of the resulting spectra difficult — a small change in the intensity or position of a peak is unlikely to be detected. Because of these limitations, interpretation of the resulting spectra is more difficult than for absorption — all analysis presented here is qualitative.

Once spectra have been recorded for an array of incident energies, they must be

normalized such that they may be compared against each other — thanks to enormous variations in the intensity of fluorescence as energy is varied (see any TFY spectrum), spectra have very different count rates for similar operating parameters. The most obvious way to normalize the spectra would involve dividing by the TFY signal at that incident energy, or by the raw incident intensity. The TFY channeltron was shut off, and I_0 was not recorded (in fact, some of these spectra take long enough to measure that the storage ring's beam current can change noticably over the course of a measurement, requiring an integrated I_0). This leaves several options for normalizing emission spectra — by peak height, by integral, or by background level. Of these three schemes, only peak height produced reliable results. The three normalization schemes are contrasted in Appendix B. Normalized by peak height, the spectra's features and their energies may be readily compared, but any information that may have been contained in the overall intensity of the spectrum is lost. In any case, the identification of peaks and their energies would not benefit from this overall intensity information.

Figure 6.1 displays a set of copper L edge XES spectra. Aside from the main peak, there may be a small shoulder around 927 eV, for scans up to 932.6 eV, and there is evidence of a peak on the high-energy side in the 934.5 and 935.6 eV spectra, and possibly also 933.6 eV. The latter feature disperses, suggesting that it might be an elastic peak — it was not possible to calibrate the energy scale on this plot, due to the absence of a clear elastic peak. If the energy axis is to be believed, however, 934–936 eV corresponds to a peak in the c-axis spectra (also present as a shoulder in the ab-plane spectra).

Figure 6.2 shows oxygen K edge XES spectra. Most of these scans (and all Cu L edge emission scans) used entrance and exit slit widths of 70 and 100 μm respectively, while spectra at 534.5 and 535.5 eV were collected with slits of 80 and 300 μm. One scan at the highest energy used slits of 80 and 450 μm to reduce the noise level, and one at 530 eV used slits of 50 μm. Wider slits allow more signal at the cost of energy resolution, but direct comparisons of the spectra at the lowest energy and of the spectra at the highest energy do not show any loss of information due to increased slit sizes. The oxygen-edge spectra are richer than those recorded on the copper edge.

Aside from the main peak at 529 eV, one clear feature in almost all spectra is a peak around 525 eV. An additional hump appears around 527 eV, possibly moving upward in energy from a start nearer 526 eV. The peak at 526 eV could also be a distinct feature — it appears in the 530 eV and 539 eV scans, and most other scans, especially at high energy, are consistent with a shoulder at that energy. The

Figure 6.1: Copper L edge XES spectra, normalized by peak height, for a $T_{\mathrm{c}} = 60$ K crystal. The a and c axis XAS spectra are reproduced in panel **a)** for convenience; vertical lines mark energies at which spectra were taken. The in-plane spectrum in panel **a)** has been corrected for c-axis weight as was done in Figure 5.9.

Figure 6.2: Oxygen K edge XES spectra, normalized by peak height, for a $T_c = 60$ K crystal. The in-plane and c axis XAS spectra are reproduced in panel **a)** for reference; vertical lines indicate energies at which spectra were taken.

spectra at incident energies of 530 eV and 531.5 eV show an elastic peak, which allows calibration of the horizontal axis (no shift was required).

6.2 Band Structure Calculations

As in Section 5.5, LDA band structure calculations have been used to determine the densities of states near the Fermi level from each site and, where relevant, orbital. In this chapter, the filled states immediately below the Fermi energy are of interest.

Figure 6.3: Densities of states at and below the Fermi level for the various atomic species and sites, from LDA band structure calculations.

Figure 6.3 shows the contribution to the density of states from each atomic species and crystallographic site in Tl-2201, and Figure 6.4 shows the weight arising from thallium, barium and copper orbitals. Figure 6.5 shows the thallium and copper contributions by orbital, and 6.6 divides the oxygen contribution by site and symmetry. These figures will serve primarily as a reference for a discussion on the features observed experimentally.

The copper L edge X-ray emission spectra, shown in Figure 6.1, display little conclusive evidence of features other than the main peak. For lower excitation energies, there may be a shoulder about 2 eV below the peak. The LDA band structure calculations show a copper band 1.2 eV below the Fermi level (composed of $3d_{xy}$, $3d_{xz}$ and $3d_{yz}$ orbitals), with another strong band 2.4 eV below (of $3d_{3z^2-r^2}$ character). Given that whether a shoulder is even visible is unclear, nothing can be concluded here.

The other feature visible is what appears at first to be an elastic peak. If this were the case, and the emission energies were miscalibrated, the excitation energy would be on the high-energy side of a peak identified as copper atoms on the thallium site,
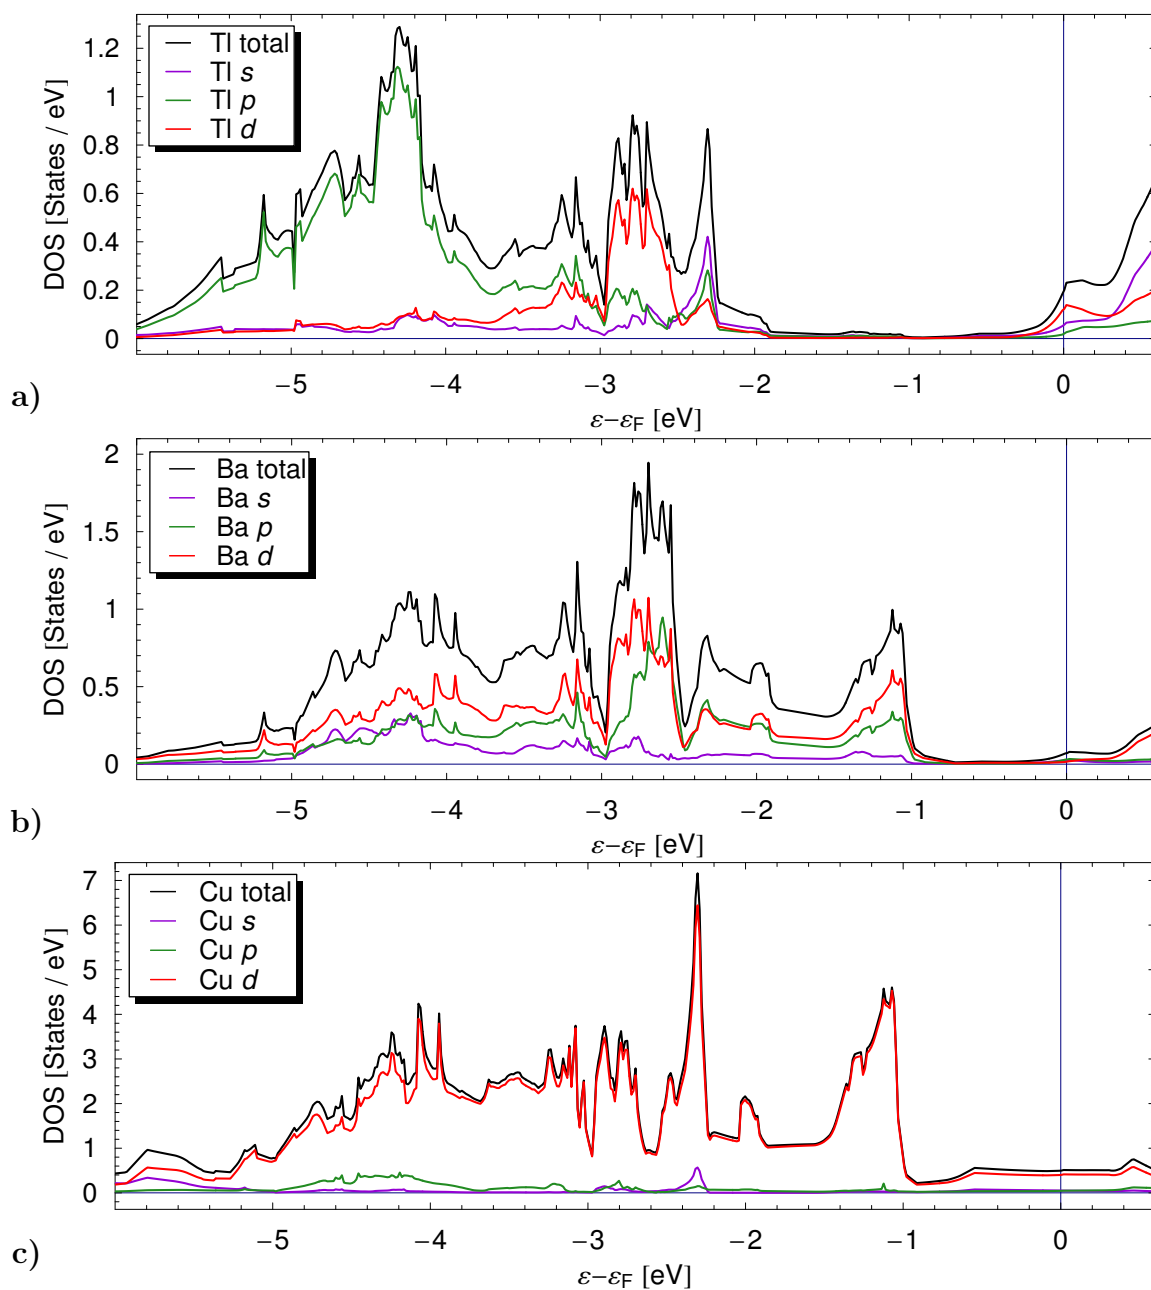
Figure 6.4: Densities of states at and below the Fermi level for **a)** thallium, **b)** barium and **c)** copper, from LDA band structure calculations.

Figure 6.5: Densities of states at and below the Fermi level for different symmetry a) thallium and b) copper orbitals, from LDA band structure calculations.

rather than the peak from ligand hole initial states. However, even if it's assumed that the lowest-energy traces are peaked at the incident energy, the apparent elastic peak is at least 1 eV too low to actually be an elastic peak (ruling out energy miscalibrations). The shift is consistent, however — the feature appears 3 eV below the incident energy on all three spectra.

Peaks that disperse with incident energy in emission spectra are generally identified with Raman processes, whereby a forbidden $d - d$ transition occurs as the X-ray impinges upon the sample [132]. Raman features can be dominant on copper L edges, as the initial excitation is $3d^9 \rightarrow 3d^{10}\underline{C}$, $\underline{C}$ referring to a core hole. The emitted X-ray has its energy reduced by the 3 eV required for the $d - d$ excitation. The copper orbitals at the Fermi level are $3d_{x^2-y^2}$, and the LDA spectra show a strong $3d_{3z^2-r^2}$ band 2.5-3 eV below the Fermi level, suggesting this to be the band observed.

Overdoped LSCO ($p = 0.27$) has been found [128] to exhibit weak remnants of a dispersive feature for energy transfers of 2.6 eV, which was attributed to $d - d$

excitations, and a similar feature in Tl-2212 2 eV below the incoming energy has been attributed to primarily transitions out of $3d_{xz}$ and $3d_{yz}$ states [134]. Which d orbitals contribute and at which energies depends on the details of the chemical environment.

The oxygen K edge emission spectra show many more features than the copper edge. Elastic peaks are observed for 530 and 531.5 eV, which correspond to the Zhang-Rice singlet band and the thallium oxide electron pocket respectively. An elastic peak suggests that the feature being excited into has excitonic character — this leads to an increased likelihood that the excited electron will fall back to annihilate its own core hole and release the same energy it absorbed. Unfortunately, data were not taken between these peaks or below the former, so it is not possible to determine whether these independently have exciton character or whether the exciton character in one is due to overlap with the other.

At 525 eV, 5 eV below the Fermi energy, a peak or hump is visible in almost every spectrum. Weak bands incorporating all three oxygen sites (indeed, every atom in the unit cell) are seen in the LDA calculations 4.3 eV below the Fermi level. It is possible to tentatively assign this peak to the weak band, but the LDA results are broad and weak, and closer to the Fermi surface than measured.

Around 527–527.5 eV, 2.5–3 eV below the Fermi level, a strong feature is visible when exciting at energies at or below the main edge, but not when exciting on the Zhang-Rice singlet band. This suggests that the excitations need to be in the thallium oxide layer for this band to fluoresce. In Figures 6.6b and 6.6c, strong bands may be seen in the thallium oxide and barium oxide layers 2.8 eV and 2.6 eV below the Fermi level respectively, using primarily $2p_x$ and $2p_y$ oxygen orbitals. Given that this is a shoulder on a strong peak, these energies are consistent, and this peak can be identified with the blocking layers.

The main peak itself is at 529 eV in emission, while the lowest peak in the absorption is at 530 eV. A possible explanation can be found in Figure 6.3 or 6.6 — there is no significant oxygen weight within 1 eV of the Fermi level, then a strong band appears between 1 eV and 1.3 eV.

Finally, there appears to be a feature at 526 eV in scans taken with incident X-ray energies of 530 eV, 539 eV, and possibly 531.5 eV. An XAS feature at 539 eV, just above the main edge, was visibile in the spectra but was not discussed, high-precision data were not taken on it, and its identity is not clear. The Zhang-Rice singlet band at 530 eV is associated with the CuO_2 planes, and may have some weight remaining at 531.5 eV. This feature, then, ought to be associated with the planes. However,
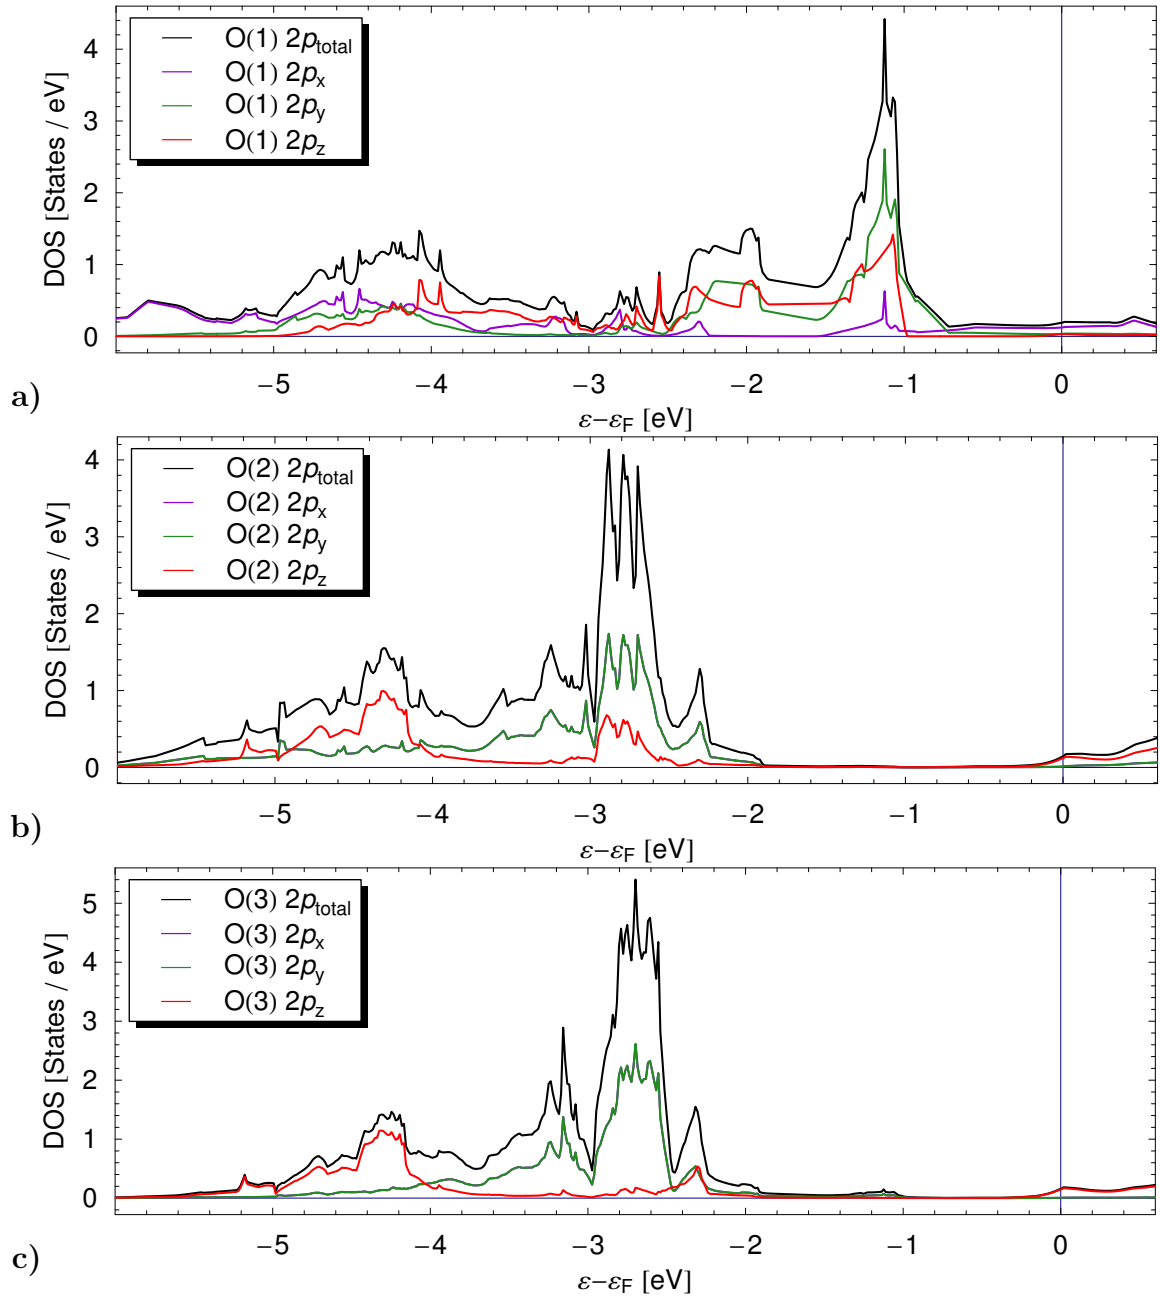
Figure 6.6: Densities of states at and below the Fermi level for different symmetry oxygen orbitals at each major oxygen site, from LDA band structure calculations: **a)** O(1), **b)** O(2) and **c)** O(3) sites refer to the plane, TlO-layer and apical oxygens respectively.

no obvious plane-related feature is observed 3.5–4 eV below the Fermi level. This feature remains unexplained.

As for the absorption spectra, the oxygen emission spectra appear remarkably unremarkable. The oxygen bands may be somewhat closer to the Fermi level than expected, but all but one possible feature are consistent with band structure (although the quality of the data do not permit the level of confidence available on the absorption data). No lower Hubbard band is observed. In the context of the high-T_c cuprates, this relative normalcy is almost without precedent.

The copper L edge shows evidence of $d-d$ excitations, most likely from a $3d_{3z^2-r^2}$ band expected to be 2.4-3 eV below the Fermi energy. A possible shoulder on the low-energy side of the main edge could not be identified.

That Tl-2201's emission spectra, like its absorption spectra, can be explained by band structure calculations and that no lower Hubbard band is observed again suggest that a microscopic understanding of Tl-2201's electronic structure may be within our grasp. Limitations of the data do not permit conclusions as strong as those in Chapter 5, but the data do corroborate its findings. The lack of signatures of strong correlation effects in the electronic structure of overdoped Tl-2201 suggests a new approach to an understanding of the cuprate superconductors — research on the overdoped regime might solve the puzzle that two decades' work on the underdoped regime have been unable to.

7. Conclusion

High-quality platelet single crystals of $Tl_2Ba_2CuO_{6\pm\delta}$ have been grown using a novel time-varying encapsulation scheme, minimizing the loss of thallium oxide that has plagued other attempts to grow crystals of this material. With thallium loss controlled, other growth parameters were then optimized. The composition settled upon was Tl:Ba:Cu = 2.2:2:1.8, and a slow cooling from the liquidus toward the eutectic of -3°C/h was used. The melt was decanted from the crystals, a step not previously possible, and cation substitution was allowed to homogenize with a high-temperature post-growth anneal.

Annealing schemes were developed to produce transition temperatures from 5 to 85 K without damaging the crystals, necessitating control not only over the oxygen partial pressure and annealing temperature, but also over the thallium partial pressure for some dopings.

The resulting crystals were highly homogeneous, had high crystalline perfection, and exhibited what are thought to be particularly low levels of cation substitution. The superconducting transitions normally had narrow widths ΔT_c of 1-2 K, while X-ray rocking curve widths were typically 0.025-0.050°, comparable to high-quality $YBa_2Cu_3O_{6+\delta}$ crystals grown in YSZ crucibles. By both metrics, these crystals still fall short of $BaZrO_3$-grown YBCO, but those are the highest-quality high-T_c crystals in existence, so it would be unrealistic to hold these to that high standard.

To characterize the cation substitution level, two techniques were applied — electron probe microanalysis (EPMA) and an X-ray diffraction structure refinement. EPMA found the cation substitution level to vary as the cation ratio in the melt varied, and determined the crystals' composition as grown from the optimum melt composition to be $Tl_{1.920(2)}Ba_{1.96(2)}Cu_{1.080(2)}O_{6+\delta}$. The X-ray diffraction study found the composition for a crystal with a T_c of 75 K to be $Tl_{1.914(14)}Ba_2Cu_{1.086(14)}O_{6.07(5)}$, in excellent agreement with the EPMA result.

The X-ray diffraction study found a crystal structure and lattice parameter behaviour that closely resembled those published by Wagner [47] for similar cation defect levels. Orthorhombicity matching Wagner's results was also observed. These Tl-2201 crystals become orthorhombic on overdoping, ceasing to be tetragonal for

T_cs between 75 K and 55 K. There can no longer be doubt that crystals may be orthorhombic, that the cation-substituted phase may be orthorhombic, and that the orthorhombic phase superconducts. All three results differ from the picture provided by the bulk of the previous literature.

In addition, the X-ray diffraction study produced evidence for an unexpected commensurate superlattice modulation, most likely in the Tl_2O_2 and BaO blocking layers. It was not possible to determine the exact nature of this modulation, so other techniques such as neutron diffraction or resonant X-ray diffraction will need to be applied. Insofar as the material is thought to cleave between the TlO layers, the nature of this distortion would best be determined quickly, particularly if scanning tunnelling microscopy or spectroscopy are to be applied to it — modulations have been found in other materials which have drawn a great deal of attention over the past few years, and distinguishing this modulation from any modulation of charge may be important. The temperature at which the X-ray diffraction data were taken makes it unlikely that the modulation is related to a charge density wave.

The Tl-2201 crystals thus grown and characterized were studied by X-ray absorption and emission spectroscopies (XAS/XES), to investigate their electronic structure. The features observed on the absorption side correspond very closely with those expected from LDA band structure calculations.

The oxygen K edge shows slight increases with overdoping in a peak associated with the Zhang-Rice singlet, and a thallium oxide-derived band that exhibits less doping dependence than expected. The increase in Zhang-Rice singlet weight is significantly less than would be expected given the changes in hole doping, but this should be unsurprising — Zhang-Rice singlets are postulated as a basis state only where the holes added to the CuO_2 plane are dilute, and one extra hole for every four plane copper atoms is manifestly not dilute. The lack of dispersion in the thallium oxide band as doping is increased (and the lack of an electron pocket) is more surprising, but no such electron pocket has ever been observed experimentally, which implies that this band is not fully understood.

The copper L edge exhibits doping-dependence, but in a feature primarily attributed to copper atoms on the thallium site. Barely-discernible changes to a shoulder on the high-energy side of the edge in the in-plane spectrum may be attributable to hole doping.

In the emission spectra, features on both the oxygen K and copper L edges can be successfully explained by band structure calculations, although some are displaced from their expected energies. The quality of data available on emission spectra is

lower, so less confidence in these results is possible than on the absorption side.

In the absorption and emission spectra respectively, neither an upper nor a lower Hubbard band is observed, although the former could possibly be weak, doping-independent and swamped by other features, while the latter could simply be weak. Although band structure calculations would not predict their existence, correlation effects would, and the absence of Hubbard bands here would be significant, given their ubiquitous presence in the underdoped regime. The decline observed in YBCO in the upper Hubbard band's intensity suggests its elimination somewhere above optimal doping, and the Tl-2201 results do not detect it at any doping studied, suggesting a rapid decline in the importance of Mott-Hubbard physics as holes are added. A lack of doping dependence would be similarly significant, since it would indicate that they have become irrelevant for holes near the Fermi level. The Hubbard bands' disappearance or abrupt irrelevance not far above optimal doping coincides with the location of a commonly postulated quantum phase transition and the doping range where the pseudogap collapses.

In short, Tl-2201's electronic structure is characterized by unprecedented (in the high-T_c cuprates), breathtaking normalcy. The degree to which the observed electronic structure can be predicted by LDA band structure calculations is unanticipated. The possible existence of a region in the phase diagram where our usual understanding of materials applies is an exciting prospect. These fundamental changes in the electronic structure may lend credence to the many theories predicting a quantum phase transition near or just above optimal doping, and the availability of high-quality crystals in this key region should trigger a renewed effort aimed at sorting out the nature of this transition.

To complete the work described in this thesis, aside from the superlattice study mentioned above, X-ray emission data should be collected with a finer point spacing, to ensure that all features are being observed and reduce ambiguity. The cation substitution could possibly be further reduced, particularly if a scheme were developed whereby the furnace could be pressurized to prevent volatile monovalent thallium oxide from forming. However, most of the work outstanding is for our collaborators.

Given that Tl-2201 is known or believed to be clean enough for high-sensitivity bulk transport probes such as microwave susceptibility, thermal conductivity, AMRO, and de Haas-van Alphen, yet cleaves well enough for ARPES and potentially STM and STS, Tl-2201 is already an excellent candidate for finally joining the sensitive bulk techniques with the surface-sensitive electronic structure probes. That Tl-2201 is also apparently free from many complicating correlated electron effects makes it

a particularly good candidate for thorough study. The availability of high-quality single crystals suitable for these techniques should be taken full advantage of.

Tl-2201 and its surprising simplicity may well offer a key toehold of understanding in an otherwise confusing landscape of new and unexplained phases.

References

[1] H. Kamerlingh-Onnes, Commun. Phys. Lab. Univ. Leiden **12**, 120 (1911).

[2] W. Meissner and R. Ochsenfeld, Naturwiss. **21**, 787 (1933), doi:10.1007/BF01504252.

[3] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, Phys. Rev. **108**, 1175 (1957), available at aps.org.

[4] J. G. Bednorz and K. A. Müller, Z. Physik B **64**, 189 (1986), doi:10.1007/BF01303701.

[5] G. F. Sun, K. W. Wong, B. R. Xu, Y. Lin, and D. F. Lu, Phys. Lett. A **192**, 122 (1994), doi:10.1016/0375-9601(94)91026-X.

[6] J. Orenstein and A. J. Millis, Science **288**, 468 (2000), doi:10.1126/science.288.5465.486.

[7] Nature Physics **2**, 133 (2006), doi:10.1038/nphys258, March 2006, special issue on high-temperature superconductors.

[8] J. L. Tallon and J. W. Loram, Physica C **349**, 53 (2001), doi:10.1016/S0921-4543(00)01524-0.

[9] W. N. Hardy, D. A. Bonn, D. C. Morgan, R. Liang, and K. Zhang, Phys. Rev. Lett. **70**, 3999 (1993), available at aps.org.

[10] D. M. Broun *et al.*, Phys. Rev. B **56**, R11443 (1997), available at aps.org.

[11] D. J. Scalapino, Physics Reports **250**, 329 (1995), doi:10.1016/0370-1573(94)00086-I.

[12] C. C. Tsuei and J. R. Kirtley, Rev. Mod. Phys. **72**, 969 (2000), available at aps.org.

[13] S. Özcan *et al.*, Phys. Rev. B **73**, 064506 (2006), available at aps.org.

[14] N. E. Hussey, M. Abdel-Jawad, A. Carrington, A. P. Mackenzie, and L. Balicas, Nature **425**, 814 (2003), doi:10.1038/nature01981.

[15] M. Platé *et al.*, Phys. Rev. Lett. **95**, 077001 (2005), available at aps.org.

[16] D. C. Peets *et al.*, New J. Phys. **9**, 28 (2007), available at iop.org.

[17] N. Doiron-Leyraud *et al.*, Nature **447**, 565 (2007), doi:10.1038/nature05872.

[18] A. F. Bangura *et al.*, Phys. Rev. Lett. **100**, 047004 (2008), available at aps.org.

[19] D. LeBoeuf *et al.*, Nature **450**, 533 (2007), doi:10.1038/nature06332.

[20] C. Jaudet *et al.*, Phys. Rev. Lett. **100**, 187005 (2008), available at aps.org.

[21] M. R. Norman *et al.*, Nature **392**, 157 (1998), doi:10.1038/32366.

[22] K. M. Shen *et al.*, Science **307**, 901 (2005), doi:10.1126/science.1103627.

[23] A. Kanigel *et al.*, Nature Physics **2**, 447 (2006), doi:10.1038/nphys334.

[24] C. Proust, E. Boaknin, R. W. Hill, L. Taillefer, and A. P. Mackenzie, Phys. Rev. Lett. **89**, 147003 (2002), available at aps.org.

[25] A.-W. D. Tyler, *An Investigation into the Magnetotransport Properties of Layered Superconducting Perovskites*, Ph.D. thesis, University of Cambridge, 1998.

[26] A. P. Mackenzie, S. R. Julian, D. C. Sinclair, and C. T. Lin, Phys. Rev. B **53**, 5848 (1996), available at aps.org.

[27] S. Nakamae *et al.*, Phys. Rev. B **68**, 100502 (2003), available at aps.org.

[28] T. Hanaguri *et al.*, Nature **430**, 1001 (2004), doi:10.1038/nature02861.

[29] R. W. Hill *et al.*, Phys. Rev. Lett. **92**, 027001 (2004), available at aps.org.

[30] A. Hosseini, S. Kamal, D. A. Bonn, R. Liang, and W. N. Hardy, Phys. Rev. Lett. **81**, 1298 (1998), available at aps.org.

[31] D. H. Lu *et al.*, Phys. Rev. Lett. **86**, 4370 (2001), available at aps.org.

[32] D. J. Derro *et al.*, Phys. Rev. Lett. **88**, 097002 (2002), available at aps.org.

[33] K. McElroy *et al.*, Nature **422**, 592 (2003), doi:10.1038/nature01496.

[34] C. Howald, H. Eisaki, N. Kaneko, and A. Kapitulnik, Proc. Natl Acad. Sci. USA **100**, 9705 (2003), doi:10.1073/pnas.1233768100.

[35] M. Vershinin *et al.*, Science **303**, 1995 (2004), doi:10.1126/science.1093384.

[36] A. Damascelli, Z. Hussain, and Z.-X. Shen, Rev. Mod. Phys. **75**, 473 (2003), available at aps.org.

[37] M. Chaio *et al.*, Phys. Rev. B **62**, 3554 (2000), available at aps.org.

[38] H. Eisaki *et al.*, Phys. Rev. B **69**, 064512 (2004), available at aps.org, and references therein.

[39] I. Chong *et al.*, Science **276**, 770 (1997), doi:10.1126/science.276.5313.770.

[40] A. Erb, E. Walker, and R. Flükiger, Physica C **245**, 245 (1995), doi:10.1016/0921-4543(95)00123-9.

[41] R. Liang, D. A. Bonn, and W. N. Hardy, Physica C **304**, 105 (1998), doi:10.1016/S0921-4534(98)00275-5.

[42] P. Manca *et al.*, Phys. Rev. B **63**, 134512 (2001), available at aps.org.

[43] R. Liang, D. A. Bonn, and W. N. Hardy, Physica C **336**, 57 (2000), doi:10.1016/S0921-4543(00)00091-5.

[44] M. A. Hossain *et al.*, *In situ* doping control of the surface of high-temperature superconductors, 2008, doi:10.1038/nphys998.

[45] R. Liang, D. A. Bonn, and W. N. Hardy, Phys. Rev. B **73**, 180505(R) (2006), available at aps.org.

[46] M. R. Presland, J. L. Tallon, R. G. Buckley, R. S. Liu, and N. E. Flower, Physica C **176**, 95 (1991), doi:10.1016/0921-4543(91)90700-9.

[47] J. L. Wagner *et al.*, Physica C **277**, 170 (1997), doi:10.1016/S0921-4543(97)00062-2.

[48] Y. Shimakawa *et al.*, Phys. Rev. B **42**, 10165 (1990), available at aps.org.

[49] D. G. Hawthorn *et al.*, Phys. Rev. B **75**, 104518 (2007), available at aps.org.

[50] A. P. Mackenzie *et al.*, Phys. Rev. Lett. **76**, 3786 (1996), available at aps.org.

[51] C. Bergemann, S. R. Julian, A. P. Mackenzie, S. NishiZaki, and Y. Maeno, Phys. Rev. Lett. **84**, 2662 (2000), available at aps.org.

[52] A. Damascelli *et al.*, Phys. Rev. Lett. **85**, 5194 (2000), available at aps.org.

[53] K. M. Shen *et al.*, Phys. Rev. B **64**, 180502 (2001), available at aps.org.

[54] N. J. C. Ingle *et al.*, Phys. Rev. B **72**, 205114 (2005), available at aps.org.

[55] J. L. Jorda *et al.*, Physica C **205**, 177 (1993), doi:10.1016/0921-4543(93)90185-S.

[56] M. P. Siegal, E. L. Venturini, B. Morosin, and T. L. Aselage, J. Mater. Res. **12**, 2825 (1997), doi:10.1557/JMR.1997.0378.

[57] F. Lichtenberg, A. Catana, J. Mannhart, and D. G. Schlom, Appl. Phys. Lett. **60**, 1138 (1992), available at aip.org.

[58] Y. Maeno *et al.*, Nature **372**, 532 (1994), doi:10.1038/372532a0.

[59] T. Manako, Y. Kubo, and Y. Shimakawa, Phys. Rev. B **46**, 11019 (1992), available at aps.org.

[60] H. He *et al.*, Science **295**, 1045 (2002), doi:10.1126/science.1067877.

[61] B. Raveau, C. Michel, B. Mercey, J. F. Hamet, and M. Hervieu, J. Alloy Compd. **229**, 134 (1995), doi:10.1016/0925-8388(95)01679-1, and references therein.

[62] M. v. Zimmermann *et al.*, Phys. Rev. B **68**, 104515 (2003), available at aps.org.

[63] N. H. Andersen *et al.*, Physica C **317–318**, 259 (1999), doi:10.1016/S0921-4543(99)00066-0.

[64] B. O. Wells *et al.*, Z. Phys. B **100**, 535 (1996), doi:10.1007/s002570050158.

[65] B. O. Wells *et al.*, Science **277**, 1067 (1997), doi:10.1126/science.277.5329.1067.

[66] Y. Shimakawa, Physica C **204**, 247 (1993), doi:10.1016/0921-4543(93)91006-H.

[67] M. Tinkham, *Introduction to Superconductivity*, 2nd ed. (Dover Publications, New York, 2004).

[68] D. Achkir, M. Poirier, D. A. Bonn, R. Liang, and W. N. Hardy, Phys. Rev. B **48**, 13184 (1993), available at aps.org.

[69] K. Zhang, D. A. Bonn, R. Liang, D. J. Baar, and W. N. Hardy, Appl. Phys. Lett. **62**, 3019 (1993), available at aip.org.

[70] D. A. Bonn *et al.*, Phys. Rev. B **50**, 4051 (1994), available at aps.org.

[71] M. Grüninger, *Of Spin and Charge in the Cuprates*, Ph.D. thesis, Rijksuniversiteit Groningen, 1999, available from the University of Groningen.

[72] W. Brenig, Physics Reports **251**, 153 (1995), doi:10.1016/0370-1573(94)00072-B.

[73] P. W. Anderson, Phys. Rev. **115**, 2 (1959), available at aps.org.

[74] P. W. Anderson, Science **235**, 1196 (1987), doi:10.1126/science.235.4793.1196.

[75] M. Ogata and H. Fukuyama, Rep. Prog. Phys. **71**, 036501 (2008), doi:10.1088/0034-4885/71/3/036501.

[76] E. Dagotto, Rev. Mod. Phys. **66**, 763 (1994), available at aps.org.

[77] F. C. Zhang and T. M. Rice, Phys. Rev. B **37**, 3759 (1988), available at aps.org.

[78] J. Zaanen, G. A. Sawatzky, and J. W. Allen, Phys. Rev. Lett. **55**, 418 (1985), available at aps.org.

[79] R. S. Liu *et al.*, Physica C **198**, 203 (1992), doi:10.1016/0921-4543(92)90192-F.

[80] N. N. Kolesnikov *et al.*, Physica C **195**, 219 (1992), doi:10.1016/0921-4543(92)90343-B.

[81] N. N. Kolesnikov *et al.*, Physica C **242**, 385 (1995), doi:10.1016/0921-4543(94)02420-0.

[82] O. M. Vyaselev, N. N. Kolesnikov, M. P. Kulakov, and I. F. Schegolev, Physica C **199**, 50 (1992), doi:10.1016/0921-4543(92)90539-O.

[83] H. Takei, F. Sakai, M. Hasegawa, S. Nakajima, and M. Kikuchi, Jpn. J. Appl. Phys. **32**, L1403 (1993), doi:10.1143/JJAP.32.L1403.

[84] M. Hasegawa, H. Takei, K. Izawa, and Y. Matsuda, J. Low Temp. Phys. **117**, 741 (1999), doi:10.1023/A:1022549329077.

[85] M. Hasegawa, H. Takei, K. Izawa, and Y. Matsuda, J. Cryst. Growth **229**, 401 (2001), doi:10.1016/S0022-0248(01)01189-7.

[86] M. Hasegawa, Y. Matsushita, Y. Iye, and H. Takei, Physica C **231**, 161 (1994), doi:10.1016/0921-4543(94)90156-2.

[87] D. C. Peets, Growth of $Tl_2Ba_2CuO_{6+\delta}$ single crystals by the self-flux method, M. Sc. thesis, University of British Columbia, 2002.

[88] D. C. Peets, R. Liang, M. Raudsepp, D. A. Bonn, and W. N. Hardy, Physica C **460–462**, 1335 (2007), doi:10.1016/j.physc.2007.04.014.

[89] C. C. Torardi *et al.*, Phys. Rev. B **38**, 225 (1988), available at aps.org.

[90] C. Opagiste *et al.*, Physica C **213**, 17 (1993), doi:10.1016/0921-4543(93)90753-D.

[91] D. R. Lide, *Handbook of Chemistry and Physics*, 88$^{\text{th}}$ ed. (CRC Press, 2007), http://www.hbcpnetbase.com.

[92] Chemical-Elements.info, 2007, http://tl-thallium.info/energy.html, accessed 2008-08-08.

[93] Q. Design, *Magnetic moment measurement using Quantum Design's magnetic property measurement system: Application Note #1* (Quantum Design, San Diego, 1989), .

[94] D. C. Peets *et al.*, J. Supercond. **15**, 531 (2002), doi:10.1023/A:1021230704882.

[95] J. L. Pouchou and F. Pichoir, "PAP" $\varphi(\rho z)$ procedure for improved quantitative microanalysis, in *Microbeam Analysis 1985*, edited by J. T. Armstrong, pp. 104–106, San Francisco Press, 1985.

[96] Y. Shimakawa, Y. Kubo, T. Manako, and H. Igarashi, Phys. Rev. B **40**, 11400 (1989), available at aps.org.

[97] M. A. G. Aranda, D. C. Sinclair, J. P. Attfield, and A. P. Mackenzie, Phys. Rev. B **51**, 12747 (1995), available at aps.org.

[98] G. M. Sheldrick, SHELXL-97 — A program for crystal structure refinement, 1997, shelx.uni-ac.gwdg.de/SHELX/, Release 97-2.

[99] G. M. Sheldrick, Acta Cryst. A **64**, 112, doi:10.1107/S0108767307043930.

[100] A. W. Hewat *et al.*, Physica C **156**, 369 (1988), doi:10.1016/0921-4543(88)90760-5.

[101] R. Beyers *et al.*, Appl. Phys. Lett. **53**, 432 (1988), available at aip.org.

[102] F. M. F. deGroot, J. Electron. Spectrosc. Relat. Phenom. **67**, 529 (1994), doi:10.1016/0368-2048(93)02041-J.

[103] E. Pellegrin *et al.*, Phys. Rev. B **47**, 3354 (1993), available at aps.org.

[104] C. T. Chen *et al.*, Phys. Rev. Lett. **68**, 2543 (1992), available at aps.org.

[105] S. M. Heald, J. M. Tranquada, A. R. Moodenbaugh, and Y. Xu, Phys. Rev. B **38**, 761 (1988), available at aps.org.

[106] N. Nücker *et al.*, Phys. Rev. B **51**, 8529 (1995), available at aps.org.

[107] J. Guo, D. E. Ellis, E. E. Alp, and G. L. Goodman, Phys. Rev. B **42**, 251 (1990), available at aps.org.

[108] A. Krol *et al.*, Phys. Rev. B **45**, 2581 (1992), available at aps.org.

[109] E. Pellegrin *et al.*, Phys. Rev. B **48**, 10520 (1993), available at aps.org.

[110] H. Romberg *et al.*, Phys. Rev. B **41**, 2609 (1990), available at aps.org.

[111] J. J. Jia *et al.*, Rev. Sci. Instrum. **66**, 1394 (1995), doi:10.1063/1.1145985.

[112] A. C. Thompson *et al.*, *X-Ray Data Booklet*, 2nd ed. (Lawrence Berkeley National Laboratory, 2001), xdb.lbl.gov.

[113] J. D. Denlinger, private communication.

[114] A. Kotani and S. Shin, Rev. Mod. Phys. **73**, 203 (2001), available at aps.org.

[115] J. S. Bobowski *et al.*, Oxygen chain disorder as the weak scattering source in $YBa_2Cu_3O_{6.50}$, 2006, arxiv:cond-mat/0612344.

[116] J. S. Bobowski *et al.*, Oxygen chain disorder as a weak-limit scatterer in $YBa_2Cu_3O_{6.50}$, 2007, doi:10.1016/j.physc.2007.03.272.

[117] N. Nücker *et al.*, Synthetic Metals **71**, 1563 (1995), doi:10.1016/0379-6779(94)02951-T.

[118] F. deGroot, Coord. Chem. Rev. **249**, 31 (2005), doi:10.1016/j.ccr.2004.03.018.

[119] Y. J. Uemura *et al.*, Phys. Rev. Lett. **62**, 2317 (1989), available at aps.org.

[120] D. G. Hawthorn, I. S. Elfimov, and G. A. Sawatzky, unpublished.

[121] A. Bianconi *et al.*, Solid State Commun. **63**, 1009 (1987), doi:10.1016/0038-1098(87)90650-8.

[122] D. D. Sarma *et al.*, Phys. Rev. B **37**, 9784 (1988), available at aps.org.

[123] D. J. Singh and W. E. Pickett, Physica C **203**, 193 (1992), doi:10.1016/0921-4543(92)90526-I.

[124] D. R. Hamann and L. F. Mattheiss, Phys. Rev. B **38**, 5138 (1988), available at aps.org.

[125] A. Bansil *et al.*, Physica C **460–462**, 222 (2007), doi:10.1016/j.physc.2007.03.281.

[126] S. Sahrakorpi, H. Lin, R. S. Markiewicz, and A. Bansil, Physica C **460–462**, 428 (2007), doi:10.1016/j.physc.2007.03.109.

[127] A. Kotani, Eur. Phys. J. B **47**, 3 (2005), doi:10.1140/epjb/e2005-00303-4.

[128] G. Ghiringhelli, N. B. Brookes, C. Dallera, A. Tagliaferri, and L. Braicovich, Phys. Rev. B **76**, 085116 (2007), available at aps.org.

[129] P. G. Steeneken *et al.*, Phys. Rev. Lett. **90**, 247005 (2003), available at aps.org.

[130] Y.-J. Kim *et al.*, Phys. Rev. B **70**, 094524 (2004), available at aps.org.

[131] Y.-J. Kim *et al.*, Phys. Rev. B **70**, 205128 (2004), available at aps.org.

[132] G. Ghiringhelli *et al.*, Phys. Rev. Lett. **92**, 117406 (2004), available at aps.org.

[133] L. Lu *et al.*, Phys. Rev. Lett. **95**, 217003 (2005), available at aps.org.

[134] W. A. Little *et al.*, Physica C **460–462**, 40 (2007), doi:10.1016/j.physc.2007.03.031.

A. Full XAS Results on YBCO

This appendix covers results on YBCO that were not included in Chapter 4. The results are not discussed to any significant degree; they're presented in graphical form.

A.1 XAS Results on YBa$_2$Cu$_3$O$_{6.0}$ Crystals

The copper L edges as measured on a YBa$_2$Cu$_3$O$_{6.0}$ crystal are shown in Figure A.1. The $L\alpha$ and $L\beta$ edges should, in principle, have identical features, but the former's higher intensity makes some of its features more clearly resolved. Similarly, oxygen K edge spectra for YBa$_2$Cu$_3$O$_{6.0}$ are shown in Figure A.2. This doping was not included in the body of the thesis, because YBCO was used primarily as context for Tl-2201.

A.2 XAS Results on YBa$_2$Cu$_3$O$_{6.99}$ Crystals

The copper L edge for YBa$_2$Cu$_3$O$_{6.99}$ is shown in Figure A.3 and the oxygen edge in Figure A.4. Figures A.3b and A.4b also appeared earlier in Figures 4.3 and 4.4 respectively.

A.3 XAS Results on YBa$_2$Cu$_3$O$_{6.35}$ Crystals

The Cu L and O K edges for oxygen-disordered YBa$_2$Cu$_3$O$_{6.35}$ are shown in Figures A.5 and A.6 respectively. As mentioned earlier in this section, after these crystals were measured fully disordered, they were allowed to order at room temperature for several hours, then remeasured.

The YBa$_2$Cu$_3$O$_{6.35}$ crystals were allowed to order at room temperature for three hours, then a further seventeen hours, to look for changes in their electronic structure. The results on both the copper L and oxygen K edges are shown in Figure A.7. Changes in T_c are initially rapid under such conditions, having a time constant on the order of a day. However, the slope of T_c with doping (dT_c/dp) is thought to be quite steep in this doping range. While changes appear after the first three hours, data after the full twenty hour anneal more closely resemble the initial, unannealed data
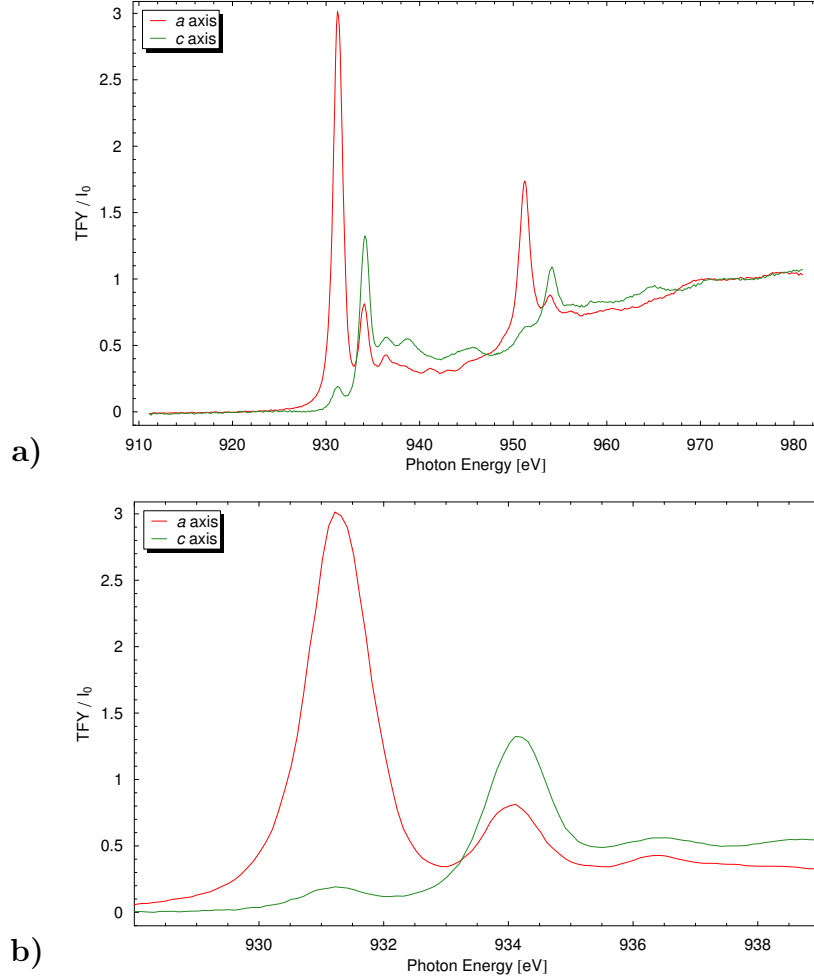
Figure A.1: Copper L edge for $YBa_2Cu_3O_{6.0}$ crystals: **a)** Entire spectrum, and **b)** detail of the Cu $L\alpha$ peak. Since YBCO is tetragonal at this doping, the b axis is identical to the a axis.

than data from the three-hour anneal. The evidence here for any electronic structure changes within the twenty-hour window is not strong, and the experiment has not been repeated, so it is not possible to conclude that any changes are occurring. Given that the characteristic time for this process is on the order of a day, this result may not be surprising. Further, annealing these crystals at room temperature under pressure produces better ordering, ordering which relaxes at room temperature to reduce T_c. These pressure-induced improvements to ordering were not fully investigated before the XAS data were taken. Fully disordering this more-ordered state could be done in under an hour at the beamline, likely on the cryostat itself, and would have higher odds of producing a noticable difference in the spectra. This represents a possible
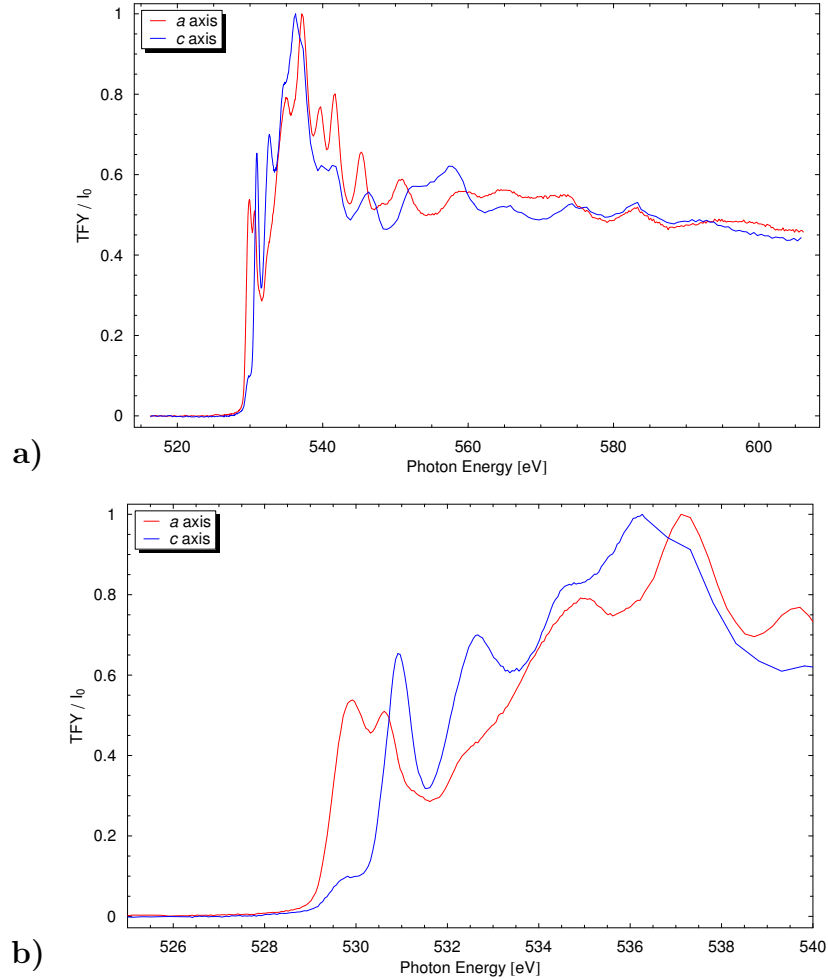
Figure A.2: Oxygen K edge for $YBa_2Cu_3O_{6.0}$: **a)** Entire spectrum, and **b)** detail of the O K edge and prepeaks. Recall that the b axis is identical to the a axis.

future phase of the research.

A.4 Dependence on Oxygen Chain Order

Figures A.8 and A.9 show the Cu L and O K edges respectively for Ortho-II ordered and Ortho-I (disordered) $YBa_2Cu_3O_{6.5}$. The difference is plotted more clearly in Figure A.10. Selected plots were presented in Section 4.3. Figures A.11 and A.12 show the absolute and relative changes on $YBa_2Cu_3O_{6.5}$ respectively between the Ortho-II and disordered crystals. In these two figures, Δ refers to the excess spectral weight due to Ortho-II order, in absolute or relative terms as appropriate.
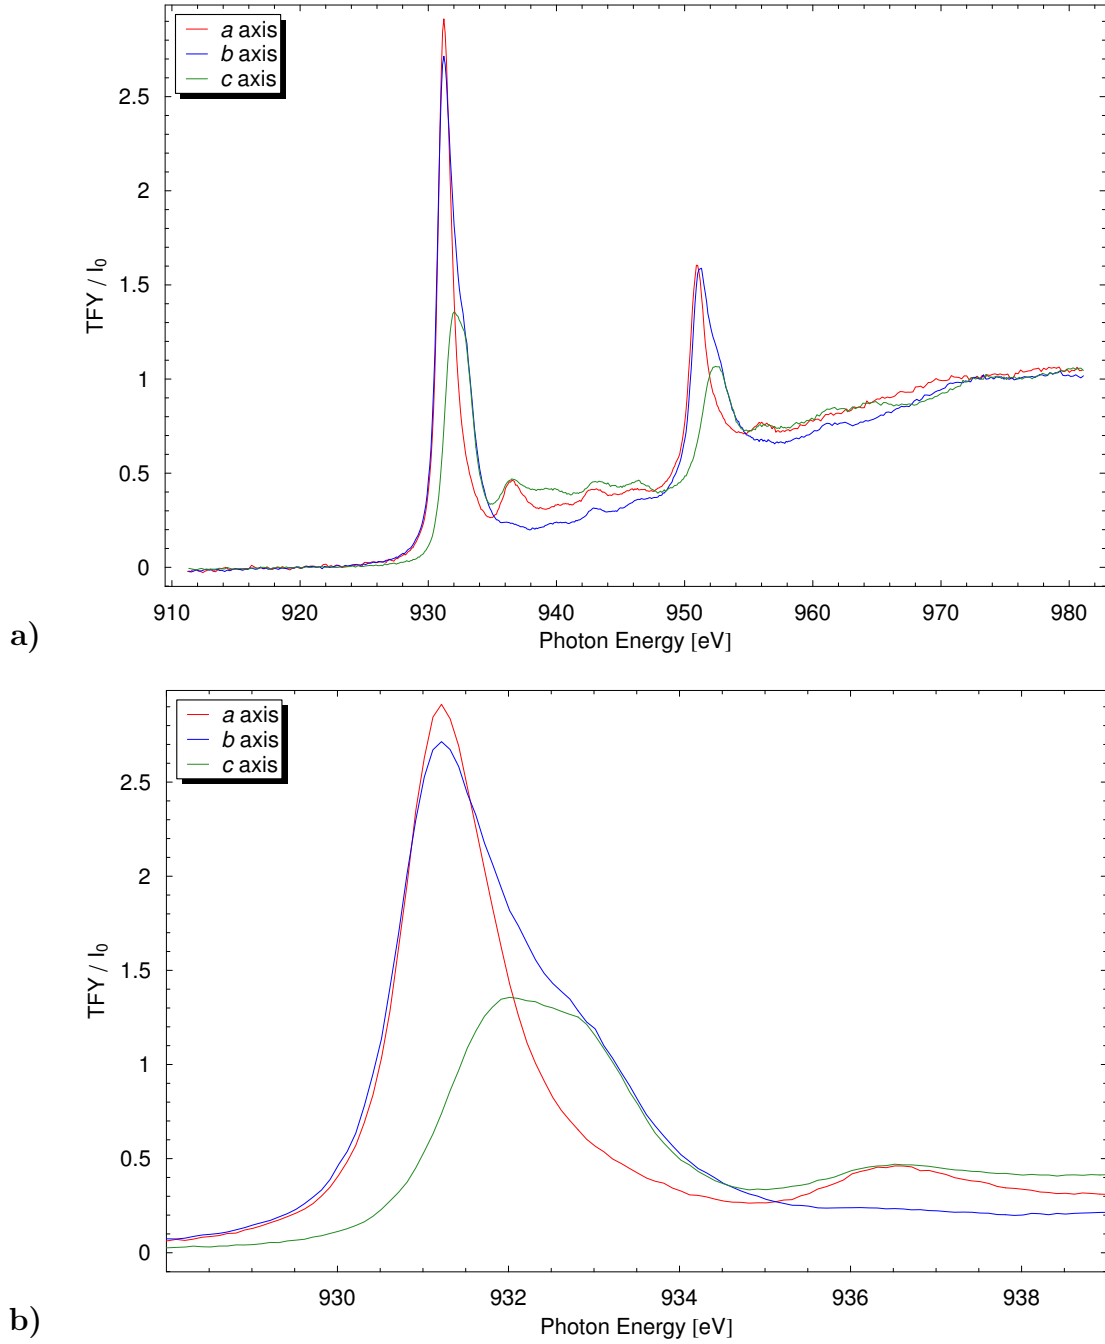
Figure A.3: Copper L edge for $YBa_2Cu_3O_{6.99}$ crystals as a function of the incident X-ray beam's polarization. **a)** shows the entire edge, while **b)** shows the $L\alpha$ edge alone.

Figure A.4: Oxygen K edge for $YBa_2Cu_3O_{6.99}$: **a)** Full spectrum, and **b)** detail of the O K edge and prepeaks.
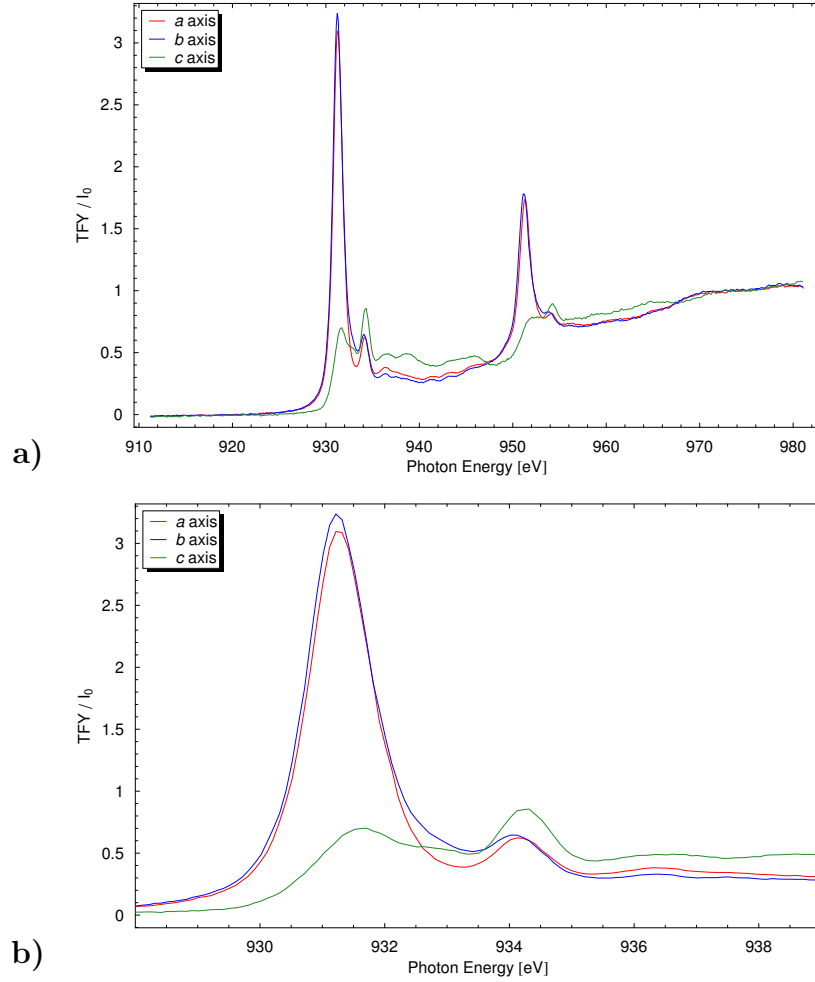
Figure A.5: Copper L edge for $YBa_2Cu_3O_{6.35}$ crystals: **a)** Entire spectrum, and **b)** detail of the Cu $L\alpha$ peak.

A.5 Doping Dependence

Collecting these results for a comparison of doping dependence (also discussed briefly in Section 4.2), the Cu L edge is compared in Figure A.13 and the O K edge in Figure A.14. Comparisons against Nücker's results have been included in Figures A.15 and A.16.
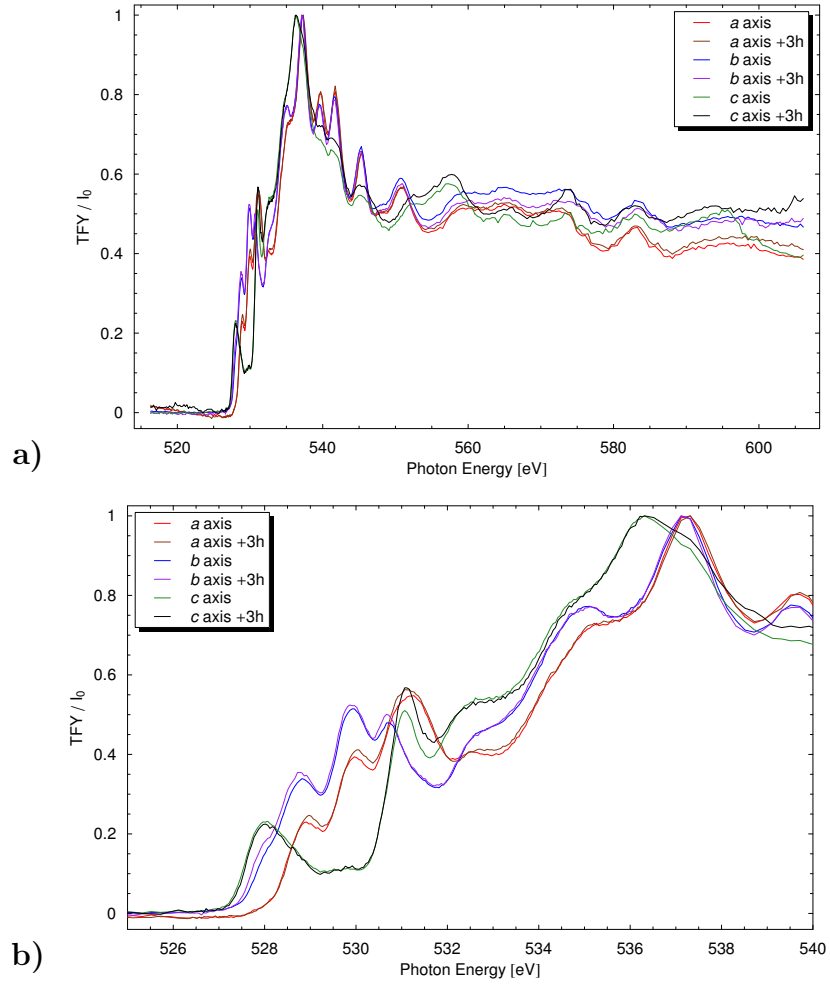
Figure A.6: Oxygen K edge for $YBa_2Cu_3O_{6.35}$ crystals: **a)** Entire spectrum, and **b)** the O K edge and prepeaks.
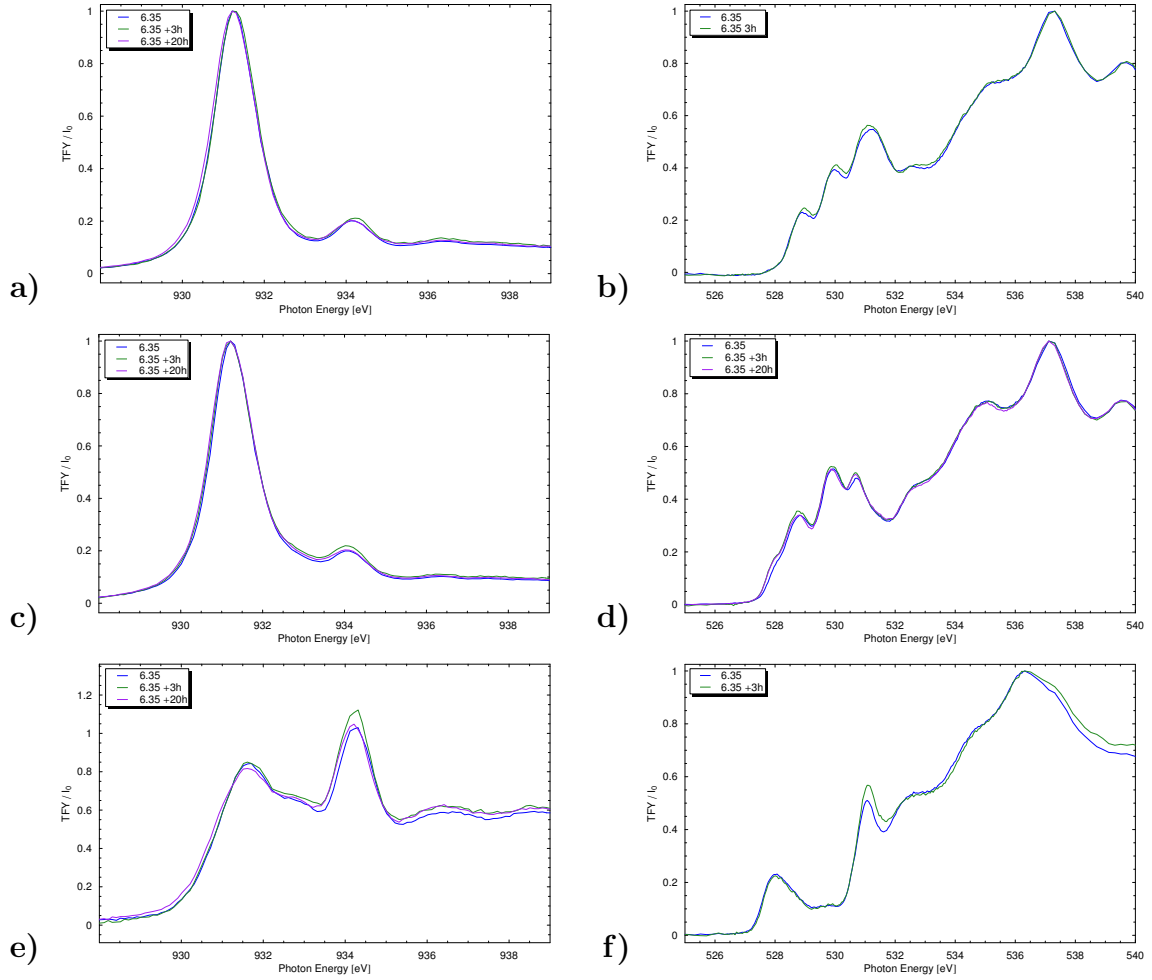
Figure A.7: Effect of oxygen dopant ordering on $YBa_2Cu_3O_{6.35}$ crystals: **a)** and **b)** show changes to the *a*-axis Cu *L* and O *K* edges respectively, **c)** and **d)** show changes to the *b*-axis Cu *L* and O *K* edge spectra respectively, and **e)** and **f)** show changes to the *c*-axis Cu *L* and O *K* edges respectively. Due to time constraints, it was not possible to ripen the crystals for more than a day.
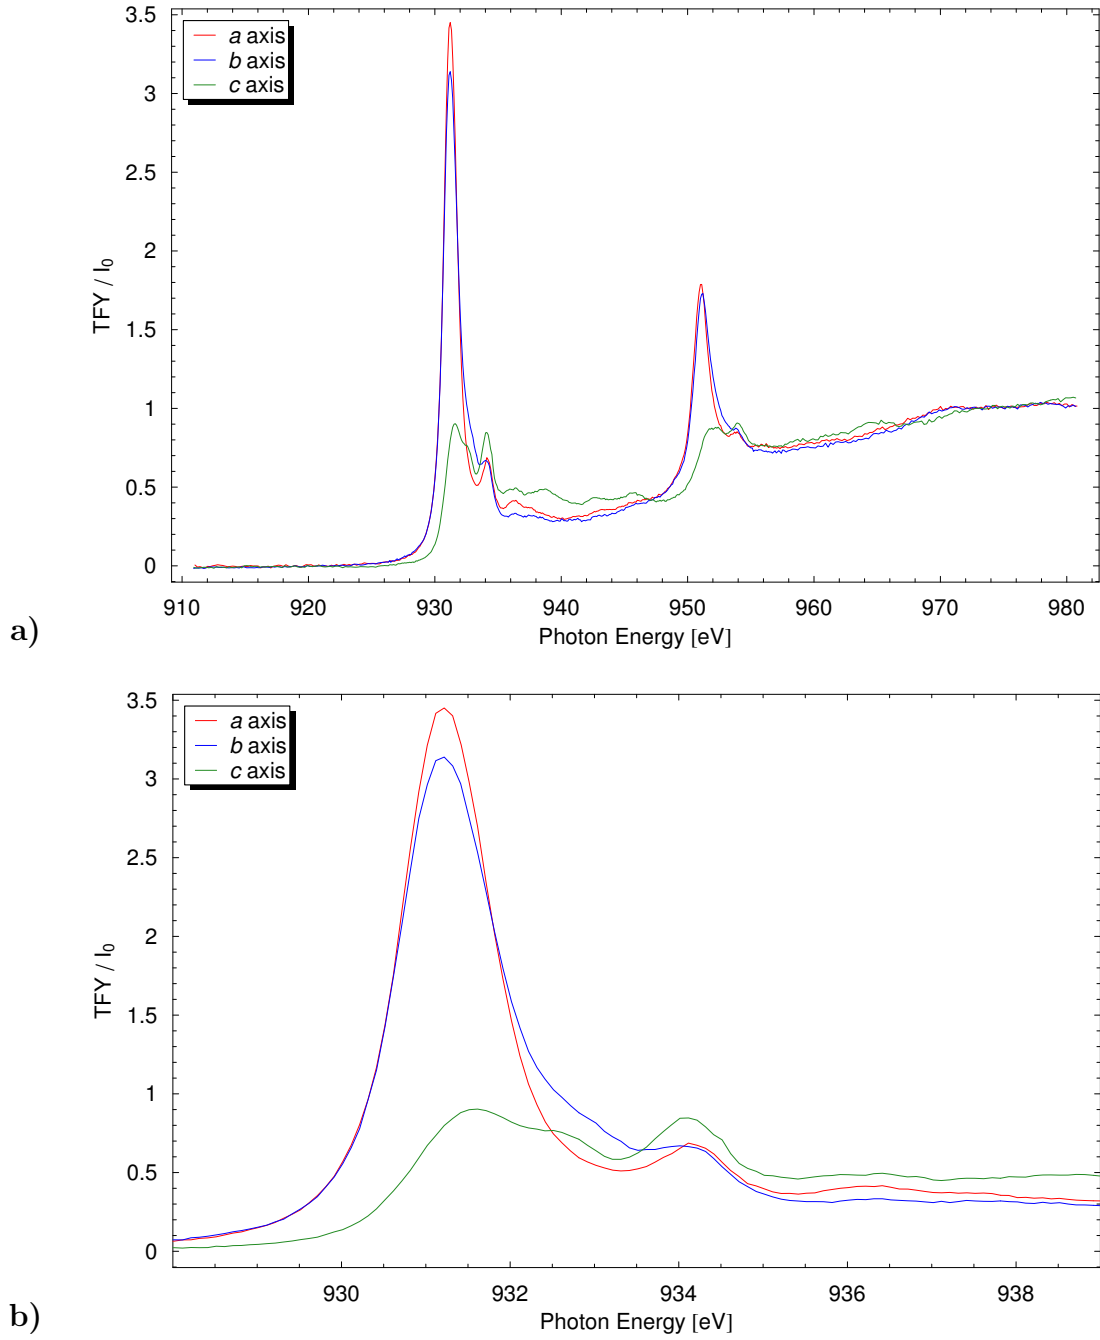
Figure A.8: Copper L edge for Ortho-II ordered $YBa_2Cu_3O_{6.5}$ crystals: **a)** Entire spectrum, and **b)** detail of the Cu $L\alpha$ peak.
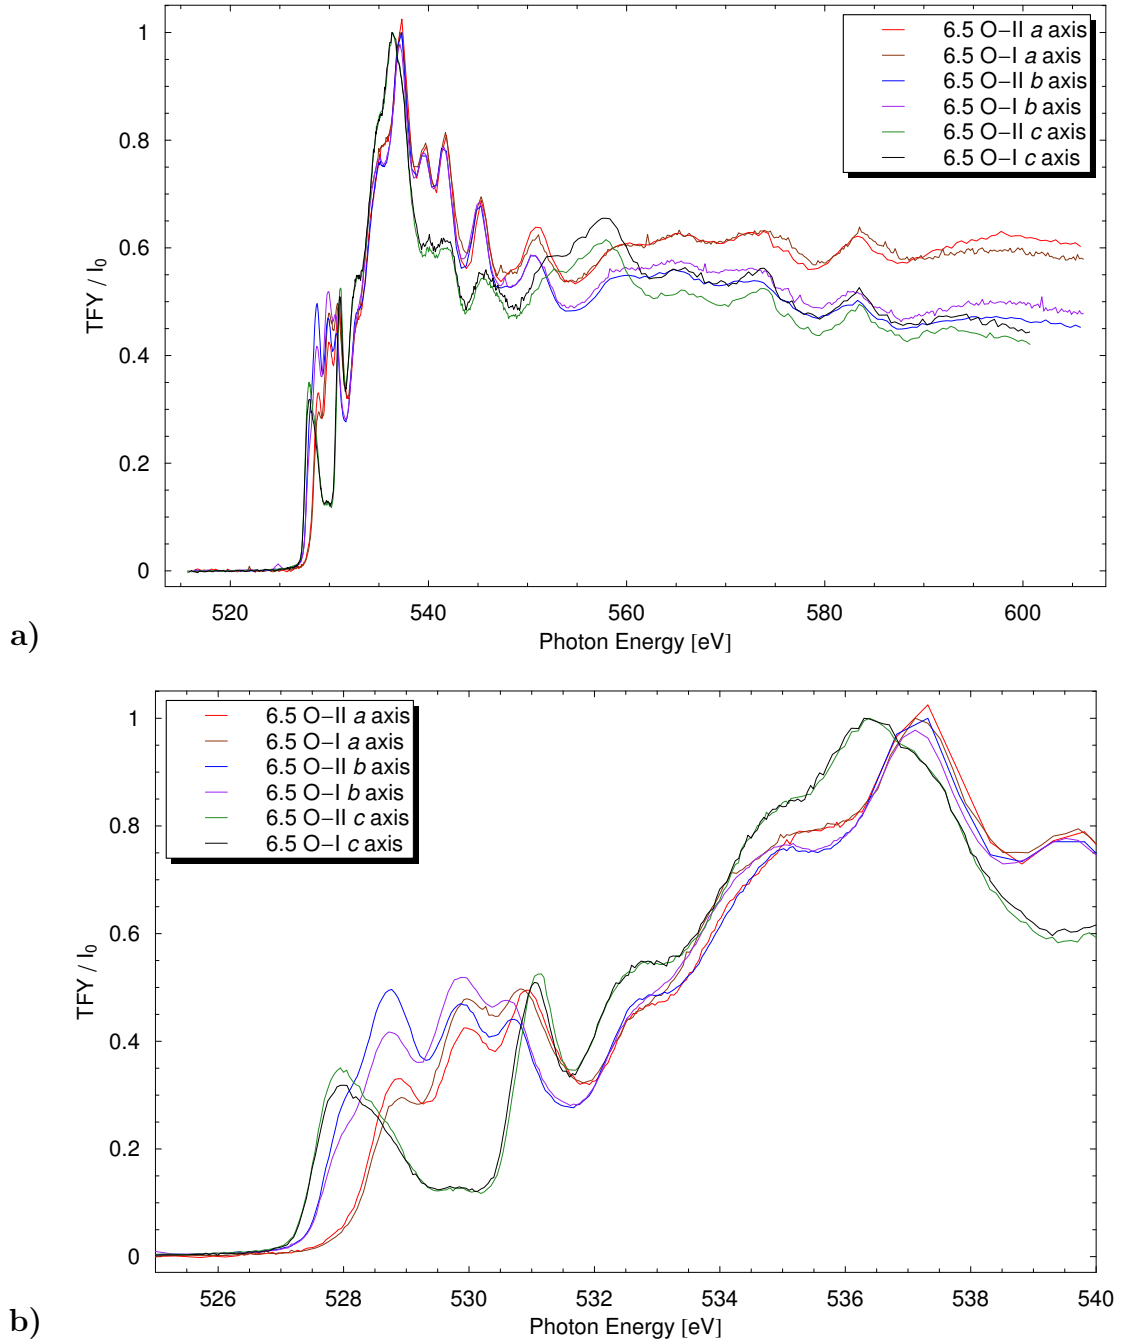
Figure A.9: Oxygen K edge for Ortho-II ordered and disordered (Ortho-I) $YBa_2Cu_3O_{6.5}$ crystals. The prepeaks previously appeared in Figure 4.7.
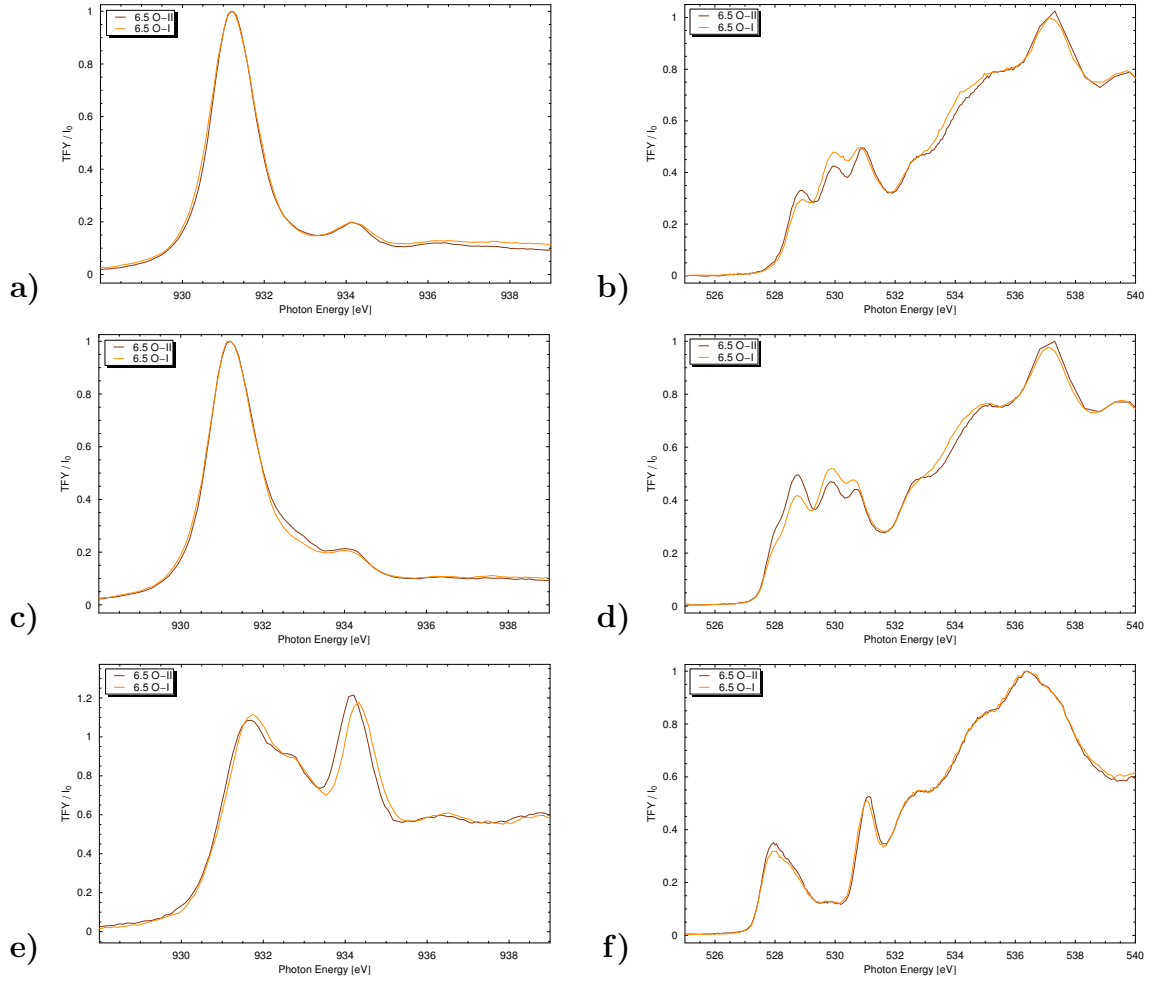
Figure A.10: Effect of oxygen dopant disorder on $YBa_2Cu_3O_{6.5}$ crystals: **a)** and **b)** show changes to the *a*-axis Cu *L* and O *K* edges respectively, **c)** and **d)** show changes to the *b*-axis Cu *L* and O *K* edge spectra respectively, and **e)** and **f)** show changes to the *c*-axis Cu *L* and O *K* edges respectively, as the Ortho-II oxygen order is scrambled to Ortho-I.
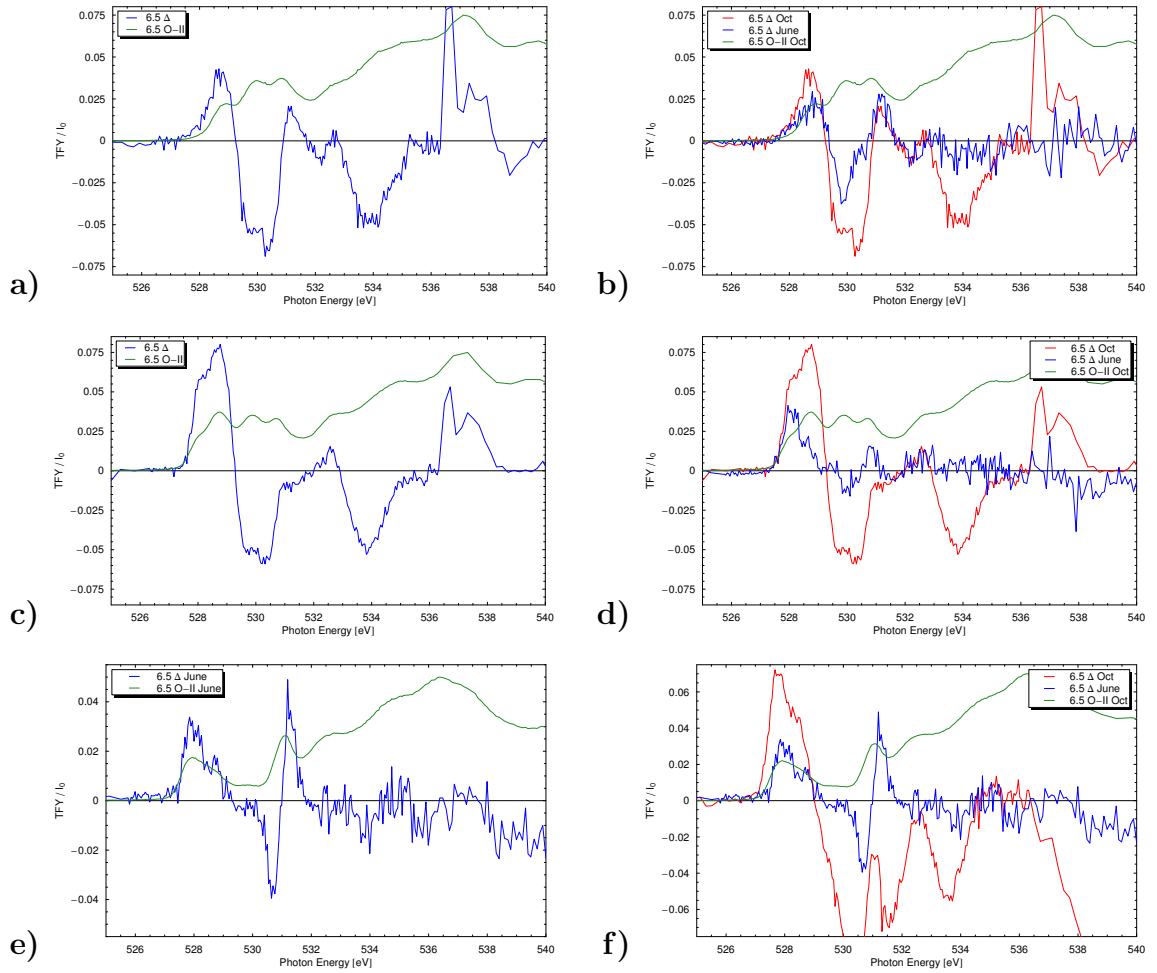
Figure A.11: Change due to oxygen ordering, subtraction of spectra: **a)** and **b)** show the *a*-axis, **c)** and **d)** the *b*-axis, and **e)** and **f)** the *c*-axis spectra.
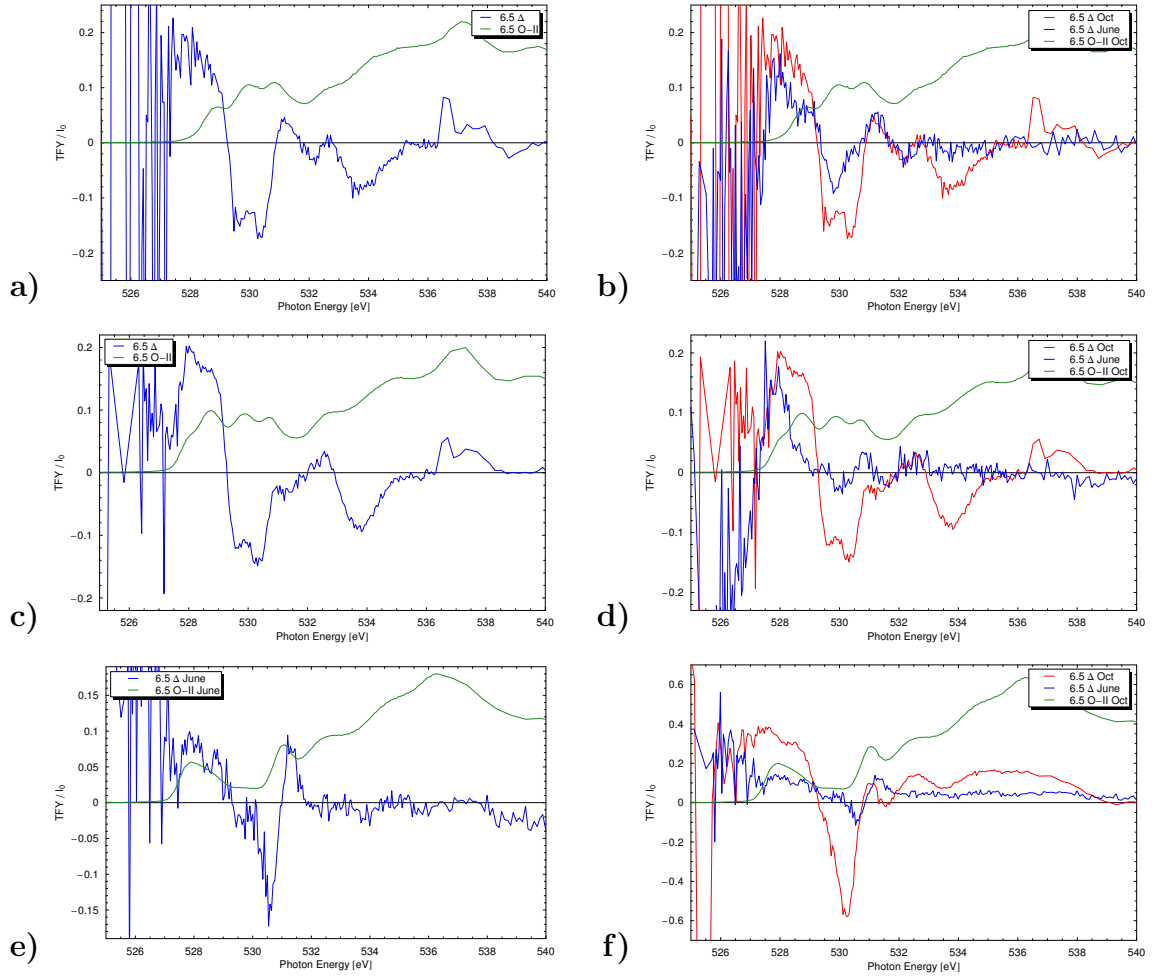
Figure A.12: Relative change due to oxygen ordering: **a)** and **b)** show the *a*-axis, **c)** and **d)** the *b*-axis, and **e)** and **f)** the *c*-axis spectra.
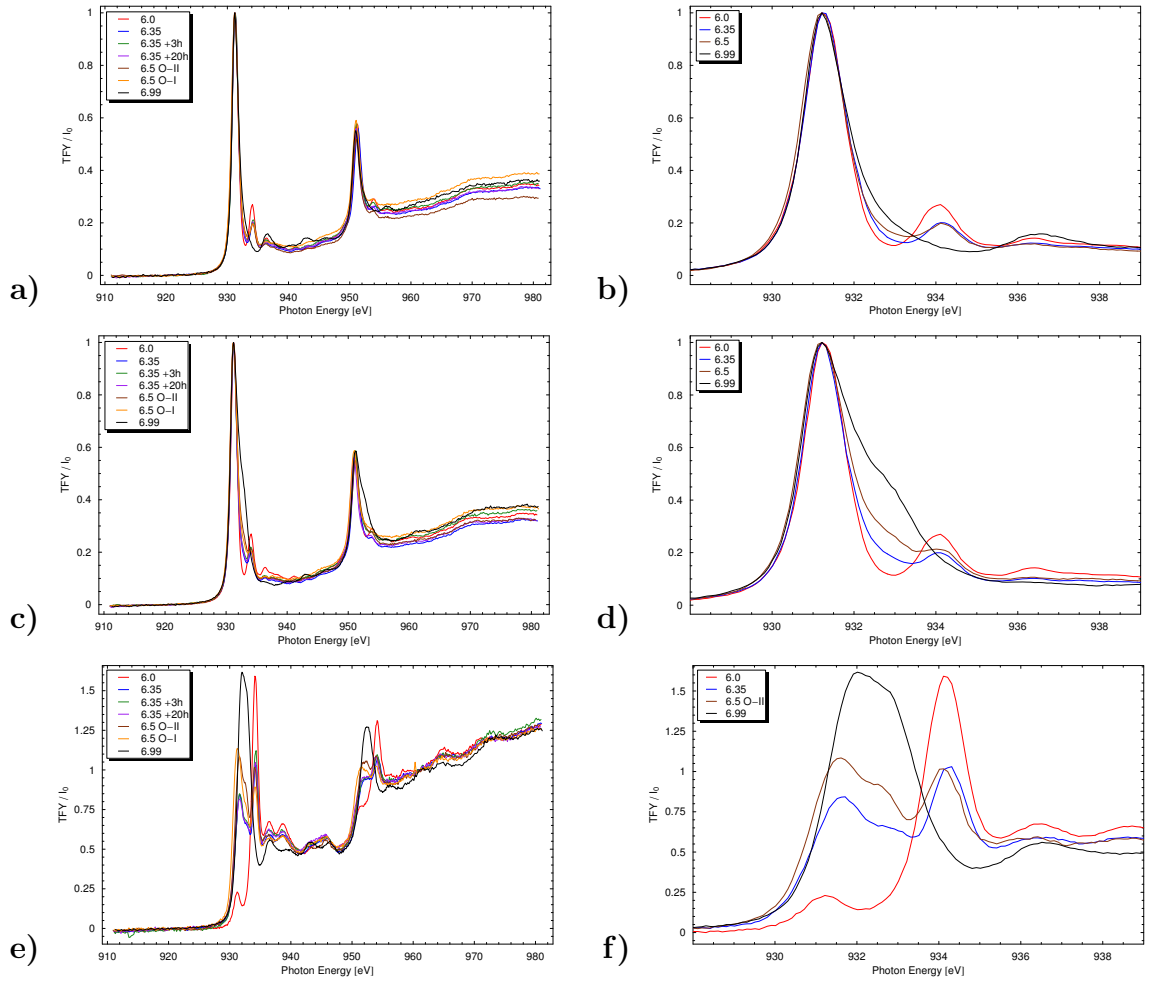
Figure A.13: Doping dependence of copper L edge in YBCO: **a)** and **b)** show the a-axis, **c)** and **d)** the b-axis, and **e)** and **f)** the c-axis spectra. **a)**, **c)** and **e)** show the entire spectrum, while **b)**, **d)** and **f)** show the edge only, and for selected traces.
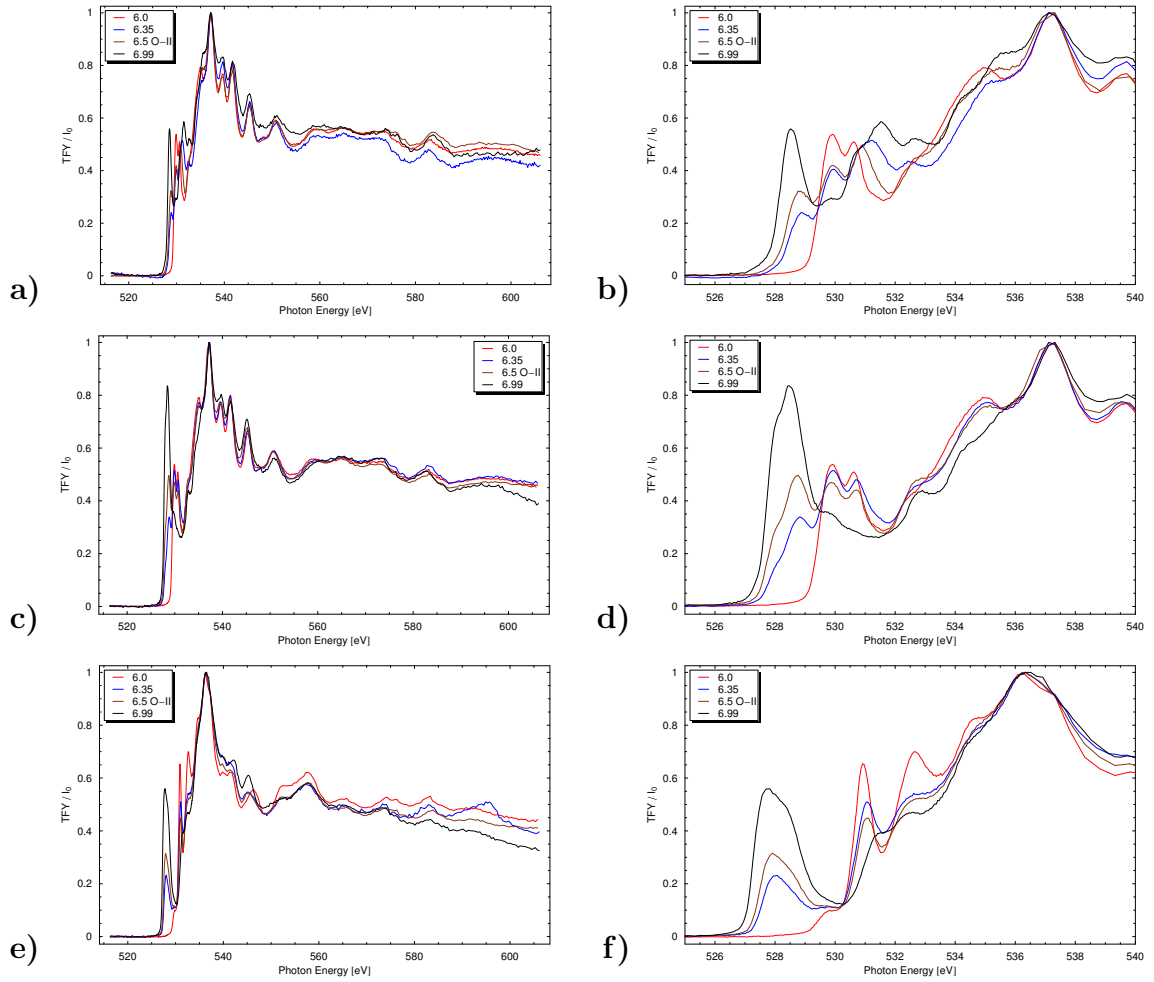
Figure A.14: Doping dependence of O K edge in YBCO: **a)** and **b)** show the a-axis, **c)** and **d)** the b-axis, and **e)** and **f)** the c-axis spectra. **a)**, **c)** and **e)** show the entire spectrum, while **b)**, **d)** and **f)** show only the prepeaks and edge.
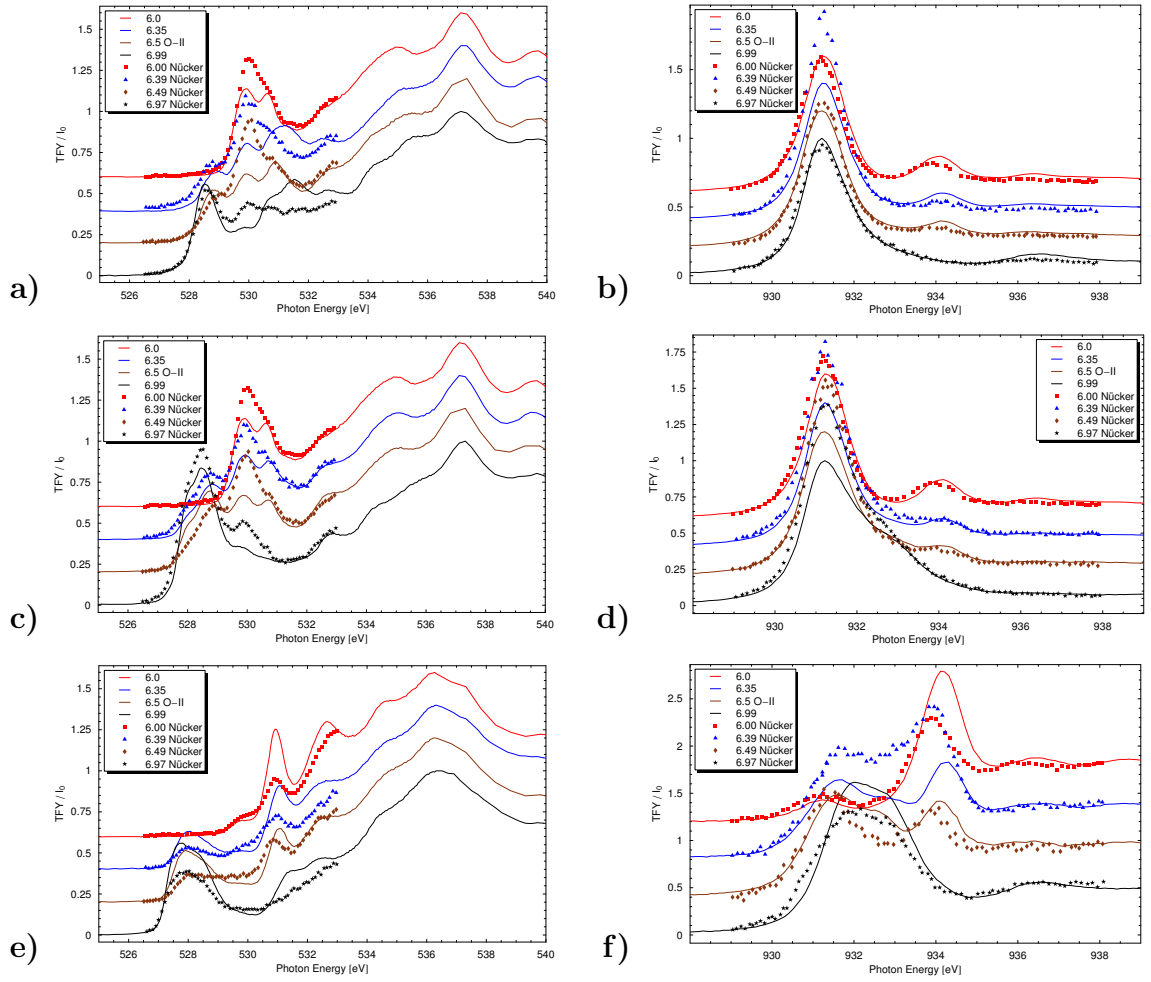
Figure A.15: Doping dependence of each edge, compared against Nücker's results [106]: **a)** and **b)** show the *a*-axis, **c)** and **d)** the *b*-axis, and **e)** and **f)** the *c*-axis spectra. Panels **a)**, **c)** and **e)** show the oxygen *K* edge, while **b)**, **d)** and **f)** show the copper *L* edge.
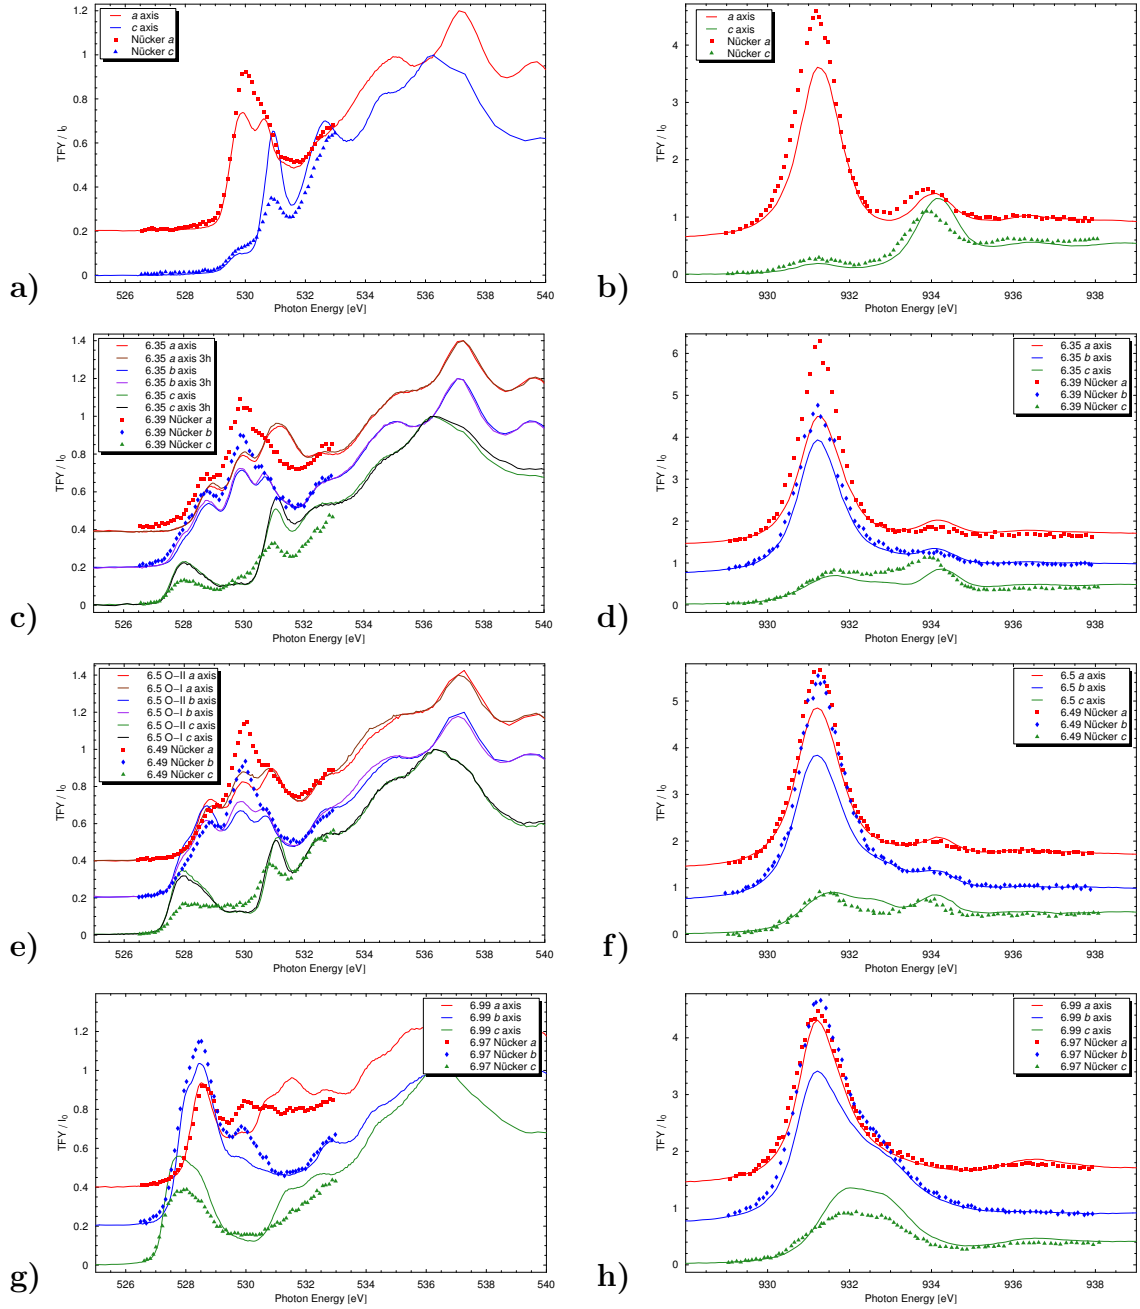
Figure A.16: Polarization dependence of each edge, compared against Nücker's results [106]: **a)** and **b)** show $7 - \delta = 6.0$, **c)** and **d)** 6.35, **e)** and **f)** 6.5, and **g)** and **h)** the 6.99 spectra. Panels **a)**, **c)**, **e)**, and **g)** present the O K edge results, while panels **b)**, **d)**, **f)**, and **h)** show the Cu L edge.

B. Full XAS and XES Results on Tl-2201

As Appendix A did for $YBa_2Cu_3O_{6+\delta}$, this appendix covers Tl-2201 results that were not presented in the text of this thesis. The results are not discussed to any significant degree; they're presented in graphical form. In the context of XAS, angle-dependent scans amount to raw data for c-axis spectra, while alternate normalization schemes are included here for completeness' sake.

B.1 9.5 K Critical Temperature

This doping is presented and discussed in Section 5.1.

The copper L edges measured as a function of incident photon angle are shown in Figure B.1 for a Tl-2201 crystal with $T_c = 9.5$ K. A clear evolution can be seen from low to high angle, with the edges themselves being strongly suppressed, features above the edge being enhanced, and a shoulder near 933.5 eV separating out into a distinct peak. The $L\alpha$ and $L\beta$ edges should, in principle, behave identically, but the former is more clearly resolved because of its higher intensity. The oxygen K edge equivalent is Figure B.2. A full comparison of the in-plane and c-axis spectra, normalized by different methods, is presented in B.3.

As an internal consistency check, high-resolution spectra were repeated with samples at zero and fifty degrees to the beam. These spectra are shown in Figure B.4. At normal incidence, three scans on sample #5 and two on sample #4 are shown — #5 is the crystal featured in Figure 3.7 and was used for all analysis. The spectra for sample #5 are taken with different point spacings and entrance slit settings (the latter changes energy resolution and intensity); the spectra on sample #4 were taken eight hours apart at slightly different positions on the sample, separated by measurements on other samples and at other angles. The spectra recorded at 50° are from very different locations on the same crystal, and were separated in time by measurements at other angles.

The most serious reproducibility issue is the background. Variations from sample to sample and among distinct locations on the same sample are otherwise minor. The third spectrum at 0°, which appears to exhibit slightly stronger in the vicinity
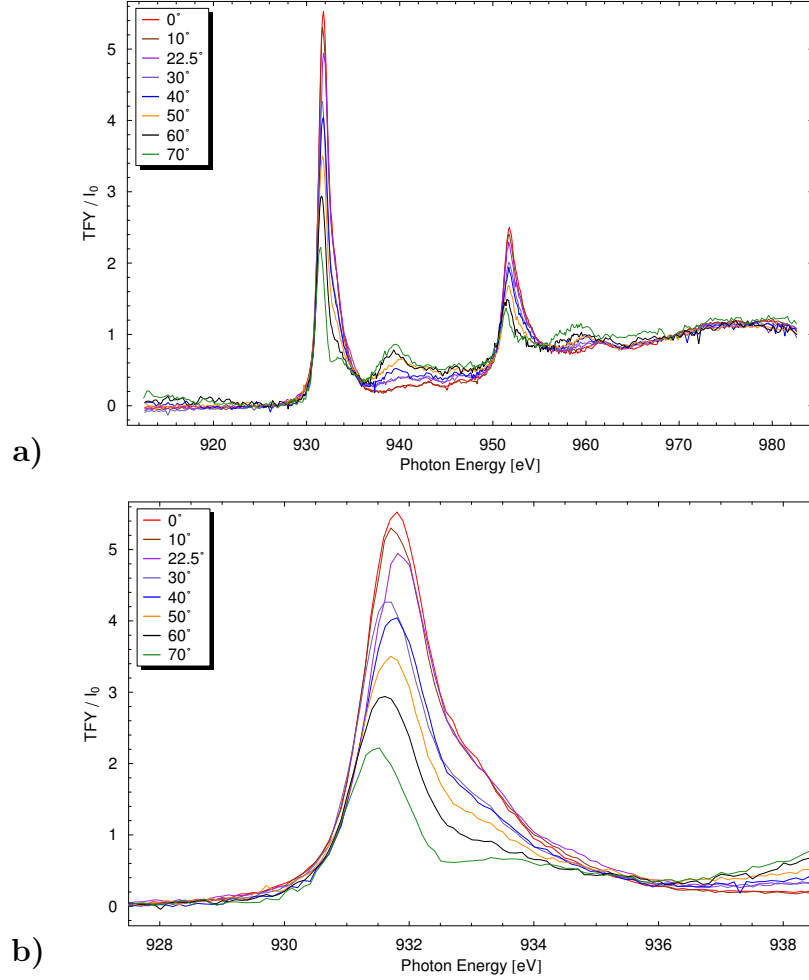
Figure B.1: Evolution of Cu L edge with angle for a $T_c = 9.5$ K Tl-2201 crystal. **a)** Entire spectrum for all angles, and **b)** detail of the Cu $L\alpha$ peak.

of the edge and prepeaks, was taken with narrower exit and entrance slits on the monochromator, leading to improved energy resolution and roughly a factor of three decrease in X-ray flux.

Additionally, Figure B.5 is included. This figure shows the effect of subtracting $\frac{1}{4}$ of the c-axis weight from the Cu L edge, along with the Lorentzian fit that was used for subtraction.

B.2   60 K Critical Temperature

As for the $T_c = 9.5$ K crystal presented in Section B.1, this section includes reference data for the $T_c = 60$ K crystal discussed in Section 5.2. Figures B.6 and B.7 show
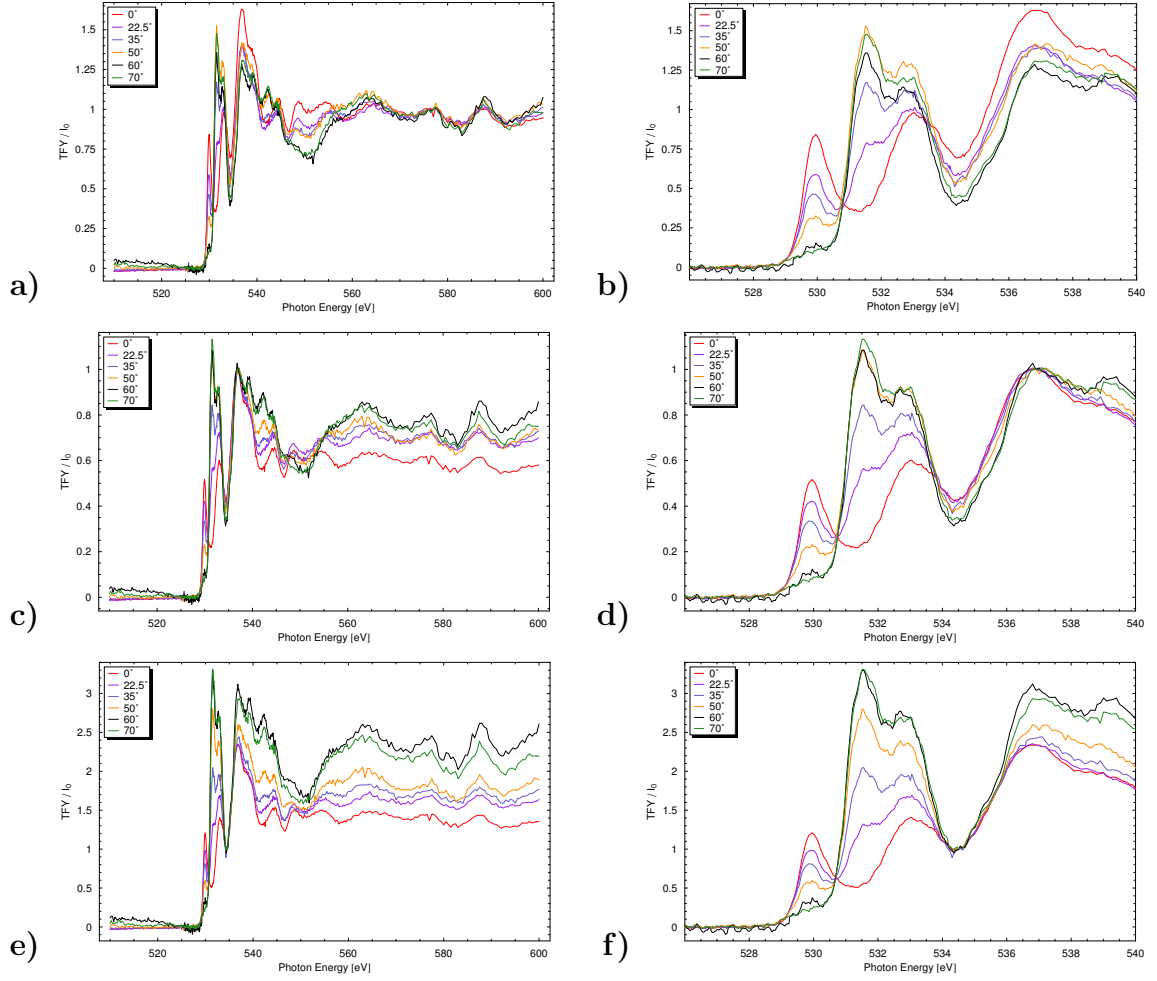
Figure B.2: Evolution of oxygen K edge with angle for $T_{\mathrm{c}} = 9.5$ K Tl-2201 crystal, normalized several different ways. Panels **a)**, **c)** and **e)** are normalized well above the edge, on the edge, and in the notch between the prepeaks and edge respectively, while panels **b)**, **d)** and **f)** show the prepeak and oxygen K edge in more detail for each respective spectrum.
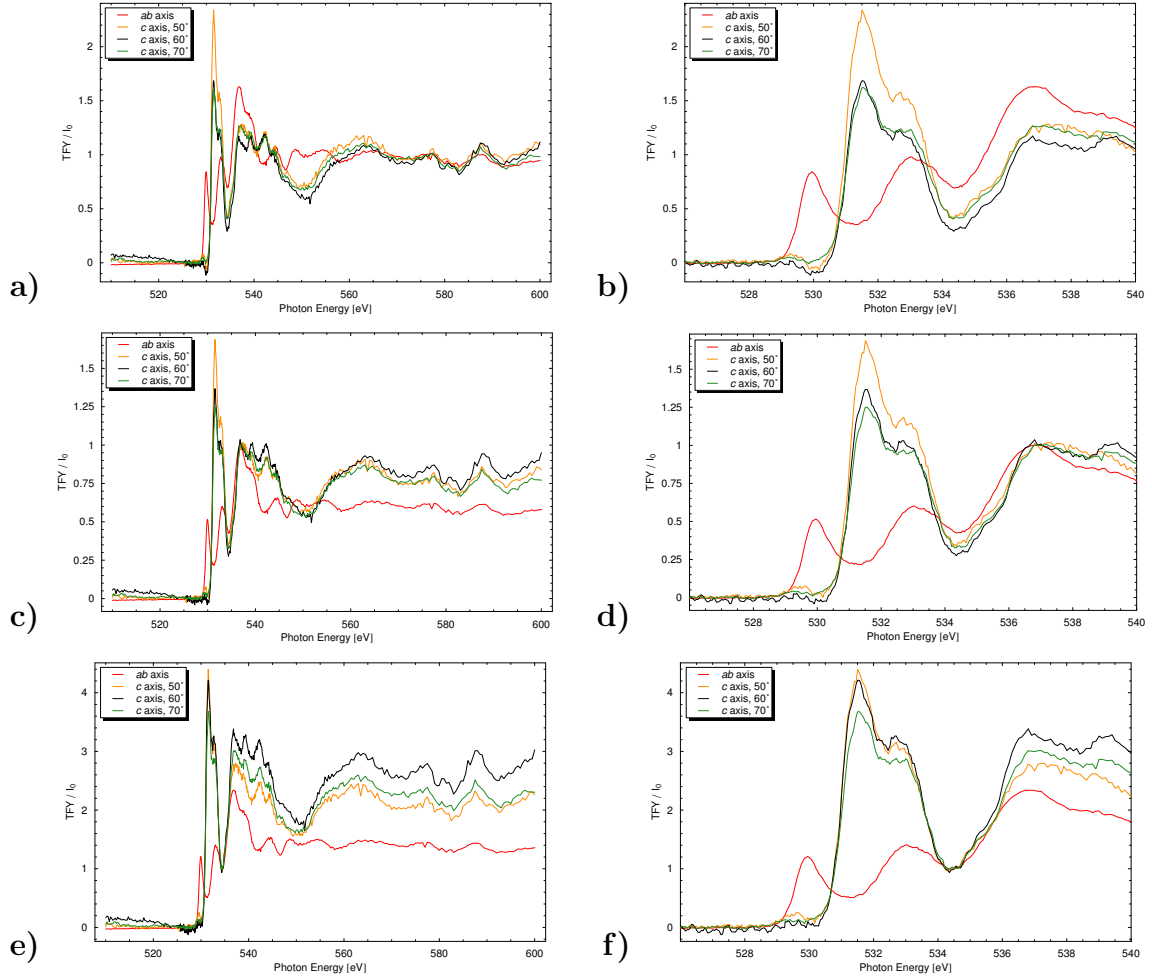
Figure B.3: A comparison of the *ab*-plane and *c*-axis O *K* edge spectra for a $T_c = 9.5$ K crystal, normalized three different ways. Panels **a)**, **c)** and **e)** are extracted from spectra normalized well above the edge, on the edge, and in the notch between the prepeaks and edge respectively, while panels **b)**, **d)** and **f)** show the prepeak and oxygen *K* edge in more detail for each respective spectrum. The latter panel is also shown as Figure 5.2 above.
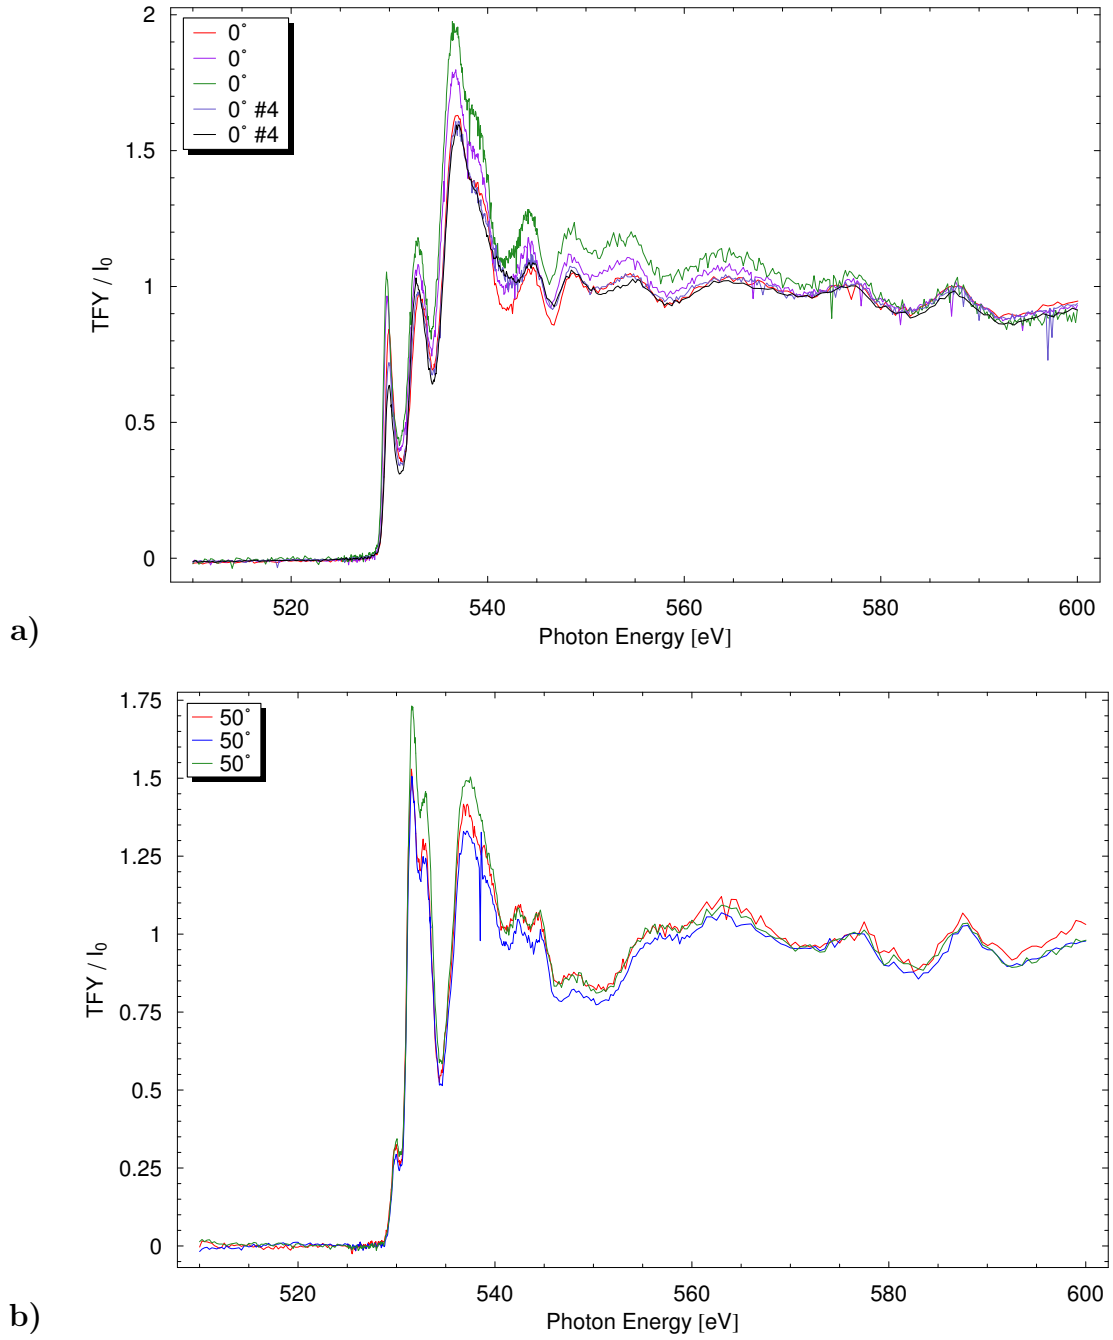
Figure B.4: Reproducibility of XAS spectra: Multiple oxygen K edge spectra obtained on $T_c = 9.5$ K crystals at incident angles of **a)** 0° and **b)** 50°.
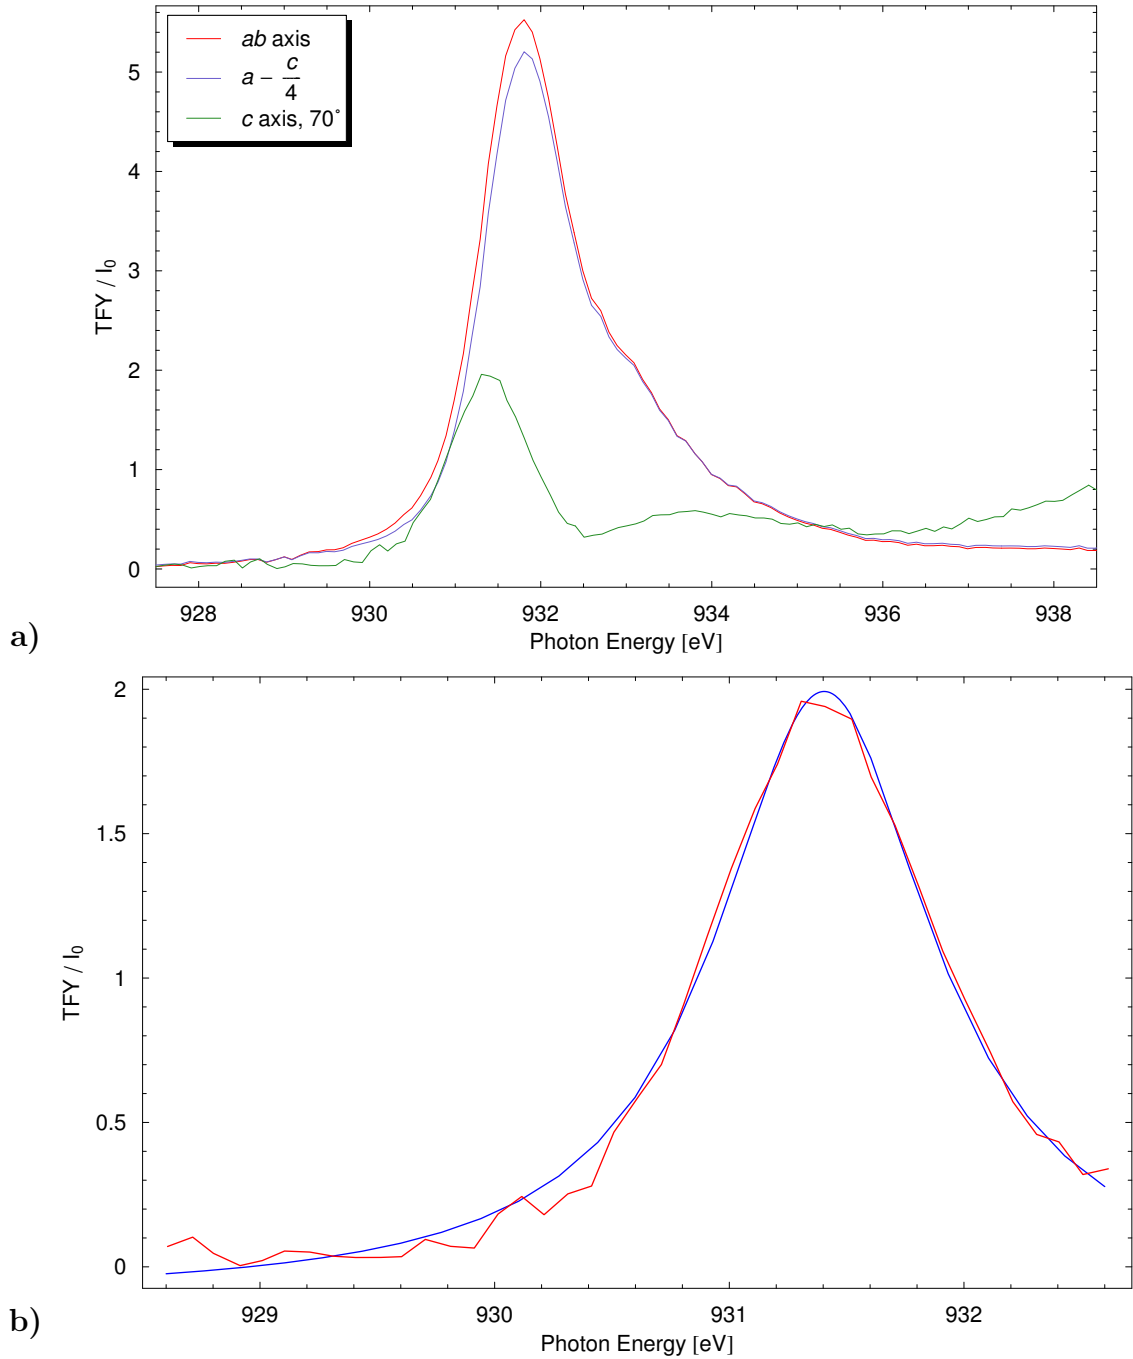
Figure B.5: Effect of subtracting $c/4$ from the ab-plane spectrum. **a)** The $T_c = 9.5$ K
crystal's ab-plane and c-axis spectra are shown, along with the ab-plane
spectrum as corrected by subtracting one-quarter of the Lorentzian fit to
the first peak in the c-axis scan. **b)** The Lorentzian fit.

the angle dependence of the Cu L and O K edges respectively, while Figures B.8 and B.8 show the polarization dependence for the Cu L and O K edges respectively.

B.3   69 K Critical Temperature

As for the other two dopings, this section contains reference data for the $T_{\mathrm{c}} = 69$ K crystal discussed in Section 5.3. Figures B.10 and B.11 show the angle dependence of the Cu L and O K edges respectively, while Figures B.12 and B.12 show the polarization dependence for the Cu L and O K edges respectively.

B.4   Doping Dependence

This section contains a more complete set of traces for the edges' doping dependence presented in Section 5.4. Figures B.14 and B.15 show the oxygen K edge's doping evolution for in-plane and c-axis polarization respectively, using multiple normalization schemes. There was no need to include matching copper data here.

Of the assorted normalization techniques, the one which gives the most consistent results on the oxygen prepeaks is normalizing on the notch between them and the main edge. This is somewhat surprising, as this energy range could in principle have doping- or angle-dependent features, making it susceptible to changes. In Tl-2201, however, this does not seem to be the case. On the copper edge, normalizing by peak height produced the best results for comparisons among dopings (if there were a systematic change in peak height with doping, this technique would not work). For polarization-dependence, this is not an option, so only the region well above both edges is suitable.

B.5   X-ray Emission

This section presents a few alternate normalization schemes and display methods for the Tl-2201 XES data presented in Chapter 6. Figures B.16 and B.17 show the Cu L and O K edge emission data respectively, normalized three distinct ways.

Another, particularly colourful, way of presenting these data is in the form of a density plot. Density plots of the copper and oxygen edge emission spectra are displayed in Figures B.18 and B.19 respectively.
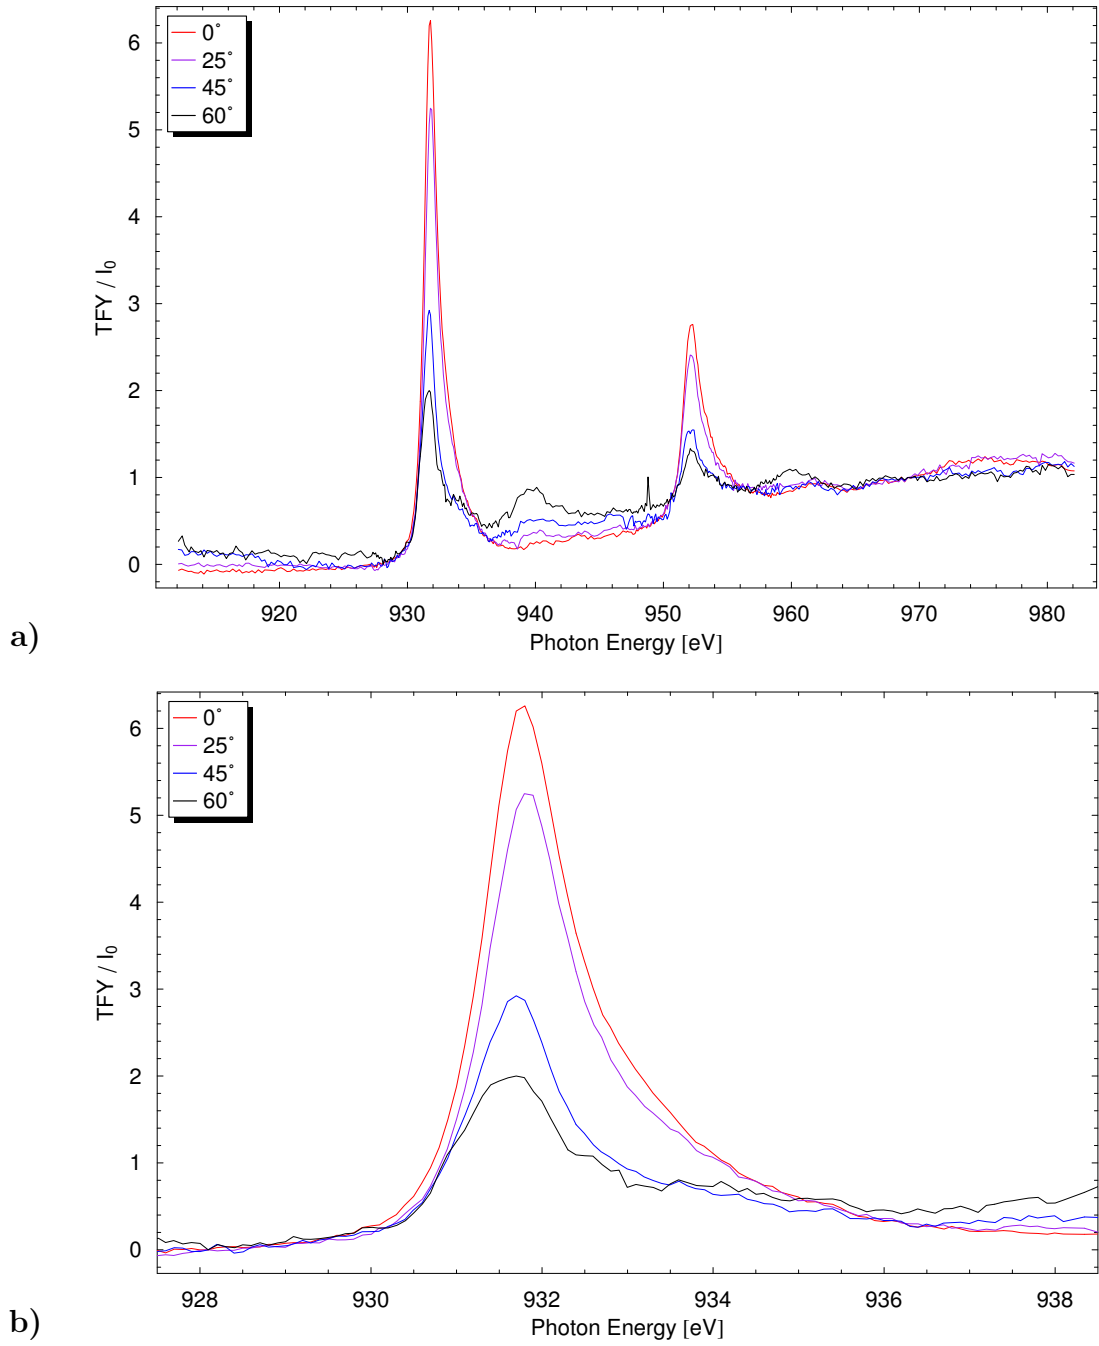
Figure B.6: Evolution of Cu L edge with angle for $T_c = 60$ K Tl-2201 crystal. **a)** Entire spectrum for all angles, and **b)** detail of the Cu $L\alpha$ peak.
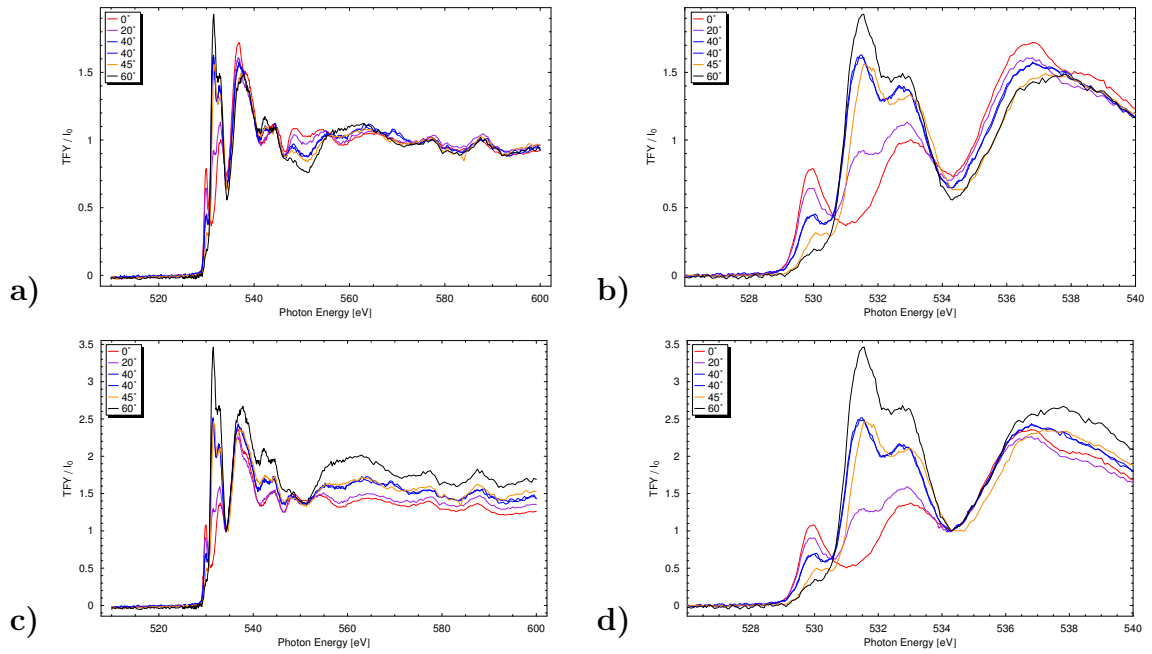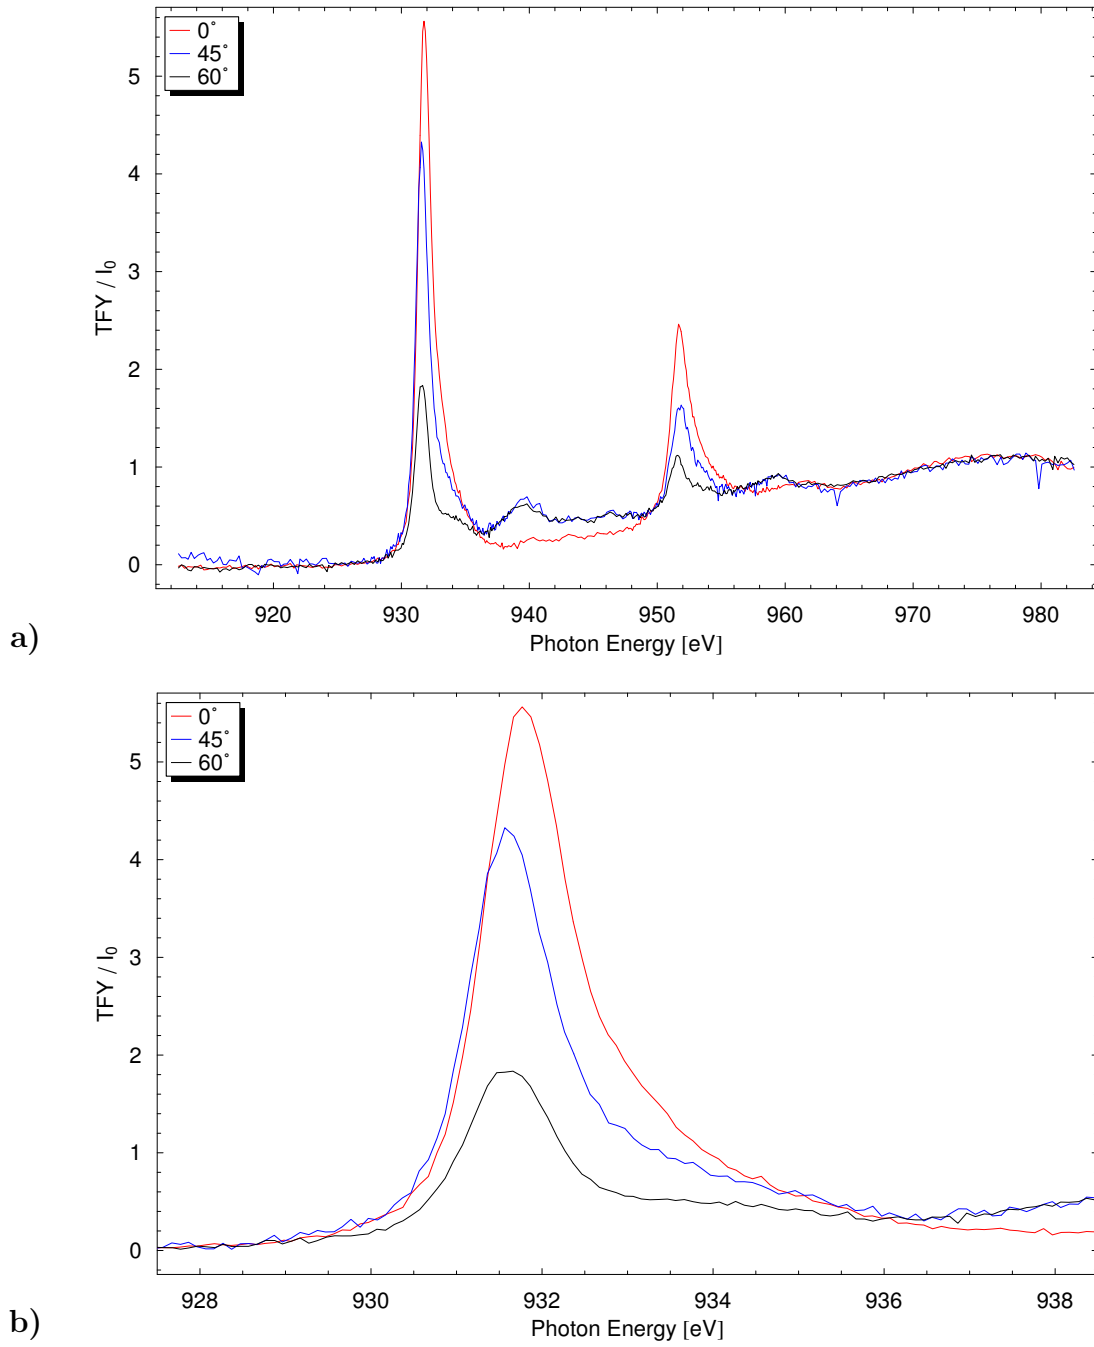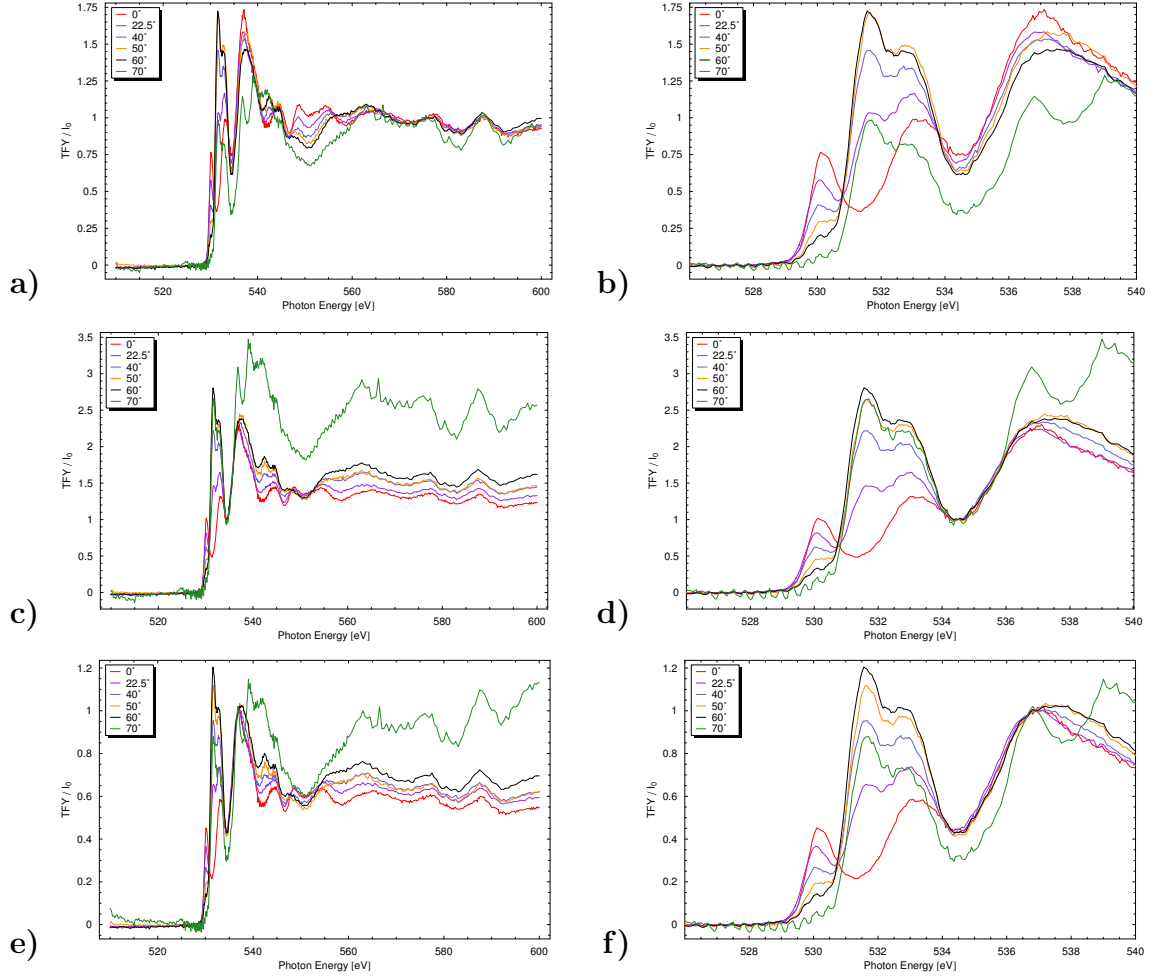
Figure B.7: Evolution of oxygen K edge with angle for $T_c = 60$ K Tl-2201 crystal, normalized two different ways. Panels **a)** and **c)** are normalized far above the edge and in the notch between the prepeaks and edge respectively, while panels **b)** and **d)** show the prepeak and oxygen K edge in more detail for each respective spectrum.
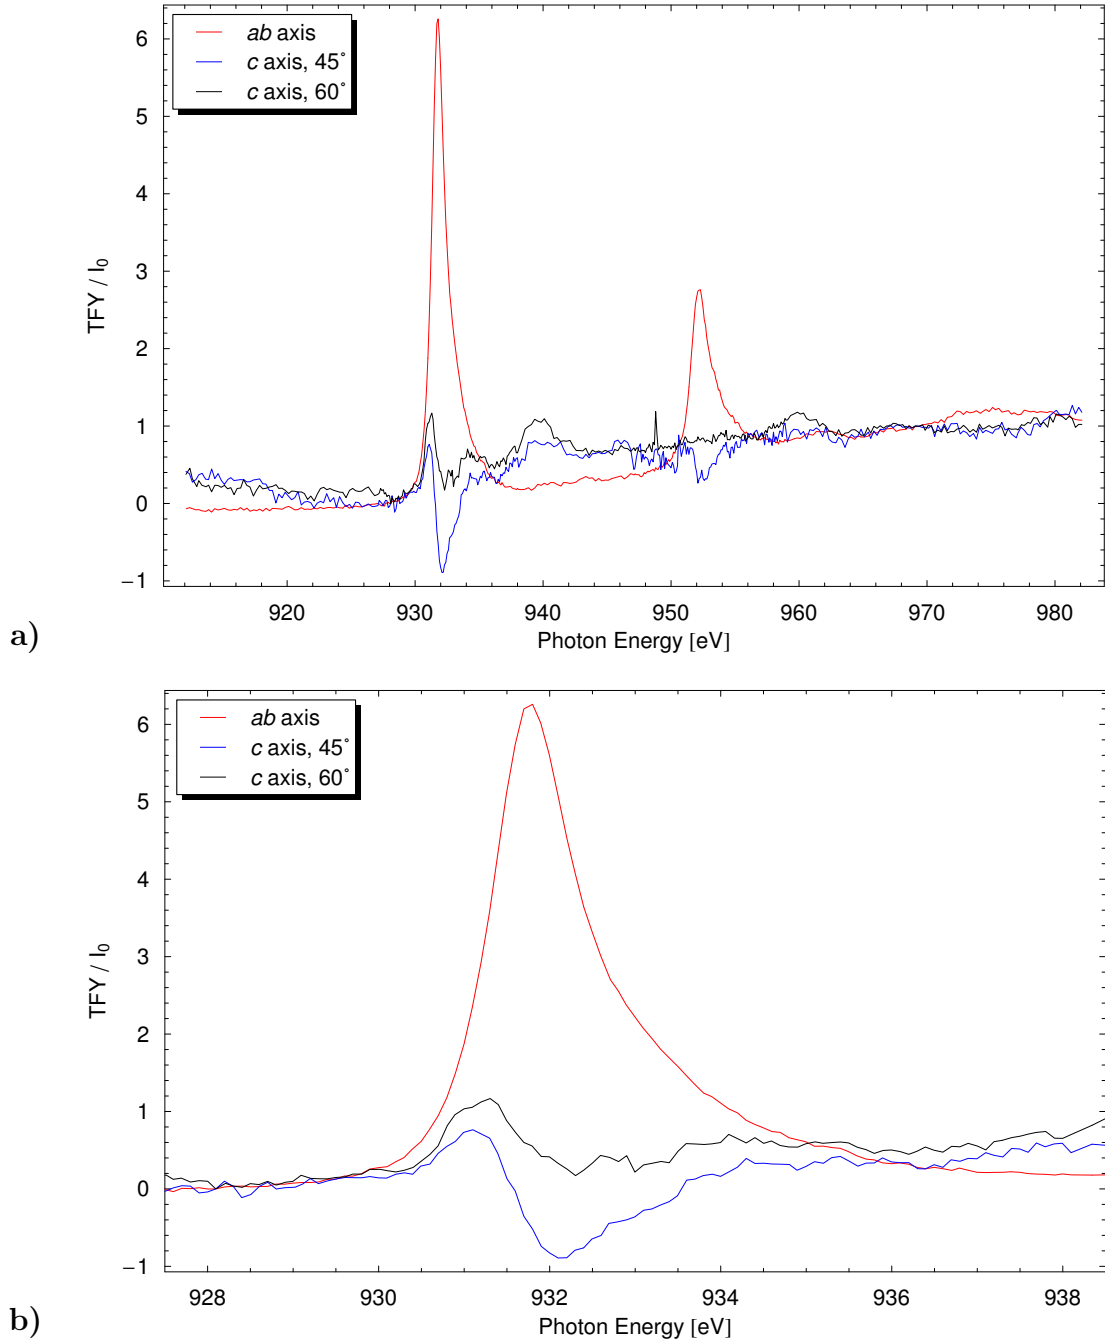
Figure B.8: A comparison of the *ab*-plane and *c*-axis Cu *L* edge spectra for a $T_c = 60$ K crystal; **b)** is a detail of the *Lα* region of **a)**. The *c*-axis component is again quite different from the *ab* component.
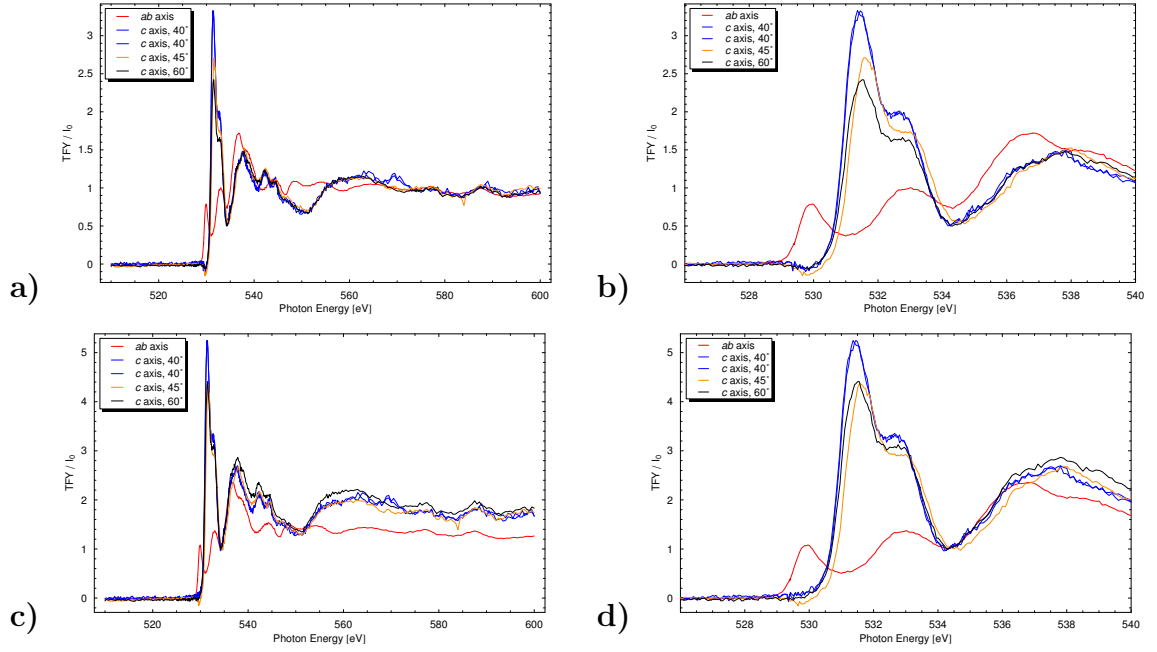
Figure B.9: A comparison of the *ab*-plane and *c*-axis O K edge spectra for a $T_\mathrm{c} = 9.5$ K crystal, normalized two different ways. Panels **a)** and **c)** are extracted from spectra normalized well above the edge and in the notch between the prepeaks and edge respectively, while panels **b)** and **d)** show the prepeak and oxygen K edge in more detail for each respective spectrum. The *c*-axis component extracted for higher angles converges, and is again different from the *ab* component.

Figure B.10: Evolution of Cu L edge with angle for $T_{\mathrm{c}} = 69$ K Tl-2201 crystal. **a)** Entire spectrum for all angles, and **b)** detail of the Cu $L\alpha$ peak. These scans originally had backgrounds that prevented them from being effectively compared. A parabola was fit to the difference between each higher-angle scan and the 0° scan (well away from the edges) and subtracted.

Figure B.11: Evolution of oxygen K edge with angle for $T_c = 69$ K Tl-2201 crystal, normalized three different ways. Panels **a)**, **c)** and **e)** are normalized far above the edge, in the notch between the prepeaks and edge, and on the edge itself respectively, while panels **b)**, **d)** and **f)** show the prepeak and oxygen K edge in more detail for each respective spectrum.
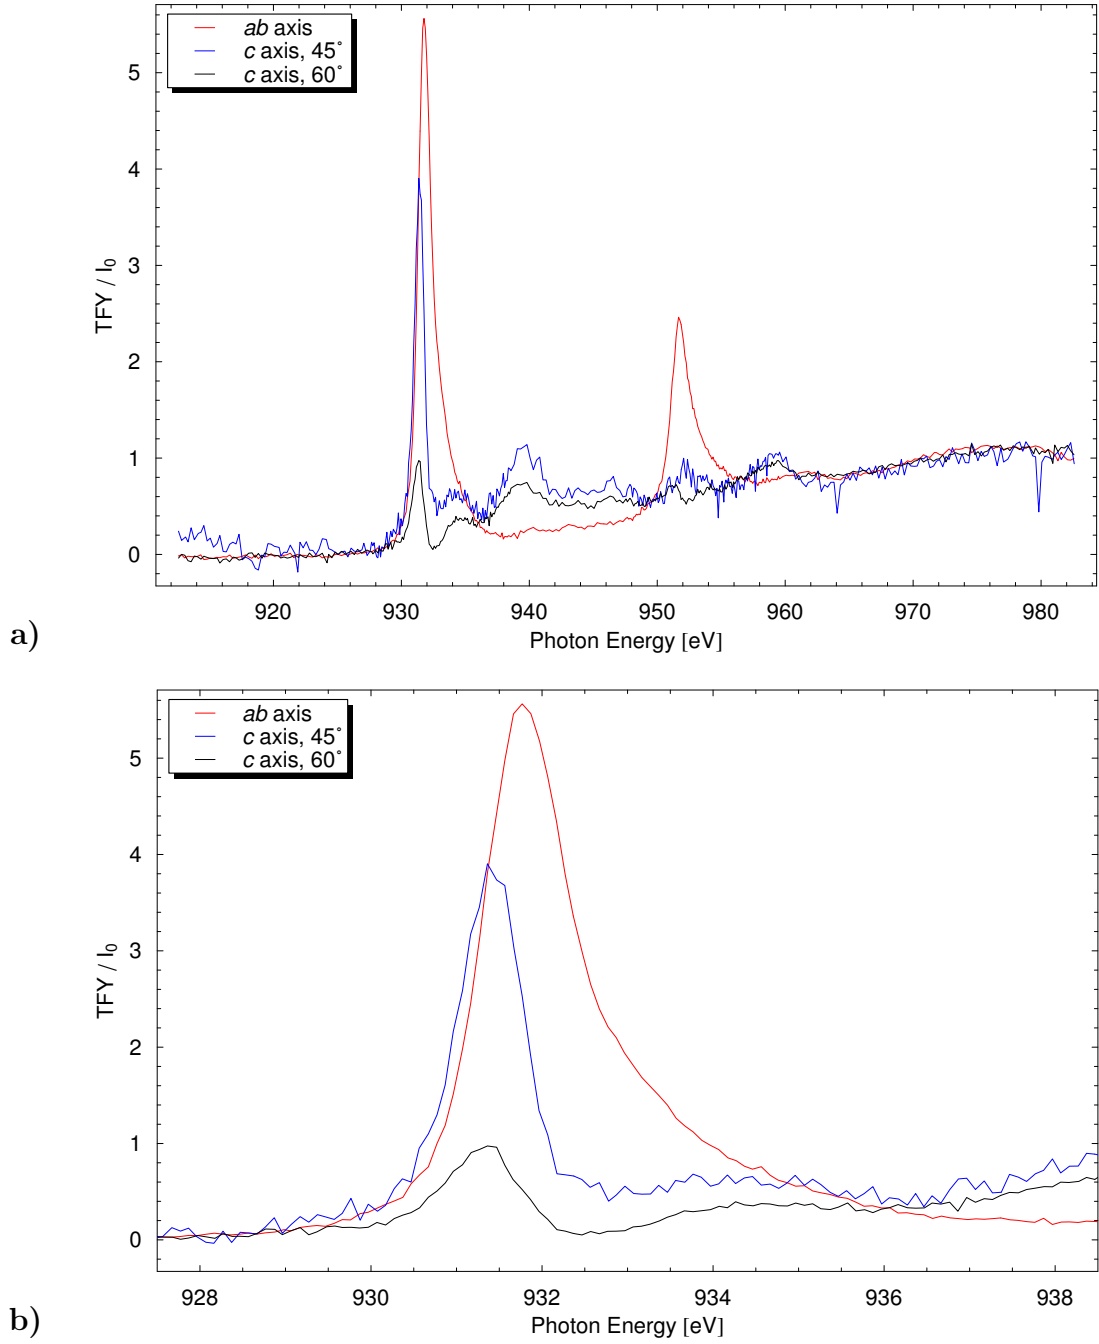
Figure B.12: A comparison of the *ab*-plane and *c*-axis Cu L edge spectra for a $T_c = 69$ K crystal; **b)** is a detail of the $L\alpha$ region of **a)**.
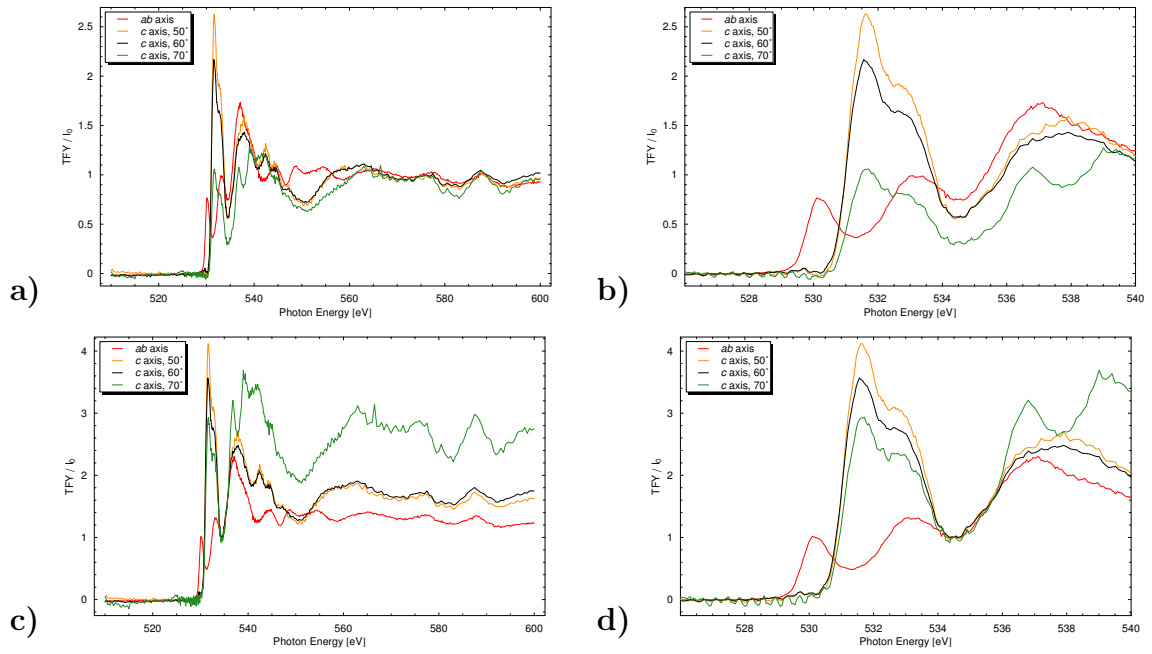
Figure B.13:   A comparison of the *ab*-plane and *c*-axis O K edge spectra for a $T_c = 69$ K crystal, as obtained using two distinct normalization schemes. Panels **a)** and **c)** are extracted from spectra normalized well above the edge and in the notch between the prepeaks and edge respectively, while panels **b)** and **d)** show the prepeak and oxygen K edge in more detail for each respective spectrum.
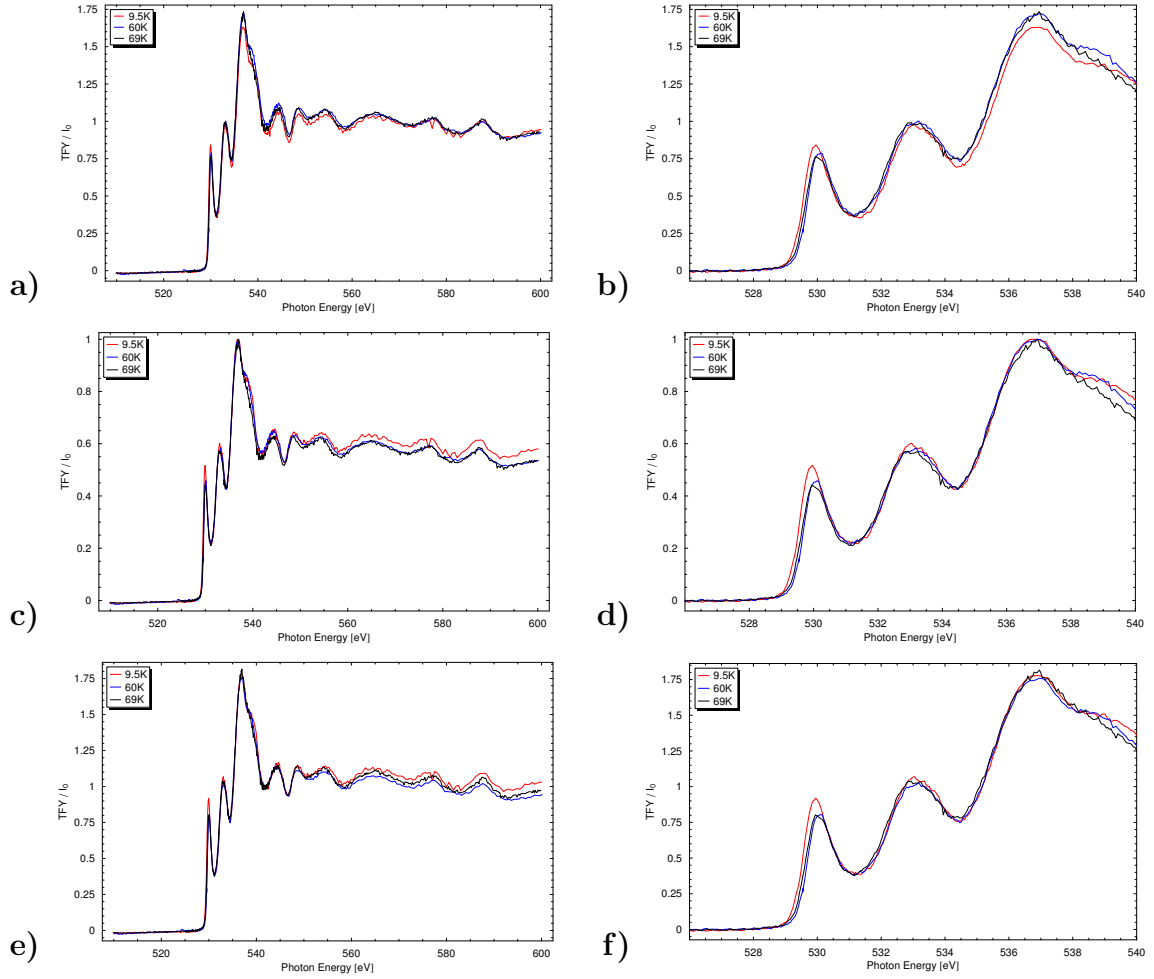
Figure B.14: Doping dependence of the O K edge in-plane spectra. The ab-plane spectra are shown, normalized using **a)** the energy range above the edge, **c)** the oxygen K edge, and **e)** the notch between the prepeak and main edge. Panels **b**, **d** and **f** show the prepeaks and edge in more detail for each respective normalization scheme.
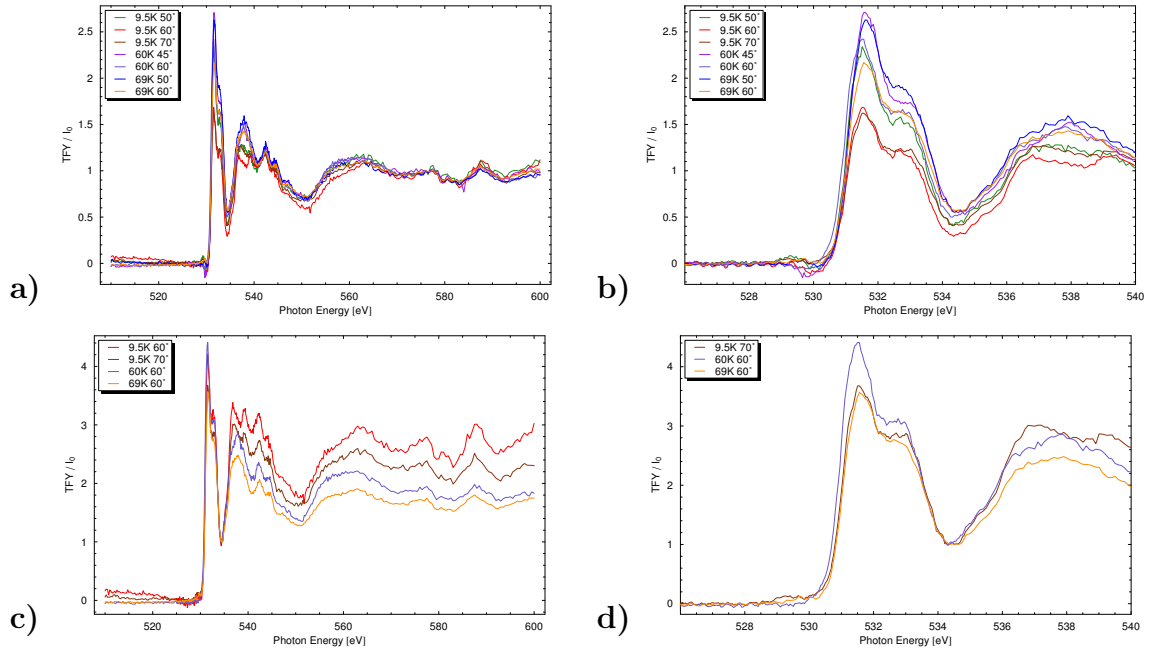
Figure B.15: Doping dependence of the O K edge c-axis spectra, extracted from spectra normalized using **a)** an energy well above the main edge and **c)** the notch between the prepeaks and main edge. Panels **b)** and **d)** detail the region around their respective prepeaks and edge.
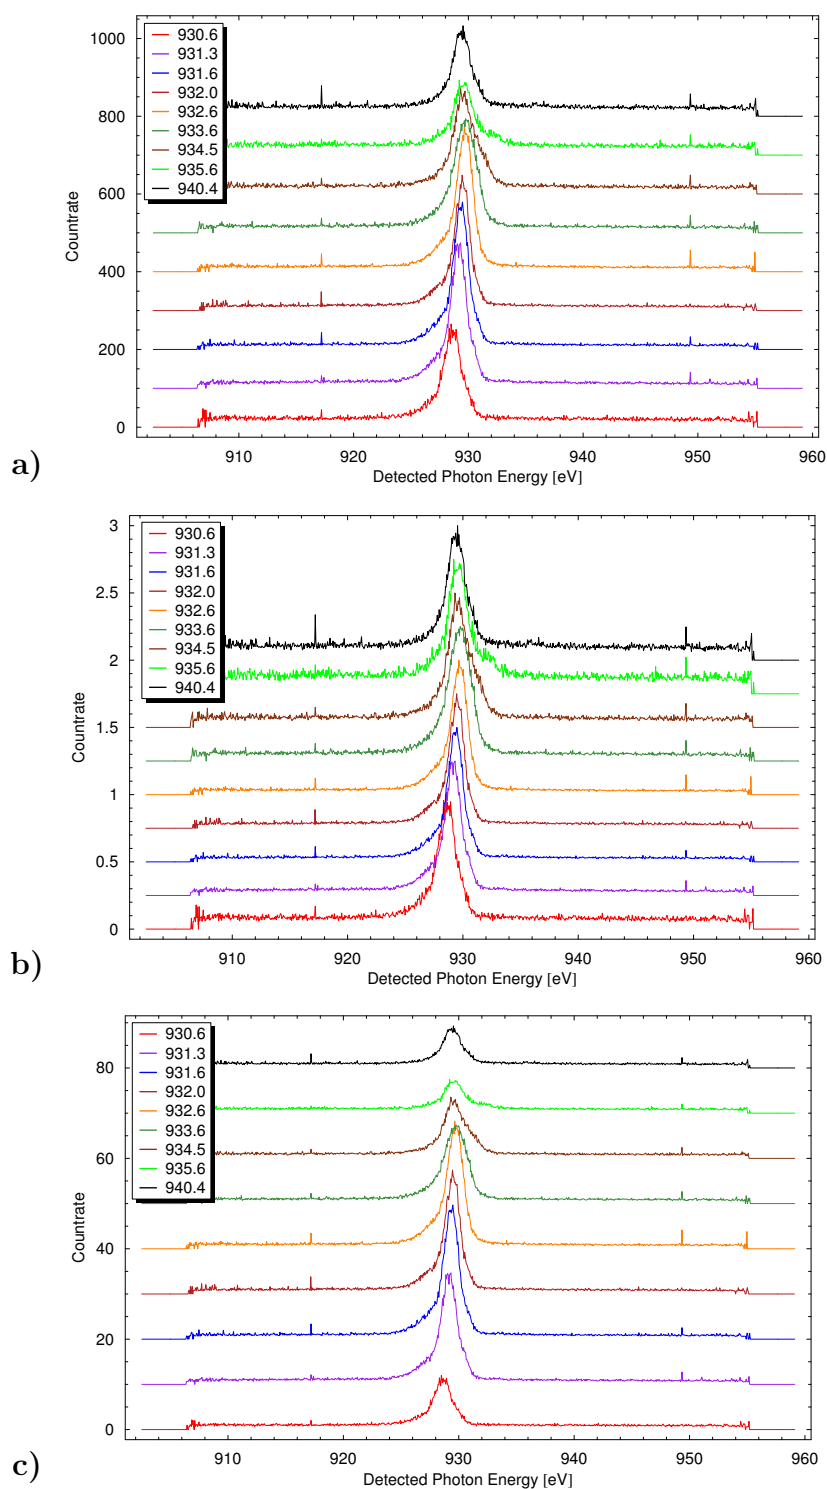
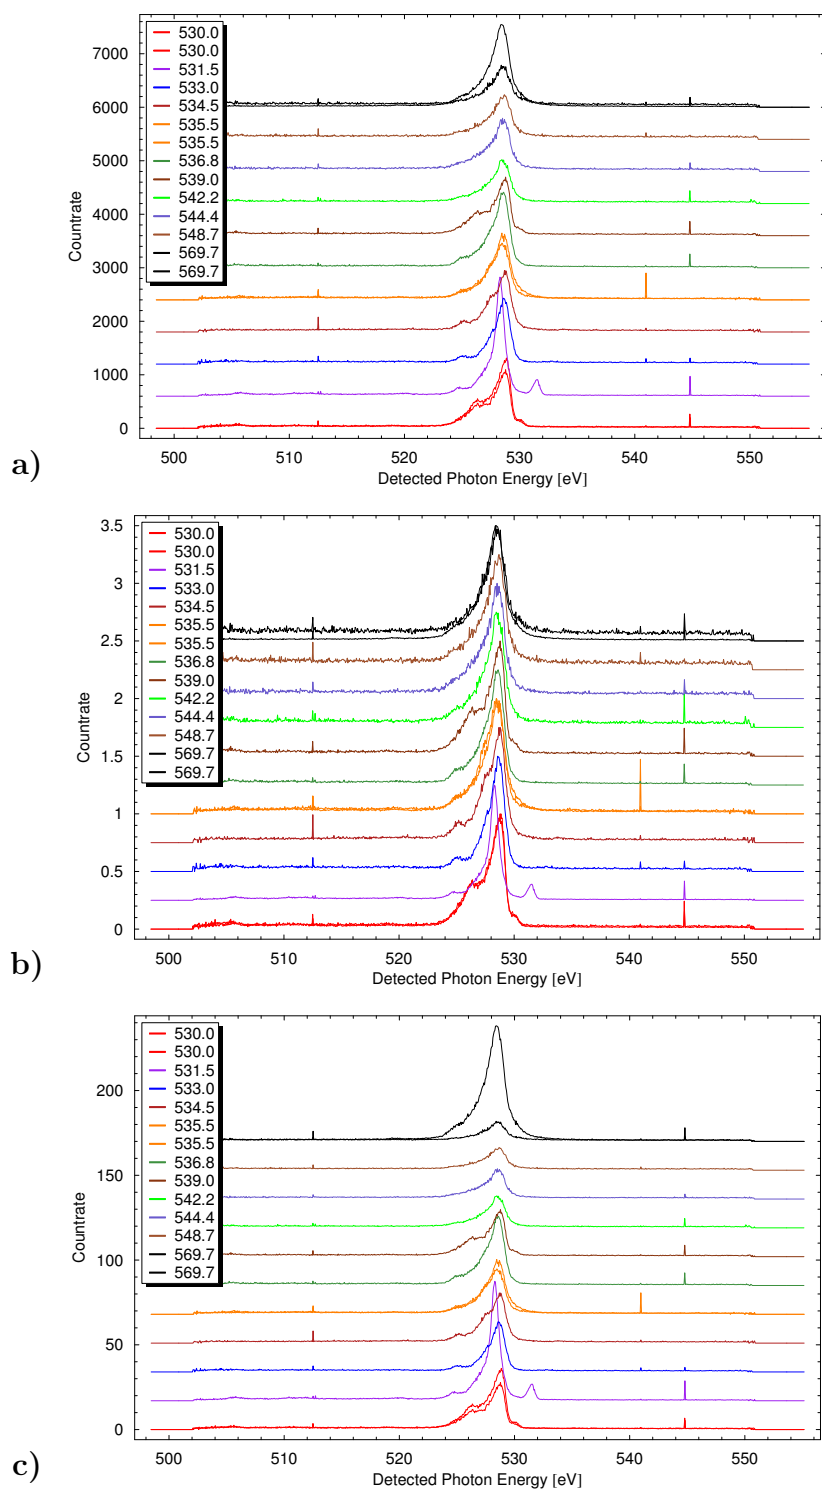Figure B.16: Copper L edge XES spectra, normalized by **a)** integral, **b)** peak height and **c)** background count rate.

Figure B.17: Oxygen K edge XES spectra, normalized by **a)** integral, **b)** peak height and **c)** background count rate.

Figure B.18: Copper *L* edge XES spectra, normalized by peak height, displayed in the form of a density plot.

Figure B.19: Oxygen *K* edge XES spectra, normalized by peak height, displayed in the form of a density plot.

C. Amusing Photo

"If that photo turns up on a thesis it'll make my whole year!" — Max Read, Communications and Thesis Coordinator, Faculty of Graduate Studies


Figure C.1: Fire Hydrant, a very nearly but not quite successful candidate for the UBC Board of Governors, meets with Dr. Martha Piper outside the Board and Senate room in Old Admin. Back row, left to right: Dr. Martha Piper (UBC President, 1997-2006), Brian Duong (student Governor, 2004-05), the author (student Governor, 2007-08). Front row, left to right: Quinn Omori (student Governor, 2005-06), Fire Hydrant (candidate, 2004, 2005 and 2006), and Tim Louman-Gardiner (student Governor, 2005-06).